

BRIDGING CAUSAL REVERSIBILITY AND TIME REVERSIBILITY: A STOCHASTIC PROCESS ALGEBRAIC APPROACH

MARCO BERNARDO AND CLAUDIO A. MEZZINA

Dipartimento di Scienze Pure e Applicate, Università di Urbino, Italy
e-mail address: marco.bernardo@uniurb.it, claudio.mezzina@uniurb.it

ABSTRACT. Causal reversibility blends reversibility and causality for concurrent systems. It indicates that an action can be undone provided that all of its consequences have been undone already, thus making it possible to bring the system back to a past consistent state. Time reversibility is instead considered in the field of stochastic processes, mostly for efficient analysis purposes. A performance model based on a continuous-time Markov chain is time reversible if its stochastic behavior remains the same when the direction of time is reversed. We bridge these two theories of reversibility by showing the conditions under which causal reversibility and time reversibility are both ensured by construction. This is done in the setting of a stochastic process calculus, which is then equipped with a variant of stochastic bisimilarity accounting for both forward and backward directions.

1. INTRODUCTION

The interest into computation reversibility dates back to the 60's, when it was observed that irreversible computations cause heat dissipation into circuits [Lan61]. More precisely, Landauer's principle states that any logically irreversible manipulation of information, such as the erasure of bits or the merging of computation paths, must be accompanied by a corresponding entropy increase in non-information-bearing degrees of freedom of the information processing apparatus or its environment [Ben03]. Hence, according to this principle, which has been recently verified in [BAP⁺12] and given a physical foundation in [Fra18], any logically reversible computation, in which no information is erased, may be potentially carried out without releasing any heat. This suggested that low energy consumption could be achieved by resorting to *reversible computing*, in which there is no information loss [Ben73]. Nowadays, reversible computing has several applications such as biochemical reaction modeling [PUY13, Pin17], parallel discrete-event simulation [PP14, SOJB18], robotics [LES18], control theory [SPP19], fault tolerant systems [DK05, dKH10, LLM⁺13, VS18], and concurrent program debugging [GLM14, LNPV18a].

In a reversible system, we see two directions of computation: a *forward* one, coinciding with the normal way of computing, and a *backward* one, which is able to undo the effects of the forward one when needed. In the literature, there are different variants of reversibility. For instance, in a Petri net reversibility means that one can always reach the initial marking [BKMP18], while in a distributed system it stands for the capability of returning to a past

Key words and phrases: Reversibility, Causality, Markov Chains, Process Calculi, Bisimilarity.




Figure 1: Making a system causally reversible and time reversible

consistent state [DK04, DK05]. In contrast, in the performance evaluation field reversibility is related to time and is instrumental to develop efficient analysis methods [Kel79].

Our focus is on integrating *causal reversibility* and *time reversibility*. Causal reversibility describes the capability of going back to a past state that is consistent with the computational history of a system; in this setting, quantitative aspects have been totally disregarded. On the other hand, the theory of time reversibility studies the conditions under which the stochastic behavior of a system remains the same when the direction of time is reversed; unfortunately, it has been applied to concurrent systems without explicitly taking causality into account. In this paper, we aim at bridging these two theories by showing how causal reversibility and time reversibility can be jointly achieved. To this purpose, since process algebra constitutes a common ground for concurrency theory and probability theory [LS91, Hil96], we develop our proposal by considering a stochastic process calculus, in which actions are equipped with positive real numbers. Each of these numbers expresses the rate at which the corresponding action is executed and uniquely identifies the exponential distribution that quantifies the duration of the action, so that the stochastic process underlying such a calculus turns out to be a continuous-time Markov chain (CTMC) [KS60].

Ensuring that a system is both causally reversible and time reversible is not a trivial task. Consider for instance a system that can perform either action a at rate λ followed by action b at rate μ , or action c at rate γ . In stochastic process algebra terms, it would be represented as $\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle c, \gamma \rangle . \underline{0}$ where “.” is the action prefix operator, “+” is the choice operator, and $\underline{0}$ is the terminated process. The underlying labeled transition system is the action-labeled CTMC depicted in the leftmost part of Figure 1. This system is not reversible because, given a pair of states connected by a transition, it is not possible to go from the target state back to the source state. So the first step towards reversibility is to add a backward transition for each such pair of states, which is labeled with the same action as the forward transition together with a suitable backward rate. The resulting labeled transition system is the action-labeled CTMC depicted in the central part of Figure 1.

However, this system is not causally reversible because after performing a and b (resp. c) the system could go back by undoing c (resp. b and a), i.e., the actions in the backward direction could be different from those in the forward direction. In order for this system to be time reversible, the product of the rates along any cycle should be the same when changing direction [Kel79]. This holds for cycles of length 2 or 4 involving only actions a and b , whereas it is not the case in general with cycles of length 3 involving all the actions, unless $\lambda \cdot \mu \cdot \bar{\gamma} = \gamma \cdot \bar{\mu} \cdot \bar{\lambda}$ as happens, e.g., when $\lambda = \bar{\lambda}$, $\mu = \bar{\mu}$, and $\gamma = \bar{\gamma}$, i.e., when every backward rate is equal to the corresponding forward rate. The version depicted in the rightmost part of Figure 1, which is obtained from $\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle c, \gamma \rangle . \underline{0}$ by applying the approach we propose, is both causally reversible and time reversible.

The contribution of this paper is threefold. Firstly, we apply for the first time the general methodology for reversing process calculi of [PU07] to a stochastic process calculus. In particular, we provide forward and backward operational semantic rules – featuring forward and backward actions and rates – and then we show that the resulting calculus is causally reversible. The latter is accomplished by importing from the reduction semantics setting of [DK04] the notion of concurrent transitions, which is new in the structural operational semantics framework of [PU07] and is then handled through the recent technique of [LPU20].

Secondly, we prove that time reversibility can be obtained by using in the operational semantic rules backward rates equal to the corresponding forward rates or by restricting the syntax in such a way that parallel composition cannot occur within the scope of action prefix or choice. This is quite different from the approach adopted for example in [Har03, MR15], where time reversibility is verified a posteriori, as we instead produce a calculus in which time reversibility can be guaranteed by construction. The difference in the approach does not prevent us from importing in our setting compositionality results from those works.

Thirdly, we address behavioral equivalences for our reversible stochastic process calculus, in order to provide a means to identify systems possessing the same functional and performance properties. In particular, we focus on Markovian bisimilarity [Hil96], which equates systems capable of stepwise mimicking each other’s functional and performance behavior. We extend a result of [DMV90] to our stochastic setting by showing that a variant of Markovian bisimilarity on computations accounting for both forward and backward directions coincides with Markovian bisimilarity on states, thus inheriting the nice properties of the latter [Ber07].

This paper, which is a revised and enriched version of [BM20], is organized as follows. In Sections 2 and 3 we recall background notions about causal reversibility and time reversibility, respectively. In Section 4 we develop our proposal of integration for these two forms of reversibility in the setting of a CTMC-based process calculus, which is then equipped in Section 5 with a forward and backward Markovian bisimilarity. In Section 6 we provide some examples. Finally, in Section 7 we conclude with some directions for future work.

2. CAUSAL REVERSIBILITY OF CONCURRENT SYSTEMS

Reversibility in a computing system has to do with the possibility of reverting the last performed action. In a sequential system, this is very simple as there exists just one last action, hence the only challenge is how to store the information needed to reverse this last action. After Bennett [Ben73], several techniques have been developed to reverse the computations of a sequential program [Lee86, YG07, HGCH21, PV15].

In a concurrent system, the situation is much more complex as the last action may not be uniquely identifiable. Indeed, there might be several concurrent last actions. One could resort to timestamps to decide which action is the last one, but having synchronized clocks in a distributed system is rather difficult. A good approximation is to consider as last action every action that has not caused any other action yet. This is at the basis of the so called *causal reversibility*, which relates reversibility with causality [DK04, DK05]. Intuitively, the definition states that, in a concurrent system, any executed action can be undone provided that all of its consequences, if any, have been undone beforehand.

In the process algebra literature, specifically for CCS [Mil89], two approaches have been developed to reverse a computation and to keep track of past actions: the *dynamic* approach of [DK04] and the *static* approach of [PU07]. Despite these two approaches being quite different, they have been recently shown to be equivalent in terms of labeled transition

system isomorphism [LMM21]. On the application side, the former approach is more suitable for systems whose operational semantics is given in terms of reduction semantics, hence it is to be preferred in the case of very expressive calculi [LMS10, CKV13, LM20] as well as programming languages [LLMS12, LNPV18b]. On the other hand, the latter approach is very handy when it comes to deal with labeled transition systems and basic process calculi.

The dynamic approach relies on RCCS [DK04], a variant of CCS that uses external memories attached to processes to record all the actions executed by the processes themselves. Every RCCS process is thus of the form $m \triangleright P$, where m is a memory whilst P is a CCS process. At the beginning of the computation, if P is not sequential, i.e., it contains occurrences of the parallel composition operator, then an empty memory is distributed among all the sequential components. Every time that one of the sequential components performs an action, the action is pushed on top of the memory stack of the component as it is its last executed action. This construction is shown to allow a concurrent system to backtrack along any causally equivalent path in such a way that no previously inaccessible state is reached, which is the essence of causal reversibility.

In contrast, the static approach of [PU07] proposes a general method, of which CCSK is a result, to reverse process calculi whose operational semantic rules are expressed in the path format of [BV93]. The basic idea of this method is to make all the operators of the calculus static and each executed action univocally identified by a communication key. In particular, since dynamic process algebraic operators like action prefix “.” and choice “+” would disappear after a transition – thus causing information loss – all of them are kept within target processes of transitions. This is the approach that we will follow for defining our reversible stochastic process calculus in Section 4.

For example, consider the CCS process $a.P + b.Q$, which either performs action a and then behaves as process P , or performs action b and then behaves as process Q . In the CCS transition $a.P + b.Q \xrightarrow{a} P$, both the executed action a and the discarded process $b.Q$ are lost after performing the transition, hence it is not possible to get back to the initial state. In CCSK the same process evolves instead as $a.P + b.Q \xrightarrow{a[i]} a[i].P + b.Q$, where $a[i].P + b.Q$ behaves like P in further forward transitions while the colored parts can be seen as decorations of P to be used only in backward transitions. Alternatively, $a[i].P + b.Q$ can be seen as $\mathcal{C}[P]$, where P is the active part and $\mathcal{C} = a[i].\bullet + b.Q$ is its dead context. In this way the use of external memories of [DK04] is avoided, because all the necessary information for enabling reversibility is syntactically maintained within processes. Note that action a is decorated with a fresh key $[i]$ – e.g., a natural number – so as to distinguish among several actions executed in the past that have the same name. As we will see in Section 4, this is important for correctly managing communications when going backward.

3. TIME REVERSIBILITY OF MARKOV CHAINS

In the performance evaluation field, a different notion of reversibility, called *time reversibility*, is considered. We illustrate it in the specific case of continuous-time Markov chains, which are discrete-state stochastic processes characterized by the *memoryless property* [KS60].

A *stochastic process* describes the evolution of some random phenomenon over time through a set of random variables, one for each time instant. A stochastic process $X(t)$ taking values from a discrete state space \mathcal{S} for $t \in \mathbb{R}_{\geq 0}$ is a *continuous-time Markov chain (CTMC)* iff for all $n \in \mathbb{N}$, time instants $t_0 < t_1 < \dots < t_n < t_{n+1} \in \mathbb{R}_{\geq 0}$, and states $s_0, s_1, \dots, s_n, s_{n+1} \in \mathcal{S}$ 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 can be equivalently represented 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 s to the target state s' of the transition itself. In the second case, the same information is stored into an entry, indexed by those two states, of a square matrix. The value of this information is, in general, a function of time.

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 information considered above is given by $\lim_{t' \rightarrow 0} \frac{\Pr\{X(t+t')=s' \mid X(t)=s\}}{t'}$. This limit yields a number called the *rate* at which the CTMC moves from state s to state s' and characterizes the exponentially distributed random time taken by the considered move. It can be shown that the sojourn time in any state $s \in \mathcal{S}$ is exponentially distributed with rate given by the sum of the rates of the moves of s . The average sojourn time in s is the inverse of such a sum and the probability of moving from s to s' is the ratio of the corresponding rate to the aforementioned sum.

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 called *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, as they form a finite strongly connected component.

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 $t_1 < \dots < t_n, t' \in \mathbb{R}_{\geq 0}$. In this case, $X(t)$ has a unique *steady-state probability distribution* $\boldsymbol{\pi} = (\pi(s))_{s \in \mathcal{S}}$ that fulfills $\pi(s) = \lim_{t \rightarrow \infty} \Pr\{X(t) = s \mid X(0) = s'\}$ for any $s' \in \mathcal{S}$ because the CTMC has reached equilibrium. These probabilities are 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* $\mathbf{Q} = (q_{s,s'})_{s,s' \in \mathcal{S}}$ 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. Therefore $\boldsymbol{\pi} \cdot \mathbf{Q} = \mathbf{0}$ means that, once reached equilibrium, for every state the incoming probability flux is equal to the outgoing probability flux.

Due to state space explosion and numerical stability problems, the calculation of the solution of the global balance equation system is not always feasible [Ste94]. 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 $t_1 < \dots < t_n, t' \in \mathbb{R}_{\geq 0}$. In this case $X(t)$ and its reversed version $X^r(t) = X(-t)$, $t \in \mathbb{R}_{\geq 0}$, 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 $q_{s_1,s_2} \cdot \dots \cdot q_{s_{n-1},s_n} \cdot q_{s_n,s_1} = q_{s_1,s_n} \cdot q_{s_n,s_{n-1}} \cdot \dots \cdot q_{s_2,s_1}$ for all $n \in \mathbb{N}_{\geq 2}$ and distinct $s_1, \dots, s_n \in \mathcal{S}$ [Kel79]. Note that the sum of the partial balance equations for $s \in \mathcal{S}$ yields the global balance equation $\pi(s) \cdot |q_{s,s}| = \sum_{s' \neq s} \pi(s') \cdot q_{s',s}$.

$ \begin{array}{l} P, Q ::= \underline{0} \mid \langle a, \lambda \rangle . P \mid P + Q \mid P \parallel_L Q \\ R, S ::= P \mid \langle a, \lambda \rangle [i] . R \mid R + S \mid R \parallel_L S \end{array} $

Table 1: Syntax of forward processes (top) and reversible processes (bottom)

The time-reversed version $X^r(t)$ of a stationary CTMC $X(t)$ can be defined even when $X(t)$ is not reversible. As shown in [Kel79, Har03], this is accomplished by using the steady-state probability distribution π of $X(t)$, with $X^r(t)$ turning out to be a CTMC too and having the same steady-state probability distribution π . More precisely, $q_{s_j, s_i}^r = \frac{\pi(s_i)}{\pi(s_j)} \cdot q_{s_i, s_j}$ 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 $X^r(t)$ is $X(t)$.

4. INTEGRATING CAUSAL REVERSIBILITY AND TIME REVERSIBILITY

In this section, we integrate the two concepts of causal reversibility and time reversibility recalled in the previous two sections by means of a simple calculus named RMPC – Reversible Markovian Process Calculus. We start with its syntax – where actions are paired with rates – and its semantics, which are inspired by [PU07]. We then show that the reversibility induced by RMPC is consistent with causality by adapting the notion of concurrent transitions from [DK04] and exploiting the technique of [LPU20]. Finally, we exhibit the conditions under which time reversibility is achieved too and we compare our setting, in which time reversibility is ensured by construction, with those of [Har03, MR15], in which time reversibility is verified a posteriori, and import from them some results enabling an efficient analysis of performance.

4.1. Syntax of RMPC. The syntax of RMPC is shown in Table 1. A standard *forward process* P can be one of the following: the terminated process $\underline{0}$; the action-prefixed process $\langle a, \lambda \rangle . P$, which is able to perform an action a at rate λ – called an exponentially timed action as its duration follows an exponential distribution of parameter λ – and then continues as process P ; the choice $P + Q$ between processes P and Q based on the rates of their initial actions; or the parallel composition $P \parallel_L Q$, indicating that processes P and Q execute in parallel and must synchronize only on actions belonging to L .

A *reversible process* R is built on top of forward processes. As in [PU07], the syntax of reversible processes differs from the one of forward processes due to the fact that in the former each prefix $\langle a, \lambda \rangle$ can be decorated with a *communication key* i thus becoming $\langle a, \lambda \rangle [i]$. A process of the form $\langle a, \lambda \rangle [i] . R$ expresses that in the past the process synchronized with the environment and this synchronization was identified by key i . Keys are thus attached only to already executed actions; they are not needed in the absence of parallel composition, provided that executed actions are marked in some way.

Let \mathcal{A} be a countable set of actions (ranged over by a, b), $\mathcal{R} = \mathbb{R}_{>0}$ be a set of rates (ranged over by λ, μ), and \mathcal{K} be a countable set of keys (ranged over by i, j). Let $\mathcal{L} = \mathcal{A} \times \mathcal{R} \times \mathcal{K}$ be a set of labels each formed by an action, a rate, and a communication key (ranged over by ℓ). Given a forward label $\ell = \langle a, \lambda \rangle [i]$, we write $\bar{\ell} = \langle a, \bar{\lambda} \rangle [i]$ for the corresponding backward label, where $\bar{\lambda}$ stands for the backward rate of the action – i.e., the rate at which the action is undone – whose value will be discussed later on.

$\text{ACT1} \frac{\text{std}(R)}{\langle a, \lambda \rangle . R \xrightarrow{\langle a, \lambda \rangle [i]} \langle a, \lambda \rangle [i] . R}$	$\text{ACT1}^\bullet \frac{\text{std}(R)}{\langle a, \lambda \rangle [i] . R \dashrightarrow \langle a, \lambda \rangle . R}$
$\text{ACT2} \frac{R \xrightarrow{\langle b, \mu \rangle [j]} R' \quad j \neq i}{\langle a, \lambda \rangle [i] . R \xrightarrow{\langle b, \mu \rangle [j]} \langle a, \lambda \rangle [i] . R'}$	$\text{ACT2}^\bullet \frac{R \dashrightarrow \langle b, \bar{\mu} \rangle [j] R' \quad j \neq i}{\langle a, \lambda \rangle [i] . R \dashrightarrow \langle a, \lambda \rangle [i] . R'}$
$\text{CHO} \frac{R \xrightarrow{\langle a, \lambda \rangle [i]} R' \quad \text{std}(S)}{R + S \xrightarrow{\langle a, \lambda \rangle [i]} R' + S}$	$\text{CHO}^\bullet \frac{R \dashrightarrow \langle a, \bar{\lambda} \rangle [i] R' \quad \text{std}(S)}{R + S \dashrightarrow R' + S}$
$\text{PAR} \frac{R \xrightarrow{\langle a, \lambda \rangle [i]} R' \quad a \notin L \quad i \notin \text{key}(S)}{R \parallel_L S \xrightarrow{\langle a, \lambda \rangle [i]} R' \parallel_L S}$	$\text{PAR}^\bullet \frac{R \dashrightarrow \langle a, \bar{\lambda} \rangle [i] R' \quad a \notin L \quad i \notin \text{key}(S)}{R \parallel_L S \dashrightarrow R' \parallel_L S}$
$\text{COO} \frac{R \xrightarrow{\langle a, \lambda \rangle [i]} R' \quad S \xrightarrow{\langle a, \mu \rangle [i]} S' \quad a \in L}{R \parallel_L S \xrightarrow{\langle a, \lambda \cdot \mu \rangle [i]} R' \parallel_L S'}$	$\text{COO}^\bullet \frac{R \dashrightarrow \langle a, \bar{\lambda} \rangle [i] R' \quad S \dashrightarrow \langle a, \bar{\mu} \rangle [i] S' \quad a \in L}{R \parallel_L S \dashrightarrow R' \parallel_L S'}$

Table 2: Structural operational semantic rules for RMPC

We denote by \mathcal{P} the set of processes generated by the production for R in Table 1, while we use predicate $\text{std}(\cdot)$ to identify the standard forward processes that can be derived from the production for P in the same table. For example, $\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0}$ is a standard forward process that can perform action a at rate λ followed by action b at rate μ , while $\langle a, \lambda \rangle [i] . \langle b, \mu \rangle . \underline{0}$ is a non-standard reversible process that can either undo action a at rate $\bar{\lambda}$ or perform action b at rate μ . Note that $\langle a, \lambda \rangle . \langle b, \mu \rangle [j] . \underline{0} \notin \mathcal{P}$ because a future action cannot precede a past action in the description of the behavior of a process.

The set $\text{key}(R)$ of keys occurring in a process $R \in \mathcal{P}$ is inductively defined as follows:

$$\begin{aligned} \text{key}(P) &= \emptyset \\ \text{key}(\langle a, \lambda \rangle [i] . R) &= \{i\} \cup \text{key}(R) \\ \text{key}(R + S) &= \text{key}(R) \cup \text{key}(S) \\ \text{key}(R \parallel_L S) &= \text{key}(R) \cup \text{key}(S) \end{aligned}$$

4.2. Semantics for RMPC. The semantics for RMPC is the labeled transition system $(\mathcal{P}, \mathcal{L}, \mapsto)$. As in [PU07], the transition relation $\mapsto \subseteq \mathcal{P} \times \mathcal{L} \times \mathcal{P}$ is given by $\mapsto = \longrightarrow \cup \dashrightarrow$, where the *forward transition relation* \longrightarrow and the *backward transition relation* \dashrightarrow are the least relations respectively induced by the forward rules in the left part of Table 2 and by the backward rules in the right part of the same table.

Rule ACT1 handles processes of the form $\langle a, \lambda \rangle . P$, where P is written as R subject to $\text{std}(R)$. In addition to transforming the action prefix into a transition label, it generates a key i that is bound to the action $\langle a, \lambda \rangle$ thus yielding the label $\langle a, \lambda \rangle [i]$. As can be noted, the prefix is not discarded by the application of the rule, instead it becomes a key-storing part of the target process that is necessary to offer again that action after coming back. Rule ACT1 $^\bullet$ reverts the action $\langle a, \lambda \rangle [i]$ of the process $\langle a, \lambda \rangle [i] . R$ provided that R is a

standard process, which ensures that $\langle a, \lambda \rangle [i]$ is the only past action that is left to undo. One of the main design choices of the entire framework is how the rate $\bar{\lambda}$ of the backward action is calculated. For the time being, we leave it unspecified in ACT1[•] as the value of this rate is not needed to prove the causal consistency part of reversibility; we will see later on that it may be important to achieve time reversibility.

The presence of rules ACT2 and ACT2[•] is motivated by the fact that rule ACT1 does not discard the executed prefix from the process it generates. In particular, rule ACT2 allows a process $\langle a, \lambda \rangle [i].R$ to execute if R itself can execute provided that the action performed by R picks a key j different from i , so that all the action prefixes in a sequence are decorated with distinct keys. Rule ACT2[•] simply propagates the execution of backward actions from inner subprocesses that are not standard as long as key uniqueness is preserved, in such a way that past actions are overall undone from the most recent one to the least recent one.

Unlike the classical rules for the choice operator [Mil89], rule CHO does not discard the part of the overall process that has not contributed to the executed action. More in detail, if process R does an action, say $\langle a, \lambda \rangle [i]$, and becomes R' , then the entire process $R + S$ becomes $R' + S$ as the information about $+ S$ is necessary for offering again the original choice after coming back. Once the choice is made, only the non-standard process R' can proceed further, with the standard process S constituting a dead context of R' . Note that, in order to apply rule CHO, at least one of R and S has to be standard, meaning that it is impossible for two non-standard processes to execute if they are composed by a choice operator. Rule CHO[•] has precisely the same structure as rule CHO, but deals with the backward transition relation; if R' is standard, then the dead context S will come into play again. For both rules, we omit their symmetric variants in which it is S to move.

The semantics of parallel composition is inspired by [Hoa85]. Rule PAR allows process R within $R \parallel_L S$ to individually perform an action $\langle a, \lambda \rangle [i]$ provided that $a \notin L$. It is also checked that the executing action is bound to a fresh key $i \notin \text{key}(S)$, thus ensuring the uniqueness of communication keys across parallel composition too. Rule COO instead allows both R and S to move by synchronizing on any action in the set L , provided that the communication key is the same on both sides, i.e., $i \in \text{key}(R') \cap \text{key}(S')$. The resulting cooperation action has the same name and the same key and is assumed to be exponentially distributed with rate given by the product of the rates of the two involved actions [Hil94]. Rules PAR[•] and COO[•] respectively have the same structure as PAR and COO; the symmetric variants of PAR and PAR[•] are omitted.

In our stochastic setting, choice and parallel composition are not subject to non-determinism. The decision of which action is chosen or which process advances is made probabilistically based on the rates of the actions that are ready to execute. This mechanism is called the *race policy* [Hil96] because the higher the rate, i.e., the faster the action, the higher its execution probability. This is consistent with the fact that the minimum of a set of exponentially distributed random variables is still exponentially distributed, from which it follows that, as already mentioned in Section 3, the sojourn time in any state of a CTMC is exponentially distributed with rate given by the sum of the rates of the outgoing transitions. The average sojourn time in that state is the inverse of such a sum and the probability of moving from that state to another one is the ratio of the rate of the corresponding transition to the aforementioned sum.

We further recall from [Hil96] that the race policy and the memoryless property of CTMCs allow the *interleaving view of concurrency* – typically employed in nondeterministic process calculi – to be adopted in the setting of a stochastic process calculus like RMPC.

Consider the standard process $\langle a, \lambda \rangle.\underline{0} \parallel_{\emptyset} \langle b, \mu \rangle.\underline{0}$. It can initially execute either action a at rate λ , which happens with probability $\frac{\lambda}{\lambda+\mu}$, or action b at rate μ , which happens with probability $\frac{\mu}{\lambda+\mu}$; the average sojourn time in the corresponding state is $\frac{1}{\lambda+\mu}$. Observing that the probability of simultaneous termination of the two actions is zero, suppose that a terminates first. In the state reached by the corresponding forward transition labeled with a and λ the execution of b is still in progress, so there will be an outgoing forward transition labeled with b and a parameter quantifying the residual duration of b . That parameter is still μ because, thanks to the memoryless property, the residual time to the termination of b is still exponentially distributed with rate μ (see the left part of the forthcoming Figure 3).

4.3. Reachable Processes. The syntax in Table 1 prevents future actions from preceding past actions. However, this is not the only necessary limitation, because not all the processes generated by the considered grammar are semantically meaningful. On the one hand, in the case of a choice at least one of the two subprocesses must be standard, hence for instance $\langle a, \lambda \rangle.[i].\underline{0} + \langle b, \mu \rangle.[j].\underline{0}$ is not admissible as it indicates that both branches have been selected. On the other hand, key uniqueness must be enforced within non-standard processes, so for example $\langle a, \lambda \rangle.[i].\langle b, \mu \rangle.[i].\underline{0}$ and $\langle a, \lambda \rangle.[i].\underline{0} \parallel_{\emptyset} \langle b, \mu \rangle.[i].\underline{0}$ are not admissible either.

In the rest of the paper, we thus consider only *reachable processes*, whose set we denote by \mathbb{P} . They include processes from which a computation can start, i.e., standard forward processes, as well as processes that can be derived from the previous ones via finitely many applications of the rules for \longrightarrow in Table 2. Given a reachable process $R \in \mathbb{P}$, we observe that if $\text{std}(R)$ then $\text{key}(R) = \emptyset$ while any other process R' reachable from R is such that $\text{key}(R') \neq \emptyset$, as at least one of the actions occurring in R has been executed and hence it has been equipped with a communication key inside R' .

4.4. Preliminary Reversibility Properties. A basic property in order for RMPC to be reversible, both in the causal sense and in the time one, is the so called *loop lemma* [DK04, PU07]. This property states that each executed action can be undone and that any undone action can be redone. In other words, when considering the states associated with two arbitrary reachable processes, either there is no transition between them, or there is a pair of transitions such that one is a forward transition from the first state to the second state while the other is a backward transition from the second state to the first state.

Lemma 4.1 (*loop lemma*). *Let $R, S \in \mathbb{P}$. Then $R \xrightarrow{\langle a, \lambda \rangle.[i]} S$ iff $S \xrightarrow{\langle a, \bar{\lambda} \rangle.[i]} R$.*

Proof. By induction on the depth of the derivation of the considered transition by noting that each forward (resp. backward) rule has a corresponding backward (resp. forward) rule. \square

Given a sequence $\sigma = \ell_1 \dots \ell_n$ of $n \in \mathbb{N}_{>0}$ labels, $R \xrightarrow{\sigma} S$ denotes a sequence of forward transitions $R \xrightarrow{\ell_1} R_1 \xrightarrow{\ell_2} \dots \xrightarrow{\ell_n} S$ labeled with that sequence. If $\bar{\sigma} = \bar{\ell}_n \dots \bar{\ell}_1$ is the corresponding backward sequence, then for each ℓ_i occurring in σ it holds that $R_{i-1} \xrightarrow{\ell_i} R_i$ iff $R_i \xrightarrow{\bar{\ell}_i} R_{i-1}$. The loop lemma thus generalizes as follows.

Corollary 4.2. *Let $R, S \in \mathbb{P}$. Then $R \xrightarrow{\sigma} S$ iff $S \xrightarrow{\bar{\sigma}} R$.* \blacksquare

4.5. Causal Reversibility of RMPC. We now prove the causal reversibility of RMPC, which means that each of its reachable processes is able to backtrack *correctly*, i.e., without encountering previously inaccessible states, and *flexibly*, i.e., along any causally equivalent path. To this purpose, we start by borrowing some machinery from [DK04], in particular the notion of concurrent transitions, that needs to be adapted because the reversing method of [PU07] we have followed is different from the one of [DK04] we are going to exploit.

Given a transition $\theta = R \xrightarrow{\ell} S$ with $R, S \in \mathbb{P}$, we call R the *source* of θ and S its *target*. If θ is a forward transition, i.e., $\theta = R \xrightarrow{\ell} S$, we denote by $\bar{\theta} = S \xrightarrow{\bar{\ell}} R$ the corresponding backward transition. Two transitions are said to be *cointial* if they have the same source and *cofinal* if they have the same target. Two transitions are *composable* when the target of the first transition coincides with the source of the second transition. A finite sequence of pairwise composable transitions is called a *computation*. We use ϵ for the empty computation and ω to range over computations, with $|\omega|$ denoting the length of ω expressed as the number of transitions constituting it. When ω is a forward computation, we denote by $\bar{\omega}$ the corresponding backward computation. The notions of source, target, cointiality, cofinality, and composability naturally extend to computations. We indicate with $\omega_1\omega_2$ the composition of the two computations ω_1 and ω_2 when they are composable.

Before specifying when two transitions are concurrent, we need to present the set of causes – identified by keys – that lead to a given communication key, along with the notion of process context.

The *causal set* $\text{cau}(R, i)$ of process $R \in \mathbb{P}$ until key $i \in \mathcal{K}$ is inductively defined as:

$$\begin{aligned} \text{cau}(P, i) &= \emptyset \\ \text{cau}(\langle a, \lambda \rangle [j].R, i) &= \begin{cases} \emptyset & \text{if } j = i \text{ or } i \notin \text{key}(R) \\ \{j\} \cup \text{cau}(R, i) & \text{otherwise} \end{cases} \\ \text{cau}(R + S, i) &= \text{cau}(R, i) \cup \text{cau}(S, i) \\ \text{cau}(R \parallel_L S, i) &= \text{cau}(R, i) \cup \text{cau}(S, i) \end{aligned}$$

If $i \in \text{key}(R)$, then $\text{cau}(R, i)$ represents the set of keys in R that caused i , with $\text{cau}(R, i) \subset \text{key}(R)$ because on the one hand $i \notin \text{cau}(R, i)$ and on the other hand keys that are not causally related to i are not considered. A key j causes i if it appears syntactically before i in R or, said otherwise, i is inside the scope of j .

A *process context* is a process with a hole in it, generated by the grammar:

$$\mathcal{C} ::= \bullet \mid \langle a, \lambda \rangle [i].\mathcal{C} \mid R + \mathcal{C} \mid \mathcal{C} + R \mid R \parallel_L \mathcal{C} \mid \mathcal{C} \parallel_L R$$

We write $\mathcal{C}[R]$ to denote the process obtained from \mathcal{C} by replacing its hole with R .

We are now in a position to define what we mean by concurrent transitions on the basis of the notion of conflicting transitions. The first condition below tells that a forward transition is in conflict with a backward one whenever the latter tries to undo a cause of the key of the former. The second condition below deems as conflictual two transitions respectively generated by the two subprocesses of a choice.

Definition 4.3 (*conflicting and concurrent transitions*). Two cointial transitions θ_1 and θ_2 from a process $R \in \mathbb{P}$ are in conflict iff one of the following two conditions holds:

- (1) $\theta_1 = R \xrightarrow{\langle a, \lambda \rangle [i]} S_1$ and $\theta_2 = R \xrightarrow{\langle b, \bar{\mu} \rangle [j]} S_2$ with $j \in \text{cau}(S_1, i)$.
- (2) $R = \mathcal{C}[P_1 + P_2]$ with θ_k deriving from $P_k \xrightarrow{\langle a_k, \lambda_k \rangle [i_k]} S_k$ for $k = 1, 2$.

Two cointial transitions are concurrent when they are not in conflict. ■

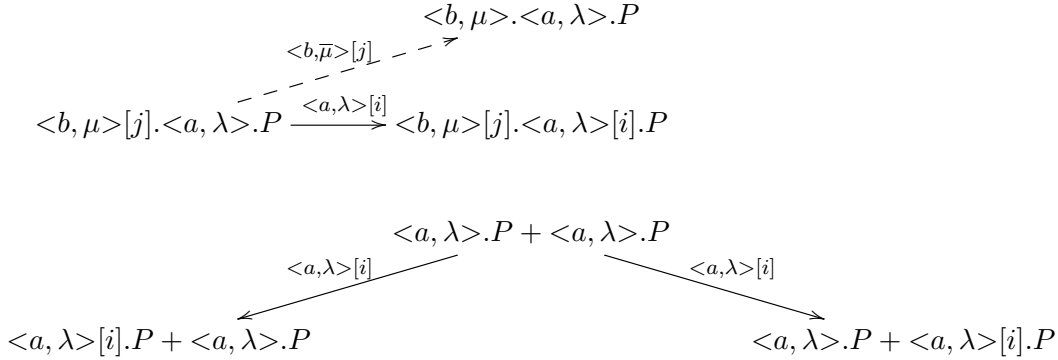


Figure 2: Two examples of conflicting transitions

Figure 2 shows two related examples. In the former, the process $\langle b, \mu \rangle[j].\langle a, \lambda \rangle.P$ can perform two transitions: a backward one and a forward one. They are in conflict according to the first condition of Definition 4.3 as the backward transition removes the key j that is in the causal set of i . In the latter, we have that process $\langle a, \lambda \rangle.P + \langle a, \lambda \rangle.P$ is able to trigger two forward transitions. Since they arise from the same choice operator, they are in conflict according to the second condition of Definition 4.3.

Remark 4.4. In a stochastic process calculus like RMPC the semantic treatment of the choice operator is not trivial because for a process of the form $\langle a, \lambda \rangle.P + \langle a, \lambda \rangle.P$ the operational rules should produce either a single a -transition whose rate is $\lambda + \lambda$, or two a -transitions each having rate λ that do not collapse into a single one [Hil96]. In our reversible framework, two distinct transitions are naturally generated thanks to the fact that the key associated with the executed action and the discarded alternative subprocess are stored into the reached process, as shown in the bottom part of Figure 2. ■

We finally show that reversibility is causally consistent in our concurrent framework. This can be done in two ways: either by adapting the original proof of [DK04], as we did in [BM20], or by using the general technique recently provided by [LPU20]. We opt for the latter, according to which causal consistency stems from the *diamond property* – which amounts to concurrent transitions being confluent – *backward transition independence* – which generalizes the concept of backward determinism used for reversible sequential languages [YG07] – and *past well foundedness* – which ensures that reachable processes have a finite past.

Concurrent transitions can commute as formally stated below by the diamond property (as an example see the left part of the forthcoming Figure 3), while conflicting ones cannot.

Lemma 4.5 (*diamond property*). *Let $\theta_1 = R \xrightarrow{\ell_1} S_1$ and $\theta_2 = R \xrightarrow{\ell_2} S_2$ be two coinitial transitions from a process $R \in \mathbb{P}$. If θ_1 and θ_2 are concurrent, then there exist two cofinal transitions $\theta'_2 = S_1 \xrightarrow{\ell_2} S$ and $\theta'_1 = S_2 \xrightarrow{\ell_1} S$ with $S \in \mathbb{P}$.*

Proof. The proof is by case analysis on the direction of θ_1 and θ_2 . We distinguish three cases according to whether the two transitions are both forward, both backward, or one forward and the other backward. Suppose that $\ell_1 = \langle a, \lambda \rangle[i]$ and $\ell_2 = \langle b, \mu \rangle[j]$, with $i \neq j$ otherwise θ_1 and θ_2 would be generated by the two subprocesses of a choice operator and hence could not be concurrent:

- Suppose that θ_1 and θ_2 are both forward. Since θ_1 and θ_2 are concurrent, by virtue of Definition 4.3 the two transitions cannot originate from a choice operator. They must thus be generated by a parallel composition, but not through the COO rule because θ_1 and θ_2 have different keys and hence cannot synchronize. Without loss of generality, we can assume that $R = R_1 \parallel_L R_2$ with $R_1 \xrightarrow{\langle a, \lambda \rangle [i]} S_1$ and $R_2 \xrightarrow{\langle b, \mu \rangle [j]} S_2$ and $a, b \notin L$. By applying the PAR rule, we have that $R_1 \parallel_L R_2 \xrightarrow{\langle a, \lambda \rangle [i]} S_1 \parallel_L R_2 \xrightarrow{\langle b, \mu \rangle [j]} S_1 \parallel_L S_2$ as well as $R_1 \parallel_L R_2 \xrightarrow{\langle b, \mu \rangle [j]} R_1 \parallel_L S_2 \xrightarrow{\langle a, \lambda \rangle [i]} S_1 \parallel_L S_2$.
- The case in which θ_1 and θ_2 are both backward is similar to the previous one.
- Suppose that θ_1 is forward and θ_2 is backward. Since θ_1 and θ_2 are concurrent, by virtue of Definition 4.3 the backward transition cannot remove a cause of the forward one. Since either subprocess of a choice operator or a parallel composition cannot perform a forward transition and a backward transition without preventing the backward one from removing a cause of the forward one, and in the case of the choice operator only one of the two subprocesses can perform transitions, without loss of generality we can assume that $R = R_1 \parallel_L R_2$ with $R_1 \xrightarrow{\langle a, \lambda \rangle [i]} S_1$ and $R_2 \xrightarrow{\langle b, \bar{\mu} \rangle [j]} S_2$ as well as $a, b \notin L$ so as to preserve causes. By applying the PAR rule, we have that $R_1 \parallel_L R_2 \xrightarrow{\langle a, \lambda \rangle [i]} S_1 \parallel_L R_2 \xrightarrow{\langle b, \bar{\mu} \rangle [j]} S_1 \parallel_L S_2$ as well as $R_1 \parallel_L R_2 \xrightarrow{\langle b, \bar{\mu} \rangle [j]} R_1 \parallel_L S_2 \xrightarrow{\langle a, \lambda \rangle [i]} S_1 \parallel_L S_2$. \square

Lemma 4.6 (*backward transition independence*). *Let $R \in \mathbb{P}$. Then any two cointial backward transitions $\theta_1 = R \xrightarrow{\langle a, \bar{\lambda} \rangle [i]} S_1$ and $\theta_2 = R \xrightarrow{\langle b, \bar{\mu} \rangle [j]} S_2$ are concurrent.*

Proof. Since by Definition 4.3 there is no case in which two backward transitions are conflicting, the property trivially holds. \square

Lemma 4.7 (*past well foundedness*). *Let $R_0 \in \mathbb{P}$. Then there is no infinite sequence of backward transitions such that $R_i \xrightarrow{\ell_i} R_{i+1}$ for all $i \in \mathbb{N}$.*

Proof. It can be easily proved by induction on $|\text{key}(R_0)|$ by observing that a backward transition decreases by one the total number of keys of R_0 , with this number being finite. \square

Following [DK04, Lév76], we also define a notion of *causal equivalence* over computations, which abstracts from the order of concurrent transitions. In this way, computations obtained by swapping the order of their concurrent transitions are identified with each other (see for instance the left part of the forthcoming Figure 3) and the composition of a computation with its inverse is identified with the empty computation.

Definition 4.8 (*causal equivalence*). Causal equivalence is the smallest equivalence relation \simeq on computations that is closed under composition and satisfies the following:

- (1) $\theta_1 \theta'_2 \simeq \theta_2 \theta'_1$ for any two cointial concurrent transitions $\theta_1 = R \xrightarrow{\ell_1} R_1$ and $\theta_2 = R \xrightarrow{\ell_2} R_2$ and any two cofinal transitions $\theta'_2 = R_1 \xrightarrow{\ell'_2} S$ and $\theta'_1 = R_2 \xrightarrow{\ell'_1} S$ respectively composable with the previous ones.
- (2) $\theta \bar{\theta} \simeq \epsilon$ and $\bar{\theta} \theta \simeq \epsilon$. \blacksquare

The further property below, called *parabolic lemma* in [LPU20], states that any computation can be seen as a backward computation followed by a forward computation. As observed in [DK04], up to causal equivalence one can always reach for the maximum freedom of choice among transitions by going backward and only then going forward (not the other

way around). Intuitively, one could depict computations as parabolas: the system first draws potential energy from its memory, by undoing all the executed actions, and then restarts.

Lemma 4.9 (*parabolic lemma*). *For any computation ω , there exist two forward computations ω_1 and ω_2 such that $\omega \asymp \bar{\omega}_1\omega_2$ and $|\omega_1| + |\omega_2| \leq |\omega|$.*

Proof. It follows from the diamond property and backward transition independence thanks to [LPU20]. \square

Theorem 4.10 (*causal consistency*). *Let ω_1 and ω_2 be two computations. Then $\omega_1 \asymp \omega_2$ iff ω_1 and ω_2 are both coinital and cofinal.*

Proof. It follows from past well foundedness and the parabolic lemma thanks to [LPU20]. \square

Theorem 4.10 shows that causal equivalence characterizes a space for admissible rollbacks that are (i) correct as they do not lead to states not reachable by some forward computation and (ii) flexible enough to allow undo operations to be rearranged with respect to the order in which the undone concurrent transitions were originally performed. This implies that the states reached by any backward computation could be reached by performing forward computations only. Therefore, we can conclude that RMPC meets causal reversibility.

4.6. Time Reversibility of RMPC. The rules in Table 2 associate with any process $R \in \mathbb{P}$ a labeled transition system $\llbracket R \rrbracket = (\mathbb{P}, \mathcal{L}, \mapsto, R)$ whose initial state corresponds to R . To investigate time reversibility, we have to derive from $\llbracket R \rrbracket$ the underlying CTMC $\mathcal{M}[\llbracket R \rrbracket]$ and we have to specify the value of any rate $\bar{\lambda}$ labeling a backward transition.

First of all, we observe that every state of $\llbracket R \rrbracket$ with an outgoing forward transition actually has infinitely many copies of that transition. The motivation is that rules ACT1 and ACT2 generate a transition for each possible admissible key, with the key being part of both the label and the reached process. For example, $\langle a, \lambda \rangle.P$ has an outgoing forward transition towards $\langle a, \lambda \rangle[i].P$ for each $i \in \mathcal{K}$; note that from each $\langle a, \lambda \rangle[i].P$ there is instead a single outgoing backward transition to $\langle a, \lambda \rangle.P$. When building the CTMC underlying $\langle a, \lambda \rangle.P$, only one of those outgoing forward transitions has to be considered, otherwise the corresponding state would have an infinite exit rate.

In the construction of $\mathcal{M}[\llbracket R \rrbracket]$ we thus need to reason in terms of transition bundles. A *transition bundle* collects all the transitions departing from the same state and labeled with the same action, the same rate, and different keys, whose target states are syntactically identical up to keys in the same positions. To this purpose, we make use of contexts to introduce $\equiv_{\mathcal{K}}$ as the smallest equivalence relation on \mathbb{P} that satisfies the following:

- If \mathcal{C} does not contain occurrences of $\|_L$ with $a \in L$ and i and j do not occur in \mathcal{C} and R , then $\mathcal{C}[\langle a, \lambda \rangle[i].R] \equiv_{\mathcal{K}} \mathcal{C}[\langle a, \lambda \rangle[j].R]$.
- If $a \in L_1 \cap \dots \cap L_{n-1}$ for $n \geq 2$, i and j do not occur in $\mathcal{C}, \mathcal{C}_1, \dots, \mathcal{C}_n, R_1, \dots, R_n$, and no further a with key i occurs in \mathcal{C} , then $\mathcal{C}[\mathcal{C}_1[\langle a, \lambda_1 \rangle[i].R_1] \|_{L_1} \dots \|_{L_{n-1}} \mathcal{C}_n[\langle a, \lambda_n \rangle[i].R_n]] \equiv_{\mathcal{K}} \mathcal{C}[\mathcal{C}_1[\langle a, \lambda_1 \rangle[j].R_1] \|_{L_1} \dots \|_{L_{n-1}} \mathcal{C}_n[\langle a, \lambda_n \rangle[j].R_n]]$.

As for the second case, for instance $\langle a, \lambda_1 \rangle.P_1 \|_{\{a\}} \langle a, \lambda_2 \rangle.P_2$ evolves to $\langle a, \lambda_1 \rangle[i].P_1 \|_{\{a\}} \langle a, \lambda_2 \rangle[i].P_2$ for all $i \in \mathcal{K}$, but we cannot replace $\langle a, \lambda_1 \rangle[i].P_1$ with $\langle a, \lambda_1 \rangle[j].P_1$, unless we replace $\langle a, \lambda_2 \rangle[i].P_2$ with $\langle a, \lambda_2 \rangle[j].P_2$ too, because $\langle a, \lambda_1 \rangle[j].P_1 \|_{\{a\}} \langle a, \lambda_2 \rangle[i].P_2$ has no backward transition to $\langle a, \lambda_1 \rangle.P_1 \|_{\{a\}} \langle a, \lambda_2 \rangle.P_2$ when $j \neq i$.

Among the representations mentioned in Section 3, below we choose to formalize $\mathcal{M}[\llbracket R \rrbracket]$ as a labeled transition system because this is closer to $\llbracket R \rrbracket$ obtained from the rules in Table 2.

Definition 4.11 (*underlying CTMC*). The CTMC underlying a process $R \in \mathbb{P}$ is the labeled transition system $\mathcal{M}[[R]] = (\mathbb{P}/\equiv_{\mathcal{K}}, \mathcal{A} \times \mathcal{R}, \mapsto_{\mathcal{K}}, [R]_{\equiv_{\mathcal{K}}})$ where:

- $\mathbb{P}/\equiv_{\mathcal{K}}$ is the quotient set of $\equiv_{\mathcal{K}}$ over \mathbb{P} , i.e., the set of classes of $\equiv_{\mathcal{K}}$ -equivalent processes, representing the set of states.
- $[R]_{\equiv_{\mathcal{K}}}$ is the $\equiv_{\mathcal{K}}$ -equivalence class of R , which simply is $\{R\}$ when R is standard and hence contains no keys, representing the initial state.
- $\mapsto_{\mathcal{K}} \subseteq (\mathbb{P}/\equiv_{\mathcal{K}}) \times (\mathcal{A} \times \mathcal{R}) \times (\mathbb{P}/\equiv_{\mathcal{K}})$ is the transition relation given by $\mapsto_{\mathcal{K}} = \longrightarrow_{\mathcal{K}} \cup \dashrightarrow_{\mathcal{K}}$:
 - $[S]_{\equiv_{\mathcal{K}}} \xrightarrow{\langle a, \lambda \rangle}_{\mathcal{K}} [S']_{\equiv_{\mathcal{K}}}$ whenever $S \xrightarrow{\langle a, \lambda \rangle [i]} S'$ for some $i \in \mathcal{K}$.
 - $[S]_{\equiv_{\mathcal{K}}} \dashrightarrow_{\mathcal{K}} [S']_{\equiv_{\mathcal{K}}}$ whenever $S \dashrightarrow_{\mathcal{K}} S'$ for some $i \in \mathcal{K}$. ■

When moving from $[[R]]$ to $\mathcal{M}[[R]]$, individual states are thus replaced by classes of states that are syntactically identical up to keys in the same positions; moreover, keys are removed from transition labels. As a consequence, every state of $\mathcal{M}[[R]]$ turns out to have finitely many outgoing transitions. We also note that $\mathcal{M}[[R]]$ is an action-labeled CTMC, as each of its transitions is labeled not only with a rate but also with an action.

As a preliminary step towards time reversibility, we have to show that $\mathcal{M}[[R]]$ is stationary. It holds that $\mathcal{M}[[R]]$ is even ergodic thanks to the loop lemma.

Lemma 4.12. *Let $R \in \mathbb{P}$. Then $\mathcal{M}[[R]]$ is time homogeneous and ergodic.*

Proof. The time homogeneity of $\mathcal{M}[[R]]$ is a straightforward consequence of the fact that its rates simply are positive real numbers, not time-dependent functions. The ergodicity of $\mathcal{M}[[R]]$ stems from the fact that the graph representing $\mathcal{M}[[R]]$ is a finite strongly connected component due to Corollary 4.2. □

The proof of time reversibility exploits the necessary and sufficient condition based on partial balance equations [Kel79] (see Section 3). The loop lemma and the assumption $\bar{\lambda} = \lambda$ for all $\lambda \in \mathbb{R}_{>0}$ ensure that the steady-state probability distribution of $\mathcal{M}[[R]]$ is the uniform distribution. From this it immediately follows that the partial balance equations are satisfied, i.e., that time reversibility holds.

Lemma 4.13. *Let $R \in \mathbb{P}$, \mathcal{S} be the set of states of $\mathcal{M}[[R]]$, and $n = |\mathcal{S}|$. If every backward rate is equal to the corresponding forward rate, then the steady-state probability distribution π of $\mathcal{M}[[R]]$ satisfies $\pi(s) = 1/n$ for all $s \in \mathcal{S}$.*

Proof. If $n = 1$, i.e., R is equal to $\underline{0}$ or to the parallel composition of several standard processes whose initial actions have to synchronize but are different from each other, then trivially $\pi(s) = 1/n = 1$ for the only state $s \in \mathcal{S}$.

Suppose now that $n \geq 2$. From Lemma 4.12 it follows that $\mathcal{M}[[R]]$ has a unique steady-state probability distribution π . Due to Lemma 4.1, each outgoing forward/backward transition of an arbitrary state $s \in \mathcal{S}$ has a corresponding incoming backward/forward transition, hence the global balance equation for s reformulated in terms of transitions looks as follows, where s' (resp., s'') is the target state of a forward (resp., backward) transition departing from s :

$$\pi(s) \cdot \left(\sum_{s \xrightarrow{\langle a', \lambda' \rangle}_{\mathcal{K}} s'} \lambda' + \sum_{s \dashrightarrow_{\mathcal{K}} s''} \bar{\lambda}'' \right) = \sum_{s' \dashrightarrow_{\mathcal{K}} s} \pi(s') \cdot \bar{\lambda}' + \sum_{s'' \xrightarrow{\langle a'', \lambda'' \rangle}_{\mathcal{K}} s} \pi(s'') \cdot \lambda''$$

Since every backward rate is equal to the corresponding forward rate, the global balance equation for s boils down to:

$$\pi(s) \cdot \left(\sum_{s \xrightarrow{\langle a', \lambda' \rangle}_{\mathcal{K}} s'} \lambda' + \sum_{s \dashrightarrow_{\mathcal{K}} s''} \lambda'' \right) = \sum_{s' \dashrightarrow_{\mathcal{K}} s} \pi(s') \cdot \lambda' + \sum_{s'' \xrightarrow{\langle a'', \lambda'' \rangle}_{\mathcal{K}} s} \pi(s'') \cdot \lambda''$$

Since the two summations related to s' as well as the two summations related to s'' have the same number of summands, the equation above is satisfied when $\pi(s) = \pi(s') = \pi(s'')$ for each s' reached by a forward transition of s and each s'' reached by a backward transition of s . All global balance equations are thus satisfied when $\pi(s) = 1/n$ for all $s \in \mathcal{S}$. \square

Theorem 4.14 (*time reversibility 1*). *Let $R \in \mathbb{P}$. If every backward rate is equal to the corresponding forward rate, then $\mathcal{M}[[R]]$ is time reversible.*

Proof. Let \mathcal{S} be the set of states of $\mathcal{M}[[R]]$ and $n = |\mathcal{S}|$; to avoid trivial cases, suppose $n \geq 2$. From Lemma 4.12 it follows that $\mathcal{M}[[R]]$ has a unique steady-state probability distribution π . Now consider $s, s' \in \mathcal{S}$ with $s \neq s'$ connected by transitions. The proof resembles the one of Lemma 4.13, but focuses on partial balance equations rather than on global ones.

Due to Lemma 4.1, the partial balance equation for s and s' reformulated in terms of transitions looks as follows:

$$\pi(s) \cdot \left(\sum_{s \xrightarrow{\langle a', \lambda' \rangle} \mathcal{K}s'} \lambda' + \sum_{s \xrightarrow{\langle a'', \bar{\lambda}'' \rangle} \mathcal{K}s'} \bar{\lambda}'' \right) = \pi(s') \cdot \left(\sum_{s' \xrightarrow{\langle a', \bar{\lambda}' \rangle} \mathcal{K}s} \bar{\lambda}' + \sum_{s' \xrightarrow{\langle a'', \lambda'' \rangle} \mathcal{K}s} \lambda'' \right)$$

Since every backward rate is equal to the corresponding forward rate, the partial balance equation for s and s' boils down to:

$$\pi(s) \cdot \left(\sum_{s \xrightarrow{\langle a', \lambda' \rangle} \mathcal{K}s'} \lambda' + \sum_{s \xrightarrow{\langle a'', \lambda'' \rangle} \mathcal{K}s'} \lambda'' \right) = \pi(s') \cdot \left(\sum_{s' \xrightarrow{\langle a', \lambda' \rangle} \mathcal{K}s} \lambda' + \sum_{s' \xrightarrow{\langle a'', \lambda'' \rangle} \mathcal{K}s} \lambda'' \right)$$

Since $\pi(s) = \pi(s') = 1/n$ due to Lemma 4.13 and the two left summations on both sides as well as the two right summations on both sides have the same number of summands, the equation above is satisfied. The result then follows from the fact that s and s' are two arbitrary distinct states connected by transitions. \square

The result above holds under the assumption $\bar{\lambda} = \lambda$ for all $\lambda \in \mathbb{R}_{>0}$. This constraint on rates can be relaxed if we exploit the structure of the graph representing $\mathcal{M}[[R]]$. For example, it is well known that CTMCs in the form of stationary birth-death processes are time reversible [Kel79], where a birth-death process comprises a totally ordered set of states such that every state different from the final one has a (birth) transition to the next state and every state different from the initial one has a (death) transition to the previous state. Time reversibility extends to tree-like birth-death processes [Kel79], where each such variant comprises a partially ordered set of states such that every non-final state may have several birth transitions and every non-initial state has one death transition to its only parent state.

Lemma 4.15. *Let $R \in \mathbb{P}$. If parallel composition does not occur in R , then $\mathcal{M}[[R]]$ is a tree-like birth-death process.*

Proof. The fragment of graph representing $\mathcal{M}[[R]]$ that includes only forward transitions, i.e., generated by applying only the rules in the left part of Table 2, is a tree in the absence of parallel composition inside R . In fact, observing that if R is $\underline{0}$ then the graph contains a single state with no transitions, in any state other than $\underline{0}$ there are two options:

- The application of rule ACT1 or ACT2 yields a new transition towards a new state due to the generation of a fresh communication key that is stored within the new state.
- The application of rule CHO yields a single new transition or two new transitions – depending on whether a single subprocess is standard or both subprocesses are standard – each of which reaches a new state due to the generation of a fresh communication key that is stored within the new state and the fact that the dead context – i.e., the standard subprocess that stays idle – is stored within the new state too (this is important in order

for the two new states to be different when the two subprocesses are standard and identical, see the bottom part of Figure 2).

For each forward transition generated by one of the three forward rules mentioned above, the corresponding backward rule generates the corresponding backward transition. Therefore $\mathcal{M}[[R]]$ is a tree-like birth-death process. \square

Theorem 4.16 (*time reversibility 2*). *Let $R \in \mathbb{P}$. If parallel composition does not occur in R , then $\mathcal{M}[[R]]$ is time reversible.*

Proof. A straightforward consequence of Lemma 4.15 above and Lemma 1.5 of [Kel79] establishing the time reversibility of any stationary tree-like birth-death process. Note that unlike $\mathcal{M}[[R]]$ the graph mentioned in the latter lemma is not directed, i.e., each of its edges can be traversed in both directions, but this is the same as having two directed edges in opposite directions like a forward transition and the corresponding backward transition. \square

The time reversibility of CTMC-based compositional models of concurrent systems has already been investigated in [Har03]. That work examines conditions relying on the conservation of total exit rates of states in addition to the conservation of products of rates around cycles [Kel79] (see Section 3), which support the hierarchical and compositional reversal of stochastic process algebra terms. These conditions also lead to the efficient calculation of steady-state probability distributions in a product form typical of queueing theory [Kle75], thus avoiding the need of solving the global balance equations of the entire system. More recently, in [MR15] similar conditions have been employed to characterize the class of ρ -reversible stochastic automata, which allow for state permutations. Under certain constraints, the joint steady-state probability distribution of the composition of two such automata is the product of the steady-state probability distributions of the two automata.

The main difference between our approach to time reversibility and the aforementioned ones is twofold. Firstly, our approach is part of a more general framework in which also causal reversibility is addressed. Secondly, our approach is inspired by the idea of [PU07] of developing a formalism in which it is possible to express models whose reversibility is guaranteed by construction, instead of building a posteriori the time-reversed version of a certain model like in [Har03] or verifying a posteriori whether a given model is time reversible or not like in [MR15].

It is worth noting that these methodological differences do not prevent us from importing in our setting some results from [Har03, MR15], although a few preliminary observations about notational differences are necessary. Both [Har03] and [MR15] make a distinction between active actions, each of which is given a rate, and passive actions, each of which is given a weight, with the constraint that, in case of synchronization, the rate of the active action is multiplied by the weight of the corresponding passive action. In RMPC there is no such distinction, however the same operation, i.e., multiplication, is applied to the rates of two synchronizing actions. A passive action can thus be rendered as an action with rate 1, while a set of alternative passive actions can be rendered as a set of actions whose rates sum up to 1. Moreover, in [MR15] synchronization is enforced between any active-passive pair of identical actions, whereas in RMPC the parallel composition operator is enriched with an explicit synchronization set L , which yields as a special case the aforementioned synchronization discipline when L is equal to the set \mathcal{A} of all the actions. We can therefore conclude that our parallel composition operator is a generalization of those used in [Har03, MR15], hence the recalled notational differences do not hamper the transferral of results.

In [Har03] the compositionality of a CTMC-based stochastic process calculus is exploited to prove RCAT – Reversed Compound Agent Theorem, which establishes the conditions under which the time-reversed version of the cooperation of two processes is equal to the cooperation of the time-reversed versions of those two processes. The application of RCAT leads to product-form solution results from stochastic process algebraic models, including a new different proof of Jackson’s theorem for product-form queueing networks [Jac63].

In [MR15] the notion of ρ -reversibility is introduced for stochastic automata, which are essentially action-labeled CTMCs. Function ρ is a state permutation that ensures (i) for each action the equality of the total exit rate of any state s and $\rho(s)$ and (ii) the conservation of action-related rate products across cycles when considering states in the forward direction and their ρ -counterparts in the backward direction. For any ergodic ρ -reversible automaton, it turns out that $\pi(s) = \pi(\rho(s))$ for every state s . Moreover, the synchronization inspired by [Pla85] of two ρ -reversible stochastic automata is still ρ -reversible and, in case of ergodicity, under certain conditions the steady-state probability of any compound state is the product of the steady-state probabilities of the two constituent states.

Our time reversibility result for RMPC can be rephrased in the setting of [MR15] in terms of ρ -reversibility with ρ being the identity function over states. As a consequence, the following two results stem from Theorems 4.14 and 4.16 of the present paper and, respectively, Theorems 2 and 3 of [MR15].

Corollary 4.17 (*time reversibility closure*). *Let $R_1, R_2 \in \mathbb{P}$ and $L \subseteq \mathcal{A}$. Under the assumptions of Theorem 4.14 or Theorem 4.16, $\mathcal{M}[[R_1 \parallel_L R_2]]$ is time reversible too. ■*

Corollary 4.18 (*product form*). *Let $R_1, R_2 \in \mathbb{P}$ and $L \subseteq \mathcal{A}$. Under the assumptions of Theorem 4.14 or Theorem 4.16, if the set of states \mathcal{S} of $\mathcal{M}[[R_1 \parallel_L R_2]]$ is equal to $\mathcal{S}_{R_1} \times \mathcal{S}_{R_2}$ where \mathcal{S}_{R_k} is the set of states of $\mathcal{M}[[R_k]]$ for $k \in \{1, 2\}$, then $\pi(s_1, s_2) = \pi_{R_1}(s_1) \cdot \pi_{R_2}(s_2)$ for all $(s_1, s_2) \in \mathcal{S}_{R_1} \times \mathcal{S}_{R_2}$. ■*

The product-form result above avoids the calculation of the global balance equations over $\mathcal{M}[[R_1 \parallel_L R_2]]$, as $\pi(s_1, s_2)$ can simply be obtained by multiplying $\pi_{R_1}(s_1)$ with $\pi_{R_2}(s_2)$. However, the condition $\mathcal{S} = \mathcal{S}_{R_1} \times \mathcal{S}_{R_2}$ requires to check that every state in the full Cartesian product is reachable from $R_1 \parallel_L R_2$. This means that no compound state is such that its constituent states enable some action but none of these enabled actions can be executed due to the constraints imposed by the synchronization set L . The condition $\mathcal{S} = \mathcal{S}_{R_1} \times \mathcal{S}_{R_2}$ implies that $\mathcal{M}[[R_1 \parallel_L R_2]]$ is ergodic over the full Cartesian product of the two original state spaces, which is the condition used in [MR15]. Although implicit in the statement of the corollary, the time reversibility of $\mathcal{M}[[R_1 \parallel_L R_2]]$ is essential for the product-form result.

We conclude by observing that an important consequence of time reversibility for processes without parallel composition, i.e., the result of Theorem 4.16, and time reversibility closure with respect to parallel composition, i.e., the result of Corollary 4.17, is the fact that time reversibility extends to the entire sublanguage \mathbb{P}' of RMPC in which parallel composition cannot occur within the scope of action prefix or choice. This is quite useful because systems are typically modeled as the parallel composition of a number of sequential processes, i.e., processes in which only action prefix and choice can occur.

Corollary 4.19 (*time reversibility 3*). *Let $R \in \mathbb{P}'$. Then $\mathcal{M}[[R]]$ is time reversible. ■*

5. FORWARD AND BACKWARD MARKOVIAN BISIMILARITY

In this section, we equip RMPC with a notion of equivalence capable of identifying syntactically different processes that expose the same observable behavior in terms of executable actions and their rates. A well known behavioral equivalence for forward-only stochastic process calculi is Markovian bisimilarity [Hil96]. It equates systems that stepwise mimic each other's functional and performance behavior and enjoys nice properties in terms of compositional reasoning as well as equational and logical characterizations [Ber07]. We show below how it can be adapted to our reversible setting.

In the following, we work at the state space level, rather than at the linguistic level, because we address different variants of Markovian bisimilarity that are defined not only over states, with which process terms are naturally associated, but also over computations, which are more easily expressed through states and transitions. We thus consider an action-labeled CTMC $(\mathcal{S}, \mathcal{A} \times \mathcal{R}, \mapsto)$, where \mathcal{A} is a countable set of actions while $\mathcal{R} = \mathbb{R}_{>0}$ is a set of rates, which in the case of RMPC is obtained by collecting transitions into bundles as formalized in Definition 4.11. Moreover, we use symbols $\{\}$ and $\}\}$ as multiset parentheses.

Definition 5.1 (*Markovian bisimilarity*). Two states $s_1, s_2 \in \mathcal{S}$ are Markovian bisimilar, written $s_1 \sim_{\text{MB}} s_2$, iff there exists a Markovian bisimulation \mathcal{B} such that $(s_1, s_2) \in \mathcal{B}$. An equivalence relation \mathcal{B} over the set of states \mathcal{S} is a Markovian bisimulation iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all actions $a \in \mathcal{A}$ and equivalence classes $C \in \mathcal{S}/\mathcal{B}$:

$$\text{rate}(s_1, a, C) = \text{rate}(s_2, a, C)$$

where $\text{rate}(s, a, C) = \sum \{\lambda \in \mathcal{R} \mid \exists s' \in C. s \xrightarrow{\langle a, \lambda \rangle} s'\}$. ■

Following [PU07], we may adapt Markovian bisimilarity to our reversible calculus by means of two conditions like the one in the definition above, with the former referring to forward transitions and the latter referring to backward transitions. By so doing, we would end up with a very restrictive equivalence, as for instance $\langle a, \lambda \rangle.0 \parallel_{\emptyset} \langle b, \mu \rangle.0$ and $\langle a, \lambda \rangle.\langle b, \mu \rangle.0 + \langle b, \mu \rangle.\langle a, \lambda \rangle.0$ would be told apart. As shown in Figure 3, in the former process from $\langle a, \lambda \rangle[i].0 \parallel_{\emptyset} \langle b, \mu \rangle[j].0$ both a backward a -transition and a backward b -transition are enabled, whilst in the latter process from $\langle a, \lambda \rangle[i].\langle b, \mu \rangle[j].0 + \langle b, \mu \rangle.\langle a, \lambda \rangle.0$ only a backward b -transition is enabled as well as from $\langle a, \lambda \rangle.\langle b, \mu \rangle.0 + \langle b, \mu \rangle[i].\langle a, \lambda \rangle[j].0$ only a backward a -transition is enabled. In other words, the so-called expansion law [Mil89], which transforms a parallel composition into a choice among all of its possible interleaved computations, would not hold, because the two aforementioned transitions in the former process are concurrent and hence when going backward there is no obligation to follow the path traversed in the forward direction, whereas this is not the case with the latter process.

As recognized in [DMV90], in order to set up a more useful equivalence in a reversible setting, it is necessary to enforce not only causality but also history preservation. This means that, when going backward, a process can only move along the path representing the history that brought the process to the current state. For example, if $\langle a, \lambda \rangle.0 \parallel_{\emptyset} \langle b, \mu \rangle.0$ performs a before b , then a and b can be undone in any order from $\langle a, \lambda \rangle[i].0 \parallel_{\emptyset} \langle b, \mu \rangle[j].0$ because there is no causality constraint between the two actions, but history is preserved only if b is undone before a . To accomplish this, bisimilarity has to be defined over computations, not over states, and the notion of transition has to be suitably revised. We start by adapting the notation of the nondeterministic setting of [DMV90] to our stochastic setting.




Figure 3: $\mathcal{M}[\langle a, \lambda \rangle . \underline{0} \parallel_{\emptyset} \langle b, \mu \rangle . \underline{0}]$ and $\mathcal{M}[\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle b, \mu \rangle . \langle a, \lambda \rangle . \underline{0}]$

Definition 5.2 (*path*). A sequence $\xi = (s_0, \langle a_1, \lambda_1 \rangle, s_1) \dots (s_{n-1}, \langle a_n, \lambda_n \rangle, s_n) \in \mapsto^*$ is called a path of length n from state s_0 . We let $\text{first}(\xi) = s_0$ and $\text{last}(\xi) = s_n$; the empty path is indicated with ϵ . We denote by $\text{path}(s)$ the set of paths from state s . ■

Definition 5.3 (*run*). A pair $\rho = (s, \xi)$ is called a run from state s iff $\xi \in \text{path}(s)$, in which case we let $\text{path}(\rho) = \xi$, $\text{first}(\rho) = \text{first}(\xi)$, $\text{last}(\rho) = \text{last}(\xi)$, with $\text{first}(\rho) = \text{last}(\rho) = s$ when $\xi = \epsilon$. We denote by $\text{run}(s)$ the set of runs from state s . Given $\rho = (s, \xi) \in \text{run}(s)$ and $\rho' = (s', \xi') \in \text{run}(s')$, their composition $\rho\rho' = (s, \xi\xi') \in \text{run}(s)$ is defined iff $\text{last}(\rho) = \text{first}(\rho')$. We write $\rho \xrightarrow{\langle a, \lambda \rangle} \rho'$ iff there exists $\rho'' = (s, (s, \langle a, \lambda \rangle, s'))$ with $s = \text{last}(\rho)$ such that $\rho' = \rho\rho''$; note that $\text{first}(\rho) = \text{first}(\rho')$. ■

In the considered action-labeled CTMC $(\mathcal{S}, \mathcal{A} \times \mathcal{R}, \mapsto)$, we work with the set \mathcal{U} of runs in lieu of \mathcal{S} . Furthermore, we view the transition relation \mapsto as being symmetric with respect to source and target states, so that every transition can be traversed in both directions. In the setting of RMPC, this amounts to considering only the forward transition relation thanks to Lemma 4.1. Given a run ρ , an action a , and a bisimulation equivalence class C , based on [DMV90] we distinguish between the total rate of *outgoing* and *incoming* run transitions, respectively, when moving between ρ and C via a . Forward and backward Markovian bisimilarity thus relies on checking both outgoing and incoming rate equalities.

Definition 5.4 (*forward and backward Markovian bisimilarity*). Two states $s_1, s_2 \in \mathcal{S}$ are forward and backward Markovian bisimilar, written $s_1 \sim_{\text{FBMB}} s_2$, iff there exists a forward and backward Markovian bisimulation \mathcal{B} such that $((s_1, \epsilon), (s_2, \epsilon)) \in \mathcal{B}$. An equivalence relation \mathcal{B} over the set of runs \mathcal{U} is a forward and backward Markovian bisimulation iff, whenever $(\rho_1, \rho_2) \in \mathcal{B}$, then for all actions $a \in \mathcal{A}$ and equivalence classes $C \in \mathcal{U}/\mathcal{B}$:

$$\begin{aligned} \text{rate}_0(\rho_1, a, C) &= \text{rate}_0(\rho_2, a, C) \\ \text{rate}_1(\rho_1, a, C) &= \text{rate}_1(\rho_2, a, C) \end{aligned}$$

where:

$$\begin{aligned} \text{rate}_0(\rho, a, C) &= \sum \{ \lambda \in \mathcal{R} \mid \exists \rho' \in C. \rho \xrightarrow{\langle a, \lambda \rangle} \rho' \} \\ \text{rate}_1(\rho, a, C) &= \sum \{ \lambda \in \mathcal{R} \mid \exists \rho' \in C. \rho' \xrightarrow{\langle a, \lambda \rangle} \rho \} \end{aligned} \quad \blacksquare$$

Theorem 5.5. *Let $s_1, s_2 \in \mathcal{S}$. Then $s_1 \sim_{\text{FBMB}} s_2 \implies s_1 \sim_{\text{MB}} s_2$.*

Proof. Suppose that $s_1 \sim_{\text{FBMB}} s_2$ and let \mathcal{B} be a forward and backward Markovian bisimulation on \mathcal{U} such that $((s_1, \epsilon), (s_2, \epsilon)) \in \mathcal{B}$. We show that $\mathcal{B}' = \{(\text{last}(\rho_1), \text{last}(\rho_2)) \mid (\rho_1, \rho_2) \in \mathcal{B}\}$ is a Markovian bisimulation on \mathcal{S} , from which $s_1 \sim_{\text{MB}} s_2$ will follow.

Consider $(\text{last}(\rho_1), \text{last}(\rho_2)) \in \mathcal{B}'$. By definition of \mathcal{B}' , we have that $(\rho_1, \rho_2) \in \mathcal{B}$. Since \mathcal{B} is

a forward and backward Markovian bisimulation, for all $a \in \mathcal{A}$ and $C \in \mathcal{U}/\mathcal{B}$ it holds in particular that $\text{rate}_0(\rho_1, a, C) = \text{rate}_0(\rho_2, a, C)$. Since $\rho_k \xrightarrow{\langle a, \lambda \rangle} \rho'_k$ iff $\text{last}(\rho_k) \xrightarrow{\langle a, \lambda \rangle} \text{last}(\rho'_k)$ for $k \in \{1, 2\}$ and – provided that function last is lifted from runs to sets of runs – any equivalence class $C' \in \mathcal{S}/\mathcal{B}'$ is of the form $[\text{last}(\rho)]_{\mathcal{B}'} = \{\text{last}(\rho') \in \mathcal{S} \mid (\text{last}(\rho), \text{last}(\rho')) \in \mathcal{B}'\} = \text{last}(\{\rho' \in \mathcal{U} \mid (\rho, \rho') \in \mathcal{B}\}) = \text{last}([\rho]_{\mathcal{B}})$, i.e., $C' = \text{last}(C)$ for some equivalence class $C \in \mathcal{U}/\mathcal{B}$, it follows that for all $a \in \mathcal{A}$ and $C' \in \mathcal{S}/\mathcal{B}'$ such that $C' = \text{last}(C)$ for $C \in \mathcal{U}/\mathcal{B}$ it holds that $\text{rate}(\text{last}(\rho_1), a, C') = \text{rate}_0(\rho_1, a, C) = \text{rate}_0(\rho_2, a, C) = \text{rate}(\text{last}(\rho_2), a, C')$. \square

The behavioral equivalence \sim_{FBMB} is strictly finer than \sim_{MB} . Indeed, it turns out to be exceedingly discriminating. For example, a CTMC with the only two transitions $s_1 \xrightarrow{\langle a, \lambda \rangle} s'_1$ and $s_1 \xrightarrow{\langle a, \mu \rangle} s''_1$ would be distinguished from a CTMC having only transition $s_2 \xrightarrow{\langle a, \lambda + \mu \rangle} s'_2$. Observing that in terms of runs the considered transitions are reformulated as $\rho_1 \xrightarrow{\langle a, \lambda \rangle} \rho'_1$, $\rho_1 \xrightarrow{\langle a, \mu \rangle} \rho''_1$, and $\rho_2 \xrightarrow{\langle a, \lambda + \mu \rangle} \rho'_2$ where $\rho_1 = (s_1, \epsilon)$ and $\rho_2 = (s_2, \epsilon)$, the reflexive, symmetric, and transitive closure of the relation $\{(\rho_1, \rho_2), (\rho'_1, \rho'_2), (\rho''_1, \rho'_2)\}$ would work well when moving forward, as for instance $\text{rate}_0(\rho_1, a, C) = \lambda + \mu = \text{rate}_0(\rho_2, a, C)$ for $C = \{\rho'_1, \rho''_1, \rho'_2\}$, whereas this would not be the case when moving backward, as for instance $\text{rate}_I(\rho'_1, a, C) = \lambda$, $\text{rate}_I(\rho''_1, a, C) = \mu$, and $\text{rate}_I(\rho'_2, a, C) = \lambda + \mu$ for $C = \{\rho_1, \rho_2\}$.

This example suggests that rate-based quantitative aspects should be neglected when going backward. Indeed, summing up rates when considering transitions departing from the same state is consistent with the fact that the sojourn time in that state is exponentially distributed with rate given by the sum of the rates of the outgoing transitions. Operationally, as mentioned at the end of Section 4.2, this can be interpreted as if there were a race among those transitions to decide which one will be executed, with each transition having a winning probability proportional to its rate. This race interpretation no longer applies in the backward direction, as we sum up rates of transitions that may depart from different states. We therefore consider a variant of \sim_{FBMB} that abstracts from time when going backward. It is worth noting that, if it abstracted from time also when going forward, then we would precisely obtain the strong back and forth bisimilarity of [DMV90].

Definition 5.6 (*forward and time-abstract backward Markovian bisimilarity*). Two states $s_1, s_2 \in \mathcal{S}$ are forward and time-abstract backward Markovian bisimilar, written $s_1 \sim_{\text{FTABMB}} s_2$, iff there exists a forward and time-abstract backward Markovian bisimulation \mathcal{B} such that $((s_1, \epsilon), (s_2, \epsilon)) \in \mathcal{B}$. An equivalence relation \mathcal{B} over the set of runs \mathcal{U} is a forward and time-abstract backward Markovian bisimulation iff, whenever $(\rho_1, \rho_2) \in \mathcal{B}$, then for all actions $a \in \mathcal{A}$ and equivalence classes $C \in \mathcal{U}/\mathcal{B}$:

$$\begin{aligned} \text{rate}_0(\rho_1, a, C) &= \text{rate}_0(\rho_2, a, C) \\ \text{trans}_I(\rho_1, a, C) &= \text{trans}_I(\rho_2, a, C) \end{aligned}$$

where $\text{trans}_I(\rho, a, C) = 1$ if there exist $\rho' \in C$ and $\lambda \in \mathcal{R}$ such that $\rho' \xrightarrow{\langle a, \lambda \rangle} \rho$, otherwise $\text{trans}_I(\rho, a, C) = 0$. \blacksquare

We conclude by showing that \sim_{FTABMB} – which is defined on runs, compares for any action both outgoing total rates and the existence of incoming transitions, and preserves history when going backward even in the presence of concurrent transitions – coincides with the standard \sim_{MB} – which is defined on states and compares for any action only outgoing total rates – thus generalizing the first result of [DMV90] to our stochastic setting. As a consequence, when lifted to \mathbb{P} , the former equivalence inherits the compositionality properties

and the equational and logical characterizations of the latter [Ber07], including in particular the expansion law as well as the stochastic variant of idempotency according to which $\langle a, \lambda \rangle.P + \langle a, \mu \rangle.P$ is identified with $\langle a, \lambda + \mu \rangle.P$.

Theorem 5.7. *Let $s_1, s_2 \in \mathcal{S}$. Then $s_1 \sim_{\text{FTABMB}} s_2 \iff s_1 \sim_{\text{MB}} s_2$.*

Proof. The proof is divided into two parts:

- The proof of $s_1 \sim_{\text{FTABMB}} s_2 \implies s_1 \sim_{\text{MB}} s_2$ is identical to the proof of Theorem 5.5 as only outgoing total rates are considered.
- Suppose that $s_1 \sim_{\text{MB}} s_2$. Let ct be the mapping that associates with each path ξ its colored trace, i.e., the path obtained from ξ by replacing each state with its Markovian bisimulation equivalence class:

$$\text{ct}((s_0, \langle a_1, \lambda_1 \rangle, s_1)(s_1, \langle a_2, \lambda_2 \rangle, s_2) \dots (s_{n-1}, \langle a_n, \lambda_n \rangle, s_n)) = \\ ([s_0]_{\sim_{\text{MB}}}, \langle a_1, \lambda_1 \rangle, [s_1]_{\sim_{\text{MB}}})([s_1]_{\sim_{\text{MB}}}, \langle a_2, \lambda_2 \rangle, [s_2]_{\sim_{\text{MB}}}) \dots ([s_{n-1}]_{\sim_{\text{MB}}}, \langle a_n, \lambda_n \rangle, [s_n]_{\sim_{\text{MB}}})$$

Let $\mathcal{B} = \{(\rho_1, \rho_2) \mid \rho_1 \in \text{Run}(s_1), \rho_2 \in \text{Run}(s_2), \text{ct}(\text{path}(\rho_1)) = \text{ct}(\text{path}(\rho_2))\}$, which contains in particular $((s_1, \epsilon), (s_2, \epsilon))$. We show that its reflexive, symmetric, and transitive closure \mathcal{B}' is a forward and time-abstract backward Markovian bisimulation.

Given $(\rho_1, \rho_2) \in \mathcal{B}$, i.e., $\rho_1 \in \text{Run}(s_1)$ and $\rho_2 \in \text{Run}(s_2)$ such that $\text{ct}(\text{path}(\rho_1)) = \text{ct}(\text{path}(\rho_2))$, with l being the length of $\text{path}(\rho_1)$ and $\text{path}(\rho_2)$, let us examine the forward and backward directions respectively:

- By virtue of $\text{ct}(\text{path}(\rho_1)) = \text{ct}(\text{path}(\rho_2))$, it holds in particular that $\text{last}(\rho_1) \sim_{\text{MB}} \text{last}(\rho_2)$, hence for all $a \in \mathcal{A}$ and $C \in \mathcal{S}/\sim_{\text{MB}}$ we have that $\text{rate}(\text{last}(\rho_1), a, C) = \text{rate}(\text{last}(\rho_2), a, C)$. Since $\text{last}(\rho_k) \xrightarrow{\langle a, \lambda \rangle} \text{last}(\rho'_k)$ iff $\rho_k \xrightarrow{\langle a, \lambda \rangle} \rho'_k$ for $k \in \{1, 2\}$, with ρ'_k still belonging to $\text{Run}(s_k)$, and any equivalence class $C' \in (\text{Run}(s_1) \cup \text{Run}(s_2))/\mathcal{B}'$ is made of runs with paths of the same length that traverse the same sequence of Markovian bisimulation equivalence classes of states and perform the same sequence of exponentially timed actions, it follows that for all $a \in \mathcal{A}$ and $C' \in (\text{Run}(s_1) \cup \text{Run}(s_2))/\mathcal{B}'$:
 - * if the length of all the runs of C' is less than $l + 1$, then it trivially holds that $\text{rate}_0(\rho_1, a, C') = 0 = \text{rate}_0(\rho_2, a, C')$;
 - * otherwise, assuming that the states reached after $l + 1$ transitions by all of the runs of C' belong to some $C \in \mathcal{S}/\sim_{\text{MB}}$, it holds that $\text{rate}_0(\rho_1, a, C') = \text{rate}(\text{last}(\rho_1), a, C) = \text{rate}(\text{last}(\rho_2), a, C) = \text{rate}_0(\rho_2, a, C')$.
- Since $\text{ct}(\text{path}(\rho_1)) = \text{ct}(\text{path}(\rho_2))$, it cannot be the case that $\text{path}(\rho_1) = \epsilon$ and $\text{path}(\rho_2) \neq \epsilon$, or vice versa. If $\text{path}(\rho_1) = \text{path}(\rho_2) = \epsilon$, then for all $a \in \mathcal{A}$ and $C' \in (\text{Run}(s_1) \cup \text{Run}(s_2))/\mathcal{B}'$ it trivially holds that $\text{trans}_{\text{I}}(\rho_1, a, C') = 0 = \text{trans}_{\text{I}}(\rho_2, a, C')$. Suppose that $\text{path}(\rho_1) \neq \epsilon \neq \text{path}(\rho_2)$, with $\rho_k = \rho'_k \rho''_k$ and $\rho''_k = (s'_k, (s'_k, \langle a, \lambda \rangle, s''_k))$, so that $s'_k \xrightarrow{\langle a, \lambda \rangle} s''_k$ is the last transition in $\text{path}(\rho_k)$ and hence $\rho'_k \xrightarrow{\langle a, \lambda \rangle} \rho_k$ with ρ'_k still belonging to $\text{Run}(s_k)$, for $k \in \{1, 2\}$. From $\text{ct}(\text{path}(\rho_1)) = \text{ct}(\text{path}(\rho_2))$, it follows in particular that $s'_1 \sim_{\text{MB}} s'_2$. As a consequence, for all $a \in \mathcal{A}$ and $C' \in (\text{Run}(s_1) \cup \text{Run}(s_2))/\mathcal{B}'$:
 - * if the length of all runs of C' is less than l , then it trivially holds that $\text{trans}_{\text{I}}(\rho_1, a, C') = 0 = \text{trans}_{\text{I}}(\rho_2, a, C')$;
 - * otherwise, assuming that the states reached after $l - 1$ transitions by all runs of C' belong to some $C \in \mathcal{S}/\sim_{\text{MB}}$, it holds that $\text{trans}_{\text{I}}(\rho_1, a, C') = 1 = \text{trans}_{\text{I}}(\rho_2, a, C')$ or $\text{trans}_{\text{I}}(\rho_1, a, C') = 0 = \text{trans}_{\text{I}}(\rho_2, a, C')$ depending on whether $s'_1, s'_2 \in C$ or $s'_1, s'_2 \notin C$. \square

6. EXAMPLES

In this section, we show how to use RMPC for modeling scenarios coming from two different application domains: distributed systems and systems biology.

6.1. Consensus Protocols. The two-phase commit protocol (2PC) [Mul93] is a distributed algorithm that coordinates all the processes participating in a distributed transaction on whether to commit or abort the transaction itself. In the first phase, a process acting as the coordinator initiates a voting procedure to decide whether the transaction can be committed or not. All the participants in the distributed transaction receive a request to vote and then reply either yes for commit or no for abort depending on their local state of the transaction. In the second phase, the coordinator decides according to the received votes. If all the participants have agreed to commit, then the coordinator decides to commit, otherwise to abort, and notifies its decision to the participants.

In RMPC the considered distributed system can be modeled as the parallel composition of the coordinator with $m \geq 2$ processes participating in the transaction:

$$2PC = \text{COORD} \parallel_L P_1 \parallel_L P_2 \parallel_L \dots \parallel_L P_m$$

where $L = \{vt, cmt, abt\} \cup \{y_i, n_i \mid i \in I\}$, for $I = \{1, 2, \dots, m\}$, is the set of actions on which the coordinator and all the participating processes have to synchronize.

The coordinator sends the request to vote at rate λ , then it waits for a yes/no answer from each process and decides to commit only if all processes have replied yes; the commit/abort decision is broadcast at rate δ . The coordinator has a component for each possible process, with all these components being in parallel among them with empty synchronization set (rendered as \prod for brevity), and a component counting the positive answers regardless of the order in which they arrive:

$$\begin{aligned} \text{COORD} &= \langle vt, \lambda \rangle. (\text{COORD}' \parallel_Y \text{COORD}'') \\ \text{COORD}' &= \prod_{i \in I} (\langle y_i, 1 \rangle. \underline{0} + \langle n_i, 1 \rangle. \langle abt, \delta \rangle. \underline{0}) \\ \text{COORD}'' &= \langle y_1, 1 \rangle. \langle y_2, 1 \rangle. \dots. \langle y_m, 1 \rangle. \langle cmt, \delta \rangle. \underline{0} \end{aligned}$$

where $Y = \{y_i \mid i \in I\}$ and the use of rate 1 in certain actions indicates the passivity of the coordinator with respect to the synchronization on the corresponding activities.

After performing its part of transaction at rate μ_i , the i -th participating process waits for a vote request. Then at rate ρ_i it answers yes with probability p_i , in which case it waits for the commit or abort decision of the coordinator, or no with probability $q_i = 1 - p_i$, in which case it waits for the abort decision of the coordinator (let $\gamma_i = \rho_i \cdot p_i$ and $\eta_i = \rho_i \cdot q_i$):

$$\begin{aligned} P_i &= \langle t_i, \mu_i \rangle. \langle vt, 1 \rangle. P'_i \\ P'_i &= \langle y_i, \gamma_i \rangle. (\langle cmt, 1 \rangle. \underline{0} + \langle abt, 1 \rangle. \underline{0}) + \langle n_i, \eta_i \rangle. \langle abt, 1 \rangle. \underline{0} \end{aligned}$$

Even if a process chooses to commit, the coordinator may decide to abort. On the other hand, if a process chooses to abort, it is granted that the overall transaction will be aborted.

Assuming $m = 2$ for simplicity, a standard forward computation may start with the following three transitions and reach a term that we call $2PC'$:

$$2PC \xrightarrow{\langle t_1, \mu_1 \rangle [1]} \xrightarrow{\langle t_2, \mu_2 \rangle [2]} \xrightarrow{\langle vt, \lambda \rangle [3]} \langle vt, \lambda \rangle [3]. (\text{COORD}' \parallel_Y \text{COORD}'') \parallel_L \langle t_1, \mu_1 \rangle [1]. \langle vt, 1 \rangle [3]. P'_1 \parallel_L \langle t_2, \mu_2 \rangle [2]. \langle vt, 1 \rangle [3]. P'_2$$

The two participating processes can now choose independently of each other. If process 1 votes yes while process 2 votes no, so that the coordinator decides to abort the transaction, then the computation may continue as follows and reach a term that we call $2PC''$:


$2PC' \xrightarrow{\langle y_1, \gamma_1 \rangle [4]} \xrightarrow{\langle n_2, \eta_2 \rangle [5]} \xrightarrow{\langle abt, \delta \rangle [6]}$
 $\langle vt, \lambda \rangle [3].(((\langle y_1, 1 \rangle [4].\underline{0} + \langle n_1, 1 \rangle .\langle abt, \delta \rangle .\underline{0}) \parallel_{\emptyset}$
 $(\langle y_2, 1 \rangle .\underline{0} + \langle n_2, 1 \rangle [5].\langle abt, \delta \rangle [6].\underline{0})) \parallel_L$
 $(\langle y_1, 1 \rangle [4].\langle y_2, 1 \rangle .\langle cmt, \delta \rangle .\underline{0}) \parallel_L$
 $\langle t_1, \mu_1 \rangle [1].\langle vt, 1 \rangle [3].(\langle y_1, \gamma_1 \rangle [4].(\langle cmt, 1 \rangle .\underline{0} + \langle abt, 1 \rangle [6].\underline{0}) + \langle n_1, \eta_1 \rangle .\langle abt, 1 \rangle .\underline{0}) \parallel_L$
 $\langle t_2, \mu_2 \rangle [2].\langle vt, 1 \rangle [3].(\langle y_1, \gamma_1 \rangle .(\langle cmt, 1 \rangle .\underline{0} + \langle abt, 1 \rangle .\underline{0}) + \langle n_1, \eta_1 \rangle [5].\langle abt, 1 \rangle [6].\underline{0})$
 where the coordinator and the two processes could only synchronize on the abort action.

At this point, the computation can only continue as a distributed rollback like the one below, which takes place with no need to model it due to the reversibility of RMPC:

$$2PC'' \xrightarrow{\langle abt, \bar{\delta} \rangle [6]} \xrightarrow{\langle n_2, \bar{\eta}_2 \rangle [5]} \xrightarrow{\langle y_1, \bar{\gamma}_1 \rangle [4]} \xrightarrow{\langle vt, \bar{\lambda} \rangle [3]} \xrightarrow{\langle t_2, \bar{\mu}_2 \rangle [2]} \xrightarrow{\langle t_1, \bar{\mu}_1 \rangle [1]} 2PC$$

Note that this is not the only possibility, as there are other backward computations causally equivalent to the one above obtained by exchanging the order in which actions y_1 and n_2 on the one hand and actions t_1 and t_2 on the other hand are undone.

6.2. Protein Interaction Networks. Consider a prototypical model from systems biology where a molecule A has multiple binding sites to which a molecule B can bind reversibly [CFG08]. Since the number of reactions grows exponentially with the number of binding sites, as in [STTV22] we only examine the case of two binding sites. We denote by $A_{1,0}$ (resp., $A_{0,1}$) a molecule A having a binding with a single molecule B on the first (resp., second) site and by $A_{1,1}$ a molecule A in the case that the binding is with two molecules B , one on each site. In the description of the reaction network of such molecules, the left column below indicates the binding (i.e., forward) rules while the right column below indicates the unbinding (i.e., backward) rules where, according to the standard biochemical notation, $A + B$ stands for the presence of at least one molecule A and at least one molecule B as well as κ_{b_i} and κ_{u_i} represent the rates of the corresponding binding and unbinding rules:



In RMPC we can model the same network as follows:

$$\begin{aligned}
 \text{NET} &= A \parallel_L (B \parallel_{\emptyset} B) \\
 A &= \langle b_1, 1 \rangle .\underline{0} \parallel_{\emptyset} \langle b_2, 1 \rangle .\underline{0} \\
 B &= \langle b_1, \kappa_{b_1} \rangle .\underline{0} + \langle b_2, \kappa_{b_2} \rangle .\underline{0}
 \end{aligned}$$

where $L = \{b_1, b_2\}$. Note that there is no need to model the unbinding rules thanks to the reversibility of RMPC. So for instance the forward computation:

$$\begin{aligned}
 \text{NET} &\xrightarrow{\langle b_1, \kappa_{b_1} \rangle [1]} (\langle b_1, 1 \rangle [1].\underline{0} \parallel_{\emptyset} \langle b_2, 1 \rangle .\underline{0}) \parallel_L \\
 &\quad ((\langle b_1, \kappa_{b_1} \rangle [1].\underline{0} + \langle b_2, \kappa_{b_2} \rangle .\underline{0}) \parallel_{\emptyset} (\langle b_1, \kappa_{b_1} \rangle .\underline{0} + \langle b_2, \kappa_{b_2} \rangle .\underline{0})) \\
 &\xrightarrow{\langle b_2, \kappa_{b_2} \rangle [2]} (\langle b_1, 1 \rangle [1].\underline{0} \parallel_{\emptyset} \langle b_2, 1 \rangle [2].\underline{0}) \parallel_L \\
 &\quad ((\langle b_1, \kappa_{b_1} \rangle [1].\underline{0} + \langle b_2, \kappa_{b_2} \rangle .\underline{0}) \parallel_{\emptyset} (\langle b_1, \kappa_{b_1} \rangle .\underline{0} + \langle b_2, \kappa_{b_2} \rangle [2].\underline{0}))
 \end{aligned}$$

has the following two backward counterparts as the two unbindings can happen in any order from the reached term that we call NET' :

$$\begin{aligned}
 \text{NET}' &\xrightarrow{\langle b_2, \bar{\kappa}_{b_2} \rangle [2]} \xrightarrow{\langle b_1, \bar{\kappa}_{b_1} \rangle [1]} \text{NET} \\
 \text{NET}' &\xrightarrow{\langle b_1, \bar{\kappa}_{b_1} \rangle [1]} \xrightarrow{\langle b_2, \bar{\kappa}_{b_2} \rangle [2]} \text{NET}
 \end{aligned}$$

7. CONCLUSIONS

After Landauer [Lan61] and Bennett [Ben73], reversible computing has attracted a growing attention mainly for the possibility of building energy efficient circuits. Nowadays, the interest has spread into many application areas and hence reversible computing requires a deep investigation of its theoretical foundations in computer science. There exist different interpretations of reversibility in the literature. In this paper, we have addressed our research quest towards bridging causal reversibility [DK04, DK05] – developed in concurrency theory – and time reversibility [Kel79] – originated for the efficient analysis of stochastic processes.

We have accomplished this in the setting of process algebra because it constitutes a common ground for concurrency theory and probability theory [LS91, Hil96]. Specifically, we have introduced the stochastic process calculus RMPC, whose syntax and semantics follow the approach of [PU07], with the aim of paving the way to concurrent system models that are causally reversible and time reversible *by construction*. Causal reversibility has been proved by exploiting the technique of [LPU20] after importing in the setting of [PU07] some notions coming from [DK04]. Time reversibility has been shown under the constraint that every backward rate is equal to the corresponding forward rate – regardless of the syntactical structure of processes – or that parallel composition cannot occur within the scope of action prefix or choice – regardless of the values of backward rates – and has allowed us to inherit from [MR15] a product form result that enables the efficient calculation of performance measures. Finally, as far as Markovian bisimilarity [Hil96] is concerned, we have developed two forward and backward variants inspired by [DMV90], with the former preserving the expansion law and the latter preserving a stochastic variant of the idempotency law too.

There are several lines of research that we plan to undergo, ranging from the application of our results to case studies modeled with RMPC to the development of further theoretical results possibly admitting the presence of irreversible actions. In particular, we would like to investigate further conditions under which time reversibility is achieved. Our conjecture is that time reversibility should hold for the entire set of RMPC reachable processes regardless of their backward rates and syntactical structure. This would imply that the methodology of [PU07] for reversing process calculi – which we have followed in this paper – is fully robust not only with respect to causal reversibility, but also with respect to time reversibility.

Moreover, we wish to find a suitable way of adding recursion to the syntax of RMPC. From the point of view of the ergodicity of the underlying CTMC, the absence of recursion is not a problem because every forward transition has the corresponding backward transition by construction. However, there might be situations in which recursion is necessary to appropriately describe the behavior of a system. Because of the use of communication keys, a simple process of the form $P \triangleq \langle a, \lambda \rangle . P$, whose standard labeled transition system features a single state with a self-looping transition, produces a sequence of infinitely many distinct states even if we resort to transition bundles. Our claim is that the specific cooperation operator that we have considered may require a mechanism lighter than communication keys to keep track of past actions, which may avoid the generation of an infinite state space in the presence of recursion.

ACKNOWLEDGMENT

We would like to thank Andrea Marin and Sabina Rossi for the valuable discussions on time reversibility and the anonymous referees for their constructive comments and suggestions.

This research has been supported by the Italian MUR PRIN 2020 project *NiRvAna – Noninterference and Reversibility Analysis in Private Blockchains*, the Italian INdAM GNCS 2022 project *Proprietà Qualitative e Quantitative di Sistemi Reversibili*, and the French ANR 2018 project *DCore – Causal Debugging for Concurrent Systems*.

REFERENCES

- [BAP⁺12] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. Experimental verification of Landauer’s principle linking information and thermodynamics. *Nature*, 483:187–189, 2012.
- [Ben73] C.H. Bennett. Logical reversibility of computations. *IBM Journal of Research and Development*, 17:525–532, 1973.
- [Ben03] C.H. Bennett. Notes on Landauer’s principle, reversible computation, and Maxwell’s demon. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34:501–510, 2003.
- [Ber07] M. Bernardo. A survey of Markovian behavioral equivalences. In *Formal Methods for Performance Evaluation*, volume 4486 of *LNCS*, pages 180–219. Springer, 2007.
- [BKMP18] K. Barylska, M. Koutny, L. Mikulski, and M. Piatkowski. Reversible computation vs. reversibility in Petri nets. *Science of Computer Programming*, 151:48–60, 2018.
- [BM20] M. Bernardo and C.A. Mezzina. Towards bridging time and causal reversibility. In *Proc. of FORTE 2020*, volume 12136 of *LNCS*, pages 22–38. Springer, 2020.
- [BV93] J.C.M. Baeten and C. Verhoef. A congruence theorem for structured operational semantics with predicates. In *Proc. of CONCUR 1993*, volume 715 of *LNCS*, pages 477–492. Springer, 1993.
- [CFG08] H. Conzelmann, D. Fey, and E.D. Gilles. Exact model reduction of combinatorial reaction networks. *BMC Systems Biology*, 2(78):1–25, 2008.
- [CKV13] I. Cristescu, J. Krivine, and D. Varacca. A compositional semantics for the reversible π -calculus. In *Proc. of LICS 2013*, pages 388–397. IEEE-CS Press, 2013.
- [DK04] V. Danos and J. Krivine. Reversible communicating systems. In *Proc. of CONCUR 2004*, volume 3170 of *LNCS*, pages 292–307. Springer, 2004.
- [DK05] V. Danos and J. Krivine. Transactions in RCCS. In *Proc. of CONCUR 2005*, volume 3653 of *LNCS*, pages 398–412. Springer, 2005.
- [dKH10] E. de Vries, V. Koutavas, and M. Hennessy. Communicating transactions. In *Proc. of CONCUR 2010*, volume 6269 of *LNCS*, pages 569–583. Springer, 2010.
- [DMV90] R. De Nicola, U. Montanari, and F.W. Vaandrager. Back and forth bisimulations. In *Proc. of CONCUR 1990*, volume 458 of *LNCS*, pages 152–165. Springer, 1990.
- [Fra18] M.P. Frank. Physical foundations of Landauer’s principle. In *Proc. of RC 2018*, volume 11106 of *LNCS*, pages 3–33. Springer, 2018.
- [GLM14] E. Giachino, I. Lanese, and C.A. Mezzina. Causal-consistent reversible debugging. In *Proc. of FASE 2014*, volume 8411 of *LNCS*, pages 370–384. Springer, 2014.
- [Har03] P.G. Harrison. Turning back time in Markovian process algebra. *Theoretical Computer Science*, 290:1947–1986, 2003.
- [HGCH21] L. Hay-Schmidt, R. Glück, M.H. Cservenka, and T. Haulund. Towards a unified language architecture for reversible object-oriented programming. In *Proc. of RC 2021*, volume 12805 of *LNCS*, pages 96–106. Springer, 2021.
- [Hil94] J. Hillston. The nature of synchronisation. In *Proc. of PAPM 1994*, pages 51–70. University of Erlangen, Technical Report 27-4, 1994.
- [Hil96] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
- [Hoa85] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [Jac63] J.-R. Jackson. Jobshop-like queueing systems. *Management Science*, 10:131–142, 1963.
- [Kel79] F.P. Kelly. *Reversibility and Stochastic Networks*. John Wiley & Sons, 1979.
- [Kle75] L. Kleinrock. *Queueing Systems*. John Wiley & Sons, 1975.
- [KS60] J.G. Kemeny and J.L. Snell. *Finite Markov Chains*. Van Nostrand, 1960.

- [Lan61] R. Landauer. Irreversibility and heat generated in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.
- [Lee86] G.B. Leeman Jr. A formal approach to undo operations in programming languages. *ACM Transactions on Programming Languages and Systems*, 8:50–87, 1986.
- [LES18] J.S. Laursen, L.-P. Ellekilde, and U.P. Schultz. Modelling reversible execution of robotic assembly. *Robotica*, 36:625–654, 2018.
- [Lév76] J.-J. Lévy. An algebraic interpretation of the $\lambda\beta$ K-calculus; and an application of a labelled λ -calculus. *Theoretical Computer Science*, 2:97–114, 1976.
- [LLM⁺13] I. Lanese, M. Lienhardt, C.A. Mezzina, A. Schmitt, and J.-B. Stefani. Concurrent flexible reversibility. In *Proc. of ESOP 2013*, volume 7792 of *LNCS*, pages 370–390. Springer, 2013.
- [LLMS12] M. Lienhardt, I. Lanese, C.A. Mezzina, and J.-B. Stefani. A reversible abstract machine and its space overhead. In *Proc. of FMOODS/FORTE 2012*, volume 7273 of *LNCS*, pages 1–17. Springer, 2012.
- [LM20] I. Lanese and D. Medić. A general approach to derive uncontrolled reversible semantics. In *Proc. of CONCUR 2020*, volume 171 of *LIPICs*, pages 33:1–33:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [LMM21] I. Lanese, D. Medić, and C.A. Mezzina. Static versus dynamic reversibility in CCS. *Acta Informatica*, 58:1–34, 2021.
- [LMS10] I. Lanese, C.A. Mezzina, and J.-B. Stefani. Reversing higher-order π . In *Proc. of CONCUR 2010*, volume 6269 of *LNCS*, pages 478–493. Springer, 2010.
- [LNPV18a] I. Lanese, N. Nishida, A. Palacios, and G. Vidal. CauDEr: A causal-consistent reversible debugger for Erlang. In *Proc. of FLOPS 2018*, volume 10818 of *LNCS*, pages 247–263. Springer, 2018.
- [LNPV18b] I. Lanese, N. Nishida, A. Palacios, and G. Vidal. A theory of reversibility for Erlang. *Journal of Logical and Algebraic Methods in Programming*, 100:71–97, 2018.
- [LPU20] I. Lanese, I.C.C. Phillips, and I. Ulidowski. An axiomatic approach to reversible computation. In *Proc. of FOSSACS 2020*, volume 12077, pages 442–461. Springer, 2020.
- [LS91] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
- [Mil89] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [MR15] A. Marin and S. Rossi. Quantitative analysis of concurrent reversible computations. In *Proc. of FORMATS 2015*, volume 9268 of *LNCS*, pages 206–221. Springer, 2015.
- [Mul93] S. Mullender. *Distributed Systems*. Addison-Wesley, 1993.
- [Pin17] G.M. Pinna. Reversing steps in membrane systems computations. In *Proc. of CMC 2017*, volume 10725 of *LNCS*, pages 245–261. Springer, 2017.
- [Pla85] B. Plateau. On the stochastic structure of parallelism and synchronization models for distributed algorithms. In *Proc. of SIGMETRICS 1985*, pages 147–154. ACM Press, 1985.
- [PP14] K.S. Perumalla and A.J. Park. Reverse computation for rollback-based fault tolerance in large parallel systems - Evaluating the potential gains and systems effects. *Cluster Computing*, 17:303–313, 2014.
- [PU07] I.C.C. Phillips and I. Ulidowski. Reversing algebraic process calculi. *Journal of Logic and Algebraic Programming*, 73:70–96, 2007.
- [PUY13] I.C.C. Phillips, I. Ulidowski, and S. Yuen. A reversible process calculus and the modelling of the ERK signalling pathway. In *Proc. of RC 2012*, volume 7581 of *LNCS*, pages 218–232. Springer, 2013.
- [PV15] A. Palacios and G. Vidal. Concolic execution in functional programming by program instrumentation. In *Proc. of LOPSTR 2015*, volume 9527 of *LNCS*, pages 277–292. Springer, 2015.
- [SOJB18] M. Schordan, T. Opielstrup, D.R. Jefferson, and P.D. Barnes Jr. Generation of reversible C++ code for optimistic parallel discrete event simulation. *New Generation Computing*, 36:257–280, 2018.
- [SPP19] H. Siljak, K. Psara, and A. Philippou. Distributed antenna selection for massive MIMO using reversing Petri nets. *IEEE Wireless Communication Letters*, 8:1427–1430, 2019.
- [Ste94] W.J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1994.

- [STTV22] G. Squillace, M. Tribastone, M. Tschaikowski, and A. Vandin. An algorithm for the formal reduction of differential equations as over-approximations. In *Proc. of QEST 2022*, volume 13479 of *LNCS*, pages 173–191. Springer, 2022.
- [VS18] M. Vassor and J.-B. Stefani. Checkpoint/rollback vs causally-consistent reversibility. In *Proc. of RC 2018*, volume 11106 of *LNCS*, pages 286–303. Springer, 2018.
- [YG07] T. Yokoyama and R. Glück. A reversible programming language and its invertible self-interpreter. In *Proc. of PEPM 2007*, pages 144–153. ACM Press, 2007.