Math. Struct. in Comp. Science (2011), *vol.* 21, *pp.* 383–416.
doi:10.1017/S0960129510000514

On monotone modalities and adjointness

MARTA BÍLKOVÁ[†], JIŘÍ VELEBIL[‡] and YDE VENEMA[§]

[†]*Department of Logic, Charles University in Prague*
E-mail: marta.bilkova@ff.cuni.cz
[‡]*Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic*
E-mail: velebil@math.feld.cvut.cz
[§]*Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands*
E-mail: Y.Venema@uva.nl

Received 26 April 2010; revised 20 September 2010

We fix a logical connection ($Stone \dashv Pred : \mathsf{Set}^{op} \longrightarrow \mathsf{BA}$ given by 2 as a schizophrenic object) and study coalgebraic modal logic that is induced by a functor $T : \mathsf{Set} \longrightarrow \mathsf{Set}$ that is finitary and standard and preserves weak pullbacks and finite sets. We prove that for any such T, the cover modality nabla is a left (and its dual delta is a right) adjoint relative to $\mathbb{P}_\omega$. We then consider monotone unary modalities arising from the logical connection and show that they all are left (or right) adjoints relative to $\mathbb{P}_\omega$.

1. Introduction

We are going to study universal properties of modalities in coalgebraic modal logic, considered as monotone operations on the set of modal formulas, preordered by the semantic consequence relation.

In coalgebraic logic, there are essentially two approaches to modalities: modalities are given by *predicate liftings*, which can be viewed as the modalities described in Section 2.2, (Pattinson 2003), or, in the case of set-coalgebras, by *cover modalities* that identify the modalities with the coalgebra functor (Moss 1999).

In either case, we are naturally interested in the adjointness properties of the modalities that would entail their 'nice' behaviour, for example, with respect to suprema/infima in the consequence preorder. However, it can be seen almost immediately that a monotone modality can rarely be a left or right adjoint, since, typically, a modality preserves only some and not all suprema/infima.

In this paper, we show that all cover modalities and all monotone unary modalities do indeed enjoy adjointness properties in a *weaker* sense: the desired left/right adjoints do exist if we require the adjointness property to hold only *relative* to the doctrine $\mathbb{P}_\omega$

[†] The work of Marta Bílková is supported by the grant P202/11/P304 of the Grant Agency of the Czech Republic, and is part of the research plan MSM 0021620839 of the Ministry of Education of the Czech Republic.
[‡] Jiří Velebil acknowledges the support of the grant MSM 6840770014 of the Ministry of Education of the Czech Republic.
[§] The research of Yde Venema was made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO)

of free join-semilattices – see Definition 2.11 for details. In fact, as we argue below, this weak notion of adjointness is, when proper adjunction is not available, the 'second best' one can hope for in case of finitary languages.

Moreover, such adjointness has a proof-theoretic significance: proper adjointness is closely related to the possibility of formulating a sound and invertible rule for the operator in question. The above weaker adjointness property indicates a possibility of formulating a weakly invertible rule. The rule, read backwards, gives finitely many possible continuations of the proof search – a situation well known in the proof theory of modal logics.

Adjointness properties of modalities only make sense, of course, for monotone modalities. Monotone modalities yield expressive languages for coalgebraic functors preserving weak pullbacks (Kurz and Leal 2009). Furthermore, monotone modalities having an expressive language allow one to add fixpoint operators to the language.

Our propositional setting for coalgebraic modal logic is *classical*, that is, we work over Boolean algebras. It is not hard to see, however, that one can obtain the same results for *positive fragments* of the logics in question, that is, for the case when the propositional part of the logic is given by distributive lattices. We indicate below how such a generalisation can be made.

1.1. *Organisation of the paper*

After recalling the notions required for the rest of the paper in Section 2, we prove in Section 3 that every nabla cover modality is a left adjoint and every delta cover modality is a right adjoint in the weak sense. In this way we generalise the results of Santocanale and Venema (2007) for coalgebras for the finitary powerset functor to a rather wide class of functors. In Section 4 we analyse monotone unary modalities and prove that *every* such modality is a left adjoint in the weak sense. We also prove a structural result: every unary monotone modality is a finite join of unary modalities that are *both* left and right adjoints in the weak sense. For that our coalgebraic behaviour functor needs to fulfil certain side conditions – see Section 2.1.

1.2. *Related work*

We should certainly mention the paper Schröder and Venema (2010), which was written in parallel with our work. While the main purpose of Schröder and Venema (2010) was to generalise the completeness result of Santocanale and Venema (2007) to a more general coalgebraic setting, the notion of O-adjointness also plays a supporting but crucial role[†]. Working in a setting of modalities obtained from predicate liftings, Schröder and Venema isolate a fairly large syntactic fragment of the language, consisting of so-called *admissible* formulas, which they prove to be O-adjoints. We will discuss the relation between the two papers in more detail in Remark 5.1.

[†] The notion of O-adjointness is *exactly* our notion of adjointness relative to $\mathbb{P}_\omega$ – see Definition 2.11.

2. Preliminaries

In this section we gather together the notions we will need in the rest of the paper. Most of the material presented here is standard, and the details can be found in the references.

2.1. *The basic setting of coalgebraic modal logic*

We study the coalgebraic modal logic induced by a functor $T : \mathsf{Set} \longrightarrow \mathsf{Set}$ (thus, we study a logic of T-coalgebras) where T has the following properties:

(1) T is finitary, standard and preserves weak pullbacks.

(2) T preserves finite sets.

The preservation of weak pullbacks is crucial for the whole setting to work – it enables one to pass from the category Set to the category PreOrd of preorders and monotone maps. The concepts we use for defining the semantics live naturally in PreOrd, the basic example being the concept of *relation lifting* used to define the semantics of the nabla modality.

A functor $T : \mathsf{Set} \longrightarrow \mathsf{Set}$ is said to be *standard* if it preserves inclusions and each distinguished point of T is standard. Here, a point $x \in TX$ is *distinguished* if $Tf(x) = Tg(x)$ holds for every pair $f, g : X \longrightarrow Y$. And a point $x \in TX$ is *standard* if $x = T!_X(x_0)$ for some $x_0 \in T\varnothing$ and the unique map $!_X : \varnothing \longrightarrow X$.

We restrict ourselves to *standard* functors (the concept goes back to Trnková (1969)) since they behave well with respect to the preservation of set inclusions and finite intersections. Namely, every standard functor preserves inclusions (by definition) and finite intersections of sets (see, for example, Adámek (1983, Theorem 5.9)).

The requirement to be finitary produces finitary logic (which is not so important *per se* if one is ready to use infinitary languages), but also means that the adjointness result is to be obtained relative to $\mathbb{P}_\omega$ (and it is crucial for this). For the latter result, we also need the requirement that T preserves finite sets (see Example 3.1).

We consider two different, though closely related, ways of defining a modal language for T-coalgebras. The first is based strictly on a *logical connection* between a category of *spaces*, where the coalgebras are studied, and a category of *algebras*, where the logic is studied, induced by a *schizophrenic object* (which represents the external 'truth values'), see, for example, Bonsangue and Kurz (2005) or Pavlović *et al.* (2006). Throughout this paper, we fix spaces to be the category Set of sets and mappings and algebras to be the category BA of Boolean algebras, and the schizophrenic object to be the two-element set (Boolean algebra) 2. However, we want to be as open as possible to possible generalisations of all three ingredients, so we will continue to distinguish the three levels. In this paper we will concentrate on unary monotone modalities arising from the logical connection.

The other approach to coalgebraic modal logic is based on Moss' *cover* modality nabla (Moss 1999), which comes naturally with a functor T as its 'arity', and its semantics arises from the notion of *relation lifting* given by T.

For the propositional part of the logics, we fix a countable set of propositional variables $Prop$ and consider the following propositional language $\mathscr{L}_0$:

$$a ::= p \mid \neg a \mid \bigwedge \varphi \mid \bigvee \varphi$$

where $p \in Prop$ and φ is in $P_\omega \mathscr{L}_0$. We will use the abbreviations $\top := \bigwedge \varnothing$ and $\bot := \bigvee \varnothing$.

We will also consider the positive fragment $\mathscr{L}_0^+$ of the language dropping the clause for negation from the definition, and a variant of $\mathscr{L}_0$ with negations restricted to act on propositional variables only. We will consider extensions of the propositional language $\mathscr{L}_0$ and its variants with various modalities. If no confusion arises, we will simply denote all the resulting modal languages by $\mathscr{L}$ to help keep the proofs readable.

In our approach, propositional letters are not a part of the coalgebraic functor. T-coalgebras are therefore the frames, and the semantics is completed by adding a valuation map in the usual manner (in the setting of a logical connection of Section 2.2, given a T coalgebra $c : X \longrightarrow TX$, a valuation is an algebraic homomorphism $\|.\| : \mathscr{L}_0 \longrightarrow Pred\,X$).

2.2. *Modalities arising from a logical connection*

The T-coalgebras considered in this paper live in the category Set, and we assume that the propositional part of the modal logic is classical, therefore we fix a logical connection

$$Stone \dashv Pred \ : \mathsf{Set}^{op} \longrightarrow \mathsf{BA} \tag{2.1}$$

between the category Set and the category BA of Boolean algebras and their homomorphisms that is induced by a two-element schizophrenic object.

This means that, for every set X, the algebra $Pred\,X$ of predicates on X is the Boolean algebra $[X, 2]$ of characteristic functions on X, and for every Boolean algebra A, the set $Stone\,A$ is the set $\mathsf{BA}(A, 2)$ of all ultrafilters on A. Predicates on X can be seen as possible meanings of formulas, while the set $\mathsf{BA}(A, 2)$ of all ultrafilters on A is the set of theories of states.

By a well-known procedure (Bonsangue and Kurz 2005), there is a way of constructing the Boolean algebra $\mathscr{L}$ of formulas of the modal logic corresponding to T. Namely, for every natural number n, the elements $\heartsuit$ of the set $\mathsf{Set}(T(2^n), 2)$ are the *n-ary modalities* of the logic and every formula of the form $\heartsuit(a_0, \dots, a_{n-1})$ has an *interpretation*

$$\|\heartsuit(a_0, \dots, a_{n-1})\|_c : X \longrightarrow 2, \quad x \mapsto 1 \text{ iff } x \Vdash_c \heartsuit(a_0, \dots, a_{n-1})$$

in a coalgebra $c : X \longrightarrow TX$. The interpretation of $\heartsuit(a_0, \dots, a_{n-1})$ is defined inductively as the composite

$$X \xrightarrow{\ c\ } TX \xrightarrow{T\langle \|a_0\|_c, \dots, \|a_{n-1}\|_c \rangle} T(2^n) \xrightarrow{\ \heartsuit\ } 2$$

One can see each modality $\heartsuit : T(2^n) \longrightarrow 2$ as coming with its 'truth table', where objects in $T(2^n)$ are its 'rows' coding the 'type of future' of the state of a coalgebra with respect to the validity of an n-tuple of formulas, each such 'type of future' returning a value in 2. The modalities we consider here are in one-to-one correspondence with *predicate liftings* (Pattinson 2003).

We slightly abuse the notation and use the same ♡ sign to denote the corresponding operator on the free modal algebra $\mathscr{L}$ maping an n-tuple of formulas to a formula.

The logical connection taken together with T automatically induces all possible modalities, amongst which we will concentrate on the unary monotone ones. A modality is *monotone* if the underlying map $\heartsuit : T(2^n) \longrightarrow 2$ is monotone. Here we consider $T(2^n)$ with the preorder *lifted* from that on 2^n – see Section 2.3.

It is not difficult to see that if ♡ is monotone, the corresponding operator on formulas is monotone with respect to the (local) consequence preorder given by the coalgebraic semantics.

We will study the basic unary monotone modalities that arise naturally from the logical connection in Section 4 and are given by members of $T2$ (they can be related to the singleton predicate liftings of Kurz and Leal (2009), only here we consider them to be monotone). We show that they are essentially equivalent to simple nablas or deltas defined in Section 2.3. This is one way that nablas appear in the setting given by the logical connection.

Moreover, the above logical connection (2.1) can be replaced by one having the category DL of distributive lattices instead of Boolean algebras, thereby yielding results on positive fragments of modal languages.

2.3. *Cover modalities nabla and delta*

There is a general single modality ∇, called the *cover modality*, that corresponds directly to the functor T. Rather than n-tuples, ∇ can be applied to 'T-tuples', that is, for every $\alpha \in T\mathscr{L}$ we have a formula $\nabla\alpha$[†]. Thus the language we consider is

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

where p is a propositional variable, φ is in $P_\omega\mathscr{L}$ and α is in $T\mathscr{L}$.

In order to define the semantics of the cover modality, we need to exploit the fact that T preserves weak pullbacks: it is known, see, for example, Moss (1999), that the preservation of weak pullbacks implies that T can be *lifted* to the functor $\overline{T}$ on the category Rel of binary relations in the following sense.

For a binary relation $R \subseteq A \times B$, represented as a span

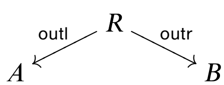

where outl and outr are the left-hand and right-hand projections, respectively, we form the span

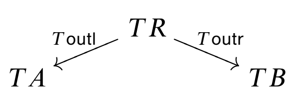

[†] Note that here we are not quite in the setting of the logical connection fixed above, since the functor T is also freely used on the 'algebra' side. However, the idea of T providing the arity as well as the semantics is natural, and everything works well in the setting we are in (based on Set as the basic category).

and denote the corresponding binary relation between TA and TB by $\overline{R}$. Thus, $(a, b) \in \overline{R}$ if and only if there exists a *witness* $w \in TR$ such that

$$T\mathsf{outl}(w) = a$$
$$T\mathsf{outr}(w) = b.$$

The lifted functor $\overline{T}\colon \mathsf{Rel} \longrightarrow \mathsf{Rel}$ then acts as follows:

(1) For every object A of Rel, that is, for every set A, we put $\overline{T}\,A = TA$.
(2) For every morphism $R : A \longrightarrow B$ in Rel, that is, for every binary relation $R \subseteq A \times B$, we define $\overline{T}\,R$ to be the binary relation $\overline{R} \subseteq TA \times TB$, described above.

The lifting of relations allows us to define the semantics $\|\nabla\alpha\|_c : X \longrightarrow 2$ in the coalgebra $c : X \longrightarrow TX$ by putting

$$x \Vdash_c \nabla\alpha \quad \text{iff} \quad c(x)\, \overline{\Vdash_c}\, \alpha$$

Example 2.1. For the powerset functor $P : \mathsf{Set} \longrightarrow \mathsf{Set}$, the lifted binary relation $\overline{R} \subseteq PA \times PB$ of $R \subseteq A \times B$ can be described in the Egli–Milner manner, that is, $(a, b) \in \overline{R}$ holds if and only if the following two conditions hold:

(1) For all $i \in a$ there exists $j \in b$ such that $(i, j) \in R$.
(2) For all $j \in b$ there exists $i \in a$ such that $(i, j) \in R$.

Thus the semantics of the cover modality ∇ for P takes the following form: given a coalgebra $c : X \longrightarrow PX$ (that is, given a Kripke frame c), we define

$$\begin{aligned} x \Vdash_c \nabla\alpha \ \text{ iff } & \text{ for every } x' \in c(x) \text{ there exists } a \in \alpha \text{ such that } x' \Vdash_c a, \text{ and} \\ & \text{ for every } a \in \alpha \text{ there exists } x' \in c(x) \text{ such that } x' \Vdash_c a. \end{aligned}$$

for every $\alpha \in P\mathscr{L}$.

In fact, the lifted functor $\overline{T}$ induces an endofunctor of the category PreOrd of preorders and monotone maps, which we again denote by $\overline{T}$. More precisely, for a preorder $\langle X, \leqslant \rangle$, we put $\overline{T}\langle X, \leqslant \rangle = \langle TX, \overline{\leqslant} \rangle$ (where $\overline{\leqslant}$ is the lifted relation $\leqslant$) and, for a monotone map $f : \langle X, \leqslant \rangle \longrightarrow \langle Y, \leqslant \rangle$, we put $\overline{T}f = Tf$ (this is correct since it is easy to verify that Tf is monotone with respect to the lifted preorders).

Moreover, $\overline{T} : \mathsf{PreOrd} \longrightarrow \mathsf{PreOrd}$ is *locally monotone*, that is, it preserves the preorder on hom-sets.

In order to access the subformulas of $\nabla\alpha$, we need, for each $\alpha \in T\mathscr{L}$, the notion of its *base*. We will now give a general definition and exploit some of its properties.

Definition 2.2 (Venema 2006). We define, for every X, the mapping

$$\mathsf{Base}_X : TX \longrightarrow P_\omega X$$

by putting

$$\mathsf{Base}_X(z) = \bigcap \{W \mid W \in P_\omega X \text{ such that } z \in TW\}.$$

Note that, for a fixed X and $z \in TX$, the system

$$\{W \mid W \in P_\omega X \text{ such that } z \in TW\}$$

is always non-empty since T is finitary. Thus $\mathsf{Base}_X(z)$ is always a finite set. The set $\mathsf{Base}_X(z)$ may be empty, but the following result shows that $\mathsf{Base}_X(z)$ is the smallest finite set W such that $z \in TW$.

Lemma 2.3. For every $z \in TX$, we have $z \in T\mathsf{Base}_X(z)$.

Proof. Since T is assumed to be finitary, there exists a finite set W_0 such that $z \in TW_0$. We define the finite set Z by

$$Z = \bigcap \{W \mid W \subseteq W_0 \text{ such that } z \in TW\}.$$

We then claim the following:

(1) $z \in TZ$.
(2) $Z \subseteq \mathsf{Base}_X(z)$.

The first assertion follows immediately from the fact that T is standard and hence preserves finite intersections (see, for example, Adámek (1983)). The second assertion then follows from the first as follows. First observe that if $z \in TW$ for a finite set W, then $Z \subseteq W$. This is because if $z \in TW$, then $z \in TW \cap TZ = T(W \cap Z)$, so $Z \subseteq W \cap Z$. But the last inclusion implies that $Z \subseteq W$. Thus $Z \subseteq \mathsf{Base}_X(z)$.

To complete the proof, we use the fact that T, being standard, preserves inclusions, so $z \in TZ \subseteq T\mathsf{Base}_X(z)$. □

The following technical lemma shows that lifted relations can be restricted to bases.

Lemma 2.4. Suppose $R \subseteq X \times Y$ and $x\,\overline{R}\,y$ holds. Then there exists

$$S \subseteq \mathsf{Base}_X(x) \times \mathsf{Base}_Y(y)$$

such that $S \subseteq R$ and $x\,\overline{S}\,y$.

Proof. Consider the diagram

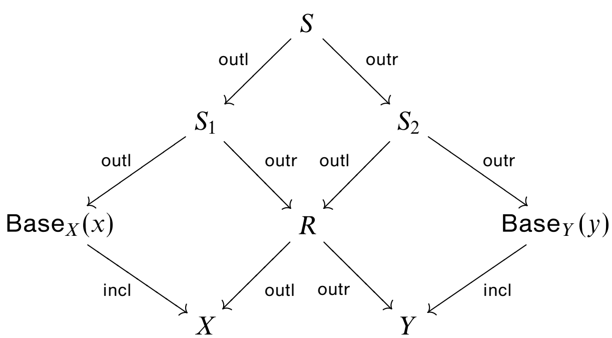

where incl denotes the inclusions and all the squares are pullbacks.

Since injective mappings are stable under pulling back, we know that S is a subset of R through the injective diagonal $f : S \longrightarrow R$ of the middle pullback.

Let $w \in TR$ be the witness of $x\,\overline{R}\,y$. Since T preserves all the above pullbacks weakly, and since $x \in T\mathsf{Base}_X(x)$ and $y \in T\mathsf{Base}_Y(y)$, we conclude that there exists $w' \in TS$ with $Tf(w') = w$ and, moreover, w' witnesses $x\,\overline{S}\,y$. □

The notion of a base extends naturally to objects A of type $P_\omega T\mathscr{L}$ as follows:

$$\mathsf{Base}[A] \;=\; \bigcup\{\mathsf{Base}(\alpha) \mid \alpha \in A\}. \tag{2.2}$$

Another technicality we use when working with nabla is the notion of *slim redistribution* (Kupke *et al.* 2008; Bílková *et al.* 2008).

Definition 2.5. An element $\Phi \in TP_\omega\mathscr{L}$ is a *slim redistribution* of $A \in P_\omega T\mathscr{L}$ (notation $\Phi \in \mathsf{SRD}(A)$) if the following two conditions hold:

(1) $A \subseteq \{\alpha \in T\mathscr{L} \mid \alpha \,\overline{\in}\, \Phi\}$.
(2) $\Phi \in TP_\omega\mathsf{Base}[A]$.

Example 2.6. If T is the power set functor, a set $\Phi \in P_\omega P_\omega X$ is a slim redistribution of a set $A \in P_\omega P_\omega X$ if and only if $\bigcup A = \bigcup \Phi$ (by condition (2) of Definition 2.5) and $\varphi \cap \alpha \neq \emptyset$ for all $\varphi \in \Phi$ and $\alpha \in A$ (by condition (1) of 2.5).

In addition to ∇, there is always its boolean dual modality $\Delta : T\mathscr{L} \longrightarrow \mathscr{L}$ satisfying

$$\Delta\alpha \equiv \neg\nabla(T\neg)\alpha$$

(where $T\neg : T\mathscr{L} \longrightarrow T\mathscr{L}$ is the lifted negation operation). Observe that the following lemma holds for its coalgebraic interpretation.

Lemma 2.7. For any coalgebra $c : X \longrightarrow TX$ and any $x \in X$, we have $x \not\Vdash_c \Delta\beta$ if and only if $c(x)\,\overline{\not\Vdash_c}\,\beta$.

Here, we will just mention those properties of this delta modality that we will need later – see the paper Kissig and Venema (2009) for a detailed discussion. In the following we consider $\mathscr{L}$ to be a boolean language with negations restricted to occur at atoms only, and with both ∇ and Δ (we can clearly do this without loss of generality). We will also consider its positive variant without negation.

We define a dualisation map $d : \mathscr{L} \longrightarrow \mathscr{L}$ and its lifting $Td : T\mathscr{L} \longrightarrow T\mathscr{L}$ inductively:

$$\begin{aligned} d(p) &= \neg p & d(\neg p) &= p \\ d(a \wedge b) &= d(a) \vee d(b) & d(a \vee b) &= d(a) \wedge d(b) \\ d(\nabla\alpha) &= \Delta(Td)\alpha & d(\Delta\alpha) &= \nabla(Td)\alpha. \end{aligned} \tag{2.3}$$

It is clear that through dualisation, we can define nabla using delta and *vice versa*, though this mutual definition only works in the presence of negation. In the positive case we will need another definition – see below. However, it is easy to see that both in the full boolean case and in the positive case d and Td satisfy the following lemma.

Lemma 2.8. For every $\alpha, \beta \in T\mathscr{L}$ and $a, b \in \mathscr{L}$, we have:

(1) $(Td)\alpha \,\overline{\leqslant}\, \beta$ if and only if $(Td)\beta \,\overline{\leqslant}\, \alpha$.
(2) $d(a) \leqslant b$ if and only if $d(b) \leqslant a$.
(3) $(Td)(Td)\alpha = \alpha$.
(4) $d(d(a)) = a$.

Not only is Δ definable from nabla using Boolean negation, but we also have the following definability result, which holds independently of there being a negation in the language and thus also applies to the positive case:

$$\Delta\beta = \bigvee\{\nabla\gamma \mid \gamma \in Q(\beta)\}, \tag{2.4}$$

where

$$Q(\beta) = \{(T\bigwedge)\Phi \mid \ \Phi \in TP_\omega\mathsf{Base}(\beta) \text{ and not } \beta \,\overline{\notin}\, \Phi\ \}$$

where

$$T(\bigwedge) : TP_\omega\mathscr{L} \longrightarrow T\mathscr{L}$$

is the image of the finitary conjunction

$$\bigwedge : P_\omega\mathscr{L} \longrightarrow \mathscr{L}$$

under T. Observe that the set $Q(\beta)$ is *finite* since we assume that T preserves finite sets.

Dually, nabla can be defined as a conjunction of deltas. We can therefore restrict consideration to either of the two modalities as the only modality, even in the positive case.

Given a functor T satisfying our assumptions (that is, standard and finitary and preserving weak pullbacks and finite sets), one can axiomatise the modal logic of T-coalgebras in the boolean language with nabla as the only modality (or dually with delta as the only modality) and prove *completeness* with respect to coalgebraic semantics (Kupke *et al.* 2008). Everything restricts to the language not containing boolean negation and thus to the positive fragment of modal logic for T-coalgebras. We spell out the axioms explicitly for illustration:

$$\begin{array}{ll} (\nabla 0) & \text{Axioms and rules for Classical Propositional Logic} \\ (\nabla 1) & \text{From } \alpha \,\overline{\leqslant}\, \beta \text{ infer } \nabla\alpha \leqslant \nabla\beta \\ (\nabla 2) & \bigwedge\{\nabla\alpha \mid \alpha \in A\} \leqslant \bigvee\{\nabla(T\bigwedge)\Phi \mid \Phi \in \mathsf{SRD}(A)\} \\ (\nabla 3) & \nabla(T\bigvee)\Phi \leqslant \bigvee\{\nabla\alpha \mid \alpha \,\overline{\in}\, \Phi\} \end{array} \tag{2.5}$$

In fact, rules (∇2) and (∇3) above are equalities and they can be thought of as certain modal distributive laws.

Clearly, by the rule (∇1), the mapping $\nabla : \langle T\mathscr{L}, \overline{\leqslant}\rangle \longrightarrow \langle \mathscr{L}, \leqslant\rangle$ is monotone. Using either of the definitions of Δ by ∇, it is clear that $\Delta : \langle T\mathscr{L}, \overline{\leqslant}\rangle \longrightarrow \langle \mathscr{L}, \leqslant\rangle$ is also monotone.

The axiomatisation of the logic based on nabla (or the dual of the axiomatisation in the language with delta) gives a disjunctive (or dually a conjunctive) normal form – every formula is equivalent to a formula in one of the following restricted languages, where π

denotes any non-modal formula:

$$a ::= \pi \mid \bigvee \varphi \mid \pi \wedge \nabla\alpha \tag{2.6}$$

$$a ::= \pi \mid \bigwedge \varphi \mid \pi \vee \Delta\alpha. \tag{2.7}$$

Mutual translations between languages given by predicate liftings, that is, by the modalities we consider, and languages with the nabla modality have been given in Kurz and Leal (2009). In particular, for T satisfying our requirements, there is a one-step translation from modalities to nablas – see Kurz and Leal (2009, Theorem 5.2).

2.4. *Adjunction relative to a doctrine*

We will study the adjointness properties of modalities that are *monotone* operations on the free modal algebra $\mathscr{L}$ that is (pre)ordered by the (semantic) consequence relation. In general, we cannot expect these modalities to be left/right adjoints in the usual sense since they need not preserve suprema/infima in general. Therefore we confine ourselves to adjointness *relative* a certain doctrine on the category PreOrd of preorders and monotone maps. The doctrine will then serve as a measure of adjointness.

In general, a *doctrine* $(\mathbb{D}, \eta)$ consists of a locally monotone functor $\mathbb{D} : \mathsf{PreOrd} \longrightarrow \mathsf{PreOrd}$, together with a natural collection $\eta_{\langle X,\leqslant\rangle} : \langle X,\leqslant\rangle \longrightarrow \mathbb{D}\langle X,\leqslant\rangle$ of fully faithful dense monotone maps. Being fully faithful has the usual meaning when we consider preorders as categories and hence monotone maps as functors. Thus a monotone map $f : \langle X,\leqslant_X\rangle \longrightarrow \langle Y,\leqslant_Y\rangle$ is fully faithful if, given $f(x) \leqslant_Y f(x')$, it follows that $x \leqslant_X x'$. Density means that the natural map $\widetilde{\eta_{\langle X,\leqslant\rangle}}$ from

$$\mathbb{D}\langle X,\leqslant\rangle = \langle DX, \leqslant^{\mathbb{D}}\rangle$$

to the free complete join-semilattice $\mathbb{P}\langle X,\leqslant\rangle$ sending $A \in DX$ to the downset

$$\{x \in X \mid \eta_{\langle X,\leqslant\rangle}(x) \leqslant^{\mathbb{D}} A\}$$

is fully faithful.

Every doctrine $(\mathbb{D}, \eta)$ then allows us to define, for a monotone map

$$L : \langle X,\leqslant_X\rangle \longrightarrow \langle Y,\leqslant_Y\rangle,$$

its *right adjoint relative to* $(\mathbb{D}, \eta)$ as a monotone map

$$R : \langle Y,\leqslant_Y\rangle \longrightarrow \mathbb{D}\langle X,\leqslant_X\rangle$$

with the property

$$Lx \leqslant_Y y \quad \text{iff} \quad \eta_{\langle X,\leqslant_X\rangle}(x) \leqslant^{\mathbb{D}} Ry$$

for all x and y.

Remark 2.9. We will work with 'joins' in preorders in the following. What we mean by this is the notion of a *colimit* known from category theory. Such a colimit is determined uniquely only up to the equivalence $x \sim y$ if and only if $x \leqslant y$ and $y \leqslant x$.

Thus, if we write $x = \bigvee_{i\in I} x_i$ in a preorder, we mean a choice of x such that the following conditions are satisfied:

(1) The inequality $x_i \leqslant x$ holds for all $i \in I$.
(2) Whenever $x_i \leqslant y$ holds for all $i \in I$, we have $x \leqslant y$.

When we speak about join-semilattices, and so on, in the following, the 'joins' are to be understood in the above sense.

Example 2.10. The 'least' possible doctrine consists of the identity functor on PreOrd, and adjointness relative to the identity doctrine is the usual concept of adjointness. The 'largest' possible doctrine is the doctrine $(\mathbb{P}, \eta)$ of complete join-semilattices, and the concept of adjointness relative to this doctrine is void, since then *every* monotone map $L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$ has a right adjoint relative to the doctrine $(\mathbb{P}, \eta)$ of free *complete* join-semilattices, so it suffices to define $R(y) = X$, for every $y \in Y$.

For our applications we choose a doctrine $(\mathbb{P}_\omega, \eta)$ of *free join-semilattices*, which provides us with a concept 'in-between' the usual adjointness and the void concept. In fact, as we argue in Remark 2.13, this doctrine is the 'best possible' when proper adjoints are not available. More precisely, we let

$$\mathbb{P}_\omega : \mathsf{PreOrd} \longrightarrow \mathsf{PreOrd}$$

denote the locally monotone functor of free join-semilattices, that is, for a preorder $\langle X, \leqslant \rangle$, we use $\mathbb{P}_\omega \langle X, \leqslant \rangle$ to denote the free join-semilattice on $\langle X, \leqslant \rangle$. Hence $\mathbb{P}_\omega \langle X, \leqslant \rangle$ is the preorder of all finitely generated downsets of $\langle X, \leqslant \rangle$, ordered by inclusion. The map $\eta_{\langle X, \leqslant \rangle}$ is the usual inclusion of $\langle X, \leqslant \rangle$ in $\mathbb{P}_\omega \langle X, \leqslant \rangle$ given by $x \mapsto \downarrow x$.

Alternatively, and for us, more conveniently, we can describe $\mathbb{P}_\omega \langle X, \leqslant \rangle$ as the set $P_\omega X$ of all finite subsets of X, preordered by the *Hoare preorder* $\leqslant^H$ defined as follows (see Stoltenberg-Hansen *et al.* (1994)):

$$W \leqslant^H Z \text{ if and only if for each } w \in W \text{ there exists } z \in Z \text{ such that } w \leqslant z.$$

We use $\eta_{\langle X, \leqslant \rangle}$ to denote the universal map $x \mapsto \{x\}$ exhibiting $\mathbb{P}_\omega \langle X, \leqslant \rangle$ as a free join-semilattice on $\langle X, \leqslant \rangle$.

We will now spell out explicitly what we mean by the concept of adjointness relative to $(\mathbb{P}_\omega, \eta)$ – from now on, we will refer to the doctrine as $\mathbb{P}_\omega$, and omit η from the notation.

Definition 2.11. We say that a monotone map $L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$ in PreOrd is a *left adjoint relative to* $\mathbb{P}_\omega$ to a monotone map $R : \langle Y, \leqslant_Y \rangle \longrightarrow \mathbb{P}_\omega \langle X, \leqslant_X \rangle$ if we have an equivalence

$$Lx \leqslant_Y y \text{ iff } \{x\} \leqslant^H R(y) \tag{2.8}$$

for all $x \in X$ and $y \in Y$. Or, after unravelling the definition of the preorder $\leqslant^H$ on $\mathbb{P}_\omega \langle X, \leqslant \rangle$, we have the equivalence

$$Lx \leqslant_Y y \text{ iff } x \leqslant_X z \text{ for some } z \in R(y).$$

Remark 2.12. Weak concepts of adjunctions similar to the above were studied in Tholen (1984) and, in the context of modal logic, in Santocanale (2007) and Santocanale and Venema (2007).

The following comments refer to the properties of adjunctions relative to a doctrine, and all follow easily by considering preorders as special (enriched) categories. The above concept of adjunction relative to $\mathbb{P}_\omega$ is then an instance of an adjunction relative to a doctrine in enriched category theory – see Karazeris and Velebil (2009).

(1) Karazeris and Velebil (2009, Theorem 3.7) gives us an 'Adjoint Functor Theorem' for adjointness relative to $\mathbb{P}_\omega$:

A monotone map $L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$ has a right adjoint relative to $\mathbb{P}_\omega$ if and only if the suprema[†]

$$R(y) = \bigsqcup \Big\{ \{x\} \mid L(x) \leqslant_Y y \Big\} \tag{2.9}$$

exist in $\mathbb{P}_\omega \langle X, \leqslant \rangle$ for every $y \in Y$ and are all preserved by the monotone map

$$\widetilde{\eta_{\langle X, \leqslant \rangle}} : A \mapsto \{x \mid \ x \leqslant_X a \text{ for some } a \in A\}$$

from $\mathbb{P}_\omega \langle X, \leqslant \rangle$ to the preorder $\mathbb{P}\langle X, \leqslant \rangle)$ of all lower sets of $\langle X, \leqslant \rangle$.

Moreover, the desired adjoint $R : \langle Y, \leqslant \rangle \longrightarrow \mathbb{P}_\omega \langle X, \leqslant \rangle$ then has the above suprema $R(y)$ as values.

(2) Provided the preorder $\langle Y, \leqslant_Y \rangle$ has enough suprema, one can prove that the monotone map $L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$ has a right adjoint relative to $\mathbb{P}_\omega$ if and only if the monotone map

$$\widehat{L} : \mathbb{P}_\omega \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$$

sending a finite set A to the supremum

$$\bigsqcup \{ Lx \mid x \leqslant_X a \text{ for some } a \in A \}$$

has a right adjoint in the usual sense.

(3) Every left adjoint $L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle$ relative to $\mathbb{P}_\omega$ preserves existing *directed* suprema. This is proved in Santocanale (2007, Lemma 6.2), or it can be derived easily from the fact that $\mathbb{P}_\omega$ freely adds finite suprema to a preorder.

The following remark explains that studying the existence of right adjoints relative to $\mathbb{P}_\omega$ is the 'next best thing to do' when the proper right adjoints are not available.

Remark 2.13. Clearly, the doctrine $\mathbb{P}_\omega$ is an instance of the doctrine $\mathbb{P}_\lambda$, where λ is an infinite regular cardinal. More precisely, $\mathbb{P}_\lambda \langle X, \leqslant_X \rangle$ is the preorder of λ-*generated* downsets in $\langle X, \leqslant_X \rangle$. It is easy to see that $\mathbb{P}_\lambda$ is a locally monotone endofunctor of PreOrd and that, for each λ, there exists a unique 'comparison' natural transformation $\iota : \mathbb{P}_\omega \longrightarrow \mathbb{P}_\lambda$ such that the diagram

$$\begin{array}{ccccc}
\mathbb{P}_\omega & \xrightarrow{\ \iota\ } & \mathbb{P}_\lambda & \longrightarrow & \mathbb{P} \\
{\scriptstyle \eta}\uparrow & & {\scriptstyle \eta}\uparrow & & {\scriptstyle \eta}\uparrow \\
Id & = & Id & = & Id
\end{array}$$

[†] Suprema are to be understood here as colimits in preorders, and are thus determined uniquely up to isomorphisms.

commutes, where the η's are the corresponding canonical maps of the doctrines in question and the upper horizontal unnamed arrow is the obvious comparison from $\mathbb{P}_\lambda$ to $\mathbb{P}$.

In other words, the doctrine $\mathbb{P}_\omega$ is the 'least' among the doctrines that add joins freely and is above the identity doctrine. This can be formulated as the slogan "adjoints relative to $\mathbb{P}_\omega$ are the 'closest' ones to adjoints in the ordinary sense".

3. Cover modalities are adjoints relative to $\mathbb{P}_\omega$

In this section we prove that $\nabla : T\mathscr{L} \longrightarrow \mathscr{L}$, as a monotone map

$$\nabla : \langle T\mathscr{L}, \overline{\leqslant} \rangle \longrightarrow \langle \mathscr{L}, \leqslant \rangle,$$

is a left adjoint relative to $\mathbb{P}_\omega$. In this way we can generalise the results of Santocanale and Venema (2007) from the finite powerset functor to general T (that satisfies our standing assumptions). The adjointness means that for ∇ there exists a monotone map

$$G : \langle \mathscr{L}, \leqslant \rangle \longrightarrow \mathbb{P}_\omega \langle T\mathscr{L}, \overline{\leqslant} \rangle$$

such that for any $b \in \mathscr{L}$,

$$\nabla\alpha \leqslant b \quad \text{iff} \quad \alpha \,\overline{\leqslant}\, \gamma \text{ for some } \gamma \in G(b).$$

Although the considerations in Remark 2.12 give the desired right adjoint explicitly, it may be hard to see straight away that for a given monotone map

$$L : \langle X, \leqslant_X \rangle \longrightarrow \langle Y, \leqslant_Y \rangle,$$

the relevant suprema (2.9) in $\mathbb{P}_\omega\langle X, \leqslant\rangle$ exist. Since this is the case for ∇, we will use a strategy of defining the right adjoint G relative to $\mathbb{P}_\omega$ inductively, using *conjunctive normal forms* of formulas in $\mathscr{L}$ using Δ.

We stress that the side condition that T preserves finite sets will be crucial for our result. The following example shows that one cannot hope to obtain such a result generally.

Example 3.1. Consider the following functor $N : \mathsf{Set} \longrightarrow \mathsf{Set}$ that sends every set to the set $\mathbb{N}$ of natural numbers, and sends any set map to the identity map. It is easy to see that it is a functor, and that it preserves weak pullbacks. Notice that objects in $N\mathscr{L}$ are natural numbers and that $\alpha \,\overline{\leqslant}\, \beta$ means $\alpha = \beta$ as natural numbers, since any lifted relation is equality on $\mathbb{N}$. We claim that the cover modality for this functor is not a left adjoint relative to $\mathbb{P}_\omega$. Consider the case $\nabla\alpha \leqslant \top$ (which always holds). The only candidate for $G(\top)$ such that for any α we have some $\gamma \in G(\top)$ such that $\alpha = \gamma$ as natural numbers clearly equals the whole set $\mathbb{N}$, and no *finite* subset of it would do the job.

We can now state and prove our main theorem.

Theorem 3.2. $\nabla : \langle T\mathscr{L}, \overline{\leqslant} \rangle \longrightarrow \langle \mathscr{L}, \leqslant \rangle$ is a left adjoint relative to $\mathbb{P}_\omega$.

Proof. We have to define a monotone map $G : \langle \mathscr{L}, \leqslant \rangle \longrightarrow \mathbb{P}_\omega \langle T\mathscr{L}, \overline{\leqslant} \rangle$ and prove that, for any $b \in \mathscr{L}$,

$$\nabla\alpha \leqslant b \quad \text{iff} \quad \alpha \,\overline{\leqslant}\, \gamma \text{ for some } \gamma \in G(b). \tag{3.10}$$

We use the fact that we can consider b in a conjunctive normal form given by (2.6). So $b = \bigwedge \varphi$ for some finite set φ, where each $b \in \varphi$ is of the form $\pi \vee \Delta\beta$ and π is a non-modal formula. We will reason by induction on the complexity of b in terms of this normal form.

(I) We first consider the clauses $\pi \vee \Delta\beta$ and will use the following lemma.

Lemma 3.3. Suppose π is a non-modal formula and β is in $T\mathscr{L}$. Then the following are equivalent for every α in $T\mathscr{L}$:

(1) $\nabla\alpha \leqslant \pi \vee \Delta\beta$.

(2) $\top \leqslant \pi$ or $\nabla\alpha \leqslant \Delta\beta$ or $\nabla\alpha \leqslant \bot$.

Proof. The fact that (2) implies (1) is trivial, so we will just prove that (1) implies (2). In condition (2), the third case implies the second, but we will need to distinguish the three cases in the definition of G, so we mention all three cases here. Suppose $\top \nleqslant \pi$ and $\nabla\alpha \nleqslant \Delta\beta$ (thus, clearly, $\nabla\alpha \nleqslant \bot$). Hence the (non-modal!) formula $\neg\pi$ is satisfiable and we can choose a maximal consistent set Γ containing $\neg\pi$. Furthermore, there exists a coalgebra $c : X \longrightarrow TX$ and $x_0 \in X$ such that $x_0 \Vdash_c \nabla\alpha$ and $x_0 \nVdash_c \Delta\beta$. By abuse of notation, we use $x_0 : 1 \longrightarrow X$ to denote the mapping that has x_0 as the value.

We define a new coalgebra

$$c' : X + 1 \longrightarrow T(X+1)$$

as follows:

$$\begin{array}{ccc} X+1 & \xrightarrow{c'} & T(X+1) \\ \uparrow{\scriptstyle \mathsf{inl}} & & \uparrow{\scriptstyle T\mathsf{inl}} \\ X & \xrightarrow[c]{} & TX \end{array} \qquad \begin{array}{ccc} X+1 & \xrightarrow{c'} & T(X+1) \\ \uparrow{\scriptstyle \mathsf{inr}} & & \uparrow{\scriptstyle c} \\ 1 & \xrightarrow[x_0]{} & X \end{array}$$

We need to define the theory map $th_{c'} : X + 1 \longrightarrow Stone(\mathscr{L})$, which we do as follows:

$$\begin{array}{ccc} X+1 & \xrightarrow{th_{c'}} & Stone(\mathscr{L}) \\ \uparrow{\scriptstyle \mathsf{inl}} & \nearrow{\scriptstyle th_c} & \\ X & & \end{array} \qquad \begin{array}{ccc} X+1 & \xrightarrow{th_{c'}} & Stone(\mathscr{L}) \\ \uparrow{\scriptstyle \mathsf{inr}} & \nearrow{\scriptstyle \Gamma} & \\ 1 & & \end{array}$$

Now observe that

$$t \Vdash_{c'} \neg\pi, \quad t \Vdash_{c'} \nabla\alpha, \quad t \nVdash_{c'} \Delta\beta$$

for t being the unique element of 1 in $X+1$, which completes the proof of Lemma 3.3. $\square$

From this lemma, it follows that we can define

$$G(\pi \vee \Delta\beta) := G(\pi) \cup G(\Delta\beta). \tag{3.11}$$

It is therefore clear that we need to distinguish the cases

(a) $\nabla\alpha \leqslant \Delta\beta$, and define $G(\Delta\beta)$.

(b) $\top \leqslant \pi$, and define $G(\top)$.

(c) $\nabla\alpha \leqslant \bot$, and define $G(\bot)$.

We consider each of these in turn:

(a) We first consider the case $\nabla\alpha \leqslant \Delta\beta$:
We define

$$G(\Delta\beta) := \{(T \bigwedge)\Phi \mid \Phi \in TP_\omega\mathsf{Base}(\beta) \text{ and not } \beta \,\overline{\not\in}\, \Phi\} \tag{3.12}$$

where

$$T(\bigwedge) : TP_\omega\mathscr{L} \longrightarrow T\mathscr{L}$$

is the image of finitary conjunction

$$\bigwedge : P_\omega\mathscr{L} \longrightarrow \mathscr{L}$$

under T.
We will use the following result.

Lemma 3.4. The following are equivalent:

(i) $\nabla\alpha \leqslant \Delta\beta$.

(ii) It is not the case that $\alpha \,\overline{\not\leqslant}\, \beta$.

Proof. To prove (i) implies (ii), we will suppose that $\alpha \,\overline{\not\leqslant}\, \beta$ holds and show a contradiction. By Lemma 2.4, we can consider the span

$$\mathsf{Base}(\alpha) \xleftarrow{\mathsf{outl}} \not\leqslant \xrightarrow{\mathsf{outr}} \mathsf{Base}(\beta)$$

By assumption, there exists $x \in T(\not\leqslant)$ such that $T\mathsf{outl}(x) = \alpha$ and $T\mathsf{outr}(x) = \beta$. Now, for every $(a,b) \in \not\leqslant$ there exists a coalgebra $c_{a,b} : X_{a,b} \longrightarrow TX_{a,b}$ and $x_{a,b} \in X_{a,b}$ such that $x_{a,b} \Vdash_{c_{a,b}} a \wedge \neg b$.
We define

$$X = 1 + \coprod_{(a,b)\in\not\leqslant} X_{a,b}$$

and observe that there exists a map

$$f : \not\leqslant \longrightarrow X, \quad (a,b) \mapsto x_{a,b}$$

We define $c : X \longrightarrow TX$ as follows:

$$\begin{array}{ccc} 1 + \coprod_{(a,b)} X_{a,b} & \xrightarrow{\;c\;} & T(1 + \coprod_{(a,b)} X_{a,b}) \\ \uparrow{\scriptstyle \mathsf{in}_1} & & \uparrow{\scriptstyle Tf} \\ 1 & \xrightarrow[t \mapsto x]{} & T(\not\leqslant) \end{array}$$

$$\begin{array}{ccc} 1 + \coprod_{(a,b)} X_{a,b} & \xrightarrow{c} & T(1 + \coprod_{(a,b)} X_{a,b}) \\ {\scriptstyle \mathsf{in}_{(a,b)}} \uparrow & & \uparrow {\scriptstyle T\mathsf{in}_{(a,b)}} \\ X_{a,b} & \xrightarrow[c_{a,b}]{} & TX_{a,b} \end{array}$$

where we have used t to denote the unique element of 1.
We will prove that $t \Vdash_c \nabla\alpha$ and $t \nVdash_c \Delta\beta$, which will give the desired contradiction. Consider the mapping

$$g : \not\leqslant \longrightarrow \Vdash, \quad (a, b) \mapsto (x_{a,b}, a).$$

Then the diagram

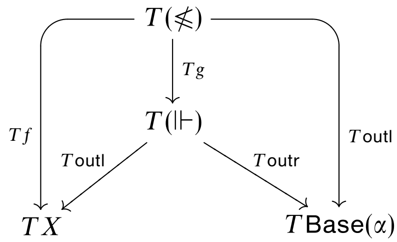

commutes. Put $w = Tg(x)$. Clearly, $T\mathsf{outl}(w) = Tf(x) = c(t)$ and $T\mathsf{outr}(w) = \alpha$, and we have now proved that $t \Vdash_c \nabla\alpha$ holds.
To prove that $t \nVdash_c \Delta\beta$, consider the mapping $h : \not\leqslant \longrightarrow \nVdash$, defined by $(a, b) \mapsto (x_{a,b}, b)$. Then the diagram

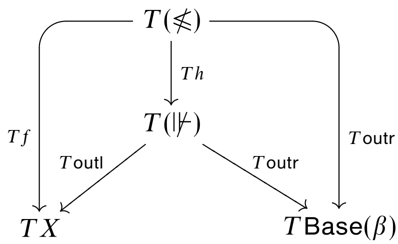

commutes. For $w = Th(x)$, we have $T\mathsf{outl}(w) = Tf(x) = c(t)$ and $T\mathsf{outr}(w) = \beta$, proving that $t \nVdash_c \Delta\beta$.
To prove (ii) implies (i), suppose $\nabla\alpha \not\leqslant \Delta\beta$ and let $c : X \longrightarrow TX$ be a coalgebra and $x \in X$ be such that $x \Vdash_c \nabla\alpha$ and $x \nVdash_c \Delta\beta$. We use $c(x)$ to find a witness of $\alpha \,\overline{\leqslant}\, \beta$.
Consider the diagram

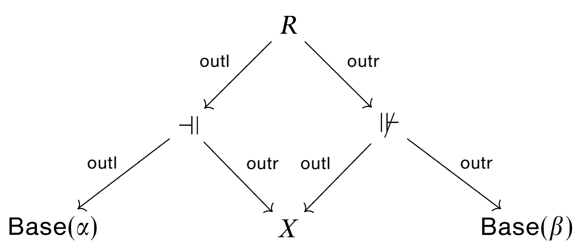

where the square is a pullback. Then $R \subseteq \nleqslant$, since $(a,b) \in R$ means that there exists $x' \in X$ such that $x' \Vdash_c a$ and $x' \nVdash_c b$.
Consider the diagram

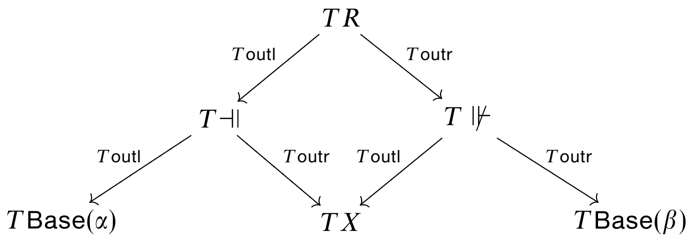

and observe that the square is a weak pullback. We now take w_1 witnessing $c(x)\,\overline{\Vdash}\,\alpha$ and w_2 witnessing $c(x)\,\overline{\nVdash}\,\beta$ and use the weak pullback property to produce an element $w \in TR$. Since T is standard, we have an inclusion $TR \subseteq T(\nleqslant)$ so w is a witness of $\alpha\,\overline{\nleqslant}\,\beta$, which completes the proof of Lemma 3.4. □

To derive (3.10) from the above, we define $f : \mathsf{Base}(\alpha) \longrightarrow P\mathsf{Base}(\beta)$ by

$$f(a) := \{b \in \mathsf{Base}(\beta) \mid a \leqslant b\}.$$

Then for all $a \in \mathsf{Base}(\alpha), b \in \mathsf{Base}(\beta)$, we have that $a \leqslant b$ implies $b \in f(a)$, so $b \notin f(a)$ implies $a \nleqslant b$. By the properties of relation lifting, this means that

$$\beta\,\overline{\notin}\,(Tf)\alpha \text{ implies } \alpha\,\overline{\nleqslant}\,\beta,$$

so it follows from Lemma 3.4 that $(\beta, (Tf)\alpha) \notin \overline{\notin}$. Hence we find that

$$\gamma := (T\textstyle\bigwedge)(Tf)\alpha \in G(\Delta\beta). \tag{3.13}$$

We still need to verify that

$$\alpha\,\overline{\leqslant}\,\gamma. \tag{3.14}$$

To see why this holds, observe that by the definition of f, we have that $a \leqslant \bigwedge f(a)$ for all $a \in \mathsf{Base}(\alpha)$. Thus, in the diagram

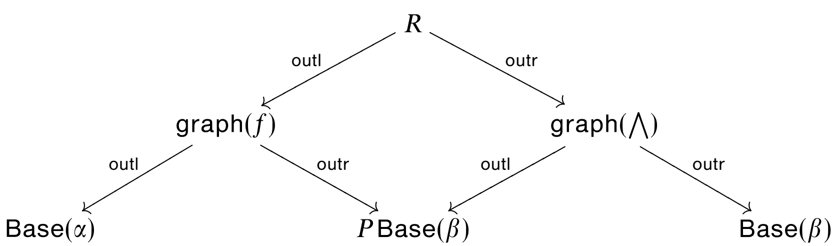

where the square is a pullback, we have $R \subseteq \leqslant$.

Hence the diagram

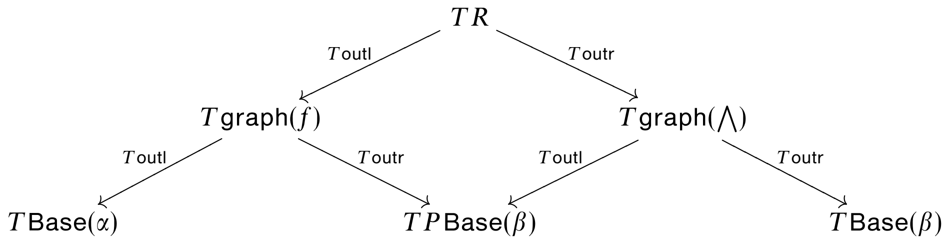

produces a witness $w \in TR$ of $\gamma = (T\bigwedge)(Tf)\alpha$. Since T is standard, $TR \subseteq T(\leqslant)$, from which (3.14) follows.

For the right to left direction, we suppose $\alpha \overline{\leqslant} (T\bigwedge)\Phi$ for some $\Phi \in TP\mathsf{Base}(\beta)$ such that not $\beta \overline{\notin} \Phi$. We need to show that $\nabla\alpha \leqslant \Delta\beta$ holds. By Lemma 3.4 it is sufficient to show that $\alpha \overline{\nleqslant} \beta$ does not hold.

From the assumption $\beta \overline{\notin} \Phi$, we will show by contradiction that $T(\bigwedge)\Phi \overline{\nleqslant} \beta$ does not hold. So we ssume $T(\bigwedge)\Phi \overline{\nleqslant} \beta$ holds. Consider the diagram

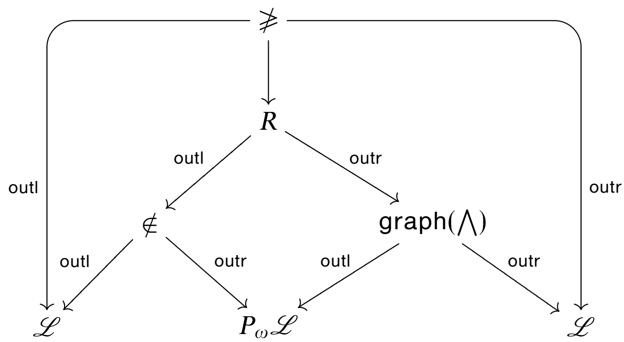

where the middle square is a pullback (hence $R = \{(a, \bigwedge\varphi) \mid a \notin \varphi\}$) and the unnamed arrow is an inclusion due to the fact that $b \ngeqslant \bigwedge\varphi$ implies $b \notin \varphi$.

By lifting the above diagram, we get that $\beta \overline{\notin} \Phi$, which gives a contradiction. Therefore $T(\bigwedge)\Phi \overline{\nleqslant} \beta$ does not hold.

But if $T(\bigwedge)\Phi \overline{\nleqslant} \beta$ does not hold, then $\alpha \overline{\nleqslant} \beta$ does not hold, since $\alpha \overline{\leqslant} T(\bigwedge)\Phi$ holds by assumption and $\overline{\leqslant}$ is transitive.

(b) Next we consider the case $\top \leqslant \pi$:

So $\nabla\alpha \leqslant \pi$ holds for every α. We define

$$G(\top) := \{(T\bigwedge)\Phi \mid \Phi \in TP\{\top\}\}. \tag{3.15}$$

Notice that $TP\{\top\}$ is finite, so $(T\bigwedge) : TP\{\top\} \longrightarrow T\{\top\}$, and thus, in particular, $G(\top)$ is finite and a subset of $T\{\top\}$.

The right to left direction of the theorem is now immediate since $\nabla\alpha \leqslant \top$ holds for free.

We will now show the left to right direction. As in the previous case, we define $f : \mathsf{Base}(\alpha) \longrightarrow P\{\top\}$ as a constant map assigning $\{\top\}$ to each $a \in \mathsf{Base}(\alpha)$

(morally, $f(a)$ is $\{c \in \{\top\} \mid a \leqslant c\}$ as before). Now, for each α and for the same reasons as before, $\alpha \mathrel{\overline{\leqslant}} (T \bigwedge)(Tf)\alpha$, where clearly $(Tf)\alpha$ is in $TP\{\top\}$.

(c) We now consider the case $\nabla\alpha \leqslant \bot$:
We define

$$G(\bot) := \{(T \bigwedge)\Phi \mid \Phi \in TP\{\bot\} \text{ and for all } \beta \in T\{\bot\}, \text{ it is not the case that } \beta \mathrel{\overline{\notin}} \Phi\}. \tag{3.16}$$

The definition of $G(\bot)$ is motivated by the following fact.

Lemma 3.5. $\nabla\alpha \leqslant \bot$ holds if and only if, for all $\beta \in T\{\bot\}$, it is not the case that $\alpha \mathrel{\overline{\not\leqslant}} \beta$.

Proof. We first prove the left to right direction, which we will need later on. We suppose that there exists $\beta \in T\{\bot\}$ such that $\alpha \mathrel{\overline{\not\leqslant}} \beta$ and show that $\nabla\alpha$ is then satisfiable.
Suppose $\alpha \mathrel{\overline{\not\leqslant}} \beta$ is witnessed by an element $x \in T{\not\leqslant}$ such that $(T\mathsf{outl})x = \alpha$ and $(T\mathsf{outr})x = \beta$. Let $Z \subseteq \mathsf{Base}(\alpha) \times \{\bot\}$ denote the obvious restriction of the relation $\not\leqslant$, so $Z = \{(a, \bot) \mid a \text{ satisfiable }\}$. For each $(a, \bot) \in Z$ there is a coalgebra $c_a : X_a \longrightarrow TX_a$ and a state s_a such that $s_a \Vdash_{c_a} a$. We consider a coalgebra having the disjoint union of the family $\{X_a \mid (a, \bot) \in Z\}$ with a new root s_0 added as its carrier. Thus we put

$$X := \coprod\{X_a \mid (a, \bot) \in Z\} + \{s_0\}$$

and the coalgebra map $c : X \longrightarrow TX$ will have an obvious definition for $s \neq s_0$.
We define a mapping $f : Z \longrightarrow S$ by $f(a, \bot) := s_a$ and define $c(s_0) := (Tf)x$.
For each $z = (a, \bot) \in Z$, we have that $(f(z), \mathsf{outl} z) = (s_a, a) \in \Vdash$. By the properties of the relation lifting, it follows that

$$((Tf)x, (T\mathsf{outl})x) = (c(s_0), \alpha) \in \overline{\Vdash},$$

which means that $s_0 \Vdash \nabla\alpha$.
We now suppose $\nabla\alpha$ is satisfiable and will prove that there exists $\beta \in T\{\bot\}$ such that $\alpha \mathrel{\overline{\not\leqslant}} \beta$ holds.
Let $c : X \longrightarrow TX$ be a coalgebra and $x \in X$ be such that $x \Vdash_c \nabla\alpha$ holds. Consider the diagram

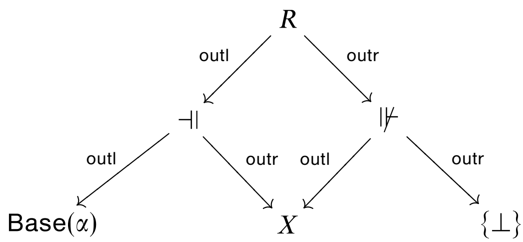

where the square is a pullback. We clearly have $R \subseteq \not\leqslant$. Moreover, R is the graph of the function $g : \dashv\mid \longrightarrow \Vdash$ sending (a, y) to $(y, \bot)$. Hence TR is the graph of Tg since T preserves weak pullbacks and we can find $w \in T(\not\leqslant)$ witnessing

$\alpha \mathrel{\overline{\not\leqslant}} \beta$ as follows. We first find $w_1 \in T(\dashv\!\mid)$ witnessing $c(x) \mathrel{\overline{\Vdash}} \alpha$, and then put $w_2 := Tg(w_1) \in T(\nVdash)$ to produce $w := (w_1, w_2) \in TR$. Since T is standard, $TR \subseteq T(\not\leqslant)$ holds, which concludes the proof of Lemma 3.5. □

To continue the proof of case (c), we first deal with the left to right direction of (3.10). We suppose, therefore, that $\nabla\alpha \leqslant \bot$ holds.
We define $f : \mathsf{Base}(\alpha) \longrightarrow P\{\bot\}$ by putting $f(a) = \{c \in \{\bot\} \mid a \leqslant c\}$. For the same reasons as before, $\alpha \mathrel{\overline{\leqslant}} (T \bigwedge)(Tf)\alpha$, and $(Tf)\alpha$ is in $TP\{\bot\}$.
To prove that $(Tf)\alpha$ is a possible Φ in (3.16), we need to show that the condition $\beta \mathrel{\overline{\notin}} (Tf)\alpha$ holds for no $\beta \in T\{\bot\}$.
Since we assume that $\nabla\alpha \leqslant \bot$ holds, and since by Lemma 3.5 this is equivalent to $\alpha \mathrel{\overline{\not\leqslant}} \beta$ holding for no $\beta \in T\{\bot\}$, it suffices to prove that

$$\text{for all } \beta \in T\{\bot\} \text{ not } \alpha \mathrel{\overline{\not\leqslant}} \beta \text{ implies for all } \beta \in T\{\bot\} \text{ not } \beta \mathrel{\overline{\notin}} (Tf)\alpha \tag{3.17}$$

holds.
To justify (3.17), suppose there is $\beta \in T\{\bot\}$ such that $\beta \mathrel{\overline{\notin}} (Tf)\alpha$ holds, and is witnessed by an element w_1 in $T{\notin}$. (Here we use $\notin$ to denote the obvious restriction to the subset of $\{\bot\} \times P\{\bot\}$.)
We show there is a witness w in $T{\not\leqslant}$ of $\alpha \mathrel{\overline{\not\leqslant}} \beta$. We again write

$$\not\leqslant \subseteq \mathsf{Base}(\alpha) \times \{\bot\}$$

to mean the obvious restriction of the full $\not\leqslant$ relation.
Now consider the opposite of the graph of f:

$$\mathsf{graph}(f)^{op} \subseteq P\{\bot\} \times \mathsf{Base}(\alpha),$$

equipped with the mapping

$$\mathsf{fibre}(f) : \mathsf{Base}(\alpha) \longrightarrow \mathsf{graph}(f)^{op}$$

that assigns to each a the pair $(f(a), a)$. We use R to denote the pullback of the relations $\mathsf{graph}(f)^{op}$ and $\notin$. Thus, we have the following diagram

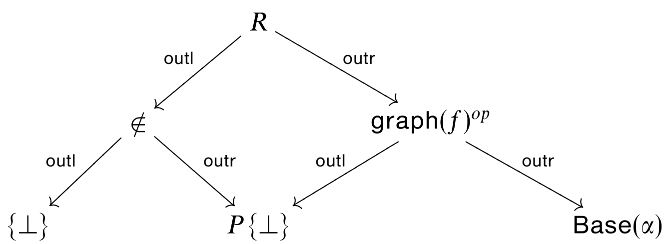

where the square is a pullback.
By lifting the above diagram and using the properties of T and the witness w_1 in $T \notin$ from assumption, such that

$$(T\mathsf{outl})w_1 = \beta$$
$$(T\mathsf{outr})w_1 = (Tf)\alpha,$$

and using the element

$$w_2 := (T\mathsf{fibre}(f))\alpha$$

in $T(\mathsf{graph}(f)^{op})$, we obtain an element w in the pullback TR. Since $R \subseteq \not\geq$, we have $TR \subseteq T \not\geq$. Thus w witnesses that $\beta \mathrel{\overline{\not\geq}} \alpha$, meaning that $\alpha \mathrel{\overline{\not\leq}} \beta$ (since $\overline{\not\geq} = (\overline{\not\leq})^{op}$), which concludes the proof of (3.17).

The right to left direction of (3.10) for case (c) can be reformulated as follows. We suppose that $\alpha \mathrel{\overline{\leqslant}} (T \bigwedge)\Phi$ for some $\Phi \in TP\{\bot\}$, and then want to show that if $\nabla\alpha$ were satisfiable, then there would exist β in $T\{\bot\}$ such that $\beta \mathrel{\overline{\notin}} \Phi$.

So we suppose that there is a coalgebra $c : X \longrightarrow TX$ and a state $x \in X$ satisfying $\nabla\alpha$. Thus $c(s) \mathrel{\overline{\Vdash}} \alpha$. Since $\alpha \mathrel{\overline{\leqslant}} (T \bigwedge)\Phi$, we also have $c(s) \mathrel{\overline{\Vdash}} (T \bigwedge)\Phi$. We define a map g from $\Vdash \subseteq X \times \{\top, \bot\}$ to $\not\ni \subseteq P\{\bot\} \times \{\bot\}$ by $g : (x, \top) \longrightarrow (\varnothing, \bot)$.

Since $c(s) \mathrel{\overline{\Vdash}} (T \bigwedge)\Phi$, we have a witness w in $T\Vdash$ of this fact. We now suppose outr is the right projection of $\not\ni$, and define $\beta = (T\mathsf{outr})(Tg)w$. Now $(Tg)w$ witnesses that $\Phi \mathrel{\overline{\not\ni}} \beta$ – see the following diagram

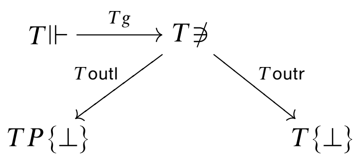

This concludes the proof of case (c), and thus of part (I) of the theorem.

(II) Finally, we consider the general case where b is a conjunction $\bigwedge \varphi$.

Let $\mathscr{B}_\varphi$ be the collection of sets of the form $\{\beta_b \mid b \in \varphi\}$ such that for each $b \in \varphi$, the set $\beta_b \in G(b)$.

We now define

$$G(\bigwedge \varphi) := \{(T \bigwedge)\Phi \mid \Phi \in \mathsf{SRD}(B) \text{ for some } B \in \mathscr{B}_\varphi\}. \tag{3.18}$$

For the left to right direction in (3.10), we assume that $\nabla\alpha \leqslant \bigwedge \varphi$. Then $\nabla\alpha \leqslant b$ for all $b \in \varphi$, so for all $b \in \varphi$, there is a $\beta_b \in G(b)$ such that $\alpha \mathrel{\overline{\leqslant}} \beta_b$. We define $B := \{\beta_b \mid b \in \varphi\}$, and it is then clear that $B \in \mathscr{B}_\varphi$.

With the abbreviation

$$\mathsf{Base}[B] := \bigcup_{b \in \varphi} \mathsf{Base}(\beta_b),$$

we define the map

$$f : \mathsf{Base}(\alpha) \longrightarrow P\mathsf{Base}[B]$$

by putting

$$f(a) := \{c \in \mathsf{Base}[B] \mid a \leqslant c\},$$

and let

$$\Phi := (Tf)\alpha$$

$$\gamma := (T \bigwedge)\Phi.$$

Then it suffices to prove that

$$\Phi \in \mathsf{SRD}(B), \tag{3.19}$$

from which it follows that $\gamma \in G(\bigwedge \varphi)$, and that

$$\alpha \mathrel{\overline{\leqslant}} \gamma. \tag{3.20}$$

The proof of (3.20) is analogous to that of (3.14), so we confine our attention to the proof of (3.19). Since $\Phi \in TP\mathsf{Base}[B]$, it suffices to show that every element $\beta \in B$ is a lifted member of Φ. We take such a β; by the definition of B, we know β is of the form β_b for some $b \in \varphi$. From this it follows that $\alpha \mathrel{\overline{\leqslant}} \beta$. Also, for any $a \in \mathsf{Base}(a)$ and $c \in \mathsf{Base}[B]$, we have by the definition of f that $a \leqslant c$ implies $c \in f(a)$. So, by the properties of relation lifting, it follows from $\alpha \mathrel{\overline{\leqslant}} \beta$ that $\beta \mathrel{\overline{\in}} (Tf)\alpha = \Phi$. This completes the proof of (3.19).

Conversely, suppose $\alpha \mathrel{\overline{\leqslant}} (T \bigwedge)\Phi$ for some $\Phi \in \mathsf{SRD}(B)$ and some $B \in \mathscr{B}_\varphi$. Since $\Phi \in \mathsf{SRD}(B)$, we have $\beta_b \mathrel{\overline{\in}} \Phi$ for all $\beta_b \in B$.

Consider the following diagram

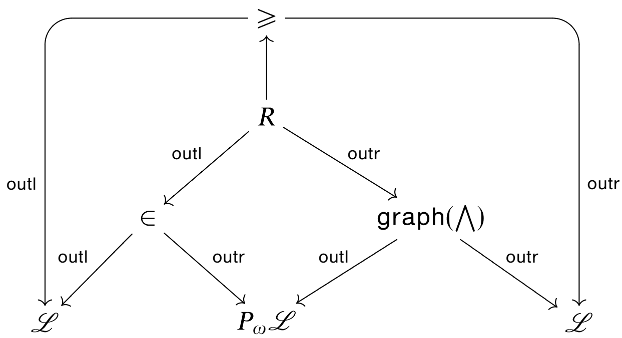

where the middle square is a pullback (hence $R = \{(a, \bigwedge \varphi) \mid a \in \varphi\}$) and the unnamed arrow is an inclusion due to the fact that $a \geqslant \bigwedge \varphi$ holds for every pair $(a, \bigwedge \varphi) \in R$. By lifting the above diagram by T, we get that $T(\bigwedge)\Phi \mathrel{\overline{\leqslant}} \beta_b$ holds for every $\beta_b \in B$. Thus, in particular, by the definition of B, we have proved that for every $b \in \varphi$ there exists $\beta_b \in G(b)$ such that $\alpha \mathrel{\overline{\leqslant}} \beta_b$. Hence, for every $b \in \varphi$ we have $\nabla\alpha \leqslant b$, so $\nabla\alpha \leqslant \bigwedge \varphi$ as desired.

This completes the proof of part (II), and thus the proof of Theorem 3.2 – our main theorem. □

Corollary 3.6. $\Delta : \langle T\mathscr{L}, \overline{\leqslant}\rangle \longrightarrow \langle \mathscr{L}, \leqslant\rangle$ is a right adjoint relative to $\mathbb{P}_\omega$.

Proof. For Δ this means showing that there is a monotone map

$$G' : \langle \mathscr{L}, \leqslant\rangle \longrightarrow \mathbb{P}_\omega\langle T\mathscr{L}, \overline{\leqslant}\rangle$$

such that for any $a \in \mathscr{L}$,

$$a \leqslant \Delta\beta \quad \text{iff} \quad \gamma \mathrel{\overline{\leqslant}} \beta \text{ for some } \gamma \in G'(a).$$

The following are equivalent (the second equivalence follows from Theorem 3.2, and the first and last from Lemma 2.8):

$$\begin{aligned} a &\leqslant \Delta\beta \\ \nabla(Td)\beta &\leqslant d(a) \\ (Td)\beta &\overline{\leqslant}\ \gamma \text{ for some } \gamma \in G(d(a)) \\ (Td)\gamma &\overline{\leqslant}\ \beta \text{ for some } \gamma \in G(d(a)). \end{aligned}$$

Now we can define $G'(a) = \{(Td)\gamma \mid \gamma \in G(d(a))\}$, which is clearly monotone. □

4. Unary monotone modalities

We start with the definition of some 'basic' unary monotone modalities. The idea is that we define for each 'mode of future' (with respect to a meaning $\|a\|$ of a formula a) given by an element $r \in T2$, its own pair of monotone modalities $\oplus_r a$, $\ominus_r a$ with the intended interpretation of 'being satisfied (refuted) by futures of type at least (at most) r'. Thus defined, the two sets turn out to be mutually definable through the usual dual boolean laws, and, moreover, they are mutually definable even in the positive case, analogously to nabla and delta modalities.

In fact, in doing this we will have systematically covered all the upper sets in the preorder $\langle T2, \overline{\leqslant}\rangle$, and thus all unary monotone modalities are definable using (disjunctions of) the basic ones.

We will show that $\oplus_r$ are essentially nablas, and $\ominus_r$ are essentially deltas, which, using the result obtained in the previous section, leads to the conclusion that they are left and right adjoints, respectively, relative to $\mathbb{P}_\omega$. Moreover, because of their mutual definability, they are all *both* right and left adjoints relative to $\mathbb{P}_\omega$ (a similar argument does *not* apply to nabla and delta, and it is not known whether they are both right and left relative adjoints or not). One might relate these unary modalities to the singleton liftings of Kurz and Leal (2009), except that here we consider the modalities to be monotone. The relationship to nablas that we prove below is related to the results on mutual translations between predicates liftings and nabla given in Kurz and Leal (2009). However, we prove the relationship explicitly in the setting we fixed in Section 2.2.

Definition 4.1. For each $r \in T2$, we define unary modalities $\oplus_r : T2 \longrightarrow 2$ and $\ominus_r : T2 \longrightarrow 2$ as follows:

$$\begin{aligned} \oplus_r s = 1 &\text{ iff } r \overline{\leqslant} s, \\ \ominus_r s = 0 &\text{ iff } s \overline{\leqslant} r. \end{aligned}$$

The modalities are obviously monotone, the sets of those $s \in T2$ returning value 1 are in both cases upper subsets of $T2$.

Lemma 4.2. Given a coalgebra $c : X \longrightarrow TX$, we have

$$x \Vdash_c \oplus_r a \quad \text{iff} \quad T\|a\|c(x) \overline{\geqslant} r$$

and, analogously,

$$x \not\Vdash_c \ominus_r a \quad \text{iff} \quad T\|a\|c(x) \mathrel{\overline{\leqslant}} r.$$

Moreover, the following properties hold:

(1) If $a \leqslant b$, then $\oplus_r a \leqslant \oplus_r b$ and $\ominus_r a \leqslant \ominus_r b$.
(2) If $s \mathrel{\overline{\leqslant}} r$, then $\oplus_r a \leqslant \oplus_s a$ and $\ominus_r a \leqslant \ominus_s a$.
(3) $\oplus_r a \equiv \neg \ominus_{r'} \neg a$ and $\ominus_r a \equiv \neg \oplus_{r'} \neg a$, where r' is the image of $r \in T2$ under $T\mathsf{swap} : T2 \longrightarrow T2$ where $\mathsf{swap} : 2 \longrightarrow 2$ swaps 0 and 1.
(4) $c \leqslant \ominus_r a$ if and only if $\oplus_{r'} d(a) \leqslant d(c)$, where r' is as in the previous condition and d is the dualisation map $d : \mathscr{L} \longrightarrow \mathscr{L}$, see (2.3) above.
(5) $\ominus_r a \equiv \bigvee_{s \not\overline{\leqslant} r} \oplus_s a$ and, dually, $\oplus_r a \equiv \bigwedge_{s \not\overline{\geqslant} r} \ominus_s a$.
(6) Given a unary monotone modality $\odot : T2 \longrightarrow 2$, there are $r_1, \ldots, r_k$ such that $\odot a \equiv \bigvee_{i=1}^{k} \oplus_{r_i} a$.

Proof. The satisfaction relations for $\oplus_r$ and $\ominus_r$ follow immediately from the definition.

Assertions (1), (2), (3) and (4) are straightforward. To prove (5), suppose $\ominus_r t = 1$, for $t \in T2$. Then $t \not\overline{\leqslant} r$ and $\oplus_t t = 1$, hence

$$\bigvee_{s \not\overline{\leqslant} r} \oplus_s t = 1.$$

Conversely, if

$$\bigvee_{s \not\overline{\leqslant} r} \oplus_s t = 1,$$

there exists $s_0 \not\overline{\leqslant} r$ such that $\oplus_{s_0} t = 1$. From the first fact we get that $\ominus_r s_0 = 1$, and from the second fact we know that $s_0 \mathrel{\overline{\leqslant}} t$. We now use the fact that $\ominus_r$ is monotone to conclude that $\ominus_r t = 1$. For assertion (6), we use the fact that $T2$ is a finite set: the upper set $\{s \mid \odot s = 1\}$ in $T2$ must be finitely generated, hence

$$\odot t = \bigvee_{i=1}^{k} \oplus_{r_i} t,$$

for some $r_1, \ldots, r_k$ in $T2$. □

Notice that properties (5) and (6) in this lemma rely heavily on the fact that $T2$ is a finite set.

Example 4.3. For $T = P_\omega$, observe that $\langle T2, \overline{\leqslant} \rangle$ is the following poset (in the general case $T2$ is only a preorder):

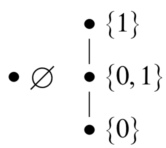

And, in the language of $\Box$ and $\Diamond$, we have the following equivalences:

$$
\begin{array}{lll@{\qquad}lll}
\oplus_{\varnothing} a & \equiv & \Box\bot & \ominus_{\varnothing} a & \equiv & \Diamond\top \\
\oplus_{\{0\}} a & \equiv & \Diamond\top & \ominus_{\{0\}} a & \equiv & \Box a \vee \Diamond a \\
\oplus_{\{0,1\}} a & \equiv & \Diamond a & \ominus_{\{0,1\}} a & \equiv & \Box a \\
\oplus_{\{1\}} a & \equiv & \Box a \wedge \Diamond a & \ominus_{\{1\}} a & \equiv & \Box\bot.
\end{array}
$$

Recalling the definition of nabla in the language of box and diamond,

$$\nabla\alpha \equiv \Box \bigvee \alpha \wedge \bigwedge \Diamond\alpha,$$

the pattern of nablas (deltas) is immediately visible behind the modalities in the previous example: for example, $\oplus_{\{0,1\}} a \equiv \nabla\{\top, a\}$ or $\oplus_{\{1\}} a \equiv \nabla\{a\}$, so $0 \in r$ means consider $\top \in \alpha$, while $1 \in r$ means consider $a \in \alpha$.

We will now prove that this is true for *any* T we may consider – unary modalities $\oplus_r$ are essentially ∇ (dually, $\ominus_r$ are essentially Δ). For the following it is instructive to look first at the proofs of Santocanale and Venema (2007, Theorem 6.10 and Corollary 6.12), where analogous ideas are used to prove the relative adjointness of nabla for the power set functor.

We will also need the following auxiliary notation for the proof.

Notation 4.4. For every formula a in $\mathscr{L}$, define its *name* to be the function

$$\ulcorner a \urcorner : 2 \longrightarrow \mathscr{L}, \quad 0 \mapsto \top, \quad 1 \mapsto a.$$

Observe that, for each formula a, the map $\ulcorner a \urcorner$ reverses the order, thus it is a map

$$\ulcorner a \urcorner : \langle 2, \leqslant \rangle^{op} \longrightarrow \langle \mathscr{L}, \leqslant \rangle.$$

Therefore, it can be lifted to a monotone map

$$\overline{T} \ulcorner a \urcorner : \langle T2, \overline{\leqslant} \rangle^{op} \longrightarrow \langle T\mathscr{L}, \overline{\leqslant} \rangle$$

and, finally, we have a monotone map

$$\langle \mathscr{L}, \leqslant \rangle \longrightarrow [\langle T2, \overline{\leqslant} \rangle^{op}, \langle T\mathscr{L}, \overline{\leqslant} \rangle], \quad a \mapsto \overline{T} \ulcorner a \urcorner.$$

The transpose of the above map is therefore a monotone map

$$\langle T2, \overline{\leqslant} \rangle^{op} \longrightarrow [\langle \mathscr{L}, \leqslant \rangle, \langle T\mathscr{L}, \overline{\leqslant} \rangle], \quad r \mapsto f_r.$$

Observe that $f_r(a)$ is an element of $T\mathscr{L}$.

Observe that the following holds for each $f_r(a)$.

Lemma 4.5. $\mathsf{Base}(f_r(a)) \subseteq \{a, \top\}$.

Lemma 4.6. For each $r \in T2$ and each $a \in \mathscr{L}$, we have the equivalence $\oplus_r a \equiv \nabla f_r(a)$.

Proof. Fix $r \in T2$ and $a \in \mathscr{L}$. Suppose $c : X \longrightarrow TX$ is any coalgebra and $x \in X$ is arbitrary.

(1) Suppose $x \Vdash_c \oplus_r a$ holds. By Lemma 4.2, this means that $T\|a\|c(x) \,\overline{\geqslant}\, r$ holds. First consider a pullback

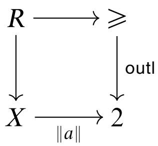

Thus $R = \{(x, a, b) \mid \|a\|(x) = a \geqslant b\}$. Observe that the diagram

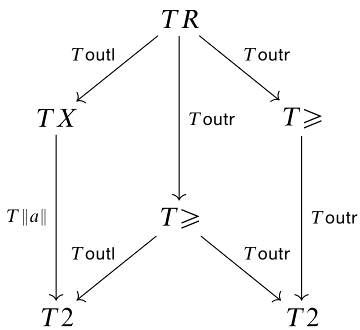

(4.21)

commutes and that the left-hand square is a weak pullback.
Since we assume that $T\|a\|c(x) \,\overline{\geqslant}\, r$ holds, we have a witness $z \in T{\geqslant}$ with

$$T\mathsf{outl}(z) = \|a\|c(x),$$

and by the property of weak pullbacks, there exists $w \in TR$ such that $T\mathsf{outl}(w) = c(x)$ and $T\mathsf{outr}(T\mathsf{outr}(w)) = r$.
Define $g : R \longrightarrow \Vdash_c$ by putting

$$g(x, 0, 0) = g(x, 1, 0) = (x, \top)$$
$$g(x, 1, 1) = (x, a).$$

Then the diagram

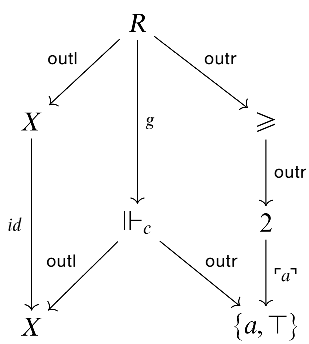

(4.22)

commutes (where we have slightly abused the notation for the restricted relation $\Vdash_c$ and $\ulcorner a \urcorner$).
By applying $Tg : TR \longrightarrow T{\Vdash_c}$ to the element $w \in TR$ we have found above, we get $Tg(w) \in T{\Vdash_c}$ witnessing $x \Vdash_c \nabla f_r(a)$.

(2) Suppose $x \Vdash_c \nabla f_r(a)$ holds. We will prove that $x \Vdash_c \oplus_r a$ holds.
Consider the diagram

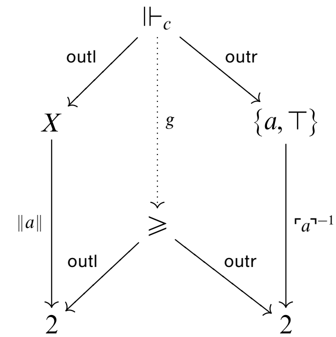

where we have slightly abused the notation for the restricted relation $\Vdash_c$ and $\ulcorner a \urcorner$. The perimeter of the diagram obviously commutes, and it is clear that we can define the dotted arrow g. Thus, the diagram

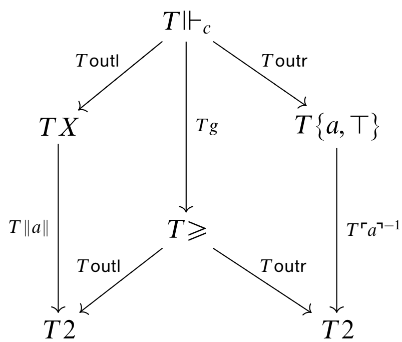

commutes. Hence, if $x \Vdash_c \nabla f_r(a)$ holds, we have a witness of the relation

$$T\|a\|c(x) \mathrel{\overline{\geqslant}} r,$$

and this is exactly the statement $x \Vdash_c \oplus_r a$ by Lemma 4.2, which completes the proof. □

Proposition 4.7. Every $f_r : \langle \mathscr{L}, \leqslant \rangle \longrightarrow \langle T\mathscr{L}, \overline{\leqslant} \rangle$ is a left adjoint relative to $\mathbb{P}_\omega$.

Proof. We will define a monotone map $G : \langle T\mathscr{L}, \overline{\leqslant} \rangle \longrightarrow \mathbb{P}_\omega \langle \mathscr{L}, \leqslant \rangle$ and prove that for every $a \in \mathscr{L}$ and every $\alpha \in T\mathscr{L}$,

$$f_r(a) \mathrel{\overline{\leqslant}} \alpha \quad \text{iff} \quad a \leqslant c \text{ for some } c \in G(\alpha).$$

We define an auxiliary map $t : \mathscr{L} \longrightarrow 2$ by putting

$$t(\top) = 0 \text{ and } t(c) = 1 \text{ for } c \neq \top.$$

We claim that the map $G : T\mathscr{L} \longrightarrow P_\omega \mathscr{L}$ defined by

$$G(\alpha) = \begin{cases} \varnothing & \text{iff } Tt(\alpha) \mathrel{\overline{\not\leqslant}} r \\ \{\bigwedge \mathsf{Base}(\alpha)\} & \text{iff } Tt(\alpha) \mathrel{\overline{\leqslant}} r \end{cases}$$

is the desired right adjoint relative to $\mathbb{P}_\omega$.

We divide the proof into two cases according to the definition of $G(\alpha)$:

(1) The case of the empty set.
The main idea is that there is an a such that $f_r(a) \mathrel{\overline{\leqslant}} \alpha$ if and only if $Tt(\alpha) \mathrel{\overline{\leqslant}} r$. The above equivalence then proves that the adjointness formula for G works for α with $Tt(\alpha) \mathrel{\not\overline{\leqslant}} r$.
We will prove:

$$\text{There is an } a \text{ such that } f_r(a) \mathrel{\overline{\leqslant}} \alpha \text{ holds if and only if } Tt(\alpha) \mathrel{\overline{\leqslant}} r. \tag{4.23}$$

Suppose there exists a such that $f_r(a) \mathrel{\overline{\leqslant}} \alpha$ holds. Consider the pullback

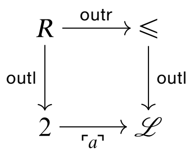

This means that

$$R = \{(0, \top, \top), (1, a, \top)\} \cup \{(1, a, c) \mid a \leqslant c \neq \top\}.$$

Define $g : R \longrightarrow \geqslant$ by putting

$$g(0, \top, \top) = (0, 0)$$
$$g(1, a, \top) = (1, 0)$$
$$g(1, a, c) = (1, 1).$$

Then the diagram

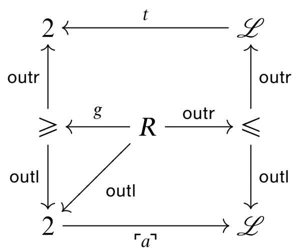

clearly commutes, and so does its image under T:

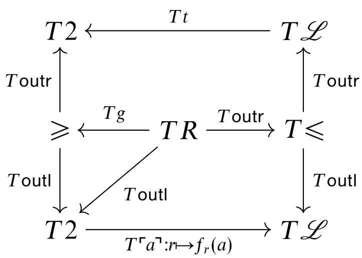

Since $f_r(a) \mathrel{\overline{\leqslant}} \alpha$ holds, there is $z \in T\leqslant$ such that $T\mathsf{outl}(z) = f_r(a)$ and $T\mathsf{outr}(z) = \alpha$. Since T weakly preserves pullbacks, there is a witness $w \in TR$ such that $T\mathsf{outl}(w) = r$

and $T\mathsf{outr}(w) = z$. Then $Tg(w) \in T\geqslant$ witnesses $r \mathrel{\overline{\geqslant}} Tt(\alpha)$, which is what we set out to prove.

Conversely, suppose $Tt(\alpha) \mathrel{\overline{\leqslant}} r$ holds. We claim that for $a = \bot$, the relation $f_r(a) \mathrel{\overline{\leqslant}} r$ holds.

Consider the pullback

$$\begin{array}{ccc} Q & \xrightarrow{\mathsf{outr}} & \geqslant \\ {\scriptstyle \mathsf{outl}}\downarrow & & \downarrow{\scriptstyle \mathsf{outr}} \\ \mathscr{L} & \xrightarrow[t]{} & 2 \end{array} \tag{4.24}$$

Thus

$$Q = \{(\top, 0, 0), (\top, 1, 0)\} \cup \{(c, 1, 1) \mid c \neq \top\}.$$

We define $h : Q \longrightarrow \leqslant$ by putting

$$h(\top, 0, 0) = (\top, \top)$$
$$h(\top, 1, 0) = (\bot, \top)$$
$$h(c, 1, 1) = (\bot, c) \text{ for } c \neq \top.$$

Then the diagram

$$\begin{array}{ccccc} \mathscr{L} & & \xleftarrow{\ulcorner\bot\urcorner} & & 2 \\ {\scriptstyle \mathsf{outl}}\uparrow & & & & \uparrow{\scriptstyle \mathsf{outl}} \\ \leqslant & \xleftarrow{h} & Q & \xrightarrow{\mathsf{outr}} & \geqslant \\ {\scriptstyle \mathsf{outr}}\downarrow & \swarrow{\scriptstyle \mathsf{outl}} & & & \downarrow{\scriptstyle \mathsf{outr}} \\ \mathscr{L} & & \xrightarrow[t]{} & & 2 \end{array} \tag{4.25}$$

clearly commutes, and so does its image under T:

$$\begin{array}{ccccc} T\mathscr{L} & & \xleftarrow{T\ulcorner\bot\urcorner : r \mapsto f_r(\bot)} & & T2 \\ {\scriptstyle T\mathsf{outl}}\uparrow & & & & \uparrow{\scriptstyle T\mathsf{outl}} \\ T\leqslant & \xleftarrow{Th} & TQ & \xrightarrow{T\mathsf{outr}} & T\geqslant \\ {\scriptstyle T\mathsf{outr}}\downarrow & \swarrow{\scriptstyle T\mathsf{outl}} & & & \downarrow{\scriptstyle T\mathsf{outr}} \\ T\mathscr{L} & & \xrightarrow[Tt]{} & & T2 \end{array} \tag{4.26}$$

Since $Tt(\alpha) \mathrel{\overline{\leqslant}} r$ holds, there exists $z \in T\geqslant$ such that $T\mathsf{outr}(z) = Tt(\alpha)$ and $T\mathsf{outl}(z) = r$. Since T weakly preserves pullbacks, there exists $w \in TQ$ such that $T\mathsf{outr}(w) = z$ and $T\mathsf{outl}(w) = \alpha$. Then $Th(w) \in T\leqslant$ witnesses $f_r(\bot) \mathrel{\overline{\leqslant}} \alpha$, which is what we set out to prove.

(2) For the other cases, we need to prove

$$f_r(a) \mathrel{\overline{\leqslant}} \alpha \text{ iff } a \leqslant \bigwedge \mathsf{Base}(\alpha) \tag{4.27}$$

Suppose $f_r(a)\ \overline{\leqslant}\ \alpha$ holds. By Lemma 4.5 we know that $\mathsf{Base}(f_r(a)) \subseteq \{a, \top\}$. We will use the denotation

$$d(a) = \{c \in \mathsf{Base}(\alpha) \mid a \leqslant c\}.$$

We will distinguish three cases:

(a) $\mathsf{Base}(f_r(a))$ contains a:
It suffices to show that $\alpha \in T(d(a))$, since then, from the definition, each $b \in \mathsf{Base}(\alpha)$ would be in $d(a)$, thus proving that $a \leqslant \bigwedge \mathsf{Base}(\alpha)$.
By Lemma 2.4, it follows that $w \in T\leqslant$ witnessing $f_r(a)\ \overline{\leqslant}\ \alpha$ is, in fact, in the vertex of the span

$$\begin{array}{ccc} & T\leqslant & \\ {\scriptstyle T\mathsf{outl}}\swarrow & & \searrow{\scriptstyle T\mathsf{outr}} \\ T\mathsf{Base}(f_r(a)) & & T\mathsf{Base}(\alpha) \end{array}$$

A more detailed inspection shows that the above span is, in fact, the lift under T of the span

$$\begin{array}{ccc} & \leqslant & \\ {\scriptstyle \mathsf{outl}}\swarrow & & \searrow{\scriptstyle \mathsf{outr}} \\ \mathsf{Base}(f_r(a)) & & d(a) \end{array} \qquad (4.28)$$

Thus $T\mathsf{outr} : T\leqslant \longrightarrow Td(a)$ and $T\mathsf{outr}(w) = \alpha \in Td(a)$.

(b) $\mathsf{Base}(f_r(a)) = \{\top\}$:
We proceed as above with $d(\top)$ instead of $d(a)$ to conclude that $\alpha \in T(d(\top))$. Thus every $b \in \mathsf{Base}(\alpha)$ is in $d(\top) = \{\top\}$, hence $a \leqslant \bigwedge \mathsf{Base}(\alpha)$.

(c) $\mathsf{Base}(f_r(a)) = \varnothing$.
So $\leqslant$ in (4.28) is empty, and we can work with $\varnothing$ instead of $d(a)$ to conclude that $\alpha \in T\varnothing$. Hence $\mathsf{Base}(\alpha) = \varnothing$ and $\bigwedge \mathsf{Base}(\alpha) = \top$.

Conversely, suppose $a \leqslant \bigwedge \mathsf{Base}(\alpha)$. Also, recall that we are working under the assumption that $Tt(\alpha)\ \overline{\leqslant}\ r$ holds.
We now proceed similarly to part (1).

— When $\top \in \mathsf{Base}(\alpha)$, instead of the pullback (4.24), we consider

$$\begin{array}{ccc} Q & \xrightarrow{\ \mathsf{outr}\ } & \geqslant \\ {\scriptstyle \mathsf{outl}}\downarrow & & \downarrow{\scriptstyle \mathsf{outr}} \\ \mathsf{Base}(\alpha) & \xrightarrow[\ t\ ]{} & 2 \end{array}$$

where $t : \mathsf{Base}(\alpha) \longrightarrow 2$ sends $\top$ to 0 and everything else to 1. Hence

$$Q = \{(\top, 0, 0), (\top, 1, 0)\} \cup \{(c, 1, 1) \mid c \in \mathsf{Base}(\alpha), c \neq \top\}.$$

Define $h : Q \longrightarrow \leqslant$ by putting

$$h(\top, 0, 0) = (\top, \top)$$
$$h(\top, 1, 0) = (a, \top)$$
$$h(c, 1, 1) = (a, c) \text{ for } c \neq \top.$$

Since $a \leqslant c$ for all $c \in \mathsf{Base}(\alpha)$, we obtain similar diagrams to (4.25) and (4.26), and we can conclude that $f_r(a) \overline{\leqslant} \alpha$.

— When $\top \notin \mathsf{Base}(\alpha)$, we let $t : \mathsf{Base}(\alpha) \longrightarrow 2$ be constantly 1 and define Q as a pullback again. We can again conclude that $f_r(a) \overline{\leqslant} \alpha$. □

Corollary 4.8. Every $\oplus_r : T\mathscr{L} \longrightarrow \mathscr{L}$ is a left adjoint relative to $\mathbb{P}_\omega$. Every $\ominus_r : T\mathscr{L} \longrightarrow \mathscr{L}$ is a right adjoint relative to $\mathbb{P}_\omega$.

Proof. By Lemma 4.6 we have $\oplus_r a \equiv \nabla f_r(a)$ for every formula a. Since both f_r and ∇ are left adjoints relative to $\mathbb{P}_\omega$, so is their composite.

To prove that $\ominus_r$ is a right adjoint relative to $\mathbb{P}_\omega$, we suppose $c \leqslant \ominus_r a$, but by Lemma 4.2, this is equivalent to $\oplus_{r'} d(a) \leqslant d(c)$. We use $G : \langle \mathscr{L}, \leqslant \rangle \longrightarrow \mathbb{P}_\omega \langle \mathscr{L}, \leqslant \rangle$ to denote the left relative adjoint of $\oplus_{r'}$. Then $\oplus_{r'} d(a) \leqslant d(c)$ is equivalent to $d(a) \leqslant b$ for some $b \in G(d(c))$, and the latter is equivalent to $d(b) \leqslant a$ by Lemma 2.8. Thus the desired right adjoint of $\ominus_r$ can be defined as $H(c) = \{d(b) \mid b \in G(d(c))\}$. □

We can show even more in that $\oplus_r$ and $\ominus_r$ are in fact *both* left and right adjoints relative to $\mathbb{P}_\omega$. And, any other monotone unary modality is a *left* adjoint relative to $\mathbb{P}_\omega$. To show this, we need the following result.

Lemma 4.9.

(1) Let $\odot : T\mathscr{L} \longrightarrow \mathscr{L}$ be given such that $\odot a \equiv \bigvee_{i=1}^{k} \oplus_{r_i} a$ for some $r_1, \ldots, r_k \in T2$. Then $\odot$ is a left adjoint relative to $\mathbb{P}_\omega$.

(2) Let $\oslash : T\mathscr{L} \longrightarrow \mathscr{L}$ be given such that $\oslash a \equiv \bigwedge_{i=1}^{k} \ominus_{r_i} a$ for some $r_1, \ldots, r_k \in T2$. Then $\odot$ is a right adjoint relative to $\mathbb{P}_\omega$.

Proof. We will just prove the first assertion since the second will follow by duality. We use $G_{r_i} : \mathscr{L} \longrightarrow P_\omega T\mathscr{L}$ to denote the right adjoints of $\oplus_{r_i}$ relative to $\mathbb{P}_\omega$, $i = 1, \ldots, k$. Then $\odot\alpha \leqslant a$ holds if and only if $\oplus_{r_i}\alpha \leqslant a$ holds if and only if $\alpha \overline{\leqslant} \beta$ holds for some $i = 1, \ldots, k$ and some $\beta \in G_{r_i}(a)$. It is then enough to put

$$G(a) = \bigcup_{i=1}^{k} G_{r_i}(a)$$

to define a left adjoint to $\odot$ relative to $\mathbb{P}_\omega$. □

The following corollary now follows immediately by Lemma 4.2.

Corollary 4.10.

(1) Every $\oplus_r : T\mathscr{L} \longrightarrow \mathscr{L}$ is a left and a right adjoint relative to $\mathbb{P}_\omega$.

(2) Every $\ominus_r : T\mathscr{L} \longrightarrow \mathscr{L}$ is a right and a left adjoint relative to $\mathbb{P}_\omega$.
(3) Every monotone $\odot : T\mathscr{L} \longrightarrow \mathscr{L}$ is a left adjoint relative to $\mathbb{P}_\omega$.

5. Conclusions

We have proved that for any finitary standard functor $T : \mathsf{Set} \longrightarrow \mathsf{Set}$ preserving weak pullbacks and finite sets, the corresponding nabla modality is a left adjoint relative to $\mathbb{P}_\omega$ and, analogously, the dual delta modality is a right adjoint relative to $\mathbb{P}_\omega$. This is not only true for the propositional part of the logic being Boolean, but also for the positive fragment of the logic. These results generalise those of Santocanale and Venema (2007) for the finitary powerset functor. We used the results for nabla and delta to show that all unary monotone modalities are also left adjoints relative to $\mathbb{P}_\omega$. Moreover, we have identified those unary monotone modalities that are both left and right adjoints relative to $\mathbb{P}_\omega$. Left or right adjointness relative to $\mathbb{P}_\omega$ entails, for all the modalities we consider, the preservation of directed meets or joins, respectively.

We have argued that the choice of $\mathbb{P}_\omega$ as the doctrine gives the second best level of adjointness when dealing with finitary languages. We propose that the proper doctrine for the λ-ary fragment of the full infinitary logic (where λ is a regular uncountable cardinal) would be the doctrine $\mathbb{P}_\lambda$ of free semilattices with $< \lambda$-joins.

We feel that another place where the choice of the doctrine $\mathbb{P}_\omega$ is relevant is proof theory. Adjointness relative to $\mathbb{P}_\omega$ signifies the possible existence of a 'nice sequent rule' for the modality under consideration in a Gentzen-style proof system. Nice here means: the rule is invertible only in a weaker sense, that is, if we consider the rule backwards, there is a finite set of candidates for a possible prolongation of a proof search. This relates to the ongoing work of the first and third authors on proof theory for coalgebraic logics following the results of Bílková *et al.* (2008).

Remark 5.1. As promised in the introduction, we will now discuss the relation between our results and those in Schröder and Venema (2010) in a bit more detail. Recall that O-adjointness means adjointness relative to $\mathbb{P}_\omega$.

Working in the setting of modalities obtained through (monotone!) predicate liftings, and relative to a 'nice' proof system, Schröder and Venema proved as their main technical result on O-adjointness that every formula that is monotone and *uniform* in a variable x (roughly meaning that all occurrences of x are at the same depth) gives rise to an operation that is an O-adjoint. Using the well-known closure properties of O-adjoints (Santocanale 2007), they arrived at the result that all so-called *admissible* formulas provide O-adjoints. Note that there is no restriction on the functor here.

The differences and similarities between the two papers can be described by the following two observations:

(1) First, the fact that the results of Schröder and Venema (2010) are modulo a 'nice' proof system is not a restriction since every coalgebraic logic admits a sound and complete system of this form Schröder (2007, Theorem 18). Thus Schröder and Venema (2010) covers a wide range of coalgebraic logics, since they impose no restrictions on the functor such as the preservation of weak pullbacks or finite sets.

(2) Second, once we make the above restrictions, there are mutual translations in our Boolean setting between the nabla language and the language of all predicate liftings: in fact, Kurz and Leal (2009) showed that all expressive coalgebraic logics for a finitary functor that preserves weak pullbacks and finite sets are mutually translatable. However, it is not immediately clear whether the existence of these translations implies that our results can be derived from those in Schröder and Venema (2010), or *vice versa* (if we were to consider the semantically defined version of a logic based on predicate liftings). The precise connection between the two approaches remains a task for further research.

Apart from the fact that we obtained our results on the nabla modality just *before* the results of Schröder and Venema (2010), the main value of our results here, relative to Schröder and Venema (2010), is that we provide a direct, semantics-oriented proof for the O-adjointness of the nabla-modalities.

Acknowledgement

We would like to thank the referee for suggestions that improved the final version of the paper.

References

Adámek, J. (1983) *Theory of mathematical structures*, Reidel Publications.

Bílková, M., Palmigiano, A. and Venema, Y. (2008) Proof systems for the coalgebraic cover modality. In: Areces, C. and Goldblatt, R. (eds.) *Advances in Modal Logic*, Volume 7, College Publications 1–21.

Bonsangue, M. M. and Kurz, A. (2005) Duality for logics of transition systems. In: Sassone, V. (ed.) Foundations of Software Science and Computation Structures 2005. *Springer-Verlag Lecture Notes in Computer Science* **3441** 455–469.

Karazeris, P. and Velebil, J. (2009) Representability relative to a doctrine. *Cah. Topol. Géom. Différ. Catég*. **L-1** 3–22.

Kissig, C. and Venema, Y. (2009) Complementation of Coalgebra Automata. In: Kurz, A. *et al.* (eds.) Algebra and Coalgebra in Computer Science (CALCO 2009). *Springer-Verlag Lecture Notes in Computer Science* 5728 81–96.

Kupke, C., Kurz, A. and Venema, Y. (2008) Completeness of the finitary Moss logic. In: Areces, C. and Goldblatt, R. (eds.) *Advances in Modal Logic*, Volume 7, College Publications 193–217.

Kurz, A. and Leal, R. (2009) Equational coalgebraic logic. In: Abramsky, S., Mislove, M. and Palamidessi, C. (eds.) MFPS XXV. *Electronic Notes in Theoretical Computer Science* **249** 333–356.

Moss, L. (1999) Coalgebraic logic. *Ann. Pure Appl. Logic* **96** 277–317. (Erratum published in *Ann. Pure Appl. Logic* **99** 241–259.)

Pattinson, D. (2003) Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science* **309** 177–193.

Pavlović, D., Mislove, M. and Worrell, J. B. (2006) Testing semantics: Connecting processes and process logics. In: Johnson, M. and Vene, V. (eds.) Algebraic Methodology and Software Technology 2006 – Proceedings 11th International Conference, AMAST 2006. *Springer-Verlag Lecture Notes in Computer Science* **4019** 308–322.

Santocanale, L. (2007) Completions of μ-algebras. *Ann. Pure Appl. Logic* **154** 1 27–50.
Santocanale, L. and Venema, Y. (2007) Completeness for flat modal fixpoint logic (extended abstract). In: Dershowitz, N. and Voronkov, A. (eds.) Proceedings of the 14th International Conference on Logic for Programming, Artificial Inteligence and Reasoning (LPAR 2007). *Springer-Verlag Lecture Notes in Computer Science* **4790** 394–408.
Schröder, L. (2007) A finite model construction for coalgebraic modal logic. *Log. Algebr. Program.* **73** 97–110.
Schröder, L. and Venema, Y. (2010) Flat coalgebraic fixed point logics. In: Gastin, P. and Laroussinie, F. (eds.) Proceedings of 21st International Conference on Concurrency Theory, CONCUR 2010. *Springer-Verlag Lecture Notes in Computer Science* **6269** 524–538.
Stoltenberg-Hansen, V., Lindström, I. and Griffor, E. G. (1994) *Mathematical theory of domains*, Cambridge Tracts in Theoretical Computer Science **22**, Cambridge University Press.
Tholen, W. (1984) Pro-categories and multiadjoint functors. *Can. J. Math.* **36** 144–155.
Trnková, V. (1969) Some properties of set functors. *Comment. Math. Univ. Carolin.* **10** 323–352.
Venema, Y. (2006) Automata and fixed point logics: A coalgebraic perspective *Inform. and Comput.* **204** 637–678.