The Goldblatt-Thomason Theorem for Coalgebras

Alexander Kurz[1]⋆ Jiří Rosický[2]⋆⋆

[1] University of Leicester, UK
[2] Masaryk University, Brno, Czech Republic

Abstract. Goldblatt and Thomason's theorem on modally definable classes of Kripke frames and Venema's theorem on modally definable classes of Kripke models are generalised to coalgebras.

1 Introduction

The Goldblatt-Thomason theorem [11] states that a class of Kripke frames closed under ultrafilter extensions is modally definable if and only if it reflects ultrafilter extensions and is closed under generated subframes, homomorphic images and disjoint unions. The proof is based on the duality between Boolean algebras and sets

$$\mathsf{BA} \underset{\Sigma}{\overset{\Pi}{\rightleftarrows}} \mathsf{Set}^{\mathrm{op}} \tag{1}$$

where Π is powerset and Σ assigns to a BA the set of ultrafilters. Σ is left-adjoint to Π but, of course, this adjunction does *not* form a dual equivalence. The price we have to pay for this is that going from Set to BA and back leaves us with $\Sigma\Pi X$: If X is the carrier of a Kripke frame, then its ultrafilter extension has carrier $\Sigma\Pi X$, which explains why ultrafilter extensions appear in the theorem.

Our generalisation from Kripke frames to T-coalgebras works as follows. Set and BA are completions (with filtered colimits) of the categories Set_ω of finite sets and BA_ω of finite Boolean algebras, respectively. BA_ω and Set_ω are dually equivalent. Now, given a functor T on Set that preserves finite sets, we can restrict T to Set_ω. Via the dual equivalence $\mathsf{BA}_\omega \simeq \mathsf{Set}_\omega^{\mathrm{op}}$, this gives us a functor on BA_ω, which we can then lift to a functor $L : \mathsf{BA} \to \mathsf{BA}$.

$$\begin{array}{ccc} L \circlearrowright \mathsf{BA} & \underset{\Sigma}{\overset{\Pi}{\rightleftarrows}} & \mathsf{Set}^{\mathrm{op}} \circlearrowleft T \\ \uparrow & & \uparrow \\ \mathsf{BA}_\omega & \rightleftarrows & \mathsf{Set}_\omega^{\mathrm{op}} \end{array} \tag{2}$$

⋆ Partially supported by EPSRC EP/C014014/1.

⋆⋆ Supported by the Ministry of Education of the Czech Republic under the project 1M0545.

[17] showed the following: (i) L has a presentation and therefore determines a logic for T-coalgebras, (ii) Π extends to a functor $\mathsf{Coalg}(T) \to \mathsf{Alg}(L)$, (iii) if T weakly preserves cofiltered limits, then Σ extends to a map on objects $\mathsf{Alg}(L) \to \mathsf{Coalg}(T)$. This note shows that the classical Goldblatt-Thomason theorem generalises to those T-coalgebras where $\Sigma : \mathsf{BA} \to \mathsf{Set}$ can be extended to a *functor* $\mathsf{Alg}(L) \to \mathsf{Coalg}(T)$.

$$\mathsf{Alg}(L) \underset{\Sigma}{\overset{\Pi}{\leftrightarrows}} \mathsf{Coalg}(T)^{\mathrm{op}} \tag{3}$$

The same argument also generalises a similar definability result for Kripke models due to Venema [22].

Related Work An algebraic semantics for logics for coalgebras and its investigation via the adjunction between BA and Set has been given in Jacobs [13]. The idea that a logic for T-coalgebras is a functor L on BA appears in [5,15] and can be traced back to Abramsky [1,2] and Ghilardi [10]. It has been further developed in [6,16]. The general picture underlying diagram (2) has been discussed in Lawvere [19] where it is attributed to Isbell. The implications of this Isbell-conjugacy for logics for coalgebras are explained in [17]. For topological spaces, which can be seen as particular coalgebras, the Goldblatt-Thomason theorem is due to Gabelaia [9] and ten Cate et al [7].

Acknowledgements Our greatest debts are to David Gabelaia whose sharp eye allowed us to correct a serious mistake in the statement of our main theorem. He also suggested numerous other improvements and made us aware of Venema's paper on definability for Kripke models [22]. The first author would like to thank Dmitry Sustretov for the discussions during ESSLLI 2006 in Malaga, which prompted the writing of this paper; and Neil Ghani for relating Conor McBride's statement about isomorphisms, see footnote 3. We are also grateful to the referees for their suggestions.

2 Coalgebras and their logics

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

Examples of functors of interest to us in this paper are described by

Definition 2.2 (gKPF). *A generalised Kripke polynomial functor (gKPF)* $T : \mathsf{Set} \to \mathsf{Set}$ *is built according to*

$$T ::= Id \mid K_C \mid T + T \mid T \times T \mid T \circ T \mid \mathcal{P} \mid \mathcal{H}$$

where Id *is the identity functor,* K_C *is the constant functor that maps all sets to a finite set* C*,* $\mathcal{P}$ *is covariant powerset and* $\mathcal{H}$ *is* 2^{2^-}.

Remark 2.3. The term 'Kripke polynomial functor' was coined in Rößiger [20]. We add the functor $\mathcal{H}$. $\mathcal{H}$-coalgebras are known as neighbourhood frames in modal logic and are investigated, from a coalgebraic point of view, in Hansen and Kupke [12].

We describe logics for coalgebras by functors L on the category BA of Boolean algebras. Although this approach differs conceptually from Jacobs's [13], the equations appearing in the example below are the same as his.

Example 2.4. We describe functors $L : \mathsf{BA} \to \mathsf{BA}$ or $L : \mathsf{BA} \times \mathsf{BA} \to \mathsf{BA}$ by generators and relations as follows.

1. $L_{K_C}(A)$ is the free BA given by generators $c \in C$ and satisfying $c_1 \wedge c_2 = \bot$ for all $c_1 \neq c_2$ and $\bigvee_{c \in C} c = \top$.
2. $L_+(A_1, A_2)$ is generated by $[\kappa_1]a_1$, $[\kappa_2]a_2$, $a_i \in A_i$ where the $[\kappa_i]$ preserve finite joins and binary meets and satisfy $[\kappa_1]a_1 \wedge [\kappa_2]a_2 = \bot$, $[\kappa_1]\top \vee [\kappa_2]\top = \top$, $\neg[\kappa_1]a_1 = [\kappa_2]\top \vee [\kappa_1]\neg a_1$, $\neg[\kappa_2]a_2 = [\kappa_1]\top \vee [\kappa_2]\neg a_2$.
3. $L_\times(A_1, A_2)$ is generated by $[\pi_1]a_1$, $[\pi_2]a_2$, $a_i \in A_i$ where $[\pi_i]$ preserve Boolean operations.
4. $L_\mathcal{P}(A)$ is generated by $\Box a$, $a \in A$, and $\Box$ preserves finite meets.
5. $L_\mathcal{H}(A)$ is generated by $\Box a$, $a \in A$ (no equations).

Informally, the equations in the 2nd item are justified as follows. Take A_1, A_2 to be the collections of subsets of two sets X_1, X_2, take $[\kappa_i]a_i$ to be the direct image of the injection $\kappa_i : X_i \to X_1 + X_2$ and describe how the $[\kappa_i]$ interact with the Boolean operations, interpreting $\wedge$ as $\cap$, etc.

More formally, we recall that sets and Boolean algebras are related by two functors

$$\mathsf{BA} \underset{\Sigma}{\overset{\Pi}{\rightleftarrows}} \mathsf{Set}^{\mathrm{op}} \tag{4}$$

where Π maps a set to its powerset and Σ a Boolean algebra to the set of its ultrafilters. On arrows, both functors are given by inverse image.

The justification for the presentations is now given, in essence, by the following isomorphisms. For Boolean algebras A, A_1, A_2, we have $L_{K_C}(A) \cong \Pi C$; $L_+(A_1, A_2) \cong A_1 \times A_2$; $L_\times(A_1, A_2) \cong A_1 + A_2$. For finite sets X, we have $L_\mathcal{P}(\Pi X) \cong \Pi\mathcal{P}X$; $L_\mathcal{H}(\Pi X) \cong \Pi\mathcal{H}X$. We will make this more precise in Definition 2.6 and Proposition 2.8.

Definition 2.5 (L_T). *For each gKPF (see Definition 2.2) $T : \mathsf{Set} \to \mathsf{Set}$ we define L_T by the corresponding constructions of Example 2.4.*

Example 2.4 illustrates how (a presentation of) a functor on BA describes the syntax and proof system of a modal logic. The semantics is given by a natural transformation

$$L\Pi X \xrightarrow{\quad \delta_X \quad} \Pi T X, \tag{5}$$

since this is exactly what is needed to define the extension $[\![-]\!]$ of formulas via the unique morphism from the initial L-algebra $LI \to I$. In detail, given a coalgebra (X, ξ) we let $[\![-]\!]$ be as in

$$\begin{array}{ccccc} I & & \longleftarrow & & LI \\ \downarrow{\scriptstyle [\![-]\!]} & & & & \downarrow{\scriptstyle L[\![-]\!]} \\ \Pi X & \xleftarrow[\Pi\xi]{} & \Pi T X & \xleftarrow[\delta_X]{} & L \Pi X \end{array} \tag{6}$$

In our examples, for gKPFs T, we define $\delta_T : L_T \Pi \to \Pi T$ as follows.

Definition 2.6 **(δ_T).** *We define Boolean algebra morphisms*

1. $L_{K_C} \Pi X \to \Pi C$ *by* $c \mapsto \{c\}$,
2. $L_+(\Pi X_1, \Pi X_2) \to \Pi(X_1 + X_2)$ *by* $[\kappa_i]a_i \mapsto a_i$,
3. $L_\times(\Pi X, \Pi Y) \to \Pi(X_1 \times X_2)$ *by* $[\pi_1]a_1 \mapsto a_1 \times X_2$, $[\pi_2]a_2 \mapsto X_1 \times a_2$,
4. $L_{\mathcal{P}} \Pi X \to \Pi \mathcal{P} X$ *by* $\Box a \mapsto \{b \subseteq X \mid b \subseteq a\}$,
5. $L_{\mathcal{H}} \Pi X \to \Pi \mathcal{H} X$ *by* $\Box a \mapsto \{s \in \mathcal{H} X \mid a \in s\}$.

and extend them inductively to $\delta_T : L_T \Pi \to \Pi T$ *for all gKPF* T.

The definition exploits that BA-morphisms are determined by their action on the generators.

Example 2.7. Together with (6), item 4 and 5 of Definition 2.6 give rise to the Kripke and neighbourhood semantics of modal logic:

- For $\xi : X \to \mathcal{P} X$ and $\Box\varphi$ in the initial $L_{\mathcal{P}}$-algebra, we have $[\![\Box\varphi]\!] = \{x \in X \mid \xi(x) \subseteq [\![\varphi]\!]\}$;
- For $\xi : X \to \mathcal{H} X$ and $\Box\varphi$ in the initial $L_{\mathcal{H}}$-algebra, we have $[\![\Box\varphi]\!] = \{x \in X \mid [\![\varphi]\!] \in \xi(x)\}$.

The justification for the definition of L_T and δ_T is now given by the following proposition. It says that (L, δ) completely captures the action of T on finite X; and more can hardly be expected from a finitary logic of T.

Proposition 2.8. *Let* T *be a gKPF. Then* $(\delta_T)_X : L_T \Pi X \to \Pi T X$ *is an isomorphism for all finite sets* X.

Proof. For finite X, $(\delta_T)_X : L_T \Pi X \to \Pi T X$ is an isomorphism in all of the 5 cases of Definition 2.6. The result then follows by induction, using that all the functors involved restrict to finite sets and finite BAs. $\Box$

The property of Proposition 2.8, namely

$$L \Pi X \cong \Pi T X \qquad \text{for all finite sets } X, \tag{7}$$

or, equivalently, $LA \cong \Pi T \Sigma A$ for finite A, is of central importance as it sets up the relationship between the logic (=functors L given by a presentation) and the semantics (=functor T). (7) can be read in two different ways: If the logic (ie L and $L\Pi \to \Pi T$) is given, then (7) is a requirement; on the other hand, given T, we can take (7) also as a definition of L (up to isomorphism) and look for a presentation of L, which then gives us a syntax and proof system of a logic for T-coalgebras.[1]

To summarise, we might say the whole point of the paper is to show that, once we presented a functor L satisfying (7), everything else flows from this: syntax and proof system are determined by the presentation and the semantics is determined by (7). This also means that the approach presented in the next section is not restricted to gKPFs.

3 The Goldblatt-Thomason theorem for coalgebras

To clarify the relationship between L-algebras and T-coalgebras in diagram (2), we review the categorical analysis given in [17], before returning the special case of Boolean algebras and sets.

3.1 Algebras and coalgebras on Ind- and Pro-completions

The general picture[2] underlying the situation discussed in the introduction is

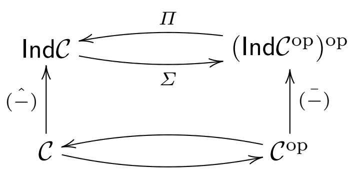

(8)

where $\mathcal{C}$ is a finitely complete and cocomplete category, $\mathsf{Ind}\mathcal{C}$ is the full subcategory of $\mathsf{Set}^{\mathcal{C}^{\mathrm{op}}}$ of finite limit preserving functors, $(\hat{-})$ and $(\bar{-})$ are the Yoneda embeddings. It is well-known that, under these assumptions, $\mathsf{Ind}\mathcal{C}$ is the completion of $\mathcal{C}$ with filtered colimits, see eg [14, Chapter VI]. Dually,

$$\mathsf{Pro}\mathcal{C} \stackrel{\mathrm{def}}{=} (\mathsf{Ind}\mathcal{C}^{\mathrm{op}})^{\mathrm{op}}$$

is the completion of $\mathcal{C}$ with cofiltered limits. Furthermore, we let Σ be the left Kan-extension of $(\bar{-})$ along $(\hat{-})$, and Π the right Kan-extension of $(\hat{-})$ along $(\bar{-})$ (in particular, $\Sigma\hat{C} \cong \bar{C}$, $\Pi\bar{C} \cong \hat{C}$). Σ is left adjoint to Π.

Example 3.1. 1. $\mathcal{C} = \mathsf{BA}_\omega$ (finite Boolean algebras = finitely presentable Boolean algebras), $\mathsf{Ind}\mathcal{C} = \mathsf{BA}$, $\mathsf{Pro}\mathcal{C} = \mathsf{Set}^{\mathrm{op}}$. ΣA is the set of ultrafilters over A and Π is (contravariant) powerset.

[1] For a general definition of 'presentation of a functor' and how presentations give rise to logics see [6]. Further investigations can be found in [17] showing, for example, that an endofunctor on BA has a finitary presentation iff it preserves filtered colimits.

[2] Actually, the general picture is even more general, see Lawvere [19, Section 7], an interesting special case of which is investigated in [18,21].

2. $\mathcal{C} = \mathsf{DL}_\omega$ (finite distributive lattices = finitely presentable distributive lattices), $\mathsf{Ind}\mathcal{C} = \mathsf{DL}$, $\mathsf{Pro}\mathcal{C} = \mathsf{Poset}^{\mathrm{op}}$. ΣA is the set of prime filters over A and Π gives the set of upsets.
3. In fact, (8) can be instantiated with any locally finite variety for $\mathsf{Ind}\mathcal{C}$. (A variety is locally finite if finitely generated free algebras are finite.)

We are interested in coalgebras over $(\mathsf{Ind}\mathcal{C}^{\mathrm{op}})$, ie, algebras over $\mathsf{Pro}\mathcal{C} = (\mathsf{Ind}\mathcal{C}^{\mathrm{op}})^{\mathrm{op}}$. Consider

$$L \circlearrowright \mathsf{Ind}\mathcal{C} \underset{\Sigma}{\overset{\Pi}{\leftrightarrows}} \mathsf{Pro}\mathcal{C} \circlearrowleft T \tag{9}$$

where we assume that L and T agree on $\mathcal{C}$, that is,

$$L\Pi\bar{C} \cong L\hat{C} \cong \Pi T\bar{C} \qquad \Sigma L\hat{C} \cong T\bar{C} \cong T\Sigma\hat{C} \tag{10}$$

Example 3.2. For $\mathsf{Ind}\mathcal{C} = \mathsf{BA}$ and $\mathsf{Pro}\mathcal{C} = \mathsf{Set}^{\mathrm{op}}$, the gKPF T and the $L = L_T$ satisfy (10) by Proposition 2.8.

Remark 3.3. We will usually denote by the same symbol a functor and its dual, writing eg $T : \mathcal{K} \to \mathcal{K}$ and $T : \mathcal{K}^{\mathrm{op}} \to \mathcal{K}^{\mathrm{op}}$.

In order to lift Π and Σ to algebras, we extend the natural isomorphisms (10) from $\mathcal{C}$ to $\mathsf{Ind}\mathcal{C}$ and $\mathsf{Pro}\mathcal{C}$, respectively. As a result of the procedure below, the lifted $L\Pi \to \Pi T$ and $T\Sigma \to \Sigma L$ will in general not be isomorphisms, the second may even fail to be natural.

The natural transformation $\delta : L\Pi \to \Pi T$. ΠX is a filtered colimit $\hat{C}_i \to \Pi X$. If L preserves filtered colimits we therefore obtain $L\Pi \to \Pi T$ as in

$$\begin{array}{ccccc} \Pi X & \qquad & L\Pi X & \xrightarrow{\delta_X} & \Pi T X \\ \uparrow{\scriptstyle c_i} & & \uparrow{\scriptstyle Lc_i} & & \uparrow{\scriptstyle \Pi T c_i^\sharp} \\ \hat{C}_i & & L\hat{C}_i & \xrightarrow{=} & \Pi T \Sigma \hat{C}_i \end{array} \tag{11}$$

where $c_i^\sharp : \Sigma\hat{C}_i \to X$ is the transpose of $c_i : \hat{C}_i \to \Pi X$. δ allows us to lift Π to a functor

$$\mathsf{Alg}(L) \xleftarrow{\tilde{\Pi}} \mathsf{Coalg}(T)^{\mathrm{op}} \tag{12}$$

mapping a T-algebra (X, ξ) to the L-algebra $(\Pi X, \xi \circ \delta_X)$.

Example 3.4. For $\mathsf{Ind}\mathcal{C} = \mathsf{BA}$, $\mathsf{Pro}\mathcal{C} = \mathsf{Set}^{\mathrm{op}}$, and T being one of $\mathcal{P}$ or $\mathcal{H}$, δ has been given explicitly in Definition 2.6.

The transformation $h : T\Sigma \to \Sigma L$. We will need that there exists h such that the following diagram commutes in $\mathsf{Pro}\mathcal{C}$ (where the d_k are the filtered colimit approximating A).

$$
\begin{array}{ccccc}
A & \quad & T\Sigma A & \xrightarrow{h_A} & \Sigma L A \\
{\scriptstyle d_k}\uparrow & & \uparrow & & \uparrow \\
\hat{A}_k & & T\Sigma \hat{A}_k & \xrightarrow{=} & \Sigma L \hat{A}_k
\end{array}
\tag{13}
$$

Remark 3.5. A sufficient condition for the existence of h is that T weakly preserves filtered colimits in $\mathsf{Pro}\mathcal{C}$, or, equivalently, weakly preserves cofiltered limits in $(\mathsf{Pro}\mathcal{C})^{\mathrm{op}}$. If T preserves these limits (non weakly) then h is natural.

Example 3.6. For gKPFs excluding $\mathcal{H}$, the maps h have been described by Jacobs [13, Definition 5.1]. We detail the definitions of the following to cases.

1. $h_A : \Sigma L_{\mathcal{P}} A \to \mathcal{P}\Sigma A$ maps $v \in \Sigma L_{\mathcal{P}} A$ to $\{u \in \Sigma A \mid \Box a \in v \Rightarrow a \in u\}$.
2. $h_A : \Sigma L_{\mathcal{H}} A \to \mathcal{H}\Sigma A$ maps $v \in \Sigma L_{\mathcal{H}} A$ to $\{\hat{a} \in 2^{\Sigma A} \mid \Box a \in v\}$.

Remark 3.7. There is a systematic way of calculating h from δ. For $A \in \mathcal{C}$, denoting the unit and counit of the adjunction $\Sigma \dashv \Pi$ by η and ε, h_A is given in $(\mathsf{Pro}\mathcal{C})^{\mathrm{op}}$ (thinking of Set) by

$$\Sigma L A \overset{(\Sigma L \eta_A)^\circ}{\longrightarrow} \Sigma L \Pi \Sigma A \overset{(\Sigma \delta_{\Sigma A})^\circ}{\longrightarrow} \Sigma \Pi T \Sigma A \overset{(\varepsilon_{T\Sigma A})^\circ}{\longrightarrow} T\Sigma A \tag{14}$$

Here we use that the arrows above are isos and we can take their inverse, denoted by $^\circ$. The calculations showing that Example 3.6 derives directly from (14) are detailed[3] in the appendix.

In general, h_A is not uniquely determined by (13) and we cannot assume it to be natural. Nevertheless, in the cases we are aware of h is natural.

Proposition 3.8. *For gKPFs T, the map*

$$h : \Sigma L_T \to T\Sigma$$

in Set *is natural.*

Proof. The type constructors $K_C, +, \times$ preserve cofiltered limits, hence the corresponding map h defined by (13) is uniquely determined and therefore natural. In the other two cases, $T = \mathcal{P}$ and $T = \mathcal{H}$, we take Example 3.6 as the definition of h and verify that it is natural and satisfies (13). We detail this for $T = \mathcal{H}$. Note first that $h_A : \Sigma L A \to \mathcal{H}\Sigma A$ is $\nu_{\Sigma A} \circ in_A^{-1}$ where in_A is the insertion of generators $A \to LA, a \mapsto \Box a$, and $\nu_X : X \to \mathcal{H}X$ maps x to $\{a \subseteq X \mid x \in a\}$. Now both the commutativity of (13) and the naturality of h follow from naturality of in and ν. $\Box$

[3] We hope these calculations show that *isomorphisms do work*. This balances Conor McBride's view, from a programming perspective, that *isomorphisms cost*.

To finish the category theoretic part of our development, we note that h allows us to lift Σ to

$$\mathsf{Alg}(L) \xrightarrow{\tilde{\Sigma}} \mathsf{Coalg}(T)^{\mathrm{op}} \tag{15}$$

via $(LA \to A) \mapsto (\Sigma A \to \Sigma LA \to T\Sigma A)$. If h is natural, then this map is a functor.

3.2 The Goldblatt-Thomason theorem for coalgebras

We used the general categorical framework to clarify the relationship between the functors T and L. We will now return to the special case discussed in the introduction. In particular, $\mathsf{Ind}\mathcal{C} = \mathsf{BA}$ and $\mathsf{Ind}\mathcal{C}^{\mathrm{op}} = \mathsf{Set}$; $\Pi : \mathsf{Set} \to \mathsf{BA}$ maps X to 2^X and $\Sigma : \mathsf{BA} \to \mathsf{Set}$ maps a Boolean algebra A to the set of ultrafilters over A.

We say that a functor $T : \mathsf{Set} \to \mathsf{Set}$ *preserves finite sets* if T maps finite sets to finite sets.

Definition 3.9 (modal logic of a functor). *The modal logic of a functor* $T : \mathsf{Set} \to \mathsf{Set}$ *is the pair* $(L, \delta : L\Pi \to \Pi T)$ *where* $L = \Pi T \Sigma$ *on finite Boolean algebras and* L *is continuously extended to all of* BA. δ *is then given as in (11).*

Remark 3.10. 1. The definition of L does not require T to preserve finite sets. This condition, which implies the right-hand side of (10), is needed for h in (13).
2. For gKPFs T, the modal logic corresponding to (L, δ) has been described explicitly in Example 2.4. But we know from [17] that any $L : \mathsf{BA} \to \mathsf{BA}$ arising from Definition 3.9 has such a presentation by modal operators and axioms.

The notion of a modal theory now arises from the initial, or free, L-algebra, see diagram (6).

Definition 3.11 (modal theory). *Consider a functor* $T : \mathsf{Set} \to \mathsf{Set}$ *with its associated modal logic* (L, δ) *and a* T*-coalgebra* (X, ξ).

1. *Let* I *be the initial* L*-algebra and* $[\![-]\!] : I \to \Pi(X, \xi)$ *be the unique morphism. Then the* variable-free modal theory *of* (X, ξ) *is* $\{\varphi \in I \mid [\![\varphi]\!] = X\}$.
2. *Let* I_P *be the free* L*-algebra over the free Boolean algebra generated by a countable set* P *of propositional variables. Let* $[\![-]\!]_v : I_P \to \Pi(X, \xi)$ *be the unique morphism extending a valuation* $v : P \to \Pi X$ *of the propositional variables. Then the* modal theory *of* (X, ξ) *is* $\{\varphi \in I_P \mid [\![\varphi]\!]_v = X$ *for all* $v : P \to \Pi X\}$.

The next proposition provides the first main ingredient to the Goldblatt-Thomason theorem, namely that modally definable classes 'reflect' ultrafilter extensions. In case of variable-free theories, definable classes are also closed under ultrafilter extensions.

Proposition 3.12. *Let $T : \mathsf{Set} \to \mathsf{Set}$ preserve finite sets and assume that h as in (13) exists. Then*

1. *(X, ξ) and $\Sigma\Pi(X, \xi)$ have the same variable-free modal theory,*
2. *(X, ξ) satisfies the modal theory of $\Sigma\Pi(X, \xi)$.*

Proof. (1): By construction of the logic from L, a formula φ is an element of the initial L-algebra and $(X, \xi) \models \varphi$ iff the unique morphism $[\![-]\!]$ from the initial L-algebra to $\Pi(X, \xi)$ maps φ to $X \in \Pi(X, \xi) = 2^X$. Therefore[4], to show that $(X, \xi) \models \varphi \;\Leftrightarrow\; \Pi(X, \xi) \models \varphi$, it suffices to establish that the map $\iota : \Pi(X, \xi) \to \Pi\Sigma\Pi(X\xi)$ is an injective algebra morphism. This follows from (the proof of) Theorem 5.3 in [17] and Stone's representation theorem for Boolean algebras. (2): Suppose there is a valuation v showing that φ does not hold in (X, ξ), that is, $[\![\varphi]\!]_v \neq X$. Then $\iota \circ [\![-]\!]_v(\varphi) \neq \Sigma\Pi X$, that is, there is a valuation showing that φ does not hold in $\Sigma\Pi(X, \xi)$. □

The second main ingredient (of the algebraic proof) of the Goldblatt-Thomason theorem is Birkhoff's variety theorem stating that a class of algebras is definable by equations iff it is closed under homomorphic images (H), subalgebras (S), and products (P). A set of equations is called *ground* if it does not contain any variables. This corresponds to the absence of propositional variables in a modal theory, or, in other words, to treating Kripke models instead of Kripke frames. The lesser expressivity of ground equations is reflected algebraically by also closing under embeddings (E). Closure under H, S, P (and E) is equivalent to closure under HSP ($EHSP$).

Theorem 3.13 (Birkhoff's variety theorem). *A class of algebras is definable by a set of*

1. *ground equations iff it is closed under $EHSP$,*
2. *equations iff it is closed under HSP.*

Proof. We sketch the proof of the less well-known 1st statement. It is routine to check that a definable class of algebras enjoys the required closure properties. Conversely, let $\mathcal{K}$ be a class of algebras closed under $EHSP$ and let Φ be the ground theory of $\mathcal{K}$. Consider an algebra A with $A \models \Phi$. We have to show that $A \in \mathcal{K}$. Since $\mathcal{K}$ is closed under SP, the quotient $Q = I/\Phi$ of the initial algebra I by Φ is in $\mathcal{K}$. $A \models \Phi$ then means that there is a morphism $Q \to A$, hence $A \in \mathcal{K}$ by closure under EH.

The dual of closure under S and E is closure under quotients and domains of quotients. This is equivalent to closure under 'co-spans' $(X, \xi) \twoheadrightarrow \bullet \twoheadleftarrow (X', \xi')$, or bisimilarity:

Definition 3.14 (bisimilarity). *Two coalgebras (X, ξ), (X', ξ') are bisimilar if there are surjective coalgebra morphisms $(X, \xi) \twoheadrightarrow \bullet \twoheadleftarrow (X', \xi')$.*

[4] Note that the top-element X of the BA $\Pi(X, \xi)$ is preserved by algebra morphisms.

We can now generalise to coalgebras the Goldblatt-Thomason theorem [11] for Kripke frames and Venema's corresponding result for Kripke models [22]. For a textbook account of the former see [3,4]. [4, Theorem 5.54] gives an excellent account of the algebraic proof that we generalise, [4, Theorem 3.19] presents an alternative model theoretic proof, and [4, Theorem 2.75] gives a version for pointed Kripke models.

We say that a class $\mathcal{K}$ of coalgebras is *closed under ultrafilter extensions* if $(X,\xi) \in \mathcal{K} \;\Rightarrow\; \Sigma\Pi(X,\xi) \in \mathcal{K}$ and that it *reflects ultrafilter extensions* if $\Sigma\Pi(X,\xi) \in \mathcal{K} \;\Rightarrow\; (X,\xi) \in \mathcal{K}$.

The first part of the theorem below is the definability result for coalgebras as generalisations of Kripke models, the second part treats definability for coalgebras as generalisations of Kripke frames. The difference in the formulation, apart from replacing closure under bisimilarity by closure under quotients, is due to the fact that all modally definable classes of Kripke models but not all modally definable classes of Kripke frames are closed under ultrafilter extensions (compare the two items of Proposition 3.12).

Theorem 3.15. *Let $T : \mathsf{Set} \to \mathsf{Set}$ preserve finite sets and assume there is a natural transformation h satisfying (13).*

1. *A class $\mathcal{K} \subseteq \mathsf{Coalg}(T)$ is definable by a variable-free modal theory iff $\mathcal{K}$ is closed under subcoalgebras, bisimilarity, coproducts and ultrafilter extensions and $\mathcal{K}$ reflects ultrafilter extensions.*
2. *A class $\mathcal{K} \subseteq \mathsf{Coalg}(T)$ closed under ultrafilter extensions is definable by a modal theory iff $\mathcal{K}$ is closed under subcoalgebras, quotients and coproducts and $\mathcal{K}$ reflects ultrafilter extensions.*

Proof. (1): For 'if' let X be a coalgebra that is a model of the theory of $\mathcal{K}$, that is, by Theorem 3.13.1, $\Pi X \in EHSP(\Pi\mathcal{K})$ where $\Pi\mathcal{K} = \{\Pi Y \mid Y \in \mathcal{K}\}$. We have to show $X \in \mathcal{K}$. ΠX embeds a quotient of a subalgebra of a product $\prod_i \Pi(X_i)$, $X_i \in \mathcal{K}$. Since Π is right adjoint, we obtain $\prod_i \Pi(X_i) \cong \Pi(\coprod_i X_i)$. Since Σ maps injective maps to surjective maps and vice versa, we have

$$\Sigma\Pi X \twoheadrightarrow \bullet \hookrightarrow \bullet \twoheadleftarrow \Sigma\Pi(\coprod X_i).$$

The stipulated closure properties now imply $X \in \mathcal{K}$. For 'only if', using that ground equationally definable classes of algebras are closed under $EHSP$, it is enough to observe (i) that Π maps surjective maps to injective maps and vice versa, (ii) that Π maps coproducts to products, (iii) Proposition 3.12.1.
(2): The proof is a straightforward variation of the previous one. For 'if' let X be a coalgebra that is a model of the theory of $\mathcal{K}$, that is, by Theorem 3.13.2, $\Pi X \in HSP(\Pi\mathcal{K})$ where $\Pi\mathcal{K} = \{\Pi Y \mid Y \in \mathcal{K}\}$. We have to show $X \in \mathcal{K}$. ΠX is a quotient of a subalgebra of a product $\prod_i \Pi(X_i)$, $X_i \in \mathcal{K}$. We have

$$\Sigma\Pi X \hookrightarrow \bullet \twoheadleftarrow \Sigma\Pi(\coprod X_i).$$

The stipulated closure properties now imply $X \in \mathcal{K}$. For 'only if', we use (i) and (ii) as in part 1 and (iii) Proposition 3.12.2. $\square$

Before deriving our main result as a corollary, let us analyse the hypotheses needed for Theorem 3.15 in terms of the general setting discussed in Section 3.1.

Remark 3.16. The following ingredients are used in the proof of Theorem 3.15.

1. $\mathcal{C}$ in diagram (8) has all finite limits and finite colimits. This is a strong requirement. But it holds if $\mathsf{Ind}\mathcal{C}$ is a locally finite variety and $\mathcal{C}$ is the subcategory of finitely presentable algebras. This includes BA and DL.
2. $A \to \Sigma\Pi A$ is injective. This is unlikely to hold in the generality of diagram (8) but it seems to be a rather mild requirement: For example, it holds for (subvarieties of) BA and DL.
3. T preserves finite sets (or, more generally, T restricts to $\mathcal{C}^{\mathrm{op}}$). This is needed in diagram 13. It excludes polynomial functors with infinite constants and the probability distribution functor. For a further discussion and the connection with strong completeness see [17].
4. h exists and is natural. The status of this requirement remains somewhat unclear. As emphasised by the corollary, it is satisfied in important examples. Let us note here that h is certainly natural if T preserves cofiltered limits. This is the case for all polynomial functors. The example we are aware of where the existence of h fails, is if T is the finite powerset functor (the ultrafilter extension of a Kripke frame is not finitely branching).

The main result of the paper is the following corollary. The second part generalises the Goldblatt-Thomason theorem from Kripke frames to all gKPF-coalgebras and the first part generalises Venema's definability theorem for Kripke models to all gKPF-coalgebras.

Corollary 3.17. *Let T be a gKPF.*

1. *A class $\mathcal{K} \subseteq \mathsf{Coalg}(T)$ is definable by a variable-free modal theory iff $\mathcal{K}$ is closed under subcoalgebras, bisimilarity, coproducts and ultrafilter extensions and $\mathcal{K}$ reflects ultrafilter extensions.*
2. *A class $\mathcal{K} \subseteq \mathsf{Coalg}(T)$ closed under ultrafilter extensions is definable by a modal theory iff $\mathcal{K}$ is closed under subcoalgebras, quotients and coproducts and $\mathcal{K}$ reflects ultrafilter extensions.*

Remark 3.18. 1. As far as we know, the special case of $\mathcal{H}$-coalgebras (neighbourhood frames) is a new result.

2. In the statement of the theorem, we can replace "$\mathsf{Coalg}(T)$" by a modally definable full subcategory of $\mathsf{Coalg}(T)$. For example, the theorem holds for monotone neighbourhood frames or topological neighbourhood frames. For topological spaces, the result is due to Gabelaia [9, Theorem 2.3.4], but see also ten Cate et al [7].

The original formulation of Venema's theorem [22] has closure under surjective bisimulations instead of closure under subcoalgebras and bisimilarity. The relationship between the two formulations is as follows. For functors T that preserve weak pullbacks, one can use 'spans' $(X, \xi) \twoheadleftarrow \bullet \twoheadrightarrow (X', \xi')$ in the definition of bisimilarity instead of co-spans $(X, \xi) \twoheadrightarrow \bullet \twoheadleftarrow (X', \xi')$ as above. Closure under subcoalgebras (or generated submodels in the parlance of [22]) is incorporated in the notion of surjective bisimulation by not forcing the left-hand projection of the span to be surjective: A class $\mathcal{K}$ is closed under *surjective bisimulations* iff for all $(X, \xi) \in \mathcal{K}$ and all $(X, \xi) \leftarrow \bullet \twoheadrightarrow (X', \xi')$ also $(X', \xi') \in \mathcal{K}$. Since $\mathcal{H}$ is the only ingredient of a gKPF that does not preserve weak pullbacks, we obtain the following generalisation of Venema's definability theorem for Kripke models.

Corollary 3.19. *Let T be a KPF (ie a gKPF built without using $\mathcal{H}$). A class $\mathcal{K} \subseteq \mathsf{Coalg}(T)$ is definable by a variable-free modal theory iff $\mathcal{K}$ is closed under coproducts and surjective bisimulations and $\mathcal{K}$ reflects ultrafilter extensions.*

4 Conclusion

The basic idea underlying this (and previous) work is to consider the logics for coalgebras as functors L on a category of propositional logics such as BA. Let us summarise some features of this approach.

- The functor L packages up all the information about modal operators and their axioms. In this way the functor L acts as an interface to the syntax, which is given by a presentation of L.
- As long as we only use abstract properties of L, such as (7), we can prove theorems about modal logics in a syntax free way, see Corollary 3.17 or the Jónsson-Tarski theorem [17, Thm 5.3] for examples. This gives rise not only to simpler proofs, but also to more general results.
- If we instantiate our abstract categories and functors with concrete presentations, we not only get back all the riches of syntax, but find that the categorical constructions actually do work for us. For example, in diagram (8), if we let $\mathcal{C} = \mathsf{BA}_\omega$ and $\mathcal{C}^{\mathrm{op}} = \mathsf{Set}_\omega$, then the fact that Π is contravariant powerset and Σ is ultrafilters, follows from the end/coend formula for right/left Kan extensions. Another example of this phenomenon is detailed in the appendix.

Another point is that the generality of our approach suggests further work. Let us mention the following:

- It is possible to replace Boolean algebras by distributive lattices. It could be of interest to look at the details.
- It should be possible to alleviate the restriction to finite constants insofar as infinite 'input' constants C as in T^{K_C} can be allowed. But the restriction to finite 'output' sets is important, see Friggens and Goldblatt [8] for a detailed discussion.
- Is it possible to generalise definability results for pointed models or frames in the same framework?
- It should be of interest to instantiate $\mathsf{Ind}\mathcal{C}$ in diagram (8) with other locally finite varieties.

– Diagram (8) can also be varied in many directions, for example, considering other completions than Ind or going to an enriched setting (for example, for posets (ie enrichment over 2) the Galois closed subsets of the adjunction $\Sigma \dashv \Pi$ describe the elements of the MacNeille completion of $\mathcal{C}$).

Finally, and from the point of view of logics for coalgebras, most importantly: Can we find a similarly nice and abstract account for functors T that do not preserve finite sets?

A Appendix

We show that the h in Example 3.6 are derived from (14). To this end we first state a lemma on ultrafilters, which is a straightforward consequence of the respective definitions.

Notation $\eta : Id \to \Pi\Sigma$ and $\varepsilon : Id \to \Sigma\Pi$ are the (co)unit of the adjunction $\Sigma \dashv \Pi$ (note that we wrote ε here as an arrow in Set and not in $\mathsf{Set}^{\mathrm{op}}$). f° denotes the converse of a bijection f. For $a \in A$ we abbreviate $\eta(a) = \{u \in \Pi\Sigma A \mid \ a \in u\}$ by $\hat{a}$. The complement $X \setminus S$ of a subset S of X is written as $-S$.

Lemma A.1. *Let A be a finite BA, Y a finite set and L one of $L_{\mathcal{P}}$, $L_{\mathcal{H}}$.*

1. *Every ultrafilter $u \in \Sigma A$ has a smallest element given by $\bigwedge_{a\in u} a$.*
2. *$(\varepsilon_Y)^\circ : \Sigma\Pi Y \to Y$ maps u to y, where $\{y\}$ is the smallest element of u.*
3. *Every ultrafilter $v \in \Sigma LA$ is determined by the set $\{\Box a \mid a \in A, \Box a \in v\}$.*
4. *The smallest element of $v \in \Sigma LA$ is given by $\bigwedge_{\Box a\in v} \Box a \wedge \bigwedge_{\Box a\notin v} \neg\Box a$.*

Also note that for isos f we have that $(\Sigma f)^\circ$ is the direct image map of f. We obtain for $L = L_{\mathcal{P}}$:

$(\Sigma L\eta_A)^\circ$ maps v to the ultrafilter determined by $\{\Box\hat{a} \mid \Box a \in v\}$,
$(\Sigma\delta_{\Sigma A})^\circ$ maps $\{\Box\hat{a} \mid \Box a \in v\}$ to $\{\{t \in \mathcal{P}\Sigma A \mid t \subseteq \hat{a}\} \mid \Box a \in v\}$,
$(\varepsilon_{\mathcal{P}\Sigma A})^\circ$ maps the ultrafilter determined by $\{\{t \in \mathcal{P}\Sigma A \mid t \subseteq \hat{a}\} \mid \Box a \in v\}$ to $\bigcap_{\Box a\in v}\{t \in \mathcal{P}\Sigma A \mid t \subseteq \hat{a}\} \cap \bigcap_{\Box a\notin v} -\{t \in \mathcal{P}\Sigma A \mid t \subseteq \hat{a}\}$.

It is now a straightforward verification to check that this last set contains exactly one t which is $\{u \in \Sigma A \mid \Box a \in v \Rightarrow a \in u\}$. Hence we obtained the h described in Example 3.6.1.

For $L = L_{\mathcal{H}}$ we have:

$(\Sigma L\eta_A)^\circ$ maps v to the ultrafilter determined by $\{\Box\hat{a} \mid \Box a \in v\}$,
$(\Sigma\delta_{\Sigma A})^\circ$ maps $\{\Box\hat{a} \mid \Box a \in v\}$ to $\{\{t \in \mathcal{H}\Sigma A \mid \hat{a} \in t\} \mid \Box a \in v\}$,
$(\varepsilon_{\mathcal{H}\Sigma A})^\circ$ maps the ultrafilter determined by $\{\{t \in \mathcal{H}\Sigma A \mid \hat{a} \in t\} \mid \Box a \in v\}$ to $\bigcap_{\Box a\in v}\{t \in \mathcal{H}\Sigma A \mid \hat{a} \in t\} \cap \bigcap_{\Box a\notin v} -\{t \in \mathcal{H}\Sigma A \mid \hat{a} \in t\}$.

It is now a straightforward verification that this last set contains exactly one t which is $\{\hat{a} \in 2^{\Sigma A} \mid \Box a \in v\}$. Hence we obtained the h described in Example 3.6.2.

References

1. S. Abramsky. Domain theory in logical form. *Annals of Pure and Applied Logic*, 51, 1991.
2. S. Abramsky. A Cook's tour of the finitary non-well-founded sets. In *We Will Show Them: Essays in Honour of Dov Gabbay*. College Publications, 2005. Presented at BCTCS 1988.
3. J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, 1983.
4. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. CUP, 2001.
5. M. Bonsangue and A. Kurz. Duality for logics of transition systems. In *FoSSaCS'05*, LNCS 3441, 2005.
6. M. Bonsangue and A. Kurz. Presenting functors by operations and equations. In *FoSSaCS'06*, LNCS 3921, 2006.
7. B. ten Cate, D. Gabelaia, and D. Sustretov. Modal languages for topology: expressivity and definability. arXiv 0610357, 2006.
8. D. Friggens and R. Goldblatt. A modal proof theory for final polynomial coalgebras. *Theoret. Comput. Sci.*, 360, 2006.
9. D. Gabelaia. Modal definability in topology. Master's thesis, University of Amsterdam, 2001.
10. S. Ghilardi. An algebraic theory of normal forms. *Ann. Pure Appl. Logic*, 71, 1995.
11. R. I. Goldblatt and S. K. Thomason. Axiomatic classes in propositional modal logic. In *Algebra and Logic*. 1974.
12. H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. In *CMCS'04*, ENTCS 106, 2004.
13. B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theor. Inform. Appl.*, 35, 2001.
14. P. Johnstone. *Stone Spaces*. Cambridge University Press, 1982.
15. C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In *CMCS'04*, ENTCS 106, 2004.
16. C. Kupke, A. Kurz, and D. Pattinson. Ultrafilter extensions of coalgebras. In *CALCO'05*, LNCS 3629, 2005.
17. A. Kurz and J. Rosický. Strongly complete logics for coalgebras. July 2006. Submitted.
18. F. Lawvere. Metric spaces, generalized logic and closed categories. *Rendiconti del Seminario Matematico e Fisico di Milano*, XLIII, 1973. Also appeared as TAC reprints No. 1.
19. F. Lawvere. Taking categories seriously. *Revista Colombiana de Matemáticas*, XX, 1986. Also appeared as TAC reprints No. 8.
20. M. Rößiger. Coalgebras and modal logic. In *CMCS'00*, volume 33 of *ENTCS*, 2000.
21. J. Rutten. Weighted colimits and formal balls in generalized metric spaces. *Topology and its Applications*, 89, 1998.
22. Y. Venema. Model definability, purely modal. In *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*. Amsterdam University Press, 1999.