

Completeness for the coalgebraic cover modality

Clemens Kupke* and Alexander Kurz† and Yde Venema‡

February 2, 2012

Abstract

We study the finitary version of the coalgebraic logic introduced by L. Moss. The syntax of this logic, which is introduced uniformly with respect to a coalgebraic type functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, extends that of classical propositional logic with the so-called coalgebraic cover modality ∇_T . The semantics of ∇_T is defined in terms of a categorically defined relation lifting operation \bar{T} .

As the main contributions of our paper we introduce a derivation system \mathbf{M} , and prove that \mathbf{M} provides a sound and complete axiomatization for the collection of coalgebraically valid inequalities. Our soundness and completeness proof is algebraic, and we employ Pattinson's stratification method, showing that our derivation system can be stratified in ω many layers, corresponding to the modal depth of the formulas involved.

In the proof of our main result we identify some new concepts and obtain some auxiliary results of independent interest. We survey properties of the notion \bar{T} of relation lifting, induced by an arbitrary but fixed set functor T . We introduce a category \mathbf{Pres} of Boolean algebra presentations, and establish an adjunction between \mathbf{Pres} and the category \mathbf{BA} of Boolean algebras. Given the fact that our derivation system \mathbf{M} involves only formulas of depth one, it can be encoded as a functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$. We show that this functor is finitary and preserves embeddings, and we prove that the Lindenbaum-Tarski algebra of our logic can be identified with the initial \mathbb{M} -algebra.

All our definitions, results and proofs are uniform in the coalgebraic type functor T , which we require to preserve weak pullbacks and inclusions.

Keywords Coalgebra, modal logic, relation lifting, completeness, cover modality, presentations by generators and relations

Classification codes: 18C50, 03B45

*Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, United Kingdom. E-mail: ckupke@doc.ic.ac.uk. Supported by grant EP/F031173/1 from the UK EPSRC.

†Department of Computer Science, University of Leicester, United Kingdom. E-mail: kurz@mcs.le.ac.uk. Partially supported by grant EP/C014014/1 from EPSRC.

‡Institute for Logic, Language and Computation, Universiteit van Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands. E-mail: y.venema@uva.nl. The research of this author has been made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO).

1 Introduction

Coalgebra, introduced to computer science by Aczel in the late 1980s [1, 2], is rapidly gaining ground as a general mathematical framework for many kinds of state-based evolving systems. Examples of coalgebras include data streams, (infinite) labelled trees, Kripke structures, finite automata, (probabilistic/weighted) transition systems, neighborhood models, and many other familiar structures. As emphasized by Rutten [46], who developed, in analogy with Universal Algebra, the theory of Universal Coalgebra as a general theory of such transition systems, the coalgebraic viewpoint combines wide applicability with mathematical simplicity. In particular, one of the main advantages of the coalgebraic approach is that a substantial part of the theory of systems can be developed *uniformly* in a functor T which represents the *type* of the coalgebras we are dealing with. Here we restrict attention to *systems*, where T is an endofunctor on the category \mathbf{Set} of sets with functions, so that a T -coalgebra is a pair of the form

$$\mathbb{X} = \langle X, \xi : X \rightarrow TX \rangle$$

with the set X being the carrier or state space of the coalgebra, and the map ξ its unfolding or transition map. Many important notions, properties, and results of systems can be explained just in terms of properties of their type functors. As a key example, any set functor T canonically induces a notion of observational or *behavioural equivalence* between T -coalgebras; this notion generalizes the natural notions of bisimilarity that were independently developed for each specific type of system.

In order to describe and reason about the kind of behaviour modelled by coalgebras, there is a clear need for the design of coalgebraic specification languages and derivation systems, respectively. The resulting research programme of *Coalgebraic Logic* naturally supplements that of Coalgebra by searching for logical formalisms that, next to meeting the usual desiderata such as striking a good balance between expressive power and computational feasibility, can be defined and studied uniformly in the functor T . Given the fact that Kripke models and frames are prime examples of coalgebras, it should come as no surprise that in search for suitable coalgebraic logics, researchers looked for inspiration to *modal logic* [16].

This research direction was initiated by Moss [41]; roughly speaking, his idea was to take the functor T *itself* as supplying a modality ∇_T , in the sense that for every element $\alpha \in T\mathcal{L}$ (where \mathcal{L} is the collection of formulas), the object $\nabla_T\alpha$ is a formula in \mathcal{L} . While Moss' work was recognized to be of seminal conceptual importance in advocating modal logic as a specification language for coalgebra, his particular formalism did not find much acclaim, for at least two reasons. First of all, the semantics of his modality is defined in terms of relation lifting, and for this to work smoothly, Moss needed to impose a restriction on the functor (the coalgebra type functor T is required to preserve weak pullbacks). Thus the scope of his work excluded some interesting and important coalgebras such as neighborhood models and frames. And second, for practical purposes, the syntax of Moss' language was considered to be rather unwieldy, with the nonstandard operator ∇_T looking strikingly different from the usual \Box and \Diamond modalities.

Following on from Moss' work, attention turned to the question how to obtain modal languages for T -coalgebras which use more standard modalities [36, 45, 28], and how to find derivation systems for these formalisms. This approach is now usually described in terms of predicate liftings [43, 49] or, equivalently, Stone duality [17, 37]. Other approaches towards coalgebraic logic, such as the one using co-equations [3] until now have received somewhat less attention. For a while, this development directed interest away from Moss' logic, and the relationship between various approaches towards coalgebraic logic was not completely clear.

In the mean time, however, it had become obvious that even in standard modal logic, a nabla-based approach has some advantages. In this setting the coalgebra type T is instantiated by the power set functor P , so that (the finitary version of) the nabla operator ∇_P , takes a (finite) *set* α of formulas and returns a single formula $\nabla_P\alpha$. The semantics of this so-called *cover modality* can be explicitly

formulated as follows, for an arbitrary Kripke structure \mathbb{X} with accessibility relation R :

$$\mathbb{X}, x \Vdash \nabla_P \alpha \quad \text{if} \quad \begin{array}{l} \text{for all } a \in \alpha \text{ there is a } t \in R[x] \text{ with } \mathbb{X}, t \Vdash a, \text{ and} \\ \text{for all } t \in R[x] \text{ there is an } a \in \alpha \text{ with } \mathbb{X}, t \Vdash a. \end{array} \quad (1)$$

In short: $\nabla_P \alpha$ holds at a state x iff the formulas in α and the set $R[x]$ of successors of x ‘cover’ one another. Readers familiar with classical first-order logic will recognize the quantification pattern underlying (1) from the theory of Ehrenfeucht-Fraïssé games, Scott sentences, and the like, see for instance [26]. In modal logic, related ideas made an early appearance in Fine’s work on normal forms [21].

Using the standard modal language, ∇_P can be seen as a defined operator:

$$\nabla_P \alpha = \Box \vee \alpha \wedge \bigwedge \diamond \alpha, \quad (2)$$

where $\diamond \alpha$ denotes the set $\{\diamond a \mid a \in \alpha\}$. But is in fact an easy exercise to prove that with ∇_P defined by (1), we have the following semantic equivalences:

$$\begin{array}{l} \diamond \alpha \equiv \nabla_P \{\alpha, \top\} \\ \Box \alpha \equiv \nabla_P \emptyset \vee \nabla_P \{\alpha\} \end{array} \quad (3)$$

In other words, the standard modalities \Box and \diamond can be defined in terms of the nabla operator (together with \vee and \top). When combined, (2) and (3) show that the language based on the nabla operator offers an alternative formulation of standard modal logic.

In fact, independently of Moss’ work, Janin & Walukiewicz [30] had already made the much stronger observation that the set of connectives $\{\Box, \diamond, \wedge, \vee\}$ may in some sense be replaced by the connectives ∇_P and \vee , that is, without the conjunction operation. This fact, 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 [20, 47]. These observations naturally led Venema [55] to introduce, parametric in the coalgebraic type functor T , a finitary version of Moss’ logic, extended with fixpoint operators, and to generalize the link between fixpoint logics and automata theory to the coalgebraic level of generality. Subsequently, Kupke & Venema [35] showed that many fundamental results in automata theory and fixpoint logics are really theorems of universal coalgebra. The key role of the nabla modality in these results revived interest in Moss’ logic.

Our paper addresses the main problem left open in the literature on ∇ -based coalgebraic logic, namely that of providing a *sound and complete derivation system* for the logic. Moss’ approach is entirely semantic, and does not provide any kind of syntactic calculus. As a first result in the direction of a derivation system for nabla modalities, Palmigiano & Venema [42] gave a complete axiomatization for the cover modality ∇_P . This calculus was streamlined into a formulation that admits a straightforward generalization to an arbitrary set functor T , by Bílková, Palmigiano & Venema [13], who also provided suitable Gentzen systems for the logic based on ∇_P . In this paper we will prove the soundness and completeness of this axiomatization in the general case.

In the remaining part of the introduction we briefly survey the paper, its main contributions, and its proof method. Throughout the paper we let T denote the coalgebraic type functor; usually we make the proviso that T preserves weak pullbacks and inclusions (all of this will be discussed further on in detail). Our key instrument in making Moss’ language more standard is to base its syntax on the *finitary* version T_ω of the functor T which is defined on objects as follows: for a set X , $T_\omega X := \bigcup \{TY \mid Y \subseteq_\omega X\}$. As we will discuss in detail, for each object $\alpha \in T_\omega X$ there is a *minimal* finite set $Base_X(\alpha) \subseteq_\omega X$ such that $\alpha \in T Base_X(\alpha)$, and the maps $Base_X$ provide a natural transformation

$$Base : T_\omega \rightarrow P_\omega.$$

The formulas of our coalgebraic language \mathcal{L} can now be defined by the following grammar:

$$a ::= \neg a \mid \bigwedge \varphi \mid \bigvee \varphi \mid \nabla_T \alpha.$$

where $\varphi \in P_\omega \mathcal{L}$ and $\alpha \in T_\omega \mathcal{L}$. That is, the propositional basis of our coalgebraic language \mathcal{L} takes the finitary conjunction (\bigwedge) and disjunction (\bigvee) connectives as primitives, and to this we add the coalgebraic modality ∇_T , which returns a formula $\nabla_T \alpha$ for every object $\alpha \in T_\omega \mathcal{L}$. The point of restricting Moss' modality to the set $T_\omega \mathcal{L}$ is that the formula $\nabla_T \alpha$ has a finite, clearly defined set of immediate subformulas, namely the set $Base(\alpha)$; thus every formula has a finite set of subformulas.

The key observation of Moss [41] was that the *semantics* (1) of ∇ can be expressed in terms of the so-called Egli-Milner *lifting* of the satisfaction relation $\Vdash \subseteq X \times \mathcal{L}$. Generalizing this observation from the Kripke functor P to the arbitrary type T , he uniformly defined the semantics of ∇_T in a T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$ as follows:

$$\mathbb{X}, x \Vdash \nabla_T \alpha \text{ iff } \xi(x) \overline{T} \Vdash \alpha.$$

Here $\overline{T} \Vdash$ denotes a categorically defined *lifting* of the satisfaction relation $\Vdash \subseteq X \times \mathcal{L}$ between states and formulas to a relation $\overline{T} \Vdash \subseteq TX \times T\mathcal{L}$. Given the importance of the *relation lifting* operation \overline{T} in Moss' logic, we include in this paper a fairly detailed survey of its properties and related concepts.

The coalgebraic *validities*, that is, the formulas that are true at every state of every T -coalgebra thus constitute a semantically defined coalgebraic *logic*, and it is this logic that we will axiomatize in this paper. Our approach will be *algebraic* in nature, and so it will be convenient to work with equations, or rather, inequalities (expressions of the form $a \preceq b$, where a and b are terms/formulas of the language).

We obtain our derivation system for Moss' logic by extending a sound and complete derivation system for propositional logic with three rules for the ∇ -operator. The first rule, denoted by ($\nabla 1$), can be seen as a combined monotonicity and congruence rule. Rule ($\nabla 2$) is a distributive law that expresses that any conjunction of ∇ -formulas is equivalent to a (possibly infinite) disjunction of ∇ -formulas built from conjunctions. Finally, rule ($\nabla 3$) expresses that ∇ distributes over disjunctions. In the case that the functor T under consideration maps finite sets to finite sets, the rules ($\nabla 2$) and ($\nabla 3$) take the form of axioms.

The proof of our soundness and completeness theorem is based on the *stratification method* of Pattinson [43]. We will show that not only the *language* of our system, but also its *semantics* and our *derivation system* can be stratified in ω many layers corresponding to the modal depth of the formulas involved. (This means for instance that if two formulas of depth n are provably equivalent, this can be demonstrated by a derivation involving only formulas of depth at most n .) What glues these layers nicely together can be formulated in terms of properties of a one-step version of the derivation system \mathbf{M} .

In our algebraic approach, this one-step version of \mathbf{M} is incarnated as a *functor* on the category of Boolean algebras:

$$\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}.$$

To mention a few interesting properties of this functor, of which the definition is uniformly parametrized by the functor T : \mathbb{M} is finitary, and preserves atomicity of Boolean algebras, and injectivity of homomorphisms. We will be interested in algebras for the functor \mathbb{M} , and in particular, we will see that the *initial* \mathbb{M} -algebra can be seen as the Lindenbaum-Tarski algebra of our derivation system \mathbf{M} .

For the definition of \mathbb{M} , we need to go into quite a bit of detail concerning the theory of *presentations* of (Boolean) algebras. In particular, we define a *category* \mathbf{Pres} of presentations by introducing a suitable notion of presentation morphism, and establish an adjunction between the categories \mathbf{Pres} and \mathbf{BA} :

$$\mathbf{BA} \begin{array}{c} \xleftarrow{B} \\ \perp \\ \xrightarrow{C} \end{array} \mathbf{Pres} \quad (4)$$

This adjunction (which is almost an equivalence) is the instrument that allows us to turn the modal rule and axioms of \mathbf{M} into the functor \mathbb{M} ; the key property that makes this work is that all modal rules and axioms of \mathbf{M} are formulated in terms of *depth-one* formulas.

What is left to do, in order to prove the soundness and completeness of our logic, is connect the algebra functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$ (that is, the ‘logic’) to the coalgebra functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ (the ‘semantics’). Here we will apply a well-known method in coalgebraic logic [17, 37] which is often described in terms of *Stone duality* because its aim is to link functors on two different base categories that are connected themselves by a Stone-type duality or adjunction.

In our case, to make the connection between \mathbb{M} and T we invoke the already existing link on the level of the base logic, provided by the (contravariant) power set functor $\check{\mathbb{P}}$ from \mathbf{Set} to \mathbf{BA} (we do not need its adjoint functor sending a Boolean algebra to its set of ultrafilters):

$$\mathbb{M} \left(\curvearrowright \mathbf{BA} \xleftarrow{\check{\mathbb{P}}} \mathbf{Set} \left(\curvearrowright \right) T \quad (5)$$

The key remaining step in the completeness proof involves the definition of a natural transformation

$$\delta : \mathbb{M}\check{\mathbb{P}} \rightarrow \check{\mathbb{P}}T.$$

As usual in the Stone duality approach towards coalgebraic logic, the *existence* of δ corresponds to the *soundness* of the logic. To get an idea of why this is the case, observe that the existence of δ enables us to see a T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$ as an \mathbb{M} -algebra, namely its *complex algebra* $\mathbb{X}^* := \langle \check{\mathbb{P}}X, \check{\mathbb{P}}\xi \circ \delta_X \rangle$. Finally, as we will see in the final part of our stratification-based proof, the *completeness* of \mathbf{M} is based on the observation that

$$\delta \text{ is injective,} \quad (6)$$

that is, for each set X , the \mathbf{BA} -homomorphism $\delta_X : \mathbb{M}\check{\mathbb{P}}X \rightarrow \check{\mathbb{P}}TX$ is an *embedding*. The proof of (6), which technically forms the heart of our proof, is based on the fact that the nabla-axioms allow us to write depth-one formulas into a certain normal form, and on the earlier mentioned properties of the functor \mathbb{M} .

This paper replaces, extends and partly corrects (c.q. clarifies, see Remark 7.8) an earlier version [34]. The main differences with respect to [34] are the following. First of all, we provide a detailed, self-contained overview of the notion of relation lifting and its properties (which was only covered as Fact 3 in the mentioned paper). Second, our categorical treatment of presentations and the algebras they present (which is novel to the best of our knowledge) clarifies and substantially extends the treatment in [34]. Third, our axiomatization simplifies the earlier one; in particular, we show here in detail that we do not need axioms or rules specifically dealing with negation (more specifically, we prove that an earlier rule ($\nabla 4$) is derivable in the system here. Fourth, we provide a more precise definition and a more detailed discussion of the functor \mathbb{M} ; for instance, the result that \mathbb{M} preserves atomicity is new. Fifth and final, we show here in much more detail and precision how the soundness and completeness of our axiomatization follows from the one-step soundness and completeness.

Overview In the next section we fix our notation, introduce the necessary basic (co-)algebraic terminology and discuss properties of functors on the category of sets that will play an important role in our paper. After that, in Section 3, we recall the notion of a relation lifting \overline{T} induced by a set functor T and give an overview of its properties. Section 4 and Section 5 introduce the terminology that we need concerning Boolean algebras and their presentations, and concerning Moss’ coalgebraic logic, respectively.

After that we move to the main results of our paper. First, in Section 6 we introduce the derivation system for Moss’ coalgebraic logic and we define the algebra functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$. In Section 7 we

prove that our derivation system is one-step sound and complete. Within the above described categorical framework this is equivalent to establishing the existence of a natural transformation $\delta : \mathbb{MIP} \rightarrow \mathbb{PT}$ (one-step soundness) and proving that this transformation δ is injective (one-step completeness). Finally, in Section 8 we prove our main result, namely soundness and completeness of our derivation system with respect to the coalgebraic semantics. We conclude with an overview of related work and open questions.

Finally, since this paper features a multitude of categories, functors and natural transformations, for the reader's convenience we list these in an appendix.

Acknowledgement We thank the anonymous referee for many useful comments.

2 Preliminaries

The purpose of this section is to fix our notation and terminology, and to introduce some concepts that underlie our work in all other parts of the paper.

2.1 Basic mathematics and category theory

First we fix some basic mathematical issues. Given a set X , we let PX and $P_\omega X$ denote the power set and the finite power set of X , respectively. We write $Y \subseteq_\omega X$ to indicate that Y is a finite subset of X .

Given a relation $R \subseteq X \times X'$, we denote the *domain* and *range* of R by $\text{dom}(R)$ and $\text{rng}(R)$, respectively, and we denote by $\pi_1^R : R \rightarrow X$ its first projection and by $\pi_2^R : R \rightarrow X'$ its second projection map. Given subsets $Y \subseteq X$, $Y' \subseteq X'$, the *restriction* of R to Y and Y' is given as

$$R \upharpoonright_{Y \times Y'} := R \cap (Y \times Y').$$

The *converse* of a relation $R \subseteq X \times X'$ is denoted as $R^\sim \subseteq X' \times X$.

The *composition* of two relations $R \subseteq X \times X'$ and $R' \subseteq X' \times X''$ is denoted by $R ; R'$, while the composition of two functions $f : X \rightarrow X'$ and $f' : X' \rightarrow X''$ is denoted by $f' \circ f$. That is, we denote function composition by \circ and write it from right to left and we denote relation composition of relations by $;$ and write it from left to right.

It is often convenient to identify a function $f : X \rightarrow X'$ with its *graph*, that is, the relation $Gr(f) = \{(x, f(x)) \mid x \in X\} \subseteq X \times X'$. For example given a relation $R \subseteq X \times X'$ and a function $f : X' \rightarrow X''$ we write $R ; f$ to denote the composition of relations $R ; Gr(f)$.

We will assume familiarity with basic notions from category theory, including those of categories, functors, natural transformations, (co-)monads and (co-)limits; see for instance [40]. We denote by **Set** the category of sets and functions, and by **Rel** the category of sets and binary relations. **BA** is the category with Boolean algebras as objects and homomorphisms as arrows.

Endofunctors on **Set** will simply be called *set functors*. We denote by P the *power set functor* which maps a set X to its power set PX and a function $f : X \rightarrow X'$ to its *direct image* $Pf : PX \rightarrow PX'$, given by $P(X) \ni Y \mapsto \{f(y) \mid y \in Y\}$. Similarly, $P_\omega X$ denotes the finite power set functor. P is in fact (part of) a monad (P, μ, η) , with $\eta_X : X \rightarrow P(X)$ denoting the singleton map $\eta_X : x \mapsto \{x\}$, and $\mu_X : PPX \rightarrow PX$ denoting union, $\mu_X(\mathcal{A}) := \bigcup \mathcal{A}$. The contravariant power set functor will be denoted as \check{P} ; this functor maps a set X to its power set $\check{P}X = PX$, and a function $f : X \rightarrow X'$ to its *inverse image* $\check{P}f : \check{P}X' \rightarrow \check{P}X$ given by $\check{P}X' \ni Y' \mapsto \{x \in X \mid fx \in Y'\}$.

2.2 (Co-)algebras

We provide some details concerning the notions of an algebra and a coalgebra for a functor. We start with coalgebras since these provide the semantic structures of the logics considered in this paper.

Definition 2.1 Given a functor T on a category \mathbf{C} , a T -coalgebra (X, ξ) is an arrow $\xi : X \rightarrow TX$ in \mathbf{C} ; a T -coalgebra morphism $f : (X, \xi) \rightarrow (X', \xi')$ is an arrow $f : X \rightarrow X'$ such that $Tf \circ \xi = \xi' \circ f$, in a diagram:

$$\begin{array}{ccc} X & \xrightarrow{f} & X' \\ \xi \downarrow & & \downarrow \xi' \\ TX & \xrightarrow{Tf} & TX' \end{array}$$

The functor T is called the *type* of the coalgebra (X, ξ) . The category of T -coalgebras is denoted by $\mathbf{Coalg}(T)$ and we denote coalgebras by capital letters $\mathbb{X}, \mathbb{Y}, \dots$ in blackboard bold.

In the case of a set coalgebra (that is, a coalgebra for a set functor), elements of the (carrier of the) coalgebra will be called *states* of the coalgebra, and a *pointed coalgebra* is a pair consisting (\mathbb{X}, x) consisting of a coalgebra $\mathbb{X} = (X, \xi)$ and a state x of \mathbb{X} . \triangleleft

Here are some simple, standard examples of coalgebras for set functors.

Example 2.2 1. We let Id denote the *identity* functor on \mathbf{Set} . Given a set C , we let C itself also denote the *constant* functor, mapping every set X to C , and every function f to the identity map id_C on C . Coalgebras for this functor are called C -*colorings*; in case C is of the form $P(\mathbf{Prop})$ for some set \mathbf{Prop} of proposition letters, we may think of a coloring $\xi : X \rightarrow C$ as a **Prop-valuation** (in the sense that ξ says of every proposition letter p and every state x whether p is true of x or not).

2. A *Kripke frame* $\langle S, R \rangle$ can be represented as a coalgebra $\langle S, \sigma_R \rangle$ for the power set functor P , with $\sigma_R : S \rightarrow PS$ mapping a point s to its collection of successors. It is left as an exercise for the reader to verify that the coalgebra morphisms for this functor precisely coincide with the *bounded morphisms* of modal logic.

3. Coalgebras for the functor $\check{P} \circ \check{P}$ (that is, the contravariant power set functor composed with itself) can be identified with the *neighborhood frames* known from the theory of modal logic as structures that generalize Kripke frames. As a special case of this, but also generalizing Kripke frames, the *monotone neighborhood functor* N maps a set X to the collection $N(X) := \{ \alpha \in \check{P}\check{P}X \mid \alpha \text{ is upward closed} \}$, and a function f to the map $\check{P}\check{P}f$.

4. For a slightly more involved example, consider the finitary *multiset* or *bag* functor B_ω . This functor takes a set X to the collection $B_\omega X$ of maps $\mu : X \rightarrow \mathbb{N}$ of finite support (that is, for which the set $Supp(\mu) := \{x \in X \mid \mu(x) > 0\}$ is finite), while its action on arrows is defined as follows. Given an arrow $f : X \rightarrow X'$ and a map $\mu \in B_\omega X$, we define $(B_\omega f)(\mu) : X' \rightarrow \mathbb{N}$ by putting

$$(B_\omega f)(\mu)(x') := \sum \{ \mu(x) \mid f(x) = x' \}.$$

5. As a variant of B_ω , consider the finitary probability functor D_ω , where $D_\omega X = \{ \delta : X \rightarrow [0, 1] \mid Supp(\delta) \text{ is finite and } \sum_{x \in X} \delta(x) = 1 \}$, while the action of D_ω on arrows is just like that of B_ω .

Example 2.3 Many examples of coalgebraically interesting set functors are obtained by *composition* of simpler functors. Inductively define the following class *EKPF* of *extended Kripke polynomial functors*:

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

where \circ , $+$ and \times denote functor composition, coproduct (or disjoint union) and product, respectively, and $(-)^D$ denotes exponentiation with respect to some set D . Examples of such functors include:

1. Given an alphabet-color set C , the C -streams are simple specimens of coalgebras for the functor $C \times Id$; similarly, C -labelled binary trees are coalgebras for the functor $B_C = C \times Id \times Id$.
2. *Labelled transition systems* over a set A of atomic actions can be seen as coalgebras for the functor $P(-)^A$.
3. *Deterministic automata* are coalgebras for the functor $(-)^{\Sigma} \times 2$ where Σ is the finite alphabet.
4. *Kripke models* over a set \mathbf{Prop} of proposition letters can be identified with coalgebras for the functor $P(\mathbf{Prop}) \times P(-) = P \circ C_{\mathbf{Prop}} \times P \circ Id$.
5. Generalizing the previous example, viewing T -coalgebra as *frames*, we can define T -models over a set \mathbf{Prop} of proposition letters as coalgebras for the functor $T_{\mathbf{Prop}} = P(\mathbf{Prop}) \times T(-)$.

As running examples through this paper we will often take the binary tree functor over a set C of colors, and the power set functor.

The key notion of equivalence in coalgebra is of two states in two coalgebras being *behaviorally equivalent*. In case the functor T admits a final coalgebra $\mathbb{Z} = \langle Z, \zeta \rangle$ the elements of Z often provide an intuitive encoding of the notion of behaviour, and the unique coalgebra homomorphism $!_{\mathbb{X}}$ can be seen as a map that assigns to a state x in \mathbb{X} its *behaviour*. In this case we call two states, x in \mathbb{X} and x' in \mathbb{X}' , *behaviorally equivalent* if $!_{\mathbb{X}}(x) = !_{\mathbb{X}'}(x')$. In the general case, when we may not assume the existence of a final coalgebra, we define the notion as follows.

Definition 2.4 Two elements (often called states) x, x' in two coalgebras \mathbb{X} and \mathbb{X}' , respectively, are *behaviorally equivalent* iff there are coalgebra morphisms f, f' with a common codomain such that $f(x) = f'(x')$. ◁

Turning to the dual notion of algebra, we shall use algebras mainly to describe logics for coalgebras, and the notion of an algebra ‘for a functor’ will provide us with an elegant way to exploit the duality with coalgebras.

Definition 2.5 Given a functor L on a category \mathbf{A} , an L -algebra (A, α) is an arrow $\alpha : LA \rightarrow A$ in \mathbf{A} and an L -algebra morphism $f : (A, \alpha) \rightarrow (A', \alpha')$ is an arrow $f : A \rightarrow A'$ such that $f \circ \alpha = \alpha' \circ Lf$. The category of L -algebras is denoted by $\mathbf{Alg}(L)$. ◁

Example 2.6 1. If $\mathbf{A} = \mathbf{Set}$, then every signature (or similarity type) induces a functor $LX = \coprod_{n < \omega} Op_n \times X^n$ where Op_n is the set of operation symbols of arity n . Then $\mathbf{Alg}(L)$ is (isomorphic to) the category of algebras for the signature.

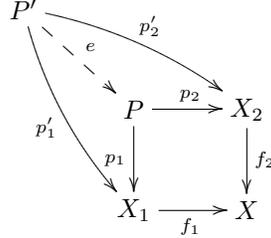
2. If $\mathbf{A} = \mathbf{BA}$, then we can define a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ to map an algebra A to the algebra LA generated by $\square a, a \in A$, and quotiented by the relation stipulating that \square preserves finite meets. Then $\mathbf{Alg}(L)$ is isomorphic to the category of modal algebras [33].

As the second example above shows, functors on \mathbf{BA} give rise to modal logics extending Boolean algebras with operators.

2.3 Properties of set functors

As mentioned in the introduction, in this paper we will restrict our attention to set functors satisfying certain properties. The first one of these is crucial.

Weak pullback preservation Recall that a set P together with functions $p_1 : P \rightarrow X_1$ and $p_2 : P \rightarrow X_2$ is a *pullback* of two functions $f_1 : X_1 \rightarrow X$ and $f_2 : X_2 \rightarrow X$ if $f_1 \circ p_1 = f_2 \circ p_2$ and for all sets P' and all functions $p'_1 : P' \rightarrow X_1$, $p'_2 : P' \rightarrow X_2$ such that $f_1 \circ p'_1 = f_2 \circ p'_2$ there exists a *unique* function $e : P' \rightarrow P$ such that $p_i \circ e = p'_i$ for $i = 1, 2$.



If the function e is not necessarily unique we call (P, p_1, p_2) a *weak pullback*. Furthermore we call a relation $R \subseteq X_1 \times X_2$ a (weak) pullback of f_1 and f_2 if R together with the projection maps π_1^R and π_2^R is a (weak) pullback of f_1 and f_2 .

In the category of sets, (weak) pullbacks have a straightforward characterization

Fact 2.7 [22]. Given two functions $f_1 : X_1 \rightarrow X_3$ and $f_2 : X_2 \rightarrow X_3$, let

$$pb(f_1, f_2) := \{(x_1, x_2) \mid f_1(x_1) = f_2(x_2)\}.$$

Furthermore, given a set P with functions $p_1 : P \rightarrow X_1$ and $p_2 : P \rightarrow X_2$, let

$$e : y \mapsto (p_1(y), p_2(y)).$$

define a function $e : P \rightarrow pb(f_1, f_2)$. Then

- (1) (P, p_1, p_2) is a pullback of f_1 and f_2 iff $f_1 \circ p_1 = f_2 \circ p_2$ and e is an isomorphism.
- (2) (P, p_1, p_2) is a weak pullback of f_1 and f_2 iff $f_1 \circ p_1 = f_2 \circ p_2$ and e is surjective.

A functor T *preserves weak pullbacks* if it transforms every weak pullback (P, p_1, p_2) for f_1 and f_2 into a weak pullback (TP, Tp_1, Tp_2) for Tf_1 and Tf_2 . An equivalent characterization is to require T to *weakly preserve pullbacks*, that is, to turn pullbacks into weak pullbacks. Further on in Corollary 3.7, we will see yet another, and probably more motivating, characterization of this property.

Example 2.8 All the functors of Example 2.2 preserve weak pullbacks, except for the neighborhood functor and its monotone variant. It can be shown that the property of preserving weak pullbacks is preserved under the operations $\circ, +, \times$ and $(-)^D$, so that all extended polynomial Kripke functors (Example 2.3) preserve weak pullbacks.

Standard functors The second property that we will impose on our set functors is that of standardness. Given two sets X and X' such that $X \subseteq X'$, let $\iota_{X, X'}$ denote the inclusion map from X into X' . A weak pullback-preserving set functor T is *standard* if it *preserves inclusions*, that is, if $T\iota_{X, X'} = \iota_{TX, TX'}$ for every inclusion map $\iota_{X, X'}$.

Remark 2.9 Unfortunately the definition of standardness is not uniform throughout the literature. Our definition of standardness is taken from Moss [41], while for instance Adámek & Trnková [7] have an additional condition involving so-called distinguished points. Fortunately, the two definitions are equivalent in case the functor preserves weak pullbacks, see Kupke [32, Lemma A.2.12]. Since we almost exclusively consider standard functors that also preserve weak pullbacks, we have opted for the simpler definition.

For readers who are interested in some more details, fix sets 0,1 and 2 of of the corresponding sizes (0,1 and 2), respectively, and let e, o denote the two maps $e, o : 1 \rightarrow 2$. Then the second condition of

standardness in the sense of [7] can be phrased as the requirement that $T0 = \{x \in T1 \mid Ti(x) = To(x)\}$, in words: all distinguished points are standard.

In any case the restriction to standard functors is for convenience only, since every set functor is ‘almost standard’ [7, Theorem III.4.5]. That is, given an arbitrary set functor T , we may find a standard set functor T' such that the restriction of T and T' to all non-empty sets and non-empty functions are naturally isomorphic. The important observation about T' is that $\text{Alg}(T) \cong \text{Alg}(T')$ and $\text{Coalg}(T) \cong \text{Coalg}(T')$. Consequently, in our work we can assume without loss of generality that our functors are standard and we will do so whenever convenient.

Example 2.10 The finitary bag functor B_ω of Example 2.2 is not standard, but we may ‘standardize’ it by representing any map $\mu : X \rightarrow \mathbb{N}$ of finite support by its ‘positive graph’ $\{(x, \mu x) \mid \mu x > 0\}$. Similarly, the finite distribution functor D_ω can be standardized by identifying a probability distribution $\mu : X \rightarrow [0, 1] \in D_\omega X$ with the (finite) set $\{(x, \mu x) \mid \mu x > 0\}$.

Finitary functors Let T be a set functor that preserves inclusions. Then T is *finitary* or ω -*accessible* if, for all sets X ,

$$TX = \bigcup\{TY \mid Y \subseteq X, \text{finite}\}.$$

Generalizing the construction of P_ω from P , we can define, for any set functor T that preserves inclusions, its *finitary version* $T_\omega : \text{Set} \rightarrow \text{Set}$ by putting

$$\begin{aligned} T_\omega(X) &:= \bigcup\{TY \mid Y \subseteq_\omega X\}, \\ T_\omega(f) &:= Tf. \end{aligned}$$

It is easy to verify that T_ω preserves inclusions, is finitary and a subfunctor of T as we have a natural transformation $\tau_X : T_\omega X \hookrightarrow TX$. Given the definition of the action of T_ω on arrows, we shall often write Tf instead of $T_\omega f$.

In order to avoid confusion, we already mention the following fact, but we postpone its proof until subsection 3.3.

Proposition 2.11 *Let T be a standard set functor that preserves weak pullbacks. Then T_ω is also a standard functor that preserves weak pullbacks.*

The reason that we are interested in finitary functors is that we want our language to be *finitary*, in the sense that a formula has only finitely many subformulas. The key property of finitary functors that will make this possible, is that every $\alpha \in TX$ is supported by a finite subset of X , and in fact, there will always be a *minimal* such set.

Definition 2.12 Given a finitary functor T and an element $\alpha \in TX$, we define

$$\text{Base}_X^T(\alpha) := \bigcap\{Y \subseteq_\omega X \mid \alpha \in TY\}.$$

◁

We write Base^T rather than Base^{T_ω} , and in fact omit the superscript whenever possible.

Example 2.13 The following examples are easy to check: $\text{Base}_X^{\text{Id}} : X \rightarrow P_\omega X$ is the singleton map, $\text{Base}_X^P : P_\omega X \rightarrow P_\omega X$ is the identity map on $P_\omega X$, $\text{Base}_X^{B^C} : C \times X \times X \rightarrow P_\omega X$ maps the triple (c, x_1, x_2) to the set $\{x_1, x_2\}$, and Base^{D_ω} maps a finitary distribution to its support.

Proposition 2.14 *Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard functor that preserves weak pullbacks.*

(1) *For any $\alpha \in T_\omega X$, $Base_X^T(\alpha)$ is the smallest set Y such that $\alpha \in TY$.*

(2) *$Base^T$ provides a natural transformation $Base : T_\omega \rightarrow P_\omega$.*

Proof. Part (1) is proved in [55].

For the second part, consider a map $f : X \rightarrow X'$. We have to show $P_\omega f \circ Base_X = Base_{X'} \circ T_\omega f$. Fix $\alpha \in T_\omega X$ and write $B = Base_X(\alpha)$ and $B' = Base_{X'}(T_\omega f(\alpha))$. We need to prove $B' = f[B]$.

For the inclusion “ \subseteq ”, from

$$\begin{array}{ccc} T_\omega B & \longrightarrow & T_\omega(f[B]) \\ \downarrow & & \downarrow \\ T_\omega X & \xrightarrow{T_\omega f} & T_\omega X' \end{array}$$

we see that $f[B]$ supports $T_\omega f(\alpha)$ and, as B' is the smallest such, $B' \subseteq f[B]$ follows.

For the opposite inclusion “ \supseteq ”, since T_ω preserves weak pullbacks, the dotted arrow in

$$\begin{array}{ccc} 1 & \xrightarrow{T_\omega f(\alpha)} & T_\omega(B') \\ \downarrow \alpha & \searrow & \downarrow \\ T_\omega(f^{-1}(B')) & \longrightarrow & T_\omega(B') \\ \downarrow & & \downarrow \\ T_\omega X & \xrightarrow{T_\omega f} & T_\omega X' \end{array}$$

exists and shows that $\alpha \in T_\omega(f^{-1}(B'))$. By minimality of the base, it follows $B \subseteq f^{-1}(B')$, that is, $B' \supseteq f[B]$. QED

Remark 2.15 A stronger version of the previous proposition follows from results in [23]. Let us briefly sketch the details using the terminology of [23]. First of all note that it is not difficult to see that all finitary set functors preserve intersections. Therefore [23, Theorem 7.4] implies that $Base$ is sub-cartesian (not necessarily natural) and this implies together with [23, Theorem 8.1] that T preserves preimages iff $Base$ is natural. Any weak pullback preserving functor preserves preimages and thus this statement implies Proposition 2.14.

3 Relation Lifting

Given the key role that the lifting of binary relations plays in the semantics of Moss’ logic, we need to discuss the notion in some detail. After giving the formal definition, we mention some of the basic properties of relation lifting: first the ones that hold for any functor, then the ones for which we require the functor to preserve weak pullbacks, and finally, we see important technical properties of relation lifting that rest on the fact that the set functor under consideration is standard. We discuss the connection of the relation lifting with categorical distributive laws: as we will see later on, this connection plays an important role in the axiomatization of ∇ . Finally we introduce the notion of a slim redistribution, which is needed to formulate one of our axioms.

3.1 Basics

First we give the formal definition of relation lifting.

Definition 3.1 Let T be a set functor. Given a binary relation R between two sets X_1 and X_2 , we define the relation $\overline{TR} \subseteq TX_1 \times TX_2$ as follows:

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

The relation \overline{TR} will be called the T -lifting of R . ◁

In other words, we apply the functor T to the relation R , seen as a *span* $X_1 \xleftarrow{\pi_1} R \xrightarrow{\pi_2} X_2$, and define \overline{TR} as the image of TR under the product map $\langle T\pi_1, T\pi_2 \rangle$ obtained from the lifted projection maps $T\pi$ and $T\pi'$. In a diagram:

$$\begin{array}{ccccc} X_1 & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & X_2 \\ & & \downarrow & & \\ TX_1 & \xleftarrow{T\pi_1} & TR & \xrightarrow{T\pi_2} & TX_2 \\ & \searrow \langle T\pi_1, T\pi_2 \rangle & \downarrow \overline{TR} & \nearrow & \\ & & TX_1 \times TX_2 & & \end{array}$$

Let us first see some concrete examples.

Example 3.2 Fix two sets X and X' , and a relation $R \subseteq X \times X'$. For the identity and constant functors, we find, respectively:

$$\begin{aligned} \overline{Id}R &= R \\ \overline{C}R &= id_C. \end{aligned}$$

The relation lifting associated with the power set functor P can be defined concretely as follows:

$$\overline{P}R = \{(A, A') \in PX \times PX' \mid \forall a \in A \exists a' \in A'. aRa' \text{ and } \forall a' \in A' \exists a \in A. aRa'\}.$$

This relation is known under many names, of which we mention that of the *Egli-Milner* lifting of R . Relation lifting for the finitary multiset functor is slightly more involved: given two maps $\mu \in B_\omega X, \mu' \in B_\omega X'$, we put

$$\begin{aligned} \mu \overline{B}_\omega R \mu' \text{ iff there is some map } \rho : R \rightarrow \mathbb{N} \text{ such that } & \forall x \in X. \sum \{\rho(x, x') \mid x' \in X'\} = \mu(x) \\ & \text{and } \forall x' \in X'. \sum \{\rho(x, x') \mid x \in X\} = \mu'(x'). \end{aligned}$$

The definition of \overline{D}_ω is similar.

Finally, relation lifting interacts well with various operations on functors [25]. In particular, we have

$$\begin{aligned} \overline{T_0 \circ T_1}R &= \overline{T_0}(\overline{T_1}R) \\ \overline{T_0 + T_1}R &= \overline{T_0}R \cup \overline{T_1}R \\ \overline{T_0 \times T_1}R &= \{((\xi_0, \xi_1), (\xi'_0, \xi'_1)) \mid (\xi_i, \xi'_i) \in \overline{T_i}R, \text{ for } i \in \{0, 1\}\} \\ \overline{T^D}R &= \{(\varphi, \varphi') \mid (\varphi(d), \varphi'(d)) \in \overline{TR} \text{ for all } d \in D\}. \end{aligned}$$

From this one may easily calculate the relation lifting of all extended Kripke polynomial functors of Example 2.3.

Remark 3.3 Strictly speaking, when defining the T -lifting of a relation $R \subseteq X_1 \times X_2$, we should explicitly mention the type of R , that is, the pair of sets X_1 and X_2 .

To see this, let X_1, X_2, Y_1 and Y_2 be sets such that $Y_i \subseteq X_i$, for $i \in \{1, 2\}$. Now any relation $R \subseteq Y_1 \times Y_2$ can also be seen as a relation between X_1 and X_2 . But in general we do not have $TY_i \subseteq TX_i$, and so the relation $\overline{TR} \subseteq Y_1 \times Y_2$ is not necessarily a relation between X_1 and X_2 . It is easy to see that if T preserves inclusions, then this problem evaporates. Since we will assume T to be standard almost throughout the paper, we ignore this subtlety for the time being. Readers who are worried about this may add the condition that T preserves inclusions throughout the subsections 3.1 and 3.2.

Remark 3.4 Relation lifting can be used to define the notion of a *bisimulation* between two coalgebras. Recall that, given two coalgebras $\mathbb{X}_1 = \langle X_1, \xi_1 \rangle$ and $\mathbb{X}_2 = \langle X_2, \xi_2 \rangle$, a relation $Z \times X_1 \times X_2$ is a bisimulation if there is a coalgebra map $\zeta : Z \rightarrow TZ$ making the two projection functions $\pi_1 : Z \rightarrow X_1$ and $\pi_2 : Z \rightarrow X_2$ into coalgebra morphisms. It can be shown that this is equivalent to requiring that $\xi_1(x_1) \overline{TZ} \xi(x_2)$ whenever $x_1 Z x_2$.

As mentioned, in this section we will discuss some important properties of relation lifting. We start with listing a number of properties that T -lifting has for *any* given set functor T . The proof of the fact below is elementary.

Fact 3.5 *Let T be an arbitrary set functor. Then the relation lifting \overline{T}*

1. *extends T : $\overline{T}f = Tf$ for all functions $f : X_1 \rightarrow X_2$,*
2. *preserves the diagonal: $\overline{T}Id_X = Id_{TX}$ for any set X ;*
3. *is monotone: $R \subseteq Q$ implies $\overline{TR} \subseteq \overline{TQ}$ for all relations $R, Q \subseteq X_1 \times X_2$;*
4. *commutes with taking converse: $\overline{TR}^\vee = (\overline{TR})^\vee$ for all relations $R \subseteq X_1 \times X_2$.*

3.2 Weak pullback preserving functors

Fact 3.5 states a number of operations on relations that interact well with relation lifting. Conspicuously absent in that list is *relational composition*: observe that \overline{T} would be a *functor* on the category Rel if it would satisfy $\overline{T}(R; Q) = \overline{TR}; \overline{TQ}$. Here we arrive at the main reason why we are interested in functors that preserve weak pullbacks: as we will see now, that property is a necessary and sufficient condition on T for \overline{T} to be functorial.

In fact, given the characterisation of (weak) pullbacks in the category Set , in terms of the relation pb (see Fact 2.7), it is easy to formulate the composition $R; Q$ of two relations R and Q as a pullback of the projection maps π_2^R and π_1^Q . Therefore it is not surprising that the question whether the T -lifting of a relation commutes with the composition of relations is tightly connected with the preservation of weak pullbacks by T . The following fact was first proved in [54].

Fact 3.6 *A functor $T : \text{Set} \rightarrow \text{Set}$ weakly preserves pullbacks iff for all relations $R \subseteq X_1 \times X_2$ and $Q \subseteq X_2 \times X_3$ we have*

$$\overline{T}(R; Q) = \overline{TR}; \overline{TQ}. \quad (7)$$

Proof. First, assume that T preserves weak pullbacks and let $R \subseteq X_1 \times X_2$ and $Q \subseteq X_2 \times X_3$ be two binary relations. The pullback of π_2^R and π_1^Q is given by the following set:

$$pb := \{ \langle (x_1, x_2), (x_3, x_4) \rangle \mid (x_1, x_2) \in R, (x_3, x_4) \in Q \text{ and } x_2 = x_3 \},$$

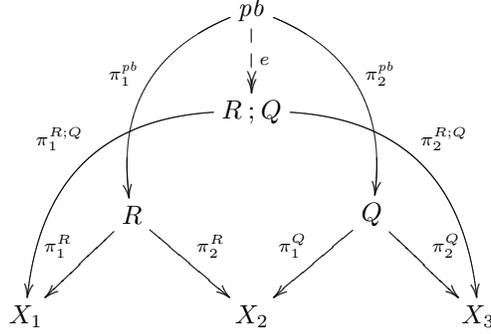


Figure 1: Composition of relations & pullback

and there is a surjective map $e : pb(\pi_2^R, \pi_1^Q) \rightarrow R; Q$ given by $e(\langle (x_1, x_2), (x_3, x_4) \rangle) = (x_1, x_4)$ with the property that

$$\pi_1^{R;Q} \circ e = \pi_1^R \circ \pi_1^{pb} \quad \text{and} \quad \pi_2^{R;Q} \circ e = \pi_2^Q \circ \pi_2^{pb}. \quad (8)$$

The situation is depicted in Figure 1.

We now prove (7). For the inclusion " \subseteq ", let $(x, y) \in \overline{T}(R; Q)$. By definition there exists some $z \in T(R; Q)$ such that $T\pi_1^{R;Q}(z) = x$ and $T\pi_2^{R;Q}(z) = y$. We know that e and thus also Te is surjective. Therefore there exists some $z' \in T(pb)$ such that $Te(z') = z$, and using (8) we obtain $T\pi_1^R(T\pi_1^{pb}(z')) = T\pi_1^{R;Q}(Te(z')) = T\pi_1^{R;Q}(z) = x$ and similarly $T\pi_2^Q(T\pi_2^{pb}(z')) = y$. On the other hand, by the definition of pb , we have $T\pi_2^R(T\pi_1^{pb}(z')) = T\pi_1^Q(T\pi_2^{pb}(z')) = u$. This implies that $(x, u) \in \overline{T}(R)$ and $(u, y) \in \overline{T}(Q)$ and we proved $(x, y) \in \overline{T}(R); \overline{T}(Q)$ as required.

For the converse inclusion suppose that $(x, y) \in \overline{T}(R); \overline{T}(Q)$. We want to prove that this implies $(x, y) \in \overline{T}(R; Q)$. It follows from $(x, y) \in \overline{T}(R); \overline{T}(Q)$ that there is some $u \in TX_2$ such that $(x, u) \in \overline{T}(R)$ and $(u, y) \in \overline{T}(Q)$; spelling out the definitions we find a $u_x \in TR$ and a $u_y \in TQ$ such that $T\pi_1^R(u_x) = x$, $T\pi_2^Q(u_y) = y$ and $T\pi_2^R(u_x) = T\pi_1^Q(u_y) = u$. By our assumption that T is weak pullback preserving we have that $T(pb)$, together with the maps $T\pi_1^{pb}$, $T\pi_2^{pb}$ is the weak pullback of $T\pi_2^R$ and $T\pi_1^Q$. Therefore there must be some $z \in T(pb)$ such that $T\pi_1^{pb}(z) = u_x$ and $T\pi_2^{pb}(z) = u_y$. This implies

$$T\pi_1^{R;Q}(Te(z)) = T\pi_1^R(T\pi_1^{pb}(z)) = T\pi_1^R(u_x) = x$$

and likewise $T\pi_2^{R;Q}(Te(z)) = y$. By definition this means that $(x, y) \in \overline{T}(R; Q)$ as required.

For the converse implication of the statement of the proposition, suppose that T does not preserve weak pullbacks and let the following be a pullback that is not weakly preserved by T :

$$\begin{array}{ccc} P & \xrightarrow{p_2} & X_2 \\ p_1 \downarrow & & g \downarrow \\ X_1 & \xrightarrow{f} & X_3 \end{array}$$

Then it is not difficult to see that the following isomorphic diagram, is also a pullback diagram that is not weakly preserved by T :

$$\begin{array}{ccc} R & \xrightarrow{\pi_2^R} & Gr(g)^\vee \\ \pi_1^R \downarrow & & \downarrow \pi_1^{g^\vee} \\ Gr(f) & \xrightarrow{\pi_2^f} & X_3 \end{array}$$

where $Gr(f)$ and $Gr(g)^\smile$ denote the graph of f and the converse of the graph of g , respectively, and $R \subseteq Grf \times Grg^\smile$ is the pullback of π_2^f and π_1^g . We will show the existence of a pair $(x, y) \in \overline{T}f; \overline{T}g^\smile \setminus \overline{T}(f; g^\smile)$, which is a clear counterexample to (7).

As before there is a surjection $e' : R \rightarrow f; g^\smile$ satisfying

$$\pi_1^{f;g^\smile} \circ e' = \pi_1^f \circ \pi_1^R \text{ and } \pi_2^{f;g^\smile} \circ e' = \pi_2^g \circ \pi_2^R \quad (9)$$

By assumption, $(TR, T\pi_1^R, \pi_2^R)$ is not a weak pullback of $T\pi_2^f$ and $T\pi_1^g$. Hence by Fact 2.7(2), there must be a $z_1 \in TGr(f)$ and a $z_2 \in TGr(g)^\smile$ such that $T\pi_2^f(z_1) = T\pi_1^g(z_2) = u$, while

$$\text{there is no } z \in TR \text{ such that } T\pi_1^R(z) = z_1 \text{ and } T\pi_2^R(z) = z_2. \quad (10)$$

Define $x := T\pi_1^f(z_1)$ and $y := T\pi_2^g(z_2)$. Since $\pi_2^f = f \circ \pi_1^f$, we have $T\pi_2^f = Tf \circ T\pi_1^f$, and so we find $u = (Tf)x$; likewise, we obtain $u = (Tg)y$. From this it is clear that $(x, y) \in \overline{T}f; \overline{T}g^\smile$. Now suppose for a contradiction that $(x, y) \in \overline{T}(f; g^\smile)$. By definition this entails the existence of some $z' \in T(f; g^\smile)$ such that $T\pi_1^{f;g^\smile}(z') = x$ and $T\pi_2^{f;g^\smile}(z') = y$. By surjectivity of e' , and hence, of Te' , then there must be some $z'' \in TR$ such that $Te'(z'') = z'$. Furthermore it follows from (9) that

$$x = T\pi_1^{f;g^\smile}(z') = T\pi_1^{f;g^\smile}(Te'(z'')) = T\pi_1^f(T\pi_1^R(z''))$$

and, similarly, $y = T\pi_2^g(T\pi_2^R(z''))$. Both $T\pi_1^f$ and $T\pi_2^g$ are isomorphisms and thus we obtain $T\pi_1^R(z'') = z_1$ and $T\pi_2^R(z'') = z_2$ - a contradiction to (10) above. QED

Putting this together with Fact 3.5(2,3) we immediately obtain the following.

Corollary 3.7 *Let T be a set functor and let \overline{T} be the operation that maps a set X to $\overline{T}X := TX$ and a relation R to the T -lifting $\overline{T}R$ of R . Then the following are equivalent:*

1. T preserves weak pullbacks;
2. \overline{T} is a functor on the category Rel of sets and relations;
3. \overline{T} is a relator, that is, a monotone functor on the category Rel .

Closely related to this is an important consequence of the functor preserving weak pullbacks, namely that the notions of bisimilarity and behavioral equivalence coincide.

Remark 3.8 In [46] it is proved that if T preserves weak pullbacks then for any pair of coalgebras $\mathbb{X} = \langle X, \xi \rangle$ and $\mathbb{X}' = \langle X', \xi' \rangle$, two states x and x' are behaviorally equivalent iff there is a bisimulation (see Remark 3.4) linking x to x' .

3.3 Standard functors

As mentioned earlier on we will almost exclusively work with Set -functors that are standard. In Remark 3.3 we saw that this will ensure that the definition of the lifting of a relation R is independent of the type of R . Now we will see some further nice consequences of standardness for the notions of relation lifting.

To start with, in case T is standard, \overline{T} commutes with the domain and range of a function; and if T preserves weak pullbacks in addition, then \overline{T} also commutes with restrictions.

Proposition 3.9 *Let T be a standard set functor. Then*

1. \overline{T} commutes with taking domains: $\text{dom}(\overline{T}R) = T(\text{dom}R)$ for all relations $R \subseteq X_1 \times X_2$.

2. \bar{T} commutes with taking range: $\text{rng}(\bar{T}R) = T(\text{rng}R)$ for all relations $R \subseteq X_1 \times X_2$.

3. If T preserves weak pullbacks, then \bar{T} commutes with taking restrictions:

$$\bar{T}(R \upharpoonright_{Y_1 \times Y_2}) = (\bar{T}R) \upharpoonright_{TY_1 \times TY_2}$$

for all sets X_1, X_2, Y_1 and Y_2 , with $Y_1 \subseteq X_1$ and $Y_2 \subseteq X_2$, and for all relations $R \subseteq X_1 \times X_2$.

Proof. For part 1, we first consider the inclusion $\text{dom}(\bar{T}R) = T(\text{dom}R)$. Let $R \subseteq X_1 \times X_2$ be a relation and take an element $\alpha \in \text{dom}(\bar{T}R)$. Then $(\alpha, \beta) \in \bar{T}R$, for some $\beta \in TX_2$. We denote by $\iota : \text{dom}(R) \rightarrow X_1$ the inclusion of $\text{dom}(R)$ into X_1 and by $\pi'_1 : R \rightarrow \text{dom}(R)$ the restriction of the projection map $\pi_1 : R \rightarrow X_1$; then we have $\pi_1 = \iota \circ \pi'_1$. By definition of \bar{T} there exists some $\rho \in TR$ such that $T\pi_1(\rho) = \alpha$ and hence $T\iota(T\pi'_1(\rho)) = \alpha$. As T is standard this shows that $\alpha = T\pi'_1(\rho) \in T\text{dom}(R)$ as required.

For the opposite inclusion, let $f : \text{dom}(R) \rightarrow \text{rng}(R)$ be any map such that $f \subseteq R$; then it follows that $Tf \subseteq \bar{T}R$. In other words, for all $\alpha \in T(\text{dom}R)$ we have $\alpha \bar{T}R Tf(\alpha)$. From this it is immediate that $T(\text{dom}R) \subseteq \text{dom}(\bar{T}R)$.

The proof of part 2 is completely analogous. For part 3, we refer to [35, Prop. 6.4]. QED

Proposition 3.9 is particularly useful for linking the relation lifting of T to that of its finitary version T_ω .

Proposition 3.10 *Let T be a standard and weak pullback preserving set functor, let T_ω be its finitary version and let $R \subseteq X_1 \times X_2$ be a relation. Then*

$$\bar{T}_\omega R = \bar{T}R \cap (T_\omega X_1 \times T_\omega X_2).$$

Proof. Let $R \subseteq X_1 \times X_2$ be a relation and take a pair $(\alpha, \beta) \in T_\omega X_1 \times T_\omega X_2$. By definition of T_ω there must be finite sets $X'_1 \subseteq_\omega X_1$ and $X'_2 \subseteq_\omega X_2$ such that $\alpha \in T_\omega X'_1 = TX'_1$ and $\beta \in T_\omega X'_2 = TX'_2$.

In order to prove the inclusion \supseteq , assume that $(\alpha, \beta) \in \bar{T}R$. By Proposition 3.9 we have

$$(\alpha, \beta) \in \bar{T}R \quad \text{iff} \quad (\alpha, \beta) \in \bar{T}(R \upharpoonright_{X'_1 \times X'_2}) \tag{11}$$

and because $\bar{T}_\omega(R \upharpoonright_{X'_1 \times X'_2}) \subseteq \bar{T}_\omega(R)$ the inclusion holds if we can prove that $(\alpha, \beta) \in \bar{T}_\omega R'$ with $R' := R \upharpoonright_{X'_1 \times X'_2}$. The following diagram commutes:

$$\begin{array}{ccccc} T_\omega X'_1 & \longleftarrow & T_\omega R' & \longrightarrow & T_\omega X'_2 \\ \parallel & & \parallel & & \parallel \\ TX'_1 & \longleftarrow & TR' & \longrightarrow & TX'_2 \end{array}$$

Therefore we have that $(\alpha, \beta) \in \bar{T}R'$ iff $(\alpha, \beta) \in \bar{T}_\omega R'$. By (11) we have $(\alpha, \beta) \in \bar{T}R'$ and hence $(\alpha, \beta) \in \bar{T}_\omega R'$ as required. The proof of the opposite inclusion is similar. QED

On the basis of Proposition 3.10 we will often be sloppy and write $(\alpha, \beta) \in \bar{T}R$ instead of $(\alpha, \beta) \in \bar{T}_\omega R$, for elements $\alpha \in T_\omega X_1$ and $\beta \in T_\omega X_2$. More importantly, Proposition 3.10 allow us to prove our earlier claim, that T_ω inherits the properties of standardness and weak pullback preservation from T .

Proof of Proposition 2.11. Let T be a standard, weak pullback preserving set functor. In order to see that T_ω is standard consider two sets X, X' with $X' \subseteq X$ and let $\iota : X' \rightarrow X$ be the inclusion of X' into X . By the definition of T_ω for every set X we have that $T_\omega X$ is a subset of TX and that the inclusion $\tau_X : T_\omega X \rightarrow TX$ is natural. It follows by naturality that $T_\omega \iota$ is also an inclusion:

$$\begin{array}{ccc} T_\omega X' \xrightarrow{\tau_{X'}} TX' \\ \downarrow T_\omega \iota \quad \downarrow T\iota \\ T_\omega X \xrightarrow{\tau_X} TX \end{array}$$

More precisely, for all $\alpha \in T_\omega X$ we have

$$T_\omega \iota(\alpha) = \tau_X(T_\omega \iota(\alpha)) \stackrel{(\text{nat. of } \tau)}{=} T\iota(\tau'_X(\alpha)) = T\iota(\alpha) \stackrel{T \text{ standard}}{=} \alpha$$

which demonstrates that $T_\omega \iota$ is the inclusion map from $T_\omega X'$ into $T_\omega X$, and shows that T_ω is standard indeed.

We now prove that T_ω preserves weak pullbacks. By Fact 3.6 it suffices to prove that for arbitrary relations $R \subseteq X_1 \times X_2$ and $Q \subseteq X_2 \times X_3$ we have $\overline{T_\omega}(R; Q) = \overline{T_\omega}(R); \overline{T_\omega}(Q)$. In order to see this we use Proposition 3.10. We have

$$\begin{aligned} (\alpha, \beta) \in \overline{T_\omega}(R; Q) & \text{ iff } (\alpha, \beta) \in \overline{T}(R; Q) \upharpoonright_{T_\omega X_1 \times T_\omega X_3} \\ & \text{ iff } (\alpha, \beta) \in \overline{T}(R; Q) \upharpoonright_{TX'_1 \times TX'_3} \text{ for some } X'_1 \subseteq_\omega X_1, X'_3 \subseteq_\omega X_3 \\ & \text{ iff } (\alpha, \beta) \in \overline{T}((R; Q) \upharpoonright_{X'_1 \times X'_3}) \text{ for some } X'_1 \subseteq_\omega X_1, X'_3 \subseteq_\omega X_3 \\ & \text{ iff } (\alpha, \beta) \in \overline{T}(R \upharpoonright_{X'_1 \times X'_2}; Q \upharpoonright_{X'_2 \times X'_3}) \\ & \quad \text{for some } X'_1 \subseteq_\omega X_1, X'_2 \subseteq_\omega X_2, X'_3 \subseteq_\omega X_3 \\ & \text{ iff } (\alpha, \beta) \in \overline{T}(R \upharpoonright_{X'_1 \times X'_2}); \overline{T}(Q \upharpoonright_{X'_2 \times X'_3}) \\ & \quad \text{for some } X'_1 \subseteq_\omega X_1, X'_2 \subseteq_\omega X_2, X'_3 \subseteq_\omega X_3 \\ & \text{ iff } (\alpha, \beta) \in \overline{T_\omega}(R); \overline{T_\omega}(Q) \end{aligned}$$

QED

Finally, we finish this subsection with noting that relation lifting interacts well with the natural transformation $Base : T_\omega \rightarrow P_\omega$.

Proposition 3.11 *Let T be a standard functor that preserves weak pullbacks. Given a relation $R \subseteq X_1 \times X_2$ and elements $\alpha_i \in TX_i$, $i \in \{1, 2\}$, it follows from $\alpha_1 \overline{TR} \alpha_2$ that $Base(\alpha_1) \overline{PR} Base(\alpha_2)$. In particular, we have that $Base(\alpha_1) \subseteq \text{dom}(R)$ and $Base(\alpha_2) \subseteq \text{rng}(R)$.*

Proof. Let π_i^R be the projection of R to X_i , then it follows from $\alpha_1 \overline{TR} \alpha_2$ that $\alpha_i = T\pi_i^R(\rho)$ for some $\rho \in TR$. But then by naturality of $Base$ we find that $Base(\alpha_i) = Base(T\pi_i^R(\rho)) = (P\pi_i^R)(Base(\rho))$, and so $Base(\rho) \in PR$ is a witness to the fact that $Base(\alpha_1) \overline{PR} Base(\alpha_2)$. QED

3.4 Relation Lifting & distributive laws

A relation that plays an important role in our paper is the T -lifting of the membership relation \in . If needed, we will denote the element relation, restricted to a given set X , as the relation $\in_X \subseteq X \times PX$.

Definition 3.12 Given a standard functor T that preserves weak pullbacks, we define, for every set X , a function $\lambda_X^T : TPX \rightarrow PTX$ by putting

$$\lambda_X^T(\Phi) := \{\alpha \in TX \mid \alpha \overline{T} \in_X \Phi\}.$$

Elements of $\lambda_X^T(\Phi)$ will be referred to as *lifted members* of Φ . The family $\lambda^T = \{\lambda_X^T\}_{X \in \text{Set}}$ will be called the T -transformation. \triangleleft

Properties of \overline{T} are intimately related to those of λ^T . In order to express the connection, we need to introduce the concept of a distributive law.

Definition 3.13 Let T be a covariant set functor. A *distributive law* of T over a (co- or contravariant) set functor M is a natural transformation $\theta : TM \rightarrow MT$; that is, the following diagram commutes, for every map $f : X \rightarrow Y$:

$$\begin{array}{ccc} TMX & \xrightarrow{\theta_X} & MTX \\ TMf \downarrow & & \downarrow MTf \\ TMY & \xrightarrow{\theta_Y} & MTY \end{array}$$

(Clearly, in case M is a contravariant functor the downward arrows have to be reversed.) For θ to be *distributive law* of T over a set monad (M, η, μ) , we require in addition that θ is compatible with the monad structure, in the sense that the following diagrams commute, for every set X :

$$\begin{array}{ccc} TX & \xrightarrow{T\eta_X} & TMX \\ \eta_{TX} \searrow & & \downarrow \theta_X \\ & & MTX \end{array} \qquad \begin{array}{ccccc} TMMX & \xrightarrow{\theta_{MX}} & MTMX & \xrightarrow{M\theta_X} & MMTX \\ \downarrow T\mu_X & & & & \downarrow \mu_X \\ TMX & \xrightarrow{\theta_X} & & \xrightarrow{\theta_X} & MTX \end{array} \quad (12)$$

◁

If the functor T preserves weak pullbacks, the T -transformation λ^T provides a distributive laws of T over the power set monad $\mathcal{P} = (P, \{\cdot\}, \cup)$. A detailed proof of this fact can be found in [29, Sec. 4].

Fact 3.14 *If T preserves weak pullbacks, $\lambda^T = \{\lambda_X^T\}_{X \in \text{Set}}$ is a distributive law of T over the power set monad \mathcal{P} .*

What it means, set-theoretically, for λ^T to be a distributive law of T over \mathcal{P} is the following. The fact that λ^T is a natural transformation from TP to PT is another way of saying that for every map $f : X \rightarrow Y$, and every object $\Phi \in TPX$, we obtain the lifted members of $TP\Phi$ by applying the operation Tf to the lifted members of Φ . The diagram on the left of (12), relating the singleton map $\eta_X : X \rightarrow PX$ to the T -transformation, states that an object $\alpha \in TX$ is always the *unique* lifted member of the lifted set $T\eta_X(\alpha)$. To understand the diagram on the right, recall that the multiplication μ of \mathcal{P} is the union map $\cup_X : PPX \rightarrow PX$. Applying the functor to this we obtain a map $T\cup_X : TPPX \rightarrow TPX$. Observe that given an object $\Phi \in TPPX$, we may thus take lifted members of $(T\cup_X)(\Phi)$; however, we may also take lifted members of Φ itself, and since each of these will belong to the set TPX , we may repeat the operation of taking lifted members. Now the right diagram in (12) states that the lifted members of $(T\cup_X)(\Phi)$ coincide with the objects we may obtain as lifted members of lifted members of Φ .

Remark 3.15 The existence of a distributive law of a set functor T over the power set monad \mathcal{P} corresponds to an extension of the functor T to the Kleisli category $\text{Kl}(\mathcal{P})$ of \mathcal{P} . Furthermore it is easy to see that $\text{Kl}(\mathcal{P})$ is isomorphic to the category Rel of sets with relations. Putting these facts together it is clear that any distributive law of a set functor T over \mathcal{P} corresponds to an extension of T to a functor on the category Rel . We saw in Corollary 3.7 that the T -lifting of a relation can be used to extend T to a functor $\bar{T} : \text{Rel} \rightarrow \text{Rel}$ iff T preserves weak pullbacks. In this case λ^T is the corresponding distributive law. Further remarks and references can be found in Section 3.6.

Perhaps somewhat surprisingly, the T -transformation can be also seen as a distributive law over the *contravariant* power set functor.

Proposition 3.16 *Let $T : \text{Set} \rightarrow \text{Set}$ be a functor that preserves weak pullbacks. Then λ^T is a distributive law of T over the contravariant power set functor.*

Proof. Let $f : X \rightarrow Y$ be a function. We have to show that the following diagram commutes:

$$\begin{array}{ccc} T\check{P}Y & \xrightarrow{\lambda_Y^T} & \check{P}TY \\ T\check{P}f \downarrow & & \downarrow \check{P}Tf \\ T\check{P}X & \xrightarrow{\lambda_X^T} & \check{P}TX \end{array}$$

This can be verified by a straightforward calculation:

$$\begin{aligned} \alpha \in \lambda_X^T((T\check{P}f)(\Phi)) & \text{ iff } \Phi(T\check{P}f; \bar{T}\ni_X)\alpha & \text{ iff } \Phi(\bar{T}(\check{P}f; \ni_X))\alpha \\ & \text{ iff } \Phi(\bar{T}(\ni_Y; f^\sim))\alpha & \text{ iff } \Phi(\bar{T}\ni_Y; Tf^\sim)\alpha \\ & \text{ iff } Tf(\alpha) \in \lambda_Y(\Phi) & \text{ iff } \alpha \in (\check{P}Tf)(\lambda_Y(\Phi)) \end{aligned}$$

Here we freely apply properties of relation lifting, and in the third equivalence we use the easily verified fact that $\check{P}f; \ni_X = \ni_Y; f^\sim$. QED

In our paper both distributive laws play an important role. The fact that λ^T is a distributive law over \check{P} is essential for proving that the semantics of Moss' logic is bisimulation invariant, and the distributivity of T over the monad \mathcal{P} is crucial for the soundness of our axiomatization.

To finish this subsection, we gather some elementary facts on the T -transformation.

Proposition 3.17 *Let T be a standard, weak pullback-preserving functor, let X be some set and let $\Phi \in T_\omega PX$.*

1. *If $\emptyset \in \text{Base}(\Phi)$ then $\lambda^T(\Phi) = \emptyset$.*
2. *If $\text{Base}(\Phi) \subseteq \{Y\}$ for some $Y \subseteq X$, then $\lambda^T(\Phi) \subseteq TY$.*
3. *If $\text{Base}(\Phi)$ consists of singletons only, then $|\lambda^T(\Phi)| = 1$.*
4. *If T maps finite sets to finite sets, then for all $\Phi \in T_\omega P_\omega X$, $|\lambda^T(\Phi)| < \omega$.*
5. *If $\Phi \in T_\omega P_\omega X$, then $\lambda^T(\Phi) \in PT_\omega X$.*

Proof. For part 1, assume that $\emptyset \in \text{Base}(\Phi)$ and assume for contradiction that α is a lifted member of Φ . It follows by Proposition 3.11 that $\text{Base}(\alpha) \bar{P} \in \text{Base}(\Phi)$. But from this it would follow, if $\emptyset \in \text{Base}(\Phi)$, that $\text{Base}(\alpha)$ contains a member of \emptyset , which is clearly impossible. Consequently, the set $\lambda^T(\Phi)$ must be empty.

In order to prove part 2, assume that $\Phi \in T\{Y\}$, for some subset Y of X , and suppose that $\alpha \bar{T} \in \Phi$. Then by Proposition 3.9(3) we have $\alpha \bar{T} \in \downarrow_{X \times \{Y\}} \Phi$ and so by part 1 of the same Proposition we find $\alpha \in T\text{dom}(\in \downarrow_{X \times \{Y\}}) = TY$.

For part 3, observe that another way of saying that $\text{Base}(\Phi)$ consists of singletons only, is that $\Phi \in T_\omega S_X$, where $S_X \subseteq PX$ is the collection of singletons from X . Let $\theta_X : S_X \rightarrow X$ be the inverse of η_X , that is, θ_X is the bijection mapping a singleton $\{x\}$ to x . Clearly then, the map $T_\omega \theta_X : T_\omega X \rightarrow T_\omega S_X$ is a bijection as well. In addition, we have $\theta_X^\smile = \in_X$, from which it follows by elementary properties of relation lifting that $(T\theta_X)^\smile = \bar{T}\in_X$. From this it is immediate that if $\Phi \in T_\omega S_X$, then $(T\theta_X)(\Phi)$ is the unique lifted member of Φ .

Concerning part 4, assume that $\Phi \in T_\omega P_\omega X$. Then by definition, $\Phi \in T\mathcal{Y}$ for some $\mathcal{Y} \subseteq_\omega P_\omega X$. From this it follows that $\mathcal{Y} \subseteq PY$ for some finite $Y \subseteq X$, and this implies that $\text{Base}(\Phi) \subseteq PY$. If α is a lifted member of Φ , then by Proposition 3.11 we obtain $\text{Base}(\alpha) \bar{P} \in \text{Base}(\Phi)$, and so in particular we find $\text{Base}(\alpha) \subseteq \bigcup \text{Base}(\Phi) \subseteq Y$. From this it follows that $\lambda^T(\Phi) \subseteq TY$, and so by the assumption on T , the set $\lambda^T(\Phi)$ must be finite.

Finally, we consider part 5. Take an object $\Phi \in T_\omega P_\omega X$ and let $\alpha \in TX$ be an arbitrary lifted member of Φ . Reasoning just as for part 4, we obtain that $\alpha \in TY$ for some finite $Y \subseteq X$, and so by definition of T_ω we find that $\alpha \in T_\omega X$. QED

3.5 Slim redistributions

The syntax of Moss' logic is built using negations, conjunctions, disjunctions and the ∇ -operator. An axiomatisation of the logic has to specify the interaction of these operations. As we will see, so-called *slim redistributions* are the key to understand how conjunction interacts with the ∇ -operator.

Definition 3.18 Let T be a set functor. A set $\Phi \in TPX$ is a *redistribution* of a set $A \in PTX$ if $A \subseteq \lambda_X^T(\Phi)$, that is, every element of A is a lifted member of Φ . In case $A \in P_\omega T_\omega X$, we call a redistribution Φ *slim* if $\Phi \in T_\omega P_\omega(\bigcup_{\alpha \in A} \text{Base}(\alpha))$. The set of slim redistributions of A is denoted as $SRD(A)$. \triangleleft

Intuitively, redistributions of A are ways to reorganize the material of A . The slimness condition $\Phi \in T_\omega P_\omega(\bigcup_{\alpha \in A} \text{Base}(\alpha))$ should be seen as a minimality requirement, ensuring that Φ is 'built from the ingredients of A '.

Example 3.19 First we consider the binary C -labelled tree functor B_C of Example 2.3. Let π_C, π_1 and π_2 denote the respective projections from $B_C X$ to C , X and X , respectively. An object $\Phi \in B_C PX$ is of the form (c, Y, Z) with $c \in C$ and $Y, Z \in PX$. Such a Φ is a redistribution of a set $A = \{(c_i, y_i, z_i) \mid i \in I\} \subseteq_\omega B_C X$ iff for all $i \in I$ we have $c_i = c$, $y_i \in Y$ and $z_i \in Z$, and such a redistribution is slim if in addition, $Y \subseteq \{y_i \mid i \in I\}$ and $Z \subseteq \{z_i \mid i \in I\}$. On this basis it is not hard to derive that

$$SRD(A) = \begin{cases} \{(c, \emptyset, \emptyset) \mid c \in C\} & \text{if } A = \emptyset \\ \emptyset & \text{if } \pi_C[A] \text{ has more than two elements} \\ \{(c_A, \pi_1[A], \pi_2[A])\} & \text{if } c_A \text{ is the unique element of } \pi_C[A] \end{cases}$$

In case we are dealing with the power set functor P , first observe that given a set X , the relation $\overline{P} \in_X \subseteq PX \times PPX$ is given by

$$\alpha \overline{P} \in \Phi \text{ iff } \alpha \subseteq \bigcup \Phi \text{ and } \alpha \cap \beta \neq \emptyset \text{ for all } \beta \in \Phi.$$

On the basis of this observation it is easy to check that $\Phi \in PX$ is a redistribution of $A \in PPX$ if $\bigcup A \subseteq \bigcup \Phi$ and $\alpha \cap \beta \neq \emptyset$ for all $\alpha \in A$ and $\beta \in \Phi$. Furthermore, we obtain

$$\Phi \in SRD(A) \text{ iff } \bigcup A = \bigcup \Phi \text{ and } \alpha \cap \beta \neq \emptyset \text{ for all } \alpha \in A, \beta \in \Phi.$$

Hence, in the case of the power set functor we are dealing with a *symmetric* relation: $\Phi \in SRD(A)$ iff $A \in SRD(\Phi)$.

The following observation, which is due to M. Bílková, shows that slim redistributions naturally occur in the context of distributive lattices.

Example 3.20 Let \mathbb{D} be a distributive lattice. The distributive law for \mathbb{D} can be formulated as follows: for any set $A \in P_\omega P_\omega D$, we have

$$\bigwedge_{\alpha \in A} \bigvee \alpha = \bigvee_{\gamma \in CF(A)} \bigwedge \text{rng}(\gamma),$$

where $CF(A)$ is the set of *choice functions* on A , that is, $CF(A)$ is the set of maps $\gamma : A \rightarrow D$ such that $\gamma(\alpha) \in \alpha$, for all $\alpha \in A$. Then it is straightforward to verify that the set $\{\text{rng}(\gamma) \mid \gamma \in CF(A)\}$ is in fact a slim redistribution of A .

In fact, we may prove that

$$\bigwedge_{\alpha \in A} \bigvee \alpha = \bigvee_{\Phi \in SRD(A)} \bigvee_{\varphi \in \Phi} \bigwedge \varphi. \quad (13)$$

Later on we will see that our axiom governing the interaction of ∇ with conjunctions, generalizes (13).

We finish the section with a proposition for future reference.

Proposition 3.21 $SRD(\emptyset) = T\{\emptyset\}$.

Proof. If Φ is a slim redistribution of the empty set, then by definition $\Phi \in TP_\omega(\emptyset) = T\{\emptyset\}$. Conversely, any $\Phi \in T\{\emptyset\}$ satisfies the condition that $\emptyset \subseteq \mathcal{X}^T(\Phi)$, and so $\Phi \in SRD(\emptyset)$. QED

3.6 Notes

The relation lifting via spans as in Definition 3.1 was defined by Barr in [10, Section 2]. Without stating it explicitly, he also proves that the relation lifting \bar{T} is a functor on \mathbf{Rel} iff T preserves weak pullbacks; see also Trnková [54] and, for a generalisation beyond set functors, Carboni, Kelly and Wood [18, 4.3] and Hermida [24, Theorem 2.3]. [18] also studies the question which functors $\mathbf{Rel} \rightarrow \mathbf{Rel}$ arise from functors $\mathbf{Set} \rightarrow \mathbf{Set}$. Closely related notions of relator, also accounting for simulation as opposed to only bisimulation, are studied by Thijs [53] and in the context of coalgebraic logic by [9, 19, 27]. The connection between coalgebraic logic and relation lifting goes back to the original paper by Moss [41] which introduced ∇ and defined its semantics by using relation liftings, albeit without making this notion explicit. Independently, essentially the same notion of relation lifting was studied in a fibrational setting by Hermida and Jacobs [25]. For a comparison of the notions of bisimulation arising from relation lifting and related definitions see Staton [51].

The relation lifting can also be obtained via a distributive law between a functor and a monad as in Definition 3.13, which is a slight, commonly used variant of the notion of a distributive law between monads [11]. As shown in [11], there is a 1-1 correspondence between distributive laws and liftings of functors to the category of algebras. Similarly, distributive laws $\lambda : TM \rightarrow MT$ between a functor T and a monad M , or monad op-functors $(T, \lambda) : (\mathbf{Set}, M) \rightarrow (\mathbf{Set}, M)$ in the terminology of Street [52], are in 1-1 correspondence with liftings \bar{T} of T to the Kleisli category of M .

We thank Dirk Hofmann, Jiří Velebil and Steve Vickers for pointing out various references and their significance.

4 Boolean algebras and their presentations

4.1 Boolean-type algebras

It will be convenient for us to work with a syntax for Boolean logic and Boolean algebras, in which the finitary meet and join symbols, \bigwedge and \bigvee , respectively, are the *primitive* symbols for the conjunction and disjunction operation, respectively.

Definition 4.1 Given a set X , we let $\mathcal{L}_0(X)$ denote the set of Boolean terms/formulas over X , defined by the following grammar:

$$a ::= x \in X \mid \neg a \mid \bigvee \varphi \mid \bigwedge \varphi,$$

where φ is a finite set of Boolean terms. We abbreviate $\perp := \bigvee \emptyset$ and $\top := \bigwedge \emptyset$, and if no confusion is likely we will write $\mathcal{L}_0 := \mathcal{L}_0(\emptyset)$. ◁

Observe that each $\mathcal{L}_0(X)$ is non-empty, always containing the elements \top and \perp .

The above definition can be brought in coherence with the categorical perspective of section 2, as follows.

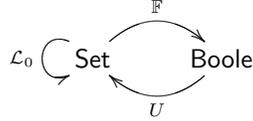
Definition 4.2 We define the category **Boole** of Boolean-type algebras as the algebras for the functor $\mathbf{Set} \rightarrow \mathbf{Set}$, $X \mapsto X + P_\omega X + P_\omega X$. A Boolean-type algebra will usually be introduced as a quadruple $\mathbb{B} = \langle B, \neg^{\mathbb{B}}, \bigwedge^{\mathbb{B}}, \bigvee^{\mathbb{B}} \rangle$, where B is the *carrier* of the algebra, and $\neg^{\mathbb{B}} : B \rightarrow B$, and $\bigwedge^{\mathbb{B}}, \bigvee^{\mathbb{B}} : P_\omega(B) \rightarrow B$ the Boolean *operations*. ◁

Note that this perspective has built in that both conjunction and disjunction are commutative, associative and have a neutral element.

We let $U : \mathbf{Boole} \rightarrow \mathbf{Set}$ denote the forgetful functor, and $\mathbb{F} : \mathbf{Set} \rightarrow \mathbf{Boole}$ its left adjoint; that is, given a set X , $\mathbb{F}X$ denotes the absolutely free Boolean-type algebra, or Boolean term algebra, over X . Note that $\mathbb{F}X$ is not a Boolean algebra. Given a set X , observe that $U\mathbb{F}(X)$ consists of the set $\mathcal{L}_0(X)$ of all Boolean terms/formulas using the elements of X as variables. In fact, we may extend \mathcal{L}_0 to the set functor $\mathcal{L}_0 : \mathbf{Set} \rightarrow \mathbf{Set}$ given by

$$\mathcal{L}_0 := U\mathbb{F}. \quad (14)$$

In this way we obtain the well-known term monad for the Boolean signature with the usual unit $\eta : Id \rightarrow \mathcal{L}_0$ ('variables are terms') and multiplication $\mu : \mathcal{L}_0\mathcal{L}_0 \rightarrow \mathcal{L}_0$ ('terms built from terms are terms').



In particular, for any $f : X \rightarrow \mathcal{L}_0Y$ there is $\hat{f} : \mathcal{L}_0X \rightarrow \mathcal{L}_0Y$ which extends f and can be defined as the composition $\mu_Y \circ \mathcal{L}_0f$. Logicians will recognise \hat{f} as the *substitution* induced by f .

Definition 4.3 Given a set X and a Boolean-type algebra \mathbb{B} , a map $f : X \rightarrow U\mathbb{B}$ is called an *assignment*. Because of the adjunction $\mathbb{F} \dashv U$, such an assignment has a unique extension to a Boole-homomorphism, denoted by

$$\tilde{f} : \mathbb{F}X \rightarrow \mathbb{B}.$$

This map \tilde{f} is the *meaning* function induced by f . ◁

Definition 4.4 A Boole-type algebra \mathbb{B} is a *Boolean algebra* if it satisfies the inequalities of Table 2.

We let $\check{\mathbb{P}} : \mathbf{Set} \rightarrow \mathbf{BA}^{op}$ denote the contravariant power set algebra functor. That is, given a set X , we let $\check{\mathbb{P}}X$ denote the power set algebra of X , and for a map $f : X \rightarrow Y$, the homomorphism $\check{\mathbb{P}}f : \check{\mathbb{P}}Y \rightarrow \check{\mathbb{P}}X$ is provided by the map $f^{-1} = \check{P}f$. ◁

4.2 Presentations of Boolean algebras

It has become a standard tool in mathematics to define an algebraic structure by means of a *presentation* by generators and relations. Usually, these definitions are given in the category-theoretic sense, and in particular do not distinguish isomorphic structures. Our proof-theoretic analysis of the logic requires us to be very precise here, and for this purpose we have developed a small piece of theory on 'concrete presentations'. We want to stress the fact that whereas we only talk about Boolean algebras here, the results in this section in fact apply to a wide universal algebraic setting.

Definition 4.5 A *presentation* is a pair $\langle G; R \rangle$ consisting of a set G of *generators* and a set $R \subseteq \mathcal{L}_0(G) \times \mathcal{L}_0(G)$. Given such a relation R , let $\equiv_R \subseteq \mathcal{L}_0(G) \times \mathcal{L}_0(G)$ be the least congruence relation on the term algebra $\mathbb{F}G$ extending R such that the quotient $\mathbb{F}G/\equiv_R$ is a Boolean algebra. We say that this quotient is the Boolean algebra *presented by* $\langle G; R \rangle$, and denote it as $B\langle G; R \rangle$. Given a presentation $\langle G; R \rangle$, we let

$$\eta_{\langle G; R \rangle} : g \mapsto [g]. \quad (15)$$

define a map $\eta_{\langle G; R \rangle} : G \rightarrow UB\langle G; R \rangle$. ◁

It is straightforward to verify that $\widehat{\eta}_{\langle G; R \rangle}$ is the quotient morphism from $\mathbb{F}G$ to $B\langle G; R \rangle$, with kernel $\ker(\widehat{\eta}_{\langle G; R \rangle}) = \equiv_R$.

Relating this definition of presentations to the more usual one, first observe that a ‘relation’ is nothing but an equation over the set of generators (but note that generators should not be seen as *variables*). Accordingly, given a presentation $\langle G; R \rangle$, a Boolean algebra \mathbb{B} , and an assignment $f : G \rightarrow U\mathbb{B}$, we say that a relation $(s, t) \in R$ is *true* in \mathbb{B} under f , notation: $\mathbb{B}, f \models s \approx t$, if $\widehat{f}(s) = \widehat{f}(t)$. \mathbb{B} is a *model* for R under f if $\mathbb{B}, f \models s \approx t$ for all $(s, t) \in R$. It is straightforward to verify that $B\langle G; R \rangle$ is a model for R under $\eta_{\langle G; R \rangle}$. We can now formulate the following proposition, of which we omit the (straightforward) proof.

Proposition 4.6 *Let $\langle G; R \rangle$ be a presentation, and let \mathbb{B} be a model for R under the assignment $f : G \rightarrow U\mathbb{B}$. Then there is a unique homomorphism $f' : B\langle G; R \rangle \rightarrow \mathbb{B}$ that extends f in the sense that $f'([g]) = f(g)$. In a diagram:*

$$\begin{array}{ccc} G & \xrightarrow{\eta_{\langle G; R \rangle}} & UB\langle G; R \rangle \\ & \searrow f & \downarrow f' \\ & & U\mathbb{B} \end{array}$$

The universal property of $B\langle G; R \rangle$ expressed by the above proposition is usually taken as the definition of the Boolean algebra presented by a presentation.

In order to turn the class of presentations into a category we need to define a notion of morphism between two presentations.

Definition 4.7 A *presentation morphism* from one presentation $\langle G; R \rangle$ to another $\langle G'; R' \rangle$ is a map $f : G \rightarrow \mathcal{L}_0(G')$ satisfying $\widehat{f}(s) \equiv_{R'} \widehat{f}(t)$ for all $s, t \in \mathcal{L}_0(G)$ such that $(s, t) \in R$. Given two presentation morphisms $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ and $g : \langle G'; R' \rangle \rightarrow \langle G''; R'' \rangle$, we define their *composition* $g \circ f : G \rightarrow \mathcal{L}_0(G'')$ as the map given by

$$g \circ f(x) := \widehat{g}(f(x)),$$

and the *identity presentation* on $\langle G; R \rangle$ as the function $id_{\langle G; R \rangle} : G \rightarrow \mathcal{L}_0(G)$ mapping a generator $x \in G$ to the term $x \in \mathcal{L}_0 G$. \triangleleft

The verification that the above defines a category is routine. Category theorists will note that identity and composition are those of the *Kleisli category* associated with the monad \mathcal{L}_0 .

Definition 4.8 We will let Pres denote the category with presentations as objects and presentation morphisms as arrows. \triangleleft

We will now extend the construction B of a Boolean algebra out of a presentation to a functor $B : \text{Pres} \rightarrow \text{BA}$, and define a functor $C : \text{BA} \rightarrow \text{Pres}$ in the opposite direction.

Definition 4.9 Given a presentation morphism $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$, it is easy to see that the map $Bf : \mathbb{F}G / \equiv_R \rightarrow \mathbb{F}G' / \equiv_{R'}$ given by

$$Bf : [s]_{\langle G; R \rangle} \mapsto [\widehat{f}(s)]_{\langle G'; R' \rangle}$$

is well-defined.

Conversely, given a Boolean algebra \mathbb{B} , define its *canonical presentation* as the pair $C\mathbb{B} := \langle U\mathbb{B}; \Delta_{\mathbb{B}} \rangle$. Here $U\mathbb{B}$ is the underlying set of \mathbb{B} , and $\Delta_{\mathbb{B}}$ is the *diagram* of \mathbb{B} , defined as follows:

$$\begin{aligned} \Delta_{\mathbb{B}} \quad := \quad & \{(a, \neg b) \mid a, b \in U\mathbb{B} \text{ with } a = \neg^{\mathbb{B}} b\} \\ & \cup \{(a, \bigwedge \varphi) \mid \{a\} \cup \varphi \subseteq_{\omega} U\mathbb{B} \text{ with } a = \bigwedge^{\mathbb{B}} \varphi\} \\ & \cup \{(a, \bigvee \varphi) \mid \{a\} \cup \varphi \subseteq_{\omega} U\mathbb{B} \text{ with } a = \bigvee^{\mathbb{B}} \varphi\}. \end{aligned}$$

Given a homomorphism $f : \mathbb{B} \rightarrow \mathbb{B}'$ between two Boolean algebras, we let

$$Cf : b \mapsto f(b)$$

define a map $Cf : U\mathbb{B} \rightarrow \mathcal{L}_0(U\mathbb{B}')$. ◁

Proposition 4.10 $B : \text{Pres} \rightarrow \text{BA}$ and $C : \text{BA} \rightarrow \text{Pres}$ are functors.

Further on we will make good use of the following definition.

Definition 4.11 A presentation morphism $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ is a *pre-isomorphism* if there is a morphism $g : \langle G'; R' \rangle \rightarrow \langle G; R \rangle$ such that $\widehat{g}\widehat{f}(s) \equiv_R s$ and $\widehat{f}\widehat{g}(s') \equiv_{R'} s'$, for all terms $s \in \mathcal{L}_0G$ and $s' \in \mathcal{L}_0G'$. This g is called a *pre-inverse* of f . ◁

Proposition 4.12 Let $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ be a presentation morphism. Then f is a pre-isomorphism iff Bf is an isomorphism.

Proof. For the direction from left to right, let f be a pre-isomorphism. We confine ourselves to proving that Bf is injective. For this purpose assume that $Bf([s]_{\langle G; R \rangle}) = Bf([t]_{\langle G; R \rangle})$. Then by definition we have $[\widehat{f}s]_{\langle G'; R' \rangle} = [\widehat{f}t]_{\langle G'; R' \rangle}$, or equivalently, $\widehat{f}s \equiv_{R'} \widehat{f}t$. From this it follows by the assumption that $s \equiv_R \widehat{g}\widehat{f}s \equiv_R \widehat{g}\widehat{f}t \equiv_R t$, and so it is immediate that $[s]_{\langle G; R \rangle} = [t]_{\langle G; R \rangle}$.

Conversely, assume that Bf is an isomorphism between $B\langle G; R \rangle$ and $B\langle G'; R' \rangle$. Let $g : G' \rightarrow \mathcal{L}_0G$ be such that $g(x') \in (Bf)^{-1}[x']$ for every generator $x' \in G'$. We claim that $Bg = (Bf)^{-1}$. To see this, note that it is straightforward to check that $g(s') \in (Bf)^{-1}[s']$; from this it follows that $(Bf)^{-1}([s']_{\langle G'; R' \rangle}) = [\widehat{g}s']_{\langle G; R \rangle}$.

In order to see that g is a pre-inverse of f , consider an arbitrary term $s \in \mathcal{L}_0G$. Clearly we have $[s]_{\langle G; R \rangle} = (Bf)^{-1}(Bf)[s]_{\langle G; R \rangle}$, and so by definition and the above observation, we find $[s]_{\langle G; R \rangle} = (Bf)^{-1}[\widehat{f}s]_{\langle G'; R' \rangle} = [\widehat{g}\widehat{f}s]_{\langle G; R \rangle}$. This means that $s \equiv_R \widehat{g}\widehat{f}s$, as required. Conversely, let $s' \in \mathcal{L}_0G'$. Then we have $[s']_{\langle G'; R' \rangle} = (Bf)(Bf)^{-1}[s']_{\langle G'; R' \rangle} = Bf[\widehat{g}s']_{\langle G; R \rangle} = [\widehat{f}\widehat{g}s']_{\langle G'; R' \rangle}$, or equivalently, $s' \equiv_{R'} \widehat{f}\widehat{g}s'$. QED

The functors B and C are very close to forming an equivalence between the categories Pres and BA . More precisely, we can formulate the following connections. Given a presentation $\langle G; R \rangle$, it is not hard to verify that the insertion of generators $\eta_{\langle G; R \rangle} : G \rightarrow UB\langle G; R \rangle$ defined in (15) is in fact a presentation morphism

$$\eta_{\langle G; R \rangle} : \langle G; R \rangle \rightarrow CB\langle G; R \rangle.$$

Conversely, given a Boolean algebra \mathbb{B} , let id_B denote the identity map on $B := U\mathbb{B}$, and recall that $\widetilde{\text{id}}_B$ denotes the unique homomorphism $\widetilde{\text{id}}_B : \mathbb{F}U\mathbb{B} \rightarrow \mathbb{B}$ extending id_B . It is not difficult to show that $\widetilde{\text{id}}_B(t(b_1, \dots, b_n)) = t^{\mathbb{B}}(b_1, \dots, b_n)$, and so we may think of $\widetilde{\text{id}}$ as an *evaluation map*. We leave it for the reader to verify that for all $s, t \in \mathbb{F}U\mathbb{B}$, we have

$$s \equiv_{C\mathbb{B}} t \text{ iff } \widetilde{\text{id}}_B(s) = \widetilde{\text{id}}_B(t). \quad (16)$$

From this it follows that the map $\epsilon_{\mathbb{B}} : BC\mathbb{B} \rightarrow \mathbb{B}$ given by putting, for any $t(b_1, \dots, b_n) \in \mathcal{L}_0(U\mathbb{B})$:

$$\epsilon_{\mathbb{B}} : [t(b_1, \dots, b_n)] \mapsto t^{\mathbb{B}}(b_1, \dots, b_n) \quad (17)$$

is a well-defined homomorphism from $BC\mathbb{B}$ to \mathbb{B} .

Theorem 4.13 *The functors B and C form an adjoint pair $B \dashv C$, with unit $\eta : Id_{\mathbb{P}res} \dashrightarrow CB$ and counit $\epsilon : BC \dashrightarrow Id_{\mathbb{B}\mathbb{A}}$ given by (15) and (17), respectively. Furthermore, each arrow $\eta_{\langle G; R \rangle} : \langle G; R \rangle \rightarrow CB\langle G; R \rangle$ is a pre-isomorphism, and each arrow $\epsilon_{\mathbb{B}} : BC\mathbb{B} \rightarrow \mathbb{B}$ is an isomorphism.*

Proof. Let us start with showing that $\eta : Id_{\mathbb{P}res} \dashrightarrow CB$ is indeed a natural transformation. That is, given an presentation morphism $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ we have to show that the following diagram commutes.

$$\begin{array}{ccc} \langle G; R \rangle & \xrightarrow{\eta_{\langle G; R \rangle}} & CB\langle G; R \rangle \\ f \downarrow & & CBf \downarrow \\ \langle G'; R' \rangle & \xrightarrow{\eta_{\langle G'; R' \rangle}} & CB\langle G'; R' \rangle \end{array}$$

For this purpose it suffices to check that the two compositions, $CBf \circ \eta_{\langle G; R \rangle}(x)$ and $\eta_{\langle G'; R' \rangle} \circ f(x)$ agree on an arbitrary generator $x \in G$. But this is immediate:

$$CBf \circ \eta_{\langle G; R \rangle}(x) = CBf[x] = [\widehat{f}x] = \widehat{\eta}_{\langle G'; R' \rangle}(\widehat{f}x) = \widehat{\eta}_{\langle G'; R' \rangle}(fx) = (\eta_{\langle G'; R' \rangle} \circ f)(x).$$

In order to prove that $\eta_{\langle G; R \rangle}$ is a pre-isomorphism, let $g : UB\langle G; R \rangle \rightarrow \mathcal{L}_0G$ be any map such that $g([s]) \in [s]$ for any element $[s] \in UB\langle G; R \rangle$. It is easy to check that g is a presentation morphism and that $\eta_{\langle G; R \rangle}$ and g are pre-inverses of each other. From this it is immediate that $\eta_{\langle G; R \rangle}$ is a pre-isomorphism.

Turning to the counit of the adjunction, let $f : \mathbb{B} \rightarrow \mathbb{B}'$ be a homomorphism between Boolean algebras. Let $[t(b_1, \dots, b_n)]$, with each b_i in \mathbb{B} , be an arbitrary element of $BC\mathbb{B}$. Then we compute

$$\begin{aligned} f \circ \epsilon_{\mathbb{B}}[t(b_1, \dots, b_n)] &= f(t^{\mathbb{B}}(b_1, \dots, b_n)) && \text{(definition of } \epsilon) \\ &= t^{\mathbb{B}'}(fb_1, \dots, fb_n) && \text{(} f \text{ is a homomorphism)} \\ &= \epsilon_{\mathbb{B}'}[t(fb_1, \dots, fb_n)] && \text{(definition of } \epsilon) \\ &= \epsilon_{\mathbb{B}'}[\widehat{f}(t(b_1, \dots, b_n))] && \text{(definition of } \widehat{f}) \\ &= \epsilon_{\mathbb{B}'}(BCf)[t(b_1, \dots, b_n)] && \text{(definition of } B \text{ and } C) \end{aligned}$$

This shows that the following diagram commutes:

$$\begin{array}{ccc} BC\mathbb{B} & \xrightarrow{\epsilon_{\mathbb{B}}} & \mathbb{B} \\ BCf \downarrow & & f \downarrow \\ BC\mathbb{B}' & \xrightarrow{\epsilon_{\mathbb{B}'}} & \mathbb{B}' \end{array}$$

and thus proves that ϵ is a natural transformation.

To show that $\epsilon_{\mathbb{B}}$ is an isomorphism, it suffices to check injectivity. But by a straightforward term induction it is easy to prove that every term $t(b_1, \dots, b_k)$ in $\mathcal{L}_0U\mathbb{B}$ satisfies

$$t(b_1, \dots, b_n) \equiv_{CB} t^{\mathbb{B}}(b_1, \dots, b_n).$$

Hence if $\epsilon_{\mathbb{B}}[s(a_1, \dots, a_k)] = \epsilon_{\mathbb{B}}[t(b_1, \dots, b_n)]$, then by $s(a_1, \dots, a_k) \equiv_{CB} s^{\mathbb{B}}(a_1, \dots, a_k) = t^{\mathbb{B}}(b_1, \dots, b_n) \equiv_{CB} t(b_1, \dots, b_n)$, we immediately find that $[s(a_1, \dots, a_k)] = [t(b_1, \dots, b_n)]$, as required.

Finally, in order to prove that $B \dashv C$, by [40, Theorem IV.1.2] it suffices to prove that (i) for any Boolean algebra \mathbb{A} , the composition

$$C\mathbb{A} \xrightarrow{\eta_{C\mathbb{A}}} CBC\mathbb{A} \xrightarrow{C\epsilon_{\mathbb{A}}} C\mathbb{A}$$

is the identity on $C\mathbb{A}$, and that (ii) for any presentation $\langle G; R \rangle$, the composition

$$B\langle G; R \rangle \xrightarrow{B\eta_{\langle G; R \rangle}} BCB\langle G; R \rangle \xrightarrow{\epsilon_{B\langle G; R \rangle}} B\langle G; R \rangle$$

is the identity on $B\langle G; R \rangle$. Both of these facts can be checked by a straightforward unravelling of the definitions, which we will leave as an exercise for the reader. QED

Remark 4.14 What keeps B and C from forming an *equivalence* of categories is that the unit η is a ‘natural pre-isomorphism’ rather than a natural isomorphism. We could remedy this by changing the notion of arrow in the category of presentations but this would be disadvantageous in our completeness proof, when we construct a stratification of our logic.

Remark 4.15 We indicate how the present section generalises beyond Boolean algebras, as suggested by a referee. We have been working with three categories, **BA**, **Boole**, and **Pres**. Instead of **Boole** consider a category \mathcal{B} with forgetful functor $U : \mathcal{B} \rightarrow \mathbf{Set}$ and left-adjoint F of U . Instead of **BA** consider a category \mathcal{A} and a full inclusion $I : \mathcal{A} \rightarrow \mathcal{B}$ with a left-adjoint L of I . Now, we can define a category **Pres**. **Pres** has as objects pairs $\langle G, R \rangle$ where G is a set and R is a relation given by a pair of arrows $R \rightrightarrows UFG$, or equivalently, by $FR \rightrightarrows FG$. A presentation morphism $f : \langle G, R \rangle \rightarrow \langle G', R' \rangle$ is then an algebra morphism $f : FG \rightarrow FG'$ such that for all $A \in \mathcal{A}$ and all $v : FG' \rightarrow IA$, if v equalises $FR' \rightrightarrows FG'$ then $v \circ f$ equalises $FR \rightrightarrows FG$. The functors $B : \mathbf{Pres} \rightarrow \mathcal{A}$ and $C : \mathcal{A} \rightarrow \mathbf{Pres}$ can then be defined as above. Indeed, for $A \in \mathcal{A}$ we let the canonical presentation CA be the kernel pair of the map $UFUIA \rightarrow UIA$, given by the counit of $F \dashv U$ at IA ; and $B\langle G, R \rangle$ is given by the coequaliser of $LFR \rightrightarrows LFG$. As in Theorem 4.13, one can now show that $B \dashv C$ and that the counit $BC \rightarrow Id$ is an iso. Moreover, the proofs do not depend on the base category **Set** and only require rather general assumptions about kernel pairs and coequalisers (which are certainly fulfilled whenever \mathcal{A} and \mathcal{B} are varieties, that is, classes of algebras given by operations of finite arity and equations).

5 Moss’ coalgebraic logic

In this section we will recall the definitions of Moss’ coalgebraic logic and its semantics [41], or rather, the finitary version thereof developed by Venema [55].

5.1 Syntax

As mentioned in the introduction, the key idea underlying the syntax of Moss’ language for reasoning about T -coalgebras is to include a modal operator ∇ into the language whose ‘arity’ is given by the functor T itself, in the same way that P_ω is the ‘arity’ of our conjunction and disjunctions. In the *finitary* version of the language, the arity of ∇ is given by the finitary version T_ω of T . In brief, the language \mathcal{L} will be defined by the following grammar:

$$a ::= \neg a \mid \bigwedge \varphi \mid \bigvee \varphi \mid \nabla \alpha$$

where $\varphi \in P_\omega \mathcal{L}$ and $\alpha \in T_\omega \mathcal{L}$. For the purpose of this paper we need some further syntactic definitions.

Definition 5.1 Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving set functor and let T_ω be the finitary version of T . The language \mathcal{L} of the finitary Moss language for T is defined inductively. We first define \mathcal{L}_0 as the set $\mathcal{L}_0(\emptyset)$ of closed Boolean formulas (see Definition 4.1). For the inductive step, we start with introducing the set functor T_ω^∇ defined by, for a given set X and function $f : X \rightarrow Y$,

$$\begin{aligned} T_\omega^\nabla X &:= \{\nabla \alpha \mid \alpha \in T_\omega X\}, \\ T_\omega^\nabla f(\nabla \alpha) &:= \nabla T f(\alpha). \end{aligned}$$

We continue the inductive definition by putting

$$\mathcal{L}_{i+1} := \mathcal{L}_0 T_\omega^\nabla \mathcal{L}_i.$$

Finally, we define \mathcal{L} as the union $\mathcal{L} := \bigcup_{i \in \omega} \mathcal{L}_i$, and fix the *rank* or *depth* of a formula $a \in \mathcal{L}$ is the smallest natural number n such that $a \in \mathcal{L}_n$. \triangleleft

Using BNF notation, we can recast the above definition as

$$\mathcal{L}_0 \ni a ::= \neg a \mid \bigwedge \varphi \mid \bigvee \varphi$$

where $\varphi \subseteq_\omega \mathcal{L}_0$, and

$$\mathcal{L}_{i+1} \ni a ::= \nabla \alpha \mid \neg a \mid \bigwedge \varphi \mid \bigvee \varphi$$

where $\alpha \in T_\omega \mathcal{L}_i$ and $\varphi \in P_\omega \mathcal{L}_{i+1}$.

Despite its unconventional appearance, the language \mathcal{L} admits fairly standard definitions of most syntactical notions. As an example we mention the notion of a subformula.

Definition 5.2 We define the set $Sfor(a)$ of *subformulas* of a by the following induction:

$$\begin{aligned} Sfor(\neg a) &:= \{\neg a\} \cup Sfor(a) \\ Sfor(\bigwedge \varphi) &:= \{\bigwedge \varphi\} \cup \bigcup_{a \in \varphi} Sfor(a) \\ Sfor(\bigvee \varphi) &:= \{\bigvee \varphi\} \cup \bigcup_{a \in \varphi} Sfor(a) \\ Sfor(\nabla \alpha) &:= \{\nabla \alpha\} \cup \bigcup_{a \in Base(\alpha)} Sfor(a). \end{aligned}$$

The elements of $Base(\alpha) \subseteq Sfor(\nabla \alpha)$ will be called the *immediate* subformulas of $\nabla \alpha$. \triangleleft

On the basis of this definition it is not difficult to prove that every formula in \mathcal{L} has only *finitely* many subformulas. This is in fact the reason why we call our language the *finitary* version of Moss'.

Remark 5.3 In order to formulate and understand the interaction principles between nabla and the Boolean operations, we need to think of the propositional connectives as *functions* on formulas. Taking disjunction as an example, observe that we may think of it as a map $\bigvee : P_\omega \mathcal{L} \rightarrow \mathcal{L}$. Thus we may apply the functor T_ω to this map, obtaining $T\bigvee : T_\omega P_\omega \mathcal{L} \rightarrow T_\omega \mathcal{L}$. (Recall from our discussion on the finitary version of a functor that to simplify notation we will write $T\bigvee$ rather than $T_\omega \bigvee$.) Hence, for $\Phi \in T_\omega P_\omega \mathcal{L}$, we find $(T\bigvee)\Phi \in T_\omega \mathcal{L}$, which means that $\nabla(T\bigvee)\Phi$ is a well-formed formula. The same applies to the formula $\nabla(T\bigwedge)\Phi$, and similarly, we may think of negation as a map $\neg : \mathcal{L} \rightarrow \mathcal{L}$, and obtain $T\neg : T\mathcal{L} \rightarrow T\mathcal{L}$; thus for any formula $\nabla \alpha$, we may also consider the formula $\nabla(T\neg)\alpha$.

Remark 5.4 The reader may be surprised that we did not include propositional variables in our language. The reason for this is that we may *encode* these into the functor. More precisely, given a functor T and a set \mathbf{Prop} of proposition letters, recall from Example 2.2(5) that the T -models over \mathbf{Prop} can be identified with the coalgebras for the functor $T_{\mathbf{Prop}} = P(\mathbf{Prop}) \times T$. Hence we may use the language \mathcal{L} associated with $T_{\mathbf{Prop}}$ to describe the \mathbf{Prop} -models based on T -coalgebras, see Example 5.10(3).

Convention 5.5 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 some kind of *naming convention*, see Table 1 below.

Set	Elements
\mathcal{L}	a, b, \dots
$T_\omega \mathcal{L}$	α, β, \dots
$P_\omega \mathcal{L}$	φ, ψ, \dots
$P_\omega T_\omega \mathcal{L}$	A, B, \dots
$T_\omega P_\omega \mathcal{L}$	Φ, Ψ, \dots

Table 1: Naming convention

It will be useful later on to have a more categorical description of the finitary Moss language for a functor T . For this purpose we need the following definition.

Definition 5.6 We define the category Boole_∇ of *Moss algebras* as the algebras for the Moss functor $A_M : \text{Set} \rightarrow \text{Set}$, given as:

$$A_M := Id + P_\omega + P_\omega + T_\omega,$$

That is, for a set S , $A_M S$ is the disjoint union of S , two (disjoint copies) of $P_\omega S$, and $T_\omega S$; for a map f , $A_M f$ is defined accordingly.

A Moss algebra will usually be introduced as a quadruple $\mathbb{B} = \langle B, \neg^\mathbb{B}, \wedge^\mathbb{B}, \vee^\mathbb{B}, \nabla^\mathbb{B} \rangle$, where $\langle B, \neg^\mathbb{B}, \wedge^\mathbb{B}, \vee^\mathbb{B} \rangle$ is a Boole-type algebra, called the *Boolean reduct* of \mathbb{B} , and $\nabla^\mathbb{B} : T_\omega B \rightarrow B$ is the *nabla operator* of \mathbb{B} . \triangleleft

Given a Moss algebra \mathbb{B} , there is a unique, natural way to interpret \mathcal{L} -terms as elements of the carrier B of \mathbb{B} . This *meaning function* $mng_\mathbb{B} : \mathcal{L} \rightarrow U\mathbb{B}$ can be defined by a straightforward induction on the complexity of formulas. For instance, the clauses for \wedge and ∇ are

$$\begin{aligned} mng_\mathbb{B}(\wedge \varphi) &:= \wedge^\mathbb{B}(Pmng_\mathbb{B})(\varphi) \\ mng_\mathbb{B}(\nabla \alpha) &:= \nabla^\mathbb{B}(Tmng_\mathbb{B})(\alpha) \end{aligned}$$

Categorically speaking, this means the following. We may view Moss' language itself as a Moss algebra, by interpreting the function symbols as the corresponding syntactic operation, as usual in universal algebra. Note that in order to prove that $\nabla^\mathcal{L} \alpha$ belongs to \mathcal{L} , it is crucial that ∇ is a finitary operation: from $\alpha \in T_\omega \mathcal{L}$ it follows that $\alpha \in T_\omega \mathcal{L}_n$ for some finite n , and then we may proceed with $\nabla \alpha \in \mathcal{L}_{n+1} \subseteq \mathcal{L}$. The arising algebra, that we will also denote as \mathcal{L} , is a rather special Moss algebra, namely, the *initial* one. Apart from the fact that the syntax of \mathcal{L} is slightly unusual, the proof of the proposition below is standard universal algebra, and so we omit it.

Proposition 5.7 \mathcal{L} is the initial Moss algebra: given an arbitrary Moss algebra \mathbb{B} , the meaning function $mng_\mathbb{B}$ is the unique homomorphism from \mathcal{L} to \mathbb{B} .

Before moving on to the coalgebraic semantics of \mathcal{L} , we finish our discussion of its syntax with the following definition, for future reference.

Definition 5.8 Let $T : \text{Set} \rightarrow \text{Set}$ be a set functor and let T_ω be the finitary version of T . We define the functor $\mathcal{L}_1 : \text{Set} \rightarrow \text{Boole}$ by putting

$$\mathcal{L}_1 := \mathcal{L}_0 \circ T_\omega^\nabla \circ \mathcal{L}_0.$$

\triangleleft

On occasion, we will consider \mathcal{L}_1 also as a Boole valued functor allowing us to write $\mathcal{L}_1 = \mathbb{F}T_\omega^\nabla \mathcal{L}_0$. The notation \mathcal{L}_1 is in accordance with the definition of \mathcal{L}_1 as the fragment of rank one formulas in \mathcal{L} , by the observation that $\mathcal{L}_1 = (\mathcal{L}_0 \circ T_\omega^\nabla)(\mathcal{L}_0) = \mathcal{L}_0 T_\omega^\nabla \mathcal{L}_0(\emptyset)$.

5.2 Semantics

Given all the preparations we have made in the previous sections, the definition of the semantics of the language is completely straightforward.

Definition 5.9 Let $T : \text{Set} \rightarrow \text{Set}$ be a standard, weak pullback preserving functor, and let $\mathbb{X} = \langle X, \xi \rangle$ be a T -coalgebra. The satisfaction relation $\Vdash_{\mathbb{X}} \subseteq X \times \mathcal{L}$ is defined by the following induction on the complexity of formulas:

$$\begin{aligned} x \Vdash_{\mathbb{X}} \neg a & \quad \text{if } x \not\Vdash_{\mathbb{X}} a, \\ x \Vdash_{\mathbb{X}} \bigwedge \varphi & \quad \text{if } x \Vdash_{\mathbb{X}} a \text{ for all } a \in \varphi, \\ x \Vdash_{\mathbb{X}} \bigvee \varphi & \quad \text{if } x \Vdash_{\mathbb{X}} a \text{ for some } a \in \varphi, \\ x \Vdash_{\mathbb{X}} \nabla \alpha & \quad \text{if } \xi(x) \bar{T} \Vdash_{\mathbb{X}} \alpha. \end{aligned}$$

If $x \Vdash_{\mathbb{X}} a$ we say that a is *true*, or *holds* at x in \mathbb{X} . We may omit the superscript when no confusion is likely, writing \Vdash instead of $\Vdash_{\mathbb{X}}$.

In case a holds throughout \mathbb{X} , that is, at every state of \mathbb{X} , we write $\mathbb{X} \Vdash a$. \triangleleft

Before we turn to look at some examples, we should argue for the *well-definedness* of the relation \Vdash . In particular, when looking at the clause for the nabla modality, the reader might be worried whether this is an inductive definition at all, since the defining clause, ' $\xi(x) \bar{T} \Vdash \alpha$ ', refers to the *full* forcing relation. The point is that because of our assumptions, \bar{T} commutes with restrictions, and so we have

$$(\xi(x), \alpha) \in \bar{T}(\Vdash) \iff (\xi(x), \alpha) \in \bar{T}(\Vdash|_{X \times \text{Base}(\alpha)}). \quad (18)$$

Thus, in order to determine whether $\nabla \alpha$ holds at x or not, we only have to know the interpretation of the *immediate subformulas* of α (that is, the elements of $\text{Base}(\alpha)$). In other words, if using the right hand side of (18) rather than the left hand side, we would have an equivalent, inductive, definition of the semantics.

- Example 5.10**
1. Let T be the C -stream functor given by $TX = C \times X$ for some set C . Then ∇_T takes as its argument a pair (c, a) where $c \in C$ and a is a formula in \mathcal{L} . The formula $\nabla(c, a)$ is true in a T -coalgebra (X, ξ) at a state x if $\xi(x) = (c', y)$ with $c = c'$ and $y \Vdash a$.
 2. The nabla operator ∇_P associated with the power set functor P is the *cover modality* discussed in the introduction.
 3. If T_{Prop} is the T -model functor of Example 2.3(5), associated with a functor T and a set Prop of proposition letters, then $\nabla_{T_{\text{Prop}}}$ takes as its argument a pair (π, α) consisting of a set $\pi \subseteq \text{Prop}$ and a set $\alpha \subseteq_{\omega} \mathcal{L}^T$. The meaning of the formula $\nabla_K(\pi, \alpha)$ can be expressed as

$$\nabla_{T_{\text{Prop}}}(\pi, \alpha) \equiv \left(\bigwedge_{p \in \pi} p \wedge \bigwedge_{p \notin \pi} \neg p \right) \wedge \nabla_T \alpha.$$

4. Finally, let $T = D_{\omega}$ be the finitary distribution functor. In this case, $\nabla_{D_{\omega}}$ takes as argument a distribution $\mu : \mathcal{L} \rightarrow [0, 1]$ of finite support. Given a T -coalgebra $\mathbb{X} = (X, \xi)$ and some $x \in X$ we have $x \Vdash_{\mathbb{X}} \nabla_{D_{\omega}} \mu$ if for all $y \in X$ and all $a \in \mathcal{L}$ there are real numbers $\rho_{y,a} \in [0, 1]$ such that

$$\begin{aligned} \rho_{y,a} \neq 0 & \quad \text{implies} \quad y \Vdash a, \xi(x)(y) \neq 0, \mu(a) \neq 0 & \quad \text{and} \\ \sum_{a' \in \mathcal{L}} \rho_{y,a'} & = \xi(x)(y) \quad \text{for all } y \in X & \quad \text{and} \\ \sum_{y \in X} \rho_{y,a} & = \mu(a) \quad \text{for all } a \in \mathcal{L}. \end{aligned}$$

The state-based semantics of the logics as presented in Definition 5.9 can be brought in accordance with the earlier algebraic perspective by the observation that every T -coalgebra naturally induces a Moss algebra, namely its *complex algebra*.

Definition 5.11 Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor, and let $\mathbb{X} = \langle X, \xi \rangle$ be a T -coalgebra. The *complex algebra* \mathbb{X}^+ of \mathbb{X} is defined as the Moss algebra \mathbb{B} which has the power set algebra $\mathbb{P}(X)$ as its Boolean reduct, while

$$\nabla^{\mathbb{X}^+} := \check{P}\xi \circ \lambda_{\mathbb{X}}^T$$

defines the nabla operation of \mathbb{X}^+ . ◁

In words: the Boolean function symbols \neg , \bigvee and \bigwedge are interpreted as the complementation, union and intersection operations on the power set of X . To understand the definition of the nabla operation, observe that applying the contravariant power set functor to the coalgebra map ξ , we obtain a function $\check{P}\xi : \check{P}TX \rightarrow \check{P}X$, so if we compose this map with the T -transformation $\lambda_{\mathbb{X}}^T : T\check{P}X \rightarrow \check{P}TX$, we obtain a map $\check{P}\xi \circ \lambda_{\mathbb{X}}^T : T\check{P}X \rightarrow \check{P}X$ of the right shape.

It follows by Proposition 5.7 that every \mathcal{L} -formula a can uniquely be assigned a meaning $mng_{\mathbb{X}^+}(a) \in \mathbb{P}X$ in the complex algebra of a T -coalgebra \mathbb{X} — in the sequel we will write $mng_{\mathbb{X}}$ rather than $mng_{\mathbb{X}^+}$. The Proposition below states that the two approaches to the coalgebraic semantics of \mathcal{L} coincide, so that we can speak without hesitation of ‘the’ meaning of a formula in a T -coalgebra.

Proposition 5.12 Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor, and let $\mathbb{X} = \langle X, \xi \rangle$ be a T -coalgebra. Then we have

$$mng_{\mathbb{X}}(a) = \{x \in X \mid x \Vdash a\},$$

for every formula $a \in \mathcal{L}$.

Proof. The proof of this proposition proceeds by a routine formula induction. QED

5.3 First observations

In this subsection we gather first observations on \mathcal{L} . First we show that Moss’ logic is *adequate*; that is, it cannot distinguish behaviorally equivalent states.

Theorem 5.13 (Adequacy) Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor, and let $f : X \rightarrow Z$ be a coalgebra morphism between the T -coalgebras (X, ξ) and (Z, ζ) . For all formulas $a \in \mathcal{L}$ and all states $x \in X$ we have

$$x \Vdash_{\mathbb{X}} a \text{ iff } f(x) \Vdash_{\mathbb{Z}} a. \tag{19}$$

We leave it as an exercise for the reader to give a *direct* proof of Theorem 5.13 — a straightforward induction will suffice, using the fact that \bar{T} distributes over relation composition in the case of a formula $a = \nabla\alpha$. We will give a proof based on the algebraic approach, involving the initiality of \mathcal{L} (Proposition 5.7), and the following result.

Proposition 5.14 Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor, and let $f : X \rightarrow Z$ be a coalgebra morphism between the T -coalgebras $\mathbb{X} = (X, \xi)$ and $\mathbb{Z} = (Z, \zeta)$. Then $\check{P}f$ is an algebraic homomorphism from \mathbb{Z}^+ to \mathbb{X}^+ .

Proof. It is well-known that $\check{P}f$ is a homomorphism from the power set algebra $\check{\mathbb{P}}(Z)$ to $\check{\mathbb{P}}(X)$. Thus it is left to show that $\check{P}f$ also is a homomorphism with respect to the nabla operators. For that purpose, consider the following diagram:

$$\begin{array}{ccccc}
T\check{P}Z & \xrightarrow{\lambda_Z^T} & \check{P}TZ & \xrightarrow{\check{P}\zeta} & \check{P}Z \\
\downarrow T\check{P}f & & \downarrow \check{P}Tf & & \downarrow \check{P}f \\
T\check{P}X & \xrightarrow{\lambda_X^T} & \check{P}TX & \xrightarrow{\check{P}\zeta} & \check{P}X
\end{array}$$

The left rectangle commutes since λ^T is a distributive law of T over \check{P} (see Proposition 3.16), and the right rectangle commutes by functoriality of \check{P} and the assumption that f is a coalgebra morphism. As a corollary, the outer diagram commutes, but by definition of $\nabla^{\mathbb{X}^+}$ and $\nabla^{\mathbb{Z}^+}$ this just means that $\check{P}f$ is a homomorphism for ∇ . QED

On the basis of the previous proposition, the proof of the Theorem is almost immediate.

Proof of Theorem 5.13. By initiality of \mathcal{L} as a Moss algebra, $mng_{\mathbb{X}}$ is the unique homomorphism $mng_{\mathbb{X}} : \mathcal{L} \rightarrow \mathbb{X}^+$. But it follows from Proposition 5.14 that $\check{P}f \circ mng_{\mathbb{Z}}$ is also a homomorphism from \mathcal{L} to \mathbb{X}^+ , so that we may conclude that

$$mng_{\mathbb{X}} = \check{P}f \circ mng_{\mathbb{Z}}. \quad (20)$$

Now let x and a be as in the statement of the theorem, then we have

$$\begin{aligned}
x \Vdash_{\mathbb{X}} a & \text{ iff } x \in mng_{\mathbb{X}}(a) && \text{(Proposition 5.12)} \\
& \text{ iff } x \in \check{P}f(mng_{\mathbb{Z}}(a)) && (20) \\
& \text{ iff } fx \in mng_{\mathbb{Z}}(a) && \text{(definition of } \check{P}f) \\
& \text{ iff } fx \Vdash_{\mathbb{Z}} a && \text{(Proposition 5.12)}
\end{aligned}$$

From this the Theorem is immediate. QED

5.4 Logic

The purpose of this paper is to provide a sound and complete axiomatization of the set of *coalgebraically valid* formulas in this language, that is, the set of \mathcal{L} -formulas that are true in every state of every coalgebra. Since our completeness proof will be algebraic in nature, for our purposes it will be convenient to formulate our results in terms of *equations*, or rather, *inequalities*.

Definition 5.15 An *inequality* is an expression of the form $a \preceq b$, where a and b are formulas in \mathcal{L} . Similarly, an *equation* is an expression of the form $a \approx b$. ◁

One may think of the inequality $a \preceq b$ as abbreviating the equation $a \wedge b \approx a$, and we will see the equation $a \approx b$ as representing the set $\{a \preceq b, b \preceq a\}$ of inequations. (In fact, in our Boolean setting, we could even represent the equation $a \approx b$ by the single inequality $(a \wedge \neg b) \vee (\neg a \wedge b) \preceq \perp$.) Thus it does not really matter whether we base our logic on equations or on inequalities, and in the sequel we will move from one perspective to the other if we deem it useful.

Definition 5.16 An inequality $a \preceq b$ holds in a Moss algebra \mathbb{A} , notation: $\mathbb{A} \models a \preceq b$, if $mng_{\mathbb{A}}(a) \leq_{\mathbb{A}} mng_{\mathbb{A}}(b)$. ◁

Given the Boolean basis of our logics, we can express coalgebraic validity in terms of equational validity, and vice versa. More precisely, given a T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$, it is easy to see that

$$\mathbb{X} \Vdash a \iff \mathbb{X}^+ \models \top \preceq a$$

and, conversely,

$$\mathbb{X}^+ \models a \preceq b \iff \mathbb{X} \Vdash \neg a \vee b.$$

As a consequence, in order to axiomatize the coalgebraically valid formulas, we may just as well find a derivation system for the inequalities that are valid in all complex algebras.

Definition 5.17 An inequality $a \preceq b$ is (T -coalgebraically) *valid*, notation: $a \models_T b$, if it holds in every complex algebra \mathbb{X}^+ . \triangleleft

As an example of a validity, we mention the following, for an arbitrary $\Phi \in T_\omega P_\omega \mathcal{L}$:

$$\nabla(T \vee) \Phi \preceq \bigvee \left\{ \nabla \beta \mid \beta \bar{T} \in \Phi \right\} \quad (\nabla 3_f)$$

(see Remark 5.3 for an explanation of the syntax). Note that the right hand side of $(\nabla 3_f)$ is a well-defined formula only if the disjunction is finite; we can guarantee this by requiring T to map finite sets to finite sets. (We will come back to this issue in the next section.)

Proposition 5.18 *If T is a weak pullback preserving, standard set functor that maps finite sets to finite sets, then the formula $(\nabla 3_f)$ is valid for every $\Phi \in T_\omega P_\omega \mathcal{L}$.*

Proof. In order to understand the validity of $(\nabla 3_f)$, fix some T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$.

First observe that for any $\varphi \subseteq_\omega \mathcal{L}$ we have $\mathbb{X}, x \Vdash \bigvee \varphi$ iff $\mathbb{X}, x \Vdash a$, for some $a \in \varphi$. Putting it differently, the relations $\Vdash; \in$ and $\Vdash; \bigvee^\sim$ coincide. From this it follows that

$$\bar{T}(\Vdash; \in) = \bar{T}(\Vdash; \bigvee^\sim). \quad (21)$$

Now fix some object $\Phi \in T_\omega P_\omega \mathcal{L}$, and suppose that x is a state in \mathbb{X} such that $x \Vdash \nabla(T \vee) \Phi$. From this it follows that the pair $(\xi(x), (T \vee)(\Phi))$ belongs to the relation $\bar{T} \Vdash$, and so $(\xi(x), \Phi)$ belongs to $(\bar{T} \Vdash); (T \vee)^\sim = \bar{T}(\Vdash; \bigvee^\sim)$. But then by (21), we find $(\xi(x), \Phi) \in \bar{T}(\Vdash; \in) = \bar{T} \Vdash; \bar{T} \in$. In other words, there is some object β such that $\xi(x) \bar{T} \Vdash \beta$ and $\beta \bar{T} \in \Phi$. Clearly then $x \Vdash \nabla \beta$, and so we have $x \Vdash \bigvee \{ \nabla \beta \mid \beta \bar{T} \in \Phi \}$, as required. QED

6 The derivation system

6.1 Introduction

In this section we introduce our derivation system \mathbf{M} for the finitary version of Moss' logic, as given in the previous section. First we fix some general notation and terminology concerning derivations.

Definition 6.1 Given a derivation system \mathbf{D} , we let each of $\vdash_{\mathbf{D}} a \preceq b$, $a \sqsubseteq_{\mathbf{D}} b$ and $b \sqsupseteq_{\mathbf{D}} a$ denote the fact that the inequality $a \preceq b$ is derivable in \mathbf{D} , and we write $a \equiv_{\mathbf{D}} b$ if both $a \sqsubseteq_{\mathbf{D}} b$ and $b \sqsubseteq_{\mathbf{D}} a$. \triangleleft

In other words, where $a \preceq b$ and $a \approx b$ are syntactic expressions in an object language, the expressions $a \sqsubseteq_{\mathbf{D}} b$ and $a \equiv_{\mathbf{D}} b$ denote statements, in the metalanguage, *about* the derivability of such expressions $a \preceq b$ and $b \preceq a$. In case no confusion is likely concerning the derivation system at hand, we will drop subscripts, simply writing $a \equiv b$ and $a \sqsubseteq b$.

$\frac{}{a \preccurlyeq a}$	$\frac{a \preccurlyeq b \quad b \preccurlyeq c}{a \preccurlyeq c}$
$\frac{\{a \preccurlyeq b \mid a \in \varphi\}}{\bigvee \varphi \preccurlyeq b}$	$\frac{a \preccurlyeq b}{a \preccurlyeq \bigvee \psi} \quad b \in \psi$
$\frac{\{a \preccurlyeq b \mid b \in \psi\}}{a \preccurlyeq \bigwedge \psi}$	$\frac{a \preccurlyeq b}{\bigwedge \varphi \preccurlyeq b} \quad a \in \varphi$
$\frac{}{\bigwedge \{\bigvee \varphi \mid \varphi \in X\} \preccurlyeq \bigvee \{\bigwedge \gamma[X] \mid \gamma \in \text{Choice}(X)\}}$	
$\frac{\bigwedge (X \cup \{-a\}) \preccurlyeq \bigvee Y}{\bigwedge X \preccurlyeq \bigvee (Y \cup \{a\})}$	$\frac{\bigwedge (X \cup \{a\}) \preccurlyeq \bigvee Y}{\bigwedge X \preccurlyeq \bigvee (Y \cup \{-a\})}$

Table 2: Axioms and rules for classical propositional logic

In principle, the derivation system that we are looking for, should have axioms and rules of three kinds. First of all, it will have a propositional core taking care of the Boolean basis of our setting. For this purpose, any sound and complete set of axioms and derivation rules would do; for concreteness, we propose the set given in Table 2. Recall that our language has \bigvee and \bigwedge as primitive connectives.

Second, our system will need some kind of *congruence rule* for the nabla modality. Since ∇ has a rather unusual form, perhaps it is not a priori clear what such a rule would look like. The naive way to formulate a congruence rule for ∇ would be as

$$\text{from } \alpha \bar{T} \equiv \beta \text{ infer } \nabla \alpha \equiv \nabla \beta \quad (22)$$

Problem is that the premiss of (22) is not itself an equation, or a set of equations. This problem can be remedied by invoking some properties of relation lifting. More precisely, note that from Proposition 3.9 we may derive the equivalence $\alpha \bar{T} \equiv \beta \iff \alpha \bar{T}Z \beta$, for some $Z \subseteq \text{Base}(\alpha) \times \text{Base}(\beta)$. This would lead to the following formulation of a congruence rule:

$$\frac{\{a \approx b \mid (a, b) \in Z\}}{\nabla \alpha \approx \nabla \beta} \quad (\alpha, \beta) \in \bar{T}Z$$

The above rule is supposed to have a set of premisses: $\{a \approx b \mid (a, b) \in Z\}$, where $Z \subseteq \text{Base}(\alpha) \times \text{Base}(\beta)$ is a relation such that $(\alpha, \beta) \in \bar{T}Z$ — the latter condition is formulated as a *side condition* of the rule.

As it turns out, however, we also want ∇ to be *order-preserving*, and the most straightforward way to formulate that would be by strengthening (22) to

$$\text{from } \alpha \bar{T} \sqsubseteq \beta \text{ infer } \nabla \alpha \sqsubseteq \nabla \beta. \quad (23)$$

If we want to turn this into a syntactically well-formed derivation rule again, we obtain our first derivation rule ($\nabla 1$):

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

which can be read as a congruence and monotonicity rule in one. It has the additional advantage of being formulated in terms of our primitive symbol, \preccurlyeq .

Example 6.2 First, consider the C -labelled binary tree functor $B_C = C \times Id \times Id$ of Example 2.3. Here, an application of rule $(\nabla 1)$ looks as follows:

$$\frac{\{a_1 \preceq b_1, a_2 \preceq b_2\}}{\nabla(c, a_1, a_2) \preceq \nabla(c, b_1, b_2)}$$

where c is an arbitrary element of C . Note that no inequality of the form $\nabla(c, a_1, a_2) \preceq \nabla(d, b_1, b_2)$ with $c \neq d$ can be derived using $(\nabla 1)$ because $(\nabla(c, a_1, a_2), \nabla(d, b_1, b_2)) \notin \overline{T}(Z)$ for any relation Z .

In the case of the power set functor P , an application of the rule $(\nabla 1)$ looks as follows:

$$\frac{\{a \preceq b \mid (a, b) \in Z\}}{\nabla\alpha \preceq \nabla\beta} (\alpha, \beta) \in \overline{P}Z$$

where $\alpha, \beta \in P_\omega \mathcal{L}$ are finite sets of formulas. It can be easily seen that the premiss of the rule can be satisfied iff for all $a \in \alpha$ there is a $b \in \beta$ such that $a \preceq b$, and vice versa.

In addition, any complete derivation system for Moss' language will need some *interaction principles* describing the interaction between the nabla modality and the Boolean connectives. As we will see, the interaction principles between ∇ and the Boolean connectives \vee and \wedge will take the form of two *distributive laws* (in the logical meaning of the word). We postpone discussing the role of negation in our system until subsection 6.5, and before giving the general formulation of the laws for \wedge and \vee , we first discuss a simple, special, case.

6.2 Functors restricting to finite sets

For a gentle introduction of our derivation system we first consider the special case where the functor restricts to finite sets.

Turning to the interaction principles, we first consider the interaction between the coalgebraic modality and *conjunctions*. More specifically, the purpose of axiom $(\nabla 2)$ will be to rewrite a conjunction of nabla formulas as an equivalent 'disjunction of nablas of conjunctions', and we think of this axiom as a distributive law (in the logical sense). Formally, recall from Definition 3.18 that given a finite set $A \in P_\omega T_\omega \mathcal{L}$, the set $SRD(A) \subseteq T_\omega P_\omega \mathcal{L}$ denotes the set of *slim redistributions* of A . Also recall that given an object $\Phi \in T_\omega P_\omega \mathcal{L}$, we find $(T \wedge)\Phi \in T_\omega \mathcal{L}$, which means that $\nabla(T \wedge)\Phi$ is a well-formed formula. We can now formulate the axiom $(\nabla 2)$ as the following inequality:

$$\bigwedge \{ \nabla\alpha \mid \alpha \in A \} \preceq \bigvee \{ \nabla(T \wedge)\Phi \mid \Phi \in SRD(A) \} \quad (\nabla 2_f)$$

Example 6.3 First consider the case of the C -labelled binary tree functor B_C of Example 2.3. In Example 3.19 we discussed the shape of the collection of slim redistributions of a collection $A \subseteq_\omega T_\omega \mathcal{L}$. From this it should be clear that we obtain the following three instances of $(\nabla 2_f)$.

1. If $A = \emptyset$, we obtain

$$\top \preceq \bigvee \{ \nabla(c, \top, \top) \mid c \in C \}$$

2. If A contains two elements (c, a_1, a_2) and (c', a'_1, a'_2) with $c \neq c'$, then we obtain

$$\bigwedge \{ \nabla\alpha \mid \alpha \in A \} \preceq \perp.$$

3. If $\pi_C[A]$ contains a unique element c_A , then we obtain

$$\bigwedge \{ \nabla\alpha \mid \alpha \in A \} \preceq \nabla(c_A, \bigwedge \pi_1[A], \bigwedge \pi_2[A]),$$

where π_C , π_1 and π_2 are the projection functions, as in Example 3.19.

Second, in the case of the power set functor, $T = P$, an instance of $(\nabla 2_f)$ looks as follows

$$\bigwedge_{\alpha \in A} \nabla \alpha \preceq \bigvee \left\{ \nabla \{ \bigwedge \beta \mid \beta \in \Phi \} \mid \bigcup A = \bigcup \Phi \text{ and } \alpha \cap \beta \neq \emptyset \text{ for all } \alpha \in A, \beta \in \Phi \right\} \quad (24)$$

Remark 6.4 In fact, we could have formulated this principle as an *equation* rather than as an inequality, since the opposite inequality of $(\nabla 2_f)$ can be derived on the basis of $(\nabla 1)$. To see this, observe that for any formula $a \in \mathcal{L}$ and any set $\varphi \in P_\omega \mathcal{L}$ it holds that $a \in \varphi$ implies that $a \sqsupseteq \bigwedge \varphi$. Reformulating this as $(\in; \bigwedge) \subseteq \sqsupseteq$, and using the properties of relation lifting we find that $\overline{T} \in; T \bigwedge \subseteq \overline{T} \sqsupseteq$. From this it follows that, whenever $\alpha \in T_\omega \mathcal{L}$ is a lifted member of $\Phi \in T_\omega P_\omega \mathcal{L}$, we find that $(T \bigwedge) \Phi \overline{T} (\sqsupseteq) \alpha$. From this, one application of $(\nabla 1)$ yields the existence of a derivation for the inequality $\nabla (T \bigwedge) \Phi \preceq \nabla \alpha$. Since this holds for any α and Φ with $\alpha \overline{T} \in \Phi$, we may conclude that

$$\bigvee \left\{ \nabla (T \bigwedge) \Phi \mid \Phi \in SRD(A) \right\} \sqsubseteq \bigwedge \left\{ \nabla \alpha \mid \alpha \in A \right\}.$$

That is, the opposite inequality of $(\nabla 2_f)$ is indeed derivable.

Our second interaction principle, $(\nabla 3)$, involves the interaction between ∇ and the *disjunction* operation. And again, we think of this axiom as a distributive law (in the logical sense), stating that the coalgebraic modality distributes over disjunctions. More precisely, the rule reads as follows:

$$\nabla (T \vee) \Phi \preceq \bigvee \left\{ \nabla \beta \mid \beta \overline{T} (\in) \Phi \right\} \quad (\nabla 3_f)$$

Example 6.5 In the case of the functor $B_C = C \times Id \times Id$, axiom $(\nabla 3_f)$ is of the following shape:

$$\nabla (c, \vee A, \vee B) \preceq \bigvee \left\{ \nabla (c, a, b) \mid a \in A, b \in B \right\}.$$

For the power set functor P , an instance of axiom $(\nabla 3_f)$ looks as follows

$$\nabla \{ \vee \beta \mid \beta \in \Phi \} \preceq \bigvee \left\{ \nabla \alpha \mid \alpha \subseteq \bigcup \Phi \text{ and } \alpha \cap \beta \neq \emptyset \text{ for all } \beta \in \Phi \right\}.$$

Remark 6.6 In this case the opposite inequality can be derived on the basis of $(\nabla 1)$ as well. Here we use the fact that $a \in \varphi$ implies $a \sqsubseteq \bigvee \varphi$, or in other words, that $\in; \bigvee \subseteq \sqsubseteq$. This implies that $\overline{T} \in; T \vee \subseteq \overline{T} \sqsubseteq$, and hence, whenever β is a lifted member of Φ , we find that $\beta \overline{T} \sqsubseteq (T \vee) \Phi$. Thus an application of $(\nabla 1)$ shows the derivability of the inequality $\nabla \beta \preceq \nabla (T \vee) \Phi$. And since this applies to every lifted member of Φ , we may conclude that

$$\bigvee \left\{ \nabla \beta \mid \beta \overline{T} (\in) \Phi \right\} \sqsubseteq \nabla (T \vee) \Phi,$$

meaning that, indeed, the opposite inequality of $(\nabla 3_f)$ is derivable.

Summarizing, in the case of a set functor T that preserves finite sets, our derivation system \mathbf{M}_f extends that of classical proposition logic (Table 2) with one congruence/monotonicity rule, and two axioms that take the form of distributive laws, see Table 3. The point of restricting to this case is to ensure that the axioms $(\nabla 2_f)$ and $(\nabla 3_f)$ are well-formed pieces of syntax, in the sense that the disjunctions on the right hand side are *finite*.

Remark 6.7 The requirement on the given set functor T to preserve finite sets is obviously sufficient in order to ensure that the axioms $(\nabla 2_f)$ and $(\nabla 3_f)$ are well-formed. Note, however, that there are set functors that do not restrict to finite sets and for which the axioms $(\nabla 2_f)$ and $(\nabla 3_f)$ are nevertheless syntactically well-formed.

$(\nabla 1) \quad \frac{\{a \preceq b \mid (a, b) \in Z\}}{\nabla \alpha \preceq \nabla \beta} \quad (\alpha, \beta) \in \bar{T}Z$
$(\nabla 2_f) \quad \bigwedge \{ \nabla \alpha \mid \alpha \in A \} \preceq \bigvee \{ \nabla (T \wedge) \Phi \mid \Phi \in SRD(A) \}$
$(\nabla 3_f) \quad \nabla (T \vee) \Phi \preceq \bigvee \{ \nabla \beta \mid \beta \bar{T}(\in) \Phi \}$

Table 3: Rules and axioms of the system \mathbf{M} (in case T preserves finite sets)

Consider for example the bag functor B_ω from Example 2.2. In order to show that $(\nabla 2_f)$ and $(\nabla 3_f)$ are well-formed we have to prove that the sets

$$\{\Phi \in B_\omega P_\omega X \mid \Phi \in SRD(A)\} \quad \text{for } A \in P_\omega B_\omega X \text{ and} \quad (25)$$

$$\{\beta \in B_\omega X \mid \beta(\bar{B}_\omega \in) \Phi\} \quad \text{for } \Phi \in B_\omega P_\omega X \quad (26)$$

are finite. Using the characterisation of the relation lifting for B_ω in Example 3.2 this is not difficult to see: Let us consider first the set in (25), ie., we consider some $A \in P_\omega B_\omega X$ and we want to prove that the set $\{\Phi \in B_\omega P_\omega X \mid \Phi \in SRD(A)\}$ is finite. If $\Phi \in SRD(A)$ then by the definition of slim redistributions we have $(\alpha, \Phi) \in (\bar{B}_\omega \in)$ for all $\alpha \in A$ and $\Phi \in B_\omega P_\omega (\bigcup_{\alpha' \in A} Base(\alpha'))$. Therefore, using Proposition 3.9, we get that

$$(\alpha, \Phi) \in \bar{B}_\omega \left(\in \uparrow_{Base(\alpha) \times P_\omega (\bigcup_{\alpha' \in A} Base(\alpha'))} \right) \quad \text{for all } \alpha \in A.$$

This implies, by the definition of \bar{B}_ω from Example 3.2, that there exists a function

$$\rho : \in \uparrow_{Base(\alpha) \times P_\omega (\bigcup_{\alpha' \in A} Base(\alpha'))} \rightarrow \mathbb{N}$$

such that for all $\alpha \in A$, all $x \in Base(\alpha)$ and all $U \in P_\omega (\bigcup_{\alpha' \in A} Base(\alpha'))$ we have

$$\Phi(U) = \sum_{x' \in Base(\alpha), x' \in U} \rho(x', U) \quad \text{and} \quad \rho(x, U) \leq \alpha(x)$$

Therefore we have $\Phi(U) \leq \sum_{x \in U} \alpha(x)$. This shows that the range of Φ has an upper bound and thus, as Φ is determined by its values on the finite set $P_\omega (\bigcup_{\alpha' \in A} Base(\alpha'))$, there can only be finitely many Φ 's that satisfy the requirement of a slim redistribution for the set A . In a similar way one can show that the set $\{\beta \in B_\omega X \mid \beta(\bar{B}_\omega \in) \Phi\}$ in (26) is finite for all $\Phi \in B_\omega P_\omega X$. We leave the details of the argument as an exercise to the reader.

One example for a set functor for which the finitary axioms $(\nabla 2_f)$ and $(\nabla 3_f)$ are not well-formed is provided by the finitary probability functor D_ω in Example 2.2.

6.3 The derivation system \mathbf{M}

In the case that we are dealing with an arbitrary set functor T (not necessarily preserving finite sets), we would like to use the same derivation system as given in Table 3. Unfortunately however, in this case the axioms $(\nabla 2_f)$ and $(\nabla 3_f)$ are no longer well-formed syntactic expressions, since we cannot

$(\nabla 1) \quad \frac{\{a \preceq b \mid (a, b) \in Z\}}{\nabla \alpha \preceq \nabla \beta} \quad (\alpha, \beta) \in \bar{T}Z$
$(\nabla 2) \quad \frac{\{\nabla(T \wedge)(\Phi) \preceq b \mid \Phi \in SRD(A)\}}{\wedge \{\nabla \alpha \mid \alpha \in A\} \preceq b}$
$(\nabla 3) \quad \frac{\{\nabla \alpha \preceq b \mid \alpha \bar{T} \in \Phi\}}{\nabla(T \vee)(\Phi) \preceq b}$

Table 4: Rules of the system \mathbf{M}

guarantee that the disjunctions on the right hand sides are taken over a *finite* set. In order to deal with this problem, we use the following trick: we replace an axiom of the form

$$a \preceq \bigvee \{a_i \mid i \in I\}$$

with the derivation rule

$$\frac{\{a_i \preceq b \mid i \in I\}}{a \preceq b}$$

The price that we have to pay for this transformation is that our derivation system will be *infinitary*.

Definition 6.8 The derivation system \mathbf{M} is given by the axioms and derivation rules of Table 4, together with the complete set of axioms and rules for classical propositional logic given in Table 2. \triangleleft

Our notions of derivation and derivability are completely standard.

Definition 6.9 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 $\frac{\mathcal{D}}{a \preceq b}$ or $\mathcal{D} : a \sqsubseteq b$. If we want to suppress the actual derivation, we write $\vdash_{\mathbf{M}} a \preceq b$ or (in accordance with Definition 6.1) $a \sqsubseteq_{\mathbf{M}} b$. \triangleleft

Note that \mathbf{M} is not a Gentzen-style derivation system; in particular, we do not have left- and right introduction- and elimination rules for ∇ . Readers who are interested to see a detailed development of the *proof theory* of nabla-style coalgebraic logic, are referred to Bílková, Palmigiano & Venema [13] (for the power set case).

6.4 Soundness and completeness

We can now very concisely formulate the main result of this paper as the following soundness and completeness result:

Theorem 6.10 *Let T be a standard set functor that preserves weak pullbacks. For all formulas $a, b \in \mathcal{L}$ we have*

$$\vdash_{\mathbf{M}} a \preceq b \quad \text{iff} \quad a \models_T b. \quad (27)$$

In words, Theorem 6.10 states that for any two \mathcal{L} -formulas a and b , the inequality $a \preceq b$ is derivable in our derivation system \mathbf{M} iff it is valid in all T -coalgebras. Our proof of this result will be based on many auxiliary results, which we will discuss in the next two sections. The final proof will be given at the end of section 8.

6.5 The role of negation

At this point, the reader may be surprised or even worried that we have formulated our derivation system for a Boolean-based coalgebraic modal logic, without mentioning the negation connective (or the implication, for that matter) in relation to the nabla modality at all. Surely there must be some validities involving both ∇ and \neg ? The point is that indeed there are such interaction principles, but we do not need to formulate them explicitly as axioms or derivation rules since they are already *derivable* in the system \mathbf{M} . The intuition underlying this fact is that in a bounded distributive lattice, all existing complementations are completely determined by the lattice operations: the complement $\neg a$ of an element a , if existing, is the unique element b such that $a \wedge b = \perp$ and $a \vee b = \top$.

Nevertheless, the key principle relating ∇ to \neg will be needed in our proofs below, and so we discuss it in some detail. For a smooth formulation we need the following definition.

Definition 6.11 Given an element $\alpha \in T_\omega \mathcal{L}$, let $Q(\alpha) \subseteq T_\omega \mathcal{L}$ be the set defined by

$$Q(\alpha) := \left\{ T(\bigwedge \circ P\neg)\Psi \mid \Psi \in T_\omega P_\omega \text{Base}(\alpha) \text{ and } (\alpha, \Psi) \notin \bar{T} \right\}.$$

◁

To unravel this definition, observe that $P\neg : P_\omega \mathcal{L} \rightarrow P_\omega \mathcal{L}$, and so we have $\bigwedge \circ P\neg : P_\omega \mathcal{L} \rightarrow \mathcal{L}$. Thus we find that for $\Psi \in T_\omega P_\omega \text{Base}(\alpha) \subseteq T_\omega P_\omega \mathcal{L}$ we have $(T(\bigwedge \circ P\neg))\Psi \in T_\omega \mathcal{L}$ indeed.

In case T preserves finite sets, $Q(\alpha)$ is a finite set, and we can express the principle relating ∇ and \neg as follows:

$$\neg \nabla \alpha \approx \bigvee \left\{ \nabla \beta \mid \beta \in Q(\alpha) \right\}. \quad (\nabla 4_f)$$

In other words: the negation of a nabla is equivalent to a disjunction of nablas of conjunctions of negations of the base formulas. Putting it yet differently, in the case of T preserving finite sets, we can define the *Boolean dual* Δ of ∇ , just in terms of ∇ and \bigvee . For more information on this dual modality Δ the reader is referred to Kissig & Venema [31].

In the general case, that is, if the functor T does not necessarily take finite sets to finite sets, we can express the interaction between ∇ and \neg in the form of a derivation rule,

$$\frac{\{ \nabla \beta \preceq b \mid \beta \in Q(\alpha) \}}{\neg \nabla \alpha \preceq b} \quad (\nabla 4_L)$$

and a collection of axioms:

$$\{ \nabla \beta \preceq \neg \nabla \alpha \mid \beta \in Q(\alpha) \}, \quad (\nabla 4_R)$$

corresponding to the directions \preceq and \succeq of $(\nabla 4_f)$, respectively. The point to make is that *both* $(\nabla 4_L)$ and $(\nabla 4_R)$ are *derivable* in \mathbf{M} . We will prove this in detail for $(\nabla 4_L)$. Given our completeness result, the derivability of $(\nabla 4_R)$ is an immediate consequence of its validity [31]. The actual derivation of $\nabla \beta \preceq \neg \nabla \alpha$ for $\beta \in Q(\alpha)$ is rather involved, so we refrain from giving the details here.

In any case, the key instruments in the derivability of both $(\nabla 4_L)$ and $(\nabla 4_R)$ are the following two rules.

Proposition 6.12 *For any finite set φ of formulas, the following rules are \mathbf{M} -derivable:*

$$(\nabla 4a) \quad \frac{\top \preceq \bigvee \varphi \quad \left\{ \nabla \alpha \preceq b \mid \alpha \in T\varphi \right\}}{\top \preceq b}$$

$$(\nabla 4b) \quad \frac{\left\{ a \wedge a' \preceq \perp \mid a \neq a' \in \varphi \right\}}{\nabla \alpha \wedge \nabla \alpha' \preceq \perp} \quad \alpha \neq \alpha' \in T\varphi$$

Proof. In the proof below, the following principle will be used a few times:

$$\text{Given } f : S \rightarrow S', \text{ for } s \in S, Tf \text{ restricts to a bijection } Tf : T\{s\} \rightarrow T\{f(s)\} \quad (28)$$

We first show the derivability of $(\nabla 4a)$. Assume that we have a derivation \mathcal{D}_\top of $\top \preceq \bigvee \varphi$, and a derivation \mathcal{D}_α of $\nabla \alpha \preceq b$, for each $\alpha \in T\varphi$.

Consider an arbitrary element $\Phi \in T\{\varphi\}$. By Proposition 3.17(2), each lifted member α of Φ belongs to $T\varphi$. If we apply $(\nabla 3)$ to the set $\{\mathcal{D}_\alpha \mid \alpha \bar{T} \in \Phi\}$, we obtain a derivation

$$\mathcal{D}_\Phi : \frac{\{\mathcal{D}_\alpha : \nabla \alpha \preceq b \mid \alpha \bar{T} \in \Phi\}}{\nabla(T\bigvee)(\Phi) \preceq b}$$

for each $\Phi \in T(\{\varphi\})$.

Applying our principle (28) to the map $\bigvee : P_\omega \mathcal{L} \rightarrow \mathcal{L}$, we find that each $\beta \in T(\{\bigvee \varphi\})$ is of the form $\beta = (T\bigvee)(\Phi_\beta)$ for some $\Phi_\beta \in T(\{\varphi\})$. Thus in fact for each such β we have a derivation

$$\mathcal{D}_\beta : \nabla \beta \preceq b$$

On the other hand, we may continue the derivation \mathcal{D}_\top as follows. Consider the bijection $f : \{\top\} \rightarrow \{\bigvee \varphi\}$, which induces a bijection $Tf : T\{\top\} \rightarrow T\{\bigvee \varphi\}$. Clearly we find that $f \subseteq \{\sqsubseteq\}$, so that $Tf \subseteq \overline{T\sqsubseteq}$. From this it follows that we may apply the rule $(\nabla 1)$ to the inequality $\top \preceq \bigvee \varphi$ and obtain, for each $\gamma \in T\{\top\}$, the derivation

$$\nabla 1 \frac{\frac{\mathcal{D}_\top}{\top \preceq \bigvee \varphi}}{\nabla \gamma \preceq \nabla(Tf)\gamma}$$

Combining the observations until now, we obtain the following derivation \mathcal{D}_γ for each $\gamma \in T\{\top\}$:

$$\mathcal{D}_\gamma : \text{cut} \frac{\nabla 1 \frac{\frac{\mathcal{D}_\top}{\top \preceq \bigvee \varphi}}{\nabla \gamma \preceq \nabla(Tf)\gamma} \quad \frac{\mathcal{D}_{(Tf)\gamma}}{\nabla(Tf)\gamma \preceq b}}{\nabla \gamma \preceq b}$$

Since $(T\wedge)(\Psi) \in T\{\top\}$ for each $\Psi \in T\{\emptyset\}$, this means that above we have obtained a derivation

$$\mathcal{D}_\Psi : \nabla(T\wedge)(\Psi) \preceq b$$

for each $\Psi \in T\{\emptyset\}$.

Finally, consider the instantiation of $(\nabla 2)$ with $A = \emptyset$. By Proposition 3.21 we have $SRD(\emptyset) = T\{\emptyset\}$, so that the set $\{\nabla(T\wedge)(\Psi) \preceq b \mid \Psi \in T\{\emptyset\}\}$ is exactly the set of premises of this instantiation of $(\nabla 2)$. Hence we may simply take the set of all derivations \mathcal{D}_Ψ , with $\Psi \in T\{\emptyset\}$, and continue as follows:

$$\nabla 2 \frac{\{\mathcal{D}_\Psi \mid \Psi \in T\{\emptyset\}\}}{\top \preceq b}$$

This finishes the proof of the derivability of $(\nabla 4a)$.

In the case of $(\nabla 4b)$ we will proceed a bit faster, leaving the details as to why our argumentation yields derivability rather than admissibility, as an exercise for the reader. Let φ be a finite set of formulas such that $a \wedge a' \equiv \perp$ for all distinct $a, a' \in \varphi$, and let α and α' be two distinct elements of $T\varphi$. We will derive the inequality $\nabla \alpha \wedge \nabla \alpha' \preceq \perp$. By $(\nabla 2)$ it suffices to show that

$$\vdash_{\mathbf{M}} \nabla(T\wedge)(\Phi) \preceq \perp,$$

where Φ is an arbitrary slim redistribution of the set $\{\alpha, \alpha'\}$.

But if $\Phi \in SRD(\{\alpha, \alpha'\})$, and both α and α' belong to $T\varphi$, then first of all we have $Base(\Phi) \subseteq P\varphi$, because $\Phi \in T_\omega P_\omega(Base(\alpha) \cup Base(\alpha'))$ by the definition of a slim redistribution and thus $Base(\Phi) \subseteq P(Base(\alpha) \cup Base(\alpha')) \subseteq P\varphi$. In addition, it follows by Proposition 3.17(1) that $\emptyset \notin Base(\Phi)$, and then by Proposition 3.17(3) that $Base(\Phi)$ contains some set $\psi \subseteq \varphi$ with $|\psi| > 1$. Define the following function $d : Base(\Phi) \rightarrow P(\varphi) \cup \{\{\top\}\}$:

$$d(\chi) := \begin{cases} \emptyset & \text{if } |\chi| > 1 \\ \chi & \text{if } |\chi| = 1 \\ \{\top\} & \text{if } |\chi| = 0 \end{cases}$$

On the basis of our set of premisses $\{a \wedge a' \preceq \perp \mid a \neq a' \in T\varphi\}$, for each $\chi \in Base(\Phi) \subseteq P\varphi$ we can find a derivation for the inequality $\bigwedge \chi \preceq \bigvee d(\chi)$. Putting these derivations together, and applying $(\nabla 1)$ with $Z = \{\{\bigwedge \chi, \bigvee d(\chi) \mid \chi \in Base(\Phi)\}\}$, we obtain a derivation \mathcal{D}_Φ for the inequality $\nabla(T \bigwedge)(\Phi) \preceq \nabla(T \bigvee)(Td(\Phi))$.

We also claim that we can derive the inequality $\nabla(T \bigvee)(Td(\Phi)) \preceq \perp$. Since $Base : T_\omega \rightarrow P_\omega$ is a natural transformation, we have that $Base(Td(\Phi)) = (Pd)(Base(\Phi)) = d[Base(\Phi)]$. Now recall that above we found a $\psi \in Base(\Phi)$ with $|\psi| > 1$; it follows that $\emptyset = d(\psi) \in Base(Td(\Phi))$, so that on the basis of Proposition 3.17(1) we may conclude that $Td(\Phi)$ has *no* lifted members. But then one single application of $(\nabla 3)$, with the *empty* set of premisses, provides the desired derivation for $\nabla(T \bigvee)(Td(\Phi)) \preceq \perp$.

Finally then, an application of the cut rule gives $\nabla(T \bigwedge)(\Phi) \preceq \perp$, as required. QED

As a corollary to this we can now prove the derivability of $(\nabla 4_L)$.

Proposition 6.13 *The rule $(\nabla 4_L)$ is derivable in \mathbf{M} .*

Proof. Let $\alpha \in T_\omega \mathcal{L}$ and $b \in \mathcal{L}$ be arbitrary, and assume that for all $\beta \in Q(\alpha)$ we have $\nabla \beta \sqsubseteq b$. We will show that $\neg \nabla \alpha \sqsubseteq b$.

Consider the map $t : P_\omega Base(\alpha) \rightarrow \mathcal{L}$ given by

$$t : \psi \mapsto \bigwedge \{a \in Base(\alpha) \mid a \notin \psi\} \wedge \bigwedge \{-b \mid b \in \psi\}.$$

Then for all $\psi \subseteq Base(\alpha)$ it is straightforward to verify that (i) $t(\psi) \sqsubseteq (\bigwedge \circ P \neg)\psi$, and (ii) if $a \notin \psi$ then $t(\psi) \sqsubseteq a$.

Define φ to be the *range* of t . Intuitively, think of φ as the set of atoms of a Boolean algebra; then it is not hard to see that

$$\top \sqsubseteq \bigvee \varphi. \tag{29}$$

We claim that

$$\text{for all } \gamma \in T\varphi : \nabla \gamma \sqsubseteq b \vee \nabla \alpha. \tag{30}$$

For the proof of (30), take an arbitrary $\gamma \in T\varphi$. By definition of φ , the map Tt is surjective when seen as $Tt : T_\omega P_\omega Base(\alpha) \rightarrow T_\omega \varphi$, and so we may fix an element $\Psi \in T_\omega P_\omega Base(\alpha)$ such that $\gamma = (Tt)\Psi$. Now distinguish cases.

First assume that $(\alpha, \Psi) \notin \overline{T} \not\sqsubseteq$. It follows from (i) that $\gamma = (Tt)\Psi \overline{T} \sqsubseteq (T(\bigwedge \circ P \neg))\Psi$, and so an application of $(\nabla 1)$ shows that $\nabla \gamma \sqsubseteq \nabla(T(\bigwedge \circ P \neg))\Psi$. Now by assumption we have $(T(\bigwedge \circ P \neg))\Psi \in Q(\alpha)$, and so there is a derivation of the inequality $\nabla(T(\bigwedge \circ P \neg))\Psi \preceq b$. Then an application of the cut rule shows that $\nabla \gamma \sqsubseteq b$.

If, on the other hand, the pair (α, Ψ) *does* belong to the relation $\overline{T} \not\sqsubseteq$, then by (ii) we obtain that $\gamma = (Tt)\Psi \overline{T} \sqsubseteq \alpha$. Now an application of $(\nabla 1)$ yields a derivation for $\nabla \gamma \preceq \nabla \alpha$.

In either case, a simple propositional continuation of the derivation shows that $\nabla\gamma \sqsubseteq b \vee \nabla\alpha$, which proves (30).

Finally, applying the derived rule ($\nabla 4a$) to the premisses given by (29) and (30), we obtain a derivation of the inequality $\top \preceq b \vee \nabla\alpha$. But from this it follows by some straightforward classical propositional manipulations that $\neg\nabla\alpha \sqsubseteq b$, as required. QED

7 One-step soundness and completeness

As mentioned in the introduction, our completeness proof is based on Pattinson's stratification method [43], which consists of stratifying the logic in ω many layers which are nicely glued together by means of a so-called one-step version of the derivation system. The main technical hurdle in this method consists of showing that this one-step derivation system is sound and complete with respect to a natural one-step semantics. In this section we will first properly introduce our version of these notions, and then prove the one-step soundness and completeness result.

7.1 One-step semantics and one-step axiomatics

Starting with the one-step semantics, fix a set X and think of $\check{P}X$ as a set of formal objects or *propositions*. Recall from Section 5 that $\mathcal{L}_0\check{P}X$ and $\mathcal{L}_1\check{P}X$ are the sets of formulas of depth zero and depth one over this language, respectively. The point underlying the one-step semantics is that there is a natural interpretation of the formulas in $\mathcal{L}_1\check{P}X$ as sets of elements of TX , or, expressed more accurately, as elements of the Boolean algebra $\mathbb{P}TX$. To explain this, first note that we may see the identity map

$$\iota : \check{P}X \rightarrow \check{P}X$$

as a natural *valuation* interpreting variables of $\check{P}X$ as subsets of X , and then extend this valuation to a unique homomorphism

$$\llbracket \cdot \rrbracket_0^X := \tilde{\iota} : \mathbb{F}\check{P}X \rightarrow \check{P}X.$$

We find it convenient to denote $U\tilde{\iota} : \mathcal{L}_0\check{P}X \rightarrow \check{P}X$ by the same symbol $\llbracket \cdot \rrbracket_0^X$ and also to occasionally drop the superscript X . We may associate a relation $\Vdash_X^0 \subseteq X \times \mathcal{L}_0\check{P}X$ with this map, which we define inductively by putting

$$\begin{aligned} x \Vdash_X^0 p & \quad \text{if } x \in p, \text{ where } p \in \check{P}X, \\ x \Vdash_X^0 \bigvee \varphi & \quad \text{if } x \Vdash_X^0 a \text{ for some } a \in \varphi, \\ x \Vdash_X^0 \bigwedge \varphi & \quad \text{if } x \Vdash_X^0 a \text{ for all } a \in \varphi. \end{aligned}$$

Clearly the relation between $\llbracket \cdot \rrbracket_0$ and \Vdash_X^0 is given by

$$x \in \llbracket a \rrbracket_0 \text{ iff } x \Vdash_X^0 a,$$

for all $x \in X$ and all $a \in \mathcal{L}_0\check{P}X$.

We note for future reference that $\llbracket \cdot \rrbracket_0$ gives rise to a natural transformation.

Proposition 7.1 *The family of homomorphisms $\{\llbracket \cdot \rrbracket_0^X\}_{X \in \text{Set}}$ is a natural transformation $\mathbb{F}\check{P} \rightarrow \check{P}$ and, therefore, also a natural transformation $\llbracket \cdot \rrbracket_0 : \mathcal{L}_0\check{P} \rightarrow \check{P}$.*

Proof. Naturality of $\llbracket \cdot \rrbracket_0$ is a matter of routine checking. The key for the proof is that for any function $f : X \rightarrow Y$, $\check{P}f : \check{P}Y \rightarrow \check{P}X$ is a Boolean homomorphism. QED

Turning our attention to depth-one formulas, perhaps the easiest way to explain their one-step semantics is to introduce a similar relation $\Vdash_X^1 \subseteq TX \times \mathcal{L}_1\check{P}X$:

$$\begin{aligned} TX, \xi \Vdash_X^1 \nabla \alpha & \text{ if } (\xi, \alpha) \in \overline{T}(\Vdash_X^0), \\ TX, \xi \Vdash_X^1 \bigvee \varphi & \text{ if } TX, \xi \Vdash_X^1 c \text{ for some } c \in \varphi, \\ TX, \xi \Vdash_X^1 \bigwedge \varphi & \text{ if } TX, \xi \Vdash_X^1 c \text{ for all } c \in \varphi. \end{aligned}$$

Remark 7.2 It is instructive to have a look at the relationship between the one-step semantics of depth-one formulas and the coalgebraic semantics for arbitrary formulas from Definition 5.9. Roughly, the definition of the one-step semantics of a formula captures precisely what is needed to inductively define the semantics of the logic.

More precisely, let (X, ξ) be some T -coalgebra and let, for $i < \omega$, $mng_i : \mathcal{L}_i \rightarrow \check{P}X$ be the map, that maps any formula $a \in \mathcal{L}_i$ of modal rank i to its coalgebraic meaning, that is, for all $a \in \mathcal{L}_i$ and all $x \in X$ we let $x \in mng_i(a)$ if $x \Vdash a$. Now we claim that for any $k < \omega$ and any $\nabla \alpha \in \mathcal{L}_{k+1}$ we have

$$x \Vdash_{\mathbb{X}} \nabla \alpha \text{ iff } TX, \xi(x) \Vdash_X^1 \nabla (Tmng_k)\alpha. \quad (31)$$

To see this, first observe that by induction on the Boolean structure of \mathcal{L}_k -formulas, we may show that for any $a \in \mathcal{L}_k$ and any $x \in X$, we have $x \Vdash_{\mathbb{X}} a$ iff $x \Vdash_X^0 mng_k(a)$. In other words, we have

$$(\Vdash_{\mathbb{X}}) \upharpoonright_{X \times \mathcal{L}_k} = \Vdash_X^0 ; mng_k^\checkmark. \quad (32)$$

Based on this, we may reason as follows:

$$\begin{aligned} x \Vdash_{\mathbb{X}} \nabla \alpha & \iff \xi(x) \overline{T} \Vdash_{\mathbb{X}} \alpha && \text{(definition of } \Vdash) \\ & \iff \xi(x) \overline{T}((\Vdash_{\mathbb{X}}) \upharpoonright_{X \times \mathcal{L}_k}) \alpha \\ & \iff \xi(x) \overline{T} \Vdash_X^0 ; mng_k^\checkmark \alpha && \text{(equation (32))} \\ & \iff \xi(x) \overline{T} \Vdash_X^0 (Tmng_k)\alpha && \text{(properties of relation lifting)} \\ & \iff TX, \xi(x) \Vdash_X^1 \nabla (Tmng_k)\alpha. && \text{(definition of } \Vdash^1) \end{aligned}$$

In words: if we assume that we have already defined the interpretation of all formulas of modal rank k then the one-step semantics allows us to extend this interpretation to formulas of rank $k + 1$.

The relation \Vdash_X^1 provides a natural semantics for terms of depth one, and induces a natural semantic equivalence relation.

Definition 7.3 Given a set X , we define the one-step semantics $\llbracket a' \rrbracket_1$ of a formula $a' \in \mathcal{L}_1(\check{P}X)$ as

$$\llbracket a' \rrbracket_1 := \{\xi \in TX \mid TX, \xi \Vdash_X^1 a'\}.$$

We say that two formulas $a', b' \in \mathcal{L}_1(\check{P}X)$ are *semantically one-step equivalent*, notation: $a' \equiv_{sem} b'$, if $\llbracket a' \rrbracket_1 = \llbracket b' \rrbracket_1$. \triangleleft

Remark 7.4 Alternatively but equivalently, we can define the $\llbracket \cdot \rrbracket_1$ as follows. Apply T to the map $\llbracket \cdot \rrbracket_0$, and compose with the function λ_X^T to obtain

$$\lambda_X^T \circ T[\llbracket \cdot \rrbracket_0] : T_\omega \mathcal{L}_0 \check{P}X \rightarrow \check{P}TX.$$

This map then provides us with an interpretation of the basic formulas in $\mathcal{L}_1\check{P}X = \mathcal{L}_0 T_\omega^\nabla \mathcal{L}_0 \check{P}X$, namely the ones of the form $\nabla \alpha \in T_\omega^\nabla \mathcal{L}_0 \check{P}X$:

$$\mu_X(\nabla \alpha) := (\lambda_X^T \circ T[\llbracket \cdot \rrbracket_0])(\alpha).$$

Now $\llbracket \cdot \rrbracket_1$ may be identified with $U\tilde{\mu}_X : \mathcal{L}_0 T_\omega^\nabla \mathcal{L}_0 \check{P}X \rightarrow \check{P}TX$. Occasionally, we will write $\llbracket \cdot \rrbracket_1$ also for the Boole-morphism $\tilde{\mu}_X : \mathbb{F}T_\omega^\nabla \mathcal{L}_0 \check{P}X \rightarrow \check{\mathbb{P}}TX$.

To match the semantic notions of equivalence between $\mathcal{L}_1\check{P}X$ -formulas, we introduce a one-step version of the derivation system \mathbf{M} , associated with the presentation $C\check{P}X$ of the power set Boolean algebra $\check{P}X$. Formal definitions will be given below, but the basic idea is straightforward: modify \mathbf{M} by (i) restricting attention to the depth-0 and depth-1 formulas over the set PX of (formal) variables, and (ii) adding the ‘true facts about $\check{P}X$ ’ as additional axioms. The resulting derivation system naturally induces an interderivability relation on $\mathcal{L}_1\check{P}X$ -formulas that we shall denote as $\equiv_{\mathbf{M}C\check{P}X}$ for reasons that we will clarify in Remark 7.7 further on. This then raises the question whether the two equivalence relations are the same or not, and the main aim of this section is to provide an affirmative answer to this question.

Theorem 7.5 (1-step soundness and completeness) *For any set X , and for any pair of formulas $c, d \in \mathcal{L}_1PX$ we have*

$$c \equiv_{sem} d \text{ iff } c \equiv_{\mathbf{M}C\check{P}X} d. \quad (33)$$

Our proof of this result will be algebraic, and before we can move to the details of the proof, we need to set up the appropriate framework for this.

We now define the one-step derivation system $\mathbf{M}\langle G; R \rangle$ associated with a presentation $\langle G; R \rangle$. Recall that \mathcal{L}_0G and $\mathcal{L}_1G = \mathcal{L}_0T_\omega^\nabla \mathcal{L}_0(G)$ are the set of depth zero and depth one formulas in G , respectively. In this section if we want to stress the difference between the two kinds of formulas, we shall use a, b, \dots for formulas in $\mathcal{L}_0(G)$, and c, d, \dots for formulas in $\mathcal{L}_1(G)$. An \mathcal{L}_0G -inequality is an inequality of the form $a \preccurlyeq b$, with $a, b \in \mathcal{L}_0G$; and likewise for \mathcal{L}_1G . Intuitively, we obtain $\mathbf{M}\langle G; R \rangle$ from \mathbf{M} by restricting attention to \mathcal{L}_0G - and \mathcal{L}_1G -inequalities, and adding the (in)equalities of R as additional axioms.

Definition 7.6 Given a presentation $\langle G; R \rangle$, we let $\mathbf{M}\langle G; R \rangle$ denote the *one-step derivation system* associated with $\langle G; R \rangle$. The language of $\mathbf{M}\langle G; R \rangle$ consists of \mathcal{L}_0G -inequalities, and \mathcal{L}_1G -inequalities, and its axioms and rules are those of \mathbf{M} , together with the set

$$R^\preccurlyeq := \{a \preccurlyeq b, b \preccurlyeq a \mid (a, b) \in R\}.$$

A $\mathbf{M}\langle G; R \rangle$ -*derivation* is a well-founded tree, labelled with \mathcal{L}_0G - and \mathcal{L}_1G -inequalities, such that (i) the leaves of the tree are labelled with axioms of \mathbf{M} or with inequalities in R^\preccurlyeq , (ii) 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. \triangleleft

We will leave it for the reader to verify that in $\mathbf{M}\langle G; R \rangle$ -derivations, a parent node is generally labelled with the same type of inequality (i.e. \mathcal{L}_0G versus \mathcal{L}_1G) as its children; the single exception is the rule $(\nabla 1)$ which links \mathcal{L}_0 -inequalities of the premisses to an \mathcal{L}_1 -inequality in the conclusion. As a corollary, $\mathbf{M}\langle G; R \rangle$ -derivation trees can be divided into a (possibly empty) upper \mathcal{L}_0G -part and a (possibly empty) lower \mathcal{L}_1G -part.

Remark 7.7 We can now clarify the syntactic interderivability notion of our one-step soundness and completeness theorem. Given a set X , recall that $C\check{P}X$ is the canonical presentation of the Boolean algebra $\check{P}X$, and observe that $\equiv_{\mathbf{M}C\check{P}X}$ is the associated relation of derivable equivalence of $\mathcal{L}_1\check{P}X$ -terms in the one-step derivation system $\mathbf{M}C\check{P}X$. It is *this* derivation system that Theorem 7.5, stating that the semantic equivalence relation is the same as the relation $\equiv_{\mathbf{M}C\check{P}X}$, is concerned with.

Remark 7.8 Definition 7.6 corrects and clarifies the corresponding definition in this paper’s earlier incarnation, where the one-step proof system $\mathbf{M}\langle G; R \rangle$ was not properly specified. In particular, the sentence in [34, Definition 22], ‘in which *only* elements of X and $\mathfrak{L}(X)$ may be used’ (where X denotes the set of generators) was not only rather vague, but in fact *mistaken*: it would not permit nontrivial applications of the derivation rules $(\nabla 2)$ and $(\nabla 3)$, since these require the use of more terms in $\mathcal{L}_0(X)$ than only the generators in X themselves.

7.2 The functor M on presentations

As we will see now, the notion of a one-step derivation system induces a functor on the category of presentations.

Definition 7.9 Given a presentation $\langle G; R \rangle$, we let $M\langle G; R \rangle$ denote the presentation given as

$$M\langle G; R \rangle := \langle T_\omega^\nabla \mathcal{L}_0(G); \equiv_{\mathbf{M}\langle G; R \rangle} \rangle.$$

For a presentation morphism $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$, the definition

$$Mf : \nabla\alpha \mapsto \nabla(T_\omega \widehat{f})\alpha$$

provides us with a map $Mf : T_\omega^\nabla \mathcal{L}_0(G) \rightarrow T_\omega^\nabla \mathcal{L}_0(G')$. \triangleleft

In other words, Mf maps generators of the presentation $M\langle G; R \rangle$ to generators of the presentation $M\langle G'; R' \rangle$. We will now show that Mf is in fact a presentation morphism from $M\langle G; R \rangle$ to $M\langle G'; R' \rangle$.

Remark 7.10 To be more precise, we need to compose Mf with the unit $\eta_{T_\omega^\nabla \mathcal{L}_0(G')}$ of the monad \mathcal{L}_0 , instantiated at $T_\omega^\nabla \mathcal{L}_0(G')$, in order to obtain a map with the right codomain, $\mathcal{L}_0 T_\omega^\nabla \mathcal{L}_0(G')$. In the sequel we will suppress this subtlety.

Our key tool in the proof that Mf is a presentation morphism, consists of a natural way to transform $\mathbf{M}\langle G; R \rangle$ -derivations into $\mathbf{M}\langle G'; R' \rangle$ -proofs.

Proposition 7.11 *If $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ is a presentation morphism, then there is a map $(\cdot)^f$ transforming $\mathbf{M}\langle G; R \rangle$ -derivations into $\mathbf{M}\langle G'; R' \rangle$ -derivations such that*

$$\mathcal{D} : c \preceq d \Rightarrow \mathcal{D}^f : \widehat{Mf}c \preceq \widehat{Mf}d.$$

for every \mathcal{L}_1G -inequality $c \preceq d$.

Proof. As an easy auxiliary result we need that for any two terms $a, b \in \mathcal{L}_0G$,

$$a \sqsubseteq_{\mathbf{M}\langle G; R \rangle} b \iff a \sqsubseteq_R b, \quad (34)$$

where $a \sqsubseteq_R b$ means that $a \equiv_R a \wedge b$. From (34) and the fact that f is a presentation morphism it is easy to derive that

$$a \sqsubseteq_{\mathbf{M}\langle G; R \rangle} b \text{ only if } \widehat{f}a \sqsubseteq_{\mathbf{M}\langle G'; R' \rangle} \widehat{f}b. \quad (35)$$

We now turn to the proof of the Proposition proper, which will be based on a straightforward induction on the complexity of $\mathcal{D} : c \preceq d$, where c and d are \mathcal{L}_1 -formulas. We make a case distinction as to the last rule applied in \mathcal{D} .

First assume that the last applied rule in \mathcal{D} was $(\nabla 1)$. That is, the formulas c and d in $\mathcal{D} : c \preceq d$ are of the form $c = \nabla\alpha$ and $d = \nabla\beta$, for some α and β in $T_\omega \mathcal{L}_0G$, respectively, and we may assume that \mathcal{D} is of the following form:

$$\mathcal{D} : \frac{\{\mathcal{D}_{ab} : a \preceq b \mid (a, b) \in Z\}}{\nabla\alpha \preceq \nabla\beta}$$

Here $Z \subseteq \text{Base}(\alpha) \times \text{Base}(\beta)$ is some set with $(\alpha, \beta) \in \overline{T}Z$, and such that for every pair $(a, b) \in Z$, there is a depth zero derivation $\mathcal{D}_{ab} : a \preceq b$.

Define $Z' := \{(\widehat{f}a, \widehat{f}b) \mid (a, b) \in Z\}$, or, equivalently, $Z' := (\widehat{f})^\circ; Z; \widehat{f}$. Then it follows from (35) that for each $(a', b') \in Z'$, there is a derivation $\mathcal{D}_{a'b'}^f : a' \preceq b'$. Using the properties of relation lifting

we find that $\overline{TZ}' = (T\widehat{f})^\circ; \overline{TZ}; T\widehat{f}$, and from this it is immediate that $(T\widehat{f}\alpha, T\widehat{f}\beta) \in \overline{TZ}'$. Combining these observations, we may transform the derivation \mathcal{D} into

$$\mathcal{D}^f : \frac{\{\mathcal{D}_{a'b'} : a' \preceq b' \mid (a', b') \in Z'\}}{\nabla T\widehat{f}\alpha \preceq \nabla T\widehat{f}\beta}$$

But then we are done, since $Mf(\nabla\alpha) = \nabla T\widehat{f}\alpha$, and likewise for β .

Second, suppose that the last applied rule in \mathcal{D} was $(\nabla 2)$. That is, \mathcal{D} ends with

$$\mathcal{D} : \frac{\{\mathcal{D}_\Phi : \nabla(T \wedge)\Phi \preceq d \mid \Phi \in SRD(A)\}}{\wedge\{\nabla\alpha \mid \alpha \in A\} \preceq d}$$

We are to transform \mathcal{D} into a derivation \mathcal{D}^f of the inequality $\wedge\{\nabla\alpha' \mid \alpha' \in A'\} \preceq \widehat{M}fd$, where $A' := \{T\widehat{f}\alpha \mid \alpha \in A\}$. Working towards an application of $(\nabla 2)$, we claim that

$$SRD(A') \subseteq \left\{ \Phi' \in T_\omega \left(\bigcup_{\alpha \in A} Base(T\widehat{f}\alpha) \right) \mid \exists \Phi \in SRD(A) \text{ such that } TPf(\Phi) = \Phi' \right\}. \quad (36)$$

To see why this is so, consider an arbitrary slim redistribution Φ' of A' . First observe that

$$\widehat{f}[Base[A]] = \bigcup_{\alpha \in A} (P\widehat{f})(Base(\alpha)) = \bigcup_{\alpha \in A} Base((T\widehat{f})\alpha) = Base[A'], \quad (37)$$

where the second identity is by the fact that $Base : T_\omega \rightarrow P_\omega$ is a natural transformation (cf. Fact 2.14). If we restrict \widehat{f} to the set $Base[A]$, by (37) we obtain a *surjective* map

$$g : Base[A] \rightarrow Base[A'].$$

From the surjectiveness of g it follows that $(Pg) \circ (\check{P}g) = id_{PBase[A']}$, and so we also find that $(TPg) \circ (T\check{P}g) = id_{T_\omega PBase[A']}$. Hence if we define

$$\Phi := (T\check{P}g)\Phi',$$

we see that $\Phi' = TPg(\Phi) = TP\widehat{f}(\Phi)$. Therefore, using $\in; Pf \subseteq f; \in$, it is easy to see that $\alpha(\overline{T} \in)\Phi$ implies $Tf\alpha(\overline{T} \in)\Phi'$ for all $\alpha \in T_\omega \mathcal{L}_0 G$. Thus, in order to prove (36) it suffices to prove that Φ is a slim redistribution of A . To see why this is the case, first observe that by definition of g we have that $T\check{P}g : T_\omega PBase[A'] \rightarrow T_\omega PBase[A]$, and so we find that $\Phi \in T_\omega PBase[A]$. It is left to prove that every element of A is a lifted member of Φ .

Take an arbitrary element $\alpha \in A$, then $Tg(\alpha) \in A'$ by definition of A' and g , and so $Tg(\alpha)$ is a lifted member of Φ' by the assumption that $\Phi' \in SRD(A')$. This means that $(\alpha, \Phi) \in (Tg); (\overline{T} \in); (T\check{P}g)$. The key observation now is that $(Tg); (\overline{T} \in); (T\check{P}g) \subseteq \overline{T} \in$, which is immediate from $g; \in; (\check{P}g) \subseteq \in$ by the properties of relation lifting. Applying this key observation we find that $(\alpha, \Phi) \in \overline{T} \in$ as required. This finishes the proof of (36).

Returning to the construction of our derivation \mathcal{D}^f , consider an arbitrary slim redistribution Φ' of A' , which by (36) we may assume to be of the form $(TP\widehat{f})\Phi$ with $\Phi \in SRD(A)$. Applying the inductive hypothesis to the derivation \mathcal{D}_Φ we obtain a proof $\mathcal{D}_\Phi^f : \widehat{M}f\nabla(T \wedge)(\Phi) \preceq \widehat{M}fd$. However, from $\widehat{f} \circ \wedge = \wedge \circ (P\widehat{f})$ we obtain that

$$\widehat{M}f\nabla(T \wedge)(\Phi) = \nabla(T\widehat{f})(T \wedge)(\Phi) = \nabla(T \wedge)(TP\widehat{f})(\Phi) = \nabla(T \wedge)\Phi'.$$

In other words, for any $\Phi' \in SRD(A')$ there is a derivation of the inequality $\nabla(T \wedge)\Phi' \preceq \widehat{M}f q$. Putting all these derivations together, one application of $(\nabla 2)$ gives the desired derivation

$$\mathcal{D}^f : (\widehat{M}f) \left(\bigwedge \{ \nabla \alpha \mid \alpha \in A \} \right) \preceq \widehat{M}f d.$$

Now suppose that the last applied rule in \mathcal{D} was $(\nabla 3)$. In this case \mathcal{D} has the following shape:

$$\mathcal{D} : \frac{\{ \mathcal{D}_\alpha : \nabla \alpha \preceq q \mid \alpha \bar{T}(\in) \Phi \}}{\nabla(T \vee)\Phi \preceq d}$$

In order to see which inequality we need to derive, we first compute

$$\widehat{M}f(\nabla(T \vee)\Phi) = \nabla(T \widehat{f})(T \vee)\Phi = \nabla(T \vee)(TP \widehat{f})\Phi,$$

where the latter identity follows from the fact that $\widehat{f} \circ \vee = \vee \circ P \widehat{f}$. We are looking for a derivation of the inequality $\nabla(T \vee)(TP \widehat{f})\Phi \preceq \widehat{M}f d$. Since we want to apply the rule $(\nabla 3)$, we first compute the set of lifted members of $(TP \widehat{f})\Phi$. But since $\in; (P \widehat{f})^\smile = \widehat{f}^\smile; \in$, applying relation lifting we obtain $\bar{T} \in; (TP \widehat{f})^\smile = (T \widehat{f})^\smile; \bar{T} \in$. This immediately shows that

$$(\alpha', (TP \widehat{f})\Phi) \in \bar{T} \in \text{ iff } \alpha' = T \widehat{f} \alpha \text{ for some } \alpha \bar{T} \in \Phi.$$

By the induction hypothesis, for each $\alpha \bar{T} \in \Phi$ we have a derivation $\mathcal{D}_\alpha^f : \nabla T \widehat{f} \alpha \preceq \widehat{M}f d$. In other words, for every lifted member α' of $(TP \widehat{f})\Phi$, there is a derivation $\mathcal{D}_{\alpha'} : \nabla \alpha' \preceq \widehat{M}f d$. But then by one application of $(\nabla 3)$ we are done:

$$\mathcal{D}^f : \frac{\{ \mathcal{D}_{\alpha'} : \nabla \alpha' \preceq q \mid \alpha \bar{T}(\in) (TP \widehat{f})\Phi \}}{\widehat{M}f(\nabla(T \vee)\Phi) \preceq \widehat{M}f d}$$

Finally, the cases where the last applied rule in \mathcal{D} was a propositional one, are left as exercises to the reader. QED

Given Proposition 7.11 it is not difficult to prove that M is a functor.

Theorem 7.12 $M : \text{Pres} \rightarrow \text{Pres}$ is a functor. In addition, M maps pre-isomorphisms to pre-isomorphisms.

Proof. Since it is not difficult to verify that M preserves identity arrows and distributes over composition, we confine our attention to the proof that M maps presentation morphisms to presentation morphisms.

Let $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ be a presentation morphism, and let $c, d \in \mathcal{L}_0(T_\omega^\nabla \mathcal{L}_0 G) = \mathcal{L}_1 G$ be such that $c \equiv_{\mathbf{M}\langle G; R \rangle} d$, that is, there are $\mathbf{M}\langle G; R \rangle$ -derivations $\mathcal{D}_1 : c \preceq d$ and $\mathcal{D}_2 : d \preceq c$. Proposition 7.11 provides us with $\mathbf{M}\langle G'; R' \rangle$ -derivations $\mathcal{D}_1^f : \widehat{M}f c \preceq \widehat{M}f d$ and $\mathcal{D}_2^f : \widehat{M}f d \preceq \widehat{M}f c$. This means that we have $\widehat{M}f c \equiv_{\mathbf{M}\langle G'; R' \rangle} \widehat{M}f d$, as is required for Mf to be a presentation morphism.

In order to prove that M maps pre-isomorphisms to pre-isomorphisms, a routine proof will show that M preserves pre-inverses. QED

7.3 The functor \mathbb{M} and its algebras

Given the intimate relation between Boolean algebras and their presentations, it should come as no surprise that the presentation functor M naturally induces a functor on the category of Boolean algebras.

Definition 7.13 The functor $\mathbb{M} : \mathbf{BA} \rightarrow \mathbf{BA}$ is defined as $\mathbb{M} := B \circ M \circ C$. \triangleleft

To explain this functor in words, first consider the objects. Given a Boolean algebra \mathbb{A} with carrier $A := U\mathbb{A}$, the elements of $\mathbb{M}\mathbb{A}$ are the equivalence classes of the form $[a]_{\mathbf{MC}\mathbb{A}}$, where $a \in \mathcal{L}_1 A$ is a depth-one term over the carrier of \mathbb{A} , and the equivalence relation $\equiv_{\mathbf{MC}\mathbb{A}}$ is the interderivability relation in the one-step derivation system $\mathbf{MC}\mathbb{A}$ which takes, as its additional axioms, the *diagram* $\Delta_{\mathbb{A}}$ of \mathbb{A} (listing the ‘true facts’ about \mathbb{A}). Summarizing, we find that

$$U\mathbb{M}\mathbb{A} = \mathcal{L}_1 A / \equiv_{\mathbf{MC}\mathbb{A}}.$$

In order to explain the action of \mathbb{M} on a homomorphism $f : \mathbb{A} \rightarrow \mathbb{A}'$, the upshot of Theorem 7.12 is that the map

$$\mathbb{M}f : [a]_{\mathbf{MC}\mathbb{A}} \mapsto [\mathcal{L}_1 Uf(a)]_{\mathbf{MC}\mathbb{A}'}, \quad (38)$$

correctly defines a homomorphism $\mathbb{M}f : \mathbb{M}\mathbb{A} \rightarrow \mathbb{M}\mathbb{A}'$. Here \mathcal{L}_1 is given in Definition 5.8, and the observation (38) is a direct consequence of the definitions and of the following proposition.

Proposition 7.14 Let $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ be a presentation morphism. If f maps generators to generators (in the sense that $f[G] \subseteq G'$), then

$$\widehat{Mf} = \mathcal{L}_1 f.$$

Proof. Suppose that $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ maps generators to generators, then it is immediate that $\widehat{f} = \mathcal{L}_0 f$. From this it follows that $Mf = T_{\omega}^{\nabla} \widehat{f} = T_{\omega}^{\nabla} \mathcal{L}_0 f$, and since Mf also maps generators to generators, we find that $\widehat{Mf} = \mathcal{L}_0 Mf = \mathcal{L}_0 T_{\omega}^{\nabla} \mathcal{L}_0 f = \mathcal{L}_1 f$. QED

For future reference we mention the following.

Definition 7.15 Given algebra \mathbb{B} , we shall denote with $\rho_{\mathbb{B}} : \mathcal{L}_1(U\mathbb{B}) \rightarrow \mathbb{M}\mathbb{B}$ the map

$$\rho_{\mathbb{B}} : b \mapsto [b]_{\mathbf{MC}\mathbb{B}},$$

that is, $\rho_{\mathbb{B}} = \widetilde{\eta}_{\mathbf{MC}\mathbb{B}}$ is the quotient map sending a formula b to its equivalence class under $\equiv_{\mathbf{MC}\mathbb{B}}$. \triangleleft

Proposition 7.16 The family of homomorphisms $\rho_{\mathbb{B}}$, with \mathbb{B} ranging over the class of Boolean algebras, provides a natural transformation $\rho : \mathcal{L}_1 U \rightarrow \mathbb{M}$.

Proof. Let $f : \mathbb{B} \rightarrow \mathbb{B}'$ be some Boolean homomorphism. In order to prove that ρ is a natural transformation, we need to show that the diagram below commutes:

$$\begin{array}{ccc} \mathbb{B} & \mathcal{L}_1 U\mathbb{B} & \xrightarrow{\rho_{\mathbb{B}}} & \mathbb{M}\mathbb{B} \\ f \downarrow & \mathcal{L}_1 Uf \downarrow & & \mathbb{M}f \downarrow \\ \mathbb{B}' & \mathcal{L}_1 U\mathbb{B}' & \xrightarrow{\rho_{\mathbb{B}'}} & \mathbb{M}\mathbb{B}' \end{array} \quad (39)$$

This follows from a straightforward unfolding of the definitions: For any $b \in \mathcal{L}_1 U\mathbb{B}$ we have

$$(\mathbb{M}f \circ \rho_{\mathbb{B}})(b) = \mathbb{M}f([b]_{\mathbf{MC}\mathbb{B}}) = [\mathcal{L}_1 Uf(b)]_{\mathbf{MC}\mathbb{B}'} = \rho_{\mathbb{B}'}(\mathcal{L}_1 Uf(b)) = (\rho_{\mathbb{B}'} \circ \mathcal{L}_1 Uf)(b).$$

Here the second step is by (38) above. QED

It turns out that \mathbb{M} has some nice properties that will be of use later on. In particular, we may show that \mathbb{M} is *finitary* and *preserves embeddings*. Intuitively, being finitary 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} . (Note that this is not obvious since we may be dealing with an infinitary proof system.) Formally, we need the following definition, referring to [6] for more details. Recall that a partial order is *directed* if any finite set of elements has an upper bound.

Definition 7.17 Given a category \mathcal{C} , a *directed diagram* over \mathcal{C} is a diagram which is indexed by a directed partial order. An endofunctor on \mathcal{C} is *finitary* if it preserves colimits of directed diagrams. \triangleleft

In case of an endofunctor on \mathbf{Set} this definition is equivalent to the one of Section 2.

Example 7.18 Given a Boolean algebra \mathbb{B} , let $\langle \text{Sub}(\mathbb{B}), \subseteq \rangle$ be the set of finite subalgebras of \mathbb{B} , ordered by inclusion. We can turn this poset into a diagram $S_{\mathbb{B}}$ by supplying, for each pair of finite subalgebras \mathbb{B}' and \mathbb{B}'' such that $\mathbb{B}' \subseteq \mathbb{B}''$, the (unique) inclusion $\iota_{\mathbb{B}'\mathbb{B}''}$. Since the variety \mathbf{BA} is *locally finite*, which means that every finitely generated Boolean algebra is finite, one may easily see that every Boolean algebra \mathbb{B} is the directed colimit of its associated diagram $S_{\mathbb{B}}$.

In fact, it is a routine exercise to verify that for an endofunctor on the category on Boolean algebras to be finitary, it suffices to preserve the directed colimits of the subalgebra diagrams described in Example 7.18.

Proposition 7.19 \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 $a_1, a_2 \in \mathcal{L}_1 A$, 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 $a_1, a_2 \in \mathcal{L}_1 A'$, where we define $A' := U\mathbb{A}'$.

We claim that

$$\vdash_{\mathbf{MC}\mathbb{A}} a_1 \preceq a_2 \quad \text{iff} \quad \vdash_{\mathbf{MC}\mathbb{A}'} a_1 \preceq a_2. \quad (40)$$

The interesting direction of (40) 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 $a' \in \mathbb{A}'$. (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 [50]. But if f is a homomorphism, by Proposition 7.11 it follows from $\vdash_{\mathbf{MC}\mathbb{A}} a_1 \preceq a_2$ that $\vdash_{\mathbf{MC}\mathbb{A}'} \widehat{M}f(a_1) \preceq \widehat{M}f(a_2)$. Since f restricts to the identity on \mathbb{A}' , so does $\widehat{M}f = \mathcal{L}_1 f$ on $\mathcal{L}_1 A'$. As a direct consequence we find that $\widehat{M}f(a_i) = a_i$, for both $i = 1, 2$. Thus, indeed, $\vdash_{\mathbf{MC}\mathbb{A}'} a_1 \preceq a_2$, which proves (40).

It is now easy to see that \mathbb{M} is a finitary functor. As mentioned above, it suffices to show that $\mathbb{M}\mathbb{A}$ is a directed colimit of the image $\mathbb{M}S_{\mathbb{A}}$ under \mathbb{M} of the subalgebra diagram $S_{\mathbb{A}}$ of \mathbb{A} (see Example 7.18). Given a finite subalgebra \mathbb{B} of \mathbb{A} , let $e_{\mathbb{B}}$ denote the inclusion homomorphism, $e_{\mathbb{B}} : \mathbb{B} \hookrightarrow \mathbb{A}$. We claim that

$$\langle \mathbb{M}\mathbb{A}, \mathbb{M}e_{\mathbb{B}} \rangle_{\mathbb{B} \in S_{\mathbb{A}}} \text{ is a colimit of } \mathbb{M}S_{\mathbb{A}}. \quad (41)$$

Since for every pair \mathbb{B}, \mathbb{B}' such that $\mathbb{B} \hookrightarrow \mathbb{B}' \hookrightarrow \mathbb{A}$, we have $e_{\mathbb{B}} = e_{\mathbb{B}'} \circ \iota_{\mathbb{B}\mathbb{B}'}$, it is obvious from the functoriality of \mathbb{M} that $\langle \mathbb{M}\mathbb{A}, \mathbb{M}e_{\mathbb{B}} \rangle_{\mathbb{B} \in S_{\mathbb{A}}}$ is a cocone over $\mathbb{M}S_{\mathbb{A}}$. To see why it is in fact a colimit, suppose that $\langle \mathbb{D}, d_{\mathbb{B}} \rangle_{\mathbb{B} \in S_{\mathbb{A}}}$ is another cocone over $\mathbb{M}S_{\mathbb{A}}$, and take an arbitrary element of $\mathbb{M}\mathbb{A}$. By definition, this element is of the form $[a]_{\mathbf{MC}\mathbb{A}}$ for some formula $a \in \mathcal{L}_1 A$. Let, as above, \mathbb{A}' be a finite subalgebra of \mathbb{A} such that $a \in \mathcal{L}_1 A'$, then it follows from (40) that the following provides a well-defined homomorphism $d : \mathbb{M}\mathbb{A} \rightarrow \mathbb{D}$:

$$d([a]_{\mathbf{MC}\mathbb{A}}) := d_{\mathbb{A}'}([a]_{\mathbf{MC}\mathbb{A}'}).$$

We leave it as an exercise for the reader to verify that d is the unique homomorphism $d : \mathbb{M}\mathbb{A} \rightarrow \mathbb{D}$ such that for all $\mathbb{B} \hookrightarrow \mathbb{A}$, the following diagram commutes:

$$\begin{array}{ccc} \mathbb{M}\mathbb{B} & \xrightarrow{\mathbb{M}e_{\mathbb{B}'}} & \mathbb{M}\mathbb{A} \\ & \searrow d_{\mathbb{B}'} & \downarrow d \\ & & \mathbb{D} \end{array}$$

This proves (41), and as mentioned this suffices to establish that \mathbb{M} is finitary.

For the second part of the Proposition, let $e : \mathbb{A} \rightarrow \mathbb{B}$ be an embedding. Without loss of generality we will assume that e is actually an 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 $a_1, a_2 \in A$:

$$\vdash_{\mathbb{M}\mathbb{C}\mathbb{B}} a_1 \preceq a_2 \text{ implies } \vdash_{\mathbb{M}\mathbb{C}\mathbb{A}} a_1 \preceq a_2. \quad (42)$$

But the proof of (42) simply follows from two applications of (40). QED

In the sequel we will be interested in algebras for the functor \mathbb{M} . Recall that these are pairs of the form $\langle \mathbb{A}, f \rangle$, where \mathbb{A} is some Boolean algebra, and f is a homomorphism from $\mathbb{M}\mathbb{A}$ to \mathbb{A} . First of all, we will see that such \mathbb{M} -algebras are Moss algebras in disguise.

Definition 7.20 Given an \mathbb{M} -algebra $\langle \mathbb{A}, f \rangle$, we let $V\langle \mathbb{A}, f \rangle$ denote the Moss algebra

$$V\langle \mathbb{A}, f \rangle := \langle U\mathbb{A}, \neg^{\mathbb{A}}, \bigvee^{\mathbb{A}}, \bigwedge^{\mathbb{A}}, \nabla^{V\langle \mathbb{A}, f \rangle} \rangle.$$

Here we define $\nabla^{V\langle \mathbb{A}, f \rangle} : T_{\omega}U\mathbb{A} \rightarrow U\mathbb{A}$ by recalling that $T_{\omega}U\mathbb{A}$ is a subset of $\mathcal{L}_1U\mathbb{A}$, and putting

$$\nabla^{V\langle \mathbb{A}, f \rangle} \alpha := f(\rho_{\mathbb{A}}(\nabla \alpha)),$$

where $\rho_{\mathbb{A}}$ is as in Definition 7.15. In addition, given an \mathbb{M} -morphism $g : \langle \mathbb{A}, f \rangle \rightarrow \langle \mathbb{A}', f' \rangle$, we define Vg to be the morphism $Vg : V\mathbb{A} \rightarrow V\mathbb{A}'$ given by

$$Vg := Ug.$$

That is, as a map, Vg is simply the same as g . ◁

We leave it for the reader to verify that with this definition, V actually defines a functor transforming \mathbb{M} -algebras into Moss algebras.

Proposition 7.21 *The operation V defines a functor*

$$V : \mathbf{Alg}_{\mathbf{BA}}(\mathbb{M}) \rightarrow \mathbf{Alg}_{\mathbf{Set}}(A_M).$$

Because \mathbb{M} is a finitary functor we can define the initial \mathbb{M} -algebra to be the colimit of the first ω steps of the initial sequence of \mathbb{M} .

Definition 7.22 *The initial sequence*

$$\mathbb{2} \xrightarrow{j_0} \mathbb{M}\mathbb{2} \xrightarrow{j_1} \mathbb{M}^2\mathbb{2} \xrightarrow{j_2} \dots \quad \mathbb{M}^k\mathbb{2} \xrightarrow{j_k} \mathbb{M}^{k+1}\mathbb{2} \xrightarrow{j_{k+1}} \dots \quad (43)$$

results from starting with j_0 as the unique homomorphism from $\mathbb{2}$ to $\mathbb{M}\mathbb{2}$, and defining $j_{k+1} := \mathbb{M}j_k$, for all $k \in \mathbb{N}$. We let $\mathbb{M}^{\omega}\mathbb{2}$, with the collection of maps $(i_k : \mathbb{M}^k\mathbb{2} \rightarrow \mathbb{M}^{\omega}\mathbb{2})_{k \in \omega}$, denote the colimit of this sequence. ◁

In the following Proposition we gather some facts about these structures.

- Proposition 7.23** 1. For each $k \in \omega$, the map $j_k : \mathbb{M}^k \mathfrak{D} \rightarrow \mathbb{M}^{k+1} \mathfrak{D}$ is an embedding, and so is the map $i_k : \mathbb{M}^k \mathfrak{D} \rightarrow \mathbb{M}^\omega \mathfrak{D}$.
2. There is a map $j_\omega : \mathbb{M}^\omega \mathfrak{D} \rightarrow \mathbb{M}^{\omega+1} \mathfrak{D}$ such that the following diagram commutes, for every $k \in \omega$:

$$\begin{array}{ccc} \mathbb{M}^k \mathfrak{D} & \xrightarrow{j_k} & \mathbb{M}^{k+1} \mathfrak{D} \\ i_k \downarrow & & \downarrow \mathbb{M}i_k \\ \mathbb{M}^\omega \mathfrak{D} & \xrightarrow{j_\omega} & \mathbb{M}^{\omega+1} \mathfrak{D} \end{array}$$

3. The map j_ω has an inverse $\heartsuit^{\mathcal{M}} : \mathbb{M}^{\omega+1} \mathfrak{D} \rightarrow \mathbb{M}^\omega \mathfrak{D}$.
4. The structure $\langle \mathbb{M}^\omega \mathfrak{D}, \heartsuit^{\mathcal{M}} \rangle$ is an initial \mathbb{M} -algebra.
5. For all $k \in \omega$ we have that $i_{k+1} = \heartsuit^{\mathcal{M}} \circ \mathbb{M}i_k$.

Proof. Part 1 is immediate by Proposition 7.19 and basic category theory. Part 2 follows from $\mathbb{M}^\omega \mathfrak{D}$ being a colimit of the initial sequence (43). The inverse of j_ω , mentioned in part 3, exists by the facts that the initial sequence is a chain, and hence directed, and that T preserves directed colimits.

For part 4, consider an arbitrary \mathbb{M} -algebra $\mathbb{A} \xleftarrow{\alpha} \mathbb{M}\mathbb{A}$, and define the co-cone $\langle \mathbb{A}, \alpha_k : \mathbb{M}^k \mathfrak{D} \rightarrow \mathbb{A} \rangle$ as follows: $\alpha_0 : \mathfrak{D} \rightarrow \mathbb{A}$ is given by initiality, and for $k \in \omega$ we put $\alpha_{k+1} := \lambda \circ \alpha_k$. Then by $\mathbb{M}^\omega \mathfrak{D}$ being the colimit of the initial sequence, there is a unique map $\alpha_\omega : \mathbb{M}^\omega \mathfrak{D} \rightarrow \mathbb{A}$ such that $\alpha_k = \alpha_\omega \circ i_k$, for all $k \in \omega$. Now consider the following diagram:

$$\begin{array}{ccc} \mathbb{M}^\omega \mathfrak{D} & \xrightarrow{j_\omega} & \mathbb{M}^{\omega+1} \mathfrak{D} \\ \alpha_\omega \downarrow & & \downarrow \mathbb{M}\alpha_\omega \\ \mathbb{A} & \xleftarrow{\alpha} & \mathbb{M}\mathbb{A} \end{array} \quad (44)$$

This diagram commutes by $\mathbb{M}^\omega \mathfrak{D}$ being the colimit of the initial sequence. Finally, consider the map $\heartsuit^{\mathcal{M}}$ of part 3. Then

$$\begin{aligned} \alpha_\omega \circ \heartsuit^{\mathcal{M}} &= (\alpha \circ \mathbb{M}\alpha_\omega \circ j_\omega) \circ \heartsuit^{\mathcal{M}} && \text{(diagram (44) commutes)} \\ &= \alpha \circ \alpha_\omega && (j_\omega \text{ and } \heartsuit^{\mathcal{M}} \text{ are converses)} \end{aligned}$$

and from this part 4 is immediate.

Finally, for part 5, fix $k \in \omega$. By definition, $\langle \mathcal{M}, i_n \rangle_{n \in \omega}$ is a co-cone of the initial sequence, and so we have $i_k = j_k \circ i_{k+1}$. From this it follows by (the diagram of) part 2 of this Proposition that $j_\omega \circ i_{k+1} = \mathbb{M}i_k$, and from this we easily derive by part 3 that $i_{k+1} = j_\omega^{-1} \circ \mathbb{M}i_k = \heartsuit^{\mathcal{M}} \circ \mathbb{M}i_k$. QED

The above Proposition justifies the following Definition.

Definition 7.24 We let \mathcal{M} denote the \mathbb{M} -algebra $\langle \mathbb{M}^\omega \mathfrak{D}, \heartsuit^{\mathcal{M}} \rangle$, and we will refer to this structure as the *initial \mathbb{M} -algebra*. \triangleleft

Remark 7.25 In the sequel, we will be interested in the Moss algebra $V\mathcal{M}$. Observe that the nabla operation $\nabla^{V\mathcal{M}}$ of this structure is defined as $\nabla^{V\mathcal{M}}(\alpha) = \heartsuit^{\mathcal{M}}(\rho_{\mathbb{M}^\omega \mathfrak{D}}(\nabla\alpha))$, and so by definition of $\heartsuit^{\mathcal{M}}$ we find that

$$\nabla^{V\mathcal{M}}(\alpha) = j_\omega^{-1}(\rho_{\mathbb{M}^\omega \mathfrak{D}}(\nabla\alpha)).$$

7.4 Proof of One-Step Soundness

In this subsection we will establish one-step soundness of the one-step derivation system; that is, we prove the direction from right to left of Theorem 7.5.

Proposition 7.26 *For any set X , and for any pair of formulas $c, d \in \mathcal{L}_1PX$ we have*

$$c \equiv_{sem} d \text{ if } c \equiv_{MCPX} d. \quad (45)$$

Proof. We argue by induction on derivations, so that clearly it suffices to show that each of the rules $(\nabla 1) - (\nabla 3)$ is sound. Fix a set X .

Case $(\nabla 1)$. Let $\sqsubseteq_0 \subseteq \mathcal{L}_0\check{P}X \times \mathcal{L}_0\check{P}X$ be the relation of ‘provable inequality’ $a \sqsubseteq_0 b$ if the inequality $a \preceq b$ is derivable. It is straightforward to see that for all $a, b \in \mathcal{L}_0\check{P}X$, it follows from $a \sqsubseteq b$ that $\llbracket a \rrbracket_0 \subseteq \llbracket b \rrbracket_0$. (This boils down to showing that our Boolean axioms of Table 2 are sound.) Hence it remains to show that for all $\alpha, \beta \in \mathcal{L}_1\check{P}X$, we have

$$\text{if } \alpha \bar{T}Z \beta \text{ for some } Z \subseteq \sqsubseteq_0, \text{ then } \llbracket \nabla \alpha \rrbracket_1 \subseteq \llbracket \nabla \beta \rrbracket_1. \quad (46)$$

For this purpose, assume that $\alpha \bar{T}Z \beta$ for some $Z \subseteq \sqsubseteq_0$, and take an arbitrary element $\xi \in TX$ such that $TX, \xi \Vdash_1 \nabla \alpha$. Then by definition of \Vdash_1 , we have $\xi \bar{T} \Vdash_0 \alpha$, so that by the properties of relation lifting we obtain that $\xi \bar{T}(\Vdash_0; Z) \beta$. However, it is straightforward to verify that $\Vdash_0; Z \subseteq \Vdash_0; \sqsubseteq_0 \subseteq \Vdash_0$, and so we obtain that $\xi \bar{T} \Vdash_0 \beta$. From this it is immediate that $TX, \xi \Vdash_1 \nabla \beta$.

Case $(\nabla 2)$. Given a set $A \subseteq T_\omega \mathcal{L}_0(PX)$ and an element $\xi \in TX$, assume that $TX, \xi \Vdash_1 \nabla \alpha$ for each $\alpha \in A$. We need to prove that $TX, \xi \Vdash_1 \nabla(T \wedge) \Phi$ for some $\Phi \in SRD(A)$. To come up with a suitable Φ , let $B := \bigcup Base[A]$ and consider the map $\varphi : X \rightarrow P_\omega B$ given by

$$\varphi : x \mapsto \{b \in B \mid X, x \Vdash_0 b\}.$$

We claim that the set

$$\Phi := (T\varphi)(\xi)$$

fulfills our requirements.

First of all, in order to prove that $TX, \xi \Vdash \nabla(T \wedge)(\Phi)$, observe that by definition of φ , we have $\varphi; \wedge \subseteq \Vdash_0$. Hence by the properties of relation lifting, it follows that $T\varphi; T\wedge \subseteq \bar{T} \Vdash_0$. In particular, we find that $(\xi, (T\wedge)(\Phi)) \in \bar{T} \Vdash_0$; but then it is immediate from the definitions that $TX, \xi \Vdash \nabla(T \wedge)(\Phi)$.

Second, by definition we have $\Phi \in T_\omega P_\omega B$ and so, in order to show that $\Phi \in SRD(A)$, it suffices to prove that $\alpha \in \lambda_{PX}^T(\Phi)$ for all $\alpha \in A$. For this purpose, observe that $\varphi; \in^\vee = \Vdash_0 \upharpoonright_{X \times B}$. Then by the properties of relation lifting we obtain $T\varphi; (\bar{T} \in^\vee) = \bar{T} \Vdash_0 \upharpoonright_{TX \times TB}$. In particular, since $\xi \bar{T} \Vdash_0 \upharpoonright_{TX \times TB} \alpha$ by assumption, it follows that $\alpha \bar{T} \in T\varphi(\xi) = \Phi$, as required.

Case $(\nabla 3)$. We could prove the soundness of $(\nabla 3)$ analogously to our proof of Proposition 5.18, but we prefer to give a different proof here, stressing the role of the distributive of λ^T over the power set *monad*, cf. Fact 3.14.

Fix an element $\Phi \in T_\omega P_\omega \mathcal{L}_0PX$. Given Remark 7.4, it suffices to show that

$$\llbracket \nabla(T \vee)(\Phi) \rrbracket_1 = \bigcup \{ \lambda_X^T(T[\cdot]_0(\alpha)) \mid \alpha \bar{T} \in \Phi \}. \quad (47)$$

The point is now that (47) can be read off the following diagram, where we tacitly use the fact that λ^T restricts to a natural transformation $\lambda_X^T : T_\omega P_\omega \rightarrow PT_\omega$ (see Proposition 3.17).

$$\begin{array}{ccc}
T_\omega P_\omega \mathcal{L}_0 P X & \xrightarrow{\lambda_{\mathcal{L}_0 P X}^T} & P T_\omega \mathcal{L}_0 P X \\
T_\omega \vee \downarrow & \searrow^{T_\omega P_\omega([\cdot]_0)} & \downarrow P T[\cdot]_0 \\
T_\omega \mathcal{L}_0 P X & \xrightarrow{\lambda_{P X}^T} & P T_\omega P X \\
T_\omega [\cdot]_0 \downarrow & \swarrow^{T_\omega \cup} & \downarrow P \lambda_X^T \\
T_\omega P X & & P P T_\omega P X \\
\lambda_X^T \downarrow & \xleftarrow{\cup_{T X}} & \downarrow \\
P T_\omega X & & P P T_\omega X
\end{array} \tag{48}$$

To see this, first observe that the left hand side of (47) corresponds to the left edge of the diagram, where an arbitrary element $\Phi \in T_\omega P_\omega \mathcal{L}_0 P X$ is mapped to

$$\lambda_X^T(T_\omega[\cdot]_0(T_\omega \vee(\Phi))) = [\nabla(T \vee)(\Phi)]_1.$$

Similarly, the right hand side of (47) corresponds to clockwise following $\Phi \in T_\omega P_\omega \mathcal{L}_0 P X$ along the outer edges of the diagram, from the upper left to the lower left corner, arriving at the object $\bigcup\{\lambda_X^T(T[\cdot]_0(\alpha)) \mid \alpha \bar{T} \in \Phi\}$.

Therefore in order to show (47) it suffices to show that the diagram commutes. But this is fairly straightforward. First observe that

$$[\cdot]_0 \circ \vee = \bigcup \circ P_\omega [\cdot]_0, \tag{49}$$

as a straightforward verification will reveal. After applying the functor T_ω to (49), we immediately obtain that the left quadrangle of (48) commutes. The right-hand quadrangle commutes since λ^T is natural. And finally, the pentagon commutes since λ^T is a distributive law over the power set monad, see Fact 3.14. As a consequence, the diagram (48) itself commutes. QED

7.5 Proof of One-Step Completeness

We now turn to the one-step completeness of our derivation system. Our proof is based on properties of algebras of the form $\mathbb{M}\mathbb{B}$, with \mathbb{B} an arbitrary finite Boolean algebra. With $At\mathbb{B}$ denoting the set of atoms of \mathbb{B} , we can formulate our key insight by stating that the Boolean algebra $\mathbb{M}\mathbb{B}$ is join-generated by its ‘lifted atoms’, that is, its elements of the form $[\nabla\alpha]$ with $\alpha \in T(At\mathbb{B})$. That is to say, we can prove that every element x of $\mathbb{M}\mathbb{B}$ is the join of the elements in $T(At\mathbb{B})$ below it:

$$x = \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B}), [\nabla\alpha] \leq x\}.$$

Here, as elsewhere in this subsection, the join is taken in the algebra $\mathbb{M}\mathbb{B}$, and may be happen to be taken over an infinite set; in that case, the statement should be read as saying that ‘the join on the righthandside exists, and it is equal to the lefthandside’. As we will see, in the case that the functor does not preserve finite sets, this is a convenient way of treating infinitary rules as identities.

Arriving at the proof details, in order to establish the one-step completeness of \mathbf{M} , we need to prove the direction from left to right of (33). We will reason by contraposition, showing that for arbitrary $a', b' \in \mathcal{L}_1 \dot{P} X$:

$$a' \not\equiv_{MCPX} b' \text{ implies } \llbracket a' \rrbracket_1 \neq \llbracket b' \rrbracket_1.$$

Given the fact that our logic extends classical propositional logic, we may confine ourselves to the case where $b' = \perp$.

Fix an element $a' \in \mathcal{L}_1\check{P}X$, and assume that a' is one-step consistent: $a' \not\equiv_{ax} \perp$, or, equivalently, $[a'] > \perp^{\mathbb{M}\mathbb{B}}$. We will prove that a' is one-step satisfiable: $\llbracket a' \rrbracket_1 \neq \emptyset$. Let $\{\alpha_1, \dots, \alpha_n\}$ be the (finite!) set of elements $\alpha \in T_\omega\check{P}X$ such that $\nabla\alpha$ occurs in a' , and define

$$Base(a') := \bigcup_{1 \leq i \leq n} Base(\alpha_i).$$

This is a finite subset of $\mathcal{L}_0\check{P}X$, that is, a finite set of Boolean formulas in which the subsets of X are the formal generators. Let $D \subseteq_\omega PX$ be the collection of those subsets of X that actually occur (as a formal object) in one of the formulas in $Base(a')$, and let \mathbb{B} be the subalgebra of $\check{P}X$ that is generated by D . Then both D and \mathbb{B} are finite (whereas their elements may themselves be infinite subsets of X). The point is that \mathbb{B} is a finite subalgebra of $\check{P}X$ such that $a' \in \mathcal{L}_1(U\mathbb{B})$.

It follows by the key lemma in the one-step completeness proof, Theorem 7.28 below, that

$$[a'] = \bigvee^{\mathbb{M}\mathbb{B}} \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B}), \nabla\alpha \sqsubseteq a'\}. \quad (50)$$

But since a' is consistent, we have that $[a'] > \perp$, and so we may conclude that there actually exists an $\alpha \in T_\omega At\mathbb{B}$ such that $\nabla\alpha \sqsubseteq a'$ — if there were no such α , then the righthandside of (50) would evaluate to \perp . By Proposition 7.27 we obtain for this α that $\llbracket \nabla\alpha \rrbracket_1 \neq \emptyset$, and so by soundness we may conclude that $\llbracket a' \rrbracket_1 \supseteq \llbracket \nabla\alpha \rrbracket_1 \neq \emptyset$. In other words, we find that $\llbracket a' \rrbracket_1$ is one-step satisfiable, as required.

Proposition 7.27 *Fix a set X and let $\alpha \in T_\omega(At\mathbb{B})$ for some finite subalgebra \mathbb{B} of $\mathbb{P}X$. Then $\llbracket \nabla\alpha \rrbracket_1 \neq \emptyset$.*

Proof. Clearly the set $At\mathbb{B} \subseteq PX$ forms a partition of X . Let $h : At\mathbb{B} \rightarrow X$ be a choice function, that is, $h(a) \in a$ for each $a \in At\mathbb{B}$. Using the properties of relation lifting, it is not hard to derive from this that $(Th)(\alpha)\bar{T}(\in_X)\alpha$ for each lifted atom α . It follows immediately that $(Th)(\alpha) \in \llbracket \nabla\alpha \rrbracket_1$. QED

The following is the key lemma in the one-step completeness proof.

Theorem 7.28 *Let \mathbb{B} be a finite Boolean algebra.*

1. *For any two elements $\alpha, \beta \in T_\omega(At\mathbb{B})$, we have*

$$[\nabla\alpha] \wedge [\nabla\beta] > \perp \text{ iff } \alpha = \beta. \quad (51)$$

2. *The top element of $\mathbb{M}\mathbb{B}$ satisfies*

$$\top^{\mathbb{M}\mathbb{B}} = \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B})\}. \quad (52)$$

3. *The set $\{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B})\}$ join-generates $\mathbb{M}\mathbb{B}$; that is, for all $a' \in \mathcal{L}_1 U\mathbb{B}$:*

$$[a'] = \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B}), [\nabla\alpha] \leq [a']\}. \quad (53)$$

Summarizing, the algebra $\mathbb{M}\mathbb{B}$ is atomic, with $At(\mathbb{M}\mathbb{B}) = \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B})\}$.

Proof. Throughout the proof we will abbreviate $A := At\mathbb{B}$ and $B := U\mathbb{B}$.

The proof of first two statements is immediate by Proposition 6.12 (take for φ the set A). Concerning the third statement of the Theorem, observe that the inequality ' \geq ' of (53) always holds, so it will be the opposite inequality that we need to establish. Our proof will be by induction on the complexity of a' (as a boolean formula over the set $T_\omega^\nabla \mathcal{L}_0 B$).

In the base case of the induction, a' is of the form $\nabla\beta$, with $\beta \in T_\omega \mathcal{L}_0 B$. Our first claim is that without loss of generality, we may assume that $\nabla\beta$ actually belongs to $T_\omega B$. The justification for this claim is that for any $b \in \mathcal{L}_0 B$ there is a $b_0 \in B$ such that the equation $b_0 \approx b$ is derivable in the proof system $C\mathbb{B}$ associated with the canonical presentation of \mathbb{B} : simply let $b_0 := \tilde{id}_B(b)$ be the element of B to which the term b evaluates. (For the definition of \tilde{id}_B we refer to 4.3.) Thus an application of $(\nabla 1)$ shows that for any $\beta \in T_\omega \mathcal{L}_0 B$ there is a $\beta_0 \in T_\omega B$ such that $\vdash_{MC\mathbb{B}} \nabla\beta \approx \nabla\beta_0$: simply take $\beta_0 := T\tilde{id}_B(\beta)$.

Hence, assume that indeed, $\beta \in T_\omega B$. Think of the finitary join as a map $\bigvee : P_\omega A \rightarrow B$. As such it is a bijection, and this property is inherited by the map $T\bigvee : T_\omega P_\omega A \rightarrow T_\omega B$. Furthermore, it is easy to verify that for any $\varphi \in P_\omega A$ and any $a \in A$, we have that

$$a \in_X \varphi \text{ iff } a \leq \bigvee \varphi, \quad (54)$$

which can be succinctly formulated as $\exists_X = \bigvee; \geq$ (where \bigvee now denotes the graph of the disjunction function). By the properties of relation lifting, this implies $\overline{T}(\exists_X) = T\bigvee; \overline{T}\geq$, which can again be reformulated as stating that for any $\Phi \in T_\omega P_\omega A$ and any $\alpha \in T_\omega A$ it holds that

$$\alpha \overline{T}(\in_X)\Phi \text{ iff } \alpha \overline{T} \leq (T\bigvee)\Phi. \quad (55)$$

Now consider an arbitrary element $\beta \in T_\omega B$, and let Φ be the (unique) element of $T_\omega P_\omega A$ such that $\beta = (T\bigvee)(\Phi)$. Then (55) reads that $\alpha \overline{T}(\in_X)\Phi$ iff $\alpha \overline{T}(\leq)\beta$, for all $\alpha \in T_\omega A$, and so axiom $(\nabla 3)$ instantiates to

$$[\nabla\beta] = \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega A \text{ and } \alpha \overline{T}(\leq)\beta\}. \quad (56)$$

But since by the nature of the one-step derivation system we have $\leq = \sqsubseteq$ on elements of PX , we also have $\overline{T}(\leq) = \overline{T}(\sqsubseteq)$. So if $\alpha \overline{T}(\leq)\beta$ then one application of $(\nabla 1)$ gives that $\nabla\alpha \sqsubseteq \nabla\beta$, which implies that $[\nabla\alpha] \leq [\nabla\beta]$. From this and (56) is immediate that

$$[\nabla\beta] \leq \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega(At\mathbb{B}), [\nabla\alpha] \leq [\nabla\beta]\}.$$

This finishes the base case of the inductive proof of (53).

For the inductive step of the proof there are three cases to consider. First, assume that a' is of the form $\bigvee_{i \in I} a'_i$ for some finite index set I . Then we may compute

$$\begin{aligned} [a'] &= \bigvee \{[a'_i] \mid i \in I\} && \text{(assumption)} \\ &= \bigvee \left\{ \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega A, [\nabla\alpha] \leq [a'_i]\} \mid i \in I \right\} && \text{(induction hypothesis)} \\ &= \bigvee \left\{ [\nabla\alpha] \mid \alpha \in T_\omega A, [\nabla\alpha] \leq [a'_i] \text{ for some } i \in I \right\} && \text{(associativity of } \bigvee \text{)} \\ &\leq \bigvee \left\{ [\nabla\alpha] \mid \alpha \in T_\omega A, [\nabla\alpha] \leq \bigvee_{i \in I} [a'_i] = a' \right\} && \text{(properties of } \bigvee \text{)} \end{aligned}$$

Second, consider the case that a' is a conjunction $\bigwedge_{i \in I} a'_i$ for some finite I . Now we have

$$\begin{aligned}
[a'] &= \bigwedge \{[a'_i] \mid i \in I\} && \text{(assumption)} \\
&= \bigwedge \left\{ \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega A, [\nabla\alpha] \leq [a'_i]\} \mid i \in I \right\} && \text{(induction hypothesis)} \\
&= \bigvee \left\{ \bigwedge_{i \in I} [\nabla\gamma(i)] \mid \gamma : I \rightarrow T_\omega A \text{ such that } [\nabla\gamma(i)] \leq [a'_i] \text{ for all } i \right\} && \text{(distributivity)} \\
&= \bigvee \left\{ [\nabla\gamma] \mid \gamma \in T_\omega A \text{ such that } [\nabla\gamma] \leq [a'_i] \text{ for all } i \right\} && \text{(part 1)} \\
&= \bigvee \left\{ [\nabla\gamma] \mid \gamma \in T_\omega A, [\nabla\gamma] \leq \bigwedge_{i \in I} [a'_i] = a' \right\} && \text{(properties of } \bigvee \text{)}
\end{aligned}$$

Here ‘distributivity’ refers to the fact that in any Boolean algebra, finite meets distribute over arbitrary joins, and ‘part 1’ refers to the first statement of this Theorem. The point here is that we only need to consider those meets $\bigwedge_{i \in I} [\nabla\gamma(i)]$ for which $\gamma(i) = \gamma(j)$ for all $i, j \in I$, since the other meets will reduce to \perp .

Finally, suppose that a' is a negation, say $a' = \neg b'$. We first claim that

$$\text{for all } \alpha \in T_\omega A \text{ either } \nabla\alpha \sqsubseteq b' \text{ or } \nabla\alpha \sqsubseteq \neg b'. \quad (57)$$

To see this, assume that $\nabla\alpha \not\sqsubseteq \neg b'$; then by propositional logic,

$$[\nabla\alpha] \wedge [b'] > \perp.$$

By the inductive hypothesis, we have $[b'] = \bigvee \{[\nabla\beta] \mid \beta \in T_\omega A, [\nabla\beta] \leq [b']\}$, and so by distributivity we obtain

$$\bigvee \{[\nabla\alpha] \wedge [\nabla\beta] \mid \beta \in T_\omega A, [\nabla\beta] \leq [b']\} > \perp.$$

But then there must be at least one $\beta \in T_\omega A$ with $[\nabla\alpha] \wedge [\nabla\beta] > \perp$ and $[\nabla\beta] \leq [b']$. By the first statement of this Theorem, we can only have $[\nabla\alpha] \wedge [\nabla\beta] > \perp$ if α is *identical* to β , and so indeed we find that $[\nabla\alpha] \leq [b']$. This proves (57).

Because of this we can rewrite $[\neg b']$ as follows:

$$\begin{aligned}
[\neg b'] &= [\neg b'] \wedge \bigvee \{[\nabla\alpha] \mid \alpha \in T_\omega A\} && \text{(part 2)} \\
&= \bigvee \{[\neg b'] \wedge [\nabla\alpha] \mid \alpha \in T_\omega A\} && \text{(distributivity)} \\
&= \bigvee \left(\{[\neg b' \wedge \nabla\alpha] \mid b' \sqsupseteq \nabla\alpha, \alpha \in T_\omega A\} \cup \{[\neg b' \wedge \nabla\alpha] \mid \neg b' \sqsupseteq \nabla\alpha, \alpha \in T_\omega A\} \right) && (57) \\
&= \bigvee \left(\{[\perp] \mid b' \sqsupseteq \nabla\alpha, \alpha \in T_\omega A\} \cup \{[\nabla\alpha] \mid \neg b' \sqsupseteq \nabla\alpha, \alpha \in T_\omega A\} \right) && \text{(immediate)} \\
&= \bigvee \{[\nabla\alpha] \mid [\neg b'] \geq [\nabla\alpha], \alpha \in T_\omega A\} && \text{(immediate)}
\end{aligned}$$

This settles the remaining inductive case, and thus finishes the proof of the third part of the Theorem. QED

7.6 Connecting algebra and coalgebra

Now that we have proved the one-step soundness and completeness of our logic, we will show how to connect the algebraic functor \mathbb{M} to the coalgebraic functor T by defining a natural transformation

$$\delta : \mathbb{M}\check{\mathbb{P}} \rightarrow \check{\mathbb{P}}T$$

which in fact provides an embedding δ_X for each set X .

For the definition of δ , note that given a set X , it follows from one-step soundness that $\llbracket a \rrbracket_1 = \llbracket b \rrbracket_1$ for all $a, b \in \mathcal{L}_1\check{P}X$ such that $[a]_{\mathbb{M}C\check{P}X} = [b]_{\mathbb{M}C\check{P}X}$. This ensures that the following is well-defined.

Definition 7.29 Given a set X , let

$$\delta_X([a]_{\mathbf{MCP}X}) := \llbracket a \rrbracket_1$$

define a map $\delta_X : \mathbf{MIP}X \rightarrow \check{\mathbb{P}}TX$. \triangleleft

Proposition 7.30 *The family of maps δ_X , with X ranging over the category \mathbf{Set} , provides a natural transformation $\delta : \mathbf{MIP} \rightarrow \check{\mathbb{P}}T$. Furthermore, each $\delta_X : \mathbf{MIP}X \rightarrow \check{\mathbb{P}}TX$ is an embedding.*

Proof. In order to demonstrate that δ is a natural transformation, we have to prove that for any function $f : X \rightarrow Y$ the following diagram commutes:

$$\begin{array}{ccc} \mathbf{MIP}X & \xrightarrow{\delta_X} & \check{\mathbb{P}}TX \\ \mathbf{MIP}f \uparrow & & \uparrow \check{\mathbb{P}}Tf \\ \mathbf{MIP}Y & \xrightarrow{\delta_Y} & \check{\mathbb{P}}TY \end{array}$$

In order to see that the above diagram commutes it suffices to show that it commutes on the generators of $\mathbf{MIP}Y$. Consider such a generator $\nabla\alpha \in T_\omega^\nabla \mathcal{L}_0 \check{\mathbb{P}}Y$. Then

$$\begin{aligned} \delta_X(\mathbf{MIP}(f)(\nabla\alpha)) &= \delta_X([T_\omega^\nabla \mathcal{L}_0 \check{\mathbb{P}}(f)(\nabla\alpha)]) &= \llbracket [T_\omega^\nabla \mathcal{L}_0 \check{\mathbb{P}}(f)(\nabla\alpha)] \rrbracket_1 \\ &\stackrel{\text{Remark 7.4}}{=} \lambda_X^T(T[\cdot]_0(T\mathcal{L}_0 \check{\mathbb{P}}(f)(\alpha))) &= \lambda_X^T(T([\cdot]_0 \circ \mathcal{L}_0 \check{\mathbb{P}}(f))(\alpha)) \\ &\stackrel{[\cdot]_0 \text{ natural, Lem. 7.1}}{=} \lambda_X^T(T(\check{\mathbb{P}}f \circ [\cdot]_0)(\alpha)) &= \lambda_X^T(T\check{\mathbb{P}}f \circ T[\cdot]_0(\alpha)) \\ &\stackrel{\lambda \text{ natural}}{=} \check{\mathbb{P}}Tf(\lambda_Y^T(T[\cdot]_0(\alpha))) &= \check{\mathbb{P}}Tf(\llbracket \nabla\alpha \rrbracket_1) \\ &= \check{\mathbb{P}}Tf(\delta_Y(\llbracket \nabla\alpha \rrbracket)) \end{aligned}$$

Let us finally show that δ_X is injective for an arbitrary set X . Suppose that $\delta_X([a]) = \delta_X([b])$ for some $a, b \in \mathcal{L}_0 T_\omega^\nabla \mathcal{L}_0 \check{\mathbb{P}}X$. By definition of δ_X that means that $\llbracket a \rrbracket_1 = \llbracket b \rrbracket_1$ which by one-step completeness of the logic entails that $[a] = [a']$ in $\mathbf{MIP}X$. QED

On the basis of this natural transformation we can define a second notion of complex algebra of a coalgebra, next to the Moss complex algebra of Definition 5.11.

Definition 7.31 Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor, and let $\mathbb{X} = \langle X, \xi \rangle$ be a T -coalgebra. We define the *complex \mathbb{M} -algebra* of \mathbb{X} as the pair $\mathbb{X}^* := \langle \check{\mathbb{P}}X, \delta_X \circ \check{\mathbb{P}}\xi \rangle$. \triangleleft

The link between the two kinds of complex algebras is given by the functor V from Definition 7.20 which allows us to see \mathbb{M} -algebras as Moss algebras.

Proposition 7.32 *Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard, weak pullback preserving functor. Then*

$$\mathbb{X}^+ = V\mathbb{X}^*.$$

for any T -coalgebra \mathbb{X} . Therefore, for any T -coalgebra \mathbb{X} and any formula $a \in \mathcal{L}$ we have

$$\text{mng}_{V\mathbb{X}^*}(a) = \text{mng}_{\mathbb{X}^+}(a) = \{x \in X \mid x \Vdash a\}.$$

8 Soundness and completeness

In this section we will apply Pattinson’s *stratification method* [43] in order to prove the soundness and completeness of our axiom system \mathbf{M} with respect to the coalgebraic semantics. This stratification method consists in showing that not only the *language* of our system, but also its *semantics* and our *logic* can be stratified in ω many layers. As we will see further on, the results in the previous section will then serve to glue these layers nicely together.

In order to understand the idea of the proof, first assume that a final T -coalgebra $\mathbb{Z} = \langle Z, \zeta : Z \rightarrow TZ \rangle$ exists. Then we could prove that the unique Moss morphism $mng_{\mathbb{Z}}$ from the initial Moss algebra \mathcal{L} to the algebra \mathbb{Z}^+ actually factors as $mng_{\mathbb{Z}} = Vmng_{\mathbb{Z}}^* \circ q$, where $q : \mathcal{L} \rightarrow \mathcal{M}$ is the quotient map modulo derivability (in the sense that $\ker(q)$ is the relation $\equiv_{\mathbf{M}}$ of interderivability in \mathbf{M}), and $mng_{\mathbb{Z}}^*$ is an *injective* \mathbb{M} -algebra morphism from \mathcal{M} to \mathbb{Z}^* :

$$\begin{array}{ccc} \mathcal{L} & \xrightarrow{q} & V\mathcal{M} \\ & \searrow mng_{\mathbb{Z}} & \downarrow Vmng_{\mathbb{Z}}^* \\ & & V\mathbb{Z}^* = \mathbb{Z}^+ \end{array}$$

On the basis of this we would prove that $a \not\sqsubseteq_{\mathbf{M}} b$ implies that $q(a) \not\sqsubseteq_{\mathcal{M}} q(b)$, and so by injectivity of m we would conclude that $mng_{\mathbb{Z}}(a) \not\sqsubseteq mng_{\mathbb{Z}}(b)$, providing a state $z \in Z$ such that $z \Vdash_{\mathbb{Z}} a$ and $z \not\vdash_{\mathbb{Z}} b$.

Since our set functor T generally does not admit a final coalgebra, we replace the final coalgebra with the *final sequence*.

Definition 8.1 The *final T -sequence* is defined as follows.

$$1 \xleftarrow{h_0} T1 \xleftarrow{h_1} T^2 1 \quad \dots \xleftarrow{h_n} T^n 1 \xleftarrow{h_{n+1}} T^{n+1} 1 \quad \dots \quad (58)$$

We denote by $1 = T^0 1$ the final object in \mathbf{Set} . The map $h_0 : T1 \rightarrow 1$ is given by finality and inductively, $h_{n+1} : T(T^n 1) \rightarrow T^n 1$ is defined to be the map $T^n h_0 = Th_n$. \triangleleft

The reader may think of the $T^n 1$ as approximating the final coalgebra. Indeed, if we let the final sequence run through all ordinals, we obtain the final coalgebra as a limit if it exists [4]. Intuitively, where the states of the final coalgebra provide all possible T -behaviors, the elements of $T^n 1$ represent all ‘ n -step behaviors’. Given a T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$, for each $n \in \omega$ we may canonically define a map $\xi_n : X \rightarrow T^n 1$ providing the n -step behavior of the states of \mathbb{X} .

Definition 8.2 Given a T -coalgebra $\mathbb{X} = \langle X, \xi \rangle$, we define the arrows $\xi_n : X \rightarrow T^n 1$, for $n \in \omega$, to the approximants of the final coalgebra by the following induction: $\xi_0 : X \rightarrow 1$ is given by finality of 1 in \mathbf{Set} , and $\xi_{n+1} := T\xi_n \circ \xi$. \triangleleft

Interestingly, every object $T^n 1$ in the final sequence can be equipped with coalgebra structure.

Definition 8.3 Let, for each $n \in \omega$, \mathbb{Z}_n be the coalgebra

$$\mathbb{Z}_n := (T^n 1, T^n g),$$

where g is an arbitrary but fixed map $g : 1 \rightarrow T1$. \triangleleft

As we will see in a moment, these ‘ n -final coalgebras’ display all possible n -step behaviours, and thus act as a canonical witness for all non-provable inequalities between formulas of depth n .

8.1 A stratification of the semantics

We first show how to slice the semantics of nabla formulas into layers. For that purpose we define the n -step meaning of depth- n modal formulas as a subset of the set $T^n 1$.

Definition 8.4 By induction on n we define maps $mng_n : \mathcal{L}_n \rightarrow PT^n 1$. For $n = 0$, we define mng_0 by initiality of \mathcal{L}_0 , or equivalently:

$$mng_0(a) := \begin{cases} 1 & \text{if } a \text{ is a tautology,} \\ \emptyset & \text{otherwise.} \end{cases}$$

Inductively, assuming that $mng_n : \mathcal{L}_n \rightarrow PT^n 1$ has been defined, we may compose $Tmng_n : T\mathcal{L}_n \rightarrow TPT^n 1$ with $\lambda_{PT^n 1}^T : TPT^n 1 \rightarrow PT^{n+1} 1$ to obtain

$$\lambda_{PT^n 1}^T \circ Tmng_n : T\omega \mathcal{L}_n \rightarrow PT^{n+1} 1.$$

Then we let $mng_{n+1} : \mathcal{L}_{n+1} \rightarrow PT^{n+1} 1$ be the unique Boole-homomorphism from $\mathbb{F}(T_\omega^\nabla \mathcal{L}_n)$ to $\mathbb{P}T^{n+1} 1$ that extends the mapping given by

$$\nabla \alpha \mapsto (\lambda_{PT^n 1}^T \circ Tmng_n(\alpha)) \quad \text{for } \nabla \alpha \in T_\omega^\nabla \mathcal{L}_n.$$

◁

The following proposition provides a clear link between the n -step meaning of formulas and the n -step behaviour map of a coalgebra.

Proposition 8.5 *Let \mathbb{X} be a coalgebra, and $a \in \mathcal{L}_n$ a formula of rank n . Then*

$$mng_{\mathbb{X}}(a) = (\check{P}\xi_n)(mng_n(a)).$$

Proof. The proof of the proposition is by induction on the modal depth and on the structure of the formula a . We only provide the induction case for $a = \nabla \alpha \in \mathcal{L}_{n+1}$ for some $n \in \omega$. In this case we have

$$\begin{aligned} mng_{\mathbb{X}}(\nabla \alpha) &= \check{P}\xi(\lambda_X(Tmng_{\mathbb{X}}(\alpha))) && \text{(definition of } mng_{\mathbb{X}}) \\ &= \check{P}\xi(\lambda_X(T\check{P}\xi_n(Tmng_n(\alpha)))) && \text{(induction hypothesis)} \\ &= \check{P}\xi\left(\check{P}T\xi_n(\lambda_{T^n 1}(Tmng_n(\alpha)))\right) && \text{(naturality of } \lambda) \\ &= \check{P}\xi_{n+1}(mng_{n+1}(\nabla \alpha)) && \text{(definition of } mng_{n+1} \text{ and } \xi_{n+1}) \end{aligned}$$

QED

The n -final coalgebra of Definition 8.3 has the interesting property that its n -step behaviour map is the *identity* map on $T^n 1$. As a corollary, the n -step meaning of any depth- n formula a coincides with its meaning in the n -step coalgebra.

Proposition 8.6 *Let a be a formula of depth n . Then*

$$mng_{\mathbb{Z}_n}(a) = mng_n(a).$$

Proof. It is not difficult to see that for the coalgebra \mathbb{Z}_n (and for this n), we have

$$(T^n g)_n := id_{T^n 1}. \tag{59}$$

We confine ourselves to a proof sketch. The basic idea of the proof is to prove inductively that $(T^n g)_k = h_{nk}$ for all $k \leq n$, where $h_{nk} : T^n 1 \rightarrow T^k 1$ is the map $h_{nk} := h_k \circ h_{k+1} \circ \dots \circ h_n$. Further details can be found in [43, Section 4].

The Proposition itself is immediate by Proposition 8.5 and (59).

QED

As a fairly direct corollary to the previous two propositions we can formulate our semantic stratification theorem. Basically it states that the meaning of depth- n formulas is determined at level n of the final sequence, and in the n -step final coalgebra \mathbb{Z}_n .

Theorem 8.7 (Semantic Stratification Theorem) *Let $a, b \in \mathcal{L}_n$ be formulas. Then the following are equivalent:*

1. $a \models_T b$;
2. $mng_n(a) \subseteq mng_n(b)$;
3. $mng_{\mathbb{Z}_n}(a) \subseteq mng_{\mathbb{Z}_n}(b)$.

Proof. The implication $1 \Rightarrow 3$ is immediate by the definitions, while the implication $2 \Rightarrow 1$ follows by Proposition 8.5: given a coalgebra $\mathbb{X} = \langle X, \xi \rangle$, we conclude from $mng_n(a) \subseteq mng_n(b)$ that $mng_{\mathbb{X}}(a) = (\mathbb{P}\xi_n)(mng_n(a)) \subseteq (\mathbb{P}\xi_n)(mng_n(b)) = mng_{\mathbb{X}}(b)$. The remaining implication $3 \Rightarrow 2$ follows directly by Proposition 8.6. QED

8.2 A stratification of the logic

To see in detail how our *logic* can be stratified, let us first introduce some terminology concerning the stratification of the language.

Definition 8.8 Let $G_0 := \emptyset$, and define inductively $G_{n+1} := T_\omega^\nabla \mathcal{L}_0 G_n = \{\nabla \alpha \mid \alpha \in T_\omega \mathcal{L}_0(G_n)\}$. In addition, let $e_0 : G_0 \rightarrow \mathcal{L}_0 G_1$ be the empty map, and define $e_{n+1} : G_{n+1} \rightarrow \mathcal{L}_0 G_{n+2}$ by putting $e_{n+1} := M e_n$. Finally, we let d_n denote the inclusion $d_n : \mathcal{L}_n \hookrightarrow \mathcal{L}$. ◁

Recall that \mathcal{L}_n denotes the set of formulas of rank n (see Definition 5.1), and observe that $\mathcal{L}_n = \mathcal{L}_0 G_n$, for all n , and that each \mathcal{L}_n is also the carrier of an algebra in **Boole**; this algebra will also be denoted as \mathcal{L}_n . Consequently, $\mathcal{L}_{n+1} = \mathcal{L}_1(G_n)$, which is different from $\mathcal{L}_1(\mathcal{L}_n) = \mathcal{L}_1(\mathcal{L}_0(G_n))$ since in **Boole** we do not identify terms which are equivalent in the theory of Boolean algebras. Also observe that the map $\widehat{e}_n : \mathcal{L}_0 G_n \rightarrow \mathcal{L}_0 G_{n+1}$ is in fact the embedding of \mathcal{L}_n into \mathcal{L}_{n+1} :

$$\widehat{e}_n : \mathcal{L}_n \hookrightarrow \mathcal{L}_{n+1},$$

and that the embedding $d_n : \mathcal{L}_n \hookrightarrow \mathcal{L}$ commutes with the one-step embeddings, in the sense that $d_n = d_{n+1} \circ \widehat{e}_n$.

We can now formulate our stratification theorem as follows. Recall that \mathcal{L} is the initial algebra in the category **Boole** $_\nabla$.

Theorem 8.9 (Axiomatic Stratification Theorem) *Let $m := mng_{V\mathcal{M}}$ be the unique homomorphism $m : \mathcal{L} \rightarrow V\mathcal{M}$ in the category of Moss algebras.*

1. *There are maps $q_n : \mathcal{L}_n \rightarrow \mathbb{M}^n \mathfrak{D}$, with each q_n a **Boole**-homomorphism, such that the following diagram (in the category **Boole**) commutes:*

$$\begin{array}{ccccccc}
 \mathcal{L}_0 & \xrightarrow{\widehat{e}_0} & \mathcal{L}_1 & \xrightarrow{\widehat{e}_1} & \mathcal{L}_2 & \cdots & \mathcal{L}_n & \xrightarrow{\widehat{e}_n} & \mathcal{L}_{n+1} & \cdots & \mathcal{L} \\
 q_0 \downarrow & & q_1 \downarrow & & q_2 \downarrow & & q_n \downarrow & & q_{n+1} \downarrow & & m \downarrow \\
 \mathfrak{D} & \xrightarrow{j_0} & \mathbb{M} \mathfrak{D} & \xrightarrow{j_1} & \mathbb{M}^2 \mathfrak{D} & \cdots & \mathbb{M}^n \mathfrak{D} & \xrightarrow{j_n} & \mathbb{M}^{n+1} \mathfrak{D} & \cdots & \mathbb{M}^\omega \mathfrak{D} \\
 & & & & & & & & & & i_0 \nearrow \\
 & & & & & & & & & & i_1 \nearrow \\
 & & & & & & & & & & i_2 \nearrow \\
 & & & & & & & & & & \cdots \nearrow \\
 & & & & & & & & & & i_n \nearrow \\
 & & & & & & & & & & i_{n+1} \nearrow
 \end{array} \tag{60}$$

2. In addition, $\ker(m) = \equiv_{\mathbf{M}}$; that is, $m(a) = m(b)$ iff a and b are provably equivalent in \mathbf{M} .

Before turning to the proof of this result, let us briefly summarize its meaning. Most importantly, Theorem 8.9 states that for each $n < \omega$, the Boolean algebra $\mathbb{M}^n\mathcal{D}$ coincides with the quotient of the Boole-algebra \mathcal{L}^n under the relation $\equiv_{\mathbf{M}}$ of provable equivalence in our derivation system \mathbf{M} . In addition, the quotient maps q_n commute with the inclusions \widehat{e}_n of \mathcal{L}_n into \mathcal{L}_{n+1} , and j_n from $\mathbb{M}^n\mathcal{D}$ into $\mathbb{M}^{n+1}\mathcal{D}$.

In order to prove Theorem 8.9, we will inductively define a relation \equiv_n of “ n -interderivability” between \mathcal{L}_n -formulas. We will see that for every n , the Boolean algebra $\mathbb{L}_n = \mathcal{L}_n / \equiv_n$ is isomorphic to $\mathbb{M}^n\mathcal{D}$, but also, that for formulas $a, b \in \mathcal{L}_n$, we have $a \equiv_n b$ iff $a \equiv_{\mathbf{M}} b$. The definition of \equiv_n will be such that

$$\langle G_{n+1}; \equiv_{n+1} \rangle = M \langle G_n; \equiv_n \rangle.$$

Definition 8.10 Let $\equiv_0 \subseteq \mathcal{L}_0 \times \mathcal{L}_0$ be the relation of provable equivalence between closed Boolean terms. Inductively, define the relation $\equiv_{n+1} \subseteq \mathcal{L}_{n+1} \times \mathcal{L}_{n+1}$ as the congruence relation of the presentation $M \langle G_n; \equiv_n \rangle$, and let \mathbb{L}_n denote the Boolean algebra $B \langle G_n; \equiv_n \rangle$, or equivalently, $\mathbb{L}_n = \mathcal{L}_n / \equiv_n$. Given a formula $a \in \mathcal{L}_n$, we let $[a]_n$ denote the equivalence class of a under the relation \equiv_n . \triangleleft

As we will see, the algebras \mathbb{L}_n form an intermediate row in the stratification diagram (60) (in the category **Boole**):

$$\begin{array}{ccccc}
 \mathcal{L}_0 & \xrightarrow{\widehat{e}_0} & \mathcal{L}_1 & \xrightarrow{\widehat{e}_1} & \mathcal{L}_2 & \cdots & \mathcal{L}_n & \xrightarrow{\widehat{e}_n} & \mathcal{L}_{n+1} & \cdots & (61) \\
 \tilde{\eta}_0 \downarrow & & \tilde{\eta}_1 \downarrow & & \tilde{\eta}_2 \downarrow & & \tilde{\eta}_n \downarrow & & \tilde{\eta}_{n+1} \downarrow & & \\
 \mathbb{L}_0 & \xrightarrow{Be_0} & \mathbb{L}_1 & \xrightarrow{Be_1} & \mathbb{L}_2 & \cdots & \mathbb{L}_n & \xrightarrow{Be_n} & \mathbb{L}_{n+1} & \cdots & \\
 f_0 \downarrow & & f_1 \downarrow & & f_2 \downarrow & & f_n \downarrow & & f_{n+1} \downarrow & & \\
 \mathcal{D} & \xrightarrow{j_0} & \mathbb{M}\mathcal{D} & \xrightarrow{j_1} & \mathbb{M}^2\mathcal{D} & \cdots & \mathbb{M}^n\mathcal{D} & \xrightarrow{j_n} & \mathbb{M}^{n+1}\mathcal{D} & \cdots &
 \end{array}$$

We now turn to the details of the proof of Theorem 8.9, step by step filling in diagram (61). Since we already discussed the embeddings \widehat{e}_n , $n \in \omega$, we start with the map $\tilde{\eta}_n$, which will denote the quotient map associated with the congruence \equiv_n .

Definition 8.11 Let $\eta_n : G_n \rightarrow \mathcal{L}_n / \equiv_n$ be the map given by $\eta_n : g \mapsto [g]_n$. \triangleleft

We may see the map η_n as a presentation morphism from $\langle G_n; \equiv_n \rangle$ to $C(\mathbb{L}_n)$ — as such it is the unit $\eta_{\langle G_n; \equiv_n \rangle}$ of the adjunction $B \dashv C$, and hence, a pre-isomorphism (cf. Theorem 4.13). This function extends to a homomorphism in **Boole**:

$$\tilde{\eta}_n : \mathcal{L}_n \rightarrow \mathbb{L}_n$$

which maps a formula $a \in \mathcal{L}_n$ to its n -equivalence class:

$$\tilde{\eta}_n : a \mapsto [a]_n.$$

Concerning the maps $Be_n : \mathbb{L}_n \rightarrow \mathbb{L}_{n+1}$, it is easy to see that they are indeed well-typed, but in order to prove that each Be_n is an embedding, some work will be needed. The embeddings $j_n : \mathbb{M}^n\mathcal{D} \rightarrow \mathbb{M}^{n+1}\mathcal{D}$ have been defined in Definition 7.22.

Finally, the isomorphisms f_n of diagram (61) will be defined inductively.

Definition 8.12 By induction on n we define Boolean homomorphisms $f_n : \mathbb{L}_n \rightarrow \mathbb{M}^n \mathfrak{2}$. For $n = 0$, we let f_0 be the (unique) isomorphism from \mathbb{L}_0 to $\mathfrak{2}$. For $n = k + 1$, we first define $p_{n+1} : \mathbb{L}_{n+1} \rightarrow \mathbb{M}\mathbb{L}_n$ by putting $p_{n+1} := BM\eta_n$. Then we compose the maps

$$\mathbb{L}_{n+1} \xrightarrow{p_{n+1}} \mathbb{M}\mathbb{L}_n \xrightarrow{\mathbb{M}f_n} \mathbb{M}^{n+1} \mathfrak{2},$$

and define $f_{n+1} := (\mathbb{M}f_n) \circ p_{n+1}$. \triangleleft

The following proposition gathers all the facts about the maps defined until now that are needed to prove that diagram (61) commutes:

Proposition 8.13 1. *In the category Pres of presentation each map e_n is a morphism $e_n : \langle G_n; \Xi_n \rangle \rightarrow \langle G_{n+1}; \Xi_{n+1} \rangle$, each map $\eta_n : \langle G_{n+1}; \Xi_{n+1} \rangle \rightarrow C\mathbb{L}_n$ is a pre-isomorphism, and each of the following diagrams commutes:*

$$\begin{array}{ccc} \langle G_n; \Xi_n \rangle & \xrightarrow{e_n} & \langle G_{n+1}; \Xi_{n+1} \rangle \\ \eta_n \downarrow & & \eta_{n+1} \downarrow \\ C(\mathbb{L}_n) & \xrightarrow{CB e_n} & C(\mathbb{L}_{n+1}) \end{array} \quad (62)$$

2. *In the category Boole, each of the following diagrams commutes:*

$$\begin{array}{ccc} \mathcal{L}_n & \xrightarrow{\hat{e}_n} & \mathcal{L}_{n+1} \\ \tilde{\eta}_n \downarrow & & \tilde{\eta}_{n+1} \downarrow \\ \mathbb{L}_n & \xrightarrow{B e_n} & \mathbb{L}_{n+1} \end{array} \quad (63)$$

3. *In the category BA of Boolean algebras, each map p_{n+1} is an isomorphism, and each of the following diagrams commutes:*

$$\begin{array}{ccc} \mathbb{L}_{n+1} & \xrightarrow{B e_{n+1}} & \mathbb{L}_{n+2} \\ p_{n+1} \downarrow & & p_{n+2} \downarrow \\ \mathbb{M}(\mathbb{L}_n) & \xrightarrow{\mathbb{M}B e_n} & \mathbb{M}(\mathbb{L}_{n+1}) \end{array} \quad (64)$$

4. *In the category Boole, each of the following diagrams commutes:*

$$\begin{array}{ccc} \mathcal{L}_{n+1} & \xrightarrow{\tilde{\eta}_{n+1}} & \mathbb{L}_{n+1} \\ \mathcal{L}_1 \eta_n \downarrow & & p_{n+1} \downarrow \\ \mathcal{L}_1 U(\mathbb{L}_n) & \xrightarrow{\rho_{\mathbb{L}_n}} & \mathbb{M}(\mathbb{L}_n) \end{array} \quad (65)$$

with $\rho_{\mathbb{L}_n}$ as in Definition 7.15.

5. *In the category BA of Boolean algebras, each map f_n is an isomorphism; each map $B e_n : \mathbb{L}_n \rightarrow \mathbb{L}_{n+1}$ is an embedding; and each of the following diagrams commutes:*

$$\begin{array}{ccc} \mathbb{L}_n & \xrightarrow{B e_n} & \mathbb{L}_{n+1} \\ f_n \downarrow & & f_{n+1} \downarrow \\ \mathbb{M}^n \mathfrak{2} & \xrightarrow{j_n} & \mathbb{M}^{n+1} \mathfrak{2} \end{array} \quad (66)$$

Proof.

1. It follows by a straightforward induction that every e_n is a presentation morphism. The other statements of this item follow from the fact that $\eta_n = \eta_{\langle G_n; \equiv_n \rangle}$, together with our earlier observation (cf. Theorem 4.13) that $\eta : Id_{\text{Pres}} \rightarrow CB$ is a natural transformation of which each $\eta_{\langle G; R \rangle}$ is a pre-isomorphism.
2. We claim that if $f : \langle G; R \rangle \rightarrow \langle G'; R' \rangle$ is the presentation morphism represented by one of the four arrows of the diagram (62), then the corresponding arrow \hat{f} in (63) is the *unique* Boole-morphism extending f (seen as a map between sets). For instance, if f is the presentation morphism $\eta_n : \langle G_n; \equiv_n \rangle \rightarrow C\mathbb{L}_n$, then using the fact that $\mathcal{L}_n = \mathcal{L}_0 G_n$ is the free Boole-algebra over G_n , it follows that $\hat{f} = \tilde{\eta}_n$ is the unique homomorphism in **Boole** from \mathcal{L}_n to \mathbb{L}_n . Or, to give a second example, Be_n is clearly the only homomorphism from \mathbb{L}_n to \mathbb{L}_{n+1} which “extends” $CBe_n : C\mathbb{L}_n \rightarrow C\mathbb{L}_{n+1}$.

From this it follows that both $\tilde{\eta}_{n+1} \circ \hat{e}_n$ and $Be_n \circ \tilde{\eta}_n$ are morphisms in **Boole** that extend the map $\eta_{n+1} \circ e_n = CBe_n \circ \eta_n$ (with the identity holding because diagram (62) commutes). But then, again by the freeness of \mathcal{L}_n over G_n in **Boole**, these two extensions must be equal, which is the same as to say that (63) commutes.

3. It is easy to see that our definition of the map p_{n+1} indeed provides an isomorphism, because

$$M\eta_n : \langle G_{n+1}; \equiv_{n+1} \rangle = M\langle G_n; \equiv_n \rangle \rightarrow M C\mathbb{L}_n,$$

is a pre-isomorphism in **Pres**, by Theorem 7.12 inheriting this property from $\eta_n : \langle G_n; \equiv_n \rangle \rightarrow C\mathbb{L}_n$, and B maps pre-isomorphisms to isomorphisms, see Proposition 4.12.

To prove that diagram (64) commutes it suffices to see that we may obtain it from diagram (62) by applying the functor BM .

4. Recall that the family of presentation morphisms $\eta_{\langle G; R \rangle} : \langle G; R \rangle \rightarrow CB\langle G; R \rangle$, defined by (15), constitutes a natural transformation $\eta : Id_{\text{Pres}} \rightarrow CB$. Instantiating the diagram which expresses this fact for the arrow $M\eta_n : M\langle G_n; \equiv_n \rangle \rightarrow M C\mathbb{L}_n$, we obtain the following commuting diagram:

$$\begin{array}{ccc} M\langle G_n; \equiv_n \rangle & \xrightarrow{\eta_{M\langle G_n; \equiv_n \rangle}} & CBM\langle G_n; \equiv_n \rangle = C\mathbb{L}_{n+1} \\ M\eta_n \downarrow & & \downarrow CBM\eta_n \\ M C\mathbb{L}_n & \xrightarrow{\eta_{M C\mathbb{L}_n}} & CBM C\mathbb{L}_n = C\mathbb{M}\mathbb{L}_n \end{array} \quad (67)$$

Now we can, similarly as in the proof of item 2, show that each of the arrows in (65) is the unique morphism in **Boole** that extends the corresponding map in (67). For example, consider the map $\mathcal{L}_1 \eta_n : \mathcal{L}_{n+1} \rightarrow \mathcal{L}_1 U\mathbb{L}_n$. It follows from a straightforward unravelling of the definitions that $\mathcal{L}_1 \eta_n$ extends $M\eta_n$ (see Proposition 7.14). The latter, as a function between sets, is just a map from $T_\omega^\nabla \mathcal{L}_0 G_n = G_{n+1}$ to the set of generators of the presentation $M C\mathbb{L}_n$, which is nothing but the set $T_\omega^\nabla \mathcal{L}_0 U\mathbb{L}_n$.

But then, again similar to the proof of item 2, we can prove that the maps $p_{n+1} \circ \tilde{\eta}_{n+1}$ and $\rho_{\mathbb{L}_n} \circ \mathcal{L}_1 \eta_n$ are identical, by noting that both are morphisms in **Boole** that extend the presentation morphism $CBM\eta_n \circ \eta_{M\langle G_n; \equiv_n \rangle} = \eta_{M C\mathbb{L}_n} \circ M\eta_n$ of diagram (67).

5. This part of the Proposition is proved by induction on n . For $n = 0$, the map f_0 is an isomorphism by definition, and the map Be_0 is an embedding by initiality of $\mathbb{2}$ in **BA**. Finally, the following

diagram commutes simply by the initiality of the algebra \mathbb{L}_0 in the category BA:

$$\begin{array}{ccc}
 \mathbb{L}_0 & \xrightarrow{Be_0} & \mathbb{L}_1 \\
 f_0 \downarrow & & \downarrow f_1 \\
 \mathfrak{2} & \xrightarrow{j_0} & \mathbb{M}\mathfrak{2}
 \end{array} \tag{68}$$

In the inductive case for $n + 1$, by hypothesis the map f_n is an isomorphism, and the map Be_n an embedding. From this it is immediate that $\mathbb{M}f_n$ is an isomorphism as well, and since p_{n+1} is an isomorphism by Proposition 8.13(2), it follows that the map f_{n+1} , being the composition of two isomorphisms, is an isomorphism as well.

Now consider the following diagram:

$$\begin{array}{ccccc}
 & & \mathbb{L}_{n+1} & \xrightarrow{Be_{n+1}} & \mathbb{L}_{n+2} & & \\
 & & \downarrow p_{n+1} & & \downarrow p_{n+2} & & \\
 f_{n+1} & \curvearrowright & \mathbb{M}(\mathbb{L}_n) & \xrightarrow{MBe_n} & \mathbb{M}(\mathbb{L}_{n+1}) & \curvearrowleft & f_{n+2} \\
 & & \downarrow \mathbb{M}f_n & & \downarrow \mathbb{M}f_{n+1} & & \\
 & & \mathbb{M}^{n+1}\mathfrak{2} & \xrightarrow{j_{n+1}} & \mathbb{M}^{n+2}\mathfrak{2} & &
 \end{array} \tag{69}$$

The upper rectangle of this diagram commutes by Proposition 8.13(2), and the lower rectangle, by applying the functor \mathbb{M} to the diagram (66) which commutes by the inductive hypothesis. As a consequence, the outer rectangle, which exactly corresponds to the diagram (66) for the case $n + 1$, commutes as well. Finally, then, the injectivity of Be_{n+1} is immediate by that of j_{n+1} , which was established in Lemma 7.23(1).

QED

By Proposition 8.13 it follows that the diagram (61) commutes.

For future reference we state the following technical fact, which links the quotient maps q_n and q_{n+1} to the natural transformation ρ of Definition 7.15, instantiated at the Boolean algebra $\mathbb{M}^n\mathfrak{2}$.

Proposition 8.14 *For any element $\alpha \in T_\omega\mathcal{L}_n$, we have*

$$q_{n+1}(\nabla\alpha) = \rho_{\mathbb{M}^n\mathfrak{2}}\nabla(Tq_n(\alpha)). \tag{70}$$

Proof. To see why this proposition holds, recall that $q_k = f_k \circ \tilde{\eta}_k$ for each $k \in \omega$, and consider the diagram below

$$\begin{array}{ccccc}
 T_\omega\mathcal{L}_n & \xrightarrow{\nabla_{G_n}} & \mathcal{L}_1(G_n) = \mathcal{L}_{n+1} & \xrightarrow{\tilde{\eta}_{n+1}} & \mathbb{L}_{n+1} & & \\
 T\tilde{\eta}_n \downarrow & & \mathcal{L}_1\eta_n \downarrow & & \downarrow p_{n+1} & & \\
 T_\omega U(\mathbb{L}_n) & \xrightarrow{\nabla_{U(\mathbb{L}_n)}} & \mathcal{L}_1 U(\mathbb{L}_n) & \xrightarrow{\rho_{\mathbb{L}_n}} & \mathbb{M}(\mathbb{L}_n) & \curvearrowleft & f_{n+1} \\
 Tf_n \downarrow & & \mathcal{L}_1 f_n \downarrow & & \downarrow \mathbb{M}f_n & & \\
 T_\omega U(\mathbb{M}^n\mathfrak{2}) & \xrightarrow{\nabla_{U(\mathbb{M}^n\mathfrak{2})}} & \mathcal{L}_1 U(\mathbb{M}^n\mathfrak{2}) & \xrightarrow{\rho_{\mathbb{M}^n\mathfrak{2}}} & \mathbb{M}^{n+1}\mathfrak{2} & &
 \end{array} \tag{71}$$

where, in order to simplify the diagram, we omit the forgetful functors to \mathbf{Set} on the right-hand side of the diagram and exploit our ambiguous notation allowing \mathcal{L}_1 to be considered as \mathbf{Set} -valued or \mathbf{Boole} -valued.

Here an arrow labelled ∇_G represents the function mapping an object $\alpha \in T_\omega \mathcal{L}_0 G$ to the corresponding formula $\nabla \alpha \in \mathcal{L}_1(G)$. Note that in the case that $G = U(\mathbb{L}_n)$ and $G = U\mathbb{M}^n \mathcal{L}$ we use the fact that $T_\omega G \subseteq T_\omega \mathcal{L}_0 G$.

We claim that all squares of (71) commute. To check this for the left squares this is simply a matter of unravelling the definitions, and the upper right square has been shown to commute in Proposition 8.13(4). Finally, that the lower right square commutes is a consequence of the fact that ρ is a natural transformation $\rho : \mathcal{L}_1 U \rightarrow \mathbb{M}$, cf. Proposition 7.16.

But if indeed all squares of (71) commute, then the identity (70) can simply be read off from the outer sides of the diagram. QED

Continuing the proof of the Stratification Theorem, what is left to do is link the algebras \mathcal{L} and \mathcal{M} to diagram (61). We first need a proof-theoretical result stating that on formulas in \mathcal{L}_n , the notions of n -derivability and derivability coincide.

Proposition 8.15 *Let a and b be two formula in \mathcal{L}_n .*

1. $a \equiv_n b$ iff $a \equiv_m b$ for some $m \in \omega$;
2. $a \equiv_n b$ iff $a \equiv_{\mathbb{M}} b$.

Proof. Part 1 of the proposition is a direct consequence of diagram (61) commuting. Concerning the second part, the left-to-right direction can be proved by a straightforward induction on n . For the opposite direction ‘ \Leftarrow ’, it suffices to establish that for two formulas $a, b \in \mathcal{L}_n$ we have

$$\mathcal{D} : \vdash_{\mathbb{M}} a \preceq b \text{ implies } a \sqsubseteq_n b, \tag{72}$$

where we use $a \sqsubseteq_n b$ to denote that $a \equiv_n a \wedge b$. The proof of (72) is by induction on the complexity of the derivation \mathcal{D} .

We confine ourselves to the most difficult case of the inductive step, namely where the last applied rule in \mathcal{D} is the cut rule; that is, we assume \mathcal{D} to be of the form

$$\mathcal{D} : \quad \text{cut} \frac{\frac{\mathcal{D}_1}{a \preceq c} \quad \frac{\mathcal{D}_2}{c \preceq b}}{a \preceq b}$$

(This case is the most difficult one since here we may not assume c to be in \mathcal{L}_n .) Let m be such that $c \in \mathcal{L}_m$, and put $k := \max(m, n)$. Then inductively, we have $a \sqsubseteq_k c$ and $c \sqsubseteq_k b$, from which we easily obtain that $a \sqsubseteq_k b$. But then by the first part of the Proposition, we see that $a \sqsubseteq_n b$, as required. QED

Proposition 8.16 *The relation $\equiv_{\mathbb{M}} \subseteq \mathcal{L} \times \mathcal{L}$ is the kernel of the unique \mathbf{Boole}_∇ -quotient map from \mathcal{L} to $V\mathcal{M}$.*

Proof. Define the map $q : \mathcal{L} \rightarrow \mathbb{M}^\omega \mathcal{L}$ as follows. Given a formula $a \in \mathcal{L}$, there is some $n \in \omega$ such that $a \in \mathcal{L}_n$. Now define

$$q(a) := i_n q_n(a)$$

This is well-defined by the fact that diagram (61) commutes and we have $\ker(q) = \equiv_{\mathbb{M}}$ by Proposition 8.15.

Then by initiality of \mathcal{L} in Boole_∇ it suffices to prove that q is an algebraic homomorphism. For the Boolean connectives/operators this is straightforward, and so we leave this as an exercise for the reader. For the ∇ modality we need to prove that the following diagram commutes:

$$\begin{array}{ccc} T_\omega \mathcal{L} & \xrightarrow{\nabla^\mathcal{L}} & \mathcal{L} \\ Tq \downarrow & & \downarrow q \\ T_\omega \text{UM}^\omega \mathfrak{D} & \xrightarrow{\nabla^{V\mathcal{M}}} & \text{UM}^\omega \mathfrak{D} \end{array} \quad (73)$$

In order to prove this, take an arbitrary element $\alpha \in T_\omega(\mathcal{L})$. Without loss of generality, assume that $\alpha \in T_\omega(\mathcal{L}_n)$, so that $\nabla\alpha \in \mathcal{L}_{n+1}$. Then by definition of q , we have

$$(q \circ \nabla^\mathcal{L})(\alpha) = q(\nabla\alpha) = i_{n+1}q_{n+1}(\nabla\alpha). \quad (74)$$

Computing $(\nabla^{V\mathcal{M}} \circ Tq)(\alpha)$, we first calculate

$$(Tq)(\alpha) = Ti_n((Tq_n)(\alpha)),$$

where $(Tq_n)(\alpha)$ belongs to the set $T_\omega \text{UM}^n \mathfrak{D}$. Now we claim that for all $\beta \in T_\omega \text{UM}^n \mathfrak{D}$:

$$\nabla^{V\mathcal{M}}(Ti_n)\beta = i_{n+1}\rho_{\text{M}^n \mathfrak{D}}(\nabla\beta), \quad (75)$$

with $\rho_{\text{M}^n \mathfrak{D}}$ as in Definition 7.15. To see this, consider the following calculation:

$$\begin{aligned} \nabla^{V\mathcal{M}}(Ti_n)\beta &= j_\omega^{-1}(\rho_{\text{M}^\omega \mathfrak{D}}(\nabla(Ti_n)(\beta))) && \text{(Remark 7.25)} \\ &= j_\omega^{-1}(\rho_{\text{M}^\omega \mathfrak{D}}((T_\omega^\nabla i_n)(\nabla\beta))) && \text{(definition of } T_\omega^\nabla) \\ &= j_\omega^{-1}(\rho_{\text{M}^\omega \mathfrak{D}}((\mathcal{L}_1 U i_n)(\nabla\beta))) && (\mathcal{L}_1 U i_n \upharpoonright_{T_\omega^\nabla \text{UM}^n \mathfrak{D}} = T_\omega^\nabla i_n) \\ &= j_\omega^{-1}((\mathbb{M}i_n \circ \rho_{\text{M}^n \mathfrak{D}})(\nabla\beta)) && \text{(naturality of } \rho) \\ &= i_{n+1}\rho_{\text{M}^n \mathfrak{D}}(\nabla\beta) && (\dagger) \end{aligned}$$

where the last equality (\dagger) follows by Proposition 7.23(5).

And so we obtain that

$$(\nabla^{V\mathcal{M}} \circ Tq)(\alpha) = i_{n+1}\rho_{\text{M}^n \mathfrak{D}}(\nabla(Tq_n)(\alpha)) \quad (76)$$

Thus in order to prove the commutativity of (73), by (74) and (76) it suffices to prove that

$$q_{n+1}(\nabla\alpha) = \rho_{\text{M}^n \mathfrak{D}}(\nabla(Tq_n)(\alpha)). \quad (77)$$

But this is precisely the content of Proposition 8.14. QED

We can now prove the Stratification Theorem.

Proof of Theorem 8.9. Given the Propositions 8.13, 8.15 and 8.16, all that is left to do is prove that the following diagram commutes for each $n \in \omega$:

$$\begin{array}{ccc} \mathcal{L}_n & \xrightarrow{d_n} & \mathcal{L} \\ q_n \downarrow & & \downarrow m \\ \text{M}^n \mathfrak{D} & \xrightarrow{i_n} & V\mathcal{M} \end{array} \quad (78)$$

We already saw in the proof of Proposition 8.16 that the map $q : \mathcal{L} \rightarrow \mathbb{M}^\omega \mathfrak{2}$, defined by putting, for $a \in \mathcal{L}_n$,

$$q(a) := i_n(q_n(a)),$$

is the unique Moss homomorphism from \mathcal{L} to $V\mathcal{M}$; in other words, this map q coincides with m . Reformulating this in terms that explicitize the role of the inclusion map $d_n : \mathcal{L}_n \hookrightarrow \mathcal{L}$, we obtain that $m(d_n(a)) = q(d_n(a)) = i_n(q_n(a))$. In other words, the diagram (78) commutes indeed. QED

As a corollary we obtain that the algebra $V\mathcal{M}$ is the initial algebra in the class of Moss algebras that satisfy the nabla-equations. This means that we may see \mathcal{M} as the *Lindenbaum-Tarski* algebra of our logic.

Corollary 8.17 *Let $\mathbb{B} = \langle B, \neg^{\mathbb{B}}, \wedge^{\mathbb{B}}, \vee^{\mathbb{B}}, \nabla^{\mathbb{B}} \rangle$ be a Moss algebra such that \mathbb{B} validates every instance of the axioms $(\nabla 1) - (\nabla 3)$. Then there is a unique morphism $mng_{\mathbb{B}}^* : V\mathcal{M} \rightarrow \mathbb{B}$ through which the meaning function $mng_{\mathbb{B}}$ factors:*

$$\begin{array}{ccc} \mathcal{L} & \xrightarrow{m} & V\mathcal{M} \\ & \searrow mng_{\mathbb{B}} & \downarrow mng_{\mathbb{B}}^* \\ & & \mathbb{B} \end{array}$$

Proof. An arbitrary element of (the carrier of) $V\mathcal{M}$ is of the form $m(a)$ for some formula $a \in \mathcal{L}$. We leave it as an exercise for the reader to verify that the following map

$$mng_{\mathbb{B}}^*(m(a)) := mng_{\mathbb{B}}(a)$$

is well-defined and has the right properties. QED

Remark 8.18 In fact, we can show that the functor V constitutes an *isomorphism* between the category $\text{Coalg}_{\text{BA}}(\mathbb{M})$ and the variety of Moss algebras validating the nabla axioms. We omit the details of this proof.

8.3 Proof of soundness and completeness

We are almost ready to prove our main result. What is left to do is link the final T -sequence to the initial \mathbb{M} -sequence. Recall that the elements of $T^n 1$ intuitively correspond to the n -behaviors associated with T , and that \mathcal{M} , the initial \mathbb{M} -algebra, is the colimit of the initial sequence $\langle \mathbb{M}^n \mathfrak{2}, j_n \rangle_{n < \omega}$, where elements of $\mathbb{M}^n \mathfrak{2}$ correspond to (equivalence classes of) formulas of depth n .

Definition 8.19 We define the sequence of maps $s_n : \mathbb{M}^n \mathfrak{2} \rightarrow \check{\mathbb{P}}T^n 1$ as follows. The map $s_0 : \mathfrak{2} \rightarrow \check{\mathbb{P}}1$ is given by initiality (and is actually the identity). For the definition of s_{n+1} , recall from Defintion 7.29 that $\delta_{T^n 1} : \mathbb{M}\check{\mathbb{P}}T^n 1 \rightarrow \check{\mathbb{P}}T^{n+1} 1$, and assume inductively that $s_n : \mathbb{M}^n \mathfrak{2} \rightarrow \check{\mathbb{P}}T^n 1$ has been defined, so that $\mathbb{M}s_n : \mathbb{M}^{n+1} \mathfrak{2} \rightarrow \mathbb{M}\check{\mathbb{P}}T^n 1$. Composing these two maps, we obtain $s_{n+1} := \delta_{T^n 1} \circ \mathbb{M}(s_n)$. \triangleleft

Intuitively, the reader may think of the map s_n as providing semantics of elements of $\mathbb{M}^n \mathfrak{2}$. This can be made more precise by proving that the following diagram commutes:

$$\begin{array}{ccc} \mathcal{L}_n & \xrightarrow{q_n} & \mathbb{M}^n \mathfrak{2} \\ & \searrow mng_n & \downarrow s_n \\ & & \check{\mathbb{P}}T^n 1 \end{array}$$

Here q_n is the quotient map under n -step derivability of Theorem 8.9 and mng_n is the n -step meaning function of Definition 8.4.

From this perspective, the following proposition states that the semantics of a formula with respect to the final sequence is independent of the particular approximant we choose.

Proposition 8.20 *The following diagram commutes:*

$$\begin{array}{ccccc}
\check{\mathbb{P}}1 & \xrightarrow{\check{\mathbb{P}}h_0} & \dots & & \check{\mathbb{P}}T^n 1 & \xrightarrow{\check{\mathbb{P}}h_n} & \check{\mathbb{P}}T^{n+1} 1 & \dots & \dots \\
\uparrow s_0 & & & & \uparrow s_n & & \uparrow s_{n+1} & & \\
\mathfrak{D} & \xrightarrow{j_0} & \dots & & \mathbb{M}^n \mathfrak{D} & \xrightarrow{j_n} & \mathbb{M}^{n+1} \mathfrak{D} & \dots & \dots
\end{array} \tag{79}$$

In addition, each map s_n is injective.

Proof. In order to show that diagram (79) commutes, we will prove that

$$s_{n+1} \circ j_n = \check{\mathbb{P}}h_n \circ s_n$$

for all $n \in \omega$. The proof is by induction on n . The base case $s_1 \circ j_0 = \check{\mathbb{P}}h_0 \circ s_0$ is a consequence of the fact that \mathfrak{D} is the initial object in BA. For the inductive case, where $n = k + 1$ for some $k \in \omega$, we reason as follows:

$$\begin{aligned}
s_{k+2} \circ j_{k+1} &= \delta_{T^{k+1} 1} \circ \mathbb{M}(s_{k+1}) \circ \mathbb{M}(j_k) && \text{(unfolding definitions)} \\
&= \delta_{T^{k+1} 1} \circ \mathbb{M}(s_{k+1} \circ j_k) && \text{(functoriality of } \mathbb{M} \text{)} \\
&= \delta_{T^{k+1} 1} \circ \mathbb{M}(\check{\mathbb{P}}h_k \circ s_k) && \text{(inductive hypothesis)} \\
&= \check{\mathbb{P}}Th_k \circ \delta_{T^k 1} \circ \mathbb{M}(s_k) && \text{(naturality of } \delta \text{)} \\
&= \check{\mathbb{P}}h_{k+1} \circ s_{k+1} && \text{(definition } s_{k+1} \text{)}
\end{aligned}$$

Since δ is injective (Proposition 7.30) and \mathbb{M} preserves embeddings (Proposition 7.19), a straightforward inductive proof shows that all s_n , $n \in \omega$, are injective. QED

We are now going to demonstrate that the coalgebraic semantics and the semantics via the final sequence coincide.

Proposition 8.21 *For a given coalgebra $\mathbb{X} = \langle X, \xi \rangle$ and any formula $a \in \mathcal{L}_n$, the following holds:*

$$mng_{\mathbb{X}}(a) = \xi_n^{-1}(s_n(q_n(a))), \quad \text{for all } a \in \mathcal{L}_n \text{ and } n \in \omega. \tag{80}$$

Proof. First note that $\check{\mathbb{P}}X$ together with the maps $\check{\mathbb{P}}\xi_n \circ s_n = \xi_n^{-1} \circ s_n$ form a cocone over the initial sequence of \mathbb{M} . Therefore there is a mediating arrow

$$mng_{\mathbb{X}}^* : \mathbb{M}^\omega \mathfrak{D} \rightarrow \check{\mathbb{P}}X$$

from the carrier of the initial \mathbb{M} -algebra \mathcal{M} to $\check{\mathbb{P}}X$ with the property that $mng_{\mathbb{X}}^* \circ i_n = \xi_n^{-1} \circ s_n$.

We claim that

$$\text{the map } mng_{\mathbb{X}}^* \text{ is an } \mathbb{M}\text{-algebra morphism from } \mathcal{M} \text{ to } \mathbb{X}^*. \tag{81}$$

In order to prove (81), observe that by Proposition 7.23, for all $n \in \omega$ we have

$$j_\omega \circ i_{n+1} = \mathbb{M}(i_n), \tag{82}$$

where $j_\omega : \mathbb{M}^\omega \mathfrak{D} \rightarrow \mathbb{M}\mathbb{M}^\omega \mathfrak{D}$ is the inverse of the algebra structure map $\heartsuit^{\mathcal{M}}$ of the initial \mathbb{M} -algebra. In order to prove the claim it suffices to show that the following diagram commutes

$$\begin{array}{ccc}
\mathbb{M}\mathbb{M}^\omega \mathfrak{D} & \xrightarrow{\mathbb{M}mng_{\mathbb{X}}^*} & \mathbb{M}\check{\mathbb{P}}X \\
\uparrow j_\omega & & \downarrow \delta_X \\
\mathbb{M}^\omega \mathfrak{D} & \xrightarrow{mng_{\mathbb{X}}^*} & \check{\mathbb{P}}TX \\
& & \downarrow \check{\mathbb{P}}\xi \\
& & \check{\mathbb{P}}X
\end{array}$$

We prove that the diagram commutes by showing that $f := \check{\mathbb{P}}\xi \circ \delta_X \circ \mathbb{M}(mng_{\mathbb{X}}^*) \circ j_\omega$ is a mediating arrow from $(\mathbb{M}^\omega \mathfrak{D}, \{i_n\}_{n \in \omega})$ to $(\check{\mathbb{P}}X, \{\check{\mathbb{P}}\xi_n \circ s_n\}_{n \in \omega})$. Therefore f has to be equal to $mng_{\mathbb{X}}^*$ by the universal property of the colimit $(\mathbb{M}^\omega \mathfrak{D}, \{i_n\}_{n \in \omega})$. We show that f has the claimed property by proving that for all $n \in \omega$ we have

$$\check{\mathbb{P}}(\xi_n) \circ s_n = f \circ i_n \quad (83)$$

For $n = 0$ the equation holds by initiality of \mathfrak{D} . Furthermore for an arbitrary $n \geq 0$ we have

$$\begin{aligned}
\check{\mathbb{P}}(\xi_{n+1}) \circ s_{n+1} &= \check{\mathbb{P}}(T\xi_n \circ \xi) \circ \delta_{T^{n+1}} \circ \mathbb{M}s_n && \text{(definition of } \xi_{n+1} \text{ and of } s_{n+1}) \\
&= \check{\mathbb{P}}(\xi) \circ \check{\mathbb{P}}(T\xi_n) \circ \delta_{T^{n+1}} \circ \mathbb{M}s_n && \text{(functoriality of } \check{\mathbb{P}}) \\
&= \check{\mathbb{P}}(\xi) \circ \delta_X \circ \mathbb{M}\check{\mathbb{P}}\xi_n \circ \mathbb{M}s_n && \text{(naturality of } \delta) \\
&= \check{\mathbb{P}}(\xi) \circ \delta_X \circ \mathbb{M}(\check{\mathbb{P}}\xi_n \circ s_n) && \text{(functoriality of } \mathbb{M}) \\
&= \check{\mathbb{P}}(\xi) \circ \delta_X \circ \mathbb{M}(mng_{\mathbb{X}}^* \circ i_n) && \text{(} mng_{\mathbb{X}}^* \text{ mediating arrow)} \\
&= \check{\mathbb{P}}(\xi) \circ \delta_X \circ \mathbb{M}mng_{\mathbb{X}}^* \circ \mathbb{M}i_n && \text{(functoriality of } \mathbb{M}) \\
&= \check{\mathbb{P}}(\xi) \circ \delta_X \circ \mathbb{M}mng_{\mathbb{X}}^* \circ j_\omega \circ i_{n+1}. && \text{(equation (82))}
\end{aligned}$$

Therefore equation (83) holds for all n , which finishes the proof of (81).

From this it follows that $Vmng^* : V\mathcal{M} \rightarrow V\mathbb{X}^*$ is a Moss algebra homomorphism. Recalling from Proposition 7.32 that $V\mathbb{X}^* = \mathbb{X}^+$, we obtain by initiality of \mathcal{L} as a Moss algebra, that $Vmng_{\mathbb{X}}^* \circ m = mng_{\mathbb{X}}$. Here $mng_{\mathbb{X}} : \mathcal{L} \rightarrow \mathbb{X}^+$ is the unique Moss algebra homomorphism that maps an element of \mathcal{L} to its semantics in \mathbb{X}^+ , and $m := mng_{V\mathcal{M}}$ is the unique homomorphism $m : \mathcal{L} \rightarrow V\mathcal{M}$ in the category of Moss algebras. But then by the Axiomatic Stratification Theorem 8.9, for all $n \in \omega$ and all formulas $a \in \mathcal{L}_n$ we have $mng_{\mathbb{X}}(a) = mng_{\mathbb{X}}^*(m(a)) = mng_{\mathbb{X}}^*(i_n(q_n(a))) = \check{\mathbb{P}}\xi_n \circ s_n(q_n(a))$, where the last identity holds by the definition of $mng_{\mathbb{X}}^*$ as a mediating arrow. This shows that (80) holds, and finishes the proof of the claim. QED

On the basis of the results obtained so far, the proof of our soundness and completeness results is now more or less immediate.

Proof of Theorem 6.10. Let a and b be two formulas in \mathcal{L} . Fix a natural number n such that $a, b \in \mathcal{L}_n$. Recall that $\mathbb{F}_n = \langle T^n 1, T^n g \rangle$ denotes the ‘ n -step coalgebra’ defined in Definition 8.3.

Now consider the following sequence of equivalences:

$$\begin{aligned}
a \sqsubseteq_{\mathbf{M}} b &\iff q_n(a) \subseteq q_n(b) && \text{(Axiomatic Stratification Theorem 8.9)} \\
&\iff s_n q_n(a) \subseteq s_n q_n(b) && \text{(injectivity of } s_n) \\
&\iff (\check{\mathbb{P}}(T^n g)_n)(s_n q_n(a)) \subseteq (\check{\mathbb{P}}(T^n g)_n)(s_n q_n(b)) && \text{(equation (59))} \\
&\iff \text{mng}_{\mathbb{F}_n}(a) \subseteq \text{mng}_{\mathbb{F}_n}(b) && \text{(Proposition 8.21)} \\
&\iff a \models_T b && \text{(Semantic Stratification Theorem 8.7)}
\end{aligned}$$

From this the Theorem is immediate.

QED

9 Conclusions

Summary of results Obviously, as the main contributions of this paper we see the definition of the *derivation system* \mathbf{M} for the finitary version of Moss' coalgebraic logic, the result stating that \mathbf{M} provides a *sound and complete* axiomatization for the collection of coalgebraically valid inequalities, and the fact that all of our definitions, results and our proofs are completely *uniform* in the coalgebraic type functor T

Our proof of the soundness and completeness theorem is rather elaborate and technical, but we believe that the effort has been worth the while, and that on the way we have identified some new concepts and obtained some auxiliary results that may be of independent interest. Of these we list the following:

1. a survey of the properties of the notion \overline{T} of relation lifting, induced by an arbitrary but fixed set functor T (section 3);
2. the introduction in Definition 4.8 of the category Pres of Boolean algebra presentations, and the establishment in Theorem 4.13 of an adjunction between Pres and the category BA of Boolean algebras;
3. the introduction in section 7.3 of the functor $\mathbb{M} : \text{BA} \rightarrow \text{BA}$, and the results in Proposition 7.19 that \mathbb{M} is finitary and preserves embeddings, and in Theorem 7.28 that it preserves atomicity of Boolean algebras.
4. the stratification of our logic, both semantically (Theorem 8.7) and syntactically (Theorem 8.9);
5. the identification, in Corollary 8.17, of the initial \mathbb{M} -algebra \mathcal{M} , through the functor V , as the Lindenbaum-Tarski algebra of our logic.

Related and ongoing work As mentioned in the introduction, this paper replaces, extends and partly corrects an earlier version [34]. Since the publication of the latter paper, and the preparation of the current manuscript there have been a number of developments in the area of Moss' logic that we would like to mention here. First of all, based on our one-step soundness and completeness results, Bergfeld gave a more direct version of our completeness proof in his MSc thesis [12]; as a corollary he established a strong completeness theorem for Moss' logic (modulo some restrictions on the functor T). Second, Bilková, Palmigiano & Venema generalized their earlier result on the power set nabla [13] to the general case of a standard, weak pullback preserving functor T : in [14] they provide a sound, complete, and cut-free proof system for (the finitary version of) Moss' coalgebraic logic. Systematically using Stone duality, Kurz & Leal [38] make a detailed comparison between Moss' approach towards coalgebraic logic, and the one based on associating standard modalities with predicate liftings; their main contribution is a new coalgebraic logic combining features of both approaches. Venema, Vickers

& Vosmaer [56] study a variant of the derivation system \mathbf{M} in the setting of geometric logic; their main contribution is to generalize Johnstone’s power construction on locales, to a functor V_T , parametrically defined in a set functor T , on the category of locales. Finally, Bílková, Velebil & Venema [15] prove that on the (semantically defined) Lindenbaum-Tarski algebra of our logic, the nabla modality has the interesting order-theoretic property of being a so-called \mathcal{O} -adjoint.

Future research We finish with mentioning some directions for future research. To start with, in this paper we have studied the nabla operator in the setting of the diagram (5), which is a particular instantiation of the general Stone duality diagram

$$L \left(\text{Alg} \begin{array}{c} \xleftarrow{P} \\ \xrightarrow{S} \end{array} \text{Sp}^{\text{op}} \right) T \quad (84)$$

where Alg denotes a category of algebras representing the base logic, Sp is a category of spaces representing the semantics of the logic, T is the coalgebra functor representing all one-step behaviours, and L represents the one-step version of the coalgebraic modal logic. Given the flexibility of the Stone duality approach we believe it to be of interest to consider more instances of the diagram (84) where L is some version of our nabla logic. Of particular interest are the cases where for Alg we take the variety of distributive lattices, because this could clarify the role of the negation in our setting.

Second, a clear drawback of the current nabla-based approach towards coalgebraic logic is the restriction to functors that preserve weak pullbacks. It would therefore be interesting to see whether this restriction can be removed. A first step in this direction has been made by Santocanale & Venema [48], who introduce a nabla-based version of monotone modal logic, a variant of basic modal logic that is naturally interpreted in coalgebras for the monotone neighborhood functor of Example 2.2 — a functor that does not preserve weak pullbacks.

Finally, in the introduction we mentioned that the work of Janin & Walukiewicz [30] on automata theory and modal fixpoint logics is an independent source for the introduction of the cover modality ∇_P as a primitive modality. Since ∇_P also plays a fundamental role in Walukiewicz’ completeness result for the modal μ -calculus [57], this naturally raises the question whether we can extend our completeness result to the setting with fixpoint operators.

References

- [1] P. Aczel. *Non-well-founded sets*, volume 14 of *CSLI Lecture Notes*. CSLI, 1988.
- [2] P. Aczel and N. Mendler. A final coalgebra theorem. In D.H. Pitt, A. Poigné, and D.E. Rydeheard, editors, *Category Theory and Computer Science*, volume 389 of *LNCS*. Springer, 1989.
- [3] J. Adámek. A logic of coequations. In L. Ong, editor, *Computer Science Logic, 19th International Workshop (CSL 2005), Proceedings*, volume 3634 of *LNCS*, pages 70–86. Springer, 2005.
- [4] J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theoretical Computer Science*, 150:57–75, 1995.
- [5] J. Adámek and S. Milius, editors. *Coalgebraic Methods in Computer Science (CMCS’04)*, volume 106 of *ENTCS*, 2004.
- [6] J. Adámek and J. Rosický. *Locally Presentable and Accessible Categories*, volume 189 of *LMS Lecture Note Series*. Cambridge University Press, 1994.

- [7] J. Adámek and V. Trnková. *Automata and Algebras in Categories*. Kluwer Academic Publishers, 1990.
- [8] C. Areces and R. Goldblatt, editors. *Advances in Modal Logic 7*. College Publications, 2008.
- [9] A. Baltag. A logic for coalgebraic simulation. In Reichel [44], pages 41–60.
- [10] M. Barr. Relational algebras. In *Reports of the Midwest Category Seminar IV*, volume 137 of *Lecture Notes in Mathematics*, pages 39–55, 1970.
- [11] J. Beck. Distributive Laws. In *Seminar on Triples and Categorical Homology Theory*, volume 80 of *Lecture Notes in Mathematics*, pages 119–140. Springer, 1969.
- [12] J. Bergfeld. Moss’s coalgebraic logic: Examples and completeness results. Master’s thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2009.
- [13] M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality. In Areces and Goldblatt [8], pages 1–21.
- [14] M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for Moss’ coalgebraic logic. In preparation, 2010.
- [15] M. Bílková, J. Velebil, and Y. Venema. On monotone modalities and adjointness. *Mathematical Structures in Computer Science*, to appear.
- [16] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [17] M. Bonsangue and A. Kurz. Duality for logics of transition systems. In V. Sassone, editor, *Foundations of Software Science and Computational Structures, 8th International Conference (FoSSaCS’05)*, volume 3441 of *LNCS*, pages 455–469. Springer, 2005.
- [18] A. Carboni, G. Kelly, and R. Wood. A 2-categorical approach to change of base and geometric morphisms I. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 31(4):47–95, 1991.
- [19] C. Cîrstea. On logics for coalgebraic simulation. In Adámek and Milius [5], pages 63–90.
- [20] G. D’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus: Interpolation, Lyndon and Łoś-Tarski. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [21] K. Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16:229–234, 1975.
- [22] H.P. Gumm. Functors for Coalgebras. *Algebra Universalis*, 45(2-3):135–147, 2001.
- [23] H.P. Gumm. From T-coalgebras to filter structures and transition systems. In J. L. Fiadeiro, editor, *Algebra and Coalgebra in Computer Science (CALCO 2005)*, volume 3629 of *LNCS*, pages 194–212, 2005.
- [24] C. Hermida. A categorical outlook on relational modalities and simulations. *Inf. Comput.*, 209(12):1505–1517, 2011.
- [25] C. Hermida and B. Jacobs. Structural induction and coinduction in a fibrational setting. *Information and Computation*, 145(2):107–152, 1998.
- [26] W. Hodges. *Model Theory*. Cambridge University Press, 1993.
- [27] J. Hughes and B. Jacobs. Simulations in coalgebra. *Theoretical Computer Science*, 327, 2004.

- [28] B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoretical Informatics and Applications*, 35(1):31–59, 2001.
- [29] B. Jacobs. Trace semantics for coalgebras. In Adámek and Milius [5], pages 167–184.
- [30] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In J. Wierdmann and P. Hájek, editors, *Mathematical Foundations of Computer Science 1995, 20th International Symposium (MFCS'95)*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.
- [31] C. Kissig and Y. Venema. Complementation of coalgebra automata. In Kurz et al. [39], pages 81–96.
- [32] C. Kupke. *Finitary coalgebraic logics*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2006.
- [33] C. Kupke, A. Kurz, and Y. Venema. Stone coalgebras. *Theoretical Computer Science*, 327:109–134, 2004.
- [34] C. Kupke, A. Kurz, and Y. Venema. Completeness of the finitary Moss logic. In Areces and Goldblatt [8], pages 193–217.
- [35] C. Kupke and Y. Venema. Coalgebraic automata theory: Basic results. *Logical Methods in Computer Science*, 4(4), 2008.
- [36] A. Kurz. Specifying coalgebras with modal logic. *Theoretical Computer Science*, 260:119–138, 2001.
- [37] A. Kurz. Coalgebras and their logics. *SIGACT News*, 37:57–77, 2006.
- [38] A. Kurz and R. Leal. Modalities in the stone age: a comparison of coalgebraic logics. *Theoretical Computer Science*, to appear.
- [39] A. Kurz, M. Lenisa, and A. Tarlecki, editors. *Algebra and Coalgebra in Computer Science (CALCO 2009)*, volume 5728 of *LNCS*, 2009.
- [40] S. Mac Lane. *Categories for the working mathematician*, volume 2 of *Graduate Texts in Mathematics*. Springer, 1998.
- [41] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).
- [42] A. Palmigiano and Y. Venema. Nabla algebras and Chu spaces. In T. Mossakowski, U. Montanari, and M. Haverdaen, editors, *Algebra and Coalgebra in Computer Science (CALCO 2007)*, volume 4624 of *LNCS*, pages 394–408, 2007.
- [43] D. Pattinson. Coalgebraic modal logic: soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.
- [44] H. Reichel, editor. *Coalgebraic Methods in Computer Science (CMCS'00)*, volume 33 of *ENTCS*, 2000.
- [45] M. Röbiger. Coalgebras and modal logic. In Reichel [44], pages 299–320.
- [46] J. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.

- [47] L. Santocanale and Y. Venema. Completeness for flat modal fixpoint logics. *Annals of Pure and Applied Logic*, 162:55–82, 2010.
- [48] L. Santocanale and Y. Venema. Uniform interpolation for monotone modal logic. In L. Beklemishev, V. Goranko, and V. Shehtman, editors, *Advances in Modal Logic 8*. College Publications, 2010.
- [49] Lutz Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. *Theoretical Computer Science*, 390:230–247, 2008.
- [50] R. Sikorski. A theorem on extensions of homomorphisms. *Annals of the Polish Mathematical Society*, 21, 1948.
- [51] S. Staton. Relating coalgebraic notions of bisimulation with applications to name-passing process calculi. In Kurz et al. [39], pages 191–205.
- [52] R. Street. The formal theory of monads. *Journal of Pure and Applied Algebra*, 2(2):149–168, 1972.
- [53] A. Thijs. *Simulation and Fixpoint Semantics*. PhD thesis, Rijksuniversiteit Groningen, 1996.
- [54] V. Trnková. General theory of relational automata. *Fundamenta Informaticae*, 3(2):189–234, 1980.
- [55] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.
- [56] Y. Venema, S. Vickers, and J. Vosmaer. Powerlocales via relation lifting. In preparation, 2010.
- [57] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000. (An earlier version of the paper appeared in the Proceedings of LICS 1995).

10 Appendix: overview of notation

Since this paper features a multitude of categories, functors and natural transformations, for the reader’s convenience we list these in the tables below.

Categories	
BA	section 2.1
Boole	Definition 4.2
Pres	Definition 4.8
Set, Rel	section 2.1
Boole ∇	Definition 5.6

Natural Transformations	
$Base^T : T_\omega \rightarrow P_\omega$	Definition 2.12
$\lambda^T : TP \rightarrow PT$	Definition 3.12
$\rho : \mathcal{L}_0 U \rightarrow \mathbb{M}$	Definition 7.15
$\delta : \mathbb{M}\check{P} \rightarrow \check{P}T$	Definition 7.29

Functors	
$B : \text{Pres} \rightarrow \text{BA}$	Definition 4.5, 4.9
$C : \text{BA} \rightarrow \text{Pres}$	Definition 4.9
$\mathbb{F} : \text{Set} \rightarrow \text{Boole}$	page 22
$Id, B_\omega, D_\omega : \text{Set} \rightarrow \text{Set}$	Example 2.2
$A_M : \text{Set} \rightarrow \text{Set}$	Definition 5.6
$M : \text{Pres} \rightarrow \text{Pres}$	Definition 7.9
$\mathbb{M} : \text{BA} \rightarrow \text{BA}$	Definition 7.13
$\mathcal{L}_0 : \text{Set} \rightarrow \text{Set}$	Definition 4.1 & (14)
$\mathcal{L}_1 : \text{Set} \rightarrow \text{Boole/Set}$	Definition 5.8
$P, P_\omega : \text{Set} \rightarrow \text{Set}$	section 2.1
$\check{P} : \text{Set} \rightarrow \text{Set}^{op}$	section 2.1
$\check{\mathbb{P}} : \text{Set} \rightarrow \text{BA}^{op}$	Definition 4.4
$T_\omega : \text{Set} \rightarrow \text{Set}$	page 10
$T_\omega^\nabla : \text{Set} \rightarrow \text{Set}$	Definition 5.8
$U : \text{Boole} \rightarrow \text{Set}$	page 22
$V : \text{Alg}_{\text{BA}}(\mathbb{M}) \rightarrow \text{Alg}_{\text{Set}}(A_M)$	Definition 7.20