Math. Struct. in Comp. Science (2012), *vol.* 22, *pp.* 43–68. © Cambridge University Press 2011
doi:10.1017/S0960129511000454 First published online 26 September 2011

Bisimulations for non-deterministic labelled Markov processes[†]

PEDRO R. D'ARGENIO[‡], PEDRO SÁNCHEZ TERRAF[§] and NICOLÁS WOLOVICK[¶]

[‡]*FaMAF, Universidad Nacional de Córdoba and CONICET,*
Medina Allende s/n (Ciudad Universitaria), X5000HUA – Córdoba, Argentina
Email: dargenio@famaf.unc.edu.ar
[§]*FaMAF, Universidad Nacional de Córdoba, CONICET and CIEM,*
Medina Allende s/n (Ciudad Universitaria), X5000HUA – Córdoba, Argentina
Email: sterraf@famaf.unc.edu.ar
[¶]*FaMAF, Universidad Nacional de Córdoba,*
Medina Allende s/n (Ciudad Universitaria), X5000HUA – Córdoba, Argentina
Email: nicolasw@famaf.unc.edu.ar

Received 28 May 2010; revised 25 April 2011

We extend the theory of labelled Markov processes to include *internal* non-determinism, which is a fundamental concept for the further development of a process theory with abstraction on non-deterministic continuous probabilistic systems. We define *non-deterministic labelled Markov processes (NLMP)* and provide three definitions of bisimulations: a bisimulation following a traditional characterisation; a *state*-based bisimulation tailored to our 'measurable' non-determinism; and an *event*-based bisimulation. We show the relations between them, including the fact that the largest state bisimulation is also an event bisimulation. We also introduce a variation of the Hennessy–Milner logic that characterises event bisimulation and is sound with respect to the other bisimulations for an arbitrary NLMP. This logic, however, is infinitary as it contains a denumerable $\vee$. We then introduce a finitary sublogic that characterises all bisimulations for an image finite NLMP whose underlying measure space is also analytic. Hence, in this setting, all the notions of bisimulation we consider turn out to be equal. Finally, we show that all these bisimulation notions are different in the general case. The counterexamples that separate them turn out to be *non-probabilistic* NLMPs.

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 (Danos *et al.* 2006). Many formal frameworks have been defined to study them from a process theory or process algebra perspective (Strulo 1993; Segala 1995; Desharnais 1999; D'Argenio 1999; Bravetti 2002; Desharnais *et al.* 2002; Bravetti and D'Argenio 2004; D'Argenio and Katoen 2005; Cattani 2005; Cattani *et al.* 2005; Danos *et al.* 2006).

[†] Supported by ANPCyT PICT 26135, ANPCyT PICT-PAE 2272, SeCyT-UNC and CONICET.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

A prominent and extensive topic in this area is built on the so-called labelled Markov processes (LMP) (Desharnais 1999; Desharnais *et al.* 2002). This is because it has solid and well-understood mathematical foundations. An LMP allows many transition probability functions (or Markov kernels) leaving each state (instead of just one, as in normal 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. Hence, this model *does not* consider *internal* non-determinism. From the process algebra point of view, this is a significant drawback for the theory since internal non-determinism immediately arises in the analysis of systems, for example, because of abstracting internal activity, such as in weak bisimulation (Milner 1989), or because of state abstraction techniques, such as in model checking (Clarke *et al.* 1999).

Many other papers have defined variants of continuous Markov processes that include internal non-determinism, which are mainly used as the underlying semantics of a process algebra (Strulo 1993; D'Argenio 1999; Bravetti 2002; Bravetti and D'Argenio 2004; D'Argenio and Katoen 2005). They have also defined a continuous probabilistic variant of the (strong) bisimulation. As correctly pointed out in Cattani (2005) and Cattani *et al.* (2005), these models lack enough structure to ensure that bisimilar models share the same observable behaviour. (This is due to the case where two objects may be bisimilar but it is not possible to define probabilistic executions in one of them because the transition relation is not a measurable object.) The solution proposed in Cattani (2005) and Cattani *et al.* (2005) deals with the same unstructured type of model and lifts 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 was made in Danos *et al.* (2006) with respect to the bisimulation relation for LMPs (Desharnais 1999; Desharnais *et al.* 2002). Danos *et al.* (2006) showed that there are bisimulation relations that may distinguish *beyond events*. That is, states that cannot be *separated* (that is, distinguished) by any *measurable set* (that is, any 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, Danos *et al.* (2006) defined the so-called *event bisimulation* (in opposition to the previous *state* bisimulation, a name that 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 are also Markov kernels on Λ, that is, the original LMP is also an LMP over Λ. Λ induces an equivalence relation $\mathscr{R}(\Lambda)$, which is 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 to defining a theory of LMP with internal non-determinism. Thus, we introduce *non-deterministic labelled Markov processes (NLMP)*. An NLMP has a non-deterministic transition function T_a for each label a that, given a state, returns a measurable *set* of probability measures (rather than just 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 (Giry 1981), which we also provide. We give definitions for event bisimulation and state

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

bisimulation, and prove similar properties to those in Danos *et al.* (2006), including the fact that the largest state bisimulation is also an event bisimulation. We also provide a definition of 'traditional' bisimulation that follows along the lines of Strulo (1993), D'Argenio (1999), Bravetti (2002) and D'Argenio and Katoen (2005). We prove that a traditional bisimulation is also a state bisimulation, and give sufficient conditions that make the converse holds. We also show that LMPs are just NLMPs without internal non-determinism and that state (respectively, event) bisimulation in the different models agree.

Behavioural equivalences like bisimulation have been characterised using logic with modalities, notably the Hennessy–Milner logic (see, for example, van Glabeek (2001)). We define an extension of the logic presented in the context of LMP (Desharnais 1999). In fact, the logic is similar to that of Parma and Segala (2007), which was introduced in a discrete setting. However, unlike Parma and Segala (2007), we consider two different formula levels: one that is interpreted on states and another interpreted on measures. Such a separation gives a particular insight: the actual complexity of the model lies exactly on the internal non-determinism introduced by the values of T_a (which are sets of measures). At state level, the logic is as simple as that of Desharnais (1999). We show that this logic completely characterises event bisimulation and, as a consequence, it is sound with respect to traditional and state bisimulation.

In addition, we show that a sublogic of the previous logic characterises 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 in this setting.

Nonetheless, we also show that they are different in a more general setting. In the final part of this paper, we present two counterexamples, one showing that traditional bisimulation is strictly finer than state and event bisimulation and the other showing that state bisimulation is strictly finer than event bisimulation. Both counterexamples turn out to be *non-probabilistic* NLMPs, which can be thought of as measure theoretic versions of labelled transition systems. The first example shows that traditional bisimulation distinguishes beyond measurability, and the second that event bisimulation has some weakness that has to be overcome.

This paper revises and extends our result in D'Argenio *et al.* (2009). In particular, Section 6 is new to the current paper. Most importantly, the new counterexamples presented here lead to new and different conclusions to those of D'Argenio *et al.* (2009).

2. Fundamentals and background

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

2.1. *Measure theory*

Given a set S and a collection Σ of subsets of S, we say Σ is a *σ-algebra* if and only if $S \in \Sigma$ and Σ is closed under complement and denumerable union. We use $\sigma(\mathscr{G})$ to denote

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

the *σ-algebra generated by* the family $\mathscr{G} \subseteq 2^S$, that is, the minimal σ-algebra containing $\mathscr{G}$. Each element of $\mathscr{G}$ is said to be a *generator*, and $\mathscr{G}$ is the *set of generators*. We say the pair (S, Σ) is a *measurable space*. A *measurable set* is a set $Q \in \Sigma$. A σ-additive function $\mu : \Sigma \to [0,1]$ such that $\mu(S) = 1$ is called a *probability measure*. We use δ_a to denote the Dirac probability measure concentrated in $\{a\}$. Let $\Delta(S)$ denote the *set of all probability measures* over the measurable space (S, Σ). We use the Greek letter μ, sometimes with a subscript or superscript (for example, μ', μ_1), to range over $\Delta(S)$. Let (S_1, Σ_1) and (S_2, Σ_2) be two measurable spaces. A function $f : S_1, \to S_2$ is said to be *measurable* if $\forall Q_2 \in \Sigma_2$, $f^{-1}(Q_2) \in \Sigma_1$, that is, the inverse function maps measurable sets to measurable sets. In this case we use the denotation $f : (S_1, \Sigma_1) \to (S_2, \Sigma_2)$.

A function $f : S_1 \times \Sigma_2 \to [0,1]$ is a *transition probability* (also called a *Markov kernel*) if for all $\omega_1 \in S_1$, we have $f(\omega_1, \cdot)$ is a probability measure on (S_2, Σ_2) and for all $Q_2 \in \Sigma_2$, we have $f(\cdot, Q_2)$ is measurable.

There is a standard construction in Giry (1981) for endowing $\Delta(S)$ with a σ-algebra as follows: $\Delta(\Sigma)$ is defined as the σ-algebra[†] generated by the sets of probability measures $\Delta^B(Q) \doteq \{\mu \mid \mu(Q) \in B\}$, with $Q \in \Sigma$ and $B \in \mathscr{B}([0,1])$ (where $\mathscr{B}([0,1])$ is the Borel σ-algebra on the interval $[0,1]$ generated by the open sets). When $0 \leqslant p \leqslant 1$, we will write $\Delta^{\geqslant p}(Q)$, $\Delta^{>p}(Q)$, $\Delta^{<p}(Q)$, and so on, for $\Delta^B(Q)$ with $B = [p,1], (p,1], [0,p)$, and so on, respectively. It is known that the set $\{\Delta^{\geqslant p}(Q) \mid p \in (\mathbb{Q} \cap [0,1]), Q \in \Sigma\}$ generates all $\Delta(\Sigma)$. We use the Greek letters ξ and ζ, sometimes with a subscript or superscript, to range over $\Delta(\Sigma)$.

In this setting, $f : S_1 \times \Sigma_2 \to [0,1]$ is a transition probability if and only if its curried version $f : S_1 \to \Delta(S_2)$ is measurable. (Note the notation overloading of f.) This follows from the next lemma.

Lemma 2.1. $f : S_1 \to \Delta(S_2)$ is measurable if and only if $f(\cdot, Q) : S_1 \to [0,1]$ is measurable for all $Q \in \Sigma_2$.

Proof. It is routine to show that $f^{-1}(\Delta^B(Q)) = (f(\cdot, Q))^{-1}(B)$ for all $Q \in \Sigma_2$ and $B \in \mathscr{B}([0,1])$. From this observation, $f^{-1}(\Delta^B(Q)) \in \Sigma_1$ if and only if $(f(\cdot, Q))^{-1}(B) \in \Sigma_1$. The lemma then follows since showing that $f^{-1}(\Delta^B(Q)) \in \Sigma_1$ for all generators $\Delta^B(Q)$ is sufficient for us to state that f is measurable. □

An important result for Giry's construction is that the σ-algebra of measures is *separative* (van Breugel 2005), that is, for any two elements, there is always a measurable set that contains one element but not the other.

Proposition 2.1. $\Delta(\Sigma)$ is *separative*. That is, given different $\mu, \mu' \in \Delta(S)$, there exists $\xi \in \Delta(\Sigma)$ such that $\mu \in \xi$ and $\mu' \notin \xi$.

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

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

2.2. *Relations, measures and σ-algebras*

Given a relation $R \subseteq S \times S$, we use the predicate R-closed(Q) to denote $R(Q) \subseteq Q$. Note that if R is symmetric, R-closed(Q) if and only if $\forall s, t : s\, R\, t : s \in Q \Leftrightarrow t \in Q$. Let (S, Σ) be a measurable space. For symmetric R, we 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 2.2. Let R and R' be symmetric relations such that $R \subseteq R'$. Then:

(i) $\Sigma(R) \supseteq \Sigma(R')$.
(ii) $\Delta(\Sigma(R)) \supseteq \Delta(\Sigma(R'))$.

Proof.

(i) This part follows from the fact that any measurable set that is R'-closed is also R-closed whenever $R \subseteq R'$.
(ii) Recall that $\Delta(\Sigma(R'))$ is generated by $\mathscr{G} = \{\Delta^B(Q) \mid Q \in \Sigma(R') \text{ and } B \in \mathscr{B}([0,1])\}$. Since $\Sigma(R') \subseteq \Sigma(R)$ (by (i)), we have $\mathscr{G} \subseteq \Delta(\Sigma(R))$, from which the lemma then follows. □

We can lift R to an equivalence relation in $\Delta(S)$ as follows:

$$\mu R \mu' \quad \text{if and only if} \quad \forall Q \in \Sigma(R) : \mu(Q) = \mu'(Q).$$

So the predicate R-closed can be defined on subsets of $\Delta(S)$ in just the same way as before. The following proposition will be useful later.

Proposition 2.3. If R is a symmetric relation, then every $\Delta(\Sigma(R))$-measurable set is R-closed.

Proof. Let $Q \in \Sigma(R)$ and $B \in \mathscr{B}([0,1])$. Then, if $\mu \in \Delta^B(Q)$ and $\mu R \mu'$, we have $\mu' \in \Delta^B(Q)$. 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. □

A σ-algebra Σ defines an equivalence relation $\mathscr{R}(\Sigma)$ on S as follows:

$$s\, \mathscr{R}(\Sigma)\, t \quad \text{if and only if} \quad \forall Q \in \Sigma, s \in Q \Leftrightarrow t \in Q.$$

That is, two elements are related if they cannot be separated by any measurable set. We will just state the following properties (due to Danos *et al.* (2006)) here for the sake of completeness; they relate σ-algebras and relations. In particular, (v) is a consequence of (i) and (ii).

Proposition 2.4. Let (S, Σ) be a measurable space, R be a symmetric relation on S and $\Lambda \subseteq \Sigma$ be a sub-σ-algebra of Σ. Then,

(i) $\Lambda \subseteq \Sigma(\mathscr{R}(\Lambda))$;
(ii) $R \subseteq \mathscr{R}(\Sigma(R))$;

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

(iii) if each R-equivalence class is in Σ, then $R = \mathscr{R}(\Sigma(R))$;
(iv) $\mathscr{R}(\Lambda) = \mathscr{R}(\Sigma(\mathscr{R}(\Lambda)))$; and
(v) $\Sigma(R) = \Sigma(\mathscr{R}(\Sigma(R)))$[†].

2.3. *Labelled Markov processes*

A labelled Markov process (LMP) (Desharnais 1999; Desharnais *et al.* 2002) 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 $a \in L$, we have $\tau_a : S \times \Sigma \to [0,1]$ is a transition probability. By Lemma 2.1, 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.

Desharnais (1999) and Desharnais *et al.* (2002) introduced a notion of behavioural equivalence similar to the probabilistic bisimulation of Larsen and Skou (1991).

Definition 2.1. $R \subseteq S \times S$ is a *state bisimulation on LMP* $(S, \Sigma, \{\tau_a \mid a \in L\})$ if it is symmetric[‡] and for all $s, t \in S$, $a \in L$, we have $s\ R\ t$ 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. Danos *et al.* (2006) introduced a measure-theory aware notion of behavioural equivalence.

Definition 2.2. An *event bisimulation on an LMP* $(S, \Sigma, \{\tau_a \mid a \in L\})$ is a sub-σ-algebra Λ of Σ such that $(S, \Lambda, \{\tau_a \mid a \in L\})$ is an LMP.

Danos *et al.* (2006) shows that R is a state bisimulation if and only if $\Sigma(R)$ is an event bisimulation. This is an important result that leads to a proof that the largest state bisimulation is also an event bisimulation (see Theorem 4.3 below).

3. Non-deterministic labelled Markov processes

In this section we extend the LMP model by adding internal non-determinism. That is, we allow different but equally labelled transition probabilities for leaving the same state. We also show the relation between this extension and the original LMPs.

3.1. *The model*

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 our new model. The first is that the *non-deterministic transition function* T_a now maps states to *measurable sets of* probability measures rather than to arbitrary sets as in previous approaches. This is motivated by the

[†] Proposition 2.4(v) appears in Danos *et al.* (2006), but with the unnecessary condition that R is a state bisimulation.
[‡] The requirement of symmetry is needed since otherwise $\Sigma(R)$ may not be a σ-algebra.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

fact that later on the non-determinism will have to be resolved using schedulers. If we allowed the target set of states to be an arbitrary subset, as with some continuous ones (D'Argenio 1999; Bravetti and D'Argenio 2004; Cattani *et al.* 2005), then 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, but we omit that here since it is not needed.) The second difference is inspired by the definition of LMP and Lemma 2.1 (see also the alternative definition of LMP above): for each label $a \in L$, we ask that T_a be 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 3.1. A *non-deterministic labelled 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 $a \in L$, we have $T_a : S \to \Delta(\Sigma)$ is measurable.

To ensure that the requirement that T_a is measurable is satisfied, we need to endow $\Delta(\Sigma)$ with a σ-algebra. This is a key construction for the forthcoming definitions and theorems.

Definition 3.2. $H(\Delta(\Sigma))$ 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 (Kechris 1995), and resembles the so-called hit-and-miss topologies (Naimpally 2003). Note that the generator set H_ξ contains all measurable sets that 'hit' the measurable set ξ. Also note that $T_a^{-1}(H_\xi)$ is the set of all states s such that, through label a, they 'hit' the set of measures ξ (that is, $T_a(s) \cap \xi \neq \emptyset$). This forms the basis for existentially quantifying over the non-determinism, and it is fundamental for the behavioural equivalence and the logic.

The next two examples (inspired by an example in Cattani (2005)) 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 3.1. Let $\mathscr{V} = \{\delta_q \mid q \in V\}$, where V is the non-measurable Vitali set in $[0, 1]$. It can be shown that $\mathscr{V}$ is not measurable in $\Delta(\Sigma)$. Let $T_a(s) = \mathscr{V}$ for all $s \in [0, 1]$. The resolution of the internal non-determinism by means of so-called schedulers (also adversaries or policies) (Vardi 1985; Puterman 1994), whatever its definition is, would require us to assign probabilities to all possible choices. This amounts to measuring the non-measurable set $T_a(s)$. This is why we require that T_a maps into measurable sets.

Example 3.2. 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 a Vitali set. Note that both $T_a(s)$ and $T_b(s)$ are measurable sets for every $s \in [0, 1]$. Assuming there is a scheduler that chooses to first do a and then b starting at some state s, the probability of such a set of executions cannot be measured,

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

as it requires us to apply μ to the set $T_b^{-1}(H_{\Delta(S)}) = V$, which is not measurable. Besides, we will later need the fact that sets $T_a^{-1}(H_\xi)$ are measurable so that the semantics of the logic maps into measurable sets (see Section 5).

3.2. *NLMPs as a generalisation of LMPs*

Note that an LMP is an NLMP without internal non-determinism. That is, an NLMP in which $T_a(s)$ is a singleton for all $a \in L$ and $s \in S$ is an LMP. In fact, an LMP can be encoded as an NLMP by taking $T_a(s) = \{\tau_a(s)\}$. (We will prove this formally in Proposition 3.1.) As a consequence, singletons $\{\mu\}$ must be 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 3.1. Let $\mathscr{G}$ be a denumerable π-system on S (that is, a denumerable subset of 2^S containing S and closed under finite intersection). Then, for all $\mu \in \Delta(S)$, we have $\{\mu\} \in \Delta(\sigma(\mathscr{G}))$.

Proof. It is sufficient to prove that the set

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

which is a denumerable intersection, is equal to the singleton $\{\mu\}$. By construction, μ is in the intersection. Take μ' such that $\mu \neq \mu'$. By a classical theorem of the extension of a measure (Billingsley 1995, Theorem 3.3), there must be a $Q_i \in \mathscr{G}$ 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 $\mu(Q_i) < \mu'(Q_i)$, then 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 3.1 also gives sufficient conditions to define NLMPs with *finite and denumerable non-determinism.*

Note also that asking for measurable singletons in $\Delta(\Sigma)$ does not trivialise Σ (in the sense that $\Sigma = 2^S$). A non-trivial example in which Lemma 3.1 holds is the standard Borel σ-algebra in $\mathbb{R}$. A less obvious example is as follows. Define the σ-algebra **Q-coQ** by

$$\textbf{Q-coQ} \doteq 2^{\mathbb{Q}} \cup \{\mathbb{R} \setminus \mathbb{Q} \mid Q \in 2^{\mathbb{Q}}\}.$$

Note 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 $\{\{q\} \mid q \in \mathbb{Q}\} \cup \{\varnothing\}$, it is under the conditions of Lemma 3.1, 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 3.1. Let $T_a(s) = \{\tau_a(s)\}$ for all $s \in S$ and Σ be a σ-algebra on S. Then $\tau_a : S \to \Delta(S)$ is measurable if and only if $T_a : S \to \Delta(\Sigma)$ is measurable.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Proof. Let $\xi \in \Delta(\Sigma)$. Note that $T_a(s) \in H_\xi$ if and only if $\{\tau_a(s)\} \cap \xi \neq \emptyset$ if and only if $\tau_a(s) \in \xi$. So $T_a^{-1}(H_\xi) = \tau_a^{-1}(\xi)$. 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. □

4. The bisimulations

In this section, we provide event and state bisimulations for NLMPs and show their relation to earlier definitions of bisimulation on non-deterministic and continuous probabilistic transition systems.

4.1. *Event and state 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 4.1. An *event bisimulation on an NLMP* $(S, \Sigma, \{T_a \mid a \in L\})$ is a sub-σ-algebra Λ of Σ such that $T_a : (S, \Lambda) \to (\Delta(\Sigma), H(\Delta(\Lambda)))$ is measurable for each $a \in L$.

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

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 such that $R = \mathscr{R}(\Lambda)$. More generally, we say that two states $s, t \in S$ are *event bisimilar*, denoted $s \sim_e t$, if there is an event bisimulation Λ such that $s\, \mathscr{R}(\Lambda)\, t$. The fact that $\sim_e$ is an equivalence relation is an immediate corollary of Theorem 5.5 given below. Note that, by Proposition 3.1, an event bisimulation on an 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 1989), which was lifted to discrete probabilistic models in Larsen and Skou (1991), a traditional definition of bisimulation (see Definition 4.3) asserts that, whenever $s\, R\, t$, 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 Definition 3.2 and asserts that both $T_a(s)$ and $T_a(t)$ *hit* the same measurable sets of measures.

Definition 4.2. A relation $R \subseteq S \times S$ is a *state bisimulation* if it is symmetric and for all $a \in L$, we have $s\, R\, t$ 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 4.1. Provided R is symmetric, R is a state bisimulation if and only if $\Sigma(R)$ is an event bisimulation.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Proof. By Definition 4.1, $\Sigma(R)$ is an event bisimulation if and only if 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, we have $\xi \in \Delta(\Sigma(R))$. So we have

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

if and only if (since R is symmetric)

$$s\ R\ t \Rightarrow \left(s \in T_a^{-1}(H_\xi) \Leftrightarrow t \in T_a^{-1}(H_\xi)\right)$$

if and only if (by the definition of the inverse function)

$$s\ R\ t \Rightarrow \left(T_a(s) \in H_\xi \Leftrightarrow T_a(t) \in H_\xi\right)$$

if and only if (by the definition of H_ξ)

$$s\ R\ t \Rightarrow \left(T_a(s) \cap \xi \neq \emptyset \Leftrightarrow T_a(t) \cap \xi \neq \emptyset\right).$$

This completes the proof as the last statement is the definition of state bisimulation. □

The following results are consequences of Proposition 2.4, with, for Lemma 4.2.3, the addition of Lemma 4.1 and the fact that $\mathscr{R}(\Lambda)$ is an equivalence relation. The proofs are the same as the proofs of similar results for LMP in Danos *et al.* (2006).

Lemma 4.2. Let R be a state bisimulation. Then:

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

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

Theorem 4.3. $\sim_s$ is:

(i) the largest state bisimulation;
(ii) an event bisimulation (and hence $\sim_s \subseteq \sim_e$); and
(iii) an equivalence relation.

Proof.

(i) Take $s, t \in S$ such that $s \sim_s t$. Then there is a state bisimulation R with $s\ R\ t$. Take a measurable set $\xi \in \Delta(\Sigma(\sim_s))$. Since $R \subseteq \sim_s$, by Proposition 2.2, $\Delta(\Sigma(R)) \supseteq \Delta(\Sigma(\sim_s))$. Hence, $\xi \in \Delta(\Sigma(R))$ and by Definition 4.2,

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

which proves that $\sim_s$ is a state bisimulation. By definition, it is the largest one.

(ii) Because $\sim_s$ is a state bisimulation, $\mathscr{R}(\Sigma(\sim_s))$ is a state bisimulation and an event bisimulation (Lemma 4.2.3). Since $\sim_s$ is the largest bisimulation, we have $\sim_s = \mathscr{R}(\Sigma(\sim_s))$, so it is an event bisimulation.

(iii) By definition, every event bisimulation is an equivalence relation. □

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

4.2. *A traditional view of bisimulation*

We have already stated that our definition of state bisimulation differs from a more traditional view such as those in Strulo (1993), D'Argenio (1999), Bravetti (2002), D'Argenio and Katoen (2005) and Bravetti and D'Argenio (2004). These definitions closely resemble the definition in Larsen and Skou (1991) (the only difference is that two measures are considered equivalent if they agree in every *measurable union of* equivalence classes induced by the relation). We will now give a more 'modern' variant of this definition.

Definition 4.3. A relation R is a *traditional bisimulation* if it is symmetric and for all $a \in L$, we have $s\ R\ t$ implies $T_a(s)\ R\ T_a(t)$. We say that $s, t \in S$ are *traditionally bisimilar*, denoted $s \sim_{\mathrm{t}} t$, if there is an traditional bisimulation R such that $s\ R\ t$.

Note that R is lifted this time to sets as usual: $T_a(s)\ R\ T_a(t)$ if for all $\mu \in T_a(s)$, there is $\mu' \in T_a(t)$ such that $\mu R \mu'$ and *vice versa*. (If we had explicitly included this in Definition 4.3, it would have resembled traditional definitions.)

The proof of the next proposition follows the standard strategy of classic bisimulation (Milner 1989). Apart from the probabilistic treatment, it only differs in that the composition $R \circ R'$ is a traditional bisimulation if R and R' are *reflexive* traditional bisimulations (if one of R or R' is not reflexive, $R \circ R'$ may not be a traditional bisimulation).

Proposition 4.1. $\sim_{\mathrm{t}}$ is a traditional bisimulation and an equivalence relation.

We will now discuss the relation between state bisimulation and traditional bisimulation. Lemma 4.4 states that every traditional bisimulation is a state bisimulation. Theorems 4.5 and 4.6 give sufficient conditions for strengthening Lemma 4.4 so that the converse also holds.

Lemma 4.4. If R is a traditional bisimulation, R is a state bisimulation.

Proof. Let $s\ R\ t$ and $\xi \in \Delta(\Sigma(R))$. If $T_a(s) \cap \xi \neq \emptyset$, there is $\mu \in T_a(s)$ such that $\mu \in \xi$. Since R is a traditional bisimulation, $T_a(s)\ R\ T_a(t)$, that is, there is $\mu' \in T_a(t)$ such that $\mu R \mu'$. By Proposition 2.3, 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 that ensure that a state bisimulation is also a traditional bisimulation. The first condition focuses on the NLMP, and requires the NLMP to be *image denumerable*.

Definition 4.4. An NLMP $(S, \Sigma, \{T_a \mid a \in L\})$ is *image denumerable* if and only if for all $a \in L, s \in S$, we have $T_a(s)$ is denumerable.

Theorem 4.5. Let $(S, \Sigma, \{T_a \mid a \in L\})$ be an image denumerable NLMP. Then R is a traditional bisimulation if and only if it is a state bisimulation.

Proof. The left-to-right implication is Lemma 4.4. For the other implication, we proceed as follows.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Let $s\ R\ t$ and for all $\xi \in \Delta(\Sigma(R))$, let

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

In order to show a contradiction, we suppose that $T_a(s) \not R\ T_a(t)$, that is,

$$\exists \mu \in T_a(s), \forall \mu'_i \in T_a(t) : \exists Q_i \in \Sigma(R) : \mu(Q_i) \bowtie_i \mu'_i(Q_i),$$

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 $T_a(t) \cap \xi = \emptyset$, which contradicts the assumption. □

It should be clear from this proof that we can relax the sufficient condition so that we only require that the partition $T_a(s)/R$ is denumerable for each state s and label a instead of requiring image denumerability.

Note that a state bisimulation on an LMP is a traditional bisimulation on the encoding NLMP and *vice versa* since

$$\{\tau_a(s)\} = T_a(s)\ R\ T_a(t) = \{\tau_a(t)\}$$

if and only if

$$\tau_a(s)\ R\ \tau_a(t).$$

As a consequence of Lemma 4.4 and Theorem 4.5 (a deterministic NLMP is image denumerable!), we can conclude that a state bisimulation on an 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, then R is also a traditional bisimulation.

Theorem 4.6. Let R be a symmetric relation such that $\Sigma(R)$ is generated by a denumerable set $\mathscr{G}$. Then R is a traditional bisimulation if and only if it is a state bisimulation.

Proof. As before, the left-to-right implication is Lemma 4.4. For the other implication we proceed as follows. In order to show a contradiction, we suppose that $s\ R\ t$ and $T_a(s) \not R\ T_a(t)$, that is,

$$\exists \mu \in T_a(s), \forall \mu' \in T_a(t) : \mu \not R\ \mu'.$$

By Billingsley (1995, Theorem 3.3), this implies that there exists $Q_i \in \pi(\mathscr{G})$ such that $\mu(Q_i) \neq \mu'(Q_i)$ with $i \in \mathbb{N}$. (Note that $\pi(\mathscr{G})$, the π-system generated by $\mathscr{G}$, is also denumerable and generates $\Sigma(R)$.) The rest of the proof is as in Theorem 4.5. □

5. A logic for bisimulation on NLMP

The logic presented in this section is based on the logic given by Parma and Segala (2007). The main difference is that we consider two kinds of formulas: one that is interpreted on

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

states and another that is interpreted on measures. The syntax is as follows:

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

where $a \in L$, I is a denumerable index set and $q \in \mathbb{Q} \cap [0,1]$. We use $\mathscr{L}$ to denote the set of all formulas generated by the first production and $\mathscr{L}_\Delta$ for the set of all formulas generated by the second production.

The semantics is defined with respect to an NLMP (S, Σ, T). Formulas in $\mathscr{L}$ are interpreted as sets of states in which they become true, and formulas in $\mathscr{L}_\Delta$ are interpreted as sets of measures on the state space as follows:

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

In particular, note 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 $[\varphi]_{\geqslant q}$ is valid in a measure μ whenever $\mu([\![\varphi]\!]) \geqslant q$. As a consequence, we require that sets $[\![\varphi]\!]$ and $[\![\psi]\!]$ are measurable in Σ and $\Delta(\Sigma)$, respectively. Indeed, the satisfaction of this requirement 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 this section, we fix $[\![\mathscr{L}]\!] = \{[\![\varphi]\!] \mid \varphi \in \mathscr{L}\}$ and $[\![\mathscr{L}_\Delta]\!] = \{[\![\psi]\!] \mid \psi \in \mathscr{L}_\Delta\}$.

Note that some other operators can be encoded as syntactic sugar. For instance, we can define

$$[\varphi]_{>r} \equiv \bigvee_{q \in \mathbb{Q} \cap [0,1] \wedge q > r} [\varphi]_{\geqslant q}$$

for any real $r \in [0,1]$, and $[\varphi]_{\leqslant r} \equiv \neg [\varphi]_{>r}$.

We now show that $\mathscr{L}$ characterises event bisimulation. This is an immediate consequence of the fact that $\sigma([\![\mathscr{L}]\!])$, the σ-algebra generated by the logic $\mathscr{L}$, is the smallest event bisimulation, which we will prove in this part of the section. The proof strategy resembles that of Danos *et al.* (2006, Section 5), but it is properly tailored to our two-level logic. Moreover, this separation allowed us to find an alternative to Dynkin's Theorem (which was used in Danos *et al.* (2006)).

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

The concept of a *stable family* of measurable sets is crucial to the proof of Theorem 5.5.

Definition 5.1. Given an NLMP (S, Σ, T), the family $\mathscr{C} \subseteq \Sigma$ is *stable* for (S, Σ, T) if for all $a \in L$ and $\xi \in \Delta(\mathscr{C})$, we have $T_a^{-1}(H_\xi) \in \mathscr{C}$.

Note that $\mathscr{C}$ is an event bisimulation if and only if it is a stable σ-algebra.

The key point of the proof is to show that $[\![\mathscr{L}]\!]$ is the smallest stable π-system, which is stated in Lemma 5.2. The next lemma is auxiliary to Lemma 5.2.

Lemma 5.1. $[\![\mathscr{L}_\Delta]\!] = \Delta(\mathscr{L})$

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Proof. $[\![\mathscr{L}_\Delta]\!]$ is a σ-algebra since:

(i) $\Delta(S) = [\![[\top]_{\geqslant 1}]\!] \in [\![\mathscr{L}_\Delta]\!]$;
(ii) for $\xi_i \in [\![\mathscr{L}_\Delta]\!]$, there are $\psi_i \in \mathscr{L}_\Delta$ such that $\xi_i = [\![\psi_i]\!]$, and hence $\bigcup_i \xi_i = \bigcup_i [\![\psi_i]\!] = [\![\bigvee_{i\in I} \psi_i]\!] \in [\![\mathscr{L}_\Delta]\!]$; and
(iii) for $\xi \in [\![\mathscr{L}_\Delta]\!]$, there is $\psi \in \mathscr{L}_\Delta$ such that $\xi = [\![\psi]\!]$, so $\xi^{\mathsf{c}} = [\![\psi]\!]^{\mathsf{c}} = [\![\neg\psi]\!] \in [\![\mathscr{L}_\Delta]\!]$.

Moreover, since $[\![[\varphi]_{\geqslant p}]\!] = \Delta^{\geqslant p}([\![\varphi]\!])$, every generator set of $\Delta(\mathscr{L})$ is in $[\![\mathscr{L}_\Delta]\!]$, so $\Delta(\mathscr{L}) \subseteq [\![\mathscr{L}_\Delta]\!]$.

Finally, it can be proved by induction on the depth of the formula that $[\![\mathscr{L}_\Delta]\!] \subseteq \mathscr{C}$ for any σ-algebra $\mathscr{C}$ containing all sets $[\![[\varphi]_{\geqslant p}]\!] = \Delta^{\geqslant p}([\![\varphi]\!])$ for $p \in [0,1]$ and $\varphi \in \mathscr{L}$. So $[\![\mathscr{L}_\Delta]\!]$ is the smallest σ-algebra containing all generator sets of $\Delta(\mathscr{L})$. Therefore $[\![\mathscr{L}_\Delta]\!] = \Delta(\mathscr{L})$. □

Lemma 5.2. $[\![\mathscr{L}]\!]$ is the smallest stable π-system for (S, Σ, T).

Proof. $[\![\mathscr{L}]\!]$ is a π-system since:

(i) $S = [\![\top]\!] \in [\![\mathscr{L}]\!]$;
(ii) for $Q_1, Q_2 \in [\![\mathscr{L}]\!]$, there are $\varphi_1, \varphi_2 \in \mathscr{L}$ such that $Q_1 = [\![\varphi_1]\!]$ and $Q_2 = [\![\varphi_2]\!]$, so

$$\begin{aligned} Q_1 \cap Q_2 &= [\![\varphi_1]\!] \cap [\![\varphi_2]\!] \\ &= [\![\varphi_1 \wedge \varphi_2]\!] \\ &\in [\![\mathscr{L}]\!]. \end{aligned}$$

For stability, let $\xi \in \Delta(\mathscr{L})$. By Lemma 5.1, there is $\psi \in \mathscr{L}_\Delta$ such that $[\![\psi]\!] = \xi$. So

$$\begin{aligned} T_a^{-1}(H_\xi) &= T_a^{-1}(H_{[\![\psi]\!]}) \\ &= [\![\langle a\rangle\psi]\!] \\ &\in [\![\mathscr{L}]\!]. \end{aligned}$$

Let $\mathscr{C}$ be another stable π-system for (S, Σ, T). By induction on the depth of the formula, we show simultaneously that $\mathscr{C} \supseteq [\![\mathscr{L}]\!]$ and $\Delta(\mathscr{C}) \supseteq \Delta(\mathscr{L})$. First note that $[\![\top]\!] = S \in \mathscr{C}$ since $\mathscr{C}$ is a π-system. Now suppose as induction hypothesis that $[\![\varphi]\!], [\![\varphi_1]\!], [\![\varphi_2]\!] \in \mathscr{C}$ and $[\![\psi]\!], [\![\psi_i]\!] \in \Delta(\mathscr{C})$ for $i \geqslant 0$. Then:

(i) $[\![\varphi_1 \wedge \varphi_2]\!] = [\![\varphi_1]\!] \cap [\![\varphi_2]\!] \in \mathscr{C}$, because $\mathscr{C}$ is a π-system.
(ii) $[\![\langle a\rangle\psi]\!] = T_a^{-1}(H_{[\![\psi]\!]}) \in \mathscr{C}$, because $\mathscr{C}$ is stable.
(iii) $[\![\bigvee_{i\in I} \psi_i]\!] = \bigcup_i [\![\psi_i]\!] \in \Delta(\mathscr{C})$ because $\Delta(\mathscr{C})$ is a σ-algebra.
(iv) $[\![\neg\psi]\!] = [\![\psi]\!]^{\mathsf{c}} \in \Delta(\mathscr{C})$ because $\Delta(\mathscr{C})$ is a σ-algebra.
(v) $[\![[\varphi]_{\geqslant p}]\!] = \Delta^{\geqslant p}([\![\varphi]\!]) \in \Delta(\mathscr{C})$ by definition of generator set of $\Delta(\mathscr{C})$. □

Lemma 5.3 is auxiliary to Lemma 5.4. It is also significantly simpler than the related Danos *et al.* (2006, Lemma 5.4). This is because of our definition of stability and the use of a powerful result from Viglizzo (2005).

Lemma 5.3. If $\mathscr{C}$ is a stable π-system for (S, Σ, T), then $\sigma(\mathscr{C})$ is also stable.

Proof. We first observe that $\mathscr{C}$ is stable if and only if

$$\{T_a^{-1}(H_\xi) \mid a \in L, \xi \in \Delta(\mathscr{C})\} \subseteq \mathscr{C}.$$

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

By Viglizzo (2005, Lemma 3.6), $\Delta(\mathscr{C}) = \Delta(\sigma(\mathscr{C}))$, so

$$\{T_a^{-1}(H_\xi) \mid a \in L, \xi \in \Delta(\sigma(\mathscr{C}))\} \subseteq \mathscr{C} \subseteq \sigma(\mathscr{C}),$$

which proves that $\sigma(\mathscr{C})$ is stable. □

The next lemma is central to the proof that $\mathscr{L}$ characterises event bisimulation, which is stated in Theorem 5.5.

Lemma 5.4. $\sigma(\mathscr{L})$ is the smallest stable σ-algebra included in Σ.

Proof. Let $\mathscr{F}$ be the smallest stable σ-algebra included in Σ. By Lemma 5.2, $[\![\mathscr{L}]\!] \subseteq \mathscr{F}$, since $\mathscr{F}$ is a stable π-system. Therefore $\sigma(\mathscr{L}) \subseteq \mathscr{F}$ since $\mathscr{F}$ is also a σ-algebra. For the other inclusion, we observe that $[\![\mathscr{L}]\!]$ is a stable π-system because of Lemma 5.2. Then, by Lemma 5.3, $\sigma(\mathscr{L})$ is stable, and thus contains $\mathscr{F}$. □

Theorem 5.5. The logic $\mathscr{L}$ completely characterises event bisimulation. In other words, $\mathscr{R}(\mathscr{L}) = \sim_{\mathrm{e}}$

Proof. Lemma 5.4 establishes that $\sigma(\mathscr{L})$ is stable, that is, it is an event bisimulation. Being the smallest, it implies that any other event bisimulation preserves $\mathscr{L}$ formulas. □

A consequence of this theorem, together with Theorem 4.3 and Lemma 4.4, is that both traditional and state bisimulation are sound for $\mathscr{L}$, that is, they preserve the validity of formulas.

Theorem 5.6. $\sim_{\mathrm{t}} \subseteq \sim_{\mathrm{s}} \subseteq \sim_{\mathrm{e}} = \mathscr{R}(\mathscr{L})$.

5.1. *Completeness for image finite NLMPs*

The rest of this section is devoted to showing that the logic completely characterises (all three) bisimulation(s) on NLMPs with image finite non-determinism and standing on analytic spaces. In fact, we show completeness of the sublogic of $\mathscr{L}$ defined by

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

where $\bowtie_i \in \{>, <\}$ and $q_i \in \mathbb{Q} \cap [0,1]$. We define the new modal operation using a shorthand notation:

$$\langle a\rangle[\,_{\bowtie_i q_i}\varphi_i]_{i=1}^n \equiv \langle a\rangle \bigwedge_{i=1}^n [\varphi_i]_{\bowtie_i q_i}.$$

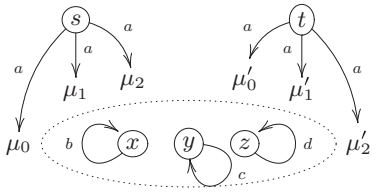
Therefore, its semantics is given by

$$[\![\langle a\rangle[\,_{\bowtie_i q_i}\varphi_i]_{i=1}^n]\!] = T_a^{-1}(H_{\cap_{i=1}^n \Delta^{\bowtie_i q_i}([\![\varphi_i]\!])}).$$

We use $\mathscr{L}_{\mathrm{f}} \subseteq \mathscr{L}$ to denote the set of all formulas defined with the grammar above. Note that $\mathscr{L}_{\mathrm{f}}$ is a denumerable set whenever the set of labels L is denumerable.

The expression $\langle a\rangle[\,_{\bowtie_i q_i}\varphi_i]_{i=1}^n$ is like a conjunction of formulas $\langle a\rangle_{\bowtie_i q_i}\varphi_i$, but the probabilistic bounds must be satisfied by the *same* non-deterministic transition. Modality $\langle a\rangle_{>q}\varphi$ suffices to characterise bisimulation on LMP (Desharnais *et al.* 2002), but, as we

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press



	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	$\frac{2}{3}$

Fig. 1. s and t are not bisimilar

see in the next example, which originates in Celayes (2006), it is not enough for the more general setting of NLMPs.

Example 5.1. Consider the discrete NLMPs shown in Figure 1. 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 behaviour after a label will not be able to distinguish between s and t. For example, $[\![\langle a \rangle_{>q} \varphi]\!] = \{w \mid T_a(w) \cap \Delta^{>q}([\![\varphi]\!]) \neq \emptyset\}$ 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).

The fact that this new modal operator is essential also shows that our σ-algebra $H(\Delta(\Sigma))$ in Definition 3.2 cannot be simplified to

$$\sigma(\{H_{\Delta^B(Q)} : B \in \mathscr{B}([0,1]), Q \in \Sigma\}).$$

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 $\{H_{\Delta^B(Q)} : B \in \mathscr{B}([0,1]), Q \in \Sigma\}$.

The logical characterisation of bisimulation is succinctly stated as

$$s \sim_{\mathrm{s}} t \Leftrightarrow s \, \mathscr{R}(\mathscr{L}_{\mathrm{f}}) \, t$$

(and similarly for $\sim_{\mathrm{t}}$). The left-to-right implication is immediate by Theorem 5.6. For the converse, we restrict the state space and the branching.

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

Definition 5.2. A topological space is *Polish* if it is separable (that is, it contains a countable dense subset) and completely metrisable. 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 $(X, \sigma(\mathscr{T}))$ where $\mathscr{T}$ is an analytic (Polish) topology on X.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Every standard Borel space is analytic, but the converse is not true. 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 Desharnais and Panangaden (2003) 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 5.7. Let (S, Σ) be an analytic Borel space. Let $\mathscr{F} \subseteq \Sigma$ be countable and assume $S \in \mathscr{F}$. Then $\Sigma(\mathscr{R}(\mathscr{F})) = \sigma(\mathscr{F})$.

The following lemma provides a general framework for proving that a logic characterises bisimulation. In fact, it has been used to prove that less expressive logics characterise traditional bisimulation in some restricted NLMPs (Celayes 2006).

Lemma 5.8. Let (S, Σ, T) be an NLMP with (S, Σ) being an analytic Borel space. Let $\mathfrak{L}$ be a logic such that:

(i) $\mathfrak{L}$ contains operators $\top$ and $\wedge$ with the usual semantics;
(ii) for every formula $\varphi \in \mathfrak{L}$, $[\![\varphi]\!]$ is Σ-measurable;
(iii) the set of all formulas in $\mathfrak{L}$ is denumerable; and
(iv) for every $s\ \mathscr{R}(\mathfrak{L})\ t$ and every $\mu \in T_a(s)$, there exists $\mu' \in T_a(t)$ such that

$$\forall \varphi \in \mathfrak{L}, \mu([\![\varphi]\!]) = \mu'([\![\varphi]\!]).$$

Then two logically equivalent states s, t are traditionally bisimilar.

Proof. Let $\mathscr{F} = \{[\![\varphi]\!] \mid \varphi \in \mathfrak{L}\}$. Because of condition (i), $[\![\top]\!] = S$ and

$$[\![\varphi_1]\!] \cap [\![\varphi_2]\!] = [\![\varphi_1 \wedge \varphi_2]\!].$$

So $\mathscr{F}$ forms a π-system. Because of condition (iv), μ, μ' agree in $\mathscr{F}$ and, by Billingsley (1995, Theorem 3.3), they also agree in $\sigma(\mathscr{F})$. Note that the hypotheses of Theorem 5.7 are met, that is, Σ is analytic, $\mathscr{F} \subseteq \Sigma$ is countable (by conditions (ii) and (iii)) such that $S \in \mathscr{F}$ (by condition (i)), and $\mathscr{R}(\mathfrak{L})$ equals $\mathscr{R}(\mathscr{F})$. Therefore, by Theorem 5.7, $\sigma(\mathscr{F}) = \Sigma(\mathscr{R}(\mathfrak{L}))$, which implies that μ and μ' agree in $\Sigma(\mathscr{R}(\mathfrak{L}))$. Since $\mathscr{R}(\mathfrak{L})$ is symmetric, $\mathscr{R}(\mathfrak{L})$ is a traditional bisimulation. □

Note that Lemma 5.8 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 $\mathscr{L}_{\mathrm{f}}$ has operators $\top$ and $\wedge$, is denumerable, and that each formula is interpreted in a Σ-measurable set. We will now show that the transfer property can be encoded using the modality.

Lemma 5.9. Let (S, Σ, T) be an image finite NLMP (that is, $T_a(s)$ is finite for all $a \in L, s \in S$). Then for every pair of states such that $s\ \mathscr{R}(\mathscr{L}_{\mathrm{f}})\ t$ and $\mu \in T_a(s)$, there is a $\mu' \in T_a(t)$ such that $\forall \varphi \in \mathscr{L}_{\mathrm{f}}, \mu([\![\varphi]\!]) = \mu'([\![\varphi]\!])$.

Proof. In order to show a contradiction, we suppose that there are s, t with $s\ \mathscr{R}(\mathscr{L}_{\mathrm{f}})\ t$ and there is a $\mu \in T_a(s)$ such that for all $\mu'_i \in T_a(t)$ there is a formula $\varphi_i \in \mathscr{L}_{\mathrm{f}}$ with $\mu([\![\varphi_i]\!]) \neq \mu'_i([\![\varphi_i]\!])$. Since $T_a(t)$ is finite, there are at most n different μ'_i. We can

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

choose $\bowtie_i \in \{>, <\}, q_i \in \mathbb{Q} \cap [0,1]$ accordingly to make $\mu([\![\varphi_i]\!]) \bowtie_i q_i \bowtie_i \mu_i'([\![\varphi_i]\!])$. Take $\psi = \langle a \rangle [_{\bowtie_i q_i} \varphi_i]_{i=1}^n$. Then $s \in [\![\psi]\!]$ but $t \notin [\![\psi]\!]$, which contradicts $s \mathscr{R}(\mathscr{L}_\mathrm{f})\, t$. □

So, finally, we can state the following theorem.

Theorem 5.10. Let (S, Σ, T) be an image finite NLMP with (S, Σ) being analytic. For all $s, t \in S$,

$$s \sim_\mathrm{t} t \quad \Leftrightarrow \quad s \sim_\mathrm{s} t \quad \Leftrightarrow \quad s \sim_\mathrm{e} t \quad \Leftrightarrow \quad s \mathscr{R}(\mathscr{L}_\mathrm{f})\, t.$$

Proof.

$$\begin{aligned} s \sim_\mathrm{t} t &\Rightarrow s \sim_\mathrm{s} t \\ &\Rightarrow s \sim_\mathrm{e} t \\ &\Leftrightarrow s \mathscr{R}(\mathscr{L})\, t && \text{(by Theorem 5.6)} \\ &\Rightarrow s \mathscr{R}(\mathscr{L}_\mathrm{f})\, t && \text{(because } \mathscr{L}_\mathrm{f} \subseteq \mathscr{L}\text{)} \\ &\Rightarrow s \sim_\mathrm{t} t. && \text{(by Lemmas 5.8 and 5.9)} \end{aligned}$$

□

6. Non-probabilistic NLMPs and counterexamples

The purpose of this section is to construct counterexamples over standard Borel spaces witnessing the fact 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, the only difference being that the state space has a σ-algebra attached.

In a way, the form of the counterexamples – non-probabilistic NLMPs over standard Borel spaces with uncountable branching – shows that our Theorems 4.5 and 4.6 are the best possible, even if we assume that our state space is the Borel space of the real numbers.

6.1. *The subspace of Dirac measures*

Since the counterexample NLMPs only run on Dirac measures over standard Borel spaces, we focus first on understanding these objects.

Let (S, Σ) be a measurable space. We define $\delta(P) = \{\delta_s : s \in P\}$ for each $P \subseteq S$. The set $\delta(S)$ inherits the measurable structure from $\Delta(S)$ by restriction: its σ-algebra is

$$\Delta(\Sigma)|\delta(S) = \{\xi \cap \delta(S) : \xi \in \Delta(\Sigma)\}.$$

Note that the elements of $\Delta(\Sigma)|\delta(S)$ are not necessarily measurable sets in $(\Delta(S), \Delta(\Sigma))$. However, it is indeed the case if Σ is Borel standard. This is stated in the following proposition.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Proposition 6.1.

(1) The map $s \mapsto \delta_s$ is a measurable embedding between S and $\Delta(S)$, that is, the function $\delta_{(\cdot)}$ is a bijection between S and its image $\delta(S)$ such that both $\delta_{(\cdot)}$ and $\delta_{(\cdot)}^{-1}$ are measurable.
(2) If (S, Σ) is a standard Borel space, $\delta(S)$ belongs to $\Delta(\Sigma)$, that is, it is a measurable set, so $\Delta(\Sigma)|\delta(S) \subseteq \Delta(\Sigma)$.
(3) If (S, Σ) is standard and $X \subseteq S$, then $\delta(X)$ is measurable if and only if X is measurable.

Proof. It is clear that δ is injective. To show that it is an embedding amounts to proving that

$$\Delta(\Sigma)|\delta(S) = \{\delta(Q) : Q \in \Sigma\}.$$

First observe that $\Delta(\Sigma)$ is the smallest family that contains

$$\mathscr{G} = \{\Delta^{<q}(Q), (\Delta(S) \setminus \Delta^{<q}(Q)) : q \in \mathbb{Q} \cap [0,1], Q \in \Sigma\}$$

and is closed under countable intersections and unions. We first show that for every $\xi \in \mathscr{G}$, we have $\xi \cap \delta(S)$ is of the form $\delta(Q)$:

$$\Delta^{<q}(Q) \cap \delta(S) = \begin{cases} \varnothing & q \leqslant 0 \\ \delta(S \setminus Q) & 0 < q \leqslant 1 \\ \delta(S) & q > 1 \end{cases}$$

$$(\Delta(S) \setminus \Delta^{<q}(Q)) \cap \delta(S) = \begin{cases} \delta(S) & q \leqslant 0 \\ \delta(Q) & 0 < q \leqslant 1 \\ \varnothing & q > 1. \end{cases}$$

Incidentally, this also proves $\Delta(\Sigma)|\delta(S) \supseteq \{\delta(Q) : Q \in \Sigma\}$. Assume $\xi_i \cap \delta(S)$ is of the form $\delta(Q_i)$ with $Q_i \in \Sigma$ for every i. Then

$$\begin{aligned} \left(\bigcup_i \xi_i\right) \cap \delta(S) &= \bigcup_i \xi_i \cap \delta(S) \\ &= \bigcup_i \delta(Q_i) \\ &= \delta\left(\bigcup_i Q_i\right). \end{aligned}$$

Since Σ is a σ-algebra, $(\bigcup_i \xi_i) \cap \delta(S)$ is of the form $\delta(Q)$. The same works for countable intersections, so we have the other inclusion.

If S is standard, then $\Delta(S)$ is standard by Kechris (1995, Theorem 17.23, 17.24). Since $s \mapsto \delta_s$ is injective, (2) follows from Kechris (1995, Corollary 15.2)[†].

By (1), we have that for $X \subseteq S$, $\delta(X)$ is $(\Delta(\Sigma)|\delta(S))$-measurable if and only if X is measurable. Using (2), we can state that $\delta(X)$ is $\Delta(\Sigma)$-measurable if and only if X is measurable. □

[†] Alternatively, (S, Σ) is isomorphic to $([0,1], \mathscr{B}([0,1]))$ and the functor Δ can be defined in the category of Polish spaces and continuous functions. It is not hard to show that δ is a *continuous* embedding. Hence $\delta([0,1])$ is compact in $\Delta([0,1])$ and *a fortiori* measurable.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

6.2. *Non-probabilistic NLMPs*

We say an NLMP $\mathbf{S} = (S, \Sigma, \{T_a : a \in L\})$ is *non-probabilistic* if for all $a \in L$ and $s \in S$, we have $T_a(s) \subseteq \delta(S)$. A non-probabilistic NLMP is essentially a labelled transition system (LTS) over a measurable space. However, as we will see, our notions of bisimulation differ from the classical notion for LTS.

We will write $\langle a \rangle Q$ for $\{s : T_a(s) \cap \delta(Q) \neq \emptyset\}$. The interpretation of this is clear: $\langle a \rangle Q$ are the states from which we can reach Q after an a-action.

Lemmas 6.2, 6.3 and 6.4 give the formulation of event, state and traditional bisimulations in the setting of non-probabilistic NLMPs over standard Borel spaces. Lemma 6.1 is the basis of the proof of Lemma 6.2 and is also used in the proof of our counterexamples.

Lemma 6.1. Assume (S, Σ) is standard and $\Lambda \subseteq \Sigma$ is a sub-σ-algebra. Let $T_a : S \to \Delta(\Sigma)$ with $T_a(s) \subseteq \delta(S)$ for all $s \in S$. Then T_a is $(\Lambda, H(\Delta(\Lambda)))$-measurable if and only if for all $Q \in \Lambda$, we have $\langle a \rangle Q \in \Lambda$ (that is, Λ is *stable* under the mapping $\langle a \rangle \cdot$).

Proof.

($\Rightarrow$) Let $Q \in \Lambda$ and take

$$\xi = (\Delta(S) \setminus \Delta^{<1}(Q)) \in \Delta(\Lambda).$$

Then

$$\begin{aligned} \langle a \rangle Q &= \{s : T_a(s) \cap \delta(Q) \neq \emptyset\} \\ &= \{s : T_a(s) \cap \delta(S) \cap \xi \neq \emptyset\} \\ &= T_a^{-1}(H_\xi) \\ &\in \Lambda. \end{aligned}$$

($\Leftarrow$) Let $\xi \in \Delta(\Lambda)$. By Proposition 6.1 (1), we have

$$\Delta(\Lambda)|\delta(S) = \{\delta(Q) : Q \in \Lambda\},$$

so $\xi \cap \delta(S) = \delta(Q)$ for some $Q \in \Lambda$. Then

$$\begin{aligned} T_a^{-1}(H_\xi) &= \{s : T_a(s) \cap \delta(Q) \neq \emptyset\} \\ &= \langle a \rangle Q \\ &\in \Lambda. \end{aligned}$$

□

Throughout the rest of this section we will assume that $\mathbf{S} = (S, \Sigma, \{T_a : a \in L\})$ is a non-probabilistic NLMP over a standard Borel space. The next lemma is a corollary of Lemma 6.1.

Lemma 6.2. A σ-algebra $\Lambda \subseteq \Sigma$ is an event bisimulation on $\mathbf{S}$ if and only if it is stable under the mapping $\langle a \rangle \cdot$.

Lemma 6.3. A symmetric relation R is a state bisimulation on $\mathbf{S}$ if and only if for all $s, t \in S$ such that $s \mathrel{R} t$, we have for all $Q \in \Sigma(R)$, that $s \in \langle a \rangle Q \Leftrightarrow t \in \langle a \rangle Q$.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

Proof.

($\Rightarrow$) Let $s\ R\ t$ and $Q \in \Sigma(R)$. Observe that $\Delta^{\geqslant 1}(Q) \in \Delta(\Sigma(R))$ and $\Delta^{\geqslant 1}(Q) \cap \delta(S) = \delta(Q)$. Then

$$\begin{aligned} s \in \langle a \rangle Q &\Leftrightarrow T(s) \cap \Delta^{\geqslant 1}(Q) \neq \emptyset \\ &\overset{\star}{\Leftrightarrow} T(t) \cap \Delta^{\geqslant 1}(Q) \neq \emptyset \\ &\Leftrightarrow t \in \langle a \rangle Q, \end{aligned}$$

where in $\star$ we use the fact that R is a state bisimulation.

($\Leftarrow$) Let $s\ R\ t$ and $\xi \in \Delta(\Sigma(R))$. Let Q be such that $\delta(Q) = \delta(S) \cap \xi$. Then

$$T_a(s) \cap \xi = T_a(s) \cap \delta(Q),$$

so

$$T_a(s) \cap \xi \neq \emptyset \Leftrightarrow s \in \langle a \rangle Q.$$

Similarly,

$$T_a(t) \cap \xi \neq \emptyset \Leftrightarrow t \in \langle a \rangle Q.$$

If $Q \in \Sigma(R)$, then

$$s \in \langle a \rangle Q \Leftrightarrow t \in \langle a \rangle Q$$

by hypothesis.

We now show that indeed $Q \in \Sigma(R)$. Since $\xi \in \Delta(\Sigma(R))$, by Proposition 6.1 (2) and (3), $Q \in \Sigma$. It only remains to show that Q is R-closed. So we let $x\ R\ y$ and $x \in Q$; hence $\delta_x \in \xi$. But for any $X \in \Sigma(R)$ and $q \in [0, 1]$, we have

$$\begin{aligned} \delta_x \in \Delta^{\geqslant q}(X) &\Leftrightarrow x \in X \\ &\Leftrightarrow y \in X \\ &\Leftrightarrow \delta_y \in \Delta^{\geqslant q}(X). \end{aligned}$$

Since δ_x and δ_y cannot be separated by any generator set of $\Delta(\Sigma(R))$, they cannot be separated by a set in $\Delta(\Sigma(R))$, so $\delta_y \in \xi$ and thus $y \in Q$. □

Lemma 6.4. A symmetric relation R is a traditional bisimulation on $\mathbf{S}$ if and only if for all $s, t \in S$ and $\delta_u \in T_a(s)$, if $s\,R\,t$ then there exists $\delta_v \in T_a(t)$ such that $u\ \mathscr{R}(\Sigma(R))\ v$.

Proof. Assume $s\,R\,t$. So $T_a(s)\,R\,T_a(t)$ if and only if for every $\mu \in T_a(s)$ there exists $\mu' \in T_a(t)$ such that $\mu\,R\,\mu'$. But since $\mathbf{S}$ is non-probabilistic, $\mu = \delta_u$, and $\mu' = \delta_v$ for some $u, v \in S$. Now $\delta_u\,R\,\delta_v$ means that for every $Q \in \Sigma(R)$, we have $\delta_u(Q) = \delta_v(Q)$, and this is equivalent to $\forall Q \in \Sigma(R) : u \in Q \Leftrightarrow v \in Q$. The last assertion is $u\ \mathscr{R}(\Sigma(R))\ v$. □

Given the last lemma, it should be clear that this 'measurable' notion of bisimulation is weaker than the standard one for LTS in Milner (1989).

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

6.3. *Traditional bisimilarity ≠ event- or state-bisimilarity.*

Consider the standard Borel space

$$(S_1, \Sigma_1) = ([0,1] \cup [2,3] \cup \{s,t,x\}, \mathscr{B}([0,1] \cup [2,3] \cup \{s,t,x\}))$$

where $\{s,t,x\} \subset \mathbb{R} \setminus [0,3]$. Let V be a non-Borel subset of $[2.5,3]$. It is clear that $[0,1]$ is equinumerous with $[2,3] \setminus V$; just pick a bijection f between them. Now, we let $L_1 = \{a\} \cup [0,1]$ be the set of labels and let $\mathbf{S_1} = (S_1, \Sigma_1, \{T_a : a \in L_1\})$ where

$$\begin{aligned} T_a(s) &= \delta([2,3]) \\ T_a(t) &= \delta([0,1]) \\ T_r(r) &= T_r(f(r)) = \{\delta_x\} && \text{if } r \in [0,1] \\ T_c(y) &= \varnothing && \text{otherwise.} \end{aligned}$$

Now, take $\mathscr{F}$ to be $\{\{s,t\}, \{x\}, \{r, f(r)\}_{r\in[0,1]}\}$ and $R = \mathscr{R}(\sigma(\mathscr{F}))$.

Lemma 6.5. $\mathbf{S_1}$ is a non-probabilistic NLMP, $\sigma(\mathscr{F})$ is an event bisimulation and R is a state bisimulation.

Proof. First note that for all c, y, we have $T_c(y) \in \Delta(\Sigma_1)$ by Proposition 6.1 (3). The proof that T_c is a measurable map for each $c \in L_1$ is routine.

We now check that $\sigma(\mathscr{F})$ is an event bisimulation. Observe first that for all $Q \in \sigma(\mathscr{F})$, we have $\langle a \rangle Q$ is empty or equal to $\{s,t\} \in \sigma(\mathscr{F})$, so $\sigma(\mathscr{F})$ is stable under T_a by Lemma 6.1. For $0 \leqslant r \leqslant 1$, we have $\langle r \rangle Q \neq \varnothing$ if and only if $x \in Q$, and in that case

$$\langle r \rangle Q = \{r, f(r)\} \in \sigma(\mathscr{F}). \tag{1}$$

We now show that R is a state bisimulation. In order to show a contradiction, using Lemma 6.3, we assume that there exists $Q \in \Sigma_1(R)$, $c \in L_1$ and $z, y \in S_1$ such that $z \, R \, y$ and $z \in \langle c \rangle Q$ but $y \notin \langle c \rangle Q$. Hence $\langle c \rangle Q$ must not be R-closed. By (1), for $0 \leqslant r \leqslant 1$ and *every* $Q \in \Sigma_1$, we have $\langle r \rangle Q$ is R-closed. So it should be the case that $\langle a \rangle Q$ is not R-closed. Observe that the only R-closed sets $Q \subseteq S_1$ such that $\langle a \rangle Q$ is not R-closed are of the form $A \cup V$ where $A \in \{\varnothing, \{s,t,x\}, \{s,t\}, \{x\}\}$. This set Q is non-measurable since V was chosen to be non-measurable. But then Q is not in $\Sigma_1(R)$, which gives a contradiction. □

Theorem 6.6. State bisimilarity (respectively, event bisimilarity) and traditional bisimilarity differ in $\mathbf{S_1}$.

Proof. Because of Lemma 6.5, it sufficies to show that s and t are not traditionally bisimilar.

It is easy to see that for $0 \leqslant r \leqslant 1$, we have $r \not\sim_{\mathrm{t}} y$ if $y \notin \{r, f(r)\}$ since $T_r(r)$ is non-empty but $T_r(y) = \varnothing$. Therefore, $\{r, f(r)\}$ is $\sim_{\mathrm{t}}$-closed for every $0 \leqslant r \leqslant 1$, so $\{r, f(r)\} \in \Sigma_1(\sim_{\mathrm{t}})$.

In order to show a contradiction, we now assume $s \sim_{\mathrm{t}} t$. Let $y \in V$. Since $\delta_y \in T_a(s)$, by Lemma 6.4, there must exist some $0 \leqslant r \leqslant 1$ such that $y \, \mathscr{R}(\Sigma_1(\sim_{\mathrm{t}})) \, r$. But $y > 1$ and is not in the image of f, so the set $\{r, f(r)\} \in \Sigma_1(\sim_{\mathrm{t}})$ separates y from r, which contradicts the fact that $y \, \mathscr{R}(\Sigma_1(\sim_{\mathrm{t}})) \, r$.

Since $\sim_{\mathrm{s}} \subseteq \sim_{\mathrm{e}}$, event bisimilarity and traditional bisimilarity also differ in $\mathbf{S_1}$. □

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

6.4. *State bisimilarity ≠ event bisimilarity*

In the final part of this section we prove that the greatest event bisimulation $\sim_e$ is not contained in $\sim_s$. We do this by slightly modifying $\mathbf{S_1}$. We now 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 by picking any bijection f between $[0, 1]$ and $[2, 2.5]$. The transition is defined in just the same way as for $\mathbf{S_1}$ except that we use the new f. We also use family $\mathscr{F}$, but defined with the new f.

Lemma 6.7. $V \notin \sigma(\mathscr{F})$.

Proof. It is clear that every member of $\sigma(\mathscr{F})$ is countable or has a countable complement, from which the lemma then follows. □

The proof of Lemma 6.5 works equally well for the following lemma.

Lemma 6.8. $\mathbf{S_2} = (S_2, \Sigma_2, \{T_a : a \in L_1\})$ is a non-probabilistic NLMP and $\sigma(\mathscr{F})$ is an event bisimulation.

In this case, relation $R = \mathscr{R}(\sigma(\mathscr{F}))$ is an event bisimulation that it is not a state bisimulation.

Theorem 6.9. Event and state bisimilarity differ in $\mathbf{S_2}$.

Proof. Since $(s, t) \in R \subseteq \sim_e$, we just have to show that $s \not\sim_s t$. Observe that $V \in \Sigma_2(R)$. If s and t were state-bisimilar, by Lemma 6.3, it would be the case that $s \in \langle a \rangle V$ if and only if $t \in \langle a \rangle V$. But this is clearly not the case since $\delta_3 \in T_a(s) \cap \delta(V)$ and $T_a(t) \cap \delta(V) = \emptyset$. □

7. Concluding remarks

In order to define a process theory that allows the verification of compositionally modelled systems against simple (possibly non-deterministic) specifications, we need a semantic relation that allows for abstraction such as weak bisimulation. In this setting, internal non-determinism is crucial.

In this paper we have introduced the model of non-deterministic labelled Markov processes that allows the modelling of continuous probabilistic systems with internal non-determinism. Unlike similar models (D'Argenio 1999; Bravetti 2002; Bravetti and D'Argenio 2004; D'Argenio and Katoen 2005; Cattani 2005), 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 (in particular, event bisimulation and logic, but it also means schedulers are definable). We have shown that NLMPs naturally extend LMPs. For the definition of the transition and the development of the whole work, Definition 3.2 is crucial, as it provides the foundation for dealing with non-determinism.

As a first step towards developing the desired process theory, we have given different definitions of bisimulations. We have proposed three possible generalisations of the two bisimulations on LMPs. The event bisimulation here responds exactly to the same

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

definition principle in both LMP and NLMP. Instead, the state bisimulation in LMPs generalises to NLMPs as both state bisimulation and traditional bisimulation. We know that traditional bisimulation is finer than state bisimulation, but, in Theorems 4.5 and 4.6, we have given sufficient conditions under which they agree.

We have also given a logical characterisation of event bisimulation (Theorem 5.5). Such a logic ($\mathscr{L}$) can be seen as a revision of the one introduced in Parma and Segala (2007) in a discrete probabilistic setting. Formulas in our setting belong to two different classes: state formulas and measure formulas. Note that negation and infinitary (but denumerable) disjunction (or conjunction) is only present for the second of these classes, meaning that the complexity of the model lies precisely in the internal non-determinism.

A consequence of the characterisation is that the logic is sound for both state and traditional bisimulations (Theorem 5.6). For the restricted case of image finite NLMPs running on analytic Borel spaces, all equivalences coincide (Theorem 5.10). Note that the logic we used to show this equivalence is in fact a sublogic of $\mathscr{L}$, which has already appeared in earlier work (Celayes 2006).

The coincidence between all these equivalences does not generalise to arbitrary NLMPs, as we showed in Theorems 6.6 and 6.9. Observe that the counterxamples presented in these theorems are non-probabilistic NLMPs over standard Borel spaces with uncountable branching. This shows that Theorems 4.5 and 4.6 are in some way the best possible for equating traditional and state bisimulation, even if we assume that the state space is the Borel space of the real numbers. Though we did not present a theorem, we mentioned a third important difference for these 'measure-theoretic' LTSs: in the general case, Park and Milner's bisimulation is strictly finer than traditional bisimulation. The latter considers the measure space of the state space, while the former does not (or, alternatively, it only considers the discrete σ-algebra 2^S).

Some additional observations on the counterexamples are in order. First, counter-example $\mathbf{S_1}$ in Theorem 6.6 relies on the fact that state bisimulation cannot distinguish a non-measurable set V while traditional bisimulation can. From our point of view, such a distinction should not be possible since V has no measure. A second observation is that counterexample $\mathbf{S_2}$ in Theorem 6.9 makes a distinction for measurable set V that the event bisimulation cannot distinguish. In our opinion, such a distinction should be observed since a possible scheduler may lead to such a 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 this sense, we argue that state bisimulation is the most appropriate definition.

This is rather disappointing since logic $\mathscr{L}$ has a natural definition but, as it completely characterises event bisimulation, it will not be able to test for the presence of states like those in V in $\mathbf{S_2}$. This is due to the fact that the logic cannot test transitions bearing continuously many labels. This means we need to add structure to the set of labels on the NLMP. In any case, this would also be necessary for the definition of schedulers and probabilistic trace semantics.

We are currently focussed on defining NLMPs with labels equipped with a σ-algebra, as well as on the study of schedulers for these objects and probabilistic trace semantics. This will allow us to contrast the two local behavioural equivalences: state

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

and traditional bisimulation. It is expected that at least one of them implies a global behavioural equivalence, like probabilistic trace equality. Schedulers would also let us define probabilistic weak transitions and their related bisimulations. We are also trying to refine the idea of event bisimulation and the logic so that they can distinguish situations like the one shown by NLMP $\mathbf{S_2}$.

If necessary, we will restrict consideration to standard Borel spaces. Confining ourselves 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 (D'Argenio 1999) in terms of NLMPs meets most of the restrictions required in this article: it runs on standard Borel spaces and is image finite. We recall that stochastic automata and similar models are used to give semantics to stochastic process algebras and specification languages (see D'Argenio (1999), Bravetti (2002), Bravetti and D'Argenio (2004), D'Argenio and Katoen (2005), Bohnenkamp *et al.* (2006), and so on), which, in turn, are used to model dynamic systems. Moreover, LMP-like models restricted to standard Borel spaces have been studied in Doberkat (2007).

Acknowledgments

The logic presented here has its roots in the Master's thesis Celayes (2006), which was jointly supervised by the first two coauthors of the current paper, and we thank Pablo Celayes for his initial colaboration. We are also grateful to Ignacio Viglizzo for fruitful discussions. In particular, he pointed out the connection between Definition 3.2 and hit-and-miss topologies.

References

Billingsley, P. (1995) *Probability and Measure*, Wiley-Interscience.

Bohnenkamp, H., D'Argenio, P. R., Hermanns, H. and Katoen, J.-P. (2006) MoDeST: A compositional modeling formalism for real-time and stochastic systems. *IEEE Trans. Softw. Eng.* **32** 812–830.

Bravetti, M. (2002) *Specification and Analysis of Stochastic Real-Time Systems*, Ph.D. thesis, Università di Bologna, Padova, Venezia.

Bravetti, M. and D'Argenio, P. R. (2004) Tutte le algebre insieme: Concepts, discussions and relations of stochastic process algebras with general distributions. In: Validation of Stochastic Systems. *Springer-Verlag Lecture Notes in Computer Science* **2925** 44–88.

van Breugel, F. (2005) The metric monad for probabilistic nondeterminism. Draft.

Cattani, S. (2005) *Trace-based Process Algebras for Real-Time Probabilistic Systems*, Ph.D. thesis, University of Birmingham.

Cattani, S., Segala, R., Kwiatkowska, M. Z. and Norman, G. (2005) Stochastic transition systems for continuous state spaces and non-determinism. In: Proceedings FoSSaCS 2005. *Springer-Verlag Lecture Notes in Computer Science* **3441** 125–139.

Celayes, P. (2006) Procesos de Markov etiquetados sobre espacios de Borel estándar. Master's thesis, FaMAF, Universidad Nacional de Córdoba.

Clarke, E. M., Grumberg, O. and Peled, D. (1999) *Model Checking*, MIT Press.

Danos, V., Desharnais, J., Laviolette, F. and Panangaden, P. (2006) Bisimulation and cocongruence for probabilistic systems. *Information and Computation* **204** 503–523.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press

D'Argenio, P. R. (1999) *Algebras and Automata for Timed and Stochastic Systems*, Ph.D. thesis, Department of Computer Science, University of Twente.

D'Argenio, P. R. and Katoen, J.-P. (2005) A theory of stochastic systems, Part I: Stochastic automata, *and* Part II: Process algebra. *Information and Computation* **203** 1–38, 39–74.

D'Argenio, P. R., Wolovick, N., Terraf, P. S. and Celayes, P. (2009) Nondeterministic labeled markov processes: Bisimulations and logical characterization. In: *Quantitative Evaluation of Systems, International Conference on*, IEEE Computer Society 11–20.

Desharnais, J. (1999) *Labeled Markov Process*, Ph.D. thesis, McGill University.

Desharnais, J. and Panangaden, P. (2003) Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *J. Log. Algebr. Program.* **56** (1-2) 99–115.

J. Desharnais, A. Edalat and P. Panangaden Bisimulation for labelled Markov processes. *Information and Computation* **179** (2) 163–193.

Doberkat, E.-E. (2007) *Stochastic relations. Foundations for Markov transition systems*, Studies in Informatics Series, Chapman and Hall.

Giry, M. (1981) A categorical approach to probability theory. In: Categorical Aspects of Topology and Analysis. *Springer-Verlag Lecture Notes in Mathematics* **915** 68–85.

van Glabeek, R. J. (2001) The linear time–branching time spectrum I. The semantics of concrete, sequential processes. In: Bergstra, J. A., Ponse, A. and Smolka, S. A. (eds.) *Handbook of Process Algebra*, North-Holland 3–99.

Kechris, A. S. (1995) *Classical Descriptive Set Theory*, Graduate Texts in Mathematics **156** Springer-Verlag.

Larsen, K. G. and Skou, A. (1991) Bisimulation through probabilistic testing. *Information and Computation* **94** (1) 1–28.

Milner, R. (1989) *Communication and Concurrency*, Prentice Hall.

Naimpally, S. (2003) What is a hit-and-miss topology? *Topological Commentary* **8** (1).

Parma, A. and Segala, R. (2007) Logical characterizations of bisimulations for discrete probabilistic systems. In: Proceedings FoSSaCS 07. *Springer-Verlag Lecture Notes in Computer Science* **4423** 287–301.

Puterman, M. L. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley-Interscience.

Segala, R. (1995) *Modeling and Verification of Randomized Distributed Real-Time Systems*, Ph.D. thesis, Massachusetts Institute of Technology.

Strulo, B. (1993) *Process Algebra for Discrete Event Simulation*, Ph.D. thesis, Department of Computing, Imperial College, University of London.

Vardi, M. Y. (1985) Automatic verification of probabilistic concurrent finite-state programs. In: *Proceedings 26th FOCS*, IEEE Computer Society 327–338.

Viglizzo, I. D. (2005) *Coalgebras on Measurable Spaces*, Ph.D. thesis, Department of Mathematics, Indiana University.

https://doi.org/10.1017/S0960129511000454 Published online by Cambridge University Press