Program Equivalence by Circular Reasoning

Dorel Lucanu, Vlad Rusu

HAL Id: hal-01065830

https://inria.hal.science/hal-01065830

Submitted on 18 Sep 2014

Program Equivalence by Circular Reasoning

Dorel Lucanu[1] and Vlad Rusu[2]

[1]Al. I. Cuza University of Iaşi, 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 written in languages having a rewriting-based operational semantics. The chosen equivalence is parametric in a so-called observation relation, and it says that two programs satisfying the observation relation will inevitably be, in the future, in the observation relation again. This notion of equivalence generalises several well-known equivalences and is appropriate for deterministic (or, at least, for confluent) programs. 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 given program-equivalence problem. We show that our approach is suitable for proving equivalence for terminating and non-terminating programs as well as for 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 the equivalence of (infinitely) many concrete programs obtained by replacing the variables by concrete statements or expressions. The approach is illustrated by proving program equivalence in two languages from different programming paradigms. The examples in the paper, as well as other examples, can be checked using an online tool.

1. Introduction

In this paper we propose a formal notion of program equivalence, together with a language-independent logic for expressing this notion and a deductive system for automatically proving it. Programs can belong to any language whose semantics, specified by a set of rewrite rules, is deterministic (or, at least, confluent). The equivalence we consider is parametric in a certain observation relation, and it requires that, for all programs satisfying the observation relation, their executions eventually lead them into satisfying the observation relation again. The proof system is circular: its conclusions can be re-used as hypotheses in a controlled way. Since the problem it tries to solve is undecidable, our proof system is not guaranteed to terminate. We prove that when it does terminate, however, it solves the program-equivalence problem as it is defined here.

The proposed framework is shown suitable for terminating and nonterminating programs as well as for concrete and for *symbolic programs*. The latter 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. Proving such equivalences is useful for ensuring the correctness of certain classes of syntax-directed source-to-source translations.

Example 1.1. We want to translate general programs with `for`-loops into programs with `while`-loops. This

---

amounts to translating the symbolic program in the left-hand side to the one in the right-hand side:

$$\texttt{for } I \texttt{ from } A \texttt{ to } B \texttt{ do } \{\, S \,\} \qquad I = A \texttt{ ;while } I \texttt{ <= } B \texttt{ do } \{\, S \,;\, I = I + 1 \,\}$$

Their symbolic variables I, A, B, S can be matched by, respectively, any identifier (I), arithmetical expressions (A, B), and program statement (S). We assume that the `for`-loop and `while`-loop statements have independent semantics (i.e., the `for` instruction is not desugared into to a `while` instruction) and the `for` loop does not modify the counter I, nor any program variables occurring in A, B (note that program variables are identifiers). If we prove the equivalence between these two symbolic programs then we also prove that every concrete instance of the `for`-loop is equivalent to its translation to the corresponding `while`-loop.

Example 1.2. The second example illustrates the equivalence of programs from another paradigm: corecursive programs. Here we consider corecursive programs over infinite sequences of integers (also called *streams*). Such a program is expressed using a set of equations; for each equation, the left-hand side is the name of a function being defined, possibly with parameters, and the right-hand side is the function's body. Let us consider the corecursive program consisting of the following equations:

$$\texttt{hd(}\; x : xs \;\texttt{)} \approx x \;;\qquad\qquad \texttt{tl(}\; x : xs \;\texttt{)} \approx xs \;;$$
$$\texttt{zero} \approx \texttt{0 : zero;} \qquad\qquad\quad \texttt{one} \approx \texttt{1 : one;}$$
$$\texttt{blink} \approx \texttt{0 : 1 : blink;} \qquad \texttt{zip(}\; xs \;,\; ys \;\texttt{)} \approx \texttt{hd(}\; xs \;\texttt{)} : \texttt{zip(}\; ys \;,\; \texttt{tl(}\; xs \;\texttt{));}$$

where x ranges over integers and xs over streams. Obviously, the complete evaluation of `zero` produces the infinite sequence $0 : 0 : 0 : \ldots$, and the evaluation of `one` produces the infinite sequence $1 : 1 : 1 : \ldots$, `blink` produces the infinite sequence $0 : 1 : 0 : 1 : \ldots$, and `zip(` xs, ys `)` produces a stream that alternates the elements of the two streams given to it as parameters. The function `hd(` xs `)` returns the first element of the stream xs (this is the only function in the languge that does not produce a stream), and `tl(` xs `)` returns the stream obtained from xs after removing the first element. A well-known equivalence over streams is that of `blink` and `zip(zero,one)`, for which many proofs can be found in the literature. We use this example to show that our notion of equivalence is general enough for being applicable to terminating programs as well as to non-terminating ones. The example also serves to illustrate the language-genericity of our approach.

A typical use of our framework consists of:

1. formally defining the syntax a programming language $\mathcal{L}$, whose concrete programs are ground terms over a certain signature, and whose symbolic programs are terms with variables over the same signature;
2. formally defining the operational semantics of $\mathcal{L}$ as a (possibly, conditional) term-rewriting system;
3. automatically constructing a new language definition $\mathcal{L} \times \mathcal{L}$, whose programs are pairs of programs of $\mathcal{L}$[1]
4. applying our deductive system to programs in $\mathcal{L} \times \mathcal{L}$.

Running the deductive system amounts essentially to symbolically executing the semantics of $\mathcal{L} \times \mathcal{L}$, which consists in applying the rewrite rules in the semantics with *unification* instead of matching; details are given in Section 4. This may lead to one 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 so-called *observationally equivalent configurations*).

---

[1] We have here developped the approach for the equivalence of programs belonging to one language $\mathcal{L}$ for simplicity reasons. However, considering two distinct languages $\mathcal{L}$ and $\mathcal{L}'$ poses no essential difficulty, as shown in [1] where we deal with another kind of program equivalence. The key step is building the *aggregation* $\mathcal{L} \amalg \mathcal{L}'$ of the two languages, which is essentially a union of the two languages that takes care of the possibly shared language elements. Then, any program in $\mathcal{L}$ or in $\mathcal{L}'$ is also a program in $\mathcal{L} \amalg \mathcal{L}'$, which reduces the equivalence of programs from the two languages $\mathcal{L}$ and $\mathcal{L}'$ to the one-language case $\mathcal{L} \amalg \mathcal{L}'$.

Contributions We propose a language-independent logic and a proof system suitable for stating and proving the equivalence of concrete and of symbolic programs as well as of terminating and non-terminating ones. Programs can be written in any deterministic (or, at least, confluent) language that has a formal operational semantics based on term rewriting. We prove the soundness and weak completeness of the proof system, which ensure that the system correctly solves the program equivalence problem as stated. The approach is illustrated on two different languages. The examples in the paper, as well as and other examples, can be tried using an online tool, available at `http://fmse.info.uaic.ro/tools/K/?tree=examples/prog-equiv/README`.

With respect to the conference paper [2]: the equivalence relation is reformulated in terms of a Linear Temporal Logic (LTL) formula, and the soundness/weak completeness proofs are simpler, thanks to an encoding of executions of our proof system as the building of proofs for the LTL formulas in question. The genericity of the approach is illustrated by considering programs in two different programming paradigms.

We also generalise (for the needs of the program-equivalence approach) a generic symbolic execution technique introduced in [3]: by executing semantical rules with unification instead of matching we also allow, e.g., the symbolic execution of symbolic statements in addition to the symbolic data considered in [3].

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 [4]. 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: [4] 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.

Other closely related recent works are [5, 6, 7]. 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, which focuses on recursive procedure definition) and proof systems, and formally prove their proof system's soundness. In [5] 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 [5] are of an input/output nature: for given inputs, programs generate equal outputs. 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 equivalence (parameterised by a certain observation relation) allows this, both for finitely and for infinitely many repeated observations. Moreover, we also show that some of the equivalences from [5] can be encoded in our approach by suitably choosing the observation relation.

Many works on program equivalence arise from the verification of compilation in a broad sense. At one end there is full compiler verification [8], and at the other end, the so-called translation validation, i.e., the individual verification of each compilation [9] (we only cite two of the most relevant recent works). As also observed by [4], 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. Another interesting application refers to information flow [10], where proving information flow properties requires to show that for any two executions of the program, the public variables cannot be distinguished. We believe that our approach can be adapted to capture such properties.

Several other works have targeted specific classes of languages: functional [11], microcode [12], CLP [13], Java [14]. The *schemata for Java code* used in [14] to validate program transformation rules are particular cases of what we call symbolic programs. There the authors adapt a specification of Java in Maude in order to execute schematic programs. In order to be less language-specific some works advocate the use of intermediate languages, such as [15], which works on the Boogie intermediate language. Only a few approaches, among which [8, 12], 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 mainstream industrial practice (see, e.g., [16] for a survey).

Our proof system is inspired by that of *circular coinduction* [17], which allows one to prove equalities of

data structures such as infinite streams and regular expressions. A notable difference between the present approach and [17] is that our specifications are essentially rewrite systems (meant to define the semantics of programming languages), whereas those of [17] are behavioural equational theories, a special class of equational specifications with visible and hidden sorts.

Symbolic linear temporal-logic model checking in term-rewriting systems, which we here use for proving program equivalence, was earlier studied in [18]. There are differences in expressiveness: we only use certain specific LTL formulas for encoding equivalence, whereas [18] handle full LTL; on the other hand, they consider unconditional term-rewriting systems only, whereas we also consider conditional term-rewriting systems. For our approach, which is based on programming-language semantics, having conditional rewriting systems is essential since unconditional rules are not expressive enough to express nontrivial languages semantics. There are also differences in the underlying deduction mechanisms: [18] rely on powerful unification-modulo-theories algorithms, while our unification algorithm delegates deduction to satisfiability modulo theory (SMT) solvers.

Organisation. After this introduction, Section 2.1 presents our running examples: IMP, a simple imperative language, and STREAM, a corecursive language for handling streams of integers. Both languages are defined in $\mathbb{K}$ [19], 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 mechanism for language definitions in Section 3, and show how $\mathbb{K}$ definitions are instances of that mechanism. We chose $\mathbb{K}$ for implementation purposes because we are familiar with it and because there are many languages defined in $\mathbb{K}$ to which our language-independent approach can be instantiated, cf. http://kframework.org for examples of $\mathbb{K}$ language definitions. In Section 4 we define a unification operation and prove some properties about it, which are used in Section 5 where we present a generic *symbolic execution* approach for languages defined in the proposed mechanism. We formally relate symbolic execution to concrete execution, which we use for proving the correctness properties (soundness and weak completeness) of our proof system.

In Section 6 we recap linear-temporal logic (LTL). This is then used in Section 7, which contains our proposed definition for program equivalence as the satisfiaction of certain LTL formulas over an execution of the transition system generated by (concretely) executing a pair of programs. The formula says that the programs will repeatedly satisfy a certain *observation relation*; this relation is a parameter of the approach. The syntax and semantics of a logic capturing the chosen notion of equivalence are also defined in Section 7.

The proof system for proving equivalence-formulas is presented in Section 8, together with its soundness and weak completeness. The properties say that, when it terminates, the proof system correctly answers to the question of whether its input (which is a set of formulas of program-equivalence logic) denotes equivalent programs. Their proofs are based on building proof witnesses for LTL formulas expressing equivalence. The witnesses are obtained by symbolically executing the pair of programs under investigation.

In Section 9 we report on a prototype implementation of the proof system in the $\mathbb{K}$ framework. This allows one to stay within the $\mathbb{K}$ environment when proving program equivalence for languages also defined in $\mathbb{K}$. Finally, the conclusion and future work are presented in Section 10.

2. Two Examples of Programming Languages and their Semantics in $\mathbb{K}$

We use two different languages as running examples: IMP, a simple imperative language, and STREAM, a language for manipulating integer streams. We present their formal definitions in the $\mathbb{K}$ framework [19], a formal environment for defining programming languages, type systems, and analysis tools. The main ingredients of a $\mathbb{K}$ definition are *computations*, *configuration*, and *rules*. Computations are sequences of elementary computational tasks, which consist of, e.g., adding two numbers, or transforming the program being executed. A configuration is a nested structure of *cells* that include all the data structures required for executing a program. The rules describe how the configurations are modified when the computational tasks are performed. For details on the theoretical background of $\mathbb{K}$ readers can consult [19].

$\mathbb{K}$ language definitions can be executed and analysed using tools from the $\mathbb{K}$ environment. Examples of language definitions and related analysis tools can be found on the web page http://kframework.org.

$$
\begin{array}{ll}
Int & ::= \text{domain of integer numbers (including operations)} \\
Bool & ::= \text{domain of boolean constants (including operations)} \\
Id & ::= \text{domain of identifiers}
\end{array}
$$

$$
\begin{array}{ll}
AExp ::= Int \mid Id & BExp ::= Bool \\
\quad \mid AExp \text{ / } AExp \text{ [strict]} & \quad \mid AExp \text{ <= } AExp \text{ [strict]} \\
\quad \mid AExp \text{ * } AExp \text{ [strict]} & \quad \mid \texttt{not } BExp \text{ [strict]} \\
\quad \mid AExp \text{ + } AExp \text{ [strict]} & \quad \mid BExp \texttt{ and } BExp \text{ [strict(1)]} \\
\quad \mid (AExp) & \quad \mid (BExp)
\end{array}
$$

$$
\begin{array}{ll}
Stmt ::= \texttt{skip} \mid Stmt \text{ ; } Stmt & \mid \{ \, Stmt \, \} \\
\quad \mid Id \text{ = } AExp & \mid \texttt{while } BExp \texttt{ do } Stmt \\
\quad \mid \texttt{if } BExp \texttt{ then } Stmt & \mid \texttt{for } Id \texttt{ from } AExp \texttt{ to } AExp \\
\qquad \texttt{else } Stmt \text{ [strict(1)]} & \qquad \texttt{do } Stmt \text{ [strict(2, 3)]}
\end{array}
$$

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

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

2.1. IMP - A Simple Imperative Language

The first 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 [19]. The syntax of IMP is described in Figure 1 and should mostly be self-explanatory. The annotation *strict*, which is not part of the syntax, means that 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 enviroment mapping identifiers to (integer) values. In $\mathbb{K}$, this is written as a nested structure of *cells*: here, a top cell cfg, having a cell k containing code and a cell env containing the environment (see Figure 3). The sort $Code^2$ contains statements and arithmetic and Boolean expressions. The empty code is denoted by $\cdot$, and code sequencing is denoted by $\curvearrowright$. Note that this is different from the (language-specific) sequencing operation ; of IMP.

The cell k includes the code to be executed, represented as a list of computation tasks $C_1 \curvearrowright C_2 \curvearrowright \ldots$, meaning that first C_1 will be executed, then C_2, etc. Computation tasks are typically the evaluation of statements and elementary expressions. An example of sequence of computations is given in Figure 3b); this sequence is obtained by applying the *heating* rules generated by the strict attribute for the statement if and the operator <. The heating/cooling rules are explained later. The cell env is an environment that binds the program variables to values; such a binding is written as a set of bindings of the form, e.g., $x \mapsto 3$.

The semantics of IMP is given by a set of rules (see Figure 2) that say how the configuration evolves when the first computation task (e..g., a statement or expression) from the k cell is executed or evaluated. Suspension 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.

In Figure 2, the operations *lookup* : $Map \times Id \to Int$ and *update* : $Map \times Id \times Int \to Map$ are not part of the IMP syntax: they belong to a semantic 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 2 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 *heating/cooling* rules (automatically generated by $\mathbb{K}$):

$$
\langle\langle \texttt{if } BE \texttt{ then } S_1 \texttt{ else } S_2 \cdots\rangle_k \cdots\rangle_{cfg} \Rightarrow \langle\langle BE \curvearrowright \texttt{if } \square \texttt{ then } S_1 \texttt{ else } S_2 \cdots\rangle_k \cdots\rangle_{cfg}
$$
$$
\langle\langle B \curvearrowright \texttt{if } \square \texttt{ then } S_1 \texttt{ else } S_2 \cdots\rangle_k \cdots\rangle_{cfg} \Rightarrow \langle\langle \texttt{if } B \texttt{ then } S_1 \texttt{ else } S_2 \cdots\rangle_k \cdots\rangle_{cfg}
$$

---

[2] In the $\mathbb{K}$ terminology the sort *Code* is called *K*. We changed its name in order to avoid confusions due to name overloading.

$$\langle\langle I_1 \text{ + } I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle I_1 +_{Int} I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle I_1 \text{ * } I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle I_1 *_{Int} I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle I_1 \text{ / } I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \wedge I_2 \neq 0 \Rightarrow \langle\langle I_1/_{Int}I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle I_1 \text{ <= } I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle I_1 \leq_{Int} I_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle true \text{ and } B \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle B \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle false \text{ and } B \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle false \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{not } true \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle false \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{not } false \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle true \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{skip } \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle S_1; S_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle S_1 \curvearrowright S_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{\{ } S \text{ \} } \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle S \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{if } true \text{ then } S_1 \text{ else } S_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle S_1 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{if } false \text{ then } S_1 \text{ else } S_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle S_2 \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{while } B \text{ do } S \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow$$
$$\langle\langle \text{if } B \text{ then\{ } S \text{ ;while } B \text{ do } S \text{ \}else skip } \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle \text{for } X \text{ from } I_1 \text{ to } I_2 \text{ do } S \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow$$
$$\langle\langle X \text{ = } I_1 \text{ ;if } X \text{ <= } I_2 \text{ then\{ } S \text{ ;for } X \text{ from } I_1 +_{Int} 1 \text{ to } I_2 \text{ do } S \text{ \} else skip } \text{ ...}\rangle_{\mathsf{k}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle X \text{ ...}\rangle_{\mathsf{k}}\langle Env\rangle_{\mathsf{env}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle lookup(Env, X) \text{ ...}\rangle_{\mathsf{k}}\langle Env\rangle_{\mathsf{env}} \text{ ...}\rangle_{\mathsf{cfg}}$$

$$\langle\langle X \text{ = } I \text{ ...}\rangle_{\mathsf{k}}\langle Env\rangle_{\mathsf{env}} \text{ ...}\rangle_{\mathsf{cfg}} \Rightarrow \langle\langle \text{ ...}\rangle_{\mathsf{k}}\langle update(Env, X, I)\rangle_{\mathsf{env}} \text{ ...}\rangle_{\mathsf{cfg}}$$

Fig. 2. $\mathbb{K}$ Semantics of IMP

$$Cfg \quad ::= \langle\langle Code\rangle_{\mathsf{k}}\langle Map\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}}$$

a) $\mathbb{K}$ Configuration of IMP

$$\left\langle \begin{array}{l} \langle \text{x } \curvearrowright \square \text{ < 0 } \curvearrowright \text{ if (} \square \text{) y = 0; else y = 1;}\rangle_{\mathsf{k}} \\ \langle \text{x } \mapsto 3 \quad \text{y } \mapsto -7\rangle_{\mathsf{env}} \end{array} \right\rangle_{\mathsf{cfg}}$$

b) An IMP configuration snapshot

Fig. 3.

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. Finally, a set of rules saying that $\cdot$ is a neutral element for $\curvearrowright$, and that $\curvearrowright$ is associative, is (conceptually) automatically included in the semantics of any language.

2.2. STREAM - a Simple Language for Corecursive Programs

Corecursive programs differ from recursive ones by the fact that their termination condition is missing. Besides functional languages, which typically use corecursion for handling infinite data structures, several other languages have been extended to support such features (see, e.g., [20] for an extension of Prolog, and [21] for an extension of Java). An example of a corecursive program was given in Section 1. Here we present a simple language for writing such programs over integer streams (= infinite sequence of integers). The standard semantics for corecursive functions is based on lazy evaluation, which delays the evaluation of expressions until their value is needed. For infinite expressions this evaluation is always partial, in the sense that only a finite part of the infinite expression is evaluated, e.g., a finite prefix of an infinite stream.

Therefore, we say that a stream expression is a *result value* if it is of the form $i : SE$, where the integer i is the first element of the stream and SE is the rest of the stream expression. Beside the constructor _ : _, two functions, often called destructors, are essential in handling streams: hd(xs), which returns the first element of the stream, and tl(xs), which returns the stream obtained after the first element is removed.

The syntax of the STREAM language is given in Figure 4. There are three expression kinds: $BExp$ - for boolean expressions, $IExp$ - integer expressions, and $SExp$ - stream expressions. Lists of stream expressions, separated by commas, are denoted by $List\{SExp, ","\}$. In $\mathbb{K}$, lists of any syntactical categories can be defined with any separators, including the empty (space) separator denoted by "". The operator $X \triangleleft B \triangleright Y$ is the

$$BExp ::= Bool$$
$$| \ IExp \ \texttt{=} \ IExp \ [\text{strict}]$$
$$| \ BExp \ \texttt{\&} \ BExp \ [\text{strict}(1)]$$
$$| \ \texttt{!} \ BExp \ [\text{strict}]$$
$$IExp ::= Int$$
$$| \ \texttt{hd} \ \texttt{(} \ SExp \ \texttt{)} \ [\text{strict}]$$
$$| \ IExp \ \texttt{+} \ IExp \ [\text{strict}]$$
$$| \ IExp \vartriangleleft BExp \vartriangleright SExp \ [\text{strict}(2)]$$

$$SExp ::= Id$$
$$| \ \texttt{tl} \ \texttt{(} \ SExp \ \texttt{)} \ [\text{strict}]$$
$$| \ Id \ \texttt{(} \ SExps \ \texttt{)}$$
$$| \ IExp : SExp$$
$$Exp ::= IExp | SExp$$

$$SSpec ::= Id \ \texttt{:=} \ IExp \ \texttt{;}$$
$$| \ Id \approx SExp \ \texttt{;}$$
$$| \ Id \ \texttt{(} \ Ids \ \texttt{)} \ \texttt{:=} \ IExp \ \texttt{;}$$
$$| \ Id \ \texttt{(} \ Ids \ \texttt{)} \approx SExp \ \texttt{;}$$
$$SPgm ::= SSpecs \ Exp$$
$$SExps ::= List\{SExp, ","\}$$
$$Ids ::= List\{Id, ","\}$$
$$SSpecs ::= List\{SSpec, ""\}$$

$$Code ::= IExp \ | \ SExp \ | \ Exp \ | \ SSpec \ | \ SSpecs \ | \ SPgm \ | \ Code \curvearrowright Code$$

Fig. 4. $\mathbb{K}$ Syntax of STREAM

conditional operator if B then X else Y written in a Hoare-like syntax, chosen here because it is more compact than its alternatives. There are two kinds of statements (specifications *SSpec*): integer function specifications, written as f := ... or $f(\ldots)$:= ... (these are the usual, recursive functions), and stream specifications, written as $s \approx \ldots$ or $s(\ldots) \approx \ldots$ (these are the corecursive functions)

A STREAM program is a sequence of function specifications, followed by an expression to be evaluated. The $\mathbb{K}$ configuration for STREAM programs is represented in Figure 5. As the snapshot suggests, the cell specs stores definitions of recursive and corecursive functions. The right-hand side of a function definition is a λ-expression, defined as follows:

$$Val ::= \lambda \ \texttt{(} \ Ids \ \texttt{)} \ \texttt{.} \ SExp$$

The cell out includes results of evaluations, which can be integers or stream result values, depending on the type of the expression to be evaluated. This cell is essential for the stream equivalence definition since it defines their observational relation. Note that the evaluation of stream expressions is an nonterminating process and the out cell includes only finite approximations of streams.

The $\mathbb{K}$ semantics of STREAM is given in Figure 6. The semantics of the boolean/integer operators that are similar to those from the IMP definition and are omitted. The function $append(OE, I)$ appends the integer I to the "computed" part of the OE stream expression, i.e., if $OE \triangleq I_1 : \ldots I_k : OE'$, where $I_1, \ldots, I_k$ are integers and OE' is an *SExp* variable, then $append(OE, I) = I_1 : \ldots I_k : I : OE'$. The semantics of a function call expression consists of replacing the expression with the function body, where the formal parameters are replaced by the actual arguments (if any). The other rules should be self-explanatory.

Example 2.1. We illustrate the semantics of STREAM on the following example. Assume that the current configuration is $\langle\langle\texttt{tl(one)}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$, where Z is a symbolic variable of sort *SExp*. In order to evaluate the expression tl(one), the above configuration is *heated*, by applying the rule generated from the corresponding strict attribute, to $\langle\langle\texttt{one}\curvearrowright\texttt{tl(}\square\texttt{)}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$. Now, the rule evaluating stream functions without parameters (the ninth one in Figure 6) is applied and generates the term $\langle\langle\texttt{1:one}\curvearrowright\texttt{tl(}\square\texttt{)}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$. The expression 1:one is a result value and the corresponding cooling rule is applied, producing $\langle\langle\texttt{tl(1:one)}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$. By applying the rule for tl, followed by the rule for function calls, we obtain $\langle\langle\texttt{1:one}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$. Since the content of the k cell consists only of 1:one, the rule writing in the out cell (the first one in Figure 6) can be applied the following configuration is obtained: $\langle\langle\texttt{one}\rangle_\text{k}\langle\texttt{one} \mapsto \lambda(). \ \texttt{1:one}\rangle_\text{specs}\langle\texttt{1} : \texttt{Z}\rangle_\text{out}\rangle_\text{cfg}$. A finite approximation of a stream can be represented in the out cell by a stream expression $i_1 : i_2 : \ldots : i_k : \texttt{Z}$, where Z is a special symbolic variable.

3. Language Definitions

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:

$$\langle\langle Code\rangle_{\mathsf{k}}\langle Map\rangle_{\mathsf{specs}}\langle SExp\rangle_{\mathsf{out}}\rangle_{\mathsf{cfg}}$$

a) $\mathbb{K}$ configuration of STREAM

$$\left\langle \begin{array}{c} \langle\mathtt{tl(blink)}\curvearrowright\mathtt{hd(\ \square\ )}\rangle_{\mathsf{k}} \\ \langle\mathtt{zero}\mapsto\lambda().\,\mathtt{0\ :\ zero}\quad\mathtt{one}\mapsto\lambda().\,\mathtt{1\ :\ one}\quad\mathtt{blink}\mapsto\lambda().\,\mathtt{0\ :\ 1\ :\ blink}\rangle_{\mathsf{specs}} \\ \langle\mathtt{1\ :\ 0\ :\ Z}\rangle_{\mathsf{out}} \end{array} \right\rangle_{\mathsf{cfg}}$$

b) A STREAM configuration snapshot

Fig. 5.

$$\langle\langle I\ :\ SE\rangle_{\mathsf{k}}\langle OE\rangle_{\mathsf{out}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle SE\rangle_{\mathsf{k}}\langle append(OE,I)\rangle_{\mathsf{out}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle I\rangle_{\mathsf{k}}\langle OE\rangle_{\mathsf{out}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle\rangle_{\mathsf{k}}\langle append(OE,I)\rangle_{\mathsf{out}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle I_1\ =\ I_2\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\wedge I_1 =_{Int} I_2\Rightarrow\langle\langle true\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle I_1\ =\ I_2\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\wedge I_1 \neq_{Int} I_2\Rightarrow\langle\langle false\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle\mathtt{hd(}\ I\ :\ \_\ \mathtt{)}\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle I\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle IE_1 \triangleleft true \triangleright IE_2\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle IE_1\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle IE_1 \triangleleft false \triangleright IE_2\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle IE_2\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle\mathtt{tl(}\ \_\ :\ SE\ \mathtt{)}\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle SE\ \cdots\rangle_{\mathsf{k}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle F\ \cdots\rangle_{\mathsf{k}}\langle\cdots\ F\mapsto\lambda().SE\ \cdots\rangle_{\mathsf{specs}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle SE\ \cdots\rangle_{\mathsf{k}}\langle\cdots\ F\mapsto\lambda().SE\ \cdots\rangle_{\mathsf{specs}}\,\cdots\rangle_{\mathsf{cfg}}$$

$$\langle\langle F(Vs)\ \cdots\rangle_{\mathsf{k}}\langle\cdots\ F\mapsto\lambda(Xs).SE\ \cdots\rangle_{\mathsf{specs}}\,\cdots\rangle_{\mathsf{cfg}}\Rightarrow\langle\langle SE[Vs/Xs]\ \cdots\rangle_{\mathsf{k}}\langle\cdots\ F\mapsto\lambda(Xs).SE\ \cdots\rangle_{\mathsf{specs}}\,\cdots\rangle_{\mathsf{cfg}}$$

where the operation $[\_/\_]$ denotes syntactical substitution.

Fig. 6. $\mathbb{K}$ Semantics of STREAM.

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 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 $\Sigma\setminus\Sigma^{Data}$, and that terms of sort *Cfg* have subterms 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*, disjoint from Σ, let $T_\Sigma(Var)$ denote the free Σ-algebra of terms with variables, $T_{\Sigma,s}(Var)$ denote the set of terms of sort s with variables, and $var(t)$ denote the set of variables occurring in the term t. For terms $t_1,\ldots,t_n$ we let $var(t_1,\ldots,t_n)\triangleq var(t_1)\cup\cdots var(t_n)$. For any *substitution* $\sigma : Var \to T_\Sigma(Var)$ and term $t\in T_\Sigma(Var)$ we denote by $t\sigma$ the term obtained by applying the substitution σ to t. We use the *diagrammatical order* for the composition of substitutions, i.e., for substitutions σ and σ', the composition $\sigma\sigma'$ consists in first applying σ then σ'.

2. A Σ-algebra $\mathcal{T}$, over which the semantics of the language is defined. $\mathcal{T}$ interprets the data sorts (those included in the subsignature Σ^{Data}) according to some Σ^{Data}-algebra $\mathcal{D}$. $\mathcal{T}$ interprets non-data sorts as ground terms over the signature of the form

$$(\Sigma\setminus\Sigma^{Data})\cup\mathcal{D} \tag{1}$$

   i.e., the elements of $\mathcal{D}$ are added to the signature $\Sigma\setminus\Sigma^{Data}$ as constants of their respective sorts. That is, a language is parametric in the way its data is implemented; it just assumes there is such an implementation $\mathcal{D}$. This is important for technical reasons (existence of a unique most general unifier, discussed below). Let $\mathcal{T}_s$ denote the elements of $\mathcal{T}$ that have the sort s; the elements of $\mathcal{T}_{Cfg}$ are called *configurations*. Any *valuation* $\rho : Var \to \mathcal{T}$ is extended to a (homonymous) Σ-algebra morphism $\rho : T_\Sigma(Var)\to\mathcal{T}$. 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 $b\rho=\mathcal{D}_{true}$, where $b\rho$ is the Boolean value obtained by applying ρ to b. For simplicity, we often write $true, false$ instead of $\mathcal{D}_{true}, \mathcal{D}_{false}$.

3. A set $\mathcal{S}$ of rewrite rules $l\wedge b\Rightarrow r$, whose formal definition is given later in the section.

We explain these concepts on the IMP example. Each nonterminal from the syntax ($Int, Bool, AExp, \ldots$) 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 $\_+\_ : AExp \times AExp \to AExp$. These operations define the constructors of the result sort. For the configuration sort Cfg, the only constructor is $\langle\langle\_\rangle_k\langle\_\rangle_{env}\rangle_{cfg} : Code \times Map_{Id,Int} \to Cfg$. 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 2) as the corresponding usual operation on integers, $Bool$ as the set of Boolean values $\{false, true\}$, the operation like $\wedge_{Bool}$ as the usual Boolean operations, the sort $Map_{Id,Int}$ as the set of maps $X \mapsto I$, where X ranges over identifiers Id and I over the integers Int. 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 the modification (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 $\mathcal{D}_{true}(= true)$.

The rewrite rules describe the transitions over configurations, whose formal definition is given below.

Definition 3.1 (pattern [22]). A *pattern* is an expression of the form $\pi \wedge \phi$, where $\pi \in T_{\Sigma,Cfg}(Var)$ is a *basic pattern* and $\phi \in T_{\Sigma,Bool}(Var)$ is a boolean term called the pattern's *condition*. If $\gamma \in \mathcal{T}_{Cfg}$ and $\rho : Var \to \mathcal{T}$ we write $(\gamma, \rho) \models \pi \wedge \phi$ for $\gamma = \pi\rho$ and $\rho \models \phi$. We let $\llbracket \pi \wedge \phi \rrbracket$ denote the set $\{\gamma \mid \text{there exists } \rho \text{ such that } (\gamma, \rho) \models \pi \wedge \phi\}$.

For any set of patterns Φ we let $\llbracket \Phi \rrbracket \triangleq \bigcup_{\varphi \in \Phi} \llbracket \varphi \rrbracket$. A basic pattern π thus defines a set of (concrete) configurations, and the condition b gives additional constraints these configurations must satisfy. In [22] 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 often 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 $\langle\langle I_1 \,/\, I_2 \curvearrowright C\rangle_k\langle Env\rangle_{env}\rangle_{cfg} \wedge I_2 \neq_{Int} 0$. An example of configuration that satisfies the second pattern is $\langle\langle(4 \,/\, 3) \curvearrowright \text{skip}\rangle_k\langle\text{a}\mapsto 5\rangle_{env}\rangle_{cfg}$.

Remark 3.1. Any pattern $\pi \wedge \phi$ can be transformed into a "semantically equivalent" pattern $\pi' \wedge \phi'$ (i.e.., $\llbracket \pi \wedge \phi \rrbracket = \llbracket \pi' \wedge \phi' \rrbracket$) such that π' is linear and all its data subterms are variables. For this, just replace all duplicated variables and all non-variable data subterms of π by fresh variables, and add constraints to equate in ϕ the fresh variables to the subterms they replaced. The transformations are presented in detail in [23].

Example 3.1. The pattern $\langle\langle X \,/\, Y\rangle_k\langle Y \mapsto A +_{Int} 1\rangle_{env}\rangle_{cfg} \wedge A \neq_{Int} -1$ with X, Y variables of sort Id and A of sort Int is nonlinear because Y occurs twice. Moreover, it contains the non-variable data terms $A +_{Int} 1$. It is transformed into the pattern $\langle\langle X \,/\, Y\rangle_k\langle Y' \mapsto A'\rangle_{env}\rangle_{cfg} \wedge Y' =_{Id} Y \wedge_{Bool} A' =_{Int} A +_{Int} 1 \wedge_{Bool} A \neq_{Int} -1$.

The proof system we propose in Section 8 uses as a basic block the testing of inclusions of the form $\llbracket \varphi \rrbracket \subseteq \llbracket \varphi' \rrbracket$. Therefore we need criteria for such inclusions. The following propositions define sufficient conditions.

Proposition 3.1. Let π' and π be two basic patterns and σ a substitution such that $\pi'\sigma = \pi$, $y\sigma = y$ for all $y \notin var(\pi')$, and $var(\pi') \cap var(\pi \wedge \phi) = \emptyset$. Then $\llbracket \pi \wedge \phi \rrbracket = \llbracket \pi' \wedge (\bigwedge_\sigma \wedge \phi) \rrbracket$, where $\bigwedge_\sigma$ denotes the conjunction $\bigwedge_{x \in var(\pi')} x = x\sigma$.

Proof. We prove the equality of the two sets by double inclusion.

($\subseteq$) Let $\gamma \in \llbracket \pi \wedge \phi \rrbracket$. Then there is $\rho : Var \to \mathcal{T}$ such that $\gamma = \pi\rho$ and $\rho \models \phi$. Let ρ' denote the valuation $\rho' : Var \to \mathcal{T}$ given by $x\rho' = x\sigma\rho$ for $x \in var(\pi')$, and $y\rho' = y\rho$ for $y \notin var(\pi')$. We have $\pi'\rho' = \pi'\sigma\rho = \pi\rho = \gamma$. Since $var(\phi) \cap var(\pi') = \emptyset$, it follows that $\rho' \models \phi$ iff $\rho \models \phi$. Finally, $var(\pi') \cap var(\pi) = \emptyset$ implies $var(\pi') \cap var(\sigma(x)) = \emptyset$ and hence $x\sigma\rho' = x\sigma\rho$ for $x \in var(\pi')$. Now, $\pi'\rho' = \gamma$ that implies $\gamma \in \llbracket \pi' \wedge (\bigwedge_\sigma \wedge \phi) \rrbracket$. Since γ was chosen arbitrarily, it follows that $\llbracket \pi \wedge \phi \rrbracket \subseteq \llbracket \pi' \wedge (\bigwedge_\sigma \wedge \phi) \rrbracket$.

($\supseteq$) Assume that $\gamma \in \llbracket (\pi' \wedge (\bigwedge_\sigma \wedge \phi) \rrbracket$. Then there is $\rho' : Var \to \mathcal{T}$ such that $\gamma = \pi'\rho'$, and $\rho' \models (\bigwedge_\sigma \wedge \phi)$. From $\rho' \models \bigwedge_\sigma$ we get $x\rho' = x\sigma\rho'$ for $x \in var(\pi')$, which implies $\gamma = \pi'\rho' = \pi'\sigma\rho = \pi\rho$. Hence $\gamma \in \llbracket \pi \wedge \phi \rrbracket$. Since γ was chosen arbitrarily, it follows that $\llbracket \pi' \wedge (\bigwedge_\sigma \wedge \phi) \rrbracket \subseteq \llbracket \pi \wedge \phi \rrbracket$. $\square$

Proposition 3.2. Let $\pi \wedge \phi$ and $\pi' \wedge \phi'$ be two patterns and σ a substitution such that $\pi'\sigma = \pi$, $y\sigma = y$ for all $y \notin var(\pi')$, and $var(\pi') \cap var(\pi \wedge \phi) = \emptyset$. If $\bigwedge_\sigma \wedge \phi$ implies ϕ', then $\llbracket \pi \wedge \phi \rrbracket \subseteq \llbracket \pi' \wedge \phi' \rrbracket$.

Proof. Let $\gamma \in [\![\pi \wedge \phi]\!]$. Then there is a valuation ρ such that $\gamma = \pi\rho$ and $\rho \models \phi$. Let ρ' be defined as in Proposition 3.1. Then $\pi'\rho' = \gamma$ and $\rho' \models \bigwedge_\sigma \wedge \phi$ that implies $\rho' \models \phi'$ by the hypotheses. Hence $\gamma \in [\![\pi' \wedge \phi']\!]$. Since γ was chosen arbitrarily the conclusion of the proposition follows. $\quad\square$

Remark 3.2. The conditions $var(\pi') \cap var(\pi \wedge \phi) = \emptyset$ and $y\sigma = y$ for all $y \notin var(\pi')$ required by Proposition 3.1 and Proposition 3.2 can be easily obtained by a variable renaming.

Proposition 3.3. Let $\pi \wedge \phi$ and $\pi' \wedge \phi'$ two patterns such that there is a substitution σ with $(\pi' \wedge \phi')\sigma = \pi \wedge \phi$. Then $[\![\pi \wedge \phi]\!] \subseteq [\![\pi' \wedge \phi']\!]$.

Proof. Let $\gamma \in [\![\pi \wedge \phi]\!]$. Then there is a valuation ρ such that $\gamma = \pi\rho$ and $\rho \models \phi$. Let ρ' be defined by $x\rho' = x\sigma\rho$ for each x in *Var*. It follows $\pi'\rho' = \pi'\sigma\rho = \pi\rho = \gamma$ and similarly $\phi'\rho' = \phi\rho$ that implies $\rho' \models \phi'$. Hence $\gamma \in [\![\pi' \wedge \phi']\!]$. Since γ was chosen arbitrarily the conclusion of the proposition follows. $\quad\square$

We are now ready to define semantical rules and the transition system that they generate.

Definition 3.2 (semantical rule and transition system [22]). A *rule* is a pair of patterns of the form $l \wedge b \Rightarrow r$ (where r is the pattern $r \wedge true$). Any set $\mathcal{S}$ of rules defines a labelled transition system $(\mathcal{T}_{Cfg}, \Rightarrow_\mathcal{S})$, where $\gamma \Rightarrow_\mathcal{S} \gamma'$ iff there exist $(l \wedge b \Rightarrow r) \in \mathcal{S}$ and $\rho : Var \to \mathcal{T}$ such that $(\gamma, \rho) \models l \wedge b$ and $(\gamma', \rho) \models r$.

Assumption 1. We assume without loss of generality that for all rules $l \wedge b \Rightarrow r \in \mathcal{S}$, l is linear and all its data subterms are variables. The generality is not restricted because the pattern $l \wedge b$ in the rule $l \wedge b \Rightarrow r \in \mathcal{S}$ can always be replaced by an equivalent one (cf. Remark 3.1) with the desired properties. This transformation of rules does not modify the transition system $(\mathcal{T}_{Cfg}, \Rightarrow_\mathcal{S})$.

Examples of semantical rules are included in Figure 2 and Figure 6, which give the semantics of the two languages IMP and, respectively, STREAM.

4. Unification

We shall be using unification for defining the symbolic execution of programs, which is used in our program-equivalence proof system. We recall that, in general, a unifier of two terms t_1, t_2 is a substituion of their variables that, when applied to the two terms, makes them equal. We shall call hereafter this unification a *symbolic* unification, in order to distinguish it from what we call *concrete* unification, introduced below.

Definition 4.1 (Unifiers). A *symbolic unifier* of two terms t_1, t_2 is any substitution $\sigma : var(t_1) \uplus var(t_2) \to T_\Sigma(Z)$ for some set Z of variables such that $t_1\sigma = t_2\sigma$. A *concrete unifier* of terms t_1, t_2 is any valuation $\rho : var(t_1) \uplus var(t_2) \to \mathcal{T}$ such that $t_1\rho = t_2\rho$. A symbolic unifier σ of two terms t_1, t_2 is a *most general unifier of t_1, t_2 with respect to concrete unification* whenever, for all concrete unifiers ρ of t_1 and t_2, there is a valuation η such that $\sigma\eta = \rho$. We often call a symbolic unifier satisfying the above a *most general unifier*[3].

Two terms are symbolically (resp. concretely) unifiable if they have a symbolic (resp. concrete) unifier.

Example 4.1. The terms $f(x, g(y))$ and $f(t, g(z))$ are symbolically unifiable, by the substitution $x \mapsto t, y \mapsto z$, extended to the identity for the other variables occuring in the terms. Assuming that $g : Int \to s$ and $f : Int \times s \to s$ and are non-Data function symbols in Σ (i.e., their codomain sort s is not a Data sort), the two terms are also concretely unifiable, e.g., by any valuation that maps all variables x, y, z, t to 1.

The next lemma gives conditions under which concretely unifiable terms are symbolically unifiable.

Lemma 4.1. All linear, concretely unifiable terms $t_1, t_2 \in T_\Sigma(Var)$, such that all their data subterms are variables, are symbolically unifiable by a most general unifier $\sigma_{t_2}^{t_1} : var(t_1) \uplus var(t_2) \to T_\Sigma(var(t_1) \uplus var(t_2))$.

Proof. By induction on the structure of, say, t_1. In the base case, $t_1 \in Var$, and we take $\sigma_{t_2}^{t_1} \triangleq (t_1 \mapsto t_2) \uplus Id|_{var(t_2)}$, i.e., $\sigma_{t_2}^{t_1}$ maps t_1 to t_2, and is the identity on $var(t_2)$. Obviously, $\sigma_{t_2}^{t_1}$ is a unifier of t_1, t_2, since $t_1\sigma_{t_2}^{t_1} = t_2$. To show that $\sigma_{t_2}^{t_1}$ is most general, consider any concrete unifier of t_1, t_2, say, ρ. Then, $t_1\sigma_{t_2}^{t_1}\rho = t_2\rho$ because $\sigma_{t_2}^{t_1}$ maps t_1 to t_2, and $t_2\rho = t_1\rho$ because ρ is a concrete unifier. Thus, $t_1\sigma_{t_2}^{t_1}\rho = t_1\rho$. Moreover,

---

[3] even though the standard notion of most general unifier in algebraic specifications and rewriting is a different one.

for all $x \in var(t_2)$, $x\sigma_{t_2}^{t_1}\rho = x\rho$ since $\sigma_{t_2}^{t_1}$ is the identity on $var(t_2)$. Thus, for all $y \in var(t_1) \uplus var(t_2)(= \{t_1\} \uplus var(t_2))$, $y\sigma_{t_2}^{t_1}\rho = y\rho$, which proves the fact that $\sigma_{t_2}^{t_1}$ is a most general unifier (by taking $\eta = \rho$ in Definition 4.1 of unifiers). The fact that the codomain of $\sigma_{t_2}^{t_1}$ is $T_\Sigma(var(t_1) \uplus var(t_2))$ results from its construction.

In the inductive step, $t_1 = f(s_1, \ldots, s_n)$ with $f \in \Sigma \setminus \Sigma^{Data}$ [4] $n \geq 0$, and $s_1, \ldots, s_n \in T_\Sigma(Var)$. For t_2 there are two subcases:

- t_2 is a variable. Then, let $\sigma_{t_2}^{t_1} \triangleq (t_2 \mapsto t_1) \uplus Id|_{var(t_1)}$, i.e., $\sigma_{t_2}^{t_1}$ maps t_2 to t_1, and is the identity on $var(t_1)$. We prove that $\sigma_{t_2}^{t_1}$ is a most general unifier with codomain $T_\Sigma(var(t_1) \uplus var(t_2))$ like in the base case.

- $t_2 = g(u_1, \ldots, u_m)$ with $g \in \Sigma$, $m \geq 0$, and $u_1, \ldots, u_m \in T_\Sigma(Var)$. Let ρ be a concrete unifier of t_1, t_2, thus, $(f\rho)(s_1\rho \ldots s_n\rho) =_{\mathcal{T}} (g\rho)(u_1\rho \ldots u_m\rho)$, where we emphasize by subscripting the equality symbol with $\mathcal{T}$ that the equality is that of the model $\mathcal{T}$. Since $\mathcal{T}$ interprets non-data terms as ground terms over the modified signature (1), we have $f\rho = f$, which implies $f = g$, $g\rho = g$, $m = n$, and $s_i\rho = u_i\rho$ for $i = 1, \ldots, n$. Since t_1 and t_2 are linear and all their data subterms are variables, the subterms s_i and u_i also have these properties. Using the induction hypothesis we build most-general-unifiers $\sigma_{u_i}^{s_i}$ of s_i and u_i, which have codomains $T_\Sigma(var(s_i) \uplus var(u_i))$, for $i = 1, \ldots, n$. Let then $\sigma_{t_2}^{t_1} \triangleq \biguplus_{i=1}^{n} \sigma_{u_i}^{s_i}$. First, $\sigma_{t_2}^{t_1}$ is a substitution of $var(t_1) \uplus var(t_2)$ into $T_\Sigma(var(t_1) \uplus var(t_2))$ since $var(t_1) = \biguplus_{i=1}^{n} var(s_i)$ and $var(t_2) = \biguplus_{i=1}^{n} var(u_i)$. Note that these equalities hold thanks to the linearity of t_1, t_2.
Second, $\sigma_{t_2}^{t_1}$ is a unifier of t_1, t_2 since all $\sigma_{u_i}^{s_i}$ are so. Third, we prove that $\sigma_{t_2}^{t_1}$ is a most general unifier of t_1, t_2. Consider any concrete unifier ρ of t_1 and t_2, thus, $s_i\rho = u_i\rho$ for $i = 1, \ldots, n$. From the fact that all the $\sigma_{u_i}^{s_i}$ are most-general-unifiers of s_i and u_i for $i = 1, \ldots, n$, we obtain the existence of valuations η_i such that $\sigma_{u_i}^{s_i}\eta_i = \rho|_{(var(s_i)\uplus var(u_i))}$, for $i = 1, \ldots, n$. Then, $\eta \triangleq \biguplus_{i=1}^{n} \eta_i$, which is also well-defined thanks to the linearity of t_1 and t_2, has the property that $\sigma_{t_2}^{t_1}\eta = \rho$, which proves that $\sigma_{t_2}^{t_1}$ is a most general unifier of t_1 and t_2 and concludes the proof.

□

Example 4.2. The terms $t_1 = f(x, g(y))$ and $t_2 = f(t, g(z))$ introduced in Example 4.1 satisfy the constraints of Lemma 4.1: they are linear and concretely unifiable, and all their Data subterms are variables (here, the variables in question are x, y, z, t, of the Data sort Int; note that the subterms $g(y)$ and $g(z)$ are not of a Data sort since we assumed the codomain sort of g to be non-Data). Their most-general (symbolic) unifier $\sigma_{t_2}^{t_1}$, built in the proof of Lemma 4.1, coincides with the substitution $x \mapsto t, y \mapsto z$ extended to the identity for all the other variables in Var. What Lemma 4.1 says is that any concrete unifier is an *instance* of the most-general unifier. For concrete unifiers ρ that map all the variables x, y, z, t to 1, noted in Example 4.1, the valuation η in Lemma 4.1, which instantiates the most-general symbolic unifier, can be taken to be ρ.

5. Symbolic Execution

In this section we present a symbolic execution approach for languages defined using the language-definition framework presented in Section 3. We prove that the transition system generated by symbolic execution forward-simulates the one generated by concrete execution, and that the transition system generated by concrete execution backward-simulates the one generated by symbolic execution (restricted to satisfiable patterns). This is used later for proving correctness results for our program-equivalance deduction system.

Symbolic execution consists of applying the semantical rules over patterns using most general unifiers. This generalises the symbolic execution approach proposed in [3], where unification was encoded using matching with modified rules, and which did not allow for symbolic statements. Symbolic execution generates a *symbolic transition system* whose states are patterns, and whose transition relation is obtained by applying rewrite rules with the most-general unifiers whose construction is given by Lemma 4.1.

Definition 5.1 (Symbolic transition relation). $\varphi \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \varphi'$ iff $\varphi \triangleq \pi \wedge \phi$, there is a rule $(l \wedge b \Rightarrow r) \in \mathcal{S}$ with $var(l) \cap var(\pi) = \emptyset$ and such that l, π are concretely unifiable, and $\varphi' = r\sigma_\pi^l \wedge (\phi \wedge b)\sigma_\pi^l$, where σ_π^l is unique, most general symbolic unifier of l, π constructed as in the proof of Lemma 4.1.

---

[4] $f \in \Sigma \setminus \Sigma^{Data}$ because the contrary would mean that t_1 has a *Data* sort, in contradiction with the lemma's hypotheses.

Note that, in order to apply Lemma 4.1 that states the existence of the most-general unifier, the terms l, π also have to be linear and all their subterms of sort Data should be variables. This is not a restriction, however, since the patterns $\varphi \triangleq \pi \wedge \phi$ and $l \wedge b$ can always be modified into equivalent patterns $\pi' \wedge \phi'$ and $l' \wedge b'$ such that the transformed terms l', π' satisfy the conditions of Lemma 4.1, as noted in Remark 3.1.

Example 5.1. Let $\varphi \triangleq \langle\langle X \,/\, Y \curvearrowright \cdot\rangle_{\mathsf{k}} \langle Y' \mapsto A'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \wedge Y' =_{Id} Y \wedge_{Bool} A' =_{Int} A +_{Int} 1 \wedge_{Bool} A \neq_{Int} -1$. φ is linear and all its subterms of sort Data are variables. Consider the rule for division from the semantics of IMP in Figure 2, which we write in full form, which means replacing the ellipses by variables: $\langle\langle I_1 \,/\, I_2 \curvearrowright C\rangle_{\mathsf{k}} \langle E\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \wedge I_2 \neq 0 \Rightarrow \langle\langle I_1/_{Int} I_2 \curvearrowright C\rangle_{\mathsf{k}} \langle E\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}}$. The left hand-side of the rule is also linear and all its subterms of sort Data are variables. The rule generates a symbolic transition from the pattern φ to $\varphi' \triangleq \langle\langle X \,/\, Y \curvearrowright \cdot\rangle_{\mathsf{k}} \langle Y' \mapsto A'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \wedge Y' =_{Id} Y \wedge_{Bool} A' =_{Int} A +_{Int} 1 \wedge_{Bool} A \neq_{Int} -1 \wedge_{Bool} Y \neq_{Int} 0$.

Definition 5.2. The *derivative* of a pattern is the set of patterns that can be obtained by one symbolic execution step: $\Delta_{\mathcal{S}}(\varphi) \triangleq \{\varphi' \mid \varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'\}$. A pattern φ is *derivable* for $\mathcal{S}$ if $\Delta_{\mathcal{S}}(\varphi)$ is a nonempty set.

Example 5.2. The pattern φ from Example 5.1 is derivable for the semantics of IMP, and its derivative is the singleton $\{\varphi'\}$, obtained by symbolically applying the rule for division in the IMP semantics.

In the rest of the paper, for patterns $\varphi \triangleq \pi \wedge \phi$ we let $var(\varphi) \triangleq var(\pi, \phi)$, and for rules $\alpha \triangleq l \wedge r \Rightarrow b$ we let $var(\alpha) \triangleq var(l, b, r)$. Moreover, for symbolic transitions $\varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'$ we assume without restriction on generality that $var(\varphi) \cap var(\alpha) = \emptyset$, which can always be obtained by variable renaming. We also omit to write the subscript $\mathcal{S}$ in the derivatives notation whenever it is understood from the context.

Lemma 5.1. If $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in [\![\varphi]\!]$ then there exists φ' such that $\gamma' \in [\![\varphi']\!]$ and $\varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'$.

Proof. Let $\varphi \triangleq \pi \wedge \phi$. From $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ we obtain the rule $\alpha \triangleq l \wedge r \Rightarrow b$ and the valuation $\rho : Var \to \mathcal{T}$ such that $\gamma = l\rho$, $b\rho = true$, and $\gamma' = r\rho$. From $\gamma \in [\![\varphi]\!]$ we obtain the valuation $\mu : Var \to \mathcal{T}$ such that $\gamma = \pi\mu$ and $\phi\mu = true$. Thus, l and π are concretely unifiable (by their concrete unifier $\rho|_{var(l)} \uplus \mu|_{var(\pi)}$). Using Lemma 4.1 we obtain their unique most-general symbolic unifier σ^l_π, whose codomain is $T_\Sigma(var(l) \uplus var(\pi))$. Let then $\eta : var(l) \uplus var(\pi) \to \mathcal{T}$ be the valuation such that $\sigma^l_\pi \eta = \rho|_{var(l)} \uplus \mu|_{var(\pi)}$. We extend σ^l_π to $var(\varphi, \alpha)$ by letting it be the identity on $var(\varphi, \alpha) \setminus var(l, \pi)$, and extend η to $var(\varphi, \alpha)$ such that $\eta|_{var(b,r) \setminus var(l)} = \rho|_{var(b,r) \setminus var(l)}$ and $\eta|_{var(\phi) \setminus var(\pi)} = \mu|_{var(\phi) \setminus var(\pi)}$. With these extensions we have $x(\sigma^l_\pi \eta) = x(\rho \uplus \mu)$ for all $x \in var(\varphi, \alpha)$.

Let $\varphi' \triangleq r\sigma^l_\pi \wedge (\phi \wedge b)\sigma^l_\pi$: we have the transition $\varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'$ by definition of the symbolic transition system. There remains to prove $\gamma' \in [\![\varphi']\!]$.

- on the one hand, $(r\sigma^l_\pi)\eta = r(\sigma^l_\pi \eta) = r(\rho \uplus \mu) = r\rho = \gamma'$; thus, $(\gamma', \eta) \models r\sigma^l_\pi$;
- on the other hand, $((\phi \wedge b)\sigma^l_\pi)\eta = \phi(\sigma^l_\pi \eta) \wedge b(\sigma^l_\pi \eta) = \phi(\rho \uplus \mu) \wedge b(\rho \uplus \mu) = \phi\mu \wedge b\rho = true$ since $\phi\mu = b\rho = true$; thus; $\eta \models ((\phi \wedge b)\sigma^l_\pi)$.

The two above items imply $(\gamma', \eta) \models r\sigma^l_\pi \wedge (\phi \wedge b)\sigma^l_\pi$, i.e., $(\gamma', \eta) \models \varphi'$, which concludes the proof. $\square$

Corollary 5.1. For every concrete execution $\gamma_0 \Rightarrow_{\mathcal{S}} \gamma_1 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n \Rightarrow_{\mathcal{S}} \cdots$ and every pattern φ_0 such that $\gamma_0 \in [\![\varphi_0]\!]$, there is a symbolic execution $\varphi_0 \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi_1 \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \cdots \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi_n \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \cdots$ such that $\gamma_i \in [\![\varphi_i]\!]$ for $i = 0, 1, \ldots$.

Lemma 5.2. If $\gamma' \in [\![\varphi']\!]$ and $\varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'$ then there exists $\gamma \in \mathcal{T}_{Cfg}$ such that $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in [\![\varphi]\!]$.

Proof. From $\varphi \Rightarrow^{\mathfrak{s}}_{\mathcal{S}} \varphi'$ with $\varphi \triangleq \pi \wedge \phi$ and $\alpha \triangleq l \wedge r \Rightarrow b$ we obtain $\varphi' = r\sigma^l_\pi \wedge (\phi \wedge b)\sigma^l_\pi$. From $\gamma' \in [\![\varphi']\!]$ we obtain $\eta : Var \to \mathcal{T}$ such that $\gamma' = (r\sigma^l_\pi)\eta$ and $((\phi \wedge b)\sigma^l_\pi)\eta = true$. We extend σ^l_π to $var(\varphi, \alpha)$ by letting it be the identity on $var(\varphi, \alpha) \setminus var(l, \pi)$. Let $\rho : Var \to \mathcal{T}$ be defined by $x\rho = x(\sigma^l_\pi \eta)$ for all $x \in var(\varphi, l)$, and $x\rho = x\eta$ for all $x \in Var \setminus var(\varphi, l)$, and let $\gamma \triangleq l\rho$. From $\gamma' = (r\sigma^l_\pi)\eta$ and the definition of ρ we obtain $\gamma' = r\rho$. From $((\phi \wedge b)\sigma^l_\pi)\eta = true$ we obtain $b(\sigma^l_\pi \eta) = true$, i.e., $b\rho = true$, which together with $\gamma \triangleq l\rho$ and $\gamma' = r\rho$ gives $\gamma \Rightarrow_{\mathcal{S}} \gamma'$. There remains to prove $\gamma \in [\![\varphi]\!]$.

- From $\gamma = l\rho$ using the definition of ρ we get $\gamma = l\rho = l(\sigma^l_\pi \eta) = (l\sigma^l_\pi)\eta = (\pi\sigma^l_\pi)\eta = \pi(\sigma^l_\pi \eta) = \pi\rho$;
- From $((\phi \wedge b)\sigma^l_\pi)\eta = true$ we obtain $(\phi\sigma^l_\pi)\eta = \phi(\sigma^l_\pi \eta) = \phi\rho = true$.

Since $\varphi \triangleq \pi \wedge \phi$, the last two items imply $(\gamma, \rho) \models \varphi$, which completes the proof. $\quad\square$

We call a symbolic execution *feasible* if all its patterns are satisfiable (a pattern φ is satisfiable if there is a configuration γ such that $\gamma \in [\![\varphi]\!]$).

Corollary 5.2. For every feasible symbolic execution $\varphi_0 \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \varphi_1 \cdots \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \varphi_n \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \cdots$ there is a concrete execution $\gamma_0 \Rightarrow_{\mathcal{S}} \gamma_1 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n \Rightarrow_{\mathcal{S}} \cdots$ such that $\gamma_i \in [\![\varphi_i]\!]$ for $i = 0, 1, \ldots$.

6. Linear Temporal Logic

The notion of program equivalence we propose requires that there is a combined execution of the two programs that infinitely repeats configurations in a certain relation. Such executions and properties are formally characterised using Linear Temporal Logic (LTL), whose definition is recapped below.

A *Kripke structure* is a tuple $(S, \Rightarrow, P, \lambda, S^i)$ where S is a set of *states*, $\Rightarrow \subseteq S \times S$ is a (total) *transition relation*, P is a set of *propositions*, $\lambda : S \to 2^P$ is the *labelling function*, and $S^i \subseteq S$ is the set of *initial states*.

An *execution* $e \triangleq s_0, \ldots, s_n, \ldots$ is a sequence of states such that $s_0 \in S^i$ and $s_j \Rightarrow s_{j+1}$ for all $j \in \mathbb{N}$. The *suffix* of an execution e from $l \in \mathbb{N}$, denoted by e^l, is the sequence such that $e^l(j) = s_{l+j}$ for all $j \in \mathbb{N}$.

LTL formulas are generated by the grammar $\psi ::= true \mid p \mid \bigcirc \psi \mid \psi \wedge \psi \mid \neg \psi \mid \psi \, \mathcal{U} \, \psi$ for all $p \in P$. Standard abbreviations are $false \triangleq \neg true$, $\psi_1 \vee \psi_2 \triangleq \neg(\neg \psi_1 \wedge \neg \psi_2)$, $\Diamond \psi \triangleq true \, \mathcal{U} \, \psi$, and $\Box \psi \triangleq \neg \Diamond \neg \psi$.

Given an execution $e \triangleq s_0, \ldots, s_n, \ldots$ of a Kripke structure $(S, \Rightarrow, P, \lambda, S^i)$ and an LTL formula ψ, the satisfaction relation $e \models \psi$ is inductively defined over the structure of ψ as follows:

- $e \models true$;
- $e \models p$ iff $p \in \lambda(s_0)$;
- $e \models \bigcirc p$ iff $e^1 \models p$
- $e \models \psi_1 \wedge \psi_2$ if $e \models \psi$ and $e \models \psi_2$;
- $e \models \neg \psi$ iff it is not the case that $e \models \psi$;
- $e \models \psi_1 \, \mathcal{U} \, \psi_2$ iff there exists $k \geq 0$ such that $e^k \models \psi_2$ and for all $0 \leq j < k$, $e^j \models \psi_1$.

We will be interested in formulas of the form $\Box \Diamond p$. Using the semantics of LTL, an execution $e = s_0, \ldots, s_n, \ldots$ satisfies $\Box \Diamond p$ iff it has an infinite subsequence $s_{i_1}, \ldots s_{i_m}, \ldots$ such that $p \in \lambda(s_{i_j})$ for all $j \in \mathbb{N}$.

7. Defining Program Equivalence

We define in this section our notion of program equivalence and a logic for stating equivalence properties.

Assumption 2. We assume without restriction of generality that the transition system $(\mathcal{T}_{Cfg}, \Rightarrow_{\mathcal{S}})$ has no terminal states (i.e., its transition relation is total). This can always be obtained by adding to $\mathcal{S}$ rules of the form $\pi \Rightarrow \pi$ for all non-derivable patterns π, which just add self-loops to terminal states of $(\mathcal{T}_{Cfg}, \Rightarrow_{\mathcal{S}})$.

Remark 7.1. Strictly speaking, it is not programs that are the subject of equivalence, but full configurations (of which programs are just one component). Indeed, program executions depend on the rest of the configuration (e.g., initial values of the variables, ...). Hence, the equivalence relation is a relation on $\mathcal{T}_{Cfg}$.

We consider a given *observation relation* $O \subseteq \mathcal{T}_{Cfg} \times \mathcal{T}_{Cfg}$, which shall serve as a parameter to our equivalence. Then, we say that two configurations are *observationally equivalent* if they are in the observation relation.

Observational equivalence should be understood as a purely local property of configuration pairs, such as, e.g., a given set of variables have the same values in both configurations. Then, our notion of program equivalence requires that starting from any two observationally equivalent configurations, by executing the programs in the configuration one will eventually encounter observationally equivalent configurations again.

This expressed by the LTL formula $O \wedge \Box \Diamond O$, which captures precisely the informal meaning given above. In order to formalise this observation, it will be convenient to consider, for a given language definition $\mathcal{L}$, the language definition denoted by $\mathcal{L}^2$, whose configurations are pairs of configurations of $\mathcal{L}$ and whose rewrite rules are those of $\mathcal{L}$, lifted at the level of configurations of $\mathcal{L}^2$; that is, each semantical rule $l \wedge b \Rightarrow r$

of $\mathcal{L}$ generates two rules of $\mathcal{L}^2$: $\langle l, X \rangle \wedge b \Rightarrow \langle r, X \rangle$ and $\langle Y, l \rangle \wedge b \Rightarrow \langle Y, r \rangle$ where $\langle \_, \_ \rangle$ is the configuration constructor for $\mathcal{L}^2$ and X, Y are variables of sort Cfg for $\mathcal{L}$ that do not occur in the rest of the rule.

Let $\mathcal{S}_l^2$ denote the set of rules of $\mathcal{L}^2$ of the form $\langle l, X \rangle \wedge b \Rightarrow \langle r, X \rangle$, and $\mathcal{S}_r^2$ denote the set of rules of $\mathcal{L}^2$ of the form $\langle Y, l \rangle \wedge b \Rightarrow \langle Y, r \rangle$. We denote by $\mathcal{S}^2$ the whole set of rules of $\mathcal{L}^2$, i.e., $\mathcal{S}^2 = \mathcal{S}_l^2 \uplus \mathcal{S}_r^2$. We transform the transition system of $\mathcal{L}^2$ into a Kripke structure by regarding the observation relation O as a proposition and by labelling the states $\langle \gamma_1, \gamma_2 \rangle$ of the transition system with O iff $(\gamma_1, \gamma_2) \in O$.

By $\mathcal{K}_{\langle \gamma_1, \gamma_2 \rangle}$ we denote the Kripke structure thus constructed, endowed with the single initial state $\langle \gamma_1, \gamma_2 \rangle$.

Definition 7.1. Two configurations γ_1, γ_2 are equivalent, written $\gamma_1 \sim \gamma_2$, if there exists an execution e of the Kripke structure $\mathcal{K}_{\langle \gamma_1, \gamma_2 \rangle}$ such that $e \models O \wedge \Box \Diamond O$.

Example 7.1. The configurations: $\gamma_1 \triangleq \langle \langle \mathtt{x = 2} \rangle_\mathsf{k} \langle \mathtt{x \mapsto 0} \rangle_\mathsf{env} \rangle_\mathsf{cfg}$ and $\gamma_1' \triangleq \langle \langle \mathtt{y = 1; \ y = y+1} \rangle_\mathsf{k} \langle \mathtt{y \mapsto 0} \rangle_\mathsf{env} \rangle_\mathsf{cfg}$ are equivalent when O is defined by requiring that $\mathtt{x = y}$. Indeed, in IMP^2, starting from $\langle \gamma_1, \gamma_2 \rangle$ there is an execution reaching the self-looping state $\langle \langle \langle \rangle_\mathsf{k} \langle \mathtt{x \mapsto 2} \rangle_\mathsf{env} \rangle_\mathsf{cfg}, \langle \langle \rangle_\mathsf{k} \langle \mathtt{y \mapsto 2} \rangle_\mathsf{env} \rangle_\mathsf{cfg} \rangle$, which is in O, hence, the execution satisfies $O \wedge \Box \Diamond O$. Note that not all executions of IMP^2 starting in $\langle \gamma_1, \gamma_2 \rangle$ satisfy $O \wedge \Box \Diamond O$, for example, an execution that applies only rules form $\mathcal{S}_l(\text{IMP}^2)$ followed by self-looping rules violates $O \wedge \Box \Diamond O$. This example also illustrates why one execution (rather than all executions) is required to satisfy $O \wedge \Box \Diamond O$.

Remark 7.2. The relation O gives us quite a lot of expressiveness for capturing various program equivalences, such as the ones classified in [5]. For example, *partial* equivalence is: two programs are equivalent if, whenever presented with the same input, if they both terminate then 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 programs are both empty and inputs are equal then outputs are equal as well. *Mutual* equivalence states that 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 exist executions starting from both configurations of the pair, and such that the programs in both configurations are forever nonempty. Finally, *reactive* equivalence requires that two programs, when presented with the same sequence of inputs, produce the same sequence of outputs. To encode this kind of equivalence we include in O all configuration pairs satisfying the property that if the input cells are equal then the output cells are equal as well.

Remark 7.3. The chosen definition of equivalence is not always adequate for nondeterministic programs. Indeed, assume a nondeterministic instruction | such that, for any statements S_1, S_2, the statement S_1 | S_2 performs either S_1 or S_2. Then, the nondeterministic program $(\mathtt{x:= 0})$ | $(\mathtt{x:= 1})$ is equivalent to the deterministic program $\mathtt{x:= 0}$ according to our definition (with O the relation requiring equality of $\mathtt{x}$ in both). Indeed, both programs have an execution that performs $\mathtt{x:= 0}$ and self-loops there. Clearly, this is inadequate because the nondeterminstic program can also end up with $\mathtt{x}$ being 1, which the deterministic one cannot. For equivalences involving nondeterministic programs, the adequate notion requires that *for all* executions e of one program, *there exists* an execution e' of the other one and an *interleaving* of e, e' satisfying $O \wedge \Box \Diamond O$. This alternation of quantifiers induces difficulties for verifying the equivalence: we do not consider it here.

Remark 7.4. If the nondeterminism is restricted to *confluent* nondeterminism, the inadequacy of our equivalence definition (cf. Remark 7.3) is not an issue any more. Confluent nondeterminism requires the rewrite system defining the semantics of the programming language considered to be confluent in the standard sense. Indeed, in this situation, for the equivalence it is enough to choose one execution of each program and to check the infinite repetition of O-related configurations, since all other executions of each program repeatedly "intersect with" the chosen ones. For example, the confluent nondeterministic program $(\mathtt{x:= 0 \ ; \ y:= 1})$ | $(\mathtt{y:= 1 \ ; \ x:= 0})$ is equivalent to the deterministic $\mathtt{x:= 0}$, when O requires the equality on $\mathtt{x}$.

Remark 7.5. As a final comment on our chosen notion of equivalence, we emphasise that, since it is parametric in a given observation relation O, one may obtain inadequate equivalences if one inadequately chooses O. For example, with the same relation as above (equality on $\mathtt{x}$), the programs $\mathtt{x:= 0 \ ; \ skip}$ and $\mathtt{x:= 0 \ ; \ while} \ true \ \mathtt{do \ skip}$ are equivalent, which may be considered inadequate since the first program terminates while the other one does not. To avoid this issue we exclude from O configuration pairs such that one eventually enters a self loop (like in the first case) and the other one does not (like in the second case).

We present in the rest of the section a logic for program equivalence. We present the logic's syntax and a notion of validity for formulas. A *derivative* operation for formulas is also defined.

Definition 7.2 (Syntax). A *formula* f is a pattern of $\mathcal{L}^2$ according to Def. 3.1 applied to $\mathcal{L}^2$, i.e., an expression of the form $\langle \pi_1, \pi_2 \rangle \wedge \phi$ where $\pi_1, \pi_2 \in T_{\Sigma, Cfg}(Var)$ are basic patterns of $\mathcal{L}$ and $C \in T_{\Sigma, Bool}(Var)$.

Example 7.2. Assume that the signature Σ for the language IMP contains a predicate $isModified : Id \times Stmt \to Bool$, 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 1.1 is

$$
\left\langle \begin{array}{c} \langle\langle \texttt{for } I \texttt{ from } A \texttt{ to } B \texttt{ do\{} S \texttt{\}}\rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle I \texttt{ = } A \texttt{ ; while } I \texttt{ <= } B \texttt{ do\{} S \texttt{ ; } I \texttt{ = } I \texttt{+1}\texttt{\}}\rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle
$$
$$
\wedge \; \neg_{Bool} \; isModified(I, S) \wedge \neg_{Bool} \; isModified(A, S) \wedge \neg_{Bool} \; isModified(B, S)
$$

where M a variable of sort *Map*. The condition says that the loop counter I is not modified in the body S, and the variables occuring in A, B are not modified by S either. The Boolean function $isModified()$ is defined by structural induction on its arguments in the expected manner.

Recall that the set $[\![f]\!]$, introduced by Definition 3.1 applied to $\mathcal{L}^2$ is the set of configurations $\langle \gamma_1, \gamma_2 \rangle$ of $\mathcal{L}^2$ for which there exists a valuation ρ such that $(\langle \gamma_1, \gamma_2 \rangle, \rho) \models f$.

Definition 7.3 (Validity). A formula φ is valid, written $\mathcal{S} \models \varphi$, if for all $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, $\gamma_1 \sim \gamma_2$.

For the deductive system we shall also be needing the following definition: the *derivative* of a formula f is the set of formulas defined by Definition. 5.2, applied to the symbolic transition of the language $\mathcal{L}^2$. We denote it by $\Delta_{\mathcal{S}^2}(f)$. We let $\Delta_l(f) \triangleq \Delta_{\mathcal{S}_l^2}(f)$ be the left-derivative and $\Delta_r(f) \triangleq \Delta_{\mathcal{S}_r^2}(f)$ be the right-derivative of f. We conclude this section by the following lemma that will be used in proofs regarding our deductive system.

Lemma 7.1. For all patterns $\varphi \triangleq \langle \pi_1, \pi_2 \rangle \wedge \phi$ of $\mathcal{L}^2$, all instances $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, and all $h \in \{l, r\}$, there exists $\varphi' \in \Delta_h(\varphi)$ and $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!]$ such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma_1', \gamma_2' \rangle$.

Proof. We prove the lemma for $h = l$, the other case being similar. By construction of the language $\mathcal{L}^2$, the pattern φ has the form $\langle \pi_1, \pi_2 \rangle \wedge \phi$, where π_1, π_2 are basic patterns of $\mathcal{L}$, i.e., terms of sort *Cfg* in $\mathcal{L}$, and ϕ is a term of sort *Bool*. Thus, $\pi_1 \wedge \phi$ is a pattern of $\mathcal{L}$. On the other hand, there exists γ_1' such that $\gamma_1 \Rightarrow_{\mathcal{S}} \gamma_1'$ because the transition system $(\mathcal{T}_{Cfg}, \Rightarrow_{\mathcal{S}})$ has no terminal states (Assumption 2). From $\gamma_1 \Rightarrow_{\mathcal{S}} \gamma_1'$ we obtain $\rho : Var \to \mathcal{T}$ and $\alpha = (l \wedge b \Rightarrow r) \in \mathcal{S}$ such that $\gamma_1 = l\rho$, $\gamma_1' = r\rho$, and $b\rho = true$. By construction of $\mathcal{L}^2$, there is a rule $\langle l, X \rangle \wedge b \Rightarrow \langle r, X \rangle \in \mathcal{S}_l^2$ for X a variable of sort *Cfg* not ocurring in the rest of the rule, i.e., satisfying $X\rho = X$. By extending ρ into a valuation ρ' such that $X\rho' = \gamma_2$ we obtain the concrete transition $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma_1', \gamma_2 \rangle$. Using Lemma 5.1 applied to the transition $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma_1', \gamma_2 \rangle$ and the pattern φ we obtain a pattern φ' such that $\langle \gamma_1', \gamma_2 \rangle \in [\![\varphi']\!]$ and $\varphi \Rightarrow_{\mathcal{S}_l^2}^{\mathfrak{s}} \varphi'$, because the rule that is symbolically applied to obtain φ' from φ is $\langle l, X \rangle \wedge b \Rightarrow \langle r, X \rangle \in \mathcal{S}_l^2$. Thus, $\varphi' \in \Delta_l(\varphi)$, which proves the lemma. $\square$

8. A Circular Proof System

In this section we define a four-rule proof system for proving program equivalence. It is inspired from *circular coinduction* [17], a coinductive proof technique for infinite data structures and coalgebras of expressions [24].

Remember that we have fixed an observation relation O. In order to be able to actually compute with it in our proof system, we assume a set of formulas Ω such that $[\![\Omega]\!] = O$. Let also $\vdash$ be an entailment relation satisfying $\mathcal{S}, F \vdash \varphi$ implies ($\mathcal{S} \models \varphi$ or there exists $f \in F$ such that $[\![\varphi]\!] \subseteq [\![f]\!]$). We intentionally leave open how the relation $\vdash$ it is implemented. It is useful for discarding "obvious" cases in our proof system, such as formulas with identical left and right-hand sides, which the rest of the proof system might take longer to prove or might not be able to prove at all. The set Ω and the relation $\vdash$ are parameters of our proof system:

Definition 8.1 (Circular Proof System).

$$[\text{Axiom}] \ \frac{}{\mathcal{S}, F \vdash^{\circlearrowleft} \emptyset}$$

$$[\text{Reduce}] \ \frac{\mathcal{S}, F \vdash^{\circlearrowleft} G}{\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}} \ \text{if } \mathcal{S}, F \vdash \varphi$$

$$[\text{Circularity}] \ \frac{\mathcal{S}, F \cup \{\varphi\} \vdash^{\circlearrowleft} G \cup \Delta_h(\varphi) \quad h \in \{l, r\}}{\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}} \ \text{if } [\![\varphi]\!] \subseteq [\![\Omega]\!]$$

$$[\text{Derive}] \ \frac{\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \Delta_h(\varphi) \quad h \in \{l, r\}}{\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}} \ \text{if } [\![\varphi]\!] \not\subseteq [\![\Omega]\!] \ \text{and } \Delta_h(\varphi) \neq \{\varphi\}$$

An *execution* of the proof system is any sequence $\bar{\delta}$ of applications of the above rules. For Γ a set of formulas (also called *goals*), a *proof* of $\mathcal{S}, \emptyset \vdash^{\circlearrowleft} \Gamma$ is an execution whose last rule is [Axiom]. Note that, for proving a set of goals, only one sequence (not all sequences) needs to end up with [Axiom].

[Axiom] says that when an empty set of goals is reached, the proof is finished. The [Reduce] rule removes from the current set of goals G any goal that can be discharged by the entailment $\vdash$. The last two rules, [Circularity] and [Derive], both say that a goal φ is replaced by either its left or right derivatives in the set of goals to be proved. However, in [Circularity], the goal φ is added as hypotheses provided that $[\![\varphi]\!] \subseteq [\![\Omega]\!]$, i.e., provided that all its instances are observationally equivalent pairs of configurations. Thus, all the hypotheses f added during executions satisfy $[\![f]\!] \subseteq [\![\Omega]\!] = O$. On the other hand, if $[\![\varphi]\!] \not\subseteq [\![\Omega]\!]$ then the [Derive] rule can be applied, which adds no hypotheses: a goal φ in the current set of goals G is just replaced by its set of left or right-derivatives, provided that theese derivatives are not φ itself (if the derivatives coincide with φ then [Derive] does not change the sequent, thus, applying it would generate a useless infinite execution).

The *soundness* of our proof system is the consequence of the following lemmas. By *sequent encountered by* $\bar{\delta}$ we mean any sequent $\mathcal{S}, F \vdash^{\circlearrowleft} G$ which is obtained by applying a prefix of the sequence $\bar{\delta}$ of rules.

Lemma 8.1. For all sequents $\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}$ encountered by $\bar{\delta}$, for all $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, there exists a sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ encountered by $\bar{\delta}$ with $[\![\varphi']\!] \subseteq O$, and $\langle \gamma'_1, \gamma'_2 \rangle \in [\![\varphi']\!]$, such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow^*_{\mathcal{S}^2} \langle \gamma'_1, \gamma'_2 \rangle$.

Proof. By (strong) induction on the length of the proof $\bar{\delta}$ and case analysis. Depending on first the rule of $\bar{\delta}$ that is applied to the sequent $\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}$:

- if the rule is [Reduce] then there are two subcases:
  - if $\mathcal{S} \vdash \varphi$ then $S \models \varphi$. Thus, $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!] \subseteq O$, and we let $F' = F$, $G' = G$, $\varphi' = \varphi$, and $\langle \gamma'_1, \gamma'_2 \rangle = \langle \gamma_1, \gamma_2 \rangle$;
  - if $[\![\varphi]\!] \subseteq [\![f]\!]$ for some $f \in F$ then $[\![\varphi]\!] \subseteq O$ since all hypotheses f are added (by [Circularity]) such that $[\![f]\!] \subseteq [\![\Omega]\!] = O$, and we can also take $F' = F$, $G' = G$, $\varphi' = \varphi$, and $\langle \gamma'_1, \gamma'_2 \rangle = \langle \gamma_1, \gamma_2 \rangle$;

- if the rule is [Derive] or [Circularity]: using Lemma 7.1, for any $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, there exists $\varphi'' \in \Delta_h(\varphi)$ and $\langle \gamma''_1, \gamma''_2 \rangle \in [\![\varphi'']\!]$ such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma''_1, \gamma''_2 \rangle$. Moreover, the goal φ'' is the current goal of a future rule application in $\bar{\delta}$ and is the origin of a proof $\bar{\delta}''$ strictly shorter than $\bar{\delta}$. Using the induction hypothesis, there exists $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ and $\langle \gamma'_1, \gamma'_2 \rangle \in [\![\varphi']\!] \subseteq O$ such that $\langle \gamma''_1, \gamma''_2 \rangle \Rightarrow^*_{\mathcal{S}^2} \langle \gamma'_1, \gamma'_2 \rangle$. By transitivity $\langle \gamma_1, \gamma_2 \rangle \in \varphi \Rightarrow^*_{\mathcal{S}^2} \langle \gamma'_1, \gamma'_2 \rangle$ holds, which proves this case and concludes the proof.

$\square$

The next lemma says that, for each instance of each hypothesis that has *actually been used for discharging a goal* during a proof, there is a *strict* successor of it satisfying the current goal of some encountered sequent.

Lemma 8.2. Let Φ denote the set of all hypotheses used for discharging a subgoal during the proof $\bar{\delta}$ (i.e., using a [Reduce] rule). Then, for all sequents $\mathcal{S}, F \vdash^{\circlearrowleft} G$ encountered by $\bar{\delta}$, for all $f \in F \cap \Phi$, and for all $\langle \gamma_1, \gamma_2 \rangle \in [\![f]\!]$, $P(f, \langle \gamma_1, \gamma_2 \rangle)$ holds, where $P(f, \langle \gamma_1, \gamma_2 \rangle) \triangleq$ *there exists* $\langle \gamma'_1, \gamma'_2 \rangle$ *and a sequent* $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ *encountered by* $\bar{\delta}$ *with* $[\![\varphi']\!] \subseteq O$, *such that* $\langle \gamma'_1, \gamma'_2 \rangle \in [\![\varphi']\!]$ *and* $\langle \gamma_1, \gamma_2 \rangle \Rightarrow^+_{\mathcal{S}^2} \langle \gamma'_1, \gamma'_2 \rangle$.

Proof. We show that the lemma's statement holds initially and that it is preserved (as an invariant) by all applications of rules in our deductive system (in particular, for the rules in the proof $\bar{\delta}$). The lemma's

statement is trivially true initially, when $F = \emptyset$. For the induction step, we assume that the lemma's statement holds for the current sequent $\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}$ and we show that it holds in the next sequent (if any) in $\bar{\delta}$.

- if the next rule is [Reduce] there are two subcases:

  - $\mathcal{S} \vdash \varphi$. The set $F \cap \Phi$ in the next sequent is the same as in the current one, since this reduction does not use hypotheses in F. With the same instance $\langle \gamma_1', \gamma_2' \rangle$ and sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ given by the inductive hypothesis, we establish that the lemma's statement still holds in the next sequent.

  - $[\![\varphi]\!] \subseteq [\![f_0]\!]$ for some $f_0 \in F$. In this case the set $F \cap \Phi$ in the next sequent is (possibly) larger than in the current one, since this may be the first time the hypothesis f_0 is used to discharge a goal (here, φ). (If $F \cap \Phi$ is the same in the next sequent as in the current one, the inductive hypothesis trivially proves, like in the previous case $\mathcal{S} \vdash \varphi$, that our lemma's statement still holds in the next sequent.) Thus, there remains to consider the case where the current rule's application is the first-time use of the hypothesis f to discharge a goal (here, φ), in which case we have to prove $P(f, \langle \gamma_1, \gamma_2 \rangle)$ for all $f \in F \cup \{f_0\}$ and $\langle \gamma_1, \gamma_2 \rangle \in [\![f]\!]$. Now, $P(f, \langle \gamma_1, \gamma_2 \rangle)$ for $f \in F$ and $\langle \gamma_1, \gamma_2 \rangle \in [\![f]\!]$ holds using the inductive hypothesis (this is proved as in the case $\mathcal{S} \vdash \varphi$). There remains to prove $P(f_0, \langle \gamma_1, \gamma_2 \rangle)$ for all $\langle \gamma_1, \gamma_2 \rangle \in [\![f_0]\!]$. For this, we note that f_0 has been added to F at an earlier proof step by [Circularity], and f_0 was replaced in the following sequent's goals by its derivatives $\Delta_h(f_0)$ for some $h \in \{l, r\}$. Using Lemma 7.1, we obtain a goal $f_0'' \in \Delta_h(f_0)$ and $\langle \gamma_1'', \gamma_2'' \rangle \in [\![f_0']\!]$ such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma_1'', \gamma_2'' \rangle$. Using Lemma 8.1 we obtain the instance $\langle \gamma_1', \gamma_2' \rangle$ and the sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ such that $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$ and $\langle \gamma_1'', \gamma_2'' \rangle \Rightarrow_{\mathcal{S}^2}^* \langle \gamma_1', \gamma_2' \rangle$. By transitivity, $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2}^+ \langle \gamma_1', \gamma_2' \rangle$, which proves that $P(f_0, \langle \gamma_1, \gamma_2 \rangle)$ holds for all $\langle \gamma_1, \gamma_2 \rangle \in [\![f_0]\!]$: the lemma's statement still holds in the next sequent.

- if the next rule is [Circularity] or [Derive]: in this case $F \cap \Phi$ in the next sequent is the same as in the current one, since this rule does not eliminate goals using circular hypotheses (even though, in the case of [Circularity] the current set of hypotheses grows). Like in the case $\mathcal{S} \vdash \varphi$ we establish that the lemma's statement still holds in the next sequent, which concludes this case and completes the proof.

  □

The last lemma used for proving our soundness result resembles Lemma 8.1, but it is stronger since it states the existence of *strict* successors in the observation relation. It can be proved thanks to Lemma 8.2.

Lemma 8.3. For all sequents $\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}$ encountered by $\bar{\delta}$:

- either $\mathcal{S} \vdash \varphi$;
- or for all $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, there exists a sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ encountered by $\bar{\delta}$, and $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$, such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2}^+ \langle \gamma_1', \gamma_2' \rangle$.

Proof. We proceed by (strong) induction on the length of the proof $\bar{\delta}$ and case analysis. Depending on the first rule of $\bar{\delta}$ that is applied to the sequent $\mathcal{S}, F \vdash^{\circlearrowleft} G \cup \{\varphi\}$:

- if the rule is [Reduce] then there are two subcases:

  - if $\mathcal{S} \vdash \varphi$ then this case is proved;
  - if $[\![\varphi]\!] \subseteq f$ for some $f \in F$: Then, $f \in \Phi$ since f is being used (by the present rule!) to discharge a goal. Thus, $f \in F \cap \Phi$. Using Lemma 8.2 we obtain the sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ and instance $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$ such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2}^+ \langle \gamma_1', \gamma_2' \rangle$, which proves this case;

- if the rule is [Derive] or [Circularity]: using Lemma 7.1, for any $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, there exists $\varphi'' \in \Delta_h(\varphi)$ and $\langle \gamma_1'', \gamma_2'' \rangle \in [\![\varphi'']\!]$ such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2} \langle \gamma_1'', \gamma_2'' \rangle$. Moreover, the goal φ'' is the current goal of a future rule application in $\bar{\delta}$ and is the origin of a proof $\bar{\delta}''$ strictly shorter than $\bar{\delta}$. Using the induction hypothesis, there exists the sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ and instance $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$ such that $\langle \gamma_1'', \gamma_2'' \rangle \Rightarrow_{\mathcal{S}^2}^+ \langle \gamma_1', \gamma_2'. \rangle$. By transitivity $\langle \gamma_1, \gamma_2 \rangle \Rightarrow_{\mathcal{S}^2}^+ \langle \gamma_1', \gamma_2' \rangle$ holds, which concludes the proof.

  □

Theorem 8.1 (soundness). Let Γ be a finite set of equivalence formulas. If $\mathcal{S} \vdash^{\circlearrowleft} \Gamma$ then $\mathcal{S} \models \Gamma$.

Proof. Pick any $\varphi \in \Gamma$ (if $\Gamma = \emptyset$ the theorem is trivially true). Applying Lemma 8.3 generates two cases:

1. either $\mathcal{S} \vdash \varphi$, which directly implies $\mathcal{S} \models \varphi$;

2. or, for all $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, there exists a sequent $\mathcal{S}, F' \vdash^{\circlearrowleft} G' \cup \{\varphi'\}$ encountered by $\overline{\delta}$, and $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$, such that $\langle \gamma_1, \gamma_2 \rangle \Rightarrow^+_{\mathcal{S}^2} \langle \gamma_1', \gamma_2' \rangle$. We apply Lemma 8.3 to the latter sequent, which generates two cases:

   (a) $\mathcal{S} \vdash \varphi'$, which implies $\mathcal{S} \models \varphi'$. Then from $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!]$, there is an execution satisfying $O \wedge \square \lozenge O$, and by adding to it the prefix $\langle \gamma_1, \gamma_2 \rangle \Rightarrow^+_{\mathcal{S}^2} \langle \gamma_1', \gamma_2' \rangle$, the resulting execution also satisfies $O \wedge \square \lozenge O$. Thus, from the (arbitrary) $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$ there is an execution satisfying $O \wedge \square \lozenge O$, meaning $\mathcal{S} \models \varphi$;

   (b) or for all $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!]$, there exists a sequent $\mathcal{S}, F'' \vdash^{\circlearrowleft} G'' \cup \{\varphi''\}$ encountered by $\overline{\delta}$, and $\langle \gamma_1'', \gamma_2'' \rangle \in [\![\varphi'']\!] \subseteq O$, such that $\langle \gamma_1', \gamma_2' \rangle \Rightarrow^+_{\mathcal{S}^2} \langle \gamma_1'', \gamma_2'' \rangle$. Applying Lemma 8.3 to the latter sequent generates two cases... It is not hard to see that in the first case we will be able to prove $\mathcal{S} \models \varphi$ like in item 2(a) above, and in the second one, another application of Lemma 8.3 will generate yet two more cases...

This repetitive process may never terminate for a goal φ, but it proves $\mathcal{S} \models \varphi$ in one of two possible ways:

- the first one assumes that, after finitely many applications of Lemma 8.3, a subgoal $\varphi^{(n)}$ satisfying $\mathcal{S} \vdash \varphi^{(n)}$ is found. Thanks to Lemma 8.3, from any $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$ there is a finite execution e than leads into some instance $\langle \gamma_1^{(n)}, \gamma_2^{(n)} \rangle \in [\![\varphi^{(n)}]\!]$. And since $\mathcal{S} \models \varphi^{(n)}$, starting from the instance $\langle \gamma_1^{(n)}, \gamma_2^{(n)} \rangle \in [\![\varphi^{(n)}]\!]$, an infinite execution e' satisfying $O \wedge \square \lozenge O$ exists. The concatenation ee' also satisfies $O \wedge \square \lozenge O$. Thus, from the arbitrarily chosen $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$ en execution satisfying $O \wedge \square \lozenge O$ exists, meaning $\mathcal{S} \models \varphi$ holds.
- the second one assumes the contrary: there is no finite number of applications of Lemma 8.3 after which a subgoal $\varphi^{(n)}$ satisfying $\mathcal{S} \vdash \varphi^{(n)}$ is found. In this case, the infinitely many applications of Lemma 8.3 build an infinite execution $\langle \gamma_1, \gamma_2 \rangle \Rightarrow^+_{\mathcal{S}^2} \langle \gamma_1', \gamma_2' \rangle \Rightarrow^+_{\mathcal{S}^2} \cdots \Rightarrow^+_{\mathcal{S}^2} \langle \gamma_1^{(n)}, \gamma_2^{(n)} \rangle \Rightarrow^+_{\mathcal{S}^2} \cdots$, starting from any arbitrarily chosen $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!]$, such that O is met infinitely many times, as $\langle \gamma_1, \gamma_2 \rangle \in [\![\varphi]\!] \subseteq O$, $\langle \gamma_1', \gamma_2' \rangle \in [\![\varphi']\!] \subseteq O$, ..., $\langle \gamma_1^{(n)}, \gamma_2^{(n)} \rangle \in [\![\varphi^{(n)}]\!] \subseteq O$, ..., which implies that our execution satisfies $O \wedge \square \lozenge O$; thus, $\mathcal{S} \models \varphi$ holds.

In both cases, this process leads to establishing $\mathcal{S} \models \varphi$, and $\varphi \in \Gamma$ was chosen arbitrarily, thus, $\mathcal{S} \models \Gamma$ holds. $\square$

Remark 8.1. For soundness it is not essential that the [Circularity] φ actually adds the current goal φ to the current set of circular hypotheses F. What does matter is that, whenever φ is added to F, then $[\![\varphi]\!] \subseteq [\![\Omega]\!]$. We use this observation in our implementation of the proof system to reduce the number of stored hypotheses.

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

Given a a semantics $\mathcal{S}$ and set of goals Γ, the proof system $\vdash^{\circlearrowleft}$ *terminates successfully* when it returns a proof. The proof system *terminates unsuccessfully* when its has a finite, maximal execution that is not a proof - we call such an execution a *disproof*. This happens when the proof system is "stuck": in the current sequent $\mathcal{S}, F \vdash^{\circlearrowleft} G$, no rule of the system can be applied because the side-conditions of the rules are not satisfied. Then, by definition, the proof system terminates on Γ if it terminates sucessfully or unsuccessfully. Weak completeness then says that if a set of goals Γ is valid, all the goals in the set are satisfiable, and the proof system terminates on Γ, then it terminates successfully.

For this we need the following adaptation to the notion of derivative: $\Delta(\varphi) = \{\varphi' \mid \varphi \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \varphi' \wedge [\![\varphi']\!] \neq \emptyset\}$, which means that only the satisfiable patterns are kept when computing derivatives. We also need:

Assumption 3. For all patterns φ of $\mathcal{L}$, if $\Delta_{\mathcal{S}}(\varphi) = \{\varphi\}$ then there is $\pi \Rightarrow \pi \in \mathcal{S}$ such that $\Delta_{\mathcal{S}}(\varphi) = \Delta_{\{\pi \Rightarrow \pi\}}(\varphi)$, and for all configurations γ, γ' of $\mathcal{L}$, if $\gamma \Rightarrow_{\mathcal{S}} \gamma$ and $\gamma' \not\Rightarrow_{\mathcal{S}} \gamma'$ then $\langle \gamma, \gamma' \rangle \notin O$.

Both assumptions regard the language $\mathcal{L}$ of interest. The first says that, whenever a the derivative of pattern is the pattern itself, then the only rule that contributes to this derivative is of the form $\pi \Rightarrow \pi$. Remember (Assumption 2) that such rules were included in the semantics $\mathcal{S}$ for technical reasons in order to transform terminal configurations into self-looping ones (ultimately, because we deal with LTL over infinite sequences). Our first assumption thus says that, except for the rules, $\pi \Rightarrow \pi$ that were added to the semantics, all the other rules "change" at least "something" in a pattern; i.e., rules that do not change anything in the semantics of a language are useless. Regarding the second of the above assumptions, it says that self-looping configurations and non self-looping ones cannot be observationally equivalent. As observed before, the self-looping configurations are (formerly) terminal configurations that were transformed into self-looping ones

by including the rules of the form $\pi \Rightarrow \pi$ in the semantics $\mathcal{S}$. Thus, our second assumption actually says that configurations where the code to be executed is finished, and configurations where there is still code to execute, cannot be observationally equivalent, which is also a reasonable constraint on equivalence.

Theorem 8.2 (weak completeness). If $\mathcal{S} \models \Gamma$ and for all $\varphi \in \Gamma$ it holds that $[\![\varphi]\!] \neq \emptyset$, and the $\vdash^{\circlearrowright}$ proof system terminates on Γ then $\mathcal{S} \vdash^{\circlearrowright} \Gamma$.

Proof. By contradiction: assume the hypotheses hold but not the conclusion, i.e., $\mathcal{S} \nvdash^{\circlearrowright} \Gamma$. Thus, the proof system terminates with a disproof $\overline{\delta}$, i.e., a sequence of rule applications that is not a proof and after which no rule can be applied. Let $\mathcal{S}, F \vdash^{\circlearrowright} G$ be the sequent resulting after $\overline{\delta}$. Thus, $G \neq \emptyset$, and for all $\varphi \in G$, $[\![\varphi]\!] \nsubseteq O$ (otherwise, [Circularity] would be applicable), and $\Delta_h(\varphi) = \{\varphi\}$ for $h \in \{l, r\}$ (otherwise, [Derive] would be applicable). We choose any $\varphi \in G$. Since both [Circularity] and [Derive] rules compute derivatives, there exists $\varphi_0 \in \Gamma$ and a symbolic execution $\varphi_0 \Rightarrow^{\mathfrak{s}}_{\mathcal{S}^2} \cdots \Rightarrow^{\mathfrak{s}}_{\mathcal{S}^2} \varphi_n = \varphi$. The symbolic execution is feasible, since we have assumed that only satisfiable patterns are kept in the derivatives. Moreover, $[\![\varphi_n]\!] \nsubseteq O$, thus, we can choose $\langle \gamma, \gamma' \rangle \in [\![\varphi]\!] \setminus O$. Hence, we can apply Corollary 5.2 and find a concrete execution $\langle \gamma_0, \gamma'_0 \rangle \Rightarrow_{\mathcal{S}^2} \cdots \Rightarrow^{\mathfrak{s}}_{\mathcal{S}^2} \langle \gamma_n, \gamma'_n \rangle = \langle \gamma, \gamma' \rangle$ such that for all $i = 0, n - 1$, $\langle \gamma_i, \gamma'_i \rangle \in [\![\varphi_i]\!]$, and $\langle \gamma_n, \gamma'_n \rangle \in [\![\varphi_n]\!] \setminus O$.

Next, due to the definition of the language $\mathcal{L}^2$, by *projecting* the above concrete execution of $\mathcal{L}^2$ on its left and right components we obtain the two executions $e \triangleq \gamma_0 \Rightarrow^*_{\mathcal{S}} \gamma$ and $e' \triangleq \gamma'_0 \Rightarrow^*_{\mathcal{S}} \gamma'$ of $\mathcal{L}$. Let $\varphi = \langle \pi, \pi' \rangle \wedge \phi$, then, $\gamma \in [\![\pi \wedge \phi]\!]$ and $\gamma' \in [\![\pi' \wedge \phi]\!]$. From $\Delta_l(\varphi) = \Delta_r(\varphi) = \{\varphi\}$ we obtain $\Delta_{\mathcal{S}}(\pi \wedge \phi) = \{\pi \wedge \phi\}$ and $\Delta_{\mathcal{S}}(\pi' \wedge \phi) = \{\pi' \wedge \phi\}$, thus, Using Assumption 3, both these derivatives were obtained by applying rules of the form $\pi \Rightarrow \pi \in \mathcal{S}$. Thus, there are transitions $\gamma \Rightarrow_{\mathcal{S}} \gamma$ and $\gamma' \Rightarrow_{\mathcal{S}} \gamma'$ in $\mathcal{L}$ (and no other transitions starting from γ and γ', otherwise, there would be rules distinct from those of the form $\pi \Rightarrow \pi$ generating those transitions, in contradiction to $\Delta_l(\varphi) = \Delta_r(\varphi) = \{\varphi\}$). The finite executions e, e' can be extended into infinite ones $\overline{e} \triangleq \gamma_0 \Rightarrow^*_{\mathcal{S}} \gamma \Rightarrow_{\mathcal{S}} \gamma \cdots \Rightarrow_{\mathcal{S}} \gamma \cdots$ and $\overline{e'} \triangleq \gamma'_0 \Rightarrow^*_{\mathcal{S}} \gamma' \Rightarrow_{\mathcal{S}} \gamma' \cdots \Rightarrow_{\mathcal{S}} \gamma' \cdots$. Since the semantics of $\mathcal{L}$ is confluent, any execution $\widehat{e}$ of $\mathcal{L}$ starting in γ_0, resp. in γ'_0 end up, like $\overline{e}$, resp. $\overline{e'}$, by self-loops on γ, resp. γ'. Indeed, all executions $\widehat{e}$ of $\mathcal{L}$ starting starting in γ_0 eventually reach γ because of confluence, and from there on only transitions of the form $\gamma \Rightarrow_{\mathcal{S}} \gamma$ exist; and simularly for executions $\widehat{e'}$ of $\mathcal{L}$ starting in γ'_0.

Moreover, *any infinite executions in $\mathcal{L}^2$ starting in $\langle \gamma_0, \gamma'_0 \rangle$ coincide with sequences obtained by interleaving transitions of $\widehat{e}$ and $\widehat{e'}$, for some execution $\widehat{e}$ of $\mathcal{L}$ starting starting in γ_0 and ending with a self-loop on γ, and some execution $\widehat{e'}$ of $\mathcal{L}$ starting in γ'_0 and ending with a self-loop on γ'.* Consider any such interleaving, denoted hereafter by $\widehat{e} \amalg \widehat{e'}$; we next show that it satisfies $\Diamond \square \neg O$. There are two cases:

- in $\widehat{e} \amalg \widehat{e'}$, both $\widehat{e}$ and $\widehat{e'}$ have reached γ, resp. γ'. Thus, $\widehat{e} \amalg \widehat{e'}$ self-loops on $\langle \gamma, \gamma' \rangle \notin O$; thus, $\widehat{e} \amalg \widehat{e'} \models \Diamond \square \neg O$;

- in $\widehat{e} \amalg \widehat{e'}$, only one of the executions, say, $\widehat{e}$, has reached γ. Thus, $\widehat{e} \amalg \widehat{e'}$ self-loops on $\langle \gamma, \gamma'' \rangle$, for some configuration $\gamma'' \neq \gamma$ *that does not have a transition $\gamma'' \Rightarrow_{\mathcal{S}} \gamma''$* (the existence of $\gamma'' \Rightarrow_{\mathcal{S}} \gamma''$ would contradict the confluence of $\mathcal{L}$: there would exist two executions starting in γ'_0, one that ends up by self-looping in γ, the other one that that ends up by self-looping in $\gamma'' \neq \gamma$, which are clearly not confluent). By Assumption 3, we have $\langle \gamma, \gamma'' \rangle \notin O$ and since $\widehat{e} \amalg \widehat{e'}$ self-loops on $\langle \gamma, \gamma'' \rangle$, we have $\widehat{e} \amalg \widehat{e'} \models \Diamond \square \neg O$.

Recapitulating, we have obtained a goal $\varphi_0 \in \Gamma$ and an instance $\langle \gamma_0, \gamma'_0 \rangle \in [\![\varphi_0]\!]$, such that any infinite execution of $\mathcal{L}^2$ starting in $\langle \gamma_0, \gamma'_0 \rangle$ satisfies $\Diamond \square \neg O$. According to Definition 7.1 this means $\gamma_0 \nsim \gamma'_0$, and by Definition 7.3, $\mathcal{S} \not\models \varphi_0$. Hence, $\mathcal{S} \not\models \Gamma$, which is in contradiction to the hypothesis $\mathcal{S} \models \Gamma$ of our theorem. The contradiction was obtained by assuming $\mathcal{S} \nvdash^{\circlearrowright} \Gamma$, hence, $\mathcal{S} \vdash^{\circlearrowright} \Gamma$ holds, which concludes the proof. $\quad\square$

Together, the soundness and weak completeness results 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 as we have stated it. Of course, termination cannot be guaranteed, because the equivalence problem is undecidable. The proof system does terminate on goals in which both programs terminate (because eventually the set of derivatives does not change the goals and no rule can be applied any more) and also for goals in which the programs does not terminate, but behave in a certain "regular" way, as shown in the examples below.

Example 8.1. We start by illustrating the use of the deductive system on the equivalence of STREAM programs since it does not require unification, hence it is a bit easier. The equivalence we want to prove is

that from Example 1.2: `blink` is equivalent to `zip(zero, one)`. This is written as the equivalence formula

$$\left\langle \begin{array}{l} \langle\langle\texttt{blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg}, \\ \langle\langle\texttt{zip(zero, one)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\, Y_1 = Y_2 \tag{2}$$

where $spec_1$ is `blink` $\mapsto \lambda().\,$`0 : 1 : blink` and $spec_2$ is the map

$$\begin{array}{l} \texttt{zero} \mapsto \lambda().\,\texttt{0:zero} \\ \texttt{one} \mapsto \lambda().\,\texttt{1:one} \\ \texttt{zip} \mapsto \lambda(xs, ys).\,\texttt{hd(}\,xs\,\texttt{)} : \texttt{zip(}\,ys\,,\,\texttt{tl(}\,xs\,\texttt{))} \end{array}$$

Note that the contents of the cells specs is not changed during the execution of the program. The observation relation is given by

$$\Omega = \{\langle\langle\langle\langle C_1\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg}, \langle\langle C_2\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg}\rangle\rangle \wedge Y_1 = Y_2\}$$

where C_1 and C_2 are two arbitrary STREAM programs. In words, two configurations are observational equivalent iff the corresponding output cells out have equal contents.

The equivalence formula (2) is the unique goal in G we start with. We first apply [Circularity] for the program `blink` (i.e., in the proof system, the derivative $\Delta_l()$ is applied), which loads the definition of `blink` in the k cell, and adds (2) to the set of circular hypotheses F. We then apply [Derive], which writes in the corresponding output cell the first head element of the stream, and produces the following goal:

$$\left\langle \begin{array}{l} \langle\langle\texttt{1 : blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(zero, one)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{3}$$

Note that the contents of the output cell in the first configuration has changed. Next, by applying Derive several times with the heating/cooling rules that compute the arguments of `zip(zero, one)`, we get

$$\left\langle \begin{array}{l} \langle\langle\texttt{1 : blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(0 : zero, one)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{4}$$

Several other applications of Derive proceed with loading the definition of `zip` in the k cell, applying heating/cooling rules for `hd(zero)`, adding content to the output cell, and computing new arguments of `zip`:

$$\left\langle \begin{array}{l} \langle\langle\texttt{1 : blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(one, zero)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{5}$$

Now we are in a situation similar to (2). The next formula is obtained in the same way (3) is obtained from (2):

$$\left\langle \begin{array}{l} \langle\langle\texttt{blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : 1 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(one, zero)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{6}$$

The next one is the result of applying the same proof rules used to derive (4) from (3):

$$\left\langle \begin{array}{l} \langle\langle\texttt{blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : 1 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(1 : one, zero)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle\texttt{0 : }\,Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{7}$$

while the last formula is obtained using the same proof rules applied to get (3) from (2):

$$\left\langle \begin{array}{l} \langle\langle\texttt{blink}\rangle_\mathsf{k}\langle spec_1\rangle_\mathsf{specs}\langle\texttt{0 : 1 : }\,Y_1\rangle_\mathsf{out}\rangle_\mathsf{cfg} \\ \langle\langle\texttt{zip(zero, one)}\rangle_\mathsf{k}\langle spec_2\rangle_\mathsf{specs}\langle\texttt{0 : 1 : }\,Y_2\rangle_\mathsf{out}\rangle_\mathsf{cfg} \end{array} \right\rangle$$
$$\wedge\;\; Y_1 = Y_2 \tag{8}$$

To conclude the proof, we note that (8) is an instance of (2) by applying the substitution $\{Y_1 \mapsto$ 0 : 1 $Y_1, Y_2 \mapsto$ 0 : 1 $Y_2\}$. Hence, $[\![(6)]\!] \subseteq [\![(2)]\!]$, and the Reduce discharges the (unique) current goal (8), and Axiom concludes the proof. Note that the first proof rule applied for (5) is Circularity and hence the following equivalence is a consequence of the above proof: 1 : blink $\approx$ zip(one, zero).

Example 8.2. We show the application of our proof system for proving the equivalence of `for` and `while` programs formalised as the validity of the following formula, with $A, B : Int$, $S : Stmt$, $I : Id$ and $M :$ *Map*. Considering A, B to be integers instead of expressions is not a restriction, since, if A and B were arithmetical expressions, the strictness attributes for the `for`, assigment, and `<=` would be applied first and would transform A, B into integers anyway. This allows us to simplify the original equivalence formula , given in Example 7.2, into the following one, based on the fact that $isModified(A, S) = isModified(B, S) = false$:

$$\left\langle \begin{array}{l} \langle\langle \texttt{for } I \texttt{ from } A \texttt{ to } B \texttt{ do\{} S \texttt{\}} \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \\ \langle\langle I = A \,; \texttt{while } I \texttt{ <= } B \texttt{ do\{} S \,; I = I+1 \texttt{\}} \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \neg_{Bool}\ isModified(I, S) \tag{9}$$

The observation relation is given by the set $\Omega = \{\langle\langle\langle\langle C_1 \rangle_{\mathsf{k}} \langle M' \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}}, \langle\langle C_2 \rangle_{\mathsf{k}} \langle M'' \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \wedge M' =_{Map} M''\}$. The relation says that two configurations are observationally equivalent iff they have equal environments.

By starting with the goal (9) only, one cannot get a (finite) proof, because the proof rules [Circularity] and [Derive] forever generate new computation tasks in the k cell. In order to avoid that, one starts with a set of goals G consisting of (9) and

$$\left\langle \begin{array}{l} \langle\langle C \curvearrowright (\texttt{for } I \texttt{ from } A \texttt{ to } B \texttt{ do\{} S \texttt{\}}) \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \\ \langle\langle C \curvearrowright (I = I+1 \texttt{;while } I \texttt{ <= } B \texttt{ do\{} S \,; I = I+1 \texttt{\}}) \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \neg_{Bool}\ isModified(I, C) \wedge_{Bool} lookup(M, I) = A \tag{10}$$

where C is a variable of sort *Code*, abstracting the additional computational tasks. Remember that *Code* is a sort that includes all statements and arithmetical and Boolean expressions, that $\cdot$ denotes the empty code, and that code sequencing is denoted by $\curvearrowright$.

In the sequel we show the application of the rules of our proof system to the chosen set of goals G. The first rule applied to (9) is [Circularity], by which (9) is added to the hypotheses H and is replaced by a goal obtained by applying the semantical rule for the `for` statement, which gives:

$$\left\langle \begin{array}{l} \langle\langle I = A \,; \texttt{if } I \texttt{ <= } B \texttt{ then } S \,; \texttt{for } I \texttt{ from } A +_{Int} 1 \texttt{ to } B \texttt{ do } \{S\} \texttt{ else skip\}} \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \\ \langle\langle I = A \,; \texttt{while } I \texttt{ <= } B \texttt{ do\{} S \,; I = I+1 \texttt{\}} \rangle_{\mathsf{k}}, \langle M \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge\ \neg_{Bool}\ isModified(I, S)$$

We now apply the sequence of rules [Circularity], [Derive], [Circularity], without adding new hypotheses[5], which replaces the above goal with the following one, obtained by applying the semantics of assignment to both sides of the formula and then the semantical rule for the `while` statement:

$$\left\langle \begin{array}{l} \langle\langle \texttt{if } I \texttt{ <= } B \texttt{ then } S \,; \texttt{for } I \texttt{ from } A +_{Int} 1 \texttt{ to } B \texttt{ do } \{S\} \texttt{ else skip\}} \rangle_{\mathsf{k}}, \langle update(M, I, A) \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \\ \langle\langle \texttt{if } I \texttt{ <= } B \texttt{ then } S \,; I = I+1 \,; \texttt{while } I \texttt{ <= } B \texttt{ do\{} S \,; I = I+1 \texttt{\}} \texttt{ else skip} \rangle_{\mathsf{k}}, \langle update(M, I, A) \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \neg_{Bool}\ isModified(I, S)$$

Next[6], the heating rules for the `if` statement and the _ `<=` _ operation, followed by the cooling rules, and finally the rules that conclude the evaluation of the `if` statement result in the two following subgoals:

$$\left\langle \begin{array}{l} \langle\langle \texttt{skip} \rangle_{\mathsf{k}}, \langle update(M, I, A) \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \\ \langle\langle \texttt{skip} \rangle_{\mathsf{k}}, \langle update(M, I, A) \rangle_{\mathsf{env}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge\ \neg_{Bool}\ isModified(I, S) \wedge_{Bool} \neg_{Bool} A \leq_{Int} B$$

---

[5] which is sound thanks to Remark 8.1. In the sequel, whenever [Circularity] is applied, by default it does not add new hypotheses.
[6] In the sequel we mention only the semantical rules used in the sequence of rules [Circularity] and [Derive] that is applied to obtain the next goal.

$$\left\langle \begin{array}{l} \langle\langle S \text{ ;for } I \text{ from } A +_{Int} 1 \text{ to } B \text{ do } \{S\}\rangle_{\mathsf{k}}, \langle update(M,I,A)\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle S \text{ ; } I = I+1 \text{ ; while } I <= B \text{ do}\{S \text{ ; } I = I+1\}\rangle_{\mathsf{k}}, \langle update(M,I,A)\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \ \neg_{Bool} \ isModified(I,S) \wedge_{Bool} \ A \leq_{Int} B$$

The first subgoal is trivially valid and is eliminated by the [Reduce] rule using the base entailment $\vdash$. By applying the semantical rule for statement sequencing, which rewrites ; to $\curvearrowright$, for the second one, we get a new goal

$$\left\langle \begin{array}{l} \langle\langle S \curvearrowright (\text{for } I \text{ from } A +_{Int} 1 \text{ to } B \text{ do } \{S\})\rangle_{\mathsf{k}}, \langle update(M,I,A)\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle S \curvearrowright (I = I+1 \text{ ; while } I <= B \text{ do } \{S \text{ ; } I = I+1\})\rangle_{\mathsf{k}}, \langle update(M,I,A)\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \ \neg_{Bool} \ isModified(I,S) \wedge_{Bool} \ A \leq_{Int} B \tag{11}$$

which is eliminated by the the [Reduce] rule since it is an instance of the goal (10) (by using the substitution $C \leftarrow S, M \leftarrow update(M,I,A)$, and by using the equality $lookup(I, update(M,I,A)) = A$).

To conclude the proof we also need to eliminate the goal (10). This elimination amounts to unifying the code C with all possible left-hand sides of rules in the semantics of IMP. We only give a subset of all the cases, since considering all cases may be overlong for the reader's patience (but not so for a computer). We first consider the case where S is unified with statements:

- $C \leftarrow \text{skip}$: by applying the semantical rules for skip (which rewrites it to $\cdot$), then the rule that consumes the empty code $\cdot$, and finally the rule sequence that evaluates I + 1 in the goal's right-hand side, the goal (10) becomes the following one, which is implied by the initial goal (9) and is eliminated by [Reduce]:

$$\left\langle \begin{array}{l} \langle\langle \text{for } I \text{ from } A \text{ to } B \text{ do}\{S\}\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle I = A \text{ ;while } I <= B \text{ do}\{S \text{ ; } I = I+1\}\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \ \neg_{Bool} \ isModified(I,S) \wedge_{Bool} \ lookup(M,I) = A$$

- $C \leftarrow \{S_1 \text{ ; } S_2\}$ for some statements S_1, S_2: the rule rewriting ; to $\curvearrowright$ produces an instance of the goal (10) itself, with the substitution $C \leftarrow S_1 \curvearrowright S_2$, which is then eliminated by [Reduce].
- $C \leftarrow \{S'\}$ for some statement S': the rule for $\{\_\}$ elimination produces an instance of the goal (10) itself, with the substitution $C \leftarrow S'$, which is then eliminated by [Reduce].
- $C \leftarrow \text{if } B' \text{ then } S_1 \text{ else } S_2$, for some Boolean expression B' and statements S_1, S_2: there are two subcases, depending on whether B' has the sort $Bool$, or does not have the sort $Bool$ but has sort $BExp$:

  - if B' has the sort $Bool$ then one can directly apply the rules for if and obtain two subgoals: one is

$$\left\langle \begin{array}{l} \langle\langle S_1 \curvearrowright (\text{for } I \text{ from } A +_{Int} 1 \text{ to } B \text{ do } \{S\})\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle S_1 \curvearrowright (I = I+1 \text{ ; while } I <= B \text{ do } \{S \text{ ; } I = I+1\})\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \ \neg_{Bool} \ isModified(I,S_1) \wedge_{Bool} \neg_{Bool} \ isModified(I,S_2) \wedge_{Bool} \ B' =_{Bool} true$$

  (where we used $isModified(I, \text{if } B' \text{ then } S_1 \text{ else } S_2) = isModified(I,S_1) \vee_{Bool} isModified(I,S_2)$). This is an instance of the goal (10) and is eliminated by [Reduce][7]. The other subgoal is similar, but with S_2 instead of S_1 and $B' =_{Bool} false$ in the condition, which is also an instance of the goal (10).
  - if B' does not have the sort $Bool$ then it has the sort $BExp$. Then, the only rule that our goal can be unified with is the heating rule for if, which generates the following goal:

$$\left\langle \begin{array}{l} \langle\langle (B' \curvearrowright \text{if } \square \text{ then } S_1 \text{ else } S_2) \curvearrowright (\text{for } I \text{ from } A +_{Int} 1 \text{ to } B \text{ do } \{S\})\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \\ \langle\langle (B' \curvearrowright \text{if } \square \text{ then } S_1 \text{ else } S_2) \curvearrowright (I = I+1 \text{ ; while } I <= B \text{ do } \{S \text{ ; } I = I+1\})\rangle_{\mathsf{k}}, \langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \end{array} \right\rangle$$
$$\wedge \ \neg_{Bool} \ isModified(I,S_1) \wedge_{Bool} \neg_{Bool} \ isModified(I,S_2)$$

  which is an instance of (10) with $S \leftarrow (B' \curvearrowright \text{if } \square \text{ then } S_1 \text{ else } S_2)$.

- $C \leftarrow \text{while } B' \text{ do } S'$. The rule for while transforms (10) into an instance of itself under the substitution $C \leftarrow (\text{if } B' \text{ then } S \text{ ; } I + 1 \text{ ; while } B' \text{ do } S' \text{ else skip})$.

---

[7] in the sequel, whenever (10) is transformed into an instance of itself, we omit the sentence "and is eliminated by [Reduce]".

- $C \leftarrow$ `for` I' `from` A' `to` B' `do` S'. The rule for `for` transforms (10) into an instance of itself under the substitution $C \leftarrow (I'$ `=` A' `;` `if` B' `then` S `;` I `+ 1` `;` `for` I' `from` A' `to` B' `do` S' `else` `skip`).
- $C \leftarrow X$ for some identifier X, which amounts to unification with the rule for program-variable lookup. That rule transforms our goal into an instance of itself with $C \leftarrow lookup(M, X, I)$.
- $C \leftarrow X' = A'$ for some identifier X' and arithmetical expression A'. Similar to the case of `if`, there are subcases depending on whether $'$ has sort Int, or does not have sort Int but has sort $AExp$.
  - in the first case the rule for variable assignment transforms (10) into an instance of itself with $C \leftarrow \cdot$ and $M \mapsto update(M, X, I)$;
  - in the second case, the heating rule for variable assignment transforms 10) into an instance of itself with $C \leftarrow (A' \curvearrowright I = \square)$.

There remain to consider the cases where C is code but is not a statement. The goal (10) can be unified with left-hand sides of semantical rules:

- $C \leftarrow C_1 \curvearrowright C_2$: then unification may be performed with both heating and cooling rules.
  - We first illustrate the situation with the cooling rule for the `if` statement, which was explicitly given in Section 2.1; the case for all the other cooling rules is completely similar. In the considered case, $C_1 \leftarrow B$ and $C_2 \leftarrow$ `if` $\square$ `then` S_1 `else` $S_2 \curvearrowright C'$ for some code C', and the cooling rule transforms (10) into an instance of itself with $C \leftarrow$ `if` B `then` S_1 `else` $S_2 \curvearrowright C'$.
  - Regarding unification with heating rules, this may only happen when the left-hand side of the rule is of the form $\langle\langle C_1' \curvearrowright C_2'\rangle_{\mathsf{k}}\langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}}$, and the right-hand side has the form $\langle\langle (C_1'' \curvearrowright C_2'') \curvearrowright C_2'\rangle_{\mathsf{k}}\langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}}$; the application of this rule transforms (10) into an instance of itself with $C \leftarrow C_1'' \curvearrowright C_2''$.
- C is an arithmetical expression or a Boolean expression. Then, again, unification may be performed with both heating and cooling rules, in a completely similar many to what has been shown above.

Thus, in all possible cases, the goal goal (10) is transformed into an instance of itself and is eliminated from the set of goals. Since the other goal (9) has been eliminated earlier, the proof system terminates sucessfully.

9. A Prototype Implementation

$\mathbb{K}$ [19] is a framework for defining the formal operational semantics of programming languages. One component of the framework is a compiler of $\mathbb{K}$ definitions to Maude [25] specifications. Programs of languages defined in $\mathbb{K}$ can thus be executed and analysed using Maude as the underlying rewriting engine. $\mathbb{K}$ also offers some support for symbolic computations, including a connection to the Z3 SMT solver [26]. We have used these components in a prototype tool implementing our deductive system for program equivalence. Here we describe how the proposed proof system is implemented for the IMP and STREAM languages. This description is generic enough and can be seen as a methodology applicable to any language defined in $\mathbb{K}$.

There are (at least) two approaches to implementing the proof system:

1. as an external procedure, which uses the $\mathbb{K}$ tool for computing derivatives of equivalence formulas only. The external procedure is then responsible all the other operations, including the searching for proofs;
2. directly in $\mathbb{K}$, by performing all the operations in the proof system using the available $\mathbb{K}$ tools (for example, the underlying Maude search engine is used in searching for proofs). This requires extending the definition of the language of interest with additional data structures, with semantical rules for storing circular hypotheses, and with rules for the entailment between these hypotheses and goals.

Since our approach is parametric in the language definition, observational relation, and basic entailment, in both cases we need a procedure that builds the definition of $\mathcal{L}^2$ for a given $\mathcal{L}$, and procedures for the basic entailment $(\mathcal{S} \vdash \varphi)$ and subsumption $(\llbracket\varphi\rrbracket \subseteq \llbracket f\rrbracket)$. The basic entailment relation $\mathcal{S}^2 \vdash \varphi$ can be specified by means of a (finite) set of equivalence formulas $\mathcal{E}$ (in the same way that Ω specifies the observation relation O), and taking $\mathcal{S}^2 \vdash \varphi$ iff there is $e \in \mathcal{E}$ such that $\llbracket\varphi\rrbracket \subseteq \llbracket e\rrbracket$. The subsumption relation can be checked using Proposition 3.2 or Proposition 3.3. For IMP the set Ω will typically consist of formulas that say that a given set of program variables have the same values in both configurations, and $\mathcal{E}$ further requires that the two contents of the k cells are the same. For STREAM, Ω says that the two out cells have the same contents.

By Proposition 3.1, the formulas $f \in F$ and $e \in \mathcal{E}$ can always be stored in the form $\pi' \wedge \bigwedge_\sigma \wedge \phi$, which

facilitates the checking of subsumptions based on Proposition 3.2 or Proposition 3.3. The validity of the implication from Proposition 3.2 is checked by calling the Z3 SMT solver. The substitution σ (ocurring in formulas of the form $\pi' \wedge \bigwedge_\sigma \wedge \phi$) is computed by inspecting the contents of the two configurations π and π'.

We chose to implement the proof system for the two languages directly in $\mathbb{K}$ since this is the most straightforward approach and allows us to benefit from tools in the $\mathbb{K}$ framework. However, we had to make some compromises. Since the current Maude backend of $\mathbb{K}$ is a rewriting engine based on matching, we had to axiomatise symbolic statements instead of using unification for them. The main axiom says how a symbolic-statement variable S affects the environment M under a current condion ϕ:

$$\langle\langle S \cdots\rangle_{\mathsf{k}}\langle M\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \wedge \phi \Rightarrow \langle\langle \cdot \cdots\rangle_{\mathsf{k}}\langle followup(S, M, \phi)\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}} \wedge \phi$$

The function $followup$ is axiomatised as well; its axiom says that S has no effect on a variables X that it does not modify: $followup(S, (X \mapsto V\ M), \phi) = X \mapsto V\ followup(S, M, \phi)$ when ϕ implies $\neg isModified(X, S)$.

An equivalence formula $\varphi \triangleq \langle \pi_1, \pi_2\rangle \wedge \phi$ for IMP is written in $\mathbb{K}$ as an IMP configuration

$$\langle\langle p_1\rangle_{\mathsf{k}}\langle M_1\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_1}\langle\langle p_2\rangle_{\mathsf{k}}\langle M_2\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_2}\langle\phi\rangle_{\mathsf{cond}}$$

where the pattern π_i is given by the contents of the cfg_i cell and the condition ϕ is stored into a new cell called cond. The circularities F are stored into a new cell hypos. For each circularity $f \in F$, the subsumption relation $[\![\varphi]\!] \subseteq [\![f]\!]$ is checked by means of two substitutions. The first one is a substitution σ from the contents of the cell k in f to the coresponding one in φ. Let

$$f \triangleq \langle\langle p_1'\rangle_{\mathsf{k}}\langle M_1'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_1}\langle\langle p_2'\rangle_{\mathsf{k}}\langle M_2'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_2}\langle\phi'\rangle_{\mathsf{cond}}$$

such an hypothesis in F and let φ be the current equivalence formula represented as above. For instance, if φ is given by (11) and f by (10), then σ is $A \leftarrow A +_{Int} 1$. The expressions from the codomain of σ are evaluated in the current configuration; in this way, e.g., the program variables are replaced by their current values. Since the cell env includes only fresh variables, we have $f\sigma$ equal to

$$\langle\langle p_1\rangle_{\mathsf{k}}\langle M_1'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_1}\langle\langle p_2\rangle_{\mathsf{k}}\langle M_2'\rangle_{\mathsf{env}}\rangle_{\mathsf{cfg}_2}\langle\phi'\sigma\rangle_{\mathsf{cond}}$$

The second substitution σ' is between the corresponding env cells such that $M_i'\sigma' = M_i$ for $i = 1, 2$. Note that $f\sigma\sigma' = \varphi$ does not hold in general, because usually $\phi'\sigma\sigma' = \phi$ does not. But if $\phi \wedge \bigwedge_{\sigma'}$ implies $\phi'\sigma$ holds, which is checked by calling the SMT solver, then we obtain $[\![\varphi]\!] \subseteq [\![f\sigma]\!]$ by Proposition 3.2 . Since $[\![f\sigma]\!] \subseteq [\![f]\!]$ by Proposition 3.3, it follows that $[\![\varphi]\!] \subseteq [\![f]\!]$.

This method for checking the subsumption relation is not specific to IMP. For any other language definition the substitution σ is defined by structural induction on the language syntax and the substitution σ' is computed by considering the rest of configurations. For instance, for the case of STREAM, only the substitution σ is required because the rest of configurations remain constant during the execution.

The efficiency of the implementation depends on how the [Circularity] and [Derive] rules are applied. In Remark 8.1 we noted that, for soundness, it is not necessary to always add the current goal φ to the hypotheses F when applying [Circularity]. Ideally, only those formulas actually subsequently used by [Reduce] rules should be added. Since there is no way of knowing in advance which circular hypotheses will be used in the future, we apply a heuristic when adding circular hypotheses. This is achieved by using labelled statements: each time two statements with the same label are at the top of the k cells, a set of rules decides which one of the following three cases holds for the current configuration and takes the corresponding action: whether it belongs to the observation relation, or it is a consequence of the circular hypotheses, or it must be added to the circular hypotheses.

The content of the hypos cell can be explored during the proving process to discover new equivalences to be proved, when the initial ones fails. We explain this for the goal $\mathtt{morse} \approx \mathtt{f(morse)}$, where

$$\mathtt{morse} \approx \mathtt{0:1:zip(tl(morse),\ not(tl(morse)));} \qquad \mathtt{not(\ xs\ )} \approx \mathtt{neg(hd(\ xs\ )):not(\ xs\ );}$$
$$\mathtt{f(\ xs\ )} \approx \mathtt{hd(\ xs\ ):neg(hd(\ xs\ )):f(tl(\ xs\ ));} \qquad \mathtt{neg(\ x\ )}\ :=\ 1 \lhd x =_{Int} 0 \rhd 0$$

If we execute the prototype for $\mathtt{morse} \approx \mathtt{f(morse)}$ only, then it forever applies the proof rules [Derive] and [Circularity], similarly to Example 8.2. Finite approximations of the infinite execution given by the prototype can be obtained using the $\mathbb{K}$ stepper. Analysing the contents of the cell storing the circularities (circular hypotheses collected by the proof rule [Circularity] in the hypos cell), we observe that it includes formulas of

the form

$$\left\langle \begin{array}{l} \langle\langle \texttt{1 : zip(tl(morse), not(tl(morse)))}\rangle_{\mathsf{k}} \langle spec_1 \rangle_{\mathsf{specs}} \langle \texttt{0 : } Z \rangle_{\mathsf{out}} \rangle_{\mathsf{cfg}} \\ \langle\langle \texttt{1 : f(1 : zip(tl(morse), not(tl(morse))))}\rangle_{\mathsf{k}} \langle spec_2 \rangle_{\mathsf{specs}} \langle \texttt{0 : } Z \rangle_{\mathsf{out}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$

$$\left\langle \begin{array}{l} \langle\langle \texttt{0 : zip}(s_1)\rangle_{\mathsf{k}} \langle spec_1 \rangle_{\mathsf{specs}} \langle \texttt{0 : 1 : 1 : } Z \rangle_{\mathsf{out}} \rangle_{\mathsf{cfg}} \\ \langle\langle \texttt{0 : f}(s_1)\rangle_{\mathsf{k}} \langle spec_2 \rangle_{\mathsf{specs}} \langle \texttt{0 : 1 : 1 : } Z \rangle_{\mathsf{out}} \rangle_{\mathsf{cfg}} \end{array} \right\rangle$$

. . .

where s_1 is the stream expression `zip(tl(morse), not(tl(morse)))`. From the specifications we have `1 : zip(tl(morse), not(tl(morse)))`=`tl(morse)`, i.e. `f(1 : zip(tl(morse), not(tl(morse))))`=`f(tl(morse)`, and hence we deduce that an abstract form of these formulas is `f(`S`)`=`zip(`S`, not(`S`))`. Running the prototype for the set $\{$`morse` $\approx$ `f(morse)`, `f(`S`)`=`zip(`S`, not(`S`))`$\}$ we get a finite proof.

We also note that the STREAM example shows that the proof system introduced in this paper includes the one defined in [17] whenever behavioural equational specifications can be encoded as programming language definitions. However, the definition for the equivalence we introduced here is more general: the equivalence considered in [17] can be defined using the LTL formula pattern $\Box O$ while the one defined here uses $\Box \Diamond O$.

10. 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 example of equivalent programs in languages from two different paradigms.

The proposed approach is generic: it does not depend on $\mathbb{K}$ and the language being defined in $\mathbb{K}$, but requires a formal semantics of the language of interest as a term-rewriting system. The chosen equivalence relation is also parametric in a certain observation relation and requires that starting from configurations in the observartion relation, configurations in the observation relation will be encountered again. We show the verification approach is applicable for concrete and symbolic programs and for terminating and non-terminating ones.

The chosen notion of equivalence is suitable for deterministic and also for confluent-nondeterminisic languages. It subsumes several equivalence relations from the literature, which can be obtained by adequately setting its parameter (the observation relation). It is based on the formal operational semantics of languages and on symbolic execution of programs based on those semantics. Currently, more and more operational semantics of "real" languages are becoming available (e.g., those published at `http://k-framework.org`), which will make our progam-equivalence approach applicable to an ever growing number of languages.

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 [22], a language-independent logic for programs written in 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.

Longer-term future work directions include the generalisation of our approach to nondeterministic languages, beyond the class of confluent-nontederministic languages that we can currently deal with.

References

[1]   Stefan Ciobaca, Dorel Lucanu, Vlad Rusu, and Grigore Rosu. A Language-Independent Proof System for Mutual Program Equivalence. In *ICFEM'14 - 16th International Conference on Formal Engineering Methods*, LNCS (to appear), Luxembourg, November 2014. Springer.

[2]   Dorel Lucanu and Vlad Rusu. Program Equivalence by Circular Reasoning. In *Integrated Formal Methods*, volume 7940 of *Lecture Notes in Computer Science*, Turku, Finland, June 2013. Springer.

[3]   Andrei Arusoaie, Dorel Lucanu, and Vlad Rusu. A Generic Framework for Symbolic Execution. In Martin Erwig, Richard F. Paige, and Eric Van Wyk, editors, *6th International Conference on Software Language Engineering*, volume 8225 of *Lecture Notes in Computer Science*, pages 281–301, Indianapolis, États-Unis, October 2013. Springer Verlag.

[4]   Sudipta Kundu, Zachary Tatlock, and Sorin Lerner. Proving optimizations correct using parameterized program equivalence. In *Programming Languages Design and Implementation*, pages 327–337, 2009.

[5]   Benny Godlin and Ofer Strichman. Inference rules for proving the equivalence of recursive procedures. *Acta Inf.*, 45(6):403–439, 2008.

[6]   Benny Godlin and Ofer Strichman. Regression verification: proving the equivalence of similar programs. *Software Testing, Verification and Reliability*, 2012. 10.1002/stvr.1472.

[7]   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.

[8]   Xavier Leroy. Formal verification of a realistic compiler. *Comm. ACM*, 52(7):107–115, 2009.

[9]   George C. Necula. Translation validation for an optimizing compiler. In Monica S. Lam, editor, *PLDI*, pages 83–94. ACM, 2000.

[10]   Torben Amtoft, Sruthi Bandhakavi, and Anindya Banerjee. A logic for information flow in object-oriented programs. In *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '06, pages 91–102, New York, NY, USA, 2006. ACM.

[11]   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.

[12]   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.

[13]   Sorin Craciunescu. Proving the equivalence of CLP programs. In *International Conference of Logic Programming, LNCS 2401*, pages 287–301, 2002.

[14]   Wolfgang Ahrendt, Andreas Roth, and Ralf Sasse. Automatic validation of transformation rules for java verification against a rewriting semantics. In *LPAR*, volume 3835 of *Lecture Notes in Computer Science*, pages 412–426. Springer, 2005.

[15]   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.

[16]   Fabio Somenzi and Andreas Kuehlmann. *Electronic Design Automation For Integrated Circuits Handbook*, volume 2, chapter 4: Equivalence Checking. Taylor & Francis, 2006.

[17]   Grigore Roşu and Dorel Lucanu. Circular coinduction – a proof theoretical foundation. In *CALCO 2009*, volume 5728 of *LNCS*, pages 127–144. Springer, 2009.

[18]   Santiago Escobar and José Meseguer. Symbolic model checking of infinite-state systems using narrowing. In *Term Rewriting and Applications, 18th International Conference, RTA 2007*, volume 4533 of *Lecture Notes in Computer Science*, pages 153–168. Springer, 2007.

[19]   G. Roşu and T.-F. Şerbănuţă. An Overview of the K Semantic Framework. *Journal of Logic and Algebraic Programming*, 79(6):397–434, 2010.

[20]   Luke Simon, Ajay Bansal, Ajay Mallya, and Gopal Gupta. Co-logic programming: Extending logic programming with coinduction. In *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 472–483. Springer, 2007.

[21]   Davide Ancona and Elena Zucca. Corecursive featherweight Java. In *Proceedings of the 14th Workshop on Formal Techniques for Java-like Programs*, FTFJP '12, pages 3–10, New York, NY, USA, 2012. ACM.

[22]   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.

[23]   Andrei Arusoaie, Dorel Lucanu, and Vlad Rusu. A Generic Approach to Symbolic Execution. Research Report RR-8189, INRIA, 2012. `http://hal.inria.fr/hal-00766220/`.

[24]   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.

[25]   M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. L. Talcott. *All About Maude, A High-Performance Logical Framework*, volume 4350 of *Lecture Notes in Computer Science*. Springer, 2007.

[26]   Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008)*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.