Relation Lifting and Coalgebraic Logic

MSc Thesis *(Afstudeerscriptie)*

written by

Ezra Schoen
(born April 17, 1997 in Utrecht, the Netherlands)

under the supervision of **prof. dr. Yde Venema**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
June 29, 2021	prof. dr. Yde Venema
	dr. Benno van den Berg
	dr. Alexandru Baltag
	dr. Jurriaan Rot

Institute for Logic, Language and Computation

Abstract

We study relation lifting in the context of universal coalgebra. In particular, we develop a family of logics based on the cover modality.

Firstly, we prove a Hennessy-Milner-style theorem, showing that on finite-branching coalgebras, logical equivalence coincides with a particular form of bisimulation. We also give a characterization of those formulas preserved under simulations.

Secondly, we present a sound and complete cut-free sequent calculus, and use it to derive sound and complete cut-free sequent calculi for modal logic and monotone modal logic.

Contents

1 Introduction

In recent years, the categorical framework of coalgebra has gained importance as a uniform way to model various kinds of state-based evolving systems. The theory of coalgebras has close ties to automata theory and modal (fixpoint) logics. Additionally, coalgebra has many applications in theoretical computer science; including concurrency theory, formal verification and semantics of programming languages.

While each of these fields have their own particular methods, they all feature some notion of behavioral equivalence between systems. The most powerful tool for establishing equivalence is by exhibiting a bisimilarity relation. Given the universality of bisimulations in coalgebraic systems, an important task for coalgebraists is to formulate a general theory of bisimulation and behavioral equivalence. An elegant way of formulating bisimulation is via *relation lifting*

Relation lifting has its roots in relational algebra [2]. At its core, relation lifting pertains to ways of 'lifting' a relation R between two sets X and Y to a relation between TX and TY, for a given coalgebraic type T. From this, we may call a relation R a bisimulation if, whenever two states are related by R, their unfoldings are related by the lift of R.

In many settings, bisimilarity between coalgebras of type T is captured exactly by a canonical relation lifting, called the *Barr* lifting. Indeed, the Barr lifting is often taken to *define* bisimilarity in universal coalgebra[20]. However, the Barr lifting only captures behavioral equivalence for functors preserving weak pullbacks. Most relational structures are coalgebras for functors of this type, but e.g. neighborhood-like functors often don't preserve weak pullbacks.

This has prompted study into relation lifting for functors not (necessarily) preserving weak pullbacks. In [18], a characterization is given of relation liftings that capture behavioral equivalence. There is also work on using relation lifting to capture coalgebraic *simulation*, rather than bisimulation [13][22]. Relation liftings have also been used to capture relationships between coalgebras weaker than behavioral equivalence [7].

On the side of coalgebraic logic, relation lifting is used to define the *nabla-modality* ∇. The specific semantics of the ∇-modality on Kripke frames was already implicitly present in work by K. Fine in modal logic [9], and explicit in work by Janin and Walukiewicz on the modal μ-calculus [14]. The general formulation using relation lifting was first given by L. Moss [19].

Since then, work on the ∇-modality has largely stuck to the Barr lifting (an early exception is [1]). The resulting logical system is certainly elegant; but given the diversity in possible relation liftings, much may be gained from moving to a more general setting. Moreover, since the Barr lifting only captures bisimilarity for functors that preserve weak pullbacks, the scope of the ∇-modality has been somewhat limited. We are also motivated by [21], where a ∇-modality is used based on the lifting $\widetilde{M}$ for the monotone neighborhood functor $\mathcal{M}$, which is not a weak pullback-preserving functor.

This thesis expands on work by A. Baltag [1] and J. Marti & Y. Venema [18], where modalities are defined based on arbitary relation liftings. We will explore these modalities in detail, both from a model-theoretic perspective as well as a proof-theoretic perspective. We will see that, remarkably, many of the standard results for the ∇-modality also hold in the more general setting. Our main results are the following:

- We show that every functor admits a minimal lifting; and that the lifting $\widetilde{M}$ arises in a natural way as the minimal lifting for the monotone neighborhood functor.
- We establish an alternative characterization of liftings in terms of weak distributive laws.
- We prove that the ∇-modality based on a lifting L fully captures L-bisimulation. We also show that in suitable circumstances, a formula is preserved under L-simulations if and only if it is equivalent to a formula featuring only the L-based nabla.
- We present a uniform family of sequent calculi which is sound and complete for logics involving any number of ∇-modalities. Moreover, we show that in suitable circumstances, the resulting sequent calculi are decidable.
- We modify the sequent calculi for the ∇-modalities to derive sound and complete cut-free sequent calculi for modal logic and monotone modal logic.

This thesis is divided into the following chapters:

1. **Introduction** Introduction and motivation.
2. **Preliminaries** This chapter consists of a collection of definitions and examples from category theory and universal coalgebra.
3. **Relation lifting** In this chapter, we define the notion of a T-lifting and give key lemmas regarding the behavior of liftings. We also give a characterization of liftings in terms of weak distributive laws.
4. **Coalgebraic logic** In this chapter, we introduce the modalities based on relation liftings, and prove a number of lemmas and propositions needed in the next two chapters.
5. **(Relative) expressivity of the ∇-modalities** In this chapter, we explore the connections between logical equivalence and bisimulation (invariance).
6. **A uniform sequent calculus** We present a uniform family of sequent calculi to one which is sound and complete for logics involving any number of ∇-modalities. We also highlight some interesting fragments for which the calculus is sound and complete.
7. **Conclusion** Summary of results and avenues for further research.

2 Preliminaries

We assume that the reader is familiar with the language of categories and functors, as well as the basic theory of coalgebras. This chapter serves only to fix some notation, used throughout the thesis.

2.1 Preliminaries on functors

In this section, we introduce a number of important functors, as well as some preservation properties of functors.

Notation 2.1. We will write **Sets** for the category of sets and functions.

Notation 2.2. For a functor $T : \mathbf{Sets} \to \mathbf{Sets}$, a function $f : X \to Y$ and an element $\alpha \in TX$, we will usually add brackets as

$$(Tf)\alpha$$

although we may in some cases write $(Tf)(\alpha)$ to avoid confusion.

Some important functors

Definition 2.3. We will denote the powerset functor with P. Explicitly, if $f : X \to Y$ is a function, then $Pf : PX \to PY$ is defined as

$$Pf : A \mapsto f[A] = \{f(a) \mid a \in A\}.$$

This makes the powerset into a covariant functor.

The powerset also has a contravariant version, which we will denote with $\breve{P}$. That is, $\breve{P}X = PX$ for all sets X; but if $f : X \to Y$ is a function, we set $\breve{P}f : PY \to PX$ to be

$$\breve{P}f : B \mapsto f^{-1}[B] = \{a \mid f(a) \in B\}$$

Two more functors of interest are the *neighborhood functor* and the *monotone neighborhood functor*.

Definition 2.4. We define the *neighborhood functor* to be the functor $\mathcal{N} := \breve{P}\breve{P}$. Since $\breve{P}$ is contravariant, $\mathcal{N}$ is a covariant functor. Explicitly, we have $\mathcal{N}X = PPX$ for a set X, and if $f : X \to Y$ is a function, then

$$\mathcal{N}f : \mathcal{A} \mapsto \{U \mid (\breve{P}f)U \in \mathcal{A}\}.$$

The *monotone neighborhood functor* is the subfunctor $\mathcal{M}$ of $\mathcal{N}$ given by

$$\mathcal{M}X = \{\mathcal{A} \in \mathcal{N}X \mid \forall U, V : \text{ if } U \in \mathcal{A} \text{ and } U \subseteq V, \text{ then } V \in \mathcal{A}\}.$$

It is easy to verify that if $f : X \to Y$ is a function, and $\mathcal{A} \in \mathcal{M}X$, then $\mathcal{N}f(\mathcal{A}) \in \mathcal{M}Y$; so, if we set

$$\mathcal{M}f = \mathcal{N}f\restriction_{\mathcal{M}X}$$

we obtain a well-defined functor $\mathcal{M}$.

We will regularly return to the monotone neighborhood functor as an important example. It will be useful to have simple notation for elements of $\mathcal{M}X$. We will write

$$\langle U_1, U_2, \ldots, U_k \rangle = \{U \in PX \mid U \supseteq U_i \text{ for some } i\}$$

for the upset generated by $\{U_1, \ldots, U_k\}$.

Notation 2.5 (Naming convention). In this thesis, we will come across elements of sets X, but also of elements of TX for T a functor, PX, as well as TPX and PTX. To avoid confusion as much as possible, we use the following conventions for elements of particular sets, following [5]:

Set	Elements
Proposition letters	$p, q, \ldots$
Formulas	$a, b, \ldots$
TX	$\alpha, \beta, \ldots$
PX	$A, B, \ldots$
PPX	$\mathcal{A}, \mathcal{B}, \ldots$
TPX	$\Phi, \Psi, \ldots$
PTX	$\Gamma, \Theta, \ldots$

Figure 1

An important notion is that of a *finitary functor*.

Definition 2.6. A functor $T : \mathbf{Sets} \to \mathbf{Sets}$ is *finitary* if for sets X, we have

$$TX = \bigcup_{\substack{X' \subseteq X \\ X' \text{ finite}}} \operatorname{im}(T\iota_{X'})$$

where $\iota : X' \to X$ denotes the inclusion map.

For an arbitrary functor $T : \mathbf{Sets} \to \mathbf{Sets}$, we can define a finitary version T_ω as

$$T_\omega = \bigcup_{\substack{X' \subseteq X \\ X' \text{ finite}}} \operatorname{im}(T\iota_{X'}),$$

where for a function $f : X \to Y$ we set

$$T_\omega f = Tf\restriction_{T_\omega X} .$$

Remark 2.7. In the case that T preserves inclusions - that is, if $X \subseteq Y$, then $TX \subseteq TY$ - definition 2.6 reduces to

$$T_\omega X = \bigcup_{\substack{X' \subseteq X \\ X' \text{ finite}}} TX'.$$

This is the case for the powerset functor, but not for, e.g., the monotone neighborhood functor.

Definition 2.8. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, let X be a set, and let $\alpha \in T_\omega X$. We define

$$\operatorname{Base}(\alpha) := \bigcap \{X' \subseteq X \mid \alpha \in \operatorname{im}(T\iota_{X'})\}.$$

Then $\operatorname{Base}(\alpha)$ is the smallest subset X' of X with the property $\alpha \in \operatorname{im}(T\iota_{X'})$.

If $\Gamma \subseteq TX$, we define

$$\mathcal{B}(\Gamma) := \bigcup_{\alpha \in \Gamma} \operatorname{Base}(\alpha).$$

We prove explicitly that the functor T_ω defined in definition 2.6 is indeed a functor, by deriving a useful 'naturality property' for Base.

Proposition 2.9. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, let X, Y be sets, and $f : X \to Y$ a function. Then for all $\alpha \in T_\omega X$, we have*

$$\operatorname{Base}((Tf)\alpha) \subseteq f[\operatorname{Base}(\alpha)]$$

Proof. Let $B = \operatorname{Base}(\alpha)$. Then let $B' = f[B]$, and let $g : B \to B'$ be the restriction of f to B. We see that

$$\begin{tikzcd} X \arrow[r, "f"] & Y \\ B \arrow[u, "\iota_B"] \arrow[r, "g"] & B' \arrow[u, "\iota_{B'}"'] \end{tikzcd}$$

commutes. Hence, if we apply T to this diagram, we see that

$$\begin{tikzcd} TX \arrow[r, "Tf"] & TY \\ TB \arrow[u, "T\iota_B"] \arrow[r, "Tg"] & TB' \arrow[u, "T\iota_{B'}"'] \end{tikzcd}$$

commutes. Now since $B = \operatorname{Base}(\alpha)$, we know that there is $\alpha' \in TB$ with $(T\iota_B)\alpha' = \alpha$. So, by commutativity, we see that

$$(Tf)\alpha = (T\iota_{B'})((Tg)\alpha') \in \operatorname{im}(T\iota_{B'})$$

which means that $\operatorname{Base}((Tf)\alpha) \subseteq B'$ by definition. $\square$

Note that Base is not a full natural transformation from T_ω to P_ω, since only one of the two inclusions is present. In [15], it is proved that for a weak pullback-preserving functor, Base *is* a natural transformation.

Remark 2.10. There is a useful criterion for Base in the case of the monotone neighborhood functor. Namely, if $p \notin \operatorname{Base}(\mathcal{U})$, then for all $U \in \mathcal{U}$, we have

$$U \in \mathcal{U} \text{ if and only if } U \setminus \{p\} \in \mathcal{U}.$$

That is, adding or removing p does not change whether U is an element of $\mathcal{U}$.

One notion which is common in coalgebraic logic is that of a *weak pullback*.

Definition 2.11. Let $\begin{tikzcd} P \arrow[r, "p_0"] \arrow[d, "p_1"] & X \arrow[d, "f"] \\ Y \arrow[r, "g"] & Z \end{tikzcd}$ be a commutative square in **Sets**. We call (P, p_0, p_1) a *weak pullback of f and g* if for every commutative square $\begin{tikzcd} P' \arrow[r, "p_0'"] \arrow[d, "p_1'"] & X \arrow[d, "f"] \\ Y \arrow[r, "g"] & Z \end{tikzcd}$ there is a (not necessarily unique!) map $m : P' \to P$ such that

$$\begin{tikzcd} P' \arrow[dr, "m"] \arrow[drr, bend left] \arrow[ddr, bend right] & & \\ & P \arrow[r, "p_0"] \arrow[d, "p_1"] & X \arrow[d, "f"] \\ & Y \arrow[r, "g"] & Z \end{tikzcd}$$

commutes.

A functor T : **Sets** $\to$ **Sets** is said to *preserve weak pullbacks* if, whenever $\begin{tikzcd} P \arrow[r, "p_0"] \arrow[d, "p_1"] & X \arrow[d, "f"] \\ Y \arrow[r, "g"] & Z \end{tikzcd}$ is a weak pullback square, so is $\begin{tikzcd} TP \arrow[r, "Tp_0"] \arrow[d, "Tp_1"] & TX \arrow[d, "Tf"] \\ TY \arrow[r, "Tg"] & TZ \end{tikzcd}$

Definition 2.12. Let T : **Sets** $\to$ **Sets** be a functor. We say that T *preserves finite sets* if TX is finite for all finite X.

2.2 Coalgebras

In this section we define T-coalgebras for a functor T, as well as behavioral equivalence for T-coalgebras.

Definition 2.13. Let T : **Sets** $\to$ **Sets** be a functor. A *T-coalgebra* is a set S together with a map $\sigma : S \to TS$. For a given set of proposition letters Prop, a *T-coalgebra model with values in* Prop, or simply a *T-coalgebra model* in case the set Prop is clear from context, is a T-coalgebra $\mathbb{S} = (S, \sigma)$ together with a map $m : S \to P(\mathsf{Prop})$.

Example 2.14. A P-coalgebra is simply a Kripke frame, presented in an unusual way. Namely, let $\mathfrak{F} = (W, R)$ be a Kripke frame, with W a set of worlds and $R \subseteq W \times W$ be an accessibility relation. Then we can define a P-coalgebra $\mathfrak{F}^\sharp = (W, \chi_R)$, where $\chi_R : W \to PW$ is defined as

$$\chi_R(x) = \{y \mid xRy\}.$$

Vice versa, if $\mathbb{S} = (S, \sigma)$ is a P-coalgebra, then we can define $\mathbb{S}^\flat = (S, \sigma^\ni)$, with

$$\sigma^\ni = \{(x, y) \mid y \in \sigma(x)\}.$$

This correspondence extends to P-coalgebra models and Kripke models. Let $\mathfrak{M} = (\mathfrak{F}, V)$ be a Kripke model, with $\mathfrak{F} = (W, R)$ a Kripke frame, and $V : \mathsf{Prop} \to PW$ a valuation. Then we can define $\mathfrak{M}^\sharp$ as $(\mathfrak{F}^\sharp, V^\sharp)$, with

$$V^\sharp(x) = \{p \in \mathsf{Prop} \mid x \in V(p)\}.$$

Vice versa, if $\mathbb{S} = (S, \sigma, m)$ is a P-coalgebra model, we can set $\mathbb{S}^\flat = (S, \sigma^\ni, m^\flat)$, with

$$m^\flat(p) = \{x \in S \mid p \in m(s)\}.$$

Definition 2.15. Let $\mathbb{S} = (S, \sigma)$ and $\mathbb{S}' = (S', \sigma')$ be T-coalgebras. A *coalgebra morphism from* $\mathbb{S}$ *to* $\mathbb{S}'$ is a function $f : S \to S'$ such that $\sigma' \circ f = Tf \circ \sigma$.

Let $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S}' = (S', \sigma', m')$ be T-coalgebra models. A *morphism of coalgebra models from* $\mathbb{S}$ *to* $\mathbb{S}'$ is a coalgebra morphism $f : \mathbb{S} \to \mathbb{S}'$ such that $m' \circ f = m$.

In cases where there can be no confusion, we will also use the term 'coalgebra morphism' for a morphism of coalgebra models.

Remark 2.16. Note that a morphism of P-coalgebras is the same as a bounded morphism between Kripke frames.

Definition 2.17. Let $\mathbb{S} = (S, \sigma)$ and $\mathbb{S}' = (S', \sigma')$ be T-coalgebras. For a given $s \in S, s' \in S'$, we call s and s' *behaviorally equivalent* if there is a T-coalgebra $\mathbb{Z} = (Z, \zeta)$, together with T-coalgebra morphisms $f : \mathbb{S} \to \mathbb{Z}$ and $g : \mathbb{S}' \to \mathbb{Z}$ such that $f(s) = g(s')$.

If s and s' are behaviorally equivalent, we write $\mathbb{S}, s \simeq \mathbb{S}', s'$.

An important class of coalgebras is given by those that are finite branching.

Definition 2.18. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let $\mathbb{S} = (S, \sigma)$ be a T-coalgebra. We call $\mathbb{S}$ *finite branching* if for all $s \in S$, we have $\sigma(s) \in T_\omega S$.

We call a T-coalgebra model $\mathbb{S} = (S, \sigma, m)$ *finite branching* if (S, σ) is finite branching.

3 Relation liftings

In this chapter, we define relation liftings and give some examples for particular functors. We show that any functor T admits a minimal lifting, and give an explicit description of the minimal lifting for the monotone neighborhood functor. Finally, we give an alternate characterization of liftings in terms of weak distributive laws.

3.1 Relations and liftings

The fundamental notion of this thesis is *relation lifting*. Before defining this, we first discuss the category of relations.

Definition 3.1. The category **Rel** of relations is defined as follows:

- Its objects are the sets.
- A morphism from X to Y is a relation from X to Y; that is, a subset of $X \times Y$.
- If R is a morphism from X to Y and S a morphism from Y to Z, then their composition is given by

$$R;S := \{(x,z) \mid \text{ there exists a } y \in Y \text{ such that } xRySz\}.$$

(Note the order in which we write composition of relations.) We will denote a morphism R in $\mathrm{Hom}_{\mathbf{Rel}}(X,Y)$ as $R : X \multimap Y$. We will denote the identity relation on X with $\Delta_X = \{(x,x) \mid x \in X\}$, and call it the *diagonal*.

The category **Rel** has a lot of structure. We note the following facts:

Fact 3.2.
1. **Rel** is enriched over the category of posets. That is, for every X the set $\mathrm{Hom}_{\mathbf{Rel}}(X,Y)$ is a partial order under the $\subseteq$-relation; and if $R, R' : X \multimap Y$ and $S, S' : Y \multimap Z$ with $R \subseteq R'$ and $S \subseteq S'$, then $R;S \subseteq R';S'$.
2. **Rel** comes equipped with an operation $(-)^\circ : \mathbf{Rel} \to \mathbf{Rel}$, given by

$$R^\circ = \{(y,x) \mid (x,y) \in R\}.$$

   This operation satisfies

$$(R^\circ)^\circ = R, \qquad (R;S)^\circ = S^\circ;R^\circ.$$

3. There is a canonical functor $(-)^{\mathrm{gr}} : \mathbf{Sets} \to \mathbf{Rel}$ defined as

$$X^{\mathrm{gr}} := X, \qquad (f : X \to Y)^{\mathrm{gr}} := \{(x, f(x)) \mid x \in X\}.$$

We now come to the central definition.

Definition 3.3. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be any functor. A *T-lifting* is an assignment L that associates to each relation $R : X \multimap Y$ a relation $LR : TX \multimap TY$, satisfying the following properties:

1. If $R \subseteq S$, then $LR \subseteq LS$.

2. For all $R : X \multimap Y$ and $S : Y \multimap Z$,
$$LR; LS \subseteq L(R; S).$$

3. If $f : X \to Y$ is a function, then
$$(Tf)^{\mathrm{gr}} \subseteq L(f^{\mathrm{gr}}) \quad \text{and} \quad ((Tf)^{\mathrm{gr}})^{\circ} \subseteq L((f^{\mathrm{gr}})^{\circ})$$

A lifting L is said to be *symmetric* if it satisfies

4. For all relations $R : X \to Y$,
$$L(R^{\circ}) = (LR)^{\circ}.$$

L is said to *preserve diagonals* if it satisfies

5. for all sets X,
$$L\Delta_X \subseteq \Delta_{TX}.$$

T-liftings in this form were first introduced in [22], where they were called *monotonic relators.*

There are a number of examples of T-liftings, for various functors T.

Example 3.4.

1. Any functor $T : \mathbf{Sets} \to \mathbf{Sets}$ admits the trivial lifting $\top_T$, which maps a relation $R : X \multimap Y$ to the maximal relation $TX \times TY$. $\top_T$ is symmetric but does not in general preserve diagonals.

2. Let $C_Q : \mathbf{Sets} \to \mathbf{Sets}$ be the constant functor with value Q. That is, for any set X we set $C_Q(X) = Q$, and for any function $f : X \to Y$ we set $C_Q(f) = \mathrm{id}_Q$. Then any preorder on Q gives rise to a corresponding C_Q-lifting, in the following way: if $\preceq$ is a preorder on Q, then
$$L_{\preceq} R := \preceq$$
defines a C_Q-lifting $L_{\preceq}$.[1]

   For, the condition that $(C_Q f)^{\mathrm{gr}} \subseteq L_{\preceq}(f^{\mathrm{gr}})$ simply states that $\Delta_Q \subseteq \preceq$, which is equivalent to $\preceq$ being reflexive; and the condition that $LR; LS \subseteq L(R; S)$ is equivalent to $\preceq; \preceq \subseteq \preceq$, which simply states that $\preceq$ is transitive.

   It is easy to see that $L_{\preceq}$ is symmetric if and only if $\preceq$ is an equivalence relation, and preserves diagonals if and only if $\preceq = \Delta_Q$.

[1] In fact, it can be shown that any C_Q-lifting is of the form $L_{\preceq}$ for some preorder $\preceq$.

3. Let $T_2 = (-)^2$ be the ordered-pair functor. Then for a relation $R : X \multimap Y$, we set
$$\overline{T}_2 R = \{(\langle x, y\rangle, \langle x', y'\rangle) \mid (x, x') \in R, (y, y') \in R\}$$
This defines a T_2-lifting $\overline{T}_2$ which is symmetric and preserves diagonals.

4. The powerset functor has two natural liftings $\overrightarrow{P}$ and $\overleftarrow{P}$. For a relation $R : X \multimap Y$, we set
$$\begin{aligned} \overrightarrow{P}R &:= \{(U, V) \mid \forall u \in U \exists v \in V : uRv\}, \\ \overleftarrow{P}R &:= \{(U, V) \mid \forall v \in V \exists u \in U : uRv\}. \end{aligned}$$
We also have the lifting $\overline{P}$ defined as $\overline{P}R = \overrightarrow{P}R \cap \overleftarrow{P}R$. $\overrightarrow{P}$ and $\overleftarrow{P}$ are neither symmetric nor diagonal-preserving, but $\overline{P}$ is both.

5. The monotone neighborhood functor has two natural liftings $\overset{\rightleftarrows}{\mathcal{M}}$ and $\overset{\leftrightarrows}{M}$. For a relation $R : X \multimap Y$, we set $\overset{\rightleftarrows}{\mathcal{M}}R := \overrightarrow{P}\overleftarrow{P}R, \overset{\leftrightarrows}{\mathcal{M}}R := \overleftarrow{P}\overrightarrow{P}R$. Explicitly, this means that
$$\begin{aligned} (\mathcal{U}, \mathcal{V}) \in \overset{\rightleftarrows}{\mathcal{M}}R &\text{ iff } \forall U \in \mathcal{U} \exists V \in \mathcal{V} : \forall v \in V \exists u \in U : uRv, \\ (\mathcal{U}, \mathcal{V}) \in \overset{\leftrightarrows}{\mathcal{M}}R &\text{ iff } \forall V \in \mathcal{V} \exists U \in \mathcal{U} : \forall u \in U \exists v \in V : uRv. \end{aligned}$$
As in the case for the powerset, we have $\widetilde{\mathcal{M}}R := \overset{\rightleftarrows}{\mathcal{M}}R \cap \overset{\leftrightarrows}{\mathcal{M}}R$. $\overset{\leftrightarrows}{\mathcal{M}}$ and $\overset{\rightleftarrows}{\mathcal{M}}$ are not symmetric and do not preserve diagonals, but $\widetilde{\mathcal{M}}$ is symmetric and preserves diagonals.

3.2 Barr lifting

Any weak pullback-preserving functor comes with a ‘canonical’ lifting.

Definition 3.5 (Barr lifting)**.** Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor. For a relation $R : X \multimap Y$, we define $\overline{T} : TX \multimap TY$
$$\overline{T}R := \{(\alpha, \beta) \mid \exists \gamma \in TR \text{ such that } (T\pi_0)\gamma = \alpha, (T\pi_1)\gamma = \beta\}$$
where $\pi_0 : R \to X$ and $\pi_1 : R \to Y$ are the natural projection maps.

This definition was first given in [2]. It is a general fact that $\overline{T}$ is a lifting if and only if T preserves weak pullbacks (see the overview article [16] for a proof; we will give the ‘if’-direction in this thesis as well). In example 3.4, we have seen two instances of the Barr lifting: the lifting $\overline{T}_2$ for the ordered-pair functor T_2, and the lifting $\overline{P}$ for the powerset functor.

Most of the existing literature on relation lifting for coalgebraic logic focuses on the Barr lifting. This is because the Barr lifting satisfies stronger versions of the properties for T-liftings.

Proposition 3.6. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a weak pullback-preserving functor.*

(i) For all functions f, we have $\overline{T} f^{gr} = (Tf)^{gr}$;

(ii) For all relations $R : X \multimap Y$ and $S : Y \multimap Z$, we have

$$\overline{T}(R;S) = \overline{T}R;\overline{T}S.$$

Proof. (i) We know that for a given relation $R : X \multimap Y$, we have

$$\overline{T}R := (T\pi_X^{\mathrm{gr}})^\circ; (T\pi_Y)$$

with π_X, π_Y the projection functions. But clearly, $\pi_X : f^{\mathrm{gr}} \to X$, which maps $(x, f(x))$ to x, is a bijection, with inverse $\rho : X \to f^{\mathrm{gr}} : x \mapsto (x, f(x))$. Since functors preserves isomorphisms, $T\pi_X$ and $T\rho$ are inverses as well. Hence,

$$\overline{T}R := (T\pi_X^{\mathrm{gr}})^\circ; (T\pi_Y) = (T\rho)^{\mathrm{gr}}; (T\pi_X)^{\mathrm{gr}} = T(\pi_Y \circ \rho)^{\mathrm{gr}} = (Tf)^{\mathrm{gr}}$$

as $\pi_Y \circ \rho = f$.

(ii) We set up the situation. Let X, Y, Z be sets, and let $R : X \multimap Y$ and $S : Y \multimap Z$ be relations. Let $P = \{(x, y, z) \in X \times Y \times Z \mid xRySz\}$. Then consider the following diagram:

$$\begin{tikzcd}
& P \arrow[dl, "\pi_{XY}"'] \arrow[d, "\pi_{XZ}"] \arrow[dr, "\pi_{YZ}"] & \\
R \arrow[d, "\pi_X"] \arrow[dr, "\pi_Y"] & R;S \arrow[dl, "\rho_X"'] \arrow[dr, "\rho_Z"] & S \arrow[dl, "\sigma_Y"'] \arrow[d, "\sigma_Z"] \\
X & Y & Z
\end{tikzcd}$$

where the functions are the obvious projection functions. We note that the 'back' square formed by $PRSY$ is a pullback square.

First, assume that $(\alpha, \beta) \in \overline{T}R;\overline{T}S$. That means that there is a $\gamma \in TY$ with $(\alpha, \gamma) \in \overline{T}R, (\gamma, \beta) \in \overline{T}S$. By definition of $\overline{T}$, that means that there are $\xi \in TR, \eta \in TS$ such that

$$(T\pi_X)\xi = \alpha, \quad (T\pi_Y)\xi = \gamma = (T\sigma_Y)\eta, \quad (T\sigma_Z)\eta = \beta$$

From this, it follows that

$$
\begin{tikzcd}
 & \{*\} \arrow[dl, "* \mapsto \xi"'] \arrow[dr, "* \mapsto \eta"] & \\
TR \arrow[dr, "\pi_Y"'] & & TS \arrow[dl, "\sigma_Y"] \\
 & TY &
\end{tikzcd}
$$

commutes. Since T preserves weak pullbacks, we know that TP is a weak pullback, and so there is a function $\theta : \{*\} \to P$ such that

$$\xi = (T\pi_{XY})\theta, \quad \eta = (T\pi_{YZ})\theta$$

We may identify θ with the element $\theta(*) \in P$. Let $\mu = (T\pi_{XZ})\theta$. Then

$$
\begin{aligned}
\alpha &= (T\pi_X)\xi \\
&= (T\pi_X)(T\pi_{XY})\theta \\
&= (T(\pi_X \circ \pi_{XY}))\theta \\
&= (T(\rho_X \circ \pi_{XZ}))\theta \\
&= (T\rho_X)\mu \\
\beta &= (T\sigma_Z)\eta \\
&= (T\sigma_Z)(T\pi_{YZ})\theta \\
&= (T(\sigma_Z \circ \pi_{YZ}))\theta \\
&= (T\rho_Z)\mu
\end{aligned}
$$

and hence $(\alpha, \beta) \in \overline{T}(R; S)$ by definition.

For the other inclusion, we draw a second diagram. We may pick, for every $(x, z) \in R; S$, a $y_{x,z} \in Y$ such that $xRy_{x,z}Sz$; this defines a function $y : R; S \to Y$. Moreover, we get functions $f : R; S \to R$ and $g : R; S \to S$ given by

$$f(x, z) = (x, y_{x,z}), \quad g(x, z) = (y_{x,z}, z)$$

This yields the following diagram:

$$
\begin{tikzcd}
R \arrow[d, "\pi_X"'] \arrow[dr, "\pi_Y"] & R; S \arrow[l, "f"'] \arrow[r, "g"] \arrow[dl, "\rho_X"'] \arrow[d, "y"] \arrow[dr, "\rho_Z"] & S \arrow[dl, "\sigma_Y"'] \arrow[d, "\sigma_Z"] \\
X & Y & Z
\end{tikzcd}
$$

Now if $(\alpha, \beta) \in \overline{T}(R; S)$, then there is a $\mu \in T(R; S)$ such that

$$(T\rho_X)\mu = \alpha, \quad (T\rho_Z)\mu = \beta$$

Then let $\xi = (Tf)\mu, \eta = (Tg)\mu, \gamma = (Ty)\mu$. Then

$$\begin{aligned}(T\pi_X)\xi &= (T\pi_X)(Tf)\mu \\ &= (T\rho_X)\mu \\ &= \alpha \\ (T\pi_Y)\xi &= (T\pi_Y)(Tf)\mu \\ &= (Ty)\mu \\ &= \gamma\end{aligned}$$

showing that $(\alpha, \gamma) \in \overline{T}R$. By a similar argument, $(\gamma, \beta) \in \overline{T}R$. Hence, $(\alpha, \beta) \in \overline{T}R; \overline{T}S$. □

3.3 Operations on liftings

In this section, we discuss some important ways of obtaining new liftings from old ones.

Definition 3.7. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and L a T-lifting. Then we define the T-lifting $L^{\smallsmile}$ by setting for $R : X \multimap Y$,

$$L^{\smallsmile}R := (L(R^\circ))^\circ$$

It follows quickly from the definitions that $L^{\smallsmile}$ is indeed a T-lifting for every T-lifting L. We note that a lifting L is symmetric if and only if $L = L^{\smallsmile}$.

Remark 3.8. In certain cases, we have separate symbols for R and R°; in particular, $\in$ versus $\ni$ and $\Vdash$ versus $\dashVv$. When dealing with these relations, we will use equalities such as

$$(L^{\smallsmile}\in)^\circ = L \ni$$

without comment.

The T-liftings are closed under arbitrary intersections.

Proposition 3.9. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let Λ be a set of T-liftings. For a relation $R : X \multimap Y$, we set*

$$(\bigwedge \Lambda)R := \bigcap_{L \in \Lambda} (LR).$$

Then $\bigwedge \Lambda$ is a lifting for T.

Proof. Immediate from the definitions. □

We finish with an important lemma.

Lemma 3.10. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let L be a T-lifting.*

(i) For any function $f : X \to Y$ and any relation $R : Y \multimap Z$, we have $L(f^{gr}; R) = (Tf)^{gr}; LR$

(ii) For any relation $R : X \multimap Y$ and any function $f : Z \to Y$, we have $L(R; (f^{gr})^\circ) = LR; ((Tf)^{gr})^\circ$

(iii) Let $R : Y \multimap Z$ be a relation, and $f : X \to Y, g : W \to Z$ be functions. Then

$$L(f^{gr}; R; (g^{gr})^\circ) = (Tf^{gr}); LR; (Tg^{gr})^\circ$$

(iv) Let $R : X \multimap Y$ be a relation, with $X \subseteq X'$ and $Y \subseteq Y'$. Let $R' : X' \multimap Y'$ be equal to R, with only the domain and codomain expanded. Then

$$LR' \supseteq \{((T\iota_X)\alpha, (T\iota_Y)\beta) \mid (\alpha, \beta) \in LR\}$$

Proof. (i) We calculate that

$$\begin{aligned} L(f^{\mathrm{gr}}; R) &= \Delta; L(f^{\mathrm{gr}}; R) \\ &\subseteq Tf^{\mathrm{gr}}; (Tf^{\mathrm{gr}})^\circ; L(f^{\mathrm{gr}}; R) \\ &\subseteq Tf^{\mathrm{gr}}; L((f^{\mathrm{gr}})^\circ); L(f^{\mathrm{gr}}; R) \\ &\subseteq Tf^{\mathrm{gr}}; L((f^{\mathrm{gr}})^\circ; f^{\mathrm{gr}}; R) \\ &= Tf^{\mathrm{gr}}; L(\Delta; R) \\ &= Tf^{\mathrm{gr}}; LR \\ &\subseteq L(f^{\mathrm{gr}}); LR \\ &\subseteq L(f^{\mathrm{gr}}; R) \end{aligned}$$

which shows that all inequalities are equalities, and hence $L(f^{\mathrm{gr}}; R) = Tf^{\mathrm{gr}}; LR$.

(ii) We simply use part (i) to calculate

$$\begin{aligned} L(R; (f^{\mathrm{gr}})^\circ) &= (L^\sim(f^{\mathrm{gr}}; R^\circ))^\circ \\ &= (Tf^{\mathrm{gr}}; L^\sim(R^\circ))^\circ \\ &= LR; (Tf^{\mathrm{gr}})^\circ \end{aligned}$$

(iii) By points (i) and (ii), we have

$$L(f^{\mathrm{gr}}; R; (g^{\mathrm{gr}})^\circ) = L(f^{\mathrm{gr}}; R); (Tg^{\mathrm{gr}})^\circ = (Tf^{\mathrm{gr}}); LR; (Tg^{\mathrm{gr}})^\circ$$

(iv) Note that

$$R' = \{(\iota_X(a), \iota_Y(b)) \mid (a, b) \in R\} = (\iota_X^{\mathrm{gr}})^\circ; R; (\iota_Y^{\mathrm{gr}})$$

and therefore

$$\begin{aligned} LR' &= L\left((\iota_X^{\mathrm{gr}})^\circ; R; (\iota_Y^{\mathrm{gr}})\right) \\ &\supseteq L(\iota_X^{\mathrm{gr}})^\circ; LR; L(\iota_Y^{\mathrm{gr}}) \\ &\supseteq T(\iota_X^{\mathrm{gr}})^\circ; LR; T(\iota_Y^{\mathrm{gr}}) \\ &= \{((T\iota_X)\alpha, (T\iota_Y)\beta) \mid (\alpha, \beta) \in LR\} \end{aligned}$$

$\square$

We will often find ourselves in the following situation:

Example 3.11. Take $\Phi \in T_\omega X$, and let $B = \mathrm{Base}(\Phi)$. Take $\Phi' \in TB$ with $\Phi = (T\iota)\Phi'$. By an inductive argument, we know $(\Phi', \alpha) \in L(R\restriction_B)$, and hence by 3.10(iv) we know that $(\Phi, \alpha) \in LR$.

We will usually omit the intermediate steps, identifying Φ and Φ', and ignoring the distinction between LR and $L(R\restriction_B)$.

3.4 Minimal liftings

Proposition 3.9 implies that any functor T has a minimal lifting $\widetilde{T}$. The following facts follow immediately from the minimality of $\widetilde{T}$.

Proposition 3.12. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor. Let $\widetilde{T}$ be the minimal T-lifting.*

(i) $\widetilde{T}$ is symmetric.

(ii) T admits a diagonal-preserving lifting L if and only if $\widetilde{T}$ preserves diagonals.

Proof. (i) Note that $\widetilde{T}^\sim$ is a T-lifting, so $\widetilde{T} \leq \widetilde{T}^\sim$. But of course, $(-)^\sim$ preserves inclusions, so $\widetilde{T}^\sim \leq (\widetilde{T}^\sim)^\sim = \widetilde{T}$. Hence we have both inequalities, showing $\widetilde{T} = \widetilde{T}^\sim$, which means that T is symmetric.

(ii) The right-to-left direction is trivial. For the left-to-right direction, let L be a T-lifting that preserves diagonals. Then we know that for any set X,

$$\widetilde{T}\Delta_X \subseteq L\Delta_X \subseteq \Delta_{TX}$$

showing that $\widetilde{T}$ preserves diagonals.

$\square$

In general, it is not easy to give an explicit description of the minimal lifting $\widetilde{T}$. In the case of weak pullback-preserving functors, however, we get a fortunate characterization.

Proposition 3.13. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a weak pullback-preserving functor. Then the minimal T-lifting $\widetilde{T}$ is equal to $\overline{T}$.*

Proof. Let L be any T-lifting, and let $R : X \multimap Y$ be a relation. Then R, being a subset of $X \times Y$, comes with two projections $\pi_X : R \to X$ and $\pi_Y : R \to Y$. Then

$$(x, y) \in R \text{ iff } \exists p \in R : \pi_X(p) = x, \pi_Y(p) = y \text{ iff } (x, y) \in (\pi_X^{\mathrm{gr}})^\circ ; (\pi_Y^{\mathrm{gr}}),$$

so $R = (\pi_X^{\mathrm{gr}})^\circ ; (\pi_Y^{\mathrm{gr}})$.

Note also that

$$\overline{T}R = \{(\alpha, \beta) \mid \exists \rho \in TR : T\pi_X(\rho) = \alpha, T\pi_Y(\rho) = \beta\} = (T\pi_X^{\mathrm{gr}})^\circ ; (T\pi_Y^{\mathrm{gr}}).$$

So, we conclude that

$$\begin{aligned} LR &= L((\pi_X^{\mathrm{gr}})^\circ; (\pi_Y^{\mathrm{gr}})) \\ &\supseteq L(\pi_X^{\mathrm{gr}})^\circ; L(\pi_Y^{\mathrm{gr}}) \\ &\supseteq (T\pi_X^{\mathrm{gr}})^\circ; (T\pi_Y^{\mathrm{gr}}) \\ &= \overline{T}R \end{aligned}$$

which shows that $\overline{T}$ is minimal. □

The monotone neighborhood functor The astute reader may recognise the notation $\widetilde{T}$ in the $\mathcal{M}$-lifting $\widetilde{\mathcal{M}}$. Indeed, we have the following proposition:

Proposition 3.14. *The minimal $\mathcal{M}$-lifting is the lifting $\widetilde{\mathcal{M}}$ from example 3.4.*

In light of this proposition, $\widetilde{\mathcal{M}}$ really is the 'canonical' $\mathcal{M}$-lifting. The proof is not quite straightforward. We proceed in two steps: First, we show that $\widetilde{\mathcal{M}}R$ is minimal whenever R is a total and surjective relation. Then, we reduce the general case to this specific case.

Lemma 3.15. *Let $R : X \multimap Y$ be a total surjective relation. Then $\widetilde{\mathcal{M}}R \leq LR$ for all liftings L.*

In [11], a similar statement appears as lemma 4.7.

Proof. Consider the two projection morphisms $\pi_X : R \to X$ and $\pi_Y : R \to Y$. Since R is total and surjective, both these functions are surjective.

We claim that $\widetilde{\mathcal{M}}R = (\mathcal{M}\pi_X^{\mathrm{gr}})^\circ; \mathcal{M}\pi_Y^{\mathrm{gr}}$. The inequality $\geq$ follows from $R = (\pi_X^{\mathrm{gr}})^\circ; \pi_Y^{\mathrm{gr}}$.

For $\leq$, let $(U, V) \in \widetilde{\mathcal{M}}R$. Then we set

$$\begin{aligned} W_0 &:= \{\{(x, y) \in R \mid x \in u\} \mid u \in U\} \\ W_1 &:= \{\{(x, y) \in R \mid y \in V\} \mid v \in V\} \\ W &:= \langle W_0 \cup W_1 \rangle \end{aligned}$$

We claim that $\mathcal{M}\pi_X(W) = U$. For this, we need to show that (1) if $u \in U$, then $\pi_X^{-1}(u) \in W$, and (2) if $\pi_X^{-1}(u) \in W$, then $u \in U$.

(1) Clearly, if $u \in U$, then $\pi_X^{-1}(u) = \{(x, y) \in R \mid x \in u\} \in W$, so $\pi_X^{-1}(u) \in W$.

(2) Assume $\pi_X^{-1}(u) \in W$. There are two cases: (i) there is a $u' \in U$ with $\{(x, y) \in R \mid x \in u'\} \subseteq \pi_X^{-1}(u)$, or (ii) there is a $v \in V$ with $\{(x, y) \in R \mid y \in v'\} \subseteq \pi_X^{-1}(u)$.

   (i) In this case, we know that $\pi_X[\{(x, y) \in R \mid x \in u'\}] \subseteq \pi_X(\pi_X^{-1}(u))$. But since R was total, we know that $\pi_X[\{(x, y) \in R \mid x \in u'\}] = u'$ and $\pi_X[\pi^{-1}(u)] = u$. So $u' \subseteq u$, and hence $u \in U$.

(ii) Clearly, $\pi_X[\{(x,y) \in R \mid y \in v\}] = \{x \mid \exists y \in v : xRy\}$. Since $(U,V) \in \widetilde{\mathcal{M}}R$, there is a $u' \in U$ such that for all $x \in u'$, there is a $y \in v$ with xRy. But this just says that $u' \subseteq \pi_X[\{(x,y) \in R \mid y \in v\}]$. So we conclude that there is a $u' \in U$ with

$$u' \subseteq \pi_X[\{(x,y) \in R \mid y \in v\}] \subseteq \pi_X(\pi_X^{-1}(u)) = u$$

and hence $u \in U$.

So in both cases, we have $u \in U$, as desired.

The proof that $\mathcal{M}\pi_Y(W) = V$ is completely symmetrical; so, we can conclude that $(U,V) \in (\mathcal{M}\pi_X^{\mathrm{gr}})^\circ ; \mathcal{M}\pi_Y^{\mathrm{gr}}$.

Now, let L be any lifting. Then

$$LR = L((\pi_X^{\mathrm{gr}})^\circ ; \pi_Y^{\mathrm{gr}}) \geq L(\pi_X^{\mathrm{gr}})^\circ ; L\pi_Y^{\mathrm{gr}} \geq (\mathcal{M}\pi_X^{\mathrm{gr}})^\circ ; \mathcal{M}\pi_Y^{\mathrm{gr}} = \widetilde{\mathcal{M}}R$$

□

With this lemma, we can prove the proposition.

Proof. Let $R : X \multimap Y$ be any relation. Let X' be the domain of R and Y' the range of R. Then we define $X_* = X \cup \{*\}, Y_* = Y \cup \{*\}$ and

$$R_* = R \cup \{(x,*) \mid x \in X \setminus X'\} \cup \{(*,y) \mid y \in Y \setminus Y'\} \cup \{(*,*)\}$$

Then $R_* : X_* \multimap Y_*$ is total and surjective.

Let $\iota_X : X \to X_*$ and $\iota_Y : Y \to Y_*$ be the natural inclusion functions. First, we note that $R = \iota_X^{\mathrm{gr}} ; R_* ; (\iota_Y^{\mathrm{gr}})^\circ$ The inequality $\leq$ is clear, since $R \subseteq R_*$. For $\geq$, notice that $*$ is not in the range of either ι_X or ι_Y.

Now by lemma 3.10, we know that for any lifting L,

$$LR = (\mathcal{M}\iota_X^{\mathrm{gr}}) ; LR* ; (\mathcal{M}\iota_Y^{\mathrm{gr}})^\circ.$$

So we can calculate that

$$\begin{aligned} LR &= \mathcal{M}\iota_X^{\mathrm{gr}} ; LR_* ; (\mathcal{M}\iota_Y^{\mathrm{gr}})^\circ \\ &\geq \mathcal{M}\iota_X^{\mathrm{gr}} ; \widetilde{\mathcal{M}}R_* ; (\mathcal{M}\iota_Y^{\mathrm{gr}})^\circ && \text{by lemma 3.15} \\ &= \widetilde{\mathcal{M}}R \end{aligned}$$

We conclude that $\widetilde{\mathcal{M}}$ is minimal. □

3.5 Lifted elements

In many situations, we will be concerned with the relation $L\in: TX \multimap TPX$.

Definition 3.16. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let L be a T-lifting. Let $\Phi \in TPX$. We define the set $\lambda^L(\Phi)$ of *lifted elements of* Φ as

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

where $\in_X: X \multimap PX$ is the element relation.

Lemma 3.17. *Let* $T : \mathbf{Sets} \to \mathbf{Sets}$ *be a weak pullback-preserving functor. Let* X *be a set, take and* $\Phi \in T_\omega P_\omega X$. *Then for all* $\alpha \in \lambda^{\overline{T}}(\Phi)$, *we have that* $\alpha \in T_\omega X$, *and* $\mathrm{Base}(\alpha) \subseteq \bigcup \mathrm{Base}(\Phi)$.

Proof. Let $B = \bigcup \mathrm{Base}(\Phi)$, and let $f : B \to X$ be the inclusion. We prove both statements by showing that $\alpha \in Tf[B]$. Note that $\Phi \in TPf[PB]$ since $Base(\Phi) \subseteq PB$. So, there is a $\Phi' \in PB$ with $TPf(\Phi') = \Phi$. We see that

$$(\alpha, \Phi') \in \overline{T}\in; (TPf^{\mathrm{gr}})^\circ = \overline{T}(\in; (Pf^{\mathrm{gr}})^\circ).$$

But $x \in Pf(A)$ for some $A \in PB$ implies that $x \in f[B]$. Vice versa, if $x \in Tf[B]$, then $x \in Pf(\{x\})$. So, $\in; (Pf^{\mathrm{gr}})^\circ = (f^{\mathrm{gr}})^\circ; \in$. Hence,

$$(\alpha, \Phi') \in \overline{T}((f^{\mathrm{gr}})^\circ; \in) = (Tf^{\mathrm{gr}})^\circ; \overline{T}\in$$

and we see that $\alpha \in \mathrm{im}\, Tf$, which means that $\mathrm{Base}(\alpha) \subseteq B$ as desired. $\square$

3.6 Simulations and bisimulations

In this section we will examine notions of (bi)similarity based on relation liftings.

Definition 3.18. Let T be a functor, and let L be a T-lifting. Let $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S}' = (S', \sigma', m')$ be T-coalgebra models. A *L-simulation from* $\mathbb{S}$ *to* $\mathbb{S}'$ is a relation $R : S \multimap S'$ with the following properties:

atom: If sRs', then $m(s) = m(s')$;

unfolding: If sRs', then $(\sigma(s), \sigma'(s')) \in LR$.

If there is an L-simulation R from $\mathbb{S}$ to $\mathbb{S}'$ with sRs', then we write $\mathbb{S}, s \rightrightarrows^L \mathbb{S}', s'$.

A *L-bisimulation from* $\mathbb{S}$ *to* $\mathbb{S}'$ is a relation $R : S \multimap S'$ such that R is an L-simulation from $\mathbb{S}$ to $\mathbb{S}'$ and R° is an L-simulation from $\mathbb{S}'$ to $\mathbb{S}$. If there is an L-bisimulation R from $\mathbb{S}$ to $\mathbb{S}'$ with sRs', then we write $\mathbb{S}, s \leftrightarrows^L \mathbb{S}', s'$.

Remark 3.19. The usual definition of simulation requires the weaker property **weak-atom**, which states that if sRs', then $m(s) \subseteq m(s')$. The reason for our deviation is that we want to have the following equivalences:

(i) $R : \mathbb{S} \multimap \mathbb{S}'$ is an L-simulation if and only if R° is an $L^\sim$-simulation;

(ii) $R : \mathbb{S} \multimap \mathbb{S}'$ is an L-bisimulation if and only if it is both an L-simulation and an $L^{\smallsmile}$-simulation;

(iii) $R : \mathbb{S} \multimap \mathbb{S}'$ is an L-bisimulation if and only if it is an $(L \cap L^{\smallsmile})$-simulation.

Each of these equivalences would fail if the condition **atom** were replaced with **weak-atom**

The following proposition can be seen as justifying the defining properties of T-liftings.

Proposition 3.20. *Let T be a functor, let L be a T-lifting, and let $\mathbb{S}, \mathbb{S}', \mathbb{S}''$ be T-coalgebra models.*

(i) If $f : \mathbb{S} \to \mathbb{S}'$ is a coalgebra morphism, then f^{gr} is an L-bisimulation.

(ii) If $\{R_i \mid i \in I\}$ is a set of relations $R_i : S \multimap S'$, such that each $R - i$ is an L-simulation, then $\bigcup_{i \in I} R_i : S \multimap S'$ is an L-simulation.

(iii) If $R : \mathbb{S} \multimap S'$ and $R' : \mathbb{S}' \multimap \mathbb{S}''$ are L-simulations, then $R; R'$ is an L-simulation.

Proof. (i) Since f is a coalgebra morphism, the property **atom** is satisfied. For the property **unfolding**, note that $(s, s') \in f^{\mathrm{gr}}$ if and only if $s' = f(s)$. So, we see that if $(s, f(s)) \in f^{\mathrm{gr}}$, then

$$\begin{aligned}(\sigma(s), \sigma'(f(s))) &= (\sigma(s), Tf(\sigma(s))) \\ &\in Tf^{\mathrm{gr}} \\ &\subseteq L(f^{\mathrm{gr}})\end{aligned}$$

by assumption on L.

So f^{gr} is an L-simulation. Since L was arbitrary, we also have that f^{gr} is an $L^{\smallsmile}$-simulation; so by point (ii) in remark 3.19, we know that f^{gr} is an L-bisimulation.

(ii) Assume $(s, s') \in \bigcup_{i \in I} R_i$. Then there is an $i \in I$ with $(s, s') \in R_i$. **atom** is clearly satisfied. For **unfolding**, we see that

$$(\sigma(s), \sigma'(s')) \in LR_i \subseteq L(\bigcup_{i \in I} R_i)$$

since L preserves the ordering of relations.

(iii) Assume $(s, s'') \in R; R'$. Then there is s' such that $(s, s') \in R$ and $(s', s'') \in R'$. Then by assumption, we have

$$m(s) = m'(s') = m''(s'')$$

showing that **atom** is satisfied.

For **unfolding**, we know that $(\sigma(s), \sigma'(s')) \in LR$ and $(\sigma'(s'), \sigma''(s'')) \in LR'$; and therefore

$$(\sigma(s), \sigma''(s'')) \in LR; LR' \subseteq L(R; R')$$

$\square$

Proposition 3.21. *Let T be a functor, let L be a T-lifting, and let $\mathbb{S}, \mathbb{S}'$ be T-coalgebra models. Then $\sqsupseteq^L : \mathbb{S} \multimap \mathbb{S}'$ is an L-simulation; moreover, it is the greatest such.*

Proof. By definition, $\sqsupseteq^L$ is the union of all L-simulations from $\mathbb{S}$ to $\mathbb{S}'$. So certainly, if $R : \mathbb{S} \to \mathbb{S}'$ is an L-simulation, then $R \leq \sqsupseteq^L$.

By point (ii), we know that L-simulations are closed under union; hence $\sqsupseteq^L$ is an L-simulation, and moreover it is the greatest such. $\square$

3.7 Relation liftings and distributive laws

In this section, we give an alternative characterization of relation lifting in terms of *distributive laws*. Distributive laws at their most general are simply natural transformations $FG \Rightarrow GF$ for two functors F, G. They can play a role in defining interactions between algebraic and coalgebraic structure. In particular, they are used in the study of coinduction and corecursion [3].

Definition 3.22. Let $\mathbb{C}$ be a category, and let $T : \mathbb{C} \to \mathbb{C}$ be a monad, with unit $\eta : 1_{\mathbb{C}} \to T$ and counit $\mu : T^2 \to T$. Then define the *Kleisli category* $\mathbb{C}_T$ to have the same objects as $\mathbb{C}$; set a morphism $X \to_\mu Y$ in $\mathbb{C}_T$ to be a morphism $X \to TY$ in $\mathbb{C}$. Morphisms are composed via

$$g \circ_\mu f = \mu \circ Tg \circ f$$

We remark that **Rel** is equivalent to the Kleisli category $\mathbf{Sets}_P$ of the powerset monad on **Sets**. The unit for P is the singleton map $\eta : X \to PX$ given by $\eta(x) = \{x\}$. The counit is the union map $\mu : PPX \to PX$ given by $\mu(\mathcal{A}) = \bigcup_{A \in \mathcal{A}} A$.

We also note that, like **Rel**, the category $\mathbf{Sets}_P$ is enriched over **Poset** by setting for $f, g : X \to_\mu Y$ that $f \leq g$ iff for all $x \in X$, we have $f(x) \subseteq g(x)$. Moreover, $\mathbf{Sets}_P$ has a 'transposition' operation $(-)^\flat$: for $f : X \to PY$, we set $f^\flat : Y \to PX$ as

$$f^\flat(y) = \{x \in X \mid y \in f(x)\}.$$

We have already seen the operation $(-)^\flat$ in example 2.14.

It is generally known that if a functor T preserves weak pullbacks, then there is a distributive law $TP \Rightarrow PT$; namely, we can map $\Phi \in TP$ to its set of lifted elements $\lambda^{\bar{T}}(\Phi)$ (see definition 3.16). There is also a partial converse; J. Beck gave a general correspondence between distributive laws $TM \Rightarrow MT$, and functors $\mathbb{C}_M \to \mathbb{C}_M$ acting as T on objects [4]. Applied to $P : \mathbf{Sets} \to \mathbf{Sets}$, we get from distributive law $TP \Rightarrow PT$ a functor $\mathbf{Rel} \to \mathbf{Rel}$; however, these are not required to be monotone (see [16] for a discussion).

For our purposes, a natural transformation is not exactly the right notion; in particular, we are not interested in *strict* functoriality $L(R;S) = LR;LS$, but only in 'weak functoriality' $L(R;S) \supseteq LR;LS$. Hence, we introduce *weak distributive laws* $\lambda : TP \rightsquigarrow PT$:

Definition 3.23. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be any functor. A *weak distributive law for* T is a collection of maps $\lambda_- : TP(-) \to PT(-)$, satisfying:

(Monotonicity) For any two functions $f, g : X \to PY$, if $f \leq g$, then

$$\lambda_Y \circ Tf \leq \lambda_Y \circ Tg$$

(Weak naturality) For any function $f : X \to PY$, we have

$$PTf \circ \lambda_X \leq \lambda_{PY} \circ TPf$$

(Weak monadicity) For any Z, we have

$$\mu_{TZ} \circ P\lambda_Z \circ \lambda_{PZ} \leq \lambda_Z \circ T\mu_Z \text{ and } \lambda_Z \circ T\eta_Z \geq \eta_{TZ}$$

There are also the optional properties

(Weak extensionality) For any Z,

$$\lambda_Z \circ T\eta_Z \leq \eta_{TZ}$$

(Symmetry) For any map $f : X \to PY$,

$$(\lambda_Y \circ Tf)^\flat = \lambda_X \circ T(f^\flat)$$

As mentioned earlier, we get a weak distributive law λ^L for each lifting L, as defined in 3.16. We also get a lifting L from a weak distributive law.

Definition 3.24. Let $\lambda : TP \rightsquigarrow PT$ be a weak distributive law. For a given relation $R : X \multimap Y$, we define $L^\lambda R$ as

$$L^\lambda R := \{(\alpha, \beta) \mid \beta \in \lambda_Y \circ T\chi_R(\alpha)\}$$

where $\chi_R : X \to PY$ is the characteristic function of R given by $\chi_R(x) = \{y \mid xRy\}$.

It is not yet clear that these operations do result in a distributive law and a T-lifting respectively. This will be the main theorem of this section.

Theorem 3.25. *Let* $T : \mathbf{Sets} \to \mathbf{Sets}$ *be a functor.*

(i) If L is a T-lifting, then λ^L is a weak distributive law. Moreover, if L preserves diagonals then λ^L is weakly extensional, and if L is symmetric, then λ^L is symmetric.

(ii) If λ is a weak distributive law, then L^λ is a T-lifting. Moreover, if λ is weakly extensional, then L^λ preserves diagonals, and if λ is symmetric, then L^λ is symmetric.

(iii) The operations $L \mapsto \lambda^{L^\smallsmile}$ and $\lambda \mapsto L^\lambda$ are inverse to each other.

There is a slight disharmony in the definitions of L^λ and λ^L, seen in point (iii). It is interesting to note that we often consider the relation $L^\smallsmile\in$, while studying the behavior of L; we will see this again in the definitions 4.12 and 4.15.

Proof. (i) First, let L be any T-lifting. We check that λ^L satisfies all the conditions.

(Monotonicity) Let $f, g : X \to PY$ with $f \leq g$. Then

$$\begin{aligned} a \in \lambda^L_Y \circ Tf(\Phi) &\iff (a, Tf(\Phi)) \in L(\in_{PY}) \\ &\iff (a, \Phi) \in L(\in_{PY}); (Tf^{\mathrm{gr}})^\circ \\ &\iff (a, \Phi) \in L(\in_{PY}; (f^{\mathrm{gr}})^\circ) \\ &\implies (a, \Phi) \in L(\in_{PY}; (g^{\mathrm{gr}})^\circ) \\ &\iff \cdots \iff a \in \lambda^L_Y \circ Tg(\Phi) \end{aligned}$$

(Weak naturality) Let $f : X \to PY$ be a function. Note that

$$((f^{\mathrm{gr}})^\circ; \in_X) \subseteq (\in_{PY}; (Pf^{\mathrm{gr}})^\circ),$$

since if $(y, A) \in (f^{\mathrm{gr}})^\circ; \in_X$, then there is an $x \in A$ with $y = f(x)$, so $y \in Pf(A)$. We can now calculate that

$$\begin{aligned} (Tf^{\mathrm{gr}})^\circ; L\in_X &\subseteq L((f^{\mathrm{gr}})^\circ; \in_X) \\ &\subseteq L(\in_{PY}; (Pf^{\mathrm{gr}})^\circ) \\ &= L\in_{PY}; (TPf^{\mathrm{gr}})^\circ \end{aligned}$$

But we know that

$$\begin{aligned} a \in PTf \circ \lambda^L_X(\Phi) &\iff \exists a' : a = Tf(a') \text{ and } a' \in \lambda^L_X(\Phi) \\ &\iff \exists a' : a = Tf(a') \text{ and } (a', \Phi) \in L\in_X \\ &\iff (a, \Phi) \in (Tf^{\mathrm{gr}})^\circ; L\in_X \\ &\implies (a, \Phi) \in \in_{PY}; (TPf^{\mathrm{gr}})^\circ \\ &\iff a \in \lambda^L_{PY} \circ TPf(\Phi) \end{aligned}$$

So $PTf \circ \lambda^L_X \leq \lambda^L_{PY} \circ TPf$.

(Weak monadicity) Let $\Gamma \in TPPZ$. Then

$$\mu \circ P\lambda_Z \circ \lambda_{PZ}(\Gamma) = \{z \mid \exists \xi : (z, \xi) \in L\in \text{ and } (\xi, \Gamma) \in L\in\}$$

Now take $z \in \mu \circ P\lambda_Z \circ \lambda_{PZ}(\Gamma)$; then $(z, \Gamma) \in L\in; L\in \subseteq L(\in;\in)$. Next, note that

$$\in; (\mu^{\mathrm{gr}})^\circ = \in; \in,$$

since $(x, \mathcal{A}) \in \ (\in; (\mu^{\mathrm{gr}})^\circ)$ if and only if $x \in \bigcup_{A \in \mathcal{A}} A$ if and only if there is $A \in \mathcal{A}$ with $x \in A$ if and only if $(x, \mathcal{A}) \in \ (\in; \in)$. So we see that $(z, \Gamma) \in L(\in; \in) = L(\in; (\mu^{\mathrm{gr}})^\circ) = L\in; (T\mu^{\mathrm{gr}})^\circ$, which says that $z \in \lambda^L_Z \circ T\mu_Z(\Gamma)$, as desired.

For the second inequality, we need to show that if $\alpha \in TZ$, then $\alpha \in \lambda^L_Z(T\eta_Z(\alpha))$. This is true if and only if $(\alpha, T\eta_Z(\alpha)) \in L\in$. But this is clear: we have

$$(\alpha, T\eta_Z(\alpha)) \in T\eta^{\mathrm{gr}}_Z \subseteq L\eta^{\mathrm{gr}}_Z \subseteq L\in_Z,$$

since $\eta_Z \subseteq \in_Z$;

(Weak extensionality) Assume that L preserves diagonals. We simply note that $\in; (\eta^{\mathrm{gr}}_Z)^\circ = \Delta_Z$. So, we know that if $\beta \in \lambda^L_Z \circ T\eta_Z(\alpha)$, then

$$(\beta, \alpha) \in L\in; (T\eta^{\mathrm{gr}}_Z)^\circ = L(\in; (\eta^{\mathrm{gr}}_Z)^\circ) = L\Delta_Z \subseteq \Delta_{TZ}.$$

So $\beta = \alpha$, so we see that $\beta \in \{\alpha\} = \eta_{TZ}(\alpha)$.

(Symmetry) Assume that L is symmetric. Then for a function $f : X \to PY$, and $\alpha \in TY$, we have

$$\begin{aligned}
\beta \in (\lambda^L_Y \circ Tf)^\flat(\alpha) &\iff \alpha \in \lambda^L_Y \circ Tf(\beta) \\
&\iff (\alpha, Tf(\beta)) \in (L\in) \\
&\iff (\alpha, \beta) \in L\in; (Tf^{\mathrm{gr}})^\circ \\
&\iff (\alpha, \beta) \in L(\in; (f^{\mathrm{gr}})^\circ) && \text{by lemma 3.10.(ii)} \\
&\iff (\beta, \alpha) \in L(f^{\mathrm{gr}}; \ni) \\
&\iff (\beta, \alpha) \in L(\in; (f^\flat)^\circ) && \text{since } f(x) \ni y \text{ iff } x \in f^\flat(y) \\
&\iff (\beta, \alpha) \in L\in; (T(f^\flat)^{\mathrm{gr}})^\circ \\
&\iff (\beta, T(f^\flat)(\alpha)) \in L\in \\
&\iff \beta \in \lambda^L_X \circ T(f^\flat)(\alpha)
\end{aligned}$$

showing that

$$(\lambda_Y \circ Tf)^\flat = \lambda_X \circ T(f^\flat)$$

(ii) Let $\lambda : TP \rightsquigarrow PT$ be a weak distributive law.

We first need that if $S \subseteq R$, then $L^\lambda(S) \subseteq L^\lambda(R)$. But this follows directly from monotonicity of λ, and the fact that $S \subseteq R$ if and only if $\chi_S \leq \chi_R$.

Next, we show that $L^\lambda(f^{\mathrm{gr}}) \geq (Tf)^{\mathrm{gr}}$ for all functions $f : X \to Y$. Let f be any function. Then $\chi_{f^{\mathrm{gr}}} = \eta_Y \circ f$. Let $\alpha \in TX$. Then we have

$$\begin{aligned} Tf(\alpha) &\in \eta_{TY} \circ Tf(\alpha) \\ &\subseteq \lambda_Y \circ T\eta_Y \circ Tf(\alpha) \\ &= \lambda_Y \circ T\chi_{f^{\mathrm{gr}}}(\alpha), \end{aligned}$$

since $\chi_{f^{\mathrm{gr}}}(x) = \{y \mid (x, y) \in f^{\mathrm{gr}}\} = \{f(x)\} = \eta_Y \circ f$. We now have that $(\alpha, Tf(\alpha)) \in L^\lambda(f^{\mathrm{gr}})$; since α was arbitrary, we conclude that $(Tf)^{\mathrm{gr}} \leq L^\lambda(f^{\mathrm{gr}})$.

For the other inclusion, we investigate $T\chi_{(f^{\mathrm{gr}})^\circ}(Tf(\alpha))$. This is equal to $T(\chi_{(f^{\mathrm{gr}})^\circ} \circ f)(\alpha)$. But $\chi_{(f^{\mathrm{gr}})^\circ} \circ f(x) = \{x' \mid f(x') = f(x)\} \supseteq \{x\} = \eta_X(x)$. Hence,

$$\lambda_X \circ T\chi_{(f^{\mathrm{gr}})^\circ}(Tf(\alpha)) \supseteq \lambda_X \circ T(\eta_X)(\alpha) \supseteq \eta_{TX}(\alpha) \ni \alpha$$

and hence we have that $(Tf(\alpha), \alpha) \in L^\lambda((f^{\mathrm{gr}})^\circ)$ by definition of L^λ. Since α was arbitrary, we conclude that $(Tf^{\mathrm{gr}})^\circ \leq L^\lambda((f^{\mathrm{gr}})^\circ)$.

Now, we show that $L^\lambda(R;S) \geq L^\lambda(R); L^\lambda(S)$. To show this, it suffices to prove that

$$\lambda_Z \circ T\chi_{R;S} \geq \mu_{TZ} \circ P(\lambda_Z \circ T\chi_S) \circ \lambda_Y \circ T\chi_R,$$

this second expression being the characteristic function of $L^\lambda(R); L^\lambda(S)$. We note that $\chi_{R;S} = \mu_Z \circ P\chi_S \circ \chi_R$. And so

$$\begin{aligned} \mu_{TZ} \circ P(\lambda_Z \circ T\chi_S) \circ \lambda_Y \circ T\chi_R &= \mu_{TZ} \circ P\lambda_Z \circ PT\chi_S \circ \lambda_Y \circ T\chi_R \\ &\leq \mu_{TZ} \circ P\lambda_Z \circ \lambda_{PZ} \circ TP\chi_S \circ T\chi_R && \text{by weak naturality} \\ &\leq \lambda_Z \circ T\mu_Z \circ TP\chi_S \circ T\chi_R && \text{by weak monadicity} \\ &= \lambda_Z \circ T(\mu_Z \circ P\chi_S \circ \chi_R) \\ &= \lambda_Z \circ T(\chi_{R;S}) \end{aligned}$$

The above chain of (in)equalities can be represented by the diagram below:

$$\begin{tikzcd} TX \arrow[r, "T\chi_{R;S}"] \arrow[d, "T\chi_R"'] \arrow[dr, phantom, "="] & TPZ \arrow[r, "\lambda_Z"] \arrow[dr, phantom, "\geq"] & PTZ \\ TPY \arrow[r, "TP\chi_S"] \arrow[d, "\lambda_Y"'] \arrow[dr, phantom, "\leq"] & TPPZ \arrow[u, "T\mu_Z"'] \arrow[r, "P\lambda_Z \circ \lambda_{PZ}"] \arrow[d, "\lambda_{PZ}"] \arrow[dr, phantom, "="] & PPTZ \arrow[u, "\mu_{TZ}"'] \arrow[d, equal] \\ PTY \arrow[r, "PT\chi_S"] & PTPZ \arrow[r, "P\lambda_Z"] & PPTZ \end{tikzcd}$$

The top left square follows from the identity $\chi_{R;S} = \mu_Z \circ P\chi_S \circ \chi_R$. The top right square is the weak monadicity property. The bottom left square is the weak naturality property.

Next, we show that if λ is weakly extensional, then L^λ preserves diagonals. But this is clear; if λ is weakly extensional, then

$$\begin{aligned}(\alpha,\beta) \in L^\lambda(\Delta_X) &\implies \beta \in \lambda_X \circ T\chi_{\Delta_X}(\alpha)\\ &\iff \beta \in \lambda_X \circ T\eta_X(\alpha)\\ &\implies \beta \in \eta_{TZ}(\alpha)\\ &\iff \beta = \alpha\end{aligned}$$

Finally, we show that if λ is symmetric, then L^λ is symmetric. We calculate

$$\begin{aligned}(x,y) \in L^\lambda(R^\circ) &\iff y \in \lambda_Y \circ T\chi_{R^\circ}(x)\\ &\iff y \in \lambda_Y \circ T\chi_R^\flat(x)\\ &\iff y \in (\lambda_X \circ T\chi_R)^\flat(x)\\ &\iff x \in \lambda_X \circ T\chi_R(y)\\ &\iff (y,x) \in L^\lambda R\\ &\iff (x,y) \in (L^\lambda R)^\circ\end{aligned}$$

where we use that $\chi_{R^\circ} = \chi_R^\flat$.

(iii) This is an easy verification. First, let L be a T-lifting, and let $R : X \multimap Y$ be a relation. Then

$$\begin{aligned}(\alpha,\beta) \in L^{\lambda^{L^\sim}} R &\iff \beta \in \lambda^{L^\sim} \circ T\chi_R(\alpha)\\ &\iff (\beta, (T\chi_R)\alpha) \in (L^\sim \in)\\ &\iff (\beta,\alpha) \in (L^\sim \in); (T\chi_R^{\mathrm{gr}})^\circ\\ &\iff (\beta,\alpha) \in L^\sim(\in; (\chi_R^{\mathrm{gr}})^\circ) && \text{by lemma 3.10.(ii)}\\ &\iff (\alpha,\beta) \in L(\chi_R^{\mathrm{gr}}; \ni) && \text{by definition of } (-)^\sim\\ &\iff (\alpha,\beta) \in LR\end{aligned}$$

since $\chi_R^{\mathrm{gr}}; \ni = R$. Hence, $L^{\lambda^{L^\sim}} = L$.

Second, let λ be a weak distributive law, and let $\Phi \in TPX$. Then for all $\alpha \in TX$, we have

$$\begin{aligned}\alpha \in \lambda_X^{(L^\lambda)^\sim}(\Phi) &\iff (\alpha,\Phi) \in ((L^\lambda)^\sim \in)\\ &\iff (\Phi,\alpha) \in (L^\lambda \ni)\\ &\iff \alpha \in \lambda_X \circ T\chi_\ni(\Phi) && \text{by definition of } L^\lambda\\ &\iff \alpha \in \lambda_X \circ T\,\mathrm{id}_{PX}(\Phi)\\ &\iff \alpha \in \lambda_X(\Phi)\end{aligned}$$

since $\chi_\ni(U) = \{x \mid x \in U\} = U$ for all $U \in PX$, and hence $\chi_\ni = \mathrm{id}_{PX}$. $\square$

One consequence of this proposition is that T-liftings are uniquely defined by their action on the $\in$-relation. This can also be seen directly: we have that

$$LR = L(\in; (\chi^{\mathrm{gr}}_{R^\circ})^\circ) = L\in; (T\chi^{\mathrm{gr}}_{R^\circ})^\circ$$

for any relation R.

4 Coalgebraic logic

The purpose of this chapter is to develop a family of logics for coalgebras based on relation lifting. We will relate these logics to modal logic and monotone modal logic, and give a number of logical laws that govern the interaction between the various logical symbols.

Throughout this thesis, we take a fixed set Prop of proposition letters.

4.1 The ∇-modalities

Definition 4.1. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let Λ be a set of T-liftings. We define the set $\mathcal{L}_T(\Lambda)$ inductively as follows:

- If $p \in \mathsf{Prop}$, then $p \in \mathcal{L}_T(\Lambda)$;
- If $a \in \mathcal{L}_T(\Lambda)$, then $\neg a \in \mathcal{L}_T(\Lambda)$;
- If $A \in P_\omega \mathcal{L}_T(\Lambda)$, then $\bigvee A, \bigwedge A \in \mathcal{L}_T(\Lambda)$;
- If $\alpha \in T_\omega \mathcal{L}_T(\Lambda)$ and $L \in \Lambda$, then

$$\nabla_L \alpha, \Delta_L \alpha \in \mathcal{L}_T(\Lambda)$$

A symbol of the form ∇_L or Δ_L will be called a *modality*. We will use the symbol $\heartsuit$ to stand in for an arbitrary modality. If p is a proposition letter, a formula of the form p or $\neg p$ will be called a *literal*.

Remark 4.2. A cautious reader may have some worries about the above 'inductive' definition. After all, the inductive case for the modalities ranges over elements of $T_\omega \mathcal{L}_T(\Lambda)$. Do we not need to have access to the entire set $L_T(\Lambda)$ for this to be well-defined?

Indeed, take for instance the monotone neighborhood functor $\mathcal{M}$. Since any (non-empty) object $\alpha \in \mathcal{M}_\omega \mathcal{L}_\mathcal{M}(\Lambda)$ contains the entire set $\mathcal{L}_\mathcal{M}(\Lambda)$, we cannot give such an object before knowing what $\mathcal{L}_\mathcal{M}(\Lambda)$ is in its entirety; and hence the undertaking seems circular.

There are two ways to waylay this doubt. The first is to note that there is no problem if T *preserves inclusions*. There is no issue with the Booleans, since in order to know what a 'finite subset of $\mathcal{L}_T(\Lambda)$' is, we do not need to have access to the entire set. Similarly, if for a finite subset $B \subseteq \mathcal{L}_T(\Lambda)$, the set TB is a *subset* of $\mathcal{L}_T(\Lambda)$, an element of $T_\omega \mathcal{L}_T(\Lambda)$ just is an element of TB for a finite subset B of $\mathcal{L}_T(\Lambda)$, and so an inductive definition is appropriate.

It so happens that every finitary functor $\mathbf{Sets} \to \mathbf{Sets}$ is equivalent to one that preserves inclusions. And so, rather than work with T_ω itself, we could replace it with an equivalent functor that preserves inclusions.[2]

[2] This is essentially what we do for $\mathcal{M}$, by giving elements of $\mathcal{M}_\omega \mathcal{L}_\mathcal{M}(\Lambda)$ in tems of a finite set of generators.

An alternative route is explored in the appendix, where the functor is retained, but the inductive construction is replaced by a colimit construction.

In either case, we are justified in arguing *as though* the set $\mathcal{L}_T(\Lambda)$ was inductively defined. In particular, we can use induction on the complexity of formulas, both in definitions and in proofs.

Remark 4.3. We will often explicitly treat $\neg, \bigwedge, \bigvee$ as functions. In particular, we will often apply the functor T to them, yielding the expressions

$$(T\neg), \quad (T\bigwedge), \quad (T\bigvee).$$

Notation 4.4. In the case for the powerset functor, we will write

$$\heartsuit_{\overrightarrow{P}} =: \overrightarrow{\heartsuit}, \quad \heartsuit_{\overleftarrow{P}} =: \overleftarrow{\heartsuit};$$

similarly, for the monotone neighborhood functor, we will write

$$\heartsuit_{\overrightarrow{\overrightarrow{\mathcal{M}}}} =: \overrightarrow{\overrightarrow{\heartsuit}}, \quad \heartsuit_{\overleftarrow{\overleftarrow{\mathcal{M}}}} =: \overleftarrow{\overleftarrow{\heartsuit}}.$$

If $L = \overline{T}$ for a weak pullback-preserving functor T, we will omit the subscript, and simply write ∇, Δ for $\nabla_{\overline{T}}, \Delta_{\overline{T}}$.

Definition 4.5. For $a \in \mathcal{L}_T(\Lambda)$, we inductively define the *set of subformulas of* a as

$$\begin{aligned}
\mathrm{Sfor}(p) &:= \{p\} \\
\mathrm{Sfor}(\neg a) &:= \{\neg a\} \cup \mathrm{Sfor}(a) \\
\mathrm{Sfor}(\odot A) &:= \{\odot A\} \cup \bigcup_{a \in A} \mathrm{Sfor}(a) && \odot \in \{\bigvee, \bigwedge\} \\
\mathrm{Sfor}(\heartsuit_L \alpha) &:= \{\heartsuit_L \alpha\} \cup \bigcup_{a \in \mathrm{Base}(\alpha)} \mathrm{Sfor}(a) && \heartsuit \in \{\nabla, \Delta\}
\end{aligned}$$

We define the *set of variables occurring in* a as $\mathrm{Var}(a) := \mathrm{Sfor}(a) \cap \mathsf{Prop}$; that is, the set of proposition letters that are subformulas of a.

Definition 4.6. A formula $a \in \mathcal{L}_T(\Lambda)$ is interpreted on a coalgebra model $\mathbb{S} = (S, \sigma, m)$ as follows:

$$\begin{array}{lll}
\mathbb{S}, s \Vdash p & \text{iff} & p \in m(s) \\
\mathbb{S}, s \Vdash \neg a & \text{iff} & \mathbb{S}, s \nVdash a \\
\mathbb{S}, s \Vdash \bigwedge A & \text{iff} & \mathbb{S}, s \Vdash a \text{ for all } a \in A \\
\mathbb{S}, s \Vdash \bigvee A & \text{iff} & \mathbb{S}, s \Vdash a \text{ for some } a \in A \\
\mathbb{S}, s \Vdash \nabla_L \alpha & \text{iff} & (\sigma(s), \alpha) \in L \Vdash \\
\mathbb{S}, s \Vdash \Delta_L \alpha & \text{iff} & (\sigma(s), \alpha) \notin L \nVdash
\end{array}$$

This is a well-defined inductive definition, since in each clause the definition of $\mathbb{S}, s \Vdash a$ depends only on $\Vdash\restriction_{S \times (\mathrm{Sfor}(a) \setminus \{a\})}$.

This ∇-logic was first introduced by L. Moss for specifically the Barr lifting in [19]. In [1], A. Baltag generalized this method to general liftings. The modalities used there are denoted $\Box_\infty$ and $\Diamond_\infty$; we have chosen to use ∇ and Δ to emphasize the connection with the original Moss logic.

We finish this preliminary section with a definition of semantic consequence.

Definition 4.7. Let $a, a' \in \mathcal{L}_T(\Lambda)$ be formulas. We say that a' is a *consequence* of a if for all T-coalgebra models $\mathbb{S}$ and all $s \in \mathbb{S}$, we have

$$\mathbb{S}, s \Vdash a \Rightarrow \mathbb{S}, s \Vdash a'$$

If a' is a consequence of a, we write $a \vDash a'$.

If $a \vDash a'$ and $a' \vDash a$, then we call a and a' *equivalent*, and write $a \equiv a'$.

4.2 Modal logic

In this section, we compare the logic $\mathcal{L}_T(\Lambda)$ with some known modal logics, which are interpreted on certain coalgebra models.

Classical modal logic As noted earlier, a P-coalgebra model is essentially a Kripke model. There is a standard logical language for reasoning about Kripke models.

Definition 4.8. We define the set $\mathcal{L}_{\Box,\Diamond}$ of *modal formulas* as follows:

- If p is a proposition letter, then p is a modal formula.
- If φ is a modal formula, then $\neg\varphi$ is a modal formula.
- If φ and ψ are modal formulas, then $\varphi \wedge \psi$ and $\varphi \vee \psi$ are modal formulas.
- If φ is a modal formula, then $\Box\varphi$ and $\Diamond\varphi$ are modal formulas.

Modal formulas are interpreted on coalgebra models as follows:

Definition 4.9. Let $\mathbb{S} = (S, \sigma, m)$ be a P-coalgebra model, and $s \in S$. Then for a modal formula φ, we define by induction when $\mathbb{S}, s \Vdash_{\mathsf{K}}$.

$$\begin{array}{lll}
\mathbb{S}, s \Vdash_{\mathsf{K}} p & \text{iff} & p \in m(s) \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \neg\varphi & \text{iff} & \mathbb{S}, s \nVdash_{\mathsf{K}} \varphi \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \varphi \wedge \psi & \text{iff} & \mathbb{S}, s \Vdash_{\mathsf{K}} \varphi \text{ and } \mathbb{S}, s \Vdash_{\mathsf{K}} \psi \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \varphi \vee \psi & \text{iff} & \mathbb{S}, s \Vdash_{\mathsf{K}} \varphi \text{ or } \mathbb{S}, s \Vdash_{\mathsf{K}} \psi \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \Box\varphi & \text{iff} & \forall t \in \sigma(s), \mathbb{S}, t \Vdash_{\mathsf{K}} \varphi \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \Diamond\varphi & \text{iff} & \exists t \in \sigma(s), \mathbb{S}, t \Vdash_{\mathsf{K}} \varphi
\end{array}$$

The definition of the modalities can also be phrased in terms of relation lifting:

$$\begin{array}{lll}
\mathbb{S}, s \Vdash_{\mathsf{K}} \Box\varphi & \text{iff} & (\sigma(s), \{\varphi\}) \in (\overrightarrow{P} \Vdash_{\mathsf{K}}) \\
\mathbb{S}, s \Vdash_{\mathsf{K}} \Diamond\varphi & \text{iff} & (\sigma(s), \{\varphi\}) \in (\overleftarrow{P} \Vdash_{\mathsf{K}})
\end{array}$$

From this, it is easy to see that the modalities $\Box$ and $\Diamond$ are expressible in terms of $\overrightarrow{\nabla}$ and $\overleftarrow{\nabla}$ respectively.

Perhaps surprisingly, they can also be expressed in terms of $\overleftarrow{\Delta}$ and $\overrightarrow{\Delta}$ respectively. For, note that

$$\begin{aligned}(U, \{x\}) \in \overrightarrow{P}R &\text{ iff for all } u \in U, uRx \\ &\text{ iff there is no } u \in U \text{ with } uR^c x \\ &\text{ iff not for all } x \in \{x\} \text{ there is } u \in U \text{ with } uR^c x \\ &\text{ iff } (U, \{x\}) \notin \overleftarrow{P}(R^c)\end{aligned}$$

and similarly

$$(U, \{x\}) \in \overleftarrow{P}R \text{ iff } (U, \{x\}) \notin \overrightarrow{P}(R^c)$$

This tells us that for singletons $\{p\}$, we have the equivalences

$$\Box p \equiv \overrightarrow{\nabla}\{p\} \equiv \overleftarrow{\Delta}\{a\}, \qquad \Diamond p \equiv \overleftarrow{\nabla}\{a\} \equiv \overrightarrow{\Delta}\{a\}$$

We see that we can regard the classical modalities as highly restricted ∇-modalities. It would therefore seem that $\mathcal{L}_P(\overrightarrow{P}, \overleftarrow{P})$ is more expressive. This is not the case however. For, we have that

$$\begin{aligned}\mathbb{S}, s \Vdash \overrightarrow{\nabla}\{p_1, \ldots, p_n\} &\text{ iff } (\sigma(s), \{p_1, \ldots, p_n\}) \in (\overrightarrow{P} \Vdash) \\ &\text{ iff } \forall t \in \sigma(s). \exists p_i : \mathbb{S}, t \Vdash p_i \\ &\text{ iff } \forall t \in \sigma(s) : \mathbb{S}, t \Vdash \bigvee_{i=1}^{n} p_i \\ &\text{ iff } \mathbb{S}, s \Vdash_{\mathsf{K}} \Box(\bigvee_{i=1}^{n} p_i)\end{aligned}$$

and

$$\begin{aligned}\mathbb{S}, s \Vdash \overleftarrow{\nabla}\{p_1, \ldots, p_n\} &\text{ iff } (\sigma(s), \{p_1, \ldots, p_n\}) \in (\overleftarrow{P} \Vdash) \\ &\text{ iff } \forall p_i . \exists t \in \sigma(s) : \mathbb{S}, t \Vdash p_i \\ &\text{ iff } \forall p_i : \mathbb{S}, s \Vdash_{\mathsf{K}} \Diamond p_i \\ &\text{ iff } \mathbb{S}, s \Vdash_{\mathsf{K}} \bigwedge_{i=1}^{n} \Diamond p_i\end{aligned}$$

showing that $\mathcal{L}_T(\overrightarrow{P}, \overleftarrow{P})$ and $\mathcal{L}_{\Box,\Diamond}$ are equally expressive on P-coalgebras.

Definition 4.10. Let $\mathbb{S} = (S, \sigma, m)$ be a P-coalgebra model. We call a modal formula φ *valid* on $\mathbb{S}$ if for all $s \in \mathbb{S}$, we have $\mathbb{S}, s \Vdash_{\mathsf{K}} \varphi$.

We write K for the set of modal formulas valid on all P-coalgebra models.

Monotone modal logic The language $\mathcal{L}_{\Box,\Diamond}$ can also be interpreted on $\mathcal{M}$-coalgebras [11]; this leads to a logic called *monotone modal logic*.

Definition 4.11. Let $\mathbb{S} = (S, \sigma, m)$ be an $\mathcal{M}$-coalgebra model. Then for a modal formula φ, we define by induction when $\mathbb{S}, s \Vdash_{\mathsf{M}} \varphi$.

$$
\begin{array}{lll}
\mathbb{S}, s \Vdash_{\mathsf{M}} p & \text{iff} & p \in m(s) \\
\mathbb{S}, s \Vdash_{\mathsf{M}} \neg\varphi & \text{iff} & \mathbb{S}, s \nVdash_{\mathsf{M}} \varphi \\
\mathbb{S}, s \Vdash_{\mathsf{M}} \varphi \wedge \psi & \text{iff} & \mathbb{S}, s \Vdash_{\mathsf{M}} \varphi \text{ and } \mathbb{S}, s \Vdash_{\mathsf{M}} \psi \\
\mathbb{S}, s \Vdash_{\mathsf{M}} \varphi \vee \psi & \text{iff} & \mathbb{S}, s \Vdash_{\mathsf{M}} \varphi \text{ or } \mathbb{S}, s \Vdash_{\mathsf{M}} \psi \\
\mathbb{S}, s \Vdash_{\mathsf{M}} \Box\varphi & \text{iff} & \exists A \in \sigma(s) : \forall t \in A : \mathbb{S}, t \Vdash_{\mathsf{M}} \varphi \\
\mathbb{S}, s \Vdash_{\mathsf{M}} \Diamond\varphi & \text{iff} & \forall A \in \sigma(s) : \exists t \in \sigma(s), \mathbb{S}, t \Vdash_{\mathsf{M}} \varphi
\end{array}
$$

Here, too, the languages $\mathcal{L}_{\Box,\Diamond}$ and $\mathcal{L}_{\mathcal{M}}(\overset{\rightsquigarrow}{\mathcal{M}}, \overset{\leftsquigarrow}{\mathcal{M}})$ are interdefinable.

It is not hard to see that

$$\Diamond p \equiv \overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle, \quad \Box p \equiv \overset{\leftsquigarrow}{\nabla}\langle\{p\}\rangle$$

In the other direction, we see that for $\alpha \in \mathcal{M}_\omega\mathsf{Prop}$,

$$
\begin{aligned}
\mathbb{S}, s \Vdash_{\mathsf{M}} \overset{\leftsquigarrow}{\nabla}\alpha &\text{ iff } \forall A \in \alpha \exists U \in \sigma(s) : \forall t \in U \exists a \in A : \mathbb{S}, t \Vdash_{\mathsf{M}} a \\
&\text{ iff } \forall A \in \alpha \exists U \in \sigma(s) : \forall t \in U : \mathbb{S}, t \Vdash \bigvee A \\
&\text{ iff } \forall A \in \alpha : \mathbb{S}, s \Vdash \Box \bigvee A \\
&\text{ iff } \mathbb{S}, s \Vdash \bigwedge_{A \in \alpha} \Box \bigvee A
\end{aligned}
$$

showing that $\overset{\leftsquigarrow}{\nabla}$ is definable in terms of $\mathcal{L}_{\Box,\Diamond}$.

$\overset{\rightsquigarrow}{\nabla}$ is more difficult; in [21] an explicit translation is given. This uses some operations on $\mathcal{M}X$ that would take us too far to define here. Instead, we will see in section 4.3 that $\overset{\rightsquigarrow}{\nabla}$ can be defined in terms of $\overset{\rightsquigarrow}{\Delta}$, which in term can be defined in terms of $\overset{\leftsquigarrow}{\nabla}$.

4.3 Dualities and distributive laws

In classical modal logic, a number of equivalences between formulas are known. The most relevant to this thesis are

Duality: $\Box\varphi \equiv \neg\Diamond\neg\varphi$

Distribution: $\Box(\varphi \wedge \psi) \equiv \Box\varphi \wedge \Box\psi$

Analogues of these laws are known for the modality $\nabla_{\overline{T}}$ for a weak pull-back preserving functor (see e.g. [5]). In this section, we prove similar laws for the languages $\mathcal{L}_T(\Lambda)$.

Duality

Definition 4.12. For a given $a \in T_\omega \mathcal{L}_T(\Lambda)$ and lifting $L \in \Lambda$, define

$$\begin{aligned} \mathcal{D}_L\alpha &:= \{\Phi \in TP\,\mathrm{Base}(\alpha) \mid (\alpha, \Phi) \notin (L^{\smile} \notin)\} \\ S_L\alpha &:= \{(T \textstyle\bigwedge)\Phi \mid \Phi \in \mathcal{D}_L\alpha\} \\ D_L\alpha &:= \{(T \textstyle\bigvee)\Phi \mid \Phi \in \mathcal{D}_L\alpha\}. \end{aligned}$$

In [5], $S_L\alpha$ and $D_L\alpha$ are called $L(\alpha)$ and $R(\alpha)$ respectively, for 'left' and 'right'. Since in this thesis the letter L already plays many roles, we have opted for S and D instead, referring to 'sinister' and 'dexter' (which are latin for 'left' and 'right').

Example 4.13. We calculate $\mathcal{D}_L(\alpha), S_L(\alpha)$ and $D_L(\alpha)$ for some specific functors.

(i) Let $T = T_2$ be the ordered-pair functor $X \mapsto X^2$, and let $L = \overline{T}_2$ be the Barr lifting. Then let $\alpha = \langle a, b\rangle$ be any element of $T_2\mathcal{L}_{T_2}(\overline{T}_2)$. We see that $\mathrm{Base}(\alpha) = \{a, b\}$. Now for a $\Phi = \langle A, B\rangle \in T_2P\,\mathrm{Base}(\alpha)$, we have $\Phi \in \mathcal{D}_{\overline{T}_2}(\alpha)$ if $(\alpha, \Phi) \notin (\overline{T} \notin)$.

By definition, $(\langle a, b\rangle, \langle A, B\rangle) \in (\overline{T}_2 \notin)$ if and only if $a \notin A$ and $b \notin B$. So $\Phi \in \mathcal{D}_{\overline{T}_2}(\alpha)$ if $a \in A$ or $b \in B$.

From this, we see that an element of $S_L\alpha$ is of the form $\langle \bigwedge A, \bigwedge B\rangle$ with $a \in A$ or $b \in B$. An element of $D_L\alpha$ is of the form $\langle \bigvee A, \bigvee B\rangle$ with $a \in A$ or $b \in B$.

(ii) Let $T = P$ be the powerset functor, and let $L = \overrightarrow{P}$. Note that $L^{\smile} = \overleftarrow{P}$.

First we consider the case $\alpha = \varnothing$. Then $\mathrm{Base}(\alpha) = \varnothing$, and so $TP\,\mathrm{Base}(\alpha) = PP\varnothing = \{\varnothing, \{\varnothing\}\}$. For $\Phi \in TP\,\mathrm{Base}(\varnothing)$, we have $(\varnothing, \Phi) \in \overleftarrow{P}$ if for all $A \in \Phi$, there is an $a \in \varnothing$ with $a \notin A$. Clearly, this is only the case if $\Phi = \varnothing$. So, $\mathcal{D}_{\overrightarrow{P}}(\varnothing)$ has one element, $\{\varnothing\}$. Hence, $S_{\overrightarrow{P}}(\varnothing) = \{\{\bigwedge \varnothing\}\}$ and $D_{\overrightarrow{P}}(\varnothing) = \{\{\bigvee \varnothing\}\}$.

In general, for a given $\alpha \in P\mathcal{L}_P(\overrightarrow{P}, \overleftarrow{P})$, we have that $\mathrm{Base}(\alpha) = \alpha$, and so an element of $\mathcal{D}_{\overrightarrow{P}}\alpha$ is a set Φ of subsets of α, such that $(\alpha, \Phi) \notin (\overleftarrow{P} \notin)$. This happens precisely when $\alpha \in \Phi$. For, if $\alpha \in \Phi$, then there is no $a \in \alpha$ with $a \notin \alpha$, so it is not the case that for all $A \in \Phi$ there is $a \in \alpha$ with $a \notin A$. And if $\alpha \notin \Phi$, then every element of Φ is a *strict* subset of α, and hence for every $A \in \Phi$ there is $a \in \alpha$ with $a \notin A$, which means that $(\alpha, \Phi) \in (\overleftarrow{P} \notin)$.

From this, we see that an element of $S_{\overrightarrow{P}}(\alpha)$ is some set of conjunctions, at least containing $\bigwedge \alpha$; and an element of $D_{\overrightarrow{P}}(\alpha)$ is a set of disjunctions, at least containing $\bigvee \alpha$.

(iii) Again, let $T = P$ be the powerset functor, and let $L = \overleftarrow{P}$ (and hence $L^{\smile} = \overrightarrow{P}$). The case $\alpha = \varnothing$ is now trivial: for all elements $\Phi \in PP\varnothing$, we have $(\varnothing, \Phi) \in (\overrightarrow{P} \not\in)$, as its condition is "for all $a \in \varnothing, \ldots$". So $\mathcal{D}_{\overleftarrow{P}}(\varnothing), S_{\overleftarrow{P}}(\varnothing), D_{\overleftarrow{P}}(\varnothing)$ are all empty.

In general, for a given $\alpha \in P\mathcal{L}_P(\overrightarrow{P}, \overleftarrow{P})$, an element of $\mathcal{D}_{\overleftarrow{P}}\alpha$ is a set Φ of subsets of α, such that $(\alpha, \Phi) \notin (\overrightarrow{P} \not\in)$. For this to happen, it should not be the case that for all $a \in \alpha$, there is an $A \in \Phi$ with $a \notin A$. In other words, $(\alpha, \Phi) \notin (\overrightarrow{P} \not\in)$ if there is an $a \in \alpha$ such that for all $A \in \Phi$, $a \in A$. Or, put concisely: if Φ has non-empty intersection.

An element of $S_{\overleftarrow{P}}\alpha$ is a set of conjunctions, such that there is one element $a \in \alpha$ occurring in each conjunction. An element of $D_{\overleftarrow{P}}\alpha$ is the same, but with disjunctions instead of conjunctions.

Proposition 4.14 (Duality lemma)**.** *For a given pointed coalgebra model $\mathbb{S}, s$, the following equivalences hold:*

$$\mathbb{S}, s \Vdash \Delta_L\alpha \textit{ iff } \mathbb{S}, s \Vdash \neg\nabla_L(T\neg)\alpha \textit{ iff } \mathbb{S}, s \Vdash \nabla_{L^{\smile}}\beta \textit{ for some } \beta \in S_L\alpha \tag{1}$$

$$\mathbb{S}, s \Vdash \nabla_L\alpha \textit{ iff } \mathbb{S}, s \Vdash \neg\Delta_L(T\neg)\alpha \textit{ iff } \mathbb{S}, s \Vdash \Delta_{L^{\smile}}\beta \textit{ for all } \beta \in D_L\alpha \tag{2}$$

Proof. Let T be a functor, Λ a set of T-liftings, $\mathbb{S} = (S, \sigma, m)$ a T-coalgebra model, and let $s \in S$. Finally, let $\alpha \in T\mathcal{L}_T(\Lambda)$.

(1) First, assume that $\mathbb{S}, s \Vdash \Delta_L\alpha$. Assume towards a contradiction that $\mathbb{S}, s \Vdash \nabla_L(T\neg)\alpha$. Then

$$(\sigma(s), \alpha) \in L \Vdash ; (T\neg^{\mathrm{gr}})^{\circ} \subseteq L(\Vdash ; (\neg^{\mathrm{gr}})^{\circ}) = L(\nVdash)$$

since $\mathbb{S}, t \Vdash \neg a$ if and only if $\mathbb{S}, t \nVdash a$. But this is a contradiction with $\mathbb{S}, s \Vdash \Delta_L\alpha$. So, $\mathbb{S}, s \Vdash \neg\nabla_L(T\neg)\alpha$.

Next, assume that $\mathbb{S}, s \Vdash \neg\nabla_L(T\neg)\alpha$. Then define $\mathrm{Th}_\alpha : S \to \mathrm{Base}(\alpha)$ as

$$\mathrm{Th}_\alpha(t) := \{a \in \mathrm{Base}(\alpha) \mid \mathbb{S}, t \Vdash a\}$$

and set $\Phi := (T\,\mathrm{Th}_\alpha)\sigma(s)$. We claim that $\Phi \in \mathcal{D}_L\alpha$. For, if $(\alpha, \Phi) \in (L^{\smile} \not\in)$, then as illustrated in figure 2,

$$(\sigma(s), (T\neg)\alpha) \in (T\,\mathrm{Th}_\alpha)^{\mathrm{gr}} ; (L \not\ni) ; (T\neg)^{\mathrm{gr}} \subseteq L(\mathrm{Th}_\alpha ; \not\ni ; \neg^{\mathrm{gr}}) \subseteq L \Vdash$$

since if $a \notin \mathrm{Th}_\alpha(t)$, then $\mathbb{S}, s \Vdash \neg a$. This is a contradiction with $\mathbb{S}, s \Vdash \neg\nabla(T\neg)\alpha$, and hence $(\alpha, \Phi) \notin (L^{\smile} \not\in)$.

$$\begin{tikzcd}
\Phi \arrow[r, "L\not\ni"] & \alpha \arrow[d, "T\neg^{\mathrm{gr}}"] \arrow[dl, phantom, "\mid\wedge"] \\
\sigma(s) \arrow[u, "T\,\mathrm{Th}_\alpha^{\mathrm{gr}}"] \arrow[r, "L\Vdash"'] & \Phi
\end{tikzcd}$$

Figure 2

Now if we set $\beta := (T\bigwedge)\Phi$, then $\beta \in S_L\alpha$. We claim that $\mathbb{S}, s \Vdash \nabla_{L^\sim}\beta$. But this is straightforward: we see that

$$(\sigma(s), \beta) \in T\,\mathrm{Th}^{\mathrm{gr}}_{\alpha}; T\bigwedge{}^{\mathrm{gr}} \subseteq L(\mathrm{Th}^{\mathrm{gr}}_{\alpha}; \bigwedge{}^{\mathrm{gr}}) \subseteq L\Vdash$$

since for all $t \in \mathbb{S}$, we have that $\mathbb{S}, t \Vdash \bigwedge \mathrm{Th}_{\alpha}(t)$ by definition.

Finally, assume that $\mathbb{S}, s \Vdash \nabla_{L^\sim}\beta$ for some $\beta \in S_L\alpha$. Let $\Phi \in \mathcal{D}_L\alpha$ such that $\beta = (T\bigwedge)\Phi$. Then assume towards a contradiction that $(\sigma(s), \alpha) \in L\nVdash$. We calculate that, as illustrated in figure 3,

$$(\alpha, \Phi) \in (L^\sim \not⫣); (L^\sim \Vdash); (T\bigwedge{}^{\mathrm{gr}})^{\circ} \subseteq L^\sim(\not⫣; \Vdash; (\bigwedge{}^{\mathrm{gr}})^{\circ}) \subseteq L^\sim \notin$$

since if $\mathbb{S}, t \nVdash a$ and $\mathbb{S}, t \Vdash \bigwedge A$, then $a \notin A$. This is in contradiction with $(\alpha, \Phi) \notin L^\sim \notin$.

$$\begin{tikzcd}
\sigma(s) \arrow[r, "L^\sim\Vdash"] \arrow[dr, phantom, "\wedge|"] & \beta \\
\alpha \arrow[u, "L^\sim\not⫣"] \arrow[r, "L^\sim\notin"'] & \Phi \arrow[u, "\bigwedge^{\mathrm{gr}}"']
\end{tikzcd}$$

Figure 3

So, we conclude that $(\sigma(s), \alpha) \notin L\nVdash$, and hence $\mathbb{S}, s \Vdash \Delta_L\alpha$.

(2) First, assume that $\mathbb{S}, s \Vdash \nabla_L\alpha$. Then

$$(\sigma(s), T\neg\alpha) \in (L\Vdash); T\neg^{\mathrm{gr}} \subseteq L(\Vdash; \neg^{\mathrm{gr}}) = L\nVdash$$

and hence $\mathbb{S}, s \Vdash \neg\Delta(T\neg)\alpha$.

Next, assume that $\mathbb{S}, s \Vdash \neg\Delta_L(T\neg)\alpha$. Let $\beta \in D_L\alpha$. Then there is $\Phi \in \mathcal{D}_L\alpha$ with $\beta = (T\bigvee)\Phi$. Assume towards a contradiction that $(\sigma(s), \beta) \in L^\sim\nVdash$. Then as illustrated in figure 4, we have

$$(\alpha, \Phi) \in T\neg^{\mathrm{gr}}; (L^\sim \not⫣); (L^\sim \nVdash); (T\bigvee{}^{\mathrm{gr}})^{\circ} \subseteq L^\sim(\neg^{\mathrm{gr}}; \not⫣; \nVdash; (\bigvee{}^{\mathrm{gr}})^{\circ}) \subseteq L^\sim \notin$$

since if $\mathbb{S}, t \nVdash \neg a$, and $\mathbb{S}, t \nVdash \bigvee A$, then $\mathbb{S}, t \Vdash a$ and hence $a \notin A$. This is in contradiction with $\Phi \in \mathcal{D}_L\alpha$, and hence $\mathbb{S}, s \Vdash \Delta_{L^\sim}\beta$.

$$\begin{tikzcd}
(T\neg)\alpha \arrow[r, "L^\sim\not⫣"] & \sigma(s) \arrow[r, "L^\sim\nVdash"] \arrow[d, phantom, "\wedge|"] & \beta \\
\alpha \arrow[u, "T\neg"] \arrow[rr, "L^\sim\notin"'] & & \Phi \arrow[u, "\bigvee^{\mathrm{gr}}"']
\end{tikzcd}$$

Figure 4

Finally, assume that $\mathbb{S}, s \Vdash \Delta_{L^{\smallsmile}}\beta$ for all $\beta \in D_L\alpha$. Let $\mathrm{Th}^{-}_{\alpha} : S \to \mathrm{Base}(\alpha)$ be defined as

$$\mathrm{Th}^{-}_{\alpha}(t) = \{a \in \mathrm{Base}(\alpha) \mid \mathbb{S}, t \nVdash a\}$$

and set $\Phi = (T\,\mathrm{Th}^{-}_{\alpha})(\sigma(s))$. Then $(\alpha, \Phi) \in L^{\smallsmile} \notin$. For, assume towards a contradiction that $(\alpha, \Phi) \notin L^{\smallsmile} \notin$. Then $\Phi \in \mathcal{D}_L\alpha$; set $\beta = (T\bigvee)\Phi$. Since $\beta \in D_L\alpha$, we know that $\mathbb{S}, s \Vdash \Delta_{L^{\smallsmile}}\beta$ by assumption. But also,

$$(\sigma(s), \beta) \in (T\,\mathrm{Th}^{-}_{\alpha})^{\mathrm{gr}}; (T\bigvee{}^{\mathrm{gr}}) \subseteq L^{\smallsmile}((\mathrm{Th}^{-}_{\alpha})^{\mathrm{gr}}; \bigvee{}^{\mathrm{gr}}) \subseteq L \nVdash$$

since certainly $\mathbb{S}, t \nVdash \bigvee \mathrm{Th}^{-}_{\alpha}(t)$. This contradicts $\mathbb{S}, s \Vdash \Delta_{L^{\smallsmile}}\beta$, and hence $(\alpha, \Phi) \in L^{\smallsmile} \notin$.

But now

$$(\sigma(s), \alpha) \in (T\,\mathrm{Th}^{-}_{\alpha})^{\mathrm{gr}}; (L \not\ni) \subseteq L(\mathrm{Th}^{-}_{\alpha}; \not\ni) \subseteq L \Vdash$$

since if $a \notin \mathrm{Th}^{-}_{\alpha}(t)$, then $\mathbb{S}, t \Vdash a$ by definition.

$\square$

Distribution

Definition 4.15. Let Λ be a set of T-liftings, and let $\Gamma \in P_\omega T_\omega \mathcal{L}_T(\Lambda)$. For every $\gamma \in \Gamma$, let $L_\gamma \in \Lambda$ be an associated lifting, and let

$$\nabla\Gamma := \{\nabla_{L_\gamma}\gamma \mid \gamma \in \Gamma\}.$$

Recall from definition 2.8 the notation $\mathcal{B}(\Gamma) = \bigcup_{\gamma\in\Gamma} \mathrm{Base}(\gamma)$. We call an object $\Phi \in TP\mathcal{B}(\Gamma)$ a *redistribution* of $\nabla\Gamma$ if for all $\gamma \in \Gamma$, we have $(\gamma, \Phi) \in L^{\smallsmile}_{\gamma}\in$. We will write

$$\mathcal{R}(\nabla\Gamma) := \{\Phi \in TP\mathcal{B}(\Gamma) \mid \Phi \text{ is a redistribution of } \nabla\Gamma\}.$$

Proposition 4.16 (Distributive law)**.** *Let Λ be a set of T-liftings, and let $\Gamma \in P_\omega T_\omega \mathcal{L}_T(\Lambda)$. Then write $\nabla\Gamma = \{\nabla_{L_\gamma}\gamma \mid \gamma \in \Gamma\}$ and assume that $L_0 \in \Lambda$ is such that $L_0 \leq L_\gamma$ for all $\gamma \in \Gamma$. Then for a pointed T-coalgebra model $\mathbb{S}, s$, the following equivalence holds:*

$$\mathbb{S}, s \Vdash \bigwedge \nabla\Gamma \text{ iff } \mathbb{S}, s \Vdash \nabla_{L_0}(T\bigwedge)\Phi \text{ for some } \Phi \in \mathcal{R}(\nabla\Gamma).$$

Proof. First, assume that $\mathbb{S}, s \Vdash \bigwedge \nabla\Gamma$. Then let $\mathrm{Th}_\Gamma : S \to \mathcal{B}(\Gamma)$ be defined as

$$\mathrm{Th}_\Gamma : t \mapsto \{a \in \mathcal{B}(\Gamma) \mid \mathbb{S}, t \Vdash a\}$$

and set $\Phi = (T\,\mathrm{Th}_\Gamma)\sigma(s)$.

We claim that $\Phi \in \mathcal{R}(\nabla\Gamma)$. For, take $\gamma \in \Gamma$. Then $\mathbb{S}, s \Vdash \nabla_{L_\gamma}\gamma$, so $(\sigma(s), \gamma) \in L_\gamma \Vdash$. Hence,

$$(\gamma, \Phi) \in L^{\smallsmile}_{\gamma}\dashv\!\!|; (T\,\mathrm{Th}_\Gamma)^{\mathrm{gr}} \subseteq L^{\smallsmile}_{\gamma}(\dashv\!\!|; \mathrm{Th}^{\mathrm{gr}}_{\Gamma}) \subseteq (L^{\smallsmile}_{\gamma}\in)$$

as desired.

Moreover, we clearly have

$$(\sigma(s), (T \bigwedge)\Phi) \in T\,\mathrm{Th}^{\mathrm{gr}}_{\Gamma}; T \bigwedge{}^{\mathrm{gr}} \subseteq L_0(\mathrm{Th}^{\mathrm{gr}}_{\Gamma}; \bigwedge{}^{\mathrm{gr}}) \subseteq (L_0 \Vdash)$$

showing that $\mathbb{S}, s \Vdash \nabla_{L_0}(T \bigwedge)\Phi$.

Now assume that there is a $\Phi \in \mathcal{R}(\nabla\Gamma)$ with $\mathbb{S}, s \Vdash \nabla_{L_0}(T \bigwedge)\Phi$. Then let $\gamma \in \Gamma$. We see that

$$\begin{aligned}(\sigma(s), \gamma) &\in (L_0 \Vdash); (T \bigwedge{}^{\mathrm{gr}})^{\circ}; (L_\gamma \ni)\\ &\subseteq (L_\gamma \Vdash); L_\gamma((\bigwedge{}^{\mathrm{gr}})^{\circ}); (L_\gamma \ni)\\ &\subseteq L_\gamma(\Vdash; (\bigwedge{}^{\mathrm{gr}})^{\circ}; \ni)\\ &\subseteq (L_\gamma \Vdash)\end{aligned}$$

since if $\mathbb{S}, t \Vdash \bigwedge A$ and $A \ni a$, then $\mathbb{S}, t \Vdash a$.

Hence, for all $\gamma \in \Gamma$, we have $\mathbb{S}, s \Vdash \nabla_{L_\gamma}\gamma$, and so $\mathbb{S}, s \Vdash \bigwedge \nabla\Gamma$. □

Remark 4.17. In order to make use of the distributive law in proposition 4.16 on some conjunction $\bigwedge \nabla\Gamma$, there must be a $L_0 \in \Lambda$ with $L_0 \leq L_\gamma$ for all γ. For this to be the case, we see that Λ needs to be *filtered*.

Definition 4.18. Let T be a functor, and let Λ be a set of T-liftings. We call Λ *filtered* if for all $L_1, \ldots, L_k \in \Lambda$, there is an $L \in \Lambda$ with $L \leq L_i$ for $i = 1, \ldots, k$.

This is not such a strong requirement; we know that T admits a minimal lifting L_0, and hence $\Lambda' = \Lambda \cup L_0$ is always filtered.

Strong distributive law One can wonder why we have chosen the distributive law for $\Box$ as our template for the distributive law for ∇. After all, the ∇-modality can be seen as a generalization of $\Diamond$ just as well as a generalization of $\Box$. And for $\Diamond$, the distributive law reads

$$\Diamond(\varphi \vee \psi) \equiv \Diamond\varphi \vee \Diamond\psi$$

It seems that the ∇-modality is closer to $\Box$ than to $\Diamond$, since both the ∇- and $\Box$-modality have a 'universal' flavor, in contrast to the 'existential' flavor of the $\Diamond$.

Still, in the specific case of $\nabla_{\overline{T}}$ for a weak pullback-preserving functor, we do obtain a distributive law for the disjunction.

Proposition 4.19 (Strong distributive law)**.** *Let T be a weak pullback-preserving functor, and let Λ be a set of T-liftings containing $\overline{T}$. Let $\Phi \in T_\omega P_\omega \mathcal{L}_T(\Lambda)$, and let $\mathbb{S}, s$ be a pointed T-coalgebra. Then we have the following equivalence:*

$$\mathbb{S}, s \Vdash \nabla(T \bigvee)\Phi \text{ iff there is } \alpha \in T_\omega \mathcal{L}_T(\Lambda) \text{ with } (\alpha, \Phi) \in (\overline{T}{\in}) \text{ and } \mathbb{S}, s \Vdash \nabla\alpha$$

Proof. First, assume that $\mathbb{S}, s \Vdash \nabla(T \bigvee)\Phi$. Then

$$\begin{aligned}(\sigma(s), \Phi) &\in (\overline{T} \Vdash); (T \bigvee{}^{\mathrm{gr}})^{\circ} \\ &= \overline{T}(\Vdash; (\bigvee{}^{\mathrm{gr}})^{\circ}) \\ &= \overline{T}(\Vdash; \in) \\ &= (\overline{T} \Vdash); (\overline{T}\in)\end{aligned}$$

and hence there is α with $(\sigma(s), \alpha) \in (\overline{T} \Vdash)$, and $(\alpha, \Phi) \in (\overline{T}\in)$, as desired.

Next, assume that $\mathbb{S}, s \Vdash \nabla\alpha$ for some α with $(\alpha, \Phi) \in (\overline{T}\in)$. Then

$$\begin{aligned}(\sigma(s), (T \bigvee \Phi)) &\in (\overline{T} \Vdash); (\overline{T}\in); (T \bigvee)^{\mathrm{gr}} \\ &= \overline{T}(\Vdash; \in; \bigvee{}^{\mathrm{gr}}) \\ &\subseteq \overline{T}(\Vdash)\end{aligned}$$

showing that $\mathbb{S}, s \Vdash \nabla(T \bigvee)\Phi$.

□

It is instructive to note that the right-to-left direction holds not just for the Barr lifting $\overline{T}$, but for general liftings L. It is only in the left-to-right direction, where we need to find the 'intermediate' point α in $\overline{T}(\Vdash; \in)$, that the strict distributivity of the Barr lifting is needed.

4.4 Normal Forms

In this section, we will discuss a number of normal forms for $\mathcal{L}_T(\Lambda)$-formulas.

Definition 4.20. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let Λ be a set of T-liftings. A formula $a \in \mathcal{L}_T(\Lambda)$ is called *clean* if all negations in a occur directly before a proposition letter.

Inductively, we define the set $\mathcal{L}^c_T(\Lambda)$ of clean (T, Λ)-formulas as follows:

- If p is a proposition letter, then p and $\neg p$ are clean formulas;
- If A is a set of clean formulas, then $\bigwedge A$ and $\bigvee A$ are clean formulas;
- If $L \in \Lambda$ and $\alpha \in T_\omega \mathcal{L}^c_T(\Lambda)$, then $\nabla_L \alpha$ and $\Delta_L \alpha$ are clean formulas.

Proposition 4.21. *Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let Λ be a set of T-liftings. Then any formula $a \in L_T(\Lambda)$ is equivalent to a formula $c(a) \in \mathcal{L}^c_T(\Lambda)$.*

Proof. We prove it by induction on the complexity of formulas. If a is not of the form $\neg a'$, then the induction is straightforward. So, assume that $a = \neg a'$ for some $a' \in \mathcal{L}_T(\Lambda)$. There are several cases to consider.

- If $a' = p$ is a proposition letter, then a is already clean.

- If $a' = \neg a''$, then we see that a is equivalent to a'', which by induction is equivalent to a clean formula $c(a'')$.
- If $a' = \bigwedge A$, then we see that a is equivalent to

$$\bigvee\{\neg a'' \mid a'' \in A\}$$

and by induction, for every a'' the formula $\neg a''$ is equivalent to a clean formula $c(\neg a'')$. So, we see

$$a \equiv \bigvee\{c(\neg a'') \mid a'' \in A\} =: c(a)$$

which shows that a is equivalent to a clean formula.

- The case $a' = \bigvee A$ is similar to the case $a' = \bigwedge A$.
- If $a' = \nabla_L \alpha$, then by proposition 4.14, we know that

$$a = \neg\nabla_L\alpha \equiv \Delta_L(T\neg)\alpha$$

and by induction, we have a function $c : \mathrm{Base}((T\neg)\alpha) \to \mathcal{L}_T^c(\Lambda)$ such that a'' is equivalent to $c(a'')$ for all $a'' \in \mathrm{Base}((T\neg)\alpha)$. Hence, if we set

$$c(a) = \Delta_L(Tc \circ T\neg)\alpha$$

we see that $c(a)$ is a clean formula, equivalent to a.

- The case $a' = \Delta_L\alpha$ is similar to the case $a' = \nabla_L\alpha$.

So now by induction, we have shown that every formula in $\mathcal{L}_T(\Lambda)$ is equivalent to a clean formula. □

If T preserves finite sets, then we can use the duality theorem to completely eliminate the Δ-modalities.

Definition 4.22. Let T : **Sets** $\to$ **Sets** be a functor, and let Λ be a set of T-liftings. We define the set $\mathcal{L}_T^\nabla(\Lambda)$ of *Δ-free (T,Λ)-formulas* as those clean formulas containing only ∇-modalities. Inductively, they are defined as follows:

- If p is a proposition letter, then p and $\neg p$ are Δ-free formulas;
- If A is a set of Δ-free formulas, then $\bigwedge A$ and $\bigvee A$ are Δ-free formulas;
- If $L \in \Lambda$ and $\alpha \in T_\omega\mathcal{L}_T^c(\Lambda)$, then $\nabla_L\alpha$ is a Δ-free formula.

Proposition 4.23. *Let T :* **Sets** $\to$ **Sets** *preserve finite sets, and let Λ be a set of T-liftings, closed under $(-)^\sim$. Then any formula $a \in \mathcal{L}_T(\Lambda)$ is equivalent to a Δ-free formula $\pi(a) \in \mathcal{L}_T^\nabla(\Lambda)$.*

Proof. First, by proposition 4.21, we may assume that a is a clean formula. The only interesting case in the induction is the case $a = \Delta_L \alpha$ for some $L \in \Lambda, \alpha \in \mathcal{L}^c_T(\Lambda)$. Then since T preserves finite sets, we know that $\mathcal{D}_L \alpha \subseteq TP\,\mathrm{Base}(\alpha)$ is finite. So, $S_L\alpha$ is finite, and hence we see that by proposition 4.14,

$$a \equiv \bigvee \{\nabla_{L^{\sim}} \beta \mid \beta \in S_L \alpha\}.$$

By induction, for every formula $b \in \mathcal{B}(S_L\alpha)$, there is an equivalent Δ-free formula $\pi(b)$. So now,

$$\bigvee \{\nabla_{L^{\sim}} (T\pi)\beta \mid \beta \in S_L \alpha\}$$

is a Δ-free formula equivalent to a. □

If Λ is filtered (see definition 4.18), we can obtain an even stronger normal form, where negations occur only before proposition letters, only ∇-modalities appear, and conjunctions only apply to sets of literals and at most one modality.

Definition 4.24. Let T : **Sets** $\to$ **Sets** be a functor, and let Λ be a set of T-liftings. We define the set $\mathcal{L}^{\mathrm{NF}}_T(\Lambda)$ of *Normal Forms* as follows:

- If P is a set of literals, then $\bigwedge P$ is a normal form.
- If P is a set of literals, $L \in \Lambda$ and $\alpha \in T_\omega \mathcal{L}^{\mathrm{NF}}_T(\Lambda)$, then $\bigwedge P \wedge \nabla_L \alpha$ is a normal form.
- If A is a finite set of normal forms, then $\bigvee A$ is a normal form.

Proposition 4.25. *Let T : **Sets** $\to$ **Sets** be a functor, and let Λ be a set of T-liftings, closed under $(-)^{\sim}$. If T preserves finite sets, and Λ is filtered, then any $\mathcal{L}_T(\Lambda)$-formula a is equivalent to a formula $\mathrm{NF}(a) \in \mathcal{L}^{\mathrm{NF}}_T(\Lambda)$.*

The idea is the following: first, we push negations down as in proposition 4.21. Then we flip each Δ as in proposition 4.23. Finally, we push conjunctions down by using the De Morgan law

$$a \wedge (b \vee c) \equiv (a \wedge b) \vee (a \wedge c)$$

and the distributive law from 4.16.

Proof. By proposition 4.23, we may assume that a is a Δ-free formula. The cases $a = \bigvee A$ and $a = \nabla_L \alpha$ are trivial. The only non-trivial case is that of $a = \bigwedge A$.

By induction, we may assume that all elements of A are normal forms. If there is any $a' \in A$ with $a' = \bigvee A'$, then

$$a \equiv \bigvee \{\bigwedge (A \setminus \{a'\}) \cup a'' \mid a'' \in A'\}$$

and every formula $\bigwedge (A \setminus \{a'\}) \cup a''$ has lower complexity than a, so by induction is equivalent to a normal form. Now a is equivalent to a disjunction of normal forms, and hence a normal form.

So, we may assume that every formula in A is either a conjunction of literals, or a conjunction of literals and one nabla formula. That is, we have $A = A_0 \cup A_1$, with

$$A_0 = \{\bigwedge P_1, \ldots, \bigwedge P_k\}, \quad A_1 = \{\bigwedge Q_1 \wedge \nabla_{L_1}\alpha_1, \ldots, \bigwedge Q_m \wedge \nabla_{L_m}\alpha_m\}$$

Now let $P = \bigcup_{i=1}^k P_i \cup \bigcup_{j=1}^m Q_j$. Clearly, we now have that

$$a \equiv \bigwedge P \wedge \bigwedge_{j=1}^{m} \nabla_{L_j}\alpha_j$$

Let $L_0 \leq L_j$ for $j = 1, \ldots, m$. Then by proposition 4.16, we know that

$$\bigwedge_{j=1}^{m} \nabla_{L_j}\alpha \equiv \bigvee\{\nabla_{L_0}(T\bigwedge)\Phi) \mid \Phi \in \mathcal{R}(\{\nabla_{L_1}\alpha_1, \ldots, \nabla_{L_m}\alpha_m\})\}$$

Hence, we can write

$$a \equiv \bigvee_{\Phi} \bigwedge P \wedge \nabla_{L_0}(T\bigwedge)\Phi.$$

Now by induction, for any $\Phi \in \mathcal{R}(\{\nabla_{L_1}\alpha_1, \ldots, \nabla_{L_m}\alpha_m\})$, every formula b in $\{T\bigwedge B \mid B \in \mathrm{Base}(\Phi)\}$ is equivalent to a normal form $\mathrm{NF}(b)$, and hence

$$a \equiv \bigvee_{\Phi} \left(\bigwedge P \wedge \nabla_{L_0}(T(\mathrm{NF} \circ \bigwedge))\Phi\right)$$

is a normal form for a.

□

We finish with a strong normal form theorem for $\mathcal{L}_T(\overline{T})$.

Definition 4.26. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and let Λ be a set of T-liftings. We define the set of *pure* (T, Λ)*-formulas* $\mathcal{L}_T^p(\Lambda)$ as follows:

- If P is a consistent set of literals, then $\bigwedge P$ is a pure formula.
- If P is a consistent set of literals, $L \in \Lambda$ a lifting, and $\alpha \in T_\omega \mathcal{L}_T^p(\Lambda)$, then $\bigwedge P \wedge \nabla_L \alpha$ is a pure formula.

We call a formula $a \in \mathcal{L}_T(\Lambda)$ a *pure normal form* if it is a disjunction of pure formulas.

Proposition 4.27. *Let* $T : \mathbf{Sets} \to \mathbf{Sets}$ *be a weak pullback-preserving functor, that additionally preserves finite sets. Then any formula* $a \in \mathcal{L}_T(\overline{T})$ *is equivalent to a pure normal form.*

Proof. By proposition 4.25, we may assume that a is a normal form. First, we consider the case $a = \bigwedge P$ for P a set of literals. If P is consistent, then

$a \equiv \bigvee\{\bigwedge P\}$ is a pure normal form. If P is inconsistent, then $a \equiv \bot \equiv \bigvee \varnothing$. So in this case, a is equivalent to a pure normal form.

Next, assume that $a = \bigvee A$ is a disjunction of normal forms. Then by induction, we can write every $a_i \in A$ as $\bigvee A_i$, with A_i a set of pure formulas. Now $a \equiv \bigvee \bigcup_i A_i$, showing that a is equivalent to a pure normal form.

Finally, assume that a is of the form $\bigwedge P \wedge \nabla\alpha$ for P a set of literals, and $\alpha \in T_\omega \mathcal{L}_T^{\mathrm{NF}}(\overline{T})$. Then again, if P is inconsistent, $a \equiv \bot$, and hence a is equivalent to a pure normal form. If P is consistent, then we continue. By induction, for every $b \in \mathrm{Base}(\alpha)$, there is a set $A(b) \subseteq \mathcal{L}_T^p(\Lambda)$ such that $b \equiv \bigvee A(b)$. Hence, we see that

$$a = \bigwedge P \wedge \nabla\alpha \equiv \bigwedge P \wedge \nabla(T \bigvee \circ TA)\alpha$$

Note that $(TA)\alpha \in T_\omega P_\omega \mathcal{L}_T^p(\Lambda)$. Let $\mathcal{A} := \{\beta \in T_\omega \mathcal{L}_T^p(\Lambda) \mid (\beta, (TA)\alpha) \in (\overline{T}\in)\}$. By lemma 3.17, we know that if $\beta \in \mathcal{A}$, then $\mathrm{Base}(\beta) \subseteq \bigcup \mathrm{Base}(TA)\alpha$. Since T preserves finite sets, we know that $T(\bigcup \mathrm{Base}(TA)\alpha)$ is finite, and hence $\mathcal{A}$ is finite as well. By proposition 4.19, we now know that

$$\nabla(T \bigvee \circ TA)\alpha \equiv \bigvee_{\beta \in \mathcal{A}} \nabla\beta.$$

Using this, we see

$$a \equiv \bigvee_{\beta \in \mathcal{A}} \bigwedge P \wedge \nabla\beta$$

showing that a is equivalent to a pure normal form. $\square$

5 (Relative) expressivity of the ∇-modalities

This chapter is devoted to investigating which coalgebraic phenomena are captured by the languages $\mathcal{L}_T(\Lambda)$ for variable Λ. We start by proving an analogue of the Hennessy-Milner theorem from modal logic[6].

The rest of this chapter explores the relative expressivity of the modalities. There are two central questions we will be concerned with:

1. If a formula a is defined from a lifting L, and $L' \leq L$, can a also be defined from L'?
2. If a formula a is preserved under L-simulations, can it be defined from L?

5.1 Adequacy

In this section, we show that two T-coalgebra models are L-bisimilar if and only if they are indistinguishable to $\mathcal{L}_T(L)$.

Definition 5.1. Let T : **Sets** $\to$ **Sets** be a functor, and let Λ be a set of T-liftings. Let $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S}' = (S', \sigma', m')$ be T-coalgebra models, and let $s \in S, s' \in S'$. We call $\mathbb{S}, s$ and $\mathbb{S}', s'$ $\mathcal{L}_T(\Lambda)$-*equivalent* if for all $a \in \mathcal{L}_T(\Lambda)$, we have

$$\mathbb{S}, s \Vdash a \text{ if and only if } \mathbb{S}', s' \Vdash a.$$

If $\mathbb{S}, s$ and $\mathbb{S}', s'$ are $\mathcal{L}_T(\Lambda)$-equivalent, we write $\mathbb{S}, s \equiv_\Lambda \mathbb{S}', s'$.

Theorem 5.2. *Let T : **Sets** $\to$ **Sets** be a functor, let Λ be a set of T-liftings, and set $L_0 := \bigwedge \Lambda$. Then for finite branching T-coalgebra models $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S}' = (S', \sigma', m')$, with $s \in S$ and $s' \in S'$, the following are equivalent:*

(i) $\mathbb{S}, s \underline{\leftrightarrow}^{L_0} \mathbb{S}', s'$;

(ii) $\mathbb{S}, s \equiv_\Lambda \mathbb{S}', s'$.

Proof. (i)$\Rightarrow$(ii): Assume that $\mathbb{S}, s \underline{\leftrightarrow}^{L_0} \mathbb{S}, s'$. Then let $R : \mathbb{S} \multimap \mathbb{S}'$ be an L_0-bisimulation. Note that since $L_0 \leq L$ for all $L \in \Lambda$, we know that for each $L \in \Lambda$, R is an L-bisimulation.

We prove by induction on $a \in \mathcal{L}_T(\Lambda)$ that $\mathbb{S}, s \Vdash a$ if and only if $\mathbb{S}', s' \Vdash a$. The only non-trivial cases are the cases $a = \nabla_L \alpha$ and $a = \Delta_L \alpha$ for some $L \in \Lambda$ and $\alpha \in T_\omega \mathcal{L}_T(\Lambda)$.

So, assume that $\mathbb{S}, s \Vdash \nabla_L \alpha$. Then $(\sigma(s), \alpha) \in (\mathcal{L} \Vdash)$. Hence,

$$(\sigma'(s'), \alpha) \in (LR^\circ); (L \Vdash) \subseteq L(R; \Vdash)$$

and by assumption, $R^\circ; \Vdash \subseteq \Vdash$ (when we restrict the codomain of $\Vdash$ to Base(α), see example 3.11). Hence, $(\sigma'(s'), \alpha) \in (L \Vdash)$, showing that $\mathbb{S}', s' \Vdash \nabla_L \alpha$. By symmetry, we now have $\mathbb{S}, s \Vdash a$ if and only if $\mathbb{S}', s' \Vdash a$.

Next, assume that $\mathbb{S}, s \Vdash \Delta_L \alpha$, and assume towards a contradiction that $(\sigma'(s'), \alpha) \in (L \nVdash)$. Then

$$(\sigma(s), \alpha) \in (LR); (L \nVdash) \subseteq L(R; \nVdash) \subseteq (L \nVdash)$$

again by the induction hypothesis. But $\mathbb{S}, s \Vdash \Delta_L \alpha$, and hence $(\sigma(s), \alpha) \notin (L \nVdash)$, and we have a contradiction. So, $(\sigma'(s'), \alpha) \notin (L \nVdash)$, and hence $\mathbb{S}', s' \Vdash \Delta_L \alpha$. By symmetry, we now have that $\mathbb{S}, s \Vdash a$ if and only if $\mathbb{S}', s' \Vdash a$.

So, by induction, we have $\mathbb{S}, s \equiv_\Lambda \mathbb{S}', s'$.

(ii)⇒(i): It suffices to show that $\equiv_\Lambda$ is an L_0-bisimulation. The condition **atom** is obvious. We prove by contraposition that if $\mathbb{S}', s' \equiv_\Lambda \mathbb{S}, s$, then $(\sigma'(s'), \sigma(s)) \in (L_0 \equiv_\Lambda)$.

Assume that $(\sigma'(s'), \sigma(s)) \notin (L_0 \equiv_\Lambda)$. By definition of L_0, this means that there is some $L \in \Lambda$ such that $(\sigma(s), \sigma'(s')) \notin (L \equiv_\Lambda)$.

Since $\mathbb{S}$ and $\mathbb{S}'$ were finite branching, we can take $B = \mathrm{Base}(\sigma(s))$ and $B' = \mathrm{Base}(\sigma'(s'))$, being finite subsets of S and S' respectively. Let $J = \{(t, t') \in B \times B' \mid \mathbb{S}, t \not\equiv_\Lambda \mathbb{S}', t'\}$. Then by definition of $\equiv_\Lambda$, for every $(t, t') \in J$, there is a formula $d_{t,t'} \in \mathcal{L}_T(\Lambda)$ such that (a) $\mathbb{S}, t \Vdash d_{t,t'}$ and $\mathbb{S}', t' \nVdash d_{t,t'}$, or (b) $\mathbb{S}, t \nVdash d_{t,t'}$ and $\mathbb{S}', t' \Vdash d_{t,t'}$.

By replacing $d_{t,t'}$ with $\neg d_{t,t'}$ we can make sure that for every (t, t') we are in case (a). Now define for $t \in B$ the formula a_t as

$$a_t := \bigcap_{(t,t') \in J} d_{t,t'}$$

Then since we had ensured that we were in case (a) for all (t, t'), we see that $\mathbb{S}, t \Vdash a_t$ for all $t \in B$. Hence, if $\mathbb{S}, t \equiv_\Lambda \mathbb{S}', t'$, then $\mathbb{S}', t' \Vdash a_t$. On the other hand, if $\mathbb{S}, t \not\equiv_\Lambda \mathbb{S}', t'$ for $t' \in B'$, then $\mathbb{S}', t' \nVdash a_t$ since $d_{t,t'}$ will be a conjunct in a_t.

We can summarize this as

$$a^{\mathrm{gr}} \subseteq \Vdash, \qquad \equiv_\Lambda \restriction_{B' \times B} = \Vdash; (a^{\mathrm{gr}})^\circ$$

where we see a as a function $B \to \mathcal{L}_T(\Lambda)$.

Now let $\beta := (Ta)\sigma(s)$, and let $b = \nabla_L \beta$. Then

$$(\sigma(s), \beta) \in (Ta)^{\mathrm{gr}} \subseteq L(a^{\mathrm{gr}}) \subseteq (L \Vdash)$$

and hence $\mathbb{S}, s \Vdash b$. On the other hand, we see that if $(\sigma'(s'), \beta) \in (L \Vdash)$, then

$$(\sigma'(s'), \sigma(s)) \in (L \Vdash); (Ta^{\mathrm{gr}})^\circ \subseteq L(\Vdash; (a^{\mathrm{gr}})^\circ) \subseteq L(\equiv_\Lambda)$$

and hence $(\sigma'(s'), \sigma(s)) \in L(\equiv_\Lambda)$, which was not the case by assumption.

Thefore, we conclude that $\mathbb{S}', s' \nVdash \nabla_L \beta$, showing that $\mathbb{S}', s' \not\equiv_\Lambda \mathbb{S}, s$, since they disagree on the formula $\nabla_L \beta$.

□

In [18], it is shown that for L a symmetric T lifting, L-bisimilarity is equivalent to behavioral equivalence if and only if L preserves diagonals. Combining this theorem with theorem 5.2, we obtain the following corollary:

Corollary 5.3. *Let T : **Sets** $\to$ **Sets** be a functor, and let L be a T-lifting such that $L \cap L^\smile$ preserves diagonals. Then for given coalgebra models $\mathbb{S}, \mathbb{S}'$ and points $s \in \mathbb{S}, s' \in \mathbb{S}'$, we have*

$$\mathbb{S}, s \simeq \mathbb{S}', s' \textit{ if and only if } \mathbb{S}, s \equiv_L \mathbb{S}', s'$$

5.2 Downward expressivity

In this section, we seek to answer question 1 from the introduction to this chapter: "If a formula a is defined from a lifting L, and $L' \leq L$, can a also be defined from L'?" It turns out that the answer is "not in general". Consider the following example:

Example 5.4. Let T be the functor $X \mapsto \mathsf{Prop}^\omega$. Then for all relations R, we have $\overline{T}R = \Delta_{\mathsf{Prop}^\omega}$, the diagonal relation (see definition 3.1).

Now consider the lifting L_ω defined as

$$L_\omega R = \{(u, v) \in \mathsf{Prop}^\omega \mid \{i \mid u_i \neq v_i\} \text{ is finite }\}$$

Then clearly, $\overline{T} \leq L_\omega$. Now let p be a proposition letter, and let a be the formula

$$a := \nabla_{L_\omega}\mathbf{p},$$

where $\mathbf{p}$ is the constant function $i \mapsto p$. We claim that a is not equivalent to any formula in $\mathcal{L}_T(\overline{T})$.

Clearly, a is not equivalent to any formula of modal depth 0. For a constant functor T, any $\mathcal{L}_T(\Lambda)$ formula is equivalent to a formula of depth ≤ 1; so it suffices to consider formulas of the form

$$b = \bigvee_{i=1}^{n} \left(\bigwedge_{j=1}^{n_i} \ell_i \right) \wedge \nabla u^i$$

with $u^i : \omega \to \mathsf{Prop}$ for every i.

Assume that $b \vDash a$ for some $\mathcal{L}_T(\overline{T})$-formula b of the above form. Then for every i, we must have $u^i_j = p$ for all but finitely many j; otherwise, we could have $\mathbb{S}, s \Vdash \nabla_{\overline{T}} u^i$ without $\mathbb{S}, s \Vdash \nabla_{L_\omega}\mathbf{p}$. Let j_0 be such that $u^i_{j_0} = p$ for all $i = 1, \ldots, n$. Then if $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S}, s \Vdash b$, we must have $\sigma(s)_{j_0} = p$. Now consider $\mathbb{S} = (\{*\}, \sigma, m)$ with

$$\sigma(*)_j = \begin{cases} p & j \neq j_0 \\ q & j = j_0 \end{cases}$$

where q is any proposition letter distinct from p.

Then clearly, $(\sigma(*), \mathbf{p}) \in L_\omega \Vdash$, and hence $\mathbb{S}, * \Vdash \nabla_{L_\omega}\mathbf{p}$. But also, $(\sigma(*), u^i) \notin \overline{T} \Vdash$ for any i; so $\mathbb{S}, * \nVdash \nabla_{\overline{T}} u^i$ for every i, and hence $\mathbb{S}, * \nVdash b$.

The obstruction in example 5.4 is the fact that only *finite* boolean combinations are allowed; otherwise, we could express $\nabla_{L_\omega}\mathbf{p}$ as

$$\bigvee_{\substack{u:\omega \to \mathsf{Prop} \\ |\{u_i \neq p\}| < \omega}} \nabla_{\overline{T}} u,$$

where the disjunction is taken over all $u : \omega \to \mathsf{Prop}$ such that u_i is different from p for only finitely many i.

It turns out that this 'finitarity problem' is, essentially, the only thing standing in our way. We can remove it by stipulating that T preserve finite sets.

The crucial property of functors that preserve finite sets, is that if $\alpha \in TX$, then $T\,\mathrm{Base}(\alpha)$ is finite, giving us strong normal form theorems, as in section 4.4. In the remainder of this chapter, we will often use this without comment.

Theorem 5.5 (Downwards expressivity)**.** *Let T be a functor preserving finite sets, and let Λ be a set of liftings closed under $(-)^{\breve{}}$. Let L_0 be a lifting such that $L_0 \leq L$ for all $L \in \Lambda$. Then any formula $a \in \mathcal{L}_T(\Lambda)$ is equivalent to a formula $a_0 \in \mathcal{L}_T(L_0)$.*

Proof. Let $\gamma \in T_\omega \mathcal{L}_T(\Lambda)$. Let $L \in \Lambda$, and set $\mathcal{S}_\gamma := \{\Phi \in TP\,\mathrm{Base}(\gamma) \mid (\gamma, \Phi) \in L^{\breve{}} \in\}$. Note that $\mathcal{S}$ is finite, as $\mathrm{Base}(\gamma)$ is finite, so $P\,\mathrm{Base}(\gamma)$ is finite, and T preserves finite sets.

Let $\mathbb{S} = (S, \sigma, m)$ be a coalgebra model, and let $s \in S$. Then consider the following equivalence:

$$\mathbb{S}, s \Vdash \nabla_L \gamma \text{ iff } \mathbb{S}, s \Vdash \bigvee_{\Phi \in \mathcal{S}} \nabla_{L_0} (T \bigwedge) \Phi.$$

This is a special case of the distributive law given in proposition 4.16, where Γ is a singleton $\{\gamma\}$. So, we can define the following partial translation:

Let $\mathcal{L}_T^{\nabla}(\Lambda)$ be the set of Δ-free (T, Λ)-formulas. Then we inductively define

$$\begin{aligned} \tau(\ell) &:= \ell && \ell \text{ is a literal} \\ \tau(\odot A) &:= \odot \tau[A] && \odot \in \{\textstyle\bigwedge, \bigvee\} \\ \tau(\nabla_L \gamma) &:= \bigvee_{\Phi \in \mathcal{S}_\gamma} \nabla_{L_0} (T \textstyle\bigwedge) \Phi \end{aligned}$$

Then by the above discussion, every $a \in \mathcal{L}_T^{\nabla}(\Lambda)$ is equivalent to $\tau(a)$. By proposition 4.23, for every $a \in \mathcal{L}_T(\Lambda)$, there is a formula $a^{\nabla} \in \mathcal{L}_T^{\nabla}(\Lambda)$ such that $a \equiv a^{\nabla}$. So, if we set

$$a_0 := \tau(a^{\nabla})$$

we have obtained the desired translation $\mathcal{L}_T(\Lambda)$ to $\mathcal{L}_T(L_0)$. $\square$

5.3 Upwards expressivity

Next, we aim to answer question 2 in the affirmative. We will prove the following theorem:

Theorem 5.6 (Upwards expressivity)**.** *Let T : **Sets** $\to$ **Sets** be a functor preserving weak pullbacks and finite sets. Let $a \in \mathcal{L}_T(\overline{T})$ be preserved under $L^{\breve{}}$-simulations. Then a is equivalent to a $\mathcal{L}_T^{\nabla}(L)$-formula.*

This theorem is a generalization of the well-known fact that a modal logic formula $a \in \mathcal{L}_{\Box,\Diamond}$ is preserved under simulations if and only if it is equivalent to a negation-free formula containing only diamonds ([6], theorem 2.78).

Our approach is similar to the method of K. Fine in [9]. In that paper, modal logic is shown to have the Final Model Property by taking an arbitrary satisfiable formula a, rewriting it into a particular normal form, and extracting an explicit model from this normal form.

Proof (Sketch). Let $a \in \mathcal{L}_T(\overline{T})$ be preserved under $L^{\sim}$-simulations. Then write a as a pure normal form $\bigvee_{i\in I} a_i$. Define a^L by replacing every ∇ in $\bigvee_{i\in I} a_i$ with ∇_L. Then $a \vDash a^L$, since at every ∇, the condition $\mathbb{S}, s \Vdash \nabla_L \alpha$ is less strict than $\mathbb{S}, s \Vdash \nabla\alpha$.

For the other direction, assume that $\mathbb{S}, s \Vdash a^L$. Then there is an $i \in I$ with $\mathbb{S}, s \Vdash a_i^L$. Now, $\mathrm{Sfor}(a_i)$ can be regarded as a T-coalgebra, such that

- $\mathrm{Sfor}(a_i), a_i \Vdash a_i$;
- $\Vdash_L := \{(s, a) \mid \mathbb{S}, s \Vdash a^L\}$ is an L-simulation from $\mathbb{S}$ to $\mathrm{Sfor}(a_i)$.

Then $\dashv\!\vert_L$ is an $L^{\sim}$-simulation from $\mathrm{Sfor}(a_i), a_i$ to $\mathbb{S}, s$. Since $\mathrm{Sfor}(a_i), a_i \Vdash a_i$, we know that $\mathrm{Sfor}(a_i), a_i \Vdash a$. Since a was preserved under $L^{\sim}$-simulations, we conclude that $\mathbb{S}, s \Vdash a$, showing that $a^L \vDash a$. $\square$

There are many technical details in this argument left to be filled in. The rest of this chapter will be devoted to smoothing out the wrinkles.

5.4 Some semantic notions

In this section, we introduce some operations on and relations between models that will be useful in the proof of theorem 5.6.

Finite simulations

Definition 5.7. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and L a T-lifting. Let $n \geq 0$ be a natural number. Let $\mathbb{S} = (S, \sigma, m)$ and $\mathbb{S} = (S', \sigma', m')$ be coalgebra models. A *(L, n)-simulation from S to S'* is a sequence of relations $(R_0, \ldots, R_n)$ with $R_i : S \multimap S'$ such that

- For all i, if sR_is', then $m(s) = m'(s')$;
- For all $i > 0$, if sR_is', then $(\sigma(s), \sigma'(s')) \in L(R_{i-1})$.

If there is an (L, n)-simulation $(R_i)_{i=0,\ldots,n}$ from $\mathbb{S}$ to $\mathbb{S}'$ with $(s, s') \in R_n$, we will write $\mathbb{S}, s \rightrightarrows_n^L \mathbb{S}', s'$.

If $\mathbb{S}, s \rightrightarrows_n^L \mathbb{S}', s'$ and $\mathbb{S}', s' \rightrightarrows_n^L \mathbb{S}, s$, then we will call $\mathbb{S}, s$ and $\mathbb{S}', s'$ (L, n)-*bisimilar* and write $\mathbb{S}, s \leftrightarrows_n^L \mathbb{S}', s'$.

Proposition 5.8. *Let $\mathbb{S}, \mathbb{S}'$ be T-coalgebra models, and assume $\mathbb{S}, s \leftrightarrows_n^L \mathbb{S}', s'$. Then for any $a \in \mathcal{L}_T(L)$ of modal depth at most n, we have*

$$\mathbb{S}, s \Vdash a \textit{ iff } \mathbb{S}', s' \Vdash a.$$

Backwards unraveling

Definition 5.9. Let $T : \mathbf{Sets} \to \mathbf{Sets}$ be a functor, and fix some default value $\delta \in T\{*\}$. Let $\mathbb{S} = (S, \sigma, m)$ be a T-coalgebra model, and let $n \geq 0$ be a natural number. The *backwards unraveling of* $\mathbb{S}$ is defined as $\rho(\mathbb{S}) = (\rho(S), \rho(\sigma), \rho(m))$ where

$$\begin{aligned}
\rho(S) &:= S \times \omega \cup \{*\} \\
\rho(\sigma)(x) &:= \begin{cases} T\iota_{i-1}(\sigma(s)) & x = (s,i), \text{ for some } s \in S, i > 0 \\ T\iota_*(\delta) & \text{otherwise} \end{cases} \\
\rho(m)(x) &:= \begin{cases} m(s) & x = (s,i) \text{ for some } s \in S, i \in \omega \\ \varnothing & x = * \end{cases}
\end{aligned}$$

where we write $\iota_i : S \to S \times \omega \cup \{*\}$ for the natural map $s \mapsto (s, i)$, and ι_* for the inclusion $\{*\} \hookrightarrow S \times \omega \cup \{*\}$.

We give an illustration of the backwards unraveling for a P-coalgebra in figure 5.

Lemma 5.10. *For every T-coalgebra model $\mathbb{S}$, T-lifting L, natural number $n \geq 0$, default value $\delta \in T\{*\}$, and $s \in \mathbb{S}$, we have*

$$\mathbb{S}, s \leftrightarroweq^L_n \rho(\mathbb{S}), (s, n)$$

Proof. Setting $R_i = \iota_i^{\mathrm{gr}}$ gives a (L, n)-bisimulation by construction. $\square$

Restricted models

Definition 5.11. Let $\mathbb{S} = (S, \sigma, m)$ be a T-coalgebra model. Let $Q \subseteq \mathsf{Prop}$ be some (usually finite) set of proposition letters. We define the *restriction of* $\mathbb{S}$ *to* Q as the T-coalgebra model $\mathbb{S}{\restriction_Q} = (S, \sigma, m^Q)$, with

$$m^Q(s) = m(s) \cap Q$$

Lemma 5.12. *If $\mathbb{S}$ is a T-coalgebra model, $a \in \mathcal{L}_T(\Lambda)$ is a formula, and $Q \subseteq \mathsf{Prop}$ a set of proposition letters with $\mathrm{Var}(a) \subseteq Q$, then for any $s \in \mathbb{S}$, we have*

$$\mathbb{S}, s \Vdash a \textit{ if and only if } \mathbb{S}{\restriction_Q}, s \Vdash a$$

Proof. Simple induction on the complexity of a. $\square$

5.5 Syntactic coalgebras

Pure formulas as coalgebras A generic pure formula has the form $a = \bigwedge P \wedge \nabla_L \alpha$ with $\alpha \in T\mathcal{L}^p_T(\Lambda)$ (see definition 4.26). It is natural to see α as the unfolding of a. In this way, we can view $\mathcal{L}^p_T(\Lambda)$ as a T-coalgebra model, where a formula $a = \bigwedge P \wedge \nabla_L \alpha$ is marked with the positive literals in P, and $\lambda(a) = \alpha$ is the unfolding.

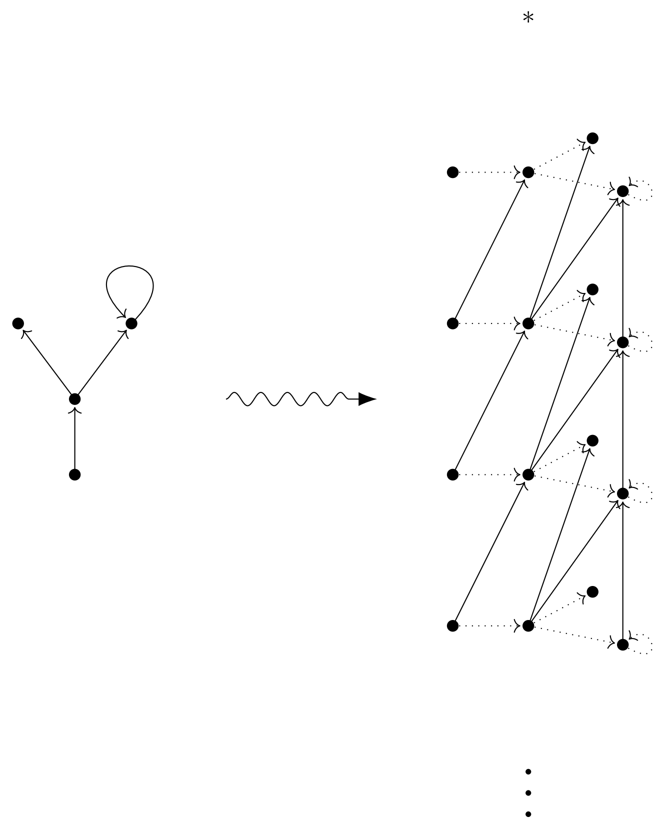

Figure 5: The frame on the right is the backward unraveling of the frame on the left, with default value $\varnothing \in P\{*\}$. It is obtained by taking ω many copies of the old frame, stacking them on top of each other, and shifting the target of each arrow up one level. To illustrate this, we have kept the arrows of the original frame in each copy, shown as dotted arrows.

There is still a gap in this definition, since there is a priori no way to unfold pure formulas of the form $\bigwedge P$. In order to give an unfolding for these formulas, we pick a default value $\delta_\top \in T\{\top\}$. This gives the following definition:

Definition 5.13. Let Λ be a set of T-liftings, and let $\delta_\top \in \{\top\}$ be a default value. Then we define a canonical (T, Λ)-coalgebra $C_T(\Lambda) = (\mathcal{L}_T^p(\Lambda), \lambda, m)$ by setting

$$\lambda(a) := \begin{cases} \alpha & a = \bigwedge P \wedge \nabla_L \alpha \text{ for some } P, \alpha \\ \delta_\top & a = \bigwedge P \text{ for some } P \end{cases}$$

$$m(a) := P \cap \mathsf{Prop} \qquad \text{where } a = \bigwedge P \text{ or } a = \bigwedge P \wedge \nabla\alpha \text{ for some } \alpha$$

Proposition 5.14. *Let a be a pure (T, Λ)-formula. Then $C_T(\Lambda), a \Vdash a$.*

Proof. We prove it by induction on the complexity of a. If a is of the form $\bigwedge P$, then since P is satisfiable, we know that it is of the form $P = P_+ \cup \{\neg p \mid p \in P_-\}$ for *disjoint* sets of proposition letters P_+, P_-. Now $m(a) = P_+$, and so $C_T(\Lambda), a \Vdash p$ for all $p \in P_+$, and since $P_+ \cap P_- = \varnothing$, we know $C_T(\Lambda), a \nVdash p$ for all $p \in P_-$. So, we know

$$C_T(\Lambda), a \Vdash \bigwedge P$$

as desired.

Let C_d be the set of pure formulas of modal depth less than d, and assume that we have proven the proposition for formulas in C_d. Then let $a = \bigwedge P \wedge \nabla\alpha$ be a pure formula of depth d. Then we know that

$$\begin{aligned}
(\alpha, \alpha) &\in \Delta_{TC_T(\lambda)} \cap (TC_d \times TC_d) \\
&= (T\iota^{\mathrm{gr}}_{C_d})^\circ; (T\iota^{\mathrm{gr}}_{C_d}) \\
&\subseteq L(\iota^{\mathrm{gr}}_{C_d})^\circ; L\iota^{\mathrm{gr}}_{C_d} \\
&\subseteq L((\iota^{\mathrm{gr}}_{C_d})^\circ; \iota^{\mathrm{gr}}_{C_d}) \\
&= L(\Delta_{C_T(\Lambda)} \cap (C_d \times C_d)) \\
&\subseteq L \Vdash
\end{aligned}$$

showing that $C_T(\Lambda), a \Vdash \nabla_L\alpha$ (since $\alpha = \lambda(a)$). By the same argument as above, we know that $C_T(\Lambda), a \Vdash \bigwedge P$, and hence we have that

$$C_T(\Lambda), a \Vdash \bigwedge P \wedge \nabla_L\alpha$$

as desired.

$\square$

Now, if $\mathbb{S}$ is a T-coalgebra, then $\Vdash$ is almost an L-simulation from $\mathbb{S}$ to $C_T(L)$. But there are two obstructions.

(1) Once we reach a formula a of modal depth 0, its unfolding is not in any relation to $\sigma(s)$, even if $s \Vdash a$. For this reason, we introduced the backwards unraveling, since this allows us to replace s with a node $\rho_n(s)$ that only has interesting behavior up to depth n. Still, if $a = \nabla\alpha$ is a pure formula of depth n, there may be $a' \in \mathrm{Base}(\alpha)$ of depth much smaller than n, which means that the 'syntactic coalgebra' $\mathrm{Sfor}(a)$ will display default behavior much earlier than $\rho_n(s)$.

(2) If $a = \bigwedge P$ is a conjunction of proposition letters, then the truth condition for $s \Vdash a$ is $P \subseteq m(s)$, while the condition for a simulation is $m(s) = P$.

We will remove these obstructions by considering classes of even more well-behaved formulas.

Definition 5.15. The set $\mathcal{L}^h_{T,d}(\Lambda)$ of *homogeneous formulas of depth d* is inductively defined as follows:

- If p is a proposition letter, then p is homogeneous of depth 0.
- If a is homogeneous of depth d, then $\neg a$ is homogeneous of depth d.
- If A is a finite set of homogeneous formulas of depth d, then $\bigwedge A$ and $\bigvee A$ are homogeneous of depth d.
- If $\alpha \in T_\omega \mathcal{L}^h_{T,d}(\Lambda)$, and $L \in \Lambda$, then $\nabla_L \alpha$ and $\Delta_L \alpha$ are homogeneous of depth $d+1$.

Notation 5.16. Let Q be a finite set of proposition letters, and let $\sigma : Q \to \{0,1\}$ be a function. We write

$$\bigwedge \sigma := \bigwedge (\{p \mid \sigma(p) = 1\} \cup \{\neg p \mid \sigma(p) = 0\})$$

Definition 5.17. Let Q be a finite set of proposition letters. The set of *Q-full formulas* $\mathcal{L}^f_T(\Lambda, Q)$ (or simply *full formulas* if Q is understood form context) is defined as follows:

- If $\sigma : Q \to \{0,1\}$ is a function, then $\bigwedge \sigma$ is full.
- If a is full, then $\neg a$ is full.
- If A is a set of full formulas, then $\bigwedge A$ and $\bigvee A$ are full.
- If $\alpha \in T_\omega \mathcal{L}^f_T(\Lambda, Q)$, and $\sigma : Q \to \{0,1\}$ is a function, then for all $L \in \Lambda$, $\bigwedge \sigma \wedge \nabla_L \alpha$ and $\bigwedge \sigma \wedge \Delta_L \alpha$ are full.

Definition 5.18. If a formula $a \in \mathcal{L}_T(\Lambda)$ is a pure formula that is (Q-)full and homogeneous, we will call a a *full homogeneous pure formula*, or *fhp formula.* We will denote the set of fhp formulas with $\mathcal{L}^{fhp}_T(\Lambda)$.

If a formula $a \in \mathcal{L}_T(\Lambda)$ is a pure normal form with every disjoint being an fhp formula, we will call a a *full homogeneous pure normal form*, or *fhp normal form.*

We will prove in several stages that every formula is equivalent to an fhp normal form.

Lemma 5.19. *Let T be a weak pullback-preserving functor, and assume that T preserves finite sets.*

(i) If $a \in \mathcal{L}^{\mathrm{NF}}_T(\overline{T})$ is a normal form of depth d, and $d' \geq d$, then a is equivalent to a homogeneous normal form of depth d'.

(ii) If $a \in \mathcal{L}^h_T(\overline{T})$ is a homogeneous normal form of depth d, and Q is a finite set of proposition letters with $\mathrm{Var}(a) \subseteq Q$, then a is equivalent to a Q-full homogeneous normal form of depth d.

(iii) Every homogeneous full normal form is equivalent to a hfp normal form.

The proof strategy will be to first pump up the modal depth of too-shallow end points by using the equivalence

$$\top \equiv \bigvee_{\gamma \in T\{\top\}} \nabla\gamma,$$

and second use the equivalence

$$p \equiv (q \wedge p) \vee (\neg q \wedge p)$$

to 'fill out' any conjunction of literals. Both these steps introduce new disjunctions, but these can be pulled out using the distributive law from proposition 4.19.

Proof. (i) As an initial case, we prove that $\top$ is equivalent to a homogeneous formula of depth d for all $d \geq 0$. As $\top \equiv \bigwedge \varnothing$, we know that it is homogeneous of depth 0.

Now assume that $\top$ is equivalent to a homogeneous normal form $\top_d$ of depth d. Then let $j : \{\top\} \to \mathcal{L}^h_{T,d}(\Lambda)$ be the function $\top \mapsto \top_d$. Note that

$$\top \equiv \bigvee_{\gamma \in T\{\top\}} \nabla\gamma$$

and hence

$$\top \equiv \bigvee_{\gamma \in T\{\top\}} \nabla(Tj)\gamma$$

is a homogeneous normal form for $\top$ of depth $d+1$.

Now we can prove by induction that every normal form a is equivalent to a normal form of depth d' whenever $d' \geq a$.

First, assume that $a = \bigwedge P$ is a conjunction of literals. Then let $d \geq 0$; above, we have seen that there is a set S such that $\top \equiv \bigvee_{\gamma \in S} \nabla\gamma$ is a homogeneous normal form for $\top$ of depth d. Now

$$a \equiv \bigwedge P \wedge \top \equiv \bigvee_{\gamma \in S} \bigwedge P \wedge \nabla\gamma$$

is a homogeneous normal form for a of depth d.

Next, assume that $a = \bigwedge P \wedge \nabla\alpha$ for some $\alpha \in T_\omega \mathcal{L}^{\mathrm{NF}}_T(\overline{T})$. Let $d' \geq d(a)$. Then for every $b \in \mathrm{Base}(\alpha)$, we know $d(b) < d(a) \leq d'$, so there is a homogeneous normal form $h(b)$ of depth $d'-1$. Now

$$a \equiv \bigwedge P \wedge \nabla(Th)\alpha$$

is a homogeneous normal form for a of depth d'.

Finally, if $a = \bigvee A$ for some set A of normal forms, then let $d' \geq d(a)$. Clearly, for every $b \in A$ we have $d(b) \leq d'$, so there is a homogeneous normal form $h(b)$ of depth d'. Now

$$a \equiv \bigvee_{b \in A} h(b)$$

is a homogeneous normal form for a of depth d'.

(ii) We proceed by induction on the complexity of the formula. First, let $a = \bigwedge P$ be a set of literals, and assume that $\mathrm{Var}(a) \subseteq Q$. Then if P is inconsistent, we know that $a \equiv \bigvee \varnothing$, which is a Q-full normal form. So, we assume that P is consistent. Now, consider the partial function

$$\sigma_0 : Q \rightharpoonup \{0,1\} : q \mapsto \begin{cases} 1 & q \in P \\ 0 & \neg q \in P \\ \text{undefined} & q \notin P \text{ and } \neg q \notin P \end{cases}$$

Let $S = \{\sigma : Q \to \{0,1\} \mid \sigma \text{ extends } \sigma_0\}$. It is now clear that

$$a \equiv \bigvee_{\sigma \in S} \bigwedge \sigma$$

showing that a is equivalent to a full homogeneous normal form.

Next, let a be homogeneous of the form $\bigwedge P \wedge \nabla\alpha$ for P a set of literals, and $\alpha \in T_\omega \mathcal{L}_T^{\mathrm{NF}}(\overline{T})$. If P is inconsistent, then $a \equiv \bigvee \varnothing$, so we may assume that P is consistent. Assume that $\mathrm{Var}(a) \subseteq Q$. Then by induction, every $a \in \mathrm{Base}(\alpha)$ is equivalent to a Q-full homogeneous normal form $f(a)$. Now, as before, let $\sigma_0 : Q \rightharpoonup \{0,1\}$ be the partial function defined as

$$\sigma_0 : p \mapsto \begin{cases} 1 & p \in P \\ 0 & \neg p \in P \\ \text{undefined} & p \notin P, \neg p \notin P \end{cases}$$

and let $S = \{\sigma : Q \to \{0,1\} \mid \sigma \text{ extends } \sigma_0\}$. Then

$$a \equiv \bigvee_{S} \left(\bigwedge \sigma \wedge \nabla (Tf)\alpha \right)$$

which is a full homogeneous normal form.

Finally, if a is of the form $\bigvee A$ for A a finite set of homogeneous normal forms, then assume $\mathrm{Var}(a) \subseteq Q$. By induction, for every $b \in A$ we have a Q-full homogeneous normal form $f(b)$, and hence

$$a \equiv \bigvee_{b \in A} f(b)$$

is a Q-full homogeneous normal form for a.

(iii) This proof is analogous to the proof of 4.27. In particular, let us look at the case $a = \bigwedge P \wedge \nabla\alpha$ for P consistent.

Let d be the modal depth of a, and let $d' \geq d$. Let Q be a finite set of proposition letters with $\mathrm{Var}(a) \subseteq Q$. By induction, we know that for every $b \in a$, there is a set $A(b) \subseteq \mathcal{L}^{hfp}_{T,d'-1}(\overline{T}, Q)$ such that $b \equiv \bigvee A(b)$. Hence, we see that

$$a \equiv \bigwedge P \wedge \nabla\alpha \equiv \bigwedge P \wedge \nabla(T \bigvee \circ TA)\alpha$$

Now let $B = \mathcal{B}((TA)\alpha)$, and let $\mathcal{A} := \{\beta \in T_\omega \mathcal{L}_T(\Lambda) \mid (\beta, (TA)\alpha) \in (\overline{T}\in)\}$. Then $\mathcal{A}$ is finite: by lemma 3.17, if $\beta \in \mathcal{A}$, then $\mathrm{Base}(\beta) \subseteq B$. Since T preserves finite sets by assumption, we know that TB is finite. Since every element of $\mathcal{A}$ is of the form $T\iota(\gamma)$ for some $\gamma \in TB$, we know that $\mathcal{A}$ must be finite as well.

Now by proposition 4.19, we know that

$$\nabla(T \bigvee \circ TA)\alpha \equiv \bigvee_{\beta \in \mathcal{A}} \nabla\beta$$

Note that $\mathcal{A} \subseteq T_\omega \mathcal{L}^{hfp}_{T,d'-1}(\overline{T}, Q)$, and hence for each $\beta \in \mathcal{A}$, we have $\beta \in T_\omega \mathcal{L}^{hfp}_{T,d'-1}$. Set $\sigma_0 : Q \rightharpoonup \{0,1\}$ given by

$$\sigma_0 : p \mapsto \begin{cases} 1 & p \in P \\ 0 & \neg p \in P \\ \text{undefined} & p \notin P, \neg p \notin P \end{cases},$$

and let $S = \{\sigma : Q \to \{0,1\} \mid \sigma \text{ extends } \sigma_0\}$. Then

$$a \equiv \bigvee_{\beta \in \mathcal{A}} \bigvee_{\sigma \in S} \left(\bigwedge \sigma \wedge \nabla\beta \right)$$

is an hfp normal form for a.

□

Note that for all sets Q of proposition letters, $\mathcal{L}^{fhp}_T(\overline{T}, Q) \cup \{\top\}$ is a subcoalgebra of the canonical T-coalgebra $C_T(\overline{T})$ (see definition 5.13). We will write this subcoalgebra as $C^{fhp}_T(Q)$. We will indicate the set of fhp formulas of depth d with $C^{fhp}_{T,d}(Q)$.

Lemma 5.20. *Let Q be a finite set of proposition letters. Let $\mathbb{S} = (S, \sigma, m)$ be a T-coalgebra model, such that $m(s) \subseteq Q$ for all s. Let*

$$R_i := \{((s,i), a) \mid s \in S \text{ and } a \in C^{fhp}_{T,i}(Q), \mathbb{S}, s \Vdash_L a\}.$$

Let $\mathbb{S}' = \rho(\mathbb{S})$ be the backwards unraveling of $\mathbb{S}$. Then

$$R = \bigcup_{i \leq n} R_i \cup \{(*, \top)\}$$

is an L-simulation from $\mathbb{S}'$ to $C^{fhp}_T(Q)$.

The definition of $\rho(\mathbb{S})$ depends on some default value $\delta \in \{*\}$, and the definition of $C_T^{fhp}(Q)$ depends on some default value $\delta_\top \in \{\top\}$. There is, of course, a unique isomorphism $! : \{*\} \to \{\top\}$. We choose δ and $\delta_\top$ such that $(T!)\delta = \delta_\top$.

Proof. First, let $(s, i) \in S'$, let $a \in C_{T,i}^{fhp}(Q)$. Then a is of the form $\bigwedge \sigma$ or $\bigwedge \sigma \wedge \nabla \alpha$ for some function $\sigma : Q \to \{0, 1\}$. We see that if $\mathbb{S}, s \Vdash_L a$, then

$$\begin{aligned} q \in m(s) &\text{ if and only if } \mathbb{S}, s \Vdash q \\ &\text{ if and only if } \sigma(q) = 1 \\ &\text{ if and only if } q \in m(a). \end{aligned}$$

Note that if $\mathbb{S}, s \Vdash q$, then $\sigma(q)$ is well-defined, as $m(s) \subseteq Q$ by assumption.

Additionally, we of course have $m(*) = \varnothing = m(\top)$.

Next, we prove that if $((s, i+1), a) \in R_{i+1}$, then $(\rho(\sigma)((s, i+1)), \lambda(a)) \in LR_i$. But this is clear: if $((s, i), \bigwedge P \wedge \nabla\alpha) \in R_{i+1}$, then $(\sigma(s), \alpha) \in L \Vdash_L$. Taking into account the restricted domain of R_i, we have

$$(\rho(\sigma)((s, i+1)), \alpha) \in (T\iota_i^{\text{gr}})^\circ; L \Vdash; Tg_i \subseteq L((\iota_i^{\text{gr}})^\circ; \Vdash; g_i) = LR_i$$

where we write g_i for the inclusion $C_{T,i}^{fhp}(Q) \hookrightarrow C_T^{fhp}(Q)$.

If $((s, 0), a) \in R_0$, then

$$(\rho(\sigma)((s, 0)), \lambda(a)) = (\delta, \delta_\top) \in T!^{\text{gr}} \subseteq L!^{\text{gr}} = L\{(*, \top)\} \subseteq LR_0 \subseteq LR,$$

which also shows that $(\rho(\sigma)(*), \lambda(\top)) \in R$.

□

5.6 Proof of the upwards expressivity theorem

Now we are ready to prove theorem 5.6.

Proof. Let a be a $\mathcal{L}_T(\overline{T})$-formula of depth d, and assume that a is invariant under $L^{\sim}$-simulations. Let Q be the set of variables occurring in a, and let a' be an fhp normal form for a of depth at least d. Then let $(a')^L$ be the formula a' with each instance of ∇ replaced with ∇_L. We claim that a is equivalent to $(a')^L$.

First, we prove that $a \vDash a^L$ for all Δ-free formulas a. This is a simple induction on the complexity - if a is of modal depth 0, then a^L is identical to a. If a is of the form $\odot A$ for $\odot \in \{\bigwedge, \bigvee\}$, then

$$a \vDash \odot\{b^L \mid b \in A\} = a^L$$

Finally, if $a = \nabla\alpha$, then by the induction hypothesis, we know $((-)^L)^{\text{gr}} \subseteq \vDash$ if we restrict to $\text{Base}(\alpha)$. Hence,

$$(\alpha, (T(-)^L)\alpha) \in L \vDash .$$

Now if $\mathbb{S}, s \Vdash \nabla\alpha$, then

$$(\sigma(s), (T(-)^L)\alpha) \in (\overline{T} \Vdash); (L \vDash) \subseteq L(\Vdash; \vDash) \subseteq L \Vdash$$

showing that $\mathbb{S}, s \Vdash \nabla_L (T(-)^L)\alpha = a^L$, as desired.

Since a' was fhp, it is Δ-free, and so

$$a \equiv a' \vDash (a')^L$$

which is one of the two directions.

Now assume that $\mathbb{S}, s \Vdash (a')^L$. Take $\mathbb{S}' = \mathbb{S}\restriction_Q$. By lemma 5.12, we know that $\mathbb{S}', s \Vdash (a')^L$. Let $\rho(\mathbb{S}')$ be the backwards unraveling of $\mathbb{S}'$. Then since $\mathbb{S}', s$ and $\rho(\mathbb{S}'), (s, d)$ are d-bisimilar, we know $\rho(\mathbb{S}'), (s, d) \Vdash (a')^L$; that is, $\rho(\mathbb{S}'), (s, d) \Vdash_L a'$.

Since a' is an fhp normal form, we know that it is of the form $\bigvee A$, for A a set of fhp formulas. So, there is an $a_0 \in A$ with $\rho(\mathbb{S}'), (s, d) \Vdash_L a_0$. By lemma 5.20, we know that $C_T^{fhp}(Q), a_0 \rightrightarrows^{L^\sim} \rho(S'), (s, d)$. By proposition 5.14, we know that $C_T^{fhp}(Q), a_0 \Vdash a_0$. As a_0 is a disjunct in a', we know that $C_T^{fhp}(Q), a_0 \Vdash a'$. Since a and a' are equivalent, we conclude that $C_T^{fhp}(Q), a_0 \Vdash a$. Since a iss preserved under $L^\sim$-simulations, we now have that $\rho(\mathbb{S}'), (s, d) \Vdash a$. And $\rho(\mathbb{S}'), (s, d)$ iss d-equivalent to $\mathbb{S}', s$, and d is at least the modal depth of a. So, $\mathbb{S}', s \Vdash a$. Finally, since $\mathrm{Var}(a) \subseteq Q$, we have by lemma 5.12 that $\mathbb{S}, s \Vdash a$.

Hence, $(a')^L \vDash a$ as desired. $\square$

Remark 5.21. We used the assumption that T preserves weak pullbacks to argue that every formula in $\mathcal{L}_T(\overline{T})$ has a pure normal form. In [21], it is shown that for the monotone neighbourhood functor, disjunctions below a $\nabla_{\widetilde{\mathcal{M}}}$ can be eliminated, yielding a distributive law similar to proposition 4.19. From this, we see that the proof of theorem 5.6 will also apply to $\mathcal{L}_T(\nabla_{\widetilde{\mathcal{M}}})$.

6 A uniform sequent calculus

In [5], a sound and complete sequent calculus is given for the $\mathcal{L}_T(\overline{T})$-formulas, where T is a weak-pullback preserving functor. In this chapter, this sequent calculus is extended to $\mathcal{L}_T(\Lambda)$-formulas for arbitrary functors T and sets of liftings Λ.

6.1 The sequent calculus

Definition 6.1. Let T be a functor and Λ a set of T-liftings. A *(T, Λ)-sequent* is an expression of the form $A \implies B$, where A and B are finite sets of $\mathcal{L}_T(\Lambda)$-formulas.

A sequent $A \implies B$ should be read as an implication

$$\bigwedge A \to \bigvee B.$$

In light of this reading, we will say that a sequent $A \implies B$ is *valid* if the following holds: For every pointed coalgebra model $\mathbb{S}, s$ such that $\mathbb{S}, s \Vdash a$ for all $a \in A$, there is a $b \in B$ such that $\mathbb{S}, s \Vdash b$. We will call a sequent *refutable* if it is not valid.

Our goal is to obtain a sequent calculus $G2_T(\Lambda)$ such that a sequent Γ is derivable in $G2_T(\Lambda)$ if and only if it is valid. We build our sequent on top of the propositional sequent calculus G2. The rules of G2 are the following:

$$\text{init}\,\frac{}{A, p \implies p, B} \qquad \textstyle\bigwedge\text{-l}\,\frac{A, A' \implies B}{A, \bigwedge A' \implies B} \qquad \bigwedge\text{-r}\,\frac{\{A \implies b, B \mid a \in B'\}}{A \implies \bigwedge B', B}$$

$$\textstyle\bigvee\text{-l}\,\frac{\{A, a \implies B \mid a \in A'\}}{A, \bigvee A' \implies B} \qquad \bigvee\text{-r}\,\frac{A \implies B', B}{A \implies \bigvee B', B}$$

$$\neg\text{-l}\,\frac{A \implies b, B}{A, \neg b \implies B} \qquad \neg\text{-r}\,\frac{A, a \implies B}{A \implies \neg a, B}$$

Note that $A \implies B$ is an initial sequent if there is a *proposition letter* p with $p \in A$ and $p \in B$.

It should be noted that weakening is not a rule in G2; though it is admissible. The exchange and contraction rules are implicit, as in this thesis sequents are considered pairs of *sets*.

Notation 6.2. If $\heartsuit$ is a modality, we will write $\heartsuit\Gamma$ for a set of the form

$$\{\heartsuit_{L_\gamma}\gamma \mid \gamma \in \Gamma\}$$

with $\Gamma \in P_\omega T_\omega \mathcal{L}$, and for every $\gamma \in \Gamma$ an associated lifting L_γ.

Separated joint slim redistributions In order to define the modal rule, we first introduce the concept of the *separated redistribution*. Recall the redistributions we had defined in proposition 4.16. A similar notion will allow us to reduce validity of a sequent to validity of sequents of smaller modal depth. However, when dealing with sequents, we will need to keep track of which side the formulas in the Base originate from.

Definition 6.3. Let $\nabla\Gamma, P \implies Q, \Delta\Theta$ be a sequent, where P, Q are sets of proposition letters, and $\nabla\Gamma, \Delta\Theta$ are of the form

$$\nabla\Gamma = \{\nabla_{L_\gamma}\gamma \mid \gamma \in \Gamma\}, \quad \Delta\Theta = \{\Delta_{L_\theta}\theta \mid \theta \in \Theta\}.$$

Then we define

$$\mathcal{B}^0(\Gamma) = \{(a,0) \mid a \in \mathcal{B}(\Gamma)\}, \quad \mathcal{B}^1(\Theta) = \{(b,1) \mid b \in \mathcal{B}(\Theta)\}$$

and let $f_0 : \mathcal{B}(\Gamma) \to \mathcal{B}^0(\Gamma), f_1 : \mathcal{B}(\Theta) \to \mathcal{B}^1(\Theta)$ be the natural maps. We identify $\mathcal{B}(\Gamma) \uplus \mathcal{B}(\Theta)$ with $\mathcal{B}^0(\Gamma) \cup \mathcal{B}^1(\Theta)$, with f_0, f_1 as canonical inclusions.

The set of *separated redistributions of* $\nabla\Gamma, P \implies Q, \Delta\Theta$ is defined as

$$\mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta) = \Big\{\Phi \in TP(\mathcal{B}(\Gamma) \uplus \mathcal{B}(\Theta)) \mid \text{ if } \gamma \in \Gamma, \text{ then } ((Tf_0)\gamma, \Phi) \in (L_\gamma^{\smile}\in), \\ \text{and if } \theta \in \Theta, \text{ then } ((Tf_1)\theta, \Phi) \in (L_\theta^{\smile}\in)\Big\}.$$

If $\Phi \in \mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta)$ is a separated redistribution, then an element $A \in \text{Base}(\Phi)$ is some subset of $\mathcal{B}(\Gamma) \uplus \mathcal{B}(\Theta)$. That is, it is of the form

$$A = \{(a,0) \mid a \in A_l\} \cup \{(b,1) \mid b \in A_r\}$$

for some sets $A_l \subseteq \mathcal{B}(\Gamma), A_r \subseteq \mathcal{B}(\Theta)$. Note that $A_l = \breve{P}f_0(A), A_r = \breve{P}f_1(A)$. This will be our formal definition:

Definition 6.4. Let $\Gamma, \Theta \in P_\omega T_\omega \mathcal{L}$, and let $A \in P(\mathcal{B}(\Gamma) \uplus \mathcal{B}(\Theta))$. We define $A_l := (\breve{P}f_0)A$ and $A_r := (\breve{P}f_1)A$.

Then we can see A as being itself a sequent $A_l \implies A_r$. This will be important in defining the modal rule in the sequent calculus $G2_T(\Lambda)$.

The sequent calculus $SC_T(\Lambda)$ To the propositional calculus $G2$, we add the following rules governing the modalities:

$$\Delta_{T,L}\text{-l}\frac{\{A, \nabla_{L^\smile}\beta \implies B \mid \beta \in S_L\alpha\}}{\Delta_L\alpha, A \implies B} \quad \nabla_{T,L}\text{-r}\frac{\{A \implies \Delta_{L^\smile}\beta, B \mid \beta \in D_L\alpha\}}{A \implies \nabla_L\alpha, B}$$

$$T(\nabla\Delta)\frac{\{A_l^\Phi \implies A_r^\Phi \mid \Phi \in \mathcal{R}^s(\Gamma,\Theta)\}}{\nabla\Gamma, P \implies Q, \Delta\Theta}\forall\Phi.A^\Phi \in \text{Base}(\Phi)$$

Recall the notations S_L, D_L from definition 4.12. The rule $T(\nabla\Delta)$ should be read as follows: If P and Q are sets of proposition letters, and Γ, Θ are as described in 6.2, the sequent $\nabla\Gamma, P \implies Q, \Delta\Theta$ can be concluded from a premise set Π if for every $\Phi \in \mathcal{R}^s(\Gamma,\Theta)$ there is some $A^\Phi \in \text{Base}(\Phi)$ with $A_l^\Phi \implies A_r^\Phi \in \Pi$.

Notation 6.5. In text, we will call the $\Delta_{T,L}$-l and $\nabla_{T,L}$-r rules the *interchange* rules, and the $T(\nabla\Delta)$ rule the *modal* rule.

Remark 6.6. The modal rule $T(\nabla\Delta)$ can be understood by thinking of proofs as two-player games, played between a *Prover* who aims to show that a sequent is provable, and a *Refuter* who aims to show that it is not. On a sequent $A \implies B$, the Prover proposes a rule with $A \implies B$ as its conclusion. Ordinarily, the Refuter then chooses a premise as the next position, at which point it is Prover's turn again, who chooses a rule, etc. The game continues until an initial sequent is reached, at which point Prover wins; or a sequent is reached that is not the conclusion of any proof rule, at which point the Refuter wins.

The behavior of the modal rule is special: if the Prover proposes the modal rule, then the Refuter does not pick a premise, but rather a separated redistribution Φ, at which point the Prover is allowed to pick an $A \in \mathrm{Base}(\Phi)$ with which to continue.

In general, a proof of $A \implies B$ in $G2_{T,L}$ proceeds as follows: first, a series of propositional rules reduces every formula to one that starts with a modal. Next, the interchange rules are applied until all modals occurring on the left are ∇'s, and those on the right are Δ's. Finally, the modal rule $T(\nabla\Delta)$ is applied, which reduces the modal depth.

Examples We give some example derivations to illustrate $G2_T(\Lambda)$.

Example 6.7. Let T be the functor $X \mapsto X^2$. Consider the $\overline{T}$-sequent

$$\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle.$$

We see that in any proof of this sequent, the last rule applied will be the modal rule. Therefore, we will want to calculate $\mathcal{R}^s(\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle)$.

An element of $\mathcal{R}^s(\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle)$ will be an element of $TP\{(p, 0), (q, 0), (p, 1)\}$; that is a pair $\langle A, B\rangle$ with $A, B \subseteq \{(p, 0), (q, 0), (p, 1)\}$. Such a pair is an element of $\mathcal{R}^s(\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle)$ if it satisfies the following conditions:

1. $(p, 0) \in A$;
2. $(q, 0) \in B$;
3. $(p, 1) \in A$;
4. $(p, 1) \in B$.

Conditions 1 and 2 together ensure $((Tf_0)\langle p, q\rangle, \langle A, B\rangle) \in \overline{T}\in$, and conditions 3 and 4 ensure that $((Tf_1)\langle p, p\rangle, \langle A, B\rangle) \in \overline{T}\in$.

For such an $\langle A, B\rangle \in \mathcal{R}^s(\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle)$, we have that $\mathrm{Base}(A, B) = \{A, B\}$. Moreover, A will be equal to either $\{(p, 0), (p, 1)\}$ or $\{(p, 0), (q, 0), (p, 1)\}$. So, we see that

$$T(\nabla\Delta)\,\frac{p \implies p \quad p, q \implies p}{\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle}$$

is a valid instance of the modal rule. Since both $p \implies p$ and $p, q \implies p$ are initial sequents, this is already a complete proof. So $\nabla(p, q) \implies \Delta(p, p)$ is a derivable sequent.

It is instructive to look at the above proof from a semantic perspective. The formula $\nabla\langle p, q\rangle$ is true at $\mathbb{S}, s$ if p is true at the left successor, and q is true at the right successor. The formula $\Delta\langle p, p\rangle$ is true at $\mathbb{S}, s$ if p is true at either the left or the right successor. So, the sequent $\nabla\langle p, q\rangle \implies \Delta\langle p, p\rangle$ is valid. The proof also reveals *why* it is so: given that $\nabla\langle p, q\rangle$ is true, we can move to the left successor to find a witness for the truth of $\Delta\langle p, p\rangle$. This is reflected in the fact that for all separated redistributions, we chose the *left* element as our base element.

Example 6.8. Consider the $P, \{\overrightarrow{P}, \overleftarrow{P}\}$-sequent $\implies \overrightarrow{\nabla}\varnothing, \overleftarrow{\nabla}\{\top\}$. This sequent expresses that every point in a Kripke frame either has zero successors or at least one successor.

We will give a derivation for this sequent. Using the calculations from example 4.13, we have that

$$
\begin{aligned}
D_{\overrightarrow{P}}(\varnothing) &= \{\{\bigvee \varnothing\}\} \\
D_{\overleftarrow{P}}(\{\top\}) &= \{\{\bigvee\{\top\}\}\}.
\end{aligned}
$$

This lets us build up a derivation

$$
\dfrac{\dfrac{\dfrac{\cdots}{\implies \overleftarrow{\Delta}\{\bigvee \varnothing\}, \overrightarrow{\Delta}\{\bigvee\{\top\}\}}}{\implies \overleftarrow{\Delta}\{\bigvee \varnothing\}, \overleftarrow{\nabla}\{\top\}}\ \overleftarrow{\nabla}\text{-r}}{\implies \overrightarrow{\nabla}\varnothing, \overleftarrow{\nabla}\{\top\}}\ \overrightarrow{\nabla}\text{-r}
$$

with two applications of the interchange rule. We are now at a point where we can apply the modal rule. An element of $\mathcal{R}^s(\implies \overrightarrow{\nabla}\varnothing, \overleftarrow{\nabla}\{\top\})$ is some set Φ of subsets of $\{(\bigvee \varnothing, 1), (\bigvee\{\top\}, 1)\}$ such that

- $(\{(\bigvee \varnothing, 0)\}, \Phi) \in (\overrightarrow{P}\in)$; that is, there is an element $A \in \Phi$ with $(\bigvee \varnothing, 1) \in A$.
- $(\{(\bigvee\{\top\}, 1)\}, \Phi) \in (\overleftarrow{P}\in)$; that is, all elements of Φ contain $(\bigvee\{\top\}, 1)$ as an element.

From this, it follows that a separated redistribution always contains $\{(\bigvee \varnothing, 1), (\bigvee\{\top\}, 1)\}$ as an element, and may additionally contain $\{(\bigvee\{\top\}, 1)\}$ as an element.

We see that both separated redistributions have $\{(\bigvee\{\top\}, 1), (\bigvee \varnothing, 1)\}$ as a base element; so, we can give our derivation as

$$\overset{\rightarrow}{\nabla}\text{-r}\,\dfrac{\overleftarrow{\nabla}\text{-r}\,\dfrac{\mathrm{P}(\nabla\Delta)\,\dfrac{\bigvee\text{-r}\,\dfrac{\bigvee\text{-r}\,\dfrac{\bigwedge\text{-r}\,\dfrac{}{\implies \top}}{\implies \bigvee\{\top\}}}{\implies \bigvee\{\top\}, \bigvee\varnothing}}{\implies \overleftarrow{\Delta}\{\bigvee\varnothing\}, \overrightarrow{\Delta}\{\bigvee\{\top\}\}}}{\implies \overleftarrow{\Delta}\{\bigvee\varnothing\}, \overleftarrow{\nabla}\{\top\}}}{\implies \overrightarrow{\nabla}\varnothing, \overleftarrow{\nabla}\{\top\}}$$

where we recall that $\top$ is an abbreviation for $\bigwedge\varnothing$.

Example 6.9. Consider the $\mathcal{M}$-sequent $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \vee \overset{\rightsquigarrow}{\nabla}\langle\{q\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle$. We first need to know what $\mathcal{D}_{\overset{\rightsquigarrow}{\mathcal{M}}}(\langle\{p\},\{q\}\rangle)$ is. It turns out that $\Phi \in \mathcal{D}_{\overset{\rightsquigarrow}{\mathcal{M}}}(\langle\{p\},\{q\}\rangle)$ if and only it satisfies the following three conditions:

(i) $\mathrm{Base}(\Phi) \subseteq P\{p,q\}$;

(ii) There is an $\alpha_p \in \Phi$ such that every $A \in \alpha_p$ contains p;

(iii) There is an $\alpha_q \in \Phi$ such that every $A \in \alpha_q$ contains q.

From this, it follows that $\nu \in D_{\overset{\rightsquigarrow}{\mathcal{M}}}(\langle\{p\},\{q\})$ if and only if $\mathrm{Base}(\nu) \subseteq \{\bigvee\varnothing, \bigvee\{p\}, \bigvee\{q\}, \bigvee\{p,q\}\}$, and there are $A_p, A_q \in \nu$ such that every disjunction in A_p contains a p, and every disjunction in A_q contains a q.

So, let $\nu \in D_{\overset{\rightsquigarrow}{\mathcal{M}}}(\langle\{p\},\{q\})$. We need to provide a proof of $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\leftsquigarrow}{\Delta}\nu$; the only applicable rule is the modal rule. If Φ is a separated redistribution of $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\leftsquigarrow}{\Delta}\nu$, then all $\alpha \in \Phi$ contain a B with $(p,0) \in B$, and there is $\alpha_p \in \Phi$ such that every $A \in \alpha_p$ contains an element of the form $(\bigvee P, 1)$ with $P \ni p$ (this the condition for $(\nu,\Phi) \in (\overset{\rightsquigarrow}{\mathcal{M}}\in)$ applied to $A_p \times \{1\}$).

Then let α'_p be equal to $\alpha_p \cap \mathrm{Base}(\Phi)$. By the criterion from remark 2.10, $\alpha'_p \in \Phi$. Then we know that there is a $A^\Phi \in \alpha'_p$ with $(p,0) \in A^\Phi$. Moreover, since $\alpha'_p \subseteq \alpha_p$, we know A^Φ has an element of the form $(\bigvee P, 1)$ with $p \in P$. Since $\alpha'_p \subseteq \mathrm{Base}(\Phi)$ by definition, we know that $A^\Phi \in \mathrm{Base}(\Phi)$.

Now for this A^Φ, the sequent $A^\Phi_l \implies A^\Phi_r$ is of the form $X, p \implies \bigvee P, Y$ with $p \in P$, and this is easily derivable by the $\bigvee$-r rule.

So, for all $\Phi \in \mathcal{R}^s(\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\leftsquigarrow}{\Delta}\nu)$, there is a derivable $A^\Phi \in \mathrm{Base}(\Phi)$, and so by the modal rule we can derive $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\leftsquigarrow}{\Delta}\nu$. Since ν was arbitrary, we can use interchange to derive $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle$. By the same argument, we can derive $\overset{\rightsquigarrow}{\nabla}\langle\{q\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle$; so, finishing off with

$$\bigvee\text{-l}\,\dfrac{\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle \quad \overset{\rightsquigarrow}{\nabla}\langle\{q\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle}{\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \vee \overset{\rightsquigarrow}{\nabla}\langle\{q\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle}$$

we have shown that the sequent $\overset{\rightsquigarrow}{\nabla}\langle\{p\}\rangle \vee \overset{\rightsquigarrow}{\nabla}\langle\{q\}\rangle \implies \overset{\rightsquigarrow}{\nabla}\langle\{p\},\{q\}\rangle$ is derivable.

From these examples, it can be seen that actually using $G2_T(\Lambda)$ can be quite complicated, even for relatively simple sequents.

6.2 Soundness

In this section, we will prove soundness of $G2_{T,L}$. That is, we will show that if a sequent $A \implies B$ is derivable, it is valid. As is standard, we proceed by proving that every rule is sound; by which we mean that if

$$\frac{\Pi}{A \implies B}$$

is any instance of a rule, where Π is a set of valid premises, then $A \implies B$ is a valid sequent.

We will not concern ourselves with the soundness of the Boolean rules, and only focus on the rules governing the modals.

Interchange rules To prove soundness of the interchange rules, we make heavy use of the duality lemma (proposition 4.14).

Proof. Let

$$\frac{\{A, \nabla_{L^\sim}\beta \implies B \mid \beta \in S_L\alpha\}}{A, \Delta_L\alpha \implies B}$$

be an instance of the $\nabla_{T,L}$-l rule for some L, and assume that for each $\beta \in S_L\alpha$, the sequent $A, \nabla_{L^\sim}\beta \implies B$ is valid.

Then let $\mathbb{S}, s$ be a pointed coalgebra model, and assume that $\mathbb{S}, s \Vdash a$ for all $a \in A \cup \{\Delta_L\alpha\}$. Then by the duality lemma, we know that there is a $\beta \in S_L\alpha$ with $\mathbb{S}, s \Vdash \nabla_{L^\sim}\beta$. So now $\mathbb{S}, s \Vdash a$ for all $a \in \{A\} \cup \nabla_{L^\sim}\beta$, and since by assumption the sequent $A, \nabla_{L^\sim}\beta \implies B$ was valid, we know that there is some $b \in B$ with $\mathbb{S}, s \Vdash b$.

We conclude that the sequent $A, \Delta_L\alpha \implies B$ is valid.

Next, let

$$\frac{\{A \implies \Delta_{L^\sim}\beta, B \mid \beta \in D_L\alpha\}}{A \implies \nabla_L\alpha, B}$$

be an instance of the $\Delta_{T,L}$-r rule for some L, and assume that for each $\beta \in D_L\alpha$, the sequent $A \implies \Delta_{L^\sim}\beta, B$ is valid.

Then let $\mathbb{S}, s$ be a pointed coalgebra model, and assume that $\mathbb{S}, s \Vdash a$ for all $a \in A$. Then since for every $\beta \in D_L\alpha$, the sequent $A \implies \Delta_{L^\sim}\beta, B$ is valid, we know that either $\mathbb{S}, s \Vdash b$ for some $b \in B$, or else $\mathbb{S}, s \Vdash \Delta_{L^\sim}\beta$ for all $\beta \in D_L\alpha$.

In the first case, we clearly have that $\mathbb{S}, s \Vdash b$ for some $b \in B \cup \{\nabla_L\alpha\}$. In the second case, we know by duality that $\mathbb{S}, s \Vdash \nabla_L\alpha$, and hence $\mathbb{S}, s \Vdash b$ for some $b \in B \cup \{\nabla_L\alpha\}$. So, we conclude that the sequent $A \implies \nabla_L\alpha, B$ is valid.

□

Remark 6.10. Note that in fact, the interchange rules are *invertible*; the conclusion is valid if *and only if* all the premises are valid.

Modal rule To prove soundness of the modal rule, we need the following key lemma:

Lemma 6.11 (Refutability lemma)**.** *Let P, Q sets of proposition letters, $\nabla\Gamma, \Delta\Theta$ of the form*

$$\nabla\Gamma = \{\nabla_{L_\gamma}\gamma \mid \gamma \in \Gamma\}, \quad \Delta\Theta = \{\Delta_{L_\theta}\theta \mid \theta \in \Theta\}.$$

The following are equivalent:

(i) The sequent $\nabla\Gamma, P \implies Q, \Delta\Theta$ is refutable;

(ii) $P \cap Q = \varnothing$, and there is a $\Phi \in \mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta)$ such that for every $A \in \mathrm{Base}(\Phi)$ the sequent $A_l \implies A_r$ is refutable.

Proof. **(i)⇒(ii):** Assume that $\mathbb{S}, s$ is a pointed coalgebra model such that $\mathbb{S}, s \Vdash a$ for all $a \in \nabla\Gamma \cup P$ but $\mathbb{S}, s \nVdash b$ for all $b \in Q \cup \Delta\Theta$. Then clearly $P \cap Q = \varnothing$; for the second part, define $A : S \to P_\omega(\mathcal{B}^0(\Gamma) \cup \mathcal{B}^1(\Theta))$ as

$$t \mapsto \{f_0 a \mid a \in \mathcal{B}(\Gamma) \text{ and } \mathbb{S}, t \Vdash a\} \cup \{f_1 b \mid b \in \mathcal{B}(\Theta) \text{ and } \mathbb{S}, t \nVdash b\}$$

and set $\Phi = TA(\sigma(s))$.

First, we show that Φ is a separated redistribution of $\nabla\Gamma, P \implies Q, \Delta\Theta$.

Let $\nabla_{L_\gamma}\gamma \in \nabla\Gamma$; then $\mathbb{S}, s \Vdash \nabla_{L_\gamma}\gamma$. Note that if $a \dashv\!\vert\, t$, then $f_0 a \in A(t)$. Hence, $(f_0^{\mathrm{gr}})^\circ; \dashv\!\vert; A^{\mathrm{gr}} \subseteq \in$. So, we can calculate that

$$(Tf_0\gamma, \Phi) \in (Tf_0^{\mathrm{gr}})^\circ; (L_\gamma \Vdash)^\circ; TA^{\mathrm{gr}} \subseteq L_\gamma^{\sim}((f_0^{\mathrm{gr}})^\circ; \dashv\!\vert; A^{\mathrm{gr}}) \subseteq L_\gamma^{\sim}\in$$

as illustrated in figure 6.

$$\begin{tikzcd}
\gamma \arrow[r, no head, "L_\gamma \Vdash"] \arrow[d, "f_0^{\mathrm{gr}}"'] \arrow[dr, phantom, "\rotatebox{90}{$\subseteq$}"] & \sigma(s) \arrow[d, "TA^{\mathrm{gr}}"] \\
(Tf_0)\gamma \arrow[r, no head, "L_\gamma^{\sim}\in"'] & \Phi
\end{tikzcd}$$

Figure 6

On the other hand, if $\Delta_{L_\theta}\theta \in \Delta\Theta$, then $\mathbb{S}, s \nVdash \Delta_{L_\theta}\theta$. Note that if $b \not\dashv\!\vert\, t$, then $f_1 b \in A(t)$. So, we have

$$(Tf_1\theta, \Phi) \in (Tf_1^{\mathrm{gr}})^\circ; (L_\theta \nVdash)^\circ; TA^{\mathrm{gr}} \subseteq L_\theta^{\sim}((f_1^{\mathrm{gr}})^\circ; \not\dashv\!\vert; A^{\mathrm{gr}}) \subseteq L_\theta^{\sim}\in$$

as illustrated in figure 7.

$$\begin{tikzcd}
\theta \arrow[r, no head, "L_\theta \nVdash"] \arrow[d, "f_1^{\mathrm{gr}}"'] \arrow[dr, phantom, "\rotatebox{90}{$\subseteq$}"] & \sigma(s) \arrow[d, "TA^{\mathrm{gr}}"] \\
(Tf_1)\theta \arrow[r, no head, "L_\theta^{\sim}\in"'] & \Phi
\end{tikzcd}$$

Figure 7

Next, we show that for every $A \in \mathrm{Base}(\Phi)$, the sequent $A_l \implies A_r$ is not valid. But this is clear: since $\Phi = TA(\sigma(s))$, we know that

$$\mathrm{Base}(\Phi) \subseteq \{A(t) \mid t \in \mathrm{Base}(s)\},$$

and for every t, the sequent $A(t)_l \implies A(t)_r$ is refuted at t.

(ii)⇒(i): Assume that $P \cap Q = \varnothing$, and that there is a $\Phi \in \mathcal{R}^s(\Gamma, \Theta)$ such that for every $A \in \mathrm{Base}(\Phi)$, the sequent $A_l \implies A_r$ is refutable.

Then fix for every $A \in \mathrm{Base}(\Phi)$ a pointed coalgebra model $\mathbb{S}_A = (S_A, \sigma_A, V_A, t_A)$, such that $A_L \implies A_R$ is refuted at $\mathbb{S}_A, t_A$. Note that t defines a map $\mathrm{Base}(\Phi) \to \biguplus_{A \in \mathrm{Base}(\Phi)} S_A$.

Then define $\mathbb{S} = (S, \sigma, V)$ as

$$S := \left(\biguplus_{A \in \mathrm{Base}(\Phi)} S_A \right) \uplus \{s_0\}$$

$$\sigma(s) := \begin{cases} (T\iota_A)(\sigma_A(s)) & s \in S_A \text{ for some } A \\ Tt(\Phi) & s = s_0 \end{cases}$$

$$V(s) := \begin{cases} V_A(s) & s \in S_A \text{ for some } A \\ P & s = s_0 \end{cases}$$

We claim that $\nabla\Gamma, P \implies Q, \Delta\Theta$ is refuted at $\mathbb{S}, s_0$.

First, we clearly have that $\mathbb{S}, s_0 \Vdash p$ for all $p \in P$ and $\mathbb{S}, s_0 \nVdash q$ for all $q \in Q$.

Let $\gamma \in \Gamma$. By assumption, for all $A \in \mathrm{Base}(\Phi)$, if $f_0 a \in A$, then $\mathbb{S}_A, t_A \Vdash a$. Since the inclusion $\mathbb{S}_A \hookrightarrow \mathbb{S}$ is a coalgebra morphism, we also have that $\mathbb{S}, t_A \Vdash a$. This means that $f_0^{\mathrm{gr}}; \in; t^{\mathrm{gr}} \subseteq \dashV$. So,

$$(\sigma(s), \gamma) \in (Tt^{\mathrm{gr}})^\circ; (L_\gamma^\sim \in)^\circ; (Tf_0^{\mathrm{gr}})^\circ \subseteq L_\gamma((t^{\mathrm{gr}})^\circ; \ni; (f_0^{\mathrm{gr}})^\circ) \subseteq L_\gamma \Vdash$$

as illustrated in figure 8

$$\begin{tikzcd}
\Phi \arrow[r, o-, "L_\gamma^\sim \in"] \arrow[d, -o, "t^{\mathrm{gr}}"'] \arrow[dr, phantom, "\rotatebox{90}{$\subseteq$}"] & (Tf_0)\gamma \\
\sigma(s)\gamma \arrow[r, -o, "L_\gamma^\sim \Vdash"'] & \gamma \arrow[u, -o, "Tf_0^{\mathrm{gr}}"']
\end{tikzcd}$$

Figure 8

Similarly, if $\theta \in \Theta_L$, then

$$(\sigma(s), \theta) \in (Tt^{\mathrm{gr}})^\circ; (L^\sim \in)^\circ; (Tf_1^{\mathrm{gr}})^\circ \subseteq L((t^{\mathrm{gr}})^\circ; \ni; (f_1^{\mathrm{gr}})^\circ) \subseteq L \nVdash .$$

We conclude that $\nabla\Gamma, P \implies Q, \Delta\Theta$ is refuted at $\mathbb{S}, s_0$, showing that it is refutable.

□

Validity of the modal rule is now simply the contrapositive of (i)⇒(ii) in lemma 6.11: If for every $\Phi \in \mathcal{R}^s(\Gamma, \Theta)$ there is an $A \in \mathrm{Base}(\Phi)$ such that $A_l \implies A_r$ is valid, then point (ii) is false, so point (i) is false, meaning that $\nabla\Gamma, P \implies Q, \Delta\Theta$ is valid.

6.3 Completeness

Our proof of completeness proceeds as follows: to every sequent $A \implies B$, we associate a measure $m(A \implies B) \in \omega^2$, where we order ω^2 lexicographically. That is, we set $(n, k) < (n', k')$ if $n < n'$ or $n = n'$ and $k < k'$. We aim to define $m(A \implies B)$ in such a way that the following claim is satisfied:

Claim 6.12. *In any instance of a rule, the measure of any assumption is strictly smaller than that of the conclusion.*

Since the lexicographic order on ω^2 is a well-order, we will be able to use induction to show that every valid sequent is derivable.

Remark 6.13. Formally, claim 6.12 only applies to *productive* instances of rules. To illustrate this, consider the following application of the $\bigwedge$-l rule:

$$\bigwedge\text{-l}\,\frac{a, b, a \wedge b \implies c}{a \wedge b \implies c}$$

If we are building a proof of $a \wedge b \implies c$, the above instance of the $\bigwedge$-l rule will not be useful in finding a proof, since we have not reduced the complexity of the sequent. Such an instance can be considered *unproductive*.

To be slightly more formal: in all rules except the initial and modal rule, there is a single formula in the conclusion that is *active*. If

$$\frac{A_1 \implies B_1 \quad \dots \quad A_k \implies B_k}{A, a \implies B}$$

is an instance of a rule, where a is the active formula, we call this instance *productive* if a does not appear in any of the A_i, and similarly if the active formula appears on the right of the conclusion. For technical reasons, we consider all instances of the initial and modal rules to be productive.

In the completeness proof, we will show that every valid sequent has a derivation where all instances are productive.

Remark 6.14. Note that if claim 6.12 holds, then the proof game from remark 6.6 always ends after finitely many turns (again, provided the Prover only chooses productive instances of rules): because the measure of the sequent chosen by Refuter is strictly decreasing, there are no infinite games, since there are no infinite descending sequences in $(\omega^2, <)$.

Measure

Definition 6.15. By induction on the complexity of a formula $a \in \mathcal{L}$, we define

$$\begin{array}{llll} k_l(p) & := 0 & k_r(p) & := 0 \\ k_l(\neg a) & := 1 + k_r(a) & k_r(\neg a) & := 1 + k_l(a) \\ k_l(\odot A) & := 1 + \sum_{a \in A} k_l(a) & k_r(\odot A) & := 1 + \sum_{a \in A} k_r(a) \quad (\odot \in \{\bigwedge, \bigvee\}) \\ k_l(\nabla_L \alpha) & := 0 & k_r(\nabla_L \alpha) & := 1 \\ k_l(\Delta_L \alpha) & := 1 & k_r(\Delta_L \alpha) & := 0 \end{array}$$

We then define

$$m(A \implies B) := \Big(\max(d[A \cup B]), \sum_{a \in A} k_l(a) + \sum_{b \in B} k_r(b) \Big)$$

It is now easy to see that claim 6.12 holds. After all, in any productive instance of a boolean rule, the boolean complexity of at least one formula decreases; in an interchange rule, a 'heavy' modal is replaced with a 'weightless' one; and in the modal rule, the maximal modal depth decreases.

We are now ready to prove the completeness of $G2_T(\Lambda)$.

Proof. As a base case: assume that $A \implies B$ is a valid sequent, and $m(A \implies B) = (0, 0)$. This can only happen if A and B are both sets of proposition letters. Therefore, $A \implies B$ is only valid if $A \cap B \neq \varnothing$, meaning that it is an initial sequent, and hence provable.

Now, assume that $A \implies B$ is a valid sequent of measure $m(A \implies B) = (d, k)$, and assume that any valid sequent $A' \implies B'$ with $m(A' \implies B') < (d, k)$ is provable.

If $k \neq 0$, then $A \implies B$ is the conclusion of a boolean or interchange rule, say

$$\frac{\Pi}{A \implies B}$$

We have previously noted that these rules are invertible, so every sequent in Π is valid. Moreover, by claim 6.12, we know that every sequent in Π has measure strictly less than (d, k), so is derivable. Hence, $A \implies B$ is derivable.

We are only left with the case that $k = 0$. This can only be true if $A \implies B$ is of the form $\nabla\Gamma, P \implies Q, \Delta\Theta$, with P, Q sets of proposition letters, and

$$\nabla\Gamma = \{\nabla_{L_\gamma}\gamma \mid \gamma \in \Gamma\}, \quad \Delta\Theta = \{\Delta_{L_\theta}\theta \mid \theta \in \Theta\}.$$

Here, we know by (ii)$\Rightarrow$(i) of the refutability lemma that since $\nabla\Gamma, P \implies Q, \Delta\Theta$ is valid, either $P \cap Q \neq \varnothing$, in which case $A \implies B$ is an initial sequent and hence derivable, or for every $\Phi \in \mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta)$, there is an $A^\Phi \in \mathrm{Base}(\Phi)$ such that $A^\Phi_l \implies A^\Phi_r$ is valid. Since $m(A^\Phi_l \implies A^\Phi_r) < (d, k)$ for all Φ, the induction hypothesis tells us that each of the sequents $A^\Phi_l \implies A^\Phi_r$ is derivable, and hence by an application of the modal rule, so is $\nabla\Gamma, P \implies Q, \Delta\Theta$.

$\square$

6.4 Finitarity and decidability

We first note that not all rules in $G2_T(\Lambda)$ are necessarily finitary. In particular, the interchange rules may introduce infinite branching, since $S_L\alpha$ and $D_L\alpha$ may be infinite sets.

At first blush, it may seem as if the modal rule may require infinitely many premises as well; this is not the case, however. If

$$\frac{\{A_l^\Phi \implies A_r^\Phi\}}{\nabla\Gamma, P \implies Q, \Delta\Theta}$$

is an instance of the modal rule, then for any Φ, we know that $A_l^\Phi \subseteq \mathcal{B}(\Gamma)$ and $A_r^\Phi \subseteq \mathcal{B}(\Theta)$. Since $\mathcal{B}(\Gamma)$ and $\mathcal{B}(\Theta)$ are finite, there are only finitely many different sequents among the $A_l^\Phi \implies A_r^\Phi$, even if there may be infinitely many distinct $\Phi \in \mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta)$.

Still, if $\mathcal{R}^s(\nabla\Gamma, P \implies Q, \Delta\Theta)$ is infinite, it may be undecidable if a particular expression

$$\frac{\Pi}{A \implies B}$$

is a valid instance of the modal rule or not. So even if we restrict our attention to sequents containing only Δ-free formulas on the left and ∇-free formulas on the right, it may not be decidable if such a sequent is valid.

Hence, if we wish to ensure that $G2_T(\Lambda)$ is decidable, it seems necessary to demand that T preserve finite sets. The question now is: is this sufficient?

A common method of establishing decidability of a sequent calculus is to show that it has the *subformula property*. A sequent calculus C has this property if in a proof of a sequent $A \implies B$, only subformulas of formulas in $A \cup B$ appear.

Unfortunately, $G2_T(\Lambda)$ does not enjoy the subformula property, since in an application of the interchange rule, new modalities are introduced. Moreover, the boolean complexity of the formulas in its Base increases.

However, these increases in complexity are not too large; we can derive a weaker form of the subformula property, from which decidability follows in favourable conditions.

Definition 6.16. Let $a \in \mathcal{L}_T(\Lambda)$ be a formula. We define by simultaneous induction two sets of supporting formulas; the *left support* $\mathrm{supp}^l(a)$ and the *right support* $\mathrm{supp}^r(a)$.

$$\begin{array}{llll}
\mathrm{supp}^l(p) & := \{p\} & \mathrm{supp}^r(p) & := \{p\} \\
\mathrm{supp}^l(\neg a) & := \{\neg a\} \cup \mathrm{supp}^r(a) & \mathrm{supp}^r(\neg a) & := \{\neg a\} \cup \mathrm{supp}^l(a) \\
\mathrm{supp}^l(\bigwedge A) & := \{\bigwedge A\} \cup \bigcup_{a \in A} \mathrm{supp}^l(a) & \mathrm{supp}^r(\bigwedge A) & := \{\bigwedge A\} \cup \bigcup_{a \in A} \mathrm{supp}^r(a) \\
\mathrm{supp}^l(\bigvee A) & := \{\bigvee A\} \cup \bigcup_{a \in A} \mathrm{supp}^l(a) & \mathrm{supp}^r(\bigvee A) & := \{\bigvee A\} \cup \bigcup_{a \in A} \mathrm{supp}^r(a) \\
\mathrm{supp}^l(\nabla_L\alpha) & := \{\nabla_L\alpha\} \cup \bigcup_{a \in \mathrm{Base}(\alpha)} \mathrm{supp}^l(a) & \mathrm{supp}^r(\nabla_L\alpha) & = \{\nabla_L\alpha\} \cup \bigcup_{\beta \in D_L\alpha} \mathrm{supp}^r(\Delta_{L^\sim}\beta) \\
\mathrm{supp}^l(\Delta_L\alpha) & := \{\Delta_L\alpha\} \cup \bigcup_{\beta \in S_L\alpha} \mathrm{supp}^l(\nabla_{L^\sim}\beta) & \mathrm{supp}^r(\Delta_L\alpha) & := \{\Delta_L\alpha\} \cup \bigcup_{a \in \mathrm{Base}(\alpha)} \mathrm{supp}^r(a)
\end{array}$$

For a sequent $A \implies B$, we write

$$\Sigma(A \implies B) := \bigcup_{a \in A} \mathrm{supp}^l(a) \cup \bigcup_{b \in B} \mathrm{supp}^r(B).$$

We will call $\Sigma(A \implies B)$ the *support* of the sequent $A \implies B$.

We now prove the following lemma:

Lemma 6.17. *Let $A \implies B$ be a $\mathcal{L}_T(\Lambda)$-sequent. If π is a proof of $A \implies B$, and $A' \implies B'$ appears in π, then*

$$A', B' \subseteq \Sigma(A \implies B)$$

Proof. First, we note that for any sequent $A' \implies B'$, we have

$$A', B' \subseteq \Sigma^l(A' \implies B')$$

Next, the key observation is that if $\frac{\Pi}{A \implies B}$ is a valid instance of a proof rule, and $A' \implies B' \in \Pi$, then

$$\Sigma(A' \implies B') \subseteq \Sigma(A \implies B)$$

From this, the statement follows by a simple induction on the length of the proof.

As an illustration, we show the key observation in two specific cases.

¬-l: Assume we have some application of the ¬-l rule

$$\frac{A \implies a, B}{A, \neg a \implies B}.$$

Then

$$\begin{aligned}
\Sigma(A, \neg a \implies B) &= \bigcup_{a' \in A} \mathrm{supp}^l(a') \cup \mathrm{supp}^l(\neg a) \cup \bigcup_{b \in B} \mathrm{supp}^r(b) \\
&= \bigcup_{a' \in A} \mathrm{supp}^l(a') \cup \{\neg a\} \cup \mathrm{supp}^r(a) \cup \bigcup_{b \in B} \mathrm{supp}^r(b) \\
&\supseteq \bigcup_{a' \in A} \mathrm{supp}^l(a') \cup \mathrm{supp}^r(a) \cup \bigcup_{b \in B} \mathrm{supp}^r(B) \\
&= \Sigma(A \implies a, B)
\end{aligned}$$

showing that the support of the assumption is included in the support of the conclusion.

∇-r: Assume we have some application of the ∇-r rule

$$\frac{\{A \implies \Delta_{L^\sim}\beta, B \mid \beta \in D_L\alpha\}}{A \implies \nabla_L \alpha, B}.$$

Then for any $\beta_0 \in D_L\alpha$, we see that

$$\begin{aligned}
\Sigma(A \implies \nabla_L\alpha, B) &= \bigcup_{a\in A} \mathrm{supp}^l(a) \cup \mathrm{supp}^r(\nabla_L\alpha) \cup \bigcup_{b\in B} \mathrm{supp}^r(B) \\
&= \bigcup_{a\in A} \mathrm{supp}^l(a) \cup \bigcup_{\beta\in D_L\alpha} \mathrm{supp}^r(\Delta_{L^\sim}\beta) \cup \bigcup_{b\in B} \mathrm{supp}^r(b) \\
&\supseteq \bigcup_{a\in A} \mathrm{supp}^l(a) \cup \mathrm{supp}^r(\Delta_{L^\sim}\beta_0) \cup \bigcup_{b\in B} \mathrm{supp}^r(b) \\
&= \Sigma(A \implies \Delta_{L^\sim}\beta_0, B)
\end{aligned}$$

showing that the support of each assumption is included in the support of the conclusion.

□

We can use this lemma to prove decidability of $G2_T(\Lambda)$.

Theorem 6.18. *Let T : **Sets** $\to$ **Sets** be a computable functor that preserves finite sets, and let Λ be a set of computable T-liftings closed under $(-)^\sim$. Then the set of valid $\mathcal{L}_T(\Lambda)$-sequents is decidable.*

Proof. Let $A \implies B$ be an $\mathcal{L}_T(\Lambda)$-sequent. Let $\Sigma = \Sigma(A \implies B)$ be the support of $A \implies B$. Since T preserves finite sets, we know that Σ is finite. In the proof of completeness, we saw that any valid sequent has a proof where the measure is strictly decreases in every step. So, let $\mathcal{T}$ be the set of trees, annotated by sequents $A' \implies B'$ for $A', B' \subseteq \Sigma$, such that for any $\pi \in \mathcal{T}$, the measure of the sequents strictly decreases along the paths. Clearly, $\mathcal{T}$ is finite as well, and computable from $A \implies B$.

Now by the completeness of $G2_T(\Lambda)$, we know that $A \implies B$ is a valid sequent if and only if at least one annotated tree $\pi \in \mathcal{T}$ is a valid proof in $G2_T(\Lambda)$. And for each such annotated tree π, it is decidable if it is a valid proof in $G2_T(\Lambda)$.

For this, we note explicitly that since T preserves finite sets, we know that for any sequent $A \implies B$, the set $\mathcal{R}^s(A \implies B)$ is finite; hence it is decidable if a given instance

$$\frac{\Pi}{A \implies B}$$

is a valid instance of the modal rule. Hence, the set of valid $\mathcal{L}_T(\Lambda)$-sequents is decidable. □

6.5 Fragments

In this section, we highlight some interesting fragments of $\mathcal{L}_T(\Lambda)$ for which $G2_T(\Lambda)$ is complete. This can be given a meaning in two different ways:

Definition 6.19. Let $\mathcal{L}' \subseteq \mathcal{L}_T(\Lambda)$ be a set of formulas. We will call $\mathcal{L}'$ a *complete fragment* if whenever $A \implies B$ is provable in $G2_T(\Lambda)$, and $A, B \subseteq \mathcal{L}'$, there is a proof containing only formulas from $\mathcal{L}'$.

Let $\mathcal{L}'_l, \mathcal{L}'_r \subseteq \mathcal{L}_T(\Lambda)$ be two sets of formulas. We will call $\mathcal{S} = \{A \implies B \mid A \subseteq \mathcal{L}'_l, B \subseteq \mathcal{L}'_r\}$ a *complete sequent-fragment* if, whenever $A \implies B \in \mathcal{S}$ is a provable sequent, there is a proof containing only sequents from $\mathcal{S}$.

We note that $\mathcal{L}'$ is a complete fragment if and only if $\{A \implies B \mid A, B \subseteq \mathcal{L}'\}$ is a complete sequent-fragment.

We can give a general sufficient criterion for some subset $\mathcal{L}' \subseteq \mathcal{L}_T(\Lambda)$ to be a complete fragment. In lemma 6.17, we showed that $G2_T(\Lambda)$ has an adjusted form of the subformula property. In particular, it tells us the following:

Proposition 6.20. *Let $\mathcal{L}' \subseteq \mathcal{L}_T(\Lambda)$ be a set of formulas, and assume that for all $a \in \mathcal{L}'$, both its supports $\mathrm{supp}^l(a), \mathrm{supp}^r(a)$ are subsets of $\mathcal{L}'$. Then $\mathcal{L}'$ is a complete fragment.*

Let $\mathcal{L}'_l, \mathcal{L}'_r \subseteq \mathcal{L}_T(\Lambda)$ be sets of formulas, and assume that for all $A \subseteq \mathcal{L}'_l, B \subseteq \mathcal{L}'_r$, if $\Sigma(A' \implies B') \subseteq \Sigma(A \implies B)$, then $A \subseteq \mathcal{L}'_l, B \subseteq \mathcal{L}'_r$. Then

$$\{A \implies B \mid A \subseteq \mathcal{L}'_l, B \subseteq \mathcal{L}'_r\}$$

is a complete sequent-fragment.

Unidirectional fragments Consider a T-lifting L. We have previously seen that the $\mathcal{L}_T(L, L^{\sim})$-formulas preserved under L-simulation are those equivalent to a clean formula only containing $\nabla_{L^{\sim}}$ and Δ_L. Let $\mathcal{L}^c_T(\nabla_{L^{\sim}}, \Delta_L)$ be the set of these formulas.

By a straightforward induction, it can be seen that if $a \in \mathcal{L}^c_T(\nabla_{L^{\sim}}, \Delta_L)$, then $\mathrm{supp}^l(a), \mathrm{supp}^r(a) \subseteq \mathcal{L}^c_T(\nabla_{L^{\sim}}, \Delta_L)$. So by proposition 6.20, we know that $\mathcal{L}^c_T(\nabla_{L^{\sim}}, \Delta_L)$ is a complete fragment.

Classical modal logic We can construct a sequent calculus for the classical modal logic K based on $G2_P(\overrightarrow{P}, \overleftarrow{P})$. After all, we have equivalences

$$\overrightarrow{\nabla}\{p\} \equiv \Box p \equiv \overleftarrow{\Delta}\{p\}, \qquad \overrightarrow{\nabla}\{p\} \equiv \Diamond p \equiv \overleftarrow{\Delta}\{p\}.$$

So, a $\mathcal{L}_{\Box,\Diamond}$-sequent can be seen as a $\mathcal{L}_P(\overrightarrow{P}, \overleftarrow{P})$-sequent, where all modalities are applied to a singleton. Moreover, we can do this in such a way that only ∇'s occur on the left, and only Δ's occur on the right.[3]

Formally, we can define by simultaneous induction two translations $\varphi \mapsto \varphi^l$ and $\varphi \mapsto \varphi^r$:

$$\begin{array}{lll lll}
p^l & := & p & p^r & := & p \\
(\neg\varphi)^l & := & \neg\varphi^r & (\neg\varphi)^r & := & \neg\varphi^l \\
(\varphi \wedge \psi)^l & := & \bigwedge\{\varphi^l, \psi^l\} & (\varphi \wedge \psi)^r & := & \bigwedge\{\varphi^r, \psi^r\} \\
(\varphi \vee \psi)^l & := & \bigvee\{\varphi^l, \psi^l\} & (\varphi \vee \psi)^r & := & \bigvee\{\varphi^r, \psi^r\} \\
(\Box\varphi)^l & := & \overrightarrow{\nabla}\{\varphi^l\} & (\Box\varphi)^r & := & \overleftarrow{\Delta}\{\varphi^r\} \\
(\Diamond\varphi)^l & := & \overleftarrow{\nabla}\{\varphi^l\} & (\Diamond\psi)^r & := & \overrightarrow{\Delta}\{\varphi^r\}
\end{array}$$

[3]this is to eliminate the interchange rules, which don't have a productive counterpart on the $\mathcal{L}_{\Box,\Diamond}$-side.

Then for sets A and B of modal formulas, the sequent $A \implies B$ is valid if and only if the sequent $A^l \implies B^r$ is valid, where $A^l = \{\varphi^l \mid \varphi \in A\}, B^r = \{\psi^r \mid \psi \in B\}$.

We can now give a proof system for K. We define a sequent calculus GK consisting of the following rules:

$$\text{init}\,\frac{}{A, p \implies p, B} \quad \wedge\text{-l}\,\frac{A, \varphi, \psi \implies B}{A, \varphi \wedge \psi \implies B} \quad \wedge\text{-r}\,\frac{A \implies \varphi, B \quad A \implies \psi, B}{A \implies \varphi \wedge \psi, B}$$

$$\vee\text{-l}\,\frac{A, \varphi \implies B \quad A, \psi \implies B}{A, \varphi \vee \psi \implies B} \quad \vee\text{-r}\,\frac{A \implies \varphi, \psi, B}{A \implies \varphi \vee \psi, B}$$

$$\neg\text{-l}\,\frac{A \implies \varphi, B}{A, \neg\varphi \implies B} \quad \neg\text{-r}\,\frac{A, \varphi \implies B}{A \implies \neg\varphi, B}$$

$$\text{modal}\,\frac{\{A_l^\Phi \implies A_r^\Phi\}}{\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2}\,\forall\Phi. A^\Phi \in \Phi$$

The modal rule should be read similarly to the $T(\nabla\Delta)$-rule in $G2_T(\Lambda)$: Call $\Phi \in P((\Gamma_1 \cup \Gamma_2) \uplus (\Theta_1 \cup \Theta_2))$ a separated redistribution if

- for all $U \in \Phi$, $\Gamma_1 \uplus \Theta_1 \subseteq U$, and
- for all $(x, i) \in \Gamma_2 \uplus \Theta_2$, there is a $U \in \Phi$ with $(x, i) \in U$,

If for all separated redistributions Φ there is an $A^\Phi \in \Phi$ such that $A_l \implies A_r$ is derivable, then $\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2$ is derivable.

This defines a sound and complete sequent calculus for K. To prove this, we argue that $\mathsf{GK} \vdash A \implies B$ if and only if $G2_P(\overrightarrow{P}, \overleftarrow{P}) \vdash A^l \implies B^r$. The crucial point is that

$$\text{R}\,\frac{A_1 \implies B_1 \quad \dots \quad A_k \implies B_k}{A \implies B}$$

is an instance of a proof rule in GK if and only if

$$\text{R}^{l,r}\,\frac{A_1^l \implies B_1^r \quad \dots \quad A_k^l \implies B_k^r}{A^l \implies B^r}$$

is an instance of a proof rule in $G2_P(\overrightarrow{P}, \overleftarrow{P})$. Moreover, $\{A^l \implies B^r \mid A, B \subseteq \mathcal{L}_{\Box,\Diamond}\}$ is a complete sequent-fragment, since it satisfies the criterion from 6.20.

So, any proof of $A \implies B$ in GK can be turned into a proof of $A^l \implies B^r$ in $G2_P(\overrightarrow{P}, \overleftarrow{P})$ by translating; and any proof of $A^l \implies B^r$ in $G2_P(\overrightarrow{P}, \overleftarrow{P})$ is equal to the translation of some proof of $A \implies B$ in GK.

We can be more explicit about the modal rule if we examine the separated redistributions. Any sequent $\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2$ has a 'canonical' redistribution

$$\Phi_0 := \{(\Gamma_1 \uplus \Theta_1) \cup \{(\gamma, 0)\} \mid \gamma \in \Gamma_2\} \cup \{(\Gamma_1 \uplus \Theta_1) \cup \{(\theta, 1)\} \mid \theta \in \Theta_2\}.$$

So, if

$$\text{modal}\,\frac{\Pi}{\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2}$$

is a valid instance of the modal rule, there is some sequent $A_l^{\Phi_0} \implies A_r^{\Phi_0} \in \Pi$ for $A^{\Phi_0} \in \Phi$.

Moreover, for any separated redistribution Φ, every $A \in \Phi_0$ is contained in some element $A' \in \Phi$. So, if there is some $A \in \Phi_0$ with $A_l \implies A_r$ derivable, then by weakening, *every* separated redistribution Φ contains an A' with $A'_l \implies A'_r$ derivable. So, we can replace our modal rule with the following two rules:

$$\Diamond\text{-l}\,\frac{\Gamma_1, \gamma \implies \Theta_1}{\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2}\,\gamma \in \Gamma_2$$

$$\Box\text{-r}\,\frac{\Gamma_1 \implies \theta, \Theta_1}{\Box\Gamma_1, \Diamond\Gamma_2, P \implies Q, \Diamond\Theta_1, \Box\Theta_2}\,\theta \in \Theta_2$$

Monotone modal logic We can perform the same process for monotone modal logic: we define

$$\begin{array}{llllll}
p^l & := & p & p^r & := & p \\
(\neg\varphi)^l & := & \neg\varphi^r & (\neg\varphi)^r & := & \neg\varphi^l \\
(\varphi \wedge \psi)^l & := & \bigwedge\{\varphi^l, \psi^l\} & (\varphi \wedge \psi)^r & := & \bigwedge\{\varphi^r, \psi^r\} \\
(\varphi \vee \psi)^l & := & \bigvee\{\varphi^l, \psi^l\} & (\varphi \vee \psi)^r & := & \bigvee\{\varphi^r, \psi^r\} \\
(\Box\varphi)^l & := & \overset{\leftsquigarrow}{\nabla}\langle\{\varphi^l\}\rangle & (\Box\varphi)^r & := & \overset{\rightsquigarrow}{\Delta}\langle\{\varphi^r\}\rangle \\
(\Diamond\varphi)^l & := & \overset{\rightsquigarrow}{\nabla}\langle\{\varphi^l\}\rangle & (\Diamond\psi)^r & := & \overset{\leftsquigarrow}{\Delta}\langle\{\varphi^r\}\rangle
\end{array}$$

and write down rules for M such that

$$\text{R}\,\frac{A_1 \implies B_1 \quad \dots \quad A_k \implies B_k}{A \implies B}$$

is an instance of a proof rule in GK if and only if

$$\text{R}^{l,r}\,\frac{A_1^l \implies B_1^r \quad \dots \quad A_k^l \implies B_k^r}{A^l \implies B^r}$$

is an instance of a proof rule in $G2_{\mathcal{M}}(\overset{\rightsquigarrow}{\mathcal{M}}, \overset{\leftsquigarrow}{\mathcal{M}})$.

In order to refine the modal rule, we investigate $\mathcal{R}^s(\nabla\Gamma^l \implies \Delta\Theta^r)$.

Let $\Gamma_1, \Gamma_2, \Theta_1, \Theta_2 \in \mathcal{ML}_{\mathcal{M}}(\overset{\rightsquigarrow}{\mathcal{M}}, \overset{\leftsquigarrow}{\mathcal{M}})$, such that

$$\Gamma_1 = \{\langle\{a_1^1\}\rangle, \dots, \langle\{a_{k_1}^1\}\rangle\}$$
$$\Gamma_2 = \{\langle\{a_1^2\}\rangle, \dots, \langle\{a_{k_2}^2\}\rangle\}$$
$$\Theta_1 = \{\langle\{b_1^1\}\rangle, \dots, \langle\{b_{m_1}^1\}\rangle\}$$
$$\Theta_2 = \{\langle\{b_1^2\}\rangle, \dots, \langle\{b_{m_2}^2\}\rangle\}$$

Let $B = \mathcal{B}(\Gamma_1 \cup \Gamma_2) \uplus \mathcal{B}(\Theta_1 \cup \Theta_2)$. A given $\Phi \in \mathcal{M}PB$ is a separated redistribution of $\overset{\rightsquigarrow}{\nabla}\Gamma_1, \overset{\leftsquigarrow}{\nabla}\Gamma_2 \implies \overset{\leftsquigarrow}{\Delta}\Theta_1, \overset{\rightsquigarrow}{\Delta}\Theta_2$ if and only if

(i) For $i = 1, \ldots, k_1$, for all $\mathcal{A} \in \Phi$ there is an $A \in \mathcal{A}$ with $(a_i^1, 0) \in A$.

(ii) For $i = 1, \ldots, k_2$, there is a $\mathcal{A}_i \in \Phi$ such that for all $A \in \mathcal{A}_i$, $(a_i^2, 0) \in A$.

(iii) For $i = 1, \ldots, m_1$, there is a $\mathcal{B}_i \in \Phi$ such that for all $B \in \mathcal{B}_i$, $(b_i^1, 1) \in B$.

(iv) For $i = 1, \ldots, m_2$, for all $\mathcal{B} \in \Phi$ there is a $B \in \mathcal{B}$ with $(b_i^2, 1) \in B$.

Again, we can consider the 'canonical' redistribution Φ_0, defined as follows: For $i = 1, \ldots, k_2$ and $j = 1, \ldots, m_1$, we set

$$\begin{aligned}
\mathcal{A}_i &:= \{\{(a_{i'}^1, 0), (a_i^2, 0)\} \mid i' = 1, \ldots, k_1\} \cup \{\{(b_{j'}^2, 1), (a_i^2, 0)\} \mid j' = 1, \ldots, m_2\} \\
\mathcal{B}_j &:= \{\{(a_{i'}^1, 0), (b_j^1, 1)\} \mid i' = 1, \ldots, k_1\} \cup \{\{(b_{j'}^2, 1), (b_j^1, 1)\} \mid j' = 1, \ldots, m_2\}
\end{aligned}$$

Then we define

$$\Phi_0 := \langle \mathcal{A}_1, \ldots, \mathcal{A}_{k_2}, \mathcal{B}_1, \ldots, \mathcal{B}_{m_1} \rangle$$

It is easy to check that Φ_0 satisfies conditions (i) - (iv), and hence is a separated redistribution. We also see that

$$\mathrm{Base}(\Phi_0) = \bigcup_{i=1}^{k_2} \mathcal{A}_i \cup \bigcup_{j=1}^{m_1} \mathcal{B}_j$$

So if $\overset{\rightsquigarrow}{\nabla}\Gamma_1, \overset{\leftsquigarrow}{\nabla}\Gamma_2 \implies \overset{\leftsquigarrow}{\Delta}\Theta_1, \overset{\rightsquigarrow}{\Delta}\Theta_2$ is valid, then one of the following four cases occurs:

1. There are $a \in \Gamma_1, a' \in \Gamma_2$ such that $a, a' \implies$ is valid.

2. There are $a \in \Gamma_1, b \in \Theta_1$ such that $a \implies b$ is valid.

3. There are $a \in \Gamma_2, b \in \Theta_2$ such that $a \implies b$ is valid.

4. There are $b \in \Theta_1, b' \in \Theta_2$ such that $\implies b, b'$ is valid.

We claim that vice versa, if any of these four cases occurs, then every separated redistribution Φ of $\overset{\rightsquigarrow}{\nabla}\Gamma_1, \overset{\leftsquigarrow}{\nabla}\Gamma_2 \implies \overset{\leftsquigarrow}{\Delta}\Theta_1, \overset{\rightsquigarrow}{\Delta}\Theta_2$ has an element $A^\Phi \in \mathrm{Base}(\Phi)$ with $A_l^\Phi \implies A_r^\Phi$ derivable.

To illustrate this, we consider case 1. Let $a \in \Gamma_1, a' \in \Gamma_2$ with $a, a' \implies$ valid. If $\Phi \in \mathcal{R}^s(\overset{\rightsquigarrow}{\nabla}\Gamma_1, \overset{\leftsquigarrow}{\nabla}\Gamma_2 \implies \overset{\leftsquigarrow}{\Delta}\Theta_1, \overset{\rightsquigarrow}{\Delta}\Theta_2)$, then since it satisfies condition (ii), there is a $\mathcal{A} \in \Phi$ such that all $A \in \mathcal{A}$ contain $(a', 0)$. Now let $\mathcal{A}' = \mathcal{A} \cap \mathrm{Base}(\Phi)$. By the criterion from remark 2.10, we know that $\mathcal{A}' \in \Phi$. Since Φ satisfies condition (i), there is an $A^\Phi \in \mathcal{A}'$ with $(a, 0) \in A^\Phi$. Now $A^\Phi \in \mathrm{Base}(\Phi)$, and $(a, 0), (a', 0) \in A^\Phi$. So, we have that $A_l^\Phi \implies A_r^\Phi$ is of the form

$$A, a, a' \implies B$$

for some sets A, B. Since $a, a' \implies$ was derivable, we know by weakening that $A_l^\Phi \implies A_r^\Phi$ is derivable.

This motivates the following sequent calculus GM for M:

$$\text{init}\,\frac{}{A, p \implies p, B} \quad \wedge\text{-l}\,\frac{A, \varphi, \psi \implies B}{A, \varphi \wedge \psi \implies B} \quad \wedge\text{-r}\,\frac{A \implies \varphi, B \quad A \implies \psi, B}{A \implies \varphi \wedge \psi, B}$$

$$\vee\text{-l}\,\frac{A, \varphi \implies B \quad A, \psi \implies B}{A, \varphi \vee \psi \implies B} \quad \vee\text{-r}\,\frac{A \implies \varphi, \psi, B}{A \implies \varphi \vee \psi, B}$$

$$\neg\text{-l}\,\frac{A \implies \varphi, B}{A, \neg\varphi \implies B} \quad \neg\text{-r}\,\frac{A, \varphi \implies B}{A \implies \neg\varphi, B}$$

$$\Diamond\text{-l},\Box\text{-l}\,\frac{\gamma, \gamma' \implies}{\Diamond\Gamma_1, \Box\Gamma_2, P \implies Q, \Box\Theta_1, \Diamond\Theta_2}\,\gamma \in \Gamma_1, \gamma' \in \Gamma_2$$

$$\Diamond\text{-l},\Box\text{-r}\,\frac{\gamma \implies \theta}{\Diamond\Gamma_1, \Box\Gamma_2, P \implies Q, \Box\Theta_1, \Diamond\Theta_2}\,\gamma \in \Gamma_1, \theta \in \Theta_1$$

$$\Box\text{-l},\Diamond\text{-r}\,\frac{\gamma \implies \theta}{\Diamond\Gamma_1, \Box\Gamma_2, P \implies Q, \Box\Theta_1, \Diamond\Theta_2}\,\gamma \in \Gamma_2, \theta \in \Theta_2$$

$$\Box\text{-r},\Diamond\text{-r}\,\frac{\implies \theta, \theta'}{\Diamond\Gamma_1, \Box\Gamma_2, P \implies Q, \Box\Theta_1, \Diamond\Theta_2}\,\theta \in \Theta_1, \theta' \in \Theta_2$$

By the discussion above, we see that for any $\mathcal{L}_{\Box,\Diamond}$-sequent $A \implies B$, $\mathsf{GM} \vdash A \implies B$ if and only if $G2_{\mathcal{M}}(\widetilde{\vec{\mathcal{M}}}, \widetilde{\overleftarrow{\mathcal{M}}}) \vdash A^l \implies B^r$. Since $G2_T(\widetilde{\vec{\mathcal{M}}}, \widetilde{\overleftarrow{\mathcal{M}}})$ was complete, we know that GM is complete for M.

7 Conclusion

7.1 Summary

In this thesis, we have studied relation lifting as a general approach to coalgebraic logic. We have defined a family of logics $\mathcal{L}_T(\Lambda)$ based on Moss' ∇-modality, and showed that on finite-branching coalgebras, logical equivalence coincides with the natural associated notion of bisimulation. It is also shown that for functors preserving weak pullbacks and finite sets, the $L^{\sim}$-simulation-invariant formulas are (up to equivalence) those defined from ∇_L.

We have further given a family of sound and complete cut-free sequent calculi $G2_T(\Lambda)$, and demonstrated how these calculi may be modified to obtain proof systems for specific coalgebraic logics of interest.

7.2 Further Research

One of the places the ∇-modality arises naturally is in the study of the modal μ-calculus [14] [8]. Indeed, coalgebras are the natural models for fixpoint logics in general. Hence, there may be a universal method to add fixpoint operators to $\mathcal{L}_T(\Lambda)$. The proof theory of the μ-calculus is notoriously complex; taking the calculi $G2_T(\Lambda)$ as a starting point may give a new perspective.

Another line of research would be to compare relation liftings to predicate liftings. In [17], it is shown that for weak pullback-preserving functors, the logic $\mathcal{L}_T(\bar{T})$ can be translated into a language based on predicate liftings. A natural question is: can this translation be extended beyond weak pullback-preserving functors? Are predicate lifting-based logics and relation lifting-based logics always equally expressive?

Here, we also mention research by Gorín and Schröder on simulations based on predicate liftings [10], which in some ways mirrors the work in this thesis.

We further note that in general, the structure of the class of T-liftings is not well understood. In this thesis, it is shown that there is always a minimal element; however, as of yet, no explicit construction is known. For the neighborhood functor $\mathcal{N}$, no non-trivial lifting is known. In [12], several notions of equivalence for the neighborhood functor are given; it may be possible to adapt one of these into a $\mathcal{N}$-lifting.

Of particular interest is the question of which functors admit a diagonal-preserving functor; since these are the functors for which behavioral equivalence is captured by L-bisimulation for some L.

Finally, in the upward expressivity theorem 5.6, we made essential use of pure normal forms. The existence of a pure normal form is guaranteed if $\overline{T}$ preserves weak pullbacks. But there are other functors that admit pure normal forms - notably, the monotone neighborhood functor. This leads to the question: can the reliance on weak pullback-preservation be eliminated from the proof of 5.6?

References

[1] Alexandru Baltag. A logic for coalgebraic simulation. *Electronic Notes in Theoretical Computer Science*, 33:42–60, 2000. CMCS'2000, Coalgebraic Methods in Computer Science.

[2] Michael Barr. Relational algebras. In S. MacLane, H. Applegate, M. Barr, B. Day, E. Dubuc, Phreilambud, A. Pultr, R. Street, M. Tierney, and S. Swierczkowski, editors, *Reports of the Midwest Category Seminar IV*, pages 39–55, Berlin, Heidelberg, 1970. Springer Berlin Heidelberg.

[3] F. Bartels. *On generalised coinduction and probabilistic specification formats: Distributive laws in coalgebraic modelling.* PhD thesis, Vrije Universiteit Amsterdam, 2004.

[4] Jon Beck. Distributive laws. In B. Eckmann, editor, *Seminar on Triples and Categorical Homology Theory*, pages 119–140, Berlin, Heidelberg, 1969. Springer Berlin Heidelberg.

[5] Marta Bilkova, Alessandra Palmigiano, and Yde Venema. Proof systems for Moss' coalgebraic logic. *Theoretical Computer Science*, 549:36–60, 09 2014.

[6] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Models*, page 50–122. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

[7] F. Bonchi, Daniela Petrisan, D. Pous, and J. Rot. A general account of coinduction up-to. *Acta Informatica*, 54:127–190, 2016.

[8] Sebastian Enqvist, Fatemeh Seifan, and Yde Venema. Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics. *Theoretical Computer Science*, 727:37–100, 2018.

[9] Kit Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2):229 – 237, 1975.

[10] Daniel Gorín and Lutz Schröder. Simulations and bisimulations for coalgebraic modal logics. In Reiko Heckel and Stefan Milius, editors, *Algebra and Coalgebra in Computer Science*, pages 253–266, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[11] Helle Hansen and Clemens Kupke. A coalgebraic perspective on monotone modal logic. *Electronic Notes in Theoretical Computer Science*, 106:121–143, 12 2004.

[12] H.H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5(2):1–38, April 2009.

[13] Jesse Hughes and Bart Jacobs. Simulations in coalgebra. *Theoretical Computer Science*, 327(1):71–108, 2004. Selected Papers of CMCS '03.

[14] David Janin and Igor Walukiewicz. Automata for the modal μ-calculus and related results. In Jiří Wiedermann and Petr Hájek, editors, *Mathematical Foundations of Computer Science 1995*, pages 552–562, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.

[15] Clemens Kupke, Alexander Kurz, and Yde Venema. Completeness for the coalgebraic cover modality. *Logical Methods in Computer Science*, 8, 06 2012.

[16] Alexander Kurz and Jiří Velebil. Relation lifting, a survey. *Journal of Logical and Algebraic Methods in Programming*, 85(4):475–499, 2016. Relational and algebraic methods in computer science.

[17] Raul Andres Leal. Predicate liftings versus nabla modalities. *Electronic Notes in Theoretical Computer Science*, 203(5):195–220, 2008. Proceedings of the Ninth Workshop on Coalgebraic Methods in Computer Science (CMCS 2008).

[18] Johannes Marti and Yde Venema. Lax extensions of coalgebra functors and their logic. *Journal of Computer and System Sciences*, 81(5):880–900, 2015. 11th International Workshop on Coalgebraic Methods in Computer Science, CMCS 2012 (Selected Papers).

[19] Lawrence S. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96(1):277–317, 1999.

[20] Jan Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 10 2000.

[21] L. Santocanale and Y. Venema. Uniform interpolation for monotone modal logic. In *Advances in Modal Logic*, 2010.

[22] Albert Marchienus Thijs. *Simulation and fixpoint semantics*. Rijksuniversiteit Groningen, 1996.

A A colimit construction for $\mathcal{L}_T(\Lambda)$

Let us recap the problem with the given construction of $\mathcal{L}_T(\Lambda)$. In the definition of formulas of the form $\nabla\alpha$, we take $\alpha \in T_\omega\mathcal{L}_T(\Lambda)$. The problem is that it is not always clear what an object in $T_\omega\mathcal{L}_T(\Lambda)$ looks like, without *first* having $\mathcal{L}_T(\Lambda)$ available.

One solution would be to make the inductive nature of the construction explicit. We could define $\mathcal{L}_0 := \mathsf{Prop}$, and

$$\mathcal{L}_{i+1} := \mathcal{L}_i \cup \{\neg a \mid a \in \mathcal{L}_i\} \cup \{\bigwedge A, \bigvee A \mid A \in P_\omega\mathcal{L}_i\} \cup \{\nabla_L\alpha, \Delta_L\alpha \mid L \in \Lambda, \alpha \in T_\omega\mathcal{L}_i\}$$

and set $\mathcal{L}_T(\Lambda) := \bigcup_{i=1}^{\infty} \mathcal{L}_i$.

This will however lead to an overdefinition of many formulas. Consider, for instance, the monotone neighborhood functor $\mathcal{M}$. At stage i, we may define some modal formula $\nabla\langle A_1, \ldots, A_k\rangle_i$. At the next stage, we can define this formula again as $\nabla\langle A_1, \ldots, A_k\rangle_{i+1}$. Here, we denote by $\langle A_1, \ldots, A_k\rangle_i$ the upwards closure of $\{A_1, \ldots, A_k\}$ *in* $\mathcal{L}_i$.

We would ideally want to identify these formulas. This identification is not completely straightforward, as their identification will also have consequences for formulas where they occur as subformulas. To coordinate the identifications, we perform a suitable colimit construction.

Define the functor $\mathcal{F} : \mathbf{Sets} \to \mathbf{Sets}$ as

$$\mathcal{F}X := X \times \{\neg\} + P_\omega X \times \{\bigwedge, \bigvee\} + \sum_{L \in \Lambda} T_\omega X \times \{\nabla_L, \Delta_L\}$$

and set $L_0 := \mathsf{Prop}$. We define L_1 to be the coproduct of L_0 and $\mathcal{F}L_0$. This yields the diagram

$$\begin{tikzcd} L_0 \arrow[r, "f_0"] & L_1 \\ \mathcal{F}L_0 \arrow[ur, "g_0"'] & \end{tikzcd}$$

Now apply $\mathcal{F}$ to f_0 to get the diagram

$$\begin{tikzcd} L_0 \arrow[r, "f_0"] & L_1 \\ \mathcal{F}L_0 \arrow[ur, "g_0"'] \arrow[r, "\mathcal{F}f_0"'] & \mathcal{F}L_1 \end{tikzcd}$$

We define L_2 as the pushout of g_0 and $\mathcal{F}f_0$. This gives us

$$\begin{tikzcd} L_0 \arrow[r, "f_0"] & L_1 \arrow[r, "f_1"] & L_2 \\ \mathcal{F}L_0 \arrow[ur, "g_0"'] \arrow[r, "\mathcal{F}f_0"'] & \mathcal{F}L_1 \arrow[ur, "g_1"'] & \end{tikzcd}$$

Now we set L_3 to be the pushout of g_1 and $\mathcal{F}f_1$, yielding

$$\begin{tikzcd} L_0 \arrow[r, "f_0"] & L_1 \arrow[r, "f_1"] & L_2 \arrow[r, "f_2"] & L_3 \\ \mathcal{F}L_0 \arrow[r, "\mathcal{F}f_0"'] \arrow[ur, "g_0"'] & \mathcal{F}L_1 \arrow[r, "\mathcal{F}f_1"'] \arrow[ur, "g_1"'] & \mathcal{F}L_2 \arrow[ur, "g_2"'] & \end{tikzcd}$$

Continuing like this, we obtain an infinite diagram

$$\begin{tikzcd} L_0 \arrow[r, "f_0"] & L_1 \arrow[r, "f_1"] & L_2 \arrow[r, "f_2"] & L_3 \arrow[r, "f_3"] & L_4 \arrow[r, "f_4"] & \dots \\ \mathcal{F}L_0 \arrow[r, "\mathcal{F}f_0"'] \arrow[ur, "g_0"'] & \mathcal{F}L_1 \arrow[r, "\mathcal{F}f_1"'] \arrow[ur, "g_1"'] & \mathcal{F}L_2 \arrow[r, "\mathcal{F}f_2"'] \arrow[ur, "g_2"'] & \mathcal{F}L_3 \arrow[r, "\mathcal{F}f_3"'] \arrow[ur, "g_3"'] & \mathcal{F}L_4 \arrow[r, "\mathcal{F}f_4"'] \arrow[ur, "g_4"'] & \dots \end{tikzcd}$$

We define $\mathcal{L}_T(\Lambda)$ to be the colimit of the above diagram. Note that this is simply the colimit of

$$L_0 \xrightarrow{f_0} L_1 \xrightarrow{f_1} L_2 \xrightarrow{f_2} \dots$$

What makes this the right construction? Consider $\mathcal{F}\mathcal{L}_T(\Lambda)$. $\mathcal{F}$ is a sum of finitary functors, and is hence itself finitary. It is a fortunate fact from category theory that a functor is finitary if and only if it commutes with directed colimits. Since $\mathcal{L}_T(\Lambda)$ is the limit of a directed diagram, we know that $\mathcal{F}\mathcal{L}_T(\Lambda)$ is the colimit of

$$\mathcal{F}L_0 \xrightarrow{\mathcal{F}f_0} \mathcal{F}L_1 \xrightarrow{\mathcal{F}f_1} \mathcal{F}L_2 \xrightarrow{\mathcal{F}f_2} \dots$$

By construction, we have a collection of maps $g_i : \mathcal{F}L_i \to L_{i+1}$, from which we get a map between the colimits $g : \mathcal{F}\mathcal{L}_T(\Lambda) \to \mathcal{L}_T(\Lambda)$.

We can read g as the following map:

$$\begin{aligned} (a, \neg) &\mapsto \neg a \\ (A, \bigwedge) &\mapsto \bigwedge A \\ (A, \bigvee) &\mapsto \bigvee A \\ (\alpha, \nabla_L) &\mapsto \nabla_L \alpha \\ (\beta, \Delta_L) &\mapsto \Delta_L \alpha \end{aligned}$$

There are a few properties that we expect g to have.

(1) g is injective, and $\mathcal{L}_T(\Lambda)$ is the disjoint union of Prop and the image of g.

(2) Set $a \prec b$ if $a \in \mathrm{Base}(\beta)$ with $b = g(\beta)$; then the transitive closure of $\prec$ is a well-founded relation.

Each of these properties is easily verified. The result is the following:

(1) Every formula $a \in \mathcal{L}_T(\Lambda)$ is of exactly one of the following forms:

- $a = p$ for a unique $p \in \mathsf{Prop}$;
- $a = \neg b$ for a unique $b \in \mathcal{L}_T(\Lambda)$;
- $a = \bigwedge A$ for a unique $A \in P_\omega \mathcal{L}_T(\Lambda)$;
- $a = \bigvee A$ for a unique $A \in P_\omega \mathcal{L}_T(\Lambda)$;
- $a = \nabla_L \alpha$ for unique $L \in \Lambda, \alpha \in T_\omega \mathcal{L}_T(\Lambda)$;
- $a = \Delta_L \alpha$ for unique $L \in \Lambda, \alpha \in T_\omega \mathcal{L}_T(\Lambda)$.

(2) If $\varphi(p)$ is true for every proposition letter p, and the implication

$$(\forall a \prec b : \varphi(a)) \implies \varphi(b)$$

holds for all formulas b, then φ holds for all formulas.

Hence, by point (2) we can perform induction on the complexity of a formula, and by point (1) we can unambiguously distinguish the cases.

This procedure may seem overly complicated to obtain something as simple as 'the set of formulas'. Yet there is a hidden merit to this construction.

For, note that both the syntax of the modalities, as well as the semantics of $\Vdash$, rely on there being an 'object of formulas' *inside the category*. This poses a problem when we move beyond the category of sets. For instance, what if we were to attempt relation lifting on the category of Stone spaces? If we simply define our 'set of formulas' by induction, we run into problems, since (a) a modal formula of the form $\nabla\alpha$ takes an $\alpha \in T\mathcal{L}_T(\Lambda)$, requiring $\mathcal{L}_T(\Lambda)$ to be something that T can be applied to, and (b) the semantics of $\Vdash$ require it to be a relation between the coalgebra and $\mathcal{L}_T(\Lambda)$, meaning that $\mathcal{L}_T(\Lambda)$ needs to live in the same category as the (carriers of the) coalgebras.

Both these problems can be solved if we view $\mathcal{L}_T(\Lambda)$ not as a collection of formulas on a meta-level, but as a colimit computed inside the category of interest.