

Definability, Canonical Models, Compactness for Finitary Coalgebraic Modal Logic

Alexander Kurz¹

CWI
P.O. Box 94079, 1090 GB Amsterdam
The Netherlands

Dirk Pattinson²

Ludwig-Maximilians-Universität München
Oettingenstr. 67, 80538 München
Germany

Abstract

This paper studies coalgebras from the perspective of the *finitary* observations that can be made of their behaviours. Based on the terminal sequence, notions of finitary behaviours and finitary predicates are introduced. A category $\text{Beh}_\omega(T)$ of coalgebras with morphisms preserving finitary behaviours is defined. We then investigate definability and compactness for finitary coalgebraic modal logic, show that the final object in $\text{Beh}_\omega(T)$ generalises the notion of a canonical model in modal logic, and study the topology induced on a coalgebra by the finitary part of the terminal sequence.

Introduction

Coalgebras for an endofunctor T on Set encompass many types of state base systems, including Kripke models and frames, labelled transition systems, Moore- and Mealy automata and deterministic systems, see e.g. Rutten [22]. The research on modal logics as specification languages for coalgebras began with Moss [15] and was taken up in e.g. [14, 21, 20, 8, 9].

The relationship between modal logic and coalgebras has been explained in [12] as follows. Denoting by Z the carrier of the final coalgebra, we can consider as the semantics of a modal formula φ the subset $\llbracket \varphi \rrbracket \subseteq Z$ satisfying φ .

The intuition here is: The elements of Z are the behaviours and the meaning of a modal formula φ is the property (i.e. set) of behaviours defined by φ . In case that the logics we are interested in are fully expressive in the sense that they allow to

¹ Email: kurz@cwi.nl

² Email: pattinson@informatik.uni-muenchen.de

define *all* subsets of Z , we can identify modal formulae and subsets of Z , resulting in an approach to algebraically investigate modal logics, see [12,13].

Unfortunately, modal logics given by a finitary syntax are in general not fully expressive. The reason for this is simply that not all properties of behaviours can be described in a finitary language. One of the main topics of this paper is the quest for a semantics of modal logic that suits *finitary* modal logic as perfectly as the semantics sketched in the previous paragraph does for fully expressive modal logics.

The key concept used to account for finitariness is the so called *terminal sequence* ($T^n 1$) of the endofunctor T . The terminal sequence ($T^n 1$) can be understood as approximating the final coalgebra, see [1]. Intuitively, the elements of the n -th approximant represent the behaviours that can be observed of a system in n steps. Following [16,17], the semantics of a finitary modal formula φ of rank n will be a subset $\llbracket \varphi \rrbracket \subseteq T^n 1$.³

In case that the functor T maps finite sets to finite sets, this approach results in a perfect match between finitary modal logics and their semantics: For all finite ordinals n and all subsets of $T^n 1$ we can assume a formula defining this subset. In case that T maps finite sets to infinite sets, which is for example the case for logics having infinitely many atomic propositions, there will be subsets of $T^n 1$ which are not defined by a single finitary formula. This is one of the reasons to introduce topologies τ_n on the approximants $T^n 1$ in Section 5, the idea being that finitary formulae correspond to clopen subsets.

The main novelty of the paper is probably, in Section 2, the introduction of the category $\mathbf{Beh}_\omega(T)$ that has coalgebras as objects and functions that preserve *finitary* behaviours as morphisms. One of the claims of this paper is that $\mathbf{Beh}_\omega(T)$ plays a role for finitary modal logics as $\mathbf{Coalg}(T)$ does for fully expressive modal logics. In Section 4, we show that $\mathbf{Beh}_\omega(T)$ always has a final object. Moreover, the final object in $\mathbf{Beh}_\omega(T)$ is compared to the canonical model as known from modal logic.

Another issue which arises when we focus on finitary logics is compactness. Compactness for coalgebraic modal logic is more complicated than for standard modal logic. In standard modal logic, compactness is inherited from the compactness of first-order logic by using van Benthem's translation of modal logic formulae to first-order formulae. The argument relies on the fact that, for an appropriate first-order language \mathcal{L} , the models of \mathcal{L} are precisely the Kripke models. It is therefore not a surprise that the argument fails if we change the class of models. For example, modal logic is *not compact* on finitely branching Kripke models. Since both Kripke models and finitely branching Kripke models appear naturally as categories of coalgebras, we cannot expect finitary modal logics for coalgebras to be compact in general. Assuming that T maps finite sets to finite sets and based on Sections 4 and 5, Section 6 characterises those functors T for which finitary modal logics for T -coalgebras are compact.

³ In this paper, we only investigate modal formulas that have a finite rank, that is, we exclude modalities speaking about an infinite number of transition steps as eg the \square^* of dynamic logic and the “always” of temporal logic.

1 Preliminaries and Notation

We consider coalgebras for a \mathbf{Set} -endofunctor T . The category of T -coalgebras and coalgebra morphisms is denoted by $\mathbf{Coalg}(T)$. The *final T -coalgebra*, denoted by $Z = (Z, \zeta)$, is – if it exists – defined up to isomorphism by the property that for all $A \in \mathbf{Coalg}(T)$ there is a unique morphism $!_A : A \rightarrow Z$. Given an element a in A , we call (A, a) a *process* and $!_A(a)$ its *behaviour*. Two processes (A, a) and (B, b) are *behaviourally equivalent* (written $(A, a) \sim (B, b)$), if they can be identified by a morphism of coalgebras, i.e. if there exists $(C, c) \in \mathbf{Coalg}(T)$ and $f : A \rightarrow C$, $g : B \rightarrow C \in \mathbf{Coalg}(T)$ such that $f(a) = g(b)$.

If the final coalgebra exists, this is clearly equivalent to $!_A(a) = !_B(b)$.

Example 1.1 [Streams] For a set D consider $TX = D \times X$. Given a coalgebra $\alpha = \langle \text{head}, \text{tail} \rangle : A \rightarrow D \times A$ the (complete) behaviour of an element $a \in A$ is the infinite list $(\text{head}(a), \text{head}(\text{tail}(a)), \text{head}(\text{tail}(\text{tail}(a))), \dots)$. Accordingly, the final coalgebra $Z = (D^{\mathbb{N}}, \langle \text{head}, \text{tail} \rangle)$ is given by the infinite lists over D .

Example 1.2 [Kripke models] Let \mathbf{Prop} be a countably infinite set. Coalgebras for the functor $T(X) = \mathcal{P}X \times \mathcal{P}\mathbf{Prop}$ are Kripke models. The coalgebraic notion of behavioural equivalence coincides with the standard notion of bisimulation in modal logic.

We have seen how the final coalgebra (if it exists) provides a notion of behaviour. In general, however, the behaviour of a process represents an infinite amount of information. This paper investigates properties of processes, which can be specified by a finite amount of information. Hence the final coalgebra (containing the infinite behaviours of all processes) has to be replaced by finitary approximations. These approximations are provided by the (finitary part) of the so-called terminal sequence of the underlying endofunctor T .

1.1 The Terminal Sequence

Terminal sequences can be thought of as approximating the final coalgebra. The following definition has been taken from [24].

The *terminal sequence* of T is an ordinal indexed sequence of sets (Z_α) together with a family $(p_\beta^\alpha)_{\beta < \alpha}$ of functions $p_\beta^\alpha : Z_\alpha \rightarrow Z_\beta$ for all ordinals $\beta \leq \alpha$ such that

- $Z_{\alpha+1} = TZ_\alpha$ and $p_{\beta+1}^{\alpha+1} = Tp_\beta^\alpha$ for all $\beta \leq \alpha$
- $p_\alpha^\alpha = id_{Z_\alpha}$ and $p_\gamma^\alpha = p_\gamma^\beta \circ p_\beta^\alpha$ for $\gamma \leq \beta \leq \alpha$
- The cone $(Z_\alpha, (p_\beta^\alpha))_{\beta < \alpha}$ is limiting whenever α is a limit ordinal.

Thinking of Z_α as the α -fold application of T to the limit 1 of the empty diagram, we write $Z_\alpha = T^\alpha 1$ in the sequel. Intuitively, $T^n 1$ represents behaviours which can be exhibited in n steps. For example, if $TX = D \times X$, then $T^n 1 \cong D^n$ contains all lists of length n .

Note that every coalgebra (C, γ) gives rise to a cone $(C, (\gamma_\alpha : C \rightarrow T^\alpha 1))$ over the terminal sequence:

Definition 1.3 Given $(C, \gamma) \in \mathbf{Coalg}(T)$, let $\gamma_0 = !_C : C \rightarrow 1$ denote the unique

mapping. For successor ordinals $\alpha = \beta + 1$ let $\gamma_\alpha : C \rightarrow T^\alpha 1 = T\gamma_\beta \circ \gamma$. If α is a limit ordinal, let γ_α be the unique map for which $\gamma_\beta = p_\beta^\alpha \circ \gamma_\alpha$ for all $\beta < \alpha$.

We will often use without further mentioning the following easy

Proposition 1.4 *Let n be an ordinal.*

- (i) *Let $f : (A, \alpha) \rightarrow (B, \beta)$ be a coalgebra morphism. Then $\beta_n \circ f = \alpha_n$.*
- (ii) *Let $(A, \alpha) \in \text{Coalg}(T)$. Then $p_n^{n+1} \circ T(\alpha_n) \circ \alpha = \alpha_n$.*

1.2 Coalgebraic Modal Logic

We explain how to extract a modal language from a given functor $T : \text{Set} \rightarrow \text{Set}$. Following [16,17], we consider modal logics for coalgebras given by predicate liftings.

Definition 1.5 *A predicate lifting for T is a natural transformation $\lambda : 2 \rightarrow 2 \circ T$, where 2 denotes the contravariant powerset functor.*

Example 1.6 Consider $TX = \mathcal{P}(X) \times \mathcal{P}(A)$, where A is a set of atomic propositions. Then T -coalgebras are in 1-1 correspondence with Kripke models over the set A of atoms. We demonstrate how to capture the interpretation of atoms and modalities using predicate liftings. Let $a \in A$ and consider the liftings λ and λ_a given by

$$\begin{aligned}\lambda(X)(\mathfrak{x}) &= \{(\mathfrak{x}', \mathfrak{a}) \in \mathcal{P}(X) \times \mathcal{P}(A) \mid \mathfrak{x}' \subseteq \mathfrak{x}\} \\ \lambda_a(X)(\mathfrak{x}) &= \{(\mathfrak{x}', \mathfrak{a}) \in \mathcal{P}(X) \times \mathcal{P}(A) \mid a \in \mathfrak{a}\}.\end{aligned}$$

Given $(C, \gamma) \in \text{Coalg}(T)$, we write $c \rightarrow c'$ if $c, c' \in C$ and $c' \in \pi_1 \circ \gamma(c)$. Also, if $c \in C$ and $a \in A$, we write $c \models a$ iff $a \in \pi_2 \circ \gamma(c)$.

For the case of modalities, consider a subset $\mathfrak{c} \subseteq C$, which we think of as the interpretation $\mathfrak{c} = \llbracket \varphi \rrbracket$ of a modal formula φ . Then

$$\gamma^{-1} \circ \lambda(C)(\mathfrak{c}) = \{c \in C \mid \forall c' \in C. c \rightarrow c' \implies c' \in \mathfrak{c}\}$$

corresponding to the interpretation of $\Box \varphi$. The same formula with λ replaced by λ_a yields

$$\gamma^{-1} \circ \lambda_a(C)(\mathfrak{c}) = \{c \in C \mid c \models a\}.$$

Hence, given any $\mathfrak{c} \subseteq C$, we can capture the set of worlds which satisfy a using the lifting λ_a . \square

Definition 1.7 (Syntax and semantics of $\mathcal{L}(T, \Lambda)$) *Given a set Λ of predicate liftings for T , we consider the language $\mathcal{L}(T, \Lambda)$, often abbreviated to $\mathcal{L}(\Lambda)$, which is given by the grammar*

$$\varphi ::= \text{ff} \mid \varphi \rightarrow \psi \mid [\lambda]\varphi \quad (\lambda \in \Lambda)$$

For a structure $(C, \gamma) \in \text{Coalg}(T)$, the semantics $\llbracket \varphi \rrbracket_\gamma \subseteq C$ is given by

$$\llbracket \text{ff} \rrbracket_\gamma = \emptyset \quad \llbracket \varphi \rightarrow \psi \rrbracket_\gamma = (C \setminus \llbracket \varphi \rrbracket_\gamma) \cup \llbracket \psi \rrbracket_\gamma \quad \llbracket [\lambda]\varphi \rrbracket_\gamma = \gamma^{-1} \circ \lambda(C)(\llbracket \varphi \rrbracket_\gamma)$$

Notation 1.8 (The modal logic \mathcal{ML}) In case of $T = \mathcal{P} \times \mathcal{P}\text{Prop}$, Prop countably infinite, we denote $\mathcal{L}(T, \Lambda)$ with Λ as in the example above by \mathcal{ML} . \mathcal{ML} is the standard modal logic for Kripke models.

In the remainder of this section we show that each formula $\varphi \in \mathcal{L}(\Lambda)$ gives rise to a predicate $t \subseteq T^n 1$ for some $n \in \mathbb{N}$. Given any structure $(C, \gamma) \in \text{Coalg}(T)$, we show that $\llbracket \varphi \rrbracket_\gamma = \gamma_n^{-1}(t)$, where γ_n is defined as in Definition 1.3.

Definition 1.9 (Formulas of rank n) Let L_0 denote the set of formulae of propositional logic, that is, L_0 is given by the grammar

$$L_0 \ni \varphi, \psi ::= \text{ff} \mid \varphi \rightarrow \psi$$

The mapping $d_0 : L_0 \rightarrow \mathcal{P}(1)$ is given by $d_0(\varphi) = 1$ iff φ is a tautology, and $d_0(\varphi) = \emptyset$, otherwise. For $n \geq 0$ the set L_{n+1} is given by the grammar

$$L_{n+1} \ni \varphi, \psi ::= \text{ff} \mid \varphi \rightarrow \psi \mid [\lambda]\varrho \quad (\lambda \in \Lambda, \varrho \in L_n)$$

The mapping $d_{n+1} : L_{n+1} \rightarrow \mathcal{P}(T^{n+1} 1)$ is given inductively by

$$\text{ff} \mapsto \emptyset \quad \varphi \rightarrow \psi \mapsto (T^{n+1} 1 \setminus d_{n+1}(\varphi)) \cup d_{n+1}(\psi) \quad [\lambda]\varrho \mapsto \lambda(T^n 1)(d_n(\varrho))$$

We call formulae in L_n formulae of rank n . For $\varphi \in L_n$, we also write $\llbracket \varphi \rrbracket_n$ or simply $\llbracket \varphi \rrbracket$ instead of $d_n(\varphi)$.

Example 1.10 In case of \mathcal{ML} , an equivalent definition of the rank of a formula is the following: $\text{rank}(\text{ff}) = 0$, $\text{rank}(\varphi \rightarrow \psi) = \max\{\text{rank}(\varphi), \text{depth}(\psi)\}$, $\text{rank}(p) = 1$ for $p \in \text{Prop}$, $\text{rank}(\Box\varphi) = \text{rank}(\varphi) + 1$. This is a slight variation of the standard definition in modal logic (where $\text{rank}(p) = 0$) but can be used to the same effect.

Intuitively, a formula $\varphi \in L_n$ describes behaviour which take n transition steps into account. When thinking of the sets $T^n 1$ as representing the behaviour which can be exhibited in n transition steps, the value $d_n(\varphi)$ is a model-independent interpretation of φ . We note that every formula is eventually contained in one of the L_n 's:

Proposition 1.11 $\mathcal{L}(\Lambda) = \bigcup_{n \in \mathbb{N}} L_n$.

The following proposition supports the intuition that $d_n(\varphi)$ is a semantic model-independent representation of the φ , for $\varphi \in L_n$.

Proposition 1.12 Suppose $(C, \gamma) \in \text{Coalg}(T)$. Then

$$\llbracket \varphi \rrbracket_\gamma = \gamma_n^{-1} \circ d_n(\varphi)$$

for all $n \in \mathbb{N}$ and all $\varphi \in L_n$.

Proof By induction on n using the naturality of predicate liftings. \square

In order to obtain definability results, we have to assume that the logic under consideration is reasonably expressive. We deal with two notions of expressiveness: logics, which allow the denotation of every $t \subseteq T^n 1$ by a formula, and those which

allow to carve out every predicate $t \subseteq T^n 1$ by a *set* of formulae. The formal definition is as follows:

Definition 1.13 (Formula-(set-)expressive) *The language $\mathcal{L}(T, \Lambda)$ is called*

- formula-expressive, if every d_n is a surjection
- formula-set-expressive, if every D_n is a surjection

where d_n is as in Definition 1.9 and D_n is defined by

$$D_n : \mathcal{P}L_n \rightarrow \mathcal{P}T^n 1, \quad \Phi \mapsto \bigcap \{d_n(\varphi) : \varphi \in \Phi\}.$$

For $\mathcal{L}(T, \Lambda)$ to be expressive, we have to put a completeness condition on the set of predicate liftings.

Definition 1.14 (Separation) (i) Suppose C is a set and $\mathcal{C} \subseteq \mathcal{P}(C)$ is a system of subsets of C . We call \mathcal{C} separating, if the map $s : C \rightarrow \mathcal{P}(\mathcal{C})$, $s(c) = \{\mathbf{c} \in \mathcal{C} \mid c \in \mathbf{c}\}$ is monic.

(ii) A set Λ of predicate liftings for C is called separating, if $\{\lambda(C)(\mathbf{c}) \mid \lambda \in \Lambda, \mathbf{c} \subseteq C\}$ is a separating set of subsets of TC , for all sets C .

The idea of separation is that the individual points of C can be distinguished by the predicates $P \in \mathcal{C}$. Passing to predicate liftings, we can distinguish individual successors $x \in TX$ by means of predicate liftings, assuming that all points of X can be distinguished.

The separation property is present in many examples, notably also in Example 1.6.

Example 1.15 Let $TX = \mathcal{P}(X) \times \mathcal{P}(A)$ for some set A of atomic propositions. Consider $\Lambda = \{\lambda\} \cup \{\lambda_a \mid a \in A\}$ as defined in Example 1.6. Then Λ is separating.

Assuming that T maps finite sets to finite sets, which corresponds to the fact of having only finitely many propositional variables in the case of Kripke models, one easily establishes

Proposition 1.16 Suppose T maps finite sets to finite sets and Λ is separating. Then $\mathcal{L}(T, \Lambda)$ is formula-expressive.

Given a structure $(C, \gamma) \in \mathbf{Coalg}(T)$, the above proposition says that every predicate $\mathbf{c} \subseteq C$ on C , which arises as $\gamma_n^{-1}(t)$ for some $t \subseteq T^n 1$ can actually be denoted by a formula.

In the case of Kripke models with a countably infinite set of propositional variables we have

Proposition 1.17 Let $T = \mathcal{P} \times \mathcal{P}\mathbf{Prop}$, \mathbf{Prop} countably infinite, and $\mathcal{L}(T, \Lambda) = \mathcal{ML}$. Then $\mathcal{L}(T, \Lambda)$ is formula-set-expressive.

Given a structure $(C, \gamma) \in \mathbf{Coalg}(\mathcal{P} \times \mathcal{P}\mathbf{Prop})$, the above proposition says that every predicate $\mathbf{c} \subseteq C$ on C , which arises as $\gamma_n^{-1}(t)$ for some $t \subseteq T^n 1$ can be denoted by a set of formulae.

2 Finitary Predicates and the Category $\text{Beh}_\omega(T)$

Behavioural predicates, i.e. predicates on coalgebras which are invariant under observational equivalence, can be considered as subsets of the carrier of the final coalgebra. Here, we are interested in finitary behavioural predicates and we propose to consider them as subsets of the $T^n 1$, $n < \omega$, given by the terminal sequence.

First, recalling Definition 1.3, we define a notion of n -step behavioural equivalence and of predicates of rank n .

Definition 2.1 (n -Behavioural equivalence) Let n be an ordinal. For two coalgebras $\mathbf{A} = (A, \alpha)$, $\mathbf{B} = (B, \beta)$ define $(\mathbf{A}, a) \sim_n (\mathbf{B}, b)$ iff $\alpha_n(a) = \beta_n(b)$. Similarly, $\mathbf{A} \sim_n \mathbf{B}$ iff $\alpha_n(A) = \beta_n(B)$ and $\mathbf{A} \sim_{<\omega} \mathbf{B}$ iff $\alpha_n(A) = \beta_n(B)$ for all $n < \omega$.

Under the assumption that the final coalgebra exists, we consider two points x and y as behaviourally equivalent, if they are identified by the unique morphism into the final coalgebra. As shown in [1], this is equivalent to $\alpha_n(x) = \alpha_n(y)$ for all ordinals n . The notion of finitary behavioural predicates, which we are about to introduce, restricts the validity of the above equation to *finite* ordinals.

Remark 2.2 While $(\mathbf{A}, a) \sim_\omega (\mathbf{B}, b) \Leftrightarrow \forall n < \omega . (\mathbf{A}, a) \sim_n (\mathbf{B}, b)$, we only have $\mathbf{A} \sim_\omega \mathbf{B} \Rightarrow \forall n < \omega . \mathbf{A} \sim_n \mathbf{B}$. For an example refuting the converse, let $TX = \{a, b\} \times X$, \mathbf{A} the final coalgebra with carrier $\{a, b\}^\omega$ and \mathbf{B} the subcoalgebra with carrier $\{s \cdot a^\omega : s \in \{a, b\}^*\}$.

Example 2.3 For $TX = \mathcal{P}X \times \mathcal{P}\text{Prop}$, n -behavioural equivalence is (a slight variation of) the bounded bisimulation of modal logic as studied in [5].

Definition 2.4 (Behavioural predicates of rank n) A set $S \subseteq T^n 1$ is called a behavioural predicate of rank n . A process (A, α, a) satisfies S , written $(A, \alpha, a) \models S$ and often abbreviated as $a \models S$, iff $\alpha_n(a) \in S$.

We also use standard notation such as $(A, \alpha) \models S \Leftrightarrow \forall a \in A . a \models S$ and $\llbracket S \rrbracket_{(A, \alpha)} = \{a \in A : a \models S\}$ and $\text{Mod}(S) = \{\mathbf{A} \in \text{Coalg}(T) : \mathbf{A} \models S\}$.

Example 2.5 Let $\varphi \in \mathcal{L}(T, \Lambda)$ be a formula of rank n . Then the semantics of φ is determined by the predicate $\llbracket \varphi \rrbracket \subseteq T^n 1$ (cf. Definition 1.9). If T maps finite sets to finite sets and Λ is separating, every predicate $S \subseteq T^n 1$ is denoted by a formula (cf. Proposition 1.16). In case of \mathcal{ML} every predicate is denoted by a set of formulae (cf. Proposition 1.17).

Remark 2.6 In [12, 11] it was proposed to investigate modal logics by considering subsets of the final coalgebra as the semantics of modal formulae. From this perspective, the approach presented here is a special case. Let (Z, ζ) be the final coalgebra. Then every $S \subseteq T^n 1$ is logically equivalent to $\zeta_n^{-1}(S) \subseteq Z$. Indeed, $(\mathbf{A}, a) \models S \Leftrightarrow \alpha_n(a) \in S \Leftrightarrow \zeta_n(!_A(a)) \in S \Leftrightarrow !_A(a) \in \zeta_n^{-1}(S)$ which was the definition of satisfaction for modal formulae as subsets of the final coalgebra. This will be used in the next section.

Predicates of rank n and n -behavioural equivalence are related by the following propositions.

Proposition 2.7 Suppose $\mathcal{L}(T, \Lambda)$ is formula-expressive. Then, for all A and a in A , there exists $\varphi_{(A,a)} \in \mathcal{L}(T, \Lambda)$ such that $(A, a) \sim_n (B, b) \Leftrightarrow (B, b) \models \varphi_{(A,a)}$.

The following is a result well-known in modal logic (cf. [5], Proposition 2.8).

Proposition 2.8 Suppose $\mathcal{L}(T, \Lambda)$ is formula-set-expressive and consider $A, B \in \text{Coalg}(T)$ and elements a in A , b in B . Then

$$(A, a) \sim_n (B, b) \Leftrightarrow (a \models \varphi \Leftrightarrow b \models \varphi \text{ for all } \varphi \in \mathcal{L}(T, \Lambda) \text{ of rank } n).$$

We now define the category $\text{Beh}_\omega(T)$ of coalgebras that has as morphisms those functions which preserve finitary behaviours.

Definition 2.9 ($\text{Beh}_\omega(T)$) The category $\text{Beh}_\omega(T)$ has T -coalgebras as objects. Morphisms $f : (A, \alpha) \rightarrow (B, \beta)$ are those functions $f : A \rightarrow B$ such that, for all $n < \omega$, $\beta_n \circ f = \alpha_n$.

Remark 2.10 For $f : A \rightarrow B$ each of the following is equivalent to f being a morphism $f : (A, \alpha) \rightarrow (B, \beta)$ in $\text{Beh}_\omega(T)$

$$\begin{aligned} \beta_\omega \circ f &= \alpha_\omega \\ \forall n < \omega . \forall S \subseteq T^n 1 . \forall a \in A . f(a) \models S &\Leftrightarrow a \models S \end{aligned}$$

Remark 2.11 Clearly, every morphism of coalgebras $f : (A, \alpha) \rightarrow (B, \beta) \in \text{Coalg}(T)$ is also a morphism $f \in \text{Beh}_\omega(T)$. We hence obtain a functorial inclusion $\text{Coalg}(T) \rightarrow \text{Beh}_\omega(T)$. There are two distinct reasons why $\text{Beh}_\omega(T)$ contains more morphisms than $\text{Coalg}(T)$. The first is that $\text{Beh}_\omega(T)$ -morphisms take only finitary behaviours into account, the second is that $\text{Beh}_\omega(T)$ -morphisms only preserve behavioural properties which do not involve colourings.

Remark 2.12 In order to explain the relationship of $\text{Beh}_\omega(T)$ to $\text{Coalg}(T)$ consider the following categories

$$\begin{array}{ccc} \text{c-Beh}(T) & \longrightarrow & \text{Beh}(T) \\ \downarrow & & \downarrow \\ \text{c-Beh}_\omega(T) & \hookrightarrow & \text{Beh}_\omega(T) \end{array}$$

which all have coalgebras as objects and morphisms as follows. $f : (A, \alpha) \rightarrow (B, \beta)$ is a $\text{Beh}(T)$ -morphism iff $\alpha_n(a) = \beta_n(f(a))$ for all ordinals n . The definitions of $\text{c-Beh}(T)$ and $\text{c-Beh}_\omega(T)$ follow the same idea, but take colourings into account: $f : (A, \alpha) \rightarrow (B, \beta)$ is a $\text{c-Beh}_\omega(T)$ -morphism iff f is a $\text{Beh}_\omega(T \times C)$ -morphism $(A, \langle \alpha, v \circ f \rangle) \rightarrow (B, \langle \beta, v \rangle)$ for all $C \in \text{Set}$ and $v : B \rightarrow C$.

If $\text{Coalg}(T)$ has cofree coalgebras then $\text{c-Beh}(T) = \text{Coalg}(T)$. If T is finitary (ie. ω -accessible) then $\text{Beh}_\omega(T) = \text{Beh}(T)$ and $\text{c-Beh}_\omega(T) = \text{c-Beh}(T)$. Whether the converse holds, that is, whether $\text{c-Beh}_\omega(T) = \text{c-Beh}(T)$ implies that T is finitary is an open question. \square

One of the claims of this paper is that studying finitary modal logics for coalgebras, it is profitable to transfer techniques and ideas known from $\text{Coalg}(T)$ to $\text{Beh}_\omega(T)$. For example, in Section 4 we will show that $\text{Beh}_\omega(T)$ always has a final object.

3 Definability Results

For the purpose of definability, we assume the existence of a final coalgebra $Z = (Z, \zeta)$. We first note the following easy proposition relating predicates over $T^n 1$ to predicates over Z .

- Proposition 3.1** (i) *For any $S \subseteq T^n 1$ it holds $A \models S$ iff $A \models \zeta_n^{-1}(S)$.*
(ii) *For any $S \subseteq Z$ it holds $A \models S$ only if $A \models \zeta_n(S)$.*
(iii) $\forall S \subseteq Z . S \subseteq \zeta_n^{-1}(\zeta_n(S))$ and $\forall S \subseteq T^n 1 . \zeta_n(\zeta_n^{-1}(S)) = S$.

We are now able to prove

Theorem 3.2 *A class $\mathcal{B} \subseteq \text{Coalg}(T)$ is definable by a subset $S \subseteq T^n 1$ iff \mathcal{B} is closed under images, coproducts, domains of morphisms, and \sim_n .*

Proof ‘only if’: Closure under images, coproducts, domains of morphisms is standard (and easy to check) and closure under \sim_n is immediate from the definitions.

For ‘if’ let $S = (S, \sigma)$ be the coalgebra given by the union of the images of all $!_B : B \rightarrow Z$, $B \in \mathcal{B}$. We show that $\mathcal{B} = \text{Mod}(\zeta_n(S))$. For $B \in \mathcal{B}$ we have, by definition of S , $B \models S$, hence $B \models \zeta_n(S)$ by Proposition 3.1.2. To show $\mathcal{B} \supset \text{Mod}(\zeta_n(S))$, define $\bar{S} = (\bar{S}, \bar{\sigma})$ as the largest subcoalgebra of $\zeta_n^{-1}(\zeta_n(S))$. It follows from Proposition 3.1.3 that $\zeta_n(S) = \zeta_n(\bar{S})$, hence $S \sim_n \bar{S}$. Since \mathcal{B} is closed under images and coproducts, \mathcal{B} is also closed under unions, hence $S \in \mathcal{B}$. Since \mathcal{B} is closed under \sim_n , $\bar{S} \in \mathcal{B}$. Now assume $A \models \zeta_n(S)$. By Proposition 3.1.1, $_A !$ factors through $\zeta_n^{-1}(\zeta(S))$ and hence through \bar{S} , i.e. there is a morphism $A \rightarrow \bar{S}$. Since \mathcal{B} is closed under domains of morphisms, $A \in \mathcal{B}$. \square

The following corollary is an immediate consequence.

- Corollary 3.3** (i) *Suppose $\mathcal{L}(T, \Lambda)$ is formula-expressive. Then a class $\mathcal{B} \subseteq \text{Coalg}(T)$ is definable by a formula of rank n iff \mathcal{B} is closed under images, coproducts, domains of morphisms, and \sim_n .*
(ii) *If $\mathcal{L}(T, \Lambda)$ is formula-set-expressive, then a class $\mathcal{B} \subseteq \text{Coalg}(T)$ is definable by a set of formulae of rank n iff \mathcal{B} is closed under images, coproducts, domains of morphisms, and \sim_n .*

A similar proof gives an expressiveness result for sets of finitary formulae.

Theorem 3.4 *A class \mathcal{B} of coalgebras is definable by a set \mathcal{S} with $S \in \mathcal{S} \Rightarrow \exists n (S \subseteq T^n 1 \ \& \ n < \omega)$ iff \mathcal{B} is closed under images, coproducts, domains of morphisms, and $\sim_{<\omega}$.*

Corollary 3.5 *If $\mathcal{L}(T, \Lambda)$ is formula-set-expressive, then a class $\mathcal{B} \subseteq \text{Coalg}(T)$ is definable by a set of formulae iff \mathcal{B} is closed under images, coproducts, domains of morphisms, and $\sim_{<\omega}$.*

For sufficient conditions ensuring formula-expressiveness and formula-set expressiveness, see Propositions 1.16 and 1.17.

4 A Canonical Model Construction for Coalgebras

Reasoning about behaviours, the final coalgebra plays a central role because, given the unique coalgebra morphism $!_A : A \rightarrow Z$ from a coalgebra A into the final coalgebra Z , for every element a of A , we can consider $!_A(a)$ as the behaviour of a . Similarly, coalgebras final in $\text{Beh}_\omega(T)$ (cf. Definition 2.9) consist of the finite behaviours. We first show that final coalgebras exist in $\text{Beh}_\omega(T)$ and then show how they generalise the canonical model construction known from modal logic.

4.1 Coalgebras Final in $\text{Beh}_\omega(T)$

A coalgebra final in $\text{Beh}_\omega(T)$ should “realise” precisely all n -behaviours, $n < \omega$. Accordingly, the carrier of a final object in $\text{Beh}_\omega(T)$ will be a subset of $T^\omega 1$.

Recall that, given any structure (C, γ) , we write γ_ω for the unique mediating map $\gamma_\omega : C \rightarrow T^\omega 1$. That is, all ω -behaviours appear as some $\gamma_\omega(c)$ in $T^\omega 1$. On the other hand, not every point $t \in T^\omega 1$ can be presented as $t = \gamma_\omega(c)$ by some structure (C, γ) and some $c \in C$. Consider for example the finite powerset functor $T = \mathcal{P}_\omega$. Worrell [24] shows, that for the final T -coalgebra (Z, ζ) the morphism $\zeta_\omega : Z \rightarrow T^\omega 1$ is not surjective.

Hence we construct the carrier of the coalgebra final in $\text{Beh}_\omega(T)$ as consisting of all $t \in T^\omega 1$, which can be “realised” by some structure, i.e. for which there are $(C, \gamma) \in \text{Coalg}(T)$ and $c \in C$ such that $\gamma_\omega(c) = t$. It then remains to find an appropriate coalgebra structure on the carrier.

Throughout, we fix the set K of “realisable” elements $t \in T^\omega 1$, which is given by

$$K = \{t \in T^\omega 1 \mid \exists(C, \gamma) \in \text{Coalg}(T) . \exists c \in C . \gamma_\omega(c) = t\}.$$

For each $k \in K$, we can now choose $(C^k, \gamma^k) \in \text{Coalg}(T)$ and $c^k \in C^k$ such that $\gamma_\omega^k(c^k) = k$. Note that K is a set, which enables us to consider

$$(C, \gamma) = \coprod_{k \in K} (C^k, \gamma^k)$$

where the coproduct is taken in $\text{Coalg}(T)$. Denoting the coproduct injections by $in_k : C_k \rightarrow C$ (which, by the construction of coproducts in $\text{Coalg}(T)$ are also coproduct injections in the category of sets), we are ready to note:

Lemma 4.1 $\gamma_\omega \circ in_k(c) = \gamma_\omega^k(c)$ for all $k \in K$ and $c \in C_k$.

Proof Since γ_ω^k is the unique mediating map into the limiting cone with vertex $T^\omega 1$, it suffices to prove that $\gamma_n \circ in_k(c) = \gamma_n^k(c)$ for all $n \in \mathbb{N}$. For $n = 0$, this is obvious. For the induction step we calculate $\gamma_{n+1} \circ in_k(c) = T\gamma_n \circ \gamma \circ in_k(c) = T\gamma_n \circ Tin_k \circ \gamma^k(c) = T\gamma_n^k \circ \gamma^k(c) = \gamma_{n+1}^k$ establishing the claim. \square

We obtain the following immediate corollary:

Corollary 4.2 For all $k \in K$ there exists $c \in C$ with $\gamma_\omega(c) = k$.

In other words, γ_ω factors through K as $\gamma_\omega = m \circ e$, m injective, e surjective. Now consider the diagram

$$\begin{array}{ccccc} TT^\omega 1 & \xleftarrow{Tm} & TK & \xleftarrow{Te} & TC \\ & \downarrow \kappa & & & \uparrow \gamma \\ T^\omega 1 & \xleftarrow{m} & K & \xleftarrow{e} & C \\ & & \searrow o & & \end{array}$$

where o is any one-sided inverse of e , i.e. $e \circ o = id_K$, the existence of which is guaranteed by e being a surjection. We let

$$\kappa = Te \circ \gamma \circ o.$$

Note that $\kappa : K \rightarrow TK$ makes K into a T -coalgebra. Denoting the limit projections by $p_n^\omega : T^\omega 1 \rightarrow T^n 1$, we obtain

Lemma 4.3 *For all $n \in \mathbb{N}$, $\kappa_n = p_n^\omega \circ m$, hence $m = \kappa_\omega$.*

Proof We proceed by induction on n , where the case $n = 0$ is evident. We calculate $\kappa_{n+1} = T\kappa_n \circ \kappa = T(p_n^\omega \circ m) \circ Te \circ \gamma \circ o = Tp_n^\omega \circ T(m \circ e) \circ \gamma \circ o = Tp_n^\omega \circ T\gamma_\omega \circ \gamma \circ o = T\gamma_n \circ \gamma \circ o = \gamma_{n+1} \circ o = p_{n+1}^\omega \circ \gamma_\omega \circ o = p_{n+1}^\omega \circ m \circ e \circ o = p_{n+1}^\omega \circ m$ for the induction step, as desired. \square

The proof of the main theorem of this section is now straightforward.

Theorem 4.4 $\text{Beh}_\omega(T)$ has a final object.

Proof We show that (K, κ) , as constructed above, is final in $\text{Beh}_\omega(T)$. Take any object $(D, \delta) \in \text{Beh}_\omega(T)$. Consider the mapping $\delta_\omega : D \rightarrow T^\omega 1$, which is the unique mediating map between the cones $(D, (\delta_n)_{n \in \mathbb{N}})$ and $(T^\omega 1, (p_n^\omega)_{n \in \mathbb{N}})$. By construction, δ_ω factors as $\delta_\omega = m \circ h$ where $m : K \rightarrow T^\omega 1$ is as above. By Lemma 4.3

$$\delta_\omega = \kappa_\omega \circ h,$$

which implies that h is a $\text{Beh}_\omega(T)$ -morphism. h is unique since κ_ω is injective. \square

Note that a coalgebra final in $\text{Beh}_\omega(T)$ is not determined uniquely up to isomorphisms in $\text{Coalg}(T)$. In case that $p_\omega^{\omega+1} : TT^\omega 1 \rightarrow T^\omega 1$ is surjective⁴, the coalgebras final in $\text{Beh}_\omega(T)$ are precisely those which are given by right-inverses of $p_\omega^{\omega+1}$.

Corollary 4.5 *Assume that $p_\omega^{\omega+1}$ is surjective. Then a coalgebra is final in $\text{Beh}_\omega(T)$ iff it is isomorphic in $\text{Coalg}(T)$ to some $(T^\omega 1, \tau)$ with $p^{\omega+1} \circ \tau = id_{T^\omega 1}$*

Proof ‘if’: To show that $(T^\omega 1, \tau)$ is final, it suffices to observe that $\tau_\omega = id_{T^\omega 1}$. This follows from $\tau_n = p_n^\omega$, $n < \omega$, the inductive case being $\tau_{n+1} = T(\tau_n) \circ \tau = T(p_n^\omega) \circ \tau = p_{n+1}^\omega \circ p_n^{\omega+1} \circ \tau = p_{n+1}^\omega$.

‘only if’: Let (A, α) be final in $\text{Beh}_\omega(T)$. Consider a final object (K, κ) as constructed in the proof of the theorem. Let $f : (A, \alpha) \rightarrow (K, \kappa)$ be the unique morphism. In particular, f is iso and $\kappa_\omega \circ f = \alpha_\omega$. Since κ_ω is injective, α_ω is as well. Since, by Proposition 1.4(ii), $\alpha_\omega = p_\omega^{\omega+1} \circ T(\alpha_\omega) \circ \alpha$, α_ω is also surjective, hence iso. Now define $\tau = T(\alpha_\omega) \circ \alpha \circ \alpha_\omega^{-1}$. \square

⁴ Which is the case for all examples in this paper with the exception of $T = \mathcal{P}_\omega$. A sufficient condition for $p_\omega^{\omega+1}$ to be surjective is that T weakly preserves limits of ω^{op} -chains.

4.2 The Canonical Model

Let M be the functor $\mathcal{P} \times \mathcal{P}\text{Prop}$, Prop a countably infinite set.

The **canonical model** (see for example [3,6]) for the modal logic \mathcal{ML} is the M -coalgebra $(L, \langle \lambda_R, \lambda_V \rangle)$

$$\begin{array}{ll} L & \{\Phi \subseteq \mathcal{ML} : \Phi \text{ is maximally consistent}\} \\ \lambda_R : L \rightarrow \mathcal{P}L & \Phi \mapsto \{\Psi : \psi \in \Psi \Rightarrow \Diamond\psi \in \Phi\} \\ \lambda_V : L \rightarrow \mathcal{P}\text{Prop} & \Phi \mapsto \Phi \cap \text{Prop} \end{array}$$

The canonical model is final in the category $\text{Th}_{\mathcal{ML}}$ which has M -coalgebras as objects and morphisms $f : (A, \alpha) \rightarrow (B, \beta)$ are functions $f : A \rightarrow B$ such that for all $a \in A$, a and $f(a)$ have the same modal theory.

From a coalgebraic viewpoint, finitary formulae correspond to subsets of $T^n 1$. Taking inverse images along the limit projections, each set of finitary formulae can be understood as a subset of $T^\omega 1$. Intuitively, maximally consistent subsets then correspond to minimal (i.e. singleton) subsets $\{t\} \subseteq T^\omega 1$ of $T^\omega 1$. We make this precise by showing that the categories $\text{Beh}_\omega(M)$ and $\text{Th}_{\mathcal{ML}}$ are actually identical.

Proposition 4.6 $\text{Beh}_\omega(M) = \text{Th}_{\mathcal{ML}}$.

Proof We have to show that for any coalgebras (A, α) , (B, β) and any function $f : A \rightarrow B$,

$$\beta_\omega \circ f(a) = \alpha_\omega(a) \Leftrightarrow \text{Th}(a) = \text{Th}(f(a)),$$

which is equivalent to $[\forall n < \omega . \beta_n \circ f(a) = \alpha_n(a)] \Leftrightarrow [\forall n < \omega . \forall \varphi \in \mathcal{ML} . \text{rank}(\varphi) = n \Rightarrow (a \models \varphi \Leftrightarrow f(a) \models \varphi)]$ which in turn is a consequence of Proposition 2.8. \square

Since the projection $p_\omega^{\omega+1} : MM^\omega 1 \rightarrow M^\omega 1$ is surjective, we know by Corollary 4.5 that all coalgebras (K, κ) final in $\text{Beh}_\omega(M)$ are given—up to coalgebra isomorphisms—by $K = M^\omega 1$ and a one-sided inverse κ of $p_\omega^{\omega+1}$. Since we have just shown that the canonical model is final in $\text{Beh}_\omega(M)$, it is indeed one of the (K, κ) constructed in the previous subsection.

5 The Topology of Finite Observations

We study the topology on a coalgebra induced by the terminal sequence and relate logical and topological properties.

Definition 5.1 (Cantor space topology) Suppose τ_n is a topology on $T^n 1$ for all finite ordinals. If $(A, \alpha) \in \text{Coalg}(T)$, then the topology τ_A induced by $(\tau_n)_{n < \omega}$ is the topology generated by the base

$$\{\alpha_n^{-1}(o) \mid n \in \omega, o \in \tau_n\}$$

of open sets. If all τ_n are discrete, we call τ_A the Cantor space topology.

Clearly, τ_A makes all projections α_n continuous. Viewing the Cantor set as the final coalgebra for $TX = 2 \times X$, the topology induced by the discrete $\tau_n = \mathcal{P}(T^n 1)$, one recovers the cantor discontinuum.

Example 5.2 Suppose $TX = 2 \times X$, where $2 = \{0, 1\}$. Consider the (final) T -coalgebra (A, α) with $A = 2^{\mathbb{N}} = \{f : \mathbb{N} \rightarrow 2\}$ and $\alpha(f) = (f(0), \lambda n \cdot f(n+1))$. Then (A, τ_A) is homeomorphic to the Cantor discontinuum C (also known as middle-third set, see e.g. [10]) via the mapping $2^{\mathbb{N}} \rightarrow C$, $f \mapsto \sum_{i=0}^{\infty} \frac{2}{3^{i+1}} \cdot f(i)$.

Remark 5.3

- (i) Suppose $f : (A, \alpha) \rightarrow (B, \beta)$ is a homomorphism of coalgebras. Then f is continuous w.r.t. the topologies on A and B . Thus the passage from a coalgebra (A, α) to the topological space (A, τ_A) induces a functor $\text{Coalg}(T) \rightarrow \text{Top}$.
- (ii) Let $(A, \alpha) \in \text{Coalg}(T)$ and let, for $a_0, a_1 \in A$, $d_A(a_0, a_1) = \inf\{2^{-n} : \forall k < n \cdot \alpha_k(a_0) = \alpha_k(a_1)\}$. Then d_A is a pseudo-ultrametric on A . d_A is a ultrametric if $\alpha_\omega : A \rightarrow T^\omega 1$ is injective. The Cantor space topology τ_A coincides with the topology induced by d_A , as studied in [2, 25].

The topologies τ_n on $T^n 1$ of interest to us are those given by a basis of ‘finitely observable properties’

$$\mathcal{B}_n = \{\llbracket \varphi \rrbracket_n : \varphi \in \mathcal{L}, \varphi \text{ a formula of rank } n\},$$

for some logic \mathcal{L} . To make precise the assumptions on \mathcal{L} that are needed in the following we make the

Convention 5.4 (Logic \mathcal{L} for T , induced topologies τ_n) *Given a functor T , a logic \mathcal{L} for T consists of sets of formulae $\mathcal{L}_n, n < \omega$, equipped with functions $\llbracket \cdot \rrbracket_n : \mathcal{L}_n \rightarrow \mathcal{P}(T^n 1)$ which assign to each formula of rank n a predicate of rank n . The semantics of \mathcal{L} is determined by the terminal sequence as in Definition 2.4. Moreover, we assume that each \mathcal{L}_n is closed under boolean operators which are interpreted in the usual way on $\mathcal{P}(T^n 1)$ and that the projections $p_m^n : T^n 1 \rightarrow T^m 1$ give rise to functions $(p_m^n)^{-1} : \mathcal{B}_m \rightarrow \mathcal{B}_n$ where*

$$\mathcal{B}_n = \{\llbracket \varphi \rrbracket_n : \varphi \in \mathcal{L}_n\}.$$

We denote by τ_n the topology given by the basic opens \mathcal{B}_n .

Note that the convention ensures that $(T^n 1, \tau_n)_{n < \omega}$ is a sequence of topological spaces and that each space has a basis of clopens, *the basic clopens being precisely the subsets definable by single formulae*. In most examples, the following stronger condition also holds (which is trivially satisfied if all the τ_n are discrete and the $T^n 1$ finite).

Condition 1 *In addition to Convention 5.4 require that the topologies τ_n are compact and Hausdorff.*⁵

Condition 1 ensures that each $(T^n 1, \tau_n)$ is Hausdorff, compact, and has a basis of clopens, that is, each $(T^n 1, \tau_n)$ is a Stone space.

Example 5.5 (i) In case that \mathcal{L} is $\mathcal{L}(T, \Lambda)$ with T mapping finite sets to finite sets and a separating set of predicate liftings Λ , Condition 1 is satisfied.

⁵ A set is *compact* iff any open cover has a finite subcover. This is sometimes called quasi-compact. A space (X, τ) is *Hausdorff* iff $\forall x, y \in X \cdot x \neq y \Rightarrow \exists U, V \in \tau \cdot x \in U \wedge y \in V \wedge U \cap V = \emptyset$.

- (ii) In case that \mathcal{L} is \mathcal{ML} , Condition 1 is satisfied. We skip the proof and only mention that compactness of the τ_n can be deduced from the compactness of \mathcal{ML} , similarly to the argument of ‘only if’ in the proof of Proposition 5.11.

We give some topological characterisations of logical properties.

Proposition 5.6 *Let $(A, \alpha) \in \text{Coalg}(T)$ and \mathcal{L} a logic for T . Then a subset of A is definable by a set of formulae iff it is closed w.r.t. τ_A .*

Proof A subset $S \subseteq A$ is closed iff there are basic opens $O_i \subseteq T^{n_i}1$, $i \in I$, such that $S = \bigcap\{\alpha_{n_i}^{-1}(T^{n_i}1 \setminus O_i) : i \in I\}$. \square

Proposition 5.7 *Suppose $(A, \alpha) \in \text{Coalg}(T)$ and (A, τ_A) compact and \mathcal{L} a logic for T . Then a subset $S \subseteq A$ is definable by a single formula in \mathcal{L} iff S is clopen.*

Proof We claim that a subset $S \subseteq A$ is clopen iff $S = \alpha_n^{-1}(O)$ for some $n \in \mathbb{N}$ and some basic clopen $O \subseteq T^n1$.

It follows from the α_n being continuous that all subsets of the form $\alpha_n^{-1}(O)$ are clopen if O clopen. Now suppose $S \subseteq A$ is clopen. Then $S = \bigcup\{\alpha_{n_i}^{-1}(O_i) : i \in I\}$ and $A \setminus S = \bigcup\{\alpha_{n_j}^{-1}(P_j) : j \in J\}$ for basic clopens $O_i \subseteq T^{n_i}1$, $P_j \subseteq T^{n_j}1$. Since A is compact, the (disjoint) union $\bigcup\{\alpha_{n_i}^{-1}(O_i) : i \in I\} \cup \bigcup\{\alpha_{n_j}^{-1}(P_j) : j \in J\}$ has a finite subcover, yielding a finite set $I' \subseteq I$ of indices such that $S = \bigcup\{\alpha_{n_i}^{-1}(O_i) : i \in I'\}$. Let $m = \max\{n_i : i \in I'\}$ and consider the projections $p_{n_i}^m : T^m1 \rightarrow T^{n_i}1$. Defining $O = \bigcup\{(p_{n_i}^m)^{-1}(O_i) : i \in I'\}$ establishes the claim. \square

Compactness is a property which is unfortunately not present in all models.

Example 5.8 Let $TX = D \times X$ and consider the final coalgebra (Z, ζ) given by $Z = D^\omega$.

- (i) (Z, ζ) is compact in the Cantor space topology iff D is finite.
- (ii) Suppose $D = \{a, b\}$. Then examples of non-compact coalgebras are given by the carriers $Z \setminus \{b^\omega\}$ and $\{s \cdot a^\omega : s \in \{a, b\}^*\}$ (and inheriting the structure from ζ).

Example 5.9 Let $TX = \{a, b\} \times X + 1$ and consider the final coalgebra (Z, ζ) with $Z = \{a, b\}^* \cup \{a, b\}^\omega$. Then Z is compact in the Cantor space topology (since the limit of compact Hausdorff spaces is compact Hausdorff, see [4], 3.2.13) and $\{a, b\}^*$ is not compact. The topology on Z is as follows. A subset of Z is open iff it is a subset of $\{a, b\}^*$ or of the form $V \cdot (\{a, b\}^* + \{a, b\}^\omega)$ for some $V \subseteq \{a, b\}^*$. In particular, every open cover of $\{a, b\}^\omega$ also covers $\{a, b\}^*$.

Example 5.10 Let $T = \mathcal{P}_\omega$. Then the final coalgebra is not compact.

Given a logic \mathcal{L} for T , we call a T -coalgebra (A, α) **logically compact**, if every set, which is finitely satisfiable in (A, α) (that is, for every finite subset $\Phi' \subseteq \Phi$ there exists $a \in A$ such that $a \models \Phi'$) is satisfiable in (A, α) (i.e. there exists $a \in A$ such that $a \models \Phi$). We are now ready to prove

Proposition 5.11 *Let $(A, \alpha) \in \text{Coalg}(T)$ and \mathcal{L} a logic for T . Then (A, α) is logically compact iff (A, α) is compact.*

Proof We use that A is compact iff every system $\mathcal{S} \subseteq \mathcal{P}(A)$ of closed subsets, which has the finite intersection property⁶, has non-empty intersection.

Assume that (A, α) is logically compact and that $\mathcal{S} \subseteq \mathcal{P}(A)$ is a system of closed sets having the finite intersection property. Every set $S \in \mathcal{S}$ is definable by some $\Phi_S \subseteq \mathcal{L}$ by Proposition 5.6. It follows from \mathcal{S} having the finite intersection property that $\bigcup\{\Phi_S \mid S \in \mathcal{S}\}$ is finitely satisfiable, thus satisfiable by logical compactness. That is, there exists $a \in A$ such that $a \models \bigcup\{\Phi_S \mid S \in \mathcal{S}\}$ which implies $a \in \bigcap \mathcal{S}$ by construction.

Now assume (A, α) is topologically compact and consider a set $\Phi \subseteq \mathcal{L}$ which is finitely satisfiable. Since $\llbracket \varphi \rrbracket \subseteq A$ is closed by Proposition 5.6, the set $\{\llbracket \varphi \rrbracket : \varphi \in \Phi\}$ is a system of closed sets having the finite intersection property. By compactness of (A, α) , there exists $a \in \bigcap\{\llbracket \varphi \rrbracket \mid \varphi \in \Phi\}$, which is equivalent to $a \models \Phi$. \square

Example 5.12 We can now easily verify Examples 5.8.2 and 5.10. For instance, in case of $T = \mathcal{P}_\omega$, it is an easy exercise to write down formulae φ_n which force any point satisfying φ_n to have at least n successors. The set $\Phi = \{\varphi_n \mid n \in \mathbb{N}\}$ is then finitely satisfiable, but not satisfiable by a \mathcal{P}_ω -coalgebra.

6 Compactness for Coalgebraic Modal Logic

In (standard) modal logic, compactness is not an issue, since it is inherited from the corresponding result in first order logic via van Benthem's standard translation [23,3]. Generalising to coalgebraic modal logic, the standard translation is no longer available. We therefore have to resort to different means in order to establish a compactness theorem. Moreover, compactness fails in the general case, for example in case of image-finite Kripke models (i.e. $T = \mathcal{P}_\omega$, cf. Examples 5.10 and 5.12).

Hence we are drawn to investigate sufficient and necessary conditions for the compactness theorem to hold. Building upon the work of Sections 5 and 4, we obtain that validity of the compactness theorem is equivalent to the existence of a *compact* canonical model. We then characterise those endofunctors T for which $\text{Beh}_\omega(T)$ has a final object compact in the Cantor space topology as those endofunctors which weakly preserve the limit of the chain $(T^n 1)_{n \in \mathbb{N}}$.

We say that a set $\Phi \subseteq \mathcal{L}(\Lambda)$ is *satisfiable*, if there exists a T -coalgebra (A, α) and $c \in C$ such that $c \models_\gamma \varphi$ for all $\varphi \in \Phi$. We call Φ *finitely satisfiable*, if every finite subset of Φ is satisfiable. Using this terminology, we are in the position to present the first version of the compactness theorem (recall Convention 5.4).

Theorem 6.1 *A logic \mathcal{L} for T -coalgebras is compact iff $\text{Beh}_\omega(T)$ has a compact final object.*

Proof ‘only if’: By Theorem 4.4 there exists a final object $(K, \kappa) \in \text{Beh}_\omega(T)$. We show that (K, κ) is logically compact, from which the result then follows by Proposition 5.11. So suppose $\Phi \subseteq \mathcal{L}$ is finitely satisfiable in (K, κ) . By compactness Φ is satisfiable. Thus there is (C, γ) and $c \in C$ such that $c \models_\gamma \Phi$. Since (K, κ) is final in $\text{Beh}_\omega(T)$, there is a mapping $u : (C, \gamma) \rightarrow (K, \kappa) \in \text{Beh}_\omega(T)$. By definition of morphisms in $\text{Beh}_\omega(T)$, we obtain $u(c) \models_\kappa \Phi$. Hence Φ is satisfiable in (K, κ) .

⁶ \mathcal{S} has the *finite intersection property* iff $\bigcap \mathcal{S}'$ is non-empty for all finite $\mathcal{S}' \subseteq \mathcal{S}$.

‘if’: Let (K, κ) be compact and final in $\text{Beh}_\omega(T)$ and suppose $\Phi \subseteq \mathcal{L}$ is finitely satisfiable. Then – by finality and by definition of morphisms in $\text{Beh}_\omega(T)$ – Φ is finitely satisfiable in (K, κ) , hence satisfiable in (K, κ) by compactness and Proposition 5.11. \square

We now proceed to characterise those endofunctors T for which $\text{Beh}_\omega(T)$ has a compact final object. It turns out that $\text{Beh}_\omega(T)$ has a compact final object iff T weakly preserves the limit of its final sequence up to ω . More precisely, consider the limiting cone $T^\omega 1$ of the sequence

$$1 \xleftarrow{p_0^1} T1 \xleftarrow{p_1^2} T^2 1 \xleftarrow{p_2^3} T^3 1 \quad \dots \quad T^\omega 1$$

with associated projections $p_n^\omega : T^\omega 1 \rightarrow T^n 1$, we say that T *weakly preserves* the limit of the sequence $(T^n 1)_{n \in \mathbb{N}}$, if the cone $(TT^\omega 1, (Tp_n^\omega)_{n \in \mathbb{N}})$ is weakly limiting.

We begin by noting that every element of an ‘approximant’ $T^n 1$ can be realised by a coalgebra.

Proposition 6.2 *Let m be any mapping $1 \rightarrow T1$ and $(C^n, \gamma^n) = (T^n 1, T^n m)$. Then $\gamma_n^n = id_{C^n}$.*

We now show that the carrier of a compact final object in $\text{Beh}_\omega(T)$ is isomorphic to $T^\omega 1$. This is the crucial step in our proof.

Lemma 6.3 *Suppose (K, κ) is compact and final in $\text{Beh}_\omega(T)$ and $u : K \rightarrow T^\omega 1$ is the unique mediating morphism between the cones $(K, (\kappa_n)_{n \in \mathbb{N}})$ and $(T^\omega 1, (p_n^\omega)_{n \in \mathbb{N}})$. Then u is iso.*

Proof It follows from the construction in Section 4 that u is mono. To see that u is epi, it suffices to show that for all $t \in T^\omega 1$ there exists $k \in K$ such that $p_n^\omega(t) = \kappa_n(k)$. Fix some $t \in T^\omega 1$ and let $u^n : (C^n, \gamma^n) \rightarrow (K, \kappa)$ denote the unique morphism into the final object (with (C^n, γ^n) as in Proposition 6.2).

Define a sequence $(k_n)_{n \in \mathbb{N}}$ by $k_n = u^n \circ p_n^\omega(t)$. It follows that $\kappa_n(k_n) = p_n^\omega(t)$ for all $n \in \mathbb{N}$. Note that (K, κ) is actually a metric space (cf. Remark 5.3), and that $(k_n)_{n \in \mathbb{N}}$ is a Cauchy-sequence. Since compact metric spaces are complete, it follows that $k = \lim k_n$ exists in K . Observing $\kappa_n(k) = \kappa_n \circ u^n \circ p_n^\omega(t) = p_n^\omega(t)$ finishes the proof. \square

Observing that T weakly preserves the limit of $(T^n 1)_{n \in \mathbb{N}}$ iff $p_\omega^{\omega+1}$ has a one-sided inverse i , $p_\omega^{\omega+1} \circ i = id_{T^\omega 1}$, we are now able to prove the following two propositions.

Proposition 6.4 *Suppose that $\text{Beh}_\omega(T)$ has a final model that is compact w.r.t. the Cantor space topology. Then T weakly preserves the limit of $(T^n 1)_{n \in \mathbb{N}}$.*

Proof Let (K, κ) be final and compact in $\text{Beh}_\omega(T)$ and $u : K \rightarrow T^\omega 1$ be the unique mediating morphism between the cones $(K, (\kappa_n)_{n \in \mathbb{N}})$ and $(T^\omega 1, (p_n^\omega)_{n \in \mathbb{N}})$. Due to the lemma above, we can define $i = Tu \circ \kappa \circ u^{-1}$. It remains to check that indeed $p_\omega^{\omega+1} \circ i = p_\omega^{\omega+1} \circ Tu \circ \kappa \circ u^{-1} = u \circ u^{-1} = id_{T^\omega 1}$. \square

Proposition 6.5 *Suppose the topologies τ_n on $T^n 1$, $n < \omega$, are compact Hausdorff. Then T weakly preserves the limit of $(T^n 1)_{n \in \mathbb{N}}$ only if the final model of $\text{Beh}_\omega(T)$ is compact in the induced topology.*

Proof Let $p_\omega^{\omega+1} \circ i = id_{T^\omega 1}$. It was shown in Corollary 4.5 that $(T^\omega 1, i)$ is final in $\text{Beh}_\omega(T)$. It is compact since $T^\omega 1$ is the limit of compact Hausdorff spaces and the induced topology on a limit of compact Hausdorff spaces is compact Hausdorff (see [4] 3.2.13). \square

We can summarise:

Theorem 6.6 *Let T map finite sets to finite sets. Then $\text{Beh}_\omega(T)$ has a final object that is compact in the Cantor space topology iff T weakly preserves the limit of $(T^n 1)_{n \in \mathbb{N}}$.*

7 Conclusions and Related Work

We have studied definability and compactness for finitary coalgebraic modal logic. The main instrument through which finitary logics have been studied is the terminal sequence and – based on the terminal sequence – the shift from the category $\text{Coalg}(T)$ to the category $\text{Beh}_\omega(T)$.

In this category, points (or states) can be distinguished iff their *finite* behaviour differs. Also, $\text{Beh}_\omega(T)$ provides the right framework in which the construction of canonical models can be generalised to a coalgebraic setting. We are not aware of any work characterising canonical models as final in a suitable category.

The main handle which allows to formalise the finitary character of the logics considered is to identify finitary predicates with subsets of $T^n 1$, where n is a finite ordinal. The idea of interpreting formulae on the elements $T^n 1$ of the terminal sequence was already used in [16]. The same idea (without the restriction to finite ordinals) also prevails in Moss [15]. There, formulae are constructed using infinitary conjunctions (which do not change the degree of the formulae) and the application of the signature functor T (increasing the degree of the constructed formulae by 1).

The signature functors (and hence the logics) which have been discussed in the present paper are all one-sorted. The standard passage to multi-sorted signatures, i.e. endofunctors $\text{Set}^n \rightarrow \text{Set}^n$ is standard and allows to include the logics discussed in [7, 19], which also rely on (syntactically defined) predicate liftings. Since the endofunctors discussed in loc. cit. are all ω -accessible, final coalgebras and canonical models coincide for these logics (which is also reflected by the fact that they are strong enough to characterise behavioural equivalence).

A coalgebraic representation of the Cantor discontinuum has also been given in [18] in the category of posets. The cantor space topology discussed in the present paper arises in a different way: We start with a final coalgebra on the category of sets, which is then equipped with a natural topology.

Acknowledgements

We would like to thank the referees for useful comments, some of which could not have been incorporated in the present version but had to be saved for future work. Also thanks to Clemens Kupke for the careful reading of a previous draft.

References

- [1] J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theoretical Computer Science*, 150:57–75, 1995.
- [2] M. Barr. Terminal coalgebras in well-founded set theory. *Theoretical Computer Science*, 114(2):299–315, June 1993.
- [3] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001. See also <http://www.mlbook.org>.
- [4] Ryszard Engelking. *General Topology*. Heldermann Verlag, 1989.
- [5] Jelle Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, 1999.
- [6] Robert Goldblatt. *Logics of Time and Computation*, volume 7 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford University, 1992. Second Edition.
- [7] B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoret. Informatics and Applications*, 35(1):31–59, 2001.
- [8] Bart Jacobs. The temporal logic of coalgebras via galois algebras. Technical Report CSI-R9906, Computing Science Institute Nijmegen, 1999.
- [9] Bart Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoretical Informatics and Applications*, 35(1):31–59, 2001.
- [10] John Kelley. *General Topology*. Springer, 1955.
- [11] Alexander Kurz. *Logics for Coalgebras and Applications to Computer Science*. PhD thesis, Ludwig-Maximilians-Universität München, 2000. <http://www.informatik.uni-muenchen.de/~kurz>.
- [12] Alexander Kurz. A co-variety-theorem for modal logic. In *Advances in Modal Logic 2*. Center for the Study of Language and Information, Stanford University, 2001. Selected Papers from AiML 2, Uppsala, 1998.
- [13] Alexander Kurz. Modal rules are co-implications. In A. Corradini, M. Lenisa, and U. Montanari, editors, *Coalgebraic Methods in Computer Science (CMCS'01)*, volume 44.1 of *ENTCS*. Elsevier, 2001.
- [14] Alexander Kurz. Specifying coalgebras with modal logic. *Theoretical Computer Science*, 260:119–138, 2001.
- [15] Lawrence Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999.
- [16] Dirk Pattinson. Semantical principles in the modal logic of coalgebras. In *Proceedings 18th International Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, volume 2010 of *LNCS*, Berlin, 2001. Springer. Also available as technical report at <http://www.informatik.uni-muenchen.de/~pattinso/>.
- [17] Dirk Pattinson. Coalgebraic Modal Logic: Soundness, Completeness and Decidability. Technical report, Institut für Informatik, LMU München, 2002.
- [18] Dusko Pavlovic and Vaughan Pratt. On coalgebra of real numbers. In B. Jacobs and J. Rutten, editors, *Electronic Notes in Theoretical Computer Science*, volume 19. Elsevier Science Publishers, 2000.
- [19] M. Rößiger. Coalgebras and Modal Logic. In H. Reichel, editor, *Coalgebraic Methods in Computer Science (CMCS'2000)*, volume 33, 2000.
- [20] Martin Rößiger. Coalgebras and modal logic. In Horst Reichel, editor, *Coalgebraic Methods in Computer Science (CMCS'00)*, volume 33 of *Electronic Notes in Theoretical Computer Science*, pages 299–320, 2000.
- [21] Martin Rößiger. From modal logic to terminal coalgebras. *Theoretical Computer Science*, 260:209–228, 2001.
- [22] J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000. First appeared as technical report CS R 9652, CWI, Amsterdam, 1996.
- [23] Johan van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Naples, 1983.
- [24] James Worrell. Terminal sequences for accessible endofunctors. In B. Jacobs and J. Rutten, editors, *Coalgebraic Methods in Computer Science (CMCS'99)*, volume 19 of *Electronic Notes in Theoretical Computer Science*, pages 39–53, 1999.
- [25] James Worrell. *On Coalgebras and Final Semantics*. PhD thesis, Oxford University Computing Laboratory, 2000.