

Duality for Labelled Markov Processes

Michael Mislove¹, Joël Ouaknine², Dusko Pavlovic³, and James Worrell¹

¹ Tulane University, Department of Mathematics, New Orleans, USA

² Computer Science Department, Carnegie Mellon University, USA

³ Kestrel Institute, Palo Alto, USA

Abstract. Labelled Markov processes (**LMPs**) are automata whose transitions are given by probability distributions. In this paper we present a ‘universal’ **LMP** as the spectrum of a commutative C^* -algebra consisting of formal linear combinations of labelled trees. We characterize the state space of the universal **LMP** as the set of homomorphisms from an ordered commutative monoid of labelled trees into the multiplicative unit interval. This yields a simple semantics for **LMPs** which is fully abstract with respect to probabilistic bisimilarity. We also consider **LMPs** with entry points and exit points in the setting of iteration theories. We define an iteration theory of **LMPs** by specifying its categorical dual: a certain category of C^* -algebras. We find that the basic operations for composing **LMPs** have simple definitions in the dual category.

1 Introduction

This paper is concerned with the semantics of certain probabilistic labelled transition systems, called labelled Markov processes (or **LMPs**) [9,11,7,8]. Probabilistic models are important for capturing quantitative aspects of process behaviour, such as performance and reliability, e.g., the average response time to a given action, or the probability with which a failure occurs. For this reason there has been a lot of research into adapting the concepts and results of classical concurrency theory to the probabilistic case. In particular, the notion of bisimilarity has been adapted to probabilistic systems [18,9,17], and its equational theory investigated in [22,4,19] amongst many others.

The bisimulation equivalence classes of **LMPs** can be gathered together into what could be termed a universal **LMP**. This object has previously been studied as the solution of a domain equation in the category of complete metric spaces [7], and in the category of coherent domains [11,8]. However, none of these domain-theoretic treatments yielded concrete representations of the elements of the universal **LMP**.

¹ Supported by the Office of Naval Research, grant No. N000149910150 and the National Science Foundation, grant No. CCR-0208743.

² Supported by the Office of Naval Research, contract N00014-95-1-0520.

³ Supported by the EHS and SGER programs of the National Science Foundation, under contract No. CCR-0209-004 and CCR-0345-397.

In this paper we exploit Stone-Gelfand-Naimark duality for commutative C^* -algebras to show that the universal **LMP** has a very straightforward characterization as a space of order-preserving monoid homomorphisms from a partially ordered monoid \mathbb{T} of trees to the multiplicative monoid $[0, 1]$.

We think of the elements of \mathbb{T} as *branching traces* or *trace trees*. Formally they are just finite trees whose edges are labelled by events from a given alphabet. The operation of grafting two such trees at the root gives the monoid multiplication in \mathbb{T} . The order on \mathbb{T} is the natural generalization of the prefix order on traces. For a given **LMP** the corresponding homomorphism maps each trace tree to the probability that it gets performed.

In an earlier paper [8] we showed that two processes are bisimilar iff they perform each trace tree¹ with the same probability. This generalized a result of Larsen and Skou [18]. The main result of this paper can be seen as extending this characterization to ‘build processes out of trace trees’. This is a natural variation of the familiar trace models of abstract machines.

The main mathematical tool that we use is the theorem of Stone asserting a dual equivalence between the category of compact Hausdorff spaces and continuous maps on the one hand, and the category of commutative real C^* -algebras (a full subcategory of the category of commutative rings) on the other hand. This duality associates to each compact Hausdorff space the ring of continuous real-valued functions on the space, and to each C^* -algebra its spectral space of *characters*: the ring homomorphisms into \mathbb{R} . We apply the duality to recover the universal **LMP** as the spectrum of a C^* -algebra consisting of formal linear combinations of trace trees.

The concrete representation of the universal **LMP**, obtained in the first part of this paper, opens a new, effective approach to composing **LMPs**. In particular, we augment the basic model of **LMPs** with entry and exit points, and study some basic combinators, such as sequential composition, probabilistic choice and iteration. Thus we obtain a category whose objects are finite sets, and where a morphism $\mathcal{S}: X \rightarrow Y$ is an **LMP** with entry points X and exit points Y . We show that this category is dual (contravariantly equivalent) to a category of commutative rings consisting of C^* -algebras of trace trees. Represented through this duality, the basic combinators for **LMPs** have remarkably simple descriptions.

1.1 Related Work

Kozen [15] presents a predicate-transformer semantics of an imperative programming language with probabilistic choice. This semantics is based on a duality between linear maps and probabilistic relations, and is formalized in the setting of iteration theories. The same themes of duality and iteration theories appear in the present paper, but our development is in the context of interactive processes rather than imperative programs. In particular, for us states are not just measurable functions on a space of variables as in [15], but have a recursively-defined structure.

¹ In fact we took the view that trace trees are types of button-pressing tests.

An iteration theory of probabilistic processes has been studied by Aceto, Ésik and Ingólfssdóttir [4], building on earlier work of Stark and Smolka [22]. These two papers treat finite-state **LMPs** as terms in a simple probabilistic process calculus. Their main contributions are soundness and completeness results for axiomatizations of the bisimilarity relation. Since these papers deal with process calculi, the basic operations of prefixing, probabilistic choice and iteration are defined at the syntactic level using an operational semantics. These operations are then lifted to bisimulation equivalence classes of terms using the fact that bisimilarity is a congruence. In contrast, we use a concrete representation of bisimulation equivalence classes of **LMPs** as maps of C^* -algebras, and define the basic operations directly on these representations.

One of the most comprehensive applications of duality in semantics can be found in the work of Abramsky on domain theory in logical form [1]. This work is based on a Stone-type duality between a category of spectral spaces (SFP domains in their Scott topologies) and a category of distributive lattices. As a case study, Abramsky considers a domain for bisimulation and computes its spectrum. In this we paper we compute the spectrum of a domain for probabilistic bisimulation. However, so far our work is much more modest in scope; in particular we have not tried to isolate a fragment of the duality for C^* -algebras that is pertinent to any reasonable category of domains.

Another paper close in spirit to the present work is Abramsky and Vickers [2]. They consider a variety of equivalences for concurrent processes in a unified framework of quantale modules—actions of quantales on sup-lattices. In particular, they present quantales of tests using generators and relations, and model transition systems as right quantale modules (where the elements of the quantale act on states of the transition system). Using the self-duality of the category of sup-lattices they obtain left quantale modules of ‘process capabilities’ which they use to build fully abstract models of processes.

Di Pierro, Hankin and Wicklicky [21] use C^* -algebras to define abstract interpretations of probabilistic transition systems. These C^* -algebras are non-commutative operator algebras, but we believe that exploring connections with this work merits further investigation.

2 Labelled Markov Processes

Below we give the formal definition of the class of probabilistic transition systems that we study in this paper. This definition extends that of Larsen and Skou [18] by including entry points and exit points as part of the basic data. For a similar treatment of labelled transition systems, see [6].

Assume a fixed finite set Act of actions or events. Given a set S , a sub-probability distribution on S is a non-negative real-valued function on S with countable support and total mass no greater than 1.

Definition 1. *Given finite sets X and Y of entry points and exit points, a labelled Markov process $\mathcal{S}: X \rightarrow Y$ is a tuple (S, ι, μ) consisting of a set S of*

states, an injective function $\iota: X \rightarrow S$ and, for each $s \in S$, a sub-probability distribution μ_s on $(\text{Act} \times S) + Y$.

Given $s \in S$ and $a \in \text{Act}$, $\mu_s(a, t)$ is the probability that the process in state s makes an a -transition to $t \in S$; $\mu_s(y)$ is the probability that it makes a transition to the exit point $y \in Y$. Note that μ_s is a *sub-probability* distribution on $(\text{Act} \times S) + Y$. We interpret the difference between the total mass of μ_s and 1 as the probability of refusing all actions. We also adopt the notation $\mu_{s,a}$ for the sub-probability distribution on S given by $\mu_{s,a}(t) = \mu_s(a, t)$.

Apart from the presence of entry and exit points, Definition 1 differs from the notion of **LMP** in [9,8] in that, for each state s , the $\mu_{s,a}$ are components of a single transition probability distribution, rather than being an arbitrary family of probability distributions. In the terminology of [14] we are using the *generative* model as opposed to the *reactive* model. Also, to keep things simple, we only define **LMPs** with discrete probabilities. However, all our results hold for the more general case where the state space of a **LMP** is a measurable space and the transitions are given by sub-probability measures (see [9,8]). This more general type of **LMP** briefly features in the definition of a universal system in Section 7.

Probabilistic bisimilarity [18] (henceforth just bisimilarity) is the probabilistic analogue of strong bisimilarity for labelled transition systems. It gives a branching-time notion of behavioural equivalence for **LMPs**.

Definition 2. Let $\mathcal{S} = (S, \iota, \mu): X \rightarrow Y$ be an **LMP**. An equivalence relation R on S is a bisimulation if sRt implies that

- for each $a \in \text{Act}$ and R -equivalence class A , $\mu_{s,a}(A) = \mu_{t,a}(A)$,
- for each $y \in Y$, $\mu_s(y) = \mu_t(y)$.

We say that two states are bisimilar if they are related by some bisimulation.

In words: an equivalence relation is a bisimulation if related states have matching probabilities of making transitions into any equivalence class and into any exit point.

3 Operations on LMPs

In this section we define some operations for composing **LMPs**. These are the counterparts on the semantic level of constructs that might be found in a typical process calculus. In particular, processes with exit points correspond to terms with free variables, composition corresponds to substitution of terms, and iteration corresponds to the application of the recursion operator. These definitions will later form the basis of a category in which **LMPs** are morphisms.

Elements. Given a finite set X , each element $x \in X$ determines an **LMP** $x: 1 \rightarrow X$ with one state $*$, where $*$ makes a transition to output $x \in X$ with probability 1. Formally $x = (\{*\}, \iota, \mu)$, where $\mu_*(x) = 1$.

Tupling. Let $\mathcal{S}: X \rightarrow Z$ and $\mathcal{S}': Y \rightarrow Z$ be **LMPs** with $\mathcal{S} = (S, \iota, \mu)$ and $\mathcal{S}' = (S', \iota', \mu')$. The tuple $\langle \mathcal{S}, \mathcal{S}' \rangle : X + Y \rightarrow Z$ is obtained by taking the coproduct of \mathcal{S} and \mathcal{S}' , and identifying the outputs. Formally, $\langle \mathcal{S}, \mathcal{S}' \rangle = (S + S', \iota'', \rho)$, where $\iota'' = \iota + \iota' : X + Y \rightarrow S + S'$, and

$$\rho_s(v) = \begin{cases} \mu_s(v) & \text{if } s \in S, v \in (\text{Act} \times S) + Z \\ \mu'_s(v) & \text{if } s \in S', v \in (\text{Act} \times S') + Z \\ 0 & \text{otherwise.} \end{cases}$$

Composition. Let $\mathcal{S}: X \rightarrow Y$ and $\mathcal{S}': Y \rightarrow Z$ be **LMPs** with $\mathcal{S} = (S, \iota, \mu)$ and $\mathcal{S}' = (S', \iota', \mu')$. The composition $(\mathcal{S} \circledast \mathcal{S}') : X \rightarrow Z$ is obtained by connecting the outputs of \mathcal{S} with the inputs of \mathcal{S}' . Formally $\mathcal{S} \circledast \mathcal{S}' = (S + S', \iota'', \rho)$, where $\iota'' = (S \rightarrow S + S') \circ \iota$, and

$$\rho_s(v) = \begin{cases} \mu_s(v) & \text{if } s \in S, v \in (\text{Act} \times S) \\ \sum_{y \in Y} \mu_s(y) \mu'_{\iota' y}(v) & \text{if } s \in S, v \in (\text{Act} \times S') + Z \\ \mu'_s(v) & \text{if } s \in S', v \in (\text{Act} \times S') + Z. \end{cases}$$

Thus composition is given by integration, just as in the category of stochastic relations studied in [3].

Iteration. Given $\mathcal{S}: X \rightarrow X + Y$ with $\mathcal{S} = (S, \iota, \mu)$, the iterate $\mathcal{S}^\dagger : X \rightarrow Y$ is obtained by connecting each exit point $x \in X$ to the corresponding entry point. Writing $\mathcal{S}^\dagger = (S, \iota, \rho)$, we define ρ so as to satisfy:

$$\rho_s(v) = \mu_s(v) + \sum_{x \in X} \mu_s(x) \rho_{\iota x}(v) \tag{1}$$

for $s \in S$ and $v \in (\text{Act} \times S) + Y$.

The definition of ρ relies on the Kleene-* operation for matrices over the semiring \mathbb{R}_+^∞ (see Appendix A). Write $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_p\}$. Let A be the $n \times n$ matrix with $A_{ij} = \mu_{\iota x_i}(x_j)$ and B the $n \times p$ matrix with $B_{ij} = \mu_{\iota x_i}(y_j)$. We first define the transition behaviour of the entry states by

$$\rho_{\iota x_i}(y_j) \stackrel{\text{def}}{=} (A^* B)_{ij} \text{ and } \rho_{\iota x_i}(a, t) \stackrel{\text{def}}{=} (A^* v)_i,$$

where v is the column vector whose j -th entry is $\mu_{\iota(x_j)}(a, t)$. Now the definition of ρ_s for $s \in S \setminus \text{Range}(\iota)$ can be read off from Equation 1.

Probabilistic Choice. Let X be a finite set, $X \xrightarrow{\text{in}_1} X + X \xleftarrow{\text{in}_2} X$ a designated coproduct, and $0 \leq r \leq 1$ a real number. We define the **LMP** $\oplus_r : X \rightarrow X + X$ by $\oplus_r = (X, \text{id}_X, \mu)$, where $\mu_x(\text{in}_1(x)) = r$ and $\mu_x(\text{in}_2(x)) = 1 - r$ for each $x \in X$ —that is, state x selects output $\text{in}_1(x)$ with probability r and output $\text{in}_2(x)$ with probability $1 - r$. Given **LMPs** $\mathcal{S}, \mathcal{S}': X \rightarrow Y$, we write $\mathcal{S} \oplus_r \mathcal{S}'$ for $\oplus_r \circledast \langle \mathcal{S}, \mathcal{S}' \rangle$.

The above operations form the basis of the iteration theory of **LMPs** defined in Section 8. More precisely, this theory is predicated on **LMPs** modulo bisimilarity, where the bisimilarity relation is extended from states to **LMPs** in the definition below.

Definition 3. Let $\mathcal{S}, \mathcal{S}': X \rightarrow Y$ be **LMPs**, and let the tuple $\langle \mathcal{S}, \mathcal{S}' \rangle$ be defined relative to a designated coproduct $X \xrightarrow{\text{in}_1} X + X \xleftarrow{\text{in}_2} X$. We say that \mathcal{S} and \mathcal{S}' are bisimilar, written $\mathcal{S} \simeq \mathcal{S}'$, if there is a bisimulation R on the tuple $\langle \mathcal{S}, \mathcal{S}' \rangle$ which relates entry points $\text{in}_1(x)$ and $\text{in}_2(x)$ for each $x \in X$.

4 A Monoid of Trace Trees

In this section we present a grammar for a class of trees corresponding to branching-time traces of an **LMP**. This language (minus exit actions) corresponds to the test languages of [18,8] which were shown to characterize, respectively, similarity in labelled transition systems, and probabilistic bisimilarity in labelled Markov processes.

Fix a finite set Y (corresponding to the exit points of an **LMP**). The language of trace trees is generated by the grammar

$$\tau ::= 1 \mid y \mid a\tau \mid \tau \cdot \tau \tag{2}$$

where $a \in \text{Act}$ and $y \in Y$.

A trace tree is either the null tree 1, an exit action $y \in Y$, a prefixing $a\tau$, or a branch point $\tau_1 \cdot \tau_2$. Note the distinction between prefixing (which is denoted by mere juxtaposition) and branching. We will typically elide the symbol 1 when denoting non-trivial trace trees, e.g., we write $a \cdot bc$ for $a1 \cdot bc1$. Without the branching construct ‘ \cdot ’ the grammar above would just specify a language of traces. In order to physically realize a branching-time trace one would need to be able to duplicate the process at any point in a run, for instance, via a save-and-restore construct.

Definition 4. Given an **LMP** $\mathcal{S}: X \rightarrow Y$, with $\mathcal{S} = (S, \iota, \mu)$, for each $s \in S$ we define $\tau_{\mathcal{S}}(s)$: the probability that s performs tree τ .

- $1_{\mathcal{S}}(s) = 1$.
- $(y)_{\mathcal{S}}(s) = \mu_s(y)$.
- $(a\tau)_{\mathcal{S}}(s) = \int \tau_{\mathcal{S}} d\mu_{s,a}$.
- $(\tau_1 \cdot \tau_2)_{\mathcal{S}}(s) = (\tau_1)_{\mathcal{S}}(s)(\tau_2)_{\mathcal{S}}(s)$.

The null tree is performed with probability 1 in any state. The probability that $a\tau$ is performed in any given state is the weighted average of the probability that τ is performed in the next state after an a -transition. The last clause says that probability of performing an immediately branching tree $\tau_1 \cdot \tau_2$ is the product of the probabilities of performing each branch.

For determining an **LMP** up to bisimilarity, only the behaviour of states reachable from the entry points matters. Given $\mathcal{S}: X \rightarrow Y$, for each $x \in X$

we define the real-valued function $\widehat{\mathcal{S}}_x$ on trace trees by $\widehat{\mathcal{S}}_x(\tau) = \tau_{\mathcal{S}}(\iota x)$. Thus $\widehat{\mathcal{S}}_x(\tau)$ is the probability that \mathcal{S} does τ on input x . The following theorem is a generalization of the main result of [8] to allow for **LMPs** with entry and exit points.

Theorem 5. *Let $\mathcal{S}, \mathcal{T}: X \rightarrow Y$ be **LMPs**. Then \mathcal{S} and \mathcal{T} are bisimilar iff $\widehat{\mathcal{S}}_x = \widehat{\mathcal{T}}_x$ for all $x \in X$.*

Having used trace trees to characterize equivalence of states, we consider the dual problem: when are two trees equivalent in that each state performs them with the same probability? More generally, we define a preorder \leq on trace trees by $\tau \leq \tau'$ iff $\tau_{\mathcal{S}} \leq \tau'_{\mathcal{S}}$ for all **LMPs** \mathcal{S} . The key to constructing a model for **LMPs** that is fully abstract with respect to bisimilarity is to axiomatize this preorder. As a first step, notice that the map $(-)_S$ preserves the equations below. These say that the set of trace trees forms a commutative monoid equipped with the smallest partial order in which 1 is the top element and prefixing and multiplication are monotone.

$$\begin{array}{ll}
 1 \cdot \tau = \tau & \tau \leq 1 \\
 \tau_1 \cdot \tau_2 = \tau_2 \cdot \tau_1 & \tau_1 \cdot \tau \leq \tau_2 \cdot \tau \text{ if } \tau_1 \leq \tau_2 \\
 \tau_1 \cdot (\tau_2 \cdot \tau_3) = (\tau_1 \cdot \tau_2) \cdot \tau_3 & a\tau_1 \leq a\tau_2 \text{ if } \tau_1 \leq \tau_2
 \end{array}$$

We denote the resulting partially-ordered monoid by $\mathbb{T}[Y]$. Thus we have a monoid for each set Y of exit actions.

5 Stone-Gelfand-Naimark Duality

Our basic reference for this section is the monograph of Johnstone [16]. We define C^* -algebras to be certain types of commutative rings. The category $C^*\text{-Alg}$ is the resulting full subcategory of CRng . Here we should emphasize that we take C^* -algebras to be algebras over \mathbb{R} as opposed to the more standard presentation as algebras over \mathbb{C} (cf. Naimark [13, Theorem III.2.1]).

Let A be a commutative ring. Since we are primarily interested in rings of functions, we use f, g to denote typical elements of A . We say that A is an *ordered ring* if it is equipped with a partial order satisfying

$$\begin{array}{l}
 f + g \leq f' + g \text{ if } f \leq f' \\
 f \cdot g \leq f' \cdot g \text{ if } f \leq f', g \geq 0 \\
 f^2 \geq 0.
 \end{array}$$

We say that an ordered commutative ring A is *Archimedean* if for all f there exists a positive integer n with $f \leq n \cdot 1$. If the additive group of A is torsion-free and divisible, so that A admits a \mathbb{Q} -algebra structure, then we may define a seminorm² on A by

$$\|f\| = \inf\{q \in \mathbb{Q} \mid -q \cdot 1_A \leq f \leq q \cdot 1_A\}. \tag{3}$$

² Non-zero elements can have norm zero.

Definition 6. A commutative ring A is a real C^* -algebra if

- the additive group of A is torsion free and divisible, and
- A possesses an Archimedean partial order such that Equation 3 defines a norm with respect to which A is complete.

In a C^* -algebra it turns out that $f \geq 0$ iff f is a square. Thus the partial order is determined by the ring structure, and ring homomorphisms between C^* -algebras are automatically order preserving.

Definition 7. A character of a C^* -algebra A is a ring homomorphism $\varphi: A \rightarrow \mathbb{R}$. The spectrum of A , denoted $\text{spec } A$, is the space of characters of A in the Zariski topology, which is generated by the cozero sets $\text{coz}(f) = \{\varphi : \varphi(f) \neq 0\}$ where $f \in A$.

The spectrum of a C^* -algebra is a compact Hausdorff space. Conversely, the ordered ring $C^*(X)$ of continuous real-valued functions on a compact Hausdorff space X is always a C^* -algebra. This association of compact Hausdorff spaces and C^* -algebras is functorial, and is in fact a dual equivalence:

Theorem 8 (Stone). The category KHaus of compact Hausdorff spaces and continuous maps is dually equivalent to $C^*\text{-Alg}$.

6 A Family of C^* -Algebras

In this section we extend the monoid of trace trees to a C^* -algebra whose spectrum is the state space of a universal LMP.

Fix a set Y of exit points. We extend the grammar (2) for trace trees to a grammar of *functional expressions* by allowing rational linear combinations. Thus functional expressions are given by

$$f ::= q \mid y \mid af \mid f \cdot f \mid f + f \tag{4}$$

where $a \in \text{Act}$, $y \in Y$ and $q \in \mathbb{Q}$.

We will use the letters f and g to denote functional expressions. We adopt the convention that a term denoted τ has been generated using only the sub-grammar (2). We reserve the phrase trace tree for such terms.

We use functional expressions as generators in a presentation of a family of ordered rings $\mathbb{O}[Y]$, where the index Y indicates the dependence on the finite set Y of exit variables. In this presentation ‘ \cdot ’ acts as multiplication, 1 is the multiplicative identity, and $+$ acts as addition in $\mathbb{O}[Y]$.

The relations in the presentation of $\mathbb{O}[Y]$ include the equations for an ordered ring: the Abelian group axioms for $+$, the commutative monoid axioms for ‘ \cdot ’, the distributive law of ‘ \cdot ’ over $+$, and axioms asserting the compatibility of the order relation with the ring structure. To these we add Equations 5–8 below. The effect of these equations is to fix the semantics of prefixing as integration against

a sub-probability measure. Note that the distributive law (8) implies that every functional expression is equal to a linear combination of trace trees.

$$0 \leq y \tag{5}$$

$$af \leq ag \text{ if } f \leq g \tag{6}$$

$$\sum_{a \in \text{Act}} a + \sum_{y \in Y} y \leq 1 \tag{7}$$

$$a(q_1 \cdot f + q_2 \cdot g) = q_1 \cdot af + q_2 \cdot ag \tag{8}$$

Definition 9. Define $\mathbb{O}[Y]$ to be the free ordered ring³ generated by the set of functional expressions and satisfying Equations 5–8.

Proposition 10. $\mathbb{O}[Y]$ is a torsion-free divisible Archimedean ordered ring.

$\mathbb{O}[Y]$ is Archimedean since each functional expression is equal to a linear combination of trace trees, and each trace tree τ satisfies $\tau \leq 1$.

Definition 11. Define the C^* -algebra $\mathbb{A}[Y]$ to be the Cauchy completion of $\mathbb{O}[Y]$ in the norm (3). The ring operations on $\mathbb{O}[Y]$ are non-expansive in this norm, so they extend to $\mathbb{A}[Y]$.

Proposition 12. $\mathbb{A}[Y]$ is the free C^* -algebra over $\mathbb{O}[Y]$ qua ordered ring.

Remark 13. Combining Definition 9 and Proposition 12 we see that in order to specify a ring homomorphism from $\mathbb{A}[Y]$ to a C^* -algebra R it suffices to give an interpretation of the functional expressions in R such that the relations in the presentation of $\mathbb{O}[Y]$ all hold. Since the interpretations of $+$ and \cdot are forced, this boils down to interpreting prefixing $a(-)$ and exit actions $y \in Y$.

Given a set S , write $C^*(S)$ for the ring of bounded real-valued functions on S . This ring is a C^* -algebra (cf. [16]).

Definition 14. Let $\mathcal{S} : X \rightarrow Y$ be an **LMP** with $\mathcal{S} = (S, \iota, \mu)$. We define a ring homomorphism

$$\mathbb{A}[Y] \xrightarrow{\mathcal{S}} C^*(S)$$

by the following clauses:

$$(af)_{\mathcal{S}}(s) = \int_S f_S d\mu_{s,a}$$

$$(y)_{\mathcal{S}}(s) = \mu_s(y).$$

Furthermore we define $\widehat{\mathcal{S}}_x \in \text{spec } \mathbb{A}[Y]$ by $\widehat{\mathcal{S}}_x(f) = f_S(\iota x)$.

Note that this extends Definition 4. Indeed, since every element of $\mathbb{O}[Y]$ is equal to a linear combination of trace trees, an element of $\text{spec } \mathbb{A}[Y]$ is determined by an order-preserving monoid homomorphism $\mathbb{T}[Y] \rightarrow [0, 1]$ satisfying Equation 7.

³ Note in passing that the existence of a free ordered ring on a given set of generators and relations follows from the existence of free algebras for Horn theories.

7 Universal LMPs

In this section we define a universal **LMP** $\mathcal{U}[Y] : \emptyset \rightarrow Y$ on outputs Y . The state space of $\mathcal{U}[Y]$ is $\text{spec } \mathbb{A}[Y]$. In order to manufacture the transition probabilities we use the Riesz representation theorem [20].

Theorem 15. (*Riesz*) *Let K be a compact Hausdorff space and $\varphi : C^*(K) \rightarrow \mathbb{R}$ a positive linear map, i.e. $\varphi(f) \geq 0$ whenever $f \geq 0$. Then there is a unique positive Borel measure μ on K such that $\varphi(f) = \int f d\mu$ for all $f \in C^*(K)$.*

The transition behaviour of $\varphi \in \text{spec } \mathbb{A}[Y]$ is given by a sub-probability measure μ_φ on $(\text{Act} \times \text{spec } \mathbb{A}[Y]) + Y$ which is defined as follows. First, $\mu_\varphi(y)$, the probability of making a transition to the exit $y \in Y$, is defined to be $\varphi(y)$. Next, given $a \in \text{Act}$, let $\varphi_a : \mathbb{A}[Y] \rightarrow \mathbb{R}$ be defined by $\varphi_a(f) = \varphi(af)$. The distributive law (8) ensures that φ_a is a linear map. Also, φ_a is positive since φ is positive and prefixing is monotone in $\mathbb{A}[Y]$. We define $\mu_{\varphi,a} = \mu_\varphi(\{a\} \times -)$ to be the Borel sub-probability measure on $\text{spec } \mathbb{A}[Y]$ corresponding by Theorem 15 to the linear map

$$C^*(\text{spec } \mathbb{A}[Y]) \cong \mathbb{A}[Y] \xrightarrow{\varphi_a} \mathbb{R}.$$

Note the application of Theorem 8 in the above isomorphism. Finally, we observe that Equations 5 and 7 guarantee that μ_φ , as defined above, is a sub-probability measure.

In order to state the universal property of $\mathcal{U}[Y]$, we define the notion of a zig-zag map [9].

Definition 16. *Let $\mathcal{S}, \mathcal{S}'$ be **LMPs** on exits Y . Suppose that $\mathcal{S} = (S, \iota, \mu)$ and $\mathcal{S}' = (S', \iota', \mu')$. A function $h : S \rightarrow S'$ is a zig-zag map iff*

- $\mu_{s,a}(h^{-1}(t)) = \mu_{h(s),a}(t)$ for all $s \in S, t \in S'$ and $a \in \text{Act}$.
- $\mu_s(y) = \mu_{h(s)}(y)$ for all $s \in S, y \in Y$.

The following proposition is proved in [9].

Proposition 17. *Let $\mathcal{S}, \mathcal{S}'$ be **LMPs** on exits Y . A function $h : S \rightarrow S'$ is a zig-zag map iff the kernel of h is a bisimulation.*

Next we show how the initiality of $\mathbb{A}[Y]$ transfers, via Stone duality, to the finality of $\mathcal{U}[Y]$.

Proposition 18. *Let $\mathcal{S} = (S, \iota, \mu)$ be an **LMP** on exits Y . Then a function $h : S \rightarrow \text{spec } \mathbb{A}[Y]$ is a zig-zag map $\mathcal{S} \rightarrow \mathcal{U}[Y]$ iff the dual map $\bar{h} : \mathbb{A}[Y] \rightarrow C^*(S)$, where $\bar{h}(f)(s) = h(s)(f)$, satisfies*

- $\bar{h}(af)(s) = \int_S \bar{h}(f) d\mu_{s,a}$, and
- $\bar{h}(y)(s) = \mu_s(y)$ for all $s \in S, y \in Y$.

Proof (Sketch). The zig-zag condition in Definition 16 says that $\int (f \circ h) d\mu_{s,a} = \int f d\mu_{h(s),a}$ for all step functions $f : S' \rightarrow \mathbb{R}$. By linearity of the integral this equation holds for all bounded real-valued f ; but this transfers through the duality to the first equation in the proposition. \square

By Remark 13 there is a unique map $\mathbb{A}[Y] \rightarrow C^*(S)$ satisfying the clauses in Proposition 18—namely the map $(-)_S$ from Definition 14. Thus we obtain:

Theorem 19. *$\mathcal{U}[Y]$ is final in the category whose objects are **LMPs** on exits Y and whose morphisms are zig-zag maps.*

In conjunction with Proposition 17, the finality of $\mathcal{U}[Y]$ implies that the relation of bisimilarity on a given **LMP** $\mathcal{S}: X \rightarrow Y$ is the kernel of the unique zig-zag map to $\mathcal{U}[Y]$. By definition, this map sends the entry state labelled by $x \in X$ to the character $\widehat{\mathcal{S}}_x$. Thus we obtain:

Corollary 20. *Two **LMPs** $\mathcal{S}, \mathcal{T}: X \rightarrow Y$ are bisimilar (cf. Definition 3) iff $\widehat{\mathcal{S}}_x = \widehat{\mathcal{T}}_x$ for each $x \in X$.*

In particular, $\mathcal{S}: X \rightarrow Y$ is represented up to bisimilarity by the X -indexed set $(\widehat{\mathcal{S}}_x)_{x \in X}$ of characters of $\mathbb{A}[Y]$.

8 An Iteration Theory of LMPs

The notion of an iteration theory arises by extending a Lawvere algebraic theory with an iteration operation on maps. In particular, an iteration theory is a category with finite coproducts. The iteration operation takes a map $\varphi: X \rightarrow X + Y$ to a map $\varphi^\dagger: X \rightarrow Y$, and is required to satisfy certain equations like the Elgot fixed point identity, the Bekič (pairing) identity and the group identities [5]. Flowchart algorithms, regular and context-free languages, synchronization trees and Floyd-Hoare logic have all been formalized in the setting of iteration theories (see [5]).

In Section 3 we presented the raw material for an iteration theory of **LMPs**; in particular we defined composition, tupling, probabilistic choice and iteration. In this section we show that one obtains an iteration theory from these operations by considering **LMPs** modulo bisimilarity. We present this iteration theory via a duality with a category of C^* -algebras and maps we call *factorizations*. In particular, we show that **LMPs** are dual to factorizations, and iterating **LMPs** is dual to taking fixed points of factorizations.

For each $x \in X$ the *projection* $\pi_x: \mathbb{A}[X] \rightarrow \mathbb{R}$ is the ring map defined by

$$\pi_x(af) = 0 \text{ and } \pi_x(x') = \begin{cases} 1 & \text{if } x = x' \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

(Recall from Remark 13 that to define a C^* -algebra map with domain $\mathbb{A}[X]$ it suffices to show how to interpret prefixing and variables in the target.) Note that $\pi_x \in \text{spec } \mathbb{A}[X]$ is the spectral counterpart to the **LMP** $x: 1 \rightarrow X$ in Section 3 according to the representation of Corollary 20. Given a ring map $\varphi: \mathbb{A}[Y] \rightarrow \mathbb{A}[X]$, the *components* of φ are the characters $\varphi_x = \pi_x \circ \varphi: \mathbb{A}[Y] \rightarrow \mathbb{R}$, $x \in X$.

Definition 21. A (tree) factorization is a ring homomorphism $\varphi: \mathbb{A}[Y] \rightarrow \mathbb{A}[X]$ satisfying the following relations:

$$\begin{aligned} \varphi(af) &= a\varphi(f) + \sum_{x \in X} \varphi_x(af) \cdot x \\ \varphi(y) &= \sum_{x \in X} \varphi_x(y) \cdot x. \end{aligned}$$

A factorization is completely determined by its components. In particular, given an X -indexed family $(\varphi_x)_{x \in X}$ in $\text{spec } \mathbb{A}[Y]$, the above equations uniquely define a factorization $\varphi: \mathbb{A}[Y] \rightarrow \mathbb{A}[X]$. To explain the terminology, note that a factorization sends a trace tree $\tau \in \mathbb{T}[Y]$ to a linear combination of all those trees in $\mathbb{T}[X]$ from which τ may be obtained by substituting trees for leaves (see the following example).

Example 22. Let $a, b, c \in \text{Act}$ and $X = \{x\}$. A factorization $\mathbb{A} \rightarrow \mathbb{A}[X]$ maps $ab \cdot c$ to a linear combination of the trees $x \cdot x, x \cdot c, ax \cdot x, ax \cdot c, ab \cdot x$ and $ab \cdot c$.

Definition 23. The category **Fact** is the subcategory of **CRng** with objects the C^* -algebras $\mathbb{A}[X]$, where X is a finite set, and morphisms the factorizations. Henceforth we simply take the objects of **Fact** to be the finite sets.

Proposition 24. **Fact** has finite products.

Proof (sketch). Given finite sets X and Y , the span $X \xleftarrow{\pi_1} X + Y \xrightarrow{\pi_2} Y$ is a binary product, where $\pi_1: \mathbb{A}[X + Y] \rightarrow \mathbb{A}[X]$ is the factorization defined by

$$\pi_1(af) = a\pi_1(f) \text{ and } \begin{cases} \pi_1(x) = x & x \in X \\ \pi_1(y) = 0 & y \in Y \end{cases}$$

and $\pi_2: \mathbb{A}[X + Y] \rightarrow \mathbb{A}[Y]$ is the factorization defined by

$$\pi_2(af) = a\pi_2(f) \text{ and } \begin{cases} \pi_2(x) = 0 & x \in X \\ \pi_2(y) = y & y \in Y. \end{cases}$$

□

Proposition 25. Each homset $\text{Fact}(X, Y)$ is a pointed dcpo in the following partial order: $\varphi \leq \psi$ iff $\varphi(\tau) \leq \psi(\tau)$ for each trace tree $\tau \in \mathbb{A}[X]$.

Proof. Recall two standard results about a C^* -algebra A . First, the partial order is a closed subset of $A \times A$, and second, the closed unit ball $\{f \in A : \|f\| \leq 1\}$ is compact in the norm topology.

Now let $\{\varphi^{(k)}\}$ be a directed set in $\text{Fact}(X, Y)$. Given $\tau \in \mathbb{A}[X]$ we have that $\|\varphi^{(k)}(\tau)\| \leq \|\tau\| \leq 1$ since all ring maps are non-expansive. By compactness of the unit ball in $\mathbb{A}[Y]$, the net $\varphi^{(k)}(\tau)$ has a limit point in the norm topology on $\mathbb{A}[Y]$. By the closedness of the partial order on $\mathbb{A}[Y]$ it easily follows that this limit point is in fact a supremum. We define a factorization $\varphi: \mathbb{A}[X] \rightarrow \mathbb{A}[Y]$ by specifying its values on trace trees: $\varphi(\tau) = \bigsqcup_k \varphi^{(k)}(\tau)$. Then φ is the supremum of the family $\{\varphi^{(k)}\}$ by construction.

Define the factorization $\perp_{XY}: X \rightarrow Y$ by $\perp_{XY}(af) = af$ and $\perp_{XY}(x) = 0$ for all $x \in X$. Then \perp_{XY} is least in $\text{Fact}(X, Y)$. □

Proposition 26. *Composition in Fact is jointly continuous:*

$$(\bigsqcup_k \varphi^{(k)}) \circ \psi = \bigsqcup_k (\varphi^{(k)} \circ \psi) \quad \text{and} \quad \chi \circ (\bigsqcup_k \varphi^{(k)}) = \bigsqcup_k (\chi \circ \varphi^{(k)}),$$

for any directed family $\varphi^{(k)}: X \rightarrow Y$, $\psi: W \rightarrow X$ and $\chi: Y \rightarrow Z$.

Moreover composition is left-strict in that $\perp_{YZ} \circ \varphi = \perp_{XZ}$ for $\varphi: X \rightarrow Y$.

Proof. Since limits of directed sets in the unit ball of a C^* -algebra are also limits in the norm topology, they are preserved by any ring map. □

We use the dcpo-enriched structure of **Fact** in a standard way (see [5, Chapter 8]) to define a dagger operation taking a morphism $\varphi: X + Y \rightarrow X$ to a morphism $\varphi^\dagger: Y \rightarrow X$ (think of parameterized fixed point). The dagger operation satisfies the Elgot fixed point identity $\varphi^\dagger = \varphi \circ \langle \varphi^\dagger, \text{id}_Y \rangle$ (where the angled brackets refer to product tupling in **Fact**).

Definition 27. *Given a factorization $\varphi: X + Y \rightarrow X$, the iterate $\varphi^\dagger: Y \rightarrow X$ is defined to be $\bigsqcup_k \varphi^{(k)}$, where $\varphi^{(0)} = \perp_{YX}$ and $\varphi^{(k+1)} = \varphi \circ \langle \varphi^{(k)}, \text{id}_Y \rangle$.*

8.1 Duality between LMPs and Factorizations

Recall that an **LMP** $\mathcal{S}: X \rightarrow Y$ is determined up to bisimilarity by the X -indexed set $(\widehat{\mathcal{S}}_x)_{x \in X}$, where $\widehat{\mathcal{S}}_x \in \text{spec } \mathbb{A}[Y]$ gives the behaviour at the entry point $x \in X$. But, by the remarks following Definition 21, this X -tuple of characters also determines a factorization, which we denote $\widehat{\mathcal{S}}: \mathbb{A}[Y] \rightarrow \mathbb{A}[X]$. We regard $\widehat{\mathcal{S}}$ as the dual of \mathcal{S} (note the reversal of direction). Each factorization is the dual of some **LMP** and this duality is faithful up to bisimilarity of **LMPs**.

Proposition 28. *Given LMPs $\mathcal{S}, \mathcal{T}: X \rightarrow Y$, $\widehat{\mathcal{S}} = \widehat{\mathcal{T}}$ iff \mathcal{S} and \mathcal{T} are bisimilar.*

The following proposition shows that composition of **LMPs** corresponds to functional composition of factorizations.

Proposition 29. *Given LMPs $\mathcal{S}: X \rightarrow Y$ and $\mathcal{T}: Y \rightarrow Z$, $(\mathcal{S} \circ \mathcal{T})^\wedge = \widehat{\mathcal{S}} \circ \widehat{\mathcal{T}}$.*

Proof (sketch). The key is to show how $\widehat{\mathcal{T}}$ acts as a ‘predicate transformer’. In particular, given $x \in X$, one shows by structural induction on $f \in \mathbb{A}[Z]$ that $(\mathcal{S} \circ \mathcal{T})^\wedge_x(f) = \widehat{\mathcal{S}}_x(\widehat{\mathcal{T}}(f))$. □

Iteration of **LMPs** and fixed points of factorizations also correspond via the duality. A nice consequence of this is the fact that the iteration theory identities for the dagger operation on **LMPs** all follow by virtue of the standard construction used in Definition 27 (see [5, Chapter 8, Theorem 2.15]). The reader may compare with the proof in [4] that, modulo bisimilarity, terms in a calculus for regular probabilistic processes form an iteration theory.

Proposition 30. *If $\mathcal{S}: X \rightarrow X + Y$ is a LMP then $(\mathcal{S}^\dagger)^\wedge = (\widehat{\mathcal{S}})^\dagger$.*

9 Summary and Future Work

In programming and semantical frameworks, there are usually many different ways to represent the same computational behaviours. In concurrency, canonical representatives of the equivalence classes of bisimilar processes are represented as elements of final coalgebras, often constructed in categories of domains. The applicability of such theories hinges on convenient representations of those elements. The final coalgebra capturing **LMPs** has been described in [7]. The domain-theoretic treatment is in [11,8]. However the issue of representability had so far not been tackled. In the present paper, a method for obtaining canonical representatives of **LMPs** has been presented. Their states are represented as simple monoid homomorphisms. The simplicity of this representation supports hope for a wider practical applicability of the **LMP** model.

Our application of Stone duality for C^* -algebras to derive canonical representatives of **LMPs** is an instance of a general approach to representing computational behaviours by lifting dualities, and adjunctions. A detailed account of this general framework, with applications to other computational structures, will be described in forthcoming work.

Unlike the papers [9,11,8] we have not emphasized the measure-theoretic aspects of **LMPs**, but instead focused on the discrete case. As we already said, the idea was to communicate the essential concepts with the minimum overhead. However, another reason for this policy is that treating **LMPs** at the level of measurable spaces sits rather uneasily with the assumption of finite sets of entry and exit points. This suggests that an interesting direction for further work would be to allow the domain and codomain of an **LMP** to be measurable spaces. This would yield a category of measurable spaces and **LMPs**. It would be interesting to compare such a category to the category of probabilistic relations studied in [3].

Finally, we intend to investigate connections between our representation of **LMPs** and the notion of formal tree series [12].

References

1. S. Abramsky. A Domain Equation for Bisimulation. *Information and Computation* 92:161–218, 1991.
2. S. Abramsky and S. Vickers. Quantales, observational logic and process semantics. *Mathematical Structures in Computer Science*, 3:161–227, 1993.
3. S. Abramsky, R. Blute and P. Panangaden. Nuclear and Trace Ideals in Tensor- $*$ -categories. *Journal of Pure and Applied Algebra*, 143:3–47, 1999.
4. L. Aceto, Z. Ésik and A. Ingólfssdóttir. Equational Axioms for Probabilistic Bisimilarity. In *Proceedings of 9th AMAST*, Lecture Notes in Computer Science, volume 2422, pages 239–253, 2002.
5. S. Bloom and Z. Esik. *Iteration Theories*. EATCS Monographs on Theoretical Computer Science. Springer, 1993.
6. S. Bloom and Z. Esik. The Equational Logic of Fixed Points. *Theoretical Computer Science*, 179:1–60.

7. F. van Breugel and J. Worrell. An Algorithm for Quantitative Verification of Probabilistic Transition Systems. In *Proceedings of CONCUR'01*, volume 2154 of *LNCS*, Springer-Verlag, 2001.
8. F. van Breugel, M. Mislove, J. Ouaknine and J. Worrell. Domains, Testing and Similarity for Labelled Markov Processes. To appear in *Proceedings of FOSSACS'03, Theoret. Comp. Sci.*, 2003.
9. J. Desharnais, A. Edalat and P. Panangaden. Bisimulation for Labelled Markov Processes. *Information and Computation*, 179(2):163–193, 2002.
10. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for Labeled Markov Processes. In *Proc. 10th International Conference on Concurrency Theory*, volume 1664 of *LNCS*, Springer-Verlag, 1999.
11. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating Labeled Markov Processes. *Information and Computation*, 184(1):160–200, 2003.
12. Z. Esik and W. Kuich. Formal Tree Series. *Journal of Automata Languages and Combinatorics*, 8(2):145–185, 2003.
13. M.A. Naimark. *Normed Rings*, 2nd ed., Nauka, Moscow 1968; reprint of the revised English edition, Wolters-Noordhoff, Groningen 1970.
14. R. van Glabbeek, S. Smolka, B. Steffen. Reactive, generative and stratified models of probabilistic processes. *Information and Computation*, 121(1):59–80, 1996.
15. D. Kozen. The Semantics of Probabilistic Programs. *Journal of Computer and System Science*, 22:328–350, 1981.
16. P. Johnstone. *Stone Spaces*. Cambridge University Press, 1982.
17. B. Jonsson, K. Larsen and W. Yi. Probabilistic Extensions of Process Algebras. In J.A. Bergstra, A. Ponse and S. Smolka, editors, *Handbook of Process Algebra*, pages 685–710, Elsevier, 2001.
18. K.G. Larsen and A. Skou. Bisimulation through Probabilistic Testing. *Information and Computation*, 94(1):1–28, 1991.
19. M. Mislove, J. Ouaknine and J. Worrell. Axioms for Probability and Nondeterminism. In *Proc. EXPRESS'03, ENTCS 91(3)*, 2003.
20. K.R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, 1967.
21. A. Di Pierro, C. Hankin, and H. Wiklicky. Quantitative Relations and Approximate Process Equivalences. In *Proceedings of CONCUR'03*, volume 2761 of *LNCS*, Springer-Verlag, 2003.
22. E.W. Stark and S.A. Smolka. A complete axiom system for finite-state probabilistic processes. In *Proof, Language, and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000.

A Semiring Facts

Let \mathbb{R}_+^∞ denote the semiring of nonnegative reals with the usual addition and multiplication, and with ∞ adjoined. In this semiring we have $\infty \cdot 0 = 0 \cdot \infty = 0$ and $\infty + a = a + \infty = \infty$. Let $\text{Mat } \mathbb{R}_+^\infty$ denote the semiring of matrices over \mathbb{R}_+^∞ . This semiring is ω -complete (it has countable sums), so there is a Kleene-* operation given by $A^* = \sum_{k=0}^{\infty} A^k$.

Proposition 31. *Let A be an $n \times n$ matrix and B an $n \times p$ matrix in $\text{Mat } \mathbb{R}_+^\infty$. If $[A \ B]$ is a sub-stochastic matrix, i.e. the entries on each row have sum no greater than 1, then A^*B is a sub-stochastic matrix.*