Model constructions for Moss' coalgebraic logic

Jort Bergfeld[1] * and Yde Venema[2]**

[1] Artificial Intelligence and Cognitive Engineering, Rijksuniversiteit Groningen, Nijenborgh 9, 9747 AG Groningen, the Netherlands.
[2] Institute for Logic, Language and Computation, Universiteit van Amsterdam, Science Park 904, 1098XH Amsterdam, the Netherlands.

Abstract. We discuss two model constructions related to the coalgebraic logic introduced by Moss. Our starting point is the derivation system $\mathbf{M}_T$ for this logic, given by Kupke, Kurz and Venema. Based on the one-step completeness of this system, we first construct a finite coalgebraic model for an arbitrary $\mathbf{M}_T$-consistent formula. This construction yields a simplified completeness proof for the logic $\mathbf{M}_T$ with respect to the intended, coalgebraic semantics. Our second main result concerns a strong completeness result for $\mathbf{M}_T$, provided that the functor T satisfies some additional constraints. Our proof for this result is based on the construction, for an $\mathbf{M}_T$-consistent set of formulas A, of a coalgebraic model in which A is satisfiable.
Keywords: coalgebra, modal logic, completeness, finite model property, strong completeness

1 Introduction

Universal Coalgebra [16] provides the notion of a *coalgebra* as the natural mathematical generalization of state-based evolving systems such as streams, (infinite) trees, Kripke models, (probabilistic) transition systems, and many others. This approach combines simplicity with generality and wide applicability: many features, including input, output, nondeterminism, probability, and interaction, can easily be encoded in the coalgebra type, which in this paper we will take to be an endofunctor T on the category Set of sets as objects with functions as arrows. Logic enters the picture if one wants to specify and reason about *behavior*, one of the most fundamental notions admitting a coalgebraic formalization. With Kripke structures constituting key examples of coalgebras, it should come as no surprise that most coalgebraic logics are some kind of modification or generalization of *modal logic* [5].

This approach was initiated by Moss [13], who generalized the so-called 'cover modality' ∇_P from Kripke structures to coalgebras of arbitrary type T. The fascinating novelty of Moss' language is that his modality has a rather non-standard

* This work is part of the VIDI research programme with number 639.072.904, which is financed by the Netherlands Organisation for Scientific Research.
** The research of this author has been made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO).

arity: Moss' syntax specifies that $\nabla_T \alpha$ is a formula for all $\alpha \in T\mathcal{L}$ (where $\mathcal{L}$ is the collection of formulas), while its semantics is given by a categorical notion of *relation lifting* $\overline{T}$. This approach is completely uniform in the functor T, but as a drawback, for $\overline{T}$ to behave well T must satisfy the category-theoretic property of preserving weak pullbacks. In order to overcome the shortcomings of Moss' logic, Kurz [11], Pattinson [15] and others considered coalgebraic modal formalisms, that use standard syntax and work for coalgebras of arbitrary type. The success of this approach, in which the semantics of each modality is determined by a so-called *predicate lifting*, directed attention away from Moss' logic.

Interest in Moss' logic revived when it became clear that an approach based on his modality could have some advantages. In particular, some key results on the modal μ-calculus were obtained by Janin & Walukiewicz [8], on the basis of proofs that crucially involve a reconstruction of the classical modal language on the basis of the nabla modality (which they introduced, independently of Moss, as a primitive connective). Kupke & Venema [10] showed that many fundamental results in the area of (fixpoint) logic and automata theory could be lifted to the abstraction level of coalgebra.

Given the nonstandard syntax of Moss' language it was not a priori clear whether the collection of coalgebraic validities would allow nice derivation systems. As a first result, Palmigiano & Venema [14] gave a complete axiomatization for the cover modality ∇_P (i.e. in the case of Kripke frames). This calculus was streamlined into a formulation that admits a straightforward generalization to a calculus $\mathbf{M}_T$ for an arbitrary set functor T, by Bílková, Palmigiano & Venema [3], who also provided suitable Gentzen systems for the logic based on ∇_P. Kupke, Kurz & Venema [10] solved the outstanding problem by proving the soundness and completeness of the calculus $\mathbf{M}_T$ with respect to the coalgebraic semantics.

In this paper, which originated in the first author's MSc thesis [2] supervised by the second author, we continue the line of investigations of [10], taking their result on one-step soundness and completeness as our starting point. (As a minor difference with [10], we add explicit proposition letters to the language.) Our main contribution is two-fold. First, based on adapting ideas from Schröder [17] to the setting of Moss' logic, we provide a coalgebraic construction that, given an $\mathbf{M}_T$-consistent formula a, yields a finite model in which a is satisfied. As a corollary, we considerably simplify the second part of the completeness proof of [10] for the logic $\mathbf{M}_T$ with respect to its intended, coalgebraic semantics. Our second main result concerns a *strong* completeness result for $\mathbf{M}_T$, provided that the functor T restricts to finite sets and weakly preserves limits of surjective ω-cochains of finite sets. Our proof for this result is based on the quasi-canonical model method of Pattinson & Schröder [18].

2 Preliminaries

Categories and Coalgebras We assume familiarity with basic notions from category theory (such as categories, functors, and natural transformations), and from coalgebra. Here we fix some notation and terminology. We restrict attention to Set-based coalgebras, where Set denotes the category with sets as objects and functions as arrows.

Convention 1 *Throughout the paper we fix a functor $T : \mathsf{Set} \to \mathsf{Set}$, which we assume to preserve inclusions and weak pullbacks.*

The restriction that T preserves inclusions is for reasons of presentation only; we motivate the other restriction in Remark 1. Many (but not all) examples of coalgebraically interesting set functors fall in the scope of our work. We mention the inductively defined class *EKPF* of *extended Kripke polynomial functors* given as follows

$$T := Id \mid C \mid P \mid B_\omega \mid D_\omega \mid T_0 \circ T_1 \mid T_0 + T_1 \mid T_0 \times T_1 \mid T^D,$$

where C is an abitrary constant functor, P is power set, B_ω is finitary multiset, D_ω is finitary probability distribution and T^D is exponentiation with respect to an arbitrary set. An example of a functor that does not preserve weak pullbacks is the (monotone) neighborhood functor.

The *finitary version* $T_\omega : \mathsf{Set} \to \mathsf{Set}$ of T is given, on objects, by $T_\omega X := \bigcup\{TY \mid Y \subseteq X, Y \text{ finite }\}$, and on arrows by $T_\omega f := Tf$. It can be proved that T_ω also preserves inclusions and weak pullbacks. Given an object $\alpha \in T_\omega A$, we let $Base_A(\alpha)$ denote the smallest finite subset of A such that $\alpha \in TBase_A(\alpha)$; in fact, the family of operations $Base_A : T_\omega A \to P_\omega A$ constitutes a natural transformation $Base : T_\omega \dot{\to} P_\omega$ [7].

Definition 1. *A* T-coalgebra *is a pair (S, σ) where S is a set and $\sigma : S \to TS$ is a function. A* morphism of T-coalgebras *from (S, σ) to (S', σ'), written $f : (S, \sigma) \to (S, \sigma')$, is a function $f : S \to S'$ such that $Tf \circ \sigma = \sigma \circ f$.*

Relation lifting The coalgebraic semantics of Moss' coalgebraic language is based on the notion of *relation lifting* that we now briefly discuss (see [10] for more information). First we introduce some notation for relations and functions. The *graph* of a function $f : X \to X'$ is the relation $Grf := \{(x, f(x)) \in X \times X' \mid x \in X\}$. The *diagonal* relation on a set X is denoted as Id_X, and the *converse* of a relation R as $R^{\smile}$. Given subsets $Y \subseteq X$, $Y' \subseteq X'$, the *restriction* of R to Y and Y' is given as $R\restriction_{Y \times Y'} := R \cap (Y \times Y')$. The composition of two relations $R \subseteq X \times X'$ and $R' \subseteq X' \times X''$ is denoted by $R \,;\, R'$, whereas the composition of two functions $f : X \to X'$ and $f' : X' \to X''$ is denoted by $f' \circ f$. Thus, we have $Gr(f' \circ f) = Grf \,;\, Grf'$.

Definition 2. *[1] Given a binary relation $R \subseteq X_1 \times X_2$ with projection functions $\pi_i : R \to X_i$, we define its T-*lifting *$\overline{T}R \subseteq TX_1 \times TX_2$ as follows:*

$$\overline{T}R := \{((T\pi_1^R)\rho, (T\pi_2^R)\rho) \mid \rho \in TR\}.$$

Throughout the paper, we will use properties of the relation $\overline{T}R$; unless explicitly stated otherwise, these can always be derived by elementary means from the following fact.

Fact 2 (Properties of Relation Lifting) *The relation lifting $\overline{T}$ satisfies the following properties, for all functions $f : X \to X'$, all relations $R, S \subseteq X \times X'$, $R' \subseteq X' \times X''$, and all subsets $Y \subseteq X$, $Y' \subseteq X'$:*

1. *$\overline{T}$ extends T: $\overline{T}(Grf) = Gr(Tf)$;*
2. *$\overline{T}$ preserves the diagonal: $\overline{T}(Id_X) = Id_{TX}$;*
3. *$\overline{T}$ commutes with relation converse: $\overline{T}(R\breve{}) = (\overline{T}R)\breve{}$;*
4. *$\overline{T}$ is monotone: if $R \subseteq S$ then $\overline{T}(R) \subseteq \overline{T}(S)$;*
5. *$\overline{T}$ distributes over composition: $\overline{T}(R\,;R') = (\overline{T}R)\,;(\overline{T}R')$;*
6. *$\overline{T}$ commutes with restriction: $\overline{T}(R\restriction_{Y \times Y'}) = \overline{T}R\restriction_{TY \times TY'}$.*
7. *$\overline{T}_\omega$ coincides with $\overline{T}$: $\overline{T}_\omega R = (\overline{T}R)\restriction_{T_\omega X \times T_\omega X'}$.*

Remark 1. The main reason why we restrict our attention to coalgebra types T that preserve weak pullbacks is that for these functors, $\overline{T}$ distributes over relation composition (Fact 2(5)) [1, 19].

Applying relation lifting to the membership relation $\in$, we obtain an interesting operation. Given a set X, we let $\in_X \subseteq X \times PX$ denote the membership relation, restricted to X. We define the map $\lambda^T_X : TPX \to PTX$ by

$$\lambda^T_X(\Phi) := \{\alpha \in TX \mid \alpha\ \overline{T}{\in_X}\ \Phi\},$$

and call elements of $\lambda^T_X(\Phi)$ *lifted members* of Φ. Related to Fact 2(5) is that the family of maps $\lambda^T_X : TPX \to PTX$ constitutes a distributive laws of T over the monad P (see [10] for a discussion). Of more immediate importance is the following distributive law relative to the *contravariant power set functor* $\breve{P}$ [4].

Fact 3 *The family of maps λ^T provides a natural transformation $\lambda^T : T\breve{P} \dot{\to} \breve{P}T$.*

The following concept is needed in the axioms describing the interaction between ∇ and conjunctions.

Definition 3. *An object $\Phi \in TPX$ is a* redistribution *of $A \in PTX$ if $A \subseteq \lambda^T_X(\Phi)$. In case $A \in P_\omega T_\omega X$, we call a redistribution Φ* slim *if $\Phi \in T_\omega P_\omega(\bigcup_{\alpha \in A} Base(\alpha))$. The set of slim redistributions of A is denoted as $SRD(A)$.*

Fact 4 *[21] Given sets X, Y, a set $\Gamma \in PTX$ and a surjection $f : X \to Y$, we have*

$$\{TPf(\Phi) \mid \Phi \in SRD(\Gamma)\} = SRD(PTf(\Gamma)).$$

Propositional logic Given a set X, we define the set $\mathcal{L}_0(X)$ of propositional formulas over X by the following grammar:

$$a\ ::=\ x \mid \neg a \mid \bigwedge A \mid \bigvee A,$$

where $x \in X$, and $A \in P_\omega \mathcal{L}_0(X)$. That is, as the primitive connectives of our propositional language we take the unary symbol $\neg$ and the finitary meet and join symbols, $\bigwedge$ and $\bigvee$. We abbreviate $\bot := \bigvee \varnothing$ and $\top := \bigwedge \varnothing$.

Given sets X and S, an *X-valuation on S* is a map $V : X \to PS$; such a map can be naturally extended to a homomorphism $\widehat{V} : \mathcal{L}_0(X) \to PS$ by putting $\widehat{V}(\bigwedge A) := \bigcap\{V_0(a) \mid a \in A\}$, etc.

3 Moss' logic and its axiomatization

In this section we briefly recall the syntax and semantics of Venema's finitary version of Moss' coalgebraic logic [13, 20], and the axiomatization of its coalgebraic valid formulas, given by Kupke, Kurz and Venema [10].

3.1 Moss' logic

The finitary version $\mathcal{L}$ of Moss' language is defined as follows.

Definition 4. *Given a set* Prop *of variables, the set* $\mathcal{L}(\mathsf{Prop})$ *of Moss formulas in* Prop *is given by the following grammar:*

$$a \ ::= \ p \mid \neg a \mid \bigwedge A \mid \bigvee A \mid \nabla \alpha$$

where $p \in \mathsf{Prop}$*,* $A \in P_\omega \mathcal{L}$ *and* $\alpha \in T_\omega \mathcal{L}$*.*

Despite its unconventional appearance, the language $\mathcal{L}$ admits fairly standard definitions of most syntactical notions. For example, we may define the (finite!) set $Sfor(a)$ of *subformulas* of a by a straightforward formula induction, of which the only nonstandard clause concerns the nabla operator:

$$Sfor(\nabla\alpha) := \{\nabla\alpha\} \cup \bigcup_{a \in Base(\alpha)} Sfor(a).$$

The elements of $Base(\alpha) \subseteq Sfor(\nabla\alpha)$ will be called the *immediate* subformulas of $\nabla\alpha$. Given a formula a, we define the *single negation* of a formula a as $\sim\!a := b$ if $a = \neg b$ for some formula b, and as $\sim\!a := \neg a$ otherwise.

Since in this paper we will not only be dealing with formulas and sets of formulas, but also with elements of the sets $T_\omega\mathcal{L}$, $P_\omega T_\omega \mathcal{L}$ and $T_\omega P_\omega \mathcal{L}$, it will be convenient to use the following *naming convention*:

Set	Prop	$\mathcal{L}$	$T_\omega\mathcal{L}$	$P_\omega\mathcal{L}$	$P_\omega T_\omega\mathcal{L}$	$T_\omega P_\omega\mathcal{L}$
Elements	$p, q, \ldots$	$a, b, \ldots$	$\alpha, \beta, \ldots$	$A, B, \ldots$	$\Gamma, \Delta, \ldots$	$\Phi, \Psi, \ldots$

We may see the boolean connectives $\bigvee$ and $\bigwedge$ as maps from finite sets of formulas to formulas, $\bigvee, \bigwedge : P_\omega\mathcal{L} \to \mathcal{L}$. Applying the functor to these maps, we obtain functions $T\bigvee, T\bigwedge : T_\omega P_\omega \mathcal{L} \to T\mathcal{L}$. In particular, for any object $\Phi \in T_\omega P_\omega \mathcal{L}$, we obtain well-formed formulas of the form $\nabla(T\bigvee)\Phi$ and $\nabla(T\bigwedge)\Phi$.

Since we consider a version of Moss' language with proposition letters, in order to interpret this language we have to introduce valuations and models.

Definition 5. *A valuation on a* T*-coalgebra* (S, σ) *is a valuation* $V : \mathsf{Prop} \to PS$*; the induced structure* (S, σ, V) *will be called a* T-model. *For such a model, the satisfaction relation* $\Vdash_{\sigma,V} \subseteq S \times \mathcal{L}$ *is defined by the following induction on the complexity of formulas:*

$$\begin{array}{lll} s \Vdash_{\sigma,V} p & \textit{if} & s \in V(p), \\ s \Vdash_{\sigma,V} \neg a & \textit{if} & s \not\Vdash_{\sigma,V} a, \\ s \Vdash_{\sigma,V} \bigwedge A & \textit{if} & s \Vdash_{\sigma,V} a \textit{ for all } a \in A, \\ s \Vdash_{\sigma,V} \bigvee A & \textit{if} & s \Vdash_{\sigma,V} a \textit{ for some } a \in A, \\ s \Vdash_{\sigma,V} \nabla\alpha & \textit{if} & \sigma(s)\ \overline{T}\Vdash_{\sigma,V} \alpha. \end{array}$$

$$(\nabla 1)\ \frac{\{a \preccurlyeq b \mid (a,b) \in Z\}}{\nabla\alpha \preccurlyeq \nabla\beta}\ (\alpha,\beta) \in \overline{T}Z$$

$$(\nabla 2)\ \frac{\{\nabla(T\bigwedge)(\Phi) \preccurlyeq a \mid \Phi \in SRD(\Gamma)\}}{\bigwedge\{\nabla\alpha \mid \alpha \in \Gamma\} \preccurlyeq a}$$

$$(\nabla 3)\ \frac{\{\nabla\alpha \preccurlyeq a \mid \alpha\ \overline{T}{\in}\ \Phi\}}{\nabla(T\bigvee)(\Phi) \preccurlyeq a}$$

Table 1. Modal derivation rules of the system **M**

When no confusion is likely, we may write $\Vdash$ *instead of* $\Vdash_{\sigma,V}$. *If* $s \Vdash_{\sigma,V} a$ *we say that* a *is* true, *or* holds *at* s *in* $\mathbb{S}$, *and we may write* $\mathbb{S}, s \Vdash a$, *where* $\mathbb{S}$ *denotes the* T*-model* (S,σ,V).

Two important observations about Moss' logic are that it is *adequate* with respect to behavioral equivalence (or, equivalently, bisimilarity), and *expressive* when we confine attention to finitely branching coalgebras.

3.2 The derivation system M

When it comes to axiomatics, following [10] we find it convenient to take an approach based on derivation systems that manipulate equations, or rather, inequalities. An *inequality* is a pair consisting of two formulas a and b, usually written $a \preccurlyeq b$. Readers may think of this as the formula $a \to b$, as is obvious from the semantics.

Definition 6. *An inequality* $a \preccurlyeq b$ *is* valid, *notation:* $\models_T a \preccurlyeq b$, *if for every coalgebraic model* $\mathbb{S} = (S,\sigma,V)$, *and any* $s \in S$, *if* $\mathbb{S}, s \Vdash a$, *then* $\mathbb{S}, s \Vdash b$.

The following axiomatization for this logic was proved to be sound and complete in [10].

Definition 7. *The derivation system* **M** *is given by the derivation rules of Table 1, together with any complete set of axioms and rules (in inequational format) for classical propositional logic.*

Observe that unless T restricts to finite sets, **M** is an infinitary derivation system, in that the rules $(\nabla 2)$ and $(\nabla 3)$ may have infinitely many premisses. To get some intuitive understanding of this derivation system, we first note that $(\nabla 1)$ functions as a combined congruence and monotonicity rule. It has a side condition expressing that it may only be applied when the set of premisses is indexed by a relation Z such that (α,β) belongs to the lifted relation $\overline{T}Z$. Each of the other two rules should be seen as a distributive law (in the logical sense of the word). To see this, first consider the case that T preserve finiteness. Then we may replace the rules ($\nabla 2$) and ($\nabla 3$) with the following *axioms*:

$$\bigwedge \left\{ \nabla\alpha \mid \alpha \in \Gamma \right\} \preccurlyeq \bigvee \left\{ \nabla (T\bigwedge)\Phi \mid \Phi \in SRD(\Gamma) \right\} \qquad (\nabla 2_f)$$

$$\nabla (T\bigvee)\Phi \preccurlyeq \bigvee \left\{ \nabla\beta \mid \beta \,\overline{T}\!\in \Phi \right\} \qquad (\nabla 3_f)$$

Roughly speaking, ($\nabla 3_f$) expresses how ∇ distributes over disjunctions, while ($\nabla 2_f$) shows how a conjunction of nabla formulas can be rewritten as a disjunction of nabla formulas of conjunctions of the collection of immediate subformulas of the nabla formulas. If T does not restrict to finite sets, we may still think of ($\nabla 2$) and ($\nabla 3$) as these identities: the only problem is that the expressions on the right hand side of ($\nabla 2_f$) and ($\nabla 3_f$) may no longer denote properly defined formulas.

The notions of derivability with respect to this system is standard. A *derivation* is a well-founded tree, labelled with inequalities, such that the leaves of the tree are labelled with axioms of **M**, whereas with each parent node we may associate a derivation rule of which the conclusion labels the parent node itself, and the premisses label its children. If there is a derivation of the inequality $a \preccurlyeq b$, we write $\vdash_T a \preccurlyeq b$. A formula a is **M**-*consistent* if the inequality $a \preccurlyeq \bot$ is not derivable in **M**; a set A of formulas is consistent if the formula $\bigwedge A_0$ is consistent for each finite subset $A_0 \subseteq A$.

The following theorem is the main result of Kupke, Kurz & Venema [10]

Fact 5 (Soundness and Completeness of M) *[10] For each pair of formulas $a, b \in \mathcal{L}$:*

$$\models_T a \preccurlyeq b \text{ iff } \vdash_T a \preccurlyeq b.$$

The completeness proof in [10] proceeds in two steps. First the authors prove a so-called one-step completeness result for their system; then they apply Pattinson's stratification method, involving the terminal sequence of the functor T, to prove Fact 5. The construction given in our paper will provide a much simpler alternative for the second part of their proof.

3.3 One-step soundness and completeness

Given a set X, we define the set $\mathcal{L}_\nabla(X)$ of *rank-1 formulas* in X by putting

$$\mathcal{L}_\nabla(X) := \mathcal{L}_0\{\nabla\alpha \mid \alpha \in T_\omega \mathcal{L}_0 X\}.$$

It will sometimes be convenient to think of $\mathcal{L}_\nabla(X)$ as a propositional language, generated from the set $T^\nabla_\omega(X) := \{\nabla\alpha \mid \alpha \in T_\omega \mathcal{L}_0 X\}$ as proposition letters.

Any valuation $V : X \to PS$, interpreting elements of X as subsets of some set S, not only extends to a propositional meaning function $\widehat{V} : \mathcal{L}_0(X) \to PS$, it also induces an interpretation $\widetilde{V} : \mathcal{L}_\nabla(X) \to PTS$ of rank-1 formulas in X as subsets of TS. For the definition of $\widetilde{V}$, observe that the map $T\widehat{V} : T_\omega\mathcal{L}_0X \to TPS$ naturally yields a $T^\nabla_\omega(X)$-valuation $\lambda^T_S \circ T\widehat{V}$ on TS given by, for $\alpha \in T_\omega\mathcal{L}_0X$:

$$\nabla\alpha \mapsto \lambda^T_S(T\widehat{V}(\alpha)).$$

Then we define $\widetilde{V} := \widehat{\lambda^T \circ T\widehat{V}}$.

We may take the set PS itself as a collection of proposition letters; then the identity map on PS becomes a special PS-valuation on PS: the *identity valuation on S*, notation $i_S : PS \to PS$. We say that an $\mathcal{L}_0(PS)$-inequality $a \preccurlyeq b$ is a *true fact on PS*, notation: $\models^S_0 a \preccurlyeq b$, if $\widehat{i_S}(a) \subseteq \widehat{i_S}(b)$; an $\mathcal{L}_\nabla(PS)$-inequality $a \preccurlyeq b$ is *one-step valid*, notation: $\models^S_1 a \preccurlyeq b$, if $\widetilde{i_S}(a) \subseteq \widetilde{i_S}(b)$.

On the axiomatic side, we modify the derivation $\mathbf{M}$ into a *one-step derivation system* $\mathbf{M}^S$, which only uses $\mathcal{L}_0(PS)$ and $\mathcal{L}_\nabla(PS)$-formulas. More precisely, a $\mathbf{M}^S$-derivation is a well-founded tree, labelled with $\mathcal{L}_0(PS)$- and $\mathcal{L}_\nabla(PS)$-inequalities, such that (1) the leaves of the tree are labelled with true facts on PS, whereas (2) with each parent node we may associate a derivation rule of which (a) the conclusion is an $\mathcal{L}_\nabla(PS)$-inequality labelling the parent node itself, (b) the premisses label its children, and (c) these premisses are either all $\mathcal{L}_\nabla(PS)$-inequalities, or all $\mathcal{L}_0(PS)$-inequalities; in the latter case the children are all leaves and the derivation rule is (∇1). Hence if we do induction on the complexity of one-step derivations, we may assume that the base case is given by an application of rule (∇1). If there is such a one-step derivation of the inequality $a \preccurlyeq b$, we write $\vdash^S_1 a \preccurlyeq b$.

Fact 6 (One-step soundness and completeness) *[10] Given a set S, for each pair of rank-1 formulas $a, b \in \mathcal{L}_\nabla(PS)$:*

$$\models^S_1 a \preccurlyeq b \text{ iff } \vdash^S_1 a \preccurlyeq b.$$

4 A finite model construction

In this section we will give the main construction of this paper, serving to prove Theorem 7 below. As a corollary we obtain Fact 5, the Soundness and Completeness Theorem of [10].

Theorem 7. *Every consistent formula is satisfied in a finite T-coalgebra.*

Our construction is based on ideas from Schröder [17]. To give a rough idea, we need to introduce some terminology. We call a set of formulas *closed* if it is closed under taking subformulas and single negations ($\sim$). Given a closed set R of formulas, we call a subset $A \subseteq R$ an *R-atom* if A is a maximal consistent subset of R. Any R-atom A has the properties, for every $a \in R$, that $a \in A$ iff $\sim a \notin A$, and for every $a \wedge b \in R$, that $a \wedge b \in A$ iff both $a \in A$ and $b \in A$, etc. As usual it is straightforward to prove a *Lindenbaum Lemma* stating that every consistent subset of R can be extended to an R-atom.

Definition 8. *Given a formula c, let $C(c)$ denote the smallest closed set containing c, and define the closed set $R(c)$ by*

$$R(c) := \{ \bigvee_{A\in\mathcal{A}} \bigwedge A, \neg \bigvee_{A\in\mathcal{A}} \bigwedge A \mid \mathcal{A} \subseteq PC(c)\}.$$

We let $S(c)$ denote the set of $R(c)$-atoms.

Clearly $S(c)$ is a finite set. Hence, by the Lindenbaum Lemma, in order to prove Theorem 7, it suffices to build a model $(S(c), \sigma, V)$ on the set $S(c)$ for which we can prove a *Truth Lemma* stating that for all atoms/states $A \in S(c)$ and all formulas $a \in C(a)$:

$$a \in A \text{ iff } A \Vdash_{\sigma,V} a. \tag{1}$$

The proof of this Truth Lemma will proceed by a formula induction. It should be obvious how to define a valuation $V : \mathsf{Prop} \to S(c)$ ensuring (1) for atomic formulas a; the earlier mentioned properties of atoms takes care of the boolean cases of the inductive step of the proof. In order to prove the ∇-case of the induction, we have to come up with a proper definition of the coalgebra map $\sigma : S(c) \to TS(c)$. This definition will be crucially based on the one-step soundness and completeness (Fact 6).

Turning to the technicalities, we fix a consistent formula c, and write C, R and S instead of $C(c)$, $R(c)$ and $S(c)$. For technical reasons, it will be convenient to see formulas in R as separate proposition letters; formally we define

$$\underline{R} := \{\underline{b} \mid b \in R\},$$

and we assume the existence of a bijection $q : \underline{R} \to R$ given by $q(\underline{b}) := b$. In order to apply the one-step soundness and completeness, we link this set with PS by defining the valuation $j : \underline{R} \to PS$ as follows:

$$j(\underline{b}) := \{A \in S \mid q(\underline{b}) = b \in A\}.$$

It is straightforward to verify that j is surjective: For each $A \in S$ we have $j(\underline{\bigwedge(A \cap C)}) = \{A\}$ and for each $Z \subseteq S$ we have $j(\underline{\bigvee_{A\in Z} \bigwedge(A \cap C)}) = Z$. We extend j to a function $j_0 : \mathcal{L}_0\underline{R} \to \mathcal{L}_0 PS$ by inductively defining $j_0(\neg b) = \neg j_0(b)$, $j_0(\bigwedge B) = \bigwedge\{j_0(b') \mid b' \in B\}$ and $j_0(\bigvee B) = \bigvee\{j_0(b') \mid b' \in B\}$ for all $b \in \mathcal{L}_0\underline{R}$ and $B \subseteq \mathcal{L}_0\underline{R}$. We now lift j_0 to a function $j_1 : \mathcal{L}_\nabla\underline{R} \to \mathcal{L}_\nabla PS$, by putting $j_1(\nabla\beta) = \nabla T j_0(\beta)$, $j_1(\neg b) = \neg j_1(b)$, $j_1(\bigwedge B) = \bigwedge\{j_1(b') \mid b' \in B\}$ and $j_1(\bigvee B) = \bigvee\{j_1(b') \mid b' \in B\}$ for all $\nabla\beta \in \mathcal{L}_\nabla\underline{R}$, $b \in \mathcal{L}_\nabla\underline{R}$ and $B \subseteq \mathcal{L}_\nabla\underline{R}$.

In the same way we obtain q_0 and q_1 from q. Note that all these functions are surjective, however, q_0 and q_1 are not necessarily injective. For example take any $a \wedge b \in R$, then $q_0(\underline{a \wedge b}) = q_0(\underline{a} \wedge \underline{b}) = a \wedge b$.

Lemma 1 below, our main technical lemma, links **M**-derivations of formulas in $\mathcal{L}_0 R$ and $\mathcal{L}_\nabla R$ to, respectively, true facts on PS, and one-step derivations of formulas in $\mathcal{L}_\nabla PS$. The $\underline{R}$-formulas serve as a bridge between R-formulas and PS-formulas.

Lemma 1. *1. For $a, b \in \mathcal{L}_0\underline{R}$ we have $\vdash_{\mathbf{M}} q_0(a) \preccurlyeq q_0(b)$ iff $\models_0^S j_0(a) \preccurlyeq j_0(b)$.*
2. For $a, b \in \mathcal{L}_\nabla\underline{R}$ we have $\vdash_{\mathbf{M}} q_1(a) \preccurlyeq q_1(b)$ iff $\vdash_1^S j_1(a) \preccurlyeq j_1(b)$.

Proof. For part 1, first observe that by some routine propositional reasoning, we may reduce the problem to the case where $a = \bigwedge A$ and $b = \bigvee B$ for some $A, B \subseteq \underline{R}$. (To see this, note that $\bigvee A' \preccurlyeq b$ corresponds to the set $\{a' \preccurlyeq b \mid a' \in A'\}$, etc.)

First suppose that $\vdash_{\mathbf{M}} q_0(\bigwedge A) \preccurlyeq q_0(\bigvee B)$, then the set $\{q(a') \mid a' \in A\} \cup \{\neg q(b') \mid b' \in B\}$ is $\mathbf{M}$-inconsistent. We claim that if $\{q(a') \mid a' \in A\} \subseteq D$ for some $D \in S$, then $q(b') \in D$ for some $b' \in B$. (If not, then since D is an R-atom, we would obtain $\neg q(b') \in D$ for all $b' \in B$, contradicting the consistency of D.) Therefore we have $\widehat{i_S} j_0(\bigwedge A) = \bigcap_{a' \in A} j(a') \subseteq \bigcup_{b' \in B} j(b') = \widehat{i_S} j_0(\bigvee B)$, thus $j_0(\bigwedge A) \preceq j_0(\bigvee B)$ is a true fact on PS.

For the other direction, suppose that $\not\vdash_{\mathbf{M}} q_0(\bigwedge A) \preccurlyeq q_0(\bigvee B)$, then the set $\{q(a') \mid a' \in A\} \cup \{\neg q(b') \mid b' \in B\}$ is $\mathbf{M}$-consistent. By the Lindenbaum Lemma there exists a $D \in S$ extending this set. Thus $D \in \widehat{i_S} j_0(\bigwedge A) = \bigcap_{a' \in A} j(a')$, but $D \notin \widehat{i_S} j_0(\bigvee B) = \bigcup_{b' \in B} j(b)$. Therefore $\bigcap_{a' \in A} j(a') \nsubseteq \bigcup_{b' \in B} j(b')$ and thus $j_0(\bigwedge A) \preceq j_0(\bigvee B)$ is not a true fact on PS.

For part 2, we only consider the direction from right to left. (The other direction, which we do not need in the remainder, is proved similarly.) By induction on the complexity of one-step derivation trees we will show that any one-step derivation tree $\mathcal{D}$, of which the root is labelled with an inequality $j_1(a) \preccurlyeq j_1(b)$, can be transformed into an $\mathbf{M}$-derivation tree for the inequality $q_1(a) \preccurlyeq q_1(b)$.

$\boxed{\text{Base case: } \nabla 1}$ By definition of our one-step derivation tree, we may just as well assume that in the base case of our inductive proof we are dealing with an instance of the rule $\nabla 1$, of which the premisses are all true facts on PS. More precisely, in this case the conclusion $j_1(a) \preccurlyeq j_1(b)$ stems from some $a = \nabla\alpha$ and $b = \nabla\beta$ (where $\alpha, \beta \in T\mathcal{L}_0\underline{R}$) in the sense that $j_1(a) = \nabla T j_0(\alpha)$ and $j_1(b) = \nabla T j_0(\beta)$, and the last applied rule was

$$(\nabla 1)\ \frac{\{a' \preceq b' \mid (a', b') \in Z\}}{\nabla T j_0(\alpha) \preceq \nabla T j_0(\beta)}\ (T j_0(\alpha), T j_0(\beta)) \in \overline{T} Z.$$

for some relation $Z \subseteq \mathcal{L}_0 PS \times \mathcal{L}_0 PS$. In fact, given the properties of relation lifting, we may assume without loss of generality that

$$Z = \{(a', b') \in \mathit{Base}(T j_0(\alpha)) \times \mathit{Base}(T j_0(\beta)) \mid a' \preccurlyeq b' \text{ is a true fact }\}.$$

By the naturality of *Base* we have for all $\delta \in T\mathcal{L}_0\underline{R}$

$$\mathit{Base}(T j_0(\delta)) = \{j_0(d) \mid d \in \mathit{Base}(\delta)\}. \tag{2}$$

Now define

$$\hat{Z} := \{(a', b') \in \mathit{Base}(\alpha) \times \mathit{Base}(\beta) \mid q_0(a') \preccurlyeq q_0(b') \text{ is derivable}\},$$

then by equation (2) and by part 1, we have for all $a' \in Base(\alpha)$, $b' \in Base(\beta)$ that $(j_0(a'), j_0(b')) \in Z$ iff $(a', b') \in \hat{Z}$. From this it follows by the properties of relation lifting that for all $\alpha' \in TBase(\alpha)$, $\beta' \in TBase(\beta)$:

$$(Tj_0(\alpha'), Tj_0(\beta')) \in \overline{T}Z \text{ iff } (\alpha', \beta') \in \overline{T}\hat{Z}.$$

In particular, we obtain that $(\alpha, \beta) \in \overline{T}\hat{Z}$. Again using the properties of relation lifting we may conclude from this that

$$(Tq_0(\alpha), Tq_0(\beta)) \in \overline{T}\Big(\{(q_0(a'), q_0(b')) \mid (a', b') \in \hat{Z}\}\Big).$$

But then, since $q_1(a) = \nabla Tq_0(\alpha)$ and $q_1(b) = \nabla Tq_0(\beta)$, we can derive the inequality $q_1(a) \preccurlyeq q_1(b)$, as follows:

$$(\nabla 1)\ \frac{\{q_0(a') \preceq q_0(b') \mid (a', b') \in \hat{Z}\}}{\nabla Tq_0(\alpha) \preceq \nabla Tq_0(\beta)},$$

where all premisses are derivable by definition of $\hat{Z}$.

In the inductive step we make a case distinction; we only consider the cases where the last applied rule was $(\nabla 3)$.

$\boxed{\text{Inductive case: } (\nabla 3)}$ Suppose that $j_1(a)$ is of the form $\nabla(T\bigvee)(\Psi)$ for some $\Psi \in TP\mathcal{L}_0PS$, and that the last applied rule is:

$$(\nabla 3) \qquad \frac{\{\nabla\beta \preceq j_1(b) \mid \beta \overline{T}{\in} \Psi\}}{\nabla(T\bigvee)(\Psi) \preceq j_1(b)} \tag{3}$$

We claim that a is of the form $\nabla(T\bigvee)(\Phi)$ for some $\Phi \in T_\omega P_\omega \mathcal{L}_0\underline{R}$ such that $\Psi = TPj_0(\Phi)$. To see this, first observe that a must obviously be of the form $\nabla\alpha$ for some $\alpha \in T_\omega\mathcal{L}_0\underline{R}$; we will show that α is of the form $T\bigvee(\Phi)$ with Φ as above. For this purpose, note that by definition of j_0, if $a' \in \mathcal{L}_0\underline{R}$ is such that $j_0(a') = \bigvee A$ for some $A \in P_\omega\mathcal{L}_0PS$, then a' must be of the form $\bigvee B$ for some $B \in P_\omega\mathcal{L}_0\underline{R}$ with $A = \{j_0(b') \mid b' \in B\}$. This condition can be expressed as $Gr(j_0)\ ;\ (Gr\bigvee)^{\smile} \subseteq (Gr\bigvee)^{\smile}\ ;\ Gr(Pj_0)$. Then by the properties of relation lifting we find that $Gr(Tj_0)\ ;\ (GrT\bigvee)^{\smile} \subseteq (GrT\bigvee)^{\smile}\ ;\ Gr(TPj_0)$. From this the existence of the required object Ψ is immediate.

In order to find an **M**-derivation for the inequality $q_1(a) \preccurlyeq q_1(b)$, we calculate $q_1(a) = q_1(\nabla(T\bigvee)(\Phi)) = \nabla(T\bigvee)(TPq_0(\Phi))$. Aiming to derive $\nabla(T\bigvee)(TPq_0(\Phi))$ $\preccurlyeq q_1(b)$ via the rule $(\nabla 3)$, let β be an arbitrary lifted member of $TPq_0(\Phi)$. From $\beta \in \lambda^T(TPq_0(\Phi)) = PTq_0(\lambda^T(\Phi))$, we obtain that β is of the form $Tq_0(\alpha')$, for some lifted member α' of Φ. But from $\alpha' \in \lambda^T(\Phi)$ it follows that $Tj_0(\alpha') \in PTj_0(\lambda^T(\Phi)) = \lambda^T(TPj_0(\Phi)) = \lambda^T(\Psi)$. Now observe that $j_1(\nabla\alpha') = \nabla Tj_0(\alpha')$, and so the inequality $j_1(\nabla\alpha') \preccurlyeq j_1(b)$ is one of the premisses of (3). Thus by the inductive hypothesis, we have a **M**-derivation for the inequality $q_1(\nabla\alpha') \preccurlyeq q_1(b)$, which is nothing but $\nabla Tq_0(\alpha') \preccurlyeq q_1(b)$, that is, $\nabla\beta \preccurlyeq q_1(b)$.

It follows that we can derive $q_1(a) = \nabla(T\bigvee)(TPq_0(\Phi)) \preceq q_1(b)$ by

$$(\nabla 3) \qquad \frac{\{\nabla\beta \preceq q_1(b) \mid \beta\, \overline{T}{\in}\, TPq_0(\Phi)\}}{\nabla(T\bigvee)(TPq_0(\Phi)) \preceq q_1(b)}$$

On the basis of Lemma 1(2) we can prove the existence of a coalgebra map $\sigma : S \to TS$ with the right properties. Note that $jq^{-1}(b) = \{B \in S \mid b \in B\}$

Lemma 2 (Existence Lemma). *There is a map $\sigma : S \to TS$ such that for all atoms $A \in S$ and all formulas of the form $\nabla\alpha \in R$:*

$$\nabla\alpha \in A \text{ iff } \sigma(A)\, \overline{T}{\in}\, T(jq^{-1})(\alpha). \tag{4}$$

Proof. Suppose towards contradiction that for some $A \in S$ there is no $\sigma(A)$ that satisfies equation (4). Define $b := \bigwedge_{\nabla\alpha\in A} \nabla Tq^{-1}(\alpha) \wedge \bigwedge_{\neg\nabla\alpha\in A} \neg\nabla Tq^{-1}(\alpha)$. By assumption we have

$$\widetilde{i_S} j_1(b) = \bigcap_{\nabla\alpha\in A} \lambda^T(T(jq^{-1})(\alpha)) \cap \bigcap_{\neg\nabla\alpha\in A} TS \setminus \lambda^T(T(jq^{-1})(\alpha)) = \emptyset.$$

In other words, $\models^S_1 j_1(b) \preceq \bot$, and so we have $\vdash^S_1 j_1(b) \preceq \bot$ by one-step completeness. Then Lemma 1 provides an **M**-derivation of $q_0(b) \preceq \bot$, contradicting the consistency of A.

Lemma 3 (Truth Lemma). *Let (S, σ, V) be a model where S is the set of $R(c)$-atoms, σ is any map satisfying condition (4) of Lemma 2, and $V : \mathsf{Prop} \to PS$ is given by $V(p) := \{A \in S \mid p \in A\}$. Then (1) holds for all $a \in C$ and all $A \in S$.*

Proof. Via a straightforward induction on the complexity of a. We only discuss the case $a = \nabla\alpha$. By definition of the semantics we have $A \Vdash \nabla\beta$ iff $\sigma(A)\, \overline{T}{\Vdash}\, \beta$, and by Lemma 2 we have $\nabla\beta \in A$ iff $\sigma(A)\, \overline{T}{\in}\, T(jq^{-1})(\beta)$. So in order to finish the proof, it suffices to show that $\overline{T}{\Vdash} = \overline{T}{\in}\, ;\, Gr(T(jq^{-1}))^{\smile}$, or, equivalently, that (*) $\overline{T}{\Vdash} = \overline{T}{\in}\, ;\, Gr(Tj)^{\smile}\, ;\, Gr(Tq)$. But inductively, if we restrict to formulas $b \in R$ of smaller complexity than $\nabla\alpha$, we have that $jq^{-1}(b) = \{B \in S \mid b \in B\} = \{B \in S \mid B \Vdash b\}$. This means that $\Vdash\, = \,\in\, ;\, Gr(j)^{\smile}\, ;\, Gr(q)$ so that (*) directly follows by the properties of relation lifting.

The proof of Theorem 7 is now straightforward. If c is a consistent formula, it belongs to some R-atom A by the Lindenbaum Lemma. Then by the Existence Lemma and the Truth Lemma, we can endow the (finite!) set S of R-atoms with a coalgebra structure σ and a valuation V such that $A \Vdash_{\sigma,V} a$.

5 Strong completeness

In this section we will prove strong completeness of the axiom system **M**. That is, we will prove that, given some restrictions on the functor T, every **M**-consistent set of formulas is satisfiable. It might be possible to see our strong completeness result as a special case of Theorem 8.1 in Kurz & Rosický [12] (see also [9]). Nevertheless, we believe our short, direct proof, which follows the ideas of Pattinson & Schröder [18], to be of value.

Our purpose is to endow the set S of maximal $\mathbf{M}$-consistent sets of formulas with a coalgebra structure $\sigma : S \to TS$ and a valuation V such that for all $A \in S$ we can prove the following Truth Lemma, stating that for all formulas a:

$$a \in A \text{ iff } A \Vdash_{\sigma,V} a. \tag{5}$$

The idea underlying the construction of σ is that, for each $A \in S$, we may *approximate* $\sigma(A)$ by considering finite versions of S. For this purpose, enumerate $\mathsf{Prop} = \{p_i \mid i \in \omega\}$, and define $\mathsf{Prop}_n := \{p_i \mid 0 \leq i < n\}$. Let $\mathcal{L}_0 = \mathcal{L}_0(\varnothing)$ and let $\mathcal{L}_n$ denote the set $\mathcal{L}_0(\mathsf{Prop}_{n+1} \cup \{\nabla\alpha \mid \alpha \in T\mathcal{L}_n\})$. Then clearly $\mathcal{L}(\mathsf{Prop}) = \bigcup_{n\in\omega} \mathcal{L}_n$. Let S_n be the set of $\mathcal{L}_n$-atoms. It is not hard to show that if T restricts to finite sets, then each $\mathcal{L}_n$ is finite modulo equivalence, whence each S_n is finite. We let $h_n : S_{n+1} \to S_n$ and $\pi_n : S \to S_n$ be defined by $h_n(A) := A \cap \mathcal{L}_n$ and $\pi_n(A) := A \cap \mathcal{L}_n$. By the Lindenbaum Lemma all h_n and π_n are surjective.

On the basis of the results in the previous section we can prove the following lemma.

Lemma 4. *For each maximal consistent set $A \in S$ there is a family $(\tau_n)_{n\in\omega}$, with $\tau_n \in TS_n$, and such that for all n:*

$$(Th_n)\tau_{n+1} = \tau_n, \tag{6}$$

and for all $\alpha \in T\mathcal{L}_n$ it holds that

$$\nabla\alpha \in \pi_{n+1}(A) \text{ iff } \alpha \,\overline{T}\!\in \tau_n. \tag{7}$$

Note that (7) requires a relation between elements of S_{n+1} and objects, not in TS_{n+1}, but in TS_n.

Proof. (Sketch) For each n let $A_n \in S_n$ denote the atom $\pi_n(A)$. Using the methods of the previous section it is straightforward to show that for the atom $A_n \in S_n$ there is an object $\rho \in TS_n$ which *works for* A_n in the sense that for all $\alpha \in T\mathcal{L}_{n-1}$ it holds that $\nabla\alpha \in A_n$ iff $\alpha \,\overline{T}\!\in \rho$. It is not hard to show that if $\rho \in TS_{n+1}$ works for A_{n+1} then $(Th_n)\rho$ works for A_n. Consider the tree with nodes $N := \bigcup_{n\in\omega}\{\rho \in TS_n \mid \rho \text{ works for } A_n\}$, and edge relation E given by $\rho E \rho'$ iff $\rho = (Th_n)\rho'$ for some n. By König's Lemma this tree has an infinite path $(\tau_n)_{n\in\omega}$, and it is a routine exercise to verify that this family satisfies the required properties.

The point of considering the sequence $(\tau_n)_{n\in\omega}$ is that, under some condition on T, they *approximate* some object $\tau \in TS$, that we can take for our $\sigma(A)$. To formulate this condition, we define a *surjective ω-cochain of finite sets* to be a sequence $(X_n)_{n\in\omega}$ of finite sets, with surjections $h_n : X_{n+1} \to X_n$ that are called *projections*. In Set, such a diagram has a limit X with limit projections $\pi_n : X \to X_n$. Clearly also in Set each endofunctor T transforms a surjective

ω-cochain of sets into a surjective ω-cochain of sets. Now, to say that T *weakly* preserve limits of these diagrams means that whenever we have a diagram as above, with limit $(X, (\pi_n)_{n\in\omega})$, the set TX with the maps $T\pi_n : TX \to TX_n$ is a *weak* limit of the diagram $((TX_n)_{n\in\omega}, (p_n)_{n\in\omega})$. An equivalent requirement is that for each so-called coherent family $(\tau_n \in TX_n)_{n\in\omega}$ (that is, satisfying (6) for all n), there is a (not necessarily unique) element $\tau \in TX$ such that $T\pi_n(\tau) = \tau_n$ for all n.

On the basis of Lemma 4 we can now prove the following.

Lemma 5. *Let T restrict to finite sets and weakly preserve limits of surjective ω-cochains of finite sets. Then there is a coalgebra map $\sigma : S \to TS$ and a valuation $V : \mathsf{Prop} \to PS$ such that for all $a \in \mathcal{L}(\mathsf{Prop})$, and all $A \in S$, the Truth Lemma* (5) *holds.*

Proof. We define $V : \mathsf{Prop} \to PS$ by putting $V(p) := \{A \in S \mid p \in A\}$. For the definition of σ, take an arbitrary $A \in S$, and consider the coherent family $(\tau_n)_{n\in\omega}$ of Lemma 4. By the assumptions on T, we may fix an element $\sigma(A) \in TS$ such that $T\pi_n(\sigma(A)) = \tau_n$ for each n.

By induction on n we prove that for all $a \in \mathcal{L}_n$ we have

$$A \Vdash_{\sigma,V} a \text{ iff } a \in \pi_n A. \tag{8}$$

Confining our attention to the inductive case where $n = k + 1$, we prove (8) by formula induction, and we only cover the case where $a = \nabla\alpha$. Note that here we have $\alpha \in T\mathcal{L}_n$, and this enables us to apply the outer induction hypothesis.

Now we prove (8) by the following chain of equivalences (writing $\Vdash$ rather than $\Vdash_{\sigma,V}$):

$$\begin{aligned}
A \Vdash \nabla\alpha \text{ iff } & (\sigma(A), \alpha) \in \overline{T}\Vdash && \text{(definition of } \Vdash\text{)}\\
\text{iff } & (\sigma(A), \alpha) \in \overline{T}\big(\Vdash\restriction_{S\times\mathcal{L}_k}\big) && \text{(Fact 2)}\\
\text{iff } & (\sigma(A), \alpha) \in \overline{T}\big(Gr(\pi_k) \,;\, (\in)^{\smile}\big) && \text{(inductive hypothesis)}\\
\text{iff } & \alpha \mathrel{\overline{T}{\in}} T\pi_k(\sigma(A)) && \text{(Fact 2)}\\
\text{iff } & \alpha \mathrel{\overline{T}{\in}} \tau_k && \text{(definition of } \sigma\text{)}\\
\text{iff } & \nabla\alpha \in \pi_{k+1}(A) && \text{(equation (7))}
\end{aligned}$$

Finally, the Truth Lemma is immediate from (8) by the fact that $\mathcal{L}(\mathsf{Prop}) = \bigcup_{n\in\omega}\mathcal{L}_n$ and the definitions.

On the basis of this the following is immediate.

Theorem 8. *Let T restrict to finite sets and weakly preserve limits of surjective ω-cochains of finite sets. Then the logic $\mathbf{M}^T$ is strongly complete with respect to its coalgebraic semantics.*

References

1. M. Barr. Relational algebras. In *Reports of the Midwest Category Seminar IV*, volume 137 of *Lecture Notes in Mathematics*, pages 39–55, 1970.
2. J. Bergfeld. Moss's coalgebraic logic: Examples and completeness results. Master's thesis, ILLC, University of Amsterdam, 2009.
3. M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality. In C. Areces and R. Goldblatt, editors, *Advances in Modal Logic 7*, pages 1–21. College Publications, 2008.
4. C. Cîrstea. A compositional approach to defining logics for coalgebras. *Theoretical Computer Science*, 327:45–69, 2004.
5. C. Cîrstea, A. Kurz, D. Pattinson, L. Schröder, and Y. Venema. Modal logics are coalgebraic. *The Computer Journal*, 54:524–538, 2011.
6. J. L. Fiadeiro, editor. *Algebra and Coalgebra in Computer Science (CALCO 2005)*, volume 3629 of *LNCS*, 2005.
7. H.P. Gumm. From T-coalgebras to filter structures and transition systems. In Fiadeiro [6], pages 194–212.
8. D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Proc. MFCS'95*, pages 552–562. Springer, 1995. LNCS 969.
9. C. Kupke, A. Kurz, and D. Pattinson. Ultrafilter extensions for coalgebras. In Fiadeiro [6], pages 263–277.
10. C. Kupke, A. Kurz, and Y. Venema. Completeness for the coalgebraic cover modality. Submitted. (An earlier version appeared in *Advances in Modal Logic 7*, College Publications, 2008), 2010.
11. A. Kurz. Specifying coalgebras with modal logic. *Theoretical Computer Science*, 260:119–138, 2001.
12. A. Kurz and J. Rosický. Strongly complete logics for coalgebras. Unpublished, 2006.
13. L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).
14. A. Palmigiano and Y. Venema. Nabla algebras and Chu spaces. In T. Mossakowski, U. Montanari, and M. Haveraaen, editors, *Algebra and Coalgebra in Computer Science (CALCO 2007)*, volume 4624 of *LNCS*, pages 394–408, 2007.
15. D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309:177–193, 2003.
16. J. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
17. L. Schröder. A finite model construction for coalgebraic modal logic. *The Journal of Logic and Algebraic Programming*, 73:97–110, 2007.
18. L. Schröder and D. Pattinson. Strong completeness of coalgebraic modal logics. In S. Albers and J.-Y. Marion, editors, *26th International Symposium on Theoretical Aspects of Computer Science (STACS 2009)*, volume 3 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 673–684, Dagstuhl, Germany, 2009.
19. V. Trnková. General theory of relational automata. *Fundamenta Informaticae*, 3(2):189–234, 1980.
20. Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.
21. Y. Venema, S. Vickers, and J. Vosmaer. Powerlocales via relation lifting. Submitted, 2010.