A Logical Characterization of Bisimulation for Labeled Markov Processes

Josée Desharnais*
School of Computer Science
McGill University
Montreal, Quebec, Canada
desharna@cs.mcgill.ca

Abbas Edalat†
Department of Computing
Imperial College
London, UK
ae@doc.ic.ac.uk

Prakash Panangaden‡
School of Computer Science
McGill University
Montreal, Quebec, Canada
prakash@cs.mcgill.ca

Abstract

This paper gives a logical characterization of probabilistic bisimulation for Markov processes introduced in [5].

- *Bisimulation can be characterized by a very weak modal logic. The most striking feature is that one has no negation or any kind of negative proposition.*
- *Bisimulation can be characterized by several inequivalent logics; we report five in this paper and there are surely many more.*
- *We do not need any finite branching assumption yet there is no need of infinitary conjunction.*
- *We give an algorithm for deciding bisimilarity of finite state systems which constructs a formula that witnesses the failure of bisimulation.*

1. Introduction

This paper gives a logical characterization of probabilistic bisimulation (henceforth we say "bisimulation" for this) for Markov processes introduced in [5]. The thrust of that work was an extension of the notion of bisimulation to systems with continuous state spaces; for example for systems where the state space is the real numbers. As far as we know there is no other treatment of bisimulation for such continuous systems. One might have expected that the logical characterization, while potentially technically demanding, would contain no surprises and all the results could be predicted by analogy with the experience gained from the analysis of discrete probabilistic systems [21]. In fact the present investigation revealed several surprises *even for discrete systems*.

- Bisimulation, for *all* systems can be characterized by a very weak modal logic, in particular much weaker than used by Larsen and Skou for discrete systems. The most striking feature is that one has no negation at all. Even for discrete systems this is an unexpected result.
- Bisimulation can be characterized by several logics, we report 5 in this paper and we have no reason to think there aren't more. Curiously enough, none of these logics are equivalent.
- We do not need any finite branching assumption yet there is no need for infinitary conjunction. One is tempted to think, by analogy with nonprobabilistic bisimulation, that one must need infinitary conjunction if one has infinite branching.
- The proofs that we give are of an entirely different character than the typical proofs of these results for nonprobabilistic systems or discrete probabilistic systems. They use quite subtle facts about analytic spaces and appear at first sight to be entirely nonconstructive. Yet one can derive an algorithm for finite state systems which constructs a formula that witnesses the failure of bisimulation or establishes bisimulation, using *the weakest of our logics*. The algorithm itself is not surprising, it is a modification of Cleaveland's algorithm [7], but it is surprising that an algorithm exists at all given the nonalgorithmic route by which we arrived at these results.

A couple of other results are described in this paper.

- We show how to construct the maximal autobisimulation on a given system. This can be viewed as a couniversal construction. In the finite state case it is a state minimization construction. This answers a question raised by Kozen at the last LICS.
- We explore whether we can characterize bisimulation equivalence classes by a single formula; here the different logics exhibit their differences.

*Research supported by NSERC, Canada.
†Research supported by EPSRC, UK.
‡Research supported by NSERC, Canada.

In this paper we use the phrase "Markov process" to stand for what are called "stationary Markov processes" in the literature. We are also simplifying the formalism from the probability literature quite a bit in order to elide several of the complexities of stochastic process theory [6, 9, 11, 4]. These are dealt with fully in [12]. We use the phrase "Markov chain" for a Markov process with a discrete state space.

We now clarify exactly what we mean by "logical characterization" of bisimulation. First we clarify the difference between "bisimilar states" and "bisimilar processes". We use the word "process" (as in Markov process) to stand for the analogue of a labeled transition system. In the literature, the word "process" is used for what we call a "state" of a transition system. We define bisimulation as a relation between the states of two processes. We can say that two states are bisimilar if they are related by a bisimulation relation.

We say that a logic $\mathcal{L}$ gives a logical characterization of bisimulation if whenever we have two states, s, s' of a probabilistic labeled transition system (labeled Markov process) then s and s' are bisimilar if and only if they obey all the same formulas of $\mathcal{L}$. If s and s' are not bisimilar then we must have a formula which one satisfies and the other does not. We give an algorithm that constructs a formula which distinguishes two nonbisimilar states in a finite-state system using one of our logics - the weakest one.

Given a system we are also interested in characterizing the bisimulation equivalence classes. In other words we might want to know if, for each equivalence class, there is a formula ϕ, in whatever logic we are focusing on at the moment, such that a state is in the given equivalence class if and only if it satisfies the formula ϕ. It can happen that a logic is powerful enough to distinguish any two nonbisimilar states but not powerful enough to characterize equivalence classes in this way. Our weakest logic is an example of this phenomenon.

The motivation of our work is largely foundational. It is premature to assess the practical importance of this notion of bisimulation from the point of view of automated verification or of modeling. Indeed it is clear that pragmatic concerns like robustness and testability will cause us to modify our thinking, but the core theoretical investigations and mathematical techniques will survive. We have, however, begun two serious investigations into practical issues - investigating the notion of robustness and a case study with a speech recognition system. We expect that eventually making contact with probabilistic (or noisy) hybrid systems will really require some of the techniques of the present work, but the case studies need to reach a more mature stage before we can comment on applications. We have begun active collaboration with Gupta and Jagadeesan who have developed probabilistic hybrid concurrent constraint programming [13, 14] and used it for modeling studies. It is already clear that one will need continuous state spaces for realistic modeling of physical systems.

In our previous paper we worked with Polish spaces; these are the topological spaces underlying complete, separable metric spaces. In the present work we had to switch to analytic spaces. Analytic spaces are continuous images of one Polish space in another. It turns out that analytic spaces are more "stable" under certain constructions, particularly under the formation of quotients. Furthermore analytic spaces seem better adapted to logical investigations. We have redone the construction of [5] to work with analytic spaces and measurable functions rather than Polish spaces and continuous functions.

We made two claims in the previous paper that are wrong. We claimed that we could actually construct pullbacks, but this is not the case. This is really not relevant for the theory we have been developing but we have provided a simple counterexample in the full paper for interested readers. This does not affect the treatment of bisimulation - all we need is to be able to compose spans in order to show that bisimulation is an equivalence - and this is indeed correct without modification. We also began a preliminary investigation into the logical characterization of bisimulation. In particular, we showed that bisimilar systems satisfy all the same formulas of a certain logic (the logic $\mathcal{L}_\neg$ of this paper). In addition we claimed that if one of the systems was discrete then the reverse was true as well, i.e. that we had a logical characterization of bisimulation. This works only when both systems are discrete. An example appears in the full paper. This example is of interest in order to see why we moved from Polish spaces but plays no logical role in the present development.

2. Review of Labeled Markov Processes and Bisimulation

We review the definitions of labeled Markov processes and bisimulation. The results mentioned in this section generalize those of [5] and, by using some powerful results about analytic spaces, we can simplify some of the arguments of that paper. The results of that paper rely on a new result in probability theory due to Edalat [12] and of course we also rely on those results here as well.

A Markov process is a transition system with the property that the transition probabilities depend only on the current state and not on the past history of the process. We will consider systems where there is an interaction with the environment described by a set of labels as in process algebra. For each fixed label the system may undergo a transition governed by a transition probability.

Definition 2.1 *A **partial, labeled, Markov process** with label set L is a structure $(S, \Sigma, \{k_l\}_{l\in\mathsf{L}})$, where S is the set of* *states, which is assumed to be an **analytic** space, and Σ is the Borel σ-field on S, and*

$$\forall l \in \mathsf{L},\ k_l : S \times \Sigma \longrightarrow [0,1]$$

is a transition sub-probability function, i.e., for each fixed $s \in S$, the set function $k_l(s,\cdot)$ is a sub-probability measure and for each fixed $A \in \Sigma$ the function $k_l(\cdot, A)$ is a measurable function (with respect to the σ-field on $[0,1]$ generated by the open intervals).

We will fix the label set to be some L once and for all. The word "partial" is to emphasize the fact that $k_l(s,\cdot)$ is any *sub*-probability measure for all $s \in S$, i.e. $k_l(s,S) \leq 1$. A *total* labeled, Markov process would satisfy that $k_l(s,\cdot)$ is either a probability measure or the null distribution (i.e. $k_l(s,S) = 1$ or 0). From now on we assume that all systems are partial and we will stop writing the adjective "partial" explicitly.

A binary relation $R \subseteq S_1 \times S_2$ is just a set of ordered pairs and as such can be conveniently viewed as a set equipped with the two projection maps $S_1 \leftarrow R \rightarrow S_2$. Such a structure is called a *span*. Indeed, for any relation R between two sets S_1 and S_2, the set of ordered pairs $\{(s_1,s_2) \in S_1 \times S_2 \mid s_1 R s_2\}$ together with the projection morphisms is a span between S_1 and S_2; conversely, given a span T, f_1, f_2, we can define the relation R as $s_1 R s_2$ if there is a $t \in T$ such that $f_1(t) = s_1$ and $f_2(t) = s_2$. If the morphisms of the span are surjective, every element of one set is related to an element of the other set. One can talk about binary relations in arbitrary categories by talking about spans. A span between an object S_1 and another object S_2 is a third object T together with morphisms from T to both S_1 and S_2. Accordingly, our definition of bisimulation will be in terms of a span of morphisms, as was initially done by Joyal, Nielsen and Winskel [20]. They observed that Milner's non-probabilistic strong bisimulation can be expressed in this formalism.

Definition 2.2 *A **zigzag morphism**[1] f between two labeled, Markov processes, $(S,\Sigma,\{k_l\}_{l\in\mathsf{L}})$ and $(S',\Sigma',\{k'_l\}_{l\in\mathsf{L}})$ is a (Borel) measurable function $f : (S,\Sigma) \longrightarrow (S',\Sigma')$ such that*

- *f is surjective*
- *$\forall l \in \mathsf{L}.\forall x \in S.\forall \sigma' \in \Sigma'$.*

$$k_l(x, f^{-1}(\sigma')) = k'_l(f(x),\sigma').$$

By demanding that zigzag morphisms be surjective we avoid situations like the following. Suppose that we have a system S and another system S' which consists of a disjoint union of a copy of S and some other completely unrelated subsystem, say S''. Now if we did not insist on surjectivity we could have a zigzag from S to S', a situation that we want to avoid.

[1] The curious word "zigzag" comes from modal logic; the idea is that they capture the "back and forth" nature of bisimulation.

Definition 2.3 *The category **LMP** has labeled Markov processes as objects and the zigzag morphisms defined above as morphisms.*

Note that a diagram of morphisms in **LMP** will commute if and only if the corresponding diagram commutes in **Sets**. One can immediately check that the identity function is a zigzag.

We often refer to a labeled Markov process by its set of states.

Definition 2.4 *Let T and T' be two labeled Markov processes. T is **probabilistically bisimilar to** T' (written $T \sim T'$) if there is a span of zigzag morphisms between them, i.e. there exists a labeled Markov process S and zigzag morphisms $f : S \longrightarrow T$ and $g : S \longrightarrow T'$.*

It is proven in [5] (using the result of [12]) that bisimulation is an equivalence relation. The basic construction used from [12] is the *semi-pullback* construction - a construction which produces a span from a cospan but without the universal property of pullbacks.

3. Modal Logics

We now describe five logics that will each be proven to characterize bisimulation. Thus *all* these logics play the role of Hennessy-Milner logic for nonprobabilistic bisimulation.

We use letters like a or b for actions. The simplest logic will be called $\mathcal{L}_0$ and has as syntax the following formulas:

$$\mathsf{T} | \phi_1 \wedge \phi_2 | \langle a \rangle_q \phi$$

where a is an action from the fixed set of actions L and q is a rational number. Given a labeled Markov process $(S,\Sigma,\{k_a\}_{a\in\mathsf{L}})$ we write $s \models \phi$ to mean that the state s satisfies the formula ϕ. The definition of the relation $\models$ is given by induction on formulas, the only nontrivial case is the modal formula. We say $s \models \langle a \rangle_q \phi$ iff $\exists A \in \Sigma.(\forall s' \in A.s' \models \phi) \wedge (k_a(s,A) \geq q)$. In other words, the system in state s can make an a-move to a state that satisfies ϕ with probability greater than q. We write $[\![\phi]\!]_S$ for the set $\{s \in S | s \models \phi\}$. We often omit the subscript when no confusion can arise.

In the following table we define four additional logics. They are all syntactic extensions of $\mathcal{L}_0$.

$$\begin{array}{rcl} \mathcal{L}_{\text{Can}} & := & \mathcal{L}_0 | \text{Can}(a) \\ \mathcal{L}_{\Delta} & := & \mathcal{L}_0 | \Delta_a \\ \mathcal{L}_{\neg} & := & \mathcal{L}_0 | \neg\phi \\ \mathcal{L}_{\bigwedge} & := & \mathcal{L}_{\neg} | \bigwedge_{i\in\mathbf{N}} \phi_i \end{array}$$

Given a labeled Markov process $(S, \Sigma, \{k_a\}_{a \in \mathsf{L}})$ we write:

$s \models \mathrm{Can}(a)$	to mean that $k_a(s, S) > 0$;
$s \models \Delta_a$	to mean that $k_a(s, S) = 0$;
$s \models \neg\phi$	to mean that $s \not\models \phi$;
$s \models \bigwedge_{i \in \mathbf{N}} \phi_i$	to mean that $s \models \phi_i$ for all $i \in N$.

Although they all characterize bisimulation, they do not have the same expressive power. Clearly all of them are at least as expressive as $\mathcal{L}_0$, and $\mathcal{L}_\wedge$ is more expressive than all the others. $\mathcal{L}_{\mathrm{Can}}$, $\mathcal{L}_\Delta$ and $\mathcal{L}_\neg$ are incomparable. It is interesting to note that none of these differences will have any impact on the characterization of bisimulation, as we have already said.

The logic that Larsen and Skou used in [21] is $\mathcal{L}_\Delta$ with the additional formula $\phi_1 \vee \phi_2$. They show that for finitely branching systems[2], two *states* of the same system are bisimilar if and only if they satisfy the same formulas of their logic.

The fact that a logic without negation and without infinitary conjunction is sufficient for systems with infinite branching is somewhat of a surprise based on what we expect from the non-probabilistic case.

Consider the nonprobabilistic processes $a.0 + a.b.0$ and $a.b.0$ in a CCS-like notation (with 0 for **NIL**). It is well-known that they cannot be distinguished by a negation-free formula of Hennessy-Milner logic. The usual formula is $\langle a \rangle \neg \langle b \rangle \mathsf{T}$, which says that the process can perform an a action and then be in a state where it cannot perform a b-action. For no assignment of probabilities are the two processes bisimilar, unless the first branch of the first process has probability 0, in which case it is *identical* to the second process. Suppose that the two a-labeled branches of the first process are given probabilities p and q, assume that the b-labeled transitions have probability 1. Now if the second process has its a-labeled transition given a probability anything other than $p+q$, say $r > p+q$ we can immediately distinguish the two processes by the formula $\langle a \rangle_r \mathsf{T}$. If $r = p + q$ then we can use the formula $\langle a \rangle_r \langle b \rangle_1 \mathsf{T}$. The first process cannot satisfy this formula. This simple example shows that one can use the probabilities to finesse the need for negation but, as we shall see, one cannot actually encode negation with just $\mathcal{L}_0$.

The next example indicates why we do not need infinite conjunctions even if we have infinite branching. Consider a process with infinitely many a-labeled branches. The first branch ends in a state that can perform no further actions, call it a "dead state." The second branch ends in a state that can perform a single a to a dead state. Similarly the nth branch can perform a sequence of n a-actions and then reach a dead state. Call this process P. Now define a process Q which is just like P except that there is an additional transition to a state which then has an a-labeled transition back to itself. Now consider the formula

$$\langle a \rangle \left(\bigwedge_n \langle a \rangle^{(n)} \mathsf{T} \right)$$

where the notation $\langle a \rangle^{(n)}$ means n nested $\langle a \rangle$ modalities. The conjunction is over all $n \geq 1$. This formula says that the process can jump to a state from which arbitrarily many a-labeled transitions are possible. The process P does not satisfy this formula but Q does. Now if we associate probabilities with these transitions we find that we can find distinguishing formulas that do not involve infinite conjunction. To see this assume that both processes satisfy all the same $\mathcal{L}_0$ formulas. By induction it follows that each branch in P must have the same probability as the equal length branch in Q. Thus the branch to the looping state in Q must have probability 0, in which case Q is identical to P.

4. $\mathcal{L}_0$ characterizes Bisimulation

In this section we prove the main theorem of the paper. The proof relies on various properties of analytic spaces. We give an overview of the proof and give the full proof in the full paper. We have included the proofs of the most interesting propositions here.

To show that two bisimilar states satisfy all the same formulas of $\mathcal{L}_0$ is relatively easy and we had included this in our earlier paper [5]. To show the converse the general plan is to construct a cospan using logical equivalence and then to use the semi-pullback construction [12] to obtain a span. To obtain the cospan one defines an equivalence relation on states - two states are equivalent if they satisfy the same formulas - and then form the quotient. We need a general theorem to assure us that the result is analytic. It was here that we had to move to analytic spaces because taking quotients of Polish spaces may not yield Polish spaces. We then define a transition probability on this quotient system in such a way as to ensure that the morphisms are zigzags. This is the part of the construction where we need most of the measure-theoretic machinery. We use a *unique structure theorem* to show that the measurable sets defined by the formulas of the logic generate the σ-field. We use a theorem on *unique extension of measure* in order to show that the transition probability is well-defined.

The first proposition below says that sets of states definable by formulas are always measurable.

Proposition 4.1 *For all formulas ϕ we have $[\![\phi]\!] \in \Sigma$.*

The next proposition links zigzag morphisms with formulas in the logic.

[2]They actually use a stronger property, the "minimum deviation condition" which uniformly bounds the degree of branching everywhere.

Proposition 4.2 *If f is a zigzag morphism from S to S' then for all state $s \in S$ and all formulas ϕ,*

$$s \models \phi \iff f(s) \models \phi.$$

Corollary 4.3 *If S and S' are bisimilar, then they satisfy the same formulas.*

In order to show the logic gives a complete characterization of bisimulation, we also want to show the converse.

Definition 4.4 *We say that two states, say s of system S and s' of system S', are $\mathcal{L}_0$-**equivalent** (written as $(s, S) \approx (s', S')$ or just $s \approx s'$ for short) if for every formula ϕ of $\mathcal{L}_0$ we have $s \models \phi$ if and only if $s' \models \phi$.*

We first show that there is a zigzag morphism from any system S to its quotient under $\approx$. If (S, Σ) is a Borel space, the quotient $(S/_\approx, \Sigma_\approx)$ is defined as follows. $S/_\approx$ is the set of all equivalence classes. Then the function $q : S \to S/_\approx$ which assigns to each point of S the equivalence class containing it maps onto $S/_\approx$, and thus determines a Borel structure on $S/_\approx$: by definition a subset E of $S/_\approx$ is a Borel set if $q^{-1}(E)$ is a Borel set in S.

Proposition 4.5 *Let $(S, \Sigma, \{k_a\}_{a \in \mathsf{L}})$ be a **LMP**. We can define h_a so that the canonical projection q from $(S, \Sigma, \{k_a\}_{a \in \mathsf{L}})$ to $(S/_\approx, \Sigma_\approx, \{h_a\}_{a \in \mathsf{L}})$ is a zigzag morphism.*

In order to prove this proposition we need a few lemmas. The first allows us to work with direct images of q. The second is elementary, while the next two are known results about analytic spaces. The final lemma is a standard uniqueness theorem.

Lemma 4.6 *$q^{-1}q[\![\phi]\!] = [\![\phi]\!]$ for each formula ϕ of the logic.*

Lemma 4.7 *Let $(S, \Sigma, \{k_a\}_{a \in \mathsf{L}})$ and $(S', \Sigma', \{k'_a\}_{a \in \mathsf{L}})$ be two systems. Then for all formulas ϕ and all pairs (s, s') such that $s \approx s'$, we have $k_a(s, [\![\phi]\!]_S) = k'_a(s', [\![\phi]\!]_{S'})$.*

The next lemmas are Theorem 3.3.5 of [1] and one of its corollaries.

Lemma 4.8 *Let $(X, \mathcal{B})$ be an analytic Borel space and let $\mathcal{B}_0$ be a countably generated sub-σ-field of $\mathcal{B}$ which separates points in X. Then $\mathcal{B}_0 = \mathcal{B}$.*

Lemma 4.9 *Let X be an analytic Borel space and let $\sim$ be an equivalence relation in X. Assume there is a sequence $f_1, f_2, \ldots$ of real valued Borel functions on X such that for any pair of points x, y in X one has $x \sim y$ iff $f_n(x) = f_n(y)$ for all n. Then $X/_\sim$ is an analytic Borel space.*

The final lemma that we need is a result which gives a condition under which two measures are equal. It is Theorem 10.4 of Billingsley [4].

Lemma 4.10 *Let X be a set and $\mathcal{A}$ a family of subsets of X, closed under finite intersections, and such that X is a countable union of sets in $\mathcal{A}$. Let $\sigma(\mathcal{A})$ be the σ-field generated by $\mathcal{A}$. Suppose that μ_1, μ_2 are finite measures on $\sigma(\mathcal{A})$. If they agree on $\mathcal{A}$ then they agree on $\sigma(\mathcal{A})$.*

Proof of Proposition 4.5: We first show that $S/_\approx$ is an analytic space. Let $\{\phi_i | i \in \mathbf{N}\}$ be the set of all formulas. We know that $[\![\phi_i]\!]_S$ is a Borel set for each i. Therefore the characteristic functions $\chi_{\phi_i} : S \to \{0, 1\}$ are Borel measurable functions. Moreover we have

$$\begin{aligned} x \approx y \quad &\text{iff} \quad (\forall i \in \mathbf{N}.\ x \in [\![\phi_i]\!]_S \iff y \in [\![\phi_i]\!]_S) \\ &\text{iff} \quad (\forall i \in \mathbf{N}.\ \chi_{\phi_i}(x) = \chi_{\phi_i}(y)). \end{aligned}$$

It now follows by Lemma 4.9 that $S/_\approx$ is an analytic space.

Let $\mathcal{B} = \{q([\![\phi_i]\!]_S) : i \in \mathbf{N}\}$. We show that $\sigma(\mathcal{B}) = \Sigma_\approx$. We have $\mathcal{B} \subseteq \Sigma_\approx$, since, by Lemma 4.6, for any $q([\![\phi_i]\!]_S) \in \mathcal{B}$, $q^{-1}q([\![\phi_i]\!]_S) = [\![\phi_i]\!]_S$ which is in Σ by Lemma 4.1. Now $\sigma(\mathcal{B})$ separates points in $S/_\approx$, for if x and y are different states of $S/_\approx$, take states $x_0 \in q^{-1}(x)$ and $y_0 \in q^{-1}(y)$. Then since $x_0 \not\approx y_0$, there is a formula ϕ such that x_0 is in $[\![\phi]\!]_S$ and y_0 is not. By Lemma 4.6, it follows that x is in $q[\![\phi]\!]_S$, whereas y is not. Since $\sigma(\mathcal{B})$ is countably generated, it follows by Lemma 4.8, that $\sigma(\mathcal{B}) = \Sigma_\approx$.

We are now ready to define $h_a(t, \cdot)$ over $\Sigma_\approx$ for $t \in S/_\approx$. We define it so that $q : S \to S/_\approx$ is a zigzag morphism (recall that q is measurable and surjective by definition), i.e., for any $B \in \Sigma_\approx$ we put

$$h_a(t, B) = k_a(s, q^{-1}(B)),$$

where $s \in q^{-1}(t)$. Clearly, for a fixed state s, $k_a(s, q^{-1}(\cdot))$ is a sub-probability measure on $\Sigma_\approx$. We now show that the definition does not depend on the choice of s in $q^{-1}(t)$ for if $s, s' \in q^{-1}(t)$, we know that $k_a(s, q^{-1}(\cdot))$ and $k_a(s', q^{-1}(\cdot))$ agree over $\mathcal{B}$ again by the fact that $q^{-1}q([\![\phi_i]\!]_S) = [\![\phi_i]\!]_S$ and by Lemma 4.7. So, since $\mathcal{B}$ is closed under the formation of finite intersections we have, from Lemma 4.10, that $k_a(s, q^{-1}(\cdot))$ and $k_a(s', q^{-1}(\cdot))$ agree on $\sigma(\mathcal{B}) = \Sigma_\approx$.

It remains to prove that for a fixed Borel set B of $\Sigma_\approx$, $h_a(\cdot, B) : S/_\approx \to [0, 1]$ is a Borel measurable function. Let A be a Borel set of $[0, 1]$. Then $h_a(\cdot, B)^{-1}(A) = q[k_a(\cdot, q^{-1}(B))^{-1}(A)]$; we know that $C = k_a(\cdot, q^{-1}(B))^{-1}(A)$ is Borel since it is the inverse image of A under a Borel measurable function. Now we have that $q(C) \in \Sigma_\approx$, since $q^{-1}q(C) = C$: indeed, if $s_1 \in q^{-1}q(C)$, there exists $s_2 \in C$ such that $q(s_1) = q(s_2)$, and we have just proved above that then the $k_a(s_i, q^{-1}(\cdot))$'s must agree, so if $k_a(s_i, q^{-1}(B)) \in A$ for $i = 2$, then it is also true for $i = 1$, so $s_1 \in C$ as wanted. So $h_a(\cdot, B)$ is Borel measurable. This concludes the proof that $S/_\approx$ is a **LMP** and q a zigzag morphism. ■

We now state the main result of this paper.

Theorem 4.11 *Two systems are bisimilar iff they obey the same formulas of our logic.*

Proof. It remains to prove the if part. Suppose that $(S, \Sigma, \{k_a\}_{a\in\mathsf{L}})$ and $(S', \Sigma', \{k'_a\}_{a\in\mathsf{L}})$ satisfy the same formulas. Instead of defining a span of zigzags directly, we can define a cospan and use the semi-pullback property to infer that S and S' are bisimilar. We first construct a system, $(T, \Sigma_T, \{j_a\}_{a\in\mathsf{L}})$, called the *direct sum* of S and S', as follows. We set $T = S \uplus S'$ with the evident σ-field. We define the transition probabilities as follows: $j_a(s, A \uplus A') = k_a(s, A)$ if $s \in S$ and $j_a(s', A \uplus A') = k'_a(s', A')$ if $s' \in S'$ where $A \in \Sigma$ and $A' \in \Sigma'$. There are the evident canonical injections ι, ι' which are *not* zigzags because they are not surjective. We know that the quotient system $(T/_\approx, \Sigma_\approx, \{h_a\}_{a\in\mathsf{L}})$ is a **LMP** and that the canonical projection r from T to $T/_\approx$ is a zigzag morphism. Thus we have morphisms $q : S \longrightarrow T$ and $q' : S' \longrightarrow T$ given by the composites $r \circ \iota$ and $r \circ \iota'$, which are measurable and surjective. To see that q, q' are surjective we recall that an equivalence class must, by hypothesis, include members of both S and S'. It remains to prove the zigzag property for q and q'. So take a set B in $\Sigma_\approx$ and $s \in S$. Then

$$\begin{aligned} h_a(q(s), B)) &= j_a(\iota(s), r^{-1}(B)) \\ &= k_a(s, r^{-1}(B) \cap S) \\ &= k_a(s, q^{-1}(B)) \end{aligned}$$

This proves that q and similarly q' are zigzag morphisms. Thus we have defined a cospan of zigzag morphisms and using the semi-pullback theorem there is a corresponding span, hence S and S' are bisimilar. ■

Note that this implies the result for discrete systems without using the minimum deviation assumption used by Larsen and Skou.

Definition 4.12 *Given two bisimilar systems S and S', we say that two states $s \in S$ and $s' \in S'$ are bisimilar, denoted $s \sim s'$, if there is a span $f : U \longrightarrow S, g : U \longrightarrow S'$ such that for some $u \in U$ we have $f(u) = s$ and $g(u) = s'$.*

It follows from the existence of semi-pullbacks that $\sim$ is an equivalence relation and we have:

Corollary 4.13 *$s \sim s$ if and only if $s \approx s'$.*

The quotient construction has the following couniversal property. In the case of finite state systems this says that the quotienting construction gives the minimal finite state system bisimilar to the given one.

Proposition 4.14 *If $S \sim T$, then there exists a unique zigzag morphism r from S to $T/_\approx$.*

Proof. If $(S, \Sigma, \{k_a\}_{a\in\mathsf{L}}) \sim (T, \Sigma_T, \{l_a\}_{a\in\mathsf{L}})$, there is a span $(U, \Sigma_U, \{j_a\}_{a\in\mathsf{L}})$ with zigzag morphisms f and g. There is also a zigzag morphism $q : T \to (T)/_\approx$ and we show how to define $r : S \to (T)/_\approx$ as in the diagram:

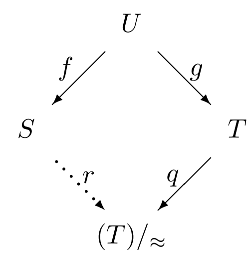

Let $s \in S$. Since f, g and q are zigzag morphisms, for every formula ϕ we have,

$$\begin{aligned} s \models \phi &\iff \forall u \in f^{-1}(s).u \models \phi \\ &\iff \forall u \in f^{-1}(s).g(u) \models \phi \\ &\iff \forall u \in f^{-1}(s).qg(u) \models \phi \end{aligned}$$

This implies that all $u \in f^{-1}(s)$ are mapped by qg to the same state $t \in T/_\approx$ and that we can set $r(s) = t$. This makes the diagram commute. Surjectivity of r is obvious; to see it is also Borel measurable, let $A \in \Sigma_\approx$. Then $r^{-1}(A) = f(g^{-1}q^{-1}A)$, $B_1 := g^{-1}q^{-1}A$ is obviously Borel in U and so is $B_2 := g^{-1}q^{-1}A^c$. Now we not only have that B_1 and B_2 are disjoint, but their images under f are also disjoint. To see this, suppose the contrary. Then there exist $u_i \in B_i$ such that $fu_1 = fu_2$; but since the diagram commutes, it implies that $qg(u_1) = qg(u_2)$, which is a contradiction to the definition of the B_i's. Thus we have that fB_1 and fB_2 are disjoint analytic sets of S; since analytic sets are separable by Borel sets and since $fB_1 \cup fB_2 = S$, fB_1 and fB_2 must be Borel sets of S, concluding the proof that r is Borel measurable. We now show that r is zigzag; let $s \in S, A \in \Sigma_\approx$ and $u \in f^{-1}(s)$. Then

$$\begin{aligned} h_a(r(s), A) &= l_a(g(u), q^{-1}A) \\ &= j_a(u, g^{-1}q^{-1}A) \\ &= j_a(u, f^{-1}r^{-1}A) \\ &= k_a(s, r^{-1}A) \end{aligned}$$

Finally, r is unique because every state s is mapped in $T/_\approx$ to the only state that satisfies the same formulas as it does (since in $T/_\approx$ there is no pair of distinct states satisfying the same formulas). ■

Corollary 4.15 *$S \sim T$ if and only if $S/_\approx \cong T/_\approx$*

Now we consider the other logics.

Proposition 4.16 *All the logics defined in the previous section characterize bisimulation.*

Proof (sketch). We must show that all results remain true if we consider any of the above four logics instead of $\mathcal{L}_0$. We begin with Proposition 4.1, i.e. we show that for all formulas ϕ of all our logics, we have $[\![\phi]\!] \in \Sigma$. $[\![\mathrm{Can}(a)]\!] = k_a(\cdot, S)^{-1}((0,1])$ which is in Σ and $[\![\Delta_a]\!]$ is its complement and hence is in Σ. Now for $\mathcal{L}_\wedge$ and $\mathcal{L}_\neg$ we only have to show that if $[\![\phi]\!] \in \Sigma$, then so is $[\![\neg\phi]\!]$ which is straightforward, and if $\forall i \in \mathbf{N}[\![\phi_i]\!] \in \Sigma$ then so is $[\![\bigwedge_{i\in\mathbf{N}} \phi_i]\!]$ which is also straightforward since Σ is a σ-field. The results follow by structural induction. We now prove the important part of Proposition 4.2, namely that for every zigzag morphism $f : S \to S'$ and every state $s \in S$, s and $f(s)$ satisfy all the same formulas. For $\mathcal{L}_{\mathrm{Can}}$ and $\mathcal{L}_\Delta$, since f is zigzag, $k_a(s,S) = k'_a(f(s),S')$, so $s \models \mathrm{Can}(a)$ if and only if $s' \models \mathrm{Can}(a)$, and $s \models \Delta_a$ if and only if $s' \models \Delta_a$. For $\mathcal{L}_\wedge$ and $\mathcal{L}_\neg$ it is obvious by structural induction. Now for Proposition 4.5, we only have to show that Lemma 4.7 remains true which is obvious. Finally it is easy to show that Theorem 4.11 is not affected by the addition of the new formulas. ∎

Corollary 4.17 *The relation $\approx$ and hence the equivalence classes in each system are the same in all the logics described above.*

Although these logics all characterize bisimulation, they do not all characterize equivalence classes, in the sense that there does not necessarily exist a formula for each equivalence class which is satisfied only by states in that class. The most powerful logic does characterize equivalence classes.

Proposition 4.18 *The logic $\mathcal{L}_\wedge$ characterizes equivalence classes of states in arbitrary Markov processes. Given a finite-state systems $\mathcal{L}_\neg$ can characterize equivalence classes of states in that system.*

This proposition is false if we consider the problem of writing a formula characterizes equivalence classes of arbitrary finite Markov chains and not just the equivalence classes of states within a fixed Markov chain. To be more precise, suppose that we look at all processes, with an initial state, and we demand that bisimulation relate the initial states. Now we want to know if, given a bisimulation equivalence class there is a formula such that the initial state of a process satisfies the formula iff it is in the bisimulation equivalence class.

Proposition 4.19 *There is no formula of $\mathcal{L}_\neg$ that characterises equivalence classes of finite Markov chains with initial states.*

The same argument as used in the proof of the above proposition can be used to show that neither $\mathcal{L}_\neg$ nor $\mathcal{L}_\Delta$ can characterize equivalence classes inside a discrete system even if they satisfy the minimal deviation assumption defined by Larsen and Skou.

5. Deciding Bisimilarity for Finite-state systems

The proof of the last section uses machinery that is unconventional in concurrency theory. The result gives a characterization of bisimulation in terms of a logic without negation, an unexpectedly weak logic. It is natural to question the constructive content of such a proof. Now a careful analysis of constructivity in measure theory is beyond the scope of this work, but we will show that in the case of finite state systems there is an algorithm which decides bisimilarity of finite Markov chains and also produces a witnessing formula *from the logic* $\mathcal{L}_0$ in case the systems are not bisimilar. The algorithm is a modification of Cleaveland's algorithm [7].

The algorithm shown in figure 1 allows us to distinguish states that do not satisfy the same formulas. Beginning with D_1 containing only the set S of all states, bisim operates as follows: for each $a \in \mathsf{L}$ and B' in D_1, it "splits" every set B of D_1 into subsets having the same probability of jumping to B' with action a; this is done until it does not modify D_1. At the end, all states satisfying the same formulas will belong to exactly the same sets in D_1.

The function split first lists in L all possible values of $P_a(b, B')$ for $b \in B$. Then, for every possible value in the list L, the subset of states in B that can jump into B' with probability at least this value are added to the set D which contains B at the beginning. Before adding a set to D, split checks if it is already a member of $D \cup D_1 \cup D_2$; if not, it defines the formula represented by the set and adds the set to D; otherwise it does not add the set to D. This ensures that we do not compute a new formula for a set that already has a shorter one.

As shown in the following proposition, given two states, we can determine if they satisfy all the same formulas by checking if they belong to the same sets of D_1. If they do not, by finding the "first" set B in D_1 that distinguishes them, we get a formula $F(B)$ that also distinguishes them. The algorithm itself doesn't record an order of creation on the sets B, but it could be easily modified to do so.

Proposition 5.1 *Two states satisfy the same formulas iff they belong to exactly the same sets in D_1 at the end of executing the algorithm above.*

```
bisim(S, L, D)
begin
  F(S) := T
  D_1 := {S}
  D_2 := ∅
  while D_1 ≠ D_2 do
    for each a ∈ L and B' ∈ D_1 do begin
      D_2 := D_1
      D_1 := ∅
      for each B ∈ D_2 do D_1 := D_1 ∪ split(B, a, B')
    end
end

split(B, a, B')
begin
  L := ∅
  D := {B}
  for each b ∈ B do L := L ∪ {P_a(b, B')}
  L := L \ {0}
  while L ≠ ∅ do begin
    l := min L
    L := L \ {l}
    B_1 := {b ∈ B : P_a(b, B') ≥ l}
    if B_1 ∉ D ∪ D_1 ∪ D_2 then begin
      F(B_1) := F(B) ∧ ⟨a⟩_l F(B')
      D := D ∪ B_1
    end
  end
  return D
end
```

Figure 1. The bisimulation algorithm

6. Related Work

There are several papers now on probabilistic analysis, modeling and verification. There are even several papers on probabilistic process algebra analyzing notions of testing and simulation, investigating model checking and exploring various other ideas [23, 24, 17, 18, 19, 8, 2, 16]. One of the most stimulating recent results on applications is the work of Hillston [15] on PEPA and its uses in compositional performance evaluation. There are also several other interesting practical developments which are worthy of attention. In particular telecommunication [22], real-time systems [3] and modeling physical systems [14] are areas where probabilistic systems are very important. It is particularly for the last type of application that we expect that the continuous space formalism developed here will be useful.

In this section we will focus the discussion on closely related papers. The fundamental work on this topic is the paper by Larsen and Skou [21] which analyzes not just bisimulation but also testing. Our work extends theirs in two fundamental ways. It applies to continuous state space systems. The mathematics is therefore entirely different and even the structure of the arguments is different. If we specialize to the discrete case we have extended their work in two ways, we characterize bisimulation with a logic that is weaker, it has no negative formulas or disjunction and we do not have any finite branching assumption or minimum deviation assumption. Of course the minimum deviation assumption is a reasonable one in the context of discrete systems but it is interesting to note that it is not needed.

The other related paper is the work of de Vink and Rutten [10]. This is also an investigation into the realm of continuous state spaces. Their coalgebraic approach is definitely attractive and it should be interesting to explore if there is any way of extending their results to metric spaces like the reals. In particular there should be interesting links between logic and coalgebras. However they work with ultrametric spaces, not with the kind of metric spaces that actually arise in physical examples. Thus, for example, the real numbers do not form an ultrametric space and their results do not apply there; nor is any easy extension likely to work. Furthermore they only are able to show that bisimulation is an equivalence relation in the case that the space is discrete.

The algorithm owes a lot to the treatment of Cleaveland [7]. His algorithm is based on working with partitions, and if the logic has negation this is an obvious strategy. Our algorithm for $\mathcal{L}_0$ is based on using a nested family of sets rather than a partition. Of course the main interest in this algorithm is that it exists at all given the apparently nonconstructive arguments used in the main theorem.

7. Conclusions

In this paper we have shown that a very weak negation-free logic characterizes probabilistic bisimulation for Markov processes with continuous or discrete state spaces. Such systems exhibit possibly infinite branching yet no infinite conjunction. The proofs are based on properties of analytic spaces and appear at first sight to be nonconstructive. However in the case of finite state systems we in fact have an algorithm for deciding bisimilarity which constructs a distinguishing formula when the processes are not bisimilar. The fact that the logic is so simple shows that the distinguishing formulas are simple. In fact, based on our limited analysis, there appears to be no complexity or size advantage to using the more complex logics.

The fundamental point that we realized is that analytic spaces - despite their richness and the fact that they include many continuous spaces - are very closely linked to logic. In our previous work we used Polish spaces and continuous maps and this accords well with physical intuitions but we did not get a very good match with logic. The second important point is that probabilistic systems are in fact very close to determinate systems. The fact that bisimilarity is so simply characterized shows that for probabilistic systems the bisimilarity is not very far from some kind of trace equivalence. In the nonprobabilistic case we know, of course, that for determinate systems trace equivalence and bisimulation coincide.

The results in this paper are encouraging from the point of view of developing algorithms. At first it seems easy to believe that measure theory on continuous state spaces is so inherently nondiscrete that there will not be any reasonable links with logical concepts. In fact, in the proper setting, namely analytic spaces, we find that the discrete and the continuous mesh very smoothly.

Acknowledgments

We have benefited from discussions with Samson Abramsky, D. Sangiorgi and Franck van Breugel. We are indebted to Jørgen Hoffman-Jørgensen for suggesting that we work with analytic spaces. Prakash Panangaden would like to thank BRICS for generous hospitality during his sabbatical year during which this work began.

References

[1] W. Arveson. *An Invitation to C^*-Algebra*. Springer-Verlag, 1976.

[2] C. Baier and M. Kwiatkowska. Domain equations for probabilistic processes. Available from URL http://www.cs.bham.ac.uk/ mzk/, March 1996.

[3] A. Benveniste, B. C. Levy, E. Fabre, and P. L. Guernic. A calculus of stochastic systems for the specification, simulation and hidden state estimation of mixed stochastic/nonstochastic systems. *Theoretical Computer Science*, 152(2):171–217, 1995.

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

[5] R. Blute, J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled markov processes. In *Proceedings of the Twelfth IEEE Symposium On Logic In Computer Science, Warsaw, Poland.*, 1997.

[6] L. Breiman. *Probability*. Addison-Wesley, 1968.

[7] R. Cleaveland. On automatically explaining bisimulation inequivalence. In E. Clarke and R. Kurshan, editors, *Computer-Aided Verification CAV 90*, number 531 in Lecture Notes in Computer Science, pages 364–372, 1990.

[8] R. Cleaveland, S. Smolka, and A. Zwarico. Testing preorders for probabilistic processes. In *Proceedings of the International Colloquium On Automata Languages And Programming 1992*, number 623 in Lecture Notes In Computer Science. Springer-Verlag, 1992.

[9] D. R. Cox and H. D. Miller. *The Theory of Stochastic Processes*. Chapman and Hall, 1965.

[10] E. de Vink and J. J. M. M. Rutten. Bisimulation for probabilistic transition systems: A coalgebraic approach. In *Proceedings of the 24th International Colloquium On Automata Languages And Programming*, 1997.

[11] J. L. Doob. *Stochastic Processes*. Wiley Classics Library. Wiley-Interscience, 1953.

[12] A. Edalat. Semi-pullbacks in cartegories of markov processes. Submitted to Mathematical Structures in Computer Science. Available in WWW from `http://theory.doc.ic.ac.uk:80/~ae`, 1997.

[13] V. Gupta, R. Jagadeesan, and V. Saraswat. Probabilistic concurrent constraint programming. In A. Mazurkiewicz and J. Winkowski, editors, *Proceedings of CONCUR97*, volume 1243 of *Lecture Notes In Computer Science*. Springer-Verlag, 1997.

[14] V. Gupta, V. Saraswat, and P. Struss. A model of a photocopier paper path. In *Proceedings of the 2nd IJCAI Workshop on Engineering Problems for Qualitative Reasoning*, 1995.

[15] J. Hillston. *A Compositional Approach to Performance Modelling*. PhD thesis, University of Edinburgh, 1994. To be published as a Distinguished Dissertation by Cambridge University Press.

[16] M. Huth and M. Kwiatkowska. On probabilistic model checking. Technical Report CSR-96-15, University of Birmingham, 1996. available from http://www.cs.bham.ac.uk/ mzk/.

[17] B. Jonsson and K. Larsen. Specification and refinement of probabilistic processes. In *Proceedings of the 6th Annual IEEE Symposium On Logic In Computer Science*, 1991.

[18] B. Jonsson and W. Yi. Compositional testing preorders for probabilistic processes. In *Proceedings of the 10th Annual IEEE Symposium On Logic In Computer Science*, pages 431–441, 1995.

[19] C. Jou and S. Smolka. Equivalences, congruences and complete axiomatizations for probabilistic processes. In *CONCUR 90*, number 458 in Lecture Notes In Computer Science, 1990.

[20] A. Joyal, M. Nielsen, and G. Winskel. Bisimulation from open maps. *Information and Computation*, 127(2):164–185, 1996.

[21] K. G. Larsen and A. Skou. Bisimulation through probablistic testing. *Information and Computation*, 94:1–28, 1991.

[22] J. J. K. Rajeev Alur, Lalita Jategaonkar Jagadeesan and J. E. von Olnhausen. Model-checking of real-time systems: A telecommunications application. In *Proceedings of the 19th International Conference on Software Engineering*, 1997.

[23] R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. In B. Jonsson and J. Parrow, editors, *Proceedings of CONCUR94*, number 836 in Lecture Notes In Computer Science, pages 481–496. Springer-Verlag, 1994.

[24] R. van Glabbeek, S. Smolka, B. Steffen, and C. Tofts. Reactive generative and stratified models for probabilistic processes. In *Proceedings of the 5th Annual IEEE Symposium On Logic In Computer Science*, 1990.

A. Definitions from Measure Theory

Definition A.1 *A σ-**field** on a set X is a family of subsets of X which includes X itself and which is closed under complementation and countable unions.*

A set equipped with a σ-field is called a *measurable space*. Given a set X and an arbitrary subset $\mathcal{A}$ of the powerset of X, we can always find a *minimal* σ-field including all the sets of $\mathcal{A}$; this is called the σ-field *generated by* $\mathcal{A}$. Many of the measurable spaces "in nature" actually arise from topological spaces. We take the open sets (or the closed sets) and look at the σ-field generated by them. In the case of $\mathbf{R}^n$ the σ-field generated by the usual open sets is called the *Borel field*. The *Lesbegue* σ-field is quite different.

The notion of *measure* generalizes the concept of "length of an interval" to arbitrary σ-fields.

Definition A.2 *Given a σ-field (X, Σ), a **(subprobability) measure** on X is a ($[0,1]$-valued) $[0,\infty]$-valued set function, μ, defined on Σ such that*

- $\mu(\emptyset) = 0$,
- *for a pairwise disjoint, countable collection of sets, $\{A_i | i \in I\}$, in Σ, we require*

$$\mu(\bigcup_{i\in I} A_i) = \sum_{i\in I} \mu(A_i).$$

In addition, for probability measures we require $\mu(X) = 1$.

The second property above is called *countable additivity*. A *probability space* is a set X equipped with a σ-field Σ and a measure P such that $P(X) = 1$. Functions between measurable spaces are called measurable functions.

Definition A.3 *A function $f : (X, \Sigma_X) \to (Y, \Sigma_Y)$ between measurable spaces is said to be **measurable** if $\forall B \in \Sigma_Y . f^{-1}(B) \in \Sigma_X$.*

Definition A.4 *A **Polish** space is the topological space underlying a complete, separable (i.e. has a countable dense subset) metric space.*

The fact that we have a countable dense subset is equivalent to saying that there is a countable basis. This basis allows many measure-theoretic arguments to go through smoothly. A typical Polish space is $\mathbf{R}^n$.

Definition A.5 *An **analytic** space is the image of a Polish space under a continuous function from one Polish space to another.*

In this definition it turns out to be equivalent to say "measurable" image and it makes no difference if we take the image of the whole Polish space or of a Borel subset of the Polish space. A typical way that analytic spaces arise is when one takes the projection of a Borel subset of $\mathbf{R}^2$ onto one of the axes. Lebesgue mistakenly thought that these had to be Borel sets but Suslin showed that they may not be and so initiated the study of analytic sets.