
Completeness of the finitary Moss logic

CLEMENS KUPKE¹, ALEXANDER KURZ², YDE VENEMA³

ABSTRACT. We give a sound and complete derivation system for the valid formulas in the finitary version of Moss' coalgebraic logic, for coalgebras of arbitrary type.

Keywords: coalgebra, modal logic, coalgebraic logic, completeness

1 Introduction

Generalizing Kripke models and frames, coalgebras provide a general, category theoretic account of state-based evolving systems. This point of view was emphasized by Rutten [22], who developed, in analogy with Universal Algebra, the basics of Universal Coalgebra as a general theory of systems. One of the strengths of the coalgebraic approach is that a substantial part of the theory of systems can be developed uniformly in a functor T (on the category \mathbf{Set} of sets and functions), which intuitively represents the type of the transition system. For example, as discovered by Aczel [2], any functor T induces a canonical notion of bisimilarity on T -coalgebras.

The research programme of Coalgebraic Logic is to extend this uniform approach to logics for specifying and reasoning about the behavior of coalgebras. This research direction was initiated by Moss [18], who described a logic for T -coalgebras uniformly for all set functors T (satisfying a mild condition). Moss' fascinating idea was, roughly, to take T itself as a modality. In the case of the power set functor \mathcal{P} , this modality, denoted as ∇ , has surfaced in modal logic from time to time, for instance in Fine's work [9] on normal forms. It can be defined using the standard box and diamond: With $\alpha \in \mathcal{P}\mathcal{L}$ a set of formulas, the formula $\nabla\alpha$ can be seen as an abbreviation: $\nabla\alpha = \Box \bigvee \alpha \wedge \bigwedge \diamond\alpha$, where $\diamond\alpha$ denotes the set $\{\diamond a \mid a \in \alpha\}$. The *semantics* of ∇ can be expressed in terms of the so-called Egli-Milner *lifting* of the satisfaction relation $\Vdash \subseteq S \times \mathcal{L}$ between states and formulas to a relation $\overline{\Vdash}$ between $\mathcal{P}S$ (sets of states) and $\mathcal{P}\mathcal{L}$ (sets of formulas):

$$\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } \sigma(s) \overline{\Vdash} \alpha, \quad (1)$$

where $\sigma : S \rightarrow \mathcal{P}S$ denotes the successor function. Since one may associate a reasonable notion of relation lifting with other set functors as well, the

¹Supported by NWO under FOCUS/BRICKS grant 642.000.502.

²Partially supported by EPSRC grant EP/C014014/1

³The research of this author has been made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO).

observation (1) paves the way for generalization to an arbitrary functor T . Moss shows that his coalgebraic logic, based on a modality ∇_T , is invariant under bisimilarity, and, in the presence of infinitary conjunctions, characterizes bisimilarity.

The operator ∇_T associated with an arbitrary functor T looks strikingly different from the usual \Box and \Diamond modalities. Following on from [18], attention turned to the question how to obtain modal languages for T -coalgebras which use more standard modalities [15, 21, 11], and how to find derivation systems for these formalisms. This approach is now usually described in terms of predicate liftings [20, 24] or, equivalently, Stone duality [6, 16]. For a while, this approach displaced the interest in Moss' logic and the relationship between the two was not completely clear.

Interest in Moss' logic revived when it became clear that even in standard modal logic, a ∇ -based approach has some advantages. In fact, independently of Moss' work, Janin & Walukiewicz [12] already observed that the connectives ∇ and \vee may in some sense replace the set $\{\Box, \Diamond, \wedge, \vee\}$. This observation, which is closely linked to fundamental automata-theoretic constructions, lies at the heart of the theory of the modal μ -calculus, and has many applications, see for instance [8, 23]. Generalizing the link between fixpoint logics and automata theory to the coalgebraic level of generality, Kupke & Venema [14] generalized some of these observations to show that many fundamental results in automata theory are really theorems of universal coalgebra.

This paper addresses the main problem left open by Moss [18]. Moss' approach focuses on semantics, and he provides only some sound logical principles which do not constitute a complete syntactic calculus. As a first result in the direction of a derivation system for ∇ modalities, Palmigiano & Venema [19] gave a complete axiomatization for the cover modality (i.e., $\nabla_{\mathcal{P}}$ for the power set functor \mathcal{P}). This calculus was streamlined by Bílková, Palmigiano & Venema [5] into a formulation that admits a straightforward generalization to an arbitrary set functor T .

Our main contribution here is a uniform completeness proof. That is, in this paper we provide, uniformly in the functor T , a derivation system \mathbf{M} which is sound and complete with respect to the semantics of the coalgebraic language based on the modality ∇_T . The main idea of the completeness proof is based on the Stone duality approach to coalgebraic logic and, as a byproduct, we also see how Moss' language fits into this approach.

In the Stone duality approach to coalgebraic logic, the relationship between logic and semantic is based on the following situation

$$\begin{array}{ccc}
 \mathbb{M}\text{-Alg} & \xleftarrow{\quad} & T\text{-Coalg} & (2) \\
 \downarrow & & \downarrow & \\
 \mathbb{M} \curvearrowright \text{BA} & \xleftarrow{\mathbb{P}} & \text{Set} \curvearrowright T &
 \end{array}$$

where \mathbb{M} is the functor on Boolean algebras given by the proof system of the logic under consideration and \mathbb{P} is the contravariant powerset functor. The

semantics of the logic appears in this setting as a natural transformation $\delta : \mathbb{M}\mathbb{P} \rightarrow \mathbb{P}T$ (using δ , \mathbb{P} lifts to a functor in the upper row, which maps a T -coalgebra to its ‘complex’ \mathbb{M} -algebra). The proof system is complete if δ is injective (Proposition 48). One advantage of this approach is its flexibility. For example, descriptive-general-frame semantics corresponds to replacing **Set** by Stone spaces. On the algebra side, one can treat positive logic by replacing **BA** by distributive lattices or infinitary logic (like in Moss’ original work) by replacing **BA** by complete atomic Boolean algebras. This paper treats the case of **BA** and **Set** which is of particular interest to us and leaves the others for future work. This means, in particular, that we will concentrate on the finitary version of Moss’ logic first introduced in [27].

2 Preliminaries

In this section we settle on notation and terminology, and we introduce the finitary version of Moss’ logic. For background on coalgebra the reader is referred to [26].

General Two categories play a major role in our paper: the category **Set** with sets as objects and functions as arrows, and the category **BA** of Boolean algebras and homomorphisms. The categories **Set** and **BA** are related by the contravariant functor $\mathbb{P} : \mathbf{Set} \rightarrow \mathbf{BA}$, by the forgetful functor $U : \mathbf{BA} \rightarrow \mathbf{Set}$, and by the left adjoint \mathbb{F} of U mapping a set X to the free Boolean algebra over X . We write P for $U\mathbb{P}$, $\mathbb{2}$ for the two-element Boolean algebra and 1 for a one-element set.

Coalgebra A coalgebra (over **Set**) for a functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, also called T -coalgebra, is a pair (S, σ) where S is a set (of “states”) and $\sigma : S \rightarrow TS$ is a function (the “transition structure”). A T -coalgebra morphism from a T -coalgebra (S_1, σ_1) to a T -coalgebra (S_2, σ_2) is a function $f : S_1 \rightarrow S_2$ such that $Tf \circ \sigma_1 = \sigma_2 \circ f$.

For a modal logician, the prime examples of coalgebras are Kripke frames and Kripke models. Bisimulations between Kripke structures also have their natural coalgebraic generalization: a relation Z between the carrier sets of two coalgebras is a bisimulation if for all $(s_1, s_2) \in Z$, the pair $(\sigma_1(s_1), \sigma_2(s_2))$ belongs to the *relation lifting* \overline{Z} of Z .

DEFINITION 1. Let T be a set functor. Given a binary relation Z between two sets S_1 and S_2 , we define the relation $\overline{Z} \subseteq TS_1 \times TS_2$ as follows:

$$\overline{Z} := \{((T\pi_1)\phi, (T\pi_2)\phi) \mid \phi \in TZ\},$$

where $\pi_i : Z \rightarrow S_i$ for $i = 1, 2$ are the projection functions.

In this paper we will confine attention to set functors that are *standard* (that is, inclusions are mapped to inclusions), and that *preserve weak pullbacks*. We will not define the latter property, but simply note that it is equivalent to requiring that the composition of two bisimulations is again a bisimulation, or, equivalently, that for all relations Z_1, Z_2 we have $\overline{Z_1 \circ Z_2} = \overline{Z_1} \circ \overline{Z_2}$ (and it will be apparent from the development below that

this property is essential to work with the Moss modality). The requirement of standardness is not essential and only serves to keep the notation a bit smoother. The class of standard and weak pullback preserving functors includes the ones that are used to model infinite words, infinite binary trees, Kripke frames and probabilistic transition systems as coalgebras. A more detailed discussion of these examples can be found in [14]. For reasons of space limitations we cannot go into further detail here.

CONVENTION 2. Throughout this paper we fix a standard and weak pullback preserving set functor T .

The following fact lists the properties of relation lifting that we use in our paper. (Here $\text{Gr}(f) \subseteq S \times S'$ denotes the graph of a function $f : S \rightarrow S'$.) For proofs we refer to [18] and references therein.

FACT 3. Let T be a set functor that is standard and weak pullback preserving. Then relation lifting

- (1) extends T : $\overline{\text{Gr}(f)} = \text{Gr}(Tf)$, and preserves the diagonal: $\overline{Id_S} = Id_{TS}$;
- (2) is monotone: $R \subseteq Q$ implies $\overline{R} \subseteq \overline{Q}$;
- (3) commutes with taking restrictions: $\overline{R|_{U \times U'}} = \overline{R}|_{TU \times TU'}$;
- (4) preserves composition: $\overline{R \circ Q} = \overline{R} \circ \overline{Q}$, and converse: $\overline{(R^\circ)} = (\overline{R})^\circ$;

We let T_ω denote the *finitary*, or, ω -*accessible*, version of T , that is, the set functor T_ω which agrees with T on finite sets, while for an infinite set X ,

$$T_\omega(X) := \bigcup \{TY \mid Y \in \mathcal{P}_\omega(X)\}.$$

On maps, T_ω simply agrees with T . It is not hard to see that T_ω is a well-defined subfunctor of T (cf. [4, p.314]) and that $T_\omega X \subseteq TX$ for all sets X . Furthermore, as any standard set functor preserves finite intersections ([4, III, Prop. 4.6]), for any set X , and any element $\alpha \in T_\omega X$, there is a *smallest, finite* subset $X_0 \subseteq X$ such that $\alpha \in T_\omega X_0$. This set X_0 is called the *base* of α , notation: $\text{Base}(\alpha)$.

Moss' language

DEFINITION 4. Given a set X of proposition letters, we define the following. $\mathcal{L}_0(X)$ is the smallest superset of X which is closed under taking negations and finitary conjunctions and disjunctions. $\mathcal{L}_{n+1}(X) := \mathcal{L}_0(\{\nabla\alpha \mid \alpha \in T_\omega \mathcal{L}_n(X)\})$ is the smallest set containing the formula $\nabla\alpha$ for each $\alpha \in T_\omega \mathcal{L}_n(X)$, which is closed under taking negations and finitary conjunctions and disjunctions. $\mathcal{L}(X) := \bigcup_{n \in \omega} \mathcal{L}_n(X)$ is the set of *formulas in X* ; in case $X = \emptyset$ we write \mathcal{L}_n and \mathcal{L} instead of $\mathcal{L}_n(\emptyset)$ or $\mathcal{L}(\emptyset)$. The *depth* of a formula a is the smallest n such that $\phi \in \mathcal{L}_n$.

We write $\top := \bigwedge \emptyset$ and $\perp := \bigvee \emptyset$. Then by definition, \top and \perp belong to every layer of the language. While it is not hard to prove that $\mathcal{L}_n \subseteq \mathcal{L}_{n+1}$, for all $n \in \omega$, it is in general not the case that $X \subseteq \mathcal{L}_n$ for $n > 0$.

It will occasionally be useful to think of $\mathcal{L}_0(X)$ as the (carrier of the) absolutely free algebra of Boolean type, or the *Boolean term algebra*, generated by X , and of \mathcal{L}_{n+1} as the Boolean term algebra generated by the set

$\{\nabla\alpha \mid \alpha \in T_\omega\mathcal{L}_n\}$.

The language can be seen as an initial algebra for a functor.

PROPOSITION 5. *Let M be the set functor $Id + Id \times Id + Id \times Id + T_\omega$. Then $(\mathcal{L}, \neg, \wedge, \vee, \nabla)$ is the initial M -algebra.*

REMARK 6. For the category theoretic minded reader we note that, identifying formulas up to Boolean equivalence, Moss' language \mathcal{L} is the initial algebra for the functor $\mathbb{L} = \mathbb{F}T_\omega U : \mathbf{BA} \rightarrow \mathbf{BA}$.

While we will refer to the above language as *Moss'* coalgebraic language, there are actually some differences. The most important of these is that by defining $\nabla\alpha$ to be a formula only for elements $\alpha \in T_\omega\mathcal{L}$ (rather than for all $\alpha \in T\mathcal{L}$), we construct a language that is *finitary* in the sense that every formula has a finite number of *subformulas*. This notion can be defined inductively, the key clause being that the subformulas of $\nabla\alpha$ are given as the closure of the set $Base(\alpha)$ under subformulas.

Concerning the semantics of \mathcal{L} , we only give the clause for the ∇ modality.

DEFINITION 7. Given a coalgebra $\mathbb{S} = (S, \sigma)$, we define $s \Vdash \nabla\alpha$ if $\sigma(s) \Vdash \alpha$.

EXAMPLE 8. Let Prop be a set of propositional variables and recall that coalgebras for the functor $K = \mathcal{P}\text{Prop} \times \mathcal{P}$ correspond to Kripke models. Then any formula $\nabla_K\alpha$ is of the form $\nabla_K\alpha = \nabla_K(P, A)$ where $P \subseteq \text{Prop}$ is a set of proposition letters and $A \subseteq \mathcal{L}$ is a finite set of formulas. If the set Prop is finite it is easy to see that one can define a translation t of formulas in \mathcal{L} into the basic modal language by putting

$$t(\nabla_K(P, A)) := \bigwedge_{p \in P} p \wedge \bigwedge_{p \notin P} \neg p \wedge \bigwedge_{a \in A} \diamond t(a) \wedge \square(\bigvee_{a \in A} t(a))$$

such that $(S, \sigma), s \Vdash a$ iff $(S, \sigma), s \Vdash t(a)$ for all $a \in \mathcal{L}$.

The semantics of a ∇ -formula can be also expressed using the following natural transformation which plays a central role in our paper.

DEFINITION 9. We define a natural transformation $\rho : TP \rightarrow PT$ by putting $\rho_X(\Phi) := \{\alpha \in TX \mid \alpha \in \bar{\Phi}\}$.

REMARK 10. ρ is natural if T preserves weak-pullbacks. This is also true if one replaces the contravariant P with the covariant \mathcal{P} .

In order to gain some intuitions about the ∇ -operator and the transformation ρ , the reader is invited to prove the following easy lemma.

LEMMA 11. *For any $\nabla\alpha \in \mathcal{L}$ we have $s \Vdash \nabla\alpha$ iff $s \in \sigma^{-1} \circ \rho_S(T\mu(\alpha))$, where $\mu : \mathcal{L} \rightarrow \mathcal{P}S$ is the function that maps a formula to its semantics.*

REMARK 12. Following on from Remark 6, freely extending ρ to Boolean algebras yields a natural transformation $\gamma : \mathbb{L}\mathbb{P} \rightarrow \mathbb{P}T$. γ allows us to associate with any coalgebra (S, σ) a 'complex \mathbb{L} -algebra' $\mathbb{L}\mathbb{P}S \xrightarrow{\gamma_S} \mathbb{P}TS \xrightarrow{\mathbb{P}\sigma} \mathbb{P}S$. Denote by \mathcal{L}' the language \mathcal{L} quotiented by Boolean equivalence. Then \mathcal{L}' is the initial \mathbb{L} -algebra. For each coalgebra (S, σ) , initiality of \mathcal{L}' gives us

a map $\llbracket \cdot \rrbracket : \mathcal{L}' \rightarrow \mathbb{P}S$ interpreting elements of \mathcal{L}' as propositions on S . This definition agrees with Definition 7 (because γ is the free extension of ρ).

3 The derivation system

In this section we will define and discuss the derivation system \mathbf{M} . Before we can provide the actual definition of \mathbf{M} , we need a few preparatory remarks and definitions.

First of all, it will be convenient for us to have the derivation system operating on *inequalities*, that is, expressions of the form $a \preceq b$, with $a, b \in \mathcal{L}$. The main reason for this is that we like our system to stay close to equational reasoning. Indeed, in any logic with an underlying algebraic semi-lattice structure, inequalities can be seen as (special) *equations*: we may for instance identify the inequality $a \preceq b$ with the equation $a \wedge b \approx a$. Conversely, we may think of an equation $a \approx b$ as a *pair* of inequalities $a \preceq b, b \preceq a$.

DEFINITION 13. An inequality $a \preceq b$ is *valid in a coalgebra* $\mathbb{S} = (S, \sigma)$, notation: $\mathbb{S} \Vdash a \preceq b$, if $\mathbb{S}, s \Vdash a$ implies $\mathbb{S}, s \Vdash b$ for all $s \in S$, and *valid simpliciter* if it is valid in every coalgebra, notation: $a \models b$.

Note that the set of valid formulas can be obtained from the set of valid inequalities: a formula a is true in every state in every coalgebra iff the inequality $\top \preceq a$ is valid.

In the sequel we will need symbols to refer to formulas (\mathcal{L}), and to elements of the sets $\mathcal{P}_\omega \mathcal{L}$, $T_\omega \mathcal{L}$, $T_\omega \mathcal{P}_\omega \mathcal{L}$ and $\mathcal{P}_\omega T_\omega \mathcal{L}$. For convenience we fix our notation for such objects as follows:

\mathcal{L}	a, b, c, \dots	$T_\omega \mathcal{L}$	$\alpha, \beta, \gamma \dots$
$\mathcal{P}_\omega \mathcal{L}$	ϕ, ψ, \dots	$T_\omega \mathcal{P}_\omega \mathcal{L}$	Φ, Ψ, \dots
$\mathcal{P}_\omega T_\omega \mathcal{L}$	$A, B, C \dots$		

The same notation will be used for variants where \mathcal{L} is replaced by an arbitrary set or $\mathcal{P}_\omega, T_\omega$ are replaced by \mathcal{P}, T .

An important role in the definition of \mathbf{M} is played by the notion of a *slim redistribution*.

DEFINITION 14. A set $\Phi \in T\mathcal{P}(X)$ is a *redistribution* of a set $A \in \mathcal{P}T(X)$ if $A \subseteq \rho_X(\Phi)$. In case $A \in \mathcal{P}_\omega T_\omega(X)$, we call a redistribution Φ *slim* if $\Phi \in T_\omega \mathcal{P}_\omega(\bigcup_{\alpha \in A} \text{Base}(\alpha))$. The set of slim redistributions of A is denoted as $\text{SRD}(A)$.

A special case

Our derivation system is given in Definition 15. It turns out, however, that we can give a somewhat simpler version in case the functor T restricts to finite sets (that is, if TX is finite whenever X is finite). This simpler system is the direct generalization of the system for $T = \mathcal{P}$ (that is, where the coalgebras are Kripke structures) given by Bílková, Palmigiano and the

third author in [5].¹

\mathbf{M} is given as follows. On top of a complete set of axioms and rules for classical propositional logic, and the cut rule (from $a \preceq b$ and $b \preceq c$ derive $a \preceq c$), it has the axioms and derivation rules given in Table 1.

($\nabla 1$)	From $\alpha \overline{\preceq} \beta$ infer $\vdash \nabla \alpha \preceq \nabla \beta$
($\nabla 2$)	$\bigwedge \{ \nabla \alpha \mid \alpha \in A \} \preceq \bigvee \{ \nabla (T \wedge)(\Phi) \mid \Phi \in SRD(A) \}$
($\nabla 3$)	$\nabla (T \vee)(\Phi) \preceq \bigvee \{ \nabla \alpha \mid \alpha \in \overline{\Phi} \}$
($\nabla 4$)	From $\vdash \top \preceq \bigvee \phi$ infer $\vdash \top \preceq \bigvee \{ \nabla \alpha \mid \alpha \in T\phi \}$

Table 1. Axioms and rules of the system \mathbf{M} , if T restricts to finite sets

Let us hasten to give some explanation of the system. To start with, the reader may be slightly puzzled by our formulation of the derivation rule ($\nabla 1$), since its premiss ‘ $\alpha \overline{\preceq} \beta$ ’ uses syntax that has not been defined as part of the object language. The proper way to read this premiss is as follows: ‘the relation $Z := \{(a, b) \in Base(\alpha) \times Base(\beta) \mid \vdash a \preceq b\}$ is such that $(\alpha, \beta) \in \overline{Z}$ ’. In order to see this, note that using Fact 3(3) one can show that for all $\alpha, \beta \in T_\omega \mathcal{L}$ and all $Z \subseteq \mathcal{L} \times \mathcal{L}$

$$(\alpha, \beta) \in \overline{Z} \quad \text{iff} \quad (\alpha, \beta) \in \overline{Z'},$$

where $Z' := Z \upharpoonright_{Base(\alpha) \times Base(\beta)}$ is the restriction of Z to the finite sets $Base(\alpha)$ and $Base(\beta)$. An alternative formulation of this rule would therefore say that ‘if there is a relation $Z \subseteq Base(\alpha) \times Base(\beta)$ such that $(\alpha, \beta) \in \overline{Z}$, and $\vdash a \preceq b$ for all $(a, b) \in Z$, then infer $\vdash \nabla \alpha \preceq \nabla \beta$ ’. But the presentation in Table 1 is shorter and reveals more clearly that the rule is in fact the inequality version of a *congruence* rule. Our discussion shows that ($\nabla 1$) is a *finitary* rule, because its set of premisses can be assumed to be contained in the finite set $Base(\alpha) \times Base(\beta)$ if we want to derive $\nabla \alpha \preceq \nabla \beta$.

The axioms ($\nabla 2$) and ($\nabla 3$) could in fact both be replaced with identities, since in both cases, the reverse inequality of the axiom can be derived as a theorem. In order to be able to *read* the axioms $\nabla 2$ and $\nabla 3$, recall that \bigwedge and \bigvee are maps from $\mathcal{P}_\omega \mathcal{L}$ to \mathcal{L} , so that $T \wedge : T\mathcal{P}_\omega \mathcal{L} \rightarrow T\mathcal{L}$, and likewise for $T \vee$. Hence for $\Phi \in T_\omega \mathcal{P}_\omega \mathcal{L}$, $(T \wedge)(\Phi)$ and $(T \vee)(\Phi)$ belong to $T_\omega \mathcal{L}$, and thus $\nabla (T \wedge)(\Phi)$ and $\nabla (T \vee)(\Phi)$ are well-formed formulas. In addition, if T restricts to finite sets, every $A \in \mathcal{P}_\omega T_\omega \mathcal{L}$ can have at most finitely many slim redistributions, and every $\Phi \in T_\omega \mathcal{P}_\omega \Phi$ can have at most finitely many lifted members. So the two axioms ($\nabla 2$) and ($\nabla 3$) are at least well-defined. What these axioms have in common further is that they can be seen as *distributive principles*. This is the clearest in the case of ($\nabla 3$), which states that ∇ distributes over certain disjunctions. In the case of ($\nabla 2$) the distributivity is a bit more involved, but basically, the axiom states that any conjunction of ∇ s can be replaced with a disjunction of ∇ s of conjunctions.

¹In [5] it was shown that for $T = \mathcal{P}$ axiom ($\nabla 4$) is derivable from ($\nabla 1$)-($\nabla 3$). We recently discovered that this is also true for the case of an arbitrary functor T .

$(\nabla 1) \quad \frac{\{b_1 \preceq b_2 \mid (b_1, b_2) \in Z\}}{\nabla \alpha \preceq \nabla \beta} \quad (\alpha, \beta) \in \bar{Z}$
$(\nabla 2) \quad \frac{\{\nabla(T \wedge)(\Phi) \preceq a \mid \Phi \in SRD(A)\}}{\wedge \{\nabla \alpha \mid \alpha \in A\} \preceq a}$
$(\nabla 3) \quad \frac{\{\nabla \alpha \preceq a \mid \alpha \in \bar{\Phi}\}}{\nabla(T \vee)(\Phi) \preceq a}$
$(\nabla 4) \quad \frac{\{a \wedge \nabla \alpha' \preceq \perp \mid \alpha' \in T_\omega(\phi), \alpha' \neq \alpha\} \quad \top \preceq \vee \phi}{a \preceq \nabla \alpha}$

Table 2. Axioms and rules of the system \mathbf{M}

Finally, although the formulation of $(\nabla 4)$ does not use the actual symbol, it is here that the interaction of the coalgebraic modality with negation is dealt with. To see why this is so, observe that the conclusion of $(\nabla 4)$ implies that $\neg \nabla \beta \preceq \vee \{\nabla \alpha \mid \beta \neq \alpha \in T\phi\}$.

The general case

In the case of a general functor, that is, one that does not necessarily restricts to finite sets, some of the axioms and rules in Table 1 above may involve ill-defined syntax. In particular, none of the disjunctions on the right-hand side of the axioms $(\nabla 2)$ and $(\nabla 3)$ will be taken over a *finite* set. (Algebraically, however, it will often be convenient to think of e.g. $(\nabla 2)$ as stating that *in case* that the least upper bound given on the right hand side *exists*, it is greater than the object denoted by the left hand side.) The solution is to replace the axioms $(\nabla 2)$ and $(\nabla 3)$ with *infinitary derivation rules* (and to do something similar for the conclusion of $(\nabla 4)$), according to the following principle. An axiom of the form $s \preceq \bigvee_{i \in I} t_i$ is replaced with the derivation rule: ‘from $\{\vdash t_i \preceq a \mid i \in I\}$, infer $\vdash s \preceq a$ ’. Applying this principle to the above axiom system, we obtain the following derivation system.

DEFINITION 15. The derivation system \mathbf{M} is given by the axioms and derivation rules of Table 2, on top of a complete set of axioms and rules for classical propositional logic, and the cut rule.

A *derivation* is a well-founded tree, labelled with inequalities, such that the leaves of the tree are labelled with axioms of \mathbf{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 \mathcal{D} is a derivation of the inequality $a \preceq b$, we write $\mathcal{D} \vdash_{\mathbf{M}} a \preceq b$. We write $\vdash_{\mathbf{M}} a \preceq b$ if we want to suppress the actual derivation and we write $a \equiv b$ if $\vdash_{\mathbf{M}} a \preceq b$ and $\vdash_{\mathbf{M}} b \preceq a$.

The main result

THEOREM 16. *Let T be a standard functor that preserves weak pullbacks. Then for any pair a and b of formulas in \mathcal{L} :*

$$\vdash_{\mathbf{M}} a \preceq b \iff a \models b.$$

4 Soundness

Soundness is the direction from left to right of Theorem 16. It is proved by induction on the complexity of derivations. The key steps are to show that the rules $(\nabla 1)$ – $(\nabla 4)$ preserve validity.

First we consider the rule $(\nabla 1)$. Suppose that $\mathbb{S} \Vdash a \preceq b$ for all pairs (a, b) belonging to some relation $Z \subseteq \mathcal{L} \times \mathcal{L}$ such that $(\alpha, \beta) \in \bar{Z}$. From the first assumption it follows that $\Vdash \circ Z \subseteq \Vdash$, and so, by the properties of relation lifting, we see that $\overline{\Vdash} \circ \bar{Z} \subseteq \overline{\Vdash}$. In order to show that $\mathbb{S} \Vdash \nabla \alpha \preceq \nabla \beta$, take an arbitrary state s such that $\mathbb{S}, s \Vdash \nabla \alpha$. Hence, by the truth definition of ∇ , we see that $\sigma(s) \overline{\Vdash} \alpha$, and so from $(\alpha, \beta) \in \bar{Z}$ we may infer that $(\sigma(s), \beta) \in \overline{\Vdash} \circ \bar{Z} \subseteq \overline{\Vdash}$. But then, again by the truth definition of ∇ , we see that, indeed, $\mathbb{S}, s \Vdash \nabla \beta$.

For the rule $(\nabla 2)$, fix a set $A \in \mathcal{P}_\omega T_\omega \mathcal{L}$, and some formula $a \in \mathcal{L}$. Suppose that \mathbb{S} validates all the premisses of the rule, that is, $\mathbb{S} \Vdash \nabla(T_\omega \wedge)(\Phi) \preceq a$, for all slim redistributions Φ of A . In order to prove that \mathbb{S} validates the conclusion of $(\nabla 2)$, assume that $\mathbb{S}, s \Vdash \wedge\{\nabla \alpha \mid \alpha \in A\}$. Clearly it suffices to come up with a slim redistribution Φ_s of A such that $\mathbb{S}, s \Vdash \nabla(T \wedge)(\Phi)$.

For the definition of Φ_s , first associate, with any state t in \mathbb{S} , the finite set

$$\phi(t) := \{b \in \bigcup_{\alpha \in A} \text{Base}(\alpha) \mid \mathbb{S}, t \Vdash b\},$$

and define $\Phi_s := (T\phi)(\sigma(s))$.

First we show that $\mathbb{S}, s \Vdash \nabla(T \wedge)(\Phi_s)$. For that purpose, observe that by definition of ϕ , the map $\wedge \circ \phi : S \rightarrow \mathcal{L}$ is such that $\text{Gr}(\wedge \circ \phi) \subset \Vdash$. From this it follows by the properties of relation lifting that $\text{Gr}((T \wedge) \circ (T\phi)) \subseteq \overline{\Vdash}$. In other words, for every element $\tau \in TS$ we have that $\tau \overline{\Vdash} ((T \wedge) \circ (T\phi))(\tau)$. Taking $\tau = \sigma(s)$, we obtain immediately by the definitions that $\mathbb{S}, s \Vdash \nabla(T \wedge)(\Phi_s)$.

In order to see that Φ_s is a slim redistribution of A , observe that by definition of ϕ , $\text{Gr}(\phi) \circ \in^\vee = \Vdash$ when restricted to elements of $\bigcup_{\alpha \in A} \text{Base}(\alpha)$. Then by the properties of relation lifting, it follows that $\text{Gr}(T\phi \circ \bar{\in}^\vee) = \overline{\Vdash}$. But then for every $\alpha \in A$ it follows from $\sigma(s) \overline{\Vdash} \alpha$ that there is some object Ψ such that the pair $(\sigma(s), \Psi)$ belongs to the relation $\text{Gr}(T\phi)$, and $\alpha \in \bar{\Psi}$. From the first fact it follows that $\Psi = \Phi_s$, and so we find that each $\alpha \in A$ is a lifted member of Φ_s . In other words, Φ_s is a redistribution of A ; but then by its definition it is slim.

In order to understand the soundness of $(\nabla 3)$, first consider the statement $\mathbb{S}, s \Vdash \bigvee \phi$. This statement can be reformulated equivalently by saying that the pair (s, ϕ) belongs to the relation $\Vdash \circ \in$, since there is some element

$a \in \phi$ such that $s \Vdash a$. Alternatively, $s \Vdash \bigvee \phi$ iff $(s, \phi) \in \Vdash \circ \text{Gr}(\bigvee)^\smile$. In other words, we find that the relations $\Vdash \circ \in$ and $\Vdash \circ \text{Gr}(\bigvee)^\smile$ coincide. From this it follows that

$$\overline{\Vdash \circ \in} = \overline{\Vdash \circ \text{Gr}(\bigvee)^\smile}. \quad (3)$$

Fix some object $\Phi \in T_\omega \mathcal{P}_\omega \mathcal{L}$ and some formula a , and suppose that the coalgebra \mathbb{S} validates all the premisses of $(\nabla 3)$, i.e., $\mathbb{S} \Vdash \nabla \alpha \preceq a$, for all $\alpha \in \overline{\Phi}$. In order to prove that \mathbb{S} also validates the conclusion of the rule, take an arbitrary state s such that $\mathbb{S}, s \Vdash \nabla(T \bigvee)(\Phi)$. From this it follows that $(\sigma(s), (T \bigvee)(\Phi))$ belongs to the relation $\overline{\Vdash}$, and so $(\sigma(s), \Phi)$ belongs to $\overline{\Vdash} \circ \text{Gr}(T \bigvee)^\smile = \overline{\Vdash} \circ \text{Gr}(\bigvee)^\smile$. But then by (3), $(\sigma(s), \Phi)$ belongs to the relation $\overline{\Vdash} \circ \in = \overline{\Vdash} \circ \overline{\Phi}$. In other words, there is some object α such that $\sigma(s) \Vdash \alpha$ and $\alpha \in \overline{\Phi}$. Clearly then $\mathbb{S}, s \Vdash \nabla \alpha$, and so by the assumption we have $\mathbb{S}, s \Vdash a$.

Finally, for the rule $(\nabla 4)$, fix some finite set ϕ of formulas. It suffices to prove that, for an arbitrary T -coalgebra $\mathbb{S} = (S, \sigma)$, if $\mathbb{S} \Vdash \top \preceq \bigvee \phi$, then for every point $s \in \mathbb{S}$ we can find an $\alpha \in T(\phi)$ such that $\mathbb{S}, s \Vdash \nabla \alpha$. From the assumption it follows that every state in \mathbb{S} satisfies some formula in ϕ . We may formulate this using a function $f : S \rightarrow \phi$ such that $\mathbb{S}, s \Vdash f(s)$, for all $s \in S$, or, equivalently, $\text{Gr}(f) \subseteq \Vdash$. But then by the properties of relation lifting, we find that $\text{Gr}(Tf) \subseteq \overline{\Vdash}$. Now consider an arbitrary state s in \mathbb{S} , and let $\alpha \in T(\phi)$ be the element $(Tf)(\sigma(s))$. Then $(\sigma(s), \alpha) \in \text{Gr}(Tf) \subseteq \overline{\Vdash}$, and so by the truth definition of ∇ , we find that $\mathbb{S}, s \Vdash \nabla \alpha$. That is, we have found our α .

5 Completeness

The completeness proof will use a standard coalgebraic technique, namely to prove completeness via one-step-completeness. This is well-known in domain theory (see e.g. Abramsky [1]) and was introduced to coalgebra by Pattinson [20]. Subsequently, it was used in for instance [7, 13, 17].

The main idea is the following. First, we show that Moss' logic (\mathcal{L}, \equiv) can be stratified into layers $(\mathcal{L}_n, \equiv_n)$, with all layers at $n + 1$ arising in a uniform way from layers at n (Proposition 37). This uniform construction can be described by means of a 'one-step version' of the derivation system \mathbf{M} . Technically, it is described by a functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$, which constructs $(\mathcal{L}_{n+1}/\equiv_{n+1})$ as $\mathbb{M}(\mathcal{L}_n/\equiv_n)$. Our main technical result consists of showing that this one-step proof system is complete in a suitable sense (Proposition 48). Then, using a standard argument, completeness follows from one-step completeness (Proposition 50).

REMARK 17. Continuing from Remark 12, the proof system \mathbf{M} defines a quotient $\mathbb{L} \rightarrow \mathbb{M}$. Then $\delta_X : \mathbb{M}PX \rightarrow \mathbb{P}TX$ is given by factoring $\gamma : \mathbb{L}P \rightarrow \mathbb{P}T$. \mathbb{M} -algebras are the Boolean algebras with operator for the Moss modality. The initial \mathbb{M} -algebra \mathcal{M} is Lindenbaum algebra of Moss' logic (Proposition 39) and δ_X is injective (Proposition 48), which then implies completeness.

5.1 A one-step proof system

Recall the definition of $\mathcal{L}_n(X)$ from Definition 4.

DEFINITION 18. Let $\mathfrak{L}(X) = \mathcal{L}_0\{\nabla\alpha \mid \alpha \in T_\omega X\}$. In the following we consider \mathfrak{L} to be a functor $\mathbf{Set} \rightarrow \mathbf{Set}$, which maps $f : X \rightarrow Y$ to the function $\mathfrak{L}(f) : \mathfrak{L}(X) \rightarrow \mathfrak{L}(Y)$ that extends the map $\nabla\alpha \mapsto \nabla(Tf)(\alpha)$ via Boolean operations.

\mathfrak{L} constructs formulas step-wise: $\mathcal{L}_n = \mathfrak{L}^n(\mathcal{L}_0)$. Next we show how to construct \equiv_n step-wise. In order to smooth our presentation, it is convenient in the following definition to assume that the generators are already closed under Boolean operations.

DEFINITION 19. Let A be the carrier of an algebra for the Boolean signature and let $R \subseteq \mathcal{L}_0(A) \times \mathcal{L}_0(A)$ be a set of pairs called *relations*. Using the laws of Boolean algebra, with pairs $(a, a') \in R$ as additional axioms $a \leq a'$, one may generate a congruence relation \equiv_R on the set $A \times A$. We say that the pair (A, R) is a *presentation of the Boolean algebra* A/\equiv_R and denote this algebra as $\mathbf{BA}\langle A, R \rangle$. A homomorphism $f : A \rightarrow B$ of algebras for the Boolean signature is a presentation morphism from (A, \equiv_R) to (B, \equiv_S) if $a_1 \equiv_R a_2$ implies $f(a_1) \equiv_S f(a_2)$ for all $a_1, a_2 \in A$. The category of presentations and presentation morphisms is denoted by \mathbf{PRS} .

The notion of a presentation morphism is motivated by the following lemma which is not difficult to prove.

LEMMA 20. *Let $f : (A, R) \rightarrow (B, S)$ be a presentation morphism. Then the function $[f] : A/\equiv_R \rightarrow B/\equiv_S$ that maps the equivalence class of an element $a \in A$ to the equivalence class of $f(a)$ is well-defined. Moreover $[f]$ is a Boolean homomorphism.*

EXAMPLE 21. The *standard presentation* of a Boolean algebra \mathbb{B} is the pair $(U\mathbb{B}, \leq)$ where \leq is the relation on terms over $U\mathbb{B}$ induced by the partial order of \mathbb{B} .

The derivation system \mathbf{M} is essentially a ‘one-step’ derivation system since in every rule involving the modality, every occurrence of α is under the scope of exactly one ∇ . The following definition makes this precise.

DEFINITION 22. Let (X, R) be a presentation. The *one-step proof system* $\mathbf{M}(X, R)$ is the version of \mathbf{M} in which all inequalities $b_1 \leq b_2$ from R (that is, with $(b_1, b_2) \in R$) are additional axioms, and in which *only* elements from X and $\mathfrak{L}(X)$ may be used.² We denote the associated notion of derivability by $\vdash_{\mathbf{M}(X, R)}$. Furthermore, for a_1, a_2 in $\mathfrak{L}(X)$, we write $a_1 \preceq_{\mathbf{M}(X, R)} a_2$ if $\vdash_{\mathbf{M}(X, R)} a_1 \leq a_2$; and $a_1 \equiv_{\mathbf{M}(X, R)} a_2$ iff $a_1 \preceq_{\mathbf{M}(X, R)} a_2$ and $a_2 \preceq_{\mathbf{M}(X, R)} a_1$. We let $\mathbb{M}(X, R)$ denote the Boolean algebra presented by $(\mathfrak{L}(X), \preceq_{\mathbf{M}(X, R)})$.

In case (X, R) is the standard representation of a Boolean algebra \mathbb{A} , we write $\mathbf{M}(\mathbb{A})$ for the one-step proof system based on the standard presentation of \mathbb{A} , and $\vdash_{\mathbf{M}(\mathbb{A})}$, $\preceq_{\mathbf{M}(\mathbb{A})}$, $\equiv_{\mathbf{M}(\mathbb{A})}$ and $\mathbb{M}(\mathbb{A})$ for the associated notions.

²In $(\nabla 1) - (\nabla 4)$, the b_i range over X , the a, a' over $\mathfrak{L}(X)$, $\alpha, \alpha' \in T_\omega X$, $\phi \subseteq X$.

The next subsection shows that $\mathbb{M}(\mathbb{A})$ is not only a Boolean algebra, but that \mathbb{M} is a functor on the category of Boolean algebras. This will allow us, in Section 5.3, to recover the Lindenbaum algebra of \mathbf{M} as the union of algebras $(\mathbb{M}^n \mathfrak{D})_{n < \omega}$, where $\mathbb{M}^{n+1} \mathfrak{D} = \mathbb{M}(\mathbb{M}^n \mathfrak{D})$.

Sections 5.2 and 5.3 are rather technical and are needed mainly³ to deduce completeness from one-step completeness. Our main result, the one-step completeness (Proposition 48), does not depend on Sections 5.2 and 5.3 and the reader might wish to go directly to Section 5.4.

5.2 Technical Interlude: \mathbb{M} as a functor of Boolean algebras

We want to define a functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$ associated with the one-step proof system. As a first step in this direction we note that \mathfrak{L} can be seen as a functor on the category PRS. The proof-theoretic content of this is that a morphism $f : (A, R) \rightarrow (B, S)$ between presentations extends to a map of derivations between the one-step proof systems $\mathbb{M}(A, R)$ and $\mathbb{M}(B, S)$.

PROPOSITION 23. *Let (A, R) and (B, S) be presentations in PRS and let $f : (A, R) \rightarrow (B, S)$ be a presentation morphism. The function $\mathfrak{L}(f)$ is a presentation morphism from $(\mathfrak{L}(A), \preceq_{\mathbb{M}(A,R)})$ to $(\mathfrak{L}(B), \preceq_{\mathbb{M}(B,S)})$, i.e., for all $a', a'' \in \mathfrak{L}(A)$,*

$$\vdash_{\mathbb{M}(A,R)} a' \preceq a'' \quad \text{implies} \quad \vdash_{\mathbb{M}(B,S)} \mathfrak{L}(f)(a') \preceq \mathfrak{L}(f)(a''). \quad (4)$$

Proof. Let $c', c'' \in \mathfrak{L}(A)$ or $c', c'' \in A$. One shows by structural induction on the derivation of $c' \preceq_{\mathbb{M}(A,R)} c''$ that substituting, in the derivation, each occurrence of $a \in A$ with $f(a) \in B$ and each occurrence of $a' \in \mathfrak{L}(A)$ with $\mathfrak{L}(f)(a') \in \mathfrak{L}(B)$ yields a proof of $\mathfrak{L}(f)(c') \preceq_{\mathbb{M}(B,S)} \mathfrak{L}(f)(c'')$ or $f(c') \preceq_{\mathbb{M}(B,S)} f(c'')$, respectively. Let us give some details of the induction argument (the case in which the derivation ends by an application of $(\nabla 2)$ is similar to the cases $(\nabla 1)$, $(\nabla 3)$ and $(\nabla 4)$ and has been omitted due to space limitations).

Case $c' \preceq_{\mathbb{M}(A,R)} c''$ is a (Boolean) axiom. If $c', c'' \in \mathfrak{L}(A)$ it is clear from the definition of \mathfrak{L} that $\mathfrak{L}(f)(c') \preceq_{\mathbb{M}(B,S)} \mathfrak{L}(f)(c'')$ is an axiom as well. If $c, c'' \in A$, $f(c') \preceq_{\mathbb{M}(B,S)} f(c'')$ follows from the fact that f is a morphism in PRS.

Case $c' \preceq_{\mathbb{M}(A,R)} c''$ is obtained by a derivation that ends by the application of some rule R for classical propositional logic. Then $\mathfrak{L}(f)(c') \preceq_{\mathbb{M}(B,S)} \mathfrak{L}(f)(c'')$ or $f(c') \preceq_{\mathbb{M}(B,S)} f(c'')$ can be easily obtained by applying the inductive hypothesis to the premises of R .

Case $(\nabla 1)$ Suppose $c' = \nabla \alpha$ and $c'' = \nabla \beta$ and suppose that there is a derivation \mathcal{D} of $c' \preceq_{\mathbb{M}(A,R)} c''$ such that $(\nabla 1)$ is the last ruled applied in \mathcal{D} , i.e., \mathcal{D} ends with

³But Propositions 26 (via 23), 28 and 37 have a proof-theoretic interpretation and are of independent interest.

$$\frac{\{b_1 \preceq b_2 \mid (b_1, b_2) \in Z\}}{\nabla\alpha \preceq \nabla\beta} \quad (\alpha, \beta) \in \bar{Z}$$

Because f is a Boolean homomorphism we get $f(b_1) \leq f(b_2)$ for all $(b_1, b_2) \in Z$. Moreover one can easily calculate that $(\nabla(Tf)(\alpha), \nabla(Tf)(\beta)) \in \bar{Z}'$ with $Z' := \{(f(b_1), f(b_2)) \mid (b_1, b_2) \in Z\}$. Therefore we can apply rule $(\nabla 1)$ again: from the premisses $\{a_1 \preceq a_2 \mid (a_1, a_2) \in Z'\}$ we obtain $\mathfrak{L}(f)(\nabla\alpha) = \nabla(Tf)(\alpha) \preceq_{\mathbf{M}(B,S)} \nabla(Tf)(\beta) = \mathfrak{L}(f)(\nabla\beta)$ as required.

Case $(\nabla 3)$ Suppose $b' = \nabla(T\bigvee)(\Phi)$ and suppose that there is a derivation \mathcal{D} of $\nabla(T\bigvee)(\Phi) \preceq_{\mathbf{M}(A,R)} b''$ that ends with the following rule:

$$\frac{\{\nabla\alpha \preceq b'' \mid \alpha \in \bar{\Phi}\}}{\nabla(T\bigvee)(\Phi) \preceq b''}$$

By the inductive hypothesis we have

$$\nabla(Tf)(\alpha) \preceq_{\mathbf{M}(B,S)} \mathfrak{L}(f)(b'') \quad \text{for all } \alpha \in \bar{\Phi}. \quad (5)$$

Furthermore we get

$$\mathfrak{L}(f)\left(\nabla(T\bigvee)(\Phi)\right) \stackrel{\text{by Def.}}{=} \nabla(Tf)\left(\nabla(T\bigvee)(\Phi)\right) \quad (6)$$

$$\stackrel{f \text{ is hom.}}{=} \nabla(T\bigvee)\left(\nabla(T\mathcal{P}f)(\Phi)\right). \quad (7)$$

Moreover the following chain of equivalences holds:

$$\begin{aligned} \alpha' \in \overline{\nabla(T\mathcal{P}f)(\Phi)} & \text{ iff } (\alpha', \Phi) \in \overline{(\in \circ \text{Gr}(\mathcal{P}f)^\circ)} = \overline{\text{Gr}(f)^\circ \circ \in} \\ & \text{ iff } \alpha' = (Tf)(\alpha) \text{ for some } \alpha \in \bar{\Phi} \end{aligned}$$

The latter equivalence together with (5) yields $\nabla\alpha' \preceq_{\mathbf{M}(B,S)} \mathfrak{L}(f)(b'')$ for all $\alpha' \in \overline{\nabla(T\mathcal{P}f)(\Phi)}$. By applying rule $(\nabla 3)$ we obtain

$$\mathfrak{L}(f)\left(\nabla(T\bigvee)(\Phi)\right) \stackrel{(7)}{=} \nabla(T\bigvee)\left(\nabla(T\mathcal{P}f)(\Phi)\right) \preceq_{\mathbf{M}(B,S)} \mathfrak{L}(f)(b'')$$

as required.

Case $(\nabla 4)$ Consider now the case that $b' = \nabla\alpha$ and that there is a derivation \mathcal{D} of $b' \preceq_{\mathbf{M}(A,R)} \nabla\alpha$ that ends as follows:

$$\frac{\{b' \wedge \nabla\alpha' \preceq \perp \mid \alpha' \in T_\omega(\phi), \alpha' \neq \alpha\}}{b' \preceq \nabla\alpha} \quad \top \preceq \bigvee \phi$$

Again we can inductively assume that

$$\mathfrak{L}(f)(b') \wedge \nabla(Tf)(\alpha') \preceq \perp \quad \text{for all } \alpha' \in T_\omega(\phi) \text{ s.t. } \alpha' \neq \alpha. \quad (8)$$

Let us put $\psi := f[\phi]$. Then by the fact that f is a homomorphism of Boolean algebras we get $\top \preceq \bigvee \psi$. It is clear that for all $\beta' \in T_\omega(\psi)$ such

that $\beta' \neq Tf(\alpha)$ there is an $\alpha' \neq \alpha$ such that $Tf(\alpha') = \beta'$. Together with (8) this implies that $\mathfrak{L}(f)(\beta') \wedge \nabla\beta' \preceq \perp$ for all $\beta' \in T_\omega(\psi)$ such that $(Tf)(\alpha) \neq \beta'$. Now we can apply rule $(\nabla 4)$ which yields $\mathfrak{L}(f)(\beta') \preceq \nabla(Tf)(\alpha)$. This is what we had to show, because $\mathfrak{L}(f)(\nabla\alpha) = \nabla(Tf)(\alpha)$. \blacksquare

As a consequence, one obtains that the value of \mathbb{M} does not depend on a choice of presentation.

PROPOSITION 24. *If (X, R) and (X', R') generate isomorphic Boolean algebras, then $\mathbb{M}(X, R) \cong \mathbb{M}(X', R')$.*

Proof. Clearly it suffices to prove that $\mathbb{M}(X, R) \cong \mathbb{M}(B, \leq)$, where (B, \leq) is the standard presentation of $\mathbb{B} := \text{BA}\langle X, R \rangle$. Recall that B consists of equivalence classes of the set $\mathcal{L}_0(X)$, let $f : \mathcal{L}_0(X) \rightarrow B$ be the quotient map, and let $m : B \rightarrow \mathcal{L}_0(X)$ be a function such that $f \circ m = \text{id}_B$. Then it is clear that f and m are presentation morphisms between (X, R) and (B, \leq) . Then by Proposition 23 also $\mathfrak{L}(f)$ and $\mathfrak{L}(m)$ are presentation morphisms and $[\mathfrak{L}(f)]$ is an surjective BA-homomorphism from $\mathbb{M}(X, R)$ to $\mathbb{M}(B, \leq)$. In order to prove the claim we show that $[\mathfrak{L}(f)]$ is also injective. Note first that for all $x \in X$ we have $m(f(x)) \equiv_R x$. Using axiom $(\nabla 1)$ it can be shown that this implies $\mathfrak{L}(m)(\mathfrak{L}(f)(\nabla\alpha)) \equiv \nabla\alpha$ for all $\nabla\alpha \in \mathfrak{L}(X)$ which can be inductively extended to $\mathfrak{L}(m)(\mathfrak{L}(f)(a)) \equiv a$ for all $a \in \mathfrak{L}(X)$. But then for all $a, a' \in \mathfrak{L}(A)$ such that $\mathfrak{L}(f)(a) \equiv \mathfrak{L}(f)(a')$ we have

$$a \equiv \mathfrak{L}(m)(\mathfrak{L}(f)(a)) \equiv \mathfrak{L}(m)(\mathfrak{L}(f)(a')) \equiv a',$$

which implies that $[\mathfrak{L}(f)]$ has to be injective as required. \blacksquare

Recall that, given a Boolean algebra \mathbb{A} with standard presentation (A, \leq) , we let $\mathbb{M}\mathbb{A}$ denote the Boolean algebra $\mathbb{M}(A, \leq)$ (Definition 22). \mathbb{M} is thus an operation on the class of Boolean algebras. We will now see that in fact it is (or can be extended to) a *functor* on the category of Boolean algebras.

DEFINITION 25. The quotient map from $\mathfrak{L}(U\mathbb{A})$ to $\mathbb{M}\mathbb{A} = (\mathfrak{L}(U\mathbb{A})/\equiv_{\mathbb{M}(\mathbb{A})})$ will be denoted by $\mathfrak{q}_{\mathbb{A}}$.

Another consequence of Proposition 23 is:

PROPOSITION 26. *\mathbb{M} is a functor on the category of Boolean algebras.*

Proof. Let $f : \mathbb{A} \rightarrow \mathbb{B}$ be a Boolean homomorphism and let A and B denote the underlying sets of \mathbb{A} and \mathbb{B} , respectively. Obviously any homomorphism $f : \mathbb{A} \rightarrow \mathbb{B}$ is a presentation morphism from (A, \leq) to (B, \leq) where (A, \leq) and (B, \leq) denote the standard presentations of \mathbb{A} and \mathbb{B} . So by Proposition 23, $\mathfrak{L}(f) : (\mathfrak{L}(A), \preceq_{\mathbb{M}(A, \leq)}) \rightarrow (\mathfrak{L}(B), \preceq_{\mathbb{M}(B, \leq)})$ is also a presentation morphism and we can define a Boolean homomorphism $\mathbb{M}f : \mathbb{M}\mathbb{A} \rightarrow \mathbb{M}\mathbb{B}$ by putting $\mathbb{M}f := [\mathfrak{L}(f)]$. It is easy to see that this \mathbb{M} satisfies the usual functor conditions, i.e., that $\mathbb{M}(\text{id}) = \text{id}$ and $\mathbb{M}(f \circ g) = \mathbb{M}f \circ \mathbb{M}g$. \blacksquare

REMARK 27. Wrt Remark 12, we note that \mathbb{M} is a quotient $\mathbb{L} \rightarrow \mathbb{M}$.

It turns out that \mathbb{M} has some nice properties, which will be of use later on. In particular, we will show that \mathbb{M} is *finitary* (or ω -accessible) which means, proof-theoretically, that for any Boolean algebra \mathbb{A} , a derivation of $\vdash_{\mathbb{M}(\mathbb{A})} a_1 \preceq a_2$ can be carried out in a *finite* subalgebra of \mathbb{A} . A fairly easy consequence of this is the second useful property given below, namely, that \mathbb{M} *preserves embeddings*.

PROPOSITION 28. *\mathbb{M} is a finitary functor that preserves embeddings.*

Proof. Fix a Boolean algebra \mathbb{A} with carrier set $A := U\mathbb{A}$. Given two elements $a_1, a_2 \in \mathcal{L}(A)$, consider the collection of elements of A that occur as *subformulas* of a_1 and a_2 . It follows from our earlier remarks on subformulas that this is a *finite* set, which then generates a finite subalgebra \mathbb{A}' of \mathbb{A} . By definition we have $a_1, a_2 \in \mathcal{L}(A')$.

We claim

$$\vdash_{\mathbb{M}(\mathbb{A})} a_1 \preceq a_2 \text{ iff } \vdash_{\mathbb{M}(\mathbb{A}')} a_1 \preceq a_2. \quad (9)$$

The interesting direction of (9) is from left to right. The key observation here is that from the fact that \mathbb{A}' is a finite subalgebra of \mathbb{A} , we may infer the existence of a *surjective* homomorphism $f : \mathbb{A} \rightarrow \mathbb{A}'$ such that $f(a') = a'$ for all $a' \in A'$. (In other words, \mathbb{A}' is a *retract* of \mathbb{A} .) There are various ways to prove this statement; here we refer to Sikorski's theorem that complete Boolean algebras are injective [25]. But if f is a homomorphism, by Proposition 23 it follows from $\vdash_{\mathbb{M}(\mathbb{A})} a_1 \preceq a_2$ that $\vdash_{\mathbb{M}(\mathbb{A}')} \mathcal{L}(f)(a_1) \preceq \mathcal{L}(f)(a_2)$. Since $a_1, a_2 \in \mathcal{L}(A')$ and f restricts to the identity on A' , we may conclude that $\mathcal{L}(f)(a_i) = a_i$, for both $i = 1, 2$. Thus, indeed, $\vdash_{\mathbb{M}(\mathbb{A}')} a_1 \preceq a_2$. Using the fact that every BA is the directed colimit (union) of finite Boolean algebras, the finitariness of \mathbb{M} follows by a standard argument.

For the second part of the proof, let $e : \mathbb{A} \rightarrow \mathbb{B}$ be an embedding. Without loss of generality we will assume that e is actually the inclusion (that is, \mathbb{A} is a subalgebra of \mathbb{B}). In order to prove that $\mathbb{M}e : \mathbb{M}\mathbb{A} \rightarrow \mathbb{M}\mathbb{B}$ is also injective, it suffices to prove the following, for all $a_1, a_2 \in A$:

$$\vdash_{\mathbb{M}(\mathbb{B})} a_1 \preceq a_2 \text{ implies } \vdash_{\mathbb{M}(\mathbb{A})} a_1 \preceq a_2. \quad (10)$$

But the proof of (10) simply follows from two applications of (9). \blacksquare

As a straightforward corollary of Proposition 28, we obtain the existence of an *initial* \mathbb{M} -algebra. Furthermore, this initial algebra is obtained as the union of the initial \mathbb{M} -sequence to be defined now.

DEFINITION 29. We define $j_0 : \mathbb{2} \rightarrow \mathbb{M}\mathbb{2}$ to be the unique embedding of the two-element Boolean algebra $\mathbb{2}$ into $\mathbb{M}\mathbb{2}$, and inductively we define $j_{n+1} : \mathbb{M}^n\mathbb{2} \rightarrow \mathbb{M}(\mathbb{M}^n\mathbb{2})$ to be $\mathbb{M}j_n$. Let \mathcal{M} be the colimit of the sequence $(\mathbb{M}^n\mathbb{2})_{n < \omega}$.

We take the liberty to consider \mathcal{M} as an \mathbb{M} -algebra, a Boolean algebra, or a set, depending on the context. Since \mathbb{M} is finitary (Proposition 28), we have the following.

COROLLARY 30. *\mathcal{M} is the initial \mathbb{M} -algebra.*

Because \mathbb{M} preserves embeddings, all maps in the initial sequence $(\mathbb{M}^n \mathfrak{2})_{n < \omega}$ are injective. This means that we can consider the initial \mathbb{M} -algebra as a *union* of its approximants $\mathbb{M}^n \mathfrak{2}$.

5.3 Technical Interlude: Stratification of the Moss logic

As we will see now, the one-step version of \mathbf{M} allows for a layer-wise construction of the (inter)derivability relation between formulas.

DEFINITION 31. For each n , we define relations \leq_n and \equiv_n on \mathcal{L}_n . For $n = 0$, we simply let \leq_0 on the set \mathcal{L}_0 of all closed Boolean formulas denote derivability (in Boolean logic). Inductively, we define \leq_{n+1} as derivability in the one-step proof system $\mathbf{M}(\mathcal{L}_n, \leq_n)$. Finally, $a \equiv_n b \Leftrightarrow (a \leq_n b \text{ and } a \leq_n b)$.

The following proposition reveals the crucial role of \mathbb{M} in this stratification. Its proof proceeds via a straightforward inductive argument, of which the inductive step is an immediate consequence of Proposition 24.

PROPOSITION 32. For all n , $\mathcal{L}_n / \equiv_n \cong \mathbb{M}^n \mathfrak{2}$.

DEFINITION 33. For every $n \in \omega$ we let q_n be the quotient map from \mathcal{L}_n onto $UM^n \mathfrak{2} \cong \mathcal{L}_n / \equiv_n$. Furthermore we let $i_0 : \mathcal{L}_0 \rightarrow \mathcal{L}_1$ be the obvious embedding, and inductively we define $i_{n+1} = \mathfrak{L}(i_n)$.

PROPOSITION 34. We have $\mathcal{L}_n = \mathfrak{L}^n(\mathcal{L}_0)$ and for all $n \in \omega$ the map i_n is the inclusion of \mathcal{L}_n into \mathcal{L}_{n+1} .

Due to lack of space we omit the simple induction argument. The next proposition establishes a connection between the embeddings $j_n : \mathbb{M}^n \mathfrak{2} \rightarrow \mathbb{M}^{n+1} \mathfrak{2}$ and the inclusions $i_n : \mathcal{L}_n \rightarrow \mathcal{L}_{n+1}$. The proof is based on the following lemma. (i) is proved using $(\nabla 1)$ and (ii) follows from Proposition 26.

LEMMA 35. For all $n < \omega$, we have (i) $q_{n+1} = q_{\mathbb{M}^n \mathfrak{2}} \circ \mathfrak{L}(q_n)$ and (ii) $Uj_n \circ q_{\mathbb{M}^{n-1} \mathfrak{2}} = q_{\mathbb{M}^n \mathfrak{2}} \circ \mathfrak{L}(Uj_{n-1})$:

$$\begin{array}{ccc}
 \mathfrak{L}(\mathcal{L}_n) & \xrightarrow{q_{n+1}} & UM^{n+1} \mathfrak{2} \\
 \searrow \mathfrak{L}(q_n) & & \nearrow q_{\mathbb{M}^n \mathfrak{2}} \\
 & & \mathfrak{L}(UM^n \mathfrak{2})
 \end{array}
 \qquad
 \begin{array}{ccc}
 \mathfrak{L}(UM^{n-1} \mathfrak{2}) & \xrightarrow{\mathfrak{L}(Uj_{n-1})} & \mathfrak{L}(UM^n \mathfrak{2}) \\
 q_{\mathbb{M}^{n-1} \mathfrak{2}} \downarrow & & \downarrow q_{\mathbb{M}^n \mathfrak{2}} \\
 UM^{n-1} \mathfrak{2} & \xrightarrow{Uj_n} & UM^n \mathfrak{2}
 \end{array}$$

PROPOSITION 36. For all $n \in \omega$ we have $Uj_n \circ q_n = q_{n+1} \circ i_n$:

$$\begin{array}{ccc}
 \mathcal{L}_n & \xrightarrow{i_n} & \mathcal{L}_{n+1} \\
 q_n \downarrow & & \downarrow q_{n+1} \\
 UM^n \mathfrak{2} & \xrightarrow{Uj_n} & UM^{n+1} \mathfrak{2}
 \end{array}$$

Proof. The case $n = 0$ can be easily checked. Consider now $n = m + 1$ and

some $\nabla\alpha \in \mathcal{L}_{m+1} = \mathfrak{L}^m(\mathcal{L}_0)$. Then

$$\begin{aligned}
j_{m+1}(\mathfrak{q}_{m+1}(\nabla\alpha)) &\stackrel{(i)}{=} j_{m+1}(\mathfrak{q}_{\mathbb{M}^m \mathfrak{Z}}(\mathfrak{L}(\mathfrak{q}_m)(\nabla\alpha))) \\
&\stackrel{(ii)}{=} \mathfrak{q}_{\mathbb{M}^{m+1} \mathfrak{Z}}(\mathfrak{L}(j_m)(\mathfrak{L}(\mathfrak{q}_m)(\nabla\alpha))) \\
&= \mathfrak{q}_{\mathbb{M}^{m+1} \mathfrak{Z}}(\mathfrak{L}(j_m \circ \mathfrak{q}_m)(\nabla\alpha)) \\
&\stackrel{\text{I.H.}}{=} \mathfrak{q}_{\mathbb{M}^{m+1} \mathfrak{Z}}(\mathfrak{L}(\mathfrak{q}_{m+1} \circ i_m)(\nabla\alpha)) \\
&\stackrel{(i)}{=} (\mathfrak{q}_{m+2} \circ i_{m+1})(\nabla\alpha) \quad \blacksquare
\end{aligned}$$

The following proposition is crucial. It shows that if we have $\vdash_{\mathbf{M}} a \preceq b$ for formulas a, b of depth n , there always is a derivation that does not employ formulas of depth greater than n . This is typical for axiomatisations where each variable is under the scope of precisely one modal operator. The situation here is slightly more complicated than usual since our rules allow infinite sets of premises.

PROPOSITION 37. *Let a and b be formulas. Then*

1. *If $a, b \in \mathcal{L}_n$, and $a \leq_m b$ for some $m > n$, then $a \leq_n b$.*
2. *If $a, b \in \mathcal{L}_n$, then $\vdash_{\mathbf{M}} a \preceq b$ iff $a \leq_n b$.*

Proof. For Part 1 of the proposition, it suffices to confine attention to the case where $m = n + 1$, which is a consequence of Proposition 36 and the fact that the injective BA-morphism j_n reflects the order.

Part 2 of the proposition is proved by induction on the complexity of derivations in \mathbf{M} . Here we discuss a sample case of the inductive step, namely, where the last applied rule was $(\nabla 4)$:

$$\frac{\{a \wedge \nabla\alpha' \preceq \perp \mid \alpha' \in T_\omega(\phi), \alpha' \neq \alpha\} \quad \top \preceq \bigvee \phi}{a \preceq \nabla\alpha}$$

Inductively, there is some natural number k such that $\top \leq_k \phi$. Let $m := \max\{d(b) \mid b \in \phi \cup \{a\}\}$, where $d(b)$ denotes the depth of the formula b . Then clearly m is (well-defined as) a finite natural number since ϕ is a finite set by assumption. Then $\phi \subseteq \mathcal{L}_m$, and so $\nabla\beta \in \mathcal{L}_{m+1}$ for all $\beta \in T_\omega\phi$. Since also $a \in \mathcal{L}_m$, by the first part of the proposition we obtain $a \wedge \nabla\alpha' \leq_{m+1} \perp$ for all $\alpha' \in T_\omega(\phi) \setminus \{\alpha\}$. Thus, with $p = \max\{m + 1, k\}$, we see that all premisses of the final rule are p -derivable. But then the conclusion is $p + 1$ -derivable, and then, by part 1, n -derivable (where we assumed that $a, \nabla\alpha \in \mathcal{L}_n$). \blacksquare

The next proposition shows that \mathcal{M} is the Lindenbaum algebra of \mathcal{L} modulo the proof system \mathbf{M} . To see this, recall from Proposition 5 that $(\mathcal{L}, \neg, \wedge, \vee, \nabla)$ is the initial algebra for the functor $M = Id + Id \times Id + Id \times Id + T_\omega$. Since \mathcal{M} can also be seen as an algebra of this kind, it follows from initiality of \mathcal{L} that there is a unique quotient $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$. The next

proposition states that the kernel of \mathfrak{q} is the interderivability relation \equiv according to the proof system \mathbf{M} .

LEMMA 38. *The quotient maps $\mathfrak{q}_n : \mathcal{L}_n \rightarrow U\mathbb{M}^n\mathfrak{2}$ are the restrictions of $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$. More precisely denote by $k_n : M^n\mathfrak{2} \rightarrow \mathcal{M}$ the embeddings of the initial sequence of \mathcal{M} and by $l_n : \mathcal{L}_n \rightarrow \mathcal{L}$ the inclusions. Then the claim is that $\mathfrak{q} \circ l_n = k_n \circ \mathfrak{q}_n$.*

Proof. To prove this we need to observe that, by definition, $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$ is the unique map for which $\mathfrak{q} \circ l_n = f_n$, where the $f_n : \mathcal{L}_n \rightarrow \mathcal{M}$ are given inductively by $f_{n+1} = \mu \circ \mathfrak{q}_{\mathcal{M}} \circ \mathfrak{L}(f_n)$ and $\mu : \mathbb{M}\mathcal{M} \rightarrow \mathcal{M}$ is the structure map of the \mathbb{M} -algebra \mathcal{M} . We then proceed to show by induction that $k_n \circ \mathfrak{q}_n = f_n$. Indeed, $f_{n+1} = \mu \circ \mathfrak{q}_{\mathcal{M}} \circ \mathfrak{L}(f_n) \stackrel{\text{indhyp}}{=} \mu \circ \mathfrak{q}_{\mathcal{M}} \circ \mathfrak{L}(k_n) \circ \mathfrak{L}\mathfrak{q}_n = \mu \circ \mathbb{M}k_n \circ \mathfrak{q}_{\mathbb{M}^n\mathfrak{2}} \circ \mathfrak{L}\mathfrak{q}_n = k_{n+1} \circ \mathfrak{q}_{\mathbb{M}^n\mathfrak{2}} \circ \mathfrak{L}\mathfrak{q}_n = k_{n+1} \circ \mathfrak{q}_{n+1}$. where the third equation holds by definition of $\mathbb{M}k_n = [\mathfrak{L}(k_n)]$ and the last equation is an instance of Lemma 35(i). ■

PROPOSITION 39. *The kernel of the quotient map $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$ is \equiv .*

Proof. Recall that \mathcal{M} can be seen as the union of the initial sequence $(\mathbb{M}^n\mathfrak{2})_{n < \omega}$, that $\mathcal{L} = \bigcup \mathcal{L}_n$, and that the maps $\mathfrak{q}_n : \mathcal{L}_n \rightarrow \mathbb{M}^n\mathfrak{2}$ map formulas to their equivalence classes. By Lemma 38, each \mathfrak{q}_n is the restriction of $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$ to $\mathcal{L}_n \rightarrow \mathbb{M}^n\mathfrak{2}$. Then, $\mathfrak{q}(a) = \mathfrak{q}(b)$ iff there is an n such that $\mathfrak{q}_n(a) = \mathfrak{q}_n(b)$. This, in turn, is equivalent to $a \equiv_n b$ (by Definition 33) and then to $a \equiv b$ by Proposition 37. ■

5.4 Semantics of \mathbb{M} and Moss algebras

Let us first summarise the two preceding sections in the following diagram:

$$\begin{array}{ccccccc}
 \mathcal{L}_0 & \xrightarrow{i_0} & \mathcal{L}_1 & \xrightarrow{i_1} & \mathcal{L}_2 & \cdots & \mathcal{L} \\
 \mathfrak{q}_0 \downarrow & & \mathfrak{q}_1 \downarrow & & \mathfrak{q}_2 \downarrow & & \mathfrak{q} \downarrow \\
 \mathfrak{2} & \xrightarrow{j_0} & \mathbb{M}\mathfrak{2} & \xrightarrow{j_1} & \mathbb{M}^2\mathfrak{2} & \cdots & \mathcal{M}
 \end{array}$$

\mathcal{L} is the union of the \mathcal{L}_n . \mathbb{M} is a functor on \mathbf{BA} , $\mathfrak{2}$ is the initial \mathbf{BA} and \mathcal{M} is the colimit of the sequence $\mathbb{M}^n\mathfrak{2}$. Since \mathbb{M} preserves injections and is finitary, \mathcal{M} is the initial \mathbb{M} -algebra and can be considered to be the union of the $\mathbb{M}^n\mathfrak{2}$. \mathfrak{q} is the quotient of \mathcal{L} wrt interderivability \equiv and the \mathfrak{q}_n are the restrictions of \mathfrak{q} .

Thus, up to interderivability, we can work with \mathbb{M} and \mathcal{M} instead of \mathcal{L} . In this section, we define the semantics of the logic directly in terms of \mathbb{M} and show that it agrees with the previously given one.

The relationship between \mathbb{M} and T is provided by a natural transformation $\delta : \mathbb{M}\mathbb{P} \rightarrow \mathbb{P}T$. For the definition of δ , recall the natural transformation ρ from Definition 9.

DEFINITION 40. Given a set S , define the map $\tilde{\rho} : \mathfrak{L}(PS) \rightarrow PTS$ as follows. For $\alpha \in T_\omega PS$, we let $\tilde{\rho}(\nabla\alpha) = \rho(\alpha)$ and then extend it freely to Boolean terms.

The soundness of the one-step proof system is enshrined in the next proposition. The proof is essentially the same as that of the soundness direction in Theorem 16.

PROPOSITION 41. $a_1 \equiv_{\mathbb{M}(\mathbb{P}X)} a_2$ implies $\tilde{\rho}(a_1) = \tilde{\rho}(a_2)$, for $a_i \in \mathcal{L}(PX)$.

By Proposition 41, the following is well-defined.

DEFINITION 42. Given a set X , let $\delta_X : \mathcal{L}(PX)/\equiv_{\mathbb{M}(\mathbb{P}X)} \rightarrow PTX$ be the map given by $\tilde{\rho}_X = \delta_X \circ \mathfrak{q}_{\mathbb{P}(X)}$.

PROPOSITION 43. The collection of maps given by Definition 42 form a natural transformation $\delta : \mathbb{M}\mathbb{P} \rightarrow \mathbb{P}T$.

Proof. We need to show that for each X , δ_X is a (Boolean) homomorphism, and that δ is natural. Both proofs are straightforward. ■

REMARK 44. Continuing from Remark 27, δ is given by factoring $\gamma : \mathbb{L}\mathbb{P} \rightarrow \mathbb{P}T$ (Remark 6) through $\mathbb{L}\mathbb{P} \rightarrow \mathbb{M}\mathbb{P}$.

The natural transformation δ allows us to associate with a coalgebra (S, σ) its ‘complex \mathbb{M} -algebra’ $\mathbb{P}\sigma \circ \delta_S : \mathbb{M}\mathbb{P}S \rightarrow \mathbb{P}S$. Recall that \mathcal{M} denotes the initial \mathbb{M} -algebra. For each coalgebra (S, σ) , initiality of \mathcal{M} gives us a map

$$\llbracket \cdot \rrbracket : \mathcal{M} \rightarrow \mathbb{P}S \quad (11)$$

interpreting elements of \mathcal{M} as propositions on S . Note that this map is an arrow in the category of Boolean algebras.

The next proposition ensures that the coalgebraic semantics of \mathcal{M} (see (11)) and of \mathcal{L} (Definition 7) agree.

PROPOSITION 45. Denote by $\mathfrak{q} : \mathcal{L} \rightarrow \mathcal{M}$ the quotient map. Given a coalgebra (S, σ) and $a \in \mathcal{L}$, we have $s \Vdash a$ iff $s \in \llbracket \mathfrak{q}(a) \rrbracket$.

Proof. The semantic map $\mu : \mathcal{L} \rightarrow \mathcal{P}S$ can be written as $\mu = f \circ \mathfrak{q}$ for some $f : \mathcal{M} \rightarrow \mathcal{P}S$ by putting $f(\mathfrak{q}(a)) := \mu(a)$. The function f is well-defined because of soundness of our logic: if $\mathfrak{q}(a) = \mathfrak{q}(a')$ then $a \equiv a'$ by Proposition 39 and therefore by soundness we get $\mu(a) = \mu(a')$. Using Lemma 11 it is not difficult to see that f is in fact an \mathbb{M} -algebra morphism from the initial \mathbb{M} -algebra \mathcal{M} to the \mathbb{M} -algebra $(\mathbb{P}S, \mathbb{P}\sigma \circ \delta_S)$. Therefore by initiality we get $f = \llbracket \cdot \rrbracket$ and thus $\mu(a) = \llbracket \mathfrak{q}(a) \rrbracket$. ■

REMARK 46. We have now finished the functorial presentation of Moss’ logic. A central role play the *Moss algebras*, that is, the algebras for the functor \mathbb{M} . In the case of $T = \mathcal{P}$, the category of Moss algebras is isomorphic to the category of Boolean algebras with operators. (11) corresponds to the fact that formulas are evaluated on a Kripke frame S by the morphism from the Lindenbaum BAO \mathcal{M} to the complex algebra $\mathbb{P}S$ of S . The completeness proof in the next section generalises the well-known fact that we have an injection (iso for finite X) $d_X : \mathbb{K}\mathbb{P}X \rightarrow \mathbb{P}\mathcal{P}X$ where \mathbb{K} is the functor $\mathbb{B}\mathbb{A} \rightarrow \mathbb{B}\mathbb{A}$ mapping \mathbb{A} to the algebra freely generated by $\Box a, a \in \mathbb{A}$, modulo the equations expressing that \Box preserves finite meets (d_X is given by $\Box a \mapsto \{b \subseteq X \mid b \subseteq a\}$).

5.5 One-step completeness

Completeness of \mathbf{M} is enshrined in the injectivity of δ_X . To show this we use the following basic fact about Boolean algebras.

LEMMA 47. *Let \mathbb{A} and \mathbb{B} be Boolean algebras and $f : \mathbb{A} \rightarrow \mathbb{B}$ be a homomorphism. Furthermore assume that \mathbb{A} is join-generated by $\mathcal{G} \subseteq A$, i.e., assume that for every $a \in A$ we have $a = \bigvee \{b \in \mathcal{G} \mid b \leq a\}$. Then $f(b) \neq \perp_{\mathbb{B}}$ for all $b \in \mathcal{G}$ implies that f is injective.*

Proof. In order to prove the claim note first that for all $a \in A$ we clearly have $\perp_{\mathbb{A}} < a$ implies $\perp_{\mathbb{B}} < f(a)$. Let now a, a' be elements of A such that $a \neq a'$. By our assumption we have w.l.o.g. that there is some $b \in \mathcal{G}$ with $b \leq a$ and $b \not\leq a'$. Therefore $\perp_{\mathbb{A}} < \neg a' \wedge b$ which implies by our first observation and the fact that f is a homomorphism that $\perp_{\mathbb{B}} < \neg f(a') \wedge f(b)$ and thus $f(b) \not\leq f(a')$. On the other hand we clearly have $f(b) \leq f(a)$ which yields $f(a) \neq f(a')$. As a, a' were assumed to be arbitrary we showed that f is injective. \blacksquare

PROPOSITION 48. *For every set X , the map $\delta_X : \mathbb{MIP}(X) \rightarrow \mathbb{PT}(X)$ is an embedding.*

Proof. The basic idea of the proof is to work with the map $T\eta : TX \rightarrow TPX$, where we write $\eta_X : X \rightarrow PX$ for the singleton map $x \mapsto \{x\}$. The crucial property is that

$$\rho_X \circ T(\eta_X) = \eta_{TX}. \quad (12)$$

The proof of (12) is based on the observation that $\text{Gr}(\eta_X) \circ \overset{\sim}{\circ} = \text{Id}_X$. From this it follows by the properties of relation lifting, that $\text{Gr}(T\eta_X) \circ \overset{\sim}{\circ} = \text{Id}_{TX}$ and thus $\rho_X(T\eta_X(\alpha)) = \{\beta \mid \beta \overset{\sim}{\circ} T\eta_X(\alpha)\} = \{\alpha\}$.

We define the set of “ T -singletons” by putting $\mathcal{G} := \{\nabla T(\eta_X)(\alpha) \mid \alpha \in T_{\omega}X\}$. In order to prove the proposition it now suffices to show that the Boolean algebra $\mathbb{MIP}(X)$ is *join-generated* by the T -singletons:

$$\forall a \in \mathbb{MIP}(X). a = \bigvee \{\nabla\beta \in \mathcal{G} \mid \nabla\beta \leq a\}. \quad (13)$$

(Note that the algebra $\mathbb{MIP}(X)$ need not be complete. The intended reading of (13) is that every element of $\mathbb{MIP}(X)$ is the join of the T -singletons below it, not that every set of T -singletons has a join.) To see why the injectivity of δ_X follows from this, note that by (12) we have $\delta_X(\nabla T\eta(\alpha)) = \rho_X(T\eta(\alpha)) = \{\alpha\} \neq \emptyset = \perp_{\mathbb{P}(TX)}$ for all $\nabla T\eta(\alpha) \in \mathcal{G}$. Therefore an application of Lemma 47 yields that f is injective.

Turning to the proof of (13), we distinguish cases as to the nature of the element a .

Case 1: Consider first an element of $\mathbb{MIP}(X)$ of the form $\nabla\beta$ with $\beta \in T_{\omega}UPX$. It can be easily shown that

$$\nabla\beta = \nabla(T_{\omega} \bigvee)(T_{\omega} \mathcal{P}_{\omega} \eta_X(\beta)) = \bigvee_{\gamma \in \rho_{PX}(T_{\omega} \mathcal{P}_{\omega} \eta(\beta))} \nabla\gamma$$

where the first equality follows from the fact that $\bigvee \circ \mathcal{P}_\omega \eta_X = \text{id}_{\mathcal{P}_\omega X}$ and the second equality is an instance of axiom $(\nabla 3)$. Furthermore one calculates (a detailed proof can be found in [10]) that

$$\rho_{\mathcal{P}X}(T_\omega \mathcal{P}_\omega \eta(\beta)) = T_\omega \eta[\rho_X(\beta)] = \{T_\omega \eta(\alpha) \mid \alpha \in \rho_X(\beta)\}.$$

By combining the latter equality with the preceding ones we obtain

$$\nabla \beta = \bigvee_{\alpha \in \rho_X(\beta)} \nabla(T_\omega \eta)(\alpha),$$

which shows that $\nabla \beta$ is the join of elements of \mathcal{G} .

Case 2: Consider now an element of $\text{MIP}(X)$ of the form $\neg \nabla \beta$. Let \mathbb{B} be the subalgebra of $\mathbb{P}(X)$ generated by $\text{Base}(\beta) \subseteq_\omega \text{UP}(X)$. We write B for the carrier of \mathbb{B} . As $\text{Base}(\beta)$ is finite, the Boolean algebra \mathbb{B} is finite as well. Let $\phi \in \mathcal{P}_\omega(X)$ be the (finite) set of atoms of \mathbb{B} . Then clearly $\bigvee \phi = \top$, while $a \wedge a' = \perp$ for any two distinct $a, a' \in \phi$. Furthermore \bigvee induces an isomorphism from $\mathcal{P}_\omega \phi$ to B that lifts to an isomorphism $T \bigvee$ between $T_\omega \mathcal{P}_\omega \phi$ and $T_\omega B$. As $\text{Base}(\beta) \subseteq B$ we have $\beta \in T_\omega B$ and thus there exists some $\Phi_\beta \in T_\omega \mathcal{P}_\omega \phi$ such that $(T \bigvee)(\Phi_\beta) = \beta$. Now axiom $(\nabla 3)$ entails that

$$\nabla \beta = \bigvee \{\nabla \gamma \mid \gamma \in T_\omega \phi, \gamma \bar{\in} \Phi_\beta\}. \quad (14)$$

Our claim is now that

$$\top \preceq \bigvee \{\nabla \gamma \mid \gamma \in T_\omega \phi\} \quad \text{and} \quad (15)$$

$$\nabla \gamma \wedge \nabla \gamma' \preceq \perp \quad \text{for all } \gamma, \gamma' \in T_\omega \phi \text{ s.t. } \gamma \neq \gamma'. \quad (16)$$

Items (14), (15) and (16) together entail that

$$\neg \nabla \beta = \bigvee \{\nabla \gamma \mid \gamma \in T_\omega \phi \text{ and not } \gamma \bar{\in} \Phi_\beta\}. \quad (17)$$

From Case 1 we know that all elements $\nabla \gamma$ that occur on the righthand side of equation (17) can be written as joins of elements of \mathcal{G} and therefore the same applies to $\neg \nabla \beta$.

We now turn to the proof of (15) and (16). First note that because $\bigvee \phi = \top$ an application of $(\nabla 4)$ shows that $\top \preceq \bigvee \{\nabla \gamma \in T_\omega \phi \mid \gamma \in T_\omega \phi\}$ which proves (15). For the proof of (16) consider $\gamma, \gamma' \in T_\omega \phi$ with $\gamma \neq \gamma'$. Let $\Phi \in \text{SRD}(C)$ be a slim redistribution of $C := \{\gamma, \gamma'\}$. We want to show that $\nabla(T \wedge)(\Phi) \preceq \perp$. As Φ is an arbitrary redistribution of C this will imply by $(\nabla 2)$ that $\nabla \gamma \wedge \nabla \gamma' \preceq \perp$ as required.

By assumption we have $(\gamma, \Phi), (\gamma', \Phi) \in \bar{\in}$. This shows that $\emptyset \notin \text{Base}(\Phi)$. Suppose now for a contradiction that $B_\Phi := \text{Base}(\Phi) \subseteq \mathcal{P}\phi$ contains only singleton sets and put $\in' := \in|_{\phi \times B_\Phi}$. Then $\in' \circ \in' \preceq \text{Id}_\phi$. As a consequence we get $(\gamma, \gamma') \in \overline{\in' \circ (\in')^\vee} \subseteq \text{Id}_{T\phi}$ which means $\gamma = \gamma'$ — a contradiction. Hence we can assume that B_Φ contains at least one set $\psi^* \subseteq \phi$ such that $|\psi^*| > 1$.

In order to prove that $\nabla(T\bigwedge)(\Phi) \preceq \perp$, define a function $d : \mathcal{P}\phi \rightarrow \mathcal{P}\phi$ by letting

$$d(\psi) := \begin{cases} \emptyset & \text{if } |\psi| \geq 1 \\ \psi & \text{if } |\psi| = 1 \\ \phi & \text{if } \psi = \emptyset \end{cases}$$

It follows from our assumptions on the set ϕ that $\vdash \bigwedge \psi \preceq \bigvee d(\psi)$, for all $\psi \in \mathcal{P}\phi$. Then an application of axiom $(\nabla 1)$ shows that $\nabla(T\bigwedge)(\Phi) \preceq \nabla(T\bigvee)(Td(\Phi))$. Because $Base(Td(\Phi)) = d[B_\Phi]$ and $d(\psi^*) = \emptyset$ we obtain $\emptyset \in Base(Td(\Phi))$. It is now a matter of routine checking that $A := \{\alpha \in T_\omega U\mathbb{P}(X) \mid \alpha \in Td(\Phi)\} = \emptyset$. By axiom $(\nabla 3)$ we have $\nabla(T\bigvee)(Td(\Phi)) \preceq \bigvee_{\alpha \in A} \nabla \alpha = \perp$, and thus $\nabla(T\bigwedge)(\Phi) \preceq \perp$ which finishes the argument for proving (16) and hence also of (17).

Case 3: Consider an element of $\mathbb{M}\mathbb{P}(X)$ of the form $\bigwedge_{i \in I} \nabla \beta_i$ for some finite set $A = \{\nabla \beta_i \in T_\omega U\mathbb{P}(X) \mid i \in I\}$. Then by axiom $(\nabla 2)$ we have

$$\begin{aligned} \bigwedge_{i \in I} \nabla \beta_i &= \bigvee_{\Phi \in SRD(A)} \nabla(T_\omega \bigwedge)(\Phi) \\ &\stackrel{\text{Case 1}}{=} \bigvee_{\Phi \in SRD(A)} \{\nabla(T\eta)(\alpha) \in \mathcal{G} \mid \nabla(T\eta)(\alpha) \leq \nabla(T_\omega \bigwedge)(\Phi)\}. \end{aligned}$$

Finally, the general case (that is, for an arbitrary element of $\mathbb{M}\mathbb{P}(X)$) can be obtained from the cases above using standard Boolean reasoning. \blacksquare

As a corollary of the proof of Proposition 48 we obtain the following one-step normal form theorem.

COROLLARY 49. *For $a \in \mathcal{L}PX$ we have*

$$a = \bigvee \{\nabla(T\eta)(\alpha) \mid \alpha \in \delta_X(a) \cap T_\omega X\}.$$

In case that T preserves finite sets, the join is finite for finite sets X and can be expressed in the language. Induction along the sequence of the \mathcal{L}_n then yields a normal form theorem for \mathcal{L} .

5.6 Completeness

The following proposition, going back to [13], is a standard result in coalgebra based on δ being injective and \mathbb{M} being finitary and preserving injective maps.

PROPOSITION 50. *Suppose $a \not\leq b$ in the initial \mathbb{M} -algebra. Then there is a T -coalgebra (S, σ) and $s \in S$ such that $s \Vdash a$ and not $s \Vdash b$.*

Proof. (Sketch) To explain the idea of the proof assume first that a final T -coalgebra $\zeta : Z \rightarrow TZ$ exists. Then we would prove that the unique \mathbb{M} -algebra morphism $[[\cdot]]$ from the initial \mathbb{M} -algebra to $\mathbb{M}\mathbb{P}Z \xrightarrow{\delta} T\mathbb{P} \xrightarrow{\zeta^{-1}} \mathbb{P}Z$ is injective. Indeed, $a \not\leq b$ then implies $[[a]] \not\leq [[b]]$, ie there is $z \in Z$ such that $(Z, \zeta), z \Vdash a$ and $(Z, \zeta), z \not\Vdash b$.

Since the assumption of the existence of a final coalgebra excludes important examples such as Kripke frames or models, we replace the final coalgebra by the corresponding final sequence $(T^n 1)_{n < \omega}$, which is defined as follows. We denote by $1 = T^0 1$ the final object in **Set**. The map $p_0 : T1 \rightarrow 1$ is given by finality and $p_{n+1} : T(T^n 1) \rightarrow T^n 1$ is defined to be Tp_n . It is easy to see that each p_n is surjective. We think of the $T^n 1$ as approximating the final coalgebra. (Indeed, if we let run the final sequence through all ordinals, we obtain the final coalgebra as a limit if it exists, see [3].) In the same way as any coalgebra $\xi : X \rightarrow TX$ has a unique arrow into the final coalgebra, there are canonical ‘ n -step behavior maps’, that is, arrows $\xi_n : X \rightarrow T^n 1$ to the approximants of the final coalgebra: $\xi_0 : X \rightarrow 1$ is given by finality and $\xi_{n+1} = T(\xi_n) \circ \xi$.

Recall that we may consider \mathcal{M} , the initial \mathbb{M} -algebra, as a union of its approximants $\mathbb{M}^n \mathfrak{D}$. Elements of $\mathbb{M}^n \mathfrak{D}$ correspond to formulas of depth n and we define their semantics wrt the final sequence of T as a BA-morphism $\llbracket - \rrbracket_n : \mathbb{M}^n \mathfrak{D} \rightarrow \mathbb{P}T^n 1$ as follows.

$$\begin{array}{ccccc}
 \mathbb{P}1 & \xrightarrow{\mathbb{P}p_0} & \dots & & \mathbb{P}T^n 1 & \xrightarrow{\mathbb{P}p_n} & \mathbb{P}T^{n+1} 1 & \dots & (18) \\
 \uparrow \llbracket - \rrbracket_0 & & & & \uparrow \llbracket - \rrbracket_n & & \uparrow \llbracket - \rrbracket_{n+1} & & \\
 \mathfrak{D} & \xrightarrow{j_0} & \dots & & \mathbb{M}^n \mathfrak{D} & \xrightarrow{j_n} & \mathbb{M}^{n+1} \mathfrak{D} & \dots &
 \end{array}$$

$\llbracket - \rrbracket_0$ is given by initiality (and is actually the identity). For the definition of $\llbracket - \rrbracket_{n+1}$, recall that $\delta_{T^n 1} : \mathbb{M}\mathbb{P}T^n 1 \rightarrow \mathbb{P}T^{n+1} 1$, and assume inductively that $\llbracket - \rrbracket_n : \mathbb{M}^n \mathfrak{D} \rightarrow \mathbb{P}T^n 1$ has been defined, so that $\mathbb{M}(\llbracket - \rrbracket_n) : \mathbb{M}^{n+1} \mathfrak{D} \rightarrow \mathbb{M}\mathbb{P}T^n 1$. Composing these two maps, we obtain $\llbracket - \rrbracket_{n+1} := \delta_{T^n 1} \circ \mathbb{M}(\llbracket - \rrbracket_n)$.

Observe that the semantics of a formula is independent of the particular approximant we choose (all squares in the diagram commute). Moreover, given a coalgebra $\xi : X \rightarrow TX$ and $a \in \mathbb{M}^n \mathfrak{D}$, the semantics via the initial \mathbb{M} -algebra and the semantics via the final sequence coincide: $\llbracket a \rrbracket_{(X, \xi)} = \xi_n^{-1}(\llbracket a \rrbracket_n)$. Since δ is injective (Proposition 48) and \mathbb{M} preserves embeddings (Proposition 28), a straightforward inductive proof shows that all $\llbracket - \rrbracket_n$, $n \in \omega$, are injective.

To show the claim now, suppose $a \not\leq b$ in the initial \mathbb{M} -algebra. We find an approximant $\mathbb{M}^n \mathfrak{D}$, in which $a \not\leq_n b$. Choosing a half-inverse h of p_0 , we let $\xi : T^n 1 \rightarrow TT^n 1$ be $T^n(h)$. ξ provides $T^n 1$ with T -coalgebra structure. Now injectivity of $\llbracket - \rrbracket_n$ shows that $(T^n 1, \xi)$ provides a counter-example for $a \leq b$. ■

The proof of Theorem 16 is now a corollary. Reasoning by contraposition, take formulas $a, b \in \mathcal{L}$ such that $\not\vdash_{\mathbb{M}} a \leq b$. By Proposition 39, $a \not\leq b$ in \mathcal{M} . Now, completeness follows from Propositions 50 and 45.

BIBLIOGRAPHY

- [1] S. Abramsky. Domain theory in logical form. *Ann. Pure Appl. Logic*, 51, 1991.
- [2] P. Aczel. *Non-Well-Founded Sets*. CSLI, Stanford, 1988.

- [3] J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theoret. Comput. Sci.*, 150, 1995.
- [4] J. Adámek and V. Trnková. *Automata and Algebras in Categories*. Kluwer, 1990.
- [5] M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality, 2008. Same volume.
- [6] M. Bonsangue and A. Kurz. Duality for logics of transition systems. In *FoSSaCS'05*.
- [7] C. Cirstea and D. Pattinson. Modular proof systems for coalgebraic logics. *Theoret. Comp. Sci.*, 338, 2007.
- [8] G. D'Agostino and G. Lenzi. On modal μ -calculus with explicit interpolants. *Journal of applied logic*, 338, 2006.
- [9] K. Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16, 1975.
- [10] B. Jacobs. Trace Semantics for Coalgebras. In *CMCS'04*.
- [11] B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theor. Inform. Appl.*, 35, 2001.
- [12] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In *Proc. MFCS'95*.
- [13] C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In *CMCS'04*.
- [14] C. Kupke and Y. Venema. Coalgebraic automata theory: basic results. *Logical Methods in Computer Science*, to appear.
- [15] A. Kurz. Specifying coalgebras with modal logic. *Theoret. Comput. Sci.*, 260, 2001.
- [16] A. Kurz. Coalgebras and their logics. *SIGACT News*, 37, 2006.
- [17] A. Kurz and D. Petrişan. Functorial coalgebraic logic: The case of many-sorted varieties. In *CMCS'08*.
- [18] L. Moss. Coalgebraic logic. *Ann. Pure Appl. Logic*, 96, 1999.
- [19] A. Palmigiano and Y. Venema. Nabla algebras and Chu spaces. In *CALCO'07*.
- [20] D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309, 2003.
- [21] M. Rößiger. From modal logic to terminal coalgebras. *Theoret. Comput. Sci.*, 260, 2001.
- [22] J. Rutten. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.*, 249, 2000.
- [23] L. Santocanale and Y. Venema. Completeness for flat modal fixpoint logics (extended abstract). In *LPAR'07*.
- [24] L. Schröder. Expressivity of Coalgebraic Modal Logic: The Limits and Beyond. In *FoSSaCS'05*.
- [25] R. Sikorski. A theorem on extensions of homomorphisms. *Annals of Polish Mathematical Society*, 21, 1948.
- [26] Y. Venema. Algebras and coalgebras. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 331–426. Elsevier, 2006.
- [27] Y. Venema. Automata and fixed point logics: a coalgebraic perspective. *Inf. Comput.*, 204, 2006.

Clemens Kupke: Centrum voor Wiskunde en Informatica (CWI), Kruislaan 413, 1098 SJ Amsterdam. *E-mail*: kupke@cwi.nl

Alexander Kurz: Department of Computer Science, University of Leicester, UK. *E-mail*: kurz@mcs.le.ac.uk

Yde Venema: Institute for Logic, Language and Computation, Universiteit van Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam. *E-mail*: yde@science.uva.nl