

A Theory for the Semantics of Stochastic and Non-deterministic Continuous Systems[★]

Carlos E. Budde^{1,2}, Pedro R. D'Argenio^{1,2},
Pedro Sánchez Terraf^{1,2}, and Nicolás Wolovick¹

¹ FaMAF, Universidad Nacional de Córdoba, Argentina

² CONICET, Córdoba, Argentina

{cbudde, dargenio, sterraf, nicolasw}@famaf.unc.edu.ar

Abstract. The description of complex systems involving physical or biological components usually requires to model complex continuous behavior induced by variables such as time, distance, speed, temperature, alkalinity of a solution, etc. Often, such variables can be quantified probabilistically to better understand the behavior of the complex systems. For example, the arrival time of events may be considered a Poisson process or the weight of an individual may be assumed to be distributed according to a log-normal distribution. However, it is also common that the uncertainty on how these variables behave makes us prefer to leave out the choice of a particular probability and rather model it as a purely non-deterministic decision, as it is the case when a system is intended to be deployed in a variety of very different computer or network architectures. Therefore, the semantics of these systems needs to be represented by a variant of probabilistic automata that involves continuous domains on the state space and the transition relation.

In this paper, we provide a survey on the theory of such kind of models. We present the theory of the so-called labeled Markov processes (LMP) and its extension with internal non-determinism (NLMP). We show that in these complex domains, the bisimulation relation can be understood in different manners. We show the relation between the different bisimulations and try to understand their expressiveness through examples. We also study variants of Hennessy-Milner logic that provides logical characterizations of some of these bisimulations.

1 Introduction

The interplay of probabilistic and non-deterministic choices in systems that live in a continuous state space is becoming more common. For example, they arise naturally on software applications for mobile devices. This type of system has discrete state (memory hierarchy) as well as continuous state (position, orientation, acceleration, battery voltage, etc.). These continuous quantities are disrupted by the environment, and such disruption may be stochastically quantifiable. Besides, many algorithms make internal decisions sampling according to discrete probabilities. Moreover, they operate in meshes of devices where the relative speeds of execution among them are not known in

[★] Supported by ANPCyT project PICT-2012-1823, SeCyT-UNC projects 05/B284 and 05/B497 and program 05/BP02, and EU 7FP grant agreement 295261 (MEALS).

advance, therefore there is no information on how these devices interleave their operations in the time-line. Observations of discrete values like enabled or disabled buttons, and observations of continuous values like displayed roll angle in a cell phone, are part of these systems.

Examples of this kind which need to interact with physical or biological components abound and they exceed the modeling capabilities of Markov processes with continuous-state spaces or continuous time evolution (or both): they also need the consideration of non-determinism. Many formal frameworks have been defined to study them from a process theory or process algebra perspective [5, 6, 8, 9, 14–16, 20, 21, 44, 45, etc.]. A prominent and extensive work on this area is the one that builds on top of the so-called *labeled Markov processes (LMP)* [14, 20, 21, 39]. This is due to its solid and well understood mathematical foundations. An LMP allows for many transition probability functions (or Markov kernels) leaving each state (instead of only one as in usual Markov processes). Each transition probability function is a measure ranging on a (possibly continuous) measurable space, and the different transition probability functions can be singled out through labels. Thus this model *does not* consider *internal* non-determinism. From the modeling point of view, this is a significant drawback for this theory since internal non-determinism immediately arises in the analysis of systems, e.g., because of abstracting internal activity (such as weak bisimulation [37]) or because of state abstraction techniques (such as in model checking [2, 12]).

Many variants of continuous Markov processes that include internal non-determinism have been defined [5, 6, 8, 9, 15, 16, 45], including a continuous probabilistic variant of the (strong) bisimulation. Contrarily to LMPs, these models lack the sufficient structure to ensure that bisimilar models share the same observable behavior. Although [8, 9] deal with the same unstructured type of model, they lift the burden of checking measurability to the semantic tools (such as bisimulation or schedulers). In particular, this results in the definition of a bisimulation as a relation between measures rather than states.

Contrarily to [8, 9] we preferred to follow the approach of Desharnais, Panangaden, et al. and extend LMPs with internal non-determinism using the power of the mathematics provided by measure theory. This led us to develop a theory of *non-deterministic labeled Markov processes (NLMP)* [7, 10, 17, 18, 48]. An NLMP has a non-deterministic transition function T_a for each label a that, given a state, it returns a measurable *set* of probability measures, rather than only one probability measure as in LMPs.

In this paper, we provide a survey to the theory of labeled Markov processes (LMP) and its extension with internal non-determinism (NLMP). Moreover, we introduce a structured version of NLMPs (SNLMP) where action labels are also endowed with a measurable structure.

The natural notion of identity on measurable spaces is given by the σ -algebra: two points can be considered indistinguishable if they cannot be separated by the σ -algebra (i.e. there is no measurable set that contains one point but not the other). As a consequence, it is expected that bisimulation respects this principle in a setting where states are endowed with a σ -algebra. However, Danos et al. [14] showed that this is not the case in the LMP model and that there are bisimulation relations that may distinguish *more* than what the underlying σ -algebra can distinguish. That is, states that cannot

be separated by any measurable set (and hence always equated in the σ -algebra) may not be related by some bisimulation relation. This is awkward since measurable sets are the smallest distinguishable objects in a σ -algebra. Therefore, this raises the question of whether this problem extends to the bisimulation equivalence. To overcome this, [14] defines the so-called *event bisimulation* (in opposition to the previously existing *state bisimulation*—name which we will use from now on). The same situation arises on NLMPs. Moreover two candidates for state bisimulation appear if internal non-determinism is considered [17, 18, 48]. We recall here the definitions of the bisimulations on the different settings. In addition, we report the relations between the different bisimulations and try to understand their expressiveness through examples.

Behavioral equivalences like bisimulation have been characterized using logics with modalities, notably the Hennessy-Milner logic [33] (see also [29]). Similarly, there are Hennessy-Milner-like logics to completely characterize the different event bisimulation equivalences in LMPs [14], NLMPs [17, 18, 48], and SNLMPs [7] which we also present in this survey and show how they relate among themselves and the other bisimulations.

The next section summarizes some preliminaries on measure theory required to understand the paper. The rest of the text is structured according to each of the models. Section 3 reviews the theory of LMPs which deals with label-deterministic models. Section 4 reviews the theory of NLMPs which deals with internal non-determinism. Section 5 introduces SNLMPs which are a restriction of NLMPs where non-determinism in general is requested to be structured in a measurable space by endowing the set of labels with a σ -algebra. The paper concludes in Section 6 by reviewing some additional results that strongly relates to these models. We remark that the proofs of all results presented in Section 3 can be found in [14], the proofs of Section 4 can be found in [17] or, with more detail, in [48], and the proofs of Section 5 can be found in [7]. The results reported in Section 5 have only appeared as part of the theses [7, 48].

2 Preliminaries on Measure Theory

In this section, we recall some fundamental notions of measure theory that will be useful throughout the paper.

Given a set S and a collection Σ of subsets of S , we call Σ a σ -algebra iff $S \in \Sigma$ and Σ is closed under complement and denumerable union. By $\sigma(\mathcal{G})$ we denote the σ -algebra generated by the family $\mathcal{G} \subseteq 2^S$, i.e., the minimal σ -algebra containing \mathcal{G} . Each element of \mathcal{G} is called a *generator* and \mathcal{G} is called the *generator set*. We call the pair (S, Σ) a *measurable space*. A *measurable set* is a set $Q \in \Sigma$. Let (L, Λ) and (S, Σ) be measurable spaces. A *measurable rectangle* is a set $A \times B$ with $A \in \Lambda$ and $B \in \Sigma$. The *product σ -algebra* on $L \times S$ is the smallest σ -algebra containing all measurable rectangles, and is denoted by $\Lambda \otimes \Sigma$.

A function $\mu : \Sigma \rightarrow [0, 1]$ is a *probability measure* if (i) it is σ -additive, i.e. $\mu(\bigcup_{i \in \mathbb{N}} Q_i) = \sum_{i \in \mathbb{N}} \mu(Q_i)$ for all countable family of pairwise disjoint measurable sets $\{Q_i \mid i \in \mathbb{N}\} \subseteq \Sigma$, and (ii) $\mu(S) = 1$. By δ_a we denote the Dirac probability measure concentrated in $\{a\}$. Let $\mathcal{A}(S)$ denote the set of all probability measures over the measurable space (S, Σ) . We let μ, μ', μ_1, \dots range over $\mathcal{A}(S)$. Let (S_1, Σ_1) and (S_2, Σ_2) be two measurable spaces. A function $f : S_1 \rightarrow S_2$ is said to be *measurable* if for all $Q_2 \in \Sigma_2$,

$f^{-1}(Q_2) \in \Sigma_1$, i.e., its inverse image maps measurable sets to measurable sets. In this case we denote $f : (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ and say that f is Σ_1 - Σ_2 measurable.

Along the article we will often set our examples on *Borel σ -algebras*. A σ -algebra is *Borel* if it is generated by the set of all open sets in a topology. Particularly, the Borel σ -algebra on the real line is $\mathcal{B}(\mathbb{R}) = \sigma(\{(a, b) \mid a, b \in \mathbb{R} \text{ and } a < b\})$. Similarly, $\mathcal{B}([0, 1])$ is the Borel σ -algebra on the interval $[0, 1]$ generated by the open sets in the interval $[0, 1]$.

There is a standard construction by Giry [28] to endow $\mathcal{A}(S)$ with a σ -algebra¹ as follows: $\mathcal{A}(\Sigma)$ is defined as the σ -algebra generated by the sets of probability measures $\mathcal{A}^B(Q) \doteq \{\mu \mid \mu(Q) \in B\}$, with $Q \in \Sigma$ and $B \in \mathcal{B}([0, 1])$. If $p \in [0, 1]$, we will write $\mathcal{A}^{\geq p}(Q), \mathcal{A}^{> p}(Q), \mathcal{A}^{< p}(Q)$, etc. for $\mathcal{A}^B(Q)$ with $B = [p, 1], (p, 1], [0, p]$, respectively. It is known that the set $\{\mathcal{A}^{\geq p}(Q) \mid p \in (\mathbb{Q} \cap [0, 1]), Q \in \Sigma\}$ generates $\mathcal{A}(\Sigma)$. We let $\xi, \zeta, \xi', \zeta', \xi_1, \zeta_1, \dots$ range over $\mathcal{A}(\Sigma)$.

3 Labeled Markov Processes

Labeled Markov processes (LMP) were developed in [20, 21] by Desharnais et al. An LMP has a labeled set of *actions* where an action represents the interaction with the environment. Thus, an LMP is a reactive model in which there are different transition probabilities for each action. In this model, uncertainty is (only) considered to be probabilistic; therefore, the LMP model can be regarded as a generalization of deterministic processes. The thesis of Josée Desharnais [20] and the book of Prakash Panangaden [39] contain a thorough study on LMPs.

Definition 1. A labeled Markov process is a triple $(S, \Sigma, \{\tau_a \mid a \in L\})$ where Σ is a σ -algebra on the set of states S , and for each label $a \in L$, the transition probability function $\tau_a : S \rightarrow \mathcal{A}(S) \cup \{\mathbf{0}\}$ is a measurable function. Here, we let $\mathbf{0} : \Sigma \rightarrow [0, 1]$ be the null measure such that $\mathbf{0}(Q) = 0$ for all $Q \in \Sigma$.

The value $\tau_a(s)(Q)$ represents the probability of making a transition to a state in Q provided that the system is in state s and action a has been accepted. Therefore, the transition probability is actually a *conditional probability*: the probability of Q is conditioned to the fact that the system is in state s and that it reacts to action a . Originally, [20, 39] allow $\tau_a(s)$ to be a *subprobability* measure (i.e., $\tau_a(s)(S) \leq 1$) where the value $1 - \tau_a(s)(S)$ represents the probability of refusing a . As we are going to review other models we prefer to deal only with full probability measures and let $\tau_a(s) = \mathbf{0}$ indicate that action a is refused at the state s with probability 1.

Example 2. Consider a computer system that measures a movement of a particle in the real line \mathbb{R} . The particle moves according to the dynamics of a Brownian motion, but the system can only measure the position of the particle at discrete time. We also want to distinguish whether the particle has passed a particular threshold $h \in \mathbb{R}$.

¹ The application $S \mapsto \mathcal{A}(S)$ gives rise to an endofunctor \mathcal{A} of the category of measurable spaces and measurable maps. The base space of $\mathcal{A}(S, \Sigma)$ is $\mathcal{A}(S)$. By an innocuous abuse of notation, we call $\mathcal{A}(\Sigma)$ the σ -algebra of this measurable space; hence $\mathcal{A}(S, \Sigma) = (\mathcal{A}(S), \mathcal{A}(\Sigma))$.

The movement of the particle is described by a Wiener process which states that if the particle is observed at position r , the new position after a delay of t time units is a random variable with distribution $N(r, t)$, i.e., a normal distribution with mean r and variance t . More precisely, this distribution is defined by

$$\mu_r^t([l, u]) = \frac{1}{\sqrt{2\pi t}} \int_l^u e^{-\frac{(x-r)^2}{2t}} dx.$$

To construct the LMP modeling the system, we let \mathbb{R} be the set of states with the usual Borel σ -algebra $\mathcal{B}(\mathbb{R})$, where each state indicates the current position of the particle. The LMP has two types of probability transitions: one that represents the measure of the position of the particle after $n \in \mathbb{N}$ time units, and the other that indicates if the system is below or above the threshold through the labels *low* and *high* respectively. Thus, for each $n \in \mathbb{N}$ and state $r \in \mathbb{R}$, $\tau_n(r) = \mu_r^n$ represents the probability at position r to read that the particle has jumped to a new position in a given interval after n (discrete) time units have elapsed. Besides, $\tau_{low}(r) = \mathbf{if } r < h \mathbf{ then } \delta_r \mathbf{ else } \mathbf{0}$ is used to indicate whether the system is below or above the threshold; if r is below the threshold $\tau_{low}(r)$ is a self-loop with probability 1, otherwise it refuses action *low*. Similarly, $\tau_{high}(r) = \mathbf{if } r < h \mathbf{ then } \mathbf{0 else } \delta_r$ indicates whether the system is above the threshold.

The system is then modeled by the LMP $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \{\tau_a \mid a \in \mathbb{N} \cup \{low, high\}\})$. □

Probabilistic bisimulation was introduced by Larsen and Skou [35] in a discrete setting very much like the LMP, only that distributions run on discrete sets. This notion has been adapted by Desharnais et al. in [20, 21] to the continuous case of LMPs. The idea behind the bisimulation equivalence is that from two equivalent states, an a -transition should lead with equal probability to any measurable aggregate of equivalence classes (properly speaking, to any measurable set that results of an arbitrary union of equivalence classes).

Given a relation $R \subseteq S \times S$, a set $Q \subseteq S$ is *R-closed* if $R(Q) \subseteq Q$. Notice that if R is symmetric, Q is *R-closed* if and only if for all $s, t \in S$ such that $s R t$, $s \in Q \Leftrightarrow t \in Q$. Using this definition, a symmetric relation R can be lifted to an equivalence relation in $\mathcal{A}(S)$ as follows: $\mu R \mu'$ iff for every *R-closed* $Q \in \Sigma$, $\mu(Q) = \mu'(Q)$.

Using this idea Desharnais et al. defined the notion of bisimulation that was called *state bisimulation* in [14] to stress the fact that the relation is defined directly on states.

Definition 3. $R \subseteq S \times S$ is a state bisimulation on the LMP $(S, \Sigma, \{\tau_a \mid a \in L\})$ if it is symmetric and for all $s, t \in S$, $a \in L$, $s R t$ implies that $\tau_a(s) R \tau_a(t)$.

We say that two states s and t are state bisimilar (or state bisimulation equivalent), denoted by $s \sim_s t$, if there is a state bisimulation R such that $s R t$.

Relation \sim_s can be proved to be a state bisimulation and also an equivalence relation [20, 21].

The definition of state bisimulation is point-wise and not event-wise as one should expect in a measure-theoretic realm, since R has no measurability restrictions. Indeed, as shown in [14], a state bisimulation can distinguish more states than what the underlying σ -algebra can distinguish. Suppose the set of states $\{1, 2, 3, 4\}$ with the σ -algebra $\{\emptyset, \{1, 2\}, \{3, 4\}, \{1, 2, 3, 4\}\}$. No matter what the transition function is, the identity relation is a state bisimulation. However, the identity relation distinguishes states that

cannot be distinguished through measurable sets (i.e., events) on the σ -algebra. Take for instance states 1 and 2: they are not related by the identity relation but they cannot be distinguished through transition probability functions because the transition probability function has to be measurable with respect to the σ -algebra.

The question is whether this problem extends to the state bisimulation equivalence \sim_s . To understand the problem, [14] introduced a measure-theory aware notion of behavioral equivalence.

Definition 4. *An event bisimulation on an LMP $(S, \Sigma, \{\tau_a \mid a \in L\})$ is a sub- σ -algebra \mathcal{E} of Σ s.t. $(S, \mathcal{E}, \{\tau_a \mid a \in L\})$ is an LMP.*

We extend the notion of event bisimulation to relations. We say that a relation R is an event bisimulation if there is an event bisimulation \mathcal{E} such that $R = \mathcal{R}(\mathcal{E})$, where $\mathcal{R}(\mathcal{E}) \doteq \{(s, t) \in S \times S \mid \forall Q \in \mathcal{E} : s \in Q \Leftrightarrow t \in Q\}$. More generally, we say that two states $s, t \in S$ are *event bisimilar*, denoted by $s \sim_e t$, if there is an event bisimulation \mathcal{E} such that $s \mathcal{R}(\mathcal{E}) t$. The fact that \sim_e is an equivalence relation is an immediate corollary of Theorem 5.

The article [14] shows that R is a state bisimulation iff $\Sigma(R)$, defined by $\Sigma(R) = \{Q \in \Sigma \mid Q \text{ is } R\text{-closed}\}$, is an event bisimulation. This is an important result that leads to prove that the largest state bisimulation \sim_s is also an event bisimulation. That is $\sim_s \subseteq \sim_e$.

Another way to understand the semantics of a process is through a modal logic. The semantics of a process is defined by the set of properties that it satisfies. Particularly, a formula in a Hennessy-Milner-like logic defines a possible observation of the execution of the system [29, 33].

Besides, it would be useful if the semantics from the point of view of the logic agrees with that defined by the bisimulation. Thus, if two states are not bisimilar there must be an observation (a formula) that distinguishes them.

In [20, 21] a variant of the Hennessy-Milner logic for LMPs is introduced. The logic, that we call \mathcal{L}^0 , is given by the following productions:

$$\varphi \equiv \top \mid \varphi_1 \wedge \varphi_2 \mid \langle a \rangle_q \varphi$$

where $a \in L$ and $q \in \mathbb{Q} \cap [0, 1]$.

Formulas in \mathcal{L}^0 are interpreted as sets of states. Thus, a formula φ is satisfied by a state s if and only if $s \in \llbracket \varphi \rrbracket$.

$$\llbracket \top \rrbracket := S \quad \llbracket \varphi_1 \wedge \varphi_2 \rrbracket := \llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket \quad \llbracket \langle a \rangle_q \varphi \rrbracket := \{s \in S : \tau_a(s, \llbracket \varphi \rrbracket) \geq q\}$$

Notice, in particular, that $\langle a \rangle_q \varphi$ is satisfied by a state s if there is an a -labeled transition from s reaching a set of states that satisfy φ with probability at least q . Therefore, for the semantics to be well defined, $\llbracket \varphi \rrbracket$ should be measurable for any formula φ . To show this, first notice that all operations involved on the definition of the semantics preserve measurability (in particular τ_a is a measurable function). Then, by structural induction on the formula φ , it is straightforward to conclude that $\llbracket \varphi \rrbracket$ is measurable. Let $\llbracket \mathcal{L}^0 \rrbracket := \{\llbracket \varphi \rrbracket : \varphi \in \mathcal{L}^0\}$. By the previous observation, $\llbracket \mathcal{L}^0 \rrbracket \subseteq \Sigma$.

It has been proved in [20, 21] that if the set of states is an analytic space² and the set of labels L is countable, then \mathcal{L}^0 characterizes state bisimulation; that is, for any two given states $s, t \in S$, $s \sim_s t$ if and only if $s \mathcal{R}(\mathcal{L}^0) t$ (i.e., for all $\varphi \in \mathcal{L}^0$, $s \in \llbracket \varphi \rrbracket \Leftrightarrow t \in \llbracket \varphi \rrbracket$). Besides, [14] showed that \mathcal{L}^0 completely characterizes the event bisimulation in general. More precisely, they proved that $\sigma(\llbracket \mathcal{L}^0 \rrbracket)$, the σ -algebra generated by $\llbracket \mathcal{L}^0 \rrbracket$, is the smallest σ -algebra that is also an event bisimulation.

Summarizing the results, we have:

Theorem 5. *For every LMP $(S, \Sigma, \{\tau_a \mid a \in L\})$, $\sim_s \subseteq \sim_e = \mathcal{R}(\mathcal{L}^0)$. Moreover, if (S, Σ) is an analytic Borel space and L is countable then $\sim_s = \sim_e = \mathcal{R}(\mathcal{L}^0)$.*

It was shown in [42] that this result does not generalize to arbitrary measurable spaces.

As we mentioned, an LMP is an inherently deterministic model in the sense that the label determines a unique transition probability function. Thus, this model *does not* consider *internal* non-determinism. From the process algebra point of view, this is a significant drawback of this theory since internal non-determinism immediately arises in the modeling and analysis of systems. For example, internal non-determinism arises by abstracting internal activity (to later use weak bisimulation [37]) or by using state abstraction techniques (such as in model checking [12]). This can be seen more clearly in the next example.

Example 6. This time we consider two computer systems, each measuring a different particle moving in the real line. One of them moves half as fast as the other, and for this reason, the system that monitors this new particle, also samples half as fast. These two systems can be modeled in a single LMP where the state space is defined by $\mathbb{R} \times \{1, 2\}$ and each state (r, k) indicates that the particle k is in position r . As in Ex. 2, we also consider a threshold h .

The complete LMP is defined by: $(\mathbb{R} \times \{1, 2\}, \mathcal{B}(\mathbb{R} \times \{1, 2\}), \{\tau_a \mid a \in \mathbb{N} \cup \{low, high\}\})$, where $\tau_{low}(r, k) = \mathbf{if } r < h \mathbf{ then } \delta_{(r,k)} \mathbf{ else } \mathbf{0}$, $\tau_{high}(r, k) = \mathbf{if } r < h \mathbf{ then } \mathbf{0} \mathbf{ else } \delta_{(r,k)}$, and for all $n \in \mathbb{N}$, $\tau_n(r, 1) = \mu_r^n \circ f_1^{-1}$, $\tau_{2n}(r, 2) = \mu_r^n \circ f_2^{-1}$, and $\tau_{2n-1}(r, 2) = \mathbf{0}$, with $f_k(x) = (x, k)$ for all $x \in \mathbb{R}$.

Notice that the probability of the first particle going beyond the threshold is the same as the probability of the second particle going beyond the threshold if both are at the same position (i.e., in states $(r, 1)$ and $(r, 2)$ respectively) and the time to reach the threshold is not important. This is easy to see since each τ_n transition of the fast particle ($k = 1$) can be matched with a τ_{2n} transition of the slow particle ($k = 2$) and vice versa. Nevertheless, it is also clear that $(r, 1)$ and $(r, 2)$ are not (event nor state) bisimilar. In particular, if $q = \mu_{h-d}^n([h, \infty))$, $(h - d, 1) \in \llbracket \langle n \rangle_q \langle high \rangle_1 \top \rrbracket$ but $(h - d, 2) \notin \llbracket \langle n \rangle_q \langle high \rangle_1 \top \rrbracket$.

The distinction occurs only because the label $n \in \mathbb{N}$ (the elapsed time) is observable. Abstracting from time would require hiding this class of label, and any reasonable hiding operation will immediately end up with an object not expressible in terms of an LMP. \square

² A topological space is *Polish* if it is separable and completely metrizable. Examples of Polish spaces are the Euclidean spaces \mathbb{R}^n and all countable discrete spaces. Polish spaces are closed under countable product, and hence $A^{\mathbb{N}}$ (with A a countable discrete space) is Polish. Finally, an *analytic* space is the continuous image of a Polish space.

4 Non-deterministic Labeled Markov Processes

Non-deterministic Labeled Markov Processes (NLMP) were introduced in [17, 18] as a generalization of LMPs that enable the modeling of internal non-determinism. That is, in an NLMP, two different but equally labeled transition probabilities are allowed to leave the same state.

There have been several attempts to define non-deterministic continuous probabilistic transition systems and all of them are straightforward extensions of (simpler) discrete versions. There are two fundamental differences in the NLMP model. The first one is that the *non-deterministic transition function* T_a now maps states to *measurable sets of probability measures* rather than arbitrary sets as previous approaches do. This is motivated by the fact that the non-determinism has to be resolved using schedulers. If we allowed the target set of states to be an arbitrary subset (as in [6, 9, 15]), the system as a whole could suffer from non-measurability issues, which would mean that it could not be quantified. (Rigorously speaking, labels should also be provided with a σ -algebra in order to define schedulers, but we omit it in this first approach.) The second difference is inspired by the definition of LMP: we ask that, for each label $a \in L$, T_a is a measurable function. One of the reasons for this restriction is to have well defined modal operators of a probabilistic Hennessy-Milner logic, like in the LMP case.

Definition 7. A non-deterministic labeled Markov process (NLMP for short) is a structure $(S, \Sigma, \{T_a \mid a \in L\})$ where Σ is a σ -algebra on the set of states S , and for each label $a \in L$, $T_a : S \rightarrow \Delta(\Sigma)$ is measurable.

For the requirement that T_a is measurable, we need to endow $\Delta(\Sigma)$ with a σ -algebra. This is a key construction for the development of the theory of NLMPs.

Definition 8. $H(\Delta(\Sigma))$ is the minimal σ -algebra containing all sets $H_\xi \doteq \{\zeta \in \Delta(\Sigma) \mid \zeta \cap \xi \neq \emptyset\}$ with $\xi \in \Delta(\Sigma)$.

This construction is similar to that of the Effros-Borel spaces [34] and resembles the so-called hit-and-miss topologies [38]. Note that the generator set H_ξ contains all measurable sets that “hit” the measurable set ξ . Also observe that $T_a^{-1}(H_\xi)$ is the set of all states that “hit” the set of measures ξ through label a (i.e., $T_a^{-1}(H_\xi) = \{s \mid T_a(s) \cap \xi \neq \emptyset\}$). This forms the basis to existentially quantify over the non-determinism, and it is fundamental for the behavioral equivalences and the logic.

The next two examples (inspired by an example in [8]) show why T_a is required to map into measurable sets and to be measurable. For these examples we let the state space and σ -algebra be the real unit interval with the standard Borel σ -algebra.

Example 9. Let $\mathcal{V} = \{\delta_q \mid q \in V\}$, where V is a non-measurable set in $[0, 1]$ (that is, $V \subseteq [0, 1]$ and $V \notin \Sigma$). It can be shown that \mathcal{V} is not measurable in $\Delta(\Sigma)$. Let $T_a(s) = \mathcal{V}$ for all $s \in [0, 1]$. The resolution of the internal non-determinism by means of so-called schedulers (also adversaries or policies) [41, 46], would require to assign probabilities to all possible choices. This amounts to measure the non-measurable set $T_a(s)$. This is why we require that T_a maps into measurable sets. \square

Example 10. Let $T_a(s) = \{\mu\}$ for a fixed measure μ , and let $T_b(s) = \mathbf{if}(s \in V) \mathbf{then} \{\delta_1\} \mathbf{else} \emptyset$, for every $s \in [0, 1]$, with V being a non-measurable set. Notice that both

$T_a(s)$ and $T_b(s)$ are measurable sets for every $s \in [0, 1]$. Assume that there is a scheduler that chooses to first do a and then b starting at some state s . The probability under such scheduler of reaching state 1 after performing both transitions cannot be measured since it requires to apply μ to the set $T_b^{-1}(H_{\Delta(s)}) = V$ which is not measurable. Besides, we will later need that sets $T_a^{-1}(H_\xi)$ are measurable so that the semantics of the logic \mathcal{L}^1 maps into measurable sets. \square

Notice that an LMP can be regarded as an NLMP without internal non-determinism, that is, an NLMP in which $T_a(s)$ is either a singleton or the empty set for all $a \in L$ and $s \in S$. In fact, an LMP can be encoded as an NLMP by taking $T_a(s) = \{\tau_a(s)\} \setminus \{\mathbf{0}\}$. For this, it is necessary that singletons $\{\mu\}$ are measurable in $\Delta(\Sigma)$ for the NLMP to be well defined. (In general, it suffices that Σ is countably generated to ensure that singleton sets are measurable [17].) Moreover, it is also necessary that function T_a is measurable, which is actually the case. Indeed, it is not difficult to verify that T_a is measurable iff τ_a is measurable [48].

Example 11. Taking the previous definition, the LMP of the two particles moving on the real line of Ex. 6, can be translated into the NLMP $(\mathbb{R} \times \{1, 2\}, \mathcal{B}(\mathbb{R} \times \{1, 2\}), \{T_a \mid a \in \mathbb{N} \cup \{low, high\}\})$, where $T_{low}(r, k) = \mathbf{if}(r < h) \mathbf{then} \{\delta_{(r,k)}\} \mathbf{else} \emptyset$, $T_{high}(r, k) = \mathbf{if}(r < h) \mathbf{then} \emptyset \mathbf{else} \{\delta_{(r,k)}\}$, and for all $n \in \mathbb{N}$, $T_n(r, 1) = \{\mu_r^n \circ f_1^{-1}\}$, $T_{2n}(r, 2) = \{\mu_r^n \circ f_2^{-1}\}$, and $T_{2n-1}(r, 2) = \emptyset$, with $f_k(x) = (x, k)$ for all $x \in \mathbb{R}$.

Notice that if we abstract the time just like in process algebra, then we obtain the NLMP $(\mathbb{R} \times \{1, 2\}, \mathcal{B}(\mathbb{R} \times \{1, 2\}), \{T_a \mid a \in \{\epsilon, low, high\}\})$, with T_{low} and T_{high} as before and $T_\epsilon(r, k) = \{\mu_r^n \circ f_k^{-1} \mid n \in \mathbb{N}\}$. Clearly this last set is measurable in $\Delta(\mathcal{B}(\mathbb{R} \times \{1, 2\}))$ since it is countable (singleton sets are measurable in $\Delta(\mathcal{B}(\mathbb{R} \times \{1, 2\}))$, see [17]). Besides, it can be proved that T_ϵ is measurable which shows that this abstraction defines a proper NLMP. \square

The original definition of bisimulation given by Larsen and Skou [35] has been generalized to a continuous setting in, e.g., [5, 6, 15, 16, 45]. These definitions closely resemble the definition of Larsen and Skou, the only difference being that two measures are considered equivalent if they agree in every measurable union of equivalence classes induced by the relation. In our setting, this definition can be instantiated as follows

Definition 12. *A relation R is a state bisimulation on an NLMP $(S, \Sigma, \{T_a \mid a \in L\})$ if it is symmetric and for all $a \in L$, $s R t$ implies that for all $\mu \in T_a(s)$, there is $\mu' \in T_a(t)$ s.t. $\mu R \mu'$. We say that $s, t \in S$ are state bisimilar, denoted by $s \sim_s t$, if there is a state bisimulation R such that $s R t$.*

The relation \sim_s is the largest state bisimulation and it is also an equivalence relation [17, 48]. The proof of this follows the standard strategy of the classic bisimulation (see [37]). Apart from the probabilistic treatment, it only differs in that the composition $R \circ R'$ is granted to be state bisimulation if R and R' are *reflexive* state bisimulations. (If one of R or R' is not reflexive, $R \circ R'$ may not be a state bisimulation.) Besides, it is easy to show that a state bisimulation on an LMP is also a state bisimulation on the encoding NLMP and vice versa.

The next example revisits Example 6. Using the timed abstracted version of Example 11, it shows that it is possible to prove that the two particles behave the same (modulo state bisimulation).

Example 13. Take the time-abstracted NLMP of Ex. 11, and let $R = \{(r, 1), (r, 2) \mid r \in \mathbb{R}\}$. Notice that any measurable R -closed set has the form $Q_B = B \times \{1, 2\}$ for some $B \in \mathcal{B}(\mathbb{R})$. Hence $\mu_r^n \circ f_1^{-1}(Q_B) = \mu_r^n(B) = \mu_r^n \circ f_2^{-1}(Q_B)$ and therefore $(\mu_r^n \circ f_1^{-1}) R (\mu_r^n \circ f_2^{-1})$. Since $\mu \in T_\epsilon(r, 1)$ implies that $\mu = \mu_r^n \circ f_1^{-1}$ for some n , then there is some $\mu' \in T_\epsilon(r, 2)$ such that $\mu R \mu'$ (in fact, $\mu' = \mu_r^n \circ f_2^{-1}$). From here (and the cases of T_{low} and T_{high} , which we omit) it follows that R is a state bisimulation. \square

In the case of an LMP $(S, \Sigma, \{\tau_a \mid a \in L\})$, an event bisimulation is a sub- σ -algebra $\mathcal{E} \subseteq \Sigma$ such that all transition probability functions are \mathcal{E} - $\mathcal{A}(\mathcal{E})$ measurable. We state our generalization following the same idea.

Definition 14. An event bisimulation on an NLMP $(S, \Sigma, \{T_a \mid a \in L\})$ is a sub- σ -algebra \mathcal{E} of Σ s.t. $T_a : (S, \mathcal{E}) \rightarrow (\mathcal{A}(\Sigma), H(\mathcal{A}(\mathcal{E})))$ is measurable for each $a \in L$.

Note that T_a is the same function from S to $\mathcal{A}(\Sigma)$ only that, for \mathcal{E} to be an event bisimulation, it should be measurable from \mathcal{E} to $H(\mathcal{A}(\mathcal{E}))$. Here, $H(\mathcal{A}(\mathcal{E}))$ is the sub- σ -algebra of $H(\mathcal{A}(\Sigma))$ generated by $\{H_\xi \mid \xi \in \mathcal{A}(\mathcal{E})\}$.

Just like for LMPs, the notion of event bisimulation can be extended to relations: R is an event bisimulation if there is an event bisimulation \mathcal{E} s.t. $R = \mathcal{R}(\mathcal{E})$. More generally, we say that two states $s, t \in S$ are *event bisimilar*, denoted by $s \sim_e t$, if there is an event bisimulation \mathcal{E} such that $s \mathcal{R}(\mathcal{E}) t$. The fact that \sim_e is an equivalence relation is an immediate corollary of Theorem 16 given below. We remark that an event bisimulation on an LMP is also an event bisimulation on the encoding NLMP and vice versa.

For NLMPs, we introduce a third kind of bisimulation that we call *hit bisimulation*³. Rather than looking point-wise at probability measures as state bisimulations do, the definition of the hit bisimulation follows the idea of Def. 8 and verifies that both $T_a(s)$ and $T_a(t)$ hit the same measurable sets of probability measures which measure only R -closed sets.

Definition 15. A relation $R \subseteq S \times S$ is a hit bisimulation on the NLMP $(S, \Sigma, \{T_a \mid a \in L\})$ if it is symmetric and for all $a \in L$, $s R t$ implies that, for all $\xi \in \mathcal{A}(\Sigma(R))$, $T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset$. We say that $s, t \in S$ are hit bisimilar, denoted by $s \sim_h t$, if there is a hit bisimulation R such that $s R t$.

The relation \sim_h is the largest hit bisimulation and an equivalence relation. Hit bisimulations relate to the event bisimulations in NLMPs very much like the the state bisimulations relate to the event bisimulations in LMPs. In particular, R is a hit bisimulation if and only if $\Sigma(R)$ is an event bisimulation. This is indeed an important result that is central to eventually prove that \sim_h is also an event bisimulation and hence $\sim_h \subseteq \sim_e$. The fact that \sim_h is an equivalence relation is actually a consequence of the fact that it is also an event bisimulation (every event bisimulation is an equivalence relation directly from its definition). The details of all these results appeared in [17, 48].

A state bisimulation R is also a hit bisimulation. The proof of this relies on the fact that if $\xi \in \mathcal{A}(\Sigma(R))$, $\mu \in \xi$ and $\mu R \mu'$, then $\mu' \in \xi$. The rest of the proof is straightforward from the definitions. An immediate consequence is that $\sim_s \subseteq \sim_h$. As we will see later, the inclusion is proper.

³ In our original works [17, 18, 48], we called the state and hit bisimulations, “traditional” and “state” respectively. We are changing the names here as we find them more appropriate.

Nevertheless both notions of bisimulation agree on NLMPs that are *image denumerable*. That is, a hit bisimulation R is also a state bisimulation on any NLMP satisfying that for all $a \in L$, $s \in S$, $T_a(s)$ is denumerable. As a consequence of this, a state bisimulation on an LMP is a hit bisimulation on the translated NLMP and vice versa, since the translated NLMP is deterministic and hence image denumerable ($|T_a(s)| \leq 1$ for all $a \in L$ and $s \in S$).

Like for LMPs, we can also provide a Hennessy-Milner-like logic for NLMPs that characterizes event bisimulation in general and all the bisimulations under some conditions. As we will see in Ex. 18, \mathcal{L}^0 is not sufficiently expressive to characterize event bisimulation in NLMPs. Therefore, we need a richer logic. The logic we present below was introduced in [17, 18] and is related to the logic of Parma and Segala [40]. The main difference is that we consider two kinds of formulas: one that is interpreted on states, and another that is interpreted on measures. The syntax is as follows,

$$\begin{aligned}\varphi &\equiv \top \mid \varphi_1 \wedge \varphi_2 \mid \langle a \rangle \psi \\ \psi &\equiv \bigvee_{i \in I} \psi_i \mid \neg \psi \mid [\varphi]_{\geq q}\end{aligned}$$

where $a \in L$, I is a denumerable index set, and $q \in \mathbb{Q} \cap [0, 1]$. We denote by \mathcal{L}^1 the set of all formulas generated by the first production and by \mathcal{L}_Δ^1 the set of all formulas generated by the second production.

The semantics is defined with respect to an NLMP (S, Σ, T) . Formulas in \mathcal{L}^1 are interpreted as sets of states, and formulas in \mathcal{L}_Δ^1 are interpreted as sets of measures on the state space as follows,

$$\begin{aligned}\llbracket \top \rrbracket &= S & \llbracket \bigvee_{i \in I} \psi_i \rrbracket &= \bigcup_i \llbracket \psi_i \rrbracket \\ \llbracket \varphi_1 \wedge \varphi_2 \rrbracket &= \llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket & \llbracket \neg \psi \rrbracket &= \llbracket \psi \rrbracket^c \\ \llbracket \langle a \rangle \psi \rrbracket &= T_a^{-1}(H_{\llbracket \psi \rrbracket}) & \llbracket [\varphi]_{\geq q} \rrbracket &= \Delta^{\geq q}(\llbracket \varphi \rrbracket)\end{aligned}$$

In particular, notice that $\langle a \rangle \psi$ is satisfied at a state s whenever there is some measure $\mu \in T_a(s)$ that satisfies ψ , and that $[\varphi]_{\geq q}$ is satisfied by a measure μ whenever $\mu(\llbracket \varphi \rrbracket) \geq q$. As in the case of LMPs, the sets $\llbracket \varphi \rrbracket$ and $\llbracket \psi \rrbracket$ are measurable in Σ and $\Delta(\Sigma)$, respectively. For the rest of the section, fix $\llbracket \mathcal{L}^1 \rrbracket = \{\llbracket \varphi \rrbracket \mid \varphi \in \mathcal{L}^1\}$.

Note that some other operators can be encoded as syntactic sugar. For instance, we can define $[\varphi]_{>r} \equiv \bigvee_{q \in \mathbb{Q} \cap [0, 1] \wedge q > r} [\varphi]_{\geq q}$ for any real $r \in [0, 1]$, and $[\varphi]_{\leq r} \equiv \neg[\varphi]_{>r}$.

It can be shown that \mathcal{L}^1 characterizes event bisimulation for NLMPs. Following the lines of the proof for the logical characterization of event bisimilarity for LMPs, it can be proved that $\sigma(\llbracket \mathcal{L}^1 \rrbracket)$ is the smallest σ -algebra that is also an event bisimulation. A mild generalization of the concept of event bisimulation, namely families of sets being *stable*⁴ plays a role in the proof; it is immediate from the definition that stable σ -algebras are exactly the event bisimulations. It is shown in [17, Sect. 5] that $\llbracket \mathcal{L}^1 \rrbracket$ is the smallest stable family of subsets that is closed under finite intersections. A key lemma that appears in [47] ensures that $\sigma(C)$ is stable whenever C is, and the result follows:

⁴ The family $C \subseteq \Sigma$ is *stable* for an NLMP (S, Σ, T) if for all $a \in L$ and $\xi \in \Delta(C)$, $T_a^{-1}(H_\xi) \in C$. This notion of stability was further generalized by Doberkat [25] to the concept of *congruence* for stochastic systems.

Theorem 16. *The logic \mathcal{L}^1 completely characterizes event bisimulation. In other words, $\mathcal{R}(\mathcal{L}^1) = \sim_e$.*

A consequence of this theorem together with the previously discussed relations between the different bisimulations, is that both state and hit bisimulation are sound for \mathcal{L}^1 , i.e., they preserve the validity of formulas.

Theorem 17. $\sim_s \subseteq \sim_h \subseteq \sim_e = \mathcal{R}(\mathcal{L}^1)$.

As we will review in Examples 20 and 21, these inclusions are proper in general. Nevertheless, for image-finite NLMPs over analytic spaces it can be proved that the same logic is complete for state bisimilarity, and hence all notions are the same. (An NLMP is image finite if $T_a(s)$ is finite for all $a \in L$ and $s \in S$.) In fact, the sub-logic of \mathcal{L}^1 defined by

$$\varphi \equiv \top \mid \varphi_1 \wedge \varphi_2 \mid \langle a \rangle \bigwedge_{i=1}^n [\varphi_i]_{\triangleright_i, q_i} \tag{1}$$

where $\triangleright_i \in \{>, <\}$ and $q_i \in \mathbb{Q} \cap [0, 1]$, is complete for \sim_s under those restriction. Let \mathcal{L}^{1-} be the set of all formulas generated by (1).

It should be noted that the expression $\langle a \rangle \bigwedge_{i=1}^n [\varphi_i]_{\triangleright_i, q_i}$ may not be expressed as a conjunction of formulas $\langle a \rangle [\varphi_i]_{\triangleright_i, q_i}$ because the probabilistic bounds must be satisfied by the *same* non-deterministic transition. The next example from [10] illustrates this fact.

Example 18. Take the discrete NLMPs depicted in Fig. 1. States s and t are not bisimilar since given a measure $\mu \in T_a(s)$, there is no $\mu' \in T_a(t)$ such that $\mu(Q) = \mu'(Q)$ for all $Q \in \{\{x\}, \{y\}, \{z\}\}$ (which are the only relevant possible R -closed sets). A logic having a modality that can only describe one behavior after a label will not be able to distinguish between s and t . For example, $\llbracket \langle a \rangle [\varphi]_{>_q} \rrbracket = \{w \mid T_a(w) \cap \mathcal{A}^{>_q}(\llbracket \varphi \rrbracket) \neq \emptyset\}$ will always have s and t together, or none of them. Observe that negation, denumerable conjunction or disjunction, do not add any distinguishing power (on an image finite setting).

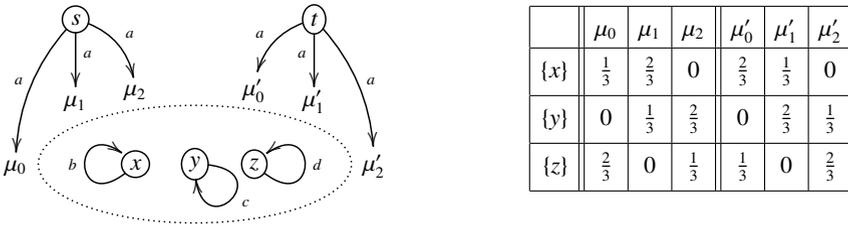


Fig. 1. s and t are not bisimilar

Notice, however, that the \mathcal{L}^{1-} formula $\langle a \rangle (\llbracket \langle b \rangle \top \rrbracket_{<_{\frac{2}{3}}} \wedge \llbracket \langle c \rangle \top \rrbracket_{>_{\frac{1}{3}}})$ is satisfied by s but not by t . □

The essential need for this new modal operator also shows that our σ -algebra $H(\mathcal{A}(\Sigma))$ in Def. 8 cannot be simplified to $\sigma(\{H_{\mathcal{A}^B(Q)} : B \in \mathcal{B}([0, 1]), Q \in \Sigma\})$. States s and t in the

example above should be observationally distinguished from each other. Formally, this amounts to saying that there must be some label a and some measurable $\Theta \in H(\Delta(\Sigma))$ such that $T_a^{-1}(\Theta)$ separates s from t . Therefore, the same must be true for some generator Θ , but this does not hold for the family $\{H_{\Delta^B(Q)} : B \in \mathcal{B}([0, 1]), Q \in \Sigma\}$.

The proof that our logic is complete for state bisimilarity follows from [17, Lemma 5.8] that states that given an NLMP (S, Σ, T) with (S, Σ) analytic, if we have a countable logic \mathcal{Q} with $\llbracket \mathcal{Q} \rrbracket \subseteq \Sigma$ satisfying some local criteria, then that logic must characterize \sim_s completely. (Logic \mathcal{Q} is countable if the number of formulas in \mathcal{Q} is countable.) In our case, \mathcal{L}^1 is not countable but the sub-logic \mathcal{L}^{1-} given by (1) satisfies all the requirements and hence $\sim_s = \mathcal{R}(\mathcal{L}^{1-})$. Since $\mathcal{R}(\mathcal{L}^1) \subseteq \mathcal{R}(\mathcal{L}^{1-})$, by Theorem 17 we have the following result.

Theorem 19. *Let (S, Σ, T) be an image finite NLMP with (S, Σ) being analytic. For all $s, t \in S$,*

$$s \sim_s t \Leftrightarrow s \sim_h t \Leftrightarrow s \sim_e t \Leftrightarrow s \mathcal{R}(\mathcal{L}^1) t$$

There are two delimiting results on possible generalizations of this theorem. First, the hypothesis of an analytic state space cannot be dropped completely (even for deterministic processes). This was seen in [42], where it is shown that state bisimilarity for LMPs is not characterized by \mathcal{L}^0 in general. Secondly, a generalization of the same arguments to image-countable processes is not feasible since there is no countable logic having formulas with measurable extensions that characterize state bisimilarity on such processes [43].

In fact, as we have already anticipated, the inclusions in Theorem 17 are proper in the general case. In the following, we construct counterexamples over standard Borel spaces witnessing that all our notions of bisimilarity are different in the case of uncountable non-determinism.

Moreover, it suffices to consider a non-probabilistic variant of NLMP, in which transitions only map into a set of Dirac measures. These structures look very much like LTSs, with the additional requirement that the set of states is endowed with a σ -algebra that the transition should respect. More formally, let (S, Σ) be a standard Borel space and $\delta(Q) = \{\delta_s \mid s \in Q\}$ for each $Q \in \Sigma$. An NLMP $\mathbf{S} = (S, \Sigma, \{T_a \mid a \in L\})$ is called *non-probabilistic* if for all $a \in L$ and $s \in S$, $T_a(s) \subseteq \delta(S)$.

Example 20. We will first construct a non-probabilistic NLMP witnessing the fact that state bisimilarity is strictly finer than the other notions. Consider the standard Borel space $(S_1, \Sigma_1) = ([0, 1] \cup [2, 3] \cup \{s, t, x\}, \mathcal{B}([0, 1] \cup [2, 3] \cup \{s, t, x\}))$ where $s, t, x \in \mathbb{R} \setminus [0, 3]$ are different. Let V be a non-Borel subset of $[2.5, 3]$. Clearly, $[0, 1]$ is equinumerous with $[2, 3] \setminus V$; pick a bijection f between them. Now, let $L_1 = \{a\} \cup [0, 1]$ be the set of labels and let $\mathbf{S}_1 = (S_1, \Sigma_1, \{T_a : a \in L_1\})$ be non-probabilistic such that

$$\begin{array}{ll} T_a(s) = \{\delta_d \mid d \in [2, 3]\} & T_r(r) = T_r(f(r)) = \{\delta_x\} \quad \text{if } r \in [0, 1] \\ T_a(t) = \{\delta_d \mid d \in [0, 1]\} & T_c(y) = \emptyset \quad \text{otherwise.} \end{array}$$

Now, take \mathcal{F} to be $\{\{s, t\}, \{r, f(r)\}_{r \in [0, 1]}\}$ and $R = \mathcal{R}(\sigma(\mathcal{F}))$. It is not hard to prove that \mathbf{S}_1 is a non-probabilistic NLMP, $\sigma(\mathcal{F})$ is an event bisimulation and R is a hit bisimulation that relate s and t . Also, it can be seen that s and t are not state bisimilar. But this shows that \sim_s differs from \sim_e and \sim_h . □

Example 21. By modifying slightly \mathbf{S}_1 we can show that the largest event bisimulation \sim_e is not contained in \sim_h . Take V to be the interval $(2.5, 3]$ and let $(S_2, \Sigma_2) = (S_1, \Sigma_1)$. We complete the construction of a non-probabilistic NLMP \mathbf{S}_2 by picking any bijection f between $[0, 1]$ and $[2, 2.5]$. The transition is defined just like for \mathbf{S}_1 only that using the new f . We also use family \mathcal{F} but defined with the new f . The same arguments for \mathbf{S}_1 go through here, showing that $s \sim_e t$ but $s \not\sim_h t$. \square

Some observations on the counterexamples are in order. First, counterexample \mathbf{S}_1 relies on the fact that hit bisimulation cannot distinguish a non-measurable set V while state bisimulation can. From our point of view, such distinction should *not* be possible since V has no measure. Second, counterexample \mathbf{S}_2 makes a difference on the measurable set V that the event bisimulation cannot distinguish. In our opinion, such distinction *should* be possible since some scheduler may lead to such set of states with certain probability. Note that in this example, states in V do not allow the system to reach state x from s , while x can always be reached from t . In particular, if the scheduler chooses uniformly the branching on the a -transition in both cases, the system starting from s will deadlock with probability 0.5 immediately while no deadlock is possible after a when starting from t . In this sense, we argue that hit bisimulation is the most appropriate definition.

Somehow, this is disappointing since logic \mathcal{L}^1 has a natural definition but, as it completely characterizes event bisimulation, it will not be able to test the presence of states like those in V in \mathbf{S}_2 . This is due to the fact that the modality $\langle a \rangle_-$ can only test one transition at a time and, together with the other operators, any \mathcal{L}^1 formula can only test countably many transitions at a time. Notice that a state in the set V can only be distinguished through a formula testing that no action (in the *uncountable* set $[0, 1]$) can be performed.

Therefore, both examples call for adding structure to the set of labels on the NLMP. In the first case, endowing the set of labels with a σ -algebra exclude the “bad behaved” NLMPs like \mathbf{S}_1 from the set of definable objects. In the second case, this will allow to define a richer logic that can test measurable sets of labels in a single formula. Regardless of these situations, a σ -algebra on the labels is also necessary for the definition of schedulers and probabilistic trace semantics [48].

5 Structured Non-deterministic Labeled Markov Processes

In view of the previous observations we developed a variant of NLMPs that requires that the set of labels has a measurable space associated. Since one of the aims of introducing structure on labels is to be able to define schedulers that resolve the (continuous) non-determinism of the model, we need to adapt the transition probability function to the new setting so that the different measurability aspects interact properly.

First, notice that a transition label is intended to represent the occurrence of a single action. Therefore, we will assume that, if L is the set of labels and Λ its associated σ -algebra, all singleton subsets of L are measurable in Λ .

Recall that a scheduler is a function that, given a particular execution history of the system, randomly selects a transition from those enabled at the last state of the execution. That is, given the fact that the execution finishes at state s in a given NLMP,

the scheduler has to randomly chose first a label a and then a measure in $T_a(s)$. More precisely, a scheduler will have to randomly choose an element in $\theta = \{(a, \mu) \mid \mu \in T_a(s)\}$. So, we actually need θ to be measurable in $\Lambda \otimes \mathcal{A}(\Sigma)$.

Therefore, a *structured* NLMP has a single transition function $T : S \rightarrow \Lambda \otimes \mathcal{A}(\Sigma)$ that assigns to each state a measurable set of pairs of label-probability measure on states. As in the case of LMP and NLMP, we will also need that the transition function is a measurable mapping. Hence, we need to endow $\Lambda \otimes \mathcal{A}(\Sigma)$ with a σ -algebra. We proceed in a similar way to Def. 8.

Definition 22. $H(\Lambda \otimes \mathcal{A}(\Sigma))$ is the smallest σ -algebra containing all sets $H_{\lambda \times \xi} = \{\theta \in \Lambda \otimes \mathcal{A}(\Sigma) \mid \theta \cap (\lambda \times \xi) \neq \emptyset\}$, with $\lambda \in \Lambda$ and $\xi \in \mathcal{A}(\Sigma)$.

Here we follow a slightly different approach to that of Def. 8 by taking only hit sets induced by rectangles rather than arbitrary measurable sets in the product σ -algebra.

Now, we can formally define the structured version of NLMPs:

Definition 23. A structured non-deterministic labeled Markov process (*SNLMP for short*) is a structure $(S, \Sigma, L, \Lambda, T)$ where Σ is a σ -algebra on the set of states S , Λ is a σ -algebra on the set of labels L so that $\{a\} \in \Lambda$ for all $a \in L$, and $T : S \rightarrow \Lambda \otimes \mathcal{A}(\Sigma)$ is measurable.

An SNLMP can be straightforwardly encoded as an NLMP by taking $T_a(\cdot) = T(\cdot)|_a$, where $\theta|_a \doteq \{\mu \in \mathcal{A}(S) \mid (a, \mu) \in \theta\}$ is the a -section of θ , known to be measurable if θ is measurable. Also, it is not difficult to see that, in our setting, the section seen as a function $(\cdot)|_a$ is a measurable function. This ensures the required properties of T_a . As it can be expected, NLMPs can not be encoded as SNLMPs in general. This is confirmed in the following example.

Example 24. Consider the NLMP \mathbf{S}_1 of Ex. 20. To translate it into an SNLMP, take $T(d) = \{(a, \mu) \mid \mu \in T_a(d)\}$ for all $d \in S_1$. Notice that

$$\begin{aligned} T(s) &= \{a\} \times \{\delta_d \mid d \in [2, 3]\} & T(r) &= T(f(r)) = \{(r, \delta_x)\} & \text{if } r \in [0, 1] \\ T(t) &= \{a\} \times \{\delta_d \mid d \in [0, 1]\} & T(y) &= \emptyset & \text{otherwise.} \end{aligned}$$

Though clearly $T(d)$ is a measurable set for any $d \in S_1$, T is not a measurable function. In effect, $T^{-1}(H_{[0,1] \times \mathcal{A}(S)}) = \{d \mid T(d) \cap ([0, 1] \times \mathcal{A}(S)) \neq \emptyset\} = [0, 1] \cup ([2, 3] \setminus V)$ which is not measurable, since V was chosen to be a non-Borel subset of $[2.5, 3]$. \square

Example 25. Notice, however, that \mathbf{S}_2 in Ex. 21 can be encoded as an SNLMP provided function f^{-1} is measurable. This is immediate after observing that

$$\begin{aligned} T^{-1}(H_{\lambda \times \xi}) &= \{s \mid a \in \lambda \wedge \{\delta_d \mid d \in [2, 3]\} \cap \xi \neq \emptyset\} \\ &\cup \{t \mid a \in \lambda \wedge \{\delta_d \mid d \in [0, 1]\} \cap \xi \neq \emptyset\} \\ &\cup \{d \mid d \in \lambda \cup f(\lambda) \wedge \delta_x \in \xi\} \end{aligned} \quad \square$$

All bisimulations introduced for NLMPs have their counterpart in SNLMPs. In fact, state bisimulation and hit bisimulation are defined exactly in the same way as for NLMPs by taking $T_a(\cdot) = T(\cdot)|_a$. For the event bisimulation, we also have to consider the fact that, in addition to states, labels are also observed through events.

Definition 26. An event bisimulation on an SNLMP $(S, \Sigma, L, \Lambda, T)$ is a sub- σ -algebra Ξ of Σ s.t. $T : (S, \Xi) \rightarrow (\Lambda \otimes \mathcal{A}(\Sigma), H(\Lambda \otimes \mathcal{A}(\Xi)))$ is measurable.

Just like for LMPs and NLMPs, the notion of event bisimulation can be extended to relations and the largest event bisimulation relation \sim_e can be analogously defined.

This way of defining event bisimulation raises the question on why not redefining also the hit bisimulation so that it considers the new hit sets containing pairs of labels and probability measures as in Def. 22. It turns out that this variant does not alter the definition of hit bisimulation as we state in the following.

Theorem 27. Consider the SNLMP $(S, \Sigma, L, \Lambda, T)$ and let $R \subseteq S \times S$ be a symmetric relation over states. The following characterizations for R are equivalent:

- (1) if $s R t$ then $T(s)|_a \cap \xi \neq \emptyset \Leftrightarrow T(t)|_a \cap \xi \neq \emptyset$, for all $a \in L$ and $\xi \in \mathcal{A}(\Sigma(R))$;
- (2) if $s R t$ then $T(s) \cap (\lambda \times \xi) \neq \emptyset \Leftrightarrow T(t) \cap (\lambda \times \xi) \neq \emptyset$, for all $\lambda \in \Lambda$ and $\xi \in \mathcal{A}(\Sigma(R))$;
- (3) if $s R t$ then $T(s) \cap \theta \neq \emptyset \Leftrightarrow T(t) \cap \theta \neq \emptyset$, for all $\theta \in \Lambda \otimes \mathcal{A}(\Sigma(R))$.

Clearly (3) implies (2) which implies (1). The proof that (1) implies (3) relies on the fact that $(\{\{a\} \times \mathcal{A}(S)\} \cap \theta)|_a \in \mathcal{A}(\Sigma(R))$. Notice that (1) is in fact the same definition of hit bisimulation as given in Def. 15 interpreting T_a as $T(\cdot)|_a$.

All results presented for the different bisimulations on NLMPs repeat on SNLMPs. In particular, it also holds that R is a hit bisimulation if and only if $\Sigma(R)$ is an event bisimulation (with the new definition of event bisimulation). Details can be found in [7].

Schedulers aside, the other reason to define SNLMPs was motivated by Ex. 21 in which the logic \mathcal{L}^1 failed to distinguish states s and t in \mathbf{S}_2 . As we saw in Ex. 25, \mathbf{S}_2 is also an SNLMP. Therefore we would like to define a new logic that can distinguish s and t . To understand the difference, notice that t can perform an a -transition and reach a state where no transition labeled with $r \in [0, 1]$ can be performed with probability 1. This behavior could be described by a formula like $\langle a \rangle \neg [\langle [0, 1] \rangle \top]_{\geq 1} \geq 1$.

Indeed, the logic \mathcal{L}^2 is the same logic as \mathcal{L}^1 where the modal construct $\langle a \rangle \psi$ has been replaced by $\langle \lambda \rangle \psi$ with $\lambda \in \Lambda$. The semantics of this new operator is given by

$$\llbracket \langle \lambda \rangle \psi \rrbracket \doteq T^{-1} \left(H_{\lambda \times \llbracket \psi \rrbracket} \right).$$

The semantics for the rest of the operations of \mathcal{L}^2 are defined just like for \mathcal{L}^1 . Again, $\llbracket \psi \rrbracket$ is measurable for all $\varphi \in \mathcal{L}^2$.

Because singletons are measurable in Λ , $\langle \{a\} \rangle \psi \in \mathcal{L}^2$ provided $a \in L$ and $\psi \in \mathcal{L}^2$ (we will use $\langle a \rangle \psi$ as a shorthand). Therefore \mathcal{L}^2 is at least as expressive as \mathcal{L}^1 . Moreover, it is strictly more expressive since $t \in \llbracket \langle a \rangle \neg [\langle [0, 1] \rangle \top]_{\geq 1} \geq 1 \rrbracket$ but $s \notin \llbracket \langle a \rangle \neg [\langle [0, 1] \rangle \top]_{\geq 1} \geq 1 \rrbracket$ and hence \mathcal{L}^2 can distinguish states s and t in \mathbf{S}_2 .

It can be shown that \mathcal{L}^2 characterizes the event bisimulation for SNLMPs. The proof follows the same strategy as that of Theorem 16.

Theorem 28. The logic \mathcal{L}^2 completely characterizes event bisimulation on SNLMPs. I.e. $\mathcal{R}(\mathcal{L}^2) = \sim_e$.

For the next result, we need to interpret logics \mathcal{L}^0 and \mathcal{L}^1 on SNLMP, but this is easy since $\langle a \rangle \psi$ of \mathcal{L}^1 corresponds to $\langle \{a\} \rangle \psi$ in \mathcal{L}^2 and $\langle a \rangle_q \varphi$ of \mathcal{L}^0 corresponds to $\langle \{a\} \rangle [\varphi]_{\geq q}$ in \mathcal{L}^2 . The following theorem summarizes the results for bisimulations and logics in SNLMPs.

Theorem 29. $\sim_s \subseteq \sim_h \subsetneq \sim_e = \mathcal{R}(\mathcal{L}^2) \subsetneq \mathcal{R}(\mathcal{L}^1) \subsetneq \mathcal{R}(\mathcal{L}^0)$.

The last inclusion is shown to be proper in Ex. 18. Besides we also showed that the inclusion $\mathcal{R}(\mathcal{L}^2) \subseteq \mathcal{R}(\mathcal{L}^1)$ is proper using SNLMP \mathbf{S}_2 . The next example shows that inclusion $\sim_h \subseteq \sim_e$ is also proper.

Example 30. Consider the SNLMP \mathbf{S}_3 which is a variant of \mathbf{S}_2 where $V = (2.5, 3]$, f is measurable, and T is redefined as follows.

$$\begin{aligned} T(s) &= \{a\} \times \{\delta_d \mid d \in [2, 3]\} \\ T(t) &= \{a\} \times \{\delta_d \mid d \in [0, 1]\} \\ T(r) &= T(f(r)) = \{(r, \delta_x) \mid r \in [0, 1] \setminus \{r\}\} && \text{if } r \in [0, 1] \\ T(d) &= \{(r, \delta_x) \mid r \in [0, 1]\} && \text{if } d \in V \\ T(y) &= \emptyset && \text{otherwise.} \end{aligned}$$

Note that the states r and $f(r)$ can perform any $[0, 1]$ -labeled transition except for the r -labeled transition whenever $r \in [0, 1]$. Instead, every $d \in V$ can perform all $[0, 1]$ -labeled transitions. Therefore, every pair of states in V are hit bisimilar, and every state $d \in V$ can be distinguished from states in $[0, 1] \cup [2, 3] \setminus V$ since $T(d)|_r \cap \mathcal{A}(S) = \{\delta_x\} \neq \emptyset = T(r)|_r \cap \mathcal{A}(S) = T(f(r))|_r \cap \mathcal{A}(S)$. Thus, $V \in \Sigma$ is \sim_h -closed and consequently $\delta_V = \{\delta_d \mid d \in V\} \in \mathcal{A}(\Sigma(\sim_h))$. From here we have that $T(s)|_a \cap \delta_V = \delta_V \neq \emptyset = T(t)|_a \cap \delta_V$, and therefore s and t are not hit bisimilar.

Now, take $\mathcal{F} = \{\{s, t\}, \{x\}, \{r, f(r)\}_{r \in [0, 1]}\}$. It is not hard to prove that $\sigma(\mathcal{F})$ is an event bisimulation. Hence $s \sim_e t$. \square

Contrarily to what happens in NLMPs with example \mathbf{S}_2 , the example above questions the hit bisimulation (rather than the event bisimulation) as it seems to distinguish sets of null measure. In fact any definable scheduler starting from state t has an “almost surely equivalent” scheduler starting from s (modulo state renaming).

Also, the question of whether state bisimilarity is strictly finer than hit bisimilarity on SNLMPs remains open. Notice that Ex. 20 is not a valid counterexample in the realm of SNLMPs because \mathbf{S}_1 is not an SNLMP.

6 Concluding Remarks

In this paper, we have presented the basic theory of LMPs and its extensions with internal non-determinism, namely NLMPs and SNLMPs. Much more research on this subject has been done. For instance, pseudometrics that behave like bisimulation in the limit have been defined and different kind of approximations for LMPs have been studied [4, 11, 13, 23, 24, 39, etc.].

When it comes to give semantics to languages or symbolic models that includes stochastic continuous behavior, NLMPs showed to be useful. Stochastic automata [15, 16] provide a symbolic framework to model soft real-timed systems. They can be seen as a non-deterministic extension of generalized semi-Markov processes that are amenable to composition. The semantics of a stochastic automaton naturally arises as a

(structured) NLMP [48]. In a similar manner NLMPs have been used to give semantics to more complex models, particularly stochastic hybrid automata [27, 31]. As a consequence, NLMPs are also the concrete underlying semantics of a process algebra like SPADES [15, 16] and modeling languages like MODEST and HMODEST [3, 31]. These languages have been used to analyze real case studies (e.g. [30, 32]). [48] presents also mappings from pGCL [36] and abstract probabilistic automata [19] into NLMPs.

We have made use of the concept of schedulers to introduce SNLMPs. In fact, we have formally defined schedulers on SNLMPs and used them to define trace distribution semantics. See [48] for these results. We remark that a still overdue result in the setting of LMPs and NLMPs is a correspondence execution theorem which states that if two states are bisimilar (in any of the senses defined here), they share the same probabilistic execution structure and hence they are also trace distribution equivalent.

Desharnais et al. [22] followed a different approach to extend LMPs with some kind of internal non-determinism. Rather than explicitly introducing the branching set of probability measures as in NLMPs, they relax the requirements on the LMP by only asking that $\tau_a(s)$ is a super-additive function on Σ (instead of a sub-probability measure). They call this new model *infLMP*. An *infLMP* can be understood as a partially specified system where a possible implementation is an LMP in which its transition probability function is greater than or equal to the transition super-additive function of the *infLMP*. It would be interesting to draw conclusions whether NLMPs can capture *infLMPs* or not. A first (but inconclusive) approach to this relation is reported in [48].

Finally, the results in [42, 43] show that the generality of the models immediately leads to unwanted results. It seems reasonable to restrict only to standard Borel spaces. Confining to standard Borel spaces is not as restricting as it seems since most natural problems arise in this setting. For example, we have that the underlying semantics of stochastic (hybrid) automata is given in terms of an NLMP on standard Borel spaces, and in the case of stochastic automata, such NLMP is also image finite. Recall that stochastic automata and similar models are used to give semantics to stochastic process algebras and specification languages, see e.g. [3, 5, 6, 15, 16, 31]. Moreover, LMP-like models restricted to standard Borel spaces have been studied in [26].

References

1. Ash, R., Doléans-Dade, C.: Probability & Measure Theory. Academic Press (2000)
2. Baier, C., Katoen, J.: Principles of Model Checking. The MIT Press (2008)
3. Bohnenkamp, H., D'Argenio, P., Hermanns, H., Katoen, J.P.: MoDeST: A compositional modeling formalism for real-time and stochastic systems. *IEEE Trans. Softw. Eng.* 32(10), 812–830 (2006)
4. Bouchard-Côté, A., Ferns, N., Panangaden, P., Precup, D.: An approximation algorithm for labelled markov processes: towards realistic approximation. In: Proc. of QEST 2005, pp. 54–62. IEEE Computer Society (2005)
5. Bravetti, M.: Specification and Analysis of Stochastic Real-Time Systems. Ph.D. thesis, Università di Bologna, Padova, Venezia (2002)
6. Bravetti, M., D'Argenio, P.R.: Tutte le algebre insieme: Concepts, discussions and relations of stochastic process algebras with general distributions. In: Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P., Siegle, M. (eds.) Validation of Stochastic Systems. LNCS, vol. 2925, pp. 44–88. Springer, Heidelberg (2004)

7. Budde, C.: No determinismo completamente medible en procesos probabilísticos continuos. Master's thesis, FaMAF, Universidad Nacional de Córdoba (2012)
8. Cattani, S.: Trace-based Process Algebras for Real-Time Probabilistic Systems. Ph.D. thesis, University of Birmingham (2005)
9. Cattani, S., Segala, R., Kwiatkowska, M., Norman, G.: Stochastic transition systems for continuous state spaces and non-determinism. In: Sassone, V. (ed.) FOSSACS 2005. LNCS, vol. 3441, pp. 125–139. Springer, Heidelberg (2005)
10. Celayes, P.: Procesos de Markov Etiquetados sobre Espacios de Borel Estándar. Master's thesis, FaMAF, Universidad Nacional de Córdoba (2006)
11. Chaput, P., Danos, V., Panangaden, P., Plotkin, G.: Approximating labelled markov processes again! In: Kurz, A., Lenisa, M., Tarlecki, A. (eds.) CALCO 2009. LNCS, vol. 5728, pp. 145–156. Springer, Heidelberg (2009)
12. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press (1999)
13. Danos, V., Desharnais, J.: Labelled markov processes: Stronger and faster approximations. In: Proc. of 18th LICS, pp. 341–350. IEEE Computer Society (2003)
14. Danos, V., Desharnais, J., Laviolette, F., Panangaden, P.: Bisimulation and cocongruence for probabilistic systems. *Inf. & Comp.* 204, 503–523 (2006)
15. D'Argenio, P.: Algebras and Automata for Timed and Stochastic Systems. Ph.D. thesis, Department of Computer Science, University of Twente (1999)
16. D'Argenio, P., Katoen, J.P.: A theory of stochastic systems, Part I: Stochastic automata, and Part II: Process algebra. *Inf. & Comp.* 203(1), 1–38, 39–74 (2005)
17. D'Argenio, P., Sánchez Terraf, P., Wolovick, N.: Bisimulations for non-deterministic labelled Markov processes. *Mathematical Structures in Comp. Sci.* 22(1), 43–68 (2012)
18. D'Argenio, P., Wolovick, N., Sánchez Terraf, P., Celayes, P.: Nondeterministic labeled Markov processes: Bisimulations and logical characterization. In: Proc. of QEST 2009, pp. 11–20. IEEE Computer Society (2009)
19. Delahaye, B., Katoen, J.-P., Larsen, K.G., Legay, A., Pedersen, M.L., Sher, F., Wąsowski, A.: Abstract probabilistic automata. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 324–339. Springer, Heidelberg (2011)
20. Desharnais, J.: Labeled Markov Process. Ph.D. thesis, McGill University (1999)
21. Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. *Inf. & Comp.* 179(2), 163–193 (2002)
22. Desharnais, J., Laviolette, F., Turgeon, A.: A logical duality for underspecified probabilistic systems. *Inf. Comput.* 209(5), 850–871 (2011)
23. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Approximating labelled markov processes. *Inf. Comput.* 184(1), 160–200 (2003)
24. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled markov processes. *Theor. Comput. Sci.* 318(3), 323–354 (2004)
25. Doberkat, E.E.: Kleisli morphisms and randomized congruences for the Giry monad. *Journal of Pure and Applied Algebra* 211(3), 638–664 (2007)
26. Doberkat, E.E.: Stochastic relations. Foundations for Markov transition systems. *Studies in Informatics Series*. Chapman & Hall/CRC (2007)
27. Fränzle, M., Hahn, E., Hermanns, H., Wolovick, N., Zhang, L.: Measurability and safety verification for stochastic hybrid systems. In: Caccamo, M., Frazzoli, E., Grosu, R. (eds.) Proc. of HSCC 2011, pp. 43–52. ACM (2011)
28. Giry, M.: A categorical approach to probability theory. In: *Categorical Aspects of Topology and Analysis*. LNM, vol. 915, pp. 68–85. Springer (1981)
29. van Glabeek, R.: The linear time–branching time spectrum I. The semantics of concrete, sequential processes. In: Bergstra, J.A., Ponse, A., Smolka, S.A. (eds.) *Handbook of Process Algebra*, pp. 3–99. North-Holland (2001)

30. Baró Graf, H., Hermanns, H., Kulshrestha, J., Peter, J., Vahldiek, A., Vasudevan, A.: A verified- wireless safety critical hard real-time design. In: Proc. of WOWMOM 2011, pp. 1–9. IEEE (2011)
31. Hahn, E., Hartmanns, A., Hermanns, H., Katoen, J.P.: A compositional modeling and analysis framework for stochastic hybrid systems. *Formal Methods in System Design* 43(2), 191–232 (2013)
32. Hartmanns, A., Hermanns, H.: Modelling and decentralised runtime control of self-stabilising power micro grids. In: Margaria, T., Steffen, B. (eds.) *ISoLA 2012, Part I. LNCS*, vol. 7609, pp. 420–439. Springer, Heidelberg (2012)
33. Hennessy, M., Milner, R.: Algebraic laws for nondeterminism and concurrency. *J. ACM* 32(1), 137–161 (1985)
34. Kechris, A.: *Classical Descriptive Set Theory*. Graduate Texts in Mathematics, vol. 156. Springer (1995)
35. Larsen, K., Skou, A.: Bisimulation through probabilistic testing. *Inf. & Comp.* 94(1), 1–28 (1991)
36. McIver, A., Morgan, C.: *Abstraction, Refinement and Proof for Probabilistic Systems*. Monographs in Computer Science. Springer (2005)
37. Milner, R.: *Communication and Concurrency*. Prentice Hall (1989)
38. Nainpally, S.: What is a hit-and-miss topology? *Topological Comment* 8(1) (2003)
39. Panangaden, P.: *Labelled Markov Processes*. Imperial College Press (2009)
40. Parma, A., Segala, R.: Logical characterizations of bisimulations for discrete probabilistic systems. In: Seidl, H. (ed.) *FOSSACS 2007. LNCS*, vol. 4423, pp. 287–301. Springer, Heidelberg (2007)
41. Puterman, M.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience (1994)
42. Sánchez Terraf, P.: Unprovability of the logical characterization of bisimulation. *Inf. & Comp.* 209(7), 1048–1056 (2011)
43. Sánchez Terraf, P.: Bisimilarity is not Borel. CoRR arXiv:1211.0967 (2012)
44. Segala, R.: *Modeling and Verification of Randomized Distributed Real-Time Systems*. Ph.D. thesis, Massachusetts Institute of Technology (1995)
45. Strulo, B.: *Process Algebra for Discrete Event Simulation*. Ph.D. thesis, Department of Computing, Imperial College, University of London (1993)
46. Vardi, M.: Automatic verification of probabilistic concurrent finite-state programs. In: 26th FOCS, pp. 327–338. IEEE (1985)
47. Viglizzo, I.: *Coalgebras on Measurable Spaces*. Ph.D. thesis, Department of Mathematics, Indiana University (2005)
48. Wolovick, N.: *Continuous probability and nondeterminism in labeled transition systems*. Ph.D. thesis, Universidad Nacional de Córdoba (2012)