2009 Sixth International Conference on Quantitative Evaluation of Systems

Nondeterministic Labeled Markov Processes: Bisimulations and Logical Characterization

Pedro R. D'Argenio*, Nicolás Wolovick, Pedro Sánchez Terraf*, and Pablo Celayes*
*FaMAF, Universidad Nacional de Córdoba (*also CONICET)*
Ciudad Universitaria - 5000, Córdoba - Argentina.
email: {dargenio,nicolasw,sterraf,celayes}@famaf.unc.edu.ar

Abstract

We extend the theory of labeled Markov processes with internal *nondeterminism, a fundamental concept for the further development of a process theory with abstraction on nondeterministic continuous probabilistic systems. We define* nondeterministic labeled Markov processes (NLMP) *and provide both a state based bisimulation and an event based bisimulation. We show the relation between them, including that the largest state bisimulation is also an event bisimulation. We also introduce a variation of the Hennessy-Milner logic that characterizes event bisimulation and that is sound w.r.t. the state base bisimulation for arbitrary NLMP. This logic, however, is infinitary as it contains a denumerable $\vee$. We then introduce a finitary sublogic that characterize both state and event bisimulation for image finite NLMP whose underlying measure space is also analytic. Hence, in this setting, all notions of bisimulation we deal with turn out to be equal.*

1. Introduction

Markov processes with continuous-state spaces or continuous time evolution (or both) arise naturally in several fields of physics, biology, economics, and computer science [11]. Many formal frameworks have been defined to study them from a process theory or process algebra perspective (see [4], [5], [7], [8], [11]–[15], [26]). A prominent and extensive work on this area is the one that builds on top of the so called labeled Markov processes (LMP) [14], [15]. This is due to its solid and well understood mathematical foundations. A 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 process algebra point of view, this is a significant drawback for this theory since internal nondeterminism immediately arises in the analysis of systems, e.g., because of abstracting internal activity (such as weak bisimulation [22]) or because of state abstraction techniques (such as in model checking [10]).

Many other works defined variants of continuous Markov processes that include internal nondeterminism and are mainly used as the underlying semantics of a process algebra [4], [5], [12], [13], [26]. They also defined a continuous probabilistic variant of the (strong) bisimulation. As correctly pointed out by Cattani et al. [7], [8], these models lack enough structure to ensure that bisimilar models share the same observable behavior. (This is due to the case in which two objects may be bisimilar but in one of them it is not possible to define probabilistic executions since the transition relation is not a measurable object.) The solution proposed in [7], [8] deals with the same unstructured type of models and 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.

A somewhat related observation has been made in [11] with respect to the bisimulation relation on LMPs [14], [15]. [11] shows that there are bisimulation relations that may distinguish *beyond events*. That is, states that cannot be *separated* (i.e., distinguished) by any *measurable set* (i.e., an event) may not be related for some bisimulation relation. This is also awkward as events (measurable sets) are the building blocks of observations (probabilistic executions). To overcome this, [11] defines the so called *event bisimulation* (in

Supported by ANPCyT PICT 26135, CONICET PIP 6391, and SeCyT-UNC.

978-0-7695-3808-2/09 $25.00 © 2009 IEEE
DOI 10.1109/QEST.2009.17

opposition to the previous *state* bisimulation—name which we will use from now on). An event bisimulation is a sub σ-algebra Λ on the set of states such that the original transition probability functions is also a Markov kernel on Λ, i.e., the original LMP is also an LMP under Λ. Λ induces an equivalence relation $\mathcal{R}(\Lambda)$ also called event bisimulation. Fortunately, it turns out that the largest state bisimulation is also an event bisimulation.

In this paper, we follow the LMP approach towards defining a theory of LMP with internal nondeterminism. Thus, we introduce *nondeterministic labeled Markov processes (NLMP)*. A NLMP has a nondeterministic 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). Moreover, T_a should be measurable. This calls for a definition of a σ-algebra on top of Giry's σ-algebra on the set of probability measures [18], which we also provide. We give a definition for event bisimulation and state bisimulation and prove similar properties to [11], including that the largest state bisimulation is also an event bisimulation. We also provide a definition of "traditional" bisimulation that follows the lines of [4], [12], [13], [26]. We prove that a traditional bisimulation is also a state bisimulation and give sufficient conditions so that the converse holds. Besides, we show that LMPs are just NLMPs without internal nondeterminism and that state (resp. event) bisimulation in the different models agree.

Behavioral equivalences like bisimulation have been characterized using logic with modalities, notably the Hennessy-Milner logic (see e.g. [19]). We define an extension of the logic presented in the context of LMP [14]. In fact, the logic is similar to that of [24], which was introduced in a discrete setting. However, unlike [24], we consider two different formula levels: one that is interpreted on states, and the other that is interpreted on measures. Such separation gives a particular insight: the actual complexity of the model lies exactly on the internal nondeterminism introduced by the target of function T_a (which is a set of measures). At state level, the logic is as simple as in [14]. We show that this logic completely characterizes event bisimulation and, as a consequence, it is sound w.r.t. traditional and state bisimulation.

In addition, we show that a sublogic of the previous logic characterizes all three bisimulations (event, state and traditional) provided certain restrictions apply, namely, NLMPs are image finite and the state space is analytic. Therefore, all bisimulation equivalences as well as logical equivalence turn out to be the same on this setting.

2. Fundamentals and Background

In this section we review some foundational theory and prove few basic results that will be of use throughout the paper.

Measure theory. Given a set Ω and a collection $\mathcal{F}$ of subsets of Ω, we call $\mathcal{F}$ a *σ-algebra* iff $\Omega \in \mathcal{F}$ and $\mathcal{F}$ is closed under complement and denumerable union. By $\sigma(\mathcal{H})$ we denote the *σ-algebra generated by* the family $\mathcal{H} \subseteq 2^{\Omega}$, i.e., the minimal σ-algebra containing $\mathcal{H}$. Each element of $\mathcal{H}$ is called *generator* and $\mathcal{H}$, the *set of generators*. We call the pair $(\Omega, \mathcal{F})$ a *measurable space*. A *measurable set* is a set $A \in \mathcal{F}$. Let $(\Omega_i, \mathcal{F}_i), i = 1, 2, \ldots, n$ be measurable spaces and $\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n$. A *measurable rectangle* in Ω is a set $A = A_1 \times A_2 \times \cdots \times A_n$, where $A_i \in \mathcal{F}_i$. The *product σ-algebra* $\mathcal{F}_1 \otimes \mathcal{F}_2 \otimes \cdots \otimes \mathcal{F}_n$ is the σ-algebra generated by measurable rectangles. A σ-additive function $\mu : \mathcal{F} \to [0, 1]$ such that $\mu(\Omega) = 1$ is called *probability measure*. By δ_a we denote the Dirac probability measure concentrated in $\{a\}$. Let $\Delta(\Omega)$ denote the *set of all probability measures* over the measurable space $(\Omega, \mathcal{F})$. A function $f : \Omega_1, \to \Omega_2$ is said to be *measurable* if $\forall A_2 \in \mathcal{F}_2,\ f^{-1}(A_2) \in \mathcal{F}_1$, i.e., the inverse function maps measurable sets to measurable sets.

A function $f : \Omega_1 \times \mathcal{F}_2 \to [0, 1]$ is a *transition probability* (also called *Markov Kernel*) if for all $\omega_1 \in \Omega_1$, $f(\omega_1, \cdot)$ is a probability measure on $(\Omega_2, \mathcal{F}_2)$ and for all $A_2 \in \mathcal{F}_2$, $f(\cdot, A_2)$ is measurable.

There is a standard construction by Giry [18] to endow $\Delta(\Omega)$ with a σ-algebra as follows: $\Delta(\mathcal{F})$ is defined as the σ-algebra generated by the sets of probability measures $\Delta^B(Q) \doteq \{\nu \mid \nu(Q) \in B\}$, with $Q \in \mathcal{F}$ and $B \in \mathcal{B}([0, 1])$. ($\mathcal{B}([0, 1])$ is the Borel σ-algebra on the interval $[0, 1]$ generated by the open sets.) When $0 \leq p \leq 1$, we will write $\Delta^{\geq p}(Q)$, $\Delta^{>p}(Q)$, $\Delta^{<p}(Q)$, etc. for $\Delta^B(Q)$ with $B = [p, 1], (p, 1], [0, p)$, etc. respectively. It is known that the set $\{\Delta^{\geq p}(Q) \mid p \in (\mathbb{Q} \cap [0, 1]), Q \in \mathcal{F}\}$ generates all $\Delta(\mathcal{F})$.

On this setting, $f : \Omega_1 \times \mathcal{F}_2 \to [0, 1]$ is a transition probability if and only if its curried version $f : \Omega_1 \to \Delta(\Omega_2)$ is measurable. (Mind the notation overloading on f.) This follows from the next lemma.

Lemma 1: $f : \Omega_1 \to \Delta(\Omega_2)$ is measurable iff $f(\cdot, Q) : \Omega_1 \to [0, 1]$ is measurable for all $Q \in \mathcal{F}_2$.

Proof: It is routine to calculate that $f^{-1}(\Delta^B(Q)) = (f(\cdot, Q))^{-1}(B)$ for all $Q \in \mathcal{F}_2$ and $B \in \mathcal{B}([0, 1])$. By this observation, $f^{-1}(\Delta^B(Q)) \in \mathcal{F}_1$ iff $(f(\cdot, Q))^{-1}(B) \in \mathcal{F}_1$. Since it is sufficient to show that $f^{-1}(\Delta^B(Q)) \in \mathcal{F}_1$ for all generators $\Delta^B(Q)$ to state that f is measurable [1], the lemma follows. $\square$

An important result on Giry's construction is that the σ-algebra of measures is *separative* [6], i.e., for any two elements, there is always a measurable set that contains one element but not the other.

Proposition 1: $\Delta(\Sigma)$ is *separative*. That is, given different $\mu, \mu' \in \Delta(S)$, there exists $\Theta \in \Delta(\Sigma)$ such that $\mu \in \Theta$ and $\mu' \notin \Theta$.

Relations, Measures, and σ-algebras. Given a relation $R \subseteq S \times S$, the predicate R-closed(Q) denotes $R(Q) \subseteq Q$. Notice that if R is symmetric, R-closed(Q) if and only if $\forall s,t : sRt : s \in Q \Leftrightarrow t \in Q$. Let (S, Σ) be a measurable space. For symmetric R, define $\Sigma(R) \doteq \{Q \in \Sigma \mid R\text{-closed}(Q)\}$. $\Sigma(R)$ is the sub-σ-algebra of Σ containing all R-closed Σ-measurable sets. The next proposition states that the inclusion order between two relations transfers inversely to the σ-algebras induced by them and to Giry's construction applied to these σ-algebras.

Proposition 2: Let R and R' be symmetric relations such that $R \subseteq R'$. Then *(i)* $\Sigma(R) \supseteq \Sigma(R')$ and *(ii)* $\Delta(\Sigma(R)) \supseteq \Delta(\Sigma(R'))$.

Proof: (i) follows from the fact that any measurable set that is R'-closed is also R-closed whenever $R \subseteq R'$. For (ii), recall that $\Delta(\Sigma(R'))$ is generated by $\mathcal{G} = \{\Delta^B(Q) \mid Q \in \Sigma(R') \text{ and } B \in \mathcal{B}([0,1])\}$. Since $\Sigma(R') \subseteq \Sigma(R)$ (by (i)), then $\mathcal{G} \subseteq \Delta(\Sigma(R))$ from which the lemma follows. $\square$

We can lift R to an equivalence relation in $\Delta(S)$ as follows: $\mu R \mu'$ iff $\forall Q \in \Sigma(R) : \mu(Q) = \mu'(Q)$. Then, the predicate R-closed can be defined on subsets of $\Delta(S)$ just like before. The following proposition will be useful.

Proposition 3: If R is a symmetric relation, every $\Delta(\Sigma(R))$-measurable set is R-closed.

Proof: Let $Q \in \Sigma(R)$ and $B \in \mathcal{B}([0,1])$. Then, if $\mu \in \Delta^B(Q)$ and $\mu R \mu'$, $\mu' \in \Delta^B(Q)$. So, each generator $\Delta^B(Q)$ of $\Delta(\Sigma(R))$ is R-closed. Moreover, for any symmetric R, the property of being R-closed is preserved by denumerable union and complement. Since the lifted R is symmetric, we can conclude that every $\Delta(\Sigma(R))$-measurable set is R-closed. $\square$

A σ-algebra Σ defines an equivalence relation $\mathcal{R}(\Sigma)$ on S as follows: $s\mathcal{R}(\Sigma)t$ iff $\forall Q \in \Sigma$, $s \in Q \Leftrightarrow t \in Q$. That is, two elements are related if they cannot be separated by any measurable set. The following properties (due to [11]) appear here for the sake of completeness; they relate σ-algebras and relations. In particular, (v) is a consequence of (i) and (ii).

Proposition 4: Let (S, Σ) be a measurable space, R a symmetric relation on S, and $\Lambda \subseteq \Sigma$ a sub-σ-algebra of Σ. Then, *(i)* $\Lambda \subseteq \Sigma(\mathcal{R}(\Lambda))$; *(ii)* $R \subseteq \mathcal{R}(\Sigma(R))$; *(iii)* if each R-equivalence class is in Σ, then $R = \mathcal{R}(\Sigma(R))$; *(iv)* $\mathcal{R}(\Lambda) = \mathcal{R}(\Sigma(\mathcal{R}(\Lambda)))$; and *(v)* $\Sigma(R) = \Sigma(\mathcal{R}(\Sigma(R)))$[1].

Labeled Markov Processes. A labeled Markov process (LMP) [14], [15] 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 $a \in L$, $\tau_a : S \times \Sigma \to [0,1]$ is a transition probability. By Lemma 1, we can say that $(S, \Sigma, \{\tau_a \mid a \in L\})$ is an LMP if every $\tau_a : S \to \Delta(S)$ is measurable.

In [14], [15], a notion of behavioral equivalence similar to Larsen-Skou probabilistic bisimulation [21] is introduced.

Definition 1: $R \subseteq S \times S$ is a *state bisimulation on LMP* $(S, \Sigma, \{\tau_a \mid a \in L\})$ if it is symmetric[2] and for all $s, t \in S$, $a \in L$, sRt implies that $\tau_a(s) R \tau_a(t)$.

This definition is pointwise and not "eventwise" as one should expect in a measure-theoretic realm, besides R has no measurability restriction. In [11] a measure-theoretic aware notion of behavioral equivalence is introduced.

Definition 2: An *event bisimulation on a LMP* $(S, \Sigma, \{\tau_a \mid a \in L\})$ is a sub-σ-algebra Λ of Σ s.t. $(S, \Lambda, \{\tau_a \mid a \in L\})$ is a LMP.

[11] shows that R is state bisimulation iff $\Sigma(R)$ is an event bisimulation. This is an important result that leads to prove that the largest state bisimulation is also an event bisimulation (see Theorem 5 below).

3. Nondeterministic Labeled Markov Processes

In this section we extend the LMP model adding internal nondeterminism. That is, we allow that different but equally labeled transition probabilities leave out the same state. We provide event and state bisimulations for this model, show the relation to LMPs and the relation to earlier definitions of bisimulation on nondeterministic and continuous probabilistic transition systems.

The model. There have been several attempts to define nondeterministic continuous probabilistic transition systems and all of them are straightforward extensions of (simpler) discrete versions. There are two fundamental differences in our new model. The first one is that the *nondeterministic transition function* T_a now maps states to *measurable sets of* probability measures rather than arbitrary sets as previous approaches.

1. Prop. 4(v) appears in [11] unnecessarily requiring that R is a state bisimulation.

2. The requirement of symmetry is needed otherwise $\Sigma(R)$ may not be a σ-algebra

This is motivated by the fact that later on the nondeterminism has to be resolved using schedulers. If we allowed the target set of states to be an arbitrary subset, (as some continuous ones [5], [8], [12]), the system as a whole could suffer from non-measurability issues and therefore it could not be quantified. (Rigorously speaking, labels should also be provided with a σ-algebra, but we omit it here since it is not needed.) The second difference is inspired by the definition of LMP and Lemma 1 (see also the alternative definition of LMP above): 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 3: A *nondeterministic 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 $a \in L$, $T_a : S \to \Delta(\Sigma)$ 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 to forthcoming definitions and theorems.

Definition 4: $H(\Delta(\Sigma))$ is the minimal σ-algebra containing all sets $H_\xi \doteq \{\Theta \in \Delta(\Sigma) \mid \Theta \cap \xi \neq \emptyset\}$ with $\xi \in \Delta(\Sigma)$.

This construction is similar to that of the Effros-Borel spaces [20] and resembles the so-called hit-and-miss topologies [23]. 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 s such that, through label a, "hit" the set of measures ξ (i.e., $T_a(s) \cap \xi \neq \emptyset$). This forms the basis to existentially quantify over the nondeterminism, and it is fundamental for the behavioral equivalence and the logic.

The next two examples (inspired by an example in [7]) show why T_a is required to map into measurable sets and to be measurable. For these examples we fix the state space and σ-algebra in the real unit interval with the standard Borel σ-algebra.

Example 1: Let $\mathcal{V} = \{\delta_q \mid q \in V\}$, where V is the non-measurable Vitali set in $[0, 1]$. 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) [25], [27], whatever its definition is, would require to assign probabilities to all possible choices. This amounts to measure the nonmeasurable set $T_a(s)$. This is why we require that T_a maps into measurable sets.

Example 2: Let $T_a(s) = \{\mu\}$ for a fixed measure μ, and let $T_b(s) = \textbf{if } (s \in V) \textbf{ then } \{\delta_1\} \textbf{ else } \emptyset$, for every $s \in [0, 1]$, with V being a Vitali set. Notice that both $T_a(s)$ and $T_b(s)$ are measurable sets for every $s \in [0, 1]$. Supposing that there is a scheduler that chooses to first do a and then b starting at some state s, the probability of such set of executions cannot be measured, as 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 maps into measurable sets (see Sec. 4).

NLMPs as a generalization of LMPs. Notice that a LMP is a NLMP without internal nondeterminism. That is, a NLMP in which $T_a(s)$ is a singleton for all $a \in L$ and $s \in S$, is a LMP. In fact, a LMP can be encoded as a NLMP by taking $T_a(s) = \{\tau_a(s)\}$. (We formally prove this in Prop. 5.) As a consequence it is necessary that singletons $\{\mu\}$ are measurable in $\Delta(\Sigma)$ for the NLMP to be well defined. The following lemma gives sufficient conditions to ensure that all singletons are measurable in $\Delta(\Sigma)$.

Lemma 2: Let $\mathcal{G}$ be a denumerable π-system on S (i.e., a denumerable subset of 2^S containing S and closed under finite intersection). Then, for all $\mu \in \Delta(S)$, $\{\mu\} \in \Delta(\sigma(\mathcal{G}))$.

Proof: It is sufficient to prove that the set

$$\cap\{\Delta^{>q_i}(Q_i) \mid Q_i \in \mathcal{G}, q_i \in \mathbb{Q} \cap [0,1], q_i < \mu(Q_i)\} \cap$$
$$\cap\{\Delta^{<q_i}(Q_i) \mid Q_i \in \mathcal{G}, q_i \in \mathbb{Q} \cap [0,1], \mu(Q_i) < q_i\},$$

which is a denumerable intersection, is equal to the singleton $\{\mu\}$. By construction μ is in the intersection. Take μ' s.t. $\mu \neq \mu'$. By a classical theorem of extension of a measure [2, Theorem 3.3], there must be a $Q_i \in \mathcal{G}$ such that $\mu(Q_i) \neq \mu'(Q_i)$. If $\mu(Q_i) > \mu'(Q_i)$ then μ' does not belong to the first intersection; if $\mu(Q_i) < \mu'(Q_i)$, μ' does not belong to the second one. □

In other words, we can guarantee that singletons are measurable in Giry's construction if the underlying σ-algebra is countably generated. Note that Lemma 2 gives also sufficient conditions to define NLMPs with *finite and denumerable nondeterminism.*

Notice also that asking for measurable singletons in $\Delta(\Sigma)$ does not trivialize Σ (in the sense that $\Sigma = 2^S$). A nontrivial example in which Lemma 2 holds is the standard Borel σ-algebra in $\mathbb{R}$. A less obvious example is as follows. Let the σ-algebra **Q-coQ** $\doteq 2^{\mathbb{Q}} \cup \{\mathbb{R} \setminus \mathbb{Q} \mid Q \in 2^{\mathbb{Q}}\}$. Notice that **Q-coQ** cannot separate one irrational from another (let alone asking for all singletons being measurable). Nevertheless, as it is generated by the denumerable π-system $\{\{q\} \mid q \in \mathbb{Q}\} \cup \{\emptyset\}$, it is under the conditions of Lemma 2 and hence for every measure μ on it, $\{\mu\}$ is measurable on $\Delta(\textbf{Q-coQ})$.

The formal connection between NLMP and LMP is an immediate consequence of the next proposition.

Proposition 5: Let $T_a(s) = \{\tau_a(s)\}$ for all $s \in S$ and let Σ be a σ-algebra on S. Then $\tau_a : S \to \Delta(S)$ is measurable iff $T_a : S \to \Delta(\Sigma)$ is measurable.

Proof: Let $\xi \in \Delta(\Sigma)$. Note that $T_a(s) \in H_\xi$ iff $\{\tau_a(s)\} \cap \xi \neq \emptyset$ iff $\tau_a(s) \in \xi$. Then $T_a^{-1}(H_\xi) = \tau_a^{-1}(\xi)$. Therefore τ_a is measurable whenever T_a is measurable. For the converse, we have that $T_a^{-1}(H_\xi)$ is measurable for all generators H_ξ. As a consequence T_a is measurable in general [1]. □

The bisimulations. Event bisimulation in NLMP is defined exactly in the same way as for LMP: an event bisimulation is a sub-σ-algebra that, together with the same set of states and transition of the original NLMP, makes a new NLMP.

Definition 5: An *event bisimulation on a NLMP* $(S, \Sigma, \{T_a \mid a \in L\})$ is a sub-σ-algebra Λ of Σ s.t. $(S, \Lambda, \{T_a \mid a \in L\})$ is a NLMP, that is, T_a is Λ-measurable for each $a \in L$.

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 Λ s.t. $R = \mathcal{R}(\Lambda)$. We remark that, by Prop. 5, an event bisimulation on a LMP is also an event bisimulation on the encoding NLMP and vice-versa.

The definition of state bisimulation is less standard. Following the original definition of Milner [22] (which was lifted to discrete probabilistic models by Larsen and Skou [21]), a traditional definition of bisimulation (see Def. 7) verifies that, whenever sRt, every measure on $T_a(s)$ has a corresponding one (modulo R) in $T_a(t)$. Rather than looking pointwise at probability measures, our definition follows the idea of Def. 4 and verifies that both $T_a(s)$ and $T_a(t)$ *hit* the same measurable sets of measures.

Definition 6: A relation $R \subseteq S \times S$ is a *state bisimulation* if it is symmetric and for all $a \in L$, sRt implies $\forall \xi \in \Delta(\Sigma(R)) : T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset$.

The following property, which also holds in LMPs, states the fundamental relation between state bisimulation and event bisimulation.

Lemma 3: Provided R is symmetric, R is a state bisimulation iff $\Sigma(R)$ is an event bisimulation.

Proof: By Def. 5, $\Sigma(R)$ is an event bisimulation iff T_a is $\Sigma(R)$-measurable. Since T_a is Σ-measurable, it suffices to prove that $T_a^{-1}(H_\xi)$ is R-closed for all labels $a \in L$ and generators H_ξ, $\xi \in \Delta(\Sigma(R))$.

R-closed$(T_a^{-1}(H_\xi))$

iff (R is symmetric)

$sRt \Rightarrow \big(s \in T_a^{-1}(H_\xi) \Leftrightarrow t \in T_a^{-1}(H_\xi)\big)$

iff (Def. inverse function)

$sRt \Rightarrow (T_a(s) \in H_\xi \Leftrightarrow T_a(t) \in H_\xi)$

iff (Def. of H_ξ)

$sRt \Rightarrow (T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset)\,.$

The last statement is the definition of state bisimulation. □

The following results are consequences of Prop. 4 and, for the case of Lemma 4.3, Lemma 3 and the fact that $\mathcal{R}(\Lambda)$ is an equivalence relation. The proofs are the same as the proofs of similar results for LMP in [11].

Lemma 4: Let R be a state bisimulation. Then:

1) R is an event bisimulation iff $R = \mathcal{R}(\Sigma(R))$.
2) If the equivalence classes of R are in Σ, R is an event bisimulation.
3) $\mathcal{R}(\Sigma(R))$ is both a state bisimulation and an event bisimulation.

Let $\sim = \bigcup\{R \mid R \text{ is a state bisimulation}\}$. In the following we show that $\sim$ is also a state bisimulation and hence the largest one. Moreover, we show that $\sim$ is also an event bisimulation and, as a consequence, an equivalence relation.

Theorem 5: $\sim$ is *(i)* the largest state bisimulation, *(ii)* an event bisimulation, and *(iii)* an equivalence relation.

Proof: (i) Take $s, t \in S$ s.t. $s \sim t$. Then there is a state bisimulation R with sRt. Take a measurable set $\xi \in \Delta(\Sigma(\sim))$. Since $R \subseteq \sim$, by Prop. 2, $\Delta(\Sigma(R)) \supseteq \Delta(\Sigma(\sim))$. Hence $\xi \in \Delta(\Sigma(R))$ and by Def. 6, $T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset$ which prove that $\sim$ is a state bisimulation. By definition, it is the largest one.
(ii) Because $\sim$ is a state bisimulation, $\mathcal{R}(\Sigma(\sim))$ is a state bisimulation and an event bisimulation (Lemma 4.3). Since $\sim$ is the largest bisimulation then $\sim = \mathcal{R}(\Sigma(\sim))$ and hence it is an event bisimulation.
(iii) By definition, every event bisimulation is an equivalence relation. □

A traditional view to bisimulation. We have already stated that our definition of state bisimulation differs from a more traditional view such as those in [4], [5], [12], [13], [26]. These definitions closely resemble Larsen & Skou's definition [21]. (The only difference is that two measures are considered equivalent if they agree in every *measurable union of* equivalence classes induced by the relation.) In the following, we give a more "modern" variant of this definition.

Definition 7: A relation R is a *traditional bisimulation* if it is symmetric and for all $a \in L$, sRt implies $T_a(s)\ R\ T_a(t)$.

Note that R is lifted this time to sets as is usual: $T_a(s)\ R\ T_a(t)$ if for all $\mu \in T_a(s)$, there is $\mu' \in T_a(t)$

s.t. $\mu R \mu'$ and vice-versa. (Had we explicitly written this definition on Def. 7, it would have resembled traditional definitions.)

In the following we discuss the relation between state bisimulation and traditional bisimulation. Lemma 6 states that every traditional bisimulation is a state bisimulation. Theorems 7 and 8 give sufficient conditions to strengthen Lemma 6 so that the converse also holds.

Lemma 6: If R is a traditional bisimulation, then R is a state bisimulation.

Proof: Let sRt and $\xi \in \Delta(\Sigma(R))$. If $T_a(s) \cap \xi \neq \emptyset$, then there is $\mu \in T_a(s)$ s.t. $\mu \in \xi$. Since R is a traditional bisimulation, $T_a(s)\ R\ T_a(t)$, i.e., there is $\mu' \in T_a(t)$ s.t. $\mu R \mu'$. By Prop. 3 R-closed(ξ), so $\mu' \in \xi$, and hence $T_a(t) \cap \xi \neq \emptyset$ as required. The other implication follows by symmetry. □

In the following we give two sufficient conditions so that a state bisimulation is also a traditional bisimulation. The first condition focuses on the NLMP. It requires the NLMP to be *image denumerable*.

Definition 8: A NLMP $(S, \Sigma, \{T_a \mid a \in L\})$ is *image denumerable* iff for all $a \in L, s \in S$, $T_a(s)$ is denumerable.

Theorem 7: Let $(S, \Sigma, \{T_a \mid a \in L\})$ be an image denumerable NLMP. Then R is a traditional bisimulation iff it is a state bisimulation.

Proof: The left-to-right implication is Lemma 6. For the other implication we proceed as follows.

Let sRt and for all $\xi \in \Delta(\Sigma(R))$, $T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset$. Suppose towards a contradiction that $T_a(s) \not R\ T_a(t)$, i.e. $\exists \mu \in T_a(s), \forall \mu_i' \in T_a(t) : \exists Q_i \in \Sigma(R) : \mu(Q_i) \bowtie_i \mu_i'(Q_i)$, where $\bowtie_i \in \{>, <\}$ and $i \in \mathbb{N}$ (the NLMP is image denumerable). By density of the rationals, there are $\{q_i\}_i \subseteq \mathbb{Q} \cap [0, 1]$ such that $\mu(Q_i) \bowtie_i q_i \bowtie_i \mu_i'(Q_i)$. Then $\mu \in \Delta^{\bowtie_i q_i}(Q_i) \not\ni \mu_i'$. Let $\xi \doteq \cap_i \Delta^{\bowtie_i q_i}(Q_i)$. This set is measurable, moreover, since every $Q_i \in \Sigma(R)$, so $\xi \in \Delta(\Sigma(R))$. Then $\mu \in T_a(s) \cap \xi$, but $T_a(t) \cap \xi = \emptyset$ hence contradicting the assumption. □

After reading the proof, it should be clear that we can relax the sufficient condition to require that the partition $T_a(s)/R$ is denumerable for each state s and label a instead of image denumerability.

Observe that a state bisimulation on a LMP is a traditional bisimulation on the encoding NLMP and vice-versa since $\{\tau_a(s)\} = T_a(s)\ R\ T_a(t) = \{\tau_a(t)\}$ iff $\tau_a(s)\ R\ \tau_a(t)$. As a consequence of Lemma 6 and Theorem 7 (a deterministic NLMP is image denumerable!), we conclude that a state bisimulation on a LMP is a state bisimulation on the encoding NLMP and vice-versa.

The second sufficient condition looks at the σ-algebra $\Sigma(R)$ induced by the state bisimulation R. It turns out that if $\Sigma(R)$ is generated by a denumerable π-system, R is also a traditional bisimulation.

Theorem 8: Let R be a symmetric relation such that $\Sigma(R)$ is generated by a denumerable set $\mathcal{G}$. Then R is a traditional bisimulation iff it is a state bisimulation.

Proof: As before, the left-to-right implication is Lemma 6. For the other implication we proceed as follows. Suppose towards a contradiction that sRt and $T_a(s) \not R\ T_a(t)$, i.e. $\exists \mu \in T_a(s), \forall \mu' \in T_a(t) : \mu \not R\ \mu'$. By [2, Theorem 3.3], this implies that there exists $Q_i \in \pi(\mathcal{G})$ s.t. $\mu(Q_i) \neq \mu'(Q_i)$ with $i \in \mathbb{N}$. (Notice that $\pi(\mathcal{G})$, the π-system generated by $\mathcal{G}$, is also denumerable and generates $\Sigma(R)$.) The rest of the proof is as in Theorem 7. □

4. A Logic for Bisimulation on NLMP

The logic we present below is based on the logic given by Parma and Segala [24]. The main difference is that we consider two kind of formulas: one that is interpreted on states, and another that is interpreted on measures. The syntax is as follows,

$$\begin{array}{rcl} \phi & \equiv & \top \mid \phi_1 \wedge \phi_2 \mid \langle a \rangle \psi \\ \psi & \equiv & \bigvee_i \psi_i \mid \neg \psi \mid [\phi]_{\geq q} \end{array}$$

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

Semantics is defined with respect to a NLMP (S, Σ, T). Formulas in $\mathcal{L}$ are interpreted as sets of states in which they become true, and formulas in $\mathcal{L}_\Delta$ are interpreted as sets of measures on states as follows,

$$\begin{array}{ll} [\![\top]\!] = S & [\![\bigvee_i \psi_i]\!] = \bigcup_i [\![\psi_i]\!] \\ [\![\phi_1 \wedge \phi_2]\!] = [\![\phi_1]\!] \cap [\![\phi_2]\!] & [\![\neg \psi]\!] = [\![\psi]\!]^c \\ [\![\langle a \rangle \psi]\!] = T_a^{-1}(H_{[\![\psi]\!]}) & [\![[\phi]_{\geq q}]\!] = \Delta^{\geq q}([\![\phi]\!]) \end{array}$$

In particular, notice that $\langle a \rangle \psi$ is valid in a state s whenever there is some measure $\mu \in T_a(s)$ that makes ψ valid, and that $[\phi]_{\geq q}$ is valid in a measure μ whenever $\mu([\![\phi]\!]) \geq q$. As a consequence, we need that sets $[\![\phi]\!]$ and $[\![\psi]\!]$ are measurable in Σ and $\Delta(\Sigma)$, respectively. Indeed, this follows straightforwardly by induction on the construction of the formula after observing that all operations involved in the definition of the semantics preserve measurability (in particular T_a is a measurable function). For the rest of the section, fix $[\![\mathcal{L}]\!] = \{[\![\phi]\!] \mid \phi \in \mathcal{L}\}$ and $[\![\mathcal{L}_\Delta]\!] = \{[\![\psi]\!] \mid \psi \in \mathcal{L}_\Delta\}$.

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

We show that $\mathcal{L}$ characterizes event bisimulation. This is an immediate consequence of the fact that $\sigma(\llbracket\mathcal{L}\rrbracket)$, the σ-algebra generated by the logic $\mathcal{L}$, is the smallest event bisimulation, which is what we aim to prove in this part of the section. The proof strategy resembles that of [11, Sec. 5] but it is properly tailored to our two level logic. Moreover, such a separation allowed us to find an alternative to Dynkin's theorem (used in [11]).

We extend the definition of $\Delta(\mathcal{C})$ to any arbitrary set $\mathcal{C} \subseteq \Sigma$ by taking $\Delta(\mathcal{C})$ to be the σ-algebra generated by $\Delta^{\geq p}(Q) \doteq \{\nu \mid \nu(Q) \geq p\}$, with $Q \in \mathcal{C}$ and $p \in [0,1]$. From now on we write $\sigma(\mathcal{L})$, $\Delta(\mathcal{L})$ and $\mathcal{R}(\mathcal{L})$ instead of $\sigma(\llbracket\mathcal{L}\rrbracket)$, $\Delta(\llbracket\mathcal{L}\rrbracket)$ and $\mathcal{R}(\llbracket\mathcal{L}\rrbracket)$, respectively.

The concept of *stable family* of measurable sets is crucial to the proof of Theorem 13.

Definition 9: Given a NLMP (S,Σ,T), the family $\mathcal{C} \subseteq \Sigma$ is *stable* for (S,Σ,T) if for all $a \in L$ and $\xi \in \Delta(\mathcal{C})$, $T_a^{-1}(H_\xi) \in \mathcal{C}$.

Notice that $\mathcal{C}$ is an event bisimulation iff it is a stable σ-algebra.

The key point of the proof is to show that $\llbracket\mathcal{L}\rrbracket$ is the smallest stable π-system, which is stated in Lemma 10. The next lemma is auxiliary to Lemma 10.

Lemma 9: $\llbracket\mathcal{L}_\Delta\rrbracket = \Delta(\mathcal{L})$

Proof: $\llbracket\mathcal{L}_\Delta\rrbracket$ is a σ-algebra since: *(i)* $\Delta(S) = \llbracket[\top]_{\geq 1}\rrbracket \in \llbracket\mathcal{L}_\Delta\rrbracket$; *(ii)* for $\xi_i \in \llbracket\mathcal{L}_\Delta\rrbracket$ there are $\psi_i \in \mathcal{L}_\Delta$ s.t. $\xi_i = \llbracket\psi_i\rrbracket$, and hence $\bigcup_i \xi_i = \bigcup_i\llbracket\psi_i\rrbracket = \llbracket\bigvee_i \psi_i\rrbracket \in \llbracket\mathcal{L}_\Delta\rrbracket$; and *(iii)* for $\xi \in \llbracket\mathcal{L}_\Delta\rrbracket$ there is $\psi \in \mathcal{L}_\Delta$ s.t. $\xi = \llbracket\psi\rrbracket$, and hence $\xi^c = \llbracket\psi\rrbracket^c = \llbracket\neg\psi\rrbracket \in \llbracket\mathcal{L}_\Delta\rrbracket$. Moreover, since $\llbracket[\phi]_{\geq p}\rrbracket = \Delta^{\geq p}(\llbracket\phi\rrbracket)$, every generator set of $\Delta(\mathcal{L})$ is in $\llbracket\mathcal{L}_\Delta\rrbracket$ and hence $\Delta(\mathcal{L}) \subseteq \llbracket\mathcal{L}_\Delta\rrbracket$.

Finally, it can be proven by induction on the depth of the formula that $\llbracket\mathcal{L}_\Delta\rrbracket \subseteq \mathcal{C}$ for any σ-algebra $\mathcal{C}$ containing all sets $\llbracket[\phi]_{\geq p}\rrbracket = \Delta^{\geq p}(\llbracket\phi\rrbracket)$ for $p \in [0,1]$ and $\phi \in \mathcal{L}$. Then $\llbracket\mathcal{L}_\Delta\rrbracket$ is the smallest σ-algebra containing all generator sets of $\Delta(\mathcal{L})$. Therefore $\llbracket\mathcal{L}_\Delta\rrbracket = \Delta(\mathcal{L})$. □

Lemma 10: $\llbracket\mathcal{L}\rrbracket$ is the smallest stable π-system for (S,Σ,T).

Proof: $\llbracket\mathcal{L}\rrbracket$ is a π-system since: *(i)* $S = \llbracket\top\rrbracket \in \llbracket\mathcal{L}\rrbracket$ and *(ii)* for $Q_1, Q_2 \in \llbracket\mathcal{L}\rrbracket$ there are $\phi_1, \phi_2 \in \mathcal{L}$ s.t. $Q_1 = \llbracket\phi_1\rrbracket$ and $Q_2 = \llbracket\phi_2\rrbracket$, and hence $Q_1 \cap Q_2 = \llbracket\phi_1\rrbracket \cap \llbracket\phi_2\rrbracket = \llbracket\phi_1 \cap \phi_2\rrbracket \in \llbracket\mathcal{L}\rrbracket$.

For stability, let $\xi \in \Delta(\mathcal{L})$. By Lemma 9, there is $\psi \in \mathcal{L}_\Delta$ s.t. $\llbracket\psi\rrbracket = \xi$. Then $T_a^{-1}(H_\xi) = T_a^{-1}(H_{\llbracket\psi\rrbracket}) = \llbracket\langle a\rangle\psi\rrbracket \in \llbracket\mathcal{L}\rrbracket$.

Let $\mathcal{C}$ be another stable π-system for (S,Σ,T). By induction in the depth of the formula we show simultaneously that $\mathcal{C} \supseteq \llbracket\mathcal{L}\rrbracket$ and $\Delta(\mathcal{C}) \supseteq \Delta(\mathcal{L})$. First notice that $\llbracket\top\rrbracket = S \in \mathcal{C}$ since $\mathcal{C}$ is a π-system. Now, suppose inductively that $\llbracket\phi\rrbracket, \llbracket\phi_1\rrbracket, \llbracket\phi_2\rrbracket \in \mathcal{C}$ and $\llbracket\psi\rrbracket, \llbracket\psi_i\rrbracket \in \Delta(\mathcal{C})$ for $i \geq 0$. Then: *(i)* $\llbracket\phi_1 \wedge \phi_2\rrbracket = \llbracket\phi_1\rrbracket \cap \llbracket\phi_2\rrbracket \in \mathcal{C}$, because $\mathcal{C}$ is a π-system; *(ii)* $\llbracket\langle a\rangle\psi\rrbracket = T_a^{-1}(H_{\llbracket\psi\rrbracket}) \in \mathcal{C}$, because $\mathcal{C}$ is stable; *(iii)* $\llbracket\bigvee_i \psi_i\rrbracket = \bigcup_i\llbracket\psi_i\rrbracket \in \Delta(\mathcal{C})$ and *(iv)* $\llbracket\neg\psi\rrbracket = \llbracket\psi\rrbracket^c \in \Delta(\mathcal{C})$ because $\Delta(\mathcal{C})$ is a σ-algebra; and finally, *(v)* $\llbracket[\phi]_{\geq p}\rrbracket = \Delta^{\geq p}(\llbracket\phi\rrbracket) \in \Delta(\mathcal{C})$ by definition of generator set of $\Delta(\mathcal{C})$. □

Lemma 11 is auxiliary to Lemma 12. It is also significantly simpler than its relative in [11, Lemma 5.4]. This is due to our definition of stability and the use of a powerful result of [28].

Lemma 11: If $\mathcal{C}$ is a stable π-system for (S,Σ,T), then $\sigma(\mathcal{C})$ is also stable.

Proof: First notice that $\mathcal{C}$ is stable iff $\{T_a^{-1}(H_\xi) \mid a \in L, \xi \in \Delta(\mathcal{C})\} \subseteq \mathcal{C}$. By [28, Lemma 3.6], $\Delta(\mathcal{C}) = \Delta(\sigma(\mathcal{C}))$. Then $\{T_a^{-1}(H_\xi) \mid a \in L, \xi \in \Delta(\sigma(\mathcal{C}))\} \subseteq \mathcal{C} \subseteq \sigma(\mathcal{C})$, which proves that $\sigma(\mathcal{C})$ is stable. □

The next lemma is central to the proof that $\mathcal{L}$ characterizes event bisimulation, which is then presented in Theorem 13.

Lemma 12: $\sigma(\mathcal{L})$ is the smallest stable σ-algebra included in Σ.

Proof: Let $\mathcal{F}$ be the smallest stable σ-algebra included in Σ. By Lemma 10, $\llbracket\mathcal{L}\rrbracket \subseteq \mathcal{F}$, since $\mathcal{F}$ is a stable π-system. Therefore $\sigma(\mathcal{L}) \subseteq \mathcal{F}$ since $\mathcal{F}$ is also a σ-algebra. For the other inclusion, notice that $\llbracket\mathcal{L}\rrbracket$ is a stable π-system because of Lemma 10. By Lemma 11, $\sigma(\mathcal{L})$ is stable, therefore it contains $\mathcal{F}$. □

Theorem 13: The logic $\mathcal{L}$ completely characterizes event bisimulation.

Proof: Lemma 12 establishes that $\sigma(\mathcal{L})$ is stable, i.e. it is an event bisimulation. Being the smallest, it implies that any other event bisimulation preserves $\mathcal{L}$ formulas. □

A consequence of this and Theorem 5 is that state bisimulation is sound for $\mathcal{L}$, i.e., it preserves the validity of formulas. This is stated in Theorem 15. We first introduce an auxiliary lemma.

Lemma 14: $\mathcal{R}(\mathcal{C}) = \mathcal{R}(\sigma(\mathcal{C}))$.

Proof: We only need to show that $\mathcal{R}(\mathcal{C}) \subseteq \mathcal{R}(\sigma(\mathcal{C}))$ since the other inclusion is obvious. Let $s\ \mathcal{R}(\mathcal{C})\ t$. Notice that $\sigma(\mathcal{C}) = \{Q \in \sigma(\mathcal{C}) \mid s \in Q \Leftrightarrow t \in Q\}$. (It is easy to see that this set is closed by complement and denumerable union and contains $\mathcal{C}$.) From this and definition of $\mathcal{R}(\sigma(\mathcal{C}))$, $s\ \mathcal{R}(\sigma(\mathcal{C}))\ t$ follows. □

Theorem 15: $\sim \subseteq \mathcal{R}(\mathcal{L})$.

Proof: By Theorem 5, $\Sigma(\sim)$ is an event bisimulation and hence a stable σ-algebra. Then $\sigma(\mathcal{L}) \subseteq \Sigma(\sim)$ by Lemma 12. Therefore, using Lemma 14, $\mathcal{R}(\mathcal{L}) = \mathcal{R}(\sigma(\mathcal{L})) \supseteq \mathcal{R}(\Sigma(\sim)) = \sim$. □

Completeness on image finite NLMPs. The rest of the section is devoted to show that the logic completely characterizes (all three) bisimulation on NLMPs with image finite nondeterminism and standing on analytic spaces. In fact, we show completeness of the sublogic of $\mathcal{L}$ defined by:

$$\phi \quad \equiv \quad \top \mid \phi_1 \wedge \phi_2 \mid \langle a \rangle [\,_{\bowtie_i q_i} \phi_i]_{i=1}^{n}$$

where $\bowtie_i \in \{>, <\}$ and $q_i \in \mathbb{Q} \cap [0, 1]$. We define the new modal operation as a shorthand notation: $\langle a \rangle [\,_{\bowtie_i q_i} \phi_i]_{i=1}^{n} \equiv \langle a \rangle \bigwedge_{i=1}^{n} [\phi]_{\bowtie_i q_i}$. Therefore, $[\![\langle a \rangle [\,_{\bowtie_i q_i} \phi_i]_{i=1}^{n}]\!] = T_a^{-1}(H_{\bigcap_{i=1}^{n} \Delta^{\bowtie_i q_i}([\![\phi_i]\!])})$. Let $\mathcal{L}_{\mathrm{f}} \subseteq \mathcal{L}$ denote the set of all formulas defined with the grammar above. Notice that $\mathcal{L}_{\mathrm{f}}$ is a denumerable set whenever the set of labels L is denumerable.

The expression $\langle a \rangle [\,_{\bowtie_i q_i} \phi_i]_{i=1}^{n}$ is like a conjunction of formulas $\langle a \rangle_{\bowtie_i q_i} \phi_i$, but the probabilistic bounds must be satisfied by the *same* nondeterministic transition. Modality $\langle a \rangle_{\bowtie q} \phi$ suffices to characterize bisimulation on LMP [15] but, as we see in the next example, it is not enough for the more general setting of NLMPs.

Example 3: Take the discrete NLMPs depicted below. States s and t are not bisimilar since given a $\mu \in T_a(s)$, 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, $[\![\langle a \rangle_{>q} \phi]\!] = \{w \mid T_a(w) \cap \Delta^{>q}([\![\phi]\!]) \neq \emptyset\}$ will always have s and t together. Observe that negation, denumerable conjunction or disjunction, do not add any distinguishing power (on an image finite setting).

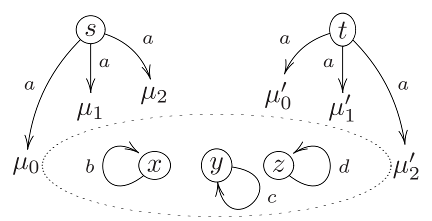

	μ_0	μ_1	μ_2	μ'_0	μ'_1	μ'_2
$\{x\}$	$\frac{1}{3}$	$\frac{2}{3}$	0	$\frac{2}{3}$	$\frac{1}{3}$	0
$\{y\}$	0	$\frac{1}{3}$	$\frac{2}{3}$	0	$\frac{2}{3}$	$\frac{1}{3}$
$\{z\}$	$\frac{2}{3}$	0	$\frac{1}{3}$	$\frac{1}{3}$	0	$\frac{2}{3}$

The essential need for this new modal operator also shows that our σ-algebra $H(\Delta(\Sigma))$ in Def. 4 can not be simplified to $\sigma(\{H_{\Delta^B(Q)} : B \in \mathcal{B}([0, 1]), Q \in \Sigma\})$. States s and t in the example above should be observationally distinguished from each other. Formally, this amounts to say that there must be some label a and some measurable Θ 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 $\{H_{\Delta^B(Q)} : B \in \mathcal{B}([0, 1]), Q \in \Sigma\}$.

Logical characterization of bisimulation is succinctly stated as $s \sim t \Leftrightarrow s\ \mathcal{R}(\mathcal{L}_{\mathrm{f}})\ t$. The left-to-right implication is immediate by Theorem 15. For the converse, we restrict the state space and the branching.

The strategy is to prove that $\mathcal{R}(\mathcal{L}_{\mathrm{f}})$ is a traditional bisimulation, that is, $s\ \mathcal{R}(\mathcal{L}_{\mathrm{f}})\ t$ implies that $\forall \mu \in T_a(s), \exists \mu' \in T_a(t), \mu\ \mathcal{R}(\mathcal{L}_{\mathrm{f}})\ \mu'$; recall this means $\mu(Q) = \mu'(Q)$ for all $Q \in \Sigma(\mathcal{R}(\mathcal{L}_{\mathrm{f}}))$. For analytic spaces this holds if it is valid for the restricted set of $Q \in \Sigma(\mathcal{R}(\mathcal{L}_{\mathrm{f}}))$ such that $Q = [\![\phi]\!]$, for some $\phi \in \mathcal{L}_{\mathrm{f}}$. We first introduce analytic spaces and a result from descriptive set theory that is fundamental for the proof.

Definition 10: A topological space is *Polish* if it is separable (i.e. it contains a countable dense subset) and completely metrizable. A topological space is *analytic* if it is the continuous image of a Polish space. A measurable space is *analytic (standard) Borel* if it is isomorphic to $(X, \sigma(\mathcal{T}))$ where $\mathcal{T}$ is an analytic (Polish) topology on X.

Every standard Borel space is analytic, but the converse is false. The real line with the usual Borel σ-algebra, and more generally, $A^{\mathbb{N}}$ with A a countable discrete space, are standard Borel and therefore, analytic.

The next theorem from [16] essentially shows that in analytic Borel spaces, the R-closed measurable sets are well-behaved when the relation R is defined in terms of a sequence of measurable sets.

Theorem 16: Let (S, Σ) be an analytic Borel space. Let $\mathcal{F} \subseteq \Sigma$ be countable and assume $S \in \mathcal{F}$. Then $\Sigma(\mathcal{R}(\mathcal{F})) = \sigma(\mathcal{F})$.

The following lemma provides a general framework to prove that a logic characterizes bisimulation. In fact we have used it to prove that less expressive logics characterize traditional bisimulation in some restricted NLMPs [9].

Lemma 17: Let (S, Σ, T) be a NLMP with (S, Σ) being an analytic Borel space. Let $\mathcal{L}'$ be a logic s.t. *(i)* $\mathcal{L}'$ contains operators $\top$ and $\wedge$ with the usual semantics; *(ii)* for every formula $\phi \in \mathcal{L}'$, $[\![\phi]\!]$ is Σ-measurable; *(iii)* the set of all formulas in $\mathcal{L}'$ is denumerable; and *(iv)* for every $s\ \mathcal{R}(\mathcal{L}')\ t$ and every $\mu \in T_a(s)$ there exists $\mu' \in T_a(t)$ such that $\forall \phi \in \mathcal{L}', \mu([\![\phi]\!]) = \mu'([\![\phi]\!])$. Then, two logically equivalent states s, t are traditionally bisimilar.

Proof: Let $\mathcal{F} = \{[\![\phi]\!] \mid \phi \in \mathcal{L}'\}$. Because of (i), $[\![\top]\!] = S$ and $[\![\phi_1]\!] \cap [\![\phi_2]\!] = [\![\phi_1 \wedge \phi_2]\!]$. Hence $\mathcal{F}$ forms a π-system. Because of (iv), μ, μ' agree in $\mathcal{F}$ and, by [2, Thm. 3.3], they also agree in $\sigma(\mathcal{F})$. Notice that hypotheses of Theorem 16 are met, i.e., Σ is analytic, $\mathcal{F} \subseteq \Sigma$ is countable (by (ii) and (iii)) such that $S \in \mathcal{F}$ (by (i)), and $\mathcal{R}(\mathcal{L}')$ equals $\mathcal{R}(\mathcal{F})$. Therefore, by Theorem 16, $\sigma(\mathcal{F}) = \Sigma(\mathcal{R}(\mathcal{L}'))$, which implies that μ and μ' agree in $\Sigma(\mathcal{R}(\mathcal{L}'))$. Since $\mathcal{R}(\mathcal{L}')$ is symmetric, $\mathcal{R}(\mathcal{L}')$ is a traditional bisimulation. □

Notice that Lemma 17 holds for any logic fulfilling the hypothesis, in particular it should encode the transfer property of the bisimulation and may not contain negation. We already know that $\mathcal{L}_\mathrm{f}$ has operators $\top$ and $\wedge$, is denumerable, and that each formula is interpreted in a Σ-measurable set. In the following, we show that the transfer property can be encoded by using the modality.

Lemma 18: Let (S, Σ, T) be an image finite NLMP (i.e. $T_a(s)$ is finite for all $a \in L, s \in S$). Then for every pair of states such that $s\,\mathcal{R}(\mathcal{L}_\mathrm{f})\,t$ and $\mu \in T_a(s)$, there is a $\mu' \in T_a(t)$ such that $\forall \phi \in \mathcal{L}_\mathrm{f}, \mu([\![\phi]\!]) = \mu'([\![\phi]\!])$.

Proof: Suppose towards a contradiction that there are s, t with $s\ \mathcal{R}(\mathcal{L}_\mathrm{f})\ t$ and there is a $\mu \in T_a(s)$, such that for all $\mu'_i \in T_a(t)$ there is a formula $\phi_i \in \mathcal{L}_\mathrm{f}$ with $\mu([\![\phi_i]\!]) \neq \mu'_i([\![\phi_i]\!])$. Since $T_a(t)$ is finite, there are at most n different μ'_i. We can choose $\bowtie_i \in \{>, <\}, q_i \in \mathbb{Q} \cap [0, 1]$ accordingly to make $\mu([\![\phi_i]\!]) \bowtie_i q_i \bowtie_i \mu'_i([\![\phi_i]\!])$. Take $\psi = \langle a \rangle [\,{\bowtie_i q_i}\,\phi_i]_{i=1}^{n}$. Then $s \in [\![\psi]\!]$ but $t \notin [\![\psi]\!]$ contradicting $s\,\mathcal{R}(\mathcal{L}_\mathrm{f})\,t$. □

So, finally, we can state the following theorem

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

$$s \sim_t t \Leftrightarrow s \sim t \Leftrightarrow s\,\mathcal{R}(\mathcal{L}_\mathrm{f})\,t$$

where $s \sim_t t$ denotes that there is a traditional bisimulation R, s.t. $s\,R\,t$.

Proof: $s \sim_t t \Rightarrow s \sim t$ (by Theorem 7) $\Rightarrow$ $s\,\mathcal{R}(\mathcal{L})\,t$ (by Theorem 15) $\Rightarrow s\,\mathcal{R}(\mathcal{L}_\mathrm{f})\,t$ (because $\mathcal{L}_\mathrm{f} \subseteq \mathcal{L}$) $\Rightarrow s \sim_t t$ (by Lemmas 17 and 18). □

5. Concluding remarks

In order to define a process theory that permits the verification of compositionally modeled systems against simple (may be nondeterministic) specifications, it is necessary to have available a semantic relation that allows for abstraction such as weak bisimulation. In this setting, internal nondeterminism is crucial.

In this paper we introduced the model of nondeterministic labeled Markov processes that allows for the modeling of continuous probabilistic systems with internal nondeterminism. Contrarily to similar models [4], [5], [7], [12], [13], [15], NLMPs are defined to have a measure theoretic structure. In particular, we require that the transition relation is a measurable function that maps on measurable sets. This was devised so that it is possible to build the rest of the theory (particularly event bisimulation and logic, but also schedulers are definable). We have shown that NLMPs extend naturally LMPs. For the definition of the transition and the development of the whole work, Def. 4 is crucial, as it provides the foundation for dealing with nondeterminism.

As a first step towards the desired process theory, we gave different definitions for the bisimulation.We proposed three possible generalizations of the two bisimulations on LMPs. The event bisimulation responds exactly to the same definition principle both in LMP and NLMP. Instead, the state bisimulation in LMPs generalizes to NLMPs as state bisimulation and as traditional bisimulation. We know that traditional bisimulation is finer than state bisimulation and, in Theorems 7 and 8, we gave sufficient conditions under which they agree. However, we do not know if they agree in general. Notice that the proofs of these theorems lie on singling out a particular distribution through a *denumerable* intersection of generator sets. Because of this observation, we are considering to restrict to standard Borel spaces to better understand the relation between the two bisimulations.

We also gave a logical characterization of event bisimulation (Theorem 13). Such logic ($\mathcal{L}$) can be seen as a revision of the one introduced by [24] in a discrete probabilistic setting. Formulas in our setting belong to two different classes: state formulas and measure formulas. Notice that negation and infinitary (but denumerable) disjunction (or conjunction) is only present on the second class, meaning that the complexity of the model lies precisely on the internal nondeterminism.

A consequence of the characterization is that the logic is sound for state and traditional bisimulations (Theorem 15). We do not have any evidence that suggests that logical equivalence (and hence event bisimulation) agrees or disagrees with state or traditional bisimulation in general. However, for the restricted case of image finite NLMPs running on analytic Borel spaces, all equivalences coincide (Theorem 19). Notice that the logic we used to show such equivalence is in fact a sublogic of $\mathcal{L}$ which has already appeared in a preliminary work [9].

In case that the bisimulations turn out not to be equivalent, the natural definition of the logic $\mathcal{L}$ suggests that event bisimulation is the most appropriate definition of all, provided one accepts that transition functions should be indeed measurable on the σ-algebra $H(\Delta(\Sigma))$.

Notice that the conditions of Lemma 17 also points to a possible restriction to standard Borel spaces, a setting in which the three bisimulation may agree. Confining to standard Borel spaces is not as restricting as it seems since most natural problems arise in this setting. For example, we have shown elsewhere that the underlying semantics of stochastic automata [12] in terms of NLMP meets most of the restrictions required in this article: it runs on standard Borel spaces and it is image finite. We recall that stochastic automata and similar models are used to give semantics to stochastic process algebras and specification languages [3]–[5], [12], [13, etc.] which, in turn, are used to model dynamic systems. Moreover, LMP-like models restricted to standard Borel spaces have been studied [17].

At the moment, we are busy on the study of schedulers for NLMPs and probabilistic trace semantics. This will allow us to contrast the two local behavioral equivalences, state and traditional bisimulation. It is expected that at least one of them implies a global behavioral equivalence, like probabilistic trace equality. Schedulers would also let us define probabilistic weak transitions and their related bisimulations.

Acknowledgments:. We thank Ignacio Viglizzo for fruitful discussions. In particular he pointed out the connection of Def. 4 to hit-and-miss topologies.

References

[1] R. Ash and C. Doléans-Dade, *Probability & Measure Theory*. Academic Press, 2000.

[2] P. Billingsley, *Probability and Measure*. Wiley-Interscience, 1995.

[3] H. Bohnenkamp, P. D'Argenio, H. Hermanns, and J.-P. Katoen, "MoDeST: A compositional modeling formalism for real-time and stochastic systems," *IEEE Trans. Softw. Eng.*, vol. 32, no. 10, pp. 812–830, 2006.

[4] M. Bravetti, "Specification and analysis of stochastic real-time systems," Ph.D. dissertation, Università di Bologna, Padova, Venezia, Feb. 2002.

[5] M. Bravetti and P. D'Argenio, "Tutte le algebre insieme: Concepts, discussions and relations of stochastic process algebras with general distributions," in *Validation of Stochastic Systems*, ser. LNCS, vol. 2925. Springer, 2004, pp. 44–88.

[6] F. v. Breugel, "The metric monad for probabilistic nondeterminism," 2005, draft.

[7] S. Cattani, "Trace-based process algebras for real-time probabilistic systems," Ph.D. dissertation, University of Birmingham, 2005.

[8] S. Cattani, R. Segala, M. Kwiatkowska, and G. Norman, "Stochastic transition systems for continuous state spaces and non-determinism." in *Procs. FoSSaCS 2005*, ser. LNCS, vol. 3441. Springer, 2005, pp. 125–139.

[9] P. Celayes, "Procesos de Markov etiquetados sobre espacios de Borel estándar," Master's thesis, FaMAF, Universidad Nacional de Córdoba, 2006.

[10] E. Clarke, O. Grumberg, and D. Peled, *Model Checking*. MIT Press, 1999.

[11] V. Danos, J. Desharnais, F. Laviolette, and P. Panangaden, "Bisimulation and cocongruence for probabilistic systems," *Inf. & Comp.*, vol. 204, pp. 503–523, 2006.

[12] P. D'Argenio, "Algebras and automata for timed and stochastic systems," Ph.D. dissertation, Department of Computer Science, University of Twente, 1999.

[13] P. D'Argenio and J.-P. Katoen, "A theory of stochastic systems, Part I: Stochastic automata, *and* Part II: Process algebra," *Inf. & Comp.*, vol. 203, no. 1, pp. 1–38, 39–74, 2005.

[14] J. Desharnais, "Labeled Markov process," Ph.D. dissertation, McGill University, 1999.

[15] J. Desharnais, A. Edalat, and P. Panangaden, "Bisimulation for labelled Markov processes," *Inf. & Comp.*, vol. 179, no. 2, pp. 163–193, 2002.

[16] J. Desharnais and P. Panangaden, "Continuous stochastic logic characterizes bisimulation of continuous-time markov processes," *J. Log. Algebr. Program.*, vol. 56, no. 1-2, pp. 99–115, 2003.

[17] E.-E. Doberkat, *Stochastic relations. Foundations for Markov transition systems.*, ser. Studies in Informatics Series. Chapman & Hall/CRC, 2007.

[18] M. Giry, "A categorical approach to probability theory," in *Categorical Aspects of Topology and Analysis*, ser. LNM, vol. 915. Springer, 1981, pp. 68–85.

[19] R. v. Glabbeek, "The linear time–branching time spectrum I. The semantics of concrete, sequential processes," in *Handbook of Process Algebra*, J. A. Bergstra, A. Ponse, and S. A. Smolka, Eds. North-Holland, 2001, pp. 3–99.

[20] A. Kechris, *Classical Descriptive Set Theory*, ser. Graduate Texts in Mathematics. Springer, 1995, vol. 156.

[21] K. Larsen and A. Skou, "Bisimulation through probabilistic testing," *Inf. & Comp.*, vol. 94, no. 1, pp. 1–28, 1991.

[22] R. Milner, *Communication and Concurrency*. Prentice Hall, 1989.

[23] S. Naimpally, "What is a hit-and-miss topology?" *Topological Comment.*, vol. 8, no. 1, 2003.

[24] A. Parma and R. Segala, "Logical characterizations of bisimulations for discrete probabilistic systems," in *FoSSaCS 07*, ser. LNCS 4423. Springer, 2007, pp. 287–301.

[25] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, 1994.

[26] B. Strulo, "Process algebra for discrete event simulation," Ph.D. dissertation, Department of Computing, Imperial College, University of London, 1993.

[27] M. Y. Vardi, "Automatic verification of probabilistic concurrent finite-state programs," in *26th FOCS*. IEEE, 1985, pp. 327–338.

[28] I. Viglizzo, "Coalgebras on measurable spaces," Ph.D. dissertation, Department of Mathematics, Indiana University, 2005.