Expressivity of finitary coalgebraic logics

C. Kupke

Abstract

We give a brief introduction to the theory of coalgebras and discuss the connection between coalgebra and modal logic. In the second more technical part of this paper - summarizing known results - we prove that the finitary version of Moss' coalgebraic logic has the Hennessy-Milner property.

Summary

The purpose of this short paper is twofold. On the one hand we want to draw the attention of the more general audience to the relatively young field of universal coalgebra and its relationship with modal logics. On the other hand we present two facts about coalgebras which might be interesting for the part of the audience that has already some background knowledge in coalgebra. Therefore in the first part of the paper we will try to explain the notion of a coalgebra and of coalgebraic bisimulations by stating the definitions and by providing several examples. Nevertheless the paper cannot be seen as a proper introduction to coalgebras. The interested reader is referred to [4, 11] for an introduction to the field of universal coalgebra. Readers whose background lies in modal logic and its algebraic semantics are recommended to consult the recently published [13].

The paper also contains two technical results: In Section 2 we establish a link between the bisimilarity game and the final sequence of a functor (cf. Theorem 2.9). In Section 4 we obtain - following the ideas of Moss in [7] - the result that finitary coalgebraic logics have the Hennessy-Milner property on coalgebras for an ω-accessible, weak pullback preserving set functor (cf. Theorem 4.1).

1 Coalgebras

A coalgebra consists of a set of *states* [1] and a coalgebra map which can be used to make *observations* about coalgebra states. The type of these observations is specified by a functor. Often it is possible to assign to every coalgebra state x its *behavior* which can be thought of as a sequence of observations generated by repeatedly applying the coalgebra map to x and its successors. For example we can think of a coalgebra in which we can observe for each state x a letter a of some alphabet Σ and some successor state x'. In this case the behaviour of x will be an infinite Σ-word. Unlike in algebra, elements of a coalgebra are not constructed from generators and operations. In fact one should think of the elements of the coalgebra as given from the outset. Using the coalgebra map one can obtain (limited) information about them. This makes it possible to use coalgebras for modeling objects with infinite or non-wellfounded "behavior" such as infinite words, trees or graphs. A reader familiar with functional programming languages such as Haskell could obtain some intuition about coalgebras by thinking of them as of lazy implementations of infinite objects.

1.1 The formal definition

When introducing coalgebras it is almost inevitable to use category theoretic terminology. It should be noted, however, that for the understanding of this paper only an intuitive understanding

[1] This is not the most general definition, but we will only consider coalgebras that have a carrier set.

of the notion of a category and of a functor between categories is necessary. In fact we will mostly work in the category Set that has sets as *objects* and functions as *arrows*. A functor is a structure preserving map between categories. A set functor $\mathrm{F} : \mathsf{Set} \to \mathsf{Set}$ maps sets to sets and functions to functions. Furthermore F preserves the composition of functions and identities, i.e. $\mathrm{F}(f \circ g) = \mathrm{F}f \circ \mathrm{F}g$ and $\mathrm{F}\mathrm{id}_X = \mathrm{id}_{\mathrm{F}X}$ where for an arbitrary set X we denote by id_X the identity function on X.

Definition 1.1 Let $\mathrm{F} : \mathsf{Set} \to \mathsf{Set}$ be a functor. Then an F-*coalgebra* is a pair (X, γ) where $X \in \mathsf{Set}$ and $\gamma : X \to \mathrm{F}X \in \mathsf{Set}$. A *pointed* F-coalgebra $(\langle X, \gamma\rangle, x)$ is an F-coalgebra $\langle X, \gamma\rangle$ together with a designated point $x \in X$. $\lhd$

Example 1.2 1. Let Σ be a set (of colors). Furthermore let B_Σ be the functor that maps a set X to the set $\Sigma \times X \times X$ and a function $f : X \to Y$ to the function that maps a triple $(c, x, x') \in \Sigma \times X \times X$ to the triple $(c, f(x), f(x'))$. Then infinite Σ-labeled binary trees are pointed B_Σ-coalgebras: a given Σ-labeled tree $t : 2^* \to \Sigma$ corresponds to the coalgebra with carrier set 2^* and coalgebra map $w \mapsto (t(w), w0, w1)$.

2. Deterministic, finite Σ-word automata can be modeled as pointed coalgebras for the functor $\mathrm{A}_\Sigma = 2 \times (\_)^\Sigma$, where Σ is the finite alphabet and 2 is the two-element set, i.e. the functor A_Σ maps a set X to the set $2 \times X^\Sigma$ and a function $f : X \to Y$ to the function that maps a pair $(o, t) \in 2 \times X^\Sigma$ to the pair $(o, f \circ t)$. More about this fundamental example can be found in [9].

3. Image-finite Kripke frames correspond to coalgebras of the power set functor $\mathcal{P}_\omega$. The functor $\mathcal{P}_\omega$ maps a set to its collection of *finite subsets*, and a function $f : X \to Y$ to its direct image function $\mathcal{P}_\omega f : \mathcal{P}_\omega X \to \mathcal{P}_\omega Y$ given by $(\mathcal{P}_\omega f)(U) := \{f(x) \in Y \mid x \in U\}$. An image-finite Kripke frame (W, R) is then modeled as the coalgebra $(W, \lambda x.R[x])$, that is, the relation R is given by the function mapping a state x to the collection $R[x]$ of its (direct) successors.[2]

4. Let Φ be a set (of propositional variables) and consider the functor $\mathcal{P}\Phi \times \mathcal{P}_\omega$, that maps a set X to the set $\mathcal{P}\Phi \times \mathcal{P}_\omega X$ and a function $f : X \to Y$ to the function $\mathrm{id}_{\mathcal{P}\Phi} \times \mathcal{P}_\omega f$. Then an image finite Kripke model $(W, R, V : \Phi \to \mathcal{P}(W))$ corresponds to the $\mathcal{P}\Phi \times \mathcal{P}_\omega$-coalgebra $\langle W, \langle V^\sharp, \lambda x.R[x]\rangle : W \to \mathcal{P}\Phi \times \mathcal{P}_\omega W\rangle$ where $p \in V^\sharp(w)$ if $w \in V(p)$.

5. Let $\mathcal{D}_\omega$ be the functor that maps a set S to the set

$$\mathcal{D}_\omega S := \left\{\rho : S \rightharpoonup [0,1] \mid \rho \text{ has finite support and } \textstyle\sum_{s \in S} \rho(s) = 1\right\}$$

where we say that a (partial) function ρ has finite support if $\rho(s) \neq 0$ for only finitely many elements s of S. Then coalgebras for the functor $1 + \mathcal{D}_\omega$ correspond to the probabilistic transition systems by Larsen and Skou in [6]. Here $1 + \mathcal{D}_\omega$ denotes the functor that maps a set S to the disjoint union of the one-element set and the set $\mathcal{D}_\omega S$. Further details about this example can be found in [7, 12].

The message of these examples should be clear: coalgebras can be used to uniformly model various types of transition systems. The advantage of such a uniform approach lies in the fact that notions, that are usually studied in isolation for different system types, can now be seen as instances of one coalgebraic concept. As examples for this phenomenon we are now going to present the notions of a coalgebra morphism and of a coalgebraic bisimulation. After that we look at some concrete instances of these coalgebraic notions.

Definition 1.3 Let $\mathbb{X} = \langle X, \gamma\rangle$ and $\mathbb{Y} = \langle Y, \delta\rangle$ be two F-coalgebras then a function $f : X \to Y$ is an F-coalgebra morphism from $\mathbb{X}$ to $\mathbb{Y}$ if $\delta \circ f = \mathrm{F}f \circ \gamma$, i.e. if the left diagram in Figure 1 commutes.

A relation $Z \subseteq X \times Y$ is an F-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ if there exists a function $\mu : Z \to \mathrm{F}Z$ such that the projection maps $\pi_1 : Z \to X$ and $\pi_2 : Z \to Y$ are coalgebra morphisms, i.e., such that the right diagram in Figure 1 commutes. Two states are *(*F*-)bisimilar* if they are linked by some bisimulation. We write $x \leftrightarroweq_{\mathrm{F}} y$ if the states x and y are F-bisimilar. $\lhd$

[2]If we drop the finiteness condition from the definition of the functor $\mathcal{P}_\omega$ we obtain the power set functor $\mathcal{P}$. $\mathcal{P}$-coalgebras are Kripke frames.

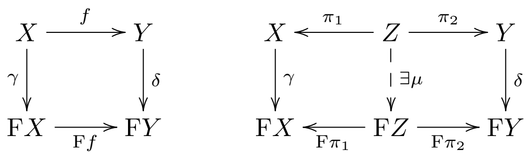

Figure 1: F-coalgebra morphism and F-bisimulation

Let us have a look at some concrete instances of F-coalgebra morphisms and F-bisimulations.

Example 1.4 1. Let $\mathrm{F} = \mathrm{B}_\Sigma$ (cf. Example 1.2(1)) for some set of labels Σ. Then $f : X \to Y$ is a B_Σ-coalgebra morphism between two B_Σ-coalgebras $\langle X, \langle c, \gamma\rangle : X \to \Sigma \times (X \times X)\rangle$ and $\langle Y, \langle d, \delta\rangle : Y \to \Sigma \times (Y \times Y)\rangle$ if for all $x \in X$ we have $c(x) = d(f(x))$ and if $\gamma(x) = (x_1, x_2)$ implies $\delta(f(x)) = (f(x_1), f(x_2))$. In words, f has to preserve the labeling and the binary tree structure. It is not difficult to check that two binary trees t_1 and t_2 are B_Σ-bisimilar iff they are identical.

2. Let $\mathrm{F} = \mathrm{A}_\Sigma = 2 \times (\_)^\Sigma$. Then pointed F-coalgebras with finite carrier set correspond to finite, deterministic automata (cf. Example 1.2(2)). Two automata accept the same language iff the corresponding pointed F-coalgebras are F-bisimilar (cf. [9]).

3. Recall from Example 1.2(4) that F-coalgebras for the functor $\mathrm{F} = \mathcal{P}\Phi \times \mathcal{P}_\omega$ correspond to image-finite Kripke models. Then F-coalgebra morphisms are precisely the bounded morphisms between Kripke models. F-bisimulations correspond exactly to the bisimulations from modal logic, i.e. a relation R between two (image-finite) Kripke models (W_1, R_1, V) and (W_2, R_2, V) is a modal bisimulation iff R is a $\mathcal{P}\Phi \times \mathcal{P}_\omega$-bisimulation between the corresponding $\mathcal{P}\Phi \times \mathcal{P}_\omega$-coalgebras.

4. Let F be the functor $1 + \mathcal{D}_\omega$ (cf. Example 1.2(5)). It can be shown that F-bisimulations between F-coalgebras coincide with the probabilistic bisimulations from [6] between the corresponding probabilistic transition systems. For the details we again refer to [7, 12].

In Section 2 we will continue our discussion of bisimilarity by showing that coalgebraic bisimilarity has also a nice game-theoretic interpretation. Furthermore in Section 3 we will recall the definition of some logical language that can be used for reasoning about F-coalgebras. In the remainder of this section we will fix the technical preliminaries.

1.2 Technicalities

In the following we will make some assumptions on the functor F under consideration. Readers who are not interested in all the technical details are advised to skip this section and to think of F in the sequel as of some functor in Example 1.2.

The first assumption we make is that all our functors are set functors. Furthermore we will assume the functors we are working with to be

- *standard:* for all set X and Y such that $X \subseteq Y$ we have that $\mathrm{F}X \subseteq \mathrm{F}Y$ and the inclusion map $\iota_{X,Y}$ from X into Y is mapped by F to the inclusion map $\iota_{\mathrm{F}X,\mathrm{F}Y}$ from $\mathrm{F}X$ into $\mathrm{F}Y$.
- *ω-accessible:* for all sets Y and all elements $t \in \mathrm{F}Y$ there is a finite set $X \subseteq_\omega Y$ such that $t \in \mathrm{F}X$.

Moreover all the functors that we are considering are *weak pullback preserving.* This is a property which often occurs in the coalgebraic literature. For details about this important, but rather technical condition we refer the reader to [11]. From now on we assume F to be a standard, ω-accessible and weak pullback preserving set functor.

1.3 Relation Lifting

A useful property of F is that it has a well-behaved *relation lifting*. We will first define this relation lifting, then list the properties we need. In the sequel this relation lifting allows us to use a concise formulation of the notion of a coalgebraic bisimulation. In addition to that it is central in the definition of the coalgebraic logic we are considering.

Definition 1.5 Given two sets X and Y, and a binary relation Z between $X \times Y$, we define the *lifted relation* $\overline{\mathrm{F}}(Z) \subseteq \mathrm{F}X \times \mathrm{F}Y$ as follows:

$$\overline{\mathrm{F}}(Z) := \{((\mathrm{F}\pi)(\phi), (\mathrm{F}\pi')(\phi)) \mid \phi \in \mathrm{F}Z\},$$

where $\pi : Z \to X$ and $\pi' : Z \to Y$ are the projection functions given by $\pi(x, y) = x$ and $\pi'(x, y) = y$. $\lhd$

Example 1.6 1. Let Σ be a set of colors and let $\mathrm{F} = \mathrm{B}_\Sigma$ and let $R \subseteq X \times Y$ be a relation. Then for arbitrary elements $(c, x_1, x_2) \in \mathrm{B}_\Sigma X$ and $(d, y_1, y_2) \in \mathrm{B}_\Sigma Y$ we have $((c, x_1, x_2), (d, y_1, y_2)) \in \overline{\mathrm{F}}R$ iff $c = d$ and $(x_i, y_i) \in R$ for $i = 1, 2$.

2. Let Φ be a set (of propositional variables), let $\mathrm{F} := \mathcal{P}\Phi \times \mathcal{P}_\omega$ and let $R \subseteq X \times Y$ be a relation. Then for arbitrary $P \subseteq \Phi$, $U \subseteq_\omega X$, $Q \subseteq \Phi$ and $V \subseteq_\omega Y$ we have

$$((P, U), (Q, V)) \in \overline{\mathrm{F}}R \quad \text{iff} \quad \begin{array}{ll} (\textsc{Prop}) & P = Q \text{ and} \\ (\textsc{Forth}) & \forall u \in U.\exists v \in V \text{ s.t.}(u, v) \in R \text{ and} \\ (\textsc{Back}) & \forall v \in V.\exists u \in U \text{ s.t.}(u, v) \in R. \end{array}$$

Readers familiar with modal logic or process algebra will recognize the similarity of the relation lifting for the functor $\mathcal{P}\Phi \times \mathcal{P}_\omega$ and the usual definition of a bisimulation between Kripke frames/transition systems. This is not a coincidence as we will see when looking at the formulation of the notion of an F-bisimulation in terms of the relation lifting (cf. Fact 2.1 below). For now we just list some properties of the relation lifting of a standard and weak pullback preserving set functor.

Fact 1.7 *The relation lifting* $\overline{\mathrm{F}}$ *satisfies the following properties, for all functions* $f : X \to Y$*, all relations* $R, Q \subseteq X \times Y$*, and all subsets* $U \subseteq X$*,* $V \subseteq Y$*:*
(1) $\overline{\mathrm{F}}$ *extends* F*:* $\overline{\mathrm{F}}(Gr(f)) = Gr(\mathrm{F}f)$*;* ($Gr(f)$ denotes the graph of f)
(2) $\overline{\mathrm{F}}$ *commutes with relation converse:* $\overline{\mathrm{F}}(R^{\smile}) = (\overline{\mathrm{F}}R)^{\smile}$*;* ($R^{\smile}$ denotes the converse of R)
(3) $\overline{\mathrm{F}}$ *is monotone: if* $R \subseteq Q$ *then* $\overline{\mathrm{F}}(R) \subseteq \overline{\mathrm{F}}(Q)$*;*
(4) $\overline{\mathrm{F}}$ *distributes over composition:* $\overline{\mathrm{F}}(R \circ Q) = \overline{\mathrm{F}}(R) \circ \overline{\mathrm{F}}(Q)$*;*
(5) $\overline{\mathrm{F}}$ *commutes with restrictions, i.e.* $\overline{\mathrm{F}}(R\restriction_{U \times V}) = (\overline{\mathrm{F}}R)\restriction_{\mathrm{F}U \times \mathrm{F}V}$ *.*

For proofs we refer to [7, 1], and references therein.

2 Finite approximations of bisimilarity

In this section we state the definition of the bisimilarity game for coalgebras. Furthermore we discuss a variation of the bisimilarity game, the n-bisimilarity game. Finally we link the game-theoretic perspective to the categorical notion of the final sequence of a functor.

As a preparation we recall that F-bisimulations can be characterized using the relation lifting of the functor F.[3]

Fact 2.1 *[10] Let* $\mathbb{X} = \langle X, \gamma \rangle$ *and* $\mathbb{Y} = \langle Y, \delta \rangle$ *be two* F*-coalgebras. A relation* $Z \subseteq X \times Y$ *is an* F*-bisimulation between* $\mathbb{X}$ *and* $\mathbb{Y}$ *if* $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}(Z)$ *for all* $(x, y) \in Z$*.*

This fact plays an important rôle in the definition of the bisimilarity game.

[3]This fact goes back to the work by Hermida and Jacobs in [3].

2.1 Bisimilarity game

The definition of the bisimilarity game for coalgebras goes back to Baltag (cf. [1]). The idea is as follows: the game is played between the two players $\exists$ (Éloise) and $\forall$ (Abélard). Given two pointed F-coalgebras $(\mathbb{X}, x)$ and $(\mathbb{Y}, y)$, $\exists$ claims that (x, y) are bisimilar, i.e. her claim is that *there exists* some F-bisimulation Z between $\mathbb{X}$ and $\mathbb{Y}$ that relates x and y. The game board consists of pairs of coalgebra states (x, y) ($\exists$'s positions) and relations $Z \subseteq X \times Y$ ($\forall$'s positions). At a position $(x, y) \in X \times Y$ it is $\exists$'s turn and she has to move to a "local bisimulation", i.e. a relation Z such that $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}Z$ (hence Z fulfills the condition from Fact 2.1 "locally"). In turn, $\forall$ has to move from a position Z to a pair of states $(x', y') \in Z$. A match of the game is lost by a player who cannot move. All infinite plays are won by $\exists$. The following definition introduces the bisimilarity game more formally - readers who are not so familiar with the game-theoretic notions are referred to [14] for a short summary of the necessary details and to [2] for a detailed introduction to infinite (parity) graph games.

Definition 2.2 Let $\mathbb{X} := \langle X, \gamma \rangle$ and $\mathbb{Y} := \langle Y, \delta \rangle$ be F-coalgebras. Then the arena of the *image-finite F-bisimilarity game* $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ is given by the bipartite graph $(B_\exists, B_\forall, E)$ that is described by the following table:

Position: b	Player	Admissible moves: $E[b]$
$(x, y) \in X \times Y$	$\exists$	$\{Z \subseteq X \times Y \mid (\gamma(x), \delta(y)) \in \overline{\mathrm{F}}Z\}$
$Z \in \mathcal{P}(X \times Y)$	$\forall$	$\{(x', y') \mid (x', y') \in Z\}$

Here the second column indicates whether a given position b belongs to player $\exists$ or $\forall$, i.e. whether $b \in B_\exists$ or $b \in B_\forall$, and $\overline{\mathrm{F}}Z$ is the relation lifting of Z. A match of $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ starts at some position $b_0 \in B_\exists \cup B_\forall$ and proceeds as follows: at position $b \in B_\exists$ player $\exists$ has to move to a position $b' \in E[b]$ and likewise at position $b \in B_\forall$ player $\forall$ has to move to some $b' \in E[b]$. A player who cannot move ("gets stuck") loses the match and all infinite matches are won by $\exists$. Two successive moves in a match are called a *round* of the game. $\lhd$

Fact 2.3 *Let $\mathbb{X} := \langle X, \gamma \rangle, \mathbb{Y} := \langle Y, \delta \rangle$ be* F*-coalgebras and let $x \in X, y \in Y$. Then $\langle X, \gamma \rangle, x \leftrightarrows \langle Y, \delta \rangle, y$ iff $\exists$ has a winning strategy in the game $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ starting at position (x, y).*

For a sketch of the proof of this fact cf. [14]. The game-theoretic analysis of F-bisimulations naturally leads to the notion of n-bisimilarity.

Definition 2.4 We say that x and y are n-bisimilar (notation: $(\langle X, \gamma \rangle, x) \leftrightarrows_n (\langle Y, \delta \rangle, y)$) if $\exists$ has a strategy in the game $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ starting at position (x, y) such that $\exists$ does not lose the game in less than n rounds. $\lhd$

The next subsection will be devoted to proving that n-bisimilarity coincides with bisimilarity given our assumption that the functor F is ω-accessible.

2.2 Final sequence

We now introduce the "finitary part" of the final sequence of a functor, i.e. its first ω elements. This final sequence plays an important rôle in the theory of coalgebras where it is used in order to compute or approximate the final F-coalgebra. We only state and motivate the basic definition. For more details about the final sequence of a functor we refer the reader to [15] and references therein.

Definition 2.5 Given a set functor F we inductively define functions $p_i : \mathrm{F}^{i+1}1 \to \mathrm{F}^i 1$ for all $i \in \mathbb{N}$ by putting $p_0 :=!_{\mathrm{F}1}$ and $p_{i+1} := \mathrm{F}p_i$. Here 1 denote the one-element set, $!_X$ denotes the (unique) function from a set X to the one-element set 1 and for a set X we write $\mathrm{F}^0 X := X$ and $\mathrm{F}^{i+1} X := \mathrm{F}(\mathrm{F}^i X)$. For all $n \in \mathbb{N}$ elements of the set $\mathrm{F}^n 1$ will be called *n-step behavior.* $\lhd$

The sequence is depicted in the lower part of Figure 2. Let us give an example which illustrates why we chose the name "n-step behavior" for members of the nth element of the sequence.

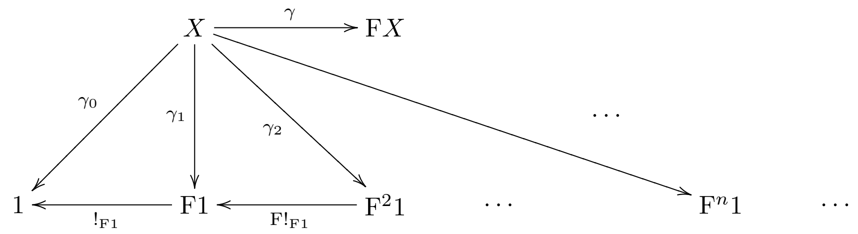

Figure 2: Final sequence and n-step behavior maps

Example 2.6 Let Σ be set of labels and let $\mathrm{F} = \mathrm{B}_\Sigma$. Then it is not difficult to see that for all $n \in \mathbb{N}$ the set $\mathrm{F}^n 1$ is isomorphic to the set of finite Σ-labeled binary trees of depth n.

We will make use of the fact that given an F-coalgebra $\langle X, \gamma \rangle$ one can easily define a sequence $\{\gamma_n\}_{n\in\mathbb{N}}$ of functions that map any coalgebra state $x \in X$ to its n-step behavior. We now give the formal definition of these γ_n's, Figure 2 provides a picture of the situation.

Definition 2.7 Given an F-coalgebra $\mathbb{X} = \langle X, \gamma \rangle$ and a state $x \in X$ we define a family of maps $\{\gamma_n : X \to \mathrm{F}^n 1\}_{n\in\mathbb{N}}$ by putting

$$\begin{array}{lcl} \gamma_0(x) & := & !_X \\ \gamma_{i+1}(x) & := & \mathrm{F}\gamma_n \circ \gamma \end{array} \qquad \lhd$$

Intuitively for each $n \in \mathbb{N}$ the map γ_n maps a state x to its n-step behavior.

Example 2.8 In the above example for $\mathrm{F} = \mathrm{B}_\Sigma$ recall from Example 1.2 that a Σ-labeled infinite binary tree t corresponds to a B_Σ-coalgebra. In this case, intuitively speaking, the map γ_n maps t to the Σ-tree of depth n that consists of the first n levels of t.

For much more information and results about the final sequence of an ω-accessible functor the reader is referred to [15]. We now make a connection between the bisimilarity game and the final sequence of F.

Theorem 2.9 *Let* $(\langle X, \gamma \rangle, x)$ *and* $(\langle Y, \delta \rangle, y)$ *be pointed* F*-coalgebras. Then for all* $n \in \mathbb{N}$ *we have*

$$\gamma_n(x) = \delta_n(y) \quad \textit{iff} \quad x \leftrightarrows_n y.$$

Proof. For $i \in \mathbb{N}$ let $Z_i := Gr(\gamma_i) \circ Gr(\delta_i)^\smile$ and put $B_i := \{(x', y') \in X \times Y \mid (\langle X, \gamma \rangle, x) \leftrightarrows_n (\langle Y, \delta \rangle, y)\}$. The claim of the proposition is equivalent to saying that

$$Z_i = B_i \quad \text{for all } n \in \mathbb{N} \tag{1}$$

In order to facilitate the proof of (1) we first show that the following holds for all $(x, y) \in X \times Y$ and all $i \in \mathbb{N}$:

$$(x, y) \in Z_{i+1} \quad \text{iff} \quad (\gamma(x), \delta(y)) \in \overline{\mathrm{F}} Z_i. \tag{2}$$

This can be seen by just spelling out the definitions and by using Fact 1.7:

$$\begin{array}{lll} (x, y) \in Z_{i+1} & \text{iff} \;\; (x, y) \in Gr(\gamma_{i+1}) \circ Gr(\delta_{i+1})^\smile & \text{(Def. of } Z_{i+1}) \\ & \text{iff} \;\; (x, y) \in Gr(\mathrm{F}\gamma_i \circ \gamma) \circ Gr(\mathrm{F}\delta_i \circ \delta)^\smile & \text{(Def. of } \gamma_{i+1}, \delta_{i+1}) \\ & \text{iff} \;\; (x, y) \in Gr(\gamma) \circ Gr(\mathrm{F}\gamma_i) \circ Gr(\mathrm{F}\delta_i)^\smile \circ Gr(\delta)^\smile & \\ & \text{iff} \;\; (\gamma(x), \delta(y)) \in Gr(\mathrm{F}\gamma_i) \circ Gr(\mathrm{F}\delta_i)^\smile & \\ & \text{iff} \;\; (\gamma(x), \delta(y)) \in \overline{\mathrm{F}}(Gr(\gamma_i) \circ Gr(\delta_i)) = \overline{\mathrm{F}} Z_i & \text{(Fact 1.7)} \end{array}$$

We now prove (1) by induction on n. For $n = 0$ there is nothing to prove, because $Z_0 = B_0 = X \times Y$. Inductively assume that $Z_i = B_i$ for some $i \in \mathbb{N}$. We have to show that $Z_{i+1} = B_{i+1}$.

$\subseteq$ Let $(x, y) \in Z_{i+1}$, i.e. $\gamma_{i+1}(x) = \delta_{i+1}(y)$. This implies by (2) that $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}Z_i$. Therefore $\exists$ can move in the bisimilarity game $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ from position (x, y) to position Z_i. By I.H. we know that all positions $(x', y') \in Z_i$ are in B_i, i.e. $\exists$ has a strategy such that for all $(x', y') \in Z_i$ she does not lose any match starting at (x', y') in less than i rounds. As a consequence by moving from (x, y) to Z_i she has a strategy in $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ such that she does not lose in less than $i + 1$ rounds which means that $(x, y) \in B_{i+1}$.

$\supseteq$ Suppose $(x, y) \in B_{i+1}$, i.e. $\exists$ has a strategy in $\mathcal{B}(\mathbb{X}, \mathbb{Y})$ such that she does not lose any match starting at position (x, y) in less than $i + 1$ rounds. In particular she can move from (x, y) to some relation $R \subseteq X \times Y$ such that $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}R$ and such that $R \subseteq B_i$. Therefore, by monotonicity of $\overline{\mathrm{F}}$, $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}B_i$ and thus by I.H. $(\gamma(x), \delta(y)) \in \overline{\mathrm{F}}Z_i$. This implies by (2) that $(x, y) \in Z_{i+1}$.

QED

As a corollary we obtain the earlier announced result that the game-theoretically defined notion of n-bisimilarity forms a good finite approximation of bisimilarity.

Corollary 2.10 *Let* F *be a set functor that meets the requirements in Section 1.2. Furthermore let* $\mathbb{X} = \langle X, \gamma \rangle, \mathbb{Y} = \langle Y, \delta \rangle$ *be two* F*-coalgebras with designated points* $x \in X$ *and* $y \in Y$*. Then*

$$x \leftrightarrows_{\mathrm{F}} y \quad \textit{iff} \quad \forall n.\ x \leftrightarrows_n y.$$

Proof. It has been shown in [15] that $\gamma_n(x) = \delta_n(y)$ for all $n \in \mathbb{N}$ implies $\mathbb{X}, x \leftrightarrows_{\mathrm{F}} \mathbb{Y}, y$ under the assumption that F is ω-accessible. Therefore the claim of the proposition is an immediate consequence of 2.9. QED

3 Modal logic in a coalgebraic shape

In the first section of this paper we showed that coalgebras provide a framework for modeling various types of transition systems. The natural question to ask is whether we can also develop logical languages to reason about coalgebras. Modal languages have been successfully employed for reasoning about transition systems. Hence these languages are a good candidate for the specification of coalgebras. The various coalgebraic modal languages that have been proposed to talk about coalgebras can roughly be split into two groups.

1. Languages whose modalities are given by so-called *predicate liftings*. For an introduction to these languages we refer the reader to [5, 8, 13] and references therein.
2. Languages that are defined via the *relation lifting* of F. This approach has originally been introduced by Moss in [7].

We will now present the definition of a version of Moss' coalgebraic logic that has a finitary syntax. This logic has been introduced in [14]. Unlike *loc.cit.* we do not consider fixed point operators. Furthermore, as mentioned before, we work under the general requirement that we are given a set functor F that is ω-accessible, standard and weak pullback preserving.

Definition 3.1 The *language* of (finitary) coalgebraic logic $\mathcal{L}^{\mathrm{F}}$ is defined inductively as follows:

$$\begin{array}{rcl} \mathcal{L}_0^{\mathrm{F}} \ni \phi & ::= & \bot \mid \top \mid \phi \wedge \phi \mid \phi \vee \phi \\ \mathcal{L}_{i+1}^{\mathrm{F}} \ni \phi & ::= & \psi \in \mathcal{L}_i^{\mathrm{F}} \mid \phi \wedge \phi \mid \phi \vee \phi \mid \nabla \pi,\ \pi \in \mathrm{F}\mathcal{L}_i^{\mathrm{F}} \\ \mathcal{L}^{\mathrm{F}} & := & \bigcup_{i \in \mathbb{N}} \mathcal{L}_i^{\mathrm{F}} \end{array}$$

The *depth* of a formula $\phi \in \mathcal{L}^{\mathrm{F}}$ is defined as the smallest natural number i_ϕ such that $\phi \in \mathcal{L}_{i_\phi}^{\mathrm{F}}$. $\lhd$

Remark 3.2 From the definition it is clear that $\mathcal{L}^{\mathrm{F}}$ is a set (in contrast to Moss' original language which consisted of a proper class of formulas). The difference with Moss' original definition is that the syntax only contains finite conjunctions and, in addition to that, finite disjunctions. Furthermore Venema defines the notion of a subformula and shows in [14] that every formula of $\mathcal{L}^{\mathrm{F}}$ has a finite number of subformulas. These facts justify to say that $\mathcal{L}^{\mathrm{F}}$ has a finitary syntax.

The semantics of the logic is defined as follows.

Definition 3.3 Let $\mathbb{X} = \langle X, \gamma \rangle$ be an F-coalgebra We inductively define a relation $\models^{\mathbb{X}} \subseteq X \times \mathcal{L}^{\mathrm{F}}$ with the intended meaning that $(x, \phi) \in \models^{\mathbb{X}}$ if ϕ is satisfied at $x \in X$. In this case we also write $x \models^{\mathbb{X}} \phi$. The inductive definition of $\models^{\mathbb{X}}$ is as follows: We put $x \not\models^{\mathbb{X}} \bot$, $x \models^{\mathbb{X}} \top$ and

$$\begin{array}{llcll} x & \models^{\mathbb{X}} & \phi_1 \wedge \phi_2 & \text{if} & x \models^{\mathbb{X}} \phi_1 \text{ and } x \models^{\mathbb{X}} \phi_2, \\ x & \models^{\mathbb{X}} & \phi_1 \vee \phi_2 & \text{if} & x \models^{\mathbb{X}} \phi_1 \text{ or } x \models^{\mathbb{X}} \phi_2, \\ x & \models^{\mathbb{X}} & \nabla \pi & \text{if} & (\gamma(x), \pi) \in \overline{\mathrm{F}}(\models^{\mathbb{X}}). \end{array}$$

If the $\mathbb{X}$ is clear from the context we simply write $\models$ for the relation $\models^{\mathbb{X}} \subseteq X \times \mathcal{L}^{\mathrm{F}}$. Furthermore we write $\mathrm{Th}(x)$ for the *theory of* x, i.e. $\mathrm{Th}(x) := \{\phi \in \mathcal{L}^{\mathrm{F}} \mid x \models \phi\}$. $\lhd$

Remark 3.4 The definition of the semantics of a formula $\nabla \pi \in \mathcal{L}^{\mathrm{F}}_{i+1}$ is not circular because

$$(\gamma(x), \pi) \in \overline{\mathrm{F}}(\models) \quad \text{iff} \quad (\gamma(x), \pi) \in \overline{\mathrm{F}}\Big((\models) \restriction_{X \times \mathcal{L}^{\mathrm{F}}_i}\Big),$$

where $(\models) \restriction_{X \times \mathcal{L}^{\mathrm{F}}_i}$ denotes the restriction of $\models^{\mathbb{X}} \subseteq X \times \mathcal{L}^{\mathrm{F}}$ to $X \times \mathcal{L}^{\mathrm{F}}_i$. This can be seen by using that $\pi \in \mathrm{F}\mathcal{L}^{\mathrm{F}}_i$ and by using Fact 1.7(5). The depth of formulas in $\mathcal{L}_i$ is strictly smaller than the depth of $\nabla \pi$ and thus we can inductively assume that $(\models) \restriction_{X \times \mathcal{L}^{\mathrm{F}}_i}$ has been already defined.

Example 3.5 Let Φ be a set of propositional variables and consider the functor $\mathrm{F} = \mathcal{P}\Phi \times \mathcal{P}_\omega\_$. Then the ∇-formulas in $\mathcal{L}^{\mathrm{F}}$ are of the form $\nabla(P, \Psi)$ where $P \subseteq \Phi$ is a set of propositional variables and Ψ is a set of formulas. Informally speaking such a formula $\nabla(P, \Psi)$ corresponds to the following formula in the usual syntax of modal logic:

$$\nabla(P, \Psi) \equiv \bigwedge_{p \in P} p \wedge \bigwedge_{p \notin P} \neg p \wedge \bigwedge_{\psi \in \Psi} \Diamond \psi \wedge \Box\Big(\bigvee_{\psi \in \Psi} \psi\Big).$$

Similarly one can translate modal formulas into formulas of $\mathcal{L}^{\mathrm{F}}$ (cf. [14, Sec. 5]).

4 Expressivity of finitary coalgebraic logics

In this section we prove that the finitary coalgebraic logic as defined in Def. 3.1 has the Hennessy-Milner property, i.e. we prove Theorem 4.1 below. Let us once more stress that despite the fact that we think that Theorem 4.1 has not been proven before, its proof can be seen as a rather straightforward adaptation of the proof of an analogue result for infinitary coalgebraic logic in [7].

Theorem 4.1 *Let* F *be an* ω*-accessible set functor. Then the language* $\mathcal{L}^{\mathrm{F}}$ *has the Hennessy-Milner property, i.e. for all pointed* F*-coalgebras* $(\langle \mathbb{X}, \gamma \rangle, x)$ *and* $(\langle \mathbb{Y}, \delta \rangle, y)$ *we have*

$$\mathrm{Th}(x) = \mathrm{Th}(y) \quad \textit{iff} \quad (\langle \mathbb{X}, \gamma \rangle, x) \leftrightarrows_{\mathrm{F}} (\langle \mathbb{Y}, \delta \rangle, y),$$

i.e. two pointed F*-coalgebras satisfy the same* $\mathcal{L}^{\mathrm{F}}$*-formulas iff they are* F*-bisimilar.*

As usual the easier part of proving the Hennessy-Milner property is to prove the invariance of the semantics of a formula under bisimilarity.

Fact 4.2 ([14]) *Let* F *be an* ω*-accessible functor and let* $(\langle \mathbb{X}, \gamma \rangle, x)$ *and* $(\langle \mathbb{Y}, \delta \rangle, y)$ *be pointed* F*-coalgebras such that* $(\langle \mathbb{X}, \gamma \rangle, x) \leftrightarrows_{\mathrm{F}} (\langle \mathbb{Y}, \delta \rangle, y)$*. Then for all formulas* $\phi \in \mathcal{L}^{\mathrm{F}}$ *we have*

$$(\langle \mathbb{X}, \gamma \rangle, x) \models \phi \quad \textit{iff} \quad (\langle \mathbb{Y}, \delta \rangle, y) \models \phi.$$

For the proof of the fact that logical equivalence implies bisimilarity we use our result from the previous section, where we saw that n-bisimilarity for all n and bisimilarity coincide if the functor is ω-accessible. The idea is to define for each $n \in \mathbb{N}$ a set of *characteristic formulas* that characterize the behavior of a given pointed F-coalgebra up-to depth n. These formulas have been introduced in [7] in order to prove the analogue of Theorem 4.1 for infinitary coalgebraic logic.

Definition 4.3 For $i \in \mathbb{N}$ we define a function $\nabla_i : \mathrm{F}^i 1 \to \mathcal{L}^{\mathrm{F}}$ by putting

$$\begin{array}{lcll} \nabla_0(*) & := & \top & \text{for } * \in 1 \\ \nabla_{i+1}(x) & := & \nabla\,(\mathrm{F}\nabla_i(x)) & \text{for } x \in \mathrm{F}^{i+1}1 \end{array} \qquad \lhd$$

The intended meaning of the ∇_i-maps is that they map a given state $x \in \mathrm{F}^i 1$, which represents some possible i-step behavior, to its characteristic formula $\nabla_i(x)$ of depth i. Let us have a look at an example.

Example 4.4 Let Φ be a set. For the functor $\mathrm{F} = \mathcal{P}\Phi \times \mathcal{P}_\omega$ the characteristic formulas look of depth ≤ 1 can be computed as follows:

$$\Psi_0 = \{\top\}, \quad \Psi_1 = \{\nabla(P, \{*\}) \mid P \in \mathcal{P}\Phi\} \cup \{\nabla(P, \emptyset) \mid P \in \Phi\}$$

where Ψ_i denotes the image of ∇_i. We can easily see that formulas in Ψ_1 precisely determine the 1-step behavior of a given pointed F-coalgebra: the formula $\nabla(P, \emptyset)$ will be true at any state in which exactly the propositional variables in P are true and which has no successors. Likewise the formula $\nabla(P, \{*\})$ expresses that precisely the p's in Φ are true and there is a successor.

The following observations about the ∇_i-maps form the justification for calling the formulas of the form $\nabla_i(\theta)$ "characterizing formulas": a coalgebra state x that has a certain n-step behavior θ satisfies precisely the formula $\nabla_n(\theta)$, i.e. no other formula of the form $\nabla_n(\theta')$ for some $\theta' \neq \theta$ is true at x.

Lemma 4.5 *Let $n \in \mathbb{N}$ and let θ be an element of $\mathrm{F}^n 1$. Furthermore let $(\langle X, \gamma\rangle, x)$. Then*

$$x \models \nabla_n \theta \qquad \textit{iff} \quad \theta = \gamma_n(x),$$

where $\gamma_n(x)$ denotes the n-step behavior of x as defined in Definition 2.7 above.

Proof. The claim of the lemma can be written in terms of relations in the following way:

$$(\models_X \circ Gr(\nabla_n)\breve{}\,) = Gr(\gamma_n). \tag{3}$$

We define $R_n := (\models_X \circ Gr(\nabla_n)\breve{}\,)$. Equation (3) will be proven by induction on n. The base case $n = 0$ is easy to check.

Inductively assume now that (3) is true for some $i \in \mathbb{N}$. We show that in this case our claim also holds for $n = i + 1$. In order to see this note that by definition of R_{i+1} we have $(x, \theta) \in R_{i+1}$ iff $x \models^{\mathbb{X}} \nabla_{i+1}(\theta)$. Unfolding the definition of ∇_{i+1} we obtain

$$\begin{array}{lll} (x,\theta) \in R_{i+1} & \text{iff} & x \models^{\mathbb{X}} \nabla\,(\mathrm{F}\nabla_i(\theta)) & \\ & \text{iff} & (\gamma(x), \mathrm{F}\nabla_i(\theta)) \in \overline{\mathrm{F}}(\models^{\mathbb{X}}) & \text{(Def. of } \models) \\ & \text{iff} & (\gamma(x), \theta) \in \overline{\mathrm{F}}(\models^{\mathbb{X}}) \circ Gr(\mathrm{F}\nabla_i)\breve{} & \\ & \text{iff} & (\gamma(x), \theta) \in \overline{\mathrm{F}}(\models^{\mathbb{X}} \circ Gr(\nabla_i)\breve{}\,) & \text{(Fact 1.7)} \\ & \text{iff} & (\gamma(x), \theta) \in \overline{\mathrm{F}}R_i = \overline{\mathrm{F}}Gr(\gamma_i) & \text{(Def. of } R_i + \text{I.H.)} \\ & \text{iff} & (x, \theta) \in Gr(\gamma) \circ Gr(\mathrm{F}\gamma_i) & \text{(Fact 1.7)} \\ & \text{iff} & \theta = (\mathrm{F}\gamma_i \circ \gamma)(x) = \gamma_{i+1}(x) & \text{(Def. of } \gamma_{i+1}) \end{array}$$

QED

We are now well prepared for finishing the proof of Theorem 4.1.

Proof of Theorem 4.1. The fact that two F-bisimilar pointed F-coalgebras are also logically equivalent was proven in Proposition 4.2. This is the implication from right to left in the theorem.

For the converse direction let $(\langle X, \gamma\rangle, x)$ and $(\langle Y, \delta\rangle, y)$ be pointed F-coalgebras. We show that $x \not\rightleftarrows y$ implies that $\mathrm{Th}(x) \neq \mathrm{Th}(y)$. By Corollary 2.10 we know that $x \not\rightleftarrows y$ entails that there is some $n \in \mathbb{N}$ such that $(\langle X, \gamma\rangle, x) \not\rightleftarrows_n (\langle Y, \delta\rangle, y)$, i.e. such that the given pointed coalgebras are not n-bisimilar. By Thm. 2.9 this implies $\gamma_n(x) \neq \delta_n(y)$. Therefore by Lemma 4.5 we have $x \models \nabla_n(\gamma_n(x))$ and $y \not\models \nabla_n(\gamma_n(x))$, i.e. $\mathrm{Th}(x) \neq \mathrm{Th}(y)$ as required. QED

5 Conclusion

We clarified the connection between the n-bisimilarity game and the final sequence of the functor. Furthermore we showed that Moss' arguments in [7] can be used to prove that finitary coalgebraic logics have the Hennessy-Milner property over coalgebras for an ω-accessible, weak pullback preserving functor.

The watchful reader will have noted that we did not make use of the game-theoretic interpretation of bisimilarity at all. We could provide a direct game-theoretic proof of Cor. 2.10 but did not do so due to lack of space. Furthermore we hope to obtain a better understanding of the bisimilarity game by further exploring the connection between the n-bisimilarity game and the final sequence that has been established in Thm. 2.9. As an example consider the functor $\mathrm{F} = \mathcal{P}_\omega(X)^A$ for some infinite set A. This functor is *not* ω-accessible but, using Worrell's results on the final sequence of F, bisimilarity can be approximated by n-bisimulations. This fact is not immediately obvious when looking only at the bisimilarity game.

References

[1] A. Baltag. A logic for coalgebraic simulation. In H. Reichel, editor, *Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS)*, volume 33 of *Electronic Notes in Theoretical Computer Science*, 2000.

[2] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.

[3] C. Hermida and B. Jacobs. Structural induction and coinduction in a fibrational setting. *Information and Computation*, 145(2):107–152, 1998.

[4] B. Jacobs and J. Rutten. A tutorial on (co)algebras and (co)induction. *EATCS Bulletin*, 62, 1997.

[5] A. Kurz. Coalgebras and modal logic. In *Course Notes for ESSLLI'2001*, 2001. Available at `http://www.cs.le.ac.uk/people/akurz/works.html`.

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

[7] L. S. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999.

[8] D. Pattinson. An introduction to the theory of coalgebras. In *North American Summer School in Logic, Language and Information (NASSLLI)*, 2003.

[9] J.J.M.M. Rutten. Automata and coinduction (an exercise in coalgebra). In D. Sangiorigi and R. de Simone, editors, *Proc. 9th International Conference on Concurrency Theory (CONCUR'98)*, volume 1466 of *LNCS*, pages 194–218. Springer, 1998.

[10] J.J.M.M. Rutten. Relators and Metric Bisimulation (Extended Abstract). *Electronic Notes in Theoretical Computer Science*, 11:1–7, 1998.

[11] J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.

[12] J.J.M.M. Rutten and E.P. de Vink. Bisimulation for probabilistic transition systems: a coalgebraic approach. *Theoretical Computer Science*, 221(1–2):271–293, 1999.

[13] Y. Venema. Algebras and coalgebras. In *Handbook of Modal Logic*, volume 3, pages 331–426. Elsevier, 2006.

[14] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.

[15] J. Worrell. On the final sequence of a finitary set functor. *Theoretical Computer Science*, 338(1–3):184–199, 2005.