

Coalgebraic Lindström Theorems

Alexander Kurz

University of Leicester, UK

Yde Venema

ILLC, Amsterdam, The Netherlands

Abstract

We study modal Lindström theorems from a coalgebraic perspective. We provide three different Lindström theorems for coalgebraic logic, one of which is a direct generalisation of de Rijke’s result for Kripke models. Both the other two results are based on the properties of bisimulation invariance, compactness, and a third property: ω -bisimilarity, and expressive closure at level ω , respectively. These also provide new results in the case of Kripke models. Discussing the relation between our work and a recent result by van Benthem, we give an example showing that only requiring bisimulation invariance together with compactness does not suffice to characterise basic modal logic.

Keywords: coalgebra, modal logic, Lindström theorem

1 Introduction

Lindström’s theorem [19,12] states that every ‘abstract logic’ extending first-order logic and satisfying Löwenheim-Skolem and compactness is equivalent to first-order logic. The notion of abstract logic is a technical one capturing some fundamental properties one expects from the set \mathcal{L} of sentences of any legitimate logic. Thus, roughly speaking, Lindström’s theorem says that first-order logic is the strongest logic satisfying Löwenheim-Skolem and compactness.

In modal logic, de Rijke’s theorem [23,7] states that every ‘abstract modal logic’ extending basic modal logic and having finite depth (or finite degree) is equivalent to basic modal logic. Later, van Benthem [5] showed that the finite depth condition can be replaced by compactness and relativisation. In collaboration with ten Cate and Väänänen, van Benthem expanded on this result in [6], where Lindström theorems are discussed for various fragments of first-order logic, including for instance graded modal logic. Inspired by this work, Otto and Piro [21] establish a Lindström type characterisation of the extension of basic modal logic by a global modality, and of the guarded fragment of first-order logic, with the corresponding bisimulation invariance, and the Tarski Union Property replacing the closure under relativisation.

Our approach to Lindström-type theorems for modal logic is coalgebraic. The notion of a coalgebra $X \rightarrow TX$ for a functor T encompasses, for particular instantiations of T , Kripke frames and models, but also many other structures of importance in computer science. Accordingly, we generalise de Rijke’s theorem to a wide range of functors T . But our analysis also gives an elegant explanation of de Rijke’s theorem and leads to new variations in the case of Kripke frames and models.

There are two main differences between our approach and that of van Benthem et al [5,6]. In the mentioned papers, abstract modal logics are *not* supposed to be bisimulation invariant, and, second, they *are* supposed to be closed under relativisation. Neither of these assumptions is very natural in the coalgebraic setting, and therefore, our starting point will be rather different from that of the cited authors. Similarly, we do not see how to translate the Tarski Union Property of Otto and Piro [21] to a coalgebraic setting.

The technique used to obtain our results is based on the final coalgebra sequence. It can be explained as follows.

$$1 \xleftarrow{p_0^1} T1 \xleftarrow{p_1^2} T^2 1 \xleftarrow{\dots} T^\alpha 1 \xleftarrow{\dots} Z \quad (1)$$

An element of Z (the final coalgebra) is an equivalence class of pointed models up to bisimilarity. Similarly, for each ordinal α , there is a notion of α -bisimilarity and the ‘approximants’ $T^\alpha 1$ classify pointed models up to α -bisimilarity. The ‘projections’ p identify elements that cannot be distinguished at a coarser level.

From our perspective, an abstract coalgebraic logic \mathcal{L} then will just be a collection of subsets of Z , see eg [11,16] for more on this perspective. \mathcal{L} will have finite depth if all these subsets are determined at a finite stage $\alpha < \omega$. Since it is known that basic modal logic (over finitely many variables) allows us to define all subsets of all finitary approximants, de Rijke’s theorem is an immediate consequence: An abstract coalgebraic logic extending basic modal logic and having a notion of finite depth must be equivalent to basic modal logic (because the formulas of such a logic, up to logical equivalence, correspond precisely to the subsets of the finitary approximants $T^n 1, n < \omega$).

In a further analysis, we investigate ways to replace the notion of finite depth by compactness. Whereas the argument in the previous paragraph was based on all subsets of all finitary approximants being definable in \mathcal{L} , we now look at higher approximants, such as $T^\omega 1$ (or Z itself) and consider the topologies generated by (extensions of) formulas. Then topological compactness and logical compactness coincide and we can combine topological results with properties of the final coalgebra sequence.

Summary and Structure of the paper. Section 7 presents our three Lindström theorems, stating that a logic invariant under bisimilarity, closed under Boolean operations and at least as expressive as the ‘basic modal logic’ \mathcal{L}_T^F is actually equivalent to \mathcal{L}_T^F if additionally one of three conditions hold:

- \mathcal{L} has finite depth,

- \mathcal{L} is compact and invariant under ω -bisimilarity,
- \mathcal{L} is compact and expressively closed at ω .

We also give an example showing that in the second item ω -bisimilarity cannot be replaced by bisimilarity.

Sections 3-5 are devoted to these three conditions, respectively. Section 6 presents two lemmas on the final coalgebra sequence, which we believe are of independent interest.

2 Preliminaries on Coalgebras and Modal Logic

2.1 Coalgebras

Coalgebras generalise Kripke frames and models.

Definition 2.1 The category $\text{Coalg}(T)$ of coalgebras for a functor T on a category \mathcal{X} has as objects arrows $\xi : X \rightarrow TX$ in \mathcal{X} and morphisms $f : (X, \xi) \rightarrow (X', \xi')$ are arrows $f : X \rightarrow X'$ such that $Tf \circ \xi = \xi' \circ f$.

The main examples of functors of interest to us in this paper are

Definition 2.2 A Kripke polynomial functor (KPF) $T : \text{Set} \rightarrow \text{Set}$ is built according to

$$T ::= Id \mid C \mid T + T \mid T \times T \mid T \circ T \mid \mathcal{P} \quad (2)$$

where Id is the identity functor, C is the constant functor that maps all sets to a finite set C , $+$ is disjoint union, \times is cartesian product, \circ is composition, and \mathcal{P} maps a set to the collection of its subsets.

Remark 2.3 • In the above definition, we insist on constants C being finite and coproducts being finite as well, hence a KPF in our sense maps finite sets to finite sets.

- A coalgebra $X \rightarrow \mathcal{P}X$ is a Kripke frame and a coalgebra $X \rightarrow 2^P \times \mathcal{P}X$ for a finite set P (of atomic propositions) is a Kripke model.
- An element x_0 of a coalgebra $\xi : X \rightarrow C \times X$ specifies a stream (infinite list) (c_1, c_2, \dots) via $(c_{n+1}, x_{n+1}) = \xi(x_n)$, $n < \omega$.
- Similarly, an element x_0 of a coalgebra $X \rightarrow C_1 + C_2 \times X \times X$ specifies a possibly infinite binary tree with leaves labelled from C_1 and the other nodes labelled from C_2 .
- Consider a coalgebra $X \rightarrow 2 \times X^C$ where $2 = \{0, 1\}$. It can be understood as a deterministic automaton over the alphabet C , where the elements in 2 are used to label states as accepting or non-accepting.
- The reader can easily extend this list, for example, non-deterministic automata are coalgebras $X \rightarrow 2 \times (\mathcal{P}X)^C$

Definition 2.4 We say that T has the properties (wp) and (fs) if T is a functor $\text{Set} \rightarrow \text{Set}$ and, respectively,

(wp) preserves weak pullbacks,

(fs) restricts to finite sets.

Example 2.5 A KPF T satisfies (wp) and (fs). Further examples of functors are obtained by extending (2) with

$$\dots | \mathcal{P}_\omega | \text{List} | \text{Mult} | \mathcal{D} | \mathcal{H} \quad (3)$$

where $\mathcal{P}_\omega X$ is the set of finite subsets of X , $\text{List}X$ the set of finite lists over X , $\text{Mult}X$ the set of finite multisets over X , $\mathcal{D}X$ the set of discrete probability distributions over X , and $\mathcal{H}X = 2^{2^X}$. They all satisfy (wp) with the exception of \mathcal{H} and they do not satisfy (fs) with the exception of \mathcal{P}_ω and \mathcal{H} . \mathcal{H} -coalgebras coincide with neighbourhood frames in modal logic and are investigated, from a coalgebraic point of view, in Hansen and Kupke [13].

Final Coalgebra and Bisimilarity. Given $T : \text{Set} \rightarrow \text{Set}$ the **final coalgebra** (Z, ζ) is determined by the property that for any coalgebra $\xi : X \rightarrow TX$ there is a unique coalgebra morphism $f : X \rightarrow Z$. Instantiating T as \mathcal{P} , we find that f identifies two elements of X iff they are bisimilar in the usual sense. So we may take this as a definition: Two elements of a coalgebra are **bisimilar** if they are identified by the unique morphism into the final coalgebra. This definition extends to elements in two different coalgebras $(X, \xi), (X', \xi')$ by considering their disjoint union (or coproduct) $X + X' \xrightarrow{\xi + \xi'} TX + TX' \rightarrow T(X + X')$. **Bisimilarity** is the smallest relation between elements of (possibly different) coalgebras containing all pairs of bisimilar elements. Clearly, bisimilarity is an equivalence relation.

Remark 2.6 (Existence of the final coalgebra) The final coalgebra does always exist if we allow its carrier to be a proper class as in [2] or if we assume the existence of an appropriate inaccessible cardinal as in [4]. Both approaches are equivalent, but we find the latter point of view more convenient as we then know that the final sequence, discussed below, always converges against the final coalgebra.

Example 2.7 (i) The final coalgebra $Z \rightarrow \mathcal{P}Z$ is Aczel's universe of non-well-founded sets [1]. The elements of Z can be understood as precisely the non-well-founded sets, that is, the equivalence classes of pointed Kripke frames under bisimilarity.

- (ii) The final coalgebra for $TX = C \times X$ is $\zeta : C^\omega \rightarrow C \times C^\omega$ where ζ maps a stream $(c_i)_{i < \omega}$ to the pair $(c_0, (c_i)_{1 \leq i < \omega})$ consisting of the 'head' and the 'tail' of the stream.
- (iii) Similarly, the final coalgebra for $TX = C_1 + C_2 \times X \times X$ consists of the set of all, possibly infinite, binary trees with leaves labelled from C_1 and the other nodes labelled from C_2 .
- (iv) In case of $TX = 2 \times X^C$, the final coalgebra Z is the set of all languages (ie subsets of C^*) and the unique coalgebra morphism $X \rightarrow Z$ maps a state of the automaton X to its accepted language [24].

Final Sequence and α -Bisimilarity. The final coalgebra is approximated by

the final coalgebra sequence

$$1 \xleftarrow{p_0^1} T1 \xleftarrow{p_1^2} T^2 1 \xleftarrow{\dots} T^\alpha 1 \xleftarrow{\dots} Z \quad (4)$$

where 1 denotes a one-element set and

- $T^{\alpha+1}1 = T(T^\alpha 1)$ and $p_{\beta+1}^{\alpha+1} = T(p_\beta^\alpha)$ for all $\beta < \alpha$,
- $p_\gamma^\alpha = p_\gamma^\beta \circ p_\beta^\alpha$ for $\gamma < \beta < \alpha$,
- if α is a limit ordinal then $(T^\alpha 1, (p_\beta^\alpha)_{\beta < \alpha})$ is a limiting cone.

Since the final coalgebra (Z, ζ) does exist [4] we know from [3, Thm 2] that the final sequence converges, that is, there is an isomorphism $p_\beta^{\beta+1} = \zeta^{-1}$. As in (4) we write Z for $T^\beta 1$ and p_α for p_α^β .

For any coalgebra $\xi : X \rightarrow TX$ there are maps $f_\alpha : X \rightarrow T^\alpha 1$

$$\begin{array}{ccccccc}
 1 & \xleftarrow{p_0^1} & T1 & \xleftarrow{p_1^2} & T^2 1 & \xleftarrow{\dots} & T^\alpha 1 & \xleftarrow{\dots} & Z \\
 & & & & & & \nearrow^{f_\alpha} & & \nearrow^f \\
 & & & & & & X & & \\
 & & \searrow_{f_0} & & & & & & \searrow_f
 \end{array} \quad (5)$$

with f_α given by $T^\alpha 1$ being a limit if α is a limit ordinal and

$$f_{\alpha+1} = T f_\alpha \circ \xi \quad (6)$$

otherwise. We say that $x, y \in X$ are **α -bisimilar** iff $f_\alpha(x) = f_\alpha(y)$. In the case of $T = \mathcal{P}$ this notion coincides with the notion of bounded bisimulation of Gerbrandy [9]. It was investigated from the point of view of coalgebraic logic in [18].

Remark 2.8 • The limit $T^\omega 1$ can be calculated explicitly as the set of sequences $\{t \in \prod_{n < \omega} T^n 1 \mid \forall m, n < \omega. p_n^m(t_m) = t_n\}$. For example, in the case of streams ($TX = C \times X$), we have that $T^n 1$ is C^n and $T^\omega 1$ is (isomorphic to) C^ω . The $p_n : C^\omega \rightarrow C^n$ are then the projections. In this case $T^\omega 1$ is also the final coalgebra and its elements classify bisimilarity.

- In general, the elements of $T^\omega 1$ classify behaviour up to ω -bisimilarity. In the case of Kripke models ($TX = 2^P \times \mathcal{P}X$), we have that $T^\omega 1$ is the canonical model of the modal logic **K**. Here, $T^\omega 1$ is not the final coalgebra (and this cannot be mended by replacing \mathcal{P} by the finite powerset).
- The final sequence for $TX = \mathcal{P}X$ has been studied eg in Ghilardi [10] and Worrell [26]. Elements can be seen as α -bisimilar classes of pointed Kripke frames.
- In case of $TX = 2 \times X^C$, the inductive definition (6) of the maps $f_n : X \rightarrow T^n 1$ provides the standard partition-refinement algorithm of minimising deterministic automata: Starting with a finite X , we are sure to find some $n < \omega$ such that the image of f_n is isomorphic to the image of f_{n+1} , which then is the minimisation of X . The coalgebraic view on this algorithm allows to generalise this, for example, to π -calculus processes [8].

2.2 Coalgebraic Logic

The aim of this section is just to give a snapshot of how modal logics for coalgebras can be set up parametric in the functor T . We cannot give full details here and it mainly serves to substantiate the notion of abstract coalgebraic logic below.

Example 2.9

- (i) We denote by \mathcal{L}_T^P the logic given by all n -ary predicate liftings, $n < \omega$, see [22,25]. An n -ary predicate lifting is a natural transformation $\lambda_X : 2^X \rightarrow 2^{TX}$. Each predicate lifting gives rise to a modal operator $[\lambda]$. It is interpreted (unary case) on a coalgebra (X, ξ) via

$$x \Vdash [\lambda]\varphi \Leftrightarrow \xi(x) \in \lambda_X(\{x \in X \mid x \Vdash \varphi\}).$$

Conversely, modal operators in the usual sense can be described by predicate liftings. For example, if $T = \mathcal{P}$, the usual modal \Box arises from the predicate lifting $(\lambda_\Box)_X(S) = \{S' \in \mathcal{P}X \mid S' \subseteq S\}$ and, similarly, $(\lambda_\Diamond)_X(S) = \{S' \in \mathcal{P}X \mid S' \cap S \neq \emptyset\}$.

- (ii) If T has (wp), we denote by \mathcal{L}_T^M the smallest logic closed under Boolean operations and the rule that $\nabla\psi \in \mathcal{L}_T^M$ whenever $\psi \in T(\Phi)$ for some finite $\Phi \subseteq \mathcal{L}_T^M$. We put $x \Vdash \nabla\psi \Leftrightarrow \exists w \in T(\Vdash). T(\pi_1)(w) = \xi(x) \ \& \ T(\pi_2)(w) = \psi$ where $\Vdash \subseteq X \times \mathcal{L}_T^M$ and π_1, π_2 are the two projections. This logic was introduced by Moss in [20], but see also [15] for the finitary version.

If T has (fs) and (wp) both logics are equally expressive, see [17] for explicit translations in both directions. For $T = \mathcal{P}$ both are equally expressive with basic modal logic (in the sense of \equiv of Definition 2.10).

2.3 Abstract coalgebraic logic

We take the point of view that any coalgebraic logic should be invariant under bisimilarity. That is, we can identify (the meaning of) a logical formula with its extension on the final coalgebra.

Definition 2.10

- An abstract coalgebraic logic for a functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ is given by a class of formulas \mathcal{L} and a function $\llbracket - \rrbracket : \mathcal{L} \rightarrow 2^Z$ to subsets of the carrier Z of the final coalgebra. We usually identify $\varphi \in \mathcal{L}$ with its extension $\llbracket \varphi \rrbracket$.
- We write $\llbracket \varphi \rrbracket_\alpha$ for the direct image $p_\alpha[\varphi]$ where p_α is as in (5).
- A subset $S \subseteq T^\alpha 1$ is definable in \mathcal{L} iff there is $\varphi \in \mathcal{L}$ such that $\llbracket \varphi \rrbracket_\alpha = S$ and $\varphi = p_\alpha^{-1}(\llbracket \varphi \rrbracket_\alpha)$.
- $\mathcal{L}_1 \leq \mathcal{L}_2$ iff for all $\varphi_1 \in \mathcal{L}_1$ there is $\varphi_2 \in \mathcal{L}_2$ such that $\llbracket \varphi_1 \rrbracket = \llbracket \varphi_2 \rrbracket$. $\mathcal{L}_1 \equiv \mathcal{L}_2$ means $\mathcal{L}_1 \leq \mathcal{L}_2$ and $\mathcal{L}_2 \leq \mathcal{L}_1$.
- Given a coalgebra $\xi : X \rightarrow TX$ and $\varphi \in \mathcal{L}$, we define

$$x \Vdash \varphi \Leftrightarrow f(x) \in \llbracket \varphi \rrbracket \tag{7}$$

where f is as in (5).

Remark 2.11 (Invariance under bisimilarity) It is immediate from (7) that any abstract coalgebraic logic is invariant under bisimilarity (for the notion of bisimilarity defined on page 4). Conversely, if a logic for coalgebras is invariant under bisimilarity then (7) holds. To summarise, invariance under bisimulation means precisely that formulas can be identified with their extensions on the final coalgebra (Z, ζ) .

In particular, we see that the logics of Example 2.9 become abstract coalgebraic logics simply by putting $\llbracket \varphi \rrbracket = \{z \in Z \mid z \Vdash \varphi\}$. The definitions of \Vdash from Example 2.9 then agree with (7).

Example 2.12 We denote by \mathcal{L}_T^F the abstract coalgebraic logic in which precisely all subsets of $T^n 1$, $n < \omega$, are definable, that is, $\mathcal{L}_T^F = \bigcup_{n < \omega} \{p_n^{-1}(S) \mid S \subseteq T^n 1\}$.

For a detailed comparison of $\mathcal{L}_T^M, \mathcal{L}_T^P, \mathcal{L}_T^F$ we refer to [17], here we only record the following.

Proposition 2.13 *We have*

$$\mathcal{L}_T^M \leq \mathcal{L}_T^P \leq \mathcal{L}_T^F,$$

and $\mathcal{L}_T^P \equiv \mathcal{L}_T^F$ if T has (fs) and $\mathcal{L}_T^M \equiv \mathcal{L}_T^P$ if T has (fs) and (wp).

Definition 2.14 \mathcal{L} has the properties (Sep), (fSep), (Full), (fFull), (Bool) if \mathcal{L} is an abstract coalgebraic logic and, respectively,

- (Sep) any two distinct behaviours in the final coalgebra are separated by two formulas, that is, for all $z \neq z' \in Z$ there are $\varphi, \varphi' \in \mathcal{L}$ such that $\varphi \cap \varphi' = \emptyset$ and $z \in \varphi$ and $z' \in \varphi'$,
- (fSep) for all $n < \omega$, any two behaviours in $T^n 1$ are separated by formulas, that is, for all $t \neq t' \in T^n 1$ there is $n < \omega$ and there are definable (see Definition 2.10) $S, S' \subseteq T^n 1$ such that $S \cap S' = \emptyset$ and $p_n^\omega(t) \in S$ and $p_n^\omega(t') \in S'$,
- (Full) for ordinals α , any subset of $T^\alpha 1$ is definable,
- (fFull) for all $n < \omega$, any subset of $T^n 1$ is definable,
- (Bool) \mathcal{L} is closed under Boolean operations.

(fSep) holds iff any two behaviours in $T^\omega 1$ can be separated by a formula. (fFull) does not imply that all subsets of $T^\omega 1$ are definable: For example, in the basic modal logic, having an infinite path is a property corresponding to a subset of $\mathcal{P}^\omega 1$ which cannot be expressed by the finitary approximants ($\bigwedge_{n < \omega} \diamond^n \top$ only implies that the lengths of paths cannot be bounded).

(Full) implies (Sep) and (fFull) implies (fSep).

All of $\mathcal{L}_T^M, \mathcal{L}_T^P, \mathcal{L}_T^F$ enjoy (fSep) and (Bool), \mathcal{L}_T^F also (fFull).

3 Finite depth

Definition 3.1 We say that an abstract coalgebraic logic \mathcal{L} has finite depth if all formulas are determined by some subset of $T^n 1$ for some $n < \omega$, that is, for all $\varphi \in \mathcal{L}$ there is $n < \omega$ such that $\llbracket \varphi \rrbracket = p_n^{-1}(\llbracket \varphi \rrbracket_n) = p_n^{-1}(p_n[\varphi])$.

Remark 3.2 (i) $\mathcal{L}_T^M, \mathcal{L}_T^P, \mathcal{L}_T^F$ have finite depth, the μ -calculus does not have finite depth.

- (ii) Every logic with finite depth is invariant under ω -bismilarity. The converse is not true.

4 Compactness

This section is needed for Theorems 7.3 and 7.8.

Notation We already overloaded the symbol \mathcal{L} to denote the logic as well as the collection of definable subsets of the final coalgebra. For example, if we say that the intersection of a set of formulas is non-empty, we refer to the extensions of the formulas on the final coalgebra.

Similarly we now write

$$\mathcal{L}_\alpha = \{p_\alpha[\varphi] \subseteq T^\alpha 1 \mid \varphi \in \mathcal{L}\}.$$

Note that \mathcal{L}_α contains (possibly strictly) the set of definable subsets $S \subseteq T^\alpha 1$ in the sense of Definition 2.10. Also note that even if \mathcal{L} is closed under Boolean operations, \mathcal{L}_α need not be (because the direct image only preserves unions but neither intersections nor complements).

Furthermore, when we treat $T^\alpha 1$ (or the carrier Z of the final coalgebra) as a topological space in the following, then we do this with respect to the topology generated by \mathcal{L}_α (or \mathcal{L}). For example, in the presence of (Bool), the following definition of compactness coincides with the topological definition of Z being compact (that is, every cover of opens has a finite subcover). Recall that a collection \mathcal{C} of subsets of the final coalgebra has the finite intersection property (f.i.p.) if all finite subcollections have non-empty intersection.

Definition 4.1 \mathcal{L} is compact if each collection of formulas with the finite intersection property has a non-empty intersection.

Later we will use results from topology to show that there is only one possible logic on $T^\omega 1$ satisfying certain restrictions. But let us first take a look at the logics we considered so far.

Proposition 4.2 (i) *If \mathcal{L} satisfies (fSep), then $T^\omega 1$ is Hausdorff.*

(ii) *Let \mathcal{L} be one of $\mathcal{L}_T^M, \mathcal{L}_T^P, \mathcal{L}_T^F$. Then $T^\omega 1$ is Hausdorff and has a basis of clopens.*

(iii) *Moreover, $(\mathcal{L}_T^M)_\omega \equiv (\mathcal{L}_T^P)_\omega \equiv (\mathcal{L}_T^F)_\omega$ are compact if T has (fs).*

Proof. (i) is immediate from the definitions. (ii) follows from (Bool). (iii) holds since under (fs) the approximants $T^n 1$, $n < \omega$, are finite (hence compact Hausdorff) and a limit of compact Hausdorff spaces is compact Hausdorff. \square

To reason about compactness we use the following standard topological facts. The first is an auxiliary statement. The second shows that a compact Hausdorff topology cannot be made smaller (without losing Hausdorff) nor bigger (without losing compactness). The third says that if compact Hausdorff spaces have a basis closed under Boolean operations, then this basis is uniquely determined.

- Proposition 4.3** (i) *In a compact space, every closed set is compact. In a Hausdorff space, every compact set is closed.*
- (ii) *Let $\tau \subseteq \tau'$ be two topologies on a set X and assume that τ is Hausdorff and τ' is compact (and hence both are compact Hausdorff). Then $\tau = \tau'$.*
- (iii) *Let (X, τ) be a Stone space, ie, τ is compact Hausdorff and has a basis of clopens. Then τ has one and only one basis closed under the Boolean operations.*

Proof. For the second statement, assume there is an open $o \in \tau'$, $o \notin \tau$. The complement o' of o is not closed in τ , hence not compact. Hence there is an open cover o'_i of o' from τ that has no finite subcover. This cover is also a cover from τ' , hence o' is not compact in τ' , hence not closed in τ' , contradicting $o \in \tau$. \square

5 Expressively closed logics

This section is needed for Theorem 7.8.

The following definition makes precise the requirement that if some property P is definable in the logic, then the weaker property P_α of ‘ P up to α -bisimilarity’ is also definable.

Definition 5.1 \mathcal{L} is expressively closed at α if for all $\varphi \in \mathcal{L}$ there is $\psi \in \mathcal{L}$ such that $p_\alpha^{-1}(\llbracket \varphi \rrbracket_\alpha) = \llbracket \psi \rrbracket$. \mathcal{L} is expressively closed if it is expressively closed at α for all ordinals α .

In other words, \mathcal{L} is expressively closed at α iff \mathcal{L}_α coincides with the definable subsets of $T^\alpha 1$ (see Definition 2.10).

Example 5.2 (i) $\mathcal{L}_T^M, \mathcal{L}_T^P, \mathcal{L}_T^F$ are expressively closed. Every logic extending \mathcal{L}_T^F is expressively closed at n for $n < \omega$.

(ii) Expressively closed at ω does not imply expressively closed at $n < \omega$. For example, in the case of $T = \mathcal{P}$, take a logic which has exactly 4 pairwise different formulas (false, true, φ , $\neg\varphi$) and where φ is definable at stage ω but not at any $n < \omega$.

(iii) The μ -calculus is not expressively closed, since (in the case of $T = \mathcal{P}$) the formula $\mu x. \Box x$ expresses termination of each execution sequence (well-foundedness) and this property cannot be defined at any ordinal level.

Remark 5.3 If \mathcal{L} is expressively closed at α then $p_\alpha : Z \rightarrow T^\alpha 1$ is continuous (with respect to the topology generated by \mathcal{L}_α).

Recall that if X is compact and $f : X \rightarrow Y$ is continuous and onto, then Y is compact. The previous remark then implies the first part of the proposition below.

Proposition 5.4 (i) *If \mathcal{L} is compact, expressively closed at α , and $p_\alpha : Z \rightarrow T^\alpha 1$ is onto, then \mathcal{L}_α is compact.*

(ii) *If \mathcal{L} is expressively closed and satisfies (Bool) and $p_\alpha : Z \rightarrow T^\alpha 1$ is onto, then \mathcal{L}_α is closed under Boolean operations.*

Proof. For the second item, because direct image preserves unions, \mathcal{L}_α is closed under unions. To see closure under complements let S be a definable subset and $\varphi = p_\alpha^{-1}(S)$. Then $\llbracket \neg\varphi \rrbracket_\alpha$ is the complement of S . \square

Proposition 5.5 *If \mathcal{L} is compact and extends \mathcal{L}_T^F , then*

- (i) $\bigcap\{(p_n^\omega)^{-1}(\llbracket \varphi \rrbracket_n) \mid n < \omega\} = \llbracket \varphi \rrbracket_\omega$
- (ii) $\bigcap\{p_n^{-1}(\llbracket \varphi \rrbracket_n) \mid n < \omega\} = p_\omega^{-1}(\llbracket \varphi \rrbracket_\omega)$
- (iii) $p_\omega : Z \rightarrow T^\omega 1$ is onto.

Proof. (ii) is immediate from (i). Putting $\varphi = \text{true}$ in (i) yields $T^\omega 1 = \llbracket \text{true} \rrbracket_\omega$, which implies (iii). To show (i), “ \supseteq ” is immediate. For “ \subseteq ”, assume that $t \in \bigcap\{(p_n^\omega)^{-1}(\llbracket \varphi \rrbracket_n) \mid n < \omega\}$, that is, for each $n < \omega$ there is $z_n \in \varphi$ with $p_n^\omega(t) = p_n(z_n)$. Since \mathcal{L} extends \mathcal{L}_T^F , we have $p_n^{-1}(\{p_n(z_n)\}) \in \mathcal{L}$. By construction, $\{p_n^{-1}(\{p_n(z_n)\}) \mid n < \omega\} \cup \{\varphi\}$ has the f.i.p. and so, by compactness, there is $z \in \varphi$ such that $p_n^\omega(t) = p_n(z)$ for all $n < \omega$. Hence $t = p_\omega(z) \in \llbracket \varphi \rrbracket_\omega$. \square

Remark 5.6 The proof also works for compact \mathcal{L} that have for each $t \in T^\omega 1$ a collection of definable $S_n \subseteq T^n 1$, $n < \omega$, such that $\{t\}$ is the intersection of the $(p_n^\omega)^{-1}(S_n)$. For example, the proposition holds if the singleton subsets of $T^n 1$, $n < \omega$ are definable, or if \mathcal{L} satisfies (fSep) and is expressively closed at n for all $n < \omega$.

The topological analysis above gives us a Lindström theorem ‘at ordinal level ω ’, roughly saying: an abstract coalgebraic logic that is compact and expressively closed at each finite n agrees with \mathcal{L}_T^F at ω . The Lindström theorems of Section 7 can be seen as variations where some further work is put into dropping the restriction ‘at ordinal level ω ’.

Theorem 5.7

- (i) *Let \mathcal{L} be an abstract coalgebraic logic that satisfies (fSep), (Bool) and is compact and expressively closed at n , $n < \omega$. Then $T^\omega 1$ is a Stone space.*
- (ii) *If, moreover, T satisfies (fs), then $\mathcal{L}_\omega = (\mathcal{L}_T^M)_\omega = (\mathcal{L}_T^P)_\omega = (\mathcal{L}_T^F)_\omega$.*

Proof. (i). By (fSep) and Proposition 4.2, \mathcal{L}_ω is Hausdorff. By compactness of \mathcal{L} and (Bool), and Propositions 5.5 and 5.4, \mathcal{L}_ω is compact and has a basis of clopens. (ii). If T preserves finite sets, then (fSep) and (Bool), together with expressively closed, imply (fFull), that is, \mathcal{L} extends \mathcal{L}_T^F . Hence (the topology generated by) \mathcal{L}_ω is a Stone topology which extends the Stone topology (Proposition 4.2) (generated by) $(\mathcal{L}_T^F)_\omega$. By Proposition 4.3(ii) the two topologies agree and by Proposition 4.3(iii) the same holds for the bases, that is, $\mathcal{L}_\omega = (\mathcal{L}_T^F)_\omega$. \square

Concerning the assumptions of the theorem, we recall that a logic that extends \mathcal{L}_T^F has (fSep) and is expressively closed at n . Also \mathcal{L}_T^M , \mathcal{L}_T^P , \mathcal{L}_T^F satisfy these two properties, as well as (Bool). But, depending on T , neither need to be compact.

6 Two lemmas on the final coalgebra sequence

For the proof of Theorem 7.8 we will need two lemmas about the final coalgebra sequence. We study the diagram

$$\begin{array}{ccccccc}
 1 & \xleftarrow{p_0^1} & T1 & \xleftarrow{p_1^2} & T^2 1 & \xleftarrow{\dots} & T^n 1 & \xleftarrow{\dots} & T^\omega 1 \\
 \uparrow k_0 & & \uparrow k_1 & & \uparrow k_2 & & \uparrow k_n & & \uparrow k_\omega \\
 0 & \xrightarrow{e_1^0} & T0 & \xrightarrow{e_2^1} & T^2 0 & \xrightarrow{\dots} & T^n 0 & \xrightarrow{\dots} & T^\omega 0
 \end{array} \quad (8)$$

where the upper row is the final coalgebra sequence (4) and the lower row is its dual, the initial algebra sequence. The arrows k_n are the unique ones induced by initiality/finality: k_0 is the empty map and $k_{n+1} = Tk_n$; for each $n < \omega$ this induces over the terminal sequence $(T^m 1)_{m < \omega}$ a cone¹ $k_{nm} : T^n 0 \rightarrow T^m 1$ given by $k_{nm} = k_m \circ e_m^n$ for $m > n$ and by $k_{nm} = p_m^n \circ k_n$ for $m \leq n$; since $T^\omega 1$ is a limit, the cones $(k_{nm})_{m < \omega}$ induce maps $k_{n\omega} : T^n 0 \rightarrow T^\omega 1$, which in turn form a co-cone² over the initial sequence $(T^n 0)_{n < \omega}$; since $T^\omega 0$ is a colimit, the co-cone $(k_{n\omega})_{n < \omega}$ induces k_ω .

In the situation above, the arrows p_n^ω induce a metric d on $T^\omega 1$ via $d(x, y) = 2^{-n}$ where n is the smallest number such that $p_n^\omega(x) \neq p_n^\omega(y)$. Moreover, $T^\omega 0$ inherits this metric via the injective k_ω .

Proposition 6.1 (Barr [4, Proposition 3.1]) *Let $T : \text{Set} \rightarrow \text{Set}$ with $T0 \neq 0$. Then k_ω in diagram (4) is a Cauchy completion.*

The topology on $T^\omega 1$ induced by the metric coincides with the limit topology over the discrete spaces $T^n 1, n < \omega$. The topology is compact, if the $T^n 1, n < \omega$, are finite. Then the collection of clopens of this topology is $(\mathcal{L}_T^F)_\omega$. Since $T^\omega 0$ is dense in $T^\omega 1$ we have

Lemma 6.2 *Let $T : \text{Set} \rightarrow \text{Set}$ with $T0 \neq 0$. Then for all $\varphi \in \mathcal{L}_T^F$, either $[\varphi]_\omega = \emptyset$ or $[\varphi]_\omega \cap T^\omega 0 \neq \emptyset$.*

For the next lemma, recall the notation $p_\omega : Z \rightarrow T^\omega 1$ from Diagram (4).

Lemma 6.3 *Let T be a weak pullback preserving functor $\text{Set} \rightarrow \text{Set}$ with $T0 \neq 0$. If for two elements z, z' of the final coalgebra satisfying $p_\omega(z) = p_\omega(z')$ we have that $p_\omega(z) = p_\omega(z')$ is in the image of k_ω in (8), then $z = z'$.*

¹ That is, $p_l^m \circ k_{nl} = k_{nl}$ for all $l \leq m < \omega$.

² That is, $k_{m\omega} \circ e_m^n = k_{n\omega}$ for all $n \leq m < \omega$.

Proof. Without loss of generality, assume that k_ω is an inclusion. Consider

$$\begin{array}{ccccccc}
 & & & & p_n & & \\
 & & & & \swarrow & & \\
 & & & & T^n 1 & \xleftarrow{p^{n+1}} & T^{n+1} 1 & \xleftarrow{p_{n+1}} & \dots & \xleftarrow{p} & Z & \xleftarrow{p} & TZ \\
 & & & & \uparrow & & \uparrow & & & & \uparrow & & \uparrow \\
 1 & \dots & T^n 1 & \xleftarrow{p^{n+1}} & T^{n+1} 1 & \dots & Z & \xleftarrow{p} & TZ \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 k_0 & & k_n & & k_{n+1} & & k & & Tk \\
 & & & & & & & & \\
 0 & \dots & T^n 0 & \xrightarrow{e_{n+1}^o} & T^{n+1} 0 & \dots & A & \xrightarrow{e} & TA \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 & & e_{n+1} & & e_{n+1} & & e & & \\
 & & & & e_n & & & &
 \end{array}$$

which is as (8) but extended transfinitely through the ordinals until reaching the initial algebra $(A, \alpha) = (A, e^{-1})$ and the final coalgebra $(Z, \zeta) = (Z, p^{-1})$ [3]. We write “;” for relational composition and converse of R as R^o , eg, if f and g are functions, $(g \circ f)^o = g^o; f^o$ and also $p^{-1} = p^o$, $e^{-1} = e^o$ (p and e are isos). The claim follows once we have shown $e_n; k = k_n; p_n^o$ for all $n < \omega$. Indeed, $p_\omega(z) = p_\omega(z') \in T^\omega 0$ means there is $n < \omega$ and $x \in T^n 0$ such that $x(k_n; p_n^o)z$ and $x(k_n; p_n^o)z'$. Since $e_n; k$ is a functional relation, we must have $z = z'$.

To show $e_n; k = k_n; p_n^o$ first observe that it holds for $n = 0$ since then both relations are the empty relation. Further, $e_{n+1}; k =$ (by definition of $e = \alpha^o$ and (A, α) being initial) $Te_n; e^o; k =$ (by definition of k) $Te_n; Tk; p =$ (by ind.hyp. and T preserving weak pullbacks) $Tk_n; (Tp_n)^o; p =$ (by definition of $p = \zeta^o$ and (Z, ζ) being final) $Tk_n; (p_{n+1} \circ p)^o; p =$ (by $p^o = \zeta$ being iso and definition of k_{n+1}) $k_{n+1}; p_{n+1}^o$. \square

7 Lindström theorems

The theorems below state that a logic \mathcal{L} invariant under bisimilarity, closed under Boolean operations and at least as expressive as \mathcal{L}_T^F is actually equivalent to \mathcal{L}_T^F if additionally one of the following conditions is satisfied

- \mathcal{L} has finite depth,
- \mathcal{L} is compact and invariant under ω -bisimilarity,
- \mathcal{L} is compact and expressively closed at ω .

7.1 Finite depth

The essence of de Rijke’s Lindström theorem [23] is coalgebraic and, in view of the final sequence, its proof is almost obvious: That \mathcal{L} has finite depth means that formulas are determined by their extensions on the approximants $T^n 1$, $n < \omega$, and that \mathcal{L} extends \mathcal{L}_T^F means, conversely, that all subsets of $T^n 1$, $n < \omega$, are definable.

Theorem 7.1 *If an abstract coalgebraic logic \mathcal{L} extends \mathcal{L}_T^F and has finite depth, then $\mathcal{L} \equiv \mathcal{L}_T^F$.*

Proof. Because \mathcal{L} is invariant under bisimilarity and has finite depth, each formula φ is determined by some $T^n 1$, $n < \omega$, that is, $\varphi = p_n^{-1}(\llbracket \varphi \rrbracket_n)$ for some $n < \omega$. Since all subsets of $T^n 1$ are definable in \mathcal{L}_T^F , it follows $\mathcal{L} \leq \mathcal{L}_T^F$. \square

- Corollary 7.2** (i) If $T = \mathcal{P}$, or $T = \mathcal{P} \times 2^P$ for a finite set P of atomic propositions, we recover de Rijke’s theorem [23], modulo the simplification that we do not consider infinitely many atomic propositions.³
- (ii) If $T = \mathcal{H}$, or $T = \mathcal{H} \times 2^P$, we obtain a Lindström theorem for non-normal modal logics and neighbourhood frames/models.
- (iii) More generally, if T satisfies (fs) [and (wp)], the conclusion of the theorem can be strengthened to $\mathcal{L} \equiv \mathcal{L}_T^P$ [and $\mathcal{L} \equiv \mathcal{L}_T^M$].

7.2 Compactness and ω -bisimilarity

Since Lindström’s theorem characterising first-order logic [19] makes crucial use of compactness, it is natural to follow van Benthem [5] and search for characterisations of modal logic involving compactness.

The proof of the following theorem is similar to Theorem 7.1, replacing the role of the sequence of approximants $(T^n 1)_{n < \omega}$ by $T^\omega 1$: That \mathcal{L} is invariant under ω -bisimilarity means that formulas are determined by their extensions on $T^\omega 1$ and that \mathcal{L} is compact and extends \mathcal{L}_T^F means that all clopen (closed and open) subsets of $T^\omega 1$ are extensions of some \mathcal{L} -formula. Finally, we use a result from topology to establish that the topologies on $T^\omega 1$ induced by the definable subsets of \mathcal{L} and \mathcal{L}_T^F coincide.

Theorem 7.3 *If T preserves finite sets and an abstract coalgebraic logic \mathcal{L} has the following properties: (i) invariant under ω -bisimilarity, (ii) closed under Boolean operations, (iii) \mathcal{L} extends \mathcal{L}_T^F , (iv) \mathcal{L} is compact, then $\mathcal{L} \equiv \mathcal{L}_T^F$.*

Proof. By (i), we can consider \mathcal{L} as a collection of subsets of $T^\omega 1$. By (ii), \mathcal{L} is a Boolean algebra. The topology τ generated by the basis \mathcal{L} is Hausdorff by (iii). By (iv), $(T^\omega 1, \tau)$ is a Stone space, by Proposition 4.3(ii) we have that τ is the topology generated by \mathcal{L}_T^F . By Proposition 4.3(iii), the basis of clopens of this space is uniquely determined, hence $\mathcal{L} \equiv \mathcal{L}_T^F$. \square

7.3 Compactness and invariance under bisimilarity is not enough

We investigate what can be said if, in the assumptions of Theorem 7.3, invariance under ω -bisimilarity is weakened to invariance under bisimilarity. The following example shows that we need to look for some additional condition.

Example 7.4 The logic \mathcal{L}^{cb} is obtained from basic modal logic (without propositional variables) by adding one constant θ , which may not appear under a \Box . Explicitly, formulas are constructed according to

$$\text{phi} ::= \theta \mid \text{psi} \mid \neg \text{phi} \mid \text{phi} \wedge \text{phi} \tag{9}$$

$$\text{psi} ::= \perp \mid \neg \text{psi} \mid \text{psi} \wedge \text{psi} \mid \Box \text{psi} \tag{10}$$

³ For infinite P the logic $\mathcal{L}_{\mathcal{P} \times 2^P}^F$ is more expressive than basic modal logic (it has formulas with infinitely many atomic propositions). This problem can be overcome in several ways, but the issues arising are orthogonal to the interests of this paper.

Given any von Neumann ordinal α , let M_α be the pointed Kripke frame which has root (distinguished point) α , carrier $\alpha + 1 = \alpha \cup \{\alpha\}$ and the accessibility relation given by the converse of \in . The semantics of θ is then given by

$$(W, R), w \Vdash \theta \Leftrightarrow (W, R), w \text{ bisimilar to } M_{\omega+2}$$

Remark 7.5 With respect to Theorem 7.3, note that $M_{\omega+2}$ and $M_{\omega+1}$ are ω -bisimilar but not bisimilar (an observation that can be found in [9]) and hence \mathcal{L}^{cb} is not invariant under ω -bisimilarity.

In our approach we identify formulas with their extension on the final coalgebra Z . In this view, θ is the singleton subset of the final coalgebra containing the equivalence class of $M_{\omega+2}$ up to bisimilarity. (If we take Z to be the non-well founded sets of Aczel [1], then $\theta = \{\omega + 2\}$.) This observation allows us to replace (9) by

$$\text{phi} ::= \text{psi} \mid \text{psi} \vee \theta \mid \text{psi} \wedge \neg\theta \tag{11}$$

To show that (9) and (11) are equivalent, we need to show that the language described by (11) is closed under Boolean operations. Closure under negation (complement) is immediate. For closure under conjunction (intersection) we have to check 6 cases. For instance, using that θ is a singleton, we calculate on the final coalgebra

$$(\psi_1 \vee \theta) \wedge (\psi_2 \vee \theta) = (\psi_1 \wedge \psi_2) \vee \theta \tag{12}$$

since for any φ we have that $\varphi \wedge \theta$ is either \emptyset or θ .

Proposition 7.6 \mathcal{L}^{cb} is compact.

Proof. To check compactness we (identify formulas with their extension on the final coalgebra and) assume that we are given a set $\Psi \cup \Psi_\vee \cup \Psi_\wedge$ of formulas enjoying the f.i.p. and where $\Psi, \Psi_\vee, \Psi_\wedge$ contain formulas of type $\text{psi}, \text{psi} \vee \theta, \text{psi} \wedge \neg\theta$, respectively.

Case 1: Ψ_\wedge is non-empty. Then from the f.i.p. we can assume without loss of generality that Ψ_\vee is empty and that $\Psi_\wedge = \{\neg\theta\}$. For a contradiction assume that $\bigcap \Psi \cup \Psi_\vee \cup \Psi_\wedge$ is empty, ie, $\bigcap \Psi \subseteq \theta$. Then, since θ is a singleton and Ψ has the f.i.p. (and $\mathcal{L}_{\mathcal{P}}^F$ is compact), it follows $\bigcap \Psi = \theta$. This contradicts the fact that θ is not definable by an infinite conjunction of formulas in $\mathcal{L}_{\mathcal{P}}^F$.

Case 2: Ψ_\wedge is empty. Then, reasoning as in (12), we see that $\bigcap \Psi_\vee$ is semantically equivalent to $C \vee \theta$ for some $C \subseteq Z$ where C is the intersection of some formulas in $\mathcal{L}_{\mathcal{P}}^F$. Note that $\Psi \cup \{C \vee \theta\}$ has the f.i.p. iff $\Psi \cup \{C\}$ or $\Psi \cup \{\theta\}$ have the f.i.p. If $\Psi \cup \{C\}$ has the f.i.p., we have $\bigcap \Psi \cap C \neq \emptyset$ since $\mathcal{L}_{\mathcal{P}}^F$ is compact; hence also $\bigcap \Psi \cap \{C \vee \theta\} \neq \emptyset$. If $\Psi \cup \{\theta\}$ has the f.i.p., we have $\theta \subseteq \psi$ for all $\psi \in \Psi$ since θ is a singleton; hence $\bigcap \Psi \cap \{C \vee \theta\} \neq \emptyset$. \square

To summarise:

Theorem 7.7 \mathcal{L}^{cb} is closed under Boolean operations, extends the basic modal logic $\mathcal{L}_{\mathcal{P}}^F$, is invariant under bisimilarity and is compact, but is not equivalent to $\mathcal{L}_{\mathcal{P}}^F$.

7.4 Compact and expressively closed at ω

The previous section showed that compactness needs to be complemented by some further condition. In Theorem 7.3 it was invariance under ω -bisimilarity, which we replace now by the weaker condition of being expressively closed at ω (Definition 5.1), meaning that if some property P is definable in the logic, then the weaker property P_ω of ‘ P up to ω -bisimilarity’ is also definable.

Theorem 7.8 *Let $T : \text{Set} \rightarrow \text{Set}$ preserve finite sets and weak pullbacks. If an abstract coalgebraic logic \mathcal{L} (i) extends \mathcal{L}_T^F , (ii) is closed under Boolean operations, (iii) is compact and (iv) is expressively closed at ω , then $\mathcal{L} \equiv \mathcal{L}_T^F$.*

Proof. Suppose \mathcal{L} extends \mathcal{L}_T^F , that is, there is a formula $\theta \in \mathcal{L}$ such that $\theta \subsetneq p_n^{-1}(\llbracket \theta \rrbracket_n)$ for all $n < \omega$. If $\theta = \bigcap \{p_n^{-1}(\llbracket \theta \rrbracket_n) \mid n < \omega\}$ we obtain a contradiction to compactness, so we assume that $\theta \subsetneq \bigcap \{p_n^{-1}(\llbracket \theta \rrbracket_n) \mid n < \omega\} = p_\omega^{-1}(\llbracket \theta \rrbracket_\omega)$, where the latter equality is due to Proposition 5.5. By (iv), there is $\psi \in \mathcal{L}$ such that (the extension of) ψ is $p_\omega^{-1}(\llbracket \theta \rrbracket_\omega)$. By (ii), we have $\psi \wedge \neg\theta \in \mathcal{L}$ and due to $\theta \subsetneq p_\omega^{-1}(\llbracket \theta \rrbracket_\omega)$ we know that $\llbracket \psi \wedge \neg\theta \rrbracket_\omega$ is non-empty. Using that \mathcal{L} has (wp), it now follows from Lemma 6.3 that $\llbracket \psi \wedge \neg\theta \rrbracket_\omega \cap T^\omega 0 = \emptyset$. Indeed, $t \in \llbracket \psi \wedge \neg\theta \rrbracket_\omega$ implies, on the one hand, $t \in \llbracket \psi \rrbracket_\omega$ and hence the existence of a $z \in \theta$ with $p_\omega(z) = t$ and, on the other hand, $t \in \llbracket \neg\theta \rrbracket_\omega$ and hence the existence of a $z' \notin \theta$ with $p_\omega(z') = t$; we thus have $p_\omega(z) = p_\omega(z')$ and, since $z \neq z'$, can apply Lemma 6.3 to conclude that $t \notin T^\omega 0$. But $\llbracket \psi \wedge \neg\theta \rrbracket_\omega \cap T^\omega 0 = \emptyset$ now contradicts $\mathcal{L}_\omega = (\mathcal{L}_T^F)_\omega$ (using that \mathcal{L} has (fs) and Theorem 5.7) and the fact that all non-empty sets in $(\mathcal{L}_T^F)_\omega$ intersect $T^\omega 0$ (Lemma 6.2). \square

Remark 7.9 One way to analyse Example 7.4 in the light of Theorems 7.3 and 7.8 is as follows. In Example 7.4, the extension $\llbracket \theta \rrbracket_\omega$ of θ on $T^\omega 1$ is closed and not open in the topology generated by \mathcal{L}_ω^{cb} . From the uniqueness of compact Hausdorff topologies (Proposition 4.3), we know that the complement of $\llbracket \theta \rrbracket_\omega$ cannot be in \mathcal{L}_ω^{cb} , which suggests a contradiction to \mathcal{L}^{cb} being closed under Boolean operations. Unfortunately, the complement of $\llbracket \theta \rrbracket_\omega$ is not definable in \mathcal{L}^{cb} (note that $\llbracket \neg\theta \rrbracket_\omega = T^\omega 1$). As we have seen in the proof of Theorem 7.8, the notion of expressive closure at ω , together with the two lemmas of Section 6, gives us a way to exploit $\neg\theta$.

The following specialises Theorem 7.8 to Kripke models.

Corollary 7.10 *Let \mathcal{L} be a logic which extends basic modal logic, is invariant under bisimilarity, is closed under Boolean operations, is compact and is expressively closed at ω , then \mathcal{L} is basic modal logic.*

Similarly the result applies to all Kripke polynomial functors.

7.5 Discussion

In [5], van Benthem presents a modal Lindström theorem stating that an *abstract modal logic* \mathcal{L} extending basic modal logic is equivalent to basic modal logic if \mathcal{L} is invariant under bisimilarity and \mathcal{L} is compact. To compare this result to ours, we recall that the definition of abstract modal logic includes the property

(rl) closed under relativisation,

which means that for every formula φ and every proposition letter p , the logic contains a formula $rel(\varphi, p)$ which is true at a state x in a Kripke model M iff φ is true at x in the Kripke model we obtain from M by throwing away all the states where p is false. This condition is also part of Lindström’s definition of an abstract logic, but from our coalgebraic perspective, it is not so natural. In particular, given a coalgebra $X \rightarrow TX$ and a subset $X' \subseteq X$, it is not a priori clear what the induced coalgebra with carrier X' would be. (In the case where $T = P$ one can simply take for $\xi'(x)$ the intersection of $\xi(x)$ with X').

In addition, there are natural logics that do not satisfy (rl). Consider for instance the diamond $\langle \star \rangle$, which is to be interpreted over the reflexive/transitive closure of the accessibility relation of the diamond \diamond . If we consider the language with $\langle \star \rangle$ but *without* \diamond , we obtain a natural logic which is bisimulation invariant and compact, but does not have the relativisation property, as can easily be verified.

Finally, there is the work of Otto and Piro [21] on Lindström theorems for the extension of basic modal logic by a global modality, and of the guarded fragment of first-order logic. This work revolves around the

(tup) Tarski Union Property,

which requires the logic L to be closed under unions of L -elementary chains. Without going into the details, we just mention that the definition of this notion also involves *substructures*, which, as opposed to *generated substructures*, do not provide a coalgebraic notion. For this reason, the property (tup) does not seem to be a natural candidate for coalgebraic generalisations.

8 Conclusion

We showed that de Rijke’s modal Lindström theorem is coalgebraic in nature and generalises from Kripke models to T -coalgebras for a large class of functors T . De Rijke’s theorem is based on the notion of finite depth, whereas Lindström’s original theorem makes crucial use of compactness. We therefore presented two coalgebraic Lindström theorems, replacing finite depth by compactness plus an additional condition. We showed that some additional condition is needed, but there may be other conditions still to be discovered (ideally such a condition would be enjoyed by all important non-compact logics extending basic modal logic). Further open questions include

- a Lindström theorem that covers basic modal logic, implies van Benthem’s result, and can be generalised to coalgebra of arbitrary type,
- a Lindström theorem for modal logics extended with fixpoint operators, in particular, modal μ -calculus,
- Lindström theorems that do not mention compactness and work for modal languages smaller than \mathcal{L}_T^F such as probabilistic modal logic [14].

References

- [1] Aczel, P., “Non-Well-Founded Sets,” CSLI, Stanford, 1988.

- [2] Aczel, P. and N. P. Mendler, *A final coalgebra theorem*, in: *Category Theory and Computer Science*, LNCS **389**, 1989.
- [3] Adámek, J. and V. Koubek, *On the greatest fixed point of a set functor*, Theoret. Comput. Sci. **150** (1995).
- [4] Barr, M., *Terminal coalgebras in well-founded set theory*, Theoret. Comput. Sci. **114** (1993).
- [5] van Benthem, J., *A new modal Lindström theorem*, Logica universalis **1** (2007).
- [6] van Benthem, J., B. ten Cate and J. A. Väänänen, *Lindström theorems for fragments of first-order logic*, Logical Methods in Computer Science **5** (2009).
- [7] Blackburn, P., M. de Rijke and Y. Venema, “Modal Logic,” CUP, 2001.
- [8] Ferrari, G., U. Montanari and M. Pistore, *Minimizing transition systems for name passing calculi: A co-algebraic formulation*, in: *FoSSaCS'02*.
- [9] Gerbrandy, J., “Bisimulations on Planet Kripke,” Ph.D. thesis, University of Amsterdam (1999).
- [10] Ghilardi, S., *An algebraic theory of normal forms*, Ann. Pure Appl. Logic **71** (1995).
- [11] Goldblatt, R., *Final coalgebras and the Hennessy-Milner property*, Ann. Pure Appl. Logic **183** (2006).
- [12] H.-D. Ebbinghaus, J. F. and W. Thomas, “Mathematical Logic,” Springer, 1994.
- [13] Hansen, H. and C. Kupke, *A coalgebraic perspective on monotone modal logic*, in: *CMCS'04, ENTCS* 106, 2004.
- [14] Heifetz, A. and P. Mongin, *Probabilistic logic for type spaces*, Games and Economic Behavior **35** (2001).
- [15] Kupke, C., A. Kurz and Y. Venema, *Completeness of the finitary Moss logic*, in: *Advances in Modal Logic*, 2008.
- [16] Kupke, C. and R. A. Leal, *Characterising behavioural equivalence: Three sides of one coin*, in: *CALCO'09*.
- [17] Kurz, A. and R. Leal, *Equational coalgebraic logic*, in: *MFPS'09*.
- [18] Kurz, A. and D. Pattinson, *Coalgebraic modal logic of finite rank*, Math. Structures Comput. Sci. **15** (2005).
- [19] Lindström, P., *On extensions of elementary logic*, Theoria **35** (1969).
- [20] Moss, L., *Coalgebraic logic*, Ann. Pure Appl. Logic **96** (1999).
- [21] Otto, M. and R. Piro, *A Lindström characterisation of the guarded fragment and of modal logic with a global modality*, in: *Advances in Modal Logic*, 2008.
- [22] Pattinson, D., *Coalgebraic modal logic: Soundness, completeness and decidability of local consequence*, Theoret. Comput. Sci. **309** (2003).
- [23] de Rijke, M., *A Lindström theorem for modal logic*, in: *Modal Logic and Process Algebra* (1995).
- [24] Rutten, J., *Automata and coinduction - an exercise in coalgebra*, in: *CONCUR'98*.
- [25] Schröder, L., *Expressivity of Coalgebraic Modal Logic: The Limits and Beyond*, in: *FoSSaCS'05*.
- [26] Worrell, J., *On the final sequence of a finitary set functor*, Theoret. Comput. Sci. **338** (2005).