Reverse Bisimilarity vs. Forward Bisimilarity

Marco Bernardo[1]([✉]) and Sabina Rossi[2]

[1] Università di Urbino, Urbino, Italy
`marco.bernardo@uniurb.it`
[2] Università Ca' Foscari di Venezia, Venice, Italy

Abstract. Reversibility is the capability of a system of undoing its own actions starting from the last performed one, in such a way that a past consistent state is reached. This is not trivial for concurrent systems, as the last performed action may not be uniquely identifiable. There are several approaches to address causality-consistent reversibility, some including a notion of forward-reverse bisimilarity. We introduce a minimal process calculus for reversible systems to investigate compositionality properties and equational characterizations of forward-reverse bisimilarity as well as of its two components, i.e., forward bisimilarity and reverse bisimilarity, so as to highlight their differences. The study is conducted not only in a nondeterministic setting, but also in a stochastic one where time reversibility and lumpability for Markov chains are exploited.

1 Introduction

Reversibility started to receive attention in computing several decades ago [15,3]. Landauer's principle states that any irreversible manipulation of information, such as bit erasure or computation path merging, must be accompanied by a corresponding entropy increase. Therefore, any reversible computation, in which no information is lost, may be potentially carried out without releasing any heat. Nowadays, *reversible computing* has many applications ranging from biochemical reaction modeling and parallel discrete-event simulation to robotics, control theory, fault tolerant systems, and concurrent program debugging.

In a reversible system, we can observe two directions of computation: 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.

In the process algebra literature, two approaches have been developed to reverse a computation based on keeping track of past actions: the dynamic one of [7] and the static one of [24]. The former yields RCCS, a variant of CCS [20] that uses stack-based memories attached to processes to record all the actions executed by those processes. 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 and dynamic operators.

In [24] *forward-reverse bisimilarity* is introduced too. Unlike standard bisimilarity [22,20], 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 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 have to occur along the path followed when going forward even in the presence of concurrency.

In this paper, we investigate compositionality properties and equational characterizations of forward-reverse bisimilarity as well as of its two components, i.e., forward bisimilarity and reverse bisimilarity, so as to highlight their differences. To this purpose, we introduce a minimal calculus including only the terminated process $\underline{0}$, the unary action prefix operator $a \,.\, \_$ where a stands for an action, and the binary alternative composition operator $\_ + \_$ also called choice. These operators are enough to compare the essential features of the three equivalences, in a neutral way with respect to interleaving view vs. true concurrency.

The paper is divided into two parts. In Section 2, we conduct our study on *nondeterministic* reversible processes, with the operational semantic rules defined in the style of [24] generating only forward transitions that are viewed as bidirectional, in lieu of a forward transition relation separated from a backward transition relation. In Section 3, we repeat our study on *stochastic* reversible processes, whose operational semantic rules in the style of [24] generate a single transition relation encompassing both forward transitions and backward transitions, by exploiting time reversibility [13] and lumpability [14] for Markov chains. In Section 4, we recap the differences between forward and reverse bisimilarities.

2   The Nondeterministic Case

In this section, we investigate forward bisimilarity, reverse bisimilarity, and forward-reverse bisimilarity over nondeterministic reversible processes. We start by introducing the syntax (Section 2.1) and the semantics (Section 2.2) for these processes through a minimal calculus, then we provide the definitions of the three equivalences (Section 2.3) and we study their congruence properties (Section 2.4) and equational characterizations (Section 2.5).

2.1   Syntax of Nondeterministic Reversible Processes

In the formalization of a process, we usually describe only its future behavior, hence the following syntax for sequential processes where $a \in A$:
$$P \ ::= \ \underline{0} \mid a \,.\, P \mid P + P$$
However, in order to support the definition of the semantics in the style of [24], we need to enrich the syntax above with information about the past, i.e., the actions that have already been executed. Due to the absence of a parallel composition operator, unlike [24] there is no need to add communication keys to executed actions. It thus suffices to mark them with some symbol, which we choose to be †. This yields the following syntax extended with information about the past:
$$P \ ::= \ \underline{0} \mid a \,.\, P \mid a^{\dagger} .\, P \mid P + P$$

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

$$initial(\underline{0})$$
$$initial(a \, . \, P) \impliedby initial(P)$$
$$initial(P_1 + P_2) \impliedby initial(P_1) \wedge initial(P_2)$$

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

$$final(\underline{0})$$
$$final(a^\dagger . \, P) \impliedby final(P)$$
$$final(P_1 + P_2) \impliedby (final(P_1) \wedge initial(P_2)) \vee$$
$$(initial(P_1) \wedge final(P_2))$$

Multiple paths arise only in the presence of alternative compositions. At each occurrence of $+$, only the subprocess chosen for execution advances, while the other one, although not selected, is kept as an initial subprocess within the overall process to support the definition of the semantics in the style of [24].

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

$$reachable(\underline{0})$$
$$reachable(a \, . \, P) \impliedby initial(P)$$
$$reachable(a^\dagger . \, P) \impliedby reachable(P)$$
$$reachable(P_1 + P_2) \impliedby (reachable(P_1) \wedge initial(P_2)) \vee$$
$$(initial(P_1) \wedge reachable(P_2))$$

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. $a^\dagger . \, P$ with $initial(P)$ and $P \neq \underline{0}$.
- 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, $a^\dagger . \, b \, . \, P \in \mathbb{P}$ if $initial(P)$ whereas $b \, . \, a^\dagger . \, P \notin \mathbb{P}$.

2.2   Semantics of Nondeterministic Reversible Processes

According to the approach of [24], 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 [24] 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*: any executed action can be undone and any 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' ($P \overset{a}{\longrightarrow} P'$). As will become clear with the definition of behavioral equivalences in Section 2.3, like in [8] when going forward we view t as an *outgoing* transition of P, while when

$$\text{ACT}_\text{f} \ \frac{initial(P)}{a \,.\, P \xrightarrow{\ a\ } a^\dagger .\, P} \qquad\qquad \text{ACT}_\text{p} \ \frac{P \xrightarrow{\ b\ } P'}{a^\dagger .\, P \xrightarrow{\ b\ } a^\dagger .\, P'}$$

$$\text{CHO}_\text{l} \ \frac{P_1 \xrightarrow{\ a\ } P_1' \quad initial(P_2)}{P_1 + P_2 \xrightarrow{\ a\ } P_1' + P_2} \qquad \text{CHO}_\text{r} \ \frac{P_2 \xrightarrow{\ a\ } P_2' \quad initial(P_1)}{P_1 + P_2 \xrightarrow{\ a\ } P_1 + P_2'}$$

**Table 1.** Operational semantic rules for nondeterministic reversible processes

going backward we view t as an *incoming* transition of P'. The semantic rules in Table 1 generate the labeled transition system $(\mathbb{P}, A, \longrightarrow)$ where $\longrightarrow \ \subseteq \mathbb{P} \times A \times \mathbb{P}$.

The first rule for action prefix (ACT_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 (ACT_p where p stands for propagation) propagates actions executed by inner initial subprocesses.

In both rules for alternative composition (CHO_l and CHO_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.

Any state corresponding to a process different from $\underline{0}$ has at least one outgoing transition and 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 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:
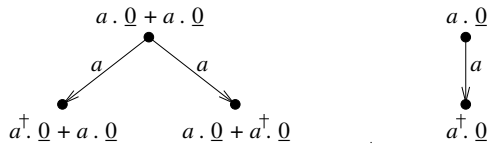
As far as the one on the left is concerned, we observe that, in the case of a standard process calculus, a single a-transition from $a \,.\, \underline{0} + a \,.\, \underline{0}$ to $\underline{0}$ would have been generated due to the absence of action decorations within processes. ∎

2.3 Bisimilarities for Nondeterministic Reversible Processes

The asymmetry between the relative positions of already executed actions and actions to be executed within reachable processes, as well as the asymmetry between the use of predicates *initial* and *final* in the operational semantic rules, determine a number of asymmetries between forward and reverse bisimilarity defined below that will become evident in Sections 2.4 and 2.5.

The difference between the definitions of forward bisimilarity and reverse bisimilarity is that the former considers only *outgoing* transitions [22,20] whereas the latter considers only *incoming* transitions. We also address forward-reverse bisimilarity [24], which considers both outgoing transitions and incoming ones. All the equivalences are strong, i.e., they do not abstract from invisible actions.

Definition 1. *We say that $P_1, P_2 \in \mathbb{P}$ are* forward bisimilar*, written $P_1 \sim_{\mathrm{FB}} P_2$, 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 $P_1 \stackrel{a}{\longrightarrow} P_1'$, then $P_2 \stackrel{a}{\longrightarrow} P_2'$ with $(P_1', P_2') \in \mathcal{B}$.* ∎

Definition 2. *We say that $P_1, P_2 \in \mathbb{P}$ are* reverse bisimilar*, written $P_1 \sim_{\mathrm{RB}} P_2$, 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 $P_1' \stackrel{a}{\longrightarrow} P_1$, then $P_2' \stackrel{a}{\longrightarrow} P_2$ with $(P_1', P_2') \in \mathcal{B}$.* ∎

Definition 3. *We say that $P_1, P_2 \in \mathbb{P}$ are* forward-reverse bisimilar*, written $P_1 \sim_{\mathrm{FRB}} P_2$, 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 $P_1 \stackrel{a}{\longrightarrow} P_1'$, then $P_2 \stackrel{a}{\longrightarrow} P_2'$ with $(P_1', P_2') \in \mathcal{B}$.*
- *Whenever $P_1' \stackrel{a}{\longrightarrow} P_1$, then $P_2' \stackrel{a}{\longrightarrow} P_2$ with $(P_1', P_2') \in \mathcal{B}$.* ∎

It holds that $\sim_{\mathrm{FRB}} \subsetneq \sim_{\mathrm{FB}} \cap \sim_{\mathrm{RB}}$. The inclusion is strict because for example the two final processes $a^\dagger . \underline{0}$ and $a^\dagger . \underline{0} + c . \underline{0}$ are identified by $\sim_{\mathrm{FB}}$ and by $\sim_{\mathrm{RB}}$, but distinguished by $\sim_{\mathrm{FRB}}$ as in the latter process action c is enabled again after undoing a. Moreover, $\sim_{\mathrm{FB}}$ and $\sim_{\mathrm{RB}}$ are incomparable because for instance:
$$a^\dagger . \underline{0} \sim_{\mathrm{FB}} \underline{0} \quad \text{but} \quad a^\dagger . \underline{0} \not\sim_{\mathrm{RB}} \underline{0}$$
$$a . \underline{0} \sim_{\mathrm{RB}} \underline{0} \quad \text{but} \quad a . \underline{0} \not\sim_{\mathrm{FB}} \underline{0}$$
The *first asymmetry* is that $\sim_{\mathrm{FRB}} = \sim_{\mathrm{FB}}$ over initial processes, with $\sim_{\mathrm{RB}}$ strictly coarser, whilst $\sim_{\mathrm{FRB}} \neq \sim_{\mathrm{RB}}$ over final processes because, after going backward, previously discarded subprocesses come into play again in the forward direction.

Example 2. The two processes shown in Example 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 $(a . \underline{0} + a^\dagger . \underline{0}, a^\dagger . \underline{0})$. ∎

2.4   Congruence Properties

In principle, it makes sense that $\sim_{\mathrm{FB}}$ identifies processes with a different past and that $\sim_{\mathrm{RB}}$ identifies processes with a different future, in particular with $\underline{0}$ that has neither past nor future. However, for $\sim_{\mathrm{FB}}$ this results in a compositionality violation with respect to alternative composition. As an example:
$$a^\dagger . b . \underline{0} \sim_{\mathrm{FB}} b . \underline{0}$$
$$a^\dagger . b . \underline{0} + c . \underline{0} \not\sim_{\mathrm{FB}} b . \underline{0} + c . \underline{0}$$

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 $\sim_{\mathrm{RB}}$ as $a^\dagger.b.\underline{0} \not\sim_{\mathrm{RB}} b.\underline{0}$ due to the incoming a-transition of $a^\dagger.b.\underline{0}$, thus yielding the *second asymmetry* between forward and reverse bisimilarity.

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

Definition 4. *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 symmetric relation $\mathcal{B}$ over $\mathbb{P}$ is a* past-sensitive forward bisimulation *iff for all $(P_1, P_2) \in \mathcal{B}$:*

- *$initial(P_1) \Longleftrightarrow initial(P_2)$.*
- *For all $a \in A$, whenever $P_1 \xrightarrow{a} P_1'$, then $P_2 \xrightarrow{a} P_2'$ with $(P_1', P_2') \in \mathcal{B}$.* ∎

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

We conclude by formalizing the congruence properties of all the considered equivalences. When present in the results below, side conditions just ensure that the overall processes are reachable.

Theorem 1. *Let $\sim \in \{\sim_{\mathrm{FB}}, \sim_{\mathrm{FB,ps}}, \sim_{\mathrm{RB}}, \sim_{\mathrm{FRB}}\}$, $\sim' \in \{\sim_{\mathrm{FB,ps}}, \sim_{\mathrm{RB}}, \sim_{\mathrm{FRB}}\}$, and $P_1, P_2 \in \mathbb{P}$:*

- *If $P_1 \sim P_2$ then for all $a \in A$:*
  - *$a.P_1 \sim a.P_2$ provided that $initial(P_1) \wedge initial(P_2)$.*
  - *$a^\dagger.P_1 \sim a^\dagger.P_2$.*
- *If $P_1 \sim' P_2$ then for all $P \in \mathbb{P}$:*
  - *$P_1 + P \sim' P_2 + P$ and $P + P_1 \sim' P + P_2$ provided that $initial(P) \vee (initial(P_1) \wedge initial(P_2))$.*
- *$\sim_{\mathrm{FB,ps}}$ is the coarsest congruence with respect to $+$ contained in $\sim_{\mathrm{FB}}$.* ∎

2.5   Equational Characterizations

We now investigate the equational characterizations of $\sim_{\mathrm{FB,ps}}$, $\sim_{\mathrm{RB}}$, and $\sim_{\mathrm{FRB}}$ so as to highlight the fundamental laws of these behavioral equivalences. In the following, by deduction system we mean a set comprising the following axioms and inference rules on $\mathbb{P}$ – possibly enriched by a set of additional axioms $\mathcal{A}$ – corresponding to the fact that $\sim_{\mathrm{FB,ps}}$, $\sim_{\mathrm{RB}}$, and $\sim_{\mathrm{FRB}}$ are equivalence relations as well as congruences with respect to action prefix and alternative composition:

$(\mathcal{A}_1)$		$(P_1 + P_2) + P_3 \;=\; P_1 + (P_2 + P_3)$	
$(\mathcal{A}_2)$		$P_1 + P_2 \;=\; P_2 + P_1$	
$(\mathcal{A}_3)$		$P + \underline{0} \;=\; P$	
$(\mathcal{A}_4)$	$[\sim_{\mathrm{FB,ps}}]$	$a^\dagger . P \;=\; P$	if $\neg initial(P)$
$(\mathcal{A}_5)$	$[\sim_{\mathrm{FB,ps}}]$	$a_1^\dagger . P \;=\; a_2^\dagger . P$	if $initial(P)$
$(\mathcal{A}_6)$	$[\sim_{\mathrm{FB,ps}}]$	$P + Q \;=\; P$	if $\neg initial(P)$, where $initial(Q)$
$(\mathcal{A}_7)$	$[\sim_{\mathrm{RB}}]$	$a . P \;=\; P$	where $initial(P)$
$(\mathcal{A}_8)$	$[\sim_{\mathrm{RB}}]$	$P + Q \;=\; P$	if $initial(Q)$
$(\mathcal{A}_9)$	$[\sim_{\mathrm{FB,ps}}]$	$P + P \;=\; P$	where $initial(P)$
$(\mathcal{A}_{10})$	$[\sim_{\mathrm{FRB}}]$	$P + Q \;=\; P$	if $initial(Q) \wedge to\_initial(P) = Q$

Table 2. Axioms characterizing bisimilarity over nondeterministic reversible processes

- Reflexivity, symmetry, transitivity: $P = P$, $\dfrac{P_1 = P_2}{P_2 = P_1}$, $\dfrac{P_1 = P_2 \quad P_2 = P_3}{P_1 = P_3}$.

- .-Substitutivity: $\dfrac{P_1 = P_2 \quad initial(P_1) \wedge initial(P_2)}{a . P_1 = a . P_2}$, $\dfrac{P_1 = P_2}{a^\dagger . P_1 = a^\dagger . P_2}$.

- +-Substitutivity: $\dfrac{P_1 = P_2 \quad initial(P) \vee (initial(P_1) \wedge initial(P_2))}{P_1 + P = P_2 + P \quad P + P_1 = P + P_2}$.

It is well known that, in the case of bisimilarity over standard nondeterministic processes, alternative composition turns out to be associative and commutative and to admit $\underline{0}$ as neutral element [11]. The same holds true for $\sim_{\mathrm{FB,ps}}$, $\sim_{\mathrm{RB}}$, and $\sim_{\mathrm{FRB}}$ because the two operational semantic rules for alternative composition are symmetric and $\underline{0}$ has no outgoing or incoming transitions. This is formalized by axioms $\mathcal{A}_1$ to $\mathcal{A}_3$ in Table 2.

Then, we have axioms specific to $\sim_{\mathrm{FB,ps}}$. Axioms $\mathcal{A}_4$ and $\mathcal{A}_5$ together establish that the past can be neglected when moving only forward, but the presence of the past cannot be ignored. Axiom $\mathcal{A}_6$ states that a previously non-selected alternative can be discarded after starting moving only forward.

Likewise, we have axioms specific to $\sim_{\mathrm{RB}}$. Axiom $\mathcal{A}_7$ means that the future can be completely canceled when moving only backward. Axiom $\mathcal{A}_8$ states that a previously non-selected alternative can be discarded when moving only backward. Since there are no constraints on P, axiom $\mathcal{A}_8$ subsumes axiom $\mathcal{A}_3$.

Finally, the idempotency of alternative composition in the case of bisimilarity over standard nondeterministic processes, i.e., $P + P = P$ [11], changes depending on the considered equivalence:

- For $\sim_{\mathrm{FB,ps}}$, idempotency is explicitly formalized by axiom $\mathcal{A}_9$, which we note to be disjoint from axiom $\mathcal{A}_6$ where P cannot be initial.
- For $\sim_{\mathrm{RB}}$, an additional axiom is not needed as idempotency follows from axiom $\mathcal{A}_8$ by taking Q equal to P. Thus, the *third asymmetry* between forward and reverse bisimilarity has to do with idempotency.

- For $\sim_{\mathrm{FRB}}$, idempotency is formalized by axiom $\mathcal{A}_{10}$, where function *to_initial* brings a process back to its initial version by removing all action decorations:
$$
\begin{aligned}
to\_initial(\underline{0}) &= \underline{0} \\
to\_initial(a \,.\, P) &= a \,.\, P \\
to\_initial(a^\dagger \,.\, P) &= a \,.\, to\_initial(P) \\
to\_initial(P_1 + P_2) &= to\_initial(P_1) + to\_initial(P_2)
\end{aligned}
$$
This axiom appeared for the first time in [16] and subsumes axioms $\mathcal{A}_9$ and $\mathcal{A}_6$ for $\sim_{\mathrm{FB,ps}}$ as well as axiom $\mathcal{A}_8$ for $\sim_{\mathrm{RB}}$.

To prove the ground completeness of the equational characterizations of the three considered bisimilarities, as usual we introduce equivalence-specific normal forms to which every process is shown to be reducible, then we work with normal forms only. All the three normal forms rely on the fact that alternative composition is associative and commutative, hence the binary $+$ can be generalized to the n-ary $\sum_{i \in I}$ for a finite nonempty index set I. In the following, we denote by $\vdash$ the deduction relation and we examine the sets of additional axioms below:

- $\mathcal{A}_{\mathrm{FB,ps}} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5, \mathcal{A}_6, \mathcal{A}_9\}$.
- $\mathcal{A}_{\mathrm{RB}} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_7, \mathcal{A}_8\}$.
- $\mathcal{A}_{\mathrm{FRB}} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_{10}\}$.

Definition 5. *We say that $P \in \mathbb{P}$ is in $\sim_{\mathrm{FB,ps}}$-normal form, written $\sim_{\mathrm{FB,ps}}$-nf, iff it is equal to one of the following:*

- $\underline{0}$.
- $\sum_{i \in I} a_i \,.\, P_i$, *where each P_i is initial and in $\sim_{\mathrm{FB,ps}}$-nf.*
- $a^\dagger \,.\, P$, *where P is initial and in $\sim_{\mathrm{FB,ps}}$-nf.* ∎

All initial processes without $\underline{0}$ summands are in $\sim_{\mathrm{FB,ps}}$-nf. We observe that, in the second case, $a_1 \,.\, P_1 \sim_{\mathrm{FB,ps}} a_2 \,.\, P_2$ trivially implies $a_1 = a_2$ and $P_1 \sim_{\mathrm{FB,ps}} P_2$. Likewise, in the third case, $a_1^\dagger \,.\, P_1 \sim_{\mathrm{FB,ps}} a_2^\dagger \,.\, P_2$ trivially implies $P_1 \sim_{\mathrm{FB,ps}} P_2$. These facts will be exploited in the proof of the forthcoming Theorem 2.

Lemma 1. *For all $P \in \mathbb{P}$ there is $Q \in \mathbb{P}$ in $\sim_{\mathrm{FB,ps}}$-nf such that $\mathcal{A}_{\mathrm{FB,ps}} \vdash P = Q$.* ∎

Theorem 2. *Let $P_1, P_2 \in \mathbb{P}$. Then $P_1 \sim_{\mathrm{FB,ps}} P_2$ iff $\mathcal{A}_{\mathrm{FB,ps}} \vdash P_1 = P_2$.* ∎

Definition 6. *We say that $P \in \mathbb{P}$ is in $\sim_{\mathrm{RB}}$-normal form, written $\sim_{\mathrm{RB}}$-nf, iff it is equal to one of the following:*

- $\underline{0}$.
- $a^\dagger \,.\, P$, *where P is in $\sim_{\mathrm{RB}}$-nf.* ∎

The normal form above boils down to a final process consisting of a possibly empty, finite sequence of already executed actions terminated by $\underline{0}$. As a consequence, $a_1^\dagger \,.\, P_1 \sim_{\mathrm{RB}} a_2^\dagger \,.\, P_2$ with P_1 and P_2 in $\sim_{\mathrm{RB}}$-nf implies $a_1 = a_2$ and $P_1 \sim_{\mathrm{RB}} P_2$, because $a_1^\dagger \,.\, P_1$ and $a_2^\dagger \,.\, P_2$ must feature the same sequence of already executed actions and the last executed action of P_1 (resp. P_2), when the process is different from $\underline{0}$, is the same as the last executed action of $a_1^\dagger \,.\, P_1$ (resp. $a_2^\dagger \,.\, P_2$). This fact will be exploited in the proof of the forthcoming Theorem 3.

Lemma 2. *For all $P \in \mathbb{P}$ there is $Q \in \mathbb{P}$ in $\sim_{\mathrm{RB}}$-nf such that $\mathcal{A}_{\mathrm{RB}} \vdash P = Q$.* ■

Theorem 3. *Let $P_1, P_2 \in \mathbb{P}$. Then $P_1 \sim_{\mathrm{RB}} P_2$ iff $\mathcal{A}_{\mathrm{RB}} \vdash P_1 = P_2$.* ■

Definition 7. *We say that $P \in \mathbb{P}$ is in $\sim_{\mathrm{FRB}}$-normal form, written $\sim_{\mathrm{FRB}}$-nf, iff it is equal to one of the following:*

- *$\underline{0}$.*
- *$\sum_{i \in I} a_i \,.\, P_i$, where each P_i is initial and in $\sim_{\mathrm{FRB}}$-nf.*
- *$a^\dagger \,.\, P$, where P is in $\sim_{\mathrm{FRB}}$-nf.*
- *$a^\dagger \,.\, P + \sum_{i \in I} a_i \,.\, P_i$, where P is in $\sim_{\mathrm{FRB}}$-nf and each P_i is initial and in $\sim_{\mathrm{FRB}}$-nf.* ■

As for the second case above, which is concerned with initial processes, we observe that $a_1 \,.\, P_1 \sim_{\mathrm{FRB}} a_2 \,.\, P_2$ trivially implies $a_1 = a_2$ and $P_1 \sim_{\mathrm{FRB}} P_2$. The last two cases together, which are concerned with non-initial processes, yield a process consisting of a finite sequence of already executed actions terminated by an initial process, such that every action in the sequence may have an initial process as an alternative. As a consequence, $a_1^\dagger \,.\, P_1 + P_1' \sim_{\mathrm{FRB}} a_2^\dagger \,.\, P_2 + P_2'$ with P_1, P_2, P_1', P_2' in $\sim_{\mathrm{FRB}}$-nf, P_1' and P_2' initial, and P_1' and P_2' moving only when going back to $\mathit{to\_initial}(a_1^\dagger \,.\, P_1)$ and $\mathit{to\_initial}(a_2^\dagger \,.\, P_2)$, implies $a_1 = a_2$, $P_1 \sim_{\mathrm{FRB}} P_2$, and $P_1' \sim_{\mathrm{FRB}} P_2'$. These facts will be exploited in the proof of the forthcoming Theorem 4.

Lemma 3. *For all $P \in \mathbb{P}$ there is $Q \in \mathbb{P}$ in $\sim_{\mathrm{FRB}}$-nf such that $\mathcal{A}_{\mathrm{FRB}} \vdash P = Q$.* ■

Theorem 4. *Let $P_1, P_2 \in \mathbb{P}$. Then $P_1 \sim_{\mathrm{FRB}} P_2$ iff $\mathcal{A}_{\mathrm{FRB}} \vdash P_1 = P_2$.* ■

3 The Markovian Case

In this section, we repeat the investigation over Markovian reversible processes. We start by recalling the theory of continuous-time Markov chains (Section 3.1) including time reversibility (Section 3.2) and lumpability (Section 3.3), then we introduce syntax and semantics for these processes (Section 3.4), we provide the definitions of the three equivalences (Section 3.5), and we study their congruence properties and equational characterizations (Section 3.6).

3.1 Markov Chains: Definition, Representation, Terminology

A Markov chain is a discrete-state stochastic process characterized by the *memoryless property* [14]. More precisely, a stochastic process $X(t)$, $t \in \mathbb{R}_{\geq 0}$, over a discrete state space $\mathcal{S}$ is a *continuous-time Markov chain (CTMC)* iff for all $n \in \mathbb{N}$, time instants $t_0 < t_1 < \cdots < t_n < t_{n+1} \in \mathbb{R}_{\geq 0}$, and states $s_0, s_1, \ldots, s_n, s_{n+1} \in \mathcal{S}$ it holds that $\Pr\{X(t_{n+1}) = s_{n+1} \mid X(t_i) = s_i, 0 \leq i \leq n\} = \Pr\{X(t_{n+1}) = s_{n+1} \mid X(t_n) = s_n\}$, i.e., the probability of moving from one

state to another does not depend on the particular path that has been followed in the past to reach the current state, hence that path can be forgotten.

A CTMC is representable as a labeled transition system or as a state-indexed matrix. In the first case, each transition is labeled with some probabilistic information describing the evolution from the source state to the target state of the transition. In the second case, the same information is stored into an entry, indexed by those two states, of a matrix. The value of this probabilistic information is a function of the time at which the state change takes place.

For the sake of simplicity, we restrict ourselves to *time-homogeneous* CTMCs, in which conditional probabilities of the form $\Pr\{X(t + t') = s' \mid X(t) = s\}$ do not depend on t, so that the considered information is simply a positive real number given by $\lim_{t' \to 0} \frac{\Pr\{X(t+t')=s'|X(t)=s\}}{t'}$. This is called the *rate* at which the CTMC moves from state s to state s' and uniquely characterizes the exponentially distributed time taken by the considered move.

A CTMC is *irreducible* iff each of its states is reachable from every other state with probability greater than 0. A state $s \in \mathcal{S}$ is *recurrent* iff the CTMC will eventually return to s with probability 1, in which case s is *positive recurrent* iff the expected number of steps until the CTMC returns to it is finite. A CTMC is *ergodic* iff it is irreducible and all of its states are positive recurrent; ergodicity coincides with irreducibility in the case that the CTMC has finitely many states.

Every time-homogeneous and ergodic CTMC $X(t)$ is *stationary*, which means that $(X(t_i + t'))_{1 \leq i \leq n}$ has the same joint distribution as $(X(t_i))_{1 \leq i \leq n}$ for all $n \in \mathbb{N}_{\geq 1}$ and $t_1 < \cdots < t_n, t' \in \mathbb{R}_{\geq 0}$. In this case, $X(t)$ has a unique *steady-state probability distribution* $\boldsymbol{\pi}$ that for all $s \in \mathcal{S}$ fulfills $\pi(s) = \lim_{t \to \infty} \Pr\{X(t) = s \mid X(0) = s'\}$ for any $s' \in \mathcal{S}$. These probabilities can be computed by solving the linear system of *global balance equations* $\boldsymbol{\pi} \cdot \mathbf{Q} = \mathbf{0}$ subject to $\sum_{s \in \mathcal{S}} \pi(s) = 1$ and $\pi(s) \in \mathbb{R}_{>0}$ for all $s \in \mathcal{S}$. The *infinitesimal generator matrix* $\mathbf{Q}$ contains for each pair of distinct states the rate of the corresponding move, which is 0 in the absence of a direct move between them, while $q_{s,s} = -\sum_{s' \neq s} q_{s,s'}$ for all $s \in \mathcal{S}$, i.e., every diagonal element contains the opposite of the total exit rate of the corresponding state, so that each row of $\mathbf{Q}$ sums up to 0.

3.2 Time Reversibility of Continuous-Time Markov Chains

Due to state space explosion and numerical stability problems [27], the calculation of the solution of the global balance equation system is not always feasible. However, it can be tackled in the case that the behavior of the considered CTMC remains the same when the direction of time is reversed. A CTMC $X(t)$ is *time reversible* iff $(X(t_i))_{1 \leq i \leq n}$ has the same joint distribution as $(X(t' - t_i))_{1 \leq i \leq n}$ for all $n \in \mathbb{N}_{\geq 1}$ and $t_1 < \cdots < t_n, t' \in \mathbb{R}_{\geq 0}$. In this case, $X(t)$ and its time-reversed version $X^{\mathrm{r}}(t) = X(t' - t)$ are stochastically identical, in particular they are stationary and share the same steady-state probability distribution $\boldsymbol{\pi}$. In order for a stationary CTMC $X(t)$ to be time reversible, it is necessary and sufficient that the *partial balance equations* $\pi(s) \cdot q_{s,s'} = \pi(s') \cdot q_{s',s}$ are satisfied for all $s, s' \in \mathcal{S}$ such that $s \neq s'$ or, equivalently, that $q_{s_1,s_2} \cdot \ldots \cdot q_{s_{n-1},s_n} \cdot q_{s_n,s_1} = q_{s_1,s_n} \cdot q_{s_n,s_{n-1}} \cdot \ldots \cdot q_{s_2,s_1}$ for all $n \in \mathbb{N}_{\geq 2}$ and distinct $s_1, \ldots, s_n \in \mathcal{S}$ [13].

The time-reversed version $X^{\mathrm{r}}(t)$ of a stationary CTMC $X(t)$ can be defined even when $X(t)$ is not reversible. As shown in [13,10], this is accomplished by using the steady-state probability distribution $\boldsymbol{\pi}$ of $X(t)$, with $X^{\mathrm{r}}(t)$ turning out to be a CTMC too and having the same steady-state probability distribution $\boldsymbol{\pi}$. More precisely, $q^{\mathrm{r}}_{s_j,s_i} = q_{s_i,s_j} \cdot \pi(s_i)/\pi(s_j)$ for all $s_i \neq s_j$, i.e., the rate from state s_j to state s_i in the time-reversed CTMC is proportional to the rate from state s_i to state s_j in the original CTMC, where the coefficient is given by the ratio of $\pi(s_i)$ to $\pi(s_j)$. Note that the time-reversed version of $X^{\mathrm{r}}(t)$ is $X(t)$.

3.3 Lumpability of Continuous-Time Markov Chains

A different approach to the state space explosion problem consists of aggregating states and transitions in a suitable way. In particular, the focus is on *exact aggregations*, i.e., partitions of the state space such that the probability of being in any of the aggregated states is equal to the sum of the probabilities of the original states it contains. In the following, we consider a time-homogeneous CTMC $X(t)$ with state space $\mathcal{S}$ and infinitesimal generator matrix $\mathbf{Q}$; the formulas for the elements of the matrix of the resulting aggregations are taken from [2].

The first notion of exact aggregation that we address is strong lumpability [14]. It was later renamed ordinary lumpability in [28,5], which we prefer to adopt so as not to generate confusion with the use of strong and weak for behavioral equivalences in concurrency theory.

Definition 8. *The partition $\mathcal{P}$ induced by an equivalence relation $\mathcal{L}$ over $\mathcal{S}$ is an* ordinary lumping *iff for all $(s_1,s_2) \in \mathcal{L}$ and $C \in \mathcal{P}$ such that $s_1, s_2 \notin C$:*
$$\sum_{s' \in C} q_{s_1,s'} = \sum_{s' \in C} q_{s_2,s'}$$
The resulting CTMC with state space $\mathcal{P}$ has infinitesimal generator matrix $\mathbf{Q}'$ defined as follows for all $C_1, C_2 \in \mathcal{P}$ such that $C_1 \neq C_2$:
$$q'_{C_1,C_2} = \sum_{s' \in C_2} q_{s,s'}$$
where $s \in C_1$. ∎

The second notion of exact aggregation is exact lumpability [25,28,5], which further enjoys the property that all the original states contained in the same aggregated state have the same probability. While ordinary lumpability considers the rates of outgoing transitions and does not check for rate equality within any class, exact lumpability considers the rates of incoming transitions and applies the rate equality check inside each class too.

Definition 9. *The partition $\mathcal{P}$ induced by an equivalence relation $\mathcal{L}$ over $\mathcal{S}$ is an* exact lumping *iff for all $(s_1,s_2) \in \mathcal{L}$ and $C \in \mathcal{P}$:*
$$\sum_{s' \in C} q_{s',s_1} = \sum_{s' \in C} q_{s',s_2}$$
The resulting CTMC with state space $\mathcal{P}$ has infinitesimal generator matrix $\mathbf{Q}'$ defined as follows for all $C_1, C_2 \in \mathcal{P}$ such that $C_1 \neq C_2$:
$$q'_{C_1,C_2} = \sum_{s' \in C_1} q_{s',s} \cdot (|C_2|/|C_1|)$$
*where $s \in C_2$.* ∎

The third notion of exact aggregation is strict lumpability [5], which is a combination of the previous two.

Definition 10. *The partition $\mathcal{P}$ induced by an equivalence relation $\mathcal{L}$ over $\mathcal{S}$ is a strict lumping iff it is both an ordinary lumping and an exact lumping.* ∎

The relationships between lumpability and time reversibility for CTMCs have been investigated in [18,19]:

- An exact lumping of a CTMC corresponds to an ordinary lumping on the time-reversed CTMC.
- An aggregation of a CTMC is a strict lumping iff it is a strict lumping for the time-reversed CTMC too.
- An exact lumping of a CTMC is also an ordinary lumping whenever the CTMC is time reversible, while the vice versa does not hold in general.

Example 3. Consider the three time-reversible, ergodic CTMCs depicted below:
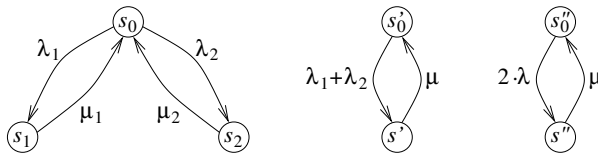
When solving the global balance equations for the first CTMC from the left, we obtain:

$$\pi(s_0) = \frac{\mu_1 \cdot \mu_2}{\mu_1 \cdot \mu_2 + \lambda_1 \cdot \mu_2 + \lambda_2 \cdot \mu_1}$$
$$\pi(s_1) = \frac{\lambda_1 \cdot \mu_2}{\mu_1 \cdot \mu_2 + \lambda_1 \cdot \mu_2 + \lambda_2 \cdot \mu_1}$$
$$\pi(s_2) = \frac{\lambda_2 \cdot \mu_1}{\mu_1 \cdot \mu_2 + \lambda_1 \cdot \mu_2 + \lambda_2 \cdot \mu_1}$$

If $\lambda_1 = \lambda_2$ but $\mu_1 \neq \mu_2$, then no exact aggregation exists for that CTMC.
If $\mu_1 = \mu_2 \triangleq \mu$ but $\lambda_1 \neq \lambda_2$, then the second CTMC from the left is an ordinary lumping of the first one, where the aggregated state s' contains the two original states s_1 and s_2 and the solution of the global balance equations is the following:

$$\pi(s_0') = \frac{\mu}{\mu + \lambda_1 + \lambda_2} = \pi(s_0)$$
$$\pi(s') = \frac{\lambda_1 + \lambda_2}{\mu + \lambda_1 + \lambda_2} = \pi(s_1) + \pi(s_2)$$

with $\pi(s_1) \neq \pi(s_2)$.
If $\lambda_1 = \lambda_2 \triangleq \lambda$ and $\mu_1 = \mu_2 \triangleq \mu$, then the third CTMC from the left is a strict – i.e., ordinary and exact – lumping of the first one, where the aggregated state s'' contains the two original states s_1 and s_2 and the solution of the global balance equations is the following:

$$\pi(s_0'') = \frac{\mu}{\mu + 2 \cdot \lambda} = \pi(s_0)$$
$$\pi(s'') = \frac{2 \cdot \lambda}{\mu + 2 \cdot \lambda} = \pi(s_1) + \pi(s_2)$$

with $\pi(s_1) = \pi(s_2)$. ∎

Example 4. The considered notions of lumpability are distinct from each other. On the one hand, in the previous example the second CTMC from the left is an ordinary lumping of the first one, but not an exact lumping as $\pi(s_1) \neq \pi(s_2)$ when $\mu_1 = \mu_2$ and $\lambda_1 \neq \lambda_2$. On the other hand, the CTMC on the right depicted below is an exact lumping of the CTMC on the left – where the aggregated state s' contains the two original states s_1 and s_2 – when $\mu' + \mu'' = \nu' + \nu''$ – corresponding to $q_{s_1,s_1} + q_{s_2,s_1} = q_{s_1,s_2} + q_{s_2,s_2}$, i.e., $-(\mu' + \mu'') + 0 = 0 - (\nu' + \nu'')$

– but it is not an ordinary lumping if $\mu' \neq \nu'$ and $\mu'' \neq \nu''$:
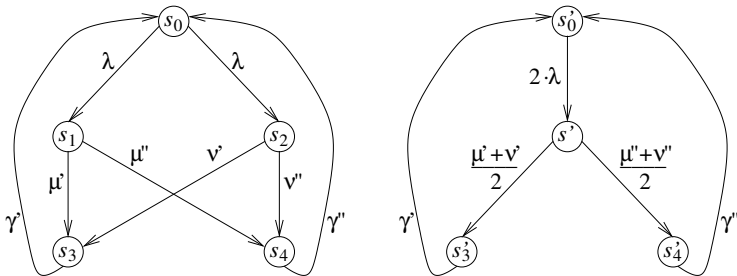
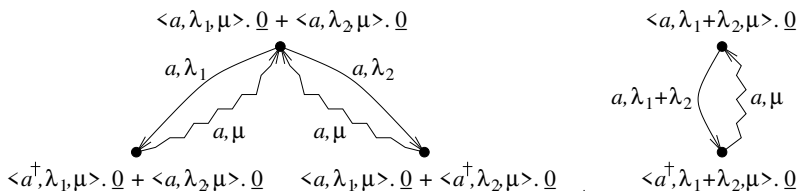Note that the two CTMCs above are ergodic, but not time reversible.   ∎

3.4   Syntax and Semantics of Markovian Reversible Processes

We have seen in Section 2 that a single forward transition relation is enough for nondeterministic processes in a reversible setting. This is due to the fact that $P \xrightarrow{a} P'$ iff $P' \xrightarrow{a} P$, where according to [24] the backward transition relation $\rightsquigarrow$ should be used in the second clause of the definition of $\sim_{\mathrm{FRB}}$ and hence in the definition of $\sim_{\mathrm{RB}}$ as well.

A transition relation in a single direction is no longer sufficient in the case of Markovian reversible processes. The reason is that every transition of these processes is also labeled with its rate, a positive real number that uniquely identifies the exponentially distributed duration of the action associated with the transition. In general, the rate may be different depending on whether the transition goes forward or backward, without necessarily affecting time reversibility.

When moving from nondeterministic reversible processes to Markovian ones, in the syntax we thus need to replace a and $a^\dagger$ with $<a, \lambda, \mu>$ and $<a^\dagger, \lambda, \mu>$ respectively, where $\lambda \in \mathbb{R}_{>0}$ is the rate of the forward a-transition whilst $\mu \in \mathbb{R}_{>0}$ is the rate of the backward a-transition. Predicates *initial*, *final*, and *reachable* are extended accordingly and the set of reachable processes is denoted by $\mathbb{P}_{\mathrm{M}}$.

In order for the semantics to be consistent with the CTMC theory recalled in Sections 3.1 to 3.3, we cannot use a transition relation $\longrightarrow$ with forward rates separated from a transition relation $\rightsquigarrow$ with backward rates, as would be the case if we applied the approach of [24]. For instance, the two Markovian processes depicted below would be identified by a Markovian variant of $\sim_{\mathrm{FRB}}$ relying on $\longrightarrow$ and $\rightsquigarrow$, but the CTMC underlying the labeled transition system of the process on the right is not an exact lumping of the CTMC underlying the labeled transition system of the process on the left if $\lambda_1 \neq \lambda_2$, i.e., this Markovian variant of $\sim_{\mathrm{FRB}}$ would not induce strict lumping:

$$\text{ACT}_\text{f} \; \frac{initial(P)}{<a,\lambda,\mu>.\,P \xrightarrow{a,\lambda}_\text{M} <a^\dagger,\lambda,\mu>.\,P} \qquad \text{ACT}_\text{r} \; \frac{initial(P)}{<a^\dagger,\lambda,\mu>.\,P \xrightarrow{a,\mu}_\text{M} <a,\lambda,\mu>.\,P}$$

$$\text{ACT}_\text{p} \; \frac{P \xrightarrow{b,\xi}_\text{M} P'}{<a^\dagger,\lambda,\mu>.\,P \xrightarrow{b,\xi}_\text{M} <a^\dagger,\lambda,\mu>.\,P'}$$

$$\text{CHO}_\text{l} \; \frac{P_1 \xrightarrow{a,\xi}_\text{M} P_1' \quad initial(P_2)}{P_1 + P_2 \xrightarrow{a,\xi}_\text{M} P_1' + P_2} \qquad \text{CHO}_\text{r} \; \frac{P_2 \xrightarrow{a,\xi}_\text{M} P_2' \quad initial(P_1)}{P_1 + P_2 \xrightarrow{a,\xi}_\text{M} P_1 + P_2'}$$

Table 3. Operational semantic rules for Markovian reversible processes

We thus keep using a single transition relation, which is $\longrightarrow_\text{M} \subseteq \mathbb{P}_\text{M} \times (A \times \mathbb{R}_{>0}) \times \mathbb{P}_\text{M}$ defined in Table 3. Unlike the one in Section 2.2, it embodies both transitions with forward rates and transitions with backward rates. This has been accomplished not only by extending all the rules in Table 1 according to the new richer syntax, but also by adding a rule for action prefix (ACT_r where r stands for reverse) that generates transitions with backward rates.

Any state corresponding to a process different from $\underline{0}$ can now have several incoming transitions too. The labeled transition system underlying an initial process turns out to be a tree-like extension of a birth-death process [23,21], with branching points corresponding to occurrences of $+$. The reason is that between any pair of connected states there can only be a transition from the former state to the latter and a transition from the latter state back to the former, with the two transitions sharing the same name as they are generated by the same action $<a,\lambda,\mu>$. The underlying CTMC, obtained by removing actions from transitions, turns out to be not only ergodic, but also time reversible due to its tree-like birth-death structure [13]. The considered calculus thus combines causality-consistent reversibility with time reversibility like in [4].

Example 5. The labeled transition systems generated by the rules in Table 3 for the two Markovian processes $<a,\lambda,\mu>.\,\underline{0} + <a,\lambda,\mu>.\,\underline{0}$ and $<a,\lambda,\mu>.\,\underline{0}$ are shown below:
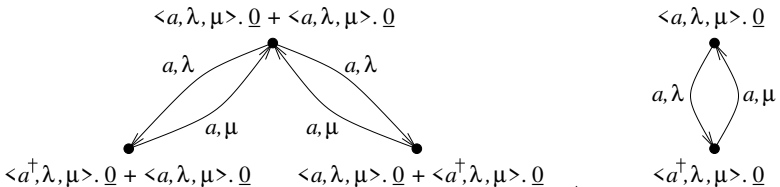
The generation of a single a-transition from $<a,\lambda,\mu>.\,\underline{0} + <a,\lambda,\mu>.\,\underline{0}$ on the left would have been wrong, as it would have not reflected the total exit rate $2\cdot\lambda$ of the source state. Several solutions to this problem have been proposed for Markovian process calculi without reversibility, while in our setting the problem is naturally prevented by action decorations within processes. ∎

3.5 Bisimilarities for Markovian Reversible Processes

We now define the Markovian variants of forward bisimilarity, reverse bisimilarity, and forward-reverse bisimilarity based on the CTMC theory recalled in Sections 3.1 to 3.3.

In the forward case, it is known that the (discrete-time) probabilistic bisimilarity of [17] and the (continuous-time) Markovian bisimilarity of [12] induce an ordinary lumping on the Markov chains underlying the considered processes, hence so does $\sim_{\mathrm{MFB}}$ below. Unlike Definition 8, in Definition 11 the rate equality check is applied inside each class too and hence not all ordinary lumpings can be induced by $\sim_{\mathrm{MFB}}$, in particular not the one identifying every pair of processes.

The reason is that while in Markov chain theory one is interested in state probabilities, in concurrency theory one experiments with processes by observing the labels of the transitions that are executed [9,1,17]. In particular, two processes with different total exit rates cannot be identified by $\sim_{\mathrm{MFB}}$ below, which is perfectly justifiable from an observational viewpoint. As an example, consider a state with a self-looping λ-transition and a state with a self-looping μ-transition. The two states would be deemed ordinarily lumpable according to Definition 8, although the more λ and μ are different, the easier it is for an observer to tell those two states apart.

In the following, $\{\!|$ and $|\!\}$ denote multiset parentheses, while $\mathbb{P}_{\mathrm{M}}/\mathcal{B}$ is the set of equivalence classes induced by the equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$.

Definition 11. *We say that $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$ are* Markovian forward bisimilar*, written $P_1 \sim_{\mathrm{MFB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some Markovian forward bisimulation $\mathcal{B}$. An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$ is a* Markovian forward bisimulation *iff for all $(P_1, P_2) \in \mathcal{B}$, $a \in A$, and $C \in \mathbb{P}_{\mathrm{M}}/\mathcal{B}$:*
$$rate_{\mathrm{out}}(P_1, a, C) \;=\; rate_{\mathrm{out}}(P_2, a, C)$$
where $rate_{\mathrm{out}}(P, a, C) \;=\; \sum \{\!| \, \xi \in \mathbb{R}_{>0} \mid \exists P' \in C. \, P \xrightarrow{a,\xi}_{\mathrm{M}} P' \, |\!\}$. ∎

In the reverse case, incoming transitions are considered instead of outgoing ones. As in [6,26], in the definition of $\sim_{\mathrm{MRB}}$ below an additional condition about total exit rate equality is needed, which in Definition 9 is naturally handled through the diagonal elements of the infinitesimal generator matrix. It is easily seen that $\sim_{\mathrm{MRB}}$ induces an exact lumping on the Markov chains underlying the considered processes, but not all exact lumpings can be induced.

Definition 12. *We say that $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$ are* Markovian reverse bisimilar*, written $P_1 \sim_{\mathrm{MRB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some Markovian reverse bisimulation $\mathcal{B}$. An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$ is a* Markovian reverse bisimulation *iff for all $(P_1, P_2) \in \mathcal{B}$ and $a \in A$:*
$$rate_{\mathrm{out}}(P_1, a, \mathbb{P}_{\mathrm{M}}) \;=\; rate_{\mathrm{out}}(P_2, a, \mathbb{P}_{\mathrm{M}})$$
and for all $C \in \mathbb{P}_{\mathrm{M}}/\mathcal{B}$:
$$rate_{\mathrm{in}}(P_1, a, C) \;=\; rate_{\mathrm{in}}(P_2, a, C)$$
where $rate_{\mathrm{in}}(P, a, C) \;=\; \sum \{\!| \, \xi \in \mathbb{R}_{>0} \mid \exists P' \in C. \, P' \xrightarrow{a,\xi}_{\mathrm{M}} P \, |\!\}$. ∎

In the forward-reverse case, $\sim_{\mathrm{MFRB}}$ below induces a strict lumping on the Markov chains underlying the considered processes.

Definition 13. *We say that $P_1, P_2 \in \mathbb{P}_M$ are* Markovian forward-reverse bisim-
*ilar, written $P_1 \sim_{\text{MFRB}} P_2$, iff $(P_1, P_2) \in \mathcal{B}$ for some Markovian forward-reverse
bisimulation $\mathcal{B}$. An equivalence relation $\mathcal{B}$ over $\mathbb{P}_M$ is a Markovian forward-
reverse bisimulation iff for all $(P_1, P_2) \in \mathcal{B}$, $a \in A$, and $C \in \mathbb{P}_M/\mathcal{B}$:*

$$rate_{\text{out}}(P_1, a, C) = rate_{\text{out}}(P_2, a, C)$$
$$rate_{\text{in}}(P_1, a, C) = rate_{\text{in}}(P_2, a, C) \qquad \blacksquare$$

It is worth noting that any aggregated state resulting from an ordinary lump-
ing is $\sim_{\text{MFB}}$-equivalent to each of the original states it contains, while this is not
necessarily the case for exact lumping and $\sim_{\text{MRB}}$, where $\sim_{\text{MRB}}$-equivalence cer-
tainly holds only among the original states contained in an aggregated state.
This is the *fourth asymmetry* between forward and reverse bisimilarity.

Example 6. The three CTMCs of Example 3 can be viewed as underlying the
labeled transition systems of the following three initial processes:

$$<a, \lambda_1, \mu_1>.\underline{0} + <a, \lambda_2, \mu_2>.\underline{0} \quad \text{corresponding to } s_0$$
$$<a, \lambda_1 + \lambda_2, \mu>.\underline{0} \qquad\qquad\quad \text{corresponding to } s_0'$$
$$<a, 2 \cdot \lambda, \mu>.\underline{0} \qquad\qquad\qquad\quad \text{corresponding to } s_0''$$

with:

$$<a^\dagger, \lambda_1, \mu_1>.\underline{0} + <a, \lambda_2, \mu_2>.\underline{0} \quad \text{corresponding to } s_1$$
$$<a, \lambda_1, \mu_1>.\underline{0} + <a^\dagger, \lambda_2, \mu_2>.\underline{0} \quad \text{corresponding to } s_2$$
$$<a^\dagger, \lambda_1 + \lambda_2, \mu>.\underline{0} \qquad\qquad\quad \text{corresponding to } s'$$
$$<a^\dagger, 2 \cdot \lambda, \mu>.\underline{0} \qquad\qquad\qquad\quad \text{corresponding to } s''$$

If $\mu_1 = \mu_2 \triangleq \mu$ but $\lambda_1 \neq \lambda_2$, then:

$$<a, \lambda_1, \mu>.\underline{0} + <a, \lambda_2, \mu>.\underline{0} \sim_{\text{MFB}} <a, \lambda_1 + \lambda_2, \mu>.\underline{0}$$
$$<a^\dagger, \lambda_1, \mu>.\underline{0} + <a, \lambda_2, \mu>.\underline{0} \sim_{\text{MFB}} <a^\dagger, \lambda_1 + \lambda_2, \mu>.\underline{0}$$
$$<a, \lambda_1, \mu>.\underline{0} + <a^\dagger, \lambda_2, \mu>.\underline{0} \sim_{\text{MFB}} <a^\dagger, \lambda_1 + \lambda_2, \mu>.\underline{0}$$

If $\lambda_1 = \lambda_2 \triangleq \lambda$ and $\mu_1 = \mu_2 \triangleq \mu$, then:

$$<a, \lambda, \mu>.\underline{0} + <a, \lambda, \mu>.\underline{0} \sim_{\text{MFB}} <a, 2 \cdot \lambda, \mu>.\underline{0}$$
$$<a^\dagger, \lambda, \mu>.\underline{0} + <a, \lambda, \mu>.\underline{0} \sim_{\text{MFB}} <a^\dagger, 2 \cdot \lambda, \mu>.\underline{0}$$
$$<a, \lambda, \mu>.\underline{0} + <a^\dagger, \lambda, \mu>.\underline{0} \sim_{\text{MFB}} <a^\dagger, 2 \cdot \lambda, \mu>.\underline{0}$$

but:

$$<a, \lambda, \mu>.\underline{0} + <a, \lambda, \mu>.\underline{0} \nsim_{\text{MRB}} <a, 2 \cdot \lambda, \mu>.\underline{0}$$
$$<a^\dagger, \lambda, \mu>.\underline{0} + <a, \lambda, \mu>.\underline{0} \nsim_{\text{MRB}} <a^\dagger, 2 \cdot \lambda, \mu>.\underline{0}$$
$$<a, \lambda, \mu>.\underline{0} + <a^\dagger, \lambda, \mu>.\underline{0} \nsim_{\text{MRB}} <a^\dagger, 2 \cdot \lambda, \mu>.\underline{0}$$

with the only exception of the following two contained in the same aggregate:

$$<a^\dagger, \lambda, \mu>.\underline{0} + <a, \lambda, \mu>.\underline{0} \sim_{\text{MRB}} <a, \lambda, \mu>.\underline{0} + <a^\dagger, \lambda, \mu>.\underline{0} \qquad \blacksquare$$

Unlike $\sim_{\text{FB}}$, it holds that $\sim_{\text{MFB}}$ is sensitive to the presence of the past,
so that in Definition 11 it is not necessary to require $initial(P_1) \Longleftrightarrow initial(P_2)$
to gain compositionality with respect to alternative composition. For example:

$$<a^\dagger, \lambda, \mu>.<b, \delta, \gamma>.\underline{0} \nsim_{\text{MFB}} <b, \delta, \gamma>.\underline{0}$$

because the process on the left has an outgoing a-transition with rate μ that
cannot be matched by the process on the right.

Furthermore, unlike $\sim_{\text{FB,ps}}$, it holds that $\sim_{\text{MFB}}$ cannot identify processes
with a different past. For instance:

$$<a^\dagger, \lambda, \mu>.\underline{0} \nsim_{\text{MFB}} <b^\dagger, \delta, \gamma>.\underline{0}$$

whenever $a \neq b$ or $\mu \neq \gamma$, as in that case the outgoing a-transition on the left cannot be matched by the outgoing b-transition on the right.

Similarly, unlike $\sim_{\mathrm{RB}}$, we have that $\sim_{\mathrm{MRB}}$ is sensitive to the presence of the future and cannot identify processes with a different future. As an example:

$$<a, \lambda, \mu> . \underline{0} \ \not\sim_{\mathrm{MRB}} \ \underline{0}$$

because the process on the left has an incoming a-transition with rate μ that cannot be matched by the process on the right. As another example:

$$<a, \lambda, \mu> . \underline{0} \ \not\sim_{\mathrm{MRB}} \ <b, \delta, \gamma> . \underline{0}$$

whenever $a \neq b$ or $\mu \neq \gamma$, as in that case the incoming a-transition on the left cannot be matched by the incoming b-transition on the right.

We conclude by showing that $\sim_{\mathrm{MFRB}}$ coincides with $\sim_{\mathrm{MRB}}$ (whilst $\sim_{\mathrm{MFB}}$ is strictly coarser) thus *extending the first asymmetry* between forward and reverse bisimilarities (see page 5). This result stems from the definition of the operational semantics and the consequent time reversibility of the underlying CTMCs.

Theorem 5. *Let $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$. Then $P_1 \sim_{\mathrm{MFRB}} P_2$ iff $P_1 \sim_{\mathrm{MRB}} P_2$.* ∎

3.6 Congruence Properties and Equational Characterizations

We start by observing that $\sim_{\mathrm{MFB}}$ is not totally sensitive to the past, in the same way as $\sim_{\mathrm{MRB}}$ is not totally sensitive to the future. For both equivalences this results in a compositionality violation with respect to $+$. As an example:

$$<a, \lambda, \lambda> . \underline{0} \ \sim_{\mathrm{MFRB}} \ <a^{\dagger}, \lambda, \lambda> . \underline{0}$$
$$<a, \lambda, \lambda> . \underline{0} + <c, \kappa_1, \kappa_2> . \underline{0} \ \not\sim_{\mathrm{MFRB}} \ <a^{\dagger}, \lambda, \lambda> . \underline{0} + <c, \kappa_1, \kappa_2> . \underline{0}$$

because in $<a^{\dagger}, \lambda, \lambda> . \underline{0} + <c, \kappa_1, \kappa_2> . \underline{0}$ action c is disabled due to the presence of the already executed action $a^{\dagger}$, while in $<a, \lambda, \lambda> . \underline{0} + <c, \kappa_1, \kappa_2> . \underline{0}$ action c is enabled as there are no past actions preventing it from occurring.

Note that $\sim_{\mathrm{MFRB}}$ would not equate the first two processes if their two rates were λ_1 and λ_2 with $\lambda_1 \neq \lambda_2$ or there were any other process in place of $\underline{0}$. Therefore, when investigating congruence with respect to alternative composition, we will consider the set of processes $\mathbb{P}'_{\mathrm{M}} = \mathbb{P}_{\mathrm{M}} \setminus \{<a, \lambda, \lambda> . \underline{0} \mid a \in A, \lambda \in \mathbb{R}_{>0}\}$.

Theorem 6. *Let $\sim_{\mathrm{M}} \in \{\sim_{\mathrm{MFB}}, \sim_{\mathrm{MRB}}\}$ and $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$:*

- *If $P_1 \sim_{\mathrm{M}} P_2$ then for all $a \in A$ and $\lambda, \mu \in \mathbb{R}_{>0}$:*
  - *$<a, \lambda, \mu> . P_1 \sim_{\mathrm{M}} <a, \lambda, \mu> . P_2$ provided that $initial(P_1) \wedge initial(P_2)$.*
  - *$<a^{\dagger}, \lambda, \mu> . P_1 \sim_{\mathrm{M}} <a^{\dagger}, \lambda, \mu> . P_2$.*
- *If $P_1 \sim_{\mathrm{M}} P_2$ with $P_1, P_2 \in \mathbb{P}'_{\mathrm{M}}$ then for all $P \in \mathbb{P}'_{\mathrm{M}}$:*
  - *$P_1 + P \sim_{\mathrm{M}} P_2 + P$ and $P + P_1 \sim_{\mathrm{M}} P + P_2$ provided that $initial(P) \vee (initial(P_1) \wedge initial(P_2))$.* ∎

With regard to equational characterizations, as expected $\sim_{\mathrm{MFB}}$ and $\sim_{\mathrm{MRB}}$ are such that alternative composition is associative and commutative and admits $\underline{0}$ as neutral element. This is formalized by axioms $\mathcal{A}_{\mathrm{M},1}$ to $\mathcal{A}_{\mathrm{M},3}$ in Table 4.

Markovian variants of axioms $\mathcal{A}_4$ to $\mathcal{A}_6$ in Table 2 are not valid for $\sim_{\mathrm{MFB}}$ because this behavioral equivalence is sensitive to the presence of the past, cannot identify processes with a different past, and views all the transitions as outgoing.

$(\mathcal{A}_{M,1})$	$(P_1 + P_2) + P_3 = P_1 + (P_2 + P_3)$
$(\mathcal{A}_{M,2})$	$P_1 + P_2 = P_2 + P_1$
$(\mathcal{A}_{M,3})$	$P + \underline{0} = P$

$(\mathcal{A}_{M,4})\ [\sim_{\mathrm{MFB}}]\quad <a,\lambda_1,\mu>.P + <a,\lambda_2,\mu>.P = <a,\lambda_1 + \lambda_2,\mu>.P$
where $initial(P)$

$(\mathcal{A}_{M,5})\ [\sim_{\mathrm{MFB}}]\ <a^\dagger,\lambda_1,\mu>.P + <a,\lambda_2,\mu>.Q = <a^\dagger,\lambda_1 + \lambda_2,\mu>.P$
if $to\_initial(P) = Q$,
where $initial(Q)$

Table 4. Axioms characterizing bisimilarity over Markovian reversible processes

Likewise, Markovian variants of axioms $\mathcal{A}_7$ and $\mathcal{A}_8$ in Table 2 are not valid for $\sim_{\mathrm{MRB}}$ because this behavioral equivalence is sensitive to the presence of the future, cannot identify processes with a different future, and views all the transitions as incoming.

As for idempotency, Markovian variants of axioms $\mathcal{A}_9$ and $\mathcal{A}_{10}$ in Table 2, which are formalized by axioms $\mathcal{A}_{M,4}$ and $\mathcal{A}_{M,5}$ in Table 4, are valid only for $\sim_{\mathrm{MFB}}$ as shown in Example 6. We further observe that in the considered example:
$$<a^\dagger,\lambda,\mu>.\underline{0} + <a,\lambda,\mu>.\underline{0} \ \sim_{\mathrm{MRB}}\ <a,\lambda,\mu>.\underline{0} + <a^\dagger,\lambda,\mu>.\underline{0}$$
can be proved via axiom $\mathcal{A}_{M,2}$.

Theorem 7. *Let $\mathcal{A}_{\mathrm{MFB}} = \{\mathcal{A}_{M,1}, \mathcal{A}_{M,2}, \mathcal{A}_{M,3}, \mathcal{A}_{M,4}, \mathcal{A}_{M,5}\}$ and $P_1, P_2 \in \mathbb{P}'_M$. Then $P_1 \sim_{\mathrm{MFB}} P_2$ iff $\mathcal{A}_{\mathrm{MFB}} \vdash P_1 = P_2$.* ∎

Theorem 8. *Let $\mathcal{A}_{\mathrm{MRB}} = \{\mathcal{A}_{M,1}, \mathcal{A}_{M,2}, \mathcal{A}_{M,3}\}$ and $P_1, P_2 \in \mathbb{P}'_M$. Then $P_1 \sim_{\mathrm{MRB}} P_2$ iff $\mathcal{A}_{\mathrm{MRB}} \vdash P_1 = P_2$.* ∎

4 Conclusions

In this paper, we have discovered the following asymmetries that shed light on forward bisimilarity, reverse bisimilarity, and forward-reverse bisimilarity:

1. In the nondeterministic case $\sim_{\mathrm{FRB}} = \sim_{\mathrm{FB}}$ over initial processes only, while in the Markovian case $\sim_{\mathrm{MFRB}} = \sim_{\mathrm{MRB}}$ over all reachable processes.
2. The insensitivity to the presence of the past breaks the compositionality of $\sim_{\mathrm{FB}}$, while the insensitivity to the presence of the future does not violate the compositionality of $\sim_{\mathrm{RB}}$. This does not happen in the Markovian case.
3. Forward bisimilarity needs explicit idempotency axioms, while reverse bisimilarity does not, especially in the nondeterministic case.
4. Any aggregated state resulting from an ordinary lumping is $\sim_{\mathrm{MFB}}$-equivalent to each of the original states it contains, while this is not necessarily the case for exact lumping and $\sim_{\mathrm{MRB}}$, where $\sim_{\mathrm{MRB}}$-equivalence certainly holds only among the original states contained in an aggregated state.

As future work, we plan to investigate logical characterizations of the same equivalences, along with what changes when admitting irreversible actions.

References

1. Abramsky, S.: Observational equivalence as a testing equivalence. Theoretical Computer Science **53**, 225–241 (1987)
2. Baarir, S., Beccuti, M., Dutheillet, C., Franceschinis, G., Haddad, S.: Lumping partially symmetrical stochastic models. Performance Evaluation **68**, 21–44 (2011)
3. Bennett, C.H.: Logical reversibility of computations. IBM Journal of Research and Development **17**, 525–532 (1973)
4. Bernardo, M., Mezzina, C.A.: Towards bridging time and causal reversibility. In: Proc. of the 40th Int. Conf. on Formal Techniques for Distributed Objects, Components, and Systems (FORTE 2020). LNCS, vol. 12136, pp. 22–38. Springer (2020)
5. Buchholz, P.: Exact and ordinary lumpability in finite Markov chains. Journal of Applied Probability **31**, 59–75 (1994)
6. Buchholz, P.: Exact performance equivalence: An equivalence relation for stochastic automata. Theoretical Computer Science **215**, 263–287 (1999)
7. Danos, V., Krivine, J.: Reversible communicating systems. In: Proc. of the 15th Int. Conf. on Concurrency Theory (CONCUR 2004). LNCS, vol. 3170, pp. 292–307. Springer (2004)
8. De Nicola, R., Montanari, U., Vaandrager, F.: Back and forth bisimulations. In: Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990). LNCS, vol. 458, pp. 152–165. Springer (1990)
9. van Glabbeek, R.J.: The linear time – branching time spectrum I. In: Handbook of Process Algebra. pp. 3–99. Elsevier (2001)
10. Harrison, P.: Turning back time in Markovian process algebra. Theoretical Computer Science **290**, 1947–1986 (2003)
11. Hennessy, M., Milner, R.: Algebraic laws for nondeterminism and concurrency. Journal of the ACM **32**, 137–162 (1985)
12. Hillston, J.: A Compositional Approach to Performance Modelling. Cambridge University Press (1996)
13. Kelly, F.: Reversibility and Stochastic Networks. John Wiley & Sons (1979)
14. Kemeny, J.G., Snell, J.L.: Finite Markov Chains. Van Nostrand (1960)
15. Landauer, R.: Irreversibility and heat generated in the computing process. IBM Journal of Research and Development **5**, 183–191 (1961)
16. Lanese, I., Phillips, I.: Forward-reverse observational equivalences in CCSK. In: Proc. of the 13th Int. Conf. on Reversible Computation (RC 2021). LNCS, vol. 12805, pp. 126–143. Springer (2021)
17. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Information and Computation **94**, 1–28 (1991)
18. Marin, A., Rossi, S.: On the relations between lumpability and reversibility. In: Proc. of the 22nd IEEE Int. Symp. on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2014). pp. 427–432. IEEE-CS Press (2014)
19. Marin, A., Rossi, S.: On the relations between Markov chain lumpability and reversibility. Acta Informatica **54**, 447–485 (2017)

20. Milner, R.: Communication and Concurrency. Prentice Hall (1989)
21. Palacios, J.L., Quiroz, D.: Birth and death chains on finite trees: Computing their stationary distribution and hitting times. Methodology and Computing in Applied Probability **18**, 487–498 (2016)
22. Park, D.: Concurrency and automata on infinite sequences. In: Proc. of the 5th GI Conf. on Theoretical Computer Science. LNCS, vol. 104, pp. 167–183. Springer (1981)
23. Pearce, L.H.: Random walks on trees. Discrete Mathematics **30**, 269–276 (1980)
24. Phillips, I., Ulidowski, I.: Reversing algebraic process calculi. Journal of Logic and Algebraic Programming **73**, 70–96 (2007)
25. Schweitzer, P.J.: Aggregation methods for large Markov chains. In: Proc. of the Int. Workshop on Computer Performance and Reliability. pp. 275–286. North Holland (1984)
26. Sproston, J., Donatelli, S.: Backward bisimulation in Markov chain model checking. IEEE Trans. on Software Engineering **32**, 531–546 (2006)
27. Stewart, W.: Introduction to the Numerical Solution of Markov Chains. Princeton University Press (1994)
28. Sumita, U., Rieders, M.: Lumpability and time reversibility in the aggregation-disaggregation method for large Markov chains. Communications in Statistics - Stochastic Models **5**, 63–81 (1989)