Modal Logic Characterizations of Forward, Reverse, and Forward-Reverse Bisimilarities

Marco Bernardo   Andrea Esposito
Dipartimento di Scienze Pure e Applicate, Università di Urbino, Urbino, Italy

Reversible systems feature both forward computations and backward computations, where the latter undo the effects of the former in a causally consistent manner. The compositionality properties and equational characterizations of strong and weak variants of forward-reverse bisimilarity as well as of its two components, i.e., forward bisimilarity and reverse bisimilarity, have been investigated on a minimal process calculus for nondeterministic reversible systems that are sequential, so as to be neutral with respect to interleaving vs. truly concurrent semantics of parallel composition. In this paper we provide logical characterizations for the considered bisimilarities based on forward and backward modalities, which reveals that strong and weak reverse bisimilarities respectively correspond to strong and weak reverse trace equivalences. Moreover, we establish a clear connection between weak forward-reverse bisimilarity and branching bisimilarity, so that the former inherits two further logical characterizations from the latter over a specific class of processes.

1 Introduction

Reversibility in computing started to gain attention since the seminal works [13, 2], where it was shown that reversible computations may achieve low levels of heat dissipation. Nowadays *reversible computing* has many applications ranging from computational biochemistry and parallel discrete-event simulation to robotics, control theory, fault tolerant systems, and concurrent program debugging.

In a reversible system, two directions of computation can be observed: a *forward* one, coinciding with the normal way of computing, and a *backward* one, along which the effects of the forward one are undone when needed in a *causally consistent* way, i.e., by returning to a past consistent state. The latter task is not easy to accomplish in a concurrent system, because the undo procedure necessarily starts from the last performed action and this may not be unique. The usually adopted strategy is that an action can be undone provided that all of its consequences, if any, have been undone beforehand [7].

In the process algebra literature, two approaches have been developed to reverse computations based on keeping track of past actions: the dynamic one of [7] and the static one of [18], later shown to be equivalent in terms of labeled transition systems isomorphism [14].

The former yields RCCS, a variant of CCS [16] that uses stack-based memories attached to processes to record all the actions executed by those processes. A single transition relation is defined, while actions are divided into forward and backward resulting in forward and backward transitions. This approach is suitable when the operational semantics is given in terms of reduction semantics, like in the case of very expressive calculi as well as programming languages.

In contrast, the latter proposes a general method, of which CCSK is a result, to reverse calculi, relying on the idea of retaining within the process syntax all executed actions, which are suitably decorated, and all dynamic operators, which are thus made static. A forward transition relation and a backward transition relation are separately defined, which are labeled with actions extended with communication keys so as to remember who synchronized with whom when going backward. This approach is very handy when it comes to deal with labeled transition systems and basic process calculi.

A. Achilleos and D. Della Monica (Eds.): Fourteenth International Symposium on Games, Automata, Logics, and Formal Verification (GandALF 2023).
EPTCS 390, 2023, pp. 67–81, doi:10.4204/EPTCS.390.5

In [18] *forward-reverse bisimilarity* was introduced too. Unlike standard forward-only bisimilarity [17, 16], it is truly concurrent as it does not satisfy the expansion law of parallel composition into a choice among all possible action sequencings. The interleaving view can be restored in a reversible setting by employing *back-and-forth bisimilarity* [8]. This is defined on computation paths instead of states, thus preserving not only causality but also history as backward moves are constrained to take place along the path followed when going forward even in the presence of concurrency. In the latter setting, a single transition relation is considered, which is viewed as bidirectional, and in the bisimulation game the distinction between going forward or backward is made by matching outgoing or incoming transitions of the considered processes, respectively.

In [4] forward-reverse bisimilarity and its two components, i.e., forward bisimilarity and reverse bisimilarity, have been investigated in terms of compositionality properties and equational characterizations, both for nondeterministic processes and Markovian processes. In order to remain neutral with respect to interleaving view vs. true concurrency, the study has been conducted over a sequential processes calculus, in which parallel composition is not admitted so that not even the communication keys of [18] are needed. Furthermore, like in [8] a single transition relation has been defined and the distinction between outgoing and incoming transitions has been exploited in the bisimulation game. In [3] the investigation of compositionality and axiomatizations has been extended to weak variants of forward, reverse, and forward-reverse bisimilarities, i.e., variants that are capable of abstracting from unobservable actions, in the case of nondeterministic processes only.

In this paper we address the logical characterization of the aforementioned strong and weak bisimilarities over nondeterministic reversibile sequential processes. The objective is to single out suitable modal logics that induce equivalences that turn out to be alternative characterizations of the considered bisimilarities, so that two processes are bisimilar iff they satisfy the same set of formulas of the corresponding logic. Starting from Hennessy-Milner logic [11], which includes forward modalities whereby it is possible to characterize the standard forward-only strong and weak bisimilarities of [16], the idea is to add backward modalities in the spirit of [8] so as to be able to characterize reverse and forward-reverse strong and weak bisimilarities. Unlike [8], where back-and-forth bisimilarities as well as modality interpretations are defined over computation paths, in our reversible setting both the considered bisimilarities and the associated modal logic interpretations are defined over states.

Our study reveals that strong and weak reverse bisimilarities do not need conjunction in their logical characterizations. In other words, they boil down to strong and weak reverse trace equivalences, respectively. Moreover, recalling that branching bisimilarity [10] is known to coincide with weak back-and-forth bisimilarity defined over computation paths [8], we show that branching bisimilarity also coincides for a specific class of processes with our weak forward-reverse bisimilarity defined over states. Based on the results in [9], this opens the way to two further logical characterizations of the latter in addition to the one based on forward and backward modalities. The first characterization replaces the aforementioned modalities with an until operator, whilst the second one is given by the temporal logic CTL* without the next operator.

The paper is organized as follows. In Section 2 we recall syntax and semantics for the considered calculus of nondeterministic reversible sequential processes as well as the strong forward, reverse, and forward-reverse bisimilarities investigated in [4] and their weak counterparts examined in [3]. In Section 3 we provide the modal logic characterizations of all the aforementioned bisimilarities based on forward and backward modalities interpreted over states. In Section 4 we establish a clear connection between branching bisimilarity and our weak forward-reverse bisimilarity defined over states. In Section 5 we conclude with final remarks and directions for future work.

2 Background

2.1 Syntax of Nondeterministic Reversible Sequential Processes

Given a countable set A of actions – ranged over by a,b,c – including an unobservable action denoted by τ, the syntax of reversible sequential processes is defined as follows [4]:

$$P \ ::= \ \underline{0} \mid a\,.\,P \mid a^{\dagger}\,.\,P \mid P + P$$

where:

- $\underline{0}$ is the terminated process.
- $a\,.\,P$ is a process that can execute action a and whose forward continuation is P.
- $a^{\dagger}\,.\,P$ is a process that executed action a and whose forward continuation is inside P.
- $P_1 + P_2$ expresses a nondeterministic choice between P_1 and P_2 as far as both of them have not executed any action yet, otherwise only the one that was selected in the past can move.

We syntactically characterize through suitable predicates three classes of processes generated by the grammar above. Firstly, we have *initial* processes, i.e., processes in which all the actions are unexecuted:

$$\begin{array}{rcl}
\mathit{initial}(\underline{0}) & & \\
\mathit{initial}(a\,.\,P) & \Longleftarrow & \mathit{initial}(P) \\
\mathit{initial}(P_1 + P_2) & \Longleftarrow & \mathit{initial}(P_1) \wedge \mathit{initial}(P_2)
\end{array}$$

Secondly, we have *final* processes, i.e., processes in which all the actions along a single path have been executed:

$$\begin{array}{rcl}
\mathit{final}(\underline{0}) & & \\
\mathit{final}(a^{\dagger}\,.\,P) & \Longleftarrow & \mathit{final}(P) \\
\mathit{final}(P_1 + P_2) & \Longleftarrow & (\mathit{final}(P_1) \wedge \mathit{initial}(P_2)) \vee \\
 & & (\mathit{initial}(P_1) \wedge \mathit{final}(P_2))
\end{array}$$

Multiple paths arise only in the presence of alternative compositions. At each occurrence of $+$, only the subprocess chosen for execution can move, while the other one, although not selected, is kept as an initial subprocess within the overall process to support reversibility.

Thirdly, we have the processes that are *reachable* from an initial one, whose set we denote by $\mathbb{P}$:

$$\begin{array}{rcl}
\mathit{reachable}(\underline{0}) & & \\
\mathit{reachable}(a\,.\,P) & \Longleftarrow & \mathit{initial}(P) \\
\mathit{reachable}(a^{\dagger}\,.\,P) & \Longleftarrow & \mathit{reachable}(P) \\
\mathit{reachable}(P_1 + P_2) & \Longleftarrow & (\mathit{reachable}(P_1) \wedge \mathit{initial}(P_2)) \vee \\
 & & (\mathit{initial}(P_1) \wedge \mathit{reachable}(P_2))
\end{array}$$

It is worth noting that:

- $\underline{0}$ is the only process that is both initial and final as well as reachable.
- Any initial or final process is reachable too.
- $\mathbb{P}$ also contains processes that are neither initial nor final, like e.g. $a^{\dagger}\,.\,b\,.\,\underline{0}$.
- The relative positions of already executed actions and actions to be executed matter; in particular, an action of the former kind can never follow one of the latter kind. For instance, $a^{\dagger}\,.\,b\,.\,\underline{0} \in \mathbb{P}$ whereas $b\,.\,a^{\dagger}\,.\,\underline{0} \notin \mathbb{P}$.

$$(\textsc{Act}_{\mathrm{f}})\ \frac{\mathit{initial}(P)}{a\,.\,P \xrightarrow{a} a^{\dagger}\,.\,P} \qquad (\textsc{Act}_{\mathrm{p}})\ \frac{P \xrightarrow{b} P'}{a^{\dagger}\,.\,P \xrightarrow{b} a^{\dagger}\,.\,P'}$$

$$(\textsc{Cho}_{\mathrm{l}})\ \frac{P_1 \xrightarrow{a} P_1' \quad \mathit{initial}(P_2)}{P_1 + P_2 \xrightarrow{a} P_1' + P_2} \qquad (\textsc{Cho}_{\mathrm{r}})\ \frac{P_2 \xrightarrow{a} P_2' \quad \mathit{initial}(P_1)}{P_1 + P_2 \xrightarrow{a} P_1 + P_2'}$$

Table 1: Operational semantic rules for reversible action prefix and choice

2.2 Operational Semantic Rules

According to the approach of [18], dynamic operators such as action prefix and alternative composition have to be made static by the semantics, so as to retain within the syntax all the information needed to enable reversibility. For the sake of minimality, unlike [18] we do not generate two distinct transition relations – a forward one $\longrightarrow$ and a backward one $\rightsquigarrow$ – but a single transition relation, which we implicitly regard as being symmetric like in [8] to enforce the *loop property*: every executed action can be undone and every undone action can be redone.

In our setting, a backward transition from P' to P ($P' \overset{a}{\rightsquigarrow} P$) is subsumed by the corresponding forward transition t from P to P' ($P \xrightarrow{a} P'$). As will become clear with the definition of behavioral equivalences, like in [8] when going forward we view t as an *outgoing* transition of P, while when going backward we view t as an *incoming* transition of P'. The semantic rules for $\longrightarrow \subseteq \mathbb{P} \times A \times \mathbb{P}$ are defined in Table 1 and generate the labeled transition system $(\mathbb{P}, A, \longrightarrow)$ [4].

The first rule for action prefix ($\textsc{Act}_{\mathrm{f}}$ where f stands for forward) applies only if P is initial and retains the executed action in the target process of the generated forward transition by decorating the action itself with $\dagger$. The second rule for action prefix ($\textsc{Act}_{\mathrm{p}}$ where p stands for propagation) propagates actions executed by inner initial subprocesses.

In both rules for alternative composition ($\textsc{Cho}_{\mathrm{l}}$ and $\textsc{Cho}_{\mathrm{r}}$ where l stands for left and r stands for right), the subprocess that has not been selected for execution is retained as an initial subprocess in the target process of the generated transition. When both subprocesses are initial, both rules for alternative composition are applicable, otherwise only one of them can be applied and in that case it is the non-initial subprocess that can move, because the other one has been discarded at the moment of the selection.

Every state corresponding to a non-final process has at least one outgoing transition, while every state corresponding to a non-initial process has exactly one incoming transition due to the decoration of executed actions. The labeled transition system underlying an initial process turns out to be a tree, whose branching points correspond to occurrences of $+$.

Example 2.1 The labeled transition systems generated by the rules in Table 1 for the two initial processes $a\,.\,\underline{0} + a\,.\,\underline{0}$ and $a\,.\,\underline{0}$ are depicted below:

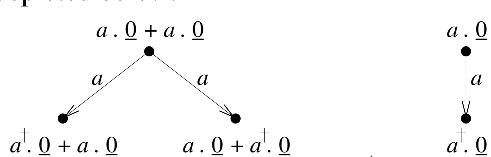

As far as the one on the left is concerned, we observe that, in the case of a standard process calculus, a single a-transition from $a\,.\,\underline{0} + a\,.\,\underline{0}$ to $\underline{0}$ would have been generated due to the absence of action decorations within processes. ■

2.3 Strong Forward, Reverse, and Forward-Reverse Bisimilarities

While forward bisimilarity considers only *outgoing* transitions [17, 16], reverse bisimilarity considers only *incoming* transitions. Forward-reverse bisimilarity [18] considers instead both outgoing transitions and incoming ones. Here are their *strong* versions studied in [4], where strong means not abstracting from τ-actions.

Definition 2.2 We say that $P_1, P_2 \in \mathbb{P}$ are *forward bisimilar*, written $P_1 \sim_{\rm FB} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some forward bisimulation $\mathcal{B}$. A symmetric relation $\mathcal{B}$ over $\mathbb{P}$ is a *forward bisimulation* iff for all $(P_1, P_2) \in \mathcal{B}$ and $a \in A$:

- Whenever $P_1 \xrightarrow{a} P'_1$, then $P_2 \xrightarrow{a} P'_2$ with $(P'_1, P'_2) \in \mathcal{B}$. ■

Definition 2.3 We say that $P_1, P_2 \in \mathbb{P}$ are *reverse bisimilar*, written $P_1 \sim_{\rm RB} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some reverse bisimulation $\mathcal{B}$. A symmetric relation $\mathcal{B}$ over $\mathbb{P}$ is a *reverse bisimulation* iff for all $(P_1, P_2) \in \mathcal{B}$ and $a \in A$:

- Whenever $P'_1 \xrightarrow{a} P_1$, then $P'_2 \xrightarrow{a} P_2$ with $(P'_1, P'_2) \in \mathcal{B}$. ■

Definition 2.4 We say that $P_1, P_2 \in \mathbb{P}$ are *forward-reverse bisimilar*, written $P_1 \sim_{\rm FRB} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some forward-reverse bisimulation $\mathcal{B}$. A symmetric relation $\mathcal{B}$ over $\mathbb{P}$ is a *forward-reverse bisimulation* iff for all $(P_1, P_2) \in \mathcal{B}$ and $a \in A$:

- Whenever $P_1 \xrightarrow{a} P'_1$, then $P_2 \xrightarrow{a} P'_2$ with $(P'_1, P'_2) \in \mathcal{B}$.
- Whenever $P'_1 \xrightarrow{a} P_1$, then $P'_2 \xrightarrow{a} P_2$ with $(P'_1, P'_2) \in \mathcal{B}$. ■

$\sim_{\rm FRB} \subsetneq \sim_{\rm FB} \cap \sim_{\rm RB}$ with the inclusion being strict because, e.g., the two final processes $a^{\dagger} . \underline{0}$ and $a^{\dagger} . \underline{0} + c . \underline{0}$ are identified by $\sim_{\rm FB}$ (no outgoing transitions on both sides) and by $\sim_{\rm RB}$ (only an incoming a-transition on both sides), but distinguished by $\sim_{\rm FRB}$ as in the latter process action c is enabled again after undoing a (and hence there is an outgoing c-transition in addition to an outgoing a-transition). Moreover, $\sim_{\rm FB}$ and $\sim_{\rm RB}$ are incomparable because for instance:

$$a^{\dagger} . \underline{0} \sim_{\rm FB} \underline{0} \quad \text{but} \quad a^{\dagger} . \underline{0} \not\sim_{\rm RB} \underline{0}$$
$$a . \underline{0} \sim_{\rm RB} \underline{0} \quad \text{but} \quad a . \underline{0} \not\sim_{\rm FB} \underline{0}$$

Note that that $\sim_{\rm FRB} = \sim_{\rm FB}$ over initial processes, with $\sim_{\rm RB}$ strictly coarser, whilst $\sim_{\rm FRB} \neq \sim_{\rm RB}$ over final processes because, after going backward, previously discarded subprocesses come into play again in the forward direction.

Example 2.5 The two processes considered in Example 2.1 are identified by all the three equivalences. This is witnessed by any bisimulation that contains the pairs $(a . \underline{0} + a . \underline{0}, a . \underline{0})$, $(a^{\dagger} . \underline{0} + a . \underline{0}, a^{\dagger} . \underline{0})$, and $(a . \underline{0} + a^{\dagger} . \underline{0}, a^{\dagger} . \underline{0})$. ■

As observed in [4], it makes sense that $\sim_{\rm FB}$ identifies processes with a different past and that $\sim_{\rm RB}$ identifies processes with a different future, in particular with $\underline{0}$ that has neither past nor future. However, for $\sim_{\rm FB}$ this breaks compositionality with respect to alternative composition. As an example:

$$a^{\dagger} . b . \underline{0} \quad \sim_{\rm FB} \quad b . \underline{0}$$
$$a^{\dagger} . b . \underline{0} + c . \underline{0} \quad \not\sim_{\rm FB} \quad b . \underline{0} + c . \underline{0}$$

because in $a^{\dagger} . b . \underline{0} + c . \underline{0}$ action c is disabled due to the presence of the already executed action $a^{\dagger}$, while in $b . \underline{0} + c . \underline{0}$ action c is enabled as there are no past actions preventing it from occurring. Note that a similar phenomenon does not happen with $\sim_{\rm RB}$ as $a^{\dagger} . b . \underline{0} \not\sim_{\rm RB} b . \underline{0}$ due to the incoming a-transition of $a^{\dagger} . b . \underline{0}$.

This problem, which does not show up for $\sim_{\rm RB}$ and $\sim_{\rm FRB}$ because these two equivalences cannot identify an initial process with a non-initial one, leads to the following variant of $\sim_{\rm FB}$ that is sensitive to the presence of the past.

Definition 2.6 We say that $P_1, P_2 \in \mathbb{P}$ are *past-sensitive forward bisimilar*, written $P_1 \sim_{\mathrm{FB:ps}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some past-sensitive forward bisimulation $\mathcal{B}$. A relation $\mathcal{B}$ over $\mathbb{P}$ is a *past-sensitive forward bisimulation* iff it is a forward bisimulation such that $initial(P_1) \Longleftrightarrow initial(P_2)$ for all $(P_1, P_2) \in \mathcal{B}$. ■

Now $\sim_{\mathrm{FB:ps}}$ is sensitive to the presence of the past:

$$a^{\dagger} . b . \underline{0} \not\sim_{\mathrm{FB:ps}} b . \underline{0}$$

but can still identify non-initial processes having a different past:

$$a_1^{\dagger} . P \sim_{\mathrm{FB:ps}} a_2^{\dagger} . P$$

It holds that $\sim_{\mathrm{FRB}} \subsetneq \sim_{\mathrm{FB:ps}} \cap \sim_{\mathrm{RB}}$, with $\sim_{\mathrm{FRB}} = \sim_{\mathrm{FB:ps}}$ over initial processes as well as $\sim_{\mathrm{FB:ps}}$ and $\sim_{\mathrm{RB}}$ being incomparable because, e.g., for $a_1 \neq a_2$:

$$a_1^{\dagger} . P \sim_{\mathrm{FB:ps}} a_2^{\dagger} . P \quad \text{but} \quad a_1^{\dagger} . P \not\sim_{\mathrm{RB}} a_2^{\dagger} . P$$
$$a_1 . P \sim_{\mathrm{RB}} a_2 . P \quad \text{but} \quad a_1 . P \not\sim_{\mathrm{FB:ps}} a_2 . P$$

In [4] it has been shown that all the considered strong bisimilarities are congruences with respect to action prefix, while only $\sim_{\mathrm{FB:ps}}$, $\sim_{\mathrm{RB}}$, and $\sim_{\mathrm{FRB}}$ are congruences with respect to alternative composition too, with $\sim_{\mathrm{FB:ps}}$ being the coarsest congruence with respect to $+$ contained in $\sim_{\mathrm{FB}}$. Moreover, sound and complete equational characterizations have been provided for the three congruences.

2.4 Weak Forward, Reverse, and Forward-Reverse Bisimilarities

In [3] *weak* variants of forward, reverse, and forward-reverse bisimilarities have been studied, which are capable of abstracting from τ-actions. In the following definitions, $P \stackrel{\tau^*}{\Longrightarrow} P'$ means that $P' = P$ or there exists a nonempty sequence of finitely many τ-transitions such that the target of each of them coincides with the source of the subsequent one, with the source of the first one being P and the target of the last one being P'. Moreover, $\stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow}$ stands for an a-transition possibly preceded and followed by finitely many τ-transitions. We further let $\bar{A} = A \setminus \{\tau\}$.

Definition 2.7 We say that $P_1, P_2 \in \mathbb{P}$ are *weakly forward bisimilar*, written $P_1 \approx_{\mathrm{FB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some weak forward bisimulation $\mathcal{B}$. A symmetric binary relation $\mathcal{B}$ over $\mathbb{P}$ is a *weak forward bisimulation* iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:

- Whenever $P_1 \xrightarrow{\tau} P_1'$, then $P_2 \stackrel{\tau^*}{\Longrightarrow} P_2'$ and $(P_1', P_2') \in \mathcal{B}$.
- Whenever $P_1 \xrightarrow{a} P_1'$ for $a \in \bar{A}$, then $P_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2'$ and $(P_1', P_2') \in \mathcal{B}$. ■

Definition 2.8 We say that $P_1, P_2 \in \mathbb{P}$ are *weakly reverse bisimilar*, written $P_1 \approx_{\mathrm{RB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some weak reverse bisimulation $\mathcal{B}$. A symmetric binary relation $\mathcal{B}$ over $\mathbb{P}$ is a *weak reverse bisimulation* iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:

- Whenever $P_1' \xrightarrow{\tau} P_1$, then $P_2' \stackrel{\tau^*}{\Longrightarrow} P_2$ and $(P_1', P_2') \in \mathcal{B}$.
- Whenever $P_1' \xrightarrow{a} P_1$ for $a \in \bar{A}$, then $P_2' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2$ and $(P_1', P_2') \in \mathcal{B}$. ■

Definition 2.9 We say that $P_1, P_2 \in \mathbb{P}$ are *weakly forward-reverse bisimilar*, written $P_1 \approx_{\mathrm{FRB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some weak forward-reverse bisimulation $\mathcal{B}$. A symmetric binary relation $\mathcal{B}$ over $\mathbb{P}$ is a *weak forward-reverse bisimulation* iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:

- Whenever $P_1 \xrightarrow{\tau} P_1'$, then $P_2 \stackrel{\tau^*}{\Longrightarrow} P_2'$ and $(P_1', P_2') \in \mathcal{B}$.

- Whenever $P_1 \xrightarrow{a} P'_1$ for $a \in \bar{A}$, then $P_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P'_2$ and $(P'_1, P'_2) \in \mathscr{B}$.
- Whenever $P'_1 \xrightarrow{\tau} P_1$, then $P'_2 \stackrel{\tau^*}{\Longrightarrow} P_2$ and $(P'_1, P'_2) \in \mathscr{B}$.
- Whenever $P'_1 \xrightarrow{a} P_1$ for $a \in \bar{A}$, then $P'_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2$ and $(P'_1, P'_2) \in \mathscr{B}$. ■

Each of the three weak bisimilarities is strictly coarser than the corresponding strong one. Similar to the strong case, $\approx_{\mathrm{FRB}} \subsetneq \approx_{\mathrm{FB}} \cap \approx_{\mathrm{RB}}$ with $\approx_{\mathrm{FB}}$ and $\approx_{\mathrm{RB}}$ being incomparable. Unlike the strong case, $\approx_{\mathrm{FRB}} \neq \approx_{\mathrm{FB}}$ over initial processes. For instance, $\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0} + b\,.\,\underline{0}$ and $\tau\,.\,a\,.\,\underline{0} + b\,.\,\underline{0}$ are identified by $\approx_{\mathrm{FB}}$ but told apart by $\approx_{\mathrm{FRB}}$: if the former performs a, the latter responds with τ followed by a and if it subsequently undoes a thus becoming $\tau^\dagger.\,a\,.\,\underline{0} + b\,.\,\underline{0}$ in which only a is enabled, the latter can only respond by undoing a thus becoming $\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0} + b\,.\,\underline{0}$ in which both a and b are enabled. An analogous counterexample with non-initial τ-actions is given by $c\,.\,(\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0} + b\,.\,\underline{0})$ and $c\,.\,(\tau\,.\,a\,.\,\underline{0} + b\,.\,\underline{0})$.

As observed in [3], $\approx_{\mathrm{FB}}$ suffers from the same compositionality problem with respect to alternative composition as $\sim_{\mathrm{FB}}$. Moreover, $\approx_{\mathrm{FB}}$ and $\approx_{\mathrm{FRB}}$ feature the same compositionality problem as weak bisimilarity for standard forward-only processes [16], i.e., for $\approx \in \{\approx_{\mathrm{FB}}, \approx_{\mathrm{FRB}}\}$ it holds that:

$$\begin{array}{rcl} \tau\,.\,a\,.\,\underline{0} & \approx & a\,.\,\underline{0} \\ \tau\,.\,a\,.\,\underline{0} + b\,.\,\underline{0} & \not\approx & a\,.\,\underline{0} + b\,.\,\underline{0} \end{array}$$

because if $\tau\,.\,a\,.\,\underline{0} + b\,.\,\underline{0}$ performs τ thereby evolving to $\tau^\dagger.\,a\,.\,\underline{0} + b\,.\,\underline{0}$ where only a is enabled in the forward direction, then $a\,.\,\underline{0} + b\,.\,\underline{0}$ can neither move nor idle in the attempt to evolve in such a way to match $\tau^\dagger.\,a\,.\,\underline{0} + b\,.\,\underline{0}$.

To solve both problems it is sufficient to redefine the two equivalences by making them sensitive to the presence of the past, exactly as in the strong case for forward bisimilarity. By so doing, $\tau\,.\,a\,.\,\underline{0}$ is no longer identified with $a\,.\,\underline{0}$: if the former performs τ thereby evolving to $\tau^\dagger.\,a\,.\,\underline{0}$ and the latter idles, then $\tau^\dagger.\,a\,.\,\underline{0}$ and $a\,.\,\underline{0}$ are told apart because they are not both initial or non-initial.

Definition 2.10 We say that $P_1, P_2 \in \mathbb{P}$ are *weakly past-sensitive forward bisimilar*, written $P_1 \approx_{\mathrm{FB:ps}} P_2$, iff $(P_1, P_2) \in \mathscr{B}$ for some weak past-sensitive forward bisimulation $\mathscr{B}$. A binary relation $\mathscr{B}$ over $\mathbb{P}$ is a *weak past-sensitive forward bisimulation* iff it is a weak forward bisimulation such that *initial*$(P_1) \Longleftrightarrow$ *initial*(P_2) for all $(P_1, P_2) \in \mathscr{B}$. ■

Definition 2.11 We say that $P_1, P_2 \in \mathbb{P}$ are *weakly past-sensitive forward-reverse bisimilar*, written $P_1 \approx_{\mathrm{FRB:ps}} P_2$, iff $(P_1, P_2) \in \mathscr{B}$ for some weak past-sensitive forward-reverse bisimulation $\mathscr{B}$. A binary relation $\mathscr{B}$ over $\mathbb{P}$ is a *weak past-sensitive forward-reverse bisimulation* iff it is a weak forward-reverse bisimulation such that *initial*$(P_1) \Longleftrightarrow$ *initial*(P_2) for all $(P_1, P_2) \in \mathscr{B}$. ■

Like in the non-past-sensitive case, $\approx_{\mathrm{FRB:ps}} \neq \approx_{\mathrm{FB:ps}}$ over initial processes, as shown by $\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0}$ and $\tau\,.\,a\,.\,\underline{0}$: if the former performs a, the latter responds with τ followed by a and if it subsequently undoes a thus becoming the non-initial process $\tau^\dagger.\,a\,.\,\underline{0}$, the latter can only respond by undoing a thus becoming the initial process $\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0}$. An analogous counterexample with non-initial τ-actions is given again by $c\,.\,(\tau\,.\,a\,.\,\underline{0} + a\,.\,\underline{0} + b\,.\,\underline{0})$ and $c\,.\,(\tau\,.\,a\,.\,\underline{0} + b\,.\,\underline{0})$.

Observing that $\sim_{\mathrm{FRB}} \subsetneq \approx_{\mathrm{FRB:ps}}$ as the former naturally satisfies the initiality condition, in [3] it has been shown that all the considered weak bisimilarities are congruences with respect to action prefix, while only $\approx_{\mathrm{FB:ps}}$, $\approx_{\mathrm{RB}}$, and $\approx_{\mathrm{FRB:ps}}$ are congruences with respect to alternative composition too, with $\approx_{\mathrm{FB:ps}}$ and $\approx_{\mathrm{FRB:ps}}$ respectively being the coarsest congruences with respect to $+$ contained in $\approx_{\mathrm{FB}}$ and $\approx_{\mathrm{FRB}}$. Sound and complete equational characterizations have been provided for the three congruences.

3 Modal Logic Characterizations

In this section we investigate modal logic characterizations for the three strong bisimilarities $\sim_{\rm FB}$, $\sim_{\rm RB}$, and $\sim_{\rm FRB}$, the three weak bisimilarities $\approx_{\rm FB}$, $\approx_{\rm RB}$, and $\approx_{\rm FRB}$, and the three past-sensitive variants $\sim_{\rm FB:ps}$, $\approx_{\rm FB:ps}$, and $\approx_{\rm FRB:ps}$.

We start by introducing a general modal logic $\mathscr{L}$ from which we will take nine fragments to characterize the nine aforementioned bisimilarities. It consists of Hennessy-Milner logic [11] extended with the proposition init, the strong backward modality $\langle a^\dagger \rangle$, the two weak forward modalities $\langle\langle \tau \rangle\rangle$ and $\langle\langle a \rangle\rangle$, and the two weak backward modalities $\langle\langle \tau^\dagger \rangle\rangle$ and $\langle\langle a^\dagger \rangle\rangle$ (where $a \in \bar{A}$ within weak modalities):

$$\phi \ ::= \ \text{true} \mid \text{init} \mid \neg\phi \mid \phi \wedge \phi \mid \langle a \rangle \phi \mid \langle a^\dagger \rangle \phi \mid \langle\langle \tau \rangle\rangle \phi \mid \langle\langle a \rangle\rangle \phi \mid \langle\langle \tau^\dagger \rangle\rangle \phi \mid \langle\langle a^\dagger \rangle\rangle \phi$$

The satisfaction relation $\models \subseteq \mathbb{P} \times \mathscr{L}$ is defined by induction on the syntactical structure of the formulas as follows:

$$\begin{array}{lcll}
P & \models & \text{true} & \text{for all } P \in \mathbb{P} \\
P & \models & \text{init} & \text{iff } \mathit{initial}(P) \\
P & \models & \neg\phi & \text{iff } P \not\models \phi \\
P & \models & \phi_1 \wedge \phi_2 & \text{iff } P \models \phi_1 \text{ and } P \models \phi_2 \\
P & \models & \langle a \rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P \xrightarrow{a} P' \text{ and } P' \models \phi \\
P & \models & \langle a^\dagger \rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P' \xrightarrow{a} P \text{ and } P' \models \phi \\
P & \models & \langle\langle \tau \rangle\rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P \stackrel{\tau^*}{\Longrightarrow} P' \text{ and } P' \models \phi \\
P & \models & \langle\langle a \rangle\rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P' \text{ and } P' \models \phi \\
P & \models & \langle\langle \tau^\dagger \rangle\rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P' \stackrel{\tau^*}{\Longrightarrow} P \text{ and } P' \models \phi \\
P & \models & \langle\langle a^\dagger \rangle\rangle \phi & \text{iff there exists } P' \in \mathbb{P} \text{ such that } P' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P \text{ and } P' \models \phi
\end{array}$$

The use of backward operators is not new in the definition of properties of programs through temporal logics [15] or modal logics [12]. In particular, in the latter work a logic with a past operator was introduced to capture interesting properties of generalized labeled transition systems where only visible actions are considered, in which setting it is proved that the equivalence induced by the considered logic coincides with a generalization of the standard forward-only strong bisimilarity of [16]. This result was later confirmed in [9] where it is shown that the addition of a strong backward modality (interpreted over computation paths instead of states) provides no additional discriminating power with respect to the Hennessy-Milner logic, i.e., the induced equivalence is again strong bisimilarity.

In contrast, in our context – in which all equivalences are defined over states – the strong forward bisimilarities $\sim_{\rm FB}$ and $\sim_{\rm FB:ps}$ do not coincide with the strong forward-reverse bisimilarity $\sim_{\rm FRB}$ and this extends to their weak counterparts. In other words, the presence of backward modalities matters. It is worth noting that our two weak backward modalities are similar to the ones considered in [8, 9] to characterize weak back-and-forth bisimilarity (defined over computation paths), which is finer than the standard forward-only weak bisimilarity of [16] and coincides with branching bisimilarity [10].

By taking suitable fragments of $\mathscr{L}$ we can characterize all the nine bisimilarities introduced in Section 2. For each of the four strong bisimilarities $\sim_B$, where $B \in \{\text{FB}, \text{FB:ps}, \text{RB}, \text{FRB}\}$, we can define the corresponding logic $\mathscr{L}_B$. The same can be done for each of the five weak bisimilarities $\approx_B$, where $B \in \{\text{FB}, \text{FB:ps}, \text{RB}, \text{FRB}, \text{FRB:ps}\}$, to obtain the corresponding logic $\mathscr{L}_B^\tau$. All the considered fragments can be found in Table 2, which indicates that the proposition init is needed only for the past-sensitive bisimilarities. The forthcoming Theorems 3.1 and 3.2 show that each such fragment induces the intended bisimilarity, in the sense that two processes are bisimilar iff they satisfy the same set of formulas of the fragment at hand.

	true	init	$\neg$	$\wedge$	$\langle a \rangle$	$\langle a^\dagger \rangle$	$\langle\langle \tau \rangle\rangle$	$\langle\langle a \rangle\rangle$	$\langle\langle \tau^\dagger \rangle\rangle$	$\langle\langle a^\dagger \rangle\rangle$
$\mathscr{L}_{\mathrm{FB}}$	✓		✓	✓	✓					
$\mathscr{L}_{\mathrm{FB:ps}}$	✓	✓	✓	✓	✓					
$\mathscr{L}_{\mathrm{RB}}$	✓					✓				
$\mathscr{L}_{\mathrm{FRB}}$	✓		✓	✓	✓	✓				
$\mathscr{L}^{\tau}_{\mathrm{FB}}$	✓		✓	✓			✓	✓		
$\mathscr{L}^{\tau}_{\mathrm{FB:ps}}$	✓	✓	✓	✓			✓	✓		
$\mathscr{L}^{\tau}_{\mathrm{RB}}$	✓								✓	✓
$\mathscr{L}^{\tau}_{\mathrm{FRB}}$	✓		✓	✓			✓	✓	✓	✓
$\mathscr{L}^{\tau}_{\mathrm{FRB:ps}}$	✓	✓	✓	✓			✓	✓	✓	✓

Table 2: Fragments of $\mathscr{L}$ characterizing the considered bisimilarities

The technique used to prove the two theorems is inspired by the one employed in [1] to show that Hennessy-Milner logic characterizes the strong forward-only bisimilarity of [16]. The two implications of either theorem are demonstrated separately. To prove that any pair of bisimilar processes P_1 and P_2 satisfy the same formulas of the considered fragment, we assume that $P_1 \models \phi$ for some formula ϕ and then we proceed by induction on the depth of ϕ to show that $P_2 \models \phi$ too, where the depth of a formula is defined by induction on the syntactical structure of the formula itself as follows:

$$\begin{aligned}
depth(\mathrm{true}) &= 1 \\
depth(\mathrm{init}) &= 1 \\
depth(\neg\phi) &= 1 + depth(\phi) \\
depth(\phi_1 \wedge \phi_2) &= 1 + \max(depth(\phi_1), depth(\phi_2)) \\
depth(\langle a \rangle \phi) &= 1 + depth(\phi) \\
depth(\langle a^\dagger \rangle \phi) &= 1 + depth(\phi) \\
depth(\langle\langle \tau \rangle\rangle \phi) &= 1 + depth(\phi) \\
depth(\langle\langle a \rangle\rangle \phi) &= 1 + depth(\phi) \\
depth(\langle\langle \tau^\dagger \rangle\rangle \phi) &= 1 + depth(\phi) \\
depth(\langle\langle a^\dagger \rangle\rangle \phi) &= 1 + depth(\phi)
\end{aligned}$$

As for the reverse implication, we show that the relation $\mathscr{B}$ formed by all pairs of processes (P_1, P_2) that satisfy the same formulas of the considered fragment is a bisimulation. More specifically, starting from $(P_1, P_2) \in \mathscr{B}$ we proceed by contradiction by assuming that, whenever P_1 has a move to/from P_1' with an action a, then there is no P_2' such that P_2 has a move to/from P_2' with a and $(P_1', P_2') \in \mathscr{B}$. This entails that, for every P_{2_i} forward/backward reachable from P_2 by performing a, by definition of $\mathscr{B}$ there exists some formula ϕ_i such that $P_1' \models \phi_i$ and $P_{2_i}' \not\models \phi_i$, which leads to a formula with a forward/backward modality on a followed by $\bigwedge_i \phi_i$ that is satisfied by P_1 but not by P_2, thereby contradicting $(P_1, P_2) \in \mathscr{B}$.

Theorem 3.1 Let $P_1, P_2 \in \mathbb{P}$ and $B \in \{\mathrm{FB}, \mathrm{FB{:}ps}, \mathrm{RB}, \mathrm{FRB}\}$. Then $P_1 \sim_B P_2 \iff \forall \phi \in \mathscr{L}_B. P_1 \models \phi \Leftrightarrow P_2 \models \phi$. ■

Theorem 3.2 Let $P_1, P_2 \in \mathbb{P}$ and $B \in \{\mathrm{FB}, \mathrm{FB{:}ps}, \mathrm{RB}, \mathrm{FRB}, \mathrm{FRB{:}ps}\}$. Then $P_1 \approx_B P_2 \iff \forall \phi \in \mathscr{L}^{\tau}_B. P_1 \models \phi \Leftrightarrow P_2 \models \phi$. ■

We conclude with the following observations:

- The fragments that characterize the four forward bisimilarities $\sim_{\mathrm{FB}}$, $\sim_{\mathrm{FB:ps}}$, $\approx_{\mathrm{FB}}$, and $\approx_{\mathrm{FB:ps}}$ are essentially identical to the Hennessy-Milner logic (first two bisimilarities) and its weak variant (last two bisimilarities). The only difference is the possible presence of the proposition init, which is needed to distinguish between initial and non-initial processes in the past-sensitive cases.

- The fragments that characterize the two reverse bisimilarities $\sim_{\mathrm{RB}}$ and $\approx_{\mathrm{RB}}$ only include true and the backward modalities $\langle a^{\dagger}\rangle$ (first bisimilarity) and $\langle\langle \tau^{\dagger}\rangle\rangle$ and $\langle\langle a^{\dagger}\rangle\rangle$ (second bisimilarity). The absence of conjunction reflects the fact that, when going backward, processes must follow exactly the sequence of actions they performed in the forward direction and hence no choice is involved, consistent with every non-initial process having precisely one incoming transition. In other words, the strong and weak reverse bisimilarities boil down to strong and weak reverse trace equivalences, respectively, which consider traces obtained when going in the backward direction.
- The fragments that characterize the three forward-reverse bisimilarities $\sim_{\mathrm{FRB}}$, $\approx_{\mathrm{FRB}}$, and $\approx_{\mathrm{FRB:PS}}$ are akin to the logic $\mathscr{L}_{\mathrm{BF}}$ introduced in [8] to characterize weak back-and-forth bisimilarity and branching bisimilarity. A crucial distinction between our three fragments and $\mathscr{L}_{\mathrm{BF}}$ is that the former are interpreted over states while $\mathscr{L}_{\mathrm{BF}}$ is interpreted over computation paths. Moreover, as already mentioned, defining a strong variant of $\mathscr{L}_{\mathrm{BF}}$ would yield a logic that characterizes strong bisimilarity, whereas in our setting forward-only bisimilarities are different from forward-reverse ones and hence different logics are needed.

4 Weak Forward-Reverse Bisimilarity and Branching Bisimilarity

In this section we establish a clear connection between weak forward-reverse bisimilarity and branching bisimilarity [10]. Unlike the standard forward-only weak bisimilarity of [16], branching bisimilarity preserves the branching structure of processes even when abstracting from τ-actions.

Definition 4.1 We say that $P_1, P_2 \in \mathbb{P}$ are *branching bisimilar*, written $P_1 \approx_{\mathrm{BB}} P_2$, iff $(P_1, P_2) \in \mathscr{B}$ for some branching bisimulation $\mathscr{B}$. A symmetric binary relation $\mathscr{B}$ over $\mathbb{P}$ is a *branching bisimulation* iff, whenever $(P_1, P_2) \in \mathscr{B}$, then for all $P_1 \xrightarrow{a} P_1'$ it holds that:

- either $a = \tau$ and $(P_1', P_2) \in \mathscr{B}$;
- or $P_2 \stackrel{\tau^*}{\Longrightarrow} \bar{P}_2 \xrightarrow{a} P_2'$ with $(P_1, \bar{P}_2) \in \mathscr{B}$ and $(P_1', P_2') \in \mathscr{B}$. ■

Branching bisimilarity is known to have some relationships with reversibility. More precisely, in [8] strong and weak back-and-forth bisimilarities have been introduced over labeled transition systems – where outgoing transitions are considered in the forward bisimulation game while incoming transitions are considered in the backward bisimulation game – and respectively shown to coincide with the standard forward-only strong bisimilarity of [16] and branching bisimilarity.

In the setting of [8], strong and weak back-and-forth bisimilarities have been defined over computation paths rather than states so that, in the presence of concurrency, any backward computation is *constrained* to follow the same path as the corresponding forward computation, which is consistent with an interleaving view of parallel composition. This is quite different from the forward-reverse bisimilarity over states defined in [18], which accounts for the fact that when going backward the order in which independent transitions are undone may be different from the order in which they were executed in the forward direction, thus leading to a truly concurrent semantics.

Since in our setting we consider only sequential processes, hence any backward computation *naturally* follows the same path as the corresponding forward computation, we are neutral with respect to interleaving vs. true concurrency. Like in [8] we define a single transition relation and then we distinguish between outgoing transitions and incoming transitions in the bisimulation game. However, unlike [8], our bisimilarities are defined over states as in [16, 10, 18], not over paths. In the rest of this section we show that our weak forward-reverse bisimilarity *over states* coincides with branching bisimilarity by following the proof strategy adopted in [8] for weak back-and-forth bisimilarity.

First of all, we prove that, like branching bisimilarity, our weak forward-reverse bisimilarity satisfies the *stuttering property* [10]. This means that, given a sequence of finitely many τ-transitions, if the source process of the first transition and the target process of the last transition are equivalent to each other, then all the intermediate processes are equivalent to them too – see $P_2 \stackrel{\tau^*}{\Longrightarrow} \bar{P}_2$ in Definition 4.1 when $P_1, P_2, \bar{P}_2$ are pairwise related by the maximal branching bisimulation $\approx_{\mathrm{BB}}$. In other words, while traversing the considered sequence of τ-transitions, we remain in the same equivalence class of processes, not only in the forward direction but – as we are talking about weak forward-reverse bisimilarity – *also in the backward direction*. This property does not hold in the case of the standard forward-only weak bisimilarity of [16].

Lemma 4.2 Let $n \in \mathbb{N}_{>0}$, $P_i \in \mathbb{P}$ for all $0 \leq i \leq n$, and $P_i \xrightarrow{\tau} P_{i+1}$ for all $0 \leq i \leq n-1$. If $P_0 \approx_{\mathrm{FRB}} P_n$ then $P_i \approx_{\mathrm{FRB}} P_0$ for all $0 \leq i \leq n$.

Proof Consider the reflexive and symmetric binary relation $\mathscr{B} = \cup_{i \in \mathbb{N}} \mathscr{B}_i$ over $\mathbb{P}$ where:

- $\mathscr{B}_0 = \approx_{\mathrm{FRB}}$.
- $\mathscr{B}_i = \mathscr{B}_{i-1} \cup \{(P, P'), (P', P) \in \mathbb{P} \times \mathbb{P} \mid \exists P'' \in \mathbb{P}. (P, P'') \in \mathscr{B}_{i-1} \wedge P \stackrel{\tau^*}{\Longrightarrow} P' \xrightarrow{\tau} P''\}$ for all $i \in \mathbb{N}_{>0}$.

We start by proving that $\mathscr{B}$ satisfies the stuttering property, i.e., given $n \in \mathbb{N}_{>0}$ and $P_i \in \mathbb{P}$ for all $0 \leq i \leq n$, if $P_i \xrightarrow{\tau} P_{i+1}$ for all $0 \leq i \leq n-1$ and $(P_0, P_n) \in \mathscr{B}$, then $(P_i, P_0) \in \mathscr{B}$ for all $0 \leq i \leq n$. We proceed by induction on n:

- If $n = 1$ then the considered computation is simply $P_0 \xrightarrow{\tau} P_1$ with $(P_0, P_1) \in \mathscr{B}$ and hence trivially $(P_i, P_0) \in \mathscr{B}$ for all $0 \leq i \leq 1$ as $\mathscr{B}$ is reflexive – $(P_0, P_0) \in \mathscr{B}$ – and symmetric – $(P_1, P_0) \in \mathscr{B}$.
- Let $n > 1$. Since $(P_0, P_n) \in \mathscr{B}$, there must exist $m \in \mathbb{N}$ such that $(P_0, P_n) \in \mathscr{B}_m$. Let us consider the smallest such m. Then $(P_0, P_{n-1}) \in \mathscr{B}_{m+1}$ by definition of $\mathscr{B}_{m+1}$, hence $(P_0, P_{n-1}) \in \mathscr{B}$. From the induction hypothesis it follows that $(P_i, P_0) \in \mathscr{B}$ for all $0 \leq i \leq n-1$, hence $(P_i, P_0) \in \mathscr{B}$ for all $0 \leq i \leq n$ because $(P_0, P_n) \in \mathscr{B}$ and $\mathscr{B}$ is symmetric so that $(P_n, P_0) \in \mathscr{B}$.

We now prove that every symmetric relation $\mathscr{B}_i$ is a weak forward-reverse bisimulation. We proceed by induction on $i \in \mathbb{N}$:

- If $i = 0$ then $\mathscr{B}_i$ is the maximal weak forward-reverse bisimulation.
- Let $i \geq 1$ and suppose that $\mathscr{B}_{i-1}$ is a weak forward-reverse bisimulation. Given $(P, P') \in \mathscr{B}_i$, assume that $P \xrightarrow{a} Q$ (resp. $Q \xrightarrow{a} P$) where $a \in A$. There are two cases:
  - If $(P, P') \in \mathscr{B}_{i-1}$ then by the induction hypothesis $a = \tau$ and $P' \stackrel{\tau^*}{\Longrightarrow} Q'$ (resp. $Q' \stackrel{\tau^*}{\Longrightarrow} P'$) or $a \neq \tau$ and $P' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} Q'$ (resp. $Q' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P'$) with $(Q, Q') \in \mathscr{B}_{i-1}$ and hence $(Q, Q') \in \mathscr{B}_i$ as $\mathscr{B}_{i-1} \subseteq \mathscr{B}_i$ by definition of $\mathscr{B}_i$.
  - If instead $(P, P') \notin \mathscr{B}_{i-1}$ then from $(P, P') \in \mathscr{B}_i$ it follows that $\exists P'' \in \mathbb{P}. (P, P'') \in \mathscr{B}_{i-1} \wedge P \stackrel{\tau^*}{\Longrightarrow} P' \xrightarrow{\tau} P''$. There are two subcases:
    - In the forward case, i.e., $P \xrightarrow{a} Q$, there are two further subcases:

- If $(Q, P'') \in \mathscr{B}_{i-1}$ and $a = \tau$, then from $P' \xrightarrow{\tau} P''$ it follows that $P' \stackrel{\tau^*}{\Longrightarrow} P''$ with $(Q, P'') \in \mathscr{B}_i$ as $\mathscr{B}_{i-1} \subseteq \mathscr{B}_i$.
- Otherwise from $(P, P'') \in \mathscr{B}_{i-1}$ and the induction hypothesis it follows that $P'' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P'''$ with $(Q, P''') \in \mathscr{B}_{i-1}$ so that $P' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P'''$ with $(Q, P''') \in \mathscr{B}_i$ as $\mathscr{B}_{i-1} \subseteq \mathscr{B}_i$.

* In the backward case, i.e., $Q \xrightarrow{a} P$, it suffices to note that from $P \stackrel{\tau^*}{\Longrightarrow} P'$ it follows that $Q \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P'$.

Since $\mathscr{B}$ is the union of countably many weak forward-reverse bisimulations, it holds that $\mathscr{B} \subseteq \approx_{\rm FRB}$. On the other hand, $\approx_{\rm FRB} \subseteq \mathscr{B}$ by definition of $\mathscr{B}_0$. In conclusion $\mathscr{B} = \approx_{\rm FRB}$ – i.e., no relation $\mathscr{B}_i$ for $i \in \mathbb{N}_{>0}$ adds further pairs with respect to $\mathscr{B}_0$ – and hence $\approx_{\rm FRB}$ satisfies the stuttering property because so does $\mathscr{B}$. ■

Note that the lemma above considers $\approx_{\rm FRB}$, not $\approx_{\rm FRB:PS}$. Indeed the stuttering property does not hold for $\approx_{\rm FRB:PS}$ when *initial*(P_0), because in that case a τ-action would be decorated inside P_1 and hence $P_1 \not\approx_{\rm FRB:ps} P_0$. Therefore $\approx_{\rm FRB:PS}$ satisfies the stuttering property only over non-initial processes.

Secondly, we prove that $\approx_{\rm FRB}$ satisfies the *cross property* [8]. This means that, whenever two processes reachable from two $\approx_{\rm FRB}$-equivalent processes can perform a sequence of finitely many τ-transitions such that each of the two target processes is $\approx_{\rm FRB}$-equivalent to the source process of the other sequence, then the two target processes are $\approx_{\rm FRB}$-equivalent to each other as well.

Lemma 4.3 Let $P_1, P_2 \in \mathbb{P}$ be such that $P_1 \approx_{\rm FRB} P_2$. For all $P_1', P_1'' \in \mathbb{P}$ reachable from P_1 such that $P_1' \stackrel{\tau^*}{\Longrightarrow} P_1''$ and for all $P_2', P_2'' \in \mathbb{P}$ reachable from P_2 such that $P_2' \stackrel{\tau^*}{\Longrightarrow} P_2''$, if $P_1' \approx_{\rm FRB} P_2''$ and $P_1'' \approx_{\rm FRB} P_2'$ then $P_1'' \approx_{\rm FRB} P_2''$.

Proof Given $P_1, P_2 \in \mathbb{P}$ with $P_1 \approx_{\rm FRB} P_2$, consider the symmetric relation $\mathscr{B} = \approx_{\rm FRB} \cup \{(P_1'', P_2''), (P_2'', P_1'') \in \mathbb{P} \times \mathbb{P} \mid \exists P_1', P_2' \in \mathbb{P}$ resp. reachable from P_1, P_2. $P_1' \stackrel{\tau^*}{\Longrightarrow} P_1'' \land P_2' \stackrel{\tau^*}{\Longrightarrow} P_2'' \land P_1' \approx_{\rm FRB} P_2'' \land P_1'' \approx_{\rm FRB} P_2'\}$. The result follows by proving that $\mathscr{B}$ is a weak forward-reverse bisimulation, because this implies that $P_1'' \approx_{\rm FRB} P_2''$ for every additional pair – i.e., $\mathscr{B}$ satisfies the cross property – as well as $\mathscr{B} = \approx_{\rm FRB}$ – hence $\approx_{\rm FRB}$ satisfies the cross property too.
Let $(P_1'', P_2'') \in \mathscr{B} \setminus \approx_{\rm FRB}$ to avoid trivial cases. Then there exist $P_1', P_2' \in \mathbb{P}$ respectively reachable from P_1, P_2 such that $P_1' \stackrel{\tau^*}{\Longrightarrow} P_1''$, $P_2' \stackrel{\tau^*}{\Longrightarrow} P_2''$, $P_1' \approx_{\rm FRB} P_2''$, and $P_1'' \approx_{\rm FRB} P_2'$. There are two cases:

- In the forward case, assume that $P_1'' \xrightarrow{a} P_1'''$, from which it follows that $P_1' \stackrel{\tau^*}{\Longrightarrow} P_1'' \xrightarrow{a} P_1'''$. Since $P_1' \approx_{\rm FRB} P_2''$, we obtain $P_2'' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2'''$, or $P_2'' \stackrel{\tau^*}{\Longrightarrow} P_2'''$ when $a = \tau$, with $P_1''' \approx_{\rm FRB} P_2'''$ and hence $(P_1''', P_2''') \in \mathscr{B}$. Starting from $P_2'' \xrightarrow{a} P_2'''$ one exploits $P_2' \stackrel{\tau^*}{\Longrightarrow} P_2''$ and $P_1'' \approx_{\rm FRB} P_2'$ instead.
- In the backward case, assume that $P_1''' \xrightarrow{a} P_1''$. Since $P_1'' \approx_{\rm FRB} P_2'$, we obtain $P_2''' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2'$, so that $P_2''' \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} P_2''$, or $P_2''' \stackrel{\tau^*}{\Longrightarrow} P_2'$ when $a = \tau$, so that $P_2''' \stackrel{\tau^*}{\Longrightarrow} P_2''$, with $P_1''' \approx_{\rm FRB} P_2'''$ and hence $(P_1''', P_2''') \in \mathscr{B}$. Starting from $P_2''' \xrightarrow{a} P_2''$ one exploits $P_1' \approx_{\rm FRB} P_2''$ and $P_1' \stackrel{\tau^*}{\Longrightarrow} P_1''$ instead. ■

We are now in a position of proving that $\approx_{\rm FRB}$ coincides with $\approx_{\rm BB}$. This only holds over initial processes though. As an example, $a_1^\dagger . b . P \approx_{\rm BB} a_2^\dagger . b . P$ but $a_1^\dagger . b . P \not\approx_{\rm FRB} a_2^\dagger . b . P$ when $a_1 \neq a_2$.

Theorem 4.4 Let $P_1, P_2 \in \mathbb{P}$ be initial. Then $P_1 \approx_{\rm FRB} P_2$ iff $P_1 \approx_{\rm BB} P_2$.

Proof Given two initial processes $P_1, P_2 \in \mathbb{P}$, we divide the proof into two parts:

- Given a weak forward-reverse bisimulation $\mathcal{B}$ witnessing $P_1 \approx_{\rm FRB} P_2$ and only containing all the pairs of $\approx_{\rm FRB}$-equivalent processes reachable from P_1 and P_2 so that Lemma 4.3 is applicable to $\mathcal{B}$, we prove that $\mathcal{B}$ is a branching bisimulation too. Let $(Q_1, Q_2) \in \mathcal{B}$, where Q_1 is reachable from P_1 while Q_2 is reachable from P_2, and assume that $Q_1 \xrightarrow{a} Q'_1$. There are two cases:
  - Suppose that $a = \tau$ and $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q'_2$ with $(Q'_1, Q'_2) \in \mathcal{B}$. This means that we have a sequence of $n \geq 0$ transitions of the form $Q_{2,i} \xrightarrow{\tau} Q_{2,i+1}$ for all $0 \leq i \leq n-1$ where $Q_{2,0}$ is Q_2 while $Q_{2,n}$ is Q'_2 so that $(Q'_1, Q_{2,n}) \in \mathcal{B}$.
    If $n = 0$ then Q'_2 is Q_2 and we are done because $(Q'_1, Q_2) \in \mathcal{B}$, otherwise from $Q_{2,n}$ we go back to $Q_{2,n-1}$ via $Q_{2,n-1} \xrightarrow{\tau} Q_{2,n}$. If Q'_1 stays idle so that $(Q'_1, Q_{2,n-1}) \in \mathcal{B}$ and $n = 1$ then we are done because $(Q'_1, Q_2) \in \mathcal{B}$, otherwise we go back to $Q_{2,n-2}$ via $Q_{2,n-2} \xrightarrow{\tau} Q_{2,n-1}$. By repeating this procedure, either we get to $(Q'_1, Q_{2,0}) \in \mathcal{B}$ and we are done because $(Q'_1, Q_2) \in \mathcal{B}$, or for some $0 < m \leq n$ such that $(Q'_1, Q_{2,m}) \in \mathcal{B}$ we have that the incoming transition $Q_{2,m-1} \xrightarrow{\tau} Q_{2,m}$ is matched by $\bar{Q}_1 \stackrel{\tau^*}{\Longrightarrow} Q_1 \xrightarrow{\tau} Q'_1$ with $(\bar{Q}_1, Q_{2,m-1}) \in \mathcal{B}$. In the latter case, since $\bar{Q}_1 \stackrel{\tau^*}{\Longrightarrow} Q_1$, $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q_{2,m-1}$, $(\bar{Q}_1, Q_{2,m-1}) \in \mathcal{B}$, and $(Q_1, Q_2) \in \mathcal{B}$, from Lemma 4.3 it follows that $(Q_1, Q_{2,m-1}) \in \mathcal{B}$. In conclusion $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q_{2,m-1} \xrightarrow{\tau} Q_{2,m}$ with $(Q_1, Q_{2,m-1}) \in \mathcal{B}$ and $(Q'_1, Q_{2,m}) \in \mathcal{B}$.
  - Suppose that $a \neq \tau$ and $Q_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2 \xrightarrow{a} \bar{Q}'_2 \stackrel{\tau^*}{\Longrightarrow} Q'_2$ with $(Q'_1, Q'_2) \in \mathcal{B}$.
    From $\bar{Q}'_2 \stackrel{\tau^*}{\Longrightarrow} Q'_2$ and $(Q'_1, Q'_2) \in \mathcal{B}$ it follows that $\bar{Q}'_1 \stackrel{\tau^*}{\Longrightarrow} Q'_1$ with $(\bar{Q}'_1, \bar{Q}'_2) \in \mathcal{B}$. Since Q'_1 already has an incoming a-transition from Q_1 and every non-initial process has exactly one incoming transition, we derive that $\bar{Q}'_1$ is Q'_1 and hence $(Q'_1, \bar{Q}'_2) \in \mathcal{B}$.
    From $\bar{Q}_2 \xrightarrow{a} \bar{Q}'_2$ and $(Q'_1, \bar{Q}'_2) \in \mathcal{B}$ it follows that $\bar{Q}_1 \stackrel{\tau^*}{\Longrightarrow} Q_1 \xrightarrow{a} Q'_1$ with $(\bar{Q}_1, \bar{Q}_2) \in \mathcal{B}$. Since $\bar{Q}_1 \stackrel{\tau^*}{\Longrightarrow} Q_1$, $Q_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2$, $(\bar{Q}_1, \bar{Q}_2) \in \mathcal{B}$, and $(Q_1, Q_2) \in \mathcal{B}$, from Lemma 4.3 it follows that $(Q_1, \bar{Q}_2) \in \mathcal{B}$.
    In conclusion $Q_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2 \xrightarrow{a} \bar{Q}'_2$ with $(Q_1, \bar{Q}_2) \in \mathcal{B}$ and $(Q'_1, \bar{Q}'_2) \in \mathcal{B}$.
- Given a branching bisimulation $\mathcal{B}$ witnessing $P_1 \approx_{\rm BB} P_2$ and only containing all the processes reachable from P_1 and P_2, we prove that $\mathcal{B}$ is a weak forward-reverse bisimulation too. Let $(Q_1, Q_2) \in \mathcal{B}$ with Q_1 reachable from P_1 and Q_2 reachable from P_2. There are two cases:
  - In the forward case, assume that $Q_1 \xrightarrow{a} Q'_1$. Then either $a = \tau$ and $(Q'_1, Q_2) \in \mathcal{B}$, hence $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q_2$ with $(Q'_1, Q_2) \in \mathcal{B}$, or $Q_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2 \xrightarrow{a} Q'_2$ with $(Q_1, \bar{Q}_2) \in \mathcal{B}$ and $(Q'_1, Q'_2) \in \mathcal{B}$, hence $Q_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} Q'_2$ with $(Q'_1, Q'_2) \in \mathcal{B}$.
  - In the backward case – which cannot be the one of $(P_1, P_2) \in \mathcal{B}$ as both processes are initial – assume that $Q'_1 \xrightarrow{a} Q_1$. There are two subcases:
    * Suppose that Q'_1 is P_1. Then either $a = \tau$ and $(Q'_1, Q_2) \in \mathcal{B}$, where Q_2 is P_2 and $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q_2$, or $Q'_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2 \xrightarrow{a} Q_2$ with $(Q'_1, \bar{Q}_2) \in \mathcal{B}$ and $(Q'_1, Q'_2) \in \mathcal{B}$, where Q'_2 is P_2 and $Q'_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} Q_2$.
    * If Q'_1 is not P_1, then P_1 reaches Q'_1 with a sequence of moves that are $\mathcal{B}$-compatible with those with which P_2 reaches some Q'_2 such that $(Q'_1, Q'_2) \in \mathcal{B}$ as $\mathcal{B}$ only contains all the processes reachable from P_1 and P_2. Therefore either $a = \tau$ and $(Q_1, Q'_2) \in \mathcal{B}$, where Q'_2 is Q_2 and $Q_2 \stackrel{\tau^*}{\Longrightarrow} Q_2$, or $Q'_2 \stackrel{\tau^*}{\Longrightarrow} \bar{Q}_2 \xrightarrow{a} Q_2$ with $(Q'_1, \bar{Q}_2) \in \mathcal{B}$ in addition to $(Q'_1, Q'_2) \in \mathcal{B}$ and $(Q_1, Q_2) \in \mathcal{B}$, where $Q'_2 \stackrel{\tau^*}{\Longrightarrow} \xrightarrow{a} \stackrel{\tau^*}{\Longrightarrow} Q_2$. ■

According to the logical characterizations of branching bisimilarity shown in [9], this result opens the way to further logical characterizations of $\approx_{\rm FRB}$ over initial processes in addition to the one of Section 3 based on forward and backward modalities:

- The first additional characterization replaces the two aforementioned modalities with an until operator $\phi_1 \langle\langle a \rangle\rangle \phi_2$. This is satisfied by a process P iff either $a = \tau$ with P satisfying ϕ_2, or $P \stackrel{\tau^*}{\Longrightarrow} \bar{P} \stackrel{a}{\longrightarrow} P'$ with every process along $P \stackrel{\tau^*}{\Longrightarrow} \bar{P}$ satisfying ϕ_1 and P' satisfying ϕ_2.
- The second additional characterization is given by the temporal logic CTL* without the next operator, thanks to a revisitation of the stuttering equivalence of [5] and the bridge between Kripke structures (in which states are labeled with propositions) and labeled transition systems (in which transitions are labeled with actions) built in [9].

5 Conclusion

In this paper we have investigated modal logic characterizations of forward, reverse, and forward-reverse bisimilarities, both strong and weak, over nondeterministic reversible sequential processes. While previous work [4, 3] has addressed compositionality and axiomatizations of those bisimilarities, here the focus has been on identifying suitable modal logics, which are essentially variants of the Hennessy-Milner logic [11], such that two processes are bisimilar iff they satisfy the same set of formulas of the corresponding modal logic.

The additional backward modalities used in this paper are inspired by those in [8], with the important difference that bisimilarities and modal interpretations in the former are defined over states – as is usual – while those in the latter are defined over computation paths. The modal logic characterizations have revealed that strong and weak reverse bisimilarities respectively boil down to strong and weak reverse trace equivalences. Moreover, we have shown that weak forward-reverse bisimilarity coincides with branching bisimilarity [10] over initial processes, thus providing two further logical characterizations for the former thanks to [9].

The study carried out in this paper can contribute, together with the results in [4, 3], to the development of a fully-fledged process algebraic theory of reversible systems. On a more applicative side, following [6] we also observe that the established modal logic characterizations are useful to provide diagnostic information because, whenever two processes are not bisimilar, then there exists at least one formula in the modal logic corresponding to the considered bisimilarity that is satisfied by only one of the two processes and hence can explain the inequivalence.

Acknowledgments. This research has been supported by the PRIN project *NiRvAna – Noninterference and Reversibility Analysis in Private Blockchains*.

References

[1] L. Aceto, A. Ingolfsdottir, K.G. Larsen & J. Srba (2007): *Reactive Systems: Modelling, Specification and Verification*. Cambridge University Press, doi:10.1017/CBO9780511814105.

[2] C.H. Bennett (1973): *Logical Reversibility of Computation*. IBM Journal of Research and Development 17, pp. 525–532, doi:10.1147/rd.176.0525.

[3] M. Bernardo & A. Esposito (2023): *On the Weak Continuation of Reverse Bisimilarity vs. Forward Bisimilarity*. In: Proc. of the 24th Italian Conf. on Theoretical Computer Science (ICTCS 2023), CEUR-WS. To appear.

[4] M. Bernardo & S. Rossi (2023): *Reverse Bisimilarity vs. Forward Bisimilarity*. In: *Proc. of the 26th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS 2023)*, *LNCS* 13992, Springer, pp. 265–284, doi:10.1007/978-3-031-30829-1_13.

[5] M.C. Browne, E.M. Clarke & O. Grümberg (1988): *Characterizing Finite Kripke Structures in Propositional Temporal Logic*. *Theoretical Computer Science* 59, pp. 115–131, doi:10.1016/0304-3975(88)90098-9.

[6] R. Cleaveland (1990): *On Automatically Explaining Bisimulation Inequivalence*. In: *Proc. of the 2nd Int. Workshop on Computer Aided Verification (CAV 1990)*, *LNCS* 531, Springer, pp. 364–372, doi:10.1007/BFb0023750.

[7] V. Danos & J. Krivine (2004): *Reversible Communicating Systems*. In: *Proc. of the 15th Int. Conf. on Concurrency Theory (CONCUR 2004)*, *LNCS* 3170, Springer, pp. 292–307, doi:10.1007/978-3-540-28644-8_19.

[8] R. De Nicola, U. Montanari & F. Vaandrager (1990): *Back and Forth Bisimulations*. In: *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, *LNCS* 458, Springer, pp. 152–165, doi:10.1007/BFb0039058.

[9] R. De Nicola & F. Vaandrager (1995): *Three Logics for Branching Bisimulation*. *Journal of the ACM* 42, pp. 458–487, doi:10.1145/201019.201032.

[10] R.J. van Glabbeek & W.P. Weijland (1996): *Branching Time and Abstraction in Bisimulation Semantics*. *Journal of the ACM* 43, pp. 555–600, doi:10.1145/233551.233556.

[11] M. Hennessy & R. Milner (1985): *Algebraic Laws for Nondeterminism and Concurrency*. *Journal of the ACM* 32, pp. 137–162, doi:10.1145/2455.2460.

[12] M. Hennessy & C. Stirling (1985): *The Power of the Future Perfect in Program Logics*. *Information and Control* 67, pp. 23–52, doi:10.1016/S0019-9958(85)80025-5.

[13] R. Landauer (1961): *Irreversibility and Heat Generation in the Computing Process*. *IBM Journal of Research and Development* 5, pp. 183–191, doi:10.1147/rd.53.0183.

[14] I. Lanese, D. Medić & C.A. Mezzina (2021): *Static versus Dynamic Reversibility in CCS*. *Acta Informatica* 58, pp. 1–34, doi:10.1007/s00236-019-00346-6.

[15] O. Lichtenstein, A. Pnueli & L. Zuck (1985): *The Glory of the Past*. In: *Proc. of the Conf. on Logics in Programs*, *LNCS* 193, Springer, pp. 196–218, doi:10.1007/3-540-15648-8_16.

[16] R. Milner (1989): *Communication and Concurrency*. Prentice Hall.

[17] D. Park (1981): *Concurrency and Automata on Infinite Sequences*. In: *Proc. of the 5th GI Conf. on Theoretical Computer Science*, *LNCS* 104, Springer, pp. 167–183, doi:10.1007/BFb0017309.

[18] I. Phillips & I. Ulidowski (2007): *Reversing Algebraic Process Calculi*. *Journal of Logic and Algebraic Programming* 73, pp. 70–96, doi:10.1016/j.jlap.2006.11.002.