



**A Language-Independent Proof System
for Mutual Program Equivalence**

**Ștefan Ciobâcă, Dorel Lucanu,
Vlad Rusu and Grigore Roșu**

TR 13-03, December 2013

ISSN 1224-9327



**Universitatea “Alexandru Ioan Cuza” Iași
Facultatea de Informatică**

Str. Berthelot 16, 6600-Iași, Romania
Tel. +40-32-201090, email: bibl@infoiasi.ro

A Language-Independent Proof System for Mutual Program Equivalence

Ștefan Ciobâcă¹, Dorel Lucanu¹, Vlad Rusu², and Grigore Roșu^{1,3}

¹ Universitatea Alexandru Ioan Cuza, Romania

² Inria Lille, France

³ University of Illinois at Urbana-Champaign, USA

Abstract. Two programs or fragments of program are mutually equivalent iff either they both diverge or they end up in similar states. Mutual equivalence is desirable in many contexts, ranging from capturing program equivalence or correctness of program transformations within the same language, to capturing correctness of compilers from one language to another. This paper introduces a language-independent proof system for mutual equivalence. The proof system is parametric in the operational semantics of the two languages and in a state similarity relation.

1 Introduction

Two terminating programs are equivalent if the final states that they reach are similar. Nontermination can be incorporated in equivalence in several ways. Two programs are said to be *mutually equivalent* iff they either both diverge or they both terminate and then the final states that they reach are similar. Mutual equivalence is thus an adequate notion of equivalence for programs written in deterministic sequential languages and is useful, e.g., in compiler verification.

In this paper we formalize the notion of mutual equivalence and propose a logic with a deductive system for stating and proving mutual equivalence of two programs that can belong to two different languages. The deductive system is language-independent, in the sense that it is parametric in the semantics of the two-languages. We prove that the proposed system is sound: when it succeeds it proves the mutual equivalence of the programs given to it as input. The key idea is to use the proof system to build a relation on configurations which is closed under the transition relations given by the corresponding operational semantics. The challenge is how to achieve that generically, where the two languages are given by their formal semantics, without relying on the specifics of any particular language. The representative rule of our proof system is CIRCULARITY, which allows us to incrementally postulate synchronization points in the two programs.

We illustrate the proof system on two programs (Fig. 5 on page 10) that both compute the Collatz sequence in two ways: one is written in an imperative language and the other one in a functional language. We prove with our system that they are mutually equivalent without knowing whether they terminate.

Section 2 introduces matching logic and shows how it can be used to give operational semantics to programming languages, on which the rest of the developments build. Section 3 then shows how to aggregate matching logic semantics, a core operation for stating and proving mutual equivalence. Section 4 presents our proof system for mutual equivalence and applies it to the two programs computing the Collatz sequence. Section 5 discusses related work and concludes.

$\begin{aligned} \text{Exp}_L &::= \text{Var} \mid \text{Int} \mid \text{Exp}_L \text{ op } \text{Exp}_L \\ \text{Stmt} &::= \text{Var} := \text{Exp}_L \\ &\quad \mid \text{skip} \mid \text{Stmt} ; \text{Stmt} \\ &\quad \mid \text{if } \text{Exp}_L \text{ then } \text{Stmt} \text{ else } \text{Stmt} \\ &\quad \mid \text{while } \text{Exp}_L \text{ do } \text{Stmt} \\ \text{Code} &::= \text{Exp}_L \mid \text{Stmt} \\ \text{Cfg}_L &::= \langle \text{Code}, \text{Map}\{\text{Var}, \text{Int}\} \rangle \end{aligned}$	$\begin{aligned} \text{Exp}_R &::= \text{Var} \mid \text{Val} \mid \text{Exp}_R \text{ op } \text{Exp}_R \\ &\quad \mid \text{Exp}_R \text{ Exp}_R \\ &\quad \mid \text{letrec } \text{Var } \text{Var} = \text{Exp}_R \text{ in } \text{Exp}_R \\ &\quad \mid \text{if } \text{Exp}_R \text{ then } \text{Exp}_R \text{ else } \text{Exp}_R \\ &\quad \mid \mu \text{Var} . \text{Exp}_R \\ \text{Val} &::= \text{Int} \mid \lambda \text{Var} . \text{Exp}_R \\ \text{Cfg}_R &::= \langle \text{Exp}_R \rangle \end{aligned}$
--	---

Fig. 1. Syntax of the IMP language on the left and of the FUN language on the right. The symbol `op` ranges over the binary arithmetic operations `+`, `%`, `/`, `==`, `!=`, etc.

2 Matching Logic and Programming Language Semantics

Here we introduce notation used throughout the paper and discuss the recently introduced *matching logic* [16, 15], a language-parametric logic for reasoning about program configurations, and its use in programming language semantics. Matching logic extends FOL with *basic patterns*, which are open terms (i.e., terms with variables) that can be used as basic formulae in the logic. The semantics of a basic pattern as a FOL formula is given by matching.

For example, if program configurations are described as pairs $\langle \text{code}, \text{env} \rangle$ of fragments of code code and state env , the matching logic formula $\langle \text{skip}, (\mathbf{x} \mapsto n, _)\rangle \wedge n > 0$ matches configurations where there is no more code to execute (the `skip` instruction does nothing) and where the program variable \mathbf{x} is mapped to a positive integer. Here the underscore $(_)$ denotes an anonymous variable matching the rest of the program state map. We only consider first-order matching logic in this paper, whose signatures are many-sorted first-order signatures with a distinguished sort for configurations. We let Σ and Π range over multi-sorted algebraic signatures and, respectively, over multi-sorted first-order predicate sets.

Definition 1. A *matching logic (ML) signature* $(\text{Cfg}, \Sigma, \Pi)$ is a first-order many-sorted signature (Σ, Π) with a distinguished sort Cfg for **configurations**.

Example 1. It is well-known that the abstract (context-free) syntax of programming languages can be defined as first-order signatures. Fig. 1 shows the syntax of a simple imperative language IMP and of a simple functional language FUN, that we use as case studies in this paper. We call the respective first-order signatures Σ_L (for IMP—the language on the left-hand side of the figure) and Σ_R (for FUN—the language on the right-hand side of the figure). Each non-terminal in the syntax context-free grammars corresponds to a sort and each production to an operation (or function) symbol. When we refer to the corresponding operation symbol of a production, we use the *mixfix* notation. For example, the operation corresponding to IMP’s assignment is $_ := _ : \text{Var} \times \text{Exp}_L \rightarrow \text{Stmt}$, and that corresponding to FUN’s λ -abstraction is $\lambda _ . _ : \text{Var} \times \text{Exp}_R \rightarrow \text{Val}$. We assume the integer numbers (sort `Int`) and variable symbols (sort `Var`), together with their usual operations and predicates, shared by both IMP and FUN. Sorts Cfg_L and respectively Cfg_R represent the configurations of each programming language, which, for uniformity, we assume constructed with a grouping operation $\langle \dots \rangle$.

$$\begin{aligned}
\langle x, env \rangle &\Rightarrow \langle env(x), env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle i_1 \text{ op } i_2, env \rangle &\Rightarrow \langle i_1 \text{ op}_{\text{Int}} i_2, env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle x := i, env \rangle &\Rightarrow \langle \text{skip}, env[x \mapsto i] \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle \text{skip}; s, env \rangle &\Rightarrow \langle s, env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle \text{if } i \text{ then } s_1 \text{ else } s_2, env \rangle \wedge i \neq 0 &\Rightarrow \langle s_1, env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle \text{if } 0 \text{ then } s_1 \text{ else } s_2, env \rangle &\Rightarrow \langle s_2, env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle \text{while } e \text{ do } s, env \rangle &\Rightarrow \langle \text{if } e \text{ then } s \text{ while } e \text{ do } s \text{ else skip}, env \rangle \in \mathcal{A}_{\text{IMP}} \\
\langle C[\text{code}], env \rangle &\Rightarrow \langle C[\text{code}'], env' \rangle \in \mathcal{A}_{\text{IMP}} \quad \text{if } \langle \text{code}, env \rangle \Rightarrow \langle \text{code}', env' \rangle \in \mathcal{A}_{\text{IMP}}
\end{aligned}$$

where $C ::= _ | C \text{ op } e | i \text{ op } C | \text{if } C \text{ then } s_1 \text{ else } s_2 | v := C | C \text{ s}$

Fig. 2. Matching logic semantics of IMP as a set \mathcal{A}_{IMP} of reachability rules (schemata). op ranges over the binary function symbols and op_{Int} is their denotation in \mathcal{T} .

We now define matching logic formulae, or patterns, and their satisfaction. For the rest of this section we fix an ML signature $(\text{Cfg}, \Sigma, \Pi)$ and use the standard notation $T_{\Sigma, s}(\mathcal{X})$ for the set of Σ -terms of sort s with variables in \mathcal{X} .

Definition 2. Matching logic formulae, or **patterns**, extend the first-order formulae with terms $\pi \in T_{\Sigma, \text{Cfg}}(\mathcal{X})$, called **basic patterns**:

$$\varphi ::= \forall x. \varphi \mid \varphi \wedge \varphi \mid \neg \varphi \mid P(t_1, \dots, t_n) \mid \pi$$

Example 2. For example, $\varphi = \langle \text{skip}, x \mapsto i \rangle \wedge i \geq 0$ is such a pattern. Intuitively, it denotes terminal configurations (the **skip**) with a state holding exactly one variable $x \in \text{Var}$ bound to an integer $i \in \mathcal{X}$ greater than or equal to 0.

Definition 3. Let \mathcal{T} be a fixed (Σ, Π) first-order model, whose elements of sort \mathcal{T}_{Cfg} are called **concrete configurations**. We define the **satisfaction relation** between pairs $(\gamma \in \mathcal{T}_{\text{Cfg}}, \rho : \mathcal{X} \rightarrow \mathcal{T})$ and patterns as follows:

$$\begin{aligned}
(\gamma, \rho) \models \forall x. \varphi &\quad \text{iff } (\gamma, \rho') \models \varphi \text{ for any } \rho' \text{ with } \rho'(y) = \rho(y) \text{ for all } y \in \mathcal{X} \setminus \{x\} \\
(\gamma, \rho) \models \varphi_1 \wedge \varphi_2 &\quad \text{iff } (\gamma, \rho) \models \varphi_1 \text{ and } (\gamma, \rho) \models \varphi_2 \\
(\gamma, \rho) \models \neg \varphi &\quad \text{iff it is not the case that } (\gamma, \rho) \models \varphi \\
(\gamma, \rho) \models P(t_1, \dots, t_n) &\quad \text{iff the predicate } P \text{ holds on } \rho(t_1), \dots, \rho(t_n) \\
(\gamma, \rho) \models \pi &\quad \text{iff } \gamma = \rho(\pi)
\end{aligned}$$

When $(\gamma, \rho) \models \varphi$ we say that “ γ matches pattern φ with witness ρ ”.

Example 3. Consider again pattern $\langle \text{skip}, x \mapsto i \rangle \wedge i \geq 0$. Then $(\langle \text{skip}, x \mapsto 42 \rangle, i \mapsto 42)$ and $(\langle \text{skip}, x \mapsto 0 \rangle, i \mapsto 0)$ satisfy it. However, $(\langle \text{skip}, x \mapsto 3 \rangle, i \mapsto -3)$ and $(\langle \text{skip}, x \mapsto 3 \rangle, i \mapsto 4)$ do not satisfy it, the former because i is negative and the latter because the *matching* fails: $\langle \text{skip}, x \mapsto 3 \rangle \neq (i \mapsto 4)(\langle \text{skip}, x \mapsto i \rangle)$.

Since matching logic patterns generalize configuration terms, the transition relation between configurations, which is at the core of programming language semantics, can be extended to a relation between matching logic patterns [16]:

Definition 4. A **reachability rule** is a pair $\varphi \Rightarrow \varphi'$, with φ and φ' patterns.

Example 4. The reachability rule $\langle x := i; s, env \rangle \Rightarrow \langle s, env[x \mapsto i] \rangle$ gives the semantics of IMP’s assignment, and $\langle \text{if } i \text{ then } s \text{ else } t, env \rangle \wedge i \neq 0 \Rightarrow \langle s, env \rangle$ gives the semantics of the conditional statement for the positive case. In these rules, $x \in \mathcal{X}_{\text{Var}}$, $i \in \mathcal{X}_{\text{Int}}$, and $s, t \in \mathcal{X}_{\text{Stmt}}$ are variables of appropriate sorts.

$$\begin{aligned}
\langle i_1 \text{ op } i_2 \rangle &\Rightarrow \langle i_1 \text{ op}_{Int} i_2 \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle \text{if } i \text{ then } e_1 \text{ else } e_2 \rangle \wedge i \neq 0 &\Rightarrow \langle e_1 \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle \text{if } 0 \text{ then } e_1 \text{ else } e_2 \rangle &\Rightarrow \langle e_2 \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle \text{letrec } f \ x = e \ \text{in } e' \rangle &\Rightarrow \langle e'[f \mapsto (\mu f. \lambda x. e)] \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle (\lambda x. e) \ v \rangle &\Rightarrow \langle e[x \mapsto v] \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle \mu x. e \rangle &\Rightarrow \langle e[x \mapsto (\mu x. e)] \rangle \in \mathcal{A}_{\text{FUN}} \\
\langle C[c] \rangle &\Rightarrow \langle C[c'] \rangle \in \mathcal{A}_{\text{FUN}} \quad \text{if } \langle c \rangle \Rightarrow \langle c' \rangle \in \mathcal{A}_{\text{FUN}}
\end{aligned}$$

where $C ::= _ \mid C \text{ op } e \mid \text{if } C \text{ then } e_1 \text{ else } e_2 \mid C \ e \mid v \ C$

Fig. 3. Matching logic semantics of FUN as a set \mathcal{A}_{FUN} of reachability rules.

Definition 5. A *matching logic(-based) semantics* for a language is a tuple $(\text{Cfg}, \Sigma, \Pi, \mathcal{A}, \mathcal{T}, \rightarrow_{\mathcal{T}})$, where $(\text{Cfg}, \Sigma, \Pi)$ is a matching logic signature and \mathcal{A} a set of reachability rules, \mathcal{T} is a configuration model like in Definition 3 and $\rightarrow_{\mathcal{T}}$ is the transition system generated by \mathcal{A} on \mathcal{T} , that is, $\rightarrow_{\mathcal{T}} \subseteq \mathcal{T}_{\text{Cfg}} \times \mathcal{T}_{\text{Cfg}}$ with $\gamma \rightarrow_{\mathcal{T}} \gamma'$ iff there exist $\varphi \Rightarrow \varphi' \in \mathcal{A}$ and ρ such that $(\gamma, \rho) \models \varphi$ and $(\gamma', \rho) \models \varphi'$.

As discussed in [16], conventional operational semantics of programming languages can be regarded as matching logic semantics: Σ includes the abstract syntax of the language as well as the syntax of the various operations in the needed mathematical domains; \mathcal{A} is the (possibly infinite) set of operational semantics rules of the language; \mathcal{T} is the model of configurations of the language together merged with the needed mathematical domains, and the relation $\rightarrow_{\mathcal{T}}$ is precisely the transition relation defined by the operational semantics. Fig. 2 and Fig. 3 show matching logic semantics of the IMP and FUN languages, respectively, obtained by mechanically representing conventional operational semantics of these languages based on reduction semantics with evaluation contexts into matching logic. The only observable difference between the original semantics of these languages and their matching logic semantics is that the side conditions have been conjuncted with the left-hand-side patterns in the positive case of the conditionals. Note that \mathcal{A}_{IMP} and \mathcal{A}_{FUN} are infinite, as the rules in Figs. 2 and 3 are schemata in meta-variable C (the evaluation context).

The following definitions allow us to state properties and to reason about the dynamic behavior of configurations denoted by matching logic patterns:

Definition 6. Let \mathcal{S} be a matching logic semantics $(\text{Cfg}, \Sigma, \Pi, \mathcal{A}, \mathcal{T}, \rightarrow_{\mathcal{T}})$. We define the following relations and predicates on patterns of \mathcal{S} :

- $\varphi \Rightarrow_{\mathcal{S}} \varphi'$ iff for all $(\gamma, \rho) \models \varphi$ there is a γ' such that $\gamma \rightarrow_{\mathcal{T}} \gamma'$ and $(\gamma', \rho) \models \varphi'$;
- $\varphi \Rightarrow_{\mathcal{S}}^* \varphi'$ iff for all $(\gamma, \rho) \models \varphi$ there is a γ' such that $\gamma \rightarrow_{\mathcal{T}}^* \gamma'$ and $(\gamma', \rho) \models \varphi'$;
- $\varphi \Rightarrow_{\mathcal{S}}^{\dagger} \varphi'$ iff for all $(\gamma, \rho) \models \varphi$ there is a γ' such that $\gamma \rightarrow_{\mathcal{T}}^{\dagger} \gamma'$ and $(\gamma', \rho) \models \varphi'$;
- $\varphi \uparrow_{\mathcal{S}}^{\infty}$, read φ diverges, iff for all $(\gamma, \rho) \models \varphi$ we have that γ does not terminate in $\rightarrow_{\mathcal{T}}$, i.e., there is an infinite sequence $\gamma \rightarrow_{\mathcal{T}} \gamma_1 \rightarrow_{\mathcal{T}} \gamma_2 \rightarrow_{\mathcal{T}} \dots$

Example 5. Let \mathcal{S}_{IMP} denote the matching logic semantics of IMP in Fig. 2. Then $\langle x := x; y := y, \emptyset \rangle \Rightarrow_{\mathcal{S}_{\text{IMP}}}^* \langle \text{skip}, (x \mapsto x, y \mapsto y) \rangle$ (if we perform the two assignments) and $\langle x := x; y := y, \emptyset \rangle \Rightarrow_{\mathcal{S}_{\text{IMP}}}^* \langle x := x; y := y, \emptyset \rangle$ (by not taking any step). The former also holds when we replace $\Rightarrow_{\mathcal{S}_{\text{IMP}}}^*$ with $\Rightarrow_{\mathcal{S}_{\text{IMP}}}^{\dagger}$, but the latter not. Also, $\langle \text{while } 1 \text{ do skip}, \emptyset \rangle \uparrow_{\mathcal{S}_{\text{IMP}}}^{\infty}$, but it is not the case that $\langle x := x; y := y, \emptyset \rangle \uparrow_{\mathcal{S}_{\text{IMP}}}^{\infty}$.

3 Aggregating Matching Logic Semantics

In order to prove mutual equivalence of patterns in two (possibly) different matching logic semantics defining two (possibly) different programming languages, it is technically convenient to first put the two matching logic semantics together into one (larger) matching logic semantics. In this section we show how to aggregate two matching logic semantics, noticing that the construction works also for more than two semantics. The idea is to add infrastructure for pairing a configuration of one matching logic semantics with a configuration of the other, into a configuration of the aggregated semantics. Care needs to be exercised to factor out the common parts of the two semantics, such as the mathematical domains that happen to be used in both semantics (integers, maps, variables, etc.). The key result that allows for matching logic semantics to be aggregated is the amalgamation theorem of first-order (multi-sorted) logic [5, 4].

Let $\mathcal{S}_L = (Cfg_L, \Sigma_L, \Pi_L, \mathcal{A}_L, \mathcal{T}_L, \rightarrow_{\mathcal{T}_L})$ and $\mathcal{S}_R = (Cfg_R, \Sigma_R, \Pi_R, \mathcal{A}_R, \mathcal{T}_R, \rightarrow_{\mathcal{T}_R})$ be two fixed matching logic semantics, for example two programming language definitions, like those of IMP and FUN defined in Section 2. Also, suppose that their signatures and predicates share a common part, (Σ_0, Π_0) , for example the domains of integers, maps, etc., which can be possibly renamed in the two semantics. Formally, we assume that there exist first-order logic morphisms

$$(\Sigma_L, \Pi_L) \xleftarrow{h_L} (\Sigma_0, \Pi_0) \xrightarrow{h_R} (\Sigma_R, \Pi_R)$$

and that the two configuration models \mathcal{T}_L and \mathcal{T}_R agree on their common part, that is, $\mathcal{T}_L \upharpoonright_{h_L} = \mathcal{T}_R \upharpoonright_{h_R}$, where \upharpoonright is the first-order model reduct operation.

We next show how to define the aggregation of \mathcal{S}_L and \mathcal{S}_R . First, compute the first-order logic pushout of the morphisms h_L and h_R above, known to exist [5, 4]:

$$(\Sigma_L, \Pi_L) \xrightarrow{h'_L} (\Sigma', \Pi') \xleftarrow{h'_R} (\Sigma_R, \Pi_R)$$

Recall that a pushout puts together the two first-order signatures, possibly renaming sorts and operations happening to have the same name, but making sure that the shared sorts and symbols are correctly preserved. In particular, $h_L; h'_L = h_R; h'_R$ (we use the diagrammatic notation for morphism composition).

Second, we add infrastructure for pairing configurations in the two object matching logic semantics. Specifically, add a new sort, say Cfg , and a new pairing operation, say $\langle -, - \rangle : Cfg_L \times Cfg_R \rightarrow Cfg$, to Σ' , and let us name the resulting first-order signature (Σ, Π) —no predicates were added to Π' , so $\Pi = \Pi'$.

The constructions above are summarized in the following diagram

$$\begin{array}{ccccc} (\Sigma_0, \Pi_0) & \xrightarrow{h_R} & (Cfg_R, \Sigma_R, \Pi_R) & & \\ \downarrow h_L & & \downarrow h'_R & & \\ (Cfg_L, \Sigma_L, \Pi_L) & \xrightarrow{h'_L} & (\Sigma', \Pi') & \xrightarrow{\iota} & (Cfg, \Sigma, \Pi) \end{array}$$

with ι an inclusion (matching logic signatures also list their configuration sorts).

Third, we can aggregate \mathcal{A}_L and \mathcal{A}_R into a set of rules \mathcal{A} , by replacing each basic pattern π in \mathcal{A}_L by a basic pattern $\langle \pi, cfg_R \rangle$ (here π is now regarded as a *term* of sort Cfg_L , not as a formula) for a fresh arbitrary but fixed variable cfg_R of sort Cfg_R , and similarly for \mathcal{A}_R . Intuitively, \mathcal{A} consists of all the semantic steps of \mathcal{A}_L and all of \mathcal{A}_R , on their respective position in the aggregated configuration.

Fourth, since $\mathcal{T}_L \upharpoonright_{h_L} = \mathcal{T}_R \upharpoonright_{h_R}$, the amalgamation theorem of (multi-sorted) first-order logic [5, 4] tells us that there exists a *unique* (Σ', Π') -model \mathcal{T}' such that $\mathcal{T}' \upharpoonright_{h'_L} = \mathcal{T}' \upharpoonright_{h'_R}$. We can further (freely) extend \mathcal{T}' into a (Σ, Π) -model \mathcal{T} by adding a new carrier $\mathcal{T}_{Cfg} = \{ \langle \gamma_L, \gamma_R \rangle \mid \gamma_L \in \mathcal{T}'_{Cfg_L}, \gamma_R \in \mathcal{T}'_{Cfg_R} \}$ and interpreting the pairing operations as expected: $\mathcal{T}_{\langle \cdot, \cdot \rangle}(\gamma_L, \gamma_R) = \langle \gamma_L, \gamma_R \rangle$.

Finally, we construct the transition relation $\rightarrow_{\mathcal{T}}$ like in Definition 5, and this way obtain our aggregated matching logic semantics $\mathcal{S} = (Cfg, \Sigma, \Pi, \mathcal{A}, \mathcal{T}, \rightarrow_{\mathcal{T}})$.

Example 6. The following aggregated matching logic pattern matches all the pairs of IMP and FUN configurations which cannot take any further steps and where the FUN configuration holds a non-negative integer that can be retrieved in the x variable in the IMP configuration:

$$\exists i. (\langle \langle \text{skip}, (x \mapsto i, _) \rangle, \langle i \rangle \rangle \wedge i > 0).$$

The underscore $(_)$ denotes an anonymous variable that matches the rest of the environment. The variable i is of sort integer and appears in both configurations (the sort **Int** is shared). The program variable x is a constant of sort **Var**.

By looking only at the left-hand side (or only at the right-hand side) component of each pair of configurations in an aggregated matching logic pattern and keeping all the other logical connectives unchanged, we obtain a pattern that specifies configurations in the corresponding matching logic semantics. We call these operations *projections*:

Definition 7. *The left projection $pr_L(\varphi)$ (resp. the right projection $pr_R(\varphi)$) of an aggregated matching logic formula φ is obtained by replacing each pattern of the form $\langle t_L, t_R \rangle$ by its left component t_L (resp. its right component t_R).*

Example 7. For $\varphi \equiv \exists i. (\langle \langle \text{skip}, (x \mapsto i, _) \rangle, \langle i \rangle \rangle \wedge i > 0)$, the aggregated matching logic pattern from the previous example, the left projection is $pr_L(\varphi) \equiv \exists i. (\langle \text{skip}, (x \mapsto i, _) \rangle \wedge i > 0)$ and the right projection is $pr_R(\varphi) \equiv \exists i. (\langle i \rangle \wedge i > 0)$.

Note that by taking the left or the right projection of an aggregated pattern, we may not end up with a matching logic pattern over the syntax of the corresponding matching logic semantics. Indeed, for a pattern φ of the aggregated matching logic semantics $\mathcal{S} = (Cfg, \Sigma, \Pi, \mathcal{A}, \mathcal{T}, \rightarrow_{\mathcal{T}})$ constructed above, the left projection $pr_L(\varphi)$ is a pattern in a matching logic semantics that extends \mathcal{S}_L to first-order syntax (Σ', Π') , and similarly for the right projection. Nevertheless, the following property holds: for any $\gamma_L \in \mathcal{T}_{Cfg_L}$, $\gamma_R \in \mathcal{T}_{Cfg_R}$, and $\rho : \mathcal{X} \rightarrow \mathcal{T}$, $(\langle \gamma_L, \gamma_R \rangle, \rho) \models \varphi$ iff $(\gamma_L, \rho) \models pr_L(\varphi)$ and $(\gamma_R, \rho) \models pr_R(\varphi)$, where for simplicity we used the same γ_L , γ_R , and ρ for their images through the model reducts.

4 Proving Mutual Program Equivalence

Here we provide a language-parametric foundation for showing equivalence of programs written in possibly different languages. Like in the previous section, we generically assume that the two languages are given as matching logic semantics $\mathcal{S}_L = (Cfg_L, \Sigma_L, \Pi_L, \mathcal{A}_L, \mathcal{T}_L, \rightarrow_{\mathcal{T}_L})$ and $\mathcal{S}_R = (Cfg_R, \Sigma_R, \Pi_R, \mathcal{A}_R, \mathcal{T}_R, \rightarrow_{\mathcal{T}_R})$ with aggregation $\mathcal{S} = (Cfg, \Sigma, \Pi, \mathcal{A}, \mathcal{T}, \rightarrow_{\mathcal{T}})$, but when we discuss examples we assume them to be the semantics \mathcal{S}_{IMP} and \mathcal{S}_{FUN} of, respectively, IMP and FUN.

4.1 Specifying Equivalent Programs

Aggregate matching logic patterns can be used to specify pairs of configurations of the two involved languages:

Definition 8. *The denotation of an aggregated matching logic pattern φ , written $\llbracket \varphi \rrbracket$, is the set of all pairs of configurations that satisfy it:*

$$\llbracket \varphi \rrbracket = \{ \langle \langle \gamma_L, \gamma_R \rangle \mid \text{there exists a valuation } \rho \text{ such that } \langle \langle \gamma_L, \gamma_R \rangle, \rho \rangle \models \varphi \}.$$

This extends sets of patterns E , written $\llbracket E \rrbracket$, as expected:

$$\llbracket E \rrbracket = \{ \langle \gamma_L, \gamma_R \rangle \mid \text{there exist } \varphi \in E \text{ such that } \langle \gamma_L, \gamma_R \rangle \in \llbracket \varphi \rrbracket \}.$$

Example 8. The following set

$$E = \{ \exists i. (\langle \langle \text{skip}, (\mathbf{x} \mapsto i, -) \rangle, \langle i \rangle \rangle) \}$$

captures all pairs of IMP and respectively FUN configurations that have terminated (since there is no more code to execute) and where the IMP variable \mathbf{x} holds the same integer as the result of the FUN program.

Suppose we have an IMP program that computes its result in a variable \mathbf{x} and suppose we want to show it computes the same integer result as a FUN program. Then the denotation $\llbracket E \rrbracket$ of set E above holds exactly the set of pairs of terminal configurations in which the two programs should end in order for them to compute the same result.

When trying to prove that two programs compute the same result, it is tempting to say that the two programs should reach the same configuration at the end. However, this is not feasible since the configuration might contain additional information (such as temporary variables) that was used in the computation but is not part of the result. When testing if the final configurations are the same in the two programs, it is important to ignore such additional information. In the example above, only the variable \mathbf{x} is inspected (the values of all other variables are ignored) when comparing final configurations. Another aspect is that, when working in a general setting where we are comparing programs from two arbitrary programming languages, the configurations of the two languages might be significantly different. This is the case above, with the configuration for IMP holding code and an environment and the configuration for FUN holding only (extended) lambda expressions. Therefore, in general, to show that two

programs end up with the same result there is a need to design such a set $\llbracket E \rrbracket$ of "base" pairs which are known to be equivalent.

Two programs are then considered *mutually equivalent* when, for all inputs, they both diverge or they both reach a pair in the base equivalence $\llbracket E \rrbracket$. This intuition is captured by the following definition:

Definition 9. We write $\models \varphi \Downarrow^\infty E$, and say that φ *reaches* E , iff for all configurations γ_L, γ_R and for all valuations ρ such that $(\langle \gamma_L, \gamma_R \rangle, \rho) \models \varphi$ we have that at least one of the following conditions holds:

1. both γ_L and γ_R diverge;
2. there are configurations γ'_L, γ'_R with $\gamma_L \rightarrow_{\mathcal{T}}^* \gamma'_L$, $\gamma_R \rightarrow_{\mathcal{T}}^* \gamma'_R$ and $(\gamma'_L, \gamma'_R) \in \llbracket E \rrbracket$.

Example 9. Let $E = \{\exists i. (\langle \text{skip}, (x \mapsto i, -) \rangle, \langle i \rangle)\}$ and let

$$\begin{aligned} \varphi_1 &= \exists n. \langle \langle \text{code}_1, n \mapsto n \rangle, \langle \text{exp}_1(n) \rangle \rangle \\ \varphi_2 &= \langle \langle \text{while } 1 \text{ do skip}, \emptyset \rangle, \langle \text{letrec } f \ x = f(x+1) \text{ in } f(1) \rangle \rangle \\ \varphi_3 &= \exists n. \langle \langle \text{code}_3, n \mapsto n \rangle, \langle \text{exp}_3(n) \rangle \rangle. \end{aligned}$$

where $\text{code}_1 \equiv i:=1; x:=0; \text{while } i \leq n \text{ do } (x:=x+i; i:=i+1)$ is the IMP program that computes the sum of the numbers from 1 to n , where $\text{exp}_1(n) \equiv \text{letrec } f \ x = \text{if } x=1 \text{ then } 1 \text{ else } x+f(x-1) \text{ in } f(n)$ is the FUN program computing the same sum and where $\text{code}_3 \equiv \text{PGM}_L$ and, resp., $\text{exp}_3(n) \equiv \text{PGM}_R(n)$ are the IMP and FUN programs in Fig. 5 that compute the Collatz function.

We have that $\models \varphi_1 \Downarrow^\infty E$ since both programs end up in a pair from $\llbracket E \rrbracket$: $\langle \text{code}_1, n \mapsto n \rangle \rightarrow_{\mathcal{T}}^* \langle \text{skip}, x \mapsto 1+2+\dots+n \rangle$ and $\langle \text{exp}_1(n) \rangle \rightarrow_{\mathcal{T}}^* \langle 1+2+\dots+n \rangle$. We also have that $\models \varphi_2 \Downarrow^\infty E$, since both configurations in φ_2 clearly diverge. We also have that $\models \varphi_3 \Downarrow^\infty E$, but this is more difficult to establish. In fact, it is currently only conjectured (not proven) that the programs terminate no matter what the input value n is. But it can be proven that if one does not terminate, the other does not terminate either and therefore $\models \varphi_3 \Downarrow^\infty E$ holds independently of the Collatz conjecture. However, $\models \varphi_3 \Downarrow^\infty E$ is more difficult to show than the previous examples since it is not clear if both programs terminate or diverge.

We next propose a proof system that allows us to derive such properties.

4.2 Proof System

In this section, we introduce a proof system that is able to derive sequents of the form $\vdash \varphi \Downarrow^\infty E$ denoting mutual equivalences that are sound in the sense that $\vdash \varphi \Downarrow^\infty E$ implies $\models \varphi \Downarrow^\infty E$. Fig. 4 contains the 5-rule proof system for proving mutual equivalence of programs.

The first rule is AXIOM. There is nothing surprising about this rule; it simply states that if an equivalence is known to be true, then it can be derived.

The second rule is STEP. This rule allows to take an arbitrary finite number of steps (zero, one or more steps) in each of the two programs. If by taking such steps from φ to φ' , we reach an equivalence φ' that is derivable, then

$$\begin{array}{c}
\text{AXIOM} \frac{\varphi \in E}{\vdash \varphi \Downarrow^\infty E} \quad \text{STEP} \frac{\varphi \Rightarrow^* \varphi' \quad \vdash \varphi' \Downarrow^\infty E}{\vdash \varphi \Downarrow^\infty E} \quad \text{CONSEQ} \frac{\models \varphi \rightarrow \varphi' \quad \vdash \varphi' \Downarrow^\infty E}{\vdash \varphi \Downarrow^\infty E} \\
\\
\text{CASE ANALYSIS} \frac{\vdash \varphi \Downarrow^\infty E \quad \vdash \varphi' \Downarrow^\infty E}{\vdash \varphi \vee \varphi' \Downarrow^\infty E} \quad \text{CIRCULARITY} \frac{\vdash \varphi' \Downarrow^\infty E \cup \{\varphi\} \quad \varphi \Rightarrow^+ \varphi'}{\vdash \varphi \Downarrow^\infty E}
\end{array}$$

Fig. 4. Mutual Equivalence Proof System. We use $\varphi \Rightarrow^* \varphi'$ as syntactic sugar for $pr_L(\varphi) \Rightarrow_{S_L}^* pr_L(\varphi')$ and $pr_R(\varphi) \Rightarrow_{S_R}^* pr_R(\varphi')$ and $\varphi \Rightarrow^+ \varphi'$ as syntactic sugar for $pr_L(\varphi) \Rightarrow_{S_L}^+ pr_L(\varphi')$ and $pr_R(\varphi) \Rightarrow_{S_R}^+ pr_R(\varphi')$.

we conclude that φ must also be derivable. The STEP rule requires an oracle to reason about reachability in operational semantics. This oracle can be, for example, the reachability proof system in [16], but any other valid reasoning will also work.

The third rule is CONSEQ(ue)nce). This rule states that if an equivalence formula φ implies another equivalence formula φ' (which means that φ' is more general than φ) and the formula φ' is derivable, then φ must also be derivable. The required implication might seem surprizing at first (we might expect it in reverse), but the intuition is that φ' is more general than φ . Therefore if we are able to prove the equivalence φ' , then φ must also hold. This rule is used in the example proof tree below (in Fig. 5) to rearrange a formula of the form $(n > 0 \vee n = 0) \wedge \dots$ into $n \geq 0 \wedge \dots$. Another possible use of CONSEQ would be, for example, to transform a more particular case, like “ $n = 20$ ”, into a more general case “ n is even” in order to be able to apply other rules.

The fourth case is CASE ANALYSIS. This allows to branch the proof depending on the different cases to consider. Typically, CASE ANALYSIS is used to branch the proof when the two programs also branch. In the proof tree below (in Fig. 5), this rule is used to perform a case analysis between the case where both programs end (because of reaching the termination condition $n = 0$) and where the programs continue ($n > 0$).

The fifth rule is CIRCULARITY. This rule is used to handle repetitive program structures such as loops or recursive functions. CIRCULARITY allows to *postulate* that the equivalence being proven (φ) holds, make progress ($\varphi \Rightarrow^+ \varphi'$) in both programs that we want to show equivalent, and then derive φ' possibly using φ as an axiom, i.e., $\vdash \varphi' \Downarrow^\infty E \cup \{\varphi\}$. We use this rule in the proof tree below to assume that at the start of the repetitive behavior (the loop for the program on the left and the recursive call for the program on the right) the two programs are equivalent; we make progress by executing the body of the loop on the left and the body of the recursive on the right and end up with the equivalence that we assumed to hold. The rule is sound because we require both programs to make progress. Therefore, intuitively, when $\vdash \varphi' \Downarrow^\infty E \cup \{\varphi\}$ is derivable, either both programs diverge because φ is applied as an axiom in the proof tree or the programs end up in E . As for the first rule, an oracle to reason about reachability in operational semantics is also needed here.

<pre> PGM_L := c := 1; LOOP_L LOOP_L := while (n != 1) c := c + 1; if (n % 2 != 0) then n := n + n + n + 1 else n := n / 2 φ := ∃i, n. (n > 0 ∧ ⟨LOOP_L, n ↦ n, c ↦ i⟩, ⟨i + LOOP_R⟩) </pre>	<pre> PGM_R(n) := letrec f n = LOOP_R in f(n) LOOP_R := 1 + if (n != 1) then if (n % 2 != 0) then f(n + n + n + 1) else f(n / 2) else 0 </pre>
<ol style="list-style-type: none"> 1. $\vdash \exists i, n. (\langle \text{skip}, n \mapsto n, c \mapsto i \rangle, \langle i \rangle)$ $\Downarrow^\infty E$ AXIOM 2. $\vdash \exists i, n. (\langle \text{skip}, n \mapsto n, c \mapsto i \rangle, \langle i \rangle)$ $\Downarrow^\infty E \cup \{\varphi\}$ AXIOM 3. $\vdash \exists i, n. (n = 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E$ STEP(1) 4. $\vdash \exists i, n. (n = 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E \cup \{\varphi\}$ STEP(2) 5. $\vdash \exists i, n. (n > 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E \cup \{\varphi\}$ AXIOM 6. $\vdash \exists i, n. (n \geq 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E \cup \{\varphi\}$ CONSEQ(CA(4, 5)) 7. $\vdash \exists i, n. (n > 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E$ CIRCULARITY (6) 8. $\vdash \exists i, n. (n \geq 0 \wedge \langle \text{LOOP}_L, n \mapsto n, c \mapsto i \rangle, \langle i + \text{LOOP}_R \rangle)$ $\Downarrow^\infty E$ CONSEQ(CA(3, 7)) 9. $\vdash \exists i, n. (n \geq 0 \wedge \langle \text{PGM}_L, n \mapsto n \rangle, \langle \text{PGM}_R(n) \rangle)$ $\Downarrow^\infty E$ STEP (8) 	

Fig. 5. Formal proof showing that the two Collatz programs are mutually equivalent. CA stands for CASE ANALYSIS. CONSEQ is used in the proof tree above to show that $n > 0 \vee n = 0$ implies $n \geq 0$. For simplicity, `letrec` is not desugared into μ .

Theorem 1 (Soundness). *For any set of aggregated matching logic patterns E and for any aggregated matching logic pattern φ , if the sequent $\vdash \varphi \Downarrow^\infty E$ is derivable using the proof system given in Fig. 4 then $\models \varphi \Downarrow^\infty E$.*

Proof. (complete proof in the appendix)

4.3 Example

Here we show the proof tree for the equivalence of the two Collatz programs in Fig. 5. As we have already discussed, in order to talk about mutual equivalence, we have to establish a “base” equivalence that contains programs that are clearly equivalent. For this case study, for the “base” equivalence, we choose to equate FUN programs that terminate by returning an integer i with IMP programs that terminate with the same integer i in variable c . The following set $E = \{\exists i. \langle \langle \text{skip}, (c \mapsto i, -) \rangle, \langle i \rangle \rangle\}$ captures exactly the intuition above. It says that an IMP configuration $\langle \text{skip}, (c \mapsto i, -) \rangle$ (describing programs that stopped (because the code cell contains `skip`) and that have the integer i in the c memory cell) is equivalent to a FUN configuration that contains exactly the integer i . The proof tree in Fig. 5 shows that the two programs are equivalent.

5 Discussion, Related Work and Conclusion

We have introduced mutual matching logic, a 5-rule proof system for proving mutual equivalence of programs. Mutual equivalence is a natural equivalence between programs: two programs are mutually equivalent if either they both diverge or if they eventually reach the same state. Mutual equivalence can be used, for example, to prove that compiler transformations preserve behavior.

Our approach is language independent. The proof system takes as input two language semantics (in the form of reachability rules) that share certain domains such as integers and produces sequents of the form $\vdash \varphi \Downarrow^\infty E$ whose semantics is that for any pair of programs that matches φ , both programs diverge or they reach a state in E . Note that in our running example (the two Collatz programs), both programs have a parameter n that is left unspecified. This shows that our approach allows parameterized programs. Although because of space limitations we do not explicitly state this, our approach can handle symbolic programs as well. For example, we can show using our proof system that in IMP extended with a `for` loop, the two programs:

```

i := 1;
while (i <= n) do
  s;
  i := i + 1;
for i := 1 to n do
  s;

```

are equivalent for a symbolic statement s . We specify the symbolic statement s as an additional statement in the signature of IMP and we render explicit the constraints on s in its semantics using reachability rules.

Related Work. It was first remarked by Hoare in [8] that program equivalence might be easier than program correctness. Among the recent works on equivalence we mention [7, 6, 2]. 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 equivalence-verification to multi-threaded programs. They use operational semantics (of a specific language they designed, called LPL) and proof systems, and formally prove their proof system's soundness. In [7] a classification of equivalence relations used in program-equivalence research is given, one of which is mutual equivalence (called *full equivalence* there). The main difference with our approach is that our proof system is language-independent, i.e., it is parametric in the semantics of the two languages in which candidate equivalent programs are written; whereas the deductive system of [7] proves equivalence for LPL programs. On the other hand, [7] propose deductive systems for several kinds of equivalences, whereas we focus on mutual (a.k.a. full) equivalence only. In [9], an implementation of a parametrized equivalence prover is presented.

A lot of work on program equivalence arise from the verification of compilation in a broad sense. One approach is full compiler verification [11], and another one is the individual verification of each compilation [13] (we only cite two of the most relevant recent works). Other work target specific classes of languages: functional [14], microcode [1], CLP [3]. In order to be less language-specific some approaches advocate the use of intermediate languages, such as [10], which works on the Boogie intermediate language. Finally, our own related work [12] gives a proof system for another equivalence relation between programs that is based on bisimulation and an observation relation and that uses other technical mechanisms. We believe that the equivalence relation that we consider in this article is more natural for certain classes of applications such as proving compilers.

Further Work. Our definition (Definition 9) of mutual equivalence is *existential* in the sense that two programs are equivalent when there exists execution

paths in each of the programs such that the paths diverge or end in configurations that are known to be equivalent. Although for deterministic languages this cannot constitute a problem (there exists exactly one execution path for each program), for non-deterministic languages stronger equivalences might be desirable. We leave such stronger equivalences as object of further study. Another issue is completeness. Although relative completeness results have been shown for matching logic based proof systems for showing partial correctness [16], it is less clear how a relevant relative-completeness result can be obtained for equivalence, since the problem is known to be Π_2^0 -complete. Another issue that we leave for further study is compositionality. Our goal here was just to obtain a sound and useful language independent proof system for reasoning about equivalence.

References

1. T. Arons, E. Elster, L. Fix, S. Mador-Haim, M. Mishaeli, J. Shalev, E. Singerman, A. Tiemeyer, M. Y. Vardi, and L. D. Zuck. Formal verification of backward compatibility of microcode. In *CAV, LNCS 3576*, pages 185–198, 2005.
2. S. Chaki, A. Gurfinkel, and O. Strichman. Regression verification for multi-threaded programs. In *VMCAI, LNCS 7148*, pages 119–135, 2012.
3. S. Craciunescu. Proving the equivalence of CLP programs. In *ICLP, LNCS 2401*, pages 287–301, 2002.
4. R. Diaconescu. *Institution-independent Model Theory*. Birkhauser Basel, 2008.
5. H. Ehrig and B. Mahr. *Fundamentals of Algebraic Specification 1: Equations und Initial Semantics*. Springer, 1985.
6. B. Godlin and O. Strichman. Regression verification: proving the equivalence of similar programs. *Software Testing, Verification and Reliability*. To appear.
7. B. Godlin and O. Strichman. Inference rules for proving the equivalence of recursive procedures. *Acta Informatica*, 45(6):403–439, 2008.
8. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969.
9. S. Kundu, Z. Tatlock, and S. Lerner. Proving optimizations correct using parameterized program equivalence. In *PLDI*, pages 327–337. ACM, 2009.
10. S. Lahiri, C. Hawblitzel, M. Kawaguchi, and H. Rebêlo. Symdiff: A language-agnostic semantic diff tool for imperative programs. In *CAV, LNCS 7358*, pages 712–717, 2012.
11. X. Leroy. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7):107–115, 2009.
12. D. Lucanu and V. Rusu. Program equivalence by circular reasoning. Technical Report RR-8116, INRIA, 2012.
13. G. Necula. Translation validation for an optimizing compiler. In *PLDI*, pages 83–94. ACM, 2000.
14. A. Pitts. Operational semantics and program equivalence. In *Applied Semantics Summer School, LNCS 2395*, pages 378–412, 2002.
15. G. Roşu, C. Ellison, and W. Schulte. Matching logic: An alternative to Hoare/Floyd logic. In *AMAST, LNCS 6486*, pages 142–162, 2010.
16. G. Roşu and A. Stefanescu. Checking reachability using matching logic. In *OOP-SLA*, pages 555–574. ACM, 2012.

A Proofs

This appendix contains the proof of the soundness theorem that was left out for space reasons.

In the following, we let $C \in \{L, R\}$ denote either left or right. By \bar{C} we denote the single element of the set $\{L, R\} \setminus \{C\}$.

Let E be a set of mutual matching logic formulae. Let \mathcal{A}_L and \mathcal{A}_R be a set of reachability formulae which describe the semantics of two languages: \mathcal{A}_L the “left” language and \mathcal{A}_R the “right” language. We extend the definition of $\models \varphi \Downarrow^\infty E$ to sets of mutual matching logic formulae as expected:

Definition 10. *If F is a set of mutual matching logic formulae, then we write*

$$\models F \Downarrow^\infty E \text{ if } \models \varphi \Downarrow^\infty E \text{ for all } \varphi \in F.$$

The following definitions will be useful in the proof of soundness. Let G denote a set of pairs of configurations.

Definition 11. *We say that a pair (γ_L, γ_R) reaches G , written $(\gamma_L, \gamma_R) \rightarrow^* G$, if there exist configurations γ'_L and γ'_R such that $\gamma_L \rightarrow_{\mathcal{A}_L}^* \gamma'_L$, $\gamma_R \rightarrow_{\mathcal{A}_R}^* \gamma'_R$ and $(\gamma'_L, \gamma'_R) \in G$.*

Definition 12. *We say that a pair (γ_L, γ_R) diverges, written $(\gamma_L, \gamma_R) \uparrow^\infty$, if both γ_L and γ_R diverge (in \mathcal{A}_L and respectively \mathcal{A}_R).*

Definition 13. *We say that a pair (γ_L, γ_R) co-reaches G , written $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} G$, if at least one of the following conditions holds:*

1. (γ_L, γ_R) diverges (i.e. $(\gamma_L, \gamma_R) \uparrow^\infty$),
2. (γ_L, γ_R) reaches G (i.e. $(\gamma_L, \gamma_R) \rightarrow^* G$).

The following utility lemma establishes the link between models of mutual matching logic formulae and the notion of co-reachability introduced above. Its proof following trivially by unrolling the above definitions.

Lemma 1. *For all sets of mutual matching logic formulae E and for any mutual matching logic formula φ , we have that:*

$$\models \varphi \Downarrow^\infty E \text{ iff for all } \gamma_L, \gamma_R \text{ such that } (\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket, (\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket.$$

The next lemma is the core of our soundness proof.

Lemma 2 (Circularity Principle).

Let F be a set of mutual matching formulae. If for each $(\gamma_L, \gamma_R) \in \llbracket F \rrbracket$ there exist γ'_L, γ'_R such that $\gamma_L \rightarrow_{\mathcal{A}_L}^+ \gamma'_L$, $\gamma_R \rightarrow_{\mathcal{A}_R}^+ \gamma'_R$, and $(\gamma'_L, \gamma'_R) \rightarrow^{,\infty} \llbracket E \cup F \rrbracket$, then $\models F \Downarrow^\infty E$.*

Proof. Let $(\gamma_L, \gamma_R) \in \llbracket F \rrbracket$ be arbitrarily chosen. We show that $\models F \Downarrow^\infty E$. By Lemma 1 it is enough to show that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.

For any natural number $i > 0$, let $P(i)$ denote the following predicate:

$P(i)$ = either there exist configurations $\gamma_L^1, \dots, \gamma_L^i, \gamma_R^1, \dots, \gamma_R^i$ such that

$$\begin{aligned} & \gamma_L \xrightarrow{\mathcal{A}_L} \gamma_L^1 \xrightarrow{\mathcal{A}_L} \gamma_L^2 \xrightarrow{\mathcal{A}_L} \dots \xrightarrow{\mathcal{A}_L} \gamma_L^i, \\ & \gamma_R \xrightarrow{\mathcal{A}_R} \gamma_R^1 \xrightarrow{\mathcal{A}_R} \gamma_R^2 \xrightarrow{\mathcal{A}_R} \dots \xrightarrow{\mathcal{A}_R} \gamma_R^i \text{ and} \\ & (\gamma_L^1, \gamma_R^1) \rightarrow^{*,\infty} E \cup F, \dots, (\gamma_L^i, \gamma_R^i) \rightarrow^{*,\infty} E \cup F \\ \text{or } & (\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket \end{aligned}$$

By induction on i , we show that for all $i > 0$, $P(i)$ holds:

1. for $i = 1$, we have that by the hypothesis, there exist γ_L^1, γ_R^1 such that $\gamma_C \xrightarrow{\mathcal{A}_C} \gamma_C^1$ (for all $C \in \{L, R\}$), and $(\gamma_L^1, \gamma_R^1) \rightarrow^{*,\infty} \llbracket E \cup F \rrbracket$.
2. for $i > 1$, we assume that $P(i-1)$ holds and we prove that $P(i)$ holds. If $P(i-1)$ holds because $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, then by definition $P(i)$ holds as well.

Otherwise, we have that there exist configurations $\gamma_L^1, \dots, \gamma_L^{i-1}, \gamma_R^1, \dots, \gamma_R^{i-1}$ such that:

- (a) $\gamma_L \xrightarrow{\mathcal{A}_L} \gamma_L^1 \xrightarrow{\mathcal{A}_L} \gamma_L^2 \xrightarrow{\mathcal{A}_L} \dots \xrightarrow{\mathcal{A}_L} \gamma_L^{i-1}$,
- (b) $\gamma_R \xrightarrow{\mathcal{A}_R} \gamma_R^1 \xrightarrow{\mathcal{A}_R} \gamma_R^2 \xrightarrow{\mathcal{A}_R} \dots \xrightarrow{\mathcal{A}_R} \gamma_R^{i-1}$ and
- (c) $(\gamma_L^1, \gamma_R^1) \rightarrow^{*,\infty} E \cup F, \dots, (\gamma_L^{i-1}, \gamma_R^{i-1}) \rightarrow^{*,\infty} E \cup F$

As $(\gamma_L^{i-1}, \gamma_R^{i-1}) \rightarrow^{*,\infty} \llbracket E \cup F \rrbracket$, we distinguish between the following two cases:

- (a) either $(\gamma_L^{i-1}, \gamma_R^{i-1}) \rightarrow^{*,\infty} \llbracket E \rrbracket$ and, because $\gamma_L \xrightarrow{\mathcal{A}_L} \gamma_L^{i-1}$, we have that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$. This implies that $P(i)$ holds (by the second condition in definition of the predicate).
- (b) or $(\gamma_L^{i-1}, \gamma_R^{i-1}) \rightarrow^{*,\infty} \llbracket F \rrbracket$. By the hypothesis, there exist γ_L^i and γ_R^i such that $\gamma_L^{i-1} \xrightarrow{\mathcal{A}_L} \gamma_L^i, \gamma_R^{i-1} \xrightarrow{\mathcal{A}_R} \gamma_R^i$ and $(\gamma_L^i, \gamma_R^i) \rightarrow^{*,\infty} \llbracket E \cup F \rrbracket$. But this implies that $P(i)$ holds (by the first condition in the definition of the predicate).

In either case, we have shown that $P(i)$ whenever $P(i-1)$ holds.

Therefore, we conclude by induction that, for all naturals $i > 0$, $P(i)$ holds.

We will now show by contradiction that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$. We will assume by contradiction that $(\gamma_L, \gamma_R) \not\rightarrow^{*,\infty} \llbracket E \rrbracket$. As $P(i)$ holds for all naturals $i > 0$, it follows that for any $i > 0$, there exist $\gamma_L^1, \dots, \gamma_L^i, \gamma_R^1, \dots, \gamma_R^i$ such that $\gamma_L \xrightarrow{\mathcal{A}_L} \gamma_L^1 \xrightarrow{\mathcal{A}_L} \dots \xrightarrow{\mathcal{A}_L} \gamma_L^i$ and $\gamma_R \xrightarrow{\mathcal{A}_R} \gamma_R^1 \xrightarrow{\mathcal{A}_R} \dots \xrightarrow{\mathcal{A}_R} \gamma_R^i$. This means that both γ_L and γ_R diverge (in \mathcal{A}_L and respectively \mathcal{A}_R) and therefore $(\gamma_L, \gamma_R) \uparrow^\infty$. But this implies $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, which yields a contradiction. As our assumption $(\gamma_L, \gamma_R) \not\rightarrow^{*,\infty} \llbracket E \rrbracket$ led to a contradiction, it follows that it must hold and therefore $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, which is what we had to show.

We are now ready to show the main result of this section, that the proof system is sound.

Theorem 2 (Soundness).

For any set of mutual matching logic formulae E , for any mutual matching logic formula φ , if the sequent $\vdash \varphi \Downarrow^\infty E$ is derivable using the proof system given in Figure 4, then $\models \varphi \Downarrow^\infty E$.

Proof. By induction on the proof tree and case analysis on the last rule applied:

1. if the last rule to be applied is **Axiom**, we have that $\varphi \in E$ and we show that $\models \varphi \Downarrow^\infty E$. This is immediate by reflexivity.
2. if the last rule to be applied is **Step**, we have by the induction hypothesis that there exists φ' such that $\mathcal{A}_L \models pr_L(\varphi) \rightarrow^* pr_L(\varphi')$, that $\mathcal{A}_R \models pr_R(\varphi) \rightarrow^* pr_R(\varphi')$ and that $\models \varphi' \Downarrow^\infty E$. We show that $\models \varphi \Downarrow^\infty E$.

Let γ_L, γ_R be arbitrary configurations such that $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$. By Lemma 1, it is sufficient to show that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.

As $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$, $\mathcal{A}_L \models pr_L(\varphi) \rightarrow^* pr_L(\varphi')$ and $\mathcal{A}_R \models pr_R(\varphi) \rightarrow^* pr_R(\varphi')$, it follows that there exist γ'_L, γ'_R such that $\gamma_R \rightarrow_{\mathcal{A}_R}^* \gamma'_R$, $\gamma_L \rightarrow_{\mathcal{A}_L}^* \gamma'_L$ and $(\gamma'_L, \gamma'_R) \models \varphi'$.

But $E \models \varphi'$ and therefore, by Lemma 1, $(\gamma'_L, \gamma'_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$. But $\gamma_L \rightarrow_{\mathcal{A}_L}^* \gamma'_L$ and $\gamma_R \rightarrow_{\mathcal{A}_R}^* \gamma'_R$ and therefore $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, which is what we had to prove.

3. if the last rule to be applied is **Consequence**, we have by the induction hypothesis that there exists φ' such that $\models \varphi \rightarrow \varphi'$ and $\models \varphi' \Downarrow^\infty E$. We show that $\models \varphi \Downarrow^\infty E$.

Let γ_L, γ_R be arbitrary configurations such that $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$. By Lemma 1, it is sufficient to show that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.

As $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$ and $\models \varphi \rightarrow \varphi'$, we have that $(\gamma_L, \gamma_R) \in \llbracket \varphi' \rrbracket$. But we have that $\models \varphi' \Downarrow^\infty E$, and therefore, by Lemma 1, $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, which is what we had to show.

4. if the last rule to be applied is **Case Analysis**, we have by the induction hypothesis that $\models \varphi \Downarrow^\infty E$ and that $\models \varphi' \Downarrow^\infty E$. We show that $\models \varphi \vee \varphi' \Downarrow^\infty E$. Let γ_L, γ_R be arbitrary configurations such that $(\gamma_L, \gamma_R) \in \llbracket \varphi \vee \varphi' \rrbracket$. By Lemma 1, it is sufficient to show that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.

As $(\gamma_L, \gamma_R) \in \llbracket \varphi \vee \varphi' \rrbracket$, we have that $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$ or that $(\gamma_L, \gamma_R) \in \llbracket \varphi' \rrbracket$:

- (a) in the first case, we have that $(\gamma_L, \gamma_R) \in \llbracket \varphi \rrbracket$. We also have that $\models \varphi \Downarrow^\infty E$ (by the induction hypothesis). By Lemma 1, we obtain that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.
- (b) in the second case, we have that $(\gamma_L, \gamma_R) \in \llbracket \varphi' \rrbracket$. We also have that $\models \varphi' \Downarrow^\infty E$ (by the induction hypothesis). By Lemma 1, we obtain that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$.

In either case we have obtained that $(\gamma_L, \gamma_R) \rightarrow^{*,\infty} \llbracket E \rrbracket$, which is what we had to prove.

5. if the last rule to be applied is **Circularity** then, by the induction hypothesis, there exists φ' such that $\mathcal{A}_L \models pr_L(\varphi) \rightarrow^+ pr_L(\varphi')$, $\mathcal{A}_R \models pr_R(\varphi) \rightarrow^+ pr_R(\varphi')$ and $\models \varphi' \Downarrow^\infty E \cup \{\varphi\}$. We show that $\models \varphi \Downarrow^\infty E$.

We let $F = \{\varphi\}$ and we apply the Circularity Principle (Lemma 2). Let γ_L, γ_R be two arbitrary configurations and let ρ be an arbitrary valuation

such that $(\langle \gamma_L, \gamma_R \rangle, \rho) \models \varphi$. As $\mathcal{A}_C \models \pi_C(\varphi) \rightarrow^+ \pi_C(\varphi')$ (for $C \in \{L, R\}$), there exists γ'_C such that $\gamma_C \rightarrow_{\mathcal{A}_C}^+ \gamma'_C$ and $(\gamma'_C, \rho) \models \pi_C(\varphi')$ (for $C \in \{L, R\}$). It follows that $(\langle \gamma'_L, \gamma'_R \rangle, \rho) \models \varphi'$. Since $\models \varphi' \Downarrow^\infty E \cup \{\varphi\}$, we get

$$(\gamma'_L, \gamma'_R) \rightarrow^{*, \infty} \llbracket E \cup \{\varphi\} \rrbracket,$$

by Lemma 1. Now the hypotheses of the Circularity Principle are satisfied and hence obtain $\models \varphi \Downarrow^\infty E$.