

A Survey of Markovian Behavioral Equivalences

Marco Bernardo

Università di Urbino “Carlo Bo” – Italy
Istituto di Scienze e Tecnologie dell’Informazione

Abstract. Markovian behavioral equivalences are a means to relate and manipulate the formal descriptions of systems with an underlying CTMC semantics. There are three fundamental approaches to their definition: bisimilarity, testing, and trace. In this paper we survey the major results appeared in the literature about Markovian bisimilarity, Markovian testing equivalence, and Markovian trace equivalence. The objective is to compare these equivalences with respect to a number of criteria such as their discriminating power, the exactness of the CTMC-level aggregations they induce, the achievement of the congruence property, the existence of sound and complete axiomatizations, the existence of logical characterizations, and the existence of efficient verification algorithms.

1 Introduction

Performance-oriented notations provide the designer with the capability of building performance-aware system models, which can be used in the early development stages to predict the satisfiability of certain performance requirements as well as to choose among alternative designs on the basis of their expected QoS guarantees. These notations range from more theoretical ones – like queueing networks [38], stochastic Petri nets [1], and stochastic process algebras [32] – to more applicative ones – like formal modeling languages (MODEST [12]), architectural description languages (\mathcal{A} EMILIA [5]), coordination languages (STOKLAIM [22]), and object-oriented modeling languages (UML SPT/MARTE [48]).

An important feature shared by most of the performance-oriented notations mentioned above is that of providing behavioral models of the systems under construction. Given two such models, establishing whether they are equivalent amounts to establishing whether the systems they represent behave the same. What is needed is thus a notion of behavioral equivalence. This would be useful not only to relate models that are syntactically different, but also to manipulate models in a way that preserves their functional and performance properties.

Among the various proposals appeared in the literature [25], there are three fundamental approaches to the definition of behavioral equivalences: bisimilarity [43, 41], testing [21], and trace [33]. In the first approach, two models are considered to be equivalent if they are able to mimic each other’s behavior step-wise. In the second approach, two models are considered to be equivalent if an external observer cannot distinguish between them, with the only way for the observer to compare their behaviors being to interact with them by means of

tests and look at their reactions. In the third approach, similarly to traditional automata theory, two models are considered to be equivalent if they are able to perform the same sequences of activities.

These three approaches, originally conceived for reasoning about functional aspects, have been subsequently extended to deal with non-functional aspects. As far as performance aspects are concerned, research has mainly concentrated on models of systems with an underlying continuous-time Markov chain (CTMC) semantics. The reason is that, due to their memoryless property, exponential distributions result in a simpler mathematical treatment without sacrificing expressiveness. In fact, besides being adequate for many real-life phenomena (like arrival processes and failure events), exponential distributions provide the most appropriate stochastic approximation if only the average duration of an activity is known, and proper combinations of them (called phase-type distributions) can approximate most of the general distributions arbitrarily closely.

This has resulted in the development of the Markovian versions of bisimilarity, testing equivalence, and trace equivalence, which will be surveyed in this paper by recalling from [32, 15, 31, 14, 18, 4, 23, 10, 30, 8, 7, 9, 49] their properties.

Although behavioral equivalences abstract from specific kinds of models, most of their properties can be better investigated and understood in a process algebraic framework [41, 33, 2, 6]. For this reason, the three Markovian behavioral equivalences mentioned above will be defined in this paper over Markovian process calculi. Many such calculi have been proposed in the literature, like TIPP [27], PEPA [32], MPA [16], EMPA_{gr} [10], S π [44], IMC [30], and PIOA [46]. They differ for the action representation – durational actions (TIPP, PEPA, MPA, EMPA_{gr}, S π , PIOA) vs. instantaneous actions separated from time passing (IMC) – as well as for the action synchronization discipline – symmetric (TIPP, MPA, S π , IMC), asymmetric (EMPA_{gr}, PIOA), or both (PEPA).

In this paper we shall start with a sequential Markovian process calculus (SMPC) with durational actions, which generates all the finite CTMCs with as few operators as possible: the null term, the action prefix operator, the alternative composition operator, and recursion. Then we shall add a parallel composition operator governed by an asymmetric action synchronization discipline, thus resulting in a concurrent Markovian process calculus (CMPC). We shall also address some syntax variations – like the inclusion of rewards, nondeterminism, and prioritized/weighted immediate actions – in order to present some useful variants of the considered equivalences.

Markovian bisimilarity, Markovian testing equivalence, and Markovian trace equivalence will be compared with respect to the following criteria:

1. *Discriminating power.* The three Markovian behavioral equivalences, together with some variants of Markovian trace equivalence, will be ordered according to a finer-than/coarser-than relation, thus providing information about the linear-time/branching-time spectrum in the Markovian case. As we shall see, similarly to what happens in the probabilistic case [36, 35], the Markovian spectrum is more condensed than the nondeterministic one [25].

2. *Exactness.* Each of the three Markovian behavioral equivalences induces an aggregation at the CTMC level. In general, a CTMC aggregation is exact whenever the transient/stationary probability of being in a macrostate of an aggregated CTMC is the sum of the transient/stationary probabilities of being in one of the constituent microstates of the original CTMC. This guarantees the preservation of the performance characteristics when going from the original CTMC to the aggregated one. As we shall see, all the three Markovian behavioral equivalences induce exact aggregations. In other words, the three approaches – bisimilarity, testing, trace – to the definition of behavioral equivalences are not only intuitively appropriate from the functional viewpoint, but also meaningful for performance evaluation purposes.
3. *Congruence.* A behavioral equivalence that is a congruence with respect to the typical process algebraic operators is particularly helpful in practice, as it supports compositional reasoning. This enables the compositional reduction of the model state space. As we shall see, Markovian bisimilarity and Markovian testing equivalence are congruences, whereas Markovian trace equivalence – unlike the nondeterministic case but similarly to the probabilistic one [36] – is not a congruence with respect to parallel composition.
4. *Axiomatization.* The axiomatization of a behavioral equivalence elucidates the fundamental equational laws on which the equivalence relies. This equational characterization is thus useful to understand what models can be related by the equivalence. Whenever it is sound and complete, the axiomatization gives rise to the specific rules of a deduction system – including reflexivity, symmetry, transitivity, and substitutivity – that can be exploited as a rewriting system to syntactically manipulate the models in a way that is consistent with the equivalence.
5. *Logical characterization.* The modal/temporal logic characterization of a behavioral equivalence shows what behavioral properties are preserved by the equivalence. This can be exploited to provide diagnostic information that explains why two models are not equivalent. As we shall see, Markovian bisimilarity preserves branching-time properties, while Markovian trace equivalence preserves linear-time properties.
6. *Verification complexity.* In order to be applicable in practice, a behavioral equivalence must be equipped with an efficient verification algorithm. As we shall see, not only Markovian bisimilarity but also Markovian testing and trace equivalences – unlike the nondeterministic case [37] but similarly to the probabilistic one [35] – are all decidable in polynomial time.

This paper is organized as follows. In Sect. 2 we introduce the syntax and the semantics for SMPC and CMPC. In Sect. 3 we study Markovian bisimilarity. In Sect. 4 we address some of its variants that include rewards, nondeterminism, and prioritized/weighted immediate actions. In Sect. 5 we present Markovian testing equivalence. In Sect. 6 we illustrate Markovian trace equivalence together with some other trace-based Markovian behavioral equivalences. Finally, in Sect. 7 we summarize the comparison of the Markovian behavioral equivalences based on the criteria explained above and we discuss some open problems.

2 Basic Markovian Process Calculi

In this section we introduce two basic Markovian process calculi with durational actions. The first one is a sequential Markovian process calculus (SMPC) that generates all the finite CTMCs with as few operators as possible: the null term, the action prefix operator, the alternative composition operator, and recursion. The second one is a concurrent Markovian process calculus (CMPC) as it additionally includes a parallel composition operator governed by an asymmetric action synchronization discipline. Then we introduce some notation concerned with the exit rates of the process terms and the attributes associated with their computations.

2.1 Syntax and Semantics for SMPC

In SMPC every action is durational, hence it is represented as a pair $\langle a, \lambda \rangle$, where $a \in Name$ is the name of the action while $\lambda \in \mathbf{R}_{>0}$ is the rate of the exponential distribution quantifying the duration of the action. The average duration of an exponentially timed action is equal to the inverse of its rate.

Whenever several exponentially timed actions are enabled, the race policy is adopted, hence the fastest action is the one that is executed. As a consequence of this generative [26] selection mechanism, the execution probability of any enabled exponentially timed action is proportional to its rate and the average sojourn time associated with a process term is exponentially distributed with rate given by the sum of the rates of the actions enabled by the term.

We denote by $Act_S = Name \times \mathbf{R}_{>0}$ the set of the actions of SMPC. Unlike standard process theory, where a distinguished symbol τ is used as the name of the invisible action, here we assume that all the actions are visible.

Definition 1. *The set \mathcal{L}_S of the process terms of SMPC is generated by the following syntax:*

$$\begin{array}{|l}
 P ::= 0 \\
 | \langle a, \lambda \rangle . P \\
 | P + P \\
 | X \\
 | \text{rec } X : P
 \end{array}$$

where X is a process variable. We denote by \mathcal{P}_S the set of the closed and guarded process terms of SMPC. ■

The semantics for SMPC is given by a state-transition model that can be defined in the usual operational style. However, unlike nondeterministic process calculi, idempotency no longer holds. In fact, a term like $\langle a, \lambda \rangle . P + \langle a, \lambda \rangle . P$ is not the same as $\langle a, \lambda \rangle . P$, as the average sojourn time associated with the latter, i.e. $1/\lambda$, is twice the average sojourn time associated with the former, i.e. $1/(\lambda + \lambda)$. To keep the two terms distinct at the semantic level, it is necessary to take into account the multiplicity of each transition, intended as the number of different proofs for the derivation of the transition.

Therefore, the behavior of each process term $P \in \mathcal{P}_S$ is given by a labeled multitransition system $\llbracket P \rrbracket$, whose states correspond to process terms and whose transitions – each of which has a multiplicity – are labeled with actions. From such a labeled multitransition system the CTMC underlying the process term can easily be retrieved by (i) discarding the action names from the transition labels and (ii) collapsing all the transitions between any two states into a single transition whose rate is the sum of the rates of the original transitions.

We now provide the semantic rules for the various operators of SMPC:

- Null term: $\underline{0}$ cannot execute any action, hence the corresponding labeled multitransition system is just a state with no transitions.
- Exponentially timed action prefix: $\langle a, \lambda \rangle.P$ can execute an action of name a and rate λ and then behaves as P :

$$\boxed{\langle a, \lambda \rangle.P \xrightarrow{a, \lambda} P}$$

- Alternative composition: $P_1 + P_2$ behaves as either P_1 or P_2 depending on whether P_1 or P_2 executes an action first:

$$\boxed{\frac{P_1 \xrightarrow{a, \lambda} P'}{P_1 + P_2 \xrightarrow{a, \lambda} P'} \quad \frac{P_2 \xrightarrow{a, \lambda} P'}{P_1 + P_2 \xrightarrow{a, \lambda} P'}}$$

- Recursion: $\text{rec } X : P$ behaves as P after replacing every occurrence of X with $\text{rec } X : P$:

$$\boxed{\frac{P\{\text{rec } X : P\}/X \xrightarrow{a, \lambda} P'}{\text{rec } X : P \xrightarrow{a, \lambda} P'}}$$

2.2 Syntax and Semantics for CMPC

CMPC extends SMPC with a parallel composition operator governed by an asymmetric action synchronization discipline, which is enforced on an explicit set of action names and makes use of passive actions. Multiway synchronizations are allowed provided that they involve at most one exponentially timed action, with all the other actions being passive.

A passive action is of the form $\langle a, *w \rangle$, where $w \in \mathbf{R}_{>0}$ is called weight and is used to quantify choices among passive actions with the same name. Every passive action has a duration that will become specified upon synchronization with an exponentially timed action having the same name.

Whenever several passive actions are enabled, the reactive [26] preselection policy is adopted. This means that, within every set of enabled passive actions with the same name, each such action is given an execution probability proportional to its weight. The choice between two enabled passive actions having different names is instead nondeterministic.

We denote by $\text{Act}_C = \text{Name} \times \text{Rate}$ the set of the actions of CMPC, where $\text{Rate} = \mathbf{R}_{>0} \cup \{*_w \mid w \in \mathbf{R}_{>0}\}$ is the set of the action rates (ranged over by $\tilde{\lambda}$). As for SMPC, we assume that all the actions are visible.

Definition 2. The set \mathcal{L}_C of the process terms of CMPC is generated by the following syntax:

$$\boxed{
\begin{array}{l}
P ::= \underline{0} \\
\quad | \langle a, \lambda \rangle . P \\
\quad | \langle a, *w \rangle . P \\
\quad | P + P \\
\quad | P \parallel_S P \\
\quad | X \\
\quad | \text{rec } X : P
\end{array}
}$$

where $S \subseteq \text{Name}$ and X is a process variable. We denote by \mathcal{P}_C the set of the closed and guarded process terms of CMPC. \blacksquare

Due to the memoryless property of exponential distributions and the fact that the probability that two concurrent exponentially timed actions terminate simultaneously is zero, the semantics for the parallel composition operator can be defined in the usual interleaving style like in the nondeterministic case. In fact, term $\langle a, \lambda \rangle . \underline{0} \parallel_{\emptyset} \langle b, \mu \rangle . \underline{0}$ behaves exactly like term $\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle b, \mu \rangle . \langle a, \lambda \rangle . \underline{0}$ as the execution of an exponentially timed action can be thought of as being started in the last state in which the action is enabled.

We now provide the semantic rules for the additional operators of CMPC:

- Passive action prefix: $\langle a, *w \rangle . P$ can execute a passive action of name a and weight w and then behaves as P :

$$\boxed{\langle a, *w \rangle . P \xrightarrow{a, *w} P}$$

- Parallel composition: $P_1 \parallel_S P_2$ behaves as P_1 in parallel with P_2 as long as actions are executed whose names do not belong to S :

$$\boxed{
\frac{P_1 \xrightarrow{a, \bar{\lambda}} P'_1 \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a, \bar{\lambda}} P'_1 \parallel_S P_2} \quad \frac{P_2 \xrightarrow{a, \bar{\lambda}} P'_2 \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a, \bar{\lambda}} P_1 \parallel_S P'_2}
}$$

Generative-reactive synchronizations are forced between any exponentially timed action executed by one term and any passive action executed by the other term that have the same name belonging to S :

$$\boxed{
\frac{P_1 \xrightarrow{a, \lambda} P'_1 \quad P_2 \xrightarrow{a, *w} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda, \frac{w}{\text{weight}(P_2, a)}} P'_1 \parallel_S P'_2}
}$$

$$\boxed{
\frac{P_1 \xrightarrow{a, *w} P'_1 \quad P_2 \xrightarrow{a, \lambda} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda, \frac{w}{\text{weight}(P_1, a)}} P'_1 \parallel_S P'_2}
}$$

while reactive-reactive synchronizations are forced between any two passive actions executed by the two terms that have the same name belonging to S :

$$\boxed{\begin{array}{c} P_1 \xrightarrow{a, *w_1} P'_1 \quad P_2 \xrightarrow{a, *w_2} P'_2 \quad a \in S \\ \hline P_1 \parallel_S P_2 \xrightarrow{a, * \frac{w_1}{\text{weight}(P_1, a)} \cdot \frac{w_2}{\text{weight}(P_2, a)} \cdot (\text{weight}(P_1, a) + \text{weight}(P_2, a))} P'_1 \parallel_S P'_2 \end{array}}$$

where the weight of a process term P with respect to the passive actions of name a that P enables is defined as follows:

$$\boxed{\text{weight}(P, a) = \sum \{ w \in \mathbf{R}_{>0} \mid \exists P' \in \mathcal{P}_C. P \xrightarrow{a, *w} P' \}}$$

We point out that the CTMC underlying a process term in \mathcal{P}_C can be retrieved only if its labeled multitransition system has no passive transitions. In this case we say that the process term is performance closed. We denote by $\mathcal{P}_{C,pc}$ the set of the performance closed process terms of \mathcal{P}_C . Note that $\mathcal{P}_{S,pc} = \mathcal{P}_S$.

2.3 Exit Rates and Computations of Process Terms

The Markovian behavioral equivalences that we shall define over SMPC and CMPC are based on concepts like the exit rates of the process terms and the traces, the probabilities, and the durations of their computations. Since these concepts will be used several times in the paper, we collect in this section the related notation.

The exit rate of a process term is the rate at which it is possible to leave the term. We distinguish among the rate at which the process term can execute actions of a given name that lead to a given set of terms, the total rate at which the process term can execute actions of a given name, and the total exit rate of the process term. The latter is the sum of the rates of all the actions that the process term can execute, and coincides with the reciprocal of the average sojourn time in the CTMC-level state corresponding to the process term whenever the process term is performance closed.

Since there are two kinds of actions – exponentially timed and passive – we consider a two-level definition of each variant of exit rate, where level 0 corresponds to exponentially timed actions and level -1 corresponds to passive actions.

Definition 3. *Let $P \in \mathcal{P}_C$, $a \in \text{Name}$, $l \in \{0, -1\}$, and $C \subseteq \mathcal{P}_C$. The exit rate of P when executing actions of name a and level l that lead to C is defined through the following non-negative real function:*

$$\boxed{\text{rate}(P, a, l, C) = \begin{cases} \sum \{ \lambda \in \mathbf{R}_{>0} \mid \exists P' \in C. P \xrightarrow{a, \lambda} P' \} & \text{if } l = 0 \\ \sum \{ w \in \mathbf{R}_{>0} \mid \exists P' \in C. P \xrightarrow{a, *w} P' \} & \text{if } l = -1 \end{cases}}$$

where each summation is taken to be zero whenever its multiset is empty. ■

Definition 4. Let $P \in \mathcal{P}_C$ and $l \in \{0, -1\}$. The total exit rate of P at level l is defined through the following non-negative real function:

$$\boxed{\text{rate}_t(P, l) = \sum_{a \in \text{Name}} \text{rate}(P, a, l, \mathcal{P}_C)}$$

where $\text{rate}(P, a, l, \mathcal{P}_C)$ is the total exit rate of P with respect to a at level l . ■

A computation of a process term is a sequence of transitions that can be executed starting from the state corresponding to the term. The length of a computation is given by the number of transitions occurring in it. We say that two computations are independent of each other if it is not the case that one of them is a proper prefix of the other one. In the following, we denote by $\mathcal{C}_f(P)$ and $\mathcal{I}_f(P)$ the multisets of the finite-length computations and independent computations of $P \in \mathcal{P}_C$. Below we inductively define the trace, the execution probability, and the duration of an element of $\mathcal{C}_f(P)$, using symbol “ \circ ” to denote the sequence concatenation operator.

Definition 5. Let $P \in \mathcal{P}_C$ and $c \in \mathcal{C}_f(P)$. The trace associated with the execution of c is the sequence of the action names labeling the transitions of c , which is defined by induction on the length of c through the following Name^* -valued function:

$$\boxed{\text{trace}(c) = \begin{cases} \varepsilon & \text{if } \text{length}(c) = 0 \\ a \circ \text{trace}(c') & \text{if } c \equiv P \xrightarrow{a, \bar{\lambda}} c' \end{cases}}$$

where ε is the empty trace. ■

Definition 6. Let $P \in \mathcal{P}_{C, \text{pc}}$ and $c \in \mathcal{C}_f(P)$. The probability of executing c is the product of the execution probabilities of the transitions of c , which is defined by induction on the length of c through the following $\mathbf{R}_{[0,1]}$ -valued function:

$$\boxed{\text{prob}(c) = \begin{cases} 1 & \text{if } \text{length}(c) = 0 \\ \frac{\lambda}{\text{rate}_t(P, 0)} \cdot \text{prob}(c') & \text{if } c \equiv P \xrightarrow{a, \lambda} c' \end{cases}}$$

We also define the probability of executing a computation of C as:

$$\boxed{\text{prob}(C) = \sum_{c \in C} \text{prob}(c)}$$

for all $C \subseteq \mathcal{I}_f(P)$. ■

Definition 7. Let $P \in \mathcal{P}_{C, \text{pc}}$ and $c \in \mathcal{C}_f(P)$. The stepwise average duration of c is the sequence of the average sojourn times in the states traversed by c , which is defined by induction on the length of c through the following $(\mathbf{R}_{>0})^*$ -valued function:

$$\boxed{\text{time}_a(c) = \begin{cases} \varepsilon & \text{if } \text{length}(c) = 0 \\ \frac{1}{\text{rate}_t(P, 0)} \circ \text{time}_a(c') & \text{if } c \equiv P \xrightarrow{a, \lambda} c' \end{cases}}$$

where ε is the empty stepwise average duration. We also define the multiset of the computations of C whose stepwise average duration is not greater than θ as:

$$C_{\leq \theta} = \{ \{ c \in C \mid \text{length}(c) \leq \text{length}(\theta) \wedge \forall i = 1, \dots, \text{length}(c). \text{time}_a(c)[i] \leq \theta[i] \} \}$$

for all $C \subseteq \mathcal{C}_f(P)$ and $\theta \in (\mathbf{R}_{>0})^*$. ■

Definition 8. Let $P \in \mathcal{P}_{C,pc}$ and $c \in \mathcal{C}_f(P)$. The stepwise duration of c is the sequence of the random variables quantifying the sojourn times in the states traversed by c , which is defined by induction on the length of c through the following random-variable-sequence-valued function:

$$\text{time}_d(c) = \begin{cases} \varepsilon & \text{if } \text{length}(c) = 0 \\ \text{Exp}_{\text{rate}_t(P,0)} \circ \text{time}_d(c') & \text{if } c \equiv P \xrightarrow{a,\lambda} c' \end{cases}$$

where ε is the empty stepwise duration while $\text{Exp}_{\text{rate}_t(P,0)}$ is the exponentially distributed random variable with rate $\text{rate}_t(P,0) \in \mathbf{R}_{>0}$. ■

Definition 9. Let $P \in \mathcal{P}_{C,pc}$, $C \subseteq \mathcal{I}_f(P)$, and $\theta \in (\mathbf{R}_{>0})^*$. The probability distribution of executing a computation of C within a sequence θ of time units is given by:

$$\text{prob}_d(C, \theta) = \sum_{c \in C}^{\text{length}(c) \leq \text{length}(\theta)} \text{prob}(c) \cdot \prod_{i=1}^{\text{length}(c)} \Pr(\text{time}_d(c)[i] \leq \theta[i])$$

where $\Pr(\text{time}_d(c)[i] \leq \theta[i]) = 1 - e^{-\theta[i]/\text{time}_a(c)[i]}$ is the cumulative distribution function of the exponentially distributed random variable $\text{time}_d(c)[i]$, whose expected value is $\text{time}_a(c)[i]$. ■

We conclude by observing that the average duration (resp. duration) of a finite-length computation has been defined as the sequence of the average sojourn times (resp. of the random variables quantifying the sojourn times) in the states traversed by the computation. The same quantity could have been defined as the sum of the same basic ingredients, but this would not have been appropriate.

Example 1. Consider the two following process terms:

$$\begin{aligned} & \langle g, \gamma \rangle . \langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle g, \gamma \rangle . \langle a, \mu \rangle . \langle d, \lambda \rangle . \underline{0} \\ & \langle g, \gamma \rangle . \langle a, \lambda \rangle . \langle d, \mu \rangle . \underline{0} + \langle g, \gamma \rangle . \langle a, \mu \rangle . \langle b, \lambda \rangle . \underline{0} \end{aligned}$$

with $\lambda \neq \mu$ and $b \neq d$. Observed that the two terms have identical non-maximal computations, we further notice that the first term has the two following maximal computations each with probability 1/2:

$$\begin{aligned} c_{1,1} & \equiv . \xrightarrow{g,\gamma} . \xrightarrow{a,\lambda} . \xrightarrow{b,\mu} . \\ c_{1,2} & \equiv . \xrightarrow{g,\gamma} . \xrightarrow{a,\mu} . \xrightarrow{d,\lambda} . \end{aligned}$$

while the second term has the two following maximal computations each with probability 1/2:

$$\begin{aligned}
c_{2,1} &\equiv . \xrightarrow{g,\gamma} . \xrightarrow{a,\lambda} . \xrightarrow{d,\mu} . \\
c_{2,2} &\equiv . \xrightarrow{g,\gamma} . \xrightarrow{a,\mu} . \xrightarrow{b,\lambda} .
\end{aligned}$$

If the average duration were defined as the sum of the average sojourn times, then $c_{1,1}$ and $c_{2,2}$ would have the same trace $g \circ a \circ b$ and the same average duration $\frac{1}{2\cdot\gamma} + \frac{1}{\lambda} + \frac{1}{\mu}$, and similarly $c_{1,2}$ and $c_{2,1}$ would have the same trace $g \circ a \circ d$ and the same average duration $\frac{1}{2\cdot\gamma} + \frac{1}{\mu} + \frac{1}{\lambda}$. This would lead to conclude that the two terms are equivalent, whereas an external observer equipped with a button-pushing machine displaying the names of the actions that are performed and the instants at which they are performed [49] would be able to distinguish between the two terms. ■

3 Markovian Bisimilarity

Markovian bisimilarity considers two process terms to be equivalent whenever they are able to mimic each other's functional and performance behavior step-wise. In this section we provide the definition of Markovian bisimilarity over \mathcal{P}_C and we recall its properties from [32, 15, 31, 14, 18, 4, 23].

3.1 Equivalence Definition

The basic idea behind Markovian bisimilarity is that, whenever a process term can perform actions with a certain name that reach a certain set of terms at a certain speed, then any process term equivalent to the given one has to be able to respond with actions with the same name that reach an equivalent set of terms at the same speed. This can be formalized through the comparison of the process term exit rates when executing actions of the same name (and level) that lead to the same set of equivalent terms.

Definition 10. *An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_C \times \mathcal{P}_C$ is a Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name}$, levels $l \in \{0, -1\}$, and equivalence classes $C \in \mathcal{P}_C/\mathcal{B}$:*

$$\text{rate}(P_1, a, l, C) = \text{rate}(P_2, a, l, C) \quad \blacksquare$$

Since the union of all the Markovian bisimulations can be proved to be the largest Markovian bisimulation, the definition below follows.

Definition 11. *Markovian bisimilarity, denoted by \sim_{MB} , is the union of all the Markovian bisimulations.* ■

Obviously, \sim_{MB} is strictly finer than classical bisimilarity [43, 41] and probabilistic bisimilarity [40]. We conclude with an easy-to-check necessary condition.

Proposition 1. *Let $P_1, P_2 \in \mathcal{P}_C$. Whenever $P_1 \sim_{\text{MB}} P_2$, then for all $a \in \text{Name}$ and $l \in \{0, -1\}$:*

$$\text{rate}(P_1, a, l, \mathcal{P}_C) = \text{rate}(P_2, a, l, \mathcal{P}_C) \quad \blacksquare$$

3.2 Exactness

Markovian bisimilarity is consistent with an exact aggregation for CTMCs that is known under the name of ordinary lumping.

Definition 12. A partition \mathcal{O} of the state space of a CTMC is an ordinary lumping iff, whenever $s_1, s_2 \in O$ for some $O \in \mathcal{O}$, then for all $O' \in \mathcal{O}$:

$$\sum \{ \lambda \in \mathbf{R}_{>0} \mid \exists s' \in O'. s_1 \xrightarrow{\lambda} s' \} = \sum \{ \lambda \in \mathbf{R}_{>0} \mid \exists s' \in O'. s_2 \xrightarrow{\lambda} s' \} \quad \blacksquare$$

Theorem 1. The CTMC-level aggregation induced by \sim_{MB} is an ordinary lumping. ■

Corollary 1. The CTMC-level aggregation induced by \sim_{MB} is exact. ■

3.3 Congruence

Markovian bisimilarity is a congruence with respect to all the operators of CMPC.

Theorem 2. Let $P_1, P_2 \in \mathcal{P}_C$. Whenever $P_1 \sim_{\text{MB}} P_2$, then:

1. $\langle a, \tilde{\lambda} \rangle . P_1 \sim_{\text{MB}} \langle a, \tilde{\lambda} \rangle . P_2$ for all $\langle a, \tilde{\lambda} \rangle \in \text{Act}_C$.
2. $P_1 + P \sim_{\text{MB}} P_2 + P$ and $P + P_1 \sim_{\text{MB}} P + P_2$ for all $P \in \mathcal{P}_C$.
3. $P_1 \parallel_S P \sim_{\text{MB}} P_2 \parallel_S P$ and $P \parallel_S P_1 \sim_{\text{MB}} P \parallel_S P_2$ for all $S \subseteq \text{Name}$ and $P \in \mathcal{P}_C$. ■

As far as recursion is concerned, we need to extend \sim_{MB} to open process terms. These are terms containing free process variables, i.e. variables not occurring within the scope of a *rec* binder.

Definition 13. Let $P_1, P_2 \in \mathcal{L}_C$ containing a free process variable X . We define $P_1 \sim_{\text{MB}} P_2$ iff $P_1\{Q/X\} \sim_{\text{MB}} P_2\{Q/X\}$ for all $Q \in \mathcal{P}_C$. ■

Theorem 3. Let $P_1, P_2 \in \mathcal{L}_C$ containing a free process variable X . Whenever $P_1 \sim_{\text{MB}} P_2$, then $\text{rec } X : P_1 \sim_{\text{MB}} \text{rec } X : P_2$. ■

3.4 Axiomatization

Markovian bisimilarity has a sound and complete axiomatization over \mathcal{P}_C , whose specific axioms are shown Table 1. This includes three basic laws for alternative composition, two laws characterizing the race policy and the reactive preselection policy, an expansion law for parallel composition, and three laws for recursion. As far as $\mathcal{A}_6^{\text{MB}}$ is concerned, I and J are finite index sets (if empty, the related summations are taken to be $\underline{0}$). The validity of this law is a consequence of the memoryless property of the exponentially distributed durations and of the fact that the probability that two concurrent exponentially timed actions terminate simultaneously is zero, which allows the semantics for the parallel composition operator to be defined in the usual interleaving style.

Theorem 4. The deduction system $\text{DED}(\mathcal{A}^{\text{MB}})$ is sound and complete for \sim_{MB} over \mathcal{P}_C , i.e. for all $P_1, P_2 \in \mathcal{P}_C$:

$$P_1 \sim_{\text{MB}} P_2 \iff \mathcal{A}^{\text{MB}} \vdash P_1 = P_2 \quad \blacksquare$$

$(\mathcal{A}_1^{\text{MB}})$ $(\mathcal{A}_2^{\text{MB}})$ $(\mathcal{A}_3^{\text{MB}})$	$P_1 + P_2 = P_2 + P_1$ $(P_1 + P_2) + P_3 = P_1 + (P_2 + P_3)$ $P + \underline{0} = P$
$(\mathcal{A}_4^{\text{MB}})$ $(\mathcal{A}_5^{\text{MB}})$	$\langle a, \lambda_1 \rangle . P + \langle a, \lambda_2 \rangle . P = \langle a, \lambda_1 + \lambda_2 \rangle . P$ $\langle a, *_{w_1} \rangle . P + \langle a, *_{w_2} \rangle . P = \langle a, *_{w_1+w_2} \rangle . P$
$(\mathcal{A}_6^{\text{MB}})$	$\sum_{i \in I} \langle a_i, \tilde{\lambda}_i \rangle . P_{1,i} \parallel_S \sum_{j \in J} \langle b_j, \tilde{\mu}_j \rangle . P_{2,j} =$ $\sum_{k \in I, a_k \notin S} \langle a_k, \tilde{\lambda}_k \rangle . \left(P_{1,k} \parallel_S \sum_{j \in J} \langle b_j, \tilde{\mu}_j \rangle . P_{2,j} \right) +$ $\sum_{h \in J, b_h \notin S} \langle b_h, \tilde{\mu}_h \rangle . \left(\sum_{i \in I} \langle a_i, \tilde{\lambda}_i \rangle . P_{1,i} \parallel_S P_{2,h} \right) +$ $\sum_{k \in I, a_k \in S, \tilde{\lambda}_k \in \mathbf{R}_{>0}} \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} \langle a_k, \tilde{\lambda}_k \cdot \frac{w_h}{\text{weight}(P_2, b_h)} \rangle . (P_{1,k} \parallel_S P_{2,h}) +$ $\sum_{h \in J, b_h \in S, \tilde{\mu}_h \in \mathbf{R}_{>0}} \sum_{k \in I, a_k = b_h, \tilde{\lambda}_k = *_{v_k}} \langle b_h, \tilde{\mu}_h \cdot \frac{v_k}{\text{weight}(P_1, a_k)} \rangle . (P_{1,k} \parallel_S P_{2,h}) +$ $\sum_{k \in I, a_k \in S, \tilde{\lambda}_k = *_{v_k}} \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} \langle a_k, *_{\frac{v_k}{\text{weight}(P_1, a_k)} \cdot \frac{w_h}{\text{weight}(P_2, b_h)} \cdot (\text{weight}(P_1, a_k) + \text{weight}(P_2, b_h))} \rangle . (P_{1,k} \parallel_S P_{2,h})$
$(\mathcal{A}_7^{\text{MB}})$ $(\mathcal{A}_8^{\text{MB}})$ $(\mathcal{A}_9^{\text{MB}})$	$\text{rec } X : P = \text{rec } Y : P\{Y/X\} \quad \text{if } Y \text{ not in } P$ $\text{rec } X : P = P\{\text{rec } X : P/X\}$ $Q = P\{Q/X\} \Rightarrow Q = \text{rec } X : P$

Table 1. Axiomatization of \sim_{MB} over \mathcal{P}_{C}

3.5 Logical Characterization

Markovian bisimilarity has a modal logic characterization over $\mathcal{P}_{\text{C,pc}}$ based on a Markovian variant of the Hennessy-Milner logic [29], in which the diamond operator is decorated with a rate lower bound.

Definition 14. *The set of the formulas of HML_{MB} is generated by the following syntax:*

$$\boxed{\phi ::= \text{true} \mid \neg\phi \mid \phi \wedge \phi \mid \nabla_a \mid \langle a \rangle_{\lambda} \phi}$$

where $a \in \text{Name}$ and $\lambda \in \mathbf{R}_{>0}$. ■

Definition 15. *The satisfaction relation \models_{MB} of HML_{MB} over $\mathcal{P}_{\text{C,pc}}$ is defined by structural induction as follows:*

$P \models_{\text{MB}} \text{true}$	
$P \models_{\text{MB}} \neg\phi$	if $P \not\models_{\text{MB}} \phi$
$P \models_{\text{MB}} \phi_1 \wedge \phi_2$	if $P \models_{\text{MB}} \phi_1$ and $P \models_{\text{MB}} \phi_2$
$P \models_{\text{MB}} \nabla_a$	if $\text{rate}(P, a, 0, \mathcal{P}_{\text{C}}) = 0$
$P \models_{\text{MB}} \langle a \rangle_{\lambda} \phi$	if $\text{rate}(P, a, 0, \text{sat}(\phi)) \geq \lambda$

where $\text{sat}(\phi) = \{P' \in \mathcal{P}_{\text{C}} \mid P' \models_{\text{MB}} \phi\}$. ■

Theorem 5. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Then:*

$$P_1 \sim_{\text{MB}} P_2 \iff (\forall \phi \in \text{HML}_{\text{MB}}. P_1 \models_{\text{MB}} \phi \iff P_2 \models_{\text{MB}} \phi) \quad \blacksquare$$

We also mention that Markovian bisimilarity has a temporal logic characterization based on the branching-time logic CSL [3]. Besides the usual propositional connectives, this logic comprises:

- a probability operator, which replaces the universal and existential computation quantifiers and allows to refer to the probability of performing a computation that satisfies a certain formula;
- a time-bounded next operator, which expresses that the next state is reached within a certain amount of time and a certain formula holds in it;
- a time-bounded until operator, which expresses that a certain formula is satisfied at some instant no later than a certain amount of time and at all preceding instants another given formula holds;
- a steady-state operator, which enables to reason about the probability to be in the long run in some state that satisfies a certain formula.

3.6 Verification Complexity

Markovian bisimilarity can be decided in polynomial time by means of a partition refinement algorithm in the style of [42]. This algorithm exploits the fact that \sim_{MB} can be characterized as a fixed point of successively finer relations. In fact we have:

$$\sim_{\text{MB}} = \bigcap_{i \in \mathbf{N}} \sim_{\text{MB},i}$$

where $\sim_{\text{MB},0} = \mathcal{P}_C \times \mathcal{P}_C$ and $\sim_{\text{MB},i}$ for $i \geq 1$ is defined as follows: whenever $(P_1, P_2) \in \sim_{\text{MB},i}$, then for all $a \in \text{Name}$, $l \in \{0, -1\}$, and $C \in \mathcal{P}_C / \sim_{\text{MB},i-1}$:

$$\text{rate}(P_1, a, l, C) = \text{rate}(P_2, a, l, C)$$

In other words, $\sim_{\text{MB},0}$ induces a trivial partition with a single equivalence class that coincides with \mathcal{P}_C , $\sim_{\text{MB},1}$ refines the previous partition by creating an equivalence class for each set of terms that satisfy the necessary condition stated by Prop. 1, and so on.

The algorithm to check whether $P_1 \sim_{\text{MB}} P_2$ thus proceeds as follows:

1. Build a partition with a single class including all the states of the disjoint union of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$, then initialize a list of splitters with this class.
2. Refine the current partition by splitting each of its classes according to the exit rates towards one of the splitters, then remove this splitter from the list.
3. For each split class, insert into the list of splitters all the resulting subclasses except for the largest one.
4. If the list of splitters is empty, return yes/no depending on whether the initial state of $\llbracket P_1 \rrbracket$ and the initial state of $\llbracket P_2 \rrbracket$ belong to the same class or not, otherwise go back to the refinement step.

The time complexity is $O(m \cdot \log n)$, where n is the number of states and m is the number of transitions of the disjoint union of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$. To achieve this complexity it is necessary to resort to a splay tree when representing the subclasses arising from the splitting of a class.

4 Variants of Markovian Bisimilarity

Due to the nice properties presented in the previous section, Markovian bisimilarity has received more attention than Markovian testing and trace equivalences. For this reason, some useful variants of it have appeared in the literature, three of which will be recalled in this section.

The first one – reward Markovian bisimilarity [10] – takes rewards into account in order to allow process terms to be compositionally manipulated in a way that is sensitive to specific performance measures. The second one – interactive Markovian bisimilarity [30] – stems from the interaction of nondeterministic process calculi and CTMCs and deals with both Markovian branchings and non-deterministic branchings. The third one – extended Markovian bisimilarity [8] – considers immediate actions à la GSPN [1] and deals with both Markovian branchings and prioritized/probabilistic branchings.

4.1 Reward Markovian Bisimilarity

Although Markovian bisimilarity is consistent with ordinary lumping, which is an exact aggregation, it may happen that specific performance measures distinguish between ordinarily lumpable states by ascribing them a different meaning.

The usual approach to the specification of performance measures for CTMC-based models relies on reward structures [34]. This requires associating real numbers with model behaviors and activities, which are then transferred to the proper states (as yield rewards) and transitions (as bonus rewards) of the underlying CTMC, respectively. A yield reward expresses the rate at which a gain (or a loss, if the number is negative) is accumulated while sojourning in the related state. By contrast, a bonus reward specifies the instantaneous gain (or loss) implied by the execution of the related transition.

The instant-of-time value of a performance measure specified through a reward structure is computed as follows for a CTMC:

$$\boxed{\sum_s yr(s) \cdot \pi(s) + \sum_{s \xrightarrow{\lambda} s'} br(s, \lambda, s') \cdot \phi(s, \lambda, s')}$$

where:

- $yr(s)$ is the yield reward associated with state s .
- $\pi(s)$ is the probability of being in state s at the considered instant of time.
- $br(s, \lambda, s')$ is the bonus reward associated with transition $s \xrightarrow{\lambda} s'$.
- $\phi(s, \lambda, s')$ is the frequency of transition $s \xrightarrow{\lambda} s'$ at the considered instant of time, which is given by $\pi(s) \cdot \lambda$.

In this setting, ascribing a different meaning to ordinarily lumpable states amounts to giving different rewards to such states or their outgoing transitions. In order to manipulate process terms in a performance-sensitive way, the definition of Markovian bisimilarity can be modified by taking rewards into account.

Before doing that, we need to extend CMPC with rewards. While bonus rewards can naturally be associated with actions, the treatment of yield rewards is not trivial because in process calculi the concept of state is implicit. An approach that can be followed is to associate yield rewards with actions too, with the yield reward of a state being given by the sum of the yield rewards associated with the actions enabled in that state (additivity assumption).

Therefore, we derive from CMPC a new calculus, which we call concurrent reward Markovian process calculus (CRMPC):

- The syntax is extended as follows:

$$\boxed{
 \begin{array}{l}
 P ::= \mathbf{0} \\
 \quad | \langle a, \lambda, yr, br \rangle . P \\
 \quad | \langle a, *w, *, * \rangle . P \\
 \quad | P + P \\
 \quad | P \parallel_S P \\
 \quad | X \\
 \quad | \text{rec } X : P
 \end{array}
 }$$

where $yr \in \mathbf{R}$ is the yield reward and $br \in \mathbf{R}$ is the bonus reward. We denote by \mathcal{P}_{CR} the set of the closed and guarded process terms of CRMPC.

- The semantic rules are modified accordingly. In particular, the synchronization rules change as follows:

$$\boxed{
 \begin{array}{c}
 \frac{P_1 \xrightarrow{a, \lambda, yr, br} P'_1 \quad P_2 \xrightarrow{a, *w, *, *} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda \cdot \frac{w}{\text{weight}(P_2, a)}, yr \cdot \frac{w}{\text{weight}(P_2, a)}, br} P'_1 \parallel_S P'_2} \\
 \\
 \frac{P_1 \xrightarrow{a, *w, *, *} P'_1 \quad P_2 \xrightarrow{a, \lambda, yr, br} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda \cdot \frac{w}{\text{weight}(P_1, a)}, yr \cdot \frac{w}{\text{weight}(P_1, a)}, br} P'_1 \parallel_S P'_2} \\
 \\
 \frac{P_1 \xrightarrow{a, *w_1, *, *} P'_1 \quad P_2 \xrightarrow{a, *w_2, *, *} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, * \frac{w_1}{\text{weight}(P_1, a)} \cdot \frac{w_2}{\text{weight}(P_2, a)} \cdot (\text{weight}(P_1, a) + \text{weight}(P_2, a)), *, *} P'_1 \parallel_S P'_2}
 \end{array}
 }$$

Note that the yield rewards are subject to the same normalization as the rates, because they are strictly related to the sojourn time in the states. By contrast, no normalization is needed for the bonus rewards, as they are earned upon the execution of the transitions.

We now introduce the concept of exit reward, on the basis of which we shall define the performance-sensitive variant of Markovian bisimilarity.

Definition 16. *Let $P \in \mathcal{P}_{\text{CR}}$, $a \in \text{Name}$, $l \in \{0, -1\}$, and $C \subseteq \mathcal{P}_{\text{CR}}$. The exit reward of P when executing actions of name a and level l that lead to C is*

defined through the following real function:

$$\text{reward}(P, a, l, C) = \begin{cases} \sum \{ \{ yr + \lambda \cdot br \in \mathbf{R} \mid \exists P' \in C. P \xrightarrow{a, \lambda, yr, br} P' \} \} & \text{if } l = 0 \\ 0 & \text{if } l = -1 \end{cases}$$

where the summation is taken to be zero whenever its multiset is empty. ■

Definition 17. An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_{\text{CR}} \times \mathcal{P}_{\text{CR}}$ is a reward Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name}$, levels $l \in \{0, -1\}$, and equivalence classes $C \in \mathcal{P}_{\text{CR}}/\mathcal{B}$:

$$\begin{aligned} \text{rate}(P_1, a, l, C) &= \text{rate}(P_2, a, l, C) \\ \text{reward}(P_1, a, l, C) &= \text{reward}(P_2, a, l, C) \end{aligned}$$

■

Definition 18. Reward Markovian bisimilarity, denoted by \sim_{RMB} , is the union of all the reward Markovian bisimulations. ■

\sim_{RMB} enjoys the same properties as \sim_{MB} . The characterizing axioms are:

$$\begin{aligned} \langle a, \lambda_1, yr_1, br_1 \rangle . P + \langle a, \lambda_2, yr_2, br_2 \rangle . P = \\ \langle a, \lambda_1 + \lambda_2, yr_1 + yr_2, \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot br_1 + \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot br_2 \rangle . P \\ \langle a, *_{w_1}, *, * \rangle . P + \langle a, *_{w_2}, *, * \rangle . P = \langle a, *_{w_1 + w_2}, *, * \rangle . P \end{aligned}$$

or equivalently:

$$\begin{aligned} \langle a, \lambda, yr, br \rangle . P = \langle a, \lambda, yr + \lambda \cdot br, 0 \rangle . P \\ \langle a, \lambda_1, yr_1, 0 \rangle . P + \langle a, \lambda_2, yr_2, 0 \rangle . P = \langle a, \lambda_1 + \lambda_2, yr_1 + yr_2, 0 \rangle . P \\ \langle a, *_{w_1}, *, * \rangle . P + \langle a, *_{w_2}, *, * \rangle . P = \langle a, *_{w_1 + w_2}, *, * \rangle . P \end{aligned}$$

or equivalently:

$$\begin{aligned} \langle a, \lambda, yr, br \rangle . P = \langle a, \lambda, 0, br + \frac{yr}{\lambda} \rangle . P \\ \langle a, \lambda_1, 0, br_1 \rangle . P + \langle a, \lambda_2, 0, br_2 \rangle . P = \langle a, \lambda_1 + \lambda_2, 0, \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot br_1 + \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot br_2 \rangle . P \\ \langle a, *_{w_1}, *, * \rangle . P + \langle a, *_{w_2}, *, * \rangle . P = \langle a, *_{w_1 + w_2}, *, * \rangle . P \end{aligned}$$

4.2 Interactive Markovian Bisimilarity

Markovian bisimilarity deals with fully probabilistic models. However, it may happen that not all the details of a model are known in the early design stages. In that case, nondeterminism is a useful abstraction.

A way to recover nondeterminism in a Markovian framework is to combine nondeterministic process calculi and CTMCs. This can be accomplished by replacing durational actions with two distinct prefix operators – one for interacting actions that take no time and one for exponential delays – with the inter-process communication being implemented through the synchronization of visible interacting actions. Since an invisible action – which we denote by τ as usual – takes no time and cannot be prevented by any synchronization constraint, we assume maximal progress, i.e. τ -actions take precedence over time passing.

We now derive from CMPC a new calculus, which we call concurrent interactive Markovian process calculus (CIMPC):

- The syntax is modified as follows:

$$\boxed{
\begin{array}{l}
P ::= 0 \\
\quad | a.P \\
\quad | (\lambda).P \\
\quad | P + P \\
\quad | P \parallel_S P \\
\quad | X \\
\quad | \text{rec } X : P
\end{array}
}$$

where $a \in \text{Name} \cup \{\tau\}$. We denote by \mathcal{P}_{CI} the set of the closed and guarded process terms of CIMPC.

- Two transition relations are necessary: one for interacting actions and one for exponential delays. Besides handling recursion, the semantic rules for interacting actions include:

$$\boxed{
\begin{array}{c}
a.P \xrightarrow{\text{I}} P \\
\\
\frac{P_1 \xrightarrow{\text{I}} P'}{P_1 + P_2 \xrightarrow{\text{I}} P'} \quad \frac{P_2 \xrightarrow{\text{I}} P'}{P_1 + P_2 \xrightarrow{\text{I}} P'} \\
\\
\frac{P_1 \xrightarrow{\text{I}} P' \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{\text{I}} P' \parallel_S P_2} \quad \frac{P_2 \xrightarrow{\text{I}} P' \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{\text{I}} P_1 \parallel_S P'_2} \\
\\
\frac{P_1 \xrightarrow{\text{I}} P' \quad P_2 \xrightarrow{\text{I}} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{\text{I}} P'_1 \parallel_S P'_2}
\end{array}
}$$

while the semantic rules for exponential delays include:

$$\boxed{
\begin{array}{c}
(\lambda).P \xrightarrow{\text{M}} P \\
\\
\frac{P_1 \xrightarrow{\text{M}} P'}{P_1 + P_2 \xrightarrow{\text{M}} P'} \quad \frac{P_2 \xrightarrow{\text{M}} P'}{P_1 + P_2 \xrightarrow{\text{M}} P'} \\
\\
\frac{P_1 \xrightarrow{\text{M}} P'_1}{P_1 \parallel_S P_2 \xrightarrow{\text{M}} P'_1 \parallel_S P_2} \quad \frac{P_2 \xrightarrow{\text{M}} P'_2}{P_1 \parallel_S P_2 \xrightarrow{\text{M}} P_1 \parallel_S P'_2}
\end{array}
}$$

Before defining the interactive variant of Markovian bisimilarity, we adapt to this setting the concept of exit rate.

Definition 19. Let $P \in \mathcal{P}_{\text{CI}}$ and $C \subseteq \mathcal{P}_{\text{CI}}$. The exit rate of P towards C is defined through the following non-negative real function:

$$\boxed{\text{rate}_{\text{M}}(P, C) = \sum \{ \lambda \in \mathbf{R}_{>0} \mid \exists P' \in C. P \xrightarrow{\text{M}} P' \}}$$

where the summation is taken to be zero whenever its multiset is empty. ■

Definition 20. An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_{\text{CI}} \times \mathcal{P}_{\text{CI}}$ is an *interactive Markovian bisimulation* iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name} \cup \{\tau\}$ and equivalence classes $C \in \mathcal{P}_{\text{CI}}/\mathcal{B}$:

- $P_1 \xrightarrow{a}_I P'_1$ implies $P_2 \xrightarrow{a}_I P'_2$ for some P'_2 with $(P'_1, P'_2) \in \mathcal{B}$.
- $P_1 \not\xrightarrow{\tau}_I$ implies $\text{rate}_{\text{M}}(P_1, C) = \text{rate}_{\text{M}}(P_2, C)$. ■

Definition 21. *Interactive Markovian bisimilarity*, denoted by \sim_{IMB} , is the union of all the interactive Markovian bisimulations. ■

\sim_{IMB} enjoys properties similar to those of \sim_{MB} . The characterizing axioms are:

$$\boxed{\begin{array}{l} a.P + a.P = a.P \\ (\lambda_1).P + (\lambda_2).P = (\lambda_1 + \lambda_2).P \\ \tau.P + (\lambda).Q = \tau.P \end{array}}$$

which encode idempotency, race policy, and maximal progress, respectively.

Since τ -actions are invisible and take no time, when comparing process terms they should be abstracted away. This can be achieved by weakening \sim_{IMB} in such a way that, after any non-pre-emptable exponential delay, all the states that can evolve via a finite sequence of τ -transitions to a given class are skipped.

Definition 22. Let $C \subseteq \mathcal{P}_{\text{CI}}$. The *internal backward closure* of C is defined as follows:

$$C_\tau = \{P' \in \mathcal{P}_{\text{CI}} \mid \exists P \in C. P' \xrightarrow{\tau^*}_I P\}$$

where $P' \xrightarrow{\tau^*}_I P$ means that P' can evolve to P after a finite sequence of zero or more τ -transitions. ■

Definition 23. An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_{\text{CI}} \times \mathcal{P}_{\text{CI}}$ is a *weak interactive Markovian bisimulation* iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name}$ and equivalence classes $C \in \mathcal{P}_{\text{CI}}/\mathcal{B}$:

- $P_1 \xrightarrow{a}_I P'_1$ implies $P_2 \xrightarrow{\tau^* a \tau^*}_I P'_2$ for some P'_2 with $(P'_1, P'_2) \in \mathcal{B}$.
- $P_1 \xrightarrow{\tau}_I P'_1$ implies $P_2 \xrightarrow{\tau^*}_I P'_2$ for some P'_2 with $(P'_1, P'_2) \in \mathcal{B}$.
- $P_1 \xrightarrow{\tau^*}_I P'_1 \not\xrightarrow{\tau}_I$ implies $P_2 \xrightarrow{\tau^*}_I P'_2 \not\xrightarrow{\tau}_I$ for some P'_2 with $\text{rate}_{\text{M}}(P'_1, C_\tau) = \text{rate}_{\text{M}}(P'_2, C_\tau)$. ■

Definition 24. *Weak interactive Markovian bisimilarity*, denoted by \approx_{IMB} , is the union of all the weak interactive Markovian bisimulations. ■

\approx_{IMB} is strictly coarser than \sim_{IMB} but it is not a congruence with respect to alternative composition. As usual, initial τ -actions need a different treatment.

Definition 25. Let $P_1, P_2 \in \mathcal{P}_{\text{CI}}$. We say that P_1 is *weakly interactive Markovian bisimulation congruent* to P_2 , written $P_1 \simeq_{\text{IMB}} P_2$, iff for all action names $a \in \text{Name} \cup \{\tau\}$ and equivalence classes $C \in \mathcal{P}_{\text{CI}}/\approx_{\text{IMB}}$:

- $P_1 \xrightarrow{a}_I P'_1$ implies $P_2 \xrightarrow{\tau^* a \tau^*}_I P'_2$ for some P'_2 with $P'_1 \approx_{\text{IMB}} P'_2$.
- $P_2 \xrightarrow{a}_I P'_2$ implies $P_1 \xrightarrow{\tau^* a \tau^*}_I P'_1$ for some P'_1 with $P'_1 \approx_{\text{IMB}} P'_2$.
- $P_1 \not\xrightarrow{\tau}_I$ iff $P_2 \not\xrightarrow{\tau}_I$.
- $P_1 \not\xrightarrow{\tau}_I$ implies $\text{rate}_M(P_1, C) = \text{rate}_M(P_2, C)$. ■

It turns out $\sim_{\text{IMB}} \subset \simeq_{\text{IMB}} \subset \approx_{\text{IMB}}$ with \simeq_{IMB} enjoying the same properties as \sim_{IMB} . Moreover, \simeq_{IMB} has the following additional characterizing axioms:

$$\boxed{\begin{array}{c} a.\tau.P = a.P \\ P + \tau.P = \tau.P \\ a.(P + \tau.Q) + \tau.Q = a.(P + \tau.Q) \\ (\lambda).\tau.P = (\lambda).P \end{array}}$$

which are called τ -laws and witness the capability of \simeq_{IMB} of abstracting from τ -actions.

4.3 Extended Markovian Bisimilarity

Markovian bisimilarity is restricted to exponential distributions. Although their combinations can approximate most of the general distributions arbitrarily closely, some useful distributions are left out, specially the one representing a zero duration. The capability of expressing zero durations would also constitute a good performance abstraction mechanism, similarly to the functional abstraction mechanism given by the invisible action name τ .

In the modeling process it often happens to deal with choices among logical events (like the reception of a message vs. its loss) with which no timing can reasonably be associated, or to encounter activities that are several orders of magnitude faster than the activities that are important for the evaluation of certain performance measures. In all of these cases, using a zero duration would be an adequate solution from the modeling standpoint.

Zero durations can be introduced by admitting the so-called immediate actions. Each of them has a name and a zero duration (or, equivalently, an infinite rate), together with a priority level $l \in \mathbf{N}_{>0}$ and a weight $w \in \mathbf{R}_{>0}$. Priority levels and weights are used to choose among several immediate actions that are simultaneously enabled. According to the generative [26] preselection policy, each of the highest priority immediate actions that are enabled is given an execution probability proportional to its weight.

It is worth noting that this extended Markovian framework is complementary with respect to the interactive Markovian framework, because it is as if nondeterminism were ruled out by augmenting each interacting action with a priority level and a weight. Here maximal progress is subsumed by pre-emption: immediate τ -actions take precedence over all the lower priority actions.

We now derive from CMPC a new calculus, which we call concurrent extended Markovian process calculus (CEMPC):

- The syntax is extended as follows:

$$\boxed{
\begin{array}{l}
P ::= 0 \\
| \langle a, \lambda \rangle . P \\
| \langle a, \infty_{l,w} \rangle . P \\
| \langle a, *_{w}^{l'} \rangle . P \\
| P + P \\
| P \parallel_S P \\
| X \\
| \text{rec } X : P
\end{array}
}$$

where $a \in \text{Name} \cup \{\tau\}$, $l \in \mathbb{N}_{>0}$, and $l' \in \mathbb{N}$. We denote by \mathcal{P}_{CE} the set of the closed and guarded process terms of CEMPC. We point out that every passive action has been augmented with a priority constraint, which is useful to keep under control the process priority interrelation. Besides synchronizing with passive actions with the same priority constraint, a passive action with priority constraint 0 can only synchronize with an exponentially timed action, while a passive action with priority constraint $l' > 0$ can only synchronize with an immediate action with priority level $l = l'$.

- The additional semantic rules specific for immediate actions are the following:

$$\boxed{
\begin{array}{c}
\langle a, \infty_{l,w} \rangle . P \xrightarrow{a, \infty_{l,w}} P \\
\\
\frac{P_1 \xrightarrow{a, \infty_{l,w}} P'_1 \quad P_2 \xrightarrow{a, *_{v}^{l'}} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \infty_{l,w} \cdot \frac{v}{\text{weight}(P_2, a, l)}} P'_1 \parallel_S P'_2} \\
\\
\frac{P_1 \xrightarrow{a, *_{v}^{l'}} P'_1 \quad P_2 \xrightarrow{a, \infty_{l,w}} P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \infty_{l,w} \cdot \frac{v}{\text{weight}(P_1, a, l)}} P'_1 \parallel_S P'_2}
\end{array}
}$$

Note that, consistently with the asymmetric action synchronization discipline, immediate actions can synchronize only with passive actions.

Before defining the extended variant of Markovian bisimilarity, we extend to immediate actions the notion of exit rate. Below, the priority level of an action is encoded through a number in \mathbb{Z} , which is 0 if the action is exponentially timed, l if the action rate is $\infty_{l,w}$, $-l - 1$ if the action rate is $*_{w}^{l}$.

Definition 26. Let $P \in \mathcal{P}_{\text{CE}}$, $a \in \text{Name} \cup \{\tau\}$, $l \in \mathbb{Z}$, and $C \subseteq \mathcal{P}_{\text{CE}}$. The exit rate of P when executing actions of name a and priority level l that lead to C is defined through the following non-negative real function:

$$\boxed{
\text{rate}(P, a, l, C) = \begin{cases} \sum \{ \lambda \in \mathbb{R}_{>0} \mid \exists P' \in C. P \xrightarrow{a, \lambda} P' \} & \text{if } l = 0 \\ \sum \{ w \in \mathbb{R}_{>0} \mid \exists P' \in C. P \xrightarrow{a, \infty_{l,w}} P' \} & \text{if } l > 0 \\ \sum \{ w \in \mathbb{R}_{>0} \mid \exists P' \in C. P \xrightarrow{a, *_{w}^{-l-1}} P' \} & \text{if } l < 0 \end{cases}
}$$

where each summation is taken to be zero whenever its multiset is empty. ■

In the following we denote by $\text{pri}_\infty^\tau(P)$ the priority level of the highest priority immediate τ -action enabled by P , and we set $\text{pri}_\infty^\tau(P) = 0$ if P does not enable any immediate τ -action. Moreover, given $l \in \mathbb{Z}$, we use $\text{no-pre}(l, P)$ to denote that no action whose priority level is l can be pre-empted in P . Formally, this is the case whenever $l \geq \text{pri}_\infty^\tau(P)$ or $-l - 1 \geq \text{pri}_\infty^\tau(P)$.

Definition 27. An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_{\text{CE}} \times \mathcal{P}_{\text{CE}}$ is an extended Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name} \cup \{\tau\}$, equivalence classes $C \in \mathcal{P}_{\text{CE}}/\mathcal{B}$, and priority levels $l \in \mathbb{Z}$ such that $\text{no-pre}(l, P_1)$ and $\text{no-pre}(l, P_2)$:

$$\text{rate}(P_1, a, l, C) = \text{rate}(P_2, a, l, C) \quad \blacksquare$$

Definition 28. Extended Markovian bisimilarity, denoted by \sim_{EMB} , is the union of all the extended Markovian bisimulations. ■

\sim_{EMB} enjoys the same properties as \sim_{MB} . The characterizing axioms are:

$\begin{aligned} &\langle a, \lambda_1 \rangle.P + \langle a, \lambda_2 \rangle.P = \langle a, \lambda_1 + \lambda_2 \rangle.P \\ &\langle a, \infty_{l, w_1} \rangle.P + \langle a, \infty_{l, w_2} \rangle.P = \langle a, \infty_{l, w_1 + w_2} \rangle.P \\ &\langle a, *_{w_1}^l \rangle.P + \langle a, *_{w_2}^l \rangle.P = \langle a, *_{w_1 + w_2}^l \rangle.P \\ &\langle \tau, \infty_{l, w} \rangle.P + \langle a, \lambda \rangle.Q = \langle \tau, \infty_{l, w} \rangle.P \\ &\langle \tau, \infty_{l, w} \rangle.P + \langle a, \infty_{l', w'} \rangle.Q = \langle \tau, \infty_{l, w} \rangle.P \quad \text{if } l > l' \\ &\langle \tau, \infty_{l, w} \rangle.P + \langle a, *_{w'}^{l'} \rangle.Q = \langle \tau, \infty_{l, w} \rangle.P \quad \text{if } l > l' \end{aligned}$
--

with the last three encoding pre-emption.

Similarly to the interactive framework, when comparing process terms the immediate τ -actions should be abstracted away, as they are unobservable both from the functional viewpoint and from the performance viewpoint. However, weakening \sim_{EMB} is harder than weakening \sim_{IMB} , because it is necessary to keep track of the priority levels and of the weights associated with the actions to be abstracted away. Furthermore, it is also necessary to take into account the degree of observability of the states.

Definition 29. Let $P \in \mathcal{P}_{\text{CE}}$ and $l \in \mathbb{N}_{>0}$. We say that P is l -unobservable iff $\text{pri}_\infty^\tau(P) = l$ and P does not enable any visible action with priority level $l' \in \mathbb{Z}$ such that $l' \geq l$ or $-l' - 1 \geq l$. ■

Definition 30. Let $n \in \mathbb{N}_{>0}$ and $P_1, P_2, \dots, P_{n+1} \in \mathcal{P}_{\text{CE}}$. A computation c of length n :

$$P_1 \xrightarrow{\tau, \infty_{l_1, w_1}} P_2 \xrightarrow{\tau, \infty_{l_2, w_2}} \dots \xrightarrow{\tau, \infty_{l_n, w_n}} P_{n+1}$$

is unobservable iff for all $i = 1, \dots, n$ process term P_i is l_i -unobservable. In that case, the probability of executing c is given by:

$$\text{prob}(c) = \prod_{i=1}^n \frac{w_i}{\text{rate}(P_i, \tau, l_i, \mathcal{P}_{\text{CE}})} \quad \blacksquare$$

Definition 31. Let $P \in \mathcal{P}_{\text{CE}}$, $a \in \text{Name} \cup \{\tau\}$, $l \in \mathbb{Z}$, and $C \subseteq \mathcal{P}_{\text{CE}}$. The weak exit rate of P when executing actions with name a and priority level l that lead to C is defined through the following non-negative real function:

$$\boxed{\text{rate}_w(P, a, l, C) = \sum_{P' \in C_w} \text{rate}(P, a, l, \{P'\}) \cdot \text{prob}_w(P', C)}$$

where:

- C_w is the weak backward closure of C :
 $C_w = C \cup \{Q \in \mathcal{P}_{\text{CE}} - C \mid Q \text{ can reach } C \text{ via unobservable computations}\}$
- prob_w is a $\mathbf{R}_{[0,1]}$ -valued function representing the sum of the probabilities of all the unobservable computations from a term in C_w to C :

$$\text{prob}_w(P', C) = \begin{cases} 1 & \text{if } P' \in C \\ \sum \{\text{prob}(c) \mid c \text{ unobservable computation from } P' \text{ to } C\} & \text{if } P' \in C_w - C \end{cases} \quad \blacksquare$$

The definition of \sim_{EMB} can be weakened by using rate_w instead of rate and by skipping the weak exit rate comparison for some equivalence classes that contain certain unobservable states:

- An observable state is a state that enables an observable action that cannot be pre-empted by any enabled unobservable action.
- An initially unobservable state is a state in which all the enabled observable actions are pre-empted by some enabled unobservable action, but at least one of the computations starting at this state with one of the higher priority enabled unobservable actions reaches an observable state.
- A fully unobservable state is a state in which all the enabled observable actions are pre-empted by some enabled unobservable action, and all the computations starting at this state with one of the higher priority enabled unobservable actions are unobservable (note that $\underline{0}$ is fully unobservable). We denote by $\mathcal{P}_{\text{CE}, \text{fu}}$ the set of the fully unobservable process terms of \mathcal{P}_{CE} .

Definition 32. An equivalence relation $\mathcal{B} \subseteq \mathcal{P}_{\text{CE}} \times \mathcal{P}_{\text{CE}}$ is a weak extended Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name} \cup \{\tau\}$ and priority levels $l \in \mathbb{Z}$ such that $\text{no-pre}(l, P_1)$ and $\text{no-pre}(l, P_2)$:

$$\begin{aligned} \text{rate}_w(P_1, a, l, C) &= \text{rate}_w(P_2, a, l, C) && \text{for all observable } C \in \mathcal{P}_{\text{CE}}/\mathcal{B} \\ \text{rate}_w(P_1, a, l, \mathcal{P}_{\text{CE}, \text{fu}}) &= \text{rate}_w(P_2, a, l, \mathcal{P}_{\text{CE}, \text{fu}}) && \blacksquare \end{aligned}$$

Definition 33. Weak extended Markovian bisimilarity, denoted by \approx_{EMB} , is the union of all the weak extended Markovian bisimulations. \blacksquare

\approx_{EMB} enjoys the same properties as \sim_{EMB} except for congruence. In fact, to recover congruence with respect to parallel composition, we have to restrict ourselves to the set $\mathcal{P}_{\text{CE}, \text{wp}}$ of the well-prioritized process terms of \mathcal{P}_{CE} . This is the smallest subset of \mathcal{P}_{CE} closed with respect to null term, action prefix, alternative composition, recursion and closed with respect to parallel composition under the

following constraint: If $P_1, P_2 \in \mathcal{P}_{\text{CE,wp}}$ and any immediate/passive transition of each of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ has priority level/constraint less than the priority level of any unobservable transition departing from an unobservable state of the other one, then $P_1 \parallel_S P_2 \in \mathcal{P}_{\text{CE,wp}}$.

The additional characterizing axioms of \approx_{EMB} over $\mathcal{P}_{\text{CE,wp}}$ are the following:

$$\boxed{\begin{array}{l} \langle a, \lambda \rangle \cdot \sum_{i \in I} \langle \tau, \infty_{l, w_i} \rangle \cdot P_i = \sum_{i \in I} \langle a, \lambda \cdot w_i / \sum_{k \in I} w_k \rangle \cdot P_i \\ \langle a, \infty_{l', w'} \rangle \cdot \sum_{i \in I} \langle \tau, \infty_{l, w_i} \rangle \cdot P_i = \sum_{i \in I} \langle a, \infty_{l', w' \cdot w_i / \sum_{k \in I} w_k} \rangle \cdot P_i \\ \langle a, *_{w'}^{l'} \rangle \cdot \sum_{i \in I} \langle \tau, \infty_{l, w_i} \rangle \cdot P_i = \sum_{i \in I} \langle a, *_{w' \cdot w_i / \sum_{k \in I} w_k}^{l'} \rangle \cdot P_i \end{array}}$$

which witness the capability of \approx_{EMB} of abstracting from immediate τ -actions in a way that correctly takes into account their weights.

5 Markovian Testing Equivalence

Markovian testing equivalence considers two process terms to be equivalent whenever an external observer, who can interact with them by means of tests, is not able to distinguish between them from the functional or performance viewpoint. In this section we provide the definition of Markovian testing equivalence over $\mathcal{P}_{\text{C,pc}}$ and we recall its properties from [7, 9].

5.1 Test Formalization

The only way that the external observer has to infer information about the behavior of the process terms is to interact with them by means of tests and look at their reactions. Was the test passed? If so, with which probability? And how long did it take to pass the test?

In our Markovian framework with asymmetric action synchronization discipline, the most convenient way to represent a test is through another process term composed of passive actions only, which interacts with the terms to be tested by means of a parallel composition operator that enforces synchronization on any action name. In this way, the parallel composition of a performance closed term to be tested and a test will still be performance closed.

From the testing viewpoint, in any of its states a process term to be tested generates the proposal of an action to be executed by means of a race among the exponentially timed actions enabled in that state. Then the test either reacts by participating in the interaction with the process term through a passive action having the same name as the proposed exponentially timed action, or blocks the interaction if it has no passive actions with the proposed name.

Since it is necessary to measure the probability with which process terms pass tests within a finite amount of time, for the test formalization we can restrict ourselves to non-recursive terms (composed of passive actions only). In other words, the expressiveness provided by labeled multitransition systems with a finite dag-like structure will be enough for the tests.

In order to represent the fact that a test is passed or not, each of the terminal nodes of the dag-like semantic model underlying a test must be suitably labeled so as to establish whether it is a success or failure state. At the process calculus level, this amounts to replace $\underline{0}$ with two zeroary operators, which we denote by “s” (for success) and “f” (for failure).

Ambiguous terms like $s + f$ will be avoided in the test syntax by replacing the action prefix operator and the binary alternative composition operator with a set of n -ary guarded alternative composition operators, with n ranging over the whole $\mathbf{N}_{>0}$.

Definition 34. *The set \mathcal{T} of the tests is generated by the following syntax:*

$$\boxed{\begin{array}{l} T ::= f \\ \quad | \quad s \\ \quad | \quad \sum_{i \in I} \langle a_i, *_{w_i} \rangle . T_i \end{array}}$$

where I is a non-empty finite index set. ■

5.2 Equivalence Definition

Markovian testing equivalence relies on comparing the process term probabilities of performing a successful test-driven computation within a given sequence of average amounts of time. A test-driven computation is a sequence of transitions in the labeled multitransition system underlying the parallel composition of a process term and a test. Due to the restrictions imposed on the test syntax, all the considered test-driven computations will turn out to have a finite length, hence the inductive definitions of Sect. 2.3 apply to them.

Definition 35. *Let $P \in \mathcal{P}_{C,pc}$ and $T \in \mathcal{T}$. The interaction system of P and T is process term $P \parallel_{Name} T \in \mathcal{P}_{C,pc}$, where we say that:*

- A configuration is a state of $\llbracket P \parallel_{Name} T \rrbracket$.
- A configuration is formed by a process part and a test part.
- A configuration is successful (resp. failed) iff its test part is “s” (resp. “f”).
- A computation is successful (resp. failed) iff so is its last configuration.
- A computation that is neither successful nor failed is interrupted.

We denote by $\mathcal{SC}(P, T)$ the multiset of the successful computations of $\mathcal{C}_f(P \parallel_{Name} T)$. ■

Note that $\mathcal{SC}(P, T) \subseteq \mathcal{I}_f(P \parallel_{Name} T)$, because of the maximality of the successful test-driven computations, and that $\mathcal{SC}(P, T)$ is finite, because of the finitely-branching structure of the considered terms.

Definition 36. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian testing equivalent to P_2 , written $P_1 \sim_{MT} P_2$, iff for all tests $T \in \mathcal{T}$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:*

$$prob(\mathcal{SC}_{\leq \theta}(P_1, T)) = prob(\mathcal{SC}_{\leq \theta}(P_2, T)) \quad \blacksquare$$

Obviously, \sim_{MT} is strictly finer than classical testing equivalence [21] and probabilistic testing equivalence [17, 19]. On the other hand, it is strictly coarser than \sim_{MB} as it is less sensitive to branching points. A consequence of this fact is that the derivatives of two Markovian testing equivalent terms are not necessarily related by \sim_{MT} . We conclude with a necessary condition.

Proposition 2. *Let $P_1, P_2 \in \mathcal{P}_{\text{C,pc}}$ and $T \in \mathcal{T}$. Whenever $P_1 \sim_{\text{MT}} P_2$, then for all $c_k \in \mathcal{SC}(P_k, T)$ with $k \in \{1, 2\}$ there exists $c_h \in \mathcal{SC}(P_h, T)$ with $h \in \{1, 2\} - \{k\}$ such that:*

$$\begin{aligned} \text{trace}(c_k) &= \text{trace}(c_h) \\ \text{time}_a(c_k) &= \text{time}_a(c_h) \end{aligned}$$

and for all $a \in \text{Name}$:

$$\text{rate}(P_{k,\text{last}}, a, 0, \mathcal{P}_{\text{C}}) = \text{rate}(P_{h,\text{last}}, a, 0, \mathcal{P}_{\text{C}})$$

with $P_{k,\text{last}}$ (resp. $P_{h,\text{last}}$) being the process part of the last configuration of c_k (resp. c_h). \blacksquare

5.3 Alternative Characterizations

We now present two alternative characterizations of \sim_{MT} . The first one is based on the probability distribution of passing a test within a certain sequence of amounts of time. A consequence of this characterization is that considering the (more accurate) stepwise durations of the test-driven computations leads to the same equivalence as considering the (easier to work with) stepwise average durations of the test-driven computations. This justifies the use of expected values instead of random variables in the definition of \sim_{MT} .

Definition 37. *Let $P_1, P_2 \in \mathcal{P}_{\text{C,pc}}$. We say that P_1 is Markovian distribution-testing equivalent to P_2 , written $P_1 \sim_{\text{MT,d}} P_2$, iff for all tests $T \in \mathcal{T}$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of amounts of time:*

$$\text{prob}_d(\mathcal{SC}(P_1, T), \theta) = \text{prob}_d(\mathcal{SC}(P_2, T), \theta) \quad \blacksquare$$

Theorem 6. *Let $P_1, P_2 \in \mathcal{P}_{\text{C,pc}}$. Then:*

$$P_1 \sim_{\text{MT,d}} P_2 \iff P_1 \sim_{\text{MT}} P_2 \quad \blacksquare$$

The second alternative characterization of \sim_{MT} is based on traces that are suitably extended with the sets of the action names permitted at each step by the environment. This means that it is possible to characterize \sim_{MT} in a way that fully abstracts from the tests. A consequence of the proof of this characterization is the identification of a set of canonical tests, i.e. a set of tests that are necessary and sufficient in order to establish whether two process terms are Markovian testing equivalent. Each such test admits a single computation leading to success, whose states can have additional computations each leading to failure in one step.

Definition 38. *An element σ of $(\text{Name} \times 2^{\text{Name}})^*$ is an extended trace iff either σ is the empty sequence or:*

$$\sigma \equiv (a_1, \mathcal{E}_1) \circ (a_2, \mathcal{E}_2) \circ \dots \circ (a_n, \mathcal{E}_n)$$

for some $n \in \mathbf{N}_{>0}$ with $a_i \in \mathcal{E}_i$ for each $i = 1, \dots, n$. We denote by \mathcal{ET} the set of the extended traces. \blacksquare

Definition 39. Let $\sigma \in \mathcal{ET}$. The trace associated with σ is defined by induction on the length of σ through the following Name^* -valued function:

$$\text{trace}(\sigma) = \begin{cases} \varepsilon & \text{if } \text{length}(\sigma) = 0 \\ a \circ \text{trace}(\sigma') & \text{if } \sigma \equiv (a, \mathcal{E}) \circ \sigma' \end{cases}$$

where ε is the empty trace. ■

Definition 40. Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\sigma \in \mathcal{ET}$. We say that c is compatible with σ iff:

$$\text{trace}(c) = \text{trace}(\sigma)$$

We denote by $\mathcal{CC}(P, \sigma)$ the multiset of the computations of $\mathcal{C}_f(P)$ that are compatible with σ . ■

Note that $\mathcal{CC}(P, \sigma) \subseteq \mathcal{I}_f(P)$ because of the compatibility of the considered computations with the same extended trace σ .

Definition 41. Let $P \in \mathcal{P}_{C,pc}$, $\sigma \in \mathcal{ET}$, and $c \in \mathcal{CC}(P, \sigma)$. The probability of executing c with respect to σ is defined by induction on the length of c through the following $\mathbf{R}_{[0,1]}$ -valued function:

$$\text{prob}^\sigma(c) = \begin{cases} 1 & \text{if } \text{length}(c) = 0 \\ \sum_{b \in \mathcal{E}} \frac{\lambda}{\text{rate}(P, b, 0, \mathcal{P}_C)} \cdot \text{prob}^{\sigma'}(c') & \text{if } c \equiv P \xrightarrow{a, \lambda} c' \\ & \text{with } \sigma \equiv (a, \mathcal{E}) \circ \sigma' \end{cases}$$

We also define the probability of executing a computation of C with respect to σ as:

$$\text{prob}^\sigma(C) = \sum_{c \in C} \text{prob}^\sigma(c)$$

for all $C \subseteq \mathcal{CC}(P, \sigma)$. ■

Definition 42. Let $P \in \mathcal{P}_{C,pc}$, $\sigma \in \mathcal{ET}$, and $c \in \mathcal{CC}(P, \sigma)$. The stepwise average duration of c with respect to σ is defined by induction on the length of c through the following $(\mathbf{R}_{>0})^*$ -valued function:

$$\text{time}_a^\sigma(c) = \begin{cases} \varepsilon & \text{if } \text{length}(c) = 0 \\ \sum_{b \in \mathcal{E}} \frac{1}{\text{rate}(P, b, 0, \mathcal{P}_C)} \circ \text{time}_a^{\sigma'}(c') & \text{if } c \equiv P \xrightarrow{a, \lambda} c' \\ & \text{with } \sigma \equiv (a, \mathcal{E}) \circ \sigma' \end{cases}$$

where ε is the empty stepwise average duration. We also define the multiset of the computations of C whose stepwise average duration with respect to σ is not greater than θ as:

$$C_{\leq \theta}^\sigma = \{ \{ c \in C \mid \text{length}(c) \leq \text{length}(\theta) \wedge \forall i = 1, \dots, \text{length}(c). \text{time}_a^\sigma(c)[i] \leq \theta[i] \} \}$$

for all $C \subseteq \mathcal{CC}(P, \sigma)$ and $\theta \in (\mathbf{R}_{>0})^*$. ■

Definition 43. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian extended-trace equivalent to P_2 , written $P_1 \sim_{\text{MTTr},e} P_2$, iff for all extended traces $\sigma \in \mathcal{ET}$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\text{prob}^\sigma(\mathcal{CC}_{\leq \theta}^\sigma(P_1, \sigma)) = \text{prob}^\sigma(\mathcal{CC}_{\leq \theta}^\sigma(P_2, \sigma)) \quad \blacksquare$$

Theorem 7. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Then:

$$P_1 \sim_{\text{MTTr},e} P_2 \iff P_1 \sim_{\text{MT}} P_2 \quad \blacksquare$$

Definition 44. The set \mathcal{T}_c of the canonical tests is generated by the following syntax:

$$\boxed{\begin{array}{l} T ::= s \\ | \langle a, * \rangle . T + \sum_{b \in \mathcal{E} - \{a\}} \langle b, * \rangle . f \end{array}}$$

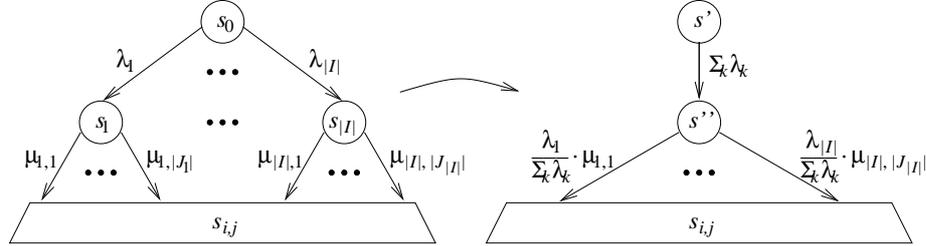
where the summation is absent whenever $\mathcal{E} - \{a\} = \emptyset$. ■

Corollary 2. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Then $P_1 \sim_{\text{MT}} P_2$ iff for all $T \in \mathcal{T}_c$ and $\theta \in (\mathbf{R}_{>0})^*$:

$$\text{prob}(\mathcal{SC}_{\leq \theta}(P_1, T)) = \text{prob}(\mathcal{SC}_{\leq \theta}(P_2, T)) \quad \blacksquare$$

5.4 Exactness

Markovian testing equivalence induces a CTMC-level aggregation that is strictly coarser than ordinary lumping and was not known in the CTMC field. This aggregation can be depicted through the following rewriting rule:



where:

- I is a finite index set with $|I| \geq 2$.
- k ranges over I .
- J_i is a non-empty finite index set for all $i \in I$.
- For all $i_1, i_2 \in I$:

$$\boxed{\sum_{j \in J_{i_1}} \mu_{i_1, j} = \sum_{j \in J_{i_2}} \mu_{i_2, j}}$$

with each summation being zero whenever its index set is empty.

Theorem 8. The CTMC-level aggregation induced by \sim_{MT} is exact. ■

5.5 Congruence

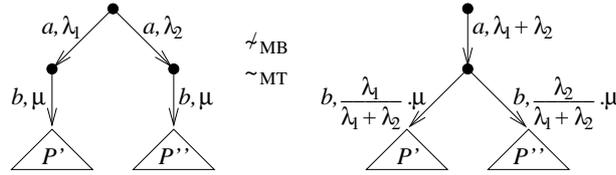
Markovian testing equivalence is a congruence with respect to all the operators of CMPC. In particular, we have what follows.

Theorem 9. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Whenever $P_1 \sim_{MT} P_2$, then:*

1. $\langle a, \lambda \rangle . P_1 \sim_{MT} \langle a, \lambda \rangle . P_2$ for all $\langle a, \lambda \rangle \in Act_S$.
2. $P_1 + P \sim_{MT} P_2 + P$ and $P + P_1 \sim_{MT} P + P_2$ for all $P \in \mathcal{P}_{C,pc}$.
3. $P_1 \parallel_S P \sim_{MT} P_2 \parallel_S P$ and $P \parallel_S P_1 \sim_{MT} P \parallel_S P_2$ for all $S \subseteq Name$ and $P \in \mathcal{P}_C$ containing only passive actions such that $P_1 \parallel_S P, P_2 \parallel_S P \in \mathcal{P}_{C,pc}$. ■

5.6 Axiomatization

Markovian testing equivalence is strictly coarser than Markovian bisimilarity, hence the axioms of Table 1 are still valid for \sim_{MT} over $\mathcal{P}_{C,pc}$, but not complete. In fact, the two process terms depicted below (with $P' \not\sim_{MB} P''$):



show that \sim_{MB} is highly sensitive to branching points. By contrast, \sim_{MT} allows choices to be deferred as long as they are related to branches starting with actions having the same name that are immediately followed by actions having the same names and the same total rates in all the branches.

The two terms above constitute the simplest instance of an axiom schema subsuming \mathcal{A}_4^{MB} that characterizes \sim_{MT} . The axiom schema is the following:

$$\boxed{\sum_{i \in I} \langle a, \lambda_i \rangle . \sum_{j \in J_i} \langle b_{i,j}, \mu_{i,j} \rangle . P_{i,j} = \langle a, \sum_{k \in I} \lambda_k \rangle . \sum_{i \in I} \sum_{j \in J_i} \langle b_{i,j}, \frac{\lambda_i}{\sum_{k \in I} \lambda_k} \cdot \mu_{i,j} \rangle . P_{i,j}}$$

where:

- I is a finite index set with $|I| \geq 2$.
- J_i is a finite index set for all $i \in I$, with the related summation being \emptyset whenever $J_i = \emptyset$.
- For all $i_1, i_2 \in I$ and $b \in Name$:

$$\boxed{\sum_{j \in J_{i_1}} \{ \mu_{i_1,j} \mid b_{i_1,j} = b \} = \sum_{j \in J_{i_2}} \{ \mu_{i_2,j} \mid b_{i_2,j} = b \}}$$

with each summation being zero whenever its index set is empty.

5.7 Logical Characterization

Markovian testing equivalence has a modal logic characterization over $\mathcal{P}_{C,pc}$ based on a Markovian variant of a restriction of the Hennessy-Milner logic, in which both negation and logical conjunction are ruled out, while the diamond operator is made dependent from the environment.

Unlike HML_{MB} , where the syntax is decorated with rate lower bounds and the satisfaction relation is qualitative, here there is no extra information in the syntax and a quantitative interpretation inspired by [39] is adopted. This establishes the probability with which a process term satisfies a formula quickly enough on average, i.e. within a given sequence of average amounts of time.

Definition 45. *The set of the formulas of HML_{MT} is generated by the following syntax:*

$$\boxed{\phi ::= \text{true} \mid \langle a | \mathcal{E} \rangle \phi}$$

where $a \in \text{Name}$ and $\mathcal{E} \subseteq \text{Name}$ such that $a \in \mathcal{E}$. ■

Definition 46. *The interpretation function $\llbracket \cdot \rrbracket_{MT}$ of HML_{MT} over $\mathcal{P}_{C,pc} \times (\mathbf{R}_{>0})^*$ is defined by structural induction as follows:*

$$\boxed{\begin{aligned} \llbracket \text{true} \rrbracket_{MT}(P, \theta) &= 1 \\ \llbracket \langle a | \mathcal{E} \rangle \phi \rrbracket_{MT}(P, \theta) &= \begin{cases} 0 & \text{if } \theta = \varepsilon \vee \frac{1}{\sum_{b \in \mathcal{E}} \text{rate}(P, b, 0, \mathcal{P}_C)} > \theta[1] \\ \sum_{P \xrightarrow{a, \lambda} P'} \frac{\lambda}{\sum_{b \in \mathcal{E}} \text{rate}(P, b, 0, \mathcal{P}_C)} \cdot \llbracket \phi \rrbracket_{MT}(P', \theta') & \text{if } \theta = t \circ \theta' \wedge \frac{1}{\sum_{b \in \mathcal{E}} \text{rate}(P, b, 0, \mathcal{P}_C)} \leq t \end{cases} \end{aligned}}$$

where the summation is taken to be zero whenever there are no a -transitions departing from P . ■

Theorem 10. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Then:*

$$P_1 \sim_{MT} P_2 \iff \forall \phi \in \text{HML}_{MT}. \forall \theta \in (\mathbf{R}_{>0})^*. \llbracket \phi \rrbracket_{MT}(P_1, \theta) = \llbracket \phi \rrbracket_{MT}(P_2, \theta) \quad \blacksquare$$

5.8 Verification Complexity

Markovian testing equivalence can be decided in polynomial time because two action-labeled CTMCs are Markovian testing equivalent iff their corresponding embedded action-labeled DTMCs (with suitably enriched labels) are probabilistic testing equivalent, with the latter coinciding with probabilistic ready equivalence and hence being decidable in polynomial time [35].

The algorithm to check whether $P_1 \sim_{MT} P_2$ thus proceeds as follows:

1. Transform $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ into their corresponding embedded discrete-time versions:
 - (a) Divide the rate of each transition by the total exit rate of its source state.

- (b) Augment the name of each transition with the total exit rate of its source state.
- 2. Compute the equivalence \mathcal{R} that relates any two states of the disjoint union of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ such that their two sets of (original) action names labeling their outgoing transitions coincide.
- 3. For each equivalence class R induced by \mathcal{R} , apply the algorithm of [47] to check the embedded discrete-time versions of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ for probabilistic language equivalence by considering R as the set of accepting states.

The time complexity is $O(n^5)$, where n is the number of states of the disjoint union of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$.

6 Markovian Trace Equivalence

Markovian trace equivalence considers two process terms to be equivalent whenever they are able to perform computations with the same functional and performance characteristics. In this section we provide the definition of Markovian trace equivalence over $\mathcal{P}_{C,pc}$ and we recall its properties from [49, 7, 9].

6.1 Equivalence Definition

Markovian trace equivalence relies on comparing the process term probabilities of performing a computation within a given sequence of average amounts of time. We emphasize that here, given a process term $P \in \mathcal{P}_{C,pc}$, we no longer have tests that interact with P . Instead, we directly consider the finite-length computations of P , to which the inductive definitions of Sect. 2.3 apply.

Definition 47. *Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\alpha \in \text{Name}^*$. We say that c is compatible with α iff:*

$$\text{trace}(c) = \alpha$$

We denote by $\mathcal{CC}(P, \alpha)$ the multiset of the finite-length computations of P that are compatible with α . ■

Note that $\mathcal{CC}(P, \alpha) \subseteq \mathcal{I}_f(P)$, because of the compatibility of the computations with the same trace α , and that $\mathcal{CC}(P, \alpha)$ is finite, because of the finitely-branching structure of the considered terms.

Definition 48. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian trace equivalent to P_2 , written $P_1 \sim_{\text{MT}_r} P_2$, iff for all traces $\alpha \in \text{Name}^*$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:*

$$\text{prob}(\mathcal{CC}_{\leq \theta}(P_1, \alpha)) = \text{prob}(\mathcal{CC}_{\leq \theta}(P_2, \alpha)) \quad \blacksquare$$

Obviously, \sim_{MT_r} is strictly finer than classical trace equivalence [33] and probabilistic trace equivalence [36]. On the other hand, it is strictly coarser than \sim_{MT} as it completely overrides branching points. Thus, similarly to \sim_{MT} , the derivatives of two Markovian trace equivalent terms are not necessarily related by \sim_{MT_r} . We conclude with a necessary condition.

Proposition 3. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$ and $\alpha \in \text{Name}^*$. Whenever $P_1 \sim_{\text{MTTr}} P_2$, then for all $c_k \in \mathcal{CC}(P_k, \alpha)$ with $k \in \{1, 2\}$ there exists $c_h \in \mathcal{CC}(P_h, \alpha)$ with $h \in \{1, 2\} - \{k\}$ such that:*

$$\begin{aligned} \text{trace}(c_k) &= \text{trace}(c_h) \\ \text{time}_a(c_k) &= \text{time}_a(c_h) \end{aligned}$$

and:

$$\text{rate}_t(P_{k,\text{last}}, 0) = \text{rate}_t(P_{h,\text{last}}, 0)$$

with $P_{k,\text{last}}$ (resp. $P_{h,\text{last}}$) being the last configuration of c_k (resp. c_h). ■

6.2 Alternative Characterizations

Similarly to \sim_{MT} , it turns out that \sim_{MTTr} has an alternative characterization based on the probability distribution of executing a trace within a certain sequence of amounts of time. A consequence of this characterization is that considering the (more accurate) stepwise durations of the computations leads to the same equivalence as considering the (easier to work with) stepwise average durations of the computations. This justifies the use of expected values instead of random variables in the definition of \sim_{MTTr} .

Definition 49. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian distribution-trace equivalent to P_2 , written $P_1 \sim_{\text{MTTr,d}} P_2$, iff for all traces $\alpha \in \text{Name}^*$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of amounts of time:*

$$\text{prob}_d(\mathcal{CC}(P_1, \alpha), \theta) = \text{prob}_d(\mathcal{CC}(P_2, \alpha), \theta) \quad \blacksquare$$

Theorem 11. *Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. Then:*

$$P_1 \sim_{\text{MTTr,d}} P_2 \iff P_1 \sim_{\text{MTTr}} P_2 \quad \blacksquare$$

6.3 Other Markovian Trace-Based Equivalences

Like for classical trace equivalence, it is possible to define some variants of \sim_{MTTr} , which are based on the notions of completed trace, failure set, ready set, failure trace, and ready trace, respectively.

These variants were originally conceived to overcome some drawbacks of classical trace equivalence. Completed traces are traces ending up in deadlock states, so considering them apart is useful to make classical trace equivalence sensitive to deadlock. A failure set is a set of names of actions that cannot be executed in a certain state, and a failure trace is a trace extended at each step with a failure set. Considering failures makes classical trace equivalence sensitive to safety properties. Likewise, a ready set is the set of the names of all the actions that must be executable in a certain state, and a ready trace is a trace extended at each step with a ready set. Considering readies makes classical trace equivalence sensitive to liveness properties.

Definition 50. *Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\alpha \in \text{Name}^*$. We say that c is a maximal computation compatible with α iff $c \in \mathcal{CC}(P, \alpha)$ and the last configuration of c is deadlocked. We denote by $\mathcal{MCC}(P, \alpha)$ the multiset of the finite-length maximal computations of P that are compatible with α . ■*

Definition 51. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian completed-trace equivalent to P_2 , written $P_1 \sim_{\text{MTr},c} P_2$, iff for all traces $\alpha \in \text{Name}^*$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\begin{aligned} \text{prob}(\mathcal{CC}_{\leq\theta}(P_1, \alpha)) &= \text{prob}(\mathcal{CC}_{\leq\theta}(P_2, \alpha)) \\ \text{prob}(\mathcal{MCC}_{\leq\theta}(P_1, \alpha)) &= \text{prob}(\mathcal{MCC}_{\leq\theta}(P_2, \alpha)) \end{aligned} \quad \blacksquare$$

Definition 52. Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\varphi \equiv (\alpha, \mathcal{F}) \in \text{Name}^* \times 2^{\text{Name}}$. We say that c is a failure computation compatible with φ iff $c \in \mathcal{CC}(P, \alpha)$ and the last configuration of c cannot execute any action whose name belongs to the failure set \mathcal{F} . We denote by $\mathcal{FCC}(P, \varphi)$ the multiset of the finite-length failure computations of P that are compatible with φ . \blacksquare

Definition 53. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian failure equivalent to P_2 , written $P_1 \sim_{\text{MF}} P_2$, iff for all traces with final failure set $\varphi \in \text{Name}^* \times 2^{\text{Name}}$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\text{prob}(\mathcal{FCC}_{\leq\theta}(P_1, \varphi)) = \text{prob}(\mathcal{FCC}_{\leq\theta}(P_2, \varphi)) \quad \blacksquare$$

Definition 54. Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\rho \equiv (\alpha, \mathcal{R}) \in \text{Name}^* \times 2^{\text{Name}}$. We say that c is a ready computation compatible with ρ iff $c \in \mathcal{CC}(P, \alpha)$ and the set of the names of all the actions executable by the last configuration of c coincides with the ready set \mathcal{R} . We denote by $\mathcal{RCC}(P, \rho)$ the multiset of the finite-length ready computations of P that are compatible with ρ . \blacksquare

Definition 55. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian ready equivalent to P_2 , written $P_1 \sim_{\text{MR}} P_2$, iff for all traces with final ready set $\rho \in \text{Name}^* \times 2^{\text{Name}}$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\text{prob}(\mathcal{RCC}_{\leq\theta}(P_1, \rho)) = \text{prob}(\mathcal{RCC}_{\leq\theta}(P_2, \rho)) \quad \blacksquare$$

Definition 56. Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\phi \in (\text{Name} \times 2^{\text{Name}})^*$. We say that c is a failure-trace computation compatible with ϕ iff c is compatible with the trace component of ϕ and each configuration of c cannot execute any action whose name belongs to the corresponding failure set in the failure component of ϕ . We denote by $\mathcal{FTCC}(P, \phi)$ the multiset of the finite-length failure-trace computations of P that are compatible with ϕ . \blacksquare

Definition 57. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian failure-trace equivalent to P_2 , written $P_1 \sim_{\text{MFTTr}} P_2$, iff for all failure traces $\phi \in (\text{Name} \times 2^{\text{Name}})^*$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\text{prob}(\mathcal{FTCC}_{\leq\theta}(P_1, \phi)) = \text{prob}(\mathcal{FTCC}_{\leq\theta}(P_2, \phi)) \quad \blacksquare$$

Definition 58. Let $P \in \mathcal{P}_{C,pc}$, $c \in \mathcal{C}_f(P)$, and $\varrho \in (\text{Name} \times 2^{\text{Name}})^*$. We say that c is a ready-trace computation compatible with ϱ iff c is compatible with the trace component of ϱ and the sets of the names of all the actions executable by the configurations of c coincide with the corresponding ready sets in the ready component of ϱ . We denote by $\mathcal{RTCC}(P, \varrho)$ the multiset of the finite-length ready-trace computations of P that are compatible with ϱ . \blacksquare

Definition 59. Let $P_1, P_2 \in \mathcal{P}_{C,pc}$. We say that P_1 is Markovian ready-trace equivalent to P_2 , written $P_1 \sim_{\text{MTr}} P_2$, iff for all ready traces $\varrho \in (\text{Name} \times 2^{\text{Name}})^*$ and sequences $\theta \in (\mathbf{R}_{>0})^*$ of average amounts of time:

$$\text{prob}(\mathcal{RTCC}_{\leq \theta}(P_1, \varrho)) = \text{prob}(\mathcal{RTCC}_{\leq \theta}(P_2, \varrho)) \quad \blacksquare$$

Similarly to the probabilistic case [36, 35], in the Markovian framework trace equivalence becomes deadlock sensitive, hence \sim_{MTr} coincides with Markovian completed-trace equivalence. Moreover, Markovian failure equivalence coincides with Markovian ready equivalence – with both coinciding with \sim_{MT} – and Markovian failure-trace equivalence coincides with Markovian ready-trace equivalence. As a consequence, the Markovian linear-time/branching-time spectrum turns out to be more condensed than the nondeterministic one [25].

Theorem 12. The Markovian linear-time/branching-time spectrum is:

$$\sim_{\text{MB}} \subset \sim_{\text{MTr}} = \sim_{\text{MFTr}} \subset \sim_{\text{MR}} = \sim_{\text{MT}} = \sim_{\text{MTr,e}} = \sim_{\text{MF}} \subset \sim_{\text{MTr,c}} = \sim_{\text{MTr}} \quad \blacksquare$$

6.4 Exactness

From the point of view of the induced CTMC-level aggregation, nothing changes when moving from \sim_{MT} to \sim_{MTr} .

Theorem 13. \sim_{MTr} induces the same CTMC-level aggregation as \sim_{MT} . ■

Corollary 3. The CTMC-level aggregation induced by \sim_{MTr} is exact. ■

6.5 Congruence

Markovian trace equivalence is a congruence with respect to all the operators of SMPC. In particular, we have what follows.

Theorem 14. Let $P_1, P_2 \in \mathcal{P}_{\text{S}}$. Whenever $P_1 \sim_{\text{MTr}} P_2$, then:

1. $\langle a, \lambda \rangle.P_1 \sim_{\text{MTr}} \langle a, \lambda \rangle.P_2$ for all $\langle a, \lambda \rangle \in \text{Act}_{\text{S}}$.
2. $P_1 + P \sim_{\text{MTr}} P_2 + P$ and $P + P_1 \sim_{\text{MTr}} P + P_2$ for all $P \in \mathcal{P}_{\text{S}}$. ■

Unfortunately, similarly to the probabilistic case [36], \sim_{MTr} is not a congruence with respect to parallel composition. Consider for instance the following two Markovian trace equivalent process terms:

$$\begin{aligned} P_1 &\equiv \langle a, \lambda_1 \rangle. \langle b, \mu \rangle.P' + \langle a, \lambda_2 \rangle. \langle c, \mu \rangle.P'' \\ P_2 &\equiv \langle a, \lambda_1 + \lambda_2 \rangle. \left(\langle b, \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot \mu \rangle.P' + \langle c, \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot \mu \rangle.P'' \right) \end{aligned}$$

where $b \neq c$. If we place each of them in the following context:

$$- \parallel_{\{a,b,c\}} \langle a, * \rangle. \langle b, * \rangle. \underline{0}$$

we obtain two performance-closed process terms – which we call Q_1 and Q_2 – that are no longer Markovian trace equivalent.

In fact, the following trace:

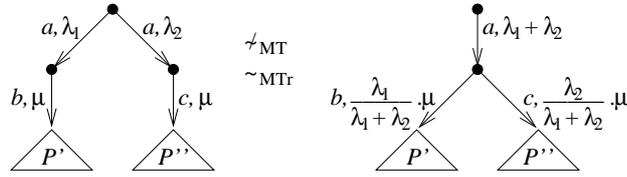
$$\alpha \equiv a \circ b$$

can distinguish between Q_1 and Q_2 . The reason is that the only computation of

Q_1 compatible with α is formed by a transition labeled with $\langle a, \lambda_1 \rangle$ followed by a transition labeled with $\langle b, \mu \rangle$, which has execution probability $\frac{\lambda_1}{\lambda_1 + \lambda_2}$ and stepwise average duration $\frac{1}{\lambda_1 + \lambda_2} \circ \frac{1}{\mu}$. By contrast, the only computation of Q_2 compatible with α is formed by a transition labeled with $\langle a, \lambda_1 + \lambda_2 \rangle$ followed by a transition labeled with $\langle b, \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot \mu \rangle$, which has execution probability 1 and stepwise average duration $\frac{1}{\lambda_1 + \lambda_2} \circ \frac{\lambda_1 + \lambda_2}{\lambda_1 \cdot \mu}$.

6.6 Axiomatization

Markovian trace equivalence is strictly coarser than Markovian testing equivalence, hence the axioms of \sim_{MT} are still valid for \sim_{MTTr} over \mathcal{P}_S , but not complete. In fact, the two process terms depicted below (with $b \neq c$):



show that, when moving from \sim_{MT} to \sim_{MTTr} , the action prefix operator tends to become left-distributive with respect to the alternative composition operator. More precisely, choices can be deferred as long as they are related to branches starting with actions having the same name that are followed by terms having the same total exit rate. Note that the names and the total rates of the initial actions of such derivative terms can be different in the various branches.

The two terms above constitute the simplest instance of an axiom schema that characterizes \sim_{MTTr} , which is more liberal than the one characterizing \sim_{MT} . The axiom schema is the following:

$$\boxed{\sum_{i \in I} \langle a, \lambda_i \rangle \cdot \sum_{j \in J_i} \langle b_{i,j}, \mu_{i,j} \rangle \cdot P_{i,j} = \langle a, \sum_{k \in I} \lambda_k \rangle \cdot \sum_{i \in I} \sum_{j \in J_i} \langle b_{i,j}, \frac{\lambda_i}{\sum_{k \in I} \lambda_k} \cdot \mu_{i,j} \rangle \cdot P_{i,j}}$$

where:

- I is a finite index set with $|I| \geq 2$.
- J_i is a finite index set for all $i \in I$, with the related summation being 0 whenever $J_i = \emptyset$.
- For all $i_1, i_2 \in I$:

$$\boxed{\sum_{j \in J_{i_1}} \mu_{i_1,j} = \sum_{j \in J_{i_2}} \mu_{i_2,j}}$$

with each summation being zero whenever its index set is empty.

6.7 Logical Characterization

Markovian trace equivalence has a modal logic characterization over $\mathcal{P}_{C,pc}$ similar to the one for Markovian testing equivalence, in which the diamond operator is no longer dependent from the environment.

Definition 60. The set of the formulas of HML_{MTr} is generated by the following syntax:

$$\boxed{\phi ::= \text{true} \mid \langle a \rangle \phi}$$

where $a \in \text{Name}$. ■

Definition 61. The interpretation function $\llbracket \cdot \rrbracket_{\text{MTr}}$ of HML_{MTr} over $\mathcal{P}_{\text{C,pc}} \times (\mathbf{R}_{>0})^*$ is defined by structural induction as follows:

$$\boxed{\begin{aligned} \llbracket \text{true} \rrbracket_{\text{MTr}}(P, \theta) &= 1 \\ \llbracket \langle a \rangle \phi \rrbracket_{\text{MTr}}(P, \theta) &= \begin{cases} 0 & \text{if } \theta = \varepsilon \vee \frac{1}{\text{rate}_t(P,0)} > \theta[1] \\ \sum_{P \xrightarrow{a,\lambda} P'} \frac{\lambda}{\text{rate}_t(P,0)} \cdot \llbracket \phi \rrbracket_{\text{MTr}}(P', \theta') & \text{if } \theta = t \circ \theta' \wedge \frac{1}{\text{rate}_t(P,0)} \leq t \end{cases} \end{aligned}}$$

where the summation is taken to be zero whenever there are no a -transitions departing from P . ■

Theorem 15. Let $P_1, P_2 \in \mathcal{P}_{\text{C,pc}}$. Then:

$$P_1 \sim_{\text{MTr}} P_2 \iff \forall \phi \in \text{HML}_{\text{MTr}}. \forall \theta \in (\mathbf{R}_{>0})^*. \llbracket \phi \rrbracket_{\text{MTr}}(P_1, \theta) = \llbracket \phi \rrbracket_{\text{MTr}}(P_2, \theta) \blacksquare$$

6.8 Verification Complexity

Markovian trace equivalence can be decided in polynomial time because two action-labeled CTMCs are Markovian trace equivalent iff their corresponding embedded action-labeled DTMCs (with suitably enriched labels) are probabilistic trace equivalent, with the latter being decidable in polynomial time [35].

Similarly to the verification of \sim_{MT} , the algorithm to check whether $P_1 \sim_{\text{MTr}} P_2$ thus proceeds as follows:

1. Transform $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ into their corresponding embedded discrete-time versions:
 - (a) Divide the rate of each transition by the total exit rate of its source state.
 - (b) Augment the name of each transition with the total exit rate of its source state.
2. Apply the algorithm of [47] to check the embedded discrete-time versions of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$ for probabilistic language equivalence by considering each of their states as being an accepting state.

The time complexity is $O(n^4)$, where n is the number of states of the disjoint union of $\llbracket P_1 \rrbracket$ and $\llbracket P_2 \rrbracket$.

7 Conclusion

In this survey we have recalled the definitions and the properties of Markovian behavioral equivalences. Besides providing information about the Markovian linear-time/branching-time spectrum, we have compared Markovian bisimilarity, Markovian testing equivalence, and Markovian trace equivalence with respect to a number of criteria, as summarized below:

	<i>exact aggregation</i>	<i>congruence property</i>	<i>sound & complete axiomatization</i>	<i>logical characteriz.</i>	<i>verification complexity</i>
\sim_{MB}	<i>OK</i>	<i>OK</i>	<i>OK</i>	<i>OK</i>	$O(m \cdot \log n)$
\sim_{MT}	<i>OK</i>	<i>OK</i>	<i>OK</i>	<i>OK</i>	$O(n^5)$
\sim_{MT_T}	<i>OK</i>	OK_{SMPC}	OK_{SMPC}	<i>OK</i>	$O(n^4)$

As can be noted, \sim_{MB} is satisfactory with respect to all the criteria, although it is often too discriminating. A good alternative may be \sim_{MT} , as it encodes the viewpoint of an external observer and is more flexible with respect to branching points. By contrast, \sim_{MT_T} cannot be considered as a valid alternative, as it fails to be a congruence with respect to parallel composition.

It is also worth emphasizing the exactness of the CTMC-level aggregations induced by each of the three considered Markovian behavioral equivalences. This means that \sim_{MB} , \sim_{MT} , and \sim_{MT_T} are all meaningful for performance evaluation purposes. Besides justifying the investigation of the other properties, this can be exploited in practice. For instance, such Markovian behavioral equivalences can be used to aggregate the state space of a model by taking advantage of symmetries within the model [24], or to reduce the state space of a model before applying analysis techniques like model checking [45], without altering the performance properties to be assessed.

We conclude by mentioning some open problems in the field of Markovian behavioral equivalences:

- In our framework based on an asymmetric action synchronization discipline, while \sim_{MB} is defined over non-performance-closed terms too, this is not the case for \sim_{MT} and \sim_{MT_T} . Finding a way to extend their definitions so that it is still possible to determine the execution probability and the stepwise average duration of the computations in the presence of passive transitions is highly desirable.
- The exactness of the CTMC-level aggregations induced by \sim_{MB} , \sim_{MT} , and \sim_{MT_T} establishes a strong connection between concurrency theory and Markov chain theory, which in particular gives rise to a process algebraic characterization of the aggregations themselves. The question here is whether the aggregation induced by \sim_{MT} and \sim_{MT_T} is the coarsest exact non-trivial one that can be obtained, or whether it can be further extended.
- The set of logical operators necessary to characterize Markovian behavioral equivalences decreases as the discriminating power of the equivalences

- decreases. However, the logical characterization of \sim_{MT} relies on a non-standard variant of the diamond operator. What if we replace the non-standard operator with the standard one and we reintroduce logical conjunction? We claim that this may result in an alternative logical characterization of \sim_{MT} if the diamonds occurring in a conjunction are independent of each other, i.e. if the names of the related actions are all different [39].
- The verification algorithm for \sim_{MB} can also be used as a minimization algorithm (with respect to \sim_{MB}), whereas this is not the case for the verification algorithms for \sim_{MT} and \sim_{MT_T} . As far as testing is concerned, it should be investigated whether the algorithm for verifying classical testing equivalence [20] can be adapted to the Markovian framework. The reason is that this algorithm reduces the verification of classical testing equivalence over standard state-transition models to the verification of (a generalization of) classical bisimilarity over transformed models inspired by acceptance trees [28], hence it can be exploited for minimization purposes as well.
 - We would like to assess whether \sim_{MT} can be used also for quantitative analysis. So far, it supports a merely qualitative analysis, in the sense that it only allows one to establish whether two models pass an arbitrary test in the same way. What we envision is the possibility of identifying classes of tests that are related to specific performability measures, which may thus be used to evaluate models with respect to certain indices of interest.
 - Finally, it would be interesting to develop weaker versions of \sim_{MB} , \sim_{MT} , and \sim_{MT_T} . In this survey we have seen \approx_{IMB} and \approx_{EMB} , which abstract from invisible immediate actions. However, one could also consider the possibility of abstracting from invisible actions that are exponentially timed, which amounts to understand to what extent exponential delays can be neglected. This issue has been tackled in [32, 11, 13, 4], but none of the proposals seems to induce an exact aggregation at the CTMC level.

References

1. M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis, “*Modelling with Generalized Stochastic Petri Nets*”, John Wiley & Sons, 1995.
2. J.C.M. Baeten and W.P. Weijland, “*Process Algebra*”, Cambridge University Press, 1990.
3. C. Baier, J.-P. Katoen, and H. Hermanns, “*Approximate Symbolic Model Checking of Continuous Time Markov Chains*”, in *Proc. of the 10th Int. Conf. on Concurrency Theory (CONCUR 1999)*, LNCS 1664:146-162, Eindhoven (The Netherlands), 1999.
4. C. Baier, J.-P. Katoen, H. Hermanns, and V. Wolf, “*Comparative Branching-Time Semantics for Markov Chains*”, in *Information and Computation* 200:149-214, 2005.
5. S. Balsamo, M. Bernardo, and M. Simeoni, “*Performance Evaluation at the Software Architecture Level*”, in *Formal Methods for Software Architectures*, LNCS 2804:207-258, 2003.
6. J.A. Bergstra, A. Ponse, and S.A. Smolka (eds.), “*Handbook of Process Algebra*”, Elsevier, 2001.

7. M. Bernardo, “*Non-Bisimulation-Based Markovian Behavioral Equivalences*”, to appear in *Journal of Logic and Algebraic Programming*.
8. M. Bernardo and A. Aldini, “*Weak Markovian Bisimilarity: Abstracting from Prioritized/Weighted Internal Immediate Actions*”, submitted for publication.
9. M. Bernardo and S. Botta, “*A Survey of Modal Logics Characterizing Behavioral Equivalences for Nondeterministic and Stochastic Systems*”, to appear in *Mathematical Structures in Computer Science*.
10. M. Bernardo and M. Bravetti, “*Performance Measure Sensitive Congruences for Markovian Process Algebras*”, in *Theoretical Computer Science* 290:117-160, 2003.
11. M. Bernardo and R. Cleaveland, “*A Theory of Testing for Markovian Processes*”, in *Proc. of the 11th Int. Conf. on Concurrency Theory (CONCUR 2000)*, LNCS 1877:305-319, State College (PA), 2000.
12. H. Bohnenkamp, P.R. D’Argenio, H. Hermanns, and J.-P. Katoen, “*MODEST: A Compositional Modeling Formalism for Hard and Softly Timed Systems*”, in *IEEE Trans. on Software Engineering* 32:812-830, 2006.
13. M. Bravetti, “*Revisiting Interactive Markov Chains*”, in *Proc. of the 3rd Int. Workshop on Models for Time-Critical Systems (MTCS 2002)*, ENTCS 68(5):1-20, Brno (Czech Republic), 2002.
14. M. Bravetti, M. Bernardo, and R. Gorrieri, “*A Note on the Congruence Proof for Recursion in Markovian Bisimulation Equivalence*”, in *Proc. of the 6th Int. Workshop on Process Algebra and Performance Modelling (PAPM 1998)*, pp. 153-164, Nice (France), 1998.
15. P. Buchholz, “*Exact and Ordinary Lumpability in Finite Markov Chains*”, in *Journal of Applied Probability* 31:59-75, 1994.
16. P. Buchholz, “*Markovian Process Algebra: Composition and Equivalence*”, in *Proc. of the 2nd Int. Workshop on Process Algebra and Performance Modelling (PAPM 1994)*, Technical Report 27-4, pp. 11-30, Erlangen (Germany), 1994.
17. I. Christoff, “*Testing Equivalences and Fully Abstract Models for Probabilistic Processes*”, in *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, LNCS 458:126-140, Amsterdam (The Netherlands), 1990.
18. G. Clark, S. Gilmore, and J. Hillston, “*Specifying Performance Measures for PEPA*”, in *Proc. of the 5th AMAST Int. Workshop on Formal Methods for Real Time and Probabilistic Systems (ARTS 1999)*, LNCS 1601:211-227, Bamberg (Germany), 1999.
19. R. Cleaveland, Z. Dayar, S.A. Smolka, and S. Yuen, “*Testing Preorders for Probabilistic Processes*”, in *Information and Computation* 154:93-148, 1999.
20. R. Cleaveland and M. Hennessy, “*Testing Equivalence as a Bisimulation Equivalence*”, in *Formal Aspects of Computing* 5:1-20, 1993.
21. R. De Nicola and M. Hennessy, “*Testing Equivalences for Processes*”, in *Theoretical Computer Science* 34:83-133, 1983.
22. R. De Nicola, D. Latella, and M. Massink, “*Formal Modeling and Quantitative Analysis of KLAIM-Based Mobile Systems*”, in *Proc. of the 20th ACM Symp. on Applied Computing (SAC 2005)*, ACM Press, pp. 428-435, Santa Fe (NM), 2005.
23. S. Derisavi, H. Hermanns, and W.H. Sanders, “*Optimal State-Space Lumping in Markov Chains*”, in *Information Processing Letters* 87:309-315, 2003.
24. S. Gilmore, J. Hillston, and M. Ribaud, “*An Efficient Algorithm for Aggregating PEPA Models*”, in *IEEE Trans. on Software Engineering* 27:449-464, 2001.
25. R.J. van Glabbeek, “*The Linear Time - Branching Time Spectrum I*”, in [6], pp. 3-99, 2001.

26. R.J. van Glabbeek, S.A. Smolka, and B. Steffen, “*Reactive, Generative and Stratified Models of Probabilistic Processes*”, in *Information and Computation* 121:59-80, 1995.
27. N. Götz, U. Herzog, and M. Rettelsbach, “*Multiprocessor and Distributed System Design: The Integration of Functional Specification and Performance Analysis Using Stochastic Process Algebras*”, in *Proc. of the 16th Int. Symp. on Computer Performance Modelling, Measurement and Evaluation (PERFORMANCE 1993)*, LNCS 729:121-146, Roma (Italy), 1993.
28. M. Hennessy, “*Acceptance Trees*”, in *Journal of the ACM* 32:896-928, 1985.
29. M. Hennessy and R. Milner, “*Algebraic Laws for Nondeterminism and Concurrency*”, in *Journal of the ACM* 32:137-162, 1985.
30. H. Hermanns, “*Interactive Markov Chains*”, LNCS 2428, 2002.
31. H. Hermanns and M. Rettelsbach, “*Syntax, Semantics, Equivalences, and Axioms for MTIPP*”, in *Proc. of the 2nd Int. Workshop on Process Algebra and Performance Modelling (PAPM 1994)*, Technical Report 27-4, pp. 71-87, Erlangen (Germany), 1994.
32. J. Hillston, “*A Compositional Approach to Performance Modelling*”, Cambridge University Press, 1996.
33. C.A.R. Hoare, “*Communicating Sequential Processes*”, Prentice Hall, 1985.
34. R.A. Howard, “*Dynamic Probabilistic Systems*”, John Wiley & Sons, 1971.
35. D.T. Huynh and L. Tian, “*On Some Equivalence Relations for Probabilistic Processes*”, in *Fundamenta Informaticae* 17:211-234, 1992.
36. C.-C. Jou and S.A. Smolka, “*Equivalences, Congruences, and Complete Axiomatizations for Probabilistic Processes*”, in *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, LNCS 458:367-383, Amsterdam (The Netherlands), 1990.
37. P.C. Kanellakis and S.A. Smolka, “*CCS Expressions, Finite State Processes, and Three Problems of Equivalence*”, in *Information and Computation* 86:43-68, 1990.
38. L. Kleinrock, “*Queueing Systems*”, John Wiley & Sons, 1975.
39. M.Z. Kwiatkowska and G.J. Norman, “*A Testing Equivalence for Reactive Probabilistic Processes*”, in *Proc. of the 2nd Int. Workshop on Expressiveness in Concurrency (EXPRESS 1998)*, ENTCS 16(2):114-132, Nice (France), 1998.
40. K.G. Larsen and A. Skou, “*Bisimulation through Probabilistic Testing*”, in *Information and Computation* 94:1-28, 1991.
41. R. Milner, “*Communication and Concurrency*”, Prentice Hall, 1989.
42. R. Paige and R.E. Tarjan, “*Three Partition Refinement Algorithms*”, in *SIAM Journal on Computing* 16:973-989, 1987.
43. D. Park, “*Concurrency and Automata on Infinite Sequences*”, in *Proc. of the 5th GI Conf. on Theoretical Computer Science*, LNCS 104:167-183, 1981.
44. C. Priami, “*Stochastic π -Calculus*”, in *Computer Journal* 38:578-589, 1995.
45. J. Sproston and S. Donatelli, “*Backward Bisimulation in Markov Chain Model Checking*”, in *IEEE Trans. on Software Engineering* 32:531-546, 2006.
46. E.W. Stark, R. Cleaveland, and S.A. Smolka, “*A Process-Algebraic Language for Probabilistic I/O Automata*”, in *Proc. of the 14th Int. Conf. on Concurrency Theory (CONCUR 2003)*, LNCS 2761:189-203, Marseille (France), 2003.
47. W.-G. Tzeng, “*A Polynomial-Time Algorithm for the Equivalence of Probabilistic Automata*”, in *SIAM Journal on Computing* 21:216-227, 1992.
48. M. Woodside, “*From Annotated Software Designs (UML SPT/MARTE) to Model Formalisms*”, to appear in *Formal Methods for Performance Evaluation*, LNCS 4486, 2007.

49. V. Wolf, C. Baier, and M. Majster-Cederbaum, “*Trace Machines for Observing Continuous-Time Markov Chains*”, in Proc. of the *3rd Int. Workshop on Quantitative Aspects of Programming Languages (QAPL 2005)*, ENTCS 153(2):259-277, Edinburgh (UK), 2005.