Maximal Traces and Path-Based Coalgebraic Temporal Logics

Corina Cîrstea

University of Southampton, UK

Abstract

This paper gives a general coalgebraic account of temporal logics whose semantics involves a notion of computation path. Examples of such logics include the logic CTL* for transition systems and the logic PCTL for probabilistic transition systems. Our path-based temporal logics are interpreted over coalgebras of endofunctors obtained as the composition of a computation type (e.g. nondeterministic or stochastic) with a general transition type. The semantics of such logics relies on the existence of execution maps similar to the trace maps introduced by Jacobs and co-authors as part of the coalgebraic theory of finite traces [1]. We consider both *finite* execution maps derived from the theory of finite traces, and a new notion of *maximal* execution map that accounts for maximal, possibly infinite executions. The latter is needed to recover the logics CTL* and PCTL as specific path-based logics.

Keywords: coalgebra, trace semantics, computation path, temporal logic, nondeterminism, probability

1. Introduction

Path-based temporal logics are commonly used as specification logics, particularly in the context of automatic verification. Instances of such logics include the logic CTL* with its fragments CTL and LTL for transition systems [2], and the logic PCTL for probabilistic transition systems [3]. In spite of the similarities shared by these logics (most notably the use of a notion of

Email address: `cc2@ecs.soton.ac.uk` (Corina Cîrstea)
URL: `http://www.ecs.soton.ac.uk/people/cc2` (Corina Cîrstea)

computation path to define their semantics), no general, unified account of path-based temporal logics exists.

Coalgebras are by now recognised as a truly general model of dynamical systems, instances of which subsume transition systems, their probabilistic counterparts, and many other interesting state-based models [4]. Moreover, the modal logics associated with coalgebraic models [5, 6, 7] are natural logics for specifying system behaviour, that also instantiate to familiar logics in particular cases. Basic coalgebraic modal languages (as considered e.g. in [5, 6]) employ modal operators whose semantics depends solely on the one-step behaviour of system states. Adding fixpoint operators (with the usual semantics) to such languages allows properties of the long-term, possibly infinite behaviour of system states to also be formalised [7, 8]. However, the use of fixpoint operators makes the formulation of application-relevant temporal properties a non-trivial task (see Example 5.2 for an illustration of this). In contrast, the syntax and semantics of temporal logics such as CTL* and PCTL make direct reference to the computation paths associated to a state in a model, thereby easing the task of formalising application-relevant temporal properties. While the relationship between CTL* and the modal μ-calculus [9] is well understood [10], that between PCTL and the fixpoint extension of the basic modal language for probabilistic systems (as considered e.g. in [8]) is not. In particular, it is unclear whether properties such as: "the likelihood of a state property p holding eventually is greater than q" can be formalised in the latter language (while this can easily be encoded in PCTL). This leads to a more general question regarding the expressive power of path-based temporal logics, and motivates the need to further investigate such logics.

The present paper makes some initial steps towards a general coalgebraic theory of path-based temporal logics: we introduce a generic syntactic format for such logics, together with a coalgebraic semantics defined in terms of execution maps. Following [11, 1], we model systems as coalgebras of a signature functor obtained as the composition of a computation type T (called *branching type* in [1]) with a transition type F, and require that T distributes over F in a suitable way. As examples, we consider nondeterministic and probabilistic systems, with the non-empty powerset functor $\mathcal{P}^+ : \mathsf{Set} \to \mathsf{Set}$ on the category of sets and respectively the probability measure functor $\mathcal{G}_1 : \mathsf{Meas} \to \mathsf{Meas}$ on the category of measurable spaces describing the computation types needed to recover the usual notions of computation path for such systems. While the transition type describes the structure of *individual* transitions (typically linear), the computation type describes how the transitions from particular states are organised (e.g. using sets, or probability distributions). Our semantics for path-based temporal logics for $T \circ F$-coalgebras relies on a notion of computation path (that is parameterised by T and F), and on the existence of so-called execution maps taking states of $T \circ F$-coalgebras to suitably-structured computation paths. The notions of *finite trace* and *finite trace map* provided by the coalgebraic theory of finite traces [1] can easily be adapted to provide notions of finite computation path and finite execution map. However, while such notions *can* be used to provide semantics for path-based coalgebraic temporal logics, their use does not allow logics such as CTL* and PCTL, whose semantics involves *infinite* computation paths, to be recovered as instances of the general framework.

The first contribution of this paper is to define notions of *maximal execution* and *maximal execution map* for deterministic, non-deterministic and stochastic computation types (and general transition types). In particular, maximal execution maps arise as instances of *maximal trace maps* (which we also define), by simply varying the transition type. Our use of the term *maximal* (instead of *infinite*) reflects the observation that, for certain choices of transition type, some of the possible maximal traces admit finite descriptions (see e.g. Example 3.3). Our approach to defining maximal trace maps is inspired by the work in [11], where *infinite trace maps* were defined for coalgebras of type $\mathcal{P} \circ F$, with $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$ the powerset functor and $F : \mathsf{Set} \to \mathsf{Set}$ a polynomial functor. At the same time, our definitions and results are not direct generalisations of those in [11] – the approach described in this paper only applies to computation types given by *affine* monads, with only the *non-empty* powerset monad $\mathcal{P}^+$ (and not the powerset monad itself) falling in this category[1]. The difference between the two approaches is more accurately summarised by the following points:

- When restricting to $\mathcal{P} \circ F$-coalgebras that are also $\mathcal{P}^+ \circ F$-coalgebras (that is, each state has at least one successor), the infinite trace maps of [11] coincide with the maximal trace maps defined in this paper. (The infinite trace maps of [11] assign an empty set of traces to states of $\mathcal{P} \circ F$-coalgebras with no successors.)

[1]The study of arbitrary (non-affine) monads is left for future work.

- Our results can be applied to arbitrary $\mathcal{P} \circ F$-coalgebras by regarding them as $\mathcal{P}^+ \circ (1 + F)$-coalgebras (where the coalgebra map takes states with no successors to $\{\iota_1(*)\}$). The resulting maximal trace maps differ from the trace maps of [11] for the original $\mathcal{P} \circ F$-coalgebras in that they also account for the maximal *finite* traces arising from the presence of states with no successors (as discussed in Example 3.3).

The second contribution of this paper is the definition of path-based coalgebraic temporal logics. These are parameterised on:

- the choice of computation and transition types, as well as the notion of execution map,
- a choice of basic modal operators (and associated one-step semantics) for both the computation type and the transition type.

The syntax of such logics distinguishes between *path* and *state* formulas, with the interpretation of the latter being defined in terms of execution maps. By instantiating our approach, we recover known temporal logics and obtain new variants of known logics. Specifically, taking T to be the non-empty powerset monad $\mathcal{P}^+ : \mathsf{Set} \to \mathsf{Set}$ and $F = \mathsf{Id} : \mathsf{Set} \to \mathsf{Set}$ sheds new light on the logic CTL* [2]. Varying F to $A \times \mathsf{Id}$ with A a set of labels yields an interesting variant of CTL* interpreted over labelled transition systems. On the other hand, taking $T = \mathcal{G}_1$ and $F = \mathsf{Id}$ allows us to recover the logic PCTL [3], as well as to obtain a version of this logic interpreted over standard Borel spaces. Specifically, the negation-free fragments of CTL* and PCTL are recovered as path-based fixpoint logics (for $\mathcal{P}^+$- and respectively $\mathcal{D}$-coalgebras, with $\mathcal{D} : \mathsf{Set} \to \mathsf{Set}$ the probability distribution monad), whereas the full logics are obtained as fragments of path-based temporal logics with Until operators (for the same functors). All of the above instantiations rely on the notion of maximal execution introduced in this paper.

This paper is structured as follows. The remainder of this section gives a brief overview of the logics CTL* and PCTL, our main examples. Section 2 recalls some basic definitions and results required later in the paper, as well as some details of the generic theory of finite traces [1]. Section 3 defines maximal traces and executions for deterministic, non-deterministic and stochastic computation types. Section 4 defines the syntax and semantics of general path-based coalgebraic logics, including fixpoint logics (with no negation operator for either the path or the state formulas) and temporal logics with Until operators. A summary of the results and an outline of future work are given in Section 5. This paper is an extended and revised version of [12].

Transition systems and the logic CTL.* The semantics of CTL* [13] is based upon the notion of computation path. Given a transition system with set of states S and accessibility relation $R \subseteq S \times S$, a *computation path* from a state $s_0 \in S$ is an infinite sequence of states $s_0 s_1 s_2 \ldots$ such that $s_i R s_{i+1}$ for $i \in \omega$. The syntax of CTL* consists of *path formulas* φ, formalising properties of computation paths, and *state formulas* Φ, formalising properties of states:

$$\begin{array}{rcl} \varphi & ::= & \Phi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U} \varphi \\ \Phi & ::= & \mathsf{tt} \mid p \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathbf{A}\varphi \end{array}$$

The path formulas of CTL* employ the temporal operators **X** (in the *neXt* state along the path) and **U** (*Until* operator). Additional temporal operators **F** (at some *Future* state along the path) and **G** (*Globally* along the path) can be defined by letting $\mathbf{F}\varphi ::= \mathsf{tt}\mathbf{U}\varphi$ and $\mathbf{G}\varphi ::= \neg\mathbf{F}\neg\varphi$. The state formulas of CTL* use atomic propositions p (interpreted as subsets of the state space of a transition system) to capture basic properties of states, and the operator **A** to quantify universally over the computation paths from a particular state. Existential quantification over paths is then captured by the derived operator **E**, defined by $\mathbf{E}\varphi ::= \neg\mathbf{A}\neg\varphi$. Every state formula is also a path formula, with the latter requiring that the first state of a path satisfies the given state formula. For example, the property "along every path, the system will eventually reach a success state" is formalised as $\mathbf{A}(\mathsf{tt}\,\mathbf{U}\,\mathit{success})$, where tt denotes the true proposition and *success* is an atomic proposition. In order to only focus on the *infinite* computation paths as defined above, an assumption is made when interpreting CTL* on a transition system, namely that each state has at least one outgoing transition[2] (and hence, all maximal paths through the transition system are infinite).

Probabilistic transition systems and the logic PCTL. In the probabilistic transition system model, the state transitions are governed by a probability distribution on the target states – this assigns a probability value to each outgoing transition from a particular state, with the values for transitions

[2]For states where this is not the case, self-loops can be added to the original transition system.

from the same state summing up to 1. The logic PCTL [3] for probabilistic transition systems is similar in spirit to CTL*, and employs the same notion of computation path as that of CTL*. Its syntax consists of path formulas φ and state formulas Φ, with operators $\mathbf{X}$ and $\mathbf{U}$ (now applied only to state formulas) for the path formulas, and with state formulas $[\varphi]_{\geq q}$ and $[\varphi]_{>q}$ stating that the likelihood of a path formula φ holding along the paths from a particular state is at least, respectively strictly greater than, q:

$$\begin{array}{rcl} \varphi & ::= & \mathbf{X}\Phi \mid \Phi\mathbf{U}\Phi \\ \Phi & ::= & \mathsf{tt} \mid p \mid \neg\Phi \mid \Phi \wedge \Phi \mid [\varphi]_{\geq q} \mid [\varphi]_{>q} \end{array}$$

For example, $[\mathsf{tt}\,\mathbf{U}\,success]_{\geq 1}$ states that the likelihood of eventually reaching a success state is 1. To interpret the state formulas of PCTL on a probabilistic transition system, one computes probability measures over the computation paths from each state (see [3] for details).

The previous examples suggest that a general account of computation paths (to be referred to as *maximal executions* in what follows) should first define the shape of a maximal execution (in the previous examples, any infinite sequence of states), and then provide a suitable structure on the maximal executions (e.g. a *subset* of all possible executions, or a *probability measure* over them), for each state of a particular model. The former should be sufficient to allow an interpretation of path formulas (of a generic path-based logic yet to be defined), whereas the latter should allow an interpretation of state formulas (of the same logic).

Acknowledgement

Thanks are due to Bart Jacobs for pointing out the notion of affine monad and its relevance to this work, and to the anonymous referees for their valuable comments and suggestions for improvement.

2. Preliminaries

We recall that a *measurable space* is given by a pair (X, Σ_X) with X a set and Σ_X a σ-algebra of (*measurable*) subsets of X, whereas a *measurable map* between (X, Σ_X) and (Y, Σ_Y) is given by a function $f : X \to Y$ with the property that $f^{-1}(V) \in \Sigma_X$ for each $V \in \Sigma_Y$. We write Meas for the category of measurable spaces and measurable maps. A measurable space (X, Σ_X) is called *discrete* if $\Sigma_X = \mathcal{P}X$. A *subprobability measure* on a measurable space (X, Σ_X) is then a function $\mu : \Sigma_X \to [0,1]$ such that $\mu(\emptyset) = 0$ and $\mu(\bigcup_{i\in\omega} X_i) = \sum_i \mu(X_i)$ for countable families $(X_i)_{i\in\omega}$ of pairwise-disjoint measurable subsets of X. Thus, $\mu(X) \leq 1$ for any subprobability measure μ on (X, Σ_X). If $\mu(X) = 1$, then μ is called a *probability measure*. Given a measurable space (X, Σ_X) and $x \in X$, the *Dirac probability measure* δ_x is defined by $\delta_x(U) = 1$ iff $x \in U$ and $\delta_x(U) = 0$ otherwise.

We write $\mathcal{G} : \mathsf{Meas} \to \mathsf{Meas}$ for the *subprobability measure functor* [14], sending a measurable space (X, Σ_X) to the set $\mathcal{M}(X, \Sigma_X)$ of subprobability measures on (X, Σ_X), equipped with the σ-algebra generated by the sets $\{\mu \mid \mu(U) \geq q\}$ with $U \in \Sigma_X$ and $q \in [0,1]$. A related functor, considered in [1], is the *subprobability distribution functor* $\mathcal{S} : \mathsf{Set} \to \mathsf{Set}$, sending a set X to the set of *subprobability distributions over* X, i.e. functions $\mu : X \to [0,1]$ with $\sum_{x\in X} \mu(x) \leq 1$[3].

For technical reasons to be discussed later (see Section 3.4), we will work in a full subcategory of Meas, namely the category SB of *standard Borel spaces* – these are the measurable spaces whose measurable sets arise as the Borel sets induced by a complete, separable metric (see [15] for further details). A notable property of this subcategory is that it is closed under countable coproducts and countable limits in Meas (see [16, Fact 1]).

Given a functor $F : \mathsf{C} \to \mathsf{C}$, an *$F$-coalgebra* is given by a pair (X, γ) with X a C-object and $\gamma : X \to FX$ a C-map, while an *F-coalgebra homomorphism* from (X, γ) to (Y, δ) is given by a C-map $f : X \to Y$ additionally satisfying $Ff \circ \gamma = \delta \circ f$. As previously mentioned, we work in the setting of coalgebras of endofunctors obtained as the composition of a computation type with a transition type. The computation type is specified by a monad T on a category C, whereas the transition type is captured by an endofunctor F on C. As in [1], a crucial assumption is the existence of a distributive law $\lambda : F \circ T \Rightarrow T \circ F$ of T over F. Such a distributive law must be compatible with the monad structure, i.e. $\lambda \circ F\eta = \eta_F$ and $\lambda \circ F\mu = \mu_F \circ T\lambda \circ \lambda_T$, where $\eta : \mathsf{Id} \Rightarrow T$ and $\mu : T^2 \Rightarrow T$ denote the unit and multiplication of the monad T.

As examples of computation types, we consider (variants of):

- the identity monad $\mathsf{Id} : \mathsf{Set} \to \mathsf{Set}$, modelling deterministic computa-

[3] Thus, a subprobability distribution can take non-zero values on at most countably-many elements of X.

tions, with unit and multiplication given by identities,

- the environment monad $\mathcal{E} := \mathsf{Id}^E : \mathsf{Set} \to \mathsf{Set}$ with E a fixed set, modelling deterministic computations with input, with unit $\eta_X : X \to X^E$ given by $\eta_X(x)(e) = x$, and multiplication $\mu_X : (X^E)^E \to X^E$ given by $\mu_X(f)(e) = f(e, e)$,

- the powerset monad $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$, modelling nondeterministic computations, with unit given by singletons and multiplication given by unions,

- the subprobability measure monad $\mathcal{G} : \mathsf{Meas} \to \mathsf{Meas}$, modelling probabilistic computations, with unit given by the Dirac measures and multiplication given by integration (see [14] for details).

All of the above monads are *strong* and *commutative*, i.e. they come equipped with a *strength map* $\mathsf{st}_{X,Y} : X \times TY \to T(X \times Y)$ as well as a *double strength map* $\mathsf{dst}_{X,Y} : TX \times TY \to T(X \times Y)$, for each choice of C-objects X, Y[4]:

- the identity monad has strength and double strength given by identities,

- the environment monad has strength given by

$$\mathsf{st}_{X,Y}(x, f)(e) = (x, f(e))$$

  for $x \in X$, $f \in Y^E$ and $e \in E$, and double strength given by given by the isomorphism $X^E \times Y^E \simeq (X \times Y)^E$,

- the powerset monad has strength and double strength given by

$$\mathsf{st}_{X,Y}(x, V) = \{x\} \times V \qquad\qquad \mathsf{dst}_{X,Y}(U, V) = U \times V$$

  for $x \in X$, $U \in \mathcal{P}X$ and $V \in \mathcal{P}Y$,

- the subprobability measure monad has strength given by

$$\mathsf{st}_{(X,\Sigma_X),(Y,\Sigma_Y)}(x, \nu)(U \times V) = \begin{cases} \nu(V) & \text{if } x \in U \\ 0 & \text{otherwise} \end{cases}$$

[4]Moreover, these are natural in X and Y.

and double strength given by

$$\mathsf{dst}_{(X,\Sigma_X),(Y,\Sigma_Y)}(\mu,\nu)(U\times V)=\mu(U)\cdot\nu(V)$$

for $x\in X$, $\mu\in\mathcal{M}(X,\Sigma_X)$, $\nu\in\mathcal{M}(Y,\Sigma_Y)$, $U\in\Sigma_X$ and $V\in\Sigma_Y$. (Note that the σ-algebra of the product $(X,\Sigma_X)\times(Y,\Sigma_Y)$ is generated by the subsets $U\times V$ with $U\in\Sigma_X$ and $V\in\Sigma_Y$.)

A particular class of transition types, namely that of *shapely polynomial functors*, is considered in [1].

Definition 1. Let C be a category with finite products and arbitrary coproducts. A functor $F:\mathsf{C}\to\mathsf{C}$ is a *shapely polynomial functor* if it is built from identity and constant functors using finite products and arbitrary coproducts.

[1, Lemma 2.3] shows that any commutative monad on Set has a canonical distributive law over any shapely polynomial functor on Set. This immediately provides examples of distributive laws of the powerset monad over shapely polynomial functors on Set.

Example 1. For $T=\mathcal{P}$ and $F=A\times\mathsf{Id}$, the canonical distributive law of T over F is defined from the canonical distributive laws of $\mathcal{P}$ over A and Id, respectively, using the double strength of the monad $\mathcal{P}$:

$$(A\times\mathsf{Id})\circ\mathcal{P}=A\times\mathcal{P}\xrightarrow{\eta_A\times id_{\mathcal{P}}}\mathcal{P}A\times\mathcal{P}\xrightarrow{\mathsf{dst}_{A,\mathsf{Id}}}\mathcal{P}\circ(A\times\mathsf{Id})$$

Here, the A-component of the unit of $\mathcal{P}$ gives the canonical distributive law of $\mathcal{P}$ over A, while the identity natural transformation provides the canonical distributive law of $\mathcal{P}$ over Id. Later in the paper, we will consider a submonad of the powerset monad, namely the non-empty powerset monad $\mathcal{P}^+:\mathsf{Set}\to\mathsf{Set}$. Its canonical distributive law over F is obtained in a similar way.

The construction of the canonical distributive law (by induction on the structure of the shapely functor) generalises straightforwardly to any category with finite products and arbitrary coproducts, thereby also providing examples of distributive laws of the subprobability measure monad over shapely polynomial functors on Meas.

As in [1], the *Kleisli category* of a monad (T,η,μ) on a category C will play an important rôle when defining notions of maximal trace and maximal execution for systems whose computation type is given by T. This category, denoted $\mathsf{Kl}(T)$, has the same objects as C, and maps from X to Y given by C-maps $f : X \to TY$. The composition of two $\mathsf{Kl}(T)$-maps $f : X \to Y$ and $g : Y \to Z$ is given by the C-map $\mu_Z \circ Tg \circ f : X \to TZ$.

We let $K : \mathsf{Kl}(T) \to \mathsf{C}$ denote the functor defined by:

- $KX = TX$,
- $Kf = \mu_Y \circ Tf$ for $f : X \to Y$ in $\mathsf{Kl}(T)$,

and write $J : \mathsf{C} \to \mathsf{Kl}(T)$ for its left adjoint, defined by:

- $JX = X$,
- $Jf = Tf \circ \eta_X = \eta_Y \circ f$ for $f : X \to Y$ in C.

Later we will make use of the following property of the functor J:

Lemma 1. *If the functor $T : \mathsf{C} \to \mathsf{C}$ (weakly) preserves the limit $(Z, (\pi_i)_{i\in\omega})$ of an ω^{op}-chain $(f_i : Z_{i+1} \to Z_i)_{i\in\omega}$*[5]*, then so does $J : \mathsf{C} \to \mathsf{Kl}(T)$.*

Proof. Assume first that T weakly preserves the limit $(Z, (\pi_i : Z \to Z_i)_{i\in\omega})$ of $(f_i : Z_{i+1} \to Z_i)_{i\in\omega}$. To show that $(JZ, (J\pi_i : JZ \to JZ_i)_{i\in\omega})$ is a weakly limiting cone for $(Jf_i : JZ_{i+1} \to JZ_i)_{i\in\omega}$ in $\mathsf{Kl}(T)$, let $(X, (\delta_i : X \to JZ_i)_{i\in\omega})$ denote an arbitrary cone for $(Jf_i)_{i\in\omega}$ in $\mathsf{Kl}(T)$. Hence, in C, $\mu_{Z_i} \circ T\eta_{Z_i} \circ Tf_i \circ \delta_{i+1} = \delta_i$, that is, $Tf_i \circ \delta_{i+1} = \delta_i$ for all $i \in \omega$. This makes $(\delta_i)_{i\in\omega}$ a cone over $(Tf_i)_{i\in\omega}$ in C, and the weak limiting property of $(TZ, (T\pi_i)_{i\in\omega})$ in C now yields a mediating map $m : X \to TZ$ such that $T\pi_i \circ m = \delta_i$ in C for all $i \in \omega$. This is equivalent to $\mu_{Z_i} \circ T\eta_{Z_i} \circ T\pi_i \circ m = \delta_i$ in C for $i \in \omega$, that is, $J\pi_i \circ m = \delta_i$ in $\mathsf{Kl}(T)$ for $i \in \omega$. The proof of the stronger statement, in the case when T preserves the limit of $(f_i)_{i\in\omega}$, is similar. □

Remark 1. The above result will later be instantiated with $T = \mathcal{P}^+ : \mathsf{Set} \to \mathsf{Set}$ and $T = \mathcal{G}_1 : \mathsf{SB} \to \mathsf{SB}$. While $T = \mathcal{G}_1$ preserves limits of ω^{op}-chains, $T = \mathcal{P}^+$ preserves such limits only weakly.

[5] T is said to *weakly preserve the limit* $(Z, (\pi_i : Z \to Z_i)_{i\in\omega})$ *of* $(f_i : Z_{i+1} \to Z_i)_{i\in\omega}$ if for any cone $(X, (g_i : X \to Z_i)_{i\in\omega})$ over $(Tf_i)_{i\in\omega}$ in C, there exists a mediating map $g : X \to TZ$ satisfying $T\pi_i \circ g = g_i$ for $i \in \omega$. If, for any such cone, the mediating map is unique, then T preserves the limit in the standard sense.

As mentioned above, we assume the existence of a distributive law λ of the monad T over the endofunctor F. Such distributive laws are known to be in one-to-one correspondence with *liftings* of the functor $F : \mathsf{C} \to \mathsf{C}$ to $\mathsf{Kl}(T)$, i.e. with functors $G : \mathsf{Kl}(T) \to \mathsf{Kl}(T)$ satisfying $G \circ J = J \circ F$ (see e.g. [1]). The lifting $\overline{F} : \mathsf{Kl}(T) \to \mathsf{Kl}(T)$ induced by a distributive law $\lambda : F \circ T \Rightarrow T \circ F$ is given by:

- $\overline{F}A = FA$,
- $\overline{F}f = \lambda_B \circ Ff$ for $f : A \to B$ in $\mathsf{Kl}(T)$.

To see that the above defines a lifting of F to $\mathsf{Kl}(T)$, note that, for $f : X \to Y$ in C, the C-maps that define the Kleisli maps $\overline{F}Jf$ and JFf are $\lambda_Y \circ F\eta_Y \circ Ff$ and respectively $\eta_{FY} \circ Ff$. By the compatibility of the distributive law λ with the monad structure, these coincide.

Finite traces and executions

In [1], the authors consider coalgebras (X, γ) of endofunctors of the form $T \circ F$ with the monad $T : \mathsf{Set} \to \mathsf{Set}$ and the endofunctor $F : \mathsf{Set} \to \mathsf{Set}$ being related by a distributive law $\lambda : F \circ T \Rightarrow T \circ F$. Moreover the Kleisli category of T is assumed to be *$DCpo_\perp$-enriched*. That is, each homset $\mathsf{Kl}(T)(X, Y)$ is a partial order with bottom element, with directed collections of maps $(f_i : X \to Y)_{i \in I}$ in $\mathsf{Kl}(T)$ admitting a join $\bigsqcup_{i \in I} f_i : X \to Y$, and with composition preserving directed joins: $g \circ (\bigsqcup_{i \in I} f_i) = \bigsqcup_{i \in I}(g \circ f_i)$ and $(\bigsqcup_{i \in I} f_i) \circ h = \bigsqcup_{i \in I}(f_i \circ h)$. In this setting, the elements of the carrier I_F of the initial F-algebra provide the potential *finite traces* of states of $T \circ F$-coalgebras[6], and a *finite trace map* $\mathsf{ftr}_\gamma : X \to T(I_F)$ is defined via finality in $\mathsf{Kl}(T)$. The crucial observation is that the initial F-algebra in Set lifts to a final $\overline{F}$-coalgebra in $\mathsf{Kl}(T)$ (where, as before, $\overline{F} : \mathsf{Kl}(T) \to \mathsf{Kl}(T)$ is the lifting of F to $\mathsf{Kl}(T)$ induced by λ). Thus, the finite trace map arises as the unique coalgebra morphism from the $\overline{F}$-coalgebra in $\mathsf{Kl}(T)$ induced by a $T \circ F$-coalgebra in Set to the final $\overline{F}$-coalgebra.

A *finite execution map* for a $T \circ F$-coalgebra (X, γ) is defined in [17], as the finite trace map obtained by regarding (X, γ) as a $T \circ F \circ (X \times \mathsf{Id})$-coalgebra. Here we propose a variant of this notion obtained by replacing the

[6]The resulting notion of trace is referred to as *fat trace* in [17], as it retains the structure specified by the transition type F and therefore may involve branching.

functor $F \circ (X \times \mathsf{Id})$ with the functor $X \times F$. The reason for this variation is that we expect finite executions to also record their initial states. This is needed if finite execution maps are used to provide semantics to path-based temporal logics (see Section 4). In order to view a $T \circ F$-coalgebra (X, γ) as a $T \circ (X \times F)$-coalgebra, we post-compose the map $\langle id_X, \gamma \rangle : X \to X \times TFX$ with the appropriate component $\mathsf{st}_{X,FX} : X \times TFX \to T(X \times FX)$ of the strength of the monad T.

Definition 2. Let $T : \mathsf{C} \to \mathsf{C}$ be a strong monad, let $F : \mathsf{C} \to \mathsf{C}$ be an endofunctor, and let $\lambda : F \circ T \Rightarrow T \circ F$ be a distributive law of T over F. Also, for a $T \circ F$-coalgebra (X, γ), let $F_X : \mathsf{C} \to \mathsf{C}$ denote the functor taking a C-object Y to $X \times FY$, let (I_X, ι_X) denote an initial F_X-algebra, and let $\lambda_X : F_X \circ T \Rightarrow T \circ F_X$ denote the natural transformation given by $(\lambda_X)_Y = \mathsf{st}_{X,FY} \circ (id_X \times \lambda_Y)$. The *finite execution map* $\mathsf{fexec}_\gamma : X \to TI_X$ is the C-map underlying the unique $\overline{F_X}$-coalgebra morphism from $(X, \mathsf{st}_{X,FX} \circ \langle id_X, \gamma \rangle)$ to the final $\overline{F_X}$-coalgebra.

Example 2. Let $T = \mathcal{P}$ and $F = 1 + A \times \mathsf{Id}$. In this case, the potential finite traces are the elements of the initial F-algebra, that is, all finite sequences of elements of A. Also, given a $T \circ F$-coalgebra (X, γ), the potential finite executions are the elements of the initial F_X-algebra, that is, all finite sequences of the form $s_0 a_1 s_1 a_2 s_2 \ldots s_n$, with $n \in \omega$, $s_i \in X$ for $i \in \{0, \ldots, n\}$ and $a_i \in A$ for $i \in \{1, \ldots, n\}$. We note that taking $F = A \times \mathsf{Id}$ results in *no* possible finite traces or executions, and consequently the finite trace/execution maps will assign the empty set to any state of any $T \circ F$-coalgebra.

Modal logics for coalgebras

Our path-based coalgebraic temporal logics will be based on the notion of predicate lifting, as introduced by Pattinson [5]. However, the semantics of these logics will differ somewhat from the standard semantics of coalgebraic modal logics induced by predicate liftings, as defined e.g. in [5]. Also, the notion of predicate lifting used here is slightly more general than the one of [5], and applies to endofunctors on both Set and Meas.

We begin by fixing a category C with forgetful functor $U : \mathsf{C} \to \mathsf{Set}$, and a contravariant functor $P : \mathsf{C} \to \mathsf{Set}^{\mathrm{op}}$ such that P is a subfunctor of $\hat{\mathcal{P}} \circ U$[7],

[7] That is, for each C-object C, $P(C) \subseteq \hat{\mathcal{P}}(UC)$, and for each C-arrow $f : C \to D$, $P(f)$ is the restriction of $\hat{\mathcal{P}}(Uf)$ to $P(C)$.

where $\hat{\mathcal{P}} : \mathsf{Set} \to \mathsf{Set}^{\mathrm{op}}$ denotes the contravariant powerset functor. For each state space X, PX specifies a set of admissible predicates. As instances of P we will consider the contravariant powerset functor $\hat{\mathcal{P}} : \mathsf{Set} \to \mathsf{Set}^{\mathrm{op}}$ in the case when $\mathsf{C} = \mathsf{Set}$, and the functor taking a measurable space to the carrier of its underlying σ-algebra in the case when $\mathsf{C} = \mathsf{Meas}$.

Given an endofunctor $F : \mathsf{C} \to \mathsf{C}$ and $n \in \omega$, an *n-ary predicate lifting for F* is a natural transformation $\lambda : P^n \Rightarrow P \circ F$. For simplicity of presentation, we assume all predicate liftings to be unary, however, our results generalise straightforwardly to predicate liftings with arbitrary finite arities. We briefly recall the syntax and semantics of coalgebraic modal logics induced by predicate liftings. Given a set Λ of predicate liftings for F, the modal language $\mathcal{L}_\Lambda$ has formulas given by the grammar:

$$\mathcal{L}_\Lambda \ni \Phi ::= \mathsf{tt} \mid \neg\Phi \mid \Phi \wedge \Phi \mid [\lambda]\Phi \qquad (\lambda \in \Lambda)$$

A coalgebraic semantics for this language is obtained by defining $[\![\Phi]\!]_\gamma \subseteq PC$ for each F-coalgebra (C, γ), by structural induction on $\Phi \in \mathcal{L}_\Lambda$. The interesting case is $[\![[\lambda]\Phi]\!]_\gamma = (P\gamma)(\lambda_C([\![\Phi]\!]_\gamma))$ for $\Phi \in \mathcal{L}_\Lambda$ and $\lambda \in \Lambda$. In Section 4, we will see a novel use of modalities arising from predicate liftings, namely to interpret state formulas in path-based temporal logics. There, we will typically require our predicate liftings to be *monotone*[8], in that $A \subseteq B$ implies $\lambda_X(A) \subseteq \lambda_X(B)$ for all X and all $A, B \in PX$.

3. Maximal Traces and Executions

Some initial steps towards a general coalgebraic treatment of maximal (possibly infinite) traces and executions were made in [11], where infinite trace maps were defined for coalgebras of type $\mathcal{P} \circ F$, with $F : \mathsf{Set} \to \mathsf{Set}$ a polynomial functor equipped with the canonical distributive law $\lambda : F \circ \mathcal{P} \Rightarrow \mathcal{P} \circ F$. Specifically, it was observed in [11] that the final F-coalgebra in Set (whose elements represent potential infinite traces) gives rise to a weakly final $\overline{F}$-coalgebra in $\mathsf{Kl}(\mathcal{P})$. Then, for a $\mathcal{P} \circ F$-coalgebra, a trace map was obtained via weak finality, by regarding this coalgebra as an $\overline{F}$-coalgebra in $\mathsf{Kl}(\mathcal{P})$. A canonical choice for the trace map was then provided by the *largest* mediating map. As mentioned earlier, our definition of maximal trace maps will only subsume that of [11] when restricting to $\mathcal{P}^+ \circ F$-coalgebras.

[8]Monotonicity in all arguments would be required in the case of predicate liftings with arbitrary finite arities.

Throughout this section, C denotes a category with countable limits, $F : \mathsf{C} \to \mathsf{C}$ is an endofunctor, $T : \mathsf{C} \to \mathsf{C}$ is a strong monad, and $\lambda : F \circ T \Rightarrow T \circ F$ is a distributive law of T over F.

3.1. Maximal traces

As in [11], the final F-coalgebra provides the potential maximal traces of states of $T \circ F$-coalgebras. We work under the assumption that F preserves the limit of the following ω^{op}-chain

$$1 \xleftarrow{\;!\;} F1 \xleftarrow{\;F!\;} F^2 1 \xleftarrow{\;F^2!\;} \ldots$$

with 1 a final object in C and $! : F1 \to 1$ the unique such map. In this case, the carrier of a final F-coalgebra is obtained as the limit in C of the above ω^{op}-chain. We let $(Z, \zeta : Z \to FZ)$ denote a final F-coalgebra, and write $\pi_i : Z \to F^i 1$ with $i \in \omega$ for the corresponding projections.

We expect the maximal trace map for a coalgebra (X, γ) to be of the form $\mathsf{tr}_\gamma : X \to TZ$. (For instance, when $T = \mathcal{P}$, the maximal trace map should assign to each state of the coalgebra, a *set* of maximal traces.) With this in mind, we define an ω-indexed sequence of maps $(\gamma_i : X \to TF^i 1)_{i \in \omega}$, which we regard as finite approximations of the maximal trace map (following the observation that the elements of $F^i 1$ provide finite approximations of potential maximal traces):

- $\gamma_0 = \eta_1 \circ !_X : X \to T1$, where $!_X : X \to 1$ is the unique such map,
- $\gamma_{i+1} = \mu_{F^{i+1}1} \circ T\lambda_{F^i 1} \circ TF\gamma_i \circ \gamma : X \to TF^{i+1}1$ for $i \in \omega$:

$$X \xrightarrow{\;\gamma\;} TFX \xrightarrow{\;TF\gamma_i\;} TFTF^i 1 \xrightarrow{\;T\lambda_{F^i 1}\;} T^2 F^{i+1} 1 \xrightarrow{\;\mu_{F^{i+1}1}\;} TF^{i+1} 1$$

That is, the maps γ_i arise by unfolding the coalgebra structure i times, and using the distributive law λ of T over F and the monad multiplication to discard inner occurrences of T from the codomain of the maps γ_i. Alternatively, the C-maps γ_i can be defined as maps in $\mathsf{Kl}(T)$ by:

- $\gamma_0 = J!_X$,
- $\gamma_{i+1} = \overline{F}\gamma_i \circ \gamma$ for $i \in \omega$.

Some additional constraints on the monad T are required for the maps $(\gamma_i)_{i\in\omega}$ to define a cone over the ω^{op}-chain $(F^i!)_{i\in\omega}$ in C:

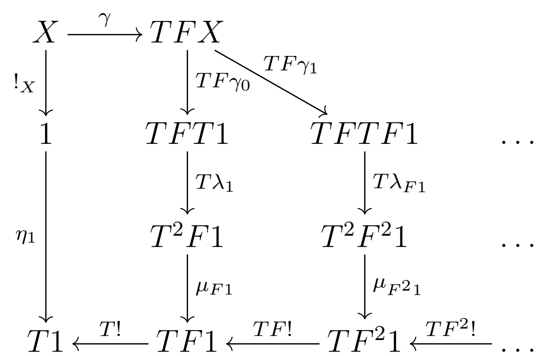

Lemma 2. *Let* $!_{TF1} : TF1 \to 1$ *be the only such map. If* $\eta_1 \circ !_{TF1} = T!$*, then the previously-defined* γ_i*s define a cone over* $(JF^i!)_{i\in\omega}$ *in* $\mathsf{Kl}(T)$*.*

Proof. The following sequence of equalities (in C) ensures $\gamma_0 = J! \circ \gamma_1$:

$$\begin{array}{rcl}
T! \circ \gamma_1 & = & \text{(definition of } \gamma_1\text{)} \\
T! \circ \mu_{F1} \circ T\lambda_1 \circ TF\gamma_0 \circ \gamma & = & \text{(definition of } \gamma_0\text{)} \\
T! \circ \mu_{F1} \circ T\lambda_1 \circ TF\eta_1 \circ TF!_X \circ \gamma & = & \text{(compatibility of } \lambda \text{ with } \eta, \mu\text{)} \\
T! \circ TF!_X \circ \gamma & = & \text{(hypothesis)} \\
\eta_1 \circ !_{TF1} \circ TF!_X \circ \gamma & = & \text{(uniqueness of } \mathsf{C}\text{-maps to 1)} \\
\eta_1 \circ !_X & = & \text{(definition of } \gamma_0\text{)} \\
\gamma_0 & &
\end{array}$$

Now assuming $\gamma_i = JF^i! \circ \gamma_{i+1}$, we immediately obtain $\overline{F}\gamma_i = \overline{F}JF^i! \circ \overline{F}\gamma_{i+1} = JF^{i+1}! \circ \overline{F}\gamma_{i+1}$, where the last equality follows by $\overline{F}$ being a lifting of F to $\mathsf{Kl}(T)$. Pre-composition with γ finally gives $\gamma_{i+1} = JF^{i+1}! \circ \gamma_{i+2}$. □

We immediately observe that the hypothesis of the above result is *not* satisfied by two of the monads identified earlier:

- for $T = \mathcal{P}$, $(\eta_1 \circ !_{TF1})(\emptyset) = 1 \neq \emptyset = (\mathcal{P}!)(\emptyset)$;
- for $T = \mathcal{G}$, $(\eta_1 \circ !_{TF1})(\nu_0) = \mu_1 \neq \mu_0 = (\mathcal{G}!)(\nu_0)$, where ν_0 is the subprobability measure on $F(1, \mathcal{P}1)$[9] which assigns the value 0 to each

[9] Note that $(1, \mathcal{P}1)$ is a final object in Meas.

measurable set, whereas μ_0 and μ_1 are the subprobability measures on $(1, \mathcal{P}1)$ given by $\mu_0(1) = 0$ and respectively $\mu_1(1) = 1$.

To remedy the situation, we will work with submonads of these two monads for which the hypothesis of Lemma 2 is true. To this end, we first note that if the monad T is such that $\eta_1 : 1 \to T1$ is an isomorphism, then the equality required by Lemma 2 is obtained immediately by finality. Strong monads with the above property are called *affine*, see e.g. [18] for an overview. Moreover, [18] shows how to construct, for any strong monad T, its *affine submonad* T_a, which is itself commutative whenever T is. Specifically, the action of T_a on a C-object X is given by the following pullback diagram:

$$\begin{array}{ccc} T_a X & \overset{\iota_X}{\dashrightarrow} & TX \\ {\scriptstyle !_{T_a X}} \downarrow & & \downarrow {\scriptstyle T!_X} \\ 1 & \underset{\eta_1}{\longrightarrow} & T1 \end{array}$$

This construction yields:

- the non-empty powerset monad $\mathcal{P}^+ : \mathsf{Set} \to \mathsf{Set}$ as the affine submonad of $\mathcal{P}$,

- the probability measure monad $\mathcal{G}_1 : \mathsf{Meas} \to \mathsf{Meas}$ (with $\mathcal{G}_1(X, \Sigma_X)$ containing only the probability measures on (X, Σ_X)) as the affine submonad of $\mathcal{G}$,

- the identity monad as the affine submonad of the *lift monad* $1 + \mathsf{Id} : \mathsf{Set} \to \mathsf{Set}$, as well as of the *finite list* and *finite multiset monads* (taking a set X to the set of finite lists, respectively finite multisets, of elements of X).

Also, it is an easy exercise to check that the identity and environment monads are affine. Thus, for $T = \mathsf{Id}$, $T = \mathcal{E}$, $T = \mathcal{P}^+$ and $T = \mathcal{G}_1$, Lemma 2 applies. We also observe that, in the case of $\mathcal{P}^+$ and $\mathcal{G}_1$, the canonical distributive laws of the original monads ($\mathcal{P}$, respectively $\mathcal{G}$) restrict to distributive laws of their affine submonads. This is a consequence of the following general result, stating that any distributive law of a strong monad T over an endofunctor F restricts to a distributive law of T_a over F.

Proposition 1. *Let $\lambda : F \circ T \Rightarrow T \circ F$ be a distributive law of T over F. Then, λ restricts to a distributive law $\lambda : F \circ T_a \Rightarrow T_a \circ F$.*

Proof. Using that $!_{F1} \circ F!_X =!_{FX}$ (by finality of 1), the pullback diagram defining T_aFX can be written as

$$\begin{array}{ccc} T_aFX & \overset{\iota_{FX}}{\dashrightarrow} & TFX \\ \vdots & & \downarrow{\scriptstyle TF!_X} \\ \vdots & & TF1 \\ \vdots & & \downarrow{\scriptstyle T!_{F1}} \\ 1 & \underset{\eta_1}{\longrightarrow} & T1 \end{array}$$

Next, note that the maps $\lambda_X \circ F\iota_X : FT_aX \to TFX$ and $!_{F1} \circ F!_{T_aX} : FT_aX \to 1$ define a cone over the diagram given by $T!_{F1} \circ TF!_X$ and η_1:

$$\begin{array}{rcl} T!_{F1} \circ TF!_X \circ \lambda_X \circ F\iota_X & = & \text{(naturality of } \lambda) \\ T!_{F1} \circ \lambda_1 \circ FT!_X \circ F\iota_X & = & \text{(definition of } T_aX) \\ T!_{F1} \circ \lambda_1 \circ F\eta_1 \circ F!_{T_aX} & = & \text{(compatibility of } \lambda \text{ with monad structure)} \\ T!_{F1} \circ \eta_{F1} \circ F!_{T_aX} & = & \text{(naturality of } \eta) \\ \eta_1 \circ !_{F1} \circ F!_{T_aX} & & \end{array}$$

The definition of T_aFX now yields a map $(\lambda_a)_X : FT_aX \to T_aFX$ that satisfies, in particular, $\iota_{FX} \circ (\lambda_a)_X = \lambda_X \circ F\iota_X$:

$$\begin{array}{ccc} FT_aX & \xrightarrow{F\iota_X} & FTX \\ {\scriptstyle (\lambda_a)_X}\downarrow & & \downarrow{\scriptstyle \lambda_X} \\ T_aFX & \xrightarrow[\iota_{FX}]{} & TFX \end{array}$$

That is, λ_a agrees with λ on FT_aX. The naturality of the resulting maps and their compatibility with the monad structure follow easily by diagram chasing. □

For our two examples ($T = \mathcal{P}^+$ and $T = \mathcal{G}_1$), assuming that F is a shapely polynomial functor, one can simply work with the canonical distributive laws. An easy induction proof shows that these coincide with the distributive laws given by the previous result. However, Proposition 1 shows how to obtain a distributive law of the affine submonad of a monad T over an *arbitrary* endofunctor F from a distributive law $\lambda : FT \Rightarrow TF$.

We now return to the definition of the maximal trace map. For this, we assume that the monad T is affine[10], and moreover, that T preserves the limit $(Z, (\pi_i)_{i\in\omega})$ of an ω^{op}-chain $(F^i! : F^{i+1}1 \to F^i1)_{i\in\omega}$ (and therefore, by Lemma 1, so does J). Since we view the maps $\gamma_i : X \to T^i1$ (with $i \in \omega$) induced by a $T \circ F$-coalgebra (X, γ) as providing finite approximations of the maximal trace map, it is natural to define the maximal trace map $\mathsf{tr}_\gamma : X \to TZ$ by exploiting the preservation by J of the limit $(Z, (\pi_i)_{i\in\omega})$ of $(F^i!)_{i\in\omega}$.

Definition 3. Assume that the monad T is affine, and that the functors F and J preserve the limit $(Z, (\pi_i)_{i\in\omega})$ of the ω^{op}-chain $(F^i!)_{i\in\omega}$. For a $T \circ F$-coalgebra (X, γ), let $(X, (\gamma_i : X \to JF^i1)_{i\in\omega})$ be the induced cone over $(JF^i!)_{i\in\omega}$. The *maximal trace map* of (X, γ) is the unique mediating map $\mathsf{tr}_\gamma : X \to JZ$ arising from the limiting property of $(JZ, (J\pi_i)_{i\in\omega})$ (regarded as a map in C).

In particular, Definition 3 applies to the identity and environment monads, as well as to the probability measure monad. It does not, however, apply to the non-empty powerset monad, since in this case the functor J does not preserve the limit of $(F^i!)_{i\in\omega}$. In Section 3.3, we will show that J *weakly* preserves this limit, which guarantees the existence (but not the uniqueness) of a maximal trace map. A canonical choice for the maximal trace map will be shown to exist in this case. The case $T = \mathcal{G}_1$ will be considered in Section 3.4.

We conclude this section by proving some properties of the maximal trace map, similar to the defining properties of the trace map in [11].

Proposition 2. *Under the assumptions of Definition 3, the maximal trace map $\mathsf{tr}_\gamma : X \to JZ$ defines an $\overline{F}$-coalgebra morphism, that is, $\overline{F}\mathsf{tr}_\gamma \circ \gamma = J\zeta \circ \mathsf{tr}_\gamma$.*

Proof. We begin by noting that the final F-coalgebra $\zeta : Z \to FZ$ satisfies $F\pi_i \circ \zeta = \pi_{i+1}$ for $i \in \omega$; hence, in $\mathsf{Kl}(T)$ we have $JF\pi_i \circ J\zeta = J\pi_{i+1}$ for $i \in \omega$. Also, the preservation by F of the limit $(Z, (\pi_i)_{i\in\omega})$ of $(F^i!)_{i\in\omega}$ results in the cone $(FZ, (F\pi_i)_{i\in\omega})$ over $(F^{i+1}!)_{i\in\omega}$ being a limiting one which, moreover, is isomorphic to the limiting cone $(Z, (\pi_{i+1})_{i\in\omega})$ over the same ω^{op}-chain. Since

[10] A treatment of monads that are *not* affine is outside the scope of this paper.

J preserves the limit of the latter, it also preserves the limit of the former. That is, $(JFZ, (JF\pi_i)_{i\in\omega})$ is a limit of $(JF^{i+1}!)_{i\in\omega}$.

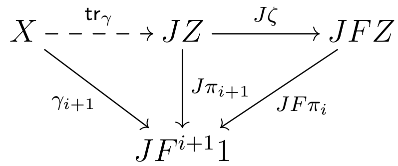

The conclusion then follows by showing that both $\overline{F}\mathsf{tr}_\gamma \circ \gamma$ and $J\zeta \circ \mathsf{tr}_\gamma$ define mediating maps for the cone $(X, (\gamma_{i+1})_{i\in\omega})$ over $(JF^{i+1}!)_{i\in\omega}$. On the one hand, we have:

$$\begin{aligned}
JF\pi_i \circ \overline{F}\mathsf{tr}_\gamma \circ \gamma \quad &= \qquad \text{(definition of } \overline{F}) \\
\overline{F}J\pi_i \circ \overline{F}\mathsf{tr}_\gamma \circ \gamma \quad &= \qquad \text{(definition of } \mathsf{tr}_\gamma) \\
\overline{F}\gamma_i \circ \gamma \quad &= \qquad \text{(definition of } \gamma_{i+1}) \\
\gamma_{i+1} &
\end{aligned}$$

On the other hand, we have:

$$JF\pi_i \circ J\zeta \circ \mathsf{tr}_\gamma = J\pi_{i+1} \circ \mathsf{tr}_\gamma = \gamma_{i+1}$$

Uniqueness of mediating maps for the cone $(X, (\gamma_{i+1})_{i\in\omega})$ over $(JF^{i+1}!)_{i\in\omega}$ now gives $\overline{F}\mathsf{tr}_\gamma \circ \gamma = J\zeta \circ \mathsf{tr}_\gamma$, that is, tr_γ is an $\overline{F}$-coalgebra morphism. □

3.2. Maximal executions

To obtain a notion of maximal execution of a state in a $T \circ F$-coalgebra, we use the approach in the previous section with a different choice of functor F. Similarly to Definition 2, for a $T \circ F$-coalgebra (X, γ), we consider the endofunctor $F_X : \mathsf{C} \to \mathsf{C}$ given by $F_X(Y) = X \times FY$ and the distributive law $\lambda_X : F_X \circ T \Rightarrow T \circ F_X$ given by $(\lambda_X)_Y = \mathsf{st}_{X,FY} \circ (id_X \times \lambda_Y)$. This choice of endofunctor captures the intuition that, in addition to the information provided by a maximal trace, a maximal execution also records the states visited during a particular computation, including the initial state of that computation; hence, the first component of the functor F_X is the state space itself. We assume that F_X preserves the limit of the initial ω^{op}-segment of its final sequence, and call an element of the carrier of the final F_X-coalgebra (Z_X, ζ_X) (obtained as the limit of the previous ω^{op}-sequence) a potential *maximal execution*, or *computation path*.

Definition 4. Let (X, γ) be a $T \circ F$-coalgebra. Assume that the monad T is affine, and that the functors F_X and J preserve the limit $(Z_X, (\pi_i)_{i\in\omega})$ of the ω^{op}-chain $({F_X}^i!)_{i\in\omega}$. Let $(X, (\gamma_i : X \to J{F_X}^i 1)_{i\in\omega})$ be the cone over $(J{F_X}^i!)_{i\in\omega}$ induced by the $T \circ F_X$-coalgebra $(X, \mathsf{st}_{X,FX} \circ \langle id_X, \gamma\rangle)$. The *maximal execution map* $\mathsf{exec}_\gamma : X \to JZ_X$ of (X, γ) is the maximal trace map of the $T \circ F_X$-coalgebra $(X, \mathsf{st}_{X,FX} \circ \langle id_X, \gamma\rangle)$.

Definition 4 yields maximal execution maps for both deterministic systems (with or without input) and probabilistic systems. The next section shows how maximal execution maps can be defined for non-deterministic systems.

3.3. Nondeterministic systems

Definitions 3 and 4 do not apply to coalgebras of type $\mathcal{P}^+ \circ F$, as the functor $J : \mathsf{Set} \to \mathsf{Kl}(\mathcal{P}^+)$ does not preserve limits of ω^{op}-chains. Crucially, J does not preserve the limit of $(F^i!)_{i\in\omega}$. In this section we show that $\mathcal{P}^+$ (and hence, by Lemma 1, also J) weakly preserves limits of ω^{op}-chains, and show how to use this property to define maximal trace and execution maps for $\mathcal{P}^+ \circ F$-coalgebras. As examples, we consider transition systems, both unlabelled and labelled – these are obtained by taking $F = \mathsf{Id}$ and respectively $F = A \times \mathsf{Id}$ with A a set of labels. We note that our use of the non-empty powerset monad agrees with the standard constraint placed on transition systems when defining the notion of computation path.

Remark 2. To see that $\mathcal{P}^+$ does not preserve limits of ω^{op}-chains, consider the final sequence $(f_i : Z_{i+1} \to Z_i)_{i\in\omega}$ of the endofunctor $1 + A \times \mathsf{Id}$, with $Z_i = \bigcup_{0\le j\le i} A^i$, and with limit object $Z = A^* \cup A^\omega$. Now define a cone $(\gamma_i : 1 \to \mathcal{P}^+ Z_i)_{i\in\omega}$ by letting $\gamma_i(*)$ consist only of the i-long sequence of a's, for some fixed $a \in A$. Then, both $m(*) = \{a\}^*$ and $m'(*) = \{a\}^* \cup \{a\}^\omega$ define mediating maps. (A similar example is discussed in [1, Section 4.2].)

Our definitions of maximal trace and execution maps for non-deterministic systems will make use of the following result.

Lemma 3. *The non-empty powerset functor* $\mathcal{P}^+ : \mathsf{Set} \to \mathsf{Set}$ *weakly preserves limits of* ω^{op}*-chains. Moreover, the set of mediating maps for the image under* $\mathcal{P}^+$ *of a limiting cone over an* ω^{op}*-chain has a maximal element (under the point-wise inclusion order).*

PROOF. Let $(Z, (\pi_i : Z \to Z_i)_{i\in\omega})$ denote the limit of an ω^{op}-chain $(f_i : Z_{i+1} \to Z_i)_{i\in\omega}$, let $(\gamma_i : X \to \mathcal{P}^+ Z_i)_{i\in\omega}$ denote a cone over $(\mathcal{P}^+ f_i)_{i\in\omega}$, and assume $X \neq \emptyset$. (If $X = \emptyset$, the existence of a mediating map is trivial.) Now define $m : X \to \mathcal{P}^+ Z$ by

$$m(x) = \{z \in Z \mid \pi_i(z) \in \gamma_i(x) \text{ for all } i \in \omega\}$$

for $x \in X$. To show that $m(x) \neq \emptyset$, observe that by using the axiom of choice one can construct a sequence $(z_i)_{i\in\omega}$ with $z_i \in \gamma_i(x)$ and $f_i(z_{i+1}) = z_i$ for $i \in \omega$ – first choose $z_0 \in \gamma_0(x)$, then for $i \in \omega$ choose $z_{i+1} \in \gamma_{i+1}(x)$ satisfying $f_i(z_{i+1}) = z_i$, by using $(\mathcal{P}^+ f_i)(\gamma_{i+1}(x)) = \gamma_i(x)$. The limiting property of Z then yields $z \in Z$ with $\pi_i(z) = z_i \in \gamma_i(x)$ for $i \in \omega$, and thus $m(x) \neq \emptyset$. It then follows using a similar line of reasoning that m is a mediating map for the cone $(X, (\gamma_i)_{i\in\omega})$. Moreover, it is clear that m is above any other mediating map (under the point-wise inclusion order). This concludes the proof. □

Using Lemma 3, notions of maximal trace and maximal execution maps for $\mathcal{P}^+ \circ F$-coalgebras can be defined by replacing mediating maps with largest mediating maps in Definitions 3 and 4.

Definition 5. Let (X, γ) be a $\mathcal{P}^+ \circ F$-coalgebra, let $(\gamma_i)_{i\in\omega}$ be the induced cone over the ω^{op}-chain $(F^i ! : F^{i+1}1 \to F^i 1)_{i\in\omega}$, and let $(Z, (\pi_i)_{i\in\omega})$ denote a limiting cone for this ω^{op}-chain. The *trace map* $\mathsf{tr}_\gamma : X \to JZ$ of (X, γ) is given by the function:

$$\mathsf{tr}_\gamma(x) = \{z \in Z \mid \pi_i(z) \in \gamma_i(x) \text{ for all } i \in \omega\}$$

The *execution map* $\mathsf{exec}_\gamma : X \to JZ_X$ of (X, γ) is the trace map of the $\mathcal{P}^+ \circ F_X$-coalgebra $(X, \mathsf{st}_{X,FX} \circ \langle id_X, \gamma\rangle)$, with Z_X the carrier of a final F_X-coalgebra.

The next example describes the resulting maximal traces and executions, as well as the trace and execution maps, for some specific choices of F.

Example 3.

1. For unlabelled transition systems subject to the requirement that every state has at least one successor ($F = \mathsf{Id}$), the maximal traces are trivial (as the final F-coalgebra has a singleton as carrier), whereas the maximal executions are exactly the computation paths, as considered in the semantics of CTL*. The maximal execution map assigns to each state of a $\mathcal{P}^+$-coalgebra the computation paths from that state.

2. For labelled transition systems subject to a similar restriction ($F = A \times \mathsf{Id}$), the maximal execution map gives, for each state s, the set of labelled computation paths from s, as infinite sequences of the form $s = s_0 a_1 s_1 a_2 s_2 \ldots$ with $s_i \xrightarrow{a_i} s_{i+1}$ for $i \in \omega$, whereas the maximal trace map gives the sequences of labels that occur along such labelled computation paths.
3. One can also vary the functor F in order to model explicit termination. This is achieved by taking $F = 1 + \mathsf{Id}$ or $F = 1 + A \times \mathsf{Id}$ as in [1], and can be used to remove the requirement of at least one successor for each state. (Note that an arbitrary transition system can be regarded as a $\mathcal{P}^+ \circ (1 + \mathsf{Id})$-coalgebra, where the coalgebra map takes states with no successors to $\{\iota_1(*)\}$.) In these cases, the maximal trace (execution) maps incorporate both finite and infinite traces (respectively executions). To illustrate this, we briefly compare the infinite trace maps of $\mathcal{P} \circ (A \times \mathsf{Id})$-coalgebras, as defined in [11], with the maximal trace maps obtained by regarding such coalgebras as $\mathcal{P}^+ \circ (1 + A \times \mathsf{Id})$-coalgebras. Consider the labelled transition system with state space $\{x, y\}$ and a single transition $x \xrightarrow{a} y$. When regarding this as a $\mathcal{P} \circ (A \times \mathsf{Id})$-coalgebra (where the coalgebra map sends x to $\{(a, y)\}$ and y to $\emptyset$), the infinite trace map of [11] assigns an empty set of traces to x, as there are no *infinite* traces (i.e. elements of A^ω, the final coalgebra of $A \times \mathsf{Id}$) for x. On the other hand, when regarding the same transition system as the $\mathcal{P}^+ \circ (1 + A \times \mathsf{Id})$-coalgebra with carrier $\{x, y\}$ and coalgebra map given by $x \mapsto \{\iota_2(a, y)\}$ and $y \mapsto \{\iota_1(*)\}$, the maximal trace map defined here assigns the maximal trace a (element of $A^* \cup A^\omega$, the final coalgebra of $1 + A \times \mathsf{Id}$) to x.

We also note that Proposition 2 does not extend to the case when $T = \mathcal{P}^+$ – its proof makes use of the preservation by J of the limit of the final sequence of F. However, a weaker statement can be proved in this case.

Proposition 3. *For $T = \mathcal{P}^+$, the maximal trace map $\mathsf{tr}_\gamma : X \to JZ$ satisfies:*

$$\overline{F}\mathsf{tr}_\gamma \circ \gamma \subseteq J\zeta \circ \mathsf{tr}_\gamma.$$

Proof. Similarly to the proof of Proposition 2, but using the *weak* preservation of limits of ω^{op}-chains by $\mathcal{P}^+$ together with Lemma 1, it follows that

$(JFZ, (JF\pi_i)_{i\in\omega})$ is a weak limit of $(JF^{i+1}!)_{i\in\omega}$. (The same notation as in Proposition 2 is used here.) We now write ζ^{-1} for the inverse of the isomorphism $\zeta : Z \to FZ$, and show that $J\zeta \circ \mathsf{tr}_\gamma : X \to JFZ$ is the largest mediating map for the cone $(X, (\gamma_{i+1})_{i\in\omega})$ over the ω^{op}-chain $(JF^{i+1}!)_{i\in\omega}$:

$$\begin{aligned} (J\zeta \circ \mathsf{tr}_\gamma)(x) &= \\ \{\zeta(z) \mid z \in Z,\ \pi_i(z) \in \gamma_i(x) \text{ for } i \in \omega\} &= \\ \{\zeta(z) \mid z \in Z,\ \pi_{i+1}(z) \in \gamma_{i+1}(x) \text{ for } i \in \omega\} &= \\ \{\zeta(\zeta^{-1}(w)) \mid w \in FZ,\ \pi_{i+1}(\zeta^{-1}(w)) \in \gamma_{i+1}(x) \text{ for } i \in \omega\} &= \\ \{w \mid w \in FZ,\ (F\pi_i)(w) \in \gamma_{i+1}(x) \text{ for } i \in \omega\} & \end{aligned}$$

The conclusion then follows after observing that, as in the proof of Proposition 2, $\overline{F}\mathsf{tr}_\gamma \circ \gamma$ is a mediating map for $(X, (\gamma_{i+1})_{i\in\omega})$. □

Remark 3. The statement of Proposition 3 is weaker than the defining property of trace maps in [11], with the latter requiring an $\overline{F}$-coalgebra morphism. We are not aware of any instances of F and λ for which the trace map is *not* an $\overline{F}$-coalgebra morphism. We conjecture that an additional assumption on the endofunctor F (possibly involving a continuity condition) would be required to strengthen the above result, and leave the study of such a condition for future work.

We conclude this section by noting that our approach does not directly apply to the case $T = \mathcal{P}^+_\omega$, with $\mathcal{P}^+_\omega : \mathsf{Set} \to \mathsf{Set}$ the non-empty, *finite* powerset functor, as this functor does not weakly preserve limits of ω^{op}-chains. This is to be expected, since states of $\mathcal{P}^+_\omega$-coalgebras will, in general, have an *infinite* number of traces. Notions of maximal trace map and maximal execution map for finitely-branching transition systems are simply obtained by regarding these as transition systems with no cardinality restrictions on the branching.

3.4. Probabilistic systems

A large variety of discrete probabilistic models have been studied, see e.g. [19] for a coalgebraic account. Among these, probabilistic transition systems (also called *Markov chains* when restricting to countable state spaces) appear as coalgebras of the endofunctor $\mathcal{D} = \mathcal{D} \circ \mathsf{Id}$ and are used to interpret the logic PCTL [3], while *generative probabilistic systems* coincide with

$\mathcal{D} \circ (A \times \mathsf{Id})$-coalgebras. Here, $\mathcal{D} : \mathsf{Set} \to \mathsf{Set}$ denotes the *probability distribution monad*, a submonad of the subprobability distribution monad defined on objects by $\mathcal{D}X = \{\mu \in \mathcal{S}X \mid \sum_{x \in X} \mu(x) = 1\}$.

We begin by observing that, although affine, the monad $\mathcal{D}$ does not satisfy the requirement of Definition 3 concerning the preservation of limits by the induced functor J. To see this, let $F : \mathsf{Set} \to \mathsf{Set}$ be given by $FX = \{a, b\} \times X$, let $(Z, (\pi_i)_{i \in \omega})$ denote the limit of the ω^{op}-chain $(F^i!)_{i \in \omega}$, and let $\mu_i \in \mathcal{D}F^i 1$ be given by $\mu_i(x) = \frac{1}{2^i}$ for $x \in \{a, b\}^i$, with $i \in \omega$. Thus, each μ_i defines a *finite* probability distribution over $F^i 1$, and we have $(\mathcal{D}^i!)(\mu_{i+1}) = \mu_i$ for $i \in \omega$. However, there is no probability distribution μ on Z (note that $Z \simeq \{a, b\}^\omega$ is uncountable) such that $(\mathcal{D}\pi_i)(\mu) = \mu_i$ for $i \in \omega$ – any such μ could only take non-zero values on countably-many elements of Z. Indeed, a state of a $\mathcal{D} \circ F$-coalgebra will in general have uncountably many infinite traces, and the emphasis when defining a maximal trace map should be on measuring *sets* of traces rather than individual traces.

A satisfactory treatment of maximal traces for discrete probabilistic models turns out to be possible by regarding such models as coalgebras of the probability measure monad $\mathcal{G}_1$. Given a $\mathcal{D} \circ F$-coalgebra γ on Set, with $F : \mathsf{Set} \to \mathsf{Set}$ a shapely polynomial functor, our approach will be to lift F to a functor $\widehat{F} : \mathsf{Meas} \to \mathsf{Meas}$ and regard γ as a $\mathcal{G}_1 \circ \widehat{F}$-coalgebra to which Definitions 3 and 4 apply. In fact, we will show more generally that Definitions 3 and 4 yield maximal trace and execution maps for coalgebras of a certain class of endofunctors on the full subcategory SB of Meas.

To this end, we let $F : \mathsf{Meas} \to \mathsf{Meas}$ denote a shapely polynomial functor, and recall that Definitions 3 and 4 require the functor $J : \mathsf{Meas} \to \mathsf{Kl}(\mathcal{G}_1)$ to preserve the limits of the initial ω^{op}-segments of the final sequences of F and $F_{(X,\Sigma)}$ (with (X, Σ) the carrier of some $\mathcal{G}_1 \circ F$-coalgebra). By Lemma 1, for this it would suffice that the functor $\mathcal{G}_1 : \mathsf{Meas} \to \mathsf{Meas}$ preserves the same limits. Unfortunately, $\mathcal{G}_1 : \mathsf{Meas} \to \mathsf{Meas}$ does not preserve ω^{op}-limits (see [20, Section 3.3]), however, its restriction to the category of standard Borel spaces does (see [20, Corollary 3.1]). For this reason, our treatment of probabilistic systems will restrict attention to the subcategory SB of Meas[11]. We show next that, under some additional constraints on the shapely polynomial functor $F : \mathsf{Meas} \to \mathsf{Meas}$, F restricts to the category SB and preserves the initial

[11] Note that the monads $\mathcal{G} : \mathsf{Meas} \to \mathsf{Meas}$ and $\mathcal{G}_1 : \mathsf{Meas} \to \mathsf{Meas}$ restrict to monads on SB, which by abuse of notation we also denote $\mathcal{G}$ and $\mathcal{G}_1$, respectively.

ω^{op}-segment of its final sequence. Moreover, the same holds for the functor $F_{(X,\Sigma)}$, with (X, Σ) is a standard Borel space.

We recall from [16, Fact 1] that the category SB is closed under countable coproducts and countable limits. This ensures the correctness of the following definition.

Definition 6. A functor $F : \mathsf{SB} \to \mathsf{SB}$ is a *restricted shapely polynomial functor* if it is built from identity and constant functors using finite products and countable coproducts.

That is, restricted shapely polynomial functors are the shapely polynomial functors on SB whose definition only involves *countable* coproducts.

[16, Fact 1] also results in SB being closed under limits of ω^{op}-chains. The next two lemmas ensure that the previously-mentioned hypotheses of Definitions 3 and 4 are satisfied by the functors J, F and $F_{(X,\Sigma)}$.

Lemma 4 ([20, Corollary 3.1]). *The functor $\mathcal{G}_1 : \mathsf{SB} \to \mathsf{SB}$ preserves limits of ω^{op}-chains.*

Hence, by Lemma 1, $J : \mathsf{SB} \to \mathsf{Kl}(\mathcal{G}_1)$ also preserves limits of ω^{op}-chains.

Lemma 5. *Restricted shapely polynomial functors preserve limits of ω^{op}-chains in SB.*

Proof. The statement follows by induction on the structure of restricted shapely polynomial functors. For constant and identity functors, the claim is immediate. Now assume that $F_i : \mathsf{SB} \to \mathsf{SB}$ preserves the limit of an ω^{op}-chain in SB, for $i \in \omega$. Preservation of the same limit by $F_1 \times F_2$ is straightforward (as limits commute with limits in any category), while its preservation by $\coprod_{i\in\omega} F_i$ is a consequence of limits of ω^{op}-chains commuting with coproducts in Set, and of the fact that all bijective SB-morphisms are isomorphisms (see [16, Fact 2 and proof of Proposition 3] for more details). □

Remark 4. As a consequence of the above, for every restricted shapely polynomial functor $F : \mathsf{SB} \to \mathsf{SB}$, the limit of the initial ω^{op}-segment of F is the carrier of a final F-coalgebra. Moreover, this also applies to the functor $F_{(X,\Sigma)} : \mathsf{SB} \to \mathsf{SB}$ defined by $F_{(X,\Sigma)}(Y, \Sigma') = (X, \Sigma) \times F(Y, \Sigma')$.

We also recall from Section 2 that commutative monads on any category with products and coproducts admit canonical distributive laws over shapely polynomial functors. This applies in particular to the monad $\mathcal{G}_1 : \mathsf{SB} \to \mathsf{SB}$ and any restricted shapely polynomial functor. With this, we can conclude that all the requirements of Definitions 3 and 4 are satisfied by the monad $\mathcal{G}_1 : \mathsf{SB} \to \mathsf{SB}$ and the restricted shapely polynomial functors $F : \mathsf{SB} \to \mathsf{SB}$ and $F_{(X,\Sigma)} : \mathsf{SB} \to \mathsf{SB}$, respectively. This yields notions of maximal trace map and maximal execution map for $\mathcal{G}_1 \circ F$-coalgebras over SB. At the same time, we note that these definitions can be applied to *any* endofunctor $F : \mathsf{SB} \to \mathsf{SB}$ which preserves the initial ω^{op}-segment of its final sequence. Our focusing on restricted shapely polynomial endofunctors was driven by the need to consider $\mathcal{D} \circ F$-coalgebras over Set, with $F : \mathsf{Set} \to \mathsf{Set}$ a shapely polynomial functor.

We now return to such $\mathcal{D} \circ F$-coalgebras. In order to lift F to a functor $\widehat{F} : \mathsf{SB} \to \mathsf{SB}$, some additional constraints on the shape of F are required.

Definition 7. A shapely polynomial functor $F : \mathsf{Set} \to \mathsf{Set}$ is a *restricted shapely polynomial functor* if it is built from identity and *countable* constant functors using finite products and *countable* coproducts.

Definition 8. Given a restricted shapely polynomial functor $F : \mathsf{Set} \to \mathsf{Set}$, its *lifting* $\widehat{F} : \mathsf{SB} \to \mathsf{SB}$ *to standard Borel spaces* is defined by structural induction on F:

- $\widehat{\mathsf{Id}}$ is the identity functor on SB,
- $\widehat{C_X}$ is the constant functor $C_{(X,\mathcal{P}X)}$, for each countable set X,
- $\widehat{F_1 \times F_2} = \widehat{F_1} \times \widehat{F_2}$,
- $\widehat{\coprod_{i\in\omega} F_i} = \coprod_{i\in\omega} \widehat{F_i}$.

The correctness of the above definition is guaranteed by the observations that a discrete measurable space $(X, \mathcal{P}X)$ is standard Borel if and only if X is countable, and that SB is closed under countable products and coproducts in Meas (see [16, Fact 1]). It then follows immediately that $\widehat{F} : \mathsf{SB} \to \mathsf{SB}$ is a restricted shapely polynomial functor. Moreover, the following hold:

Lemma 6. 1. *The sets underlying the measurable spaces* $\widehat{F_1 \times F_2}(X, \Sigma_X)$ *and* $\widehat{\coprod_{i\in\omega} F_i}(X, \Sigma_X)$ *are given by* $(F_1 \times F_2)X$ *and* $\coprod_{i\in\omega} F_i(X)$, *respectively.*

2. *The functor* $\widehat{F} : \mathsf{SB} \to \mathsf{SB}$ *preserves discrete spaces.*

Proof. We begin by recalling that (finite) products and (countable) coproducts in Meas are constructed by putting a suitable σ-algebra structure on the product, respectively coproduct of the underlying sets. Specifically, the σ-algebra on the product is generated by the cartesian products of measurable sets, whereas the σ-algebra on the coproduct is generated by the disjoint unions of measurable sets in each of the summands [20, Section 3.1]. The first statement now follows immediately, whereas the second statement follows by induction on the structure of F. □

We now show how to view a $\mathcal{D} \circ F$-coalgebra with countable carrier as a $\mathcal{G}_1 \circ \widehat{F}$-coalgebra. The restriction to countable carriers is required to stay within SB.

Proposition 4. *Let* $F : \mathsf{Set} \to \mathsf{Set}$ *be a restricted shapely polynomial functor, let* (X, γ) *be a* $\mathcal{D} \circ F$*-coalgebra with countable carrier, and let* $((X, \mathcal{P}X), \widehat{\gamma})$ *be the* $\mathcal{G}_1 \circ \widehat{F}$*-coalgebra whose structure map* $\widehat{\gamma}$ *takes* $x \in X$ *to the unique probability measure on* $\widehat{F}(X, \mathcal{P}X)$ *induced by the probability distribution* $\gamma(x)$ *on* FX [12]*. Then, the cones*[13] $(\gamma_i : X \to JF^i 1)_{i \in \omega}$ *over* $(JF^i!)_{i \in \omega}$ *induced by* γ*:*

$$\begin{array}{ccccccc} X & & & & & & \\ \downarrow{\scriptstyle \gamma_0} & \searrow{\scriptstyle \gamma_1} & & \searrow{\scriptstyle \gamma_2} & & & \\ J1 & \xleftarrow[J!]{} & JF1 & \xleftarrow[JF!]{} & JF^2 1 & \xleftarrow[JF^2!]{} & \cdots \end{array}$$

and $(\widehat{\gamma}_i : (X, \mathcal{P}X) \to J'\widehat{F}^i(1, \mathcal{P}1))_{i \in \omega}$ *over* $(J'\widehat{F}^i!)_{i \in \omega}$ *induced by* $((X, \mathcal{P}X), \widehat{\gamma})$*:*

$$\begin{array}{ccccccc} (X, \mathcal{P}X) & & & & & & \\ \downarrow{\scriptstyle \widehat{\gamma}_0} & \searrow{\scriptstyle \widehat{\gamma}_1} & & \searrow{\scriptstyle \widehat{\gamma}_2} & & & \\ J'(1, \mathcal{P}1) & \xleftarrow[J'!]{} & J'\widehat{F}(1, \mathcal{P}1) & \xleftarrow[J'\widehat{F}!]{} & J'\widehat{F}^2(1, \mathcal{P}1) & \xleftarrow[J'\widehat{F}^2!]{} & \cdots \end{array}$$

---

[12] Note that, by Lemma 6.2, we have $\widehat{F}(X, \mathcal{P}X) = (FX, \mathcal{P}FX)$.

[13] Recall that the monads $\mathcal{D}$ and $\mathcal{G}_1$ are affine, and hence Lemma 2 applies.

are such that $\widehat{\gamma}_i : (X, \mathcal{P}X) \to J'\widehat{F}^i(1, \mathcal{P}1)$ is the point-wise extension of $\gamma_i : X \to JF^i1$ to a probability measure, for $i \in \omega$. (Here, the functors $J : \mathsf{Set} \to \mathsf{Kl}(\mathcal{D})$ and $J' : \mathsf{SB} \to \mathsf{Kl}(\mathcal{G}_1)$ are as in Section 2, and the cones $(\gamma_i)_{i\in\omega}$ and $(\widehat{\gamma}_i)_{i\in\omega}$ are constructed as in Section 3.1.)

Proof. We first note that the measurability of $\widehat{\gamma}$ is an immediate consequence of $(X, \mathcal{P}X)$ being discrete. An easy induction proof then shows that, for $i \in \omega$ and $x \in X$, $\widehat{\gamma}_i(x)$ is the unique probability measure on $\widehat{F}^i(1, \mathcal{P}1)$ induced by the probability distribution $\gamma_i(x)$ on F^i1. □

As $(1, \mathcal{P}1)$ is final in Meas, the cone $\widehat{\gamma}_i$ is over the image under J' of the initial ω^{op}-segment of the final sequence of $\widehat{F}$. As a result, we can use the existence of trace maps of $\mathcal{G}_1 \circ \widehat{F}$-coalgebras to define trace maps for $\mathcal{D} \circ F$-coalgebras. Before doing so, we observe that the underlying functions defining the canonical distributive law of $\mathcal{G}_1$ over $\widehat{F}$ agree with the functions defining the canonical distributive law of $\mathcal{D}$ over F.

Proposition 5. *Let $U : \mathsf{SB} \to \mathsf{Set}$ denote the functor taking a standard Borel space to its underlying set, let $F : \mathsf{Set} \to \mathsf{Set}$ denote a restricted shapely polynomial functor, and let $\lambda : F\mathcal{D} \Rightarrow \mathcal{D}F$ and $\widehat{\lambda} : \widehat{F}\mathcal{G}_1 \Rightarrow \mathcal{G}_1\widehat{F}$ denote the canonical natural transformations of $\mathcal{D}$ over F and of $\mathcal{G}_1$ over $\widehat{F}$, respectively. Then, the following diagram commutes:*

$$\begin{array}{ccc}
F\mathcal{D}U & \xRightarrow{\lambda_U} & \mathcal{D}FU \\
{\scriptstyle F\iota}\Downarrow & & \Vert \\
FU\mathcal{G}_1 & & \mathcal{D}U\widehat{F} \\
\Vert & & \Downarrow{\scriptstyle \iota_{\widehat{F}}} \\
U\widehat{F}\mathcal{G}_1 & \xRightarrow[U\widehat{\lambda}]{} & U\mathcal{G}_1\widehat{F}
\end{array}$$

where the (X, Σ_X)-component of the natural transformation $\iota : \mathcal{D}U \Rightarrow U\mathcal{G}_1$ takes a probability distribution $\mu \in \mathcal{D}X$ to the unique probability measure on (X, Σ_X) induced by μ.

Proof. The statement follows by induction on the structure of F, using Lemma 6.1. □

We are now in a position to define probabilistic trace and execution maps for $\mathcal{D} \circ F$-coalgebras (X, γ) with countable carriers. To this end, we write (Z, Σ_Z) for the carrier of a final $\widehat{F}$-coalgebra, and (Z_X, Σ_{Z_X}) for the carrier of a final $\widehat{F}_{(X,\mathcal{P}X)}$-coalgebra. We recall that ω^{op}-limits in Meas, and hence (as SB is closed under countable limits) also in SB, are constructed from the limits of the underlying diagrams in Set (see e.g. [20, Section 3.3] for details). As a result, the state set of the coalgebra (Z, Σ_Z) is the carrier of a final F-coalgebra, whereas the σ-algebra Σ_Z is generated by the inverse images of measurable sets in $\widehat{F}^i(1, \mathcal{P}1)$ (i.e. subsets of $F^i 1$) under the maps $\pi_i : Z \to F^i 1$, for $i \in \omega$. In particular, the inverse images of singletons $\{f_i\} \subseteq F^i 1$ yield measurable subsets of Z; that is, the set of maximal traces that have the same finite prefix f_i is measurable. Similarly, Z_X is the carrier of a final F_X-coalgebra, and the set of maximal executions with the same finite prefix $e_i \in (F_X)^i 1$ is measurable.

Definition 9. Let $F : \mathsf{Set} \to \mathsf{Set}$ be a restricted shapely polynomial functor, and let (X, γ) be a $\mathcal{D} \circ F$-coalgebra with countable carrier. The *probabilistic trace map* of (X, γ) is the underlying function of the maximal trace map $\mathsf{tr}_{\widehat{\gamma}} : (X, \mathcal{P}X) \to J'(Z, \Sigma_Z)$ of the $\mathcal{G}_1 \circ \widehat{F}$-coalgebra $((X, \mathcal{P}X), \widehat{\gamma})$ of Proposition 4. The *probabilistic execution map* of (X, γ) is the probabilistic trace map of the $\mathcal{G}_1 \circ \widehat{F}_{(X,\mathcal{P}X)}$-coalgebra $((X, \mathcal{P}X), \mathsf{st}_{(X,\mathcal{P}X),\widehat{F}(X,\mathcal{P}X)} \circ \langle id_{(X,\mathcal{P}X)}, \widehat{\gamma} \rangle)$.

The above definition assumes that the canonical distributive law of $\mathcal{G}_1$ over $\widehat{F}$ is considered when defining maximal trace maps in the category SB.

As expected, the probabilistic trace map yields, for each state of a $\mathcal{D} \circ F$-coalgebra, a probability measure over (Z, Σ_Z), while the probabilistic execution map of (X, γ) yields, for each state, a probability measure over (Z_X, Σ_{Z_X}).

Example 4. In the case of Markov chains ($F = \mathsf{Id}$), the probabilistic execution map gives, for each state in a Markov chain, a probability measure over its computation paths. In particular, for each finite prefix $x_0 \ldots x_i \in {F_X}^i 1$, such a probability measure assigns a probability value to the set of computation paths that extend $x_0 \ldots x_i$. Similarly, in the case of generative probabilistic systems ($F = A \times \mathsf{Id}$), the probabilistic execution map gives, for each state, a probability measure over its labelled computation paths. As in the case of nondeterministic systems, explicit termination can be modelled by taking $F = 1 + \mathsf{Id}$ or $F = 1 + A \times \mathsf{Id}$, with the probabilistic execution maps now also incorporating finite (labelled) computation paths.

4. Path-Based Coalgebraic Temporal Logics

We now introduce coalgebraic temporal logics in the style of CTL*, whose semantics is defined in terms of execution maps. Throughout this section, we fix a monad $T : \mathsf{C} \to \mathsf{C}$, an endofunctor $F : \mathsf{C} \to \mathsf{C}$ and a $T \circ F$-coalgebra (X, γ). We let $\mathsf{exec}_\gamma : X \to TZ_X$ denote the maximal execution map given by Definition 4, with (Z_X, ζ_X) a final F_X-coalgebra.

At the same time, we note that the temporal languages defined in this section can also be interpreted by using the finite execution map $\mathsf{fexec}_\gamma : X \to TI_X$ with (I_X, ι_X) an initial F_X-algebra, as given by Definition 2, instead of the maximal execution map – the forthcoming definitions do not rely on the finality of (Z_X, ζ_X). However, this is only useful when $F0 \neq 0$, with 0 an initial object in C, as otherwise the initial F_X-algebra is trivial. In particular, modelling explicit termination via functors such as $F = 1 + \mathsf{Id}$ or $F = 1 + A \times \mathsf{Id}$ yields non-trivial finite execution maps to which the definitions in this section can be applied.

The temporal logics that we define are parameterised by sets Λ_F and Λ of monotone[14] predicate liftings for the functors F and respectively T. The category C will be instantiated to Set as well as to the full subcategory SB of Meas.

We recall that the definition of predicate liftings requires functors $U : \mathsf{C} \to \mathsf{Set}$ and $P : \mathsf{C} \to \mathsf{Set}^{\mathrm{op}}$ such that P is a subfunctor of $\hat{\mathcal{P}} \circ U$. In addition, defining the semantics of path-based fixpoint logics will require that, for each C-object X, both $(PX, \subseteq)$ and $(PX, \supseteq)$ are directed complete partial orders. This will allow us to make use of the following result.

Lemma 7 ([21, Theorem 8.22]). *Let P be a directed complete partial order and let $O : P \to P$ be order-preserving. Then, O has a least fixpoint.*

Since an ordered set is a directed complete partial order if and only if each chain has a least upper bound (see e.g. [21, Theorem 8.11]), the hypothesis of the previous result is satisfied by $(PX, \subseteq)$ as well as by $(PX, \supseteq)$, both for $P = \hat{\mathcal{P}} : \mathsf{Set} \to \mathsf{Set}$ and for $P : \mathsf{Meas} \to \mathsf{Set}$ taking a measurable space to its σ-algebra.

[14] The restriction to monotone predicate liftings is only required to define the path-based fixpoint logics of Section 4.1, and not also the path-based temporal logics with Until operators of Section 4.2. For the latter, no appeal to fixpoint existence theorems is needed.

4.1. Path-based fixpoint logics

We now proceed to define path-based coalgebraic fixpoint logics. Like CTL*, these logics are two-sorted, with *path formulas* denoted by $\varphi, \psi, \ldots$ expressing properties of executions, and *state formulas* denoted by $\Phi, \Psi, \ldots$ expressing properties of states of $T \circ F$-coalgebras.

To motivate the syntax of these logics, we recall that:

- the execution map exec_γ provides, for each state $x \in UX$ of a $T \circ F$-coalgebra (X, γ), an element of UTZ_X, that is, a T-structured observation on the possible executions,
- the coalgebra structure $\zeta_X : Z_X \to X \times FZ_X$ provides, for each execution $z \in UZ_X$, its first state, $U(\pi_1 \circ \zeta_X)(z) \in UX$, as well as an F-structured observation $U(\pi_2 \circ \zeta_X)(z) \in UFZ_X$.

Thus, it seems natural to use:

- one-step modal operators inspecting the T-structured observations on the possible executions (provided by the map exec_γ), to define state formulas,
- one-step modal operators inspecting the F-structure of executions (defined by $\pi_2 \circ \zeta_X$), to define path formulas.

At the same time, the C-map $U(\pi_1 \circ \zeta_X)$ allows a property of a state to be regarded as a property of (the first state of) an execution. These observations justify the following definition of a 2-sorted, path-based temporal language.

Definition 10. The language $\mu\mathcal{L} ::= \mu\mathcal{L}_\Lambda^{\Lambda_F}(\mathcal{U}, \mathcal{V})$ over a 2-sorted set $(\mathcal{U}, \mathcal{V})$ of propositional variables (with sorts for paths and respectively states) is defined by the grammar

$$\begin{aligned} \mu\mathcal{L}_F \ni \varphi \quad &::= \quad \mathsf{tt} \mid \mathsf{ff} \mid q \mid \Phi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid [\lambda_F]\varphi \mid \eta q.\varphi \\ \mu\mathcal{L} \ni \Phi \quad &::= \quad \mathsf{tt} \mid \mathsf{ff} \mid p \mid [\lambda]\varphi \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \end{aligned}$$

where $q \in \mathcal{U}$, $p \in \mathcal{V}$, $\eta \in \{\mu, \nu\}$, $\lambda_F \in \Lambda_F$ and $\lambda \in \Lambda$.

Thus, path formulas are constructed from propositional variables $q \in U$ and state formulas Φ using positive boolean operators, modal operators $[\lambda_F]$ and fixpoint operators, whereas state formulas are constructed from atomic propositions p and modal formulas $[\lambda]\varphi$ with φ a *path* formula, using positive boolean operators. The modal operators $[\lambda_F]$ and $[\lambda]$ with $\lambda_F \in \Lambda_F$ and $\lambda \in \Lambda$ are thus both applied to path formulas, to obtain new path formulas and respectively *state* formulas. They are, however, of very different natures: while the operators $[\lambda_F]$ quantify over the *one-step* behaviour of executions (recall that executions carry F_X-coalgebra structure, and hence F-coalgebra structure), the operators $[\lambda]$ quantify over the suitably-structured, *long-term* executions from particular states. This is made precise in the formal semantics of $\mu\mathcal{L}_\Lambda^{\Lambda_F}(\mathcal{U}, \mathcal{V})$, as defined below.

Definition 11. Given a $T \circ F$-coalgebra (X, γ) and a 2-sorted valuation $V : (\mathcal{U}, \mathcal{V}) \to (PZ_X, PX)$ (interpreting path and state variables as sets of executions and respectively of states), the semantics $(\!|\varphi|\!)_{\gamma,V} \in PZ_X$ of path formulas $\varphi \in \mu\mathcal{L}_F$ and $[\![\Phi]\!]_{\gamma,V} \in PX$ of state formulas $\Phi \in \mu\mathcal{L}$ is defined inductively on the structure of φ and Φ by:

$$\begin{aligned}
(\!|q|\!)_{\gamma,V} &= V(q)\\
(\!|\Phi|\!)_{\gamma,V} &= P(\pi_1 \circ \zeta_X)([\![\Phi]\!]_{\gamma,V})\\
(\!|[\lambda_F]\varphi|\!)_{\gamma,V} &= (P(\pi_2 \circ \zeta_X) \circ (\lambda_F)_{Z_X})((\!|\varphi|\!)_{\gamma,V})\\
(\!|\mu q.\varphi|\!)_{\gamma,V} &= \mathsf{lfp}((\varphi)_q^{\gamma,V})\\
(\!|\nu q.\varphi|\!)_{\gamma,V} &= \mathsf{gfp}((\varphi)_q^{\gamma,V})\\
[\![p]\!]_{\gamma,V} &= V(p)\\
[\![[\lambda]\varphi]\!]_{\gamma,V} &= (P\mathsf{exec}_\gamma \circ \lambda_{Z_X})((\!|\varphi|\!)_{\gamma,V})
\end{aligned}$$

and the usual clauses for the boolean operators, where, for $q \in \mathcal{U}$, $(\varphi)_q^{\gamma,V} : PX \to PX$ denotes the monotone map defined by $(\varphi)_q^{\gamma,V}(Y) = (\!|\varphi|\!)_{\gamma,V'}$ with $V'(p) = V(p)$ for $p \in \mathcal{V}$, $V'(q) = Y$ and $V'(r) = V(r)$ for $r \in \mathcal{U}$, $r \neq q$, whereas $\mathsf{lfp}(\_)$ and $\mathsf{gfp}(\_)$ construct least and respectively greatest fixpoints.

We note that the monotonicity of the predicate liftings in Λ_F and Λ together with the absence of negation in either path or state formulas ensure that the maps $(\varphi)_q^{\gamma,V} : PX \to PX$ are monotone, and hence, by Theorem 7, admit least and greatest fixpoints.

Let us now examine the definition of the semantics of $\mu\mathcal{L}_\Lambda^{\Lambda_F}(\mathcal{U}, \mathcal{V})$ in more detail:

- To define $(\!|\Phi|\!)_{\gamma,V} \in PZ_X$ from $[\![\Phi]\!]_{\gamma,V} \in PX$, one uses the image under P of the map $\pi_1 \circ \zeta_X$

$$Z_X \xrightarrow{\zeta_X} X \times FZ_X \xrightarrow{\pi_1} X$$

(which extracts the first state of an execution) to go from a set of states (those satisfying Φ) to a set of executions. This formalises the idea that a state formula Φ (regarded as a path formula) holds in a path precisely when it holds in the first state of that path.

- To define $(\![\lambda_F]\varphi]\!)_{\gamma,V} \in PZ_X$ from $(\!|\varphi|\!)_{\gamma,V} \in PZ_X$, one first applies the relevant component of the predicate lifting λ_F to obtain a set of F-structured observations on executions (as an element of PFZ_X), and then uses the image under P of the map $\pi_2 \circ \zeta_X$

$$Z_X \xrightarrow{\zeta_X} X \times FZ_X \xrightarrow{\pi_2} FZ_X$$

(which extracts the one-step F-observation of an execution) to obtain a set of executions again. This is the standard interpretation of the modal formula $[\lambda_F]\varphi$ in the F-coalgebra $\pi_2 \circ \zeta_X$.

- Finally, to define $[\![[\lambda]\varphi]\!]_{\gamma,V} \in PX$ from $(\!|\varphi|\!)_{\gamma,V} \in PZ_X$, one first applies the relevant component of the predicate lifting λ to $(\!|\varphi|\!)_{\gamma,V} \in PZ_X$ to obtain a set of suitably-structured executions (i.e. an element of PTZ_X), and then uses the image under P of the execution map to obtain a set of states:

$$PZ_X \xrightarrow{(\lambda)_{Z_X}} PTZ_X \xrightarrow{P\mathsf{exec}_\gamma} PX$$

Example 5. 1. The negation-free fragment of the logic CTL* can be recovered as a fragment of the path-based fixpoint logic obtained by taking $P = \hat{\mathcal{P}}$, $T = \mathcal{P}^+$, $F = \mathsf{Id}$, $\Lambda = \{\lambda_\Box, \lambda_\Diamond\}$ and $\Lambda_F = \{\lambda_\bigcirc\}$, with the predicate liftings $\lambda_\Box, \lambda_\Diamond : \hat{\mathcal{P}} \Rightarrow \hat{\mathcal{P}} \circ \mathcal{P}^+$ and $\lambda_\bigcirc : \hat{\mathcal{P}} \Rightarrow \hat{\mathcal{P}} \circ \mathsf{Id}$ being given by:

$$\begin{aligned}
(\lambda_\Box)_X(Y) &= \{Z \in \mathcal{P}^+X \mid Z \subseteq Y\}, \\
(\lambda_\Diamond)_X(Y) &= \{Z \in \mathcal{P}^+X \mid Z \cap Y \neq \emptyset\}, \\
(\lambda_\bigcirc)_X(Y) &= Y.
\end{aligned}$$

The choice of $\lambda_\Box$ and $\lambda_\Diamond$ as predicate liftings for $\mathcal{P}^+$ captures precisely the semantics of the path quantifiers **A** and **E** of CTL*, whereas $\lambda_\bigcirc$ captures the semantics of the temporal operator **X**. The Until operator of CTL* can then be encoded as a fixpoint path formula:

$$\varphi\,\mathbf{U}\,\psi \quad ::= \quad \mu q.(\psi \vee (\varphi \wedge \bigcirc q))$$

where we write simply $\bigcirc$ for the modal operator $[\lambda_\bigcirc]$.

2. By varying the functor F to $A \times \mathsf{Id}$, we obtain an interesting variant of CTL* interpreted over labelled transition systems. For this, we take $\Lambda_F = \{\lambda_a \mid a \in A\} \cup \{\lambda_{\bigcirc}\}$, with the predicate liftings $\lambda_a : 1 \Rightarrow \hat{\mathcal{P}} \circ (A \times \mathsf{Id})$ for $a \in A$ and $\lambda_{\bigcirc} : \hat{\mathcal{P}} \Rightarrow \hat{\mathcal{P}} \circ (A \times \mathsf{Id})$ being given by:

$$\begin{aligned}(\lambda_a)_X(*) &= \{a\} \times X,\\ (\lambda_{\bigcirc})_X(Y) &= A \times Y.\end{aligned}$$

   We write a for the nullary modality $[\lambda_a]$, and (as before) $\bigcirc$ for the unary modality $[\lambda_{\bigcirc}]$. Then, the path formula a requires the first label of a labelled computation path to be precisely a, whereas the formula $\bigcirc\varphi$ is true on a computation path $s_0a_1s_1a_2s_2\ldots$ precisely when φ is true on $s_1a_2s_2\ldots$. The property "a occurs along every computation path" can be expressed in the resulting temporal language as $\Box\, \mu X.(a \vee \bigcirc X)$ (with $\Box$ a shorthand for $[\lambda_{\Box}]$). The reader should compare this to the formulation of the same property in the language obtained by adding fixpoints to the negation-free variant of Hennessy-Milner logic, namely as $\mu X.(\langle\_\rangle\mathsf{tt} \wedge [-a]X)$. Here, the formulas $\langle\_\rangle\Phi$ and $[-a]\Phi$ should be read as "there exists a successor state (reachable by *some* label) satisfying Φ" and respectively "all states reachable by labels other than a satisfy Φ". It is easy to see that, as the required nesting depth of fixpoint operators increases, the encodings of path properties in the latter language quickly become complex, making the path-based language the preferred choice as a specification language.
3. By further varying the functor F to $F = 1 + \mathsf{Id}$ or $F = 1 + A \times \mathsf{Id}$, the resulting maximal execution maps incorporate both finite and infinite computation paths, while the finite execution maps provided by Definition 2 only account for the finite computation paths. Both maps can be used as the semantic basis for path-based languages similar to the two languages discussed above. The new languages can also contain a nullary path operator $\bot$, with the formula $\bot$ only being true on a finite path containing a single state.

Example 6. The negation-free fragment of the logic PCTL [3] can be recovered as a fragment of the path-based temporal logic obtained by taking $T = \mathcal{G}_1$ and $F = \mathsf{Id}$ on SB, and the functor $P : \mathsf{SB} \to \mathsf{Set}$ to be given by $P(X, \Sigma_X) = \Sigma_X$. The identity natural transformation $\lambda_{\bigcirc} = id_P : P \Rightarrow P$ then defines a predicate lifting for $F = \mathsf{Id}$. Also, for $q \in [0, 1]$, the natural transformations $\lambda_{\geq q}, \lambda_{>q} : P \Rightarrow P \circ \mathcal{G}_1$ given by

$$(\lambda_{\geq q})_{(X,\Sigma_X)}(Y) \;=\; \{\mu \in \mathcal{M}_1(X, \Sigma_X) \mid \mu(Y) \geq q\}$$
$$(\lambda_{>q})_{(X,\Sigma_X)}(Y) \;=\; \{\mu \in \mathcal{M}_1(X, \Sigma_X) \mid \mu(Y) > q\}$$

for $Y \in \Sigma_X$ define predicate liftings for $T = \mathcal{G}_1$. Now letting $\Lambda_F = \{\lambda_{\bigcirc}\}$ and $\Lambda = \{\lambda_{\geq q} \mid q \in [0,1]\} \cup \{\lambda_{>q} \mid q \in [0,1]\}$ yields a path-based temporal logic interpreted over standard Borel spaces. A fragment of this logic corresponds to the logic PCTL, again interpreted over standard Borel spaces: $\lambda_{\bigcirc}$, $\lambda_{\geq q}$ and $\lambda_{>q}$ capture the semantics of the PCTL operators $\mathbf{X}$, $[\_]_{\geq q}$ and respectively $[\_]_{>q}$, whereas the Until operator of PCTL can be encoded as:

$$\Phi \mathbf{U} \Psi \;\;::=\;\; \mu q.(\Psi \vee (\Phi \wedge \bigcirc q)$$

where, as before, $\bigcirc$ is shorthand for $[\lambda_{\bigcirc}]$.The interpretation of the resulting logic over Markov chains with countable state spaces is then obtained by regarding each such Markov chain as a discrete measurable space (which is also standard Borel). Moreover, by varying the transition type to $F = A \times \mathsf{Id}$ or $F = 1 + A \times \mathsf{Id}$, one automatically obtains variants of PCTL interpreted over generative probabilistic systems, possibly with explicit termination.

4.2. Path-based temporal logics with Until operators

In order to recover the full languages CTL* and PCTL as instances of general path-based logics, one needs to incorporate negation into the syntax of both path and state formulas. As a result, arbitrary fixpoints must be left out, as the operators previously used to interpret them may fail to be monotone. In what follows, we replace fixpoint formulas by Until operators similar to the ones of CTL* and PCTL. However, a similar approach can be used to define more general temporal operators.

Before defining the general syntax of path-based temporal logics with Until operators, we observe that the structure of the functor F may result in the associated notions of trace and execution involving some branching (as is for instance the case when $FX = A \times X \times X$). In such cases, Until operators can incorporate either a universal or an existential quantification over the corresponding branches. Only existential versions of branching Until operators are considered in what follows, and the reader is referred to [7] for a definition of their universal counterparts.

Path-based temporal logics with Until operators are obtained by discarding propositional variables $q \in \mathcal{U}$ from the path formulas of $\mu\mathcal{L}_F$, and replacing fixpoint formulas $\mu q.\varphi$ and $\nu q.\varphi$ by formulas $\varphi U_L \psi$, with $L \subseteq \Lambda_F$

a subset of predicate liftings for the functor F. Furthermore, negation is added to the syntax of both path and state formulas, and the restriction to monotone predicate liftings in Λ and Λ_F is dropped, as no appeal to fixpoint existence theorems is required to interpret Until operators.

Definition 12. The language $\mathcal{LU} ::= \mathcal{LU}_{\Lambda}^{\Lambda_F}(\mathcal{V})$ over a set $\mathcal{V}$ of propositional variables is defined by the grammar

$$\begin{array}{rcl} \mathcal{LU}_F \ni \varphi & ::= & \mathsf{tt} \mid \Phi \mid \neg\varphi \mid \varphi \wedge \varphi \mid [\lambda_F]\varphi \mid \varphi U_L \varphi \\ \mathcal{LU} \ni \Phi & ::= & \mathsf{tt} \mid p \mid [\lambda]\varphi \mid \neg\Phi \mid \Phi \wedge \Phi \end{array}$$

where $p \in \mathcal{V}$, $\lambda_F \in \Lambda_F$, $L \subseteq \Lambda_F$ and $\lambda \in \Lambda$.

Fixing $L \subseteq \Lambda_F$ corresponds to fixing a number of ways of inspecting the structure of executions using one-step unfoldings. Often, L just consists of a single modal operator, however, depending on the structure of the functor F, one may choose to consider non-singleton sets L of modal operators. For example, if $F = A \times \mathsf{Id} \times \mathsf{Id} : \mathsf{Set} \to \mathsf{Set}$ comes with predicate liftings $\lambda_1, \lambda_2 : \hat{\mathcal{P}} \Rightarrow \hat{\mathcal{P}} \circ F$ defined by:

$$(\lambda_i)_X(Y) \;=\; \{Z \subseteq X \times X \mid \pi_i(Z) \subseteq Y\}$$

for $i \in \{1, 2\}$, then one may choose to take $L = \{\lambda_1, \lambda_2\}$. In this case, F-coalgebras are infinite, A-labelled binary trees, and the intended meaning of an *existential Until* formula $\varphi U_L \psi$ is that φ must hold along *some* branch of the tree, starting from the root, until ψ is found to hold. In contrast, a *universal Until* formula would require this for *every* branch of the tree. Also, the existential/universal Until formula $\varphi U_{\{\lambda_1\}} \psi$ would require φ to hold along the left-most branch of the tree (as $[\lambda_1]$ inspects the first component of $X \times X$), until a state satisfying ψ is reached. More generally, an existential Until formula $\varphi U_L \psi$ should be read as "there exists a route described by the modalities in L along which φ holds until ψ holds".

The general semantics of existential Until operators is given by

$$(\!|\varphi U_L \psi|\!)_{\gamma,V} = \bigcup_{t \in \omega} (\!|\varphi U_L^{\leq t} \psi|\!)_{\gamma,V}$$

where the formulas $\varphi U_L^{\leq t} \psi$ with $t \in \omega$ are defined inductively by:

$$\begin{array}{rcl} \varphi U_L^{\leq 0} \psi & ::= & \psi \\ \varphi U_L^{\leq t+1} \psi & ::= & \psi \vee (\varphi \wedge \bigvee_{\lambda_F \in L} [\lambda_F](\varphi U_L^{\leq t} \psi)) \end{array}$$

The semantics of state formulas remains as before.

Example 7. 1. The logic CTL* can be recovered as a fragment of the path-based logic with Until operators by taking $T = \mathcal{P}^+$, $F = \mathsf{Id}$, $\Lambda = \{\lambda_\Box\}$, $\Lambda_F = \{\lambda_\bigcirc\}$ and $L = \{\lambda_\bigcirc\}$, with $\lambda_\Box$ and $\lambda_\bigcirc$ as in Example 5.1.

2. Similarly, the logic PCTL interpreted over standard Borel spaces is obtained as a fragment of the path-based logic with Until operators resulting from $T = \mathcal{G}_1$, $F = \mathsf{Id}$, $\Lambda = \{\lambda_{\geq q}\}$, $\Lambda_F = \{\lambda_\bigcirc\}$ as in Example 6, and $L = \{\lambda_\bigcirc\}$. The operators $[\_]_{>q}$ of PCTL can now be defined as:

$$[\varphi]_{>q} \quad ::= \quad \neg[\neg\varphi]_{\geq 1-q}$$

for $q \in [0,1]$. Finally, the standard interpretation of PCTL over Markov chains is again obtained by viewing a Markov chain as a standard Borel space.

5. Concluding Remarks

We have defined maximal traces and executions for systems modelled as coalgebras of functors obtained as the composition of a computational type (given by an affine monad) with a transition type (typically given by a shapely polynomial endofunctor), under the additional assumption that the computational type preserves certain ω^{op}-limits. This assumption is not satisfied by the non-empty powerset functor, which we have treated separately. As a result, we have obtained maximal trace maps and maximal execution maps for deterministic, non-deterministic and stochastic systems.

We have subsequently used (maximal) execution maps to give semantics to generic path-based coalgebraic temporal logics, instances of which subsume known path-based logics such as CTL* and PCTL. Moreover, we have shown that by simply varying the transition type, interesting variants of known logics can be obtained with little effort.

Future work will generalise the results in Section 3 to arbitrary monads. Apart from the powerset and subprobability measure monads, non-affine monads of interest include the lift, finite list and finite multiset monads (with the latter being relevant to graded temporal logic). Unlike in the case of non-deterministic or stochastic systems, working with the affine submonads of the last three monads (which, as mentioned earlier, coincide with the identity monad) is not a solution. Incorporating non-affine monads into our treatment of maximal traces is expected to involve moving from cones over the image under J of the final sequence of F to *lax cones*, with a suitable DCpo-structure on homsets in $\mathsf{Kl}(T)$. We also plan to study the relationship between finite and maximal traces.

Another direction for future work is to investigate the expressive power of path-based temporal logics (in particular, how this compares in general to the expressive power of coalgebraic fixpoint logics), and to further develop the theory of these logics.

References

[1] I. Hasuo, B. Jacobs, A. Sokolova, Generic trace semantics via coinduction, Logical Methods in Computer Science 3 (2007) 1–36.

[2] E. Clarke, O. Grumberg, D. Peled, Model Checking, MIT Press, 1999.

[3] H. Hansson, B. Jonsson, A logic for reasoning about time and reliability, Formal Aspects of Computing 6 (5) (1994) 512–535.

[4] J. Rutten, Universal coalgebra: a theory of systems, Theoretical Computer Science 249 (2000) 3–80.

[5] D. Pattinson, Expressive logics for coalgebras via terminal sequence induction, Notre Dame Journal of Formal Logic 45 (1) (2004) 19–33.

[6] C. Cîrstea, D. Pattinson, Modular construction of complete coalgebraic logics, Theoretical Computer Science 388 (2007) 83–108.

[7] C. Cîrstea, M. Sadrzadeh, Modular games for coalgebraic fixed point logics, in: Proc. CMCS 2008, Vol. 203 of Electronic Notes in Theoretical Computer Science, Elsevier, 2008, pp. 71–92.

[8] C. Cîrstea, C. Kupke, D. Pattinson, Exptime tableaux for the coalgebraic μ-calculus, in: Proc. CSL 2009, Vol. 5771 of Lecture Notes in Computer Science, Springer, 2009, pp. 179–193.

[9] D. Kozen, Results on the propositional μ-calculus, Theoretical Computer Science 27 (3) (1983) 333 – 354.

[10] M. Dam, CTL* and ECTL* as fragments of the modal mu-calculus, Theoretical Computer Science 126 (1) (1994) 77–96.

[11] B. Jacobs, Trace semantics for coalgebras, in: Proc. CMCS 2004, Vol. 106 of Electronic Notes in Theoretical Computer Science, Elsevier, 2004, pp. 167–184.

[12] C. Cîrstea, Generic infinite traces and path-based coalgebraic temporal logics, in: Proc. CMCS 2010, Vol. 264 of Electronic Notes in Theoretical Computer Science, Elsevier, 2010, pp. 83–103.

[13] E. Emerson, J. Halpern, ”sometimes” and ”not never” revisited: On branching versus linear time, Journal of the ACM 33 (1986) 151–178.

[14] M. Giry, A categorical approach to probability theory, in: Categorical Aspects of Topology and Analysis, Vol. 915 of Lecture Notes in Mathematics, Springer, 1982, pp. 68–85.

[15] J. Doob, Measure Theory, Springer, 1994.

[16] C. Schubert, Terminal coalgebras for measure-polynomial endofunctors, in: Proc. TAMC 2009, Vol. 5532 of Lecture Notes in Computer Science, Springer, 2009, pp. 325–334.

[17] B. Jacobs, A. Sokolova, Traces, executions and schedulers, coalgebraically, in: Proc. CALCO 2009, Vol. 5728 of Lecture Notes in Computer Science, Springer, 2009, pp. 206–220.

[18] B. Jacobs, Semantics of weakening and contraction, Annals of Pure and Applied Logic 69 (1994) 73–106.

[19] A. Sokolova, Coalgebraic analysis of probabilistic systems, Ph.D. thesis, TU Eindhoven (2005).

[20] I. Viglizzo, Coalgebras on measurable spaces, Ph.D. thesis, Indiana University (2005).

[21] B. Davey, H. Priestley, Introduction to Lattices and Order, Cambridge University Press, 2002.