Finitary coalgebraic logics

Clemens Kupke

Finitary coalgebraic logics

ILLC Dissertation Series DS-2006-03

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
phone: +31-20-525 6051
fax: +31-20-525 5206
e-mail: illc@science.uva.nl
homepage: http://www.illc.uva.nl/

Finitary coalgebraic logics

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof.mr. P.F. van der Heijden
ten overstaan van een door het college voor
promoties ingestelde commissie, in het openbaar
te verdedigen in de Aula der Universiteit
op donderdag 23 maart 2006, te 12.00 uur

door

Clemens Achim Kupke

geboren te Amberg, Duitsland.

Promotor: Prof.dr. J.J.M.M. Rutten

Co-promotores: Dr. A. Kurz
Dr. Y. Venema

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

The investigations were supported by the Netherlands Organization for Scientific Research (NWO) in the context of project CoMoLo (nr. 612.069.002).

Printed and bound by PrintPartners Ipskamp, Enschede.

ISBN: 90–5776–148–3

Für Mama und Papa

Acknowledgments

This thesis would have never been finished without the help and support of many people.

First of all let me thank my advisors Alexander Kurz, Jan Rutten and Yde Venema who complemented one another in a perfect way. Due to the fact that my research focused very much on the logical aspects of my project Yde became my main advisor. I learned a lot from him - his outstanding qualities as a researcher and teacher are well-known, in particular his ability to put himself in the position of the student. He knows exactly when it is necessary to put some pressure on a student, but luckily also, when the student needs encouragement and support. This was especially helpful for making sure that I finished my thesis in a reasonable time. In this context I also want to express my deep gratitude to Yde for the careful and critical reading of various drafts of my thesis. This deserves special acknowledgment as he had to proof-read two PhD thesis's at the same time - a fact which in the last stage of writing often had the result that he had to sacrifice large parts of his weekends.

Having said that Yde was my main advisor I immediately want to stress the fact that Jan's role was much more important than being just a formal promotor. In fact, at the beginning of my project, Yde's and Jan's role were very similar. In regular meetings I benefited from Jan's deep understanding of the theory of coalgebras and his profound knowledge of the literature. In addition to that Jan gave me support and advice whenever I needed it. He closely followed my first steps into the scientific world and never lost his interest in my research, despite the fact that I decided to work on topics that were not directly related to his own research.

Alexander, my second copromotor, worked as a post-doc at the CWI when I started my PhD. It is mainly thanks to him that I nowadays understand the basics of category theory. His enthusiasm for learning and understanding things is exceptional. Each of the many conversations and fruitful discussions we had were very motivating and inspiring for me. The fact that Alexander moved from Amsterdam to Leicester was definitely a big loss both for our group at the CWI and of course for me. Luckily we managed to keep up our good relation through several visits and I hope we will continue our cooperation.

I also want to thank my advisors and the members of my reading committee Johan van Benthem, Peter van Emde Boas, Bart Jacobs and Larry Moss for their willingness to be in my committee and for giving valuable comments which helped to improve the thesis. Furthermore I want to thank Merlijn Sevenster who helped me with translating my English “Summary” into the required Dutch “Samenvatting”.

Special thanks go to my coauthors with whom I enjoyed working very much: Helle Hvid Hansen, Alexander Kurz, Dirk Pattinson and Yde Venema. This thesis is based on joint papers with Alexander, Dirk and Yde and I am indebted to them for sharing their ideas with me and for helping me a lot with their profound knowledge.

I enjoyed being a member of two research groups: the ILLC and the SEN3 theme at the CWI. Many thanks to my friends and colleagues in both groups for the pleasant atmosphere both at work but also at various parties and social events. To start with let me name my office mates: Kees Blom, who shares my admiration of Bach’s music, Falk Bartels, who not only was helpful with his expertise on coalgebras but who also contributed a lot to the social life at the CWI, Nikolay Diakov, who always provided me with various types of Asian tea and Nick Bezhanishvili.

More words have to be said about Nick. He was not only a perfect office mate during the last four years but he became my best friend here in Amsterdam. I learnt a lot about the algebraic semantics of intuitionistic logic from our discussions, but often our conversations had some more serious topics, e.g. the meaning of life or why Belgium beer gives you a headache in the morning.

Special thanks goes also to Joost Joosten for organizing the “Rondje Nieuwmarkt” and the lunch for logic, language and levitation, to David Gabelaia for cooking a great Georgian dinner on my birthday, to Olivier Roy for taking the lead in organizing the Kriterion drinks and many more social events, to Gaelle Fontaine for her famous mousse au chocolat, to Helle Hansen for the new year party, to Fenrong Liu for the Chinese new year party and to Jelle Zuidema for helping me with my moving.

Furthermore I am grateful to all other colleagues for creating a nice atmosphere, in particular to: Farhad Arbab, Frank de Boer, Stefan Bold, Freek Burger, Balder ten Cate, Dave Clarke, David d’Oliviera Costa, Ulle Endrichs, Kees Everaars, Juan Guillen Scholten, Aline Honingh, Joost Jacob, Olivia Ladinig, Benedikt Löwe, Eric Pacuit, Yoav Seginer, Brian Semmes, Merlijn Sevenster, John Sloan, Leigh Smith, Reut Tsarfaty, Peter Zoeteweij. I also would like to express my gratitude to my friends in Munich: Michael, Christoph, Robert, Thomas and Jerri (with an i at the end).

For interesting and inspiring discussions I want to thank Johan van Benthem, Alexandru Baltag, Guram Bezhanishvili, Ian Hodkinson, Dick de Jongh and Larry Moss. Many thanks also to the organizers and attendees of the ACG colloquium at the CWI which serves as a platform for presenting ongoing work and which is a good exercise for defending your ideas against a sometimes tough audience.

Finally I want to thank my family, in particular my parents for their love and support and my brother Markus for stimulating my interest in mathematics. Last but not least I thank Yuliya for distracting me as often as possible from my work and for making me feeling at home in Amsterdam.

Amsterdam
February, 2006.

Clemens Kupke

Contents

Chapter 1

Introduction

This thesis studies various finitary modal languages for reasoning about coalgebras. Coalgebras can be seen as abstract state-based systems or non-well-founded structures. We first briefly explain what coalgebras are and how they are related to modal logic. After that we discuss the issues concerning coalgebraic modal logic that are addressed in this thesis.

1.1 Coalgebras

In this section we give a very short introduction to coalgebra. For a detailed introduction to universal coalgebra the reader is referred to [JR97, Rut00]. Readers whose background lies in modal logic and its algebraic semantics are recommended to consult [Ven06]. The lecture notes [Kur01b, Pat03b] also provide an introduction to coalgebra with a focus on coalgebraic modal logic.

1.1.1 Finite vs. infinite words

Instead of starting with a formal definition of what a coalgebra is, we first want to give the reader some intuition. Our starting point is the following simple inductive definition of the set A^* of finite words over some alphabet A.

$$A^* \ni w ::= \epsilon \mid a \in A \mid w \cdot w.$$

Expressed in words this means that A^* is the *least* set that contains the empty word ϵ and the one-letter word a for every letter $a \in A$, and that is closed under composition of words, i.e. if w_1 and w_2 are finite words then their concatenation $w_1 \cdot w_2$ is also a finite word.

This definition of the set of finite words should demonstrate the underlying pattern of defining algebras: finite words are *constructed* from a collection of basic constants using the algebraic *operations*. Associated with this perspective of operations as constructors comes a suitable notion of equivalence: Two finite words are equivalent if they

Algebra	Coalgebra
operations	observations
congruence	behavioural equivalence
initial algebra	final coalgebra
least fixed-point	largest fixed-point

Table 1.1: Algebra vs. Coalgebra

have been constructed out of equivalent (shorter) words. The definition of a *congruence* formalizes this notion of equivalence. Let us now move to a coalgebraic example: the set of infinite words A^ω over some alphabet A. We can define A^ω *coinductively* as the *largest* set X with the property that for every element w of X there is a letter a and some $w' \in X$ such that $w = a \cdot w'$, i.e. such that w consists of the letter a followed by w'. As a consequence we can define so-called *observations* $\mathsf{hd} : A^\omega \to A$ ("head") and $\mathsf{tl} : A^\omega \to A^\omega$ ("tail") that map an infinite word $w = a \cdot w'$ to its first letter $\mathsf{hd}(w) = a$ and to the remaining word $\mathsf{tl}(w) = w'$. The set A^ω together with these observations $\langle \mathsf{hd}, \mathsf{tl} \rangle : A^\omega \to A \times A^\omega$ is an example of a coalgebra.

Note that, unlike the definition of the algebra of finite words, this coalgebraic definition of A^ω does not carry information about how to construct elements of A^ω. Instead, we can use the observations to obtain limited information about them. This leads to the notion of *behavioural equivalence* or *bisimilarity* of two elements of a coalgebra: two elements of a coalgebra are behaviourally equivalent, if we cannot distinguish them using the observations. In our example this amounts to saying that two infinite words w_1, w_2 are behaviourally equivalent if their first letters $\mathsf{hd}(w_1)$ and $\mathsf{hd}(w_2)$ are equal and their tails $\mathsf{tl}(w_1)$ and $\mathsf{tl}(w_2)$ are again equivalent.

1.1.2 Formal definition

The formal definition of a coalgebra involves basic notions from category theory. For a brief summary of the category theory that is needed in this thesis the reader is referred to Appendix A.

1.1.1. Definition. Let C be a category and $T : \mathsf{C} \to \mathsf{C}$ be a functor. Then a *T-coalgebra* is a pair (X, γ) where $X \in \mathsf{C}$ and $\gamma : X \to TX \in \mathsf{C}$.

Throughout this thesis we will only consider coalgebras for functors over so-called *concrete categories* C, i.e. we can think of objects of C as sets, possibly together with some additional structure. Given a T-coalgebra (X, γ) we refer to X as the *set of states* and to γ as to the *coalgebra map* or *successor function*. The above example of the set of infinite words can be easily seen as a coalgebra in this formal sense.

1.1.2. Example. Let A be a set and $T : \mathsf{Set} \to \mathsf{Set}$ be the functor $(A \times \_)$ which maps a set X to the cartesian product $A \times X$. Then the set A^ω together with the map $\langle \mathsf{hd}, \mathsf{tl} \rangle : A^\omega \to A \times A^\omega$ is a T-coalgebra.

From a modal logic perspective the prime examples for coalgebras are Kripke frames or transition systems.

1.1.3. Example. A Kripke frame is a pair (W, R) such that W is a set of *states* or *worlds* and $R \subseteq W \times W$ is a binary relation. It is easy to see that Kripke frames correspond to $\mathcal{P}$-coalgebras, where $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$ denotes the power set functor: A Kripke frame (W, R) corresponds to the $\mathcal{P}$-coalgebra $(W, R[\_])$, where $R[\_] : W \to \mathcal{P}W$ denotes the function that maps a state $w \in W$ to the set $R[w] \subseteq W$ of R-successors of w. A $\mathcal{P}$-coalgebra (X, γ) on the other hand corresponds to the Kripke frame (X, R_γ) where $R_\gamma \subseteq X \times X$ is defined by putting $(x, y) \in R_\gamma$ if $y \in \gamma(x)$.

Many more examples of different types of objects with possibly infinite behaviour which are captured by the definition of a T-coalgebra can be found in the literature, see e.g. [Rut00].

It is legitimate to ask whether category-theoretic terminology is really needed for the definition of a coalgebra. And, indeed, in the book [BM96] by Barwise and Moss about non-well-founded and coalgebraic phenomena the authors do not use category theory explicitly. In our work, however, the use of category theory is mandatory as one of our claims in this thesis is that coalgebras for functors over base categories other than the category of sets are interesting from a modal logic perspective. Moreover the category-theoretic formulation of the definition of a coalgebra makes it easy to see that coalgebras are indeed, as their name suggest, the categorical dual of algebras (cf. Def. B.2.6 for the definition of an algebra for a functor).

Perhaps one of the strongest arguments for the categorical formulation of the definition of a coalgebra is that it enables us to define a very natural notion of a behaviour preserving map between coalgebras. Let us first look again at the example of Kripke frames.

1.1.4. Example. Behaviour preserving maps between Kripke frames are the so-called *bounded morphisms*. Given two Kripke frames (W_1, R_1) and (W_2, R_2) a function $f : W_1 \to W_2$ is called a bounded morphism from (W_1, R_1) to (W_2, R_2) if f satisfies the following two conditions:

(i) for all $w \in W_1$, $(w, v) \in R_1$ implies $(f(w), f(v)) \in R_2$, and

(ii) if $(f(w), v') \in R_2$ for some $w \in W_1$ and $v' \in W_2$ then there exists some $v \in W_1$ such that $(w, v) \in R_1$ and $f(v) = v'$.

These conditions can be concisely summarized in the following diagram:

$$\begin{array}{ccc} \mathcal{P}W_1 & \xrightarrow{\mathcal{P}f} & \mathcal{P}W_2 \\ {\scriptstyle R_1[\_]}\uparrow & & \uparrow{\scriptstyle R_2[\_]} \\ W_1 & \xrightarrow[f]{} & W_2 \end{array}$$

where for a function $f : W_1 \to W_2$, $\mathcal{P}f = f[\_]$ denotes the direct image function. It can be easily checked that a function $f : W_1 \to W_2$ makes the diagram commute iff f is a bounded morphism between (W_1, R_1) and (W_2, R_2).

The notion of a T-coalgebra morphism generalizes bounded morphisms in a natural way.

1.1.5. Definition. A *T-coalgebra morphism* $f : (X, \gamma) \to (Y, \delta)$ between two T-coalgebras (X, γ) and (Y, δ) is a morphism $f : X \to Y \in \mathsf{C}$ such that the following diagram commutes

$$\begin{array}{ccc} TX & \xrightarrow{Tf} & TY \\ {\scriptstyle\gamma}\uparrow & & \uparrow{\scriptstyle\delta} \\ X & \xrightarrow[f]{} & Y \end{array}$$

The category $\mathsf{Coalg}(T)$ has as objects T-coalgebras and T-coalgebra morphisms as arrows.

In particular, the category $\mathsf{Coalg}(\mathcal{P})$ consists of Kripke frames as objects and bounded morphisms as arrows. Therefore $\mathcal{P}$-coalgebra morphisms provide a good notion of a structure preserving map between Kripke frames. Again many more examples can be found in the literature (cf. e.g. [Rut00]).

When talking about infinite words we mentioned that the coalgebraic notion of equivalence between states is *behavioural equivalence* or *bisimilarity*.

1.1.6. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor and $(X_1, \gamma_1), (X_2, \gamma_2)$ be T-coalgebras. We say that two states $x_1 \in X_1$ and $x_2 \in X_2$ are *behaviourally equivalent* or *bisimilar* if there is a T-coalgebra (Y, δ) and T-coalgebra morphisms $f_1 : (X_1, \gamma_1) \to (Y, \delta)$ and $f_2 : (X_2, \gamma_2) \to (Y, \delta)$ such that $f_1(x_1) = f_2(x_2)$.

In other words, two states are bisimilar if they can be identified by coalgebra morphisms. That is, our definition of bisimilarity is based on our view that bisimilarity should be the same as *behavioural equivalence*.

1.1.7. Remark. Note that this definition is not completely standard. Usually, one sees as the definition of bisimilarity that two states are bisimilar if they are linked by a *T-bisimulation* (cf. A.3.5 for a definition). In many cases the two definitions coincide; to be precise, this applies to all functors that preserve so-called weak pullbacks. All so-called Kripke polynomial functors, which include the power set functor, have this property.

Our perspective on the matter is that behavioural equivalence is the more fundamental notion. In the case that T preserves weak pullbacks, then the existence of a T-bisimulation is a nice and concise way of capturing bisimilarity between two states.

Bisimulations also play a central role in modal logic, as modal logic can be seen as the bisimulation invariant fragment of first-order logic ([Ben76]). Again the coalgebraic notion generalizes the notion from modal logic: coalgebraic $\mathcal{P}$-bisimulations (cf. Definition A.3.5) are exactly the bisimulations from modal logic for the language without propositional variables. For a proof of this fact we refer the reader to [Rut00, Example 2.1].

We conclude this short introduction to coalgebra by mentioning the important concept of a *final coalgebra*. A final coalgebra for a functor T, if it exists, can be thought of as the T-coalgebra that contains for each element x of an arbitrary T-coalgebra exactly one state that is bisimilar or behaviourally equivalent to x. In categorical terms the final T-coalgebra is the final object in $\mathsf{Coalg}(T)$, i.e. for every T-coalgebra (X, γ) there is a unique T-coalgebra morphism into the final coalgebra. The set A^{ω} from Example 1.1.2 together with the coalgebra map $\langle \mathsf{hd}, \mathsf{tl} \rangle$ is the final $(A \times \_)$-coalgebra. The final $\mathcal{P}$-coalgebra, on the other hand, does not exist, or better, it is not a set. It is the set-theoretic universe of non-well-founded set theory (cf. [Acz88]). A final coalgebra that is very familiar to modal logicians is the so-called canonical model for the basic normal modal logic **K** without propositional variables. We will see in Chapter 3 that it is the final coalgebra of the Vietoris functor $\mathbb{V} : \mathsf{Stone} \to \mathsf{Stone}$.

1.2 Specifying coalgebras using modal logic

1.2.1 Why using modal languages?

We saw that Kripke frames, which constitute the standard semantics of modal logic, can be seen as coalgebras, and that other notions from modal logic, such as the notion of a bounded morphism and of a bisimulation, have their natural place on the level of coalgebras as well. But there are many other examples where modal languages are used to reason about coalgebras: Coalgebras generalize infinite structures, such as infinite words, infinite trees or transition systems. Various modal languages have been successfully applied to reason about these structures, e.g. **LTL** for infinite words, **CTL** and **CTL*** for infinite trees and the modal μ-calculus for transition systems.

But the close connection between coalgebras and modal logic is made manifest not only in various examples. It was observed and made precise by Kurz in [Kur00] that the duality between algebras and coalgebras can be extended to logics. His observation can be summarized by the following slogan: Modal logic and coalgebras dualize equational logic and algebras.

1.2.2 Coalgebraic modal logics

But which modal languages can be used for reasoning about coalgebras? This is in fact an issue which is still under discussion.

Research on this question goes back to work by Moss ([Mos99]). Moss' coalgebraic logic is very general. It assigns a logical language to every weak pullback preserving endofunctor on the category of sets. His syntax allows for infinite conjunctions and contains a somewhat non-standard modal operator ∇. The idea of the ∇-operator can be sketched as follows: Let Φ be the collection of formulas of the language associated to some functor $T : \mathsf{Set} \to \mathsf{Set}$. Formulas in Φ describe properties of states of a given T-coalgebra (X, γ). Then elements of the set $T\Phi$ should describe properties of successor states, i.e. of elements in TX. Therefore the language of coalgebraic logic contains for each $\pi \in T\Phi$ a formula $\nabla\pi \in \Phi$. The definition of the semantics exploits the fact, that every weak pullback preserving functor $T : \mathsf{Set} \to \mathsf{Set}$ can be extended to an endofunctor $\overline{T}$ on the category Rel. With this *relation lifting* we can lift the satisfiability relation $\models \subseteq X \times \Phi$ between states and formulas to a relation $\overline{T}(\models) \subseteq TX \times T\Phi$. The formula $\nabla\pi$ is then defined to be true at some state $x \in X$ if $\gamma(x)$ makes π true, i.e. if $(\gamma(x), \pi) \in \overline{T}(\models)$.

Moss's coalgebraic logic was followed by work by Kurz ([Kur01c]) and Rössiger ([Röß01, Röß00]). Kurz defines a finitary, multi-modal language for coalgebras for a limited class of endofunctors on Set that are of the shape

$$X \mapsto \prod_{i=1}^{n} (B_i + C_i \times X)^{A_i} \quad (A_i, B_i, C_i \in \mathsf{Set}).$$

These functors include the one that is used to describe Mealy machines as coalgebras and those functors which were employed in earlier work by Jacobs ([Jac96]) and Reichel ([Rei95]) to model certain features, in particular encapsulation, of object-oriented programming.

Rössiger's idea in [Röß01, Röß00] was to consider an inductively defined class of so-called polynomial functors and to inductively associate finitary multi-modal languages with them. A polynomial functor is a functor that can be constructed from the identity functor and constant functors by forming products and sums of functors and by allowing exponentiation of a functor with an arbitrary set. In particular all functors that were considered by Kurz in [Kur01c] are polynomial. Jacobs extended the class of polynomial functors to the class of so-called Kripke polynomial functors in [Jac01]. For the construction of these Kripke polynomial functors one can also use the power set functor, which makes it possible to model non-deterministic systems.

A slightly different approach for defining a logic for polynomial functors is provided by the work of Goldblatt (see e.g. [Gol01]). Instead of working with modal formulas he uses equations between terms that contain one variable. These variables have states as possible values. Bisimilarity is characterized by Boolean combinations of equations, i.e. two states satisfy the same equations iff they are bisimilar. Further results of Goldblatt's work are coalgebraic constructions of ultraproducts ([Gol03b]) and ultrafilter extensions ([Gol03a]).

Another line of research in the area of coalgebraic modal logic started with the work of Pattinson in the papers ([Pat01, Pat03a, Pat04]). The central idea in his approach is that the modalities for a coalgebraic modal logic should be interpreted as so-called *predicate liftings* for an endofunctor $T : \mathsf{Set} \to \mathsf{Set}$: natural transformations that map predicates over a set X (of states) to predicates over the set TX. Furthermore, Pattinson introduces a derivability relation for coalgebraic modal logic that is parametrized in a set of axioms.

What should be the criteria for deciding which modal language is suitable for reasoning about a given type of coalgebras? In this thesis we focus on three properties of a logic:

- *soundness* of a set of axioms w.r.t. the coalgebraic semantics, i.e. the property that every theorem of the logic is valid on all coalgebras of a certain type,
- *completeness* of a set of axioms w.r.t. the coalgebraic semantics, i.e. the property that every formula that is valid on all coalgebras of a certain type can be derived from the axioms, and
- *expressiveness*[1], i.e. the property that two coalgebra states satisfy the same formulas of the language iff they are bisimilar.

The inductively defined logics for Kripke polynomial functors have a sound and complete axiomatisation (cf. [Röß01, Jac01]). Furthermore these logics are expressive if the functor does not involve the power set functor. In [Pat03a] Pattinson states conditions on a set of axioms for his coalgebraic modal logics to be sound and complete. He also gives a sufficient condition for the predicate liftings to give rise to an expressive language ([Pat04]). For a version of Pattinson's language containing polyadic modalities, Schröder recently improved on Pattinson's earlier results ([Sch05]). We come back to this issue in Section 2.2.3.

Moss' coalgebraic logics are expressive for all functors $T : \mathsf{Set} \to \mathsf{Set}$ for which they are defined, i.e. for standard and weak pullback preserving functors. There is, however, neither a complete axiomatisation nor a conjecture of how a complete axiomatisation of his logic could look like.

Before we finish this survey of different coalgebraic modal logics let us mention that it is also possible to combine the different approaches. This was first pointed out by Cîrstea in [Cîr04] by showing that one can compose expressive languages associated with coalgebras for certain functors to obtain expressive languages for coalgebras for more complex functors, e.g. we can combine expressive logics for T_1- and T_2-coalgebras into an expressive logic for $T_1 \times T_2$-coalgebras. In a continuation of this line of research, Cîrstea & Pattinson demonstrated in [CP04] that also proof systems for coalgebraic modal logics can be combined in a similar fashion. This was applied in *loc.cit.* to obtain axiomatisations of complex probabilistic systems which can be modeled as coalgebras.

1.2.3 Issues concerning finitary modal languages

In this thesis we focus on modal languages for specifying coalgebras that have a finitary syntax. This restriction brings up several issues:

Expressiveness Modal languages with finitary syntax are in general not expressive with respect to the coalgebraic semantics given by $\mathsf{Coalg}(T)$ for a functor $T : \mathsf{Set} \to \mathsf{Set}$. This is in particular the case for ordinary modal logic (without

[1] This property is often also referred to as the Hennessy-Milner property. We use the term *expressiveness*, following existing work in the field of coalgebraic modal logic.

propositional variables) and $\mathcal{P}$-coalgebras, i.e. Kripke semantics. A natural question to ask is therefore, whether one can find a class of models for these logics that allows for expressive finitary languages. In this thesis we propose to resolve this issue by generalizing a well-known concept from modal logic, namely the concept of a descriptive general frame ([Gol76]). This leads to our work on what we call *Stone coalgebras*. (cf. 1.3.1 below).

Algebraic semantics Algebraic semantics is known to be useful for studying modal logic with finitary syntax (cf. e.g. [BdV01, Chapter 5], [CZ97]). A first exploitation of that idea in the coalgebraic framework can be found in [Jac01] for the inductively defined logics for Kripke polynomial functors. In this thesis we define an algebraic semantics for Pattinson's coalgebraic modal logic (cf. 1.3.2 below).

Fixed-point logics The examples of logics for specifying infinite structures in Section 1.2.1 demonstrate that finitary languages usually need some fixed-point or temporal operators in order to have the the ability to specify possibly infinite, ongoing behaviour, i.e. so-called liveness conditions. Venema ([Ven04]) addresses this issue by defining a finitary version of Moss' coalgebraic logic and by adding fixed-point operators to it. Formulas of his *coalgebraic fixed-point logic* have a natural automata-theoretic interpretation in terms of so-called coalgebra automata. In this thesis we discuss closure properties of coalgebra automata and their non-emptiness problem (cf. 1.3.3 below).

1.3 Our contributions

1.3.1 Stone coalgebras

One of our results is that *Stone coalgebras*, i.e. coalgebras for functors on the category of so-called Stone spaces, are a useful tool in studying coalgebraic modal logics. Our starting point is the observation that descriptive general frames can be represented as coalgebras for the Vietoris functor on the category of Stone spaces Stone. Although we made this observation independently in [KKV04], with hindsight it may be read in [Esa74], while we later found out that it had been made explicit in unpublished[2] lecture notes by Abramsky ([Abr88]). Hence coalgebras for endofunctors over Stone are a natural generalization of this concept. We apply this idea to the inductively defined logics for Kripke polynomial functors. For every Kripke polynomial functor we define a corresponding *Vietoris polynomial functor* on the category of Stone spaces and show that the final coalgebra for these Vietoris polynomial functors can be constructed using a generalized version of the canonical model construction from modal logic. As a corollary of this construction we obtain the result that the languages associated with Vietoris polynomial functors have the Hennessy-Milner property. Furthermore we prove

[2] Very recently, Abramsky's notes have been published as [Abr05].

that for every Vietoris polynomial functor $\mathbb{T}$ and the logic associated to it there exists an adjunction between the algebraic semantics of the logic, defined as a category of many-sorted algebras as in [Jac01], and the category of $\mathbb{T}$-coalgebras. Finally we give a characterization of those many-sorted algebras for which this adjunction is an equivalence of categories.

1.3.2 Algebraic semantics of coalgebraic modal logic

We show that the algebraic semantics of any coalgebraic modal logic that is given by a set of predicate liftings and a set of axioms of modal depth 1 can be formulated as a category of algebras for an endofunctor L on the category of Boolean algebras. Furthermore we relate the algebraic and the coalgebraic semantics via a natural transformation δ and show that certain properties of δ entail soundness, completeness and expressiveness of the logic. We will see that the soundness, completeness and expressiveness criteria are equivalent to earlier criteria given by Pattinson in [Pat03a, Pat04]. Our results improve on his work because we uniformly treat coalgebras for functors over Set and Stone and our approach can be easily generalized to other base categories. For the case of Stone coalgebras, we obtain a characterization of duality: the logic is expressive and fulfills Pattinson's soundness and completeness criteria iff the functor describing its algebraic semantics is dual to T.

1.3.3 Coalgebraic logics and coalgebra automata

Building on the work by Venema in [Ven04] on coalgebra automata we prove that certain results, that are already known about (finite) parity automata on (possibly) infinite words, trees and graphs, can be obtained at the more general level of finite automata on rooted coalgebras. The main result here is that alternating automata can always be transformed into equivalent non-deterministic ones and that non-emptiness of coalgebra automata is decidable for a large class of functors. These results relate to coalgebraic modal logics because of the fact that coalgebra automata and formulas of the so-called coalgebraic fixed-point logics are in one-to-one correspondence, a fact that can be summarized by the slogan: Coalgebra automata are formulas.

Because of this close connection, the above automata-theoretic results have logical corollaries: all coalgebraic fixed-point logics have the finite model property, and we can prove the soundness of a certain distributive law for the ∇-operator.

1.4 Origin of the presented material

This thesis consists in large parts of earlier published material. The relation between chapters and published papers is summarized in the following table.

Chapter	based on
3	[KKV04]
4	[KKP04]
5	[KV05]

The thesis is structured as follows: In Chapter 2 we give an overview of the coalgebraic modal logics that are discussed later on in this thesis. In Chapter 3 we present our work on Stone coalgebras (cf. 1.3.1 above). In Chapter 4 we introduce the algebraic semantics for coalgebraic modal logic (cf. 1.3.2 above). In Chapter 5 we prove closure properties of coalgebra automata and, in doing so, we obtain corollaries for coalgebraic fixed-point logic (cf. 1.3.3 above).

Chapter 2

Coalgebraic Modal Logics

In this chapter we introduce the logics for specifying coalgebras for a given functor T that we will study in this thesis. Three different approaches are discussed: first the inductively defined logics for Kripke polynomial functors (cf. Section 2.1), then coalgebraic modal logics given by predicate liftings (cf. Section 2.2) and finally finitary coalgebraic fixed-point logic (cf. Section 2.3). The chapter does not contain new results but provides the necessary background for this thesis.

2.1 Inductively defined logics

The main idea of this approach is that we can already capture many interesting examples of T-coalgebras by focusing on an inductively defined set of endofunctors on Set. Using the syntactic structure of these functors we can inductively define suitable logics. This idea was first exploited by Rößiger in [Röß01, Röß00]. The approach we are going to describe is slightly different from Rößiger's and is basically taken from [Jac01]. We will now first introduce the set of Kripke polynomial functors and then state the definition of sound and complete logics for any such functor.

2.1.1 Kripke polynomial functors

2.1.1. Definition. The set of *Kripke polynomial functors (KPF)* is inductively defined as follows:

$$KPF ::= Id \mid A \in \mathsf{FinSet} \mid T + T \mid T \times T \mid T^D, D \in \mathsf{Set} \mid \mathcal{P}T,$$

where $Id : \mathsf{Set} \to \mathsf{Set}$ denotes the identity functor, for a finite set A we write A for the constant functor mapping every set to A, and the $+$ and $\times$ denote disjoint union and binary product respectively. Furthermore given an arbitrary set D we write T^D for the functor mapping a set X to the D-fold product (or exponent) $(TX)^D$. The class of *polynomial functors (PF)* consists of all functors $T \in KPF$ that do not involve the power set functor.

Coalgebras for Kripke polynomial functors can be seen as abstract possibly non-deterministic transition systems, coalgebras for polynomial functors as abstract deterministic transition systems. In order to substantiate this claim we provide some examples.

2.1.2. Example.

1. Coalgebras for the functor Id^D correspond to deterministic transition systems in which every state has for each input $e \in E$ exactly one successor.
2. Coalgebras for the functor $(O \times Id)^D$ correspond to similar transition systems with the difference that when moving from one state to another the system produces an output $o \in O$ (Mealy automata).
3. Coalgebras for the power set functor $\mathcal{P}$ correspond to unlabeled graphs and therefore to non-deterministic transition systems.

In order to be able to define logics for Kripke polynomial functors we first have to analyze the syntactical structure of these functors.

2.1.3. Definition. We inductively define *paths* between *KPF*s

$$p ::= \epsilon \mid \pi_1 \cdot p \mid \pi_2 \cdot p \mid \kappa_1 \cdot p \mid \kappa_2 \cdot p \mid \mathrm{ev}(d) \cdot p \mid \mathsf{pow} \cdot p$$

and denote by *PCons* the set of *path constructors* $PCons := \{\pi_1, \pi_2, \kappa_1, \kappa_2, \mathrm{ev}(d), \mathsf{pow}\}$. Furthermore we define when two *KPF*s T_1 and T_2 are related via a path p, which will be denoted by $p : T_1 \rightsquigarrow T_2$.

$$\begin{array}{rlll}
\epsilon : & T \rightsquigarrow T & & \\
\pi_i \cdot p : & T_1 \times T_2 \rightsquigarrow T' & \text{if } p : T_i \rightsquigarrow T' & \text{for } i \in \{1,2\} \\
\kappa_i \cdot p : & T_1 + T_2 \rightsquigarrow T' & \text{if } p : T_i \rightsquigarrow T' & \text{for } i \in \{1,2\} \\
\mathrm{ev}(d) \cdot p : & T^D \rightsquigarrow T' & \text{if } p : T \rightsquigarrow T' & \text{for } d \in D \\
\mathsf{pow} \cdot p : & \mathcal{P}T \rightsquigarrow T' & \text{if } p : T \rightsquigarrow T'. &
\end{array}$$

For a functor $T \in KPF$ we define the category $\mathsf{Ing}(T)$ of ingredients of T as the category with elements of $\mathbf{Ing}(T) := \{T' \mid \exists p.p : T \rightsquigarrow T'\} \cup \{Id\}$ as objects and paths as morphisms between them.

2.1.4. Example. Let T be the functor $\mathcal{P}(O \times Id)$. Then the category $\mathsf{Ing}(T)$ can be depicted as in the diagram below (the identity arrows $\epsilon : T \rightsquigarrow T$ have been omitted):

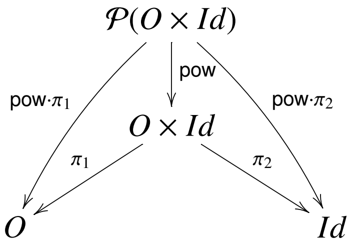

The logic which we associate with a given *KPF* T is a many-sorted modal logic, $\mathbf{Ing}(T)$ is the set of sorts and the modalities correspond to the path constructors. The semantics of a formula ϕ of type T' on a T-coalgebra (X, γ) will be given by a subset of $T'X$ as we will see in the following subsection.

$$\frac{}{\bot \in \mathbf{Form}_S}\ (\mathrm{Bool}_1) \qquad \frac{\phi_1 \in \mathbf{Form}_S \quad \phi_2 \in \mathbf{Form}_S}{\phi_1 \to \phi_2 \in \mathbf{Form}_S}\ (\mathrm{Bool}_2)$$

$$\frac{\phi \in \mathbf{Form}_{S_2} \qquad p : S_1 \rightsquigarrow S_2 \text{ and } p \in PCons}{[p]\phi \in \mathbf{Form}_{S_1}}\ (\mathrm{modal})$$

$$\frac{\phi \in \mathbf{Form}_T}{\mathbf{next}\ \phi \in \mathbf{Form}_{Id}}\ (\mathrm{next})$$

Figure 2.1: Closure rules

2.1.2 The logic MSM$\mathcal{L}$

We now turn to the definition of the syntax and semantics of the many-sorted modal logic (MSM$\mathcal{L}$) for coalgebras for Kripke polynomial functors.

2.1.5. Definition. Let T be a KPF. The set of *raw T-formulas* is defined as follows:

$$\phi ::= \bot \mid a \in A,\ A \in \mathbf{Ing}(T) \mid \phi \to \phi \mid \mathbf{next}\ \phi \mid [p]\phi,\ p \in PCons.$$

Furthermore we use the standard abbreviations $\neg\phi$, $\phi_1 \wedge \phi_2$, $\phi_1 \vee \phi_2$, and $\phi_1 \leftrightarrow \phi_2$ and we let

$$\phi_1 \dot{\vee} \phi_2 := (\phi_1 \vee \phi_2) \wedge \neg(\phi_1 \wedge \phi_2).$$

Given a functor $T \in KPF$ the set of sorted formulas **Form** is defined as a family $(\mathbf{Form}_S)_{S \in \mathbf{Ing}(T)}$ of sets of raw T-formulas such that the following closure rules are satisfied.

- for every $S \in \mathbf{Ing}(T)$, $\mathbf{Form}_S$ contains $\bot$ and is closed under implication (cf. rules (Bool_1) and (Bool_2) in Figure 2.1).
- for every $S_1, S_2 \in \mathbf{Ing}(T)$ and every path constructor $p : S_1 \rightsquigarrow S_2$ we can construct formulas of type S_1 by prefixing the modality $[p]$ to formulas of type S_2 (cf. rule (modal) in Figure 2.1).
- formulas of sort T can be transformed into formulas of type Id using the **next** modality (cf. rule (next) in Figure 2.1).

Formulas $\phi \in \mathbf{Form}_{Id}$ of sort Id are called *state formulas*.

The idea for the inductive definition of the semantics of the logic can be sketched as follows: Let $T_1, T_2 \in \mathbf{Ing}(T)$ such that $p : T_1 \rightsquigarrow T_2$ for some path constructor p and let ϕ be a formula of the logic of type T_2. Furthermore assume that the semantics of ϕ is already defined. In order to define the semantics of the formula $[p]\phi$ (which is of type T_1) on a T-coalgebra (X, γ), we lift the semantics of ϕ, i.e. a subset of T_2X, to a subset of T_1X.

To sum it up we need for every $p : T_1 \rightsquigarrow T_2$ a lifting that maps subsets of T_2X to corresponding subsets of T_1X for all $X \in \mathsf{Set}$. The next definition provides suitable liftings.

2.1.6. Definition. Let T be a KPF and let X be a set. Then we define for any two functors $T_1, T_2 \in \mathbf{Ing}(T)$ and path $p : T_1 \rightsquigarrow T_2$, a function $(\_)^p : \mathcal{P}(T_2X) \to \mathcal{P}(T_1X)$ by induction on the complexity of paths. For $\alpha \subseteq T_2X$ we put

$$\begin{array}{rcll}
\alpha^{\epsilon} & := & \alpha & \\
\alpha^{\pi_1 \cdot p} & := & \pi_1^{-1}[\alpha^p] & \\
\alpha^{\pi_2 \cdot p} & := & \pi_2^{-1}[\alpha^p] & \\
\alpha^{\kappa_1 \cdot p} & := & \kappa_1[\alpha^p] \cup \kappa_2[S_2X] & \text{for } T_2 = S_1 + S_2 \\
\alpha^{\kappa_2 \cdot p} & := & \kappa_1[S_1X] \cup \kappa_2[\alpha^p] & \text{for } T_2 = S_1 + S_2 \\
\alpha^{\mathsf{ev}(d) \cdot p} & := & \pi_d^{-1}[\alpha^p] & \\
\alpha^{\mathsf{pow} \cdot p} & := & \{\beta \in \mathcal{P}T_1X \mid \beta \subseteq \alpha^p\} & \text{for} \quad p : T_1 \rightsquigarrow T_2.
\end{array}$$

For the reader, who is familiar with modal logic, we also want to formulate the semantics of the inductively defined logics in terms of relations which will correspond to the modalities of the language.

2.1.7. Definition. Let T be a KPF, let (X, γ) be a T-coalgebra and suppose $p : S_1 \rightsquigarrow S_2 \in \mathbf{Ing}(T)$. Then we define a relation $R_p \subseteq S_1X \times S_2X$ by induction on the complexity of p. For $x \in S_1X$ and $y \in S_2X$ we let

$$\begin{array}{rcll}
xR_{\epsilon}y & :\Leftrightarrow & x = y & \\
xR_{\pi_i \cdot p}y & :\Leftrightarrow & \exists z \,.\, \pi_i(x) = z \text{ and } zR_py & i \in \{1, 2\} \\
xR_{\kappa_i \cdot p}y & :\Leftrightarrow & \exists z \,.\, \kappa_i(z) = x \text{ and } zR_py & i \in \{1, 2\} \\
xR_{\mathsf{pow} \cdot p}y & :\Leftrightarrow & \exists z \,.\, z \in x \text{ and } zR_py &
\end{array}$$

These relations R_p are used for an alternative formulation of the semantics of the logic (cf. Remark 2.1.10). But first we use the lifting functions $(\_)^p$ from Definition 2.1.6 to define the coalgebraic semantics of MSM$\mathcal{L}$.

2.1.8. Definition. Let (X, γ) be an T-coalgebra for some $T \in KPF$. For each $S \in \mathbf{Ing}(T)$ we define an interpretation function

$$[\![\_]\!]^S_{(X,\gamma)} : \mathbf{Form}_S \to \mathcal{P}(SX)$$

by induction on the structure of the formulas.

$$\begin{array}{rcl}
[\![\bot]\!]^S_{(X,\gamma)} & := & \emptyset \\
[\![\phi_1 \to \phi_2]\!]^S_{(X,\gamma)} & := & \left(SX \setminus [\![\phi_1]\!]^S_{(X,\gamma)}\right) \cup [\![\phi_2]\!]^S_{(X,\gamma)} \\
[\![a]\!]^A_{(X,\gamma)} & := & \{a\} \\
[\![\mathbf{next}\ \phi]\!]^{Id}_{(X,\gamma)} & := & \gamma^{-1}\left[[\![\phi]\!]^T_{(X,\gamma)}\right] \\
[\![[p]\phi]\!]^{S_1}_{(X,\gamma)} & := & \left([\![\phi]\!]^{S_2}_{(X,\gamma)}\right)^p \quad \text{for any } p \in PCons \text{ s.t. } p : S_1 \rightsquigarrow S_2.
\end{array}$$

Furthermore for a formula $\phi \in \mathbf{Form}_S$ and an element x of SX we write

$$\begin{aligned} (X,\gamma), x \models_S \phi \quad &\text{if} \quad x \in [\![\phi]\!]^S_{(X,\gamma)}, \\ (X,\gamma) \models_S \phi \quad &\text{if} \quad [\![\phi]\!]^S_{(X,\gamma)} = SX \text{ and} \\ \mathsf{Coalg}(T) \models_S \phi \quad &\text{if} \quad \text{for all } (X,\gamma) \in \mathsf{Coalg}(T)\ .\ (X,\gamma) \models_S \phi. \end{aligned}$$

2.1.9. Remark. The semantics of an MSM$\mathcal{L}$-formula of sort $S \in \mathbf{Ing}(T)$ is defined as a subset of SX. We want to use MSM$\mathcal{L}$-formulas to express properties of coalgebra states $x \in X$ and hence state formulas are the most important formulas of the language. The other formulas are only needed to give a uniform inductive definition of the semantics.

2.1.10. Remark. Let (X,γ) be a T-coalgebra, $p : S_1 \rightsquigarrow S_2$ a path in $\mathbf{Ing}(T)$ and ϕ be a formula of sort S_2. Then we can rephrase the definition of $[\![[p]\phi]\!]_{S_1}$ using the relations R_p from Definition 2.1.7 as follows:

$$[\![[p]\phi]\!]^{S_1}_{(X,\gamma)} = \{x \in S_1X \mid \forall y \in S_2X\ .\ xR_py \Rightarrow y \in [\![\phi]\!]^{S_2}_{(X,\gamma)}\},$$

i.e. $[p]\phi$ is satisfied at some $x \in S_1X$ if all R_p-successors of x satisfy ϕ. In other words the semantics of MSM$\mathcal{L}_F$ can be seen as a sorted variant of the Kripke semantics of the $\Box$-operator of modal logic (cf. 3.1.8).

Having seen the definition of the semantics we next discuss the axioms and derivation rules of the logic MSM$\mathcal{L}$. Certainly every Boolean tautology should be derivable and the set of derivable formulas should be closed under *modus ponens*. Furthermore every modality of the logic should be normal and therefore the **K**-axiom $\Box(\phi_1 \to \phi_2) \to (\Box\phi_1 \to \Box\phi_2)$ should be an axiom of the logic, and the necessitation rule

$$\frac{\vdash \phi}{\vdash \Box\phi}\,(N)\ .$$

should be a valid derivation rule of MSM$\mathcal{L}$. As mentioned above coalgebras for polynomial functors correspond to deterministic transition systems. Hence all modalities that correspond to polynomial functors should satisfy the determinacy axiom $\neg\Box\neg\phi \to \Box\phi$. Furthermore the modalities $[\kappa_i]$ describing the coproduct should fulfill some axiom expressing that if the successor of a state lies in $S_1X + S_2X$ then it lies in *exactly one* of them. This is expressed by the axiom (DC) below. All in all Jacobs in [Jac01] arrives at the following axioms and derivation rules.

2.1.11. Definition. Let $T \in KPF$. For every $S \in \mathbf{Ing}(T)$ we define a *derivability predicate* $\vdash_S \subseteq \mathbf{Form}_S$ such that $\vdash_S \phi$ for each Boolean tautology $\phi \in \mathbf{Form}_S$ and $\vdash_S$ is closed under *modus ponens*

$$\frac{\vdash_S \phi_1 \to \phi_2 \qquad \vdash_S \phi_1}{\vdash_S \phi_2}\ .$$

for finite sets of constants $A \in \mathbf{Ing}(T)$:

$$\vdash_A \dot{\bigvee}_{a \in A} a \qquad \text{(DC)}$$

for the **next**-operator:

$$\vdash_{Id} \textbf{next}\ \phi \leftrightarrow \neg \textbf{next}\ \neg\phi \qquad \text{(Det)}$$

$$\vdash_{Id} \textbf{next}\ (\phi_1 \to \phi_2) \to (\textbf{next}\ \phi_1 \to \textbf{next}\ \phi_2) \qquad \text{(K)}$$

$$\frac{\vdash_T \phi}{\vdash_{Id} \textbf{next}\ \phi}\ \text{(N)}$$

for the $[\pi_i]$-operator:

$$\vdash_{S_1 \times S_2} [\pi_i]\phi \leftrightarrow \neg[\pi_i]\neg\phi \qquad \text{(Det)}$$

$$\vdash_{S_1 \times S_2} [\pi_i](\phi_1 \to \phi_2) \to ([\pi_i]\phi_1 \to [\pi_i]\phi_2) \qquad \text{(K)}$$

$$\frac{\vdash_{S_i} \phi}{\vdash_{S_1 \times S_2} [\pi_i]\phi}\ \text{(N)}$$

for the $[\mathrm{ev}(d)]$-operator:

$$\vdash_{S^D} [\mathrm{ev}(d)]\phi \leftrightarrow \neg[\mathrm{ev}(d)]\neg\phi \qquad \text{(Det)}$$

$$\vdash_{S^D} [\mathrm{ev}(d)](\phi_1 \to \phi_2) \to ([\mathrm{ev}(d)]\phi_1 \to [\mathrm{ev}(d)]\phi_2) \qquad \text{(K)}$$

$$\frac{\vdash_S \phi}{\vdash_{S^D} [\mathrm{ev}(d)]\phi}\ \text{(N)}$$

for the $[\kappa_i]$-operator:

$$\vdash_{S_1 + S_2} (\neg[\kappa_1]\ \bot)\dot{\vee}(\neg[\kappa_2]\ \bot) \qquad \text{(DC)}$$

$$\vdash_{S_1 + S_2} (\neg[\kappa_i]\ \bot) \to ([\kappa_i]\phi \leftrightarrow \neg[\kappa_i]\neg\phi) \qquad \text{(Det)}$$

$$\vdash_{S_1 + S_2} [\kappa_i](\phi_1 \to \phi_2) \to ([\kappa_i]\phi_1 \to [\kappa_i]\phi_2) \qquad \text{(K)}$$

$$\frac{\vdash_{S_i} \phi}{\vdash_{S_1 \times S_2} [\kappa_i]\phi}\ \text{(N)}$$

for the $[\mathsf{pow}]$-operator:

$$\vdash_{\mathcal{P}S} [\mathsf{pow}](\phi_1 \to \phi_2) \to ([\mathsf{pow}]\phi_1 \to [\mathsf{pow}]\phi_2) \qquad \text{(K)}$$

$$\frac{\vdash_S \phi}{\vdash_{\mathsf{pow}S} [\mathsf{pow}]\phi}\ \text{(N)}$$

Table 2.1: Rules and axioms of MSM$\mathcal{L}$

In addition, the derivability predicates contain the axioms and satisfy the rules that are listed in Table 2.1. The pair

$$\big((\mathbf{Form}_S)_{S\in\mathbf{Ing}(T)}, (\vdash_S)_{S\in\mathbf{Ing}(T)}\big)$$

will be called the *many-sorted modal logic of T* and will be denoted by $\mathrm{MSM}\mathcal{L}_T$.

Given the logic of a *KPF T* it is natural to ask the question whether the logic is sound and complete with respect to the coalgebraic semantics.

2.1.12. Definition. Consider a *KPF T* and the corresponding logic $\mathrm{MSM}\mathcal{L}_T$. We say that $\mathrm{MSM}\mathcal{L}_T$ is *sound* w.r.t. the coalgebraic semantics if for all $S \in \mathbf{Ing}(T)$ and all $\phi \in \mathbf{Form}_S$ we have

$$\vdash_S \phi \text{ implies } \mathsf{Coalg}(T) \models_S \phi.$$

We say that $\mathrm{MSM}\mathcal{L}_T$ is *complete* w.r.t. the coalgebraic semantics if for all $S \in \mathbf{Ing}(T)$ and all $\phi \in \mathbf{Form}_S$

$$\mathsf{Coalg}(T) \models_S \phi \text{ implies } \vdash_S \phi.$$

Soundness of $\mathrm{MSM}\mathcal{L}_T$ is not too hard to prove.

2.1.13. Proposition ([Jac01, Lemma 3.5]). *For every KPF T the logic* $\mathrm{MSM}\mathcal{L}_T$ *is sound with respect to the coalgebraic semantics.*

Proof. The claim is proven by induction on the length of the derivation. QED

Proving completeness is more difficult. A completeness proof for the class of polynomial functors (not involving the power set functor) can be found in [Röß01]. In [Jac01] Jacobs extends this proof to Kripke polynomial functors. We obtain a completeness proof of $\mathrm{MSM}\mathcal{L}_T$ for Kripke polynomial functors in Chapter 3 as a corollary of the duality between the algebraic and coalgebraic semantics (cf. Corollary 3.4.12).

2.2 Logics given by predicate liftings

Next we are going to discuss a second family of logics, the coalgebraic modal logics that are given by a set of predicate liftings and a set of axioms. In the definition of sound and complete logics for Kripke polynomial functors, liftings of predicates played a central role (cf. Def. 2.1.6). A more abstract definition of predicate liftings for an arbitrary functor $T : \mathsf{Set} \to \mathsf{Set}$ was given in [Pat04] by Pattinson. He uses a multi-modal (but not many-sorted) logic in which each modal operator corresponds to such a lifting.

In this thesis we will be slightly more general and consider coalgebras for some functor $T : \mathsf{C} \to \mathsf{C}$. Here C will be either the category Set of sets and functions or

Stone (the category of Stone spaces, cf. 3.1.2). In both cases we have a contravariant functor

$$\mathsf{P} : \mathsf{C}^{op} \to \mathsf{BA}$$

mapping an object $X \in \mathsf{C}$ to the Boolean algebra of predicates over X. In the case $\mathsf{C} = \mathsf{Set}$, the functor P will be equal to the contravariant power set functor $\mathbf{Q} : \mathsf{Set}^{op} \to \mathsf{BA}$. In case $\mathsf{C} = \mathsf{Stone}$ we have $\mathsf{P} = \mathbb{C}\mathrm{lp} : \mathsf{Stone}^{op} \to \mathsf{BA}$, where $\mathbb{C}\mathrm{lp}$ is the functor mapping a Stone space to the Boolean algebra of its clopen subsets (cf. Section 3.1.1).

2.2.1. Remark. We will only consider the above mentioned cases $\mathsf{C} = \mathsf{Set}$ or $\mathsf{C} = \mathsf{Stone}$. In abstract terms the categorical framework can be summarized as follows:

1. C is a concrete category, i.e. there exists a faithful forgetful functor $U : \mathsf{C} \to \mathsf{Set}$.

2. C has a final object $\mathbf{1} \in \mathsf{C}$.

3. There exists a contravariant functor $\mathsf{P} : \mathsf{C}^{op} \to \mathsf{BA}$ and an injective natural transformation $\tau : \mathsf{P} \Rightarrow \mathbf{Q}U$ where $\mathbf{Q}$ denotes the contravariant power set functor.

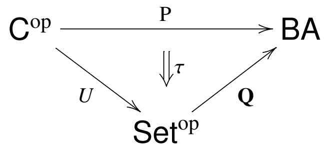

   This condition ensures in particular that P behaves on morphisms like the contravariant power set functor, i.e. $\mathsf{P}f(U) = f^{-1}[U]$ for any given $f : X \to Y \in \mathsf{C}$ and $U \in \mathsf{P}Y$.

4. There is an object $2 \in \mathsf{C}$ such that there is a natural isomorphism $\chi^{(-)} : V\mathsf{P} \Rightarrow \mathsf{C}(\_, 2)$.

Another generalization in comparison to Pattinson's earlier work is that we do not require predicate liftings to be monotone, i.e. the modalities in our logic need not be monotone. Instead they only have to satisfy the congruence rule, i.e. if ϕ and ψ are equivalent, then $\Box\phi$ and $\Box\psi$ should be also equivalent.

2.2.1 Syntax and semantics

We first state the definition of a predicate lifting for a functor $T : \mathsf{C} \to \mathsf{C}$ and then introduce the syntax and semantics of coalgebraic modal logic.

2.2.2. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor. An *n-ary predicate lifting for T* is a natural transformation

$$\lambda : V\mathsf{P}^n \Rightarrow V\mathsf{P}T.$$

Here $V : \mathsf{BA} \to \mathsf{Set}$ is the forgetful functor.

2.2.3. REMARK. The reader who is not familiar with the notion of a natural transformation between two functors T_1 and T_2 can think of a natural transformation as a family of functions $(\tau_X : T_1X \to T_2X)_{X\in\mathsf{C}}$ that is uniformly defined for all $X \in \mathsf{C}$, i.e. the definition of $\tau_{X'}$ for some $X' \in \mathsf{C}$ does not depend on special properties of X' but only on properties that all $X \in \mathsf{C}$ have in common. The formal definition can be found in the appendix (cf. Def A.1.4).

The motivation for this definition will become clear after we introduce the language of coalgebraic modal logic and its coalgebraic semantics: the fact that for every set X the lifting from $\mathrm{P}X$ to $\mathrm{P}TX$ is a function, corresponds to the fact that all the modal operators should satisfy the congruence rule

$$\frac{\vdash \phi \leftrightarrow \psi}{\vdash [\lambda]\phi \leftrightarrow [\lambda]\psi}\ (\mathrm{C})$$

and the naturality of the liftings entails invariance of coalgebraic modal logic under bisimilarity (cf. Proposition 2.2.9). The occurrence of the forgetful functor V can be explained by the fact that modal operators do in general not preserve (all the) Boolean structure.

2.2.4. DEFINITION. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor and Λ a set of predicate liftings for T. Then the *language* $\mathcal{L}(\Lambda)$ of coalgebraic modal logic is defined as follows

$$\mathcal{L}(\Lambda) \ni \phi ::= \bot \mid \phi \to \phi \mid [\lambda](\phi_1, \dots, \phi_n) \quad \text{for } \lambda \in \Lambda \quad n\text{-ary}.$$

Moreover we use the standard abbreviations $\neg\phi := \phi \to\bot$, $\phi \vee \psi := \neg\phi \to \psi$ and $\phi \wedge \psi := \neg(\neg\phi \vee \neg\psi)$.

2.2.5. DEFINITION. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ a set of predicate liftings for T, and (X, γ) a T-coalgebra. Then the semantics $[\![\phi]\!]_{(X,\gamma)} \in V\mathrm{P}X$ of a formula $\phi \in \mathcal{L}(\Lambda)$ is defined as follows

$$\begin{aligned}
[\![\bot]\!]_{(X,\gamma)} &:= \emptyset \\
[\![\phi \to \psi]\!]_{(X,\gamma)} &:= \neg[\![\phi]\!]_{(X,\gamma)} \cup [\![\psi]\!]_{(X,\gamma)} \\
\llbracket [\lambda](\phi_1, \dots, \phi_n) \rrbracket_{(X,\gamma)} &:= (V\mathrm{P}\gamma \circ \lambda_X)([\![\phi_1]\!]_{(X,\gamma)}, \dots, [\![\phi_n]\!]_{(X,\gamma)}).
\end{aligned}$$

We write

$$\begin{array}{lll}
(X, \gamma), x \models \phi & \text{if} & x \in [\![\phi]\!]_{(X,\gamma)} \text{ for } x \in X, \\
(X, \gamma) \models \phi & \text{if} & (X, \gamma), x \models \phi \text{ for all } x \in X, \text{ and} \\
\mathsf{Coalg}(T) \models \phi & \text{if} & (X, \gamma) \models \phi \text{ for all } (X, \gamma) \in \mathsf{Coalg}(T).
\end{array}$$

The theory $Th_{(X,\gamma)}(x)$ of a point x is the collection of formulas which are satisfied by that point, i.e.

$$Th_{(X,\gamma)}(x) := \{\phi \mid x \in [\![\phi]\!]_{(X,\gamma)}\}.$$

In case (X, γ) is clear from the context we drop it, i.e. we write $[\![\_]\!]$, $x \models \phi$ and $Th(x)$.

2.2.6. Remark. The interpretation of boxed formulas of the form $[\lambda]\phi$ can be explained as follows: inductively we have already defined $[\![\phi]\!] \subseteq X$ and we lift this predicate to $\lambda_X([\![\phi]\!]) \subseteq TX$. Then $[\lambda]\phi$ should be true in all points $x \in X$ with $\gamma(x) \in \lambda_X([\![\phi]\!])$.

2.2.7. Example. Let T be the functor $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$. We define two natural transformations $\lambda^{\Box}, \lambda^{\Diamond} : V\mathbf{Q} \Rightarrow V\mathbf{Q}T$ by putting for every set X

$$\begin{array}{rcl} \lambda^{\Box}_X : V\mathbf{Q}X & \to & V\mathbf{Q}\mathcal{P}X \\ U & \mapsto & \mathcal{P}(U) \end{array}$$

and

$$\begin{array}{rcl} \lambda^{\Diamond}_X : V\mathbf{Q}X & \to & V\mathbf{Q}\mathcal{P}X \\ U & \mapsto & \{V \subseteq X \mid V \cap U \neq \emptyset\}. \end{array}$$

Then the coalgebraic semantics of boxed formulas $[\lambda^{\Box}]\phi$ and $[\lambda^{\Diamond}]\psi$ on a given $\mathcal{P}$-coalgebra $(X, \gamma : X \to \mathcal{P}X)$ is calculated as follows:

$$\begin{array}{rcl} [\![[\lambda^{\Box}]\phi]\!]_{(X,\gamma)} & = & \gamma^{-1}\left[\lambda^{\Box}_X([\![\phi]\!]_{(X,\gamma)})\right] \\ & = & \gamma^{-1}\left[\mathcal{P}([\![\phi]\!]_{(X,\gamma)})\right] \\ & = & \{x \in X \mid \gamma(x) \subseteq [\![\phi]\!]_{(X,\gamma)}\} \\ [\![[\lambda^{\Diamond}]\phi]\!]_{(X,\gamma)} & = & \gamma^{-1}\left[\lambda^{\Diamond}_X([\![\phi]\!]_{(X,\gamma)})\right] \\ & = & \gamma^{-1}\left[\{U \subseteq \mathcal{P}X \mid U \cap [\![\phi]\!]_{(X,\gamma)} \neq \emptyset\}\right] \\ & = & \{x \in X \mid \gamma(x) \cap [\![\phi]\!]_{(X,\gamma)} \neq \emptyset\} \end{array}$$

If we view the $\mathcal{P}$-coalgebra (X, γ) as a Kripke frame (X, R_γ) with

$$x_1 R_\gamma x_2 \;:\Leftrightarrow\; x_2 \in \gamma(x_1),$$

we see that the coalgebraic semantics of the language given by $\Lambda = \{\lambda^{\Box}, \lambda^{\Diamond}\}$ coincides with the ordinary Kripke semantics.

The naturality of the predicate liftings ensures that the semantics of a formula is preserved under T-coalgebra morphisms.

2.2.8. Lemma. *Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ be a set of predicate liftings for T and let $\phi \in \mathcal{L}(\Lambda)$ be a formula. Furthermore let (X, γ) and (Y, δ) be T-coalgebras and let $f : (X, \gamma) \to (Y, \delta)$ be a T-coalgebra morphism. Then*

$$(X, \gamma), x \models \phi \quad \textit{iff} \quad (Y, \delta), f(x) \models \phi.$$

Proof. We prove the following claim by induction on the structure of ϕ

$$[\![\phi]\!]_{(X,\gamma)} = f^{-1}[[\![\phi]\!]_{(Y,\delta)}]. \tag{2.1}$$

The Boolean cases are trivial. Suppose that $\phi = [\lambda](\psi_1, \ldots, \psi_n)$ for some n-ary $\lambda \in \Lambda$. To prove (2.1) we use the following diagram

$$\begin{array}{ccccc}
V\mathrm{P}^n Y & \xrightarrow{\lambda_Y} & V\mathrm{P}TY & \xrightarrow{V\mathrm{P}\delta} & V\mathrm{P}Y \\
{\scriptstyle V\mathrm{P}^n f}\downarrow & & \downarrow{\scriptstyle V\mathrm{P}Tf} & & \downarrow{\scriptstyle V\mathrm{P}f} \\
V\mathrm{P}^n X & \xrightarrow[\lambda_X]{} & V\mathrm{P}TX & \xrightarrow[V\mathrm{P}\gamma]{} & V\mathrm{P}X
\end{array}$$

The left half of the diagram commutes because of the naturality of λ and the right half commutes because of f being a T-coalgebra morphism. We are now ready to finish the proof:

$$\begin{array}{rcl}
\llbracket\phi\rrbracket_{(X,\gamma)} & = & \llbracket [\lambda](\psi_1, \ldots, \psi_n)\rrbracket_{(X,\gamma)} \\
 & = & V\mathrm{P}\gamma(\lambda_X(\llbracket\phi_1\rrbracket_{(X,\gamma)}, \ldots, \llbracket\phi_n\rrbracket_{(X,\gamma)})) \\
 & \overset{\text{I.H.}}{=} & V\mathrm{P}\gamma(\lambda_X(V\mathrm{P}^n f(\llbracket\phi_1\rrbracket_{(Y,\delta)}, \ldots, \llbracket\phi_n\rrbracket_{(Y,\delta)}))) \\
 & \overset{\text{diagram}}{=} & V\mathrm{P}f(V\mathrm{P}\delta(\lambda_Y(\llbracket\phi_1\rrbracket_{(Y,\delta)}, \ldots, \llbracket\phi_n\rrbracket_{(Y,\delta)}))) \\
 & = & f^{-1}[\llbracket\phi\rrbracket_{(Y,\delta)}]
\end{array}$$

The claim of the proposition now follows easily from (2.1). QED

An immediate consequence of the proposition is the invariance of the semantics of coalgebraic logic under bisimilarity.

2.2.9. Proposition. *Let T be an endofunctor on C and Λ a set of predicate liftings for T. Furthermore let $(\mathbb{X}, x)$ and $(\mathbb{Y}, y)$ be rooted T-coalgebras. Then $(\mathbb{X}, x) \leftrightarroweq_T (\mathbb{Y}, y)$ implies $Th(x) = Th(y)$.*

Proof. Suppose $(\mathbb{X}, x)$ and $(\mathbb{Y}, y)$ are T-bisimilar (cf. Def A.3.3). Then there is a rooted T-coalgebra $(\mathbb{Z}, z)$ and $\mathsf{Coalg}(T)$-morphsims $f_1 : (\mathbb{X}, x) \to (\mathbb{Z}, z)$ and $f_2 : (\mathbb{Y}, y) \to (\mathbb{Z}, z)$ such that $f_1(x) = f_2(y) = z$. By Lemma 2.2.8 we therefore get $Th(x) = Th(z) = Th(y)$. QED

The converse of Proposition 2.2.9 is not always true. This leads to the definition of the notion of expressiveness.

2.2.10. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor and Λ a set of predicate liftings for T. Then the corresponding language $\mathcal{L}(\Lambda)$ is called *expressive* if for all $(X, \gamma), (Y, \delta) \in \mathsf{Coalg}(T)$ and $x \in X, y \in Y$

$$Th(x) = Th(y) \quad \Longrightarrow \quad (\mathbb{X}, x) \leftrightarroweq_T (\mathbb{Y}, y).$$

2.2.11. Remark. This property of a modal language is also called the *Hennessy-Milner property*.

Before we turn to the definition of the derivability relation of coalgebraic modal logic we take a look at an example.

2.2.12. Example. Let T be the power set functor $\mathcal{P}$ and let $\lambda^{\Box}, \lambda^{\Diamond}$ be defined as in Example 2.2.7. It is well known that $\mathcal{L}(\{\lambda^{\Box}, \lambda^{\Diamond}\})$ is not expressive, for a counterexample to expressivity see [BdV01, Example 2.23]. Now consider the finite power set functor $T = \mathcal{P}_\omega : \mathsf{Set} \to \mathsf{Set}$ and the same predicate lifting as in Example 2.2.7 restricted to finite sets, i.e.

$$\begin{aligned} \lambda_X^{\Box_\omega}(U) &:= \mathcal{P}_\omega(U) \\ \lambda_X^{\Diamond_\omega}(U) &:= \{V \subseteq_\omega X \mid V \cap U \neq \emptyset\}. \end{aligned}$$

Then $\mathcal{L}(\{\lambda^{\Box_\omega}, \lambda^{\Diamond_\omega}\})$ is expressive.

2.2.2 Derivability and the logic $L(\Lambda, \mathtt{Ax})$

As we will see there are rules and axioms which hold for any coalgebraic modal logic. In addition, however, one needs axioms which depend on the functor under consideration. As the language $\mathcal{L}(\Lambda)$ does not contain propositional variables, the axioms of coalgebraic modal logic are formulated as axiom schemes which contain certain meta-variables. Furthermore the axioms are rather restricted as they are only built up from formulas of modal depth not bigger than 1. In particular this means that well-known axioms from modal logic such as the transitivity axiom $\Box\phi \to \Box\Box\phi$ are not allowed.

This restriction is to be expected as we are considering logics for coalgebras for a functor and not for a comonad. Intuitively we can only axiomatize properties of lifted predicates that depend only on T and λ, i.e. properties of predicates which have been lifted by some $\lambda \in \Lambda$ once from X to TX. Properties of predicates which are lifted twice (or more) will usually depend on properties of the coalgebra map of a given T-coalgebra. This can be easily seen by spelling out the definition of the interpretation of the formula $[\lambda][\lambda]\bot$ on a T-coalgebra (X, γ). Axioms of depth bigger than 1 will therefore in general not be valid on all T-coalgebras for a certain functor T, but only on those T-coalgebras whose coalgebra maps fulfill certain extra conditions. A solution for this problem is to consider T-coalgebras for a comonad. In this thesis we will restrict ourselves to the case of coalgebras for a functor.

2.2.13. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ a set of predicate liftings for T and $\mathcal{X}$ a set (of meta-variables). An *axiom* is a pair (ϕ, ψ) with

$$\phi, \psi \in \mathbf{T}_{\Sigma_{\mathsf{BA}}}(\{[\lambda](\phi_1, \dots, \phi_n) \mid \lambda \in \Lambda \text{ and } \phi_i \in \mathbf{T}_{\Sigma_{\mathsf{BA}}}(\mathcal{X})\}),$$

where $\mathbf{T}_{\Sigma_{\mathsf{BA}}}(\_)$ denotes the term algebra for the Boolean signature defined as in Definition B.1.3. We will write axioms (ϕ, ψ) also as equations $\phi = \psi$.

2.2.14. Remark. The set $\mathbf{T}_{\Sigma_{\mathsf{BA}}}(\mathcal{X})$ can be also seen as the set of all Boolean formulas over $\mathcal{X}$. Note that we will not make an explicit distinction between Boolean formulas and Boolean terms.

The substitution instances of these "meta"-axioms are the axioms of the derivability relation of coalgebraic modal logic. We first formalize what we mean by "substitution instance" and then state the definition of the derivability relation.

2.2.15. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ a set of predicate liftings for T and $\mathcal{X}$ a set. A *substitution* is a function $\sigma : \mathcal{X} \to \mathcal{L}(\Lambda)$. Given a term $\phi \in \mathbf{T}_{\Sigma_{\mathsf{BA}}}(\mathcal{X})$ its *substitution instance* $\sigma(\phi) \in \mathcal{L}(\Lambda)$ is defined as the image of ϕ under the inductive extension of σ to terms.

2.2.16. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ a set of predicate liftings for T and $\mathtt{Ax}$ a set of axioms. We say that ϕ is modally derivable from $\mathtt{Ax}$ ($\mathtt{Ax} \vdash \phi$), if ϕ is contained in the least set Φ of formulas which

- contains $\sigma(\phi) \leftrightarrow \sigma(\psi)$ whenever σ is a substitution and $(\phi, \psi) \in \mathtt{Ax}$,
- is closed under propositional entailment,
- is for any n-ary $\lambda \in \Lambda$ closed under the congruence rule

$$\frac{\phi_1 \leftrightarrow \psi_1 \qquad \dots \qquad \phi_n \leftrightarrow \psi_n}{[\lambda](\phi_1, \dots, \phi_n) \leftrightarrow [\lambda](\psi_1, \dots, \psi_n)}\ (C)$$

After having defined the language $\mathcal{L}(\Lambda)$ of coalgebraic modal logic that corresponds to some set of predicate liftings and a notion of derivability, we can now give a formal definition of the corresponding coalgebraic modal logic. We identify the logic $L(\Lambda, \mathtt{Ax})$ with the set of all formulas that are derivable from a given set of of axioms.

2.2.17. Definition. Given a functor $T : \mathsf{C} \to \mathsf{C}$, a set Λ of predicate liftings for T and a set $\mathtt{Ax}$ of axioms. Then $L(\Lambda, \mathtt{Ax})$, the *logic given by* Λ *and* $\mathtt{Ax}$, is defined as the set of all formulas that are derivable from $\mathtt{Ax}$, i.e.

$$L(\Lambda, \mathtt{Ax}) := \{\phi \in \mathcal{L}(\Lambda) \mid \mathtt{Ax} \vdash \phi\}.$$

2.2.18. Example. Consider again the power set functor and let $\lambda^{\Box}, \lambda^{\Diamond}$ be defined as in Example 2.2.7. Furthermore let $\mathtt{Ax}$ be the set consisting of the following two axioms:

$$\begin{aligned} [\lambda^{\Box}]\top &= \top, \\ [\lambda^{\Box}](x_1 \wedge x_2) &= [\lambda^{\Box}]x_1 \wedge [\lambda^{\Box}]x_2. \end{aligned}$$

Then $L(\{\lambda^{\Box}, \lambda^{\Diamond}\}, \mathtt{Ax})$ corresponds to the basic normal modal logic **K** without propositional variables.

In [Pat03a] sufficient conditions are provided for when the logic $L(\Lambda, \mathtt{Ax})$ is sound and complete. We are now only defining these properties and continue in Chapter 4 with a categorical analysis of coalgebraic completeness proofs.

2.2.19. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, Λ a set of predicate liftings for T and $\mathtt{Ax}$ a set of axioms. Then $L(\Lambda, \mathtt{Ax})$ is called

- *sound* if for all $\phi \in \mathcal{L}(\Lambda)$

$$\mathtt{Ax} \vdash \phi \quad \Longrightarrow \quad \mathsf{Coalg}(T) \models \phi,$$

- *complete* if for all $\phi \in \mathcal{L}(\Lambda)$

$$\mathsf{Coalg}(T) \models \phi \quad \Longrightarrow \quad \mathtt{Ax} \vdash \phi.$$

The reader who is familiar with modal logic might find the restriction to axioms of modal depth 1 very strong. It should be stressed, however, that the focus of research in coalgebraic modal logic differs, at least up to now, from the one in modal logic: in modal logic one studies various logics over structures of the same type (i.e. over Kripke frames) whereas coalgebraic research is concerned with finding the basic modal logic for different types of structures. It has been proven in various places in the literature that the notion of a coalgebraic modal logic, given by a set of predicate liftings and a set of axioms in the sense of Def. 2.2.13, covers interesting examples of various types of logics, other than ordinary modal logic, see e.g. [CP04] for logics for the specification of different types probabilistic systems and [HK04] for a coalgebraic treatment of monotone modal logic.

2.2.3 The language of all liftings

Coalgebraic modal logic provides an abstract framework for studying logics for T-coalgebras whose languages are given by a set of predicate liftings. One problem is, however, that there seems to be no canonical choice for the collection of predicate liftings for a given functor T. The difficulty of finding the right collection of liftings can be avoided by considering the set of *all* predicate liftings for T. The set of all liftings for a given functor can be computed using the Yoneda Lemma (cf. Theorem A.1.6). This was observed first by Schröder in [Sch05].

2.2.20. Definition. For each $Y \in \mathsf{C}$ there is an isomorphism

$$\chi^Y_{(\_)} : V\mathrm{P}Y \to \mathsf{C}(Y, 2),$$

mapping a predicate $X \in \mathrm{P}Y$ to its so-called *characteristic function* $\chi^Y_X : Y \to 2$. Here 2 denotes the two element object in C, i.e. depending on C either the two element set or the two element Stone space.

2.2.21. Remark. The isomorphism $\chi^Y_{(\_)} : V\mathrm{P}Y \to \mathsf{C}(Y, 2)$ maps a predicate $X \in \mathrm{P}Y$ to the function

$$\begin{array}{rcl} \chi^Y_X : Y & \to & 2 \\ y & \mapsto & \begin{cases} 1 & \text{if } y \in X \\ 0 & \text{otherwise.} \end{cases} \end{array}$$

In the case $\mathsf{C} = \mathsf{Set}$ it is obvious that this defines an isomorphism $\chi^Y_{(\_)} : V\mathbf{Q}Y \to \mathsf{Set}(Y, 2)$. For $\mathsf{C} = \mathsf{Stone}$ one has to observe that for every $\mathbb{Y} \in \mathsf{Stone}$ a map $f : \mathbb{Y} \to 2$ is continuous iff $f^{-1}(0)$ and $f^{-1}(1)$ are clopen subsets of $\mathbb{Y}$. Therefore there is a one-to-one correspondence between clopen subsets of $\mathbb{Y}$ and continuous morphisms $f \in \mathsf{Stone}(\mathbb{Y}, 2)$.

2.2.22. Proposition. *There is a 1-1 correspondence*

$$\{\lambda \mid \lambda : V\mathrm{P}^n \Rightarrow V\mathrm{P}T\} \;\cong\; V\mathrm{P}T(2^n)$$

given by $U \in \mathrm{P}T(2^n) \mapsto \lambda$ *where*

$$\lambda_Y : (P_1, \ldots, P_n) \in (\mathrm{P}Y)^n \mapsto \{t \in TY \mid \chi^{2^n}_U \circ T\langle \chi^Y_{P_1}, \ldots, \chi^Y_{P_n}\rangle(t) = 1\}.$$

Proof. The crucial observation here is that the family $\{\chi^Y_{(\_)}\}_{Y \in \mathsf{C}}$ forms a natural isomorphism

$$\chi^{\cdot}_{(\_)} : V\mathrm{P} \Rightarrow \mathsf{C}(\_, 2).$$

Therefore n-ary predicate liftings $\lambda : V\mathrm{P}^n \Rightarrow V\mathrm{P}T$ are in one-to-one correspondence with natural transformations $\lambda' : \mathsf{C}(\_, 2^n) \Rightarrow \mathsf{C}(T\_, 2)$. Instantiating the Yoneda Lemma (cf. Theorem A.1.6 in the appendix) with $X := 2^n$ and $S := \mathsf{C}(T\_, 2)$ give us that there is an isomorphism

$$\Theta_{\mathsf{C}(T\_,2),2^n} : \mathsf{Cat}(\mathsf{C}(\_, 2^n), \mathsf{C}(T\_, 2)) \to \mathsf{C}(T2^n, 2).$$

Hence natural transformations $\lambda' : \mathsf{C}(\_, 2^n) \Rightarrow \mathsf{C}(T\_, 2)$ and elements of $\mathsf{C}(T2^n, 2)$ are in one-to-one correspondence. All in all this gives us together with the isomorphism $\mathsf{C}(T2^n, 2) \cong \mathrm{P}T2^n$ a one-to-one correspondence between n-ary predicate liftings λ and predicates over $T2^n$. That this one-to-one correspondence is computed as stated in the proposition is not difficult to check. One only has to spell out how the isomorphism $\Theta_{\mathsf{C}(T\_,2),2^n}$ works, which is given by the Yoneda Lemma. QED

In many cases the language corresponding to the set of all (finitary) predicate liftings is expressive. This is the content of the following fact which was proven by Schröder in [Sch05].

2.2.23. Fact. [Sch05, Corollary 38] Let $T : \mathsf{Set} \to \mathsf{Set}$ be an ω-accessible functor and Λ the set of finitary predicate liftings for T. Then the language $\mathcal{L}(\Lambda)$ is expressive.

2.2.24. Remark. In fact the statement in *loc.cit.* is more general: instead for ω one can prove a similar result for an arbitrary regular cardinal κ. The language will be, however, not finitary anymore: one has to allow predicate liftings with infinite arity α ($\alpha < \kappa$ an ordinal number) and infinite conjunctions of $< \kappa$ formulas.

2.3 Coalgebraic fixed-point logic

Next we discuss a third approach for defining a modal language for reasoning about coalgebras: coalgebraic fixed-point logic.

Coalgebraic modal logic has an obvious shortcoming [1]: we can only define a well-behaved such logic for a given functor $T : \mathsf{Set} \to \mathsf{Set}$ if we can find the right set of predicate liftings. There is no canonical choice for the right coalgebraic modal logic for T-coalgebras besides the logic given by *all* predicate liftings (cf. Section 2.2.3).

The so-called coalgebraic logic as proposed by Moss in [Mos99] is more canonical. The language of coalgebraic logic contains, independently of the given functor, only one operator ∇ and the semantics of this logical operator is defined for all functors in a uniform way. The drawback of this approach is that it yields a language that is less standard than the multi-modal language of logics given by predicate liftings. Moreover coalgebraic modal logics can be defined for arbitrary functors on Set whereas in coalgebraic logic one requires the functor to be *weak pullback preserving* (cf. Def. A.2.4).

We will now state Venema's definition of a finitary version of coalgebraic logic and its extension by a smallest and greatest fixed-point operator as presented in [Ven04].

2.3.1 Finitary coalgebraic logic

Throughout this section we assume that $T : \mathsf{Set} \to \mathsf{Set}$ is a standard and weak pullback preserving functor (definitions can be found in Appendix A). As remarked above the requirement that the functor is weak pullback preserving seems to be a real restriction. In Chapter 5 we will see that we could also define a coalgebraic logic for non-standard functors.

2.3.1. Definition. Let X be a set. We say $\pi \in TX$ has *finite T-base over X* if there is a finite set $Q \subseteq X$ such that $\pi \in TQ$ and define

$$T_\omega(X) := \{\pi \in TX \mid \pi \text{ has finite } T\text{-base over } X\}.$$

2.3.2. Remark. We chose the notation $T_\omega(X)$ to indicate that we can define a functor

$$\begin{array}{rcl} T_\omega : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & T_\omega X := T_\omega(X) \\ f & \mapsto & Tf_{\restriction T_\omega(X)}. \end{array}$$

It is not difficult to check that T_ω is well-defined if T is standard. If T is the power set functor then T_ω is the finite power set functor.

[1] under the assumption that the functor T under consideration determines the appropriate logic for reasoning about T-coalgebras

2.3.3. DEFINITION. The *language* of (finitary) coalgebraic logic $\mathcal{L}^T$ is defined inductively as follows:

$$\begin{aligned} \mathcal{L}_0^T \ni \phi &::= \bot \mid \top \mid \phi \wedge \phi \mid \phi \vee \phi \\ \mathcal{L}_{i+1}^T \ni \phi &::= \psi \in \mathcal{L}_i^T \mid \phi \wedge \phi \mid \phi \vee \phi \mid \nabla\pi,\ \pi \in T_\omega(\mathcal{L}_i^T) \\ \mathcal{L}^T &:= \bigcup_{i \in \mathbb{N}} \mathcal{L}_i^T \end{aligned}$$

The *depth* of a formula $\phi \in \mathcal{L}^T$ is defined as the smallest natural number i_ϕ such that $\phi \in \mathcal{L}_{i_\phi}^T$.

2.3.4. REMARK. From the definition it is clear that $\mathcal{L}^T$ is a set (in contrast to Moss' original language which consisted of a proper class of formulas).

The reader might wonder why negation is not included in the language $\mathcal{L}^T$. The reason for this is that we want to maintain the one-to-one correspondence between formulas of coalgebraic (fixed-point) logic and coalgebra automata which has been established in [Ven04] and which we will use in Chapter 5.

It is an open question whether we can always construct for a given coalgebra automaton an automaton that accepts precisely the complement language.

A positive answer to this question would mean that adding negation to the language would be *redundant*: given a formula ϕ and its corresponding coalgebra automaton $\mathbb{A}_\phi$ we could construct the automaton $\overline{\mathbb{A}_\phi}$ that accepts the complement language of $\mathbb{A}_\phi$. This automaton $\overline{\mathbb{A}_\phi}$, in turn, would correspond to some formula ψ. This (itself negation-free) formula ψ would express the negation of ϕ. For example in the case that T is the power set functor, it is known that negation is redundant.

A negative answer to the above question, on the other hand, would imply that we would have to adjust the notion of a coalgebra automaton if we wanted to extend the language and to keep the correspondence with coalgebra automata at the same time. It is not obvious how these "extended" coalgebra automata should look like.

The difference with Moss's original definition is that the syntax only contains finite conjunctions and, in addition to that, finite disjunctions. Moreover formulas of the form $\nabla\pi$ contain only elements $\pi \in T\mathcal{L}^T$ which have a finite T-base. We will now demonstrate that these conditions ensure that every formula ϕ has a finite construction tree and therefore a finite set of subformulas.

2.3.5. DEFINITION. Given $X \in \mathsf{Set}$ and $\pi \in T_\omega(X)$ we define the *base* of π as follows

$$Base(\pi) := \bigcap \{Q \subseteq_\omega X \mid \pi \in TQ\}.$$

2.3.6. LEMMA. *For all $X \in \mathsf{Set}$ and $\pi \in T_\omega(X)$ we have $\pi \in TBase(\pi)$, hence $Base(\pi)$ is the smallest finite subset of X with this property.*

Proof. It is easy to see that there is a finite family $\{Q_1, \ldots, Q_n\}$ of finite subsets of X such that $\pi \in TQ_i$ for all $i \in \{1, \ldots, n\}$ and

$$Base(\pi) = \bigcap_{i=1}^{n} Q_i.$$

Standard functors that are weak pullback preserving preserve finite intersections (cf. Fact A.2.13), hence

$$TBase(\pi) = \bigcap_{i=1}^{n} TQ_i.$$

Because $\pi \in TQ_i$ for all $i \in \{1, \ldots, n\}$ we can conclude that $\pi \in TBase(\pi)$. QED

2.3.7. Definition. Let $\phi \in \mathcal{L}^T$. Then the *construction tree* $\mathrm{CTree}(\phi)$ of ϕ is defined as follows

$$\mathrm{CTree}(\phi) \quad := \quad \phi \qquad \text{for } \phi \in \{\bot, \top\}$$

$$\mathrm{CTree}(\phi_1 \clubsuit \phi_2) \quad := \quad \begin{array}{c} \phi_1 \clubsuit \phi_2 \\ \swarrow \qquad \searrow \\ \mathrm{CTree}(\phi_1) \qquad \mathrm{CTree}(\phi_2) \end{array} \qquad \text{for } \clubsuit \in \{\vee, \wedge\}$$

$$\mathrm{CTree}(\nabla\pi) \quad := \quad \begin{array}{c} \nabla\pi \\ \swarrow \qquad \searrow \\ \mathrm{CTree}(\phi_1) \quad \cdots \quad \mathrm{CTree}(\phi_n) \end{array} \qquad \text{for } Base(\pi) = \{\phi_1, \ldots, \phi_n\}$$

The collection $\mathrm{Sub}(\phi)$ of *subformulas* of ϕ is defined as the set of formulas ϕ' which appear as labels in $\mathrm{CTree}(\phi)$.

The following proposition justifies that we call $\mathcal{L}^T$ a finitary language.

2.3.8. Proposition. *The construction tree of every $\phi \in \mathcal{L}^T$ is finite and therefore every formula ϕ has only finitely many subformulas.*

Proof. The proposition can be proven by an easy induction on the depth i_ϕ of ϕ. For the case $\phi = \nabla\pi$ one has to observe that $i_\psi < i_\phi$ for all $\psi \in Base(\pi)$. QED

The semantics of the ∇-operator is now defined as follows: instead of lifting predicates over some set X to predicates over TX we lift the satisfiability relation between X and $\mathcal{L}^T$ to a relation between TX and $T\mathcal{L}^T$. Then

$$x \models \nabla\pi \;:\Leftrightarrow\; \gamma(x)\overline{T}(\models)\pi. \tag{2.2}$$

Here $\overline{T} : \mathsf{Rel} \to \mathsf{Rel}$ denotes the (unique) lifting of T to the category Rel of sets and relations (cf. Appendix A.2). We postpone the formal definition of the semantics and first extend the language with fixed-point operators.

2.3.9. Remark. Note that the definition of the semantics of a formula of the form $\nabla\pi$ in (2.2) is not circular: by Lemma A.2.14 we have

$$\gamma(x)\overline{T}(\models)\pi \quad \text{iff} \quad \gamma(x)\overline{T}(\models_{\restriction TX \times Base(\pi)})\pi$$

and as the depth of formulas in $Base(\pi)$ is strictly smaller than the depth of $\nabla\pi$ we can inductively assume that $\models_{\restriction TX \times Base(\pi)}$ has been already defined.

2.3.2 Adding fixed-points

In order to be able to extend the language of coalgebraic logic with fixed-point-operators it is also necessary to add variables.

2.3.10. Definition. Let Φ be a set of variables. The *language* $\mu\mathcal{L}^T(\Phi)$ of coalgebraic fixed-point logic with variables in Φ is defined inductively as follows:

$$\begin{array}{rcl} \mu\mathcal{L}^T_0(\Phi) \ni \phi & ::= & \bot \mid \top \mid p \in \Phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \mu p.\phi,\ p \in \Phi \mid \nu p.\phi,\ p \in \Phi \\ \mu\mathcal{L}^T_{i+1}(\Phi) \ni \phi & ::= & \psi \in \mu\mathcal{L}^T_i \mid \phi \wedge \phi \mid \phi \vee \phi \mid \mu p.\phi,\ p \in \Phi \mid \nu p.\phi,\ p \in \Phi \mid \\ & & \mid \nabla\pi,\ \pi \in T_\omega(\mu\mathcal{L}^T_i(\Phi)) \\ \mu\mathcal{L}^T(\Phi) & := & \bigcup_{i \in \mathbb{N}} \mu\mathcal{L}^T_i(\Phi) \end{array}$$

The notion of a construction tree of a formula from Definition 2.3.7 can be easily extended to the case of fixed-point logic.

2.3.11. Definition. Let Φ be a set of variables and ϕ a formula in $\mu\mathcal{L}^T(\Phi)$. Then the *construction tree* CTree(ϕ) of ϕ is defined as in Definition 2.3.7 using the following additional clauses for variables and the fixed-point operators:

$$\begin{array}{rcll} \mathrm{CTree}(p) & := & p & \text{for } p \in \Phi \\ \mathrm{CTree}(\eta p.\phi) & := & \begin{array}{c} \eta p.\phi \\ \downarrow \\ \mathrm{CTree}(\phi) \end{array} & \text{for } \eta \in \{\mu, \nu\} \end{array}$$

Again we define Sub(ϕ) to be the set of formulas which occur as labels in CTree(ϕ).

Regarding the syntax of coalgebraic fixed-point logic we have to introduce some terminology.

2.3.12. Definition. Let Φ be a set of variables and $\phi \in \mu\mathcal{L}^T(\Phi)$. An *occurrence* of a variable p in ϕ is a leaf in CTree(ϕ) labeled with p. An occurrence of p in ϕ is called *bound* if it has an ancestor in CTree(ϕ) labeled with a formula $\eta p.\psi$ for $\eta \in \{\mu, \nu\}$. Otherwise the occurrence is called *free*.

We denote by $\mathrm{BVar}(\phi)$ and $\mathrm{FVar}(\phi)$ the set of bound variables of ϕ and the set of free variable of ϕ respectively, i.e. the set of all variables p that have a bound occurrence and the set of all variables that have a free occurrence in ϕ respectively.

Finally we denote by $\mathrm{Var}(p)$ the set of all variables occurring in ϕ, i.e. $\mathrm{Var}(\phi) = \mathrm{BVar}(\phi) \cup \mathrm{FVar}(\phi)$.

2.3.13. Definition. Let Φ be a set of variables. A formula $\phi \in \mu\mathcal{L}^T(\Phi)$ is called *clean* if $\mathrm{BVar}(\phi) \cap \mathrm{FVar}(\phi) = \emptyset$ and if for every $p \in \mathrm{BVar}(\phi)$ there exists a unique formula $\eta_x x.\psi_x \in \mathrm{Sub}(\phi), \eta_x \in \{\mu, \nu\}$ such that $p \in \mathrm{Sub}(\eta_x x.\psi_x)$. Furthermore we call ϕ a *closed formula* if $\mathrm{FVar}(\phi) = \emptyset$. .

The definition of the semantics of coalgebraic fixed-point logic uses so-called *valuations*, i.e. functions assigning to each variable a subset of the model.

2.3.14. Definition. Let Φ be a set of variables. A Φ-*valuation on* X is a function $V : \Phi \to \mathcal{P}(X)$. A triple (X, γ, V) where $(X, \gamma) \in \mathsf{Coalg}(T)$ and V is a Φ-valuation is called Φ-*model on* (X, γ).

Given a Φ-valuation on X, $p \in \Phi$ and $U \subseteq X$ we write $V[p \mapsto U]$ for the Φ-valuation on X defined as

$$V[x \mapsto U](q) := \begin{cases} U & \text{if } q = p \\ V(q) & \text{otherwise.} \end{cases}$$

We are now prepared to state the definition of the semantics of coalgebraic fixed-point logic.

2.3.15. Definition. Let Φ be a set of variables, V a Φ-valuation and $\mathbb{X} = (X, \gamma)$ be a T-coalgebra. We inductively define a relation $\models^V \subseteq X \times \mu\mathcal{L}^T(\Phi)$ with the intended meaning that $(x, \phi) \in \models^V$ if ϕ is satisfied at $x \in X$ under valuation V. In this case we also write $(X, \gamma), x \models^V \phi$. Furthermore we define $[\![\phi]\!]_{\mathbb{X},V} := \{x \in X \mid (X, \gamma), x \models^V \phi\}$.
The inductive definition of $\models^V$ is as follows:

$$\begin{array}{llll}
\mathbb{X}, x & \not\models^V & \bot, & \\
\mathbb{X}, x & \models^V & \top, & \\
\mathbb{X}, x & \models^V & p \in \Phi & \text{if } \; x \in V(p), \\
\mathbb{X}, x & \models^V & \phi_1 \wedge \phi_2 & \text{if } \; \mathbb{X}, x \models^V \phi_1 \text{ and } \mathbb{X}, x \models^V \phi_2, \\
\mathbb{X}, x & \models^V & \phi_1 \vee \phi_2 & \text{if } \; \mathbb{X}, x \models^V \phi_1 \text{ or } \mathbb{X}, x \models^V \phi_2, \\
\mathbb{X}, x & \models^V & \mu p.\phi & \text{if } \; x \in \bigcap\{U \subseteq X \mid [\![\phi]\!]_{\mathbb{X},V[p\mapsto U]} \subseteq U\}, \\
\mathbb{X}, x & \models^V & \nu p.\phi & \text{if } \; x \in \bigcup\{U \subseteq X \mid U \subseteq [\![\phi]\!]_{\mathbb{X},V[p\mapsto U]}\}, \\
\mathbb{X}, x & \models^V & \nabla\pi & \text{if } \; (\gamma(x), \pi) \in \overline{T}(\models^V\restriction_{Base(\pi)}),
\end{array}$$

In case for some $\phi \in \mu\mathcal{L}^T(\Phi)$ we have $\mathbb{X}, x \models^V \phi$ for all valuations V, we drop the superscript V and write $\mathbb{X}, x \models \phi$.

2.3.16. Remark. If $\phi \in \mu\mathcal{L}^T(\Phi)$ is a closed formula then $\mathbb{X}, x \models^V \phi$ for some valuation V iff $\mathbb{X}, x \models \phi$, i.e. the semantics of a closed formula does not depend on the valuation.

The semantics for the μ- and the ν-operator is the well-known smallest and largest fixed-point semantics. In order to see this one has first to observe that for an arbitrary formula $\phi \in \mu\mathcal{L}^T(\Phi)$ and an arbitrary T-coalgebra $\mathbb{X} = (X, \gamma)$, the operator

$$\begin{array}{rcl} [\![\phi]\!]_{\mathbb{X},V[p\mapsto\_]} : \mathcal{P}X & \to & \mathcal{P}X \\ U & \mapsto & [\![\phi]\!]_{\mathbb{X},V[p\mapsto U]} \end{array}$$

is monotone. Therefore we can apply the Knaster-Tarski fixed-point theorem to compute the least and the largest fixed-point of $[\![\phi]\!]_{\mathbb{X},V[p\mapsto\_]}$ as the least pre-fixed point and the largest post-fixed point of $[\![\phi]\!]_{\mathbb{X},V[p\mapsto\_]}$ respectively. For more details we refer the reader to [AN01].

2.4 Conclusion

Common features

We presented the three main approaches for assigning a finitary modal language to a given functor $T : \mathsf{Set} \to \mathsf{Set}$. Properties that all of these approaches share, are:

- formulas specify properties of coalgebra states, i.e. the interpretation of a formula ϕ on coalgebra (X, γ) corresponds to a subset of X,
- the interpretation of a formula is invariant under bisimulations, and
- finitary languages have restricted expressivity.

The last item of our list is certainly a drawback of all these languages. Unfortunately restrictions here seem to be unavoidable: Goldblatt showed in [Gol05] that a functor T admits a (possibly infinitary) expressive language which has a set of formulas iff the final T-coalgebra exists. This suggests that a similar result could be proven, showing that logics with finitary syntax can be expressive only if the final T-coalgebra exists and is not "too big".

References

The research on inductively defined coalgebraic modal logics was initiated by Kurz ([Kur01c]), who in fact did not consider inductively defined functors but polynomial functors of a certain shape, and Rößiger ([Röß01]). Their work was inspired by earlier work of Jacobs ([Jac96]) and Reichel ([Rei95]) on the question of how to use coalgebras for modeling certain features of object-oriented programming languages. Rößiger also showed how to construct the final T-coalgebra for a given polynomial functor logically, i.e. using some variant of the canonical model construction from modal logic. Furthermore he proved completeness of these logics with respect to the coalgebraic semantics. These results were extended to Kripke polynomial functors by Rößiger in [Röß00], Jacobs in [Jac01] and by the work presented in the next chapter of this thesis, which is based on the paper [KKV04]. Another aspect of research on inductively defined logics for polynomial coalgebras is the work by Goldblatt on model theory for polynomial coalgebras. In the papers [Gol01, Gol03b] he obtains a Co-Birkhoff theorem for these logics and generalizations of notions from model theory of modal logic such as ultrapower and ultrafilter extension.

The abstract formulation of a predicate lifting for a functor goes back to Pattinson's work in [Pat03a, Pat04], where he also gives sufficient criteria for a coalgebraic modal logic to be sound, complete and expressive. The polyadic version of it that we are considering was first used by Schroeder in [Sch05]. Based on results from [KKP04] we are going to define an algebraic semantics for these logics in Chapter 4.

Moss' coalgebraic logic ([Mos99]) is probably the most general approach towards coalgebraic modal logic. The somewhat unusual syntax of the modal language he is considering has lead many people to the conclusion, that Moss' approach is too abstract for being useful in applications. The connection to automata theory, which was established by Venema in [Ven04] brought new attention to Moss' logic and showed that formulas of this logic have a natural automata-theoretic interpretation. We will return to this issue in Chapter 5.

It is also possible to combine the different approaches as was pointed out by Cîrstea in [Cîr04]. She defines the notions of a language constructor and of a T-semantics of a language constructor. All the logics that we described in this chapter correspond to such language constructors and their T-coalgebraic semantics correspond to their T-semantics. The main result of [Cîr04] is that we can combine language constructors corresponding to expressive logics for coalgebras in order to obtain expressive logics for coalgebras for more complex functors, e.g. we can combine expressive logics for T_1- and T_2-coalgebras into an expressive logic for $T_1 + T_2$-coalgebras. Note that Cîstea's language constructors are certain endofunctors on a category of algebras, i.e. in her approach the *language* of a coalgebraic modal logic corresponds to such an endofunctor. This is closely related to what we are doing in Chapter 4 where the *logic* corresponds to an endofunctor on the category of Boolean algebras.

Chapter 3

Stone coalgebras

In the previous chapter we presented three different approaches of defining a finitary language for reasoning about T-coalgebras for a given functor $T : \mathsf{Set} \to \mathsf{Set}$. When further investigating these logics we are facing two problems which are well-known in modal logic:

1. The so-called *inadequacy* of Kripke semantics (cf. Section 3.1.2 below), i.e. the problem that there are consistent modal logics that are not valid on any Kripke frame, and

2. the fact that modal logic is lacking the Hennessy-Milner property with respect to the class of all Kripke frames, i.e. there are states in Kripke frames that satisfy exactly the same modal formulas and that are nevertheless not related by a bisimulation.

A source of these problems lies in the fact that we reason with finitary formulas about possibly infinite structures and, indeed, one way of solving the second problem is to consider a modal language with infinite conjunctions (cf. e.g. [Ger96] for modal logic and [Mos99] for the coalgebraic case).

Another way to tackle these problems in modal logic is to move from Kripke semantics to *general frame semantics*. A general frame can be seen as a Kripke frame (W, R) together with a (modal) algebra of admissible subsets of W. The idea is to restrict the valuation of the propositional variables and hence also the semantics of arbitrary formulas to these admissible subsets. A formula will therefore be valid on a general frame iff it is valid on the algebra of admissible subsets. Soundness and completeness of modal logic with respect to the algebraic semantics can be easily carried over to the general frame semantics. This is reflected by the categorical duality between the category DGF of descriptive general frames and the category MA of modal algebras. Furthermore the so-called *canonical frame* based on the set of maximal consistent sets of formulas can be shown to be the final object in DGF. As an immediate consequence one can prove that descriptive general frames form a Hennessy-Milner class, i.e. states that cannot be distinguished by any modal formula are related by some bisimulation (see Section 3.1.3 below).

In this chapter we propose *Stone coalgebras*, i.e. coalgebras for an endofunctor on the category Stone of Stone spaces, as a natural generalisation of descriptive general frames to the level of coalgebras.

To start with, in Section 3.2 we discuss those Stone coalgebras which correspond to the descriptive general frames of modal logic. They are the ones that are associated with the *Vietoris functor* $\mathbb{V}$: Stone $\to$ Stone, a topological analogue of the power set functor on Set. $\mathbb{V}$ is a functorial extension of a well-known topological construction which associates with a topology its Vietoris topology [Eng89]. This construction preserves a number of nice topological properties; in particular, it turns Stone spaces into Stone spaces [Joh82]. We will show that the category Coalg($\mathbb{V}$) of coalgebras for the Vietoris functor is isomorphic to the category DGF of descriptive general frames and therefore dually equivalent to the category MA. Hence $\mathbb{V}$-coalgebras provide a mathematically adequate semantics for modal logic. Furthermore the duality MA $\simeq$ Coalg$(\mathbb{V})^{op}$ has as immediate corollaries the representation of modal algebras as algebras for a functor H : BA $\to$ BA and the existence of the final $\mathbb{V}$-coalgebra.

In Sections 3.3 and 3.4 we further substantiate our case for Stone spaces as a coalgebraic base category, by considering so-called Vietoris polynomial functors as the Stone-based analogues of Kripke polynomial functors over Set (cf. Section 2.1.1). Transferring the work of Jacobs [Jac01] that we summarized in Section 2.1 from the setting of Set-coalgebras to Stone-coalgebras, we establish, for each such functor $\mathbb{T}$, a link between the category $\mathsf{BAO}_{\mathbb{T}}$ of $\mathbb{T}$-sorted Boolean algebras with operators and the category Coalg($\mathbb{T}$) of Stone coalgebras for $\mathbb{T}$. In Section 3.3 we lay the foundations of this work, introducing the notion of a Vietoris polynomial functor (VPF), the algebraic and coalgebraic categories, and functors between these categories. Section 3.4 shows that these functors form an adjunction between the categories $\mathsf{BAO}_{\mathbb{T}}$ and Coalg($\mathbb{T}$), for any *VPF* $\mathbb{T}$. Although this adjunction is not a dual equivalence in general, we will see that each coalgebra can be represented by an algebra, more precisely, Coalg$(\mathbb{T})^{op}$ is (isomorphic to) a full coreflective subcategory of $\mathsf{BAO}_{\mathbb{T}}$. We identify the full subcategory of $\mathsf{BAO}_{\mathbb{T}}$ on which the adjunction restricts to an equivalence and show that the initial $\mathbb{T}$-BAO is dual to the final $\mathbb{T}$-coalgebra.

Finally in Section 3.5 we present an alternative representation of $\mathbb{T}$-BAOs using generators and relations (cf. e.g. [Vic89]). As a result we obtain a characterization of those $\mathbb{T}$-BAOs, which form the full subcategory of $\mathsf{BAO}_{\mathbb{T}}$ that is dual to Coalg($\mathbb{T}$), in terms of free algebras.

The chapter is based on the earlier published paper [KKV04] which is joint work with Alexander Kurz and Yde Venema. Furthermore the author thanks an anonymous referee of [KKV04] whose comments lead to the material in Section 3.5.

3.1 Stone duality

3.1.1 Basic Stone duality

One cornerstone on which this and the following chapter is built is the categorical duality between Boolean algebras and Stone spaces, the so-called *Stone duality*. In this section we are going to give a short summary of Stone duality and in doing so we will fix our notation. For basic notions of general topology the reader is referred to [Eng89].

3.1.1. NOTATION. *For a topological space* $\mathbb{X}$ *we denote by* $\mathrm{O}(\mathbb{X})$ *the set of open subsets of* $\mathbb{X}$*, by* $\mathrm{Cl}(\mathbb{X})$ *the set of closed subsets of* $\mathbb{X}$ *and by* $\mathrm{Clp}(\mathbb{X})$ *the set of clopen subsets of* $\mathbb{X}$*, hence* $\mathrm{O}(\mathbb{X}) \cap \mathrm{Cl}(\mathbb{X}) = \mathrm{Clp}(\mathbb{X})$.

When working with Boolean algebras one would like to think of them in terms of the more concrete power set algebras, i.e. Boolean algebras that are based on the set of subsets of a certain set and in which the Boolean operations are interpreted as the corresponding set-theoretic operations. It is a well-known fact, however, that for each $X \in \mathsf{Set}$ the corresponding power set algebra $(\mathcal{P}(X), \cup, \cap, -, \emptyset, X)$ is a complete and atomic Boolean algebra (see e.g. [Ven06]). Therefore not every Boolean algebra can be represented as a power set algebra.

Already in the 1930s Stone found a way to represent every Boolean algebra as a subalgebra of some power set algebra (cf. [Sto36],[Sto37]). His solution was to consider the Boolean algebras of clopen subsets of certain topological spaces, which are nowadays called *Stone spaces*.

3.1.2. DEFINITION. A topological space $\mathbb{X} = (X, \tau)$ is called a *Stone space* if it is compact, Hausdorff and zero-dimensional, i.e., has a basis of clopen subsets. We denote by Stone the category of Stone spaces with continuous maps as morphisms.

In fact Stone proved not only that for every Boolean algebra $\mathbb{B}$ can we construct a Stone space $\mathbb{X}$ such that $\mathbb{B} \cong \mathrm{Clp}(\mathbb{X})$ but also that there is a correspondence between Boolean homomorphisms and continuous functions on Stone spaces. In categorical terms his results can be summarized by saying that the category Stone of Stone spaces and the category BA of Boolean algebras are dually equivalent. This dual equivalence is witnessed by two functors $\mathbb{S}\mathrm{p} : \mathsf{BA} \to \mathsf{Stone}^{\mathrm{op}}$ and $\mathbb{C}\mathrm{lp} : \mathsf{Stone}^{\mathrm{op}} \to \mathsf{BA}$, which are defined as follows:

$$\begin{array}{rcl@{\qquad}rcl} \mathbb{S}\mathrm{p} : \mathsf{BA} & \to & \mathsf{Stone}^{\mathrm{op}} & \mathbb{C}\mathrm{lp} : \mathsf{Stone}^{\mathrm{op}} & \to & \mathsf{BA} \\ \mathbb{B} & \mapsto & \mathbb{S}\mathrm{p}\mathbb{B} := (\mathsf{Uf}\mathbb{B}, \tau_{\mathbb{B}}) & \mathbb{X} & \mapsto & \mathbb{C}\mathrm{lp}\mathbb{X} := \mathrm{Clp}(\mathbb{X}) \\ f : \mathbb{B}_1 \to \mathbb{B}_2 & \mapsto & f^{-1} : \mathbb{S}\mathrm{p}\mathbb{B}_2 \to \mathbb{S}\mathrm{p}\mathbb{B}_1 & g : \mathbb{X}_1 \to \mathbb{X}_2 & \mapsto & g^{-1} : \mathbb{C}\mathrm{lp}\mathbb{X}_2 \to \mathbb{C}\mathrm{lp}\mathbb{X}_1. \end{array}$$

Here f^{-1} and g^{-1} denote the inverse image functions, i.e.

$$\begin{array}{rcll} f^{-1}(u) & := & \{b \in \mathbb{B}_1 \mid f(b) \in u\} & \text{for } u \in \mathsf{Uf}\mathbb{B}_2 \\ g^{-1}(b) & := & \{x \in X_1 \mid g(x) \in b\} & \text{for } b \in \mathrm{Clp}(\mathbb{X}_2), \end{array}$$

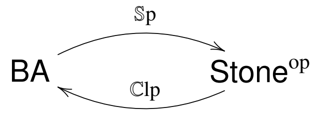

Figure 3.1: Stone duality

$\mathsf{Uf}\mathbb{B}$ is the collection of ultrafilters over $\mathbb{B}$ and $\tau_{\mathbb{B}}$ is the topology generated by the collection of sets $\{\hat{b} \mid b \in \mathbb{B}\}$, where for $b \in B$ we let $\hat{b} := \{u \in \mathsf{Uf}\mathbb{B} \mid b \in u\}$.

Furthermore we define, for every $\mathbb{B} \in \mathsf{BA}$ and every $\mathbb{X} \in \mathsf{Stone}$, morphisms

$$\begin{array}{rcl} \iota_{\mathbb{B}} : \mathbb{B} & \to & \mathbb{C}\mathrm{lp}\mathbb{S}\mathrm{p}\mathbb{B} \\ b & \mapsto & \hat{b} \end{array} \qquad \begin{array}{rcl} \epsilon_{\mathbb{X}} : \mathbb{X} & \to & \mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}\mathbb{X} \\ x & \mapsto & \{b \in \mathbb{C}\mathrm{lp}\mathbb{X} \mid x \in b\}. \end{array}$$

We are now able to state Stone's representation theorem.

3.1.3. Theorem (Stone). *The families of maps $(\epsilon_{\mathbb{X}})_{\mathbb{X} \in \mathsf{Stone}}$ and $(\iota_{\mathbb{B}})_{\mathbb{B} \in \mathsf{BA}}$ give rise to natural isomorphisms $\epsilon : Id_{\mathsf{Stone}} \to \mathbb{S}\mathrm{p} \circ \mathbb{C}\mathrm{lp}$ and $\iota : Id_{\mathsf{BA}} \to \mathbb{C}\mathrm{lp} \circ \mathbb{S}\mathrm{p}$ respectively. Therefore the categories* BA *and* Stone *are dually equivalent.*

3.1.4. Remark. Note that this implies that a given Boolean algebra $\mathbb{B}$ is isomorphic to $\mathrm{Clp}(\mathbb{S}\mathrm{p}\mathbb{B})$ and hence to a subalgebra of the power set algebra based on $\mathcal{P}(\mathsf{Uf}\mathbb{B})$.

We finish this brief introduction to Stone duality by stating how finite products and coproducts in Stone are computed[1].

3.1.5. Fact. Let $\mathbb{X}_1 = (X_1, \tau_1)$ and $\mathbb{X}_2 = (X_2, \tau_2)$ be Stone spaces. Then

- the *product* $\mathbb{X}_1 \times \mathbb{X}_2$ of $\mathbb{X}_1$ and $\mathbb{X}_2$ is the topological space that has the cartesian product $X_1 \times X_2$ as carrier set and whose topology has as base the set

$$\{U_1 \times U_2 \subseteq X_1 \times X_2 \mid U_1 \in \tau_1 \text{ and } U_2 \in \tau_2\},$$

- the *coproduct* $\mathbb{X}_1 + \mathbb{X}_2$ of $\mathbb{X}_1$ and $\mathbb{X}_2$ is the topological space that has the disjoint uion $X_1 + X_2$ as a carrier set and whose topology has as a base the set

$$\{U \subseteq X_1 + X_2 \mid U \cap X_1 \in \tau_1 \text{ and } U \cap X_2 \in \tau_2\}.$$

For a proof of the fact that $\mathbb{X}_1 \times \mathbb{X}_2$ and $\mathbb{X}_1 + \mathbb{X}_2$ are indeed the product and coproduct of the two spaces $\mathbb{X}_1$ and $\mathbb{X}_2$ the reader is referred to [Eng89]. Much more information about Stone duality can be found in [Joh82].

[1]In fact we only state what binary products and coproudcts in Stone are, but this can be of course easily generalised to arbitrary finite products and coproducts respectively.

3.1.2 Duality for modal algebras

Before we turn to the extension of Stone's representation theorem from Boolean algebras to modal algebras we briefly recall some basic facts about normal modal logic. For a more detailed introduction to the subject the reader is referred to standard textbooks about modal logic such as [BdV01, CZ97, Kra99]. In the following we restrict ourselves to the case of the basic modal language with one unary modal operator.

3.1.6. Definition. Let Φ be a set of propositional variables. The set FML(Φ) of *modal formulas* with variables in Φ is defined inductively as follows:

$$\mathrm{FML}(\Phi) ::= \bot \mid p \in \Phi \mid \neg\phi \mid \phi \wedge \phi \mid \Box\phi.$$

A logic is now identified with the set of theorems of that logic.

3.1.7. Definition. A *normal modal logic* is a set Λ of formulas such that

- Λ contains all tautologies,
- $\Box(p \to q) \to (\Box p \to \Box q) \in \Lambda$,
- Λ is closed under the rules *modus ponens*(MP) and *necessitation* (N):

$$\frac{\phi \to \psi \in \Lambda \qquad \phi \in \Lambda}{\psi \in \Lambda}\ (\mathrm{MP}) \qquad\qquad \frac{\phi \in \Lambda}{\Box\phi \in \Lambda}\ (\mathrm{N})$$

- Λ is closed under uniform substitution, i.e. if $\phi \in \Lambda$ and $\sigma : \Phi \to \mathrm{FML}(\Phi)$ is a function (a "substitution") then $\sigma(\phi) \in \Lambda$, where $\sigma(\phi)$ denotes the formula that is obtained from ϕ by replacing every occurrence of a propositional variable $p \in \Phi$ by the formula $\sigma(p)$.

The modal logic **K** is the smallest normal modal logic.

The most commonly used semantics for normal modal logic is Kripke semantics .

3.1.8. Definition. A *Kripke frame* is a pair $\mathfrak{F} := (W, R)$, where W is a set and $R \subseteq W \times W$ is a binary relation. Let Φ be a set of propositional variables. A *Kripke model* is a triple $\mathfrak{M} = (W, R, V)$ where (W, R) is a Kripke frame and $V : \Phi \to \mathcal{P}(W)$ is a so-called *valuation function*. Given a Kripke model $\mathfrak{M} = (W, R, V)$ the semantics $[\![\phi]\!]_{\mathfrak{M}} \subseteq W$ of a formula $\phi \in \mathrm{FML}(\Phi)$ is inductively defined as follows:

$$\begin{aligned}
[\![\bot]\!]_{\mathfrak{M}} &:= \emptyset \\
[\![p]\!]_{\mathfrak{M}} &:= V(p) \\
[\![\neg\phi]\!]_{\mathfrak{M}} &:= W \setminus [\![\phi]\!]_{\mathfrak{M}} \\
[\![\phi_1 \wedge \phi_2]\!]_{\mathfrak{M}} &:= [\![\phi_1]\!]_{\mathfrak{M}} \cap [\![\phi_2]\!]_{\mathfrak{M}} \\
[\![\Box\phi]\!]_{\mathfrak{M}} &:= [R]([\![\phi]\!]_{\mathfrak{M}})
\end{aligned}$$

where

$$\begin{array}{rcl} [R](\_) : \mathcal{P}(W) & \to & \mathcal{P}(W) \\ V & \mapsto & \{x \in W \mid \forall y \, . \, xRy \Rightarrow y \in V\}. \end{array}$$

Note that $[R](\_)$ preserves arbitrary intersections.

We say that a frame (W, R) *validates* a formula ϕ if for all valuation functions V we have $[\![\phi]\!]_{(W,R,V)} = W$. We say that a frame $\mathfrak{F}$ validates a logic Λ if it validates all $\phi \in \Lambda$. In this case $\mathfrak{F}$ is called a frame of the logic Λ and we denote by $\mathrm{Frm}(\Lambda)$ the set of frames of Λ.

Many modal logics Λ are sound and complete with respect to $\mathrm{Frm}(\Lambda)$, i.e.

$$\phi \in \Lambda \quad \text{iff} \quad \text{for all } \mathfrak{F} \in \mathrm{Frm}(\Lambda) \, . \, \mathfrak{F} \text{ validates } \phi.$$

There are, however, examples of modal logics that are incomplete with respect to $\mathrm{Frm}(\Lambda)$. For examples of such logics the reader is referred to [Fin74, Tho74, BdV01]. The incompleteness originates from the fact that Kripke semantics is too weak to distinguish modal logics, e.g. Thomason's logic Λ_T in [Tho74] is consistent but $\mathrm{Frm}(\Lambda_T) = \emptyset$ and therefore Kripke semantics does not distinguish it from the inconsistent logic.

Boolean algebras with additional operators provide an algebraic semantics for modal logic. The algebras corresponding to the smallest normal modal logic **K** are the so-called modal algebras.

3.1.9. Definition. A *modal algebra* is a pair $\mathfrak{B} = (\mathbb{B}, f)$ such that

$$\mathbb{B} = (B, \vee, \wedge, \neg, \bot, \top)$$

is a Boolean algebra and $f : B \to B$ preserves $\top$ and binary meets, i.e. $f(\top) = \top$ and $f(b_1 \wedge b_2) = f(b_1) \wedge f(b_2)$ for all $b_1, b_2 \in B$. The category of modal algebras as objects and homomorphisms as arrows will be denoted by MA.

3.1.10. Definition. Let Φ be a set of variables, $\mathfrak{B} = (\mathbb{B}, f)$ a modal algebra and $V : \Phi \to \mathfrak{B}$ a valuation. Then the algebraic semantics $[\![\phi]\!]_{(\mathfrak{B},V)} \in \mathbb{B}$ of a formula $\phi \in \mathrm{FML}(\Phi)$ is defined as follows

$$\begin{array}{rcl} [\![\bot]\!]_{(\mathfrak{B},V)} & := & \bot \\ [\![p]\!]_{(\mathfrak{B},V)} & := & V(p) \\ [\![\neg\phi]\!]_{(\mathfrak{B},V)} & := & \neg[\![\phi]\!]_{(\mathfrak{B},V)} \\ [\![\phi_1 \wedge \phi_2]\!]_{(\mathfrak{B},V)} & := & [\![\phi_1]\!]_{(\mathfrak{B},V)} \wedge [\![\phi_2]\!]_{(\mathfrak{B},V)} \\ [\![\Box\phi]\!]_{(\mathfrak{B},V)} & := & f([\![\phi]\!]_{(\mathfrak{B},V)}). \end{array}$$

We say $\mathfrak{B}$ validates ϕ if $[\![\phi]\!]_{(\mathfrak{B},V)} = \top$ for all $V : \Phi \to \mathfrak{B}$. Furthermore we say $\mathfrak{B}$ validates a logic Λ if $\mathfrak{B}$ validates ϕ for all $\phi \in \Lambda$. The family of all modal algebras that validate a logic Λ is denoted by V_Λ ("the variety of Λ-algebras").

In contrast to the Kripke semantics, the algebraic semantics has the advantage that it is adequate: all normal modal logics are characterized by the corresponding variety of algebras.

3.1.11. Theorem (Adequacy). *For every normal modal logic* Λ *we have*

$$\phi \in \Lambda \quad \textit{iff} \quad \textit{for all } \mathfrak{B} \in \mathrm{V}_\Lambda . \mathfrak{B} \textit{ validates } \phi.$$

3.1.12. Remark. Provided a deducibility relation $\vdash_\Lambda$ for Λ such that $\phi \in \Lambda$ iff $\vdash_\Lambda \phi$, adequacy amounts to soundness and completeness of algebraic semantics.

Modal algebras, however, are fairly abstract in nature and many modal logicians prefer the intuitive, geometric appeal of Kripke frames. *General frames*, unifying the algebraic and the Kripke semantics in one structure, provide a nice compromise.

3.1.13. Definition. A *general frame* is a structure $\mathfrak{G} = (W, R, A)$ such that (W, R) is a Kripke frame and A is a collection of so-called *admissible* subsets of W that is closed under the Boolean operations and under the operation $[R]$. A general frame $\mathfrak{G} = (W, R, A)$ is called

- *differentiated* if for all distinct $w_1, w_2 \in W$ there is a 'witness' $a \in A$ such that $w_1 \in a$ and $w_2 \notin a$,
- *tight* if whenever v is not an R-successor of w, then there is a 'witness' $a \in A$ such that $v \in a$ and $w \in [R](W \setminus a)$, and
- *compact* if $\bigcap A_0 \neq \emptyset$ for every subset A_0 of A which has the finite intersection property, i.e. for every $A_0 \subseteq A$ such that for all finite sets $B \subseteq A_0$ we have $\bigcap B \neq \emptyset$.

A general frame is *descriptive* if it is differentiated, tight and compact.

The term 'admissible' subset is explained by the semantic restriction that allows only those Kripke models on a general frame for which the extensions of the atomic formulas are admissible sets, i.e. a valuation $V : \Phi \to \mathcal{P}(W)$ on a general frame (W, R, A) is only admissible if $V(p) \in A$ for all $p \in \Phi$. This leads to the following notion of validity of a formula on a general frame.

3.1.14. Definition. A general frame (W, R, A) is said to validate a modal formula $\phi \in \mathrm{FML}(\Phi)$ if $[\![\phi]\!]_{(W,R,V)} = W$ for all valuations $V : \Phi \to A$, i.e. for all valuations that are "admissible".

3.1.15. Example.

1. Any Kripke frame (X, R) can be considered as a general frame $(X, R, \mathcal{P}(X))$.

2. If $\mathfrak{B} = (B, \vee, \wedge, \neg, \bot, \top, f)$ is a modal algebra then $(\mathsf{Uf}\mathbb{B}, R_f, \hat{B})$, where $R_f = \{(u, v) \mid \forall b \,.\, f(b) \in u \Rightarrow b \in v\}$ and $\hat{B} = \{\{u \in \mathsf{Uf}\mathbb{B} \mid b \in u\} \mid b \in B\}$, is a descriptive general frame.

3. If $\mathfrak{G} = (W, R, A)$ is a general frame then $(A, \cup, \cap, -, \emptyset, W, [R])$ is a modal algebra.

Descriptive general frames can be nicely characterized in topological terms.

3.1.16. Fact. *Let $\mathfrak{G} = (W, R, A)$ be a general frame and let τ_A be the topology on W generated by A. Then the following are equivalent:*

1. $\mathfrak{G}$ is descriptive,

2. (W, τ_A) is a Stone space and R is point-closed*, i.e. for all $w \in W$ we have $R[w] \in \mathrm{Cl}(W, \tau_A)$.*

Proof. (W, τ_A) is a Stone space because a descriptive general frame is compact, differentiated and the admissible sets are closed under the Boolean operations. The tightness condition of descriptive general frames is equivalent to the fact that the relation is point-closed. QED

Fact 3.1.16 is the key for extending Stone duality to modal algebras. Let us first introduce the categories GF of general frames and DGF of descriptive general frames before we state the duality $\mathsf{MA} \simeq \mathsf{DGF}^{\mathrm{op}}$.

3.1.17. Definition. [Gol76, Def. 5.1] A *general frame morphism* $\theta : (W, R, A) \to (W', R', A')$ is a function from W to W' such that (i) $\theta : (W, R) \to (W', R')$ is a bounded morphism and (ii) $\theta^{-1}(a') \in A$ for all $a' \in A'$. General frame morphisms will be also called *continuous bounded morphisms*. We let GF (DGF) denote the category with general frames (descriptive general frames, respectively) as its objects, and the general frame morphisms as the arrows.

In Example 3.1.15 we saw how to transform a modal algebra into a descriptive general frame and vice versa. These constructions give rise to two functors

$$
\begin{array}{llll}
D: & \mathsf{MA} & \to & \mathsf{DGF}^{\mathrm{op}} \\
 & (\mathbb{B}, f) & \mapsto & (\mathsf{Uf}\mathbb{B}, R_f, \hat{B}) \\
 & h & \mapsto & h^{-1}
\end{array}
\qquad
\begin{array}{llll}
M: & \mathsf{DGF}^{\mathrm{op}} & \to & \mathsf{MA} \\
 & (W, R, A) & \mapsto & (A, [R]) \\
 & \theta & \mapsto & \theta^{-1},
\end{array}
$$

where $\hat{B}$ and R_f are defined as in 3.1.15, (2).

3.1.18. Theorem ([Gol76]). *The categories* MA *and* DGF *are dually equivalent.*

Proof. Let (W, R, A) be a descriptive general frame and $\mathbb{W} = (W, \tau_A)$ its associated Stone space. Then the Stone isomorphism $\epsilon : \mathbb{W} \to \mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}\mathbb{W} \in \mathsf{Stone}$ is also an isomorphism in the category DGF between (W, R, A) to $D(M(W, R, A))$. Likewise, for a modal algebra $\mathfrak{B} = (\mathbb{B}, F)$ the Stone isomorphism $\iota : \mathbb{B} \to \mathbb{C}\mathrm{lp}\mathbb{S}\mathrm{p}\mathbb{B} \in \mathsf{BA}$ can be shown to be an isomorphism in MA between $\mathfrak{B}$ and $M(D\mathfrak{B})$. In this way ϵ and ι give rise to natural isomorphisms $\hat{\epsilon} : Id_{\mathsf{DGF}} \to D \circ M$ and $\hat{\iota} : Id_{\mathsf{MA}} \to M \circ D$. QED

A consequence of the duality $\mathsf{MA} \simeq \mathsf{DGF}^{op}$ is the adequacy of general frame semantics, i.e. every normal modal logic can be characterized by its class of general frames. The only thing which has to be observed is that for all modal algebras $\mathfrak{B}$ we have that

$$\mathfrak{B} \text{ validates } \phi \quad \text{iff} \quad D(\mathfrak{B}) \text{ validates } \phi \quad \text{for all } \phi \in \mathrm{FML}(\Phi).$$

The adequacy follows then directly from the duality $\mathsf{MA} \simeq \mathsf{DGF}^{op}$ and from the adequacy of algebraic semantics (cf. Theorem 3.1.11). Hence the general frame semantics combines the nice properties of both the Kripke semantics and the algebraic semantics.

3.1.3 Modal logic is expressive for descriptive general frames

Descriptive general frames also form a so-called Hennessy-Milner class as remarked in the introduction of this chapter. In this subsection we want to make this statement more precise. In the following we only consider modal formulas without proposition letters. In this case both the algebraic and the Kripke frame semantics are independent of a valuation, i.e. we can interpret a formula on a Kripke frame and on algebra respectively without mentioning a valuation function. We want to prove the following proposition, which is "folklore" among modal logicians.

3.1.19. Proposition. *Let (W_1, R_1, A_1) and (W_2, R_2, A_1) be arbitrary descriptive frames and $w_1 \in W_1, w_2 \in W_2$ such that for all (variable-free) modal formulas ϕ we have*

$$w_1 \models \phi \quad \textit{iff} \quad w_2 \models \phi.$$

Then w_1 and w_2 are bisimilar, i.e. $(W_1, R_1[\_], w_1) \leftrightarroweq_{\mathcal{P}} (W_2, R_2[\_], w_2)$.

3.1.20. Remark. We call two states $w_1 \in W_1, w_2 \in W_2$ of two Kripke frames (W_1, R_1) and (W_2, R_2) bisimilar if the corresponding rooted $\mathcal{P}$-coalgebras $(W_1, R_1[\_], w_1)$ and $(W_2, R_2[\_], w_2)$ are $\mathcal{P}$-bisimilar. This notion of bisimilarity coincides with the usual notion of bisimilarity for modal logic without proposition letters (cf. [Mos99, Section 2]).

In order to prove this proposition, we need the following facts that relate the algebraic and coalgebraic semantics of modal logic.

3.1.21. Notation. *We denote by $\mathcal{L}_{\mathsf{MA}}$ the initial object in* MA*, the so called* Lindenbaum-Tarski algebra *of modal logic. The algebra $\mathcal{L}_{\mathsf{MA}}$ is defined as the modal term algebra, which has as its carrier set the set of modal formulas* FML(∅)*, modulo the congruence relation ≡:*

$$\phi \equiv \psi :\Leftrightarrow \phi \leftrightarrow \psi \in \mathbf{K}.$$

i.e. ≡ identifies logically equivalent formulas.

Intuitively the Lindenbaum-Tarski algebra consists of the set of (variable-free) modal formulas quotiented by derivable equivalence.

3.1.22. FACT. The following are true:

1. The graph of a bounded morphism is a bisimulation.
2. Bisimulations are closed under composition.
3. Let $D\mathcal{L}_{\mathsf{MA}}$ be the dual frame of the Lindenbaum-Tarski algebra. Then the states x of $D\mathcal{L}_{\mathsf{MA}}$ are in one-to-one correspondence with the maximal consistent sets of formulas u_x and for all $x \in D\mathcal{L}_{\mathsf{MA}}$ we have

$$x \models \phi \quad \text{iff} \quad \phi \in u_x.$$

For proofs of these facts we refer the reader to [BdV01]. We are now ready to give the proof of Proposition 3.1.19.

Proof of Prop 3.1.19. Let $(W_1, R_1, A_1), (W_2, R_2, A_2) \in \mathsf{DGF}$ and $w_1 \in W_1, w_2 \in W_2$ as described in the proposition. By the duality between MA and DGF we know that $D\mathcal{L}_{\mathsf{MA}}$ is the final object in DGF, i.e. there are continuous bounded morphisms

$$f_i : (W_i, R_i, A_i) \to D\mathcal{L}_{\mathsf{MA}} \quad \text{for } i \in \{1, 2\}.$$

The states w_1 and w_2 satisfy the same formulas according to our assumption. Let u be this maximal consistent set of formulas satisfied in both w_1 and w_2. According to Fact 3.1.22 there is a (unique) state x in $D\mathcal{L}_{\mathsf{MA}}$ such that x satisfies the formulas in u. Furthermore, as the graph of a bounded morphism is a bisimulation and the truth of modal formulas is invariant under bisimulations, we get that $f_i(w_i) = x$ for $i \in 1, 2$. But then $(w_1, w_2) \in \mathrm{Gr}(f_1) \circ \mathrm{Gr}(f_2)^{\smile}$ and $\mathrm{Gr}(f_1) \circ \mathrm{Gr}(f_2)^{\smile}$ is a bisimulation according to Fact 3.1.22. Hence $w_1 \leftrightarroweq w_2$ as required. QED

3.2 From Kripke to Vietoris

Since Kripke frames (and models) form prime examples of coalgebras (cf. Example A.3.2), the question naturally arises whether (descriptive) general frames can be seen as coalgebras as well. Our positive answer to this question is based on two crucial observations from Fact 3.1.16. First, the admissible sets of a descriptive frame form a basis for a Stone topology. Second, the successor set of any point is closed in this topology. This suggests that if we are looking for a coalgebraic counterpart of a descriptive general frame $\mathfrak{G} = (W, R, A)$, it should be of the form $R[\_] : (W, \tau_A) \to (\mathrm{Cl}(W, \tau_A), \tau_?)$ where $\tau_?$ is some suitable topology on $\mathrm{Cl}(W, \tau_A)$, which turns $\mathrm{Cl}(W, \tau_A)$ again into a Stone space. A good candidate is the Vietoris topology: it is based on the closed sets of τ and it yields a Stone space if we started from one. Moreover, as we will see, choosing the Vietoris topology for $\tau_?$, continuity of the map $R[\_]$ corresponds to the admissible sets being closed under $[R]$ (cf. Remark 3.2.7). Given a topological space (X, τ), the natural question arises of what is the right notion of a hyperspace, i.e. a space which has as points subsets of the original space and which should correspond to the power set construction on Set. In [Vie22] Vietoris proposed to consider the following topological space based on the closed subsets of the original space.

3.2.1. Definition. Let $\mathbb{X} = (X, \tau)$ be a topological space. Define the operations

$$\begin{aligned} [\ni], \langle\ni\rangle : \mathcal{P}(X) &\to \mathcal{P}(\mathrm{Cl}(\mathbb{X})) \qquad \text{by} \\ [\ni]U &:= \{F \in \mathrm{Cl}(\mathbb{X}) \mid F \subseteq U\}, \\ \langle\ni\rangle U &:= \{F \in \mathrm{Cl}(\mathbb{X}) \mid F \cap U \neq \emptyset\}. \end{aligned}$$

Given a subset $Q \subseteq \mathcal{P}(X)$, define $V_Q := \{[\ni]U \mid U \in Q\} \cup \{\langle\ni\rangle U \mid U \in Q\}$. The *Vietoris space* $\mathbb{V}(\mathbb{X})$ associated with $\mathbb{X}$ is given by the topology $\upsilon_{\mathbb{X}}$ on $\mathrm{Cl}(\mathbb{X})$ which is generated by V_τ as subbasis.

In case the original topology is compact, there are other equivalent ways to generate the Vietoris topology. This has nice consequences for the case that the original topology is a Stone space.

3.2.2. Lemma. *Let $\mathbb{X} = (X, \tau)$ be a compact topological space and let $\mathcal{B}$ be a basis of τ that is closed under finite unions. Then the set $V_{\mathcal{B}}$ forms a subbasis for $\upsilon_{\mathbb{X}}$. In particular, if $\mathbb{X}$ is a Stone space, then the set $V_{\mathrm{Clp}(\mathbb{X})}$ forms a subbasis for $\upsilon_{\mathbb{X}}$.*

The Vietoris construction preserves various useful topological properties; proofs of this can be found in for instance [Mic51].

3.2.3. Lemma. *Let $\mathbb{X} = (X, \tau)$ be a topological space.*

1. *If $\mathbb{X}$ is compact then $(\mathrm{Cl}(\mathbb{X}), \upsilon_{\mathbb{X}})$ is compact.*
2. *If $\mathbb{X}$ is compact and Hausdorff, then $(\mathrm{Cl}(\mathbb{X}), \upsilon_{\mathbb{X}})$ is compact and Hausdorff.*
3. *If $\mathbb{X}$ is a Stone space, then so is $(\mathrm{Cl}(\mathbb{X}), \upsilon_{\mathbb{X}})$.*

The last item shows that the Vietoris construction can be used to map objects in Stone to objects in Stone. This gives rise to the following functor.

3.2.4. Definition. The *Vietoris functor* $\mathbb{V}$ on the category of Stone spaces is defined as follows:

$$\begin{aligned} \mathbb{V} : \mathsf{Stone} &\to \mathsf{Stone} \\ \mathbb{X} &\mapsto \mathbb{V}\mathbb{X} := (\mathrm{Cl}(\mathbb{X}), \upsilon_{\mathbb{X}}) \\ f : \mathbb{X} \to \mathbb{Y} &\mapsto f[\_] : \mathbb{V}\mathbb{X} \to \mathbb{V}\mathbb{Y} \end{aligned}$$

where $f[\_]$ denotes the direct image function.

3.2.5. Remark. The functor $\mathbb{V}$ is well-defined on morphisms, because for a continuous function $f : \mathbb{X} \to \mathbb{Y} \in \mathsf{Stone}$ and a closed subset $F \in \mathrm{Cl}(\mathbb{X})$ we have that $f[F] \in \mathrm{Cl}(\mathbb{X})$ (cf. [Eng89, Theorem 3.1.8]).

The Vietoris functor provides us with the coalgebraic representation of descriptive general frames as the categories $\mathsf{Coalg}(\mathbb{V})$ and DGF turn out to be isomorphic. Using Fact 3.1.16 it is straightforward to verify that the following definition is correct, that is, it indeed defines two functors.

3.2.6. Definition. We define the functor $\mathbb{C} : \mathsf{DGF} \to \mathsf{Coalg}(\mathbb{V})$ via

$$(W, R, A) \quad \mapsto \quad (W, \tau_A) \stackrel{R[-]}{\longrightarrow} \mathbb{V}(W, \tau_A)$$

Here τ_A denotes the Stone topology generated by taking A as a basis. Conversely, there is a functor $\mathbb{D} : \mathsf{Coalg}(\mathbb{V}) \to \mathsf{DGF}$ given by

$$(\mathbb{X}, \gamma) \quad \mapsto \quad (X, R_\gamma, \mathrm{Clp}(\mathbb{X}))$$

where R_γ is defined by $R_\gamma x_1 x_2$ iff $x_2 \in \gamma(x_1)$. On morphisms both functors act as the identity with respect to the underlying Set-functions.

3.2.7. Remark. For the well-definedness of $\mathbb{D}$ it is important that the continuity of $\gamma : \mathbb{X} \to \mathbb{V}\mathbb{X}$ implies that $\mathrm{Clp}(\mathbb{X})$ is closed under the operation $[R_\gamma] : \mathrm{Clp}(\mathbb{X}) \to \mathrm{Clp}(\mathbb{X})$. This can be seen as follows: Let U be an arbitrary clopen subset of $\mathbb{X}$. We have to show that $[R_\gamma](U) \in \mathrm{Clp}(\mathbb{X})$. Spelling out the definition of $[R_\gamma]$ we calculate that

$$\begin{array}{rcl} [R_\gamma](U) & = & \{x \mid R_\gamma[x] \subseteq U\} \\ & = & \{x \mid R_\gamma[x] \in [\ni]U\} \\ & = & \{x \mid \gamma(x) \in [\ni]U\} \\ & = & \gamma^{-1}([\ni]U). \end{array}$$

Because γ is continuous and $[\ni]U$ is a clopen subset of $\mathbb{V}\mathbb{X}$ we can conclude that $[R_\gamma](U)$ is a clopen subset of $\mathbb{X}$.

3.2.8. Theorem. *The functors $\mathbb{C}$ and $\mathbb{D}$ form an isomorphism between the categories* DGF *and* $\mathsf{Coalg}(\mathbb{V})$.

Proof. The theorem can be easily proven by just spelling out the definitions. QED

Hence descriptive general frames and $\mathbb{V}$-coalgebras are essentially the same, but we can extend this correspondence a bit further, namely to models on general frames.

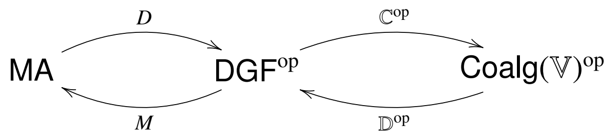

Figure 3.2: Duality for modal algebras

3.2.9. Remark. Recall that for a given set Φ of propositional variables a Kripke model is a triple $(W, R, V : \Phi \to \mathcal{P}(W))$. This can be represented as a set coalgebra

$$W \xrightarrow{\langle v, R[\_]\rangle} \left(\prod_{p\in\Phi} 2\right) \times \mathcal{P}W \ \in \mathsf{Set},$$

where $2 = \{0, 1\}$ denotes the two element set and $\pi_p(v(w)) := 1$ if $w \in V(p)$. Analogously we represent a Kripke model based on a descriptive general frame (W, R, A, V) as a Stone coalgebra

$$(W, \tau_A) \xrightarrow{\langle v, R[\_]\rangle} \left(\prod_{p\in\Phi} 2\right) \times \mathbb{V}(W, \tau_A) \ \in \mathsf{Stone},$$

where 2 is equipped with the discrete topology and we exploited the correspondence between $\mathsf{Coalg}(\mathbb{V})$ and DGF. The fact that $v : (W, \tau_A) \to \prod_{p\in\Phi} 2$ should be an arrow in Stone and therefore continuous is equivalent to the fact that $V(p) \in A$ for all $p \in \Phi$, i.e. continuity of the valuation is equivalent to the admissibility of the valuation.

Let us note some corollaries of Theorem 3.2.8. Clearly $\mathsf{Coalg}(\mathbb{V})$ is dual to MA.

3.2.10. Corollary. *The categories* MA *and* $\mathsf{Coalg}(\mathbb{V})$ *are dually equivalent.*

Proof. The claim follows immediately from the duality $\mathsf{MA} \simeq \mathsf{DGF}^{op}$ and the isomorphism $\mathsf{DGF} \cong \mathsf{Coalg}(\mathbb{V})$, cf. Figure 3.2. QED

Using the trivial duality $(\mathsf{Coalg}(\mathbb{V}))^{op} = \mathsf{Alg}(\mathbb{V}^{op})$ (cf. Def. A.1.2 and Def. A.1.3), it follows that $\mathsf{MA} \simeq \mathsf{Alg}(\mathbb{V}^{op})$. With $\mathsf{Stone}^{op} \simeq \mathsf{BA}$ we obtain the following.

3.2.11. Corollary. *There is a functor* $H : \mathsf{BA} \to \mathsf{BA}$ *such that the category of modal algebras* MA *is equivalent to the category* $\mathsf{Alg}(H)$ *of algebras for the functor* H.

Proof. With the help of the contravariant functors $\mathbb{C}\mathrm{lp} : \mathsf{Stone}^{op} \to \mathsf{BA}$, $\mathbb{S}\mathrm{p} : \mathsf{BA} \to \mathsf{Stone}^{op}$, we let $H = \mathbb{C}\mathrm{lp} \circ \mathbb{V}^{op} \circ \mathbb{S}\mathrm{p}$. The claim now follows from the observation that $\mathsf{Alg}(H)$ is dual to $\mathsf{Coalg}(\mathbb{V})$: An algebra $HA \xrightarrow{\alpha} A$ corresponds to the coalgebra $\mathbb{S}\mathrm{p}A \xrightarrow{\mathbb{S}\mathrm{p}\alpha} \mathbb{S}\mathrm{p}HA \cong \mathbb{V}\mathbb{S}\mathrm{p}A$ and a coalgebra $\mathbb{X} \xrightarrow{\xi} \mathbb{V}\mathbb{X}$ corresponds to the algebra $H\mathbb{C}\mathrm{lp}\mathbb{X} \cong \mathbb{C}\mathrm{lp}\mathbb{V}\mathbb{X} \xrightarrow{\mathbb{C}\mathrm{lp}\xi} \mathbb{C}\mathrm{lp}\mathbb{X}$. QED

An explicit description of H not involving the Vietoris functor is given by the following proposition.

3.2.12. Proposition. *Let $H : \mathsf{BA} \to \mathsf{BA}$ be the functor that assigns to a Boolean algebra the free Boolean algebra over its underlying meet-semilattice. Then $\mathsf{Alg}(H)$ is isomorphic to the category of modal algebras* MA.

Proof. Let $\mathsf{BA}_\wedge$ be the category with Boolean algebras as objects and finite meet preserving functions as morphisms, i.e. functions that preserve binary meets and the top-element. Furthermore let MPF be the following category. An object of MPF is an endofunction $A \xrightarrow{m} A \in \mathsf{BA}_\wedge$. A morphism $f : (A \xrightarrow{m} A) \longrightarrow (A' \xrightarrow{m'} A')$ is a Boolean algebra morphism $f : A \to A'$ such that $m' \circ f = f \circ m$. It is easy to see that MPF is isomorphic to the category MA and therefore we can finish our proof by showing that $\mathsf{Alg}(H)$ and MPF are isomorphic.

In order to prove that $\mathsf{Alg}(H)$ and MPF are isomorphic categories, we first show that $\mathsf{BA}(HA, A) \cong \mathsf{BA}_\wedge(A, A)$, or slightly more general and precise, $\mathsf{BA}(HA, B) \cong \mathsf{BA}_\wedge(IA, IB)$ where $I : \mathsf{BA} \hookrightarrow \mathsf{BA}_\wedge$. (We denote, for a category C and objects A, B in C, the set of morphisms between A and B by $\mathsf{C}(A, B)$, cf. A.1.1.) Indeed, consider the forgetful functors $U : \mathsf{BA} \to \mathsf{SF}$, $V : \mathsf{BA}_\wedge \to \mathsf{SF}$ to the category SF of meet-semilattices with top element, and the left adjoint F of U. Using our assumption $H = FU$, we calculate $\mathsf{BA}(HA, B) = \mathsf{BA}(FUA, B) \cong \mathsf{SF}(UA, UB) \cong \mathsf{SF}(VIA, VIB) \cong \mathsf{BA}_\wedge(IA, IB)$. Hence there exists an isomorphism

$$\phi_{A,B} : \mathsf{BA}(HA, B) \to \mathsf{BA}_\wedge(A, B)$$

that is natural in both A and B. The isomorphisms $\phi_{A,A} : \mathsf{BA}(HA, A) \to \mathsf{BA}_\wedge(A, A)$, $A \in \mathsf{BA}$, give us an isomorphism ϕ between the objects $(A, \alpha : HA \to A)$ of $\mathsf{Alg}(H)$ and the objects $(A, m : A \to A)$ of MPF. On morphisms, we define ϕ to be the identity. To sum it up, we define a functor

$$\begin{array}{rcl} \Phi : \mathsf{Alg}(H) & \to & \mathsf{BA}_\wedge \\ (A, \alpha) & \mapsto & (A, \phi_{A,A}(\alpha)) \\ f & \mapsto & f. \end{array}$$

Let us check that Φ is well-defined on morphisms. Suppose that $f : (A, \alpha_A) \to (B, \alpha_B)$ is an $\mathsf{Alg}(H)$-morphism, i.e. $f \circ \alpha_A = \alpha_B \circ Hf$. This last equation can be written as

$$\mathsf{BA}(Hf, B)(\alpha_A) = \mathsf{BA}(HA, f)(\alpha_B), \tag{3.1}$$

where $\mathsf{BA}(\_, B) : \mathsf{BA}^{\mathrm{op}} \to \mathsf{Set}$ and $\mathsf{BA}(HA, \_) : \mathsf{BA} \to \mathsf{Set}$ are defined as usual (cf. Def. A.1.5). We want to show that $f : (A, \phi_{A,A}(\alpha_A)) \to (B, \phi_{B,B}(\alpha_B))$ is a MPF-morphism, i.e. that $\phi_{B,B}(\alpha_B) \circ f = f \circ \phi_{A,A}(\alpha_A)$, i.e. we have to show that

$$\mathsf{BA}_\wedge(f, B)(\phi_{B,B}(\alpha_B)) = \mathsf{BA}_\wedge(A, f)(\phi_{A,A}(\alpha_A)). \tag{3.2}$$

where again $\mathsf{BA}_\wedge(\_, B)$ and $\mathsf{BA}_\wedge(A, \_)$ are defined as in Def. A.1.5. The following calculation shows that equation (3.2) holds and therefore that Φ is well-defined on morphisms.

$$\mathsf{BA}_\wedge(f, B)(\phi_{B,B}(\alpha_B)) \overset{\text{naturality of } \phi}{=} \phi_{A,B}(\mathsf{BA}(Hf, B)(\alpha_A))$$

$$\stackrel{\text{equ. (3.1)}}{=} \quad \phi_{A,B}(\mathsf{BA}(HA, f)(\alpha_B))$$
$$\stackrel{\text{naturality of } \phi}{=} \quad \mathsf{BA}_\wedge(A, f)(\phi_{A,A}(\alpha_A))$$

For proving that Φ is indeed a category isomorphism one can easily define a functor $\Phi^{-1} : \mathsf{MPF} \to \mathsf{Alg}(H)$ and show that Φ^{-1} is the inverse of Φ. QED

3.2.13. Remark. Spelling out the construction of a free Boolean algebra over its underlying meet-semilattice, we see that, given a Boolean algebra $\mathbb{A} = (A, \vee, \wedge, -, \bot, \top)$, $H\mathbb{A}$ is the free Boolean algebra generated by the set $\{\Box a \mid a \in A\}$ (the insertion of generators being $\Box : \mathbb{A} \to H\mathbb{A}$, $a \mapsto \Box a$) and satisfying the equations $\Box\top = \top$, $\Box(a \wedge b) = \Box a \wedge \Box b$. That is, the functor H describes how to obtain modal logic by adding an operator to Boolean logic. As observed above this functor is the Stone dual of the Vietoris functor. This observation was made earlier by Abramsky in [Abr88] and is an instance of the general relationship between syntax and semantics as laid out in his domain theory in logical form [Abr91].

We will use an abstract version of the construction of $H\mathbb{A}$ in Chapter 4 where we are going to define the algebraic semantics for an arbitrary coalgebraic modal logic in terms of a category of algebras for a functor $L : \mathsf{BA} \to \mathsf{BA}$.

As another corollary to the duality we obtain that $\mathsf{Coalg}(\mathbb{V})$ has cofree coalgebras.

3.2.14. Corollary. *The forgetful functor* $\mathsf{Coalg}(\mathbb{V}) \to \mathsf{Stone}$ *has a right adjoint.*

Proof. We only sketch the proof. Consider the forgetful functors $R : \mathsf{MA} \to \mathsf{BA}$, $U : \mathsf{MA} \to \mathsf{Set}$, $V : \mathsf{BA} \to \mathsf{Set}$ as depicted in the following diagram:

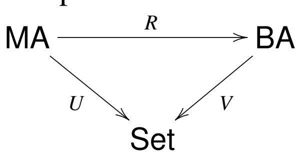

It is well-known that U and V are monadic functors (cf. e.g. [Mac71, Theorem VI.8.1]) and that the category MA has coequalizers (cf. e.g. [Man76, Lemma 3.1.31]). From this we can deduce that R has a left adjoint using Theorem 3.1.29 of [Man76]. Hence, by duality, $\mathsf{Coalg}(\mathbb{V}) \to \mathsf{Stone}$ has a right adjoint. QED

Finally, we show see how arbitrary general frames can be seen as coalgebras.

3.2.15. Remark. **(General Frames as Coalgebras)** Stone spaces provide a convenient framework to study descriptive general frames since the admissible sets can be recovered from the topology: each Stone space $\mathbb{X} = (X, \tau)$ has a *unique* basis that is closed under the Boolean operations. In order to be able to represent arbitrary general frames as coalgebras, we have to make two adjustments.

First, we work directly with admissible sets instead of with topologies: the category RBA (referential or represented Boolean algebras) has objects (X, A) where X is a set and A a set of subsets of X closed under Boolean operations. It has morphisms $f : (X, A) \to (Y, B)$ where f is a function $X \to Y$ such that $f^{-1}(b) \in A$ for all $b \in B$.

And second, in the absence of tightness, the relation of the general frame will no longer be point-closed. Hence, its coalgebraic version has the full power set as its codomain. For $\mathbb{X} = (X, A) \in \mathsf{RBA}$ let $\mathbb{W}\mathbb{X} = (\mathcal{P}(X), v_{\mathbb{X}})$ where $v_{\mathbb{X}}$ is the Boolean algebra generated by $\{\{F \in \mathcal{P}X \mid F \subseteq a\} \mid a \in A\}$. On morphisms let $\mathbb{W}f = \mathcal{P}f$. This clearly defines an endofunctor on the category RBA, and the induced category $\mathsf{Coalg}(\mathbb{W})$ is the coalgebraic version of general frames:

$$\text{There is an isomorphism between } \mathsf{GF} \text{ and } \mathsf{Coalg}(\mathbb{W}). \tag{3.3}$$

The crucial observation in the proof of (3.3) is that, for $\mathbb{X} = (X, A) \in \mathsf{RBA}$ and R a relation on X, we have that A is closed under $[R]$ iff $R[\_] : X \to \mathcal{P}X$ is a RBA-morphism $\mathbb{X} \to \mathbb{W}(\mathbb{X})$. This follows from the fact that $[R](a) = (R[\_])^{-1}(\{F \in \mathcal{P}X \mid F \subseteq a\})$.

The idea of using referential algebras instead of topological spaces for a coalgebraic representation of general frame semantics was further exploited in [KP04]. There the authors define a generalized version of the Vietoris construction which is applicable to referential algebras of arbitrary so-called selfextensional logics.

3.3 Vietoris polynomial functors

3.3.1 Definitions

In this section we introduce the notion of a Vietoris polynomial functor (*VPF*) as a natural Stone-analogue of the Kripke polynomial functors from Section 2.1.1 on Set. This section can therefore be seen as a first application of the observation that coalgebras over Stone can be used as semantics for coalgebraic modal logics. Much of the work in this section consists of transferring the work by Jacobs in [Jac01] to the topological setting. After introducing the Vietoris polynomial functors, we define, for each *VPF* $\mathbb{T}$, the category $\mathsf{BAO}_{\mathbb{T}}$ of $\mathbb{T}$-sorted Boolean algebras with operators and their morphisms.

3.3.1. Definition. The collection of *Vietoris polynomial functors* , in brief: *VPF*s, over Stone is inductively defined as follows:

$$\mathbb{T} ::= \mathbb{I} \mid \mathbb{A} \mid \mathbb{T} + \mathbb{T} \mid \mathbb{T} \times \mathbb{T} \mid \mathbb{T}^D \mid \mathbb{V}\mathbb{T}.$$

Here $\mathbb{I}$ is the identity functor on the category Stone; $\mathbb{A}$ denotes a finite Stone space with the discrete topology (that is, the functor $\mathbb{A}$ is a constant functor); '+' and '×' denote disjoint union (binary coproduct) and binary product in Stone, respectively (cf. Fact 3.1.5); and, for an arbitrary set D, $\mathbb{T}^D$ denotes the functor sending a Stone space $\mathbb{X}$ to $\mathbb{T}(\mathbb{X})^D$, the D-fold product of $\mathbb{T}\mathbb{X}$.

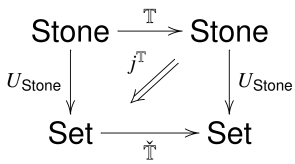

Figure 3.3: The corresponding Kripke polynomial functor $\check{\mathbb{T}}$

Intuitively we can assign to every Kripke polynomial functor $T : \mathsf{Set} \to \mathsf{Set}$ a corresponding *VPF* $\mathbb{T}$ on Stone. The next definition makes this intuition precise.

3.3.2. Definition. For every *KPF* T we define the *corresponding VPF* $\widehat{T}$ by induction:

$$\begin{array}{lcllcl} \widehat{Id} & := & \mathbb{I} & \widehat{A} & := & \mathbb{A} \\ \widehat{T_1 + T_2} & := & \widehat{T}_1 + \widehat{T}_2 & \widehat{T_1 \times T_2} & := & \widehat{T}_1 \times \widehat{T}_2 \\ \widehat{T^D} & := & \widehat{T}^D & \widehat{\mathcal{P}T} & := & \mathbb{V}\widehat{T}. \end{array}$$

The inverse of this translation assigns to a *VPF* $\mathbb{T}$ his corresponding Kripke polynomial functor denoted by $\check{\mathbb{T}}$.

This correspondence between *KPF*s and *VPF*s is witnessed by a family of natural embeddings.

3.3.3. Definition. Let $\mathbb{T} \in VPF$, $T = \check{\mathbb{T}}$ and $\mathbb{X} = (X, \tau) \in \mathsf{Stone}$. Then we define an embedding $j^{\mathbb{T}}_{\mathbb{X}} : U_{\mathsf{Stone}}\mathbb{T}\mathbb{X} \to T U_{\mathsf{Stone}}\mathbb{X}$ by induction on the structure of $\mathbb{T}$:

$$\begin{array}{lcl} j^{\mathbb{I}}_{\mathbb{X}} & := & id_{U_{\mathsf{Stone}}\mathbb{X}} \\ j^{\mathbb{A}}_{\mathbb{X}} & := & id_{U_{\mathsf{Stone}}\mathbb{A}} \\ j^{\mathbb{T}^D}_{\mathbb{X}} & := & (j^{\mathbb{T}}_{\mathbb{X}})^D \\ j^{\mathbb{T}_1+\mathbb{T}_2}_{\mathbb{X}} & := & [j^{\mathbb{T}_1}_{\mathbb{X}}, j^{\mathbb{T}_2}_{\mathbb{X}}] \\ j^{\mathbb{T}_1\times\mathbb{T}_2}_{\mathbb{X}} & := & \left\langle j^{\mathbb{T}_1}_{\mathbb{X}}, j^{\mathbb{T}_2}_{\mathbb{X}} \right\rangle \\ j^{\mathbb{V}\mathbb{T}}_{\mathbb{X}} & := & (\mathcal{P} j^{\mathbb{T}}_{\mathbb{X}})_{\restriction U_{\mathsf{Stone}}\mathbb{V}\mathbb{T}\mathbb{X}} \end{array}$$

where $[\_, \_]$ and $\langle \_, \_ \rangle$ denote cotupling and pairing of functions respectively, and for a function $f : Y \to Y'$ the function f^D maps a function $g : D \to Y$ to the function $f \circ g : D \to Y'$.

With the help of $j^{\mathbb{T}}$ it is now easy to transform $\mathbb{T}$-coalgebras into $\check{\mathbb{T}}$-coalgebras.

3.3.4. Proposition. *Let $\mathbb{T} \in VPF$. Then the family of embeddings $(j^{\mathbb{T}}_{\mathbb{X}})_{\mathbb{X} \in \mathsf{Stone}}$ gives rise to a injective natural transformation*

$$j^{\mathbb{T}} : U_{\mathsf{Stone}} \circ \mathbb{T} \Rightarrow \check{\mathbb{T}} \circ U_{\mathsf{Stone}} \quad \text{(cf. Figure 3.3).}$$

Furthermore the following definition gives rise to a functor

$$\begin{array}{rcl} K : \mathsf{Coalg}(\mathbb{T}) & \to & \mathsf{Coalg}(\check{\mathbb{T}}) \\ (\mathbb{X}, \gamma) & \mapsto & (U_{\mathsf{Stone}}\mathbb{X}, j^{\mathbb{T}}_{\mathbb{X}} \circ U_{\mathsf{Stone}}\gamma) \\ f : \mathbb{X} \to \mathbb{Y} & \mapsto & U_{\mathsf{Stone}} f. \end{array}$$

Proof. The naturality of $j^{\mathbb{T}}$ is easy to check. The fact that K is well-defined on morphisms is then a direct consequence of the naturality of $j^{\mathbb{T}}$. QED

In Definition 2.1.3 we defined paths and path constructors and used those paths to define for every functor $T \in KPF$ a category of ingredients. The category $\mathsf{Ing}(\mathbb{T})$ of ingredients of a Vietoris functor $\mathbb{T}$ is defined analogously.

3.3.5. Definition. For every path p and all $\mathbb{T}_1, \mathbb{T}_2 \in VPF$ we define

$$p : \mathbb{T}_1 \rightsquigarrow \mathbb{T}_2 \quad \text{if} \quad p : \check{\mathbb{T}}_1 \rightsquigarrow \check{\mathbb{T}}_2.$$

Furthermore we define $\mathsf{Ing}(\mathbb{T})$ to be the category with $\mathbf{Ing}(\mathbb{T}) := \{\mathbb{T}' \mid \exists p.p : \mathbb{T} \rightsquigarrow \mathbb{T}'\} \cup \{\mathbb{I}\}$ as set of objects and paths as morphisms between them.

The logic $\mathrm{MSM}\mathcal{L}_{\mathbb{T}}$ associated with a $\mathbb{T} \in VPF$ is the same as the logic associated with the corresponding KPF T. The coalgebraic semantics of $\mathrm{MSM}\mathcal{L}$ given in Section 2.1.2 can be easily adjusted to the Stone-based setting.

3.3.6. Definition. Let $\mathbb{T} \in VPF$, $T := \check{\mathbb{T}} \in KPF$ the corresponding Set-functor and let $(\mathbb{X}, \gamma) \in \mathsf{Coalg}(\mathbb{T})$. Then the *logic* $\mathrm{MSM}\mathcal{L}_{\mathbb{T}}$ *associated with* $\mathbb{T}$ is identical to the logic $\mathrm{MSM}\mathcal{L}_T$: for every ingredient $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ we define the set of formulas $\mathbf{Form}_{\mathbb{S}} := \mathbf{Form}_{\check{\mathbb{S}}}$ and a derivability predicate $\vdash_{\mathbb{S}} := \vdash_{\check{\mathbb{S}}} \subseteq \mathbf{Form}_{\mathbb{S}}$. The logic $\mathrm{MSM}\mathcal{L}_{\mathbb{T}}$ is then given as the pair

$$\Big((\mathbf{Form}_{\mathbb{S}})_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})}, (\vdash_{\mathbb{S}})_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})}\Big).$$

Furthermore we define for every $\mathbb{S}_1, \mathbb{S}_2 \in \mathbf{Ing}(\mathbb{T})$ and every path $p : \mathbb{S}_1 \rightsquigarrow \mathbb{S}_2$ a predicate lifting

$$(\_)^p : \mathrm{Clp}(\mathbb{S}_2\mathbb{X}) \to \mathrm{Clp}(\mathbb{S}_1\mathbb{X}).$$

These liftings are obtained by restricting the liftings $(\_)^p : \mathcal{P}(\check{\mathbb{S}}_2 X) \to \mathcal{P}(\check{\mathbb{S}}_1 X)$ from Definition 2.1.6 to clopen subsets and changing the clause for $(\_)^{\mathsf{pow}\cdot p}$ into

$$\alpha^{\mathsf{pow}\cdot p} := [\ni]\alpha^p.$$

For every $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ the coalgebraic semantics $[\![\_]\!]^{\mathbb{S}}_{(\mathbb{X},\gamma)} : \mathbf{Form}_{\mathbb{S}} \to \mathrm{Clp}(\mathbb{S}\mathbb{X})$ is then defined exactly like in Definition 2.1.8 for the corresponding sort $\check{\mathbb{S}} \in \mathbf{Ing}(T)$.

We will now define *the Boolean algebras with operators associated with a VPF* and then see how they provide an algebraic semantics for the logic $\mathrm{MSM}\mathcal{L}$. The definition of a so-called $\mathbb{T}$-BAO may look slightly involved, but it is based on a simple generalisation of the concept of a modal algebra. The generalisation is that instead of dealing with a single Boolean algebra, we will be working with a *family* $(\Phi(S))_{S\in\mathbf{Ing}(\mathbb{T})}$ of Boolean algebras. As before, we let $\mathsf{BA}_{\wedge}$ denote the category with Boolean algebras as objects and finite-meet preserving functions as morphisms.

3.3.7. Definition. **($\mathbb{T}$-BAO)** Let $\mathbb{T}$ be a VPF. A *$\mathbb{T}$-sorted Boolean algebra with operators*, $\mathbb{T}$-BAO, consists of a functor $\Phi : \mathsf{Ing}(\mathbb{T})^{\mathrm{op}} \longrightarrow \mathsf{BA}_\wedge$, together with an additional map $\mathsf{next} : \Phi(\mathbb{T}) \to \Phi(\mathbb{I})$ which preserves all Boolean operations. This functor is required to meet the conditions

1. $\Phi(\mathbb{A}) = \mathrm{Clp}(\mathbb{A})$,
2. the functions $\Phi(\pi_i)$ and $\Phi(\mathrm{ev}(d))$ are Boolean homomorphisms, and
3. the functions $\Phi(\kappa_i)$ induced by the injection paths satisfy
   (a) $\Phi(\kappa_1)(\bot) = \neg\Phi(\kappa_2)(\bot)$ and
   (b) $\neg\Phi(\kappa_i)(\bot) \leq (\Phi(\kappa_i)(\neg\alpha) \leftrightarrow \neg\Phi(\kappa_i)(\alpha))$.

3.3.8. Example. Let $\mathfrak{A} = (\mathbb{A}, g)$ be a modal algebra, cf. Definition 3.1.9. This algebra can be represented by two different $\mathbb{VI}$-BAOs. This functor has two ingredients $\mathbf{Ing}(\mathbb{VI}) = \{\mathbb{I}, \mathbb{VI}\}$ and one non-trivial path $\mathsf{pow} : \mathbb{VI} \rightsquigarrow \mathbb{I}$.

1. $\Phi(\mathbb{I}) := \mathbb{A}$, $\Phi(\mathbb{VI}) := \mathbb{A}$, $\Phi(\mathbb{V}) := g$, and $\mathsf{next} = id$.
2. $\Phi'(\mathbb{I}) := \mathbb{A}$, $\Phi'(\mathbb{VI}) := H\mathbb{A}$ (cf. Proposition 3.2.12), $\Phi'(\mathbb{V}) : \Phi'(\mathbb{I}) \hookrightarrow \Phi'(\mathbb{VI})$ the (meet-preserving) inclusion of generators, and next' the unique Boolean algebra morphism satisfying $\mathsf{next}' \circ \Phi'(\mathbb{V}) = g$.

We will see that (Φ', next') is the $\mathbb{VI}$-BAO obtained by considering the algebra (Φ, next) from 1. as a $\mathbb{VI}$-coalgebra and translating it back to an algebra, that is, in the notation we are about to introduce, $(\Phi', \mathsf{next}') = \mathcal{AC}(\Phi, \mathsf{next})$.

Another important example of an $\mathbb{T}$-BAO is the initial $\mathbb{T}$-BAO or Lindenbaum $\mathbb{T}$-BAO.

3.3.9. Definition. Let $\mathbb{T} \in VPF$. Then we define the *Lindenbaum* $\mathbb{T}$-BAO

$$\mathcal{L}_{\mathbb{T}} : \mathsf{Ing}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BA}_\wedge.$$

We put $\mathcal{L}_{\mathbb{T}}(\mathbb{S}) := \mathbf{Form}_{\mathbb{S}}/{\equiv_{\mathbb{S}}}$ for $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ and we define for each $p \in PCons$ (cf. Def. 2.1.3) s.t. $p : \mathbb{S}_1 \rightsquigarrow \mathbb{S}_2 \in \mathsf{Ing}(\mathbb{T})$ a function

$$\begin{aligned} \mathcal{L}_{\mathbb{T}}(p) : \mathcal{L}_{\mathbb{T}}(\mathbb{S}_2) &\to \mathcal{L}_{\mathbb{T}}(\mathbb{S}_1) \\ \phi &\mapsto [p]\phi. \end{aligned}$$

Here $\equiv_{\mathbb{S}}$ denotes the equivalence relation on $\mathbf{Form}_{\mathbb{S}}$ defined by

$$\phi_1 \equiv_{\mathbb{S}} \phi_2 \quad \text{if } \vdash_{\mathbb{S}} \phi_1 \leftrightarrow \phi_2.$$

Morphisms of $\mathbb{T}$-BAOs are families of Boolean homomorphisms satisfying certain additional conditions.

3.3.10. Definition. ($\mathsf{BAO}_{\mathbb{T}}$) A *morphism* between $\mathbb{T}$-BAOs $(\Phi', \mathsf{next}') \to (\Phi, \mathsf{next})$ is a natural transformation $t : \Phi' \to \Phi$ such that for each ingredient $\mathbb{S}$ of $\mathbb{T}$ the component $t_{\mathbb{S}} : \Phi'(\mathbb{S}) \to \Phi(\mathbb{S})$ preserves the Boolean structure, $t_{\mathbb{A}} = \mathsf{id}_{\mathrm{Clp}(\mathbb{A})}$ for all constants $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$, and $t_{\mathbb{I}}$ and $t_{\mathbb{T}}$ satisfy $\mathsf{next} \circ t_{\mathbb{T}} = t_{\mathbb{I}} \circ \mathsf{next}'$. This yields the category $\mathsf{BAO}_{\mathbb{T}}$.

3.3.11. Remark. To see that this definition is a natural generalization of the notion of a morphism between modal algebras let us look at the morphisms in the category $\mathsf{BAO}_{\mathbb{VI}}$. If we represent two modal algebras $(\mathbb{A}_1, g_1)$ and $(\mathbb{A}_1, g_2)$ as described in Example 3.3.8 (1) as $\mathbb{VI}$-BAOs $(\Phi_1, \mathsf{next}_1), (\Phi_2, \mathsf{next}_2)$ it is easy to verify that a Boolean morphism $f : \mathbb{A}_1 \to \mathbb{A}_2$ is an MA-morphism iff the natural transformation t given by $t_{\mathbb{I}} = t_{\mathbb{VI}} = f$ is a $\mathsf{BAO}_{\mathbb{VI}}$ morphism.

The following proposition states that $\mathcal{L}_{\mathbb{T}}$ is the initial object in $\mathsf{BAO}_{\mathbb{T}}$.

3.3.12. Proposition. *Let $\mathbb{T} \in VPF$ and $\mathcal{L}_{\mathbb{T}}$ the corresponding Lindenbaum algebra. Then for every $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$ there is a unique $\mathbb{T}$-BAO morphism $\mathsf{i}^{(\Phi,\mathsf{next})} : \mathcal{L}_{\mathbb{T}} \to (\Phi, \mathsf{next})$.*

Proof. By a standard argument. QED

The algebraic semantics of MSM$\mathcal{L}$ is defined in a similar fashion as the algebraic semantics of normal modal logic in Definition 3.1.10.

3.3.13. Definition. Let $\mathbb{T} \in VPF$ and $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$. Then for each $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ we define an interpretation function $[\![\_]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} : \mathbf{Form}_{\mathbb{S}} \to \Phi(\mathbb{S})$ by letting

$$
\begin{array}{lcl}
[\![\bot]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} & := & \bot, \\
[\![\phi_1 \to \phi_2]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} & := & \neg[\![\phi_1]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} \vee [\![\phi_2]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})}, \\
[\![[p]\phi]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} & := & \Phi(p)([\![\phi]\!]^{\mathbb{S}'}_{(\Phi,\mathsf{next})}) \quad \text{for} \quad p : \mathbb{S} \rightsquigarrow \mathbb{S}' \in \mathsf{Ing}(\mathbb{T}) \text{ and } p \in PCons, \\
[\![\mathbf{next}\ \phi]\!]^{\mathbb{I}}_{(\Phi,\mathsf{next})} & := & \mathsf{next}([\![\phi]\!]^{\mathbb{T}}_{(\Phi,\mathsf{next})}).
\end{array}
$$

It is a matter of routine checking to see that $[\![\_]\!]_{(\Phi,\mathsf{next})}$ factors through $\mathsf{i}^{(\Phi,\mathsf{next})}$ (cf. Prop. 3.3.12) as follows:

$$
\begin{array}{ccc}
\mathbf{Form}_{\mathbb{S}} & \xrightarrow{[\![\_]\!]^{\mathbb{S}}} & \Phi(\mathbb{S}) \\
& \searrow^{\mathsf{kan}} \quad \nearrow_{\mathsf{i}^{(\Phi,\mathsf{next})}_{\mathbb{S}}} & \\
& \mathcal{L}_{\mathbb{T}}(\mathbb{S}) &
\end{array}
\tag{3.4}
$$

Here kan denotes the canonical morphism mapping a formula to its $\equiv_{\mathbb{S}}$ -equivalence class in $\mathcal{L}_{\mathbb{T}}$. Therefore the algebraic semantics of a formula ϕ of sort $\mathbb{S}$ is the image of its $\equiv_{\mathbb{S}}$ -equivalence class under the initial map i^{Φ}. Using this observation we can deduce that $\mathbb{T}$-BAOs provide an adequate semantics for MSM$\mathcal{L}_{\mathbb{T}}$.

3.3.14. Proposition. *Let $\mathbb{T} \in VPF$ and $\phi \in \mathbf{Form}_{\mathbb{S}}$ for some $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$. Then*

$$\vdash_{\mathbb{S}} \phi \quad \textit{iff} \quad \forall (\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}. [\![\phi]\!]_{(\Phi,\mathsf{next})} = \top.$$

Proof. Let $\phi \in \mathbf{Form}_{\mathbb{S}}$ for some $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$. Suppose first that $\vdash_{\mathbb{S}} \phi$ and consider an arbitrary $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$. Then $\phi \equiv_{\mathbb{S}} \top$ and therefore $\mathsf{kan}(\phi) = \mathsf{kan}(\top)$. By the commutativity of Diagram 3.4 we get

$$\begin{aligned} [\![\phi]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} &= \mathsf{i}^{(\Phi,\mathsf{next})}_{\mathbb{S}}(\mathsf{kan}(\phi)) \\ &= \mathsf{i}^{(\Phi,\mathsf{next})}_{\mathbb{S}}(\mathsf{kan}(\top)) \\ &= [\![\top]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} = \top. \end{aligned}$$

Suppose on the other hand that $[\![\phi]\!]^{\mathbb{S}}_{(\Phi,\mathsf{next})} = \top$ for all $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$. Then in particular $[\![\phi]\!]^{\mathbb{S}}_{\mathcal{L}_{\mathbb{T}}} = \top = [\![\top]\!]^{\mathbb{S}}_{\mathcal{L}_{\mathbb{T}}}$. Together with Diagram 3.4 this implies $id_{\mathcal{L}_{\mathbb{T}}}(\mathsf{kan}(\phi)) = id_{\mathcal{L}_{\mathbb{T}}}(\mathsf{kan}(\top))$ and hence $\mathsf{kan}(\phi) = \mathsf{kan}(\top)$, i.e. $\vdash_{\mathbb{S}} \phi$. QED

3.3.2 Linking algebraic and coalgebraic semantics

The aim of this section is to establish a link between the categories $\mathsf{BAO}_{\mathbb{T}}$ and $\mathsf{Coalg}(\mathbb{T})$ by functors $\mathcal{A} : \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BAO}_{\mathbb{T}}$ and $C : \mathsf{BAO}_{\mathbb{T}} \to \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$ (cf. Prop. 3.3.22 below), i.e. between the coalgebraic and algebraic semantics of MSM$\mathcal{L}$.

It is not difficult to transform a $\mathbb{T}$-coalgebra into a $\mathbb{T}$-BAO; basically, we are dealing with a sorted version of Stone duality, together with a path-indexed predicate lifting.

3.3.15. Lemma and Definition. *($\mathcal{A}$) For each Vietoris polynomial functor $\mathbb{T}$, each $\mathbb{T}$-coalgebra $(\mathbb{X}, \gamma)$ gives rise to a $\mathbb{T}$-BAO, namely, the 'complex algebra'[2] functor $\mathcal{A}(\mathbb{X}, \gamma)$: $\mathbf{Ing}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BA}_{\wedge}$ given by*

$$\begin{aligned} \mathbb{S} &\mapsto \mathrm{Clp}(\mathbb{S}\mathbb{X}) \\ (p : \mathbb{S}_1 \rightsquigarrow \mathbb{S}_2) &\mapsto ((\_)^p : \mathrm{Clp}(\mathbb{S}_2\mathbb{X}) \to \mathrm{Clp}(\mathbb{S}_1\mathbb{X}))\,, \end{aligned}$$

accompanied by the map $\mathsf{next} : \mathrm{Clp}(\mathbb{T}\mathbb{X}) \to \mathrm{Clp}(\mathbb{X})$ *given by* $\mathsf{next} := \gamma^{-1}$.

Proof. The claim can be proven by spelling out the definitions. QED

This transformation preserves the MSM$\mathcal{L}$-semantics.

3.3.16. Proposition. *Let $\mathbb{T} \in VPF$, $(\mathbb{X}, \gamma) \in \mathsf{Coalg}(\mathbb{T})$ and $\phi \in \mathrm{MSM}\mathcal{L}_{\mathbb{T}}$ a formula of sort $\mathbb{S}$. Then*

$$[\![\phi]\!]^{\mathbb{S}}_{(\mathbb{X},\gamma)} = [\![\phi]\!]^{\mathbb{S}}_{\mathcal{A}(\mathbb{X},\gamma)}.$$

Proof. By spelling out the definitions. QED

[2]The name 'complex algebra' stems from the tradition in modal logic, cf. [BdV01, Chapter 5]

Conversely, with each $\mathbb{T}$-BAO (Φ, next) we want to associate an $\mathbb{T}$-coalgebra $C(\Phi, \mathsf{next})$. Assume that $\mathbb{T}$ has the identity functor as an ingredient; given our results in the previous section, it seems fairly obvious that we should take the dual Stone space $\mathbb{S}\mathrm{p}\Phi(\mathbb{I})$ as the carrier of this dual coalgebra. It remains to define a $\mathbb{T}$-coalgebra structure on this. Applying duality theory to the Boolean algebras obtained from Φ only seems to provide information on the spaces $\mathbb{S}\mathrm{p}\Phi(\mathbb{S})$, whereas we need to work with $\mathbb{S}(\mathbb{S}\mathrm{p}\Phi(\mathbb{I}))$ in order to define a $\mathbb{T}$-coalgebra. Fortunately, in the next lemma and definition we show that there exists a map r which produces the $\mathbb{T}$-structure. The definition of r is as in [Jac01]; what we have to show is that it works also in the topological setting.

3.3.17. Lemma and Definition. *Let $\mathbb{T}$ be a VPF and let (Φ, next) be a $\mathbb{T}$-BAO. Then the following definition by induction on the structure of ingredient functors of $\mathbb{T}$*

$$
\begin{array}{rcl}
r_\Phi(\mathbb{I})(u) & := & u \\
r_\Phi(\mathbb{A})(u) & := & (\epsilon_{\mathbb{A}})^{-1} \quad \textit{(cf. Thm. 3.1.3)} \\
r_\Phi(\mathbb{S}_1 \times \mathbb{S}_2)(u) & := & \left\langle r_\Phi(\mathbb{S}_1)(\Phi(\pi_1)^{-1}(u)), r_\Phi(\mathbb{S}_2)(\Phi(\pi_2)^{-1}(u)) \right\rangle \\
r_\Phi(\mathbb{S}_1 + \mathbb{S}_2)(u) & := & \left\{ \begin{array}{l} \kappa_1 r_\Phi(\mathbb{S}_1)(\Phi(\kappa_1)^{-1}(u)) \textit{ if } \neg\Phi(\kappa_1)(\bot) \in u \\ \kappa_2 r_\Phi(\mathbb{S}_2)(\Phi(\kappa_2)^{-1}(u)) \textit{ if } \neg\Phi(\kappa_2)(\bot) \in u \end{array} \right. \\
r_\Phi(\mathbb{S}^D)(u) & := & \lambda d \in D.\, r_\Phi(\mathbb{S})(\Phi(\mathrm{ev}(d))^{-1}(u)) \\
r_\Phi(\mathbb{V}\mathbb{S})(u) & := & \left\{ r_\Phi(\mathbb{S})(v) \mid v \in \mathsf{Uf}\Phi(\mathbb{S}) \textit{ and } \Phi(\mathsf{pow})^{-1}(u) \subseteq v \right\}
\end{array}
$$

defines, for every $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ a continuous map

$$r_\Phi(\mathbb{S}) : \mathbb{S}\mathrm{p}\Phi(\mathbb{S}) \longrightarrow \mathbb{S}\mathbb{S}\mathrm{p}\Phi(\mathbb{I}).$$

Proof. Let $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$. Both claims (i.e. the one on well-definedness and the one on the continuity of $r_\Phi(\mathbb{S})$) are proven simultaneously by induction on $\mathbb{S}$.

We only consider the case of the Vietoris functor: assume that $\mathbb{S} = \mathbb{V}\mathbb{S}'$. In order to show that $r_\Phi(\mathbb{S})$ is well-defined, take an arbitrary $u \in \mathsf{Uf}\Phi(\mathbb{V}\mathbb{S}')$ and consider the set $F := \left\{ v \mid v \in \mathsf{Uf}\Phi(\mathbb{S}') \text{ and } \Phi(\mathsf{pow})^{-1}(u) \subseteq v \right\}$. F is closed in $\mathsf{Uf}\Phi(\mathbb{S}')$), because for any $v' \in \mathsf{Uf}\Phi(\mathbb{S}') \setminus F$ there is an $a \in \Phi(\mathsf{pow})^{-1}(u)$ such that $a \notin v'$, whence $F \subseteq \hat{a}$ and $v' \notin \hat{a}$: for every $v' \notin F$ we can find an open set containing v' and disjoint from F. But from F being closed and the inductive hypothesis on $r_\Phi(\mathbb{S}')$ it follows that $r_\Phi(\mathbb{S}')[F]$ is closed as well, so by definition, $r_\Phi(\mathbb{S})(u) = r_\Phi(\mathbb{S}')[F]$ belongs to $\mathrm{Cl}(\mathbb{S}'(\mathsf{Uf}\Phi(\mathbb{I})))$. This proves that $r_\Phi(\mathbb{S})$ is well-defined.

We now turn to the continuity of $r_\Phi(\mathbb{S})$. It suffices to show that for an arbitrary clopen set $U \subseteq \mathbb{S}'(\mathsf{Uf}\Phi(\mathbb{I}))$, all sets of the form $r_\Phi(\mathbb{V}\mathbb{S}')^{-1}([\ni]U)$ and $r_\Phi(\mathbb{V}\mathbb{S}')^{-1}(\langle\ni\rangle U)$ are clopen. We only consider sets of the first kind:

$$
\begin{array}{rcl}
r_\Phi(\mathbb{V}\mathbb{S}')^{-1}([\ni](U)) & = & \{u \in \mathsf{Uf}\Phi(\mathbb{V}\mathbb{S}') \mid r_\Phi(\mathbb{V}\mathbb{S}')(u) \in [\ni]U\} \\
& = & \left\{ u \mid \{r_\Phi(\mathbb{S}')(v) \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq U \right\} \\
& = & \left\{ u \mid \{v \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq r_\Phi(\mathbb{S}')^{-1}(U) \right\}
\end{array}
$$

According to the induction hypothesis, $r_\Phi(\mathbb{S}')^{-1}(U)$ is a clopen set, say with $b \in \Phi(\mathbb{S}')$ such that $r_\Phi(\mathbb{S}')^{-1}(U) = \hat{b}$. This leads us to

$$\begin{aligned} r_\Phi(\mathbb{V}\mathbb{S}')^{-1}([\ni]U) &= \left\{u \mid \{v \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq \hat{b}\right\} \\ &= \left\{u \mid \forall v \in \mathsf{Uf}\Phi(\mathbb{S}'). \left(\Phi(\mathsf{pow})^{-1}(u) \subseteq v \rightarrow b \in v\right)\right\} \\ &\overset{(!)}{=} \{u \mid \Phi(\mathsf{pow})(b) \in u\} \end{aligned}$$

which shows that $r_\Phi(\mathbb{V}\mathbb{S}')^{-1}([\ni]U)$ is clopen. It remains to show that the equality (!) indeed holds.
$\supseteq$: trivial.
$\subseteq$: Let $u' \in \mathsf{Uf}\Phi(\mathbb{V}\mathbb{S}')$ and suppose $\Phi(\mathsf{pow})(b) \notin u'$. We will show that under this assumption there exists a $v' \in \mathsf{Uf}\Phi(\mathbb{S}')$ such that $\Phi(\mathsf{pow})^{-1}(u') \subseteq v'$ and $b \notin v'$. As an intermediate step we prove that the set $\Phi(\mathsf{pow})^{-1}(u') \cup \{\neg b\}$ has the finite intersection property. Suppose for a contradiction that there are $a_1, \ldots, a_n \in \Phi(\mathsf{pow})^{-1}(u')$ such that

$$\bigwedge_{1 \leq i \leq n} a_i \wedge \neg b = \perp .$$

Then we have $\bigwedge_{1 \leq i \leq n} a_i \leq b$ and therefore we get by monotonicity of $\Phi(\mathsf{pow})$

$$u' \ni \Phi(\mathsf{pow})\left(\bigwedge_{1 \leq i \leq n} a_i\right) \leq \Phi(\mathsf{pow})(b).$$

As u' is an ultrafilter we can conclude that $\Phi(\mathsf{pow})(b) \in u'$, which contradicts our first assumption. This means that the set $\Phi(\mathsf{pow})^{-1}(u') \cup \{\neg b\}$ has the finite intersection property and is contained in an ultrafilter v'. Clearly $b \notin v'$ and hence $u' \notin \left\{u \mid \forall v \in \mathsf{Uf}\Phi(\mathbb{S}'). \left(\Phi(\mathsf{pow})^{-1}(u) \subseteq v \rightarrow b \in v\right)\right\}$. QED

The above lemma allows us to define a $\mathbb{T}$-coalgebra for a given $\mathbb{T}$-BAO.

3.3.18. Definition. Let $\mathbb{T}$ be a *VPF* and let (Φ, next) be a $\mathbb{T}$-BAO. We define the coalgebra $C(\Phi, \mathsf{next})$ as the structure $(\mathbb{S}\mathrm{p}\Phi(\mathbb{I}), r_\Phi(\mathbb{T}) \circ \mathbb{S}\mathrm{p}(\mathsf{next}))$:

$$C(\Phi, \mathsf{next}) := \ \mathbb{S}\mathrm{p}\Phi(\mathbb{I}) \xrightarrow{\mathbb{S}\mathrm{p}(\mathsf{next})} \mathbb{S}\mathrm{p}\Phi(\mathbb{T}) \xrightarrow{r_\Phi(\mathbb{T})} \mathbb{T}\ \mathbb{S}\mathrm{p}\Phi(\mathbb{I})$$

Again we can relate the semantics of a formula on (Φ, next) to the semantics of a formula on $C(\Phi, \mathsf{next})$.

3.3.19. Proposition. *Let $\mathbb{T} \in VPF$, $(\Phi, \mathsf{next}) \in \mathsf{BAO}_\mathbb{T}$ and $\phi \in \mathbb{S}$ for some $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$. Then*

$$[\![\phi]\!]_{(\Phi,\mathsf{next})} \in u \in \mathsf{Uf}\Phi(\mathbb{S}) \quad \textit{iff} \quad r_\Phi(\mathbb{S})(u) \in [\![\phi]\!]_{C(\Phi,\mathsf{next})}.$$

The maps $\mathcal{A}$ and C from $\mathsf{BAO}_\mathbb{T}$ to $\mathsf{Coalg}(\mathbb{T})$ and back can be extended to morphisms. This will be done in the following two definitions.

3.3.20. Lemma and Definition. *Let $\mathbb{T} \in VPF$ and $f : (\mathbb{X}, \gamma) \to (\mathbb{X}', \gamma') \in \mathsf{Coalg}(\mathbb{T})$. Then we define for every $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ a map $\mathcal{A}(f)(\mathbb{S}) : \mathcal{A}(\mathbb{X}', \gamma')(\mathbb{S}) \to \mathcal{A}(\mathbb{X}, \gamma)(\mathbb{S})$ by letting $\mathcal{A}(f)(\mathbb{S}) := \mathbb{C}\mathrm{lp}(\mathbb{S}f)$. The family $\mathcal{A}(f) := (\mathcal{A}(f)(\mathbb{S}))_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})}$ is a $\mathsf{BAO}_{\mathbb{T}}$-morphism from $\mathcal{A}(\mathbb{X}', \gamma')$ to $\mathcal{A}(\mathbb{X}, \gamma)$.*

Proof. Spell out definitions. QED

3.3.21. Lemma and Definition. *Let $\mathbb{T} \in VPF$ and $t : (\Phi, \mathsf{next}) \to (\Phi', \mathsf{next}') \in \mathsf{BAO}_{\mathbb{T}}$. Then the map $C(t) := \mathbb{S}\mathrm{p}t_{\mathbb{I}} : \mathbb{S}\mathrm{p}\Phi'(\mathbb{I}) \to \mathbb{S}\mathrm{p}\Phi(\mathbb{I})$ is a $\mathbb{T}$-coalgebra morphism from $C(\Phi', \mathsf{next}')$ to $C(\Phi, \mathsf{next})$.*

Proof. The proof is not difficult and similar to the proof of Proposition 5.3 in [Jac01]. QED

3.3.22. Proposition. *If we extend $\mathcal{A}$ and C as described in Definitions 3.3.20 and 3.3.21 we obtain functors*

$$\mathcal{A} : \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BAO}_{\mathbb{T}} \quad \textit{and} \quad C : \mathsf{BAO}_{\mathbb{T}} \to \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}.$$

3.4 Duality between $\mathsf{BAO}_{\mathbb{T}}$ and $\mathsf{Coalg}(\mathbb{V})$

In the previous section we encountered the functors $\mathcal{A} : \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BAO}_{\mathbb{T}}$ and $C : \mathsf{BAO}_{\mathbb{T}} \to \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$. Here we will study these functors in more detail, and show that in fact they provide an adjunction between the categories $\mathsf{BAO}_{\mathbb{T}}$ and $\mathsf{Coalg}(\mathbb{T})$. We will define two families of morphisms, $\alpha_{\Phi} : \mathcal{A}C(\Phi, \mathsf{next}) \to (\Phi, \mathsf{next})$ in BAO_T, and $\xi_{(\mathbb{X},\gamma)} : (\mathbb{X}, \gamma) \to C\mathcal{A}(\mathbb{X}, \gamma)$ in $\mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$, and prove that these are the counit and unit witnessing the fact that $\mathcal{A}$ is left adjoint to C. Since the ξ's will turn out to be isomorphisms, this will then show that $\mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$ is (isomorphic to) a full coreflective subcategory of $\mathsf{BAO}_{\mathbb{T}}$.

In contrast to the classical case of the duality $\mathsf{MA} \simeq \mathsf{DGF}^{\mathrm{op}}$, we do not obtain a dual equivalence between $\mathsf{BAO}_{\mathbb{T}}$ and $\mathsf{Coalg}(\mathbb{T})$. This is due to the fact, which the reader might have noticed already, that the axiomatic definition of $\mathbb{T}$-BAOs does not force a $\mathbb{T}$-BAO Φ to respect $\mathbb{T}$-structure. We take a closer look at this, characterizing the largest full subcategory of $\mathsf{BAO}_{\mathbb{T}}$ on which the adjunction restricts to an equivalence. By showing that the initial algebra of $\mathsf{BAO}_{\mathbb{T}}$ is *exact*, that is, belongs to this subcategory, we obtain the final $\mathbb{T}$-coalgebra as its dual.

We start by proving that every $\mathbb{T}$-coalgebra has an 'ultrafilter representation': it is isomorphic to its double dual. Recall from Section 3.1.1 that for a Stone space $\mathbb{Y}$, $\epsilon_{\mathbb{Y}} : \mathbb{Y} \to \mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}\mathbb{Y}$ denotes the homeomorphism fixed by $\epsilon_{\mathbb{Y}}(y) := \{a \in \mathrm{Clp}(\mathbb{Y}) \mid y \in a\}$.

3.4.1. Theorem. *Let $\mathbb{T}$ be a Vietoris polynomial functor, and let $(\mathbb{X}, \gamma)$ be an $\mathbb{T}$-coalgebra. Then the map $\epsilon_{\mathbb{X}} : \mathbb{X} \to \mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}\mathbb{X}$ is a $\mathsf{Coalg}(\mathbb{T})$-isomorphism witnessing that*

$$(\mathbb{X}, \gamma) \cong C(\mathcal{A}(\mathbb{X}, \gamma)).$$

Proof. We first show that for each sort $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ the following diagram commutes:

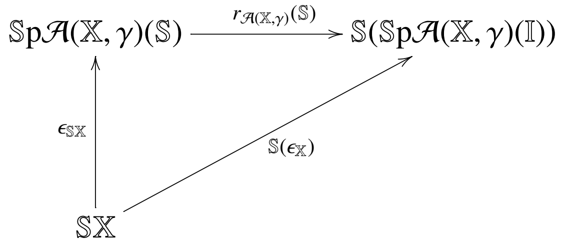

The proof is by induction on $\mathbb{S}$. We will only treat the Vietoris functor, since all other cases work exactly as in the proof of Lemma 5.6 in [Jac01]. In order to prove the commutativity of the above diagram for $\mathbb{S} = \mathbb{V}\mathbb{S}'$, take an arbitrary $F \in \mathbb{V}\mathbb{S}'(\mathbb{X})$. Then, unraveling the definitions of r, $\mathcal{A}$ and of $(\cdot)^{\mathsf{pow}}$, we find

$$\begin{aligned}
r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S})(\epsilon_{\mathbb{S}\mathbb{X}}(F)) &= r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')\left[\left\{v \mid \mathcal{A}(\mathbb{X},\gamma)(\mathsf{pow})^{-1}(\epsilon_{\mathbb{S}\mathbb{X}}(F)) \subseteq v\right\}\right] \\
&= r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')\left[\{v \mid \{\alpha \in \mathrm{Clp}(\mathbb{S}'\mathbb{X}) \mid (\alpha)^{\mathsf{pow}} \in \epsilon_{\mathbb{S}\mathbb{X}}(F)\} \subseteq v\}\right] \\
&= r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')\left[\{v \mid \{\alpha \in \mathrm{Clp}(\mathbb{S}'\mathbb{X}) \mid F \subseteq \alpha\} \subseteq v\}\right] \\
&\overset{(!)}{=} \{\mathbb{S}'(\epsilon_{\mathbb{X}})(x) \mid x \in F\} \\
&= \mathbb{S}(\epsilon_{\mathbb{X}})(F).
\end{aligned}$$

It is left to prove (!). For ($\supseteq$), take an arbitrary $x \in F$, and define $v_x := \epsilon_{\mathbb{S}'\mathbb{X}}(x)$. Then for all $a \in \mathrm{Clp}(\mathbb{S}'\mathbb{X})$ it holds that $F \subseteq a$ implies $x \in a$, which is equivalent to $a \in \epsilon_{\mathbb{X}}(x) = v_x$; in other words,$v_x$ satisfies the condition $\{a \mid F \subseteq a\} \subseteq v_x$. Also, by the inductive hypothesis we have that $\mathbb{S}'(\epsilon_{\mathbb{X}})(x) = r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')(\epsilon_{\mathbb{S}'\mathbb{X}}(x))$. Taking these observations together we see that $\mathbb{S}'(\epsilon_{\mathbb{X}})(x) \in r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')\left[\{v \mid \{a \in \mathrm{Clp}(\mathbb{S}'\mathbb{X}) \mid F \subseteq a\} \subseteq v\}\right]$.

For ($\subseteq$), let $v \in \mathsf{Uf}\mathcal{A}(\mathbb{X},\gamma)(\mathbb{S}')$ be such that $\{a \in \mathrm{Clp}(\mathbb{S}'\mathbb{X}) \mid F \subseteq a\} \subseteq v$. By Stone duality we know that $\bigcap_{a\in v} a = \{x\}$ for exactly one $x \in \mathbb{S}'\mathbb{X}$. This x must be an element of F, because $\bigcap_{a\in v} a \subseteq \bigcap\{a \mid F \subseteq a\} = F$ and we get $\epsilon_{\mathbb{S}'\mathbb{X}}(x) = v$. By the induction hypothesis this is the same as saying $r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{S}')(v) = \mathbb{S}'(\epsilon_{\mathbb{X}})(v)$, which proves the inclusion.

Now we proceed to prove the theorem: we calculate

$$\begin{aligned}
C(\mathcal{A}(\mathbb{X},\gamma)) \circ \epsilon_{\mathbb{X}} &= (r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{T}) \circ \mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}(\gamma)) \circ \epsilon_{\mathbb{X}} &&= r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{T}) \circ (\mathbb{S}\mathrm{p}\mathbb{C}\mathrm{lp}(\gamma) \circ \epsilon_{\mathbb{X}}) \\
&= r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{T}) \circ (\epsilon_{\mathbb{T}\mathbb{X}} \circ \gamma) &&= (r_{\mathcal{A}(\mathbb{X},\gamma)}(\mathbb{T}) \circ \epsilon_{\mathbb{T}\mathbb{X}}) \circ \gamma,
\end{aligned}$$

where the third step is by naturality of ϵ. Now by commutativity of the above diagram for $\mathbb{T}$ we find that $C(\mathcal{A}(\mathbb{X},\gamma)) \circ \epsilon_{\mathbb{X}} = \mathbb{T}(\epsilon_{\mathbb{X}}) \circ \gamma$, which is nothing but stating that $\epsilon_{\mathbb{X}}$ is a coalgebra homomorphism. But then since $\epsilon_{\mathbb{X}}$ is an isomorphism between Stone spaces we may conclude that it is also an isomorphism between the two given coalgebras. QED

The functor C is not faithful in general; however, when it comes to morphisms having a complex algebra $\mathcal{A}(\mathbb{X},\gamma)$ as their domain, we can prove the following.

3.4.2. Proposition. *Let* $(\mathbb{X},\gamma)$ *be a* T*-coalgebra and* (Φ,next) *be a* $\mathbb{T}$*-BAO. Furthermore let* $s, s' : \mathcal{A}(\mathbb{X},\gamma) \to (\Phi,\mathsf{next})$ *be morphisms in* $\mathsf{BAO}_{\mathbb{T}}$*. Then* $C(s) = C(s')$ *implies* $s = s'$.

Proof. Let $(\mathbb{X}, \gamma)$, (Φ, next), s and s' be as in the statement of the Proposition, and assume that $C(s) = C(s')$. Then it is clear that we have $s_{\mathbb{I}} = s'_{\mathbb{I}}$. With the help of Lemma 3.4.3 below we therefore get $s = s'$. QED

The following lemma, which forms the heart of the proof of Proposition 3.4.2, is stated separately because we need it again further on.

3.4.3. LEMMA. *Let $(\mathbb{X}, \gamma)$ be a $\mathbb{T}$-coalgebra and (Φ, next) a $\mathbb{T}$-BAO. Furthermore let $s, s' : \mathcal{A}(\mathbb{X}, \gamma) \to \Phi$ be natural transformations whose components preserve all the Boolean structure, and such that $s_{\mathbb{I}} = s'_{\mathbb{I}}$ and $s_{\mathbb{A}} = s'_{\mathbb{A}}$ for all constants $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$. Then $s = s'$.*

Proof. Assume that we have two natural transformations $s, s' : \mathcal{A}(\mathbb{X}, \gamma) \to (\Phi, \mathsf{next})$ as required in the lemma. In order to prove that $s = s'$, it suffices to show that

$$s_{\mathbb{S}} = v'_{\mathbb{S}} \text{ for all } \mathbb{S} \in \mathbf{Ing}(\mathbb{T}). \tag{3.5}$$

We will prove (3.5) by induction on $\mathbb{S}$. In the base case ($\mathbb{S} = \mathbb{I}$ or $\mathbb{S} = \mathbb{A}$ for some constant functor $\mathbb{A}$), it follows immediately that $s_{\mathbb{S}} = s'_{\mathbb{S}}$.

For the inductive step of the proof, we confine ourselves to a rough sketch of the proof idea. In each case, in order to show that $s_{\mathbb{S}}(a) = s_{\mathbb{S}'}(a)$ for every clopen a of $\mathbb{S}\mathbb{X}$, we try and find a clopen subbasis $\mathcal{B}$ such that $s_{\mathbb{S}}(b) = s'_{\mathbb{S}}(b)$ for all subbasic b. For instance, in the case that $\mathbb{S} = \mathbb{V}\mathbb{S}'$, put

$$\mathcal{B} := \{b \mid b \in (\_)^{\mathsf{pow}}[\mathrm{Clp}(\mathbb{S}'\mathbb{X})]\} \cup \{-b \mid b \in (\_)^{\mathsf{pow}}[\mathrm{Clp}(\mathbb{S}'\mathbb{X})]\},$$

and let $b \in \mathcal{B}$. Then one can easily check that we have $s_{\mathbb{S}}(b) = s'_{\mathbb{S}}(b)$ for all $b \in \mathcal{B}$ and by the fact that $\mathcal{B}$ is a clopen subbasis of the Vietoris topology one can use a straightforward argument to show that $s_{\mathbb{S}} = s'_{\mathbb{S}}$. QED

We are now ready to show that the functors $C : \mathsf{BAO}_{\mathbb{T}} \to \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$ and $\mathcal{A} : \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BAO}_{\mathbb{T}}$ form a so-called dual representation. That is, C is right adjoint to $\mathcal{A}$ and the unit of the adjunction is an isomorphism. We first define the unit ξ and the counit α of the adjunction. Recall that we proved in Theorem 3.4.1 that ξ is an isomorphism; for r_{Φ} see Definition 3.3.17 and for $\iota_{\Phi(S)}$ Theorem 3.1.3.

3.4.4. DEFINITION. **(α, ξ)** For a $\mathbb{T}$-BAO (Φ, next) and a $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ we define

$$\alpha_{(\Phi,\mathsf{next})} : \mathcal{A}C(\Phi, \mathsf{next}) \to (\Phi, \mathsf{next})$$

via $\alpha_{\Phi}(\mathbb{S}) := v_{\Phi(\mathbb{S})} \circ \mathbb{C}\mathrm{lp}(r_{\Phi}(\mathbb{S}))$, where $v_{\Phi(\mathbb{S})}$ denotes the inverse of the isomorphism $\iota_{\Phi(\mathbb{S})} : \Phi(\mathbb{S}) \to \mathbb{C}\mathrm{lp}\mathbb{S}\mathrm{p}\Phi(\mathbb{S})$. For a $\mathbb{T}$-coalgebra $(\mathbb{X}, \gamma)$, we define

$$\xi_{(\mathbb{X},\gamma)} : (\mathbb{X}, \gamma) \to C\mathcal{A}(\mathbb{X}, \gamma) \quad \text{in } \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$$

as the inverse $\xi_{(\mathbb{X},\gamma)} : C\mathcal{A}(\mathbb{X}, \gamma) \to (\mathbb{X}, \gamma)$ of the morphism $\epsilon_{(\mathbb{X},\gamma)} : (\mathbb{X}, \gamma) \to C\mathcal{A}(\mathbb{X}, \gamma)$ in $\mathsf{Coalg}(\mathbb{T})$.

Intuitively, the next theorem establishes a duality between $\mathsf{Coalg}(\mathbb{T})$ and $\mathsf{BAO}_{\mathbb{T}}$ in which every coalgebra $(\mathbb{X},\gamma)$ can be represented in a canonical way by the algebra $\mathcal{A}(\mathbb{X},\gamma)$.

3.4.5. THEOREM. *Let* $\mathbb{T}$ *be a VPF. Then* $\mathcal{A} : \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BAO}_{\mathbb{T}}$ *is a full embedding and has* $C : \mathsf{BAO}_{\mathbb{T}} \to \mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$ *as a right adjoint with* ξ *and* α *as unit and counit. That is,* $\mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$ *is (isomorphic to) a full coreflective subcategory of* $\mathsf{BAO}_{\mathbb{T}}$.

Before we turn to the proof of this theorem, we first show that α is indeed a morphism of $\mathbb{T}$-BAOs.

3.4.6. LEMMA. *The family of maps* $\alpha_{(\Phi,\mathsf{next})}(\_) : \mathcal{A}C(\Phi,\mathsf{next}) \to (\Phi,\mathsf{next})$ *is a morphism of* $\mathbb{T}$*-BAOs.*

Proof. We have to show that $\alpha_{(\Phi,\mathsf{next})}(\_)$ is a natural transformation and that $\alpha_{(\Phi,\mathsf{next})}(\_)$ fulfills an additional naturality condition with respect to the next-operator.

Concerning the first claim we must prove that for all $p : \mathbb{S} \rightsquigarrow \mathbb{S}'$ in $\mathsf{Ing}(\mathbb{T})$ we have

$$\Phi(p) \circ \alpha_{(\Phi,\mathsf{next})}(\mathbb{S}') = \alpha_{(\Phi,\mathsf{next})}(\mathbb{S}) \circ (\_)^{p}.$$

It suffices to show, by a case distinction, that this equation holds for paths of length at most one. As all of these proofs boil down to a tedious but straightforward unraveling of definitions, we confine ourselves to the case that $p = \mathsf{pow}$ and $\mathbb{S} = \mathbb{V}Gs'$. Take an arbitrary $U \in \mathrm{Clp}(\mathbb{S}'\mathbb{S}\mathrm{p}\Phi(\mathbb{I}))$ and let $a \in \Phi(\mathbb{S}')$ be such that $\mathbb{C}\mathrm{lp}(r_\Phi(\mathbb{S}))(U) = \hat{a}$. Then

$$\begin{aligned}
\alpha_{(\Phi,\mathsf{next})}(\mathbb{S})((U)^{\mathsf{pow}}) &= (\upsilon_{\Phi(\mathbb{S})} \circ \mathbb{C}\mathrm{lp}(r_{\Phi(\mathbb{S})}))((U)^{\mathsf{pow}}) \\
&= (\upsilon_{\Phi(\mathbb{S})} \circ r_\Phi(\mathbb{S})^{-1})(\{F \subseteq U \mid F \subseteq \mathbb{S}'\mathbb{S}\mathrm{p}\Phi(\mathbb{I}) \text{ closed}\}) \\
&= \upsilon_{\Phi(\mathbb{S})}(\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid r_\Phi(\mathbb{S})(u) \subseteq U\}) \\
&= \upsilon_{\Phi(\mathbb{S})}\left(\left\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid \{r_\Phi(\mathbb{S}')(v) \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq U\right\}\right) \\
&= \upsilon_{\Phi(\mathbb{S})}\left(\left\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid \{v \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq \mathbb{C}\mathrm{lp}(r_\Phi(S_1))(U)\right\}\right) \\
&= \upsilon_{\Phi(\mathbb{S})}\left(\left\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid \{v \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v\} \subseteq \hat{a}\right\}\right) \\
&= \upsilon_{\Phi(\mathbb{S})}\left(\left\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid \Phi(\mathsf{pow})^{-1}(u) \subseteq v \Rightarrow a \in v\right\}\right) \\
&= \upsilon_{\Phi(\mathbb{S})}(\{u \in \mathsf{Uf}\Phi(\mathbb{S}) \mid \Phi(\mathsf{pow})(a) \in u\}) \\
&= \Phi(\mathsf{pow})(a) \\
&= \Phi(\mathsf{pow})\,(\upsilon_{\Phi(\mathbb{S}')} \circ \mathbb{C}\mathrm{lp}(r_\Phi(\mathbb{S}'))(U)) \\
&= (\Phi(\mathsf{pow}) \circ \alpha_{(\Phi,\mathsf{next})}(\mathbb{S}'))\,(U)
\end{aligned}$$

and we get $\alpha_{(\Phi,\mathsf{next})}(\mathbb{S}) \circ (\_)^{\mathsf{pow}} = \Phi(\mathsf{pow}) \circ \alpha_{(\Phi,\mathsf{next})}(\mathbb{S}')$, as required.

Now we turn to the second claim. The 'additional naturality condition with respect to the next-operator' is the following: $\mathsf{next} \circ \alpha_{(\Phi,\mathsf{next})}(\mathbb{T}) = \alpha_{(\Phi,\mathsf{next})}(\mathbb{I}) \circ \mathbb{C}\mathrm{lp}(r_\Phi(\mathbb{T}) \circ \mathbb{S}\mathrm{p}\mathsf{next})$. This is easily shown to hold (the second identity being due the naturality of υ).

$$\begin{aligned}
\alpha_{(\Phi,\mathsf{next})}(\mathbb{I}) \circ \mathbb{C}\mathrm{lp}(r_\Phi(\mathbb{T}) \circ \mathbb{S}\mathrm{p}(\mathsf{next})) &= \upsilon_{\Phi(\mathbb{I})} \circ \mathbb{C}\mathrm{lp}(\mathbb{S}\mathrm{p}(\mathsf{next})) \circ \mathbb{C}\mathrm{lp}(r_{\Phi(\mathbb{T})}) \\
&= \mathsf{next} \circ \upsilon_{\Phi(\mathbb{T})} \circ \mathbb{C}\mathrm{lp}(r_\Phi(\mathbb{T})) \\
&= \mathsf{next} \circ \alpha_{(\Phi,\mathsf{next})}(\mathbb{T}).
\end{aligned}$$

QED

Proof of Theorem 3.4.5. For the adjunction it suffices to show ([Mac71], p. 81) that for all $(\mathbb{X}, \gamma) \in \mathsf{Stone}$ and for all $f : C(\Phi, \mathsf{next}) \to (\mathbb{X}, \gamma)$ there is a unique $s : \mathcal{A}(\mathbb{X}, \gamma) \to (\Phi, \mathsf{next})$ such that the following diagram in $\mathsf{Coalg}(\mathbb{T})$ commutes:

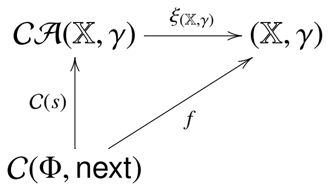

Indeed, defining $s = \alpha_{(\Phi,\mathsf{next})} \circ \mathcal{A}(f)$, we calculate

$$\begin{aligned}
\xi_{(\mathbb{X},\gamma)} \circ C(\alpha_{(\Phi,\mathsf{next})} \circ \mathcal{A}(f)) &= \xi_{(\mathbb{X},\gamma)} \circ \mathbb{S}\mathrm{p}\,(\alpha_{(\Phi,\mathsf{next})}(\mathbb{I}) \circ \mathcal{A}(f)(\mathbb{I})) \\
&= \xi_{(\mathbb{X},\gamma)} \circ \mathbb{S}\mathrm{p}(\upsilon_{\Phi(\mathbb{I})} \circ r_{\Phi}(\mathbb{I}) \circ \mathbb{C}\mathrm{lp}(f)) \\
&= \xi_{(\mathbb{X},\gamma)} \circ \mathbb{S}\mathrm{p}(\mathbb{C}\mathrm{lp}(f)) \circ \mathbb{S}\mathrm{p}(\upsilon_{\Phi(\mathbb{I})}) \\
&= f \circ \epsilon^{-1}_{\mathbb{S}\mathrm{p}(\Phi(\mathbb{I}))} \circ \mathbb{S}\mathrm{p}(\upsilon_{\Phi(\mathbb{I})}) \\
&= f
\end{aligned}$$

The last two steps use the fact that $\mathbb{S}\mathrm{p}$ and $\mathbb{C}\mathrm{lp}$ are adjoint with (co)units υ and ϵ, see Theorem 3.1.3 and Definition 3.4.4. Uniqueness of s is Proposition 3.4.2. To conclude the proof, recall that a left-adjoint is full and faithful iff the unit is an isomorphism ([Mac71], p. 88). Hence $\mathcal{A}$ is full and faithful by Theorem 3.4.1. QED

We now turn to a characterization of the largest subcategory of $\mathsf{BAO}_{\mathbb{T}}$ on which the adjunction from Theorem 3.4.5 restricts to a dual equivalence. The reader might have noticed already that our adjunction is not a dual equivalence since the definition of $\mathbb{T}$-BAOs does not force a $\mathbb{T}$-BAO (Φ, next) to respect $\mathbb{T}$-structure. For example, if $\mathbb{S}_1 \times \mathbb{S}_2$ is an ingredient of $\mathbb{T}$ then it may well be that $\Phi(\mathbb{S}_1 \times \mathbb{S}_2) \neq \Phi(\mathbb{S}_1) + \Phi(\mathbb{S}_2)$.

3.4.7. Definition. Let $\mathbb{S}$ be a functor $\mathsf{Stone} \to \mathsf{Stone}$. Then

$$\mathbb{S}^{\partial} := \mathbb{C}\mathrm{lp} \circ \mathbb{S} \circ \mathbb{S}\mathrm{p}.$$

defines a corresponding functor $\mathbb{S}^{\partial}$ on the category BA.

The following definition introduces *exact* $\mathbb{T}$-BAOs, that is, those $\mathbb{T}$-BAOs which do respect $\mathbb{T}$-structure.

3.4.8. Definition. **(exact $\mathbb{T}$-BAO)** A $\mathbb{T}$-BAO (Φ, next) is called *exact* if there is a $\mathbf{Ing}(\mathbb{S})$-indexed family of isomorphisms

$$\tau_{\mathbb{S}} : \mathbb{S}^{\partial}(\Phi(\mathbb{I})) \to \Phi(\mathbb{S})$$

with the following properties:

- $\tau : (\_)^{\partial}(\Phi(\mathbb{I})) \to \Phi$ is a natural transformation. Here the functor $(\_)^{\partial}(\Phi(\mathbb{I}))$ is defined as follows:

$$\begin{array}{rcl} (\_)^{\partial}(\Phi(\mathbb{I})) : \mathbf{Ing}(T)^{\mathrm{op}} & \to & \mathsf{BA}_{\wedge} \\ \mathbb{S} & \mapsto & \mathbb{S}^{\partial}(\Phi(\mathbb{I})) \\ p : \mathbb{S}_1 \rightsquigarrow \mathbb{S}_2 & \mapsto & (\_)^{p} : \mathbb{S}_2{}^{\partial}(\Phi(\mathbb{I})) \to \mathbb{S}_1{}^{\partial}(\Phi(\mathbb{I})) \end{array}$$

  where $(\_)^{p}$ denotes the predicate lifting from $\mathbb{S}_2{}^{\partial}(\Phi(\mathbb{I})) = \mathbb{C}\mathrm{lp}(\mathbb{S}_2(\mathbb{S}\mathrm{p}\Phi(\mathbb{I})))$ to $\mathbb{S}_1{}^{\partial}(\Phi(\mathbb{I})) = \mathbb{C}\mathrm{lp}(\mathbb{S}_1(\mathbb{S}\mathrm{p}\Phi(\mathbb{I})))$.

- $\tau_{\mathbb{I}} = \upsilon_{\Phi(\mathbb{I})}$, where again $\upsilon_{\Phi(\mathbb{I})}$ denotes the inverse of the isomorphism $\iota_{\Phi(\mathbb{I})} : \Phi(\mathbb{I}) \to \mathbb{C}\mathrm{lp}\mathbb{S}\mathrm{p}\Phi(\mathbb{I})$

- $\tau_{\mathbb{A}} = id_{\mathsf{Clp}(\mathbb{A})}$ for every constant $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$.

$\mathsf{BAO}^{e}_{\mathbb{T}}$ is the full subcategory of $\mathsf{BAO}_{\mathbb{T}}$ consisting of the exact $\mathbb{T}$-BAOs.

We will now see that exact $\mathbb{T}$-BAOs are precisely those $\mathbb{T}$-BAOs (Φ, next) for which the component $\alpha_{(\Phi,\mathsf{next})}$ of the counit of the adjunction is an isomorphism.

3.4.9. Theorem. *Let $\mathbb{T}$ be a VPF. The category $\mathsf{BAO}^{e}_{\mathbb{T}}$ is the largest subcategory of $\mathsf{BAO}_{\mathbb{T}}$ on which the adjunction of Theorem 3.4.5 restricts to a dual equivalence to $\mathsf{Coalg}(\mathbb{T})$.*

Proof. Let B be the largest subcategory of $\mathsf{BAO}_{\mathbb{T}}$ on which the adjunction $\mathcal{A} \dashv C$ restricts to an equivalence. Then for any $(\Phi, \mathsf{next}) \in \mathsf{B}$ the map $\alpha_{(\Phi,\mathsf{next})} : \mathcal{A}C(\Phi, \mathsf{next}) \to (\Phi, \mathsf{next})$ consists of a family of isomorphisms going from $\mathcal{A}C\Phi(\mathbb{S}) = \mathbb{S}^{\partial}(\Phi)(\mathbb{I})$ to $\Phi(\mathbb{S})$. Therefore we can define a family of isomorphisms $\tau_{\mathbb{S}} : \mathbb{S}^{\partial}(\Phi)(\mathbb{I}) \to \Phi(\mathbb{S})$ by letting $\tau = \alpha_{(\Phi,\mathsf{next})}$. It is straightforward to check that this family satisfies the conditions in Definition 3.4.8. Hence $(\Phi, \mathsf{next}) \in \mathsf{BAO}^{e}_{\mathbb{T}}$.

Now let $(\Phi, \mathsf{next}) \in \mathsf{BAO}^{e}_{\mathbb{T}}$. We have to show that the counit $\alpha_{(\Phi,\mathsf{next})}$ is an isomorphism. As $(\Phi, \mathsf{next}) \in \mathsf{BAO}^{e}_{\mathbb{T}}$ there is a family of isomorphisms

$$\tau_{\mathbb{S}} : (\mathcal{A}C\Phi)(\mathbb{S}) \to \Phi(\mathbb{S}).$$

which is natural in $\mathbb{S}$ and for which we have $\tau_{\mathbb{I}} = \upsilon_{\Phi(\mathbb{I})} = \alpha_{(\Phi,\mathsf{next})}(\mathbb{I})$ and $\tau_{\mathbb{A}} = id_{\mathsf{Clp}(\mathbb{A})} = \alpha_{(\Phi,\mathsf{next})}(\mathbb{A})$ for all constants $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$. Using Lemma 3.4.3 one can therefore show that $\tau_{\mathbb{S}} = \alpha_{(\Phi,\mathsf{next})}(\mathbb{S})$ for all $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$. But this means in particular that $\alpha_{(\Phi,\mathsf{next})}$ is an isomorphism. QED

We now show that the final object in $\mathsf{Coalg}(\mathbb{T})$ is obtained as the dual of the initial object in $\mathsf{BAO}_{\mathbb{T}}$. This is a direct consequence of Theorem 3.4.5 and a special case of the more general fact that the right adjoint C preserves colimits of diagrams that take values in the $\mathcal{A}$-image of $\mathsf{Coalg}(\mathbb{T})^{\mathrm{op}}$.

3.4.10. THEOREM. *Let $\mathbb{T}$ be a VPF and $\mathcal{L}_{\mathbb{T}}$ be the initial object in* BAO_T. *Then $C\mathcal{L}_{\mathbb{T}}$ is final in* $\mathsf{Coalg}(\mathbb{T})$.

Proof. We prove the theorem by showing that $\alpha_{\mathcal{L}_{\mathbb{T}}}$ is an isomorphism, i.e. $\mathcal{L}_{\mathbb{T}} \in \mathsf{BAO}^e_{\mathbb{T}}$. Finality of $C\mathcal{L}_{\mathbb{T}}$ follows then immediately from the duality between $\mathsf{Coalg}(\mathbb{T})$ and $\mathsf{BAO}^e_{\mathbb{T}}$.

Since $\mathcal{L}_{\mathbb{T}}$ is initial there is a morphism $\mathsf{i} : \mathcal{L}_{\mathbb{T}} \to \mathcal{A}C\mathcal{L}_{\mathbb{T}}$. Since $id_{\mathcal{L}_{\mathbb{T}}}$ is the unique morphism $\mathcal{L}_{\mathbb{T}} \to \mathcal{L}_{\mathbb{T}}$ it follows that $\alpha_{\mathcal{L}_{\mathbb{T}}} \circ \mathsf{i} = id_{\mathcal{L}_{\mathbb{T}}}$. We want to show that $\mathsf{i} \circ \alpha_{\mathcal{L}_{\mathbb{T}}} : \mathcal{A}C\mathcal{L}_{\mathbb{T}} \to \mathcal{A}C\mathcal{L}_{\mathbb{T}}$ is in fact the identity on $\mathcal{A}C\mathcal{L}_{\mathbb{T}}$. Since $\mathcal{A}$ is full (cf. Theorem 3.4.5) there is $f : C\mathcal{L}_{\mathbb{T}} \to C\mathcal{L}_{\mathbb{T}}$ in $\mathsf{Coalg}(\mathbb{T})$ such that $\mathcal{A}(f) = \mathsf{i} \circ \alpha_{\mathcal{L}_{\mathbb{T}}}$. We obtain $\alpha_{\mathcal{L}_{\mathbb{T}}} \circ \mathcal{A}(f) = \alpha_{\mathcal{L}_{\mathbb{T}}} \circ \mathsf{i} \circ \alpha_{\mathcal{L}_{\mathbb{T}}} = \alpha_{\mathcal{L}_{\mathbb{T}}} = \alpha_{\mathcal{L}_{\mathbb{T}}} \circ \mathcal{A}(id_{C\mathcal{L}_{\mathbb{T}}})$ and the universal property of the coreflection tells us that $f = id_{C\mathcal{L}_{\mathbb{T}}}$, hence, $id_{\mathcal{A}C\mathcal{L}_{\mathbb{T}}} = \mathsf{i} \circ \alpha_{\mathcal{L}_{\mathbb{T}}}$ and $\alpha_{\mathcal{L}_{\mathbb{T}}}$ is iso. QED

As a corollary we obtain completeness of MSM$\mathcal{L}_{\mathbb{T}}$ with respect to the coalgebraic semantics.

3.4.11. COROLLARY (COMPLETENESS OF MSM$\mathcal{L}_{\mathbb{T}}$). *Let $\mathbb{T} \in VPF$ and suppose $\phi \in$ **Form**$_{\mathbb{S}}$ is a formula of* MSM$\mathcal{L}_{\mathbb{T}}$. *Then*

$$\nvdash_{\mathbb{S}} \phi \quad \textit{implies} \quad \exists(\mathbb{X}, \gamma) \in \mathsf{Coalg}(\mathbb{T}) \textit{ s.t. } [\![\neg\phi]\!]^{\mathbb{S}}_{(\mathbb{X},\gamma)} \neq \emptyset.$$

Proof. Suppose $\nvdash_{\mathbb{S}} \phi$, then obviously $[\![\phi]\!]_{\mathcal{L}_{\mathbb{T}}} \neq \top$. Since $\mathcal{L}_{\mathbb{T}}$ was shown to be exact in the proof of the theorem we have $\mathcal{L}_{\mathbb{T}} \cong \mathcal{A}C\mathcal{L}_{\mathbb{T}}$ and therefore $[\![\phi]\!]_{\mathcal{A}C\mathcal{L}_{\mathbb{T}}} \neq \top$. But then by Proposition 3.3.16 we have $[\![\phi]\!]_{C\mathcal{L}_{\mathbb{T}}} \neq \top$, i.e. $\neg\phi$ is satisfiable in $C\mathcal{L}_{\mathbb{T}}$. QED

But completeness with respect to the Stone coalgebra semantics is not exactly what one is interested in. In general we want to specify properties of Set-based systems and thus we are aiming at completeness with respect to Set-based coalgebras. Luckily this is an easy consequence of Corollary 3.4.11.

3.4.12. COROLLARY (COMPLETENESS OF MSM$\mathcal{L}_T$). *Let $T \in KPF$ and suppose $\phi \in$ **Form**$_S$ is a formula of* MSM$\mathcal{L}_T$. *Then*

$$\nvdash_S \phi \quad \textit{implies} \quad \exists(X, \gamma) \in \mathsf{Coalg}(T) \textit{ s.t. } [\![\neg\phi]\!]^{S}_{(X,\gamma)} \neq \emptyset.$$

Proof. Let $\phi \in$ **Form**$_S$ for some $S \in$ **Ing**(T) such that $\nvdash_S \phi$. For all ingredients $S \in$ **Ing**(T) we write $\mathbb{S}$ for the corresponding VPF $\widehat{S}$. Because $\vdash_S \phi$ and $\vdash_S = \vdash_{\mathbb{S}}$ we know that $\nvdash_{\mathbb{S}} \phi$. Therefore there exists a $\mathbb{T}$-coalgebra $(\mathbb{X}, \delta)$ such that $[\![\neg\phi]\!]^{\mathbb{S}}_{(\mathbb{X},\delta)} \neq \emptyset$ by Corollary 3.3.16. This $\mathbb{T}$-coalgebra can be transformed into a T-coalgebra $(X, \gamma) := K(\mathbb{X}, \delta)$ using the functor $K : \mathsf{Coalg}(\mathbb{T}) \to \mathsf{Coalg}(T)$ from Proposition 3.3.4. Spelling out the definition of K it is not difficult to see that $[\![\neg\phi]\!]^{S}_{K(\mathbb{X},\delta)} = [\![\neg\phi]\!]^{\mathbb{S}}_{(\mathbb{X},\delta)} \neq \emptyset$. Hence $[\![\neg\phi]\!]^{S}_{(X\,\gamma)} \neq \emptyset$ and the proof is finished. QED

3.4.13. REMARK. This completeness result was already contained in [Jac01]. The canonical model construction in *loc.cit.* , however, works only for polynomial functors, i.e. for KPF's not containing the power set functor.

3.5 Alternative view: Many-sorted algebras

In this section we are going to present a slightly different view on our results in the previous sections. The algebraic semantics of MSM$\mathcal{L}$ in terms of $\mathbb{T}$-BAOs was first defined by Jacobs in [Jac01] and we sticked to his terminology, because one of the motivations of our work was to improve on Jacobs' results.

Instead of representing $\mathbb{T}$-BAOs as functors $\Phi : \mathsf{Ing}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BA}_\wedge$, however, we could have defined $\mathbb{T}$-BAOs as many-sorted algebras. We will now present this alternative representation. As a result we obtain an algebraic explanation for the definition of r_Φ in 3.3.17 and a characterization of the exact $\mathbb{T}$-BAOs from Definition 3.4.8: the exact $\mathbb{T}$-BAOs will turn out to be those $\mathbb{T}$-BAOs (Φ, next) which are freely generated from the elements in $\Phi(\mathbb{I})$.

3.5.1. Definition. For each $\mathbb{T} \in VPF$ we define the $\mathbb{T}$-*sorted algebraic theory* (Σ, E) as follows: the set of sorts is given by the ingredients $\mathbf{Ing}(\mathbb{T})$. The signature Σ contains

- for each sort $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ the Boolean operations, i.e. for all $\mathbb{S}$ we have $\vee, \wedge \in \Sigma_{\mathbb{S}\mathbb{S},\mathbb{S}}$, $\neg \in \Sigma_{\mathbb{S},\mathbb{S}}$ and $\top, \bot \in \Sigma_{1,\mathbb{S}}$,
- for each $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$ and each $a \in \mathbb{A}$ a constant $a \in \Sigma_{1,\mathbb{A}}$,
- for $\mathbb{S}^D \in \mathbf{Ing}(\mathbb{T})$ and each $d \in D$ an operation $[d, -] \in \Sigma_{\mathbb{S},\mathbb{S}^D}$,
- for $\mathbb{S}_1 \times \mathbb{S}_2 \in \mathbf{Ing}(\mathbb{T})$ an operation $- \otimes - \in \Sigma_{\mathbb{S}_1\mathbb{S}_2,\mathbb{S}_1\times\mathbb{S}_2}$,
- for $\mathbb{S}_1 + \mathbb{S}_2 \in \mathbf{Ing}(\mathbb{T})$ an operation $- \oplus - \in \Sigma_{\mathbb{S}_1\mathbb{S}_2,\mathbb{S}_1+\mathbb{S}_2}$, and
- for each $\mathbb{V}\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ an operation $\Box \in \Sigma_{\mathbb{S},\mathbb{V}\mathbb{S}}$.

The set E consists of all Boolean equations for the Boolean operations of each sort, all the equations that hold in Clp($\mathbb{A}$) for $\mathbb{A} \in \mathbf{Ing}(\mathbb{T})$ and of the following equations, that specify the behaviour of the other operators:

$$
\begin{array}{rcl}
\neg[d, x] & = & [d, \neg x] \\
\bigvee_i [d, x_i] & = & [d, \bigvee_i x_i] \\
\neg(x \oplus y) & = & \neg x \oplus \neg y \\
\bigvee_i (x_i \oplus y_i) & = & \bigvee_i x_i \oplus \bigvee_i y_i \\
(\bigvee_i x_i) \otimes y & = & \bigvee_i (x_i \otimes y) \\
x \otimes (\bigvee_i y_i) & = & \bigvee_i (x \otimes y_i) \\
\neg(x \otimes y) & = & (\neg x \otimes \top) \vee (\top \otimes \neg y) \\
\Box \bigwedge_i x_i & = & \bigwedge_i \Box x_i
\end{array}
$$

where the occurring conjunctions $\bigwedge_i$ and the disjunctions $\bigvee_i$ are required to be finite or empty. In the latter case we have $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \bot$. A (Σ, E)-algebra is called an $\mathbb{T}$-*sorted algebra* ($\mathbb{T}$ – **MAlg**). Furthermore we denote by $\mathsf{MAlg}_\mathbb{T}$ the category with $\mathbb{T}$-sorted algebras as objects and with homomorphisms as arrows.

We will now see that $\mathbb{T}$-BAOs can be seen as $\mathbb{T}$-sorted algebras with an additional **next** -operator. First we show how to transform a $\mathbb{T}$-sorted algebra into a $\mathbb{T}$-BAO.

3.5.2. Lemma and Definition. *Let $\mathbb{T} \in VPF$, $\mathfrak{A} \in \mathsf{MAlg}_{\mathbb{T}}$ with $\mathbb{T}$-sorted carrier set $(A_{\mathbb{S}})_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})}$ and* **next** *$: \mathbb{A}_{\mathbb{T}} \to \mathbb{A}_{\mathbb{I}} \in \mathsf{BA}$, where for $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ we denote by $\mathbb{A}_{\mathbb{S}}$ the Boolean algebra based on $A_{\mathbb{S}}$. We define a functor $\Phi_{\mathfrak{A}} : \mathsf{Ing}(\mathbb{T})^{\mathrm{op}} \to \mathsf{BA}_{\wedge}$ by letting*

- $\Phi(\mathbb{S}) := \mathbb{A}_{\mathbb{S}}$ *for* $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$,
- $\Phi(\mathrm{ev}(d))(b) := [d, b]$ *for* $\mathrm{ev}(d) : \mathbb{S}^D \rightsquigarrow \mathbb{S} \in \mathsf{Ing}(\mathbb{T})$,
- $\Phi(\pi_1)(b) := b \otimes \top$ *and* $\Phi(\pi_2)(b) := \top \otimes b$ *for* $\pi_i : \mathbb{S}_1 \times \mathbb{S}_2 \rightsquigarrow \mathbb{S}_i \in \mathsf{Ing}(\mathbb{T})$,
- $\Phi(\kappa_1)(b) := b \oplus \top$ *and* $\Phi(\kappa_2)(b) := \top \oplus b$ *for* $\kappa_i : \mathbb{S}_1 + \mathbb{S}_2 \rightsquigarrow \mathbb{S}_i \in \mathsf{Ing}(\mathbb{T})$*, and*
- $\Phi(\mathsf{pow})(b) := \Box b$ *for* $\mathsf{pow} : \mathbb{V}\mathbb{S} \rightsquigarrow \mathbb{S} \in \mathsf{Ing}(\mathbb{T})$.

Then the pair $(\Phi_{\mathfrak{A}}, \mathbf{next}\ )$ is an $\mathbb{T}$-BAO. Furthermore if $t : \mathfrak{A}_1 \to \mathfrak{A}_2 \in \mathsf{MAlg}_{\mathbb{T}}$ then $(t_{\restriction A_{\mathbb{S}}})_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})} : \Phi_{\mathfrak{A}_1} \to \Phi_{\mathfrak{A}_2}$ is a natural transformation.

Proof. $\Phi_{\mathfrak{A}}$ is obviously well-defined on objects $\mathbb{S} \in \mathsf{Ing}(\mathbb{T})$. The fact that the arrows $\Phi(p) \in \mathsf{BA}_{\wedge}$ for all $p \in PCons$ satisfy the $\mathbb{T}$-BAO axioms can be deduced using the equations for $\otimes, \oplus, [d, \_]$ and $\Box$ from Definition 3.5.1.

Let $t : \mathfrak{A}_1 \to \mathfrak{A}_2 \in \mathsf{MAlg}_{\mathbb{T}}$. We have to show that $s := (t_{\restriction A_{\mathbb{S}}})_{\mathbb{S}\in\mathbf{Ing}(\mathbb{T})}$ is natural. It suffices to show the naturality of s for path constructors $p \in PCons$. We only consider the case $\mathsf{pow} : \mathbb{V}\mathbb{S} \to \mathbb{S} \in \mathsf{Ing}(\mathbb{T})$, in which we have to prove that the diagram below commutes:

$$\begin{array}{ccc} \Phi_{\mathfrak{A}_1}(\mathbb{V}\mathbb{S}) & \xrightarrow{s_{\mathbb{V}\mathbb{S}}} & \Phi_{\mathfrak{A}_2}(\mathbb{V}\mathbb{S}) \\ \uparrow{\scriptstyle \Phi_{\mathfrak{A}_1}(\mathsf{pow})} & & \uparrow{\scriptstyle \Phi_{\mathfrak{A}_2}(\mathsf{pow})} \\ \Phi_{\mathfrak{A}_1}(\mathbb{S}) & \xrightarrow[s_{\mathbb{S}}]{} & \Phi_{\mathfrak{A}_2}(\mathbb{S}) \end{array}$$

But this follows from $s_{\mathbb{V}\mathbb{S}}(\Phi_{\mathfrak{A}_1}(\mathsf{pow}))(b) = t(\Box b)$ by definition, $t(\Box b) = \Box t(b)$ by $t \in \mathsf{MAlg}_{\mathbb{T}}$ and $\Box t(b) = \Phi_{\mathfrak{A}_2}(\mathsf{pow})(s_{\mathbb{S}}(b))$ by definition. QED

3.5.3. Lemma and Definition. *Let $\mathbb{T} \in VPF$ and $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$. Then the following defines a $\mathbb{T}$-sorted algebra $\mathfrak{A}_{\Phi}$:*

- *the carrier set is $(A_{\mathbb{S}})_{\mathbb{S}\in\mathsf{Ing}(\mathbb{T})}$ where for all $\mathbb{S} \in \mathsf{Ing}(\mathbb{T})$ $A_{\mathbb{S}}$ is the carrier set of $\Phi(\mathbb{S})$,*
- *for each sort $\mathbb{S}$ the Boolean operations on $A_{\mathbb{S}}$ are interpreted as in $\Phi(\mathbb{S})$, and*
- *the other operators are defined as follows:*

$$\begin{array}{lcl} [d, x] & := & \Phi(\mathrm{ev}(d))(x) \\ x_1 \otimes x_2 & := & \Phi(\pi_1)(x_1) \wedge \Phi(\pi_2)(x_2) \\ x_1 \oplus x_2 & := & \Phi(\kappa_1)(x_1) \wedge \Phi(\kappa_2)(x_2) \\ \Box x & := & \Phi(\mathsf{pow})(x). \end{array}$$

Moreover if $t : (\Phi_1, \mathsf{next}_1) \to (\Phi_2, \mathsf{next}_2)$ *is an* $\mathbb{T}$*-BAO morphism, then the map* $s : \mathfrak{A}_{\Phi_1} \to \mathfrak{A}_{\Phi_2}$ *defined by* $s(b) := t_{\mathbb{S}}(b)$ *for* $b \in A_{\mathbb{S}}$ *is an* $\mathsf{MAlg}_{\mathbb{T}}$*-morphism.*

For every functor $\mathbb{T} : \mathsf{Stone} \to \mathsf{Stone}$ there exists by Stone duality a corresponding functor $\mathbb{T}^{\partial} : \mathsf{BA} \to \mathsf{BA}$ defined by $\mathbb{T}^{\partial} := \mathbb{C}\mathrm{lp} \circ \mathbb{T} \circ \mathbb{S}\mathrm{p}$ (cf. Def. 3.4.7). For a $\mathbb{T} \in VPF$ we can, however, give an explicit description of the dual functor on BA. The idea is that for every operation on Stone spaces, such as the product and coproduct, there is a dual operation on Boolean algebras. The next definition describes these dual operations on Boolean algebras which are needed to define the dual of a VPF.

3.5.4. Definition. Let $\mathbb{B}, \mathbb{B}_1, \mathbb{B}_2$ be Boolean algebras. Then we define

$$\begin{aligned}
\mathbb{B}^D &:= \mathsf{BA}\left\langle \{[d,b] \mid d \in D, b \in \mathbb{B}\} | E_{[d,_]} \right\rangle \\
\mathbb{B}_1 \otimes \mathbb{B}_2 &:= \mathsf{BA}\left\langle \{b_1 \otimes b_2 \mid b_i \in \mathbb{B}_i\} | E_{\otimes} \right\rangle \\
\mathbb{B}_1 \oplus \mathbb{B}_2 &:= \mathsf{BA}\left\langle \{b_1 \oplus b_2 \mid b_i \in \mathbb{B}_i\} | E_{\oplus} \right\rangle \\
\mathbb{V}^{\dagger}\mathbb{B} &:= \mathsf{BA}\left\langle \{\Box b \mid b \in \mathbb{B}\} | E_{\Box} \right\rangle
\end{aligned}$$

where $E_{[d,\_]}, E_{\otimes}, E_{\oplus}$ and $E_{\Box}$ are the equations from Definition 3.5.1 for $[d, \_]$, $\otimes$, $\oplus$ and $\Box$ respectively and $\mathsf{BA}\langle G \mid E \rangle$ denotes the Boolean algebra presented by the set of generators G and equations E (cf. appendix). These operations on BA can be easily extended to operations on functors $\mathbb{T} : \mathsf{BA} \to \mathsf{BA}$, e.g. $\mathbb{T}_1 \oplus \mathbb{T}_2$ denotes the functor mapping $\mathbb{B} \in \mathsf{BA}$ to $\mathbb{T}_1\mathbb{B} \oplus \mathbb{T}_2\mathbb{B}$ and a homomorphism $f : \mathbb{B}_1 \to \mathbb{B}_2$ to the homomorphism that maps a generator $b_1 \oplus b_2 \in \mathbb{T}_1\mathbb{B}_1 \oplus \mathbb{T}_2\mathbb{B}_1$ to $f(b_1) \oplus f(b_2) \in \mathbb{T}_1\mathbb{B}_2 \oplus \mathbb{T}_2\mathbb{B}_2$.

3.5.5. Definition. For an $\mathbb{T} \in VPF$ we define the dual functor $\mathbb{T}^{\dagger}$ by induction on the structure of $\mathbb{T}$ as follows:

$$\begin{array}{lllllll}
\mathbb{A}^{\dagger} & := & \mathrm{Clp}(\mathbb{A}) & \quad & \mathbb{I}^{\dagger} & := & Id_{\mathsf{BA}} \\
(\mathbb{T}^D)^{\dagger} & := & (\mathbb{T}^{\dagger})^D & & (\mathbb{T}_1 + \mathbb{T}_2)^{\dagger} & := & {\mathbb{T}_1}^{\dagger} \oplus {\mathbb{T}_2}^{\dagger} \\
(\mathbb{T}_1 \times \mathbb{T}_2)^{\dagger} & := & {\mathbb{T}_1}^{\dagger} \otimes {\mathbb{T}_2}^{\dagger} & & (\mathbb{V}\mathbb{T})^{\dagger} & := & \mathbb{V}^{\dagger} \circ \mathbb{T}^{\dagger}
\end{array}$$

3.5.6. Remark. The definition of $(\mathbb{V}\mathbb{I})^{\dagger}$ is similar to the definition of H in Proposition 3.2.12: the equation in $E_{\Box}$ precisely expresses that the set of generators is a semi-lattice, $\mathbb{V}\mathbb{I}^{\dagger}$ is the free Boolean algebra over this semi-lattice.

The next lemma states that the functor $\mathbb{T}^{\dagger}$ indeed describes the Stone dual of a given functor $\mathbb{T}$.

3.5.7. Lemma. *Let* $\mathbb{T} \in VPF$*. Then* $\mathbb{T}^{\partial} \cong \mathbb{T}^{\dagger}$ *where* $\mathbb{T}^{\partial}$ *is defined as in Definition 3.4.7.*

Proof. This follows from the duality. QED

With the help of $\mathbb{T}^{\dagger}$ we can now describe the $\mathbb{T}$-sorted algebras that are freely generated by a given Boolean algebra.

3.5.8. LEMMA. *Let $\mathbb{T} \in VPF$ and $\mathbb{B} \in \mathsf{BA}$. Then the free $\mathbb{T}$-sorted algebra generated by $\mathbb{B}$ denoted by $\mathbf{F}_{\mathbb{T}}(\mathbb{B})$ is defined as follows:*

- *for all $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ we let $\mathbb{A}_{\mathbb{S}} := \mathbb{S}^{\partial}(\mathbb{B})$, i.e. the carrier set $A_{\mathbb{S}}$ of sort $\mathbb{S}$ is given by the carrier set of $\mathbb{S}^{\dagger}$ and the Boolean operation on $A_{\mathbb{S}}$ are interpreted as in $\mathbb{S}^{\dagger}$,*
- *for $\mathbb{S}^D, \mathbb{S}_1 \times \mathbb{S}_2, \mathbb{S}_1 + \mathbb{S}_2, \mathbb{V}\mathbb{S} \in \mathbf{Ing}(\mathbb{S})$ the operations $[d, \_], \otimes, \oplus$ and $\Box$ are interpreted as the insertion of generators, e.g. $[d, \_]$ maps an element $b \in \mathbb{S}^{\partial}(\mathbb{B})$ to the generator $[d, b] \in (\mathbb{S}^D)^{\partial}(\mathbb{B})$.*

Proof. Let $\mathfrak{A}$ be a $\mathbb{T}$-sorted algebra with sorted carrier set $(A_{\mathbb{S}})_{\mathbb{S} \in \mathbf{Ing}(\mathbb{T})}$ and $f : \mathbb{B} \to \mathbb{A}_{\mathbb{I}}$ a Boolean homomorphism. In order to show that $\mathbf{F}_{\mathbb{T}}(\mathbb{B})$ is free over $\mathbb{B}$ we have to prove that there is a unique morphism $q_{\mathfrak{A}} : \mathbf{F}_{\mathbb{T}}(\mathbb{B}) \to \mathfrak{A}$ such that $(q_{\mathfrak{A}})_{\restriction \mathbb{B}} = f$. First we define for each $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$ a function $q_{\mathbb{S}} : \mathbb{A}_{\mathbb{S}} \to \mathbb{A}'_{\mathbb{S}}$ by induction on the structure of $\mathbb{S} \in \mathbf{Ing}(\mathbb{T})$.

Case: $\mathbb{S} = \overline{\mathbb{A}}$ for some constant functor $\overline{\mathbb{A}}$. Then $q_{\overline{\mathbb{A}}}(a) := a \in$ for all $a \in \overline{\mathbb{A}}$.

Case: $\mathbb{S} = \mathbb{I}$. Then $q_{\mathbb{I}} := f$.

Case: $\mathbb{S} = \mathbb{S}_1 \times \mathbb{S}_2$. Then $q_{\mathbb{S}_1 \times \mathbb{S}_2}$ is defined as the unique BA-homomorphism that extends the mapping $(b_1 \otimes b_2) \mapsto q_{\mathbb{S}_1}(b_1) \otimes q_{\mathbb{S}_2}(b_2)$. This unique extension exists, because $\mathbb{A}_{\mathbb{S}_1 \times \mathbb{S}_2} = \mathsf{BA}\langle \{b_1 \otimes b_2 \mid b_i \in \mathbb{A}_{\mathbb{S}_i}\} | E_{\otimes} \rangle$ and $\mathbb{A}'_{\mathbb{S}_1 \times \mathbb{S}_2}$ satisfies all the equations in $E_{\otimes}$.

The remaining cases of the definition of the $q_{\mathbb{S}}$'s are analogous to the last case. Finally we define

$$\begin{aligned} q_{\mathfrak{A}} : \mathbf{F}_{\mathbb{T}}(\mathbb{B}) &\to \mathfrak{A} \\ a &\mapsto q_{\mathbb{S}}(a) \quad \text{if } a \in \mathbb{A}_{\mathbb{S}}. \end{aligned}$$

It is now not difficult to see that with this definition $q_{\mathfrak{A}}$ is the unique homomorphism from $\mathbf{F}_{\mathbb{T}}(\mathbb{B})$ to $\mathfrak{A}$ with the property required in the lemma. QED

Consider now an arbitrary $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$ and the corresponding $\mathbb{T}$-sorted algebra $\mathfrak{A}_{\Phi}$ (cf. Definition 3.5.3). Then, according to Lemma 3.5.8, there is a unique homomorphism $q_{\mathfrak{A}_{\Phi}} : \mathbf{F}_{\mathbb{T}}(\Phi(\mathbb{I})) \to \mathfrak{A}_{\Phi}$ such that $q_{\mathfrak{A}_{\Phi} \restriction \Phi(\mathbb{I})} = id_{\Phi(\mathbb{I})}$. Spelling out the definition of $\mathbf{F}_{\mathbb{T}}(\Phi(\mathbb{I}))$ it is easy to see that this homomorphism corresponds to a $\mathbf{Ing}(\mathbb{T})$-indexed family of BA-morphisms $q_{\mathbb{S}} : \mathbb{S}^{\dagger}(\Phi(\mathbb{I})) \to \Phi(\mathbb{S})$. Using the isomorphism from Lemma 3.5.7 we get a family of morphisms $\bar{q}_{\mathbb{S}} : \mathbb{S}^{\partial}(\Phi(\mathbb{I})) \to \Phi(\mathbb{S})$. Spelling out the definition one can now easily check that $\bar{q}_{\mathbb{S}} = \mathbb{C}\mathrm{lp}(r_{\Phi}(\mathbb{S}))$, i.e. the $r_{\Phi}(\mathbb{S})$-map from Definition 3.3.17 is the Stone dual of $\bar{q}_{\mathbb{S}}$.

In other words the family $(r_{\Phi}(\mathbb{S}))_{\mathbb{S} \in \mathbf{Ing}(\mathbb{T})}$ corresponds to the dual of the $\mathfrak{A}_{\Phi}$-component of the counit of the adjunction

$$\mathsf{MAlg}_{\mathbb{T}} \underset{U}{\overset{\mathbf{F}_{\mathbb{T}}(\_)}{\leftrightarrows}} \mathsf{BA} \quad (\bot)$$

Here U maps a $\mathbb{T}$-sorted algebra $\mathfrak{A}$ to its $\mathbb{I}$-component $\mathbb{A}_{\mathbb{I}} \in \mathsf{BA}$ or to some fixed (arbitrary) Boolean algebra in case $\mathbb{I} \notin \mathbf{Ing}(\mathbb{T})$.

3.5.9. Proposition. *Let* $\mathbb{T} \in VPF$. *Then* $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$ *is* exact *iff* $\mathfrak{A}_{\Phi} \in \mathbb{T} - \mathbf{MAlg}$ *is freely generated by* $\Phi(\mathbb{I})$.

Proof. If (Φ, next) is exact then by the definition of exactness there is a natural isomorphism $\tau_{\mathbb{S}} : (\_)^{\partial}(\Phi(\mathbb{I})) \to \Phi(\_)$. By Lemma 3.5.7 these isomorphisms give rise to a natural isomorphism $\bar{\tau}_{\mathbb{S}} : (\_)^{\dagger}(\Phi(\mathbb{I})) \to \Phi(\_)$, but this means that $\mathfrak{A}_{\Phi}$ is isomorphic in $\mathsf{MAlg}_{\mathbb{T}}$ to the $\mathbb{T}$-sorted algebra which is freely generated by $\Phi(\mathbb{I})$, i.e. $\mathfrak{A}_{\Phi}$ is freely generated by $\Phi(\mathbb{I})$. The other direction can be proven in a similar way: as $\mathfrak{A}_{\Phi}$ is freely generated by $\Phi(\mathbb{I})$ we can show that there is a natural isomorphism $\bar{\tau}_{\mathbb{S}} : (\_)^{\dagger}(\Phi(\mathbb{I})) \to \Phi(\_)$ that corresponds to a suitable natural isomorphism $\tau_{\mathbb{S}} : (\_)^{\partial}(\Phi(\mathbb{I})) \to \Phi(\_)$. QED

This result can be explained as follows: we know that the exact $\mathbb{T}$-BAOs form a category which is dually equivalent to $\mathsf{Coalg}(\mathbb{T})$. An element of $\mathsf{Coalg}(\mathbb{T})$ is determined by its underlying Stone space $\mathbb{X}$ and its coalgebra map $\gamma : \mathbb{X} \to \mathbb{T}\mathbb{X}$. Dually this means that an exact $\mathbb{T}$-BAO should be determined by the Boolean algebra $\Phi(\mathbb{I})$ and the **next** -operator. This is precisely the case for those $(\Phi, \mathsf{next}) \in \mathsf{BAO}_{\mathbb{T}}$ for which the corresponding $\mathbb{T}$-sorted algebra is freely generated by $\Phi(\mathbb{I})$.

3.6 Conclusions

What we have done so far can be viewed from various distinct perspectives. Here we summarise some of these, indicating possible future research directions.

Stone Coalgebras and Modal Logic Research on the relation between coalgebras and modal logic started with Moss ([Mos99]) although earlier work, e.g. by Rutten ([Rut95]) already showed that Kripke frames and models are instances of coalgebras. Kurz ([Kur01a, Kur00]) showed that modal logic for coalgebras dualises equational logic for algebras, the idea being that equations describe quotients of free algebras and modal formulas describe subsets of final (or cofree) coalgebras. Another account of the duality has been given in [KR02] where it was shown that modalities dualise algebraic operations. But whereas, usually, any quotient of a free algebra can be defined by a set of ordinary equations, one needs *infinitary* modal formulas to define all subsets of a final coalgebra. As a consequence, while we have a satisfactory description of the coalgebraic semantics of infinitary modal logics, we do not completely understand the relationship between coalgebras and finitary modal logic. The results in this chapter show that Stone coalgebras provide a natural and adequate semantics for finitary modal logics, but there is ample room for clarification here.

Another approach to a coalgebraic semantics for finitary modal logics was given in [KP05]. There, the idea is to modify coalgebra morphisms in such a way that they capture not bisimulation but only bisimulation up to rank ω. Since finitary modal logics capture precisely bisimulation up to rank ω, the resulting category Beh_ω provides a convenient framework to study the coalgebraic semantics of finitary modal logic. So an important next step is to understand the relation between both approaches.

Stone Coalgebras as Systems We investigated coalgebras over Stone spaces as models for modal logic. But what is the significance of Stone-coalgebras from the point of view of systems (that is, coalgebras over Set, cf. [Rut00])? What is the relationship between Set-coalgebras and Stone-coalgebras? An interesting observation is here that their notions of bisimilarity coincide. Recall that two elements of two coalgebras are bisimilar iff they can be identified by some coalgebra morphisms. Since Stone coalgebra morphisms have to be continuous, we expect that fewer states are identified under Stone-behavioural equivalence than under Set-behavioural equivalence. But the following holds.

Consider a Vietoris polynomial functor $\mathbb{T} : \mathsf{Stone} \to \mathsf{Stone}$ and its corresponding (Kripke polynomial) functor $\check{\mathbb{T}} : \mathsf{Set} \to \mathsf{Set}$. According to Proposition 3.3.4 there is a functor $K : \mathsf{Coalg}(\mathbb{T}) \to \mathsf{Coalg}(\check{\mathbb{T}})$. Now let $(\mathbb{X}_1, \gamma_1)$, $(\mathbb{X}_2, \gamma_2)$ be two $\mathbb{T}$-coalgebras and x_1, x_2 be two elements in $\mathbb{X}_1$, $\mathbb{X}_2$, respectively. Then $((\mathbb{X}_1, \gamma_1), x_1)$ and $((\mathbb{X}_2, \gamma_2), x_2)$ are behaviourally equivalent iff $(K(\mathbb{X}_1, \gamma_1), x_1)$ and $(K(\mathbb{X}_2, \gamma_2), x_2)$ are behavioural equivalent. — Proof: 'only if' is immediate. The converse follows from the fact that $\check{\mathbb{T}}$-bisimilarity implies that x_1 and x_2 satisfy the same formulas. Therefore x_1 and x_2 get identified by the maps from $(\mathbb{X}_i, \gamma_i)$ into the final $\mathbb{T}$-coalgebra.

Generalising Stone Coalgebras Coalgebras over Stone spaces can be generalised in different ways. We have seen that replacing the topologies by represented Boolean algebras leads to general frames. But it will also be of interest to consider other topological spaces as base categories.

From the point of view of modal logic, it is interesting to investigate the Vietoris functors on other base categories. For example, [Pal04] shows that the Vietoris functor can be defined on Priestley spaces, leading to an adequate semantics for positive modal logic. Recent results by Kurz and Bonsangue in [BK05] generalize both the Vietoris construction on Stone and on Priestley spaces to a setting in which one works with the category of $\mathcal{T}_0$-spaces as base category and the category of so-called observation frames as its algebraic dual equivalent. Also the work by Moss & Viglizzo in [MV04], which has been carried further by Viglizzo in [Vig05], fits in this context. In *loc.cit.* the authors consider polynomial functors over the category of measurable spaces in order to model Harsanyi type spaces, a notion that has its origin in the foundations of game theory.

Another generalization of our work is to move away from normal modal logics to non-normal modal logics, i.e. logics in which the modal operators do not necessarily preserve finite conjunctions. In [HK04] a functor $\mathbf{Up}\mathbb{V}$: $\mathsf{Stone} \to \mathsf{Stone}$ is defined and it is proven that $\mathbf{Up}\mathbb{V}$-coalgebras correspond to the descriptive general neighbourhood frames for monotone modal logic.

From the point of view of the theory of coalgebras, the value of the move from Set to Stone as a base category can be explained as follows. For a functor on Set the notion of behavioural equivalence is, in general, characterised by the whole terminal sequence running through all ordinals. But often, one is interested only in finitary approximations. In the examples considered in this chapter, the move from a functor on Set to its version on Stone has the consequence that the final coalgebra is the limit of the finitary approximants of the terminal sequence (and, therefore, behavioural equivalence is completely characterised by the finitary approximants of the terminal sequence). We expect that this idea of topologising a functor T in order to tailor the behaviour of T-coalgebras to meet a specific notion of observable behaviour will have further applications to universal coalgebra.

Coalgebras and Duality Theory Whereas many, or most, common dualities are induced by a schizophrenic object (see [Joh82, Section VI.4.1]), the duality of modal algebras and descriptive general frames is not. To see why this is so, write D : $\mathsf{MA} \to \mathsf{DGF}$, M : $\mathsf{DGF} \to \mathsf{MA}$ for the contravariant functors witnessing the duality and suppose, for contradiction that there is a schizophrenic object S. That is, assume that $\mathsf{MA}(\mathbb{A}, S) \cong UD(\mathbb{A})$ where U denotes the forgetful functor $\mathsf{DGF} \to \mathsf{Set}$. Then $\mathsf{Set}(1, U\mathbb{G}) \cong U\mathbb{G} \cong UDM\mathbb{G} \cong \mathsf{MA}(M\mathbb{G}, S) \cong \mathsf{DGF}(DS, DM\mathbb{G}) \cong \mathsf{DGF}(DS, \mathbb{G})$, showing that DS is a free object over one generator in DGF. But since DGF-morphisms are also bisimulations it is not hard to see that such an object cannot exist.

For suppose otherwise, i.e. suppose there is an $S \in \mathsf{MA}$ such that for all $\mathbb{G} \in \mathsf{DGF}$ we have $\mathsf{Set}(1, U\mathbb{G}) \cong \mathsf{DGF}(DS, \mathbb{G})$. It is obvious that that DS must contain at least one state. Now let $\mathbb{G}_1 \in \mathsf{DGF}$ be the general frame consisting only of one irreflexive point which we call x_1, and $\mathbb{G}_2 \in \mathsf{DGF}$ be the general frame consisting of one reflexive point x_2. Then $\mathsf{Set}(1, U\mathbb{G}_i)$ has 1 element for $i = 1, 2$ and hence by our assumption, there are continuous bounded morphisms

$$f_i : DS \to \mathbb{G}_i \quad \text{for } i = 1, 2.$$

The graphs $\mathrm{Gr}(f_1), \mathrm{Gr}(f_2)$ of the bounded morphisms f_1 and f_2 are bisimulations and bisimulations are closed under composition (cf. Fact 3.1.22). As a consequence $R := \mathrm{Gr}(f_1) \circ \mathrm{Gr}(f_2)^{\smile}$, where $\mathrm{Gr}(f_2)^{\smile}$ denotes the converse of $\mathrm{Gr}(f_2)$, is again a bisimulation. It is easy to check, however, that the pair (x_1, x_2) is an element of R. As x_2 was assumed to be reflexive and x_1 has no successors both states cannot be related by a bisimulation. Therefore we arrive at a contradiction and can conclude that the schizophrenic object S does not exist.

On the other hand, the duality between MA and DGF is an instance of the duality $\mathsf{Alg}(F^{\mathrm{op}}) \cong \mathsf{Coalg}(F)^{\mathrm{op}}$ of algebras and coalgebras, with the Vietoris functor $\mathbb{V}$ as the functor F. It seems therefore of interest to explore which dualities are instances of the algebra/coalgebra duality. As a first step in this direction, [Pal04] shows that the duality between positive modal algebras and K^+-spaces can be described in a similar way as in Section 3.2 (although the technical details are substantially more complicated).

Chapter 4
Algebraic semantics of coalgebraic modal logic

In the last chapter we saw that the algebraic semantics of normal modal logic, which is given by modal algebras, has a representation as a category of algebras for a functor $H : \mathsf{BA} \to \mathsf{BA}$ (cf. Proposition 3.2.12 on page 46). In this chapter we are going to show that this statement can be generalized to all coalgebraic modal logics that are given by a set of predicate liftings and a set of axioms (cf. Def. 2.2.17).

Before going into a more detailed discussion of the content of this chapter let us stress the fact that in our view the main contribution of this chapter lies not so much in its technical results but in the observation that any coalgebraic modal logic $L(\Lambda, \mathtt{Ax})$ that is given by a set of predicate liftings Λ and a set of axioms $\mathtt{Ax}$ can be represented by a functor $L : \mathsf{BA} \to \mathsf{BA}$. If the functor L is dual to the functor T that specifies the type of the coalgebras under consideration then the corresponding logic $L(\Lambda, \mathtt{Ax})$ is sound and complete with respect to its coalgebraic semantics and its language is expressive.

The technical content of this chapter can be summarized as follows: We take as given a functor $T : \mathsf{C} \to \mathsf{C}$, with C either equal to Set or Stone, together with a set of predicate liftings Λ for T and a set of axioms $\mathtt{Ax}$. Then we can easily define an algebraic semantics for the logic $L(\Lambda, \mathtt{Ax})$ by considering certain Boolean algebras with operators for the associated algebraic theory $\mathsf{T}(\Lambda, \mathtt{Ax})$: the signature of $\mathsf{T}(\Lambda, \mathtt{Ax})$ consists of the signature of Boolean algebras together with an additional n-ary operator for each n-ary predicate lifting $\lambda \in \Lambda$ and the equations in $\mathsf{T}(\Lambda, \mathtt{Ax})$ are the Boolean equations together with the axioms in $\mathtt{Ax}$. The logic $L(\Lambda, \mathtt{Ax})$ will be sound and complete with respect to the algebraic semantics provided by $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$, i.e. by the category of algebras for the theory $\mathsf{T}(\Lambda, \mathtt{Ax})$.

We then define a functor $L : \mathsf{BA} \to \mathsf{BA}$ such that $\mathsf{Alg}(L)$, the category of algebras for L, is isomorphic to $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$. Therefore $\mathsf{Alg}(L)$ gives us a also a sound and complete algebraic semantics for $L(\Lambda, \mathtt{Ax})$. The algebraic semantics provided by L-algebras has the advantage, that its format is very close to that of the coalgebraic semantics. Therefore it allows us to establish a close connection between algebraic and coalgebraic semantics.

This connection is given by a natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$. Here P denotes, in the case $\mathsf{C} = \mathsf{Set}$, the contravariant power set functor while in the case $\mathsf{C} = \mathsf{Stone}$, P denotes the functor $\mathbb{C}\mathrm{lp}$, which maps a Stone space to the Boolean algebra of clopen subsets. In order to be able to understand δ let us look at the following diagram.

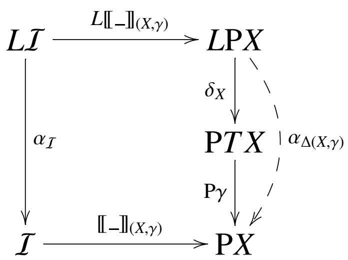

Suppose we are given a T-coalgebra (X, γ) and let $(\mathcal{I}, \alpha_{\mathcal{I}})$ be the initial L-algebra. We can think of $\mathcal{I}$ as the set of formulas of our logic modulo derivable equivalence. The natural transformation δ, if it exists, gives us a possibility to transform the T-coalgebra (X, γ) into the L-algebra $(\mathrm{P}X, \mathrm{P}\gamma \circ \delta_X)$ as depicted on the right half of the diagram. We will see that the coalgebraic semantics on a T-coalgebra (X, γ) is then obtained as the unique $\mathsf{Alg}(L)$-homomorphism from the initial L-algebra $(\mathcal{I}, \alpha_{\mathcal{I}})$ to the L-algebra $(\mathrm{P}X, \mathrm{P}\gamma \circ \delta_X)$. This observation will be used to prove that the existence of δ implies soundness of the logic w.r.t. the coalgebraic semantics: if formulas are equivalent in the logic, i.e. if they belong to the same equivalence class of formulas in $\mathcal{I}$, then they will be interpreted by the same predicate over X. Furthermore we will show that completeness w.r.t. the coalgebraic semantics is entailed by injectivity of δ and, in the case $\mathsf{C} = \mathsf{Stone}$, expressiveness of the language is a consequence of surjectivity of δ.

Sufficient conditions for soundness and completeness of a logic $L(\Lambda, \mathtt{Ax})$ and for the expressiveness of its language have been already given by Pattinson in [Pat03a, Pat04]. We will prove that our conditions formulated in terms of δ are equivalent to Pattinson's. This gives an explanation for his soundness, completeness and expressiveness results and shows that they are in fact algebraic in nature. Note, however, that our result generalizes Pattinson's because we are dealing with both coalgebras over Set and over Stone. Furthermore our setting will allow base categories that are different from these two examples.

In the case that we are dealing with coalgebras over Stone spaces we prove the main technical result of this chapter: a logic $L(\Lambda, \mathtt{Ax})$ for a functor $T : \mathsf{Stone} \to \mathsf{Stone}$ satisfies Pattinson's soundness, completeness and expressiveness conditions iff the corresponding functor $L : \mathsf{BA} \to \mathsf{BA}$ is dual to T (cf. Def. 4.3.19).

The structure of this chapter is as follows: In the first section we define the algebraic semantics, first as a category $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$ of algebras with operators, then as a category $\mathsf{Alg}(L)$ of algebras for a functor $L : \mathsf{BA} \to \mathsf{BA}$. The definition of L contains two parts: a syntactic part that operates on sets of terms and a part that operates on congruence relations. This “modular” definition of L becomes transparent by introducing the category PBA of pre-Boolean algebras. A pre-Boolean algebra consists of a term algebra for the Boolean signature and a congruence relation on it.

Having defined $L : \mathsf{BA} \to \mathsf{BA}$, we recall the definition of the initial and the final sequence of a functor in section 2. These functor sequences are the main tool for proving that the injectivity of δ implies completeness of the logic.

In section 3 we show that we can define, under certain conditions, a natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ that relates the algebraic and the coalgebraic semantics of the logic. We then match properties of δ with properties of the logic:

existence of δ	implies	soundness of $L(\Lambda, \mathtt{Ax})$	
injectivity of δ	implies	completeness of $L(\Lambda, \mathtt{Ax})$	
surjectivity of δ	implies	expressiveness of $\mathcal{L}(\Lambda)$	(for $\mathsf{C} = \mathsf{Stone}$)

In section 4 we show that Pattinson's criteria for these properties, that we mentioned above, are in fact equivalent to our conditions that are formulated in terms of δ. In this way we obtain the announced characterization of duality.

The chapter is based on the earlier published paper [KKP04] which is joint work with Alexander Kurz and Dirk Pattinson.

4.1 Definition of the algebraic semantics

Throughout this section we assume that we are given a functor $T : \mathsf{C} \to \mathsf{C}$, where $\mathsf{C} = \mathsf{Set}$ or $\mathsf{C} = \mathsf{Stone}$ (cf. Section 2.2), together with a set of predicate liftings Λ for T (cf. Def. 2.2.2) and a set of axioms $\mathtt{Ax}$ (cf. Def. 2.2.13). Furthermore we asssume that the functor T has at least one global element, i.e. an arrow from $\mathbf{1}$ to $T\mathbf{1}$. Finally we denote by

$$U : \mathsf{Alg}(\Sigma_{\mathsf{BA}}) \to \mathsf{Set}$$

the forgetful functor mapping a Σ_{BA}-algebra to its underlying set. Our aim is to describe the algebraic semantics for the coalgebraic modal logic $L(\Lambda, \mathtt{Ax})$ (cf. Definition 2.2.17) with the help of a functor $L : \mathsf{BA} \to \mathsf{BA}$.

4.1.1 Algebras for an algebraic theory

There is an obvious way of defining an algebraic semantics of coalgebraic modal logic for a functor T: extend the Boolean signature Σ_{BA} by n-ary operation symbols λ for each n-ary predicate lifting $\lambda \in \Lambda$ to a signature $\Sigma^{\Lambda}_{\mathsf{BA}}$. Then the algebras for the algebraic theory $(\Sigma^{\Lambda}_{\mathsf{BA}}, E_{\mathsf{BA}})$ that validate the axioms in $\mathtt{Ax}$ will yield a semantics for which $L(\Lambda, \mathtt{Ax})$ is sound and complete.

4.1.1. Definition. The algebraic signature $\Sigma^{\Lambda}_{\mathsf{BA}}$ is defined as follows

$$\Sigma^{\Lambda}_{\mathsf{BA}} := \Sigma_{\mathsf{BA}} \cup \{\lambda \mid \lambda \in \Lambda\},$$

and the arity of a λ is defined to be $o(\lambda) := n$ if λ is an n-ary predicate lifting.

The axioms in $\mathtt{Ax}$ can now be interpreted as equations for the extended signature.

4.1.2. DEFINITION. By $\mathsf{T}(\Lambda, \mathtt{Ax})$ we denote the algebraic theory $(\Sigma^{\Lambda}_{\mathsf{BA}}, E_{\mathsf{BA}} \cup \mathtt{Ax})$.

In this way we get a direct connection between equational logic and coalgebraic modal logic.

4.1.3. LEMMA. *Let $\phi, \psi \in \mathcal{L}(\Lambda)$ be formulas. Then*

$$\mathtt{Ax} \vdash \phi \leftrightarrow \psi \quad \textit{iff} \quad E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \approx \psi.$$

Proof. Both directions of this lemma can be proven in a standard way using induction on the length of derivations. QED

The definition of the semantics of a given formula ϕ on an $\mathsf{T}(\Lambda, \mathtt{Ax})$-algebra is an immediate adaptation of the definition of the algebraic semantics of basic modal logic (cf. Def. 3.1.10) to the multi-modal setting of coalgebraic modal logic.

4.1.4. DEFINITION. Given an $\mathsf{T}(\Lambda, \mathtt{Ax})$-algebra $\mathcal{A}$, we define the semantics of a formula inductively as follows

$$\begin{array}{rcl} [\![\bot]\!]_{\mathcal{A}} & := & \bot^{\mathcal{A}} \\ [\![\phi \to \psi]\!]_{\mathcal{A}} & := & \neg[\![\phi]\!]_{\mathcal{A}} \vee^{\mathcal{A}} [\![\psi]\!]_{\mathcal{A}} \\ [\![[\lambda](\phi_1, \ldots, \phi_n)]\!]_{\mathcal{A}} & := & \lambda^{\mathcal{A}}([\![\phi_1]\!]_{\mathcal{A}}, \ldots, [\![\phi_n]\!]_{\mathcal{A}}). \end{array}$$

We say that ϕ is true in $\mathcal{A}$ if $[\![\phi]\!]_{\mathcal{A}} = \top$. In this case we write $\mathcal{A} \models \phi$.

4.1.5. REMARK. If no confusion is possible we will drop the superscripts and write $\bot$, $\vee$ and λ instead of $\bot^{\mathcal{A}}$, $\vee^{\mathcal{A}}$ and $\lambda^{\mathcal{A}}$.

Soundness and completeness of $\mathtt{Ax}$ with respect to $\mathsf{T}(\Lambda)$-algebras are easily obtained.

4.1.6. THEOREM. *For all $\phi \in \mathcal{L}(\Lambda)$ we have*

$$\mathtt{Ax} \vdash \phi \quad \textit{iff} \quad \mathcal{A} \models \phi \textit{ for all } \mathcal{A} \in \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})).$$

Proof. From Birkhoff's completeness theorem for equational logic (cf. Theorem B.2.5) we get

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \approx \top \quad \text{iff} \quad \mathcal{A} \models \phi \text{ for all } \mathcal{A} \in \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})),$$

which is by Lemma 4.1.3 equivalent to the claim. QED

It is, however, not obvious, how we could relate this algebraic semantics to the semantics in which we are interested, namely to the coalgebraic one. Therefore we are going to bring the algebraic semantics of the logic into a more categorical format which is closer to (the dual of) the format of the coalgebraic semantics. In the remainder of this section we will see how to represent $\mathsf{T}(\Lambda, \mathtt{Ax})$-algebras as algebras for a functor L on the category BA of Boolean algebras.

4.1.2 (Pre-)Boolean algebras

Central in our definition of algebraic semantics for coalgebraic modal logic stands the notion of a pre-Boolean algebra. The basic idea is that the functor $L : \mathsf{BA} \to \mathsf{BA}$ will be built out of two components: one component working on formulas and another component working on congruences. We will introduce now the category PBA of pre-Boolean algebras. Objects of PBA consist essentially of a set of formulas together with a so-called Boolean congruence relation. Then we define a functor $\mathfrak{L} : \mathsf{PBA} \to \mathsf{PBA}$ where the above mentioned components of L are made explicit.

Our motivation for defining the category PBA will become clearer when looking at the definition of the functor L in the next subsection. It turns out to be easier and conceptually cleaner to define first a functor $\mathfrak{L} : \mathsf{PBA} \to \mathsf{PBA}$ and then the functor L using the equivalence between the categories BA and PBA. Furthermore pre-Boolean algebras will be useful when we discuss the soundness and completeness of coalgebraic modal logic, because there we often want to reason about formulas rather than equivalence classes of formulas.

4.1.7. Definition. A *pre-Boolean algebra* is a pair $(\mathbf{T}(G), \equiv)$ where $\mathbf{T}(G) := \mathbf{T}_{\Sigma_{\mathsf{BA}}}(G)$ is the term algebra for the Boolean signature over some set G and $\equiv \subseteq U\mathbf{T}(G) \times U\mathbf{T}(G)$ is a congruence relation such that $\mathbf{T}(G)/\equiv$ (cf. Definition B.1.4) is a Boolean algebra.

So a pre-Boolean algebra abstractly describes a set of formulas for the Boolean signature together with a notion of logical equivalence that is closed under the axioms and rules of propositional logic. Because we often will reason about equivalence classes of formulas we introduce some notation.

4.1.8. Notation. *Let $\mathbb{A} \in \Sigma_{\mathsf{BA}}$, $\equiv \subseteq U\mathbb{A} \times U\mathbb{A}$ a congruence and $a \in \mathbb{A}$. Then we denote by $[a]_{\mathbb{A}/\equiv}$ the equivalence class of a in $\mathbb{A}/\equiv$.*

4.1.9. Example. The Σ_{BA}-algebra based on $\mathcal{L}(\Lambda)$ together with the relation $\dashv\vdash := \{(\phi, \psi) \mid \mathtt{Ax} \vdash \phi \to \psi$ and $\mathtt{Ax} \vdash \psi \to \phi\}$ is a pre-Boolean algebra.

Every Boolean algebra can be viewed as a pre-Boolean algebra. Before we formally state this observation we first introduce some terminology.

4.1.10. Definition. For $\mathbb{A} \in \mathsf{BA}$ we define

$$\begin{array}{rcl} \mathrm{Ter}(\mathbb{A}) & := & \mathbf{T}_{\Sigma_{\mathsf{BA}}}(U\mathbb{A}) \\ \mathrm{Diag}(\mathbb{A}) & := & \{(\psi_1, \psi_2) \in \mathrm{Ter}(\mathbb{A}) \times \mathrm{Ter}(\mathbb{A}) \mid \psi_1^{\mathbb{A}} = \psi_2^{\mathbb{A}}\}, \end{array}$$

where $\psi^{\mathbb{A}}$ is inductively defined as

$$\begin{array}{rcl} \bot^{\mathbb{A}} & := & \bot \\ a^{\mathbb{A}} & := & a \quad \text{for} \quad a \in \mathbb{A} \\ (\psi_1 \to \psi_2)^{\mathbb{A}} & := & \psi_1^{\mathbb{A}} \to^{\mathbb{A}} \psi_2^{\mathbb{A}}. \end{array}$$

Furthermore for $f : \mathbb{A}_1 \to \mathbb{A}_2 \in \mathsf{BA}$ we let $\mathrm{Ter}(f) : \mathrm{Ter}(\mathbb{A}_1) \to \mathrm{Ter}(\mathbb{A}_2)$ be the unique Σ_{BA}-morphism that extends the mapping

$$U\mathbb{A}_1 \ni a \mapsto f(a) \in \mathrm{Ter}(\mathbb{A}_2).$$

With this definition Ter is a functor from BA to the category of Σ_{BA}-algebras.

4.1.11. Remark. Expressed in words $\mathrm{Ter}(\mathbb{A})$ is the set of Boolean terms generated by the carrier set of $\mathbb{A}$ and $\mathrm{Diag}(\mathbb{A})$ is the diagram of $\mathbb{A}$ encoding the equality relation on $\mathbb{A}$. It is not difficult to see that $\mathbb{A}$ is presented by $\Sigma_{\mathsf{BA}} \langle U\mathrm{Ter}(\mathbb{A}) \mid \mathrm{Diag}(\mathbb{A})\rangle$ (cf. Def. B.1.10).

4.1.12. Proposition. *Let $\mathbb{A} \in \mathsf{BA}$. Then the pair* $\mathrm{Ter}\mathbb{A} := (\mathrm{Ter}(\mathbb{A}), \mathrm{Diag}(\mathbb{A}))$ *is a pre-Boolean algebra.*

Proof. By definition we have that for all $\phi, \psi \in \mathrm{Ter}(\mathbb{A})$

$$(\phi, \psi) \in \mathrm{Diag}(\mathbb{A}) \quad \text{iff} \quad \phi^{\mathbb{A}} = \psi^{\mathbb{A}}$$

and $\mathrm{Ter}(\mathbb{A})/\mathrm{Diag}(\mathbb{A}) \cong \mathbb{A}$. Hence $\mathrm{Ter}(\mathbb{A})/\mathrm{Diag}(\mathbb{A})$ is a Boolean algebra. QED

When defining morphisms between pre-Boolean algebras we want to make sure that for two Boolean algebras $\mathbb{A}_1, \mathbb{A}_2$ there is an isomorphism between $\mathsf{BA}(\mathbb{A}_1, \mathbb{A}_2)$ and the set of morphisms from $\mathrm{Ter}\mathbb{A}_1$ to $\mathrm{Ter}\mathbb{A}_2$. By requiring that every PBA-morphism is *preserving* we ensure that every morphism between $\mathrm{Ter}\mathbb{A}_1$ and $\mathrm{Ter}\mathbb{A}_2$ corresponds to a BA-homomorphism from $\mathbb{A}_1$ to $\mathbb{A}_2$.

4.1.13. Definition. Let $\mathbb{A}_1, \mathbb{A}_2 \in \mathsf{Alg}(\Sigma_{\mathsf{BA}})$ and let $\equiv_1 \subseteq U\mathbb{A}_1 \times U\mathbb{A}_1$ and $\equiv_2 \subseteq U\mathbb{A}_2 \times U\mathbb{A}_2$ be congruences. A *preserving* map between $(\mathbb{A}_1, \equiv_1)$ and $(\mathbb{A}_2, \equiv_2)$ is a Σ_{BA}-morphism $f : \mathbb{A}_1 \to \mathbb{A}_2$ such that

$$a_1 \equiv_1 a_2 \quad \text{implies} \quad f(a_1) \equiv_2 f(a_2) \quad \text{for all } a_1, a_2 \in \mathbb{A}_1.$$

The next lemma shows that preserving maps between two pre-Boolean algebras give rise to Boolean homomorphisms between the corresponding Boolean algebras.

4.1.14. Lemma. *Suppose that $f : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2)$ is a preserving map. Then the function*

$$\begin{aligned} Qu(f) : \mathbb{A}_1/\equiv_1 &\to \mathbb{A}_2/\equiv \\ [a]_{\mathbb{A}_1/\equiv_1} &\mapsto [f(a)]_{\mathbb{A}_2/\equiv_2} \end{aligned}$$

is a Boolean homomorphism.

Proof. The fact that f is preserving ensures that $Qu(f)$ is well-defined. The proof that $Qu(f)$ is a homomorphism is straightforward and uses that f is a Σ_{BA}-homomorphism. QED

In turn Boolean homomorphisms correspond to preserving maps between pre-Boolean algebras.

4.1.15. Lemma. *Let* $f : \mathbb{A}_1 \to \mathbb{A}_2 \in \mathsf{BA}$. *Then*

$$f(\psi^{\mathbb{A}_1}) = (\mathrm{Ter}(f)(\psi))^{\mathbb{A}_2} \tag{4.1}$$

and therefore $\mathrm{Ter}(f) : \mathrm{Ter}(\mathbb{A}_1) \to \mathrm{Ter}(\mathbb{A}_2)$ *is a preserving map from* $(\mathrm{Ter}(\mathbb{A}_1), \mathrm{Diag}(\mathbb{A}_1))$ *to* $(\mathrm{Ter}(\mathbb{A}_2), \mathrm{Diag}(\mathbb{A}_2))$.

Proof. Equation 4.1 can be proven by induction on the structure of ψ. That $\mathrm{Ter}(f)$ is a preserving map follows then immediately. QED

In general however more than one such preserving map will correspond to the same BA-morphism as the following example shows.

4.1.16. Example. Let $\mathbb{A} := \mathbf{T}(\emptyset)$ be the Boolean term algebra over the empty set of generators and let $\equiv \subseteq U\mathbb{A} \times U\mathbb{A}$ denote equivalence under propositional logic, i.e. $(\mathbb{A}/\equiv) \cong 2$, where 2 is the two-element Boolean algebra. Furthermore we define the following preserving map

$$\begin{array}{rcl} f : (\mathbb{A}, \equiv) & \to & (\mathbb{A}, \equiv) \\ \mathbb{A} \ni a & \mapsto & \begin{cases} \top & \text{if } a \equiv \top \\ a & \text{otherwise.} \end{cases} \end{array}$$

Then it is easy to check that $f \neq id_{(\mathbb{A},\equiv)}$ and $Qu(f) = Qu(id_{(\mathbb{A},\equiv)}) = id_2$, i.e. Qu identifies the distinct preserving maps f and $id_{(\mathbb{A},\equiv)}$.

In order to obtain a one-to-one correspondence between BA- and PBA-morphisms we define an equivalence relation on preserving maps.

4.1.17. Definition. Let $(\mathbb{A}_1, \equiv_1), (\mathbb{A}_2, \equiv_2)$ be pre-Boolean algebras. Two preserving maps $f_1, f_2 : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2)$ are *equivalent* if

$$f_1(a) \equiv_2 f_2(a) \quad \text{for all } a \in \mathbb{A}_1.$$

In this case we write $f_1 \sim f_2$. The equivalence class of a preserving map f will be denoted by $\overline{f}$.

We are now ready to define the notion of a pre-Boolean algebra morphism.

4.1.18. Definition. Let $(\mathbb{A}_1, \equiv_1)$ and $(\mathbb{A}_2, \equiv_2)$ be pre-Boolean algebras. A morphism between $(\mathbb{A}_1, \equiv_1)$ and $(\mathbb{A}_2, \equiv_2)$ is an equivalence class $\overline{f}$ of a preserving map

$$f : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2).$$

We denote by PBA the category of pre-Boolean algebras and morphisms between them.

4.1.19. Remark. The definition of a morphism as an equivalence class of functions is unfortunately unavoidable in order to obtain an equivalence between the categories PBA and BA. This equivalence is useful for proving facts in BA by proving the corresponding result in PBA (cf. e.g. the proof that L is finitary in Section 4.2.1). Another option would have been to use relations that satisfy certain extra conditions as morphisms. We believe, however, that this alternative representation of the PBA-morphisms would not simplify the presentation.

The following class of PBA-morphisms will be of special interest as they correspond to injective BA-morphisms.

4.1.20. Definition. We call a preserving map $f : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2)$ *reflecting* if for all $a_1, a_2 \in \mathbb{A}_1$, $f(a_1) \equiv_2 f(a_2)$ implies $a_1 \equiv_1 a_2$. A PBA-morphism $\overline{f}$ is called reflecting if f and hence all $g \in \overline{f}$ are reflecting.

4.1.21. Lemma. *Let $f : (\mathbb{A}_1, \equiv) \to (\mathbb{A}_2, \equiv)$ be a preserving map. Then f is reflecting iff $Qu(f) : \mathbb{A}_1/\equiv_1 \to \mathbb{A}_1/\equiv_2$ is injective.*

Proof. Suppose $f : (\mathbb{A}_1, \equiv) \to (\mathbb{A}_2, \equiv)$ is reflecting and let $a_1, a_2 \in \mathbb{A}_1$ such that

$$Qu(f)([a_1]_{\mathbb{A}_1/\equiv_1}) = Qu(f)([a_2]_{\mathbb{A}_1/\equiv_1}).$$

Unfolding the definition of $Qu(f)$ we get $f(a_1) \equiv_2 f(a_2)$ and therefore we can conclude by the fact that f is reflecting that $a_1 \equiv_1 a_2$, i.e. $[a_1]_{\mathbb{A}_1/\equiv_1} = [a_2]_{\mathbb{A}_1/\equiv_1}$. The converse direction can be proven in a similar way. QED

Again this correspondence is mutual, i.e. every injective Boolean homomorphism corresponds to a reflecting map between pre-Boolean algebras.

4.1.22. Lemma. *Let $f : \mathbb{A}_1 \to \mathbb{A}_2 \in$ BA be an injective homomorphism. Then* $\mathrm{Ter}(f) : \mathrm{Ter}(\mathbb{A}_1) \to \mathrm{Ter}(\mathbb{A}_2)$ *is a reflecting map from* $(\mathrm{Ter}(\mathbb{A}_1), \mathrm{Diag}(\mathbb{A}_1))$ *to* $(\mathrm{Ter}(\mathbb{A}_2), \mathrm{Diag}(\mathbb{A}_2))$.

Proof. The claim can easily be proven using equation (4.1) from Lemma 4.1.15. QED

We finish our introduction of PBA-morphisms with a characterization of the PBA-isomorphisms.

4.1.23. Lemma. *Let $\overline{f} : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2)$ be a PBA-morphism. Then $\overline{f}$ is an iso iff*

1. *f is reflecting and*
2. *for all $a_2 \in \mathbb{A}_2$ there is an $a_1 \in \mathbb{A}_2$ such that $f(a_1) \equiv_2 a_2$.*

Proof. In case $\overline{f}$ fulfills the conditions 1 and 2 we first define a preserving map f^{-1} as follows:

$$\begin{array}{rcl} f^{-1} : (\mathbb{A}_2, \equiv_2) & \to & (\mathbb{A}_1, \equiv_1) \\ a & \mapsto & a' \quad \text{for some } a' \text{ s.t. } f(a') \equiv_2 a. \end{array}$$

The map f^{-1} is well-defined because of condition 2 and preserving because of condition 1. Moreover as a direct consequence of the definition of f^{-1} we have for all $a \in \mathbb{A}_2$ that $f(f^{-1}(a)) \equiv_2 a$. Therefore we get $f \circ f^{-1} \sim id_{(\mathbb{A}_2,\equiv_2)}$ and $\overline{f} \circ \overline{f^{-1}} = \overline{id}$. In the same way we can see that $\overline{f^{-1}} \circ \overline{f} = \overline{id}$ and hence f is an isomorphism.

For the other direction of the lemma suppose that $\overline{f}$ is an iso. Then there is a PBA-morphisms $\overline{g}$ such that

$$\begin{array}{rcl} \overline{g} \circ \overline{f} & = & \overline{id}_{(\mathbb{A}_1,\equiv_1)} \quad \text{and} \\ \overline{f} \circ \overline{g} & = & \overline{id}_{(\mathbb{A}_2,\equiv_2)}. \end{array}$$

We now check that f fulfills properties 1 and 2 from our claim. In order to check that f is reflecting suppose that $f(a_1) \equiv_2 f(a_2)$ for some a_1, a_2 in $\mathbb{A}_1$. Then $g(f(a_1)) \equiv_1 g(f(a_2))$ because g is preserving. As $\overline{g} \circ \overline{f} = \overline{id}$ we can conclude that $a_1 \equiv_1 a_2$. For property 2 observe that we have $f(g(a)) \equiv_2 a$ for all $a \in \mathbb{A}_2$. QED

The categories PBA and BA can be related by two functors.

4.1.24. Definition. We define the following functors:

$$\begin{array}{rcl} \mathrm{Ter} : \mathsf{BA} & \to & \mathsf{PBA} \\ \mathbb{A} & \mapsto & (\mathrm{Ter}(\mathbb{A}), \mathrm{Diag}(\mathbb{A})) \\ f & \mapsto & \{g \mid g \sim \mathrm{Ter}(f)\} \end{array}$$

$$\begin{array}{rcl} Qu : \mathsf{PBA} & \to & \mathsf{BA} \\ (\mathbb{A}, \equiv) & \mapsto & \mathbb{A}/\equiv \\ \overline{f} : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2) & \mapsto & Qu(f) : \mathbb{A}_1/\equiv_1 \to \mathbb{A}_2/\equiv_2 \end{array}$$

where f is some representant of $\overline{f}$ and $Qu(f)$ denotes the function mapping an equivalence class $[a]_{\mathbb{A}_1/\equiv_1}$ to $[f(a)]_{\mathbb{A}_2/\equiv_2}$.

4.1.25. Remark. Note that $Qu(\overline{f})$ is well-defined because the representant f of $\overline{f}$ is a preserving map. Moreover the definition does not depend on the chosen representant.

These functors witness the equivalence between the two categories BA and PBA.

4.1.26. Proposition. *The functor* Qu : PBA → BA *and the functor* Ter : BA → PBA *form an equivalence between the categories* BA *and* PBA.

Proof. For every $\mathbb{A} \in \mathsf{BA}$ we define a BA-morphism

$$\begin{array}{rcl} f_{\mathbb{A}} : \mathbb{A} & \rightarrow & Qu\mathrm{Ter}\mathbb{A} \\ a & \mapsto & [a]_{\mathrm{Ter}(\mathbb{A})/\mathrm{Diag}(\mathbb{A})} \end{array}$$

and for every pre-Boolean algebra $(\mathbf{T}(G), \equiv)$ we define a PBA-morphism

$$\overline{h}_{(\mathbf{T}(G),\equiv)} : (\mathbf{T}(G), \equiv) \rightarrow \mathrm{Ter}\,(Qu(\mathbf{T}(G), \equiv))$$

as the equivalence class of the preserving map

$$\begin{array}{rcl} h_{(\mathbf{T}(G),\equiv)} : (\mathbf{T}(G), \equiv) & \rightarrow & \mathrm{Ter}\,(Qu(\mathbf{T}(G), \equiv)) \\ \mathbf{T}(G) \ni t & \mapsto & [t]_{\mathbf{T}(G)/\equiv}\,. \end{array}$$

It is not difficult to check that the families $(f_{\mathbb{A}})_{\mathbb{A}\in\mathbb{A}}$ and $(\overline{h}_{(\mathbf{T}(G),\equiv)})_{(\mathbf{T}(G),\equiv)\in\mathsf{PBA}}$ are natural isomorphisms $Id_{\mathsf{BA}} \cong Qu \circ \mathrm{Ter}$ and $Id_{\mathsf{PBA}} \cong \mathrm{Ter} \circ Qu$ respectively. Hence the functors Qu and Ter form an equivalence. QED

4.1.27. Remark. It is possible to restrict the equivalence $\mathsf{BA} \cong \mathsf{PBA}$ to an equivalence between the category of Boolean algebras with injective homomorphisms and the category of pre-Boolean algebras with reflecting PBA-morphisms (cf. Lemma 4.1.21 and 4.1.22).

4.1.3 Liftings and the functor L

The idea for defining the functor $L : \mathsf{BA} \rightarrow \mathsf{BA}$ is the following: for a Boolean algebra $\mathbb{A}$ we think of the elements of Ter($\mathbb{A}$) as syntactic representations of the predicates over some set X. Then elements of $L\mathbb{A}$ should represent lifted predicates, i.e. predicates over TX. To define $L\mathbb{A}$ we therefore need two liftings:

- a *syntactic lifting*, which lifts the set Ter($\mathbb{A}$) whose elements describe predicates over X to a set Lift(Ter($\mathbb{A}$)) (of Boolean terms) that describe predicates over TX and
- a *congruence lifting*, which lifts the congruence $\mathrm{Diag}(\mathbb{A}) \subseteq \mathrm{Ter}(\mathbb{A}) \times \mathrm{Ter}(\mathbb{A})$ to a congruence relation $\mathsf{Lift}(\mathrm{Diag}(\mathbb{A})) \subseteq \mathsf{Lift}(\mathrm{Ter}(\mathbb{A})) \times \mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$ taking into account the equations `Ax` as they should give a description of the equations which are satisfied by the lifted predicates over TX.

We will combine these liftings and show that they give rise to a functor $\mathfrak{L} : \mathsf{PBA} \rightarrow \mathsf{PBA}$. The functor $L : \mathsf{BA} \rightarrow \mathsf{BA}$ is then the composition of functors $Qu \circ \mathfrak{L} \circ \mathrm{Ter}$. In particular this means that the Boolean algebra $L\mathbb{A}$ will be defined as Lift(Ter($\mathbb{A}$)) modulo Lift(Diag($\mathbb{A}$)).

We are now giving the definitions of the necessary liftings. They are essentially defined as similar lifting operations in [Pat03a]. First let us look at the syntactic lifting.

4.1.28. Definition. Let Ψ be a set. Then we define the Σ_{BA}-algebra $\mathsf{Lift}(\Psi)$ of lifted formulas of rank 1 as follows

$$\mathsf{Lift}(\Psi) := \mathbf{T}(\{\lambda(\psi_1, \ldots \psi_n) \mid \psi_1, \ldots, \psi_n \in \Psi \text{ and } \lambda \in \Lambda \ n\text{-ary}\}).$$

Here and in the next definition the set Ψ can be thought of as a set of Boolean formulas. The congruence lifting is defined for arbitrary relations, but we will only use it to lift Boolean congruences.

4.1.29. Definition. Let Ψ be a set and $R \subseteq \Psi \times \Psi$ a relation. Then we view R as a set of Ψ-equations and define a Σ_{BA}-congruence

$$\mathsf{Lift}(R) := \{(t, s) \in \mathsf{Lift}(\Psi) \times \mathsf{Lift}(\Psi) \mid E_{\mathsf{BA}} \cup \mathtt{Ax} \cup R \vdash^{\Sigma^{\Lambda}_{\mathsf{BA}}}_{\mathbf{EL}} t \approx s\}.$$

4.1.30. Remark. Note that in the previous definition elements of $\mathsf{Lift}(\Psi)$ are considered to be Ψ-terms for the signature $\Sigma^{\Lambda}_{\mathsf{BA}}$, i.e.

$$\mathsf{Lift}(\Psi) \subseteq \mathbf{T}_{\Sigma^{\Lambda}_{\mathsf{BA}}}(\Psi).$$

This is important because, for example, we want to be able to deduce from $(\psi_1, \psi_2) \in R$ that $(\lambda(\psi_1), \lambda(\psi_2)) \in \mathsf{Lift}(R)$.

The next lemma shows that the two liftings together can be used to define a mapping from pre-Boolean algebras to pre-Boolean algebras.

4.1.31. Lemma. *Let* $(\mathbb{A}, \equiv) \in \mathsf{PBA}$ *and recall that* U *denotes the forgetful functor from* $\mathsf{Alg}(\Sigma_{\mathsf{BA}})$ *to* Set. *Then*

$$\mathfrak{L}(\mathbb{A}, \equiv) := (\mathsf{Lift}(U\mathbb{A}), \mathsf{Lift}(\equiv))$$

is a pre-Boolean algebra.

Proof. Because $\mathsf{Lift}(\equiv)$ is a Σ_{BA}-congruence it is clear that $L\mathbb{A} := \mathsf{Lift}(U\mathbb{A})/\mathsf{Lift}(\equiv)$ is a well-defined Σ_{BA}-algebra. We have to show that $L\mathbb{A} \in \mathsf{BA}$, i.e. that for all $e_1 \approx e_2 \in E_{\mathsf{BA}}$ we have $L\mathbb{A} \models e_1 \approx e_2$. Let $e_1 \approx e_2 \in E_{\mathsf{BA}}$ and let

$$e_1^{\mathsf{Lift}(U\mathbb{A})}, e_2^{\mathsf{Lift}(U\mathbb{A})} : (\mathsf{Lift}(U\mathbb{A}))^n \to \mathsf{Lift}(U\mathbb{A})$$

be the term functions (cf. Definition B.1.5) on the Σ_{BA}-algebra $\mathsf{Lift}(U\mathbb{A})$.

Let $a_1, \ldots, a_n$ be arbitrary elements of $\mathsf{Lift}(U\mathbb{A})$. Then $e_i^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots a_n) \in \mathsf{Lift}(U\mathbb{A})$, $i = 1, 2$, and $E_{\mathsf{BA}} \vdash_{\mathbf{EL}} e_1^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots a_n) \approx e_2^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots a_n)$. Therefore we have by definition of $\mathsf{Lift}(\equiv)$ that $(e_1^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots a_n), e_2^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots a_n)) \in \mathsf{Lift}(\equiv)$ for all $a_1, \ldots, a_n \in \mathsf{Lift}(U\mathbb{A})$. Hence

$$\begin{aligned} e_1^{L\mathbb{A}}([a_1]_{L\mathbb{A}}, \ldots, [a_n]_{L\mathbb{A}}) &= \left[e_1^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots, a_n)\right]_{L\mathbb{A}} \\ &= \left[e_2^{\mathsf{Lift}(U\mathbb{A})}(a_1, \ldots, a_n)\right]_{L\mathbb{A}} \\ &= e_2^{L\mathbb{A}}([a_1]_{L\mathbb{A}}, \ldots, [a_n]_{L\mathbb{A}}). \end{aligned}$$

for all $a_1, \ldots, a_n \in \mathsf{Lift}(U\mathbb{A})$. As a consequence we get $e_1^{L\mathbb{A}} = e_2^{L\mathbb{A}}$ which in turn implies $L\mathbb{A} \models e_1 \approx e_2$. QED

With the help of the liftings we can also map PBA-morphisms to PBA-morphisms. In order to be able to do this we first show how we can lift preserving maps between pre-Boolean algebras to preserving maps.

4.1.32. Definition. Let $(\mathbb{A}_1, \equiv_1)$ and $(\mathbb{A}_2, \equiv_2)$ be pre-Boolean algebras and let $f : (\mathbb{A}_1, \equiv_1) \to (\mathbb{A}_2, \equiv_2)$ be a preserving map between them. Then we let $\mathfrak{L}(f) : \mathsf{Lift}(U\mathbb{A}_1) \to \mathsf{Lift}(U\mathbb{A}_2)$ be the Σ_{BA}-morphism that extends the mapping

$$\lambda(\psi_1, \dots, \psi_n) \mapsto \lambda(f(\psi_1), \dots, f(\psi_n)) \quad \text{for } \lambda \in \Lambda \text{ } n\text{-ary}.$$

The following lemma shows that $\mathfrak{L}(f)$ is a preserving map from $\mathfrak{L}(\mathbb{A}_1, \equiv_1)$ to $\mathfrak{L}(\mathbb{A}_2, \equiv_2)$ and that $\mathfrak{L}(f)$ is reflecting if f is.

4.1.33. Lemma. *Let $(\mathbb{A}_1, \equiv_1), (\mathbb{A}_2, \equiv_2)$ be pre-Boolean algebras and $f : \mathbb{A}_1 \to \mathbb{A}_2$ a Σ_{BA}-morphism. Then for all $t_1, t_2 \in \mathsf{Lift}(U\mathbb{A}_1)$*

1. *if f is preserving then $\mathfrak{L}(f)$ is preserving, i.e.*

$$t_1 \mathsf{Lift}(\equiv_1) t_2 \quad \textit{implies} \quad \mathfrak{L}(f)(t_1)\mathsf{Lift}(\equiv_2)\mathfrak{L}(f)(t_2), \qquad (*)$$

2. *if f is reflecting then $\mathfrak{L}(f)$ is reflecting, i.e.*

$$\mathfrak{L}(f)(t_1)\mathsf{Lift}(\equiv_2)\mathfrak{L}(f)(t_2) \quad \textit{implies} \quad t_1 \mathsf{Lift}(\equiv_1) t_2.$$

3. *if f is preserving and $g : \mathbb{A}_1 \to \mathbb{A}_2$ is another preserving map such that $f \sim g$, then also $\mathfrak{L}(f) \sim \mathfrak{L}(g)$.*

Proof. In order to prove 1, we have to show that $(*)$ holds for all $t_1, t_2 \in \mathsf{Lift}(U\mathbb{A}_1)$. So let $t_1, t_2 \in \mathsf{Lift}(U\mathbb{A}_1)$ and suppose $t_1 \mathsf{Lift}(\equiv_1) t_2$, i.e. $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_1 \vdash_{\mathbf{EL}} t_1 \approx t_2$. We prove by induction on the length of the derivation of $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_1 \vdash_{\mathbf{EL}} t_1 \approx t_2$ that $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_2 \vdash_{\mathbf{EL}} \mathfrak{L}(f)(t_1) \approx \mathfrak{L}(f)(t_2)$ and therefore $\mathfrak{L}(f)(t_1)\mathsf{Lift}(\equiv_2)\mathfrak{L}(f)(t_2)$.

Case: $t_1 \approx t_2 \in E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_1$. Since $\equiv_1$ does not relate terms in $\mathsf{Lift}(U\mathbb{A}_1) \times \mathsf{Lift}(U\mathbb{A}_2)$ we see that $t_1 \approx t_2 \in E_{\mathsf{BA}} \cup \mathtt{Ax}$, and the claim is trivial. (Note that equations in $E_{\mathsf{BA}} \cup \mathtt{Ax}$ do not contain any parameters in $U\mathbb{A}_i$ and therefore we get in this case $\mathfrak{L}(f)(t_i) = t_i$.)

Case: $t_1 = t_2$. Then obvious.

Case: $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_1 \vdash_{\mathbf{EL}} \psi_i \approx \phi_i$ for $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n \in U\mathbb{A}_1$ and by the congruence rule of equational logic we derive

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_1 \vdash_{\mathbf{EL}} \lambda(\psi_1, \dots, \psi_n) \approx \lambda(\phi_1, \dots, \phi_n).$$

Then by Fact B.2.3 in the appendix we also have

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \cup (\equiv_1)[f] \vdash_{\mathbf{EL}} (\psi_i \approx \phi_i)[f] \quad \text{for all } i \in \{1, \dots, n\}.$$

Spelling out the definition of $(\_)[f]$ (cf. B.2.3) we get $(\psi_i \approx \phi_i)[f] = f(\psi_i) \approx f(\phi_i)$. Moreover $(\equiv_1)[f] \subseteq \equiv_2$ because f is a preserving map. Therefore using Fact B.2.2 we arrive at

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_2 \vdash_{\mathbf{EL}} f(\psi_i) \approx f(\phi_i) \quad \text{for all } i \in \{1, \ldots, n\},$$

and hence by the congruence rule of equational logic we get

$$\begin{aligned} E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_2 \quad \vdash_{\mathbf{EL}} \quad & \lambda(f(\psi_1), \ldots, f(\psi_n)) \\ & \approx \lambda(f(\phi_1), \ldots, f(\phi_n)). \end{aligned}$$

The remaining cases of the induction are not difficult and can be treated using the induction hypothesis. The second item of the lemma can be proven in a similar way. To prove the last claim it suffices to show that

$$\mathfrak{L}(g)(t)\mathsf{Lift}(\equiv_2)\mathfrak{L}(f)(t) \quad \text{for all } t \in \mathsf{Lift}(U\mathbb{A}_1).$$

This can be done by an easy induction on the structure of t. QED

The next definition summarizes the results of the previous two lemmas.

4.1.34. Definition. We define a functor

$$\begin{aligned} \mathfrak{L} : \mathsf{PBA} \;&\to\; \mathsf{PBA} \\ (\mathbb{A}, \equiv) \;&\mapsto\; \mathfrak{L}(\mathbb{A}, \equiv) \\ \overline{f} \;&\mapsto\; \mathfrak{L}(\overline{f}) := \{g \mid g \sim \mathfrak{L}(f)\}, \end{aligned}$$

where $\mathfrak{L}(\mathbb{A}, \equiv)$ is defined as described in Lemma 4.1.31 and $\mathfrak{L}(f)$ as described in Definition 4.1.32.

That $\mathfrak{L}$ is well-defined on morphisms follows from items 1 and 3 of Lemma 4.1.33. Using the functors Qu and Ter from Section 4.1.2 we obtain an endofunctor on BA.

4.1.35. Definition. The *algebraic semantics functor* $L : \mathsf{BA} \to \mathsf{BA}$ is defined as $L := Qu \circ \mathfrak{L} \circ \mathrm{Ter}$ (cf. Figure 4.1).

4.1.36. Remark. In other words, in order to compute $L\mathbb{A}$ we first map $\mathbb{A}$ to its corresponding pre-Boolean algebra $(\mathrm{Ter}(\mathbb{A}), \mathrm{Diag}(\mathbb{A}))$, then lift both the syntax and the semantics by applying the functor $\mathfrak{L}$ and finally we quotient the set of lifted formulas $\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$ by the lifted congruence $\mathsf{Lift}(\mathrm{Diag}(\mathbb{A}))$ to obtain again a Boolean algebra.

It is now straightforward to define the algebraic semantics given by L-algebras (cf. Def. A.1.3 on page 147).

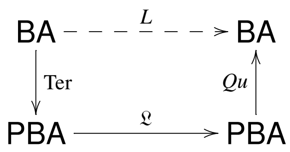

Figure 4.1: The functors L and $\mathfrak{L}$

4.1.37. Definition. For an L-algebra $(\mathbb{A}, \alpha : L\mathbb{A} \to \mathbb{A})$ we define the (L-)algebraic semantics $[\![\phi]\!]_{(\mathbb{A},\alpha)}$ of a formula ϕ inductively:

$$\begin{array}{rcl} [\![\bot]\!]_{(\mathbb{A},\alpha)} & := & \bot \\ [\![\phi \to \psi]\!]_{(\mathbb{A},\alpha)} & := & \neg[\![\phi]\!]_{(\mathbb{A},\alpha)} \cup [\![\psi]\!]_{(\mathbb{A},\alpha)} \\ [\![[\lambda](\psi_1, \dots, \psi_n)]\!]_{(\mathbb{A},\alpha)} & := & \alpha(\lambda([\![\psi_1]\!]_{(\mathbb{A},\alpha)}, \dots, [\![\psi_n]\!]_{(\mathbb{A},\alpha)})). \end{array}$$

Before we finish this section we have a detailed look at how the functor L acts on BA-morphisms. This will be useful for later calculations that involve L.

4.1.38. Lemma. *Let $f : \mathbb{A}_1 \to \mathbb{A}_2$ be a* BA*-morphism. Then Lf is the unique* BA*-morphism that extends the mapping*

$$[\lambda(\psi_1, \dots, \psi_m)]_{L\mathbb{A}_1} \mapsto [\lambda(\mathrm{Ter}(f)(\psi_1), \dots, \mathrm{Ter}(f)(\psi_n))]_{L\mathbb{A}_2}.$$

Proof. Just spell out the definition of $L = Qu \circ \mathfrak{L} \circ \mathrm{Ter}$. QED

4.1.39. Proposition. *Let $f : \mathbb{A}_1 \to \mathbb{A}_2$ be a* BA*-morphism. Then*

(i) Lf is injective if f is injective and

(ii) Lf is surjective if f is surjective.

Proof. The first half of the proposition follows from Lemma 4.1.33(2) and the fact that the functors Ter and Qu map injective BA-morphisms to reflective PBA-morphisms and reflective PBA-morphisms to injective BA-morphisms respectively (cf. Lemma 4.1.22 and Lemma 4.1.21). The second half is an easy consequence of Lemma 4.1.38. QED

4.1.4 Equivalence of Alg(T(Λ, Ax)) and Alg(L)

We presented two alternative ways of defining an algebraic semantics for coalgebraic modal logic: one using algebras for an algebraic theory $\mathsf{T}(\Lambda, \mathtt{Ax})$ and another one using algebras for the functor L. This generalizes the earlier observation that modal algebras can be represented as algebras for a functor H : BA $\to$ BA (cf. Remark 3.2.13).

To finish our definition of algebraic semantics of coalgebraic modal logic we now show that both approaches are in fact equivalent: the category Alg(L) of algebras for the functor L is isomorphic to the category $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$ and the functors which form this equivalence preserve the semantics.

4.1.40. Lemma and Definition. *The following defines a functor:*

$$\begin{array}{rcl} E : \mathsf{Alg}(L) & \to & \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})) \\ (\mathbb{A}, \alpha) & \mapsto & (\mathbb{A}, \{\lambda^{(\mathbb{A},\alpha)} \mid \lambda \in \Lambda\}) \\ f : (\mathbb{A}_1, \alpha_1) \to (\mathbb{A}_2, \alpha_2) & \mapsto & f : E(\mathbb{A}_1, \alpha_1) \to E(\mathbb{A}_2, \alpha_2). \end{array}$$

where for every n-ary $\lambda \in \Lambda$ we define $\lambda^{(\mathbb{A},\alpha)} : (U\mathbb{A})^n \to U\mathbb{A}$ to be the function mapping $(a_1, \ldots, a_n)$ to $\alpha([\lambda(a_1, \ldots, a_n)]_{L\mathbb{A}})$.

Proof. To prove that E is well defined on objects we have to show that $E(\mathbb{A}, \alpha) \models E_{\mathsf{BA}} \cup \mathtt{Ax}$. This is clear for the Boolean equations in E_{BA}. Suppose on the other hand that $e_1 \approx e_2$ is an equation from $\mathtt{Ax}$ with equation variables $\{x_1, \ldots, x_n\}$. We can easily show by induction on the structure of e_i that for all $a_1, \ldots, a_n \in \mathbb{A}$ we have

$$e_i^{E(\mathbb{A},\alpha)}(a_1, \ldots, a_n) = \alpha\left(\left[e_i^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n)\right]_{L\mathbb{A}}\right) \quad \text{for } i \in \{1, 2\}. \tag{4.2}$$

Then

$$\begin{array}{rcl} e_1^{E(\mathbb{A},\alpha)}(a_1, \ldots, a_n) & \overset{(4.2)}{=} & \alpha(\left[e_1^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n)\right]_{L\mathbb{A}}) \\ & \overset{(*)}{=} & \alpha(\left[e_2^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n)\right]_{L\mathbb{A}}) \\ & \overset{(4.2)}{=} & e_2^{E(\mathbb{A},\alpha)}(a_1, \ldots, a_n) \end{array}$$

where $(*)$ holds because

$$\mathtt{Ax} \vdash_{\mathbf{EL}} e_1^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n) \approx e_2^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n),$$

and therefore $(e_1^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n), e_2^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n)) \in \mathsf{Lift}(\mathrm{Diag}(\mathbb{A}))$. Furthermore note that because all the axioms are of depth 1 the terms $e_i[x_1/a_1] \ldots [x_n/a_n]$ can be evaluated in $\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$ and therefore $e_i^{\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))}(a_1, \ldots, a_n)$ is well-defined.

In order to show that $E(f)$ is a homomorphism if $f : (\mathbb{A}_1, \alpha_1) \to (\mathbb{A}_2, \alpha_2)$ is an L-algebra morphism it suffices to check that $Ef = f$ commutes with the λ's. Consider an n-ary predicate lifting $\lambda \in \Lambda$ and arbitrary elements $a_1, \ldots, a_n$ of $\mathbb{A}_1$. Then

$$\begin{array}{rcl} f\left(\lambda^{(\mathbb{A}_1,\alpha_1)}(a_1, \ldots, a_n)\right) & \overset{\text{by Def.}}{=} & f\left(\alpha_1([\lambda(a_1, \ldots, a_n)]_{L\mathbb{A}_1})\right) \\ & \overset{f \text{ is } L\text{-mor.}}{=} & \alpha_2\left(Lf([\lambda(a_1, \ldots, a_n)]_{L\mathbb{A}_1})\right) \\ & \overset{\text{Lem. 4.1.38}}{=} & \alpha_2\left([\lambda(\mathrm{Ter}(f)(a_1), \ldots, \mathrm{Ter}(f)(a_n))]_{L\mathbb{A}_2}\right) \\ & \overset{\mathrm{Ter}(f)(a_i) = f(a_i)}{=} & \alpha_2\left([\lambda(f(a_1), \ldots, f(a_n))]_{L\mathbb{A}_2}\right) \\ & \overset{\text{by Def.}}{=} & \lambda^{(\mathbb{A}_2,\alpha_2)}\left(f(a_1), \ldots, f(a_n)\right). \end{array}$$

QED

4.1.41. Remark. It should be stressed that the proof of the previous lemma is the point where we make essential use of the fact that we have restricted our attention to axioms of modal depth 1. The definition of the functor L would work also without this restriction, but L-algebras would no longer correspond to algebras in $\mathsf{Alg}(\Lambda, \mathtt{Ax})$ and therefore soundness and completeness of the logic $L(\Lambda, \mathtt{Ax})$ with respect to $\mathsf{Alg}(L)$ would no longer be guaranteed.

In the other direction, we can also map $\mathsf{T}(\Lambda, \mathtt{Ax})$-algebras to L-algebras.

4.1.42. Lemma and Definition. *The following defines a functor*

$$\begin{array}{rcl} A : \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})) & \to & \mathsf{Alg}(L) \\ \mathcal{A} = (\mathbb{A}, \{\lambda^{\mathcal{A}} \mid \lambda \in \Lambda\}) & \mapsto & (\mathbb{A}, \alpha_{\mathcal{A}}) \\ f : \mathcal{A}_1 \to \mathcal{A}_2 & \mapsto & f : (\mathbb{A}_1, \alpha_{\mathcal{A}_1}) \to (\mathbb{A}_2, \alpha_{\mathcal{A}_2}), \end{array}$$

where $\alpha_{\mathcal{A}} : L\mathbb{A} \to \mathbb{A} \in \mathsf{BA}$ *is defined as* $\alpha_{\mathcal{A}}([t]_{L\mathbb{A}}) := t^{\mathcal{A}}$.

Proof. We show first that for $\mathcal{A} = (\mathbb{A}, \{\lambda^{\mathcal{A}} \mid \lambda \in \Lambda\}) \in \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$ the map $\alpha_{\mathcal{A}} : L\mathbb{A} \to \mathbb{A}$ is well-defined. For this it suffices to show that for all $t_1, t_2 \in \mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$ we have

$$t_1 \mathsf{Lift}(\mathrm{Diag}(\mathbb{A})) t_2 \quad \text{implies} \quad t_1^{\mathcal{A}} = t_2^{\mathcal{A}}.$$

Suppose $t_1 \mathsf{Lift}(\mathrm{Diag}(\mathbb{A})) t_2$ for some $t_1, t_2 \in \mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$, i.e.

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \mathrm{Diag}(\mathbb{A}) \vdash_{\mathbf{EL}} t_1 \approx t_2.$$

We know that by definition $\mathcal{A} \models E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \mathrm{Diag}(\mathbb{A})$, hence $\mathcal{A} \models t_1 \approx t_2$ by the soundness of equational logic, i.e. $t_1^{\mathcal{A}} = t_2^{\mathcal{A}}$.
Now let us check whether A is well-defined on morphisms. Let $f : \mathcal{A}_1 \to \mathcal{A}_2 \in \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$. We have to prove that $f : (\mathbb{A}_1, \alpha_{\mathcal{A}_1}) \to (\mathbb{A}_2, \alpha_{\mathcal{A}_2})$ is an $\mathsf{Alg}(L)$-morphism, which means that we have to show for all $t \in \mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$ that the following diagram commutes:

$$\begin{array}{ccc} L\mathbb{A}_1 & \xrightarrow{Lf} & L\mathbb{A}_2 \\ {\scriptstyle \alpha_{\mathcal{A}_1}} \downarrow & & \downarrow {\scriptstyle \alpha_{\mathcal{A}_2}} \\ \mathbb{A}_1 & \xrightarrow[f]{} & \mathbb{A}_2 \end{array}$$

Because $L\mathbb{A}_1 = \mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))/\mathsf{Lift}(\mathrm{Diag}(\mathbb{A}))$ it is enough to check the commutativity of the diagram on the generators of $\mathsf{Lift}(\mathrm{Ter}(\mathbb{A}))$. Let $t = \lambda(\psi_1, \ldots, \psi_n)$ for some $\lambda \in \Lambda$ and $\psi_i \in \mathrm{Ter}(\mathbb{A})$. Then

$$\begin{array}{rcl} \alpha_{\mathcal{A}_2}(Lf([\lambda(\psi_1, \ldots, \psi_n)]_{L\mathbb{A}_1})) & = & \alpha_{\mathcal{A}_2}\left([\lambda\,(\mathrm{Ter}(f)(\psi_1), \ldots, \mathrm{Ter}(f)(\psi_n))]_{L\mathbb{A}_2}\right) \\ & = & \lambda^{\mathcal{A}_2}(\mathrm{Ter}(f)(\psi_1)^{\mathcal{A}_2}, \ldots, \mathrm{Ter}(f)(\psi_n)^{\mathcal{A}_2}) \end{array}$$

$$\begin{aligned}
&\overset{(*)}{=} \lambda^{\mathcal{A}_2}(f(\psi_1^{\mathcal{A}_1}), \ldots, f(\psi_n^{\mathcal{A}_1}))\\
&\overset{(*)}{=} f(\lambda^{\mathcal{A}_1}(\psi_1^{\mathcal{A}_1}, \ldots, \psi_n^{\mathcal{A}_1}))\\
&= f\left(\alpha_{\mathcal{A}_1}\left([\lambda(\psi_1, \ldots, \psi_n)]_{L\mathbb{A}_1}\right)\right)
\end{aligned}$$

where the equalities marked by $(*)$ hold because f is an $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$-morphism. QED

4.1.43. Theorem. *The categories* $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$ *and* $\mathsf{Alg}(L)$ *are isomorphic. Moreover the functors* $E : \mathsf{Alg}(L) \to \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax}))$ *and* $A : \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})) \to \mathsf{Alg}(L)$ *that witness this fact preserve the semantics in the following sense: For all formulas* $\phi \in \mathcal{L}(\Lambda)$ *we have*

$$\begin{aligned}
[\![\phi]\!]_{(\mathbb{A},\alpha)} &= [\![\phi]\!]_{E(\mathbb{A},\alpha)} && \textit{for all} \quad (\mathbb{A}, \alpha) \in \mathsf{Alg}(L)\\
[\![\phi]\!]_{\mathfrak{A}} &= [\![\phi]\!]_{A\mathfrak{A}} && \textit{for all} \quad \mathfrak{A} \in \mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})).
\end{aligned}$$

As a consequence the logic $L(\Lambda, \mathtt{Ax})$ *is sound and complete with respect to the algebraic semantics provided by* $\mathsf{Alg}(L)$*.*

Proof. Unfolding the definitions from the lemmas 4.1.40 and 4.1.42 it is easy to show that $A \circ E = Id_{\mathsf{Alg}(L)}$ and $E \circ A = Id_{\mathsf{Alg}(\mathsf{T}(\Lambda,\mathtt{Ax}))}$. The statement about the semantics can be proven by an easy induction on the structure of ϕ. QED

This isomorphism between categories allows us to give a concrete description of both the initial L-algebra and the initial $\mathfrak{L}$-algebra.

4.1.44. Definition. Let $\dashv\vdash := \{(\phi, \psi) \in \mathcal{L}(\Lambda)^2 \mid \mathtt{Ax} \vdash \phi \leftrightarrow \psi\}$ and $\mathcal{F}_\Lambda := \mathcal{L}(\Lambda)/\dashv\vdash$ the Lindenbaum algebra of $\mathsf{T}(\Lambda, \mathtt{Ax})$. We define an L-algebra $(\mathcal{I}, \alpha_\mathcal{I})$ by putting $\mathcal{I} := \mathcal{F}_\Lambda \in \mathsf{BA}$ (we forget the Λ-part of the signature) and

$$\begin{aligned}
\alpha_\mathcal{I} : L\mathcal{I} &\to \mathcal{I}\\
[t]_{L\mathcal{I}} &\mapsto t^{\mathcal{F}_\Lambda}.
\end{aligned}$$

4.1.45. Proposition. *The pair* $(\mathcal{I}, \alpha_\mathcal{I})$ *is isomorphic to the initial L-algebra. In particular this means that* $\alpha_\mathcal{I}$ *is an isomorphism.*

Proof. Immediate consequence of the isomorphism $\mathsf{Alg}(\mathsf{T}(\Lambda, \mathtt{Ax})) \cong \mathsf{Alg}(L)$ and the fact that $(\mathcal{I}, \alpha_\mathcal{I}) = A\mathcal{F}_\Lambda$. QED

The last proposition enables us to formulate the L-algebraic semantics of coalgebraic modal logic in a more categorical way.

4.1.46. Corollary. *Let* $(\mathbb{A}, \alpha) \in \mathsf{Alg}(L)$ *and* $(\mathcal{I}, \alpha_\mathcal{I})$ *the initial L-algebra. Then the interpretation function*

$$[\![\_]\!]_{(\mathbb{A},\alpha)} : \mathcal{L}(\Lambda) \to \mathbb{A}$$

that maps a formula to its L-algebraic semantics is given as the composition of the initial map $\mathsf{i}^{(\mathbb{A},\alpha)} : (\mathcal{I}, \alpha_\mathcal{I}) \to (\mathbb{A}, \alpha)$ *with the map* $[\_]_\mathcal{I} : \mathcal{L}(\Lambda) \to \mathcal{I}$ *as depicted in the* *following diagram:*

$$\begin{array}{ccccc} & & L\mathcal{I} & \xrightarrow{L\mathsf{i}^{(\mathbb{A},\alpha)}} & L\mathbb{A} \\ & & \downarrow{\scriptstyle \alpha_{\mathcal{I}}} & & \downarrow{\scriptstyle \alpha} \\ \mathcal{L}(\Lambda) & \xrightarrow{[-]_{\mathcal{I}}} & \mathcal{I} & \xrightarrow{\mathsf{i}^{(\mathbb{A},\alpha)}} & \mathbb{A} \end{array} \qquad \mathcal{L}(\Lambda) \xrightarrow{[\![-]\!]_{(\mathbb{A},\alpha)}} \mathbb{A}$$

4.1.47. Remark. In other words, the algebraic semantics of a formula ϕ on an L-algebra $(\mathbb{A}, \alpha)$ is given as the image of its equivalence class $[\phi]_{\mathcal{I}}$ in $\mathcal{I}$ under the initial map. Note that this is another way of seeing that the logic is sound with respect to the L-algebraic semantics: if $\mathtt{Ax} \vdash \phi$ then $[\phi]_{\mathcal{I}} = [\top]_{\mathcal{I}}$ and therefore $[\![\phi]\!]_{(\mathbb{A},\alpha)} = \top$. Later we will demonstrate that a similar argument works for proving soundness of the coalgebraic semantics provided that we can connect the algebraic and coalgebraic semantics via a natural transformation $\delta : L\mathsf{P} \Rightarrow \mathsf{P}T$.

In order to be able to represent the initial $\mathfrak{L}$-algebra we show that $\mathcal{L}(\Lambda)$ is in fact an absolutely free Σ_{BA}-algebra.

4.1.48. Lemma. *Let $G := \{\phi \in \mathcal{L}(\Lambda) \mid \phi \notin \mathbf{T}(\emptyset)\}$, i.e. G is the set of all formulas containing some $\lambda \in \Lambda$. Then $\mathcal{L}(\Lambda) = \mathbf{T}(G)$.*

Proof. Easy to check. QED

4.1.49. Proposition. *The pair $(\mathcal{L}(\Lambda), \dashv\vdash) \in \mathsf{PBA}$ together with the PBA-morphism $\overline{a}$: $(\mathsf{Lift}(\mathcal{L}(\Lambda)), \mathsf{Lift}(\dashv\vdash)) \to (\mathcal{L}(\Lambda), \dashv\vdash)$, which is defined as the equivalence class of the map $\phi \mapsto \phi^{\mathcal{L}(\Lambda)}$, is the initial $\mathfrak{L}$-algebra.*

Proof. First note that $(\mathcal{L}(\Lambda), \dashv\vdash)$ is indeed a pre-Boolean algebra, because we showed in Lemma 4.1.48 that $\mathcal{L}(\Lambda)$ is in fact an absolutely free Σ_{BA}-algebra. The claim follows then from the isomorphism between BA and PBA and Proposition 4.1.45. QED

4.2 Functor sequences

Before we use our algebraic semantics to study soundness and completeness conditions for coalgebraic modal logic we recall in this section the definition of the initial and the final sequence of a functor. Furthermore we show that the functors $\mathfrak{L}$ and L are *finitary*, i.e. their initial sequence converges after ω steps.

4.2.1 The initial sequence of L

Our first objective is to describe the initial L-algebra as the least fixed point of L. This fixed point can be obtained using the initial sequence (cf. e.g. [AK79]). We show

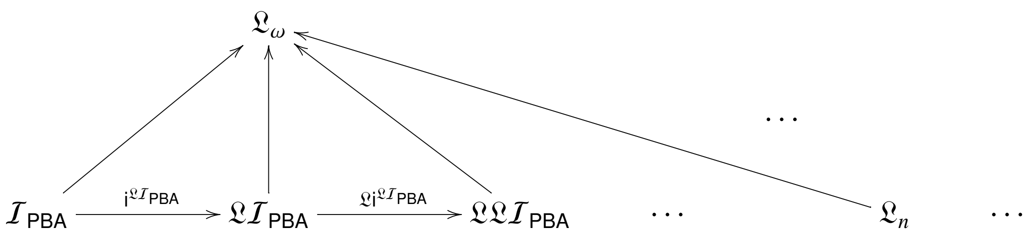

Figure 4.2: The initial ω-sequence of $\mathfrak{L} : \mathsf{PBA} \to \mathsf{PBA}$

that L is a so-called *finitary* functor, i.e. the initial sequence converges after ω steps. Therefore we only need to look at the first ω steps of the sequence.

Let $T : \mathsf{C} \to \mathsf{C}$ a functor, Λ a set of predicate liftings for T and $\mathtt{Ax}$ a set of axioms. Furthermore let $\mathfrak{L}$ and L the corresponding functors on PBA and BA respectively.

4.2.1. Definition. Let D be a category that has small colimits.[1] The ω-initial sequence of a functor $S : \mathsf{D} \to \mathsf{D}$ is given as a family of objects $\{S_i\}_{i \in \omega}$ given by

$$\begin{aligned} S_0 &:= \mathcal{I}_{\mathsf{D}} \\ S_{n+1} &:= S(S_n) \end{aligned}$$

together with a family of D-morphisms $\{f_i : S_i \to S_{i+1}\}_{i \in \omega}$, where $f_0 := \mathsf{i}^{S_1} : \mathcal{I}_{\mathsf{D}} \to S_1$ is the initial map and $f_{n+1} := S f_n$. Furthermore we define S_ω as the colimit of the sequence (cf. Figure 4.2). We say that the ω-initial sequence *converges to* $(SS_\omega, \alpha : SS_\omega \to S_\omega)$ if α is an isomorphism such that $\alpha \circ S f_i = f_{i+1}$ for all $i \in \omega$.

The ω-initial sequence of the functor $\mathfrak{L} : \mathsf{PBA} \to \mathsf{PBA}$ is useful for stratifying the language and the deducibility relation of coalgebraic modal logic. We first need the definition of the depth of a formula.

4.2.2. Definition. The *modal depth* of a formula $\phi \in \mathcal{L}(\Lambda)$ is defined inductively as follows:

$$\begin{aligned} \mathsf{d}(\bot) &:= 0 \\ \mathsf{d}(\phi \to \psi) &:= \max(\mathsf{d}(\phi), \mathsf{d}(\psi)) \\ \mathsf{d}([\lambda](\phi_1, \ldots, \phi_n)) &:= \max(\mathsf{d}(\phi_1), \ldots, \mathsf{d}(\phi_n)) + 1. \end{aligned}$$

The set of formulas of modal depth $\leq n$ will be denoted by $\mathcal{L}_n(\Lambda)$.

The n-th element of the ω-initial sequence of $\mathfrak{L}$ is based on those formulas which have modal depth less or equal to n.

4.2.3. Lemma. *For all $n \in \omega$ we have*

$$U\mathsf{Lift}^n(\mathbf{T}(\emptyset)) = \mathcal{L}_n(\Lambda).$$

[1] We will consider the cases $\mathsf{D} = \mathsf{PBA}$ and $\mathsf{D} = \mathsf{BA}$.

Proof. One only has to spell out the definition of Lift (cf. Def. 4.1.28) and to observe that $U\mathsf{Lift}^n(\mathbf{T}(\emptyset)) \subseteq U\mathsf{Lift}^{n+1}(\mathbf{T}(\emptyset))$. An easy induction on n completes the proof. QED

The ω-initial sequence of $\mathfrak{L}$ also gives rise to a stratification of the deducibility relation (cf. Definition 2.2.16).

4.2.4. Definition. For every $n \in \omega$ we define an equivalence relation $\equiv_n \subseteq \mathcal{L}_n(\Lambda) \times \mathcal{L}_n(\Lambda)$ by letting

$$\begin{aligned} \equiv_0 &:= \equiv_{\mathsf{BA}} \\ \equiv_{i+1} &:= \mathsf{Lift}(\equiv_i), \end{aligned}$$

where $\equiv_{\mathsf{BA}}$ denotes derivable equivalence of Boolean formulas in classical propositional logic.

We are now going to prove that $\equiv_n$ is equal to $\dashv\vdash$ restricted to formulas of modal depth less or equal to n. To avoid possible confusion the reader should note that our definition of $\equiv_n$ does not require that a derivation of $\phi \equiv_n \psi$ only contains formulas up to depth n. This marks a difference with the similar definition in [Pat03a].

4.2.5. Lemma. *Let* $\mathtt{Ax}$ *be consistent, i.e.* $\mathtt{Ax} \nvdash \bot$*, and* $\phi, \psi \in \mathcal{L}_n(\Lambda)$*. Then*

$$\mathtt{Ax} \vdash \phi \leftrightarrow \psi \quad \textit{iff} \quad \phi \equiv_n \psi.$$

In other words we have for $\phi, \psi \in \mathcal{L}_n(\Lambda)$ *that* $\phi \dashv\vdash \psi$ *iff* $\phi \equiv_n \psi$*.*

Proof. In order to prove that $\phi \equiv_n \psi$ implies $\mathtt{Ax} \vdash \phi \leftrightarrow \psi$ we prove the following slightly more general claim:

$$\text{for all } \phi, \psi \in \mathcal{L}(\Lambda) \quad E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_n \vdash_{\mathbf{EL}} \phi \approx \psi \quad \text{implies } E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \approx \psi.$$

The proof works by induction on n and on the height of the derivation of

$$E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_n \vdash_{\mathbf{EL}} \phi \approx \psi.$$

For $n = 0$ the claim is trivial. Suppose now $n = i + 1$, $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_{i+1} \vdash_{\mathbf{EL}} \phi \approx \psi$ and let l be the size of the derivation. In case $l = 0$ and $\phi \equiv_{i+1} \psi$ we get by the definition of $\equiv_{i+1}$ that $E_{\mathsf{BA}} \cup \mathtt{Ax} \cup \equiv_i \vdash_{\mathbf{EL}} \phi \approx \psi$ and therefore $E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \approx \psi$ by the induction hypothesis for i. In the other cases with $l = 0$ the claim is trivial. If $l = l' + 1$ we look at the last step of the derivation and use the induction hypothesis for l' to complete the proof.
For the other direction of the lemma it suffices to look at the case $n = 0$, as the other cases are trivial. Suppose that there are two Boolean formulas ϕ, ψ such that $E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \leftrightarrow \psi$ and suppose $\phi \not\equiv_0 \psi$, i.e. we cannot derive the equivalence between ϕ and ψ in classical propositional logic. Then the logic

$$\mathsf{L} := \{\phi \mid E_{\mathsf{BA}} \cup \mathtt{Ax} \vdash_{\mathbf{EL}} \phi \approx \top \text{ and } \phi \in \mathbf{T}(\emptyset)\}$$

is a proper extension of classical propositional logic that is closed under substitution and the rules of classical logic. Therefore L must be inconsistent, i.e. $\bot \in \mathsf{L}$. But this is a contradiction to our assumption that $\mathtt{Ax} \nvdash \bot$. QED

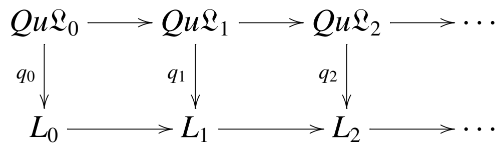

Figure 4.3: Relating L_n and $\mathfrak{L}_n$

We are now ready to show that the carrier of the initial $\mathfrak{L}$-algebra $(\mathcal{L}(\Lambda), \dashv\vdash)$ is a colimit of the ω-initial sequence of $\mathfrak{L}$ and conclude that the initial sequence of $\mathfrak{L}$ converges after ω steps.

4.2.6. PROPOSITION. *The ω-initial sequence of the functor $\mathfrak{L}$ converges to the initial $\mathfrak{L}$-algebra, i.e.*

$$\mathfrak{L}_\omega \cong (\mathcal{L}(\Lambda), \dashv\vdash).$$

Proof. We show that $(\mathcal{L}(\Lambda), \dashv\vdash)$ is the colimit of the ω-initial sequence of $\mathfrak{L}$. For every $n \in \omega$ we denote by $g_i : \mathfrak{L}_i \to (\mathcal{L}(\Lambda), \dashv\vdash)$ the inclusion map of $\mathcal{L}_i(\Lambda)$ into $\mathcal{L}(\Lambda)$, i.e. $g_i(\phi) = \phi$ for all $\phi \in \mathcal{L}_i(\Lambda)$. By Lemma 4.2.5 we know that all the g_i's are preserving and reflecting maps and therefore their equivalence classes $\overline{g}_i$ with respect to the $\sim$-relation from Definition 4.1.17 are reflecting PBA-morphisms. Furthermore $(\mathcal{L}(\Lambda), \dashv\vdash)$ together with the morphisms $\overline{g}_i$ forms a cocone over the ω-initial sequence of $\mathfrak{L}$ as depicted in the following diagram.

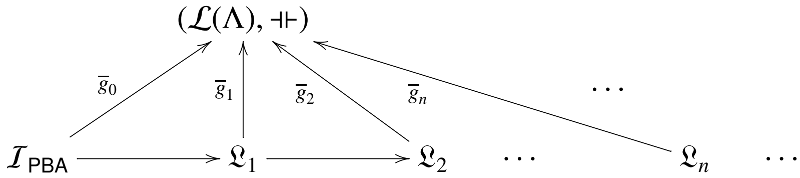

A standard argument shows that $(\mathcal{L}(\Lambda), \dashv\vdash)$ together with the family $\{\overline{g}_i\}_{i\in\omega}$ is indeed the colimit of the ω-initial sequence. Spelling out the definition of the isomorphism $\overline{a} : \mathfrak{L}(\mathcal{L}(\Lambda), \dashv\vdash) \to (\mathcal{L}(\Lambda), \dashv\vdash)$ from Proposition 4.1.49 we see that $\overline{g}_{i+1} = \overline{a} \circ \mathfrak{L}(\overline{g}_i)$ for all $i \in \omega$. Therefore we can conclude that the initial sequence of $\mathfrak{L}$ converges to the initial $\mathfrak{L}$-algebra $((\mathcal{L}(\Lambda), \dashv\vdash), \overline{a})$. QED

This result can be directly transferred to the initial sequence of L by relating the ω-initial sequence of L and the Qu-image of the ω-initial sequence of $\mathfrak{L}$ as depicted in Figure 4.3. The idea is as follows: From the equivalence between PBA and BA we deduce that the Boolean algebra $Qu\mathfrak{L}_n$ is equal to $\mathcal{L}_n(\Lambda)/\equiv_n$ which is in turn essentially the same as L_n. The difference between the two is that for calculating $Qu\mathfrak{L}_{n+1}$ we take the set of all formulas up to depth n, lift it and then quotient once by $\equiv_{n+1}$ whereas to define L_{n+1} we lift the Boolean algebra L_n (which correspond to the set of all formulas up to depth n quotiented by $\equiv_n$) and then take again a quotient, i.e. we need $n + 1$ quotient operations to construct L_{n+1} compared to 1 quotient operation for constructing $Qu\mathfrak{L}_n$. The q_n's make this connection precise.

4.2.7. Definition. For each $n \in \omega$ we define a function $q_n : Qu\mathfrak{L}_n \to L_n$, i.e. from the quotient of the n-th element of the ω-initial sequence of $\mathfrak{L}$ to the n-th element of the ω-initial sequence of L, by defining $q_0 : Qu\mathfrak{L}_0 \to L_0$ to be the isomorphism between $Qu\mathfrak{L}_0 = \mathbf{T}(\emptyset)/\equiv_{\mathsf{BA}}$ and the two-element Boolean algebra $L_0 = 2$ and by letting

$$\begin{array}{rcl} q_{n+1} : Qu\mathfrak{L}_{n+1} & \to & L_{n+1} \\ [\lambda(\psi_1, \dots, \psi_m)]_{Qu\mathfrak{L}_{n+1}} & \mapsto & \left[\lambda\left(q_n([\psi_1]_{Qu\mathfrak{L}_n}), \dots, q_n([\psi_m]_{Qu\mathfrak{L}_n})\right)\right]_{L_{n+1}}. \end{array}$$

Before we show that the q_n's are isomorphisms we first check whether our intuition about $Qu\mathfrak{L}_n$ is correct.

4.2.8. Lemma. *For all $n \in \omega$ we have $Qu\mathfrak{L}_n = \mathcal{L}_n(\Lambda)/\equiv_n$.*

Proof. The claim follows from $\mathfrak{L}_n = (\mathsf{Lift}^n(\mathbf{T}(\emptyset)), \mathsf{Lift}^n(\equiv_{\mathsf{BA}}))$ and from the fact that $\mathsf{Lift}^n(\mathbf{T}(\emptyset)) = \mathcal{L}_n(\Lambda)$ (cf. Lemma 4.2.3) and $\mathsf{Lift}^n(\equiv_{\mathsf{BA}}) = \equiv_n$ (cf. Def. 4.2.4). QED

4.2.9. Lemma. *For all $n \in \omega$ the function $q_n : Qu\mathfrak{L}_n \to L_n$ is a* BA*-isomorphism.*

Proof. An easy proof by induction on n. QED

The lemma establishes a one-to-one correspondence between equivalence classes in L_n and equivalence classes of formulas in $Qu\mathfrak{L}_n$. Later it will be convenient to directly talk about the equivalence class in L_n which corresponds to a certain formula $\phi \in \mathcal{L}_n(\Lambda)$. To that aim we introduce the following notation.

4.2.10. Definition. For a formula $\phi \in \mathcal{L}_n(\Lambda)$ we put

$$\langle\phi\rangle_{L_n} := q_n([\phi]_{Qu\mathfrak{L}_n}).$$

The close connection between $\mathfrak{L}_n$ and L_n ensures that the ω-initial sequence of L converges.

4.2.11. Corollary. *The initial sequence of L converges after ω steps and we have $L_\omega \cong Qu(\mathfrak{L}_\omega) \cong \mathcal{L}(\Lambda)/\dashv\vdash$.*

Proof. From Proposition 4.2.6 we know that the ω-initial sequence of $\mathfrak{L}$ converges after ω steps and that $Qu\mathfrak{L}_\omega = \mathcal{L}(\Lambda)/\dashv\vdash$. According to Lemma 4.2.9 the ω-initial sequence of L is connected via isomorphisms q_n to the Qu-image of the ω-initial sequence of $\mathfrak{L}$ as depicted in Figure 4.3. Because Qu is part of an equivalence of categories it preserves colimits and hence it is easy to see that L_ω is isomorphic to $Qu\mathfrak{L}_\omega = \mathcal{L}(\Lambda)/\dashv\vdash \cong \mathcal{I}$, where $\mathcal{I}$ is the carrier of the initial L-algebra. Hence $LL_\omega \cong L_\omega$, i.e. the initial sequence converges. QED

4.2.2 The final sequence of T

Dually to the construction of initial and free algebras as colimits of the initial sequence of the underlying endofunctor, the *final* or *terminal sequence* plays an important role in the coalgebraic framework. Like for the initial sequence in the previous section, it is sufficient to consider the finitary part, that is, the first ω elements of the final sequence. For a demonstration of the usefulness of the final sequence we refer the reader to Worrell [Wor05].

4.2.12. DEFINITION. Let C be a category with final object $\mathbf{1}$ and $T : \mathsf{C} \to \mathsf{C}$ a functor. Then the *final sequence* of T consists of a sequence of objects $(T_i)_{i\in\omega}$ given by

$$\begin{aligned} T_0 &:= \mathbf{1} \\ T_{i+1} &:= TT_i. \end{aligned}$$

and a sequence of morphisms $(p_i : T_{i+1} \to T_i)_{i\in\omega}$ defined as

$$\begin{aligned} p_0 &:= \;!_{T\mathbf{1}} : T\mathbf{1} \to \mathbf{1} \quad \text{(the unique morphism into the final object)} \\ p_{i+1} &:= Tp_i. \end{aligned}$$

To illustrate the importance of the final sequence for the theory of coalgebras we mention the following theorem.

4.2.13. THEOREM ([WOR05]). *Let κ be a regular cardinal and $T : \mathsf{Set} \to \mathsf{Set}$ be a κ-accessible functor*[2]. *Then the final sequence of T converges after $\kappa + \kappa$ steps to the final T-coalgebra.*

In our logical context the object T_n encodes the types of behaviours which can be described with a formula of depth n. This intuition is made formal by making use of the well-known fact that every T-coalgebra (X, γ) can be seen as a cone over the final sequence of T.

4.2.14. DEFINITION. Let $(X, \gamma) \in \mathsf{Coalg}(T)$. Then we define the sequence of *n-step behaviour maps* $(\gamma_n : X \to T_n)_{n\in\omega}$ by letting

$$\begin{aligned} \gamma_0 &:= \;!_X \\ \gamma_{i+1} &:= T\gamma_i \circ \gamma. \end{aligned}$$

4.2.15. LEMMA. *For every coalgebra (X, γ) the carrier X together with the sequence of n-step behaviour maps $(\gamma_n : X \to T_n)_{n\in\omega}$ defines a cone over the final sequence of T.*

Proof. Easy to check. QED

[2] A Set-functor T is κ-accessible iff for all sets X we have $TX = \bigcup\{TY \mid Y \subseteq X \text{ and } |Y| < \kappa\}$.

This cone consisting of the n-step behaviour maps can be used to define an n-step semantics of coalgebraic logic. This is done by induction on n using the following lifting construction, which lifts the interpretation of formulas of depth n to an interpretation of formulas of depth $n + 1$. The following definition is crucial in what follows. It defines how we lift a function, that interprets elements of a set (of formulas) Φ as predicates over some $X \in \mathsf{C}$, one level higher, i.e. to a function that interprets lifted formulas in $\mathsf{Lift}(\Phi)$ as predicates over TX.

4.2.16. Definition. Let Φ be a set (of formulas) and $d : \Phi \to \mathrm{P}X$ be a function (interpreting elements of Φ as predicates over X). Then we define a lifted function $\mathsf{Lift}(d) : \mathsf{Lift}(\Phi) \to \mathrm{P}TX \in \mathsf{Alg}(\Sigma_{\mathsf{BA}})$ as the inductively extension of the map

$$\lambda(\phi_1, \dots, \phi_n) \mapsto \lambda_X(d(\phi_1), \dots, d(\phi_n)).$$

Using this lifting of interpretation functions we define a sequence of Σ_{BA}-morphisms $(d_i : \mathcal{L}_i(\Lambda) \to \mathrm{P}T_i)_{i \in \omega}$ by letting

$$\begin{aligned} d_0 &:= \mathsf{i}^{\mathbf{P1}} : \mathbf{T}(\emptyset) \to \mathbf{P1} \text{ (the initial map)} \\ d_{i+1} &:= \mathsf{Lift}(d_i), \end{aligned}$$

and call $d_n(\phi)$ the *n-step semantics* of ϕ.

The connection with the coalgebraic semantics from Definition 2.2.5 is as follows.

4.2.17. Proposition. *Let* $(X, \gamma) \in \mathsf{Coalg}(T)$ *and* $\phi \in \mathcal{L}_n(\Lambda)$. *Then*

$$[\![\phi]\!]_{(X,\gamma)} = \mathrm{P}(\gamma_n)(d_n(\phi)),$$

i.e. $(X, \gamma), x \models \phi$ *iff* $\gamma_n(x) \in d_n(\phi)$.

Proof. This follows from an easy proof by induction on n. QED

Expressed in words the proposition means that the semantics $[\![\phi]\!] \subseteq TX$ of a formula ϕ with modal depth n is already determined by its n-step semantics, i.e. by the set $d_n(\phi) \subseteq T_n$ of n-behaviours that it specifies.

In [Pat03a] Pattinson gave sufficient conditions for proving soundness and completeness of a coalgebraic modal logic for a Set-endofunctor. His conditions are formulated in terms of the lifting $\mathsf{Lift}(h)$ of functions $h : (\mathbb{A}, \equiv) \to \mathcal{P}X$ (cf. Def. 4.2.16). We will now recall Pattinson's result and later provide later an algebraic interpretation. To state the result we have to introduce some terminology.

4.2.18. Definition. Suppose `Ax` is a set of axioms.

(i) `Ax` is *order-preserving* iff for all preserving maps $h : (\mathbb{A}, \equiv) \to \mathrm{P}X$, $(\mathbb{A}, \equiv) \in \mathsf{PBA}$ and $X \in \mathsf{C}$ we have that $\mathsf{Lift}(h) : (\mathsf{Lift}(\mathbb{A}), \mathsf{Lift}(\equiv)) \to \mathrm{P}TX$ is a preserving map.

(ii) Ax is *order-reflecting* iff for all reflecting maps $h : (\mathbb{A}, \equiv) \to PX$, $(\mathbb{A}, \equiv) \in \mathsf{PBA}$, and $X \in \mathsf{C}$ we have that $\mathsf{Lift}(h)$ is a reflecting map.

4.2.19. Remark. Note that these conditions are not exactly the same as in [Pat03a]. First Pattinson deals in *loc.cit.* only with the case $\mathsf{C} = \mathsf{Set}$. Second he formulates the condition not only for preserving maps but for arbitrary functions $h : \mathbb{A} \to PX$ that preserve the order, i.e. functions h such that $a \equiv a'$ imply $h(a) = h(a')$.

Intuitively Pattinson's condition says, that we can lift soundness and completeness of the logic step-by-step: we can look at a preserving map $h : (\mathbb{A}, \equiv) \to PX$ as a sound interpretation function because elements a_1, a_2 of $\mathbb{A}$ ("formulas") for which we have $a_1 \equiv a_2$ ("a_1 and a_2 are logically equivalent") are interpreted by the same predicate over X. Analogously, reflecting maps correspond to an interpretation that is complete. An order-preserving and -reflecting set of axioms makes it possible to lift this abstract soundness and completeness one-step higher, i.e. to formulas whose modal depth is increased by one. In [Pat03a] it was shown that these conditions are sufficient to prove soundness and completeness of the logic with respect to the coalgebraic semantics.

4.2.20. Theorem ([Pat03a]). *Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor, Λ a set of predicate liftings and* Ax *be a set of axioms.*

(i) If Ax *is order-preserving, then $L(\Lambda, \mathtt{Ax})$ is sound.*

(ii) If Ax *is order-preserving and -reflecting then $L(\Lambda, \mathtt{Ax})$ is complete.*

We will now move to the definition of the natural transformation $\delta : LP \Rightarrow PT$ that connects the algebraic and the coalgebraic semantics of the logic. After that we will relate Pattinson's result to properties of δ (cf. Theorem 4.4.1).

4.3 Coalgebraic semantics as a natural transformation

We are now going to connect the coalgebraic semantics and the algebraic semantics of coalgebraic modal logic by defining a natural transformation

$$\delta : LP \Rightarrow PT.$$

We will demonstrate which properties of this natural transformation imply soundness and completeness of the axioms and expressiveness of the language.

The functor sequences that we saw in the last section will be an important tool. Based on ideas of Pattinson from [Pat03a] our proofs will be of the following form: we have soundness and completeness for our base logic, i.e. for propositional logic. If δ has the right properties we can lift soundness and completeness along the initial sequence of L: if the logic restricted to formulas of depth at most n, i.e. the set $\mathcal{L}_n(\Lambda)$ of formulas of depth at most n together with the congruence relation $\equiv_n$, is sound and complete with respect to the n-step semantics from Definition 4.2.16, then the logic restricted to formulas of depth $n+1$ will be sound and complete with respect to the $n+1$-step semantics. In this way we obtain soundness and completeness for the logic with respect to the coalgebraic semantics.

We will see in Section 4.4 that for the case $\mathsf{C} = \mathsf{Set}$ our criteria for soundness and completeness of $L(\Lambda, \mathtt{Ax})$ and for expressiveness of the language are in fact equivalent to Pattinson's conditions. We improve on his earlier results in the following ways:

- The conditions formulated in terms of δ work both in the case $\mathsf{C} = \mathsf{Set}$ and $\mathsf{C} = \mathsf{Stone}$. Furthermore everything is formulated at a level of abstraction which makes it possible to generalize the work to similar dualities.
- Pattinson's earlier work seemed to be conceptually different from existing soundness and completeness proofs in modal logic. Our results show that his work matches with existing work on the algebraic semantics of modal logic.

4.3.1 Definition of δ

4.3.1. Definition. Let $X \in \mathsf{C}$. Then we define a function

$$d_X : (\mathsf{Lift}(\mathrm{Ter}(\mathrm{P}X)), \mathsf{Lift}(\mathrm{Diag}(\mathrm{P}X))) \to \mathrm{P}TX$$

as $d_X := \mathsf{Lift}((\_)^{\mathrm{P}X})$, where $(\_)^{\mathrm{P}X} : \mathrm{Ter}(\mathrm{P}X) \to \mathrm{P}X$ is the function which maps a term over $\mathrm{P}X$ (cf. Definition 4.1.10) to its interpretation in $\mathrm{P}X$ and $\mathsf{Lift}((\_)^{\mathrm{P}X})$ is defined as in Definition 4.2.16 on page 94. In case d_X is a preserving map, then we denote by

$$\delta_X : L\mathrm{P}X \to \mathrm{P}TX$$

the BA-homomorphism $\delta_X := Qu(d_X)$ that is defined as in Lemma 4.1.14, i.e. $\delta_X([a]_{L\mathrm{P}X}) = d_X(a)$.

4.3.2. Remark. The intuition behind the function d_X is the following: the map $(\_)^{\mathrm{P}X}$: $(\mathrm{Ter}(\mathrm{P}X), \mathrm{Diag}(\mathrm{P}X)) \to \mathrm{P}X$ can be seen as the most basic interpretation function that we can think of. Boolean terms that are built up from predicates over X are again interpreted as predicates over X, where we interpret the Boolean operators by their set-theoretic counterparts. Trivially this interpretation function is sound and complete in the sense that terms are identified if and only if they are ("logically") equivalent w.r.t. $\mathrm{Diag}(\mathrm{P}X)$. The function d_X is the lifted version of this sound and complete semantics.

Under the proviso that d_X is preserving, it factors through some Boolean homomorphism δ_X. If all components of δ_X exist we obtain a natural transformation.

4.3.3. Lemma. *Suppose that δ_X exists for all $X \in \mathsf{C}$, i.e. d_X is preserving for all $X \in \mathsf{C}$. Then the family $(\delta_X)_{X \in \mathsf{C}}$ gives rise to a natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$.*

Proof. Suppose δ_X exists for all $X \in \mathsf{C}$ and let $h : Y \to X \in \mathsf{C}$. We have to show that the following diagram commutes:

$$\begin{array}{ccc} LPX & \xrightarrow{\delta_X} & PTX \\ {\scriptstyle LPh}\downarrow & & \downarrow{\scriptstyle PTh} \\ LPY & \xrightarrow{\delta_Y} & PTY \end{array}$$

It suffices to prove the commutativity of the diagram only for the generators of LPX:

$$\begin{aligned}
PTh(\delta_X([\lambda(\psi_1,\ldots,\psi_n)]_{LPX})) &= PTh(d_X(\lambda(\psi_1,\ldots,\psi_n))) \\
&= PTh(\lambda_X(\psi_1^{PX},\ldots,\psi_n^{PX})) \\
&\overset{\text{naturality of }\lambda}{=} \lambda_Y(Ph(\psi_1^{PX}),\ldots,Ph(\psi_n^{PX})) \\
&\overset{\text{Lemma 4.1.15}}{=} \lambda_Y\left((\mathrm{Ter}(Ph)(\psi_1))^{PX},\ldots,(\mathrm{Ter}(Ph)(\psi_n))^{PX}\right) \\
&\overset{\text{Def. of }d_Y}{=} d_Y\left(\lambda\left(\mathrm{Ter}(Ph)(\psi_1),\ldots,\mathrm{Ter}(Ph)(\psi_n)\right)\right) \\
&\overset{\text{Def. of }\delta_Y}{=} \delta_Y\left([\lambda(\mathrm{Ter}(Ph)(\psi_1),\ldots,\mathrm{Ter}(Ph)(\psi_n))]_{LPY}\right) \\
&\overset{\text{Lemma 4.1.38}}{=} \delta_Y\left(LPh([\lambda(\psi_1,\ldots,\psi_n)]_{LPX})\right)
\end{aligned}$$

QED

4.3.2 A functor linking algebraic and coalgebraic semantics

Different properties of $\delta : LP \Rightarrow PT$ correspond to different properties of the logic. This can be explained as follows: We know that the logic is always sound and complete with respect to its algebraic semantics $\mathsf{Alg}(L)$. The existence of δ on the other hand implies the existence of a contravariant functor from $\mathsf{Coalg}(T)$ to $\mathsf{Alg}(L)$. This functor preserves the semantics and therefore the existence of this functor implies soundness of the logic. The proof of the fact that the injectivity of δ implies completeness is more technical and uses the final sequence of the functor T. We postpone it until the next section.

We will now present how we can use δ to define a functor $\Delta : \mathsf{Coalg}(T) \to \mathsf{Alg}(L)$ and then derive soundness from its existence. A remark on expressivity in the case $\mathsf{C} = \mathsf{Stone}$ will round up this section.

4.3.4. Definition. Suppose $\delta : LP \Rightarrow PT$ exists. Then we define a functor

$$\begin{array}{rcl} \Delta : \mathsf{Coalg}(T)^{op} & \to & \mathsf{Alg}(L) \\ (X,\gamma) & \mapsto & (PX, P\gamma \circ \delta_X) \\ f : (X,\gamma) \to (Y,\rho) & \mapsto & Pf. \end{array}$$

This functor is obviously well-defined on objects. The naturality of δ ensures that Δ is also well-defined on morphisms. In order to see that let $f : (X,\gamma) \to (Y,\xi)$ be a T-coalgebra morphism and let us take a look at the following diagram:

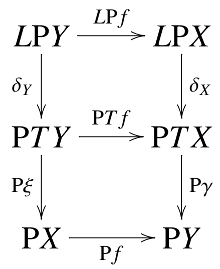

The lower half of the diagram commutes because f was assumed to be a T-coalgebra morphism and the upper half commutes by the naturality of δ. Hence the whole diagram commutes which means that $\mathrm{P}f$ is an L-morphism from $\Delta(Y, \xi)$ to $\Delta(X, \gamma)$.

The functor preserves the semantics in the following sense.

4.3.5. Lemma. *Let* $(X, \gamma) \in \mathsf{Coalg}(T)$*, then* $[\![-]\!]_{(X,\gamma)} = [\![-]\!]_{\Delta(X,\gamma)}$*, i.e. the coalgebraic semantics of a formula in* (X, γ) *is equal to its algebraic semantics in* $\Delta(X, \gamma)$*.*

Proof. Easy proof by induction on the structure of the formula. QED

We can express the statement of the lemma also with the help of the diagram in Figure 4.4: the coalgebraic semantics on (X, γ) is equal to the initial map from the initial L-algebra $(\mathcal{I}, \alpha_{\mathcal{I}})$ into $\Delta(X, \gamma)$. The proof that the existence of δ implies soundness of the logic is short and similar to what we remarked earlier about soundness with respect to the L-algebraic semantics in Remark 4.1.47.

4.3.6. Proposition. *If* $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ *exists , then* $L(\Lambda, \mathtt{Ax})$ *is sound, i.e.* $\mathsf{Coalg}(T) \models \phi$ *if* $\mathtt{Ax} \vdash \phi$ *for all* $\phi \in \mathcal{L}(\Lambda)$*.*

Proof. If δ exists we can define the functor Δ as in Definition 4.3.4. Suppose now that $\mathtt{Ax} \vdash \phi$ and let $(X, \gamma) \in \mathsf{Coalg}(T)$. Then clearly $[\phi]_{\mathcal{I}} = [\top]_{\mathcal{I}}$, i.e. ϕ and $\top$ are in the same equivalence class in $\mathcal{I}$, and therefore by the diagram in Figure 4.4 we get $[\![\phi]\!]_{(X,\gamma)} = [\![\top]\!]_{(X,\gamma)} = X$. QED

Consider now the case $\mathsf{C} = \mathsf{Stone}$. Then the functor Δ is in fact an equivalence of categories, provided that δ is a bijection.

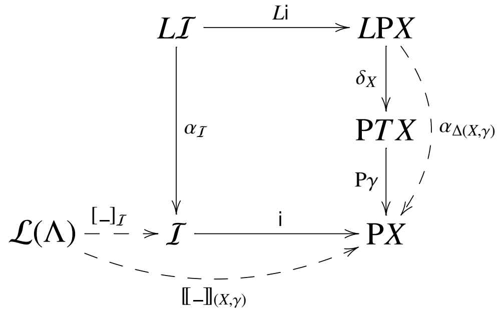

Figure 4.4: Coalgebraic semantics as an initial map

4.3.7. LEMMA. *If* $\delta : L\mathbb{C}\mathrm{lp} \Rightarrow \mathbb{C}\mathrm{lp}T$ *exists and is bijective then* $\Delta : \mathsf{Coalg}(T)^{\mathrm{op}} \to \mathsf{Alg}(L)$ *is an equivalence.*

Proof. The proof, which essentially uses Stone duality, is straightforward. One only has to observe that the δ^{-1}, the inverse of δ, enables us to define a contravariant functor $\Delta^{-1} : \mathsf{Alg}(L)^{\mathrm{op}} \to \mathsf{Coalg}(T)$. QED

This enables us to give a proof of the fact that surjectivity of δ implies expressivity of the logic.

4.3.8. PROPOSITION. *Suppose* $\delta : L\mathbb{C}\mathrm{lp} \Rightarrow \mathbb{C}\mathrm{lp}T$ *is bijective. Then* $\mathcal{L}(\Lambda)$ *is expressive.*

Proof. We only sketch the proof. As δ is bijective we know that the categories $\mathsf{Coalg}(T)$ and $\mathsf{Alg}(L)$ are dually equivalent. In particular the dual of the initial L-algebra will be the final coalgebra. It is easy to see that states of coalgebras, that satisfy the same formulas of the logic, can be mapped to the same state in the final coalgebra. Hence states that satisfy the same formulas can be identified by coalgebra morphisms and are therefore bisimilar. QED

4.3.3 Completeness

In the last section we saw that the existence of $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ implies soundness of $L(\Lambda, \mathtt{Ax})$. Relating injectivity of δ to completeness of the logic $L(\Lambda, \mathtt{Ax})$ requires a bit more technical machinery. Throughout this section we assume that the natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ exists.

Recall that in order to show that the logic is complete we have to prove that the following holds true for an arbitrary formula $\phi \in \mathcal{L}(\Lambda)$:

$$\mathsf{Coalg}(T) \models \phi \quad \text{implies} \quad \mathtt{Ax} \vdash \phi.$$

For many modal logics completeness with respect to a class of Kripke frames can be proven using a canonical model construction. This proof can be roughly sketched as follows: The set of states of the canonical model of a modal logic is equal to the set of maximal consistent sets of formulas of the logic and a formula is true at a state iff the formula is contained in the corresponding maximal consistent set. If a formula ϕ is valid on any Kripke frame of the logic then it will be in particular true in all states of the canonical model, provided that the canonical model is based on a frame of the logic. From the fact that ϕ is true in all states of the canonical model one can deduce that ϕ is contained in any maximal consistent set of the logic and therefore ϕ is a theorem of the modal logic under consideration.

There are, however, cases in which a proof along these lines doesn't work. One reason for this is the fact that not all modal logics are strongly complete with respect to Kripke semantics. Omitting the details we can think of strong completeness of a modal logic as of completeness with the additional requirement that every maximal consistent set of the logic is satisfiable in a Kripke model of that logic. Modal logics that allow for a completeness proof using the canonical model construction are strongly complete: every maximal consistent set of formulas is satisfiable in the canonical model.

The situation for coalgebraic modal logics is very similar: the canonical model of a logic for a Vietoris polynomial functor corresponds to its final coalgebra, and completeness of the logic follows from the fact that every formula that is satisfied on all states of the final coalgebra is a theorem of MSM$\mathcal{L}_T$ (cf. Corollary 3.4.11). Again not all coalgebraic modal logics are strongly complete with respect to their coalgebraic semantics.

4.3.9. Example. The standard example for this situation is probably the modal logic **K** on image-finite Kripke frames which is known to be sound and complete, but which is not strongly complete. In our coalgebraic setting this logic can be represented as the logic $L(\Lambda, \mathtt{Ax})$ for the finite power set functor $\mathcal{P}_\omega : \mathsf{Set} \to \mathsf{Set}$ that is given by the predicate liftings $\Lambda := \{\lambda^{\Box_\omega}, \lambda^{\Diamond_\omega}\}$ from Example 2.2.12 and by the set of axioms of **K** (cf. Example 2.2.18).

The argument for why $L(\Lambda, \mathtt{Ax})$ is not strongly complete with respect to $\mathcal{P}_\omega$-coalgebras can be sketched as follows: For each $i \in \omega$ let φ_i be the formula $\Diamond^i \top \wedge \Box^{i+1} \perp$, where $\Box := [\lambda^{\Box_\omega}]$ and $\Diamond := [\lambda^{\Diamond_\omega}]$. Then it is not difficult to see that

$$\mathsf{Coalg}(\mathcal{P}_\omega) \models \varphi_i \to \neg\varphi_j \quad \text{for } i \neq j,\ i, j \in \omega \tag{4.3}$$

Furthermore one can show that the set of formulas $\Phi := \{\Diamond\varphi_i \mid i \in \omega\}$ is consistent because it is satisfiable in some $\mathcal{P}$-coalgebra (note that the logic for $\mathcal{P}$- and $\mathcal{P}_\omega$-coalgebras is in both cases equal to **K**). Because of equation (4.3), however, it is easy to see that Φ is not satisfiable in a state that has only finitely many successors. Hence the set Φ is not satisfiable in a state of a $\mathcal{P}_\omega$-coalgebra. The consistent set of formulas Φ can be extended to a maximal consistent Φ' of formulas, which is of course also not satisfiable on a $\mathcal{P}_\omega$-coalgebra. This means that the logic is not strongly complete with respect to the class of $\mathcal{P}_\omega$-coalgebras.

Therefore we cannot use a canonical model argument for showing that the injectivity of δ implies completeness with respect to the coalgebraic semantics without making any additional assumptions. Instead we follow Pattinson's ideas from [Pat03a].

His observation in *loc.cit.* was that we can use the final sequence of the functor T to obtain for each n a coalgebra C_n that we call the *canonical model for formulas up to modal depth n*. This coalgebra C_n has the property that for all formulas of depth n the n-step semantics $d_n : \mathcal{L}_n(\Lambda) \to \mathrm{P}T_n$ coincides with the coalgebraic semantics, i.e.

$$[\![\phi]\!]_{C_n} = d_n(\phi) \quad \text{for all } \phi \in \mathcal{L}_n(\Lambda), \tag{4.4}$$

where the n-step semantics d_n is defined as the function that maps formulas of modal depth smaller or equal to n to predicates over T_n, the n-th element of the final sequence of T (cf. Definition 4.2.16 on page 94).

What is still missing to prove completeness of the logic using Equation 4.4 is the following fact, which we call *n-step completeness*: for all formulas $\phi, \psi \in \mathcal{L}_n(\Lambda)$ we have

$$d_n(\phi) = d_n(\psi) \quad \text{implies} \quad \phi \equiv_n \psi. \tag{4.5}$$

This equation will follow from the injectivity of the natural transformation δ.

Both equations together suffice to prove completeness of $L(\Lambda, \mathtt{Ax})$ with respect to the coalgebraic semantics: Let $\phi \in \mathcal{L}_n(\Lambda)$ be a formula of modal depth n such that $\mathsf{Coalg}(T) \models \phi$. Then in particular $C_n \models \phi$, i.e. $[\![\phi]\!]_{C_n} = \top$. By Equation 4.4 this means that $d_n(\phi) = \top$ and hence by Equation 4.5 we get $\phi \equiv_n \top$. As a consequence we get by Lemma 4.2.5 that $\mathtt{Ax} \models \phi$.

We now proceed as follows: we first construct for every n the T-coalgebra C_n and prove that Equation 4.4 is indeed true. Furthermore we prove that the injectivity of δ implies n-step completeness (Equation 4.5). Finally we obtain the result that completeness of the logic is entailed by injectivity of δ.

We start by defining for all $n \in \omega$ the T-coalgebra C_n.

4.3.10. Definition. A *global element* of a functor $T : \mathsf{C} \to \mathsf{C}$ is an arrow $\mathsf{e} : 1 \to T1$. Let $\mathsf{e}(0) : 1 \to T1$ be such a global element of T and define

$$\mathsf{e}(i+1) := T\mathsf{e}(i) : T_i \to T_{i+1},$$

where T_i denotes the set $T^i 1$, i.e. the ith element of the final sequence of T. Then for all $n \in \omega$ we define C_n, the *canonical model for formulas of modal depth up to n*, to be the T-coalgebra $(T_n, \mathsf{e}(n) : T_n \to T_{n+1}) \in \mathsf{Coalg}(T)$. Any such T-coalgebra gives rise to a cone over the final sequence of T as described in Definition 4.2.14 on page 93. We denote by $(\mathsf{e}(n)_k : T_n \to T_k)_{k \in \omega}$ the sequence of n-step behaviour maps into the final sequence.

The key of proving that Equation 4.4 holds is the following technical lemma by Pattinson.

4.3.11. Lemma ([Pat03a]). *For all $n \in \omega$ and all $k \leq n$ we have $\mathsf{e}(n)_k = T^k(!_{T_{n-k}})$. In particular $\mathsf{e}(n)_n = id_{T_n}$.*

4.3.12. Remark. Pattinson's proof of this lemma is only formulated for $\mathsf{C} = \mathsf{Set}$ but works without changes also for the case $\mathsf{C} = \mathsf{Stone}$.

We are now ready to show that Equation 4.4 holds:

4.3.13. Proposition. *Let $n \in \omega$ and let $C_n = (T_n, \mathsf{e}(n))$ be defined as in Definition 4.3.10. Then for all formulas $\phi \in \mathcal{L}_n(\Lambda)$ we have $[\![\phi]\!]_{C_n} = d_n(\phi)$, where again $d_n : \mathcal{L}_n(\Lambda) \to \mathrm{P}T_n$ is the n-step semantics of ϕ.*

Proof. Let $n \in \omega$ and $C_n = (T_n, \mathsf{e}(n))$. Furthermore recall from Lemma 4.2.17 on page 94 that the semantics $[\![\phi]\!]_{C_n}$ of a formula $\phi \in \mathcal{L}_n(\Lambda)$ is computed as follows:

$$[\![\phi]\!]_{(T_n,\mathsf{e}(n))} = \mathrm{P}(\mathsf{e}(n)_n)(d_n(\phi)).$$

But $\mathsf{e}(n)_n$ is equal to id_{T_n} according to Lemma 4.3.11 and hence also $\mathrm{P}(\mathsf{e}(n)_n) = id_{\mathrm{P}T_n}$. Putting everything together we arrive at

$$\begin{aligned} [\![\phi]\!]_{(T_n,\mathsf{e}(n))} &= \mathrm{P}(\mathsf{e}(n)_n)(d_n(\phi)) \\ &= d_n(\phi). \end{aligned}$$

QED

We now turn to the proof of the fact that the injectivity of $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ implies n-step completeness, i.e. we show that equation (4.5) from page 101 holds if δ is injective. Our first observation is that the existence of δ implies that the n-step semantics $d_n : \mathcal{L}_n(\Lambda) \to \mathrm{P}T_n$ factors through a Boolean homomorphism $\delta_n : L_n \to \mathrm{P}T_n$. We first define the δ_n's and then prove that d_n factors through δ_n.

4.3.14. Definition. For $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ we define for each $n \in \omega$ a BA-morphism $\delta_n : L_n \to \mathrm{P}T_n$ by putting

$$\begin{aligned} \delta_0 &:= \mathsf{i}^{\mathrm{P}\mathbf{1}} \quad \text{(the initial map from } L_0 = 2 \text{ to } \mathrm{P}T_0 = \mathrm{P}\mathbf{1}) \\ \delta_{n+1} &:= \delta_{T_n} \circ L\delta_n, \end{aligned}$$

where we denoted the n-th element of the initial sequence of L and the n-th element of the final sequence of T by L_n and T_n respectively.

The following observation is immediate from the fact that L preserves the injectivity and surjectivity of a morphism (cf. Proposition 4.1.39).

4.3.15. Proposition. *For all $n \in \omega$ the Boolean homomorphism δ_n is injective or surjective if δ is injective or surjective respectively.*

Proof. The claim can be proven with an easy induction on n. The map δ_0 is always an isomorphism. Inductively assume that δ_n is injective (surjective). Then $L\delta_n$ is also injective (surjective) by Prop. 4.1.39. But this implies injectivity (surjectivity) of $\delta_{n+1} = \delta_{T_n} \circ L\delta_n$ under the condition that δ is injective (surjective). QED

The n-step semantics d_n factors through δ_n.

4.3.16. Lemma. *For all $n \in \omega$ and all $\phi \in \mathcal{L}_n$ we have $d_n(\phi) = \delta_n(\langle\phi\rangle_{L_n})$, i.e.*

$$\mathcal{L}_n(\Lambda) \xrightarrow{\langle - \rangle_{L_n}} L_n \xrightarrow{\delta_n} \mathrm{P}T_n, \qquad d_n : \mathcal{L}_n(\Lambda) \to \mathrm{P}T_n$$

commutes for all $n \in \omega$.

Proof. The lemma is proven by induction on n. For $n = 0$ the n-step semantics $d_0 = \mathsf{i}^{\mathrm{P}1} : \mathbf{T}(\emptyset) \to \mathrm{P}1$ is the initial Σ_{BA}-morphism. The composition of maps $\delta_0 \circ \langle \_ \rangle_{L_0}$ is also a Σ_{BA}-morphism from $\mathbf{T}(\emptyset)$ to P1, and therefore, as $\mathbf{T}(\emptyset)$ is the initial Σ_{BA}-algebra, we get $d_0 = \delta_0 \circ \langle \_ \rangle_{L_0}$. Let now $n = i + 1$. Then $\mathcal{L}_{i+1}(\Lambda) = \mathsf{Lift}(\mathcal{L}_i(\Lambda))$ and hence by definition of Lift (cf. Def. 4.1.28) we get that $\mathcal{L}_{i+1}$ is freely generated by the set

$$G := \{\lambda(\psi_1, \dots \psi_n) \mid \psi_1, \dots, \psi_n \in \mathcal{L}_i(\Lambda) \text{ and } \lambda \in \Lambda \ n\text{-ary}\}.$$

Therefore we have to check the commutativity of the diagram only on elements of G. Let $\lambda(\psi_1, \dots, \psi_m)$ be an arbitrary element of G. Then

$$\begin{aligned}
\delta_{i+1}(\langle \lambda(\psi_1, \dots, \psi_m) \rangle_{L_{i+1}}) &= \delta_{T_i}\left(L\delta_i(\langle \lambda(\psi_1, \dots, \psi_m) \rangle_{L_{i+1}})\right) \\
&\overset{\text{Def. 4.2.10}}{=} \delta_{T_i}\left(L\delta_i(q_{i+1}([\lambda(\psi_1, \dots, \psi_m)]_{Qu\mathfrak{L}_{i+1}}))\right) \\
&\overset{\text{Def. 4.2.7}}{=} \delta_{T_i}\left(L\delta_i\left(\left[\lambda\left(q_i([\psi_1]_{Qu\mathfrak{L}_i}), \dots, q_i([\psi_m]_{Qu\mathfrak{L}_i})\right)\right]_{L_{i+1}}\right)\right) \\
&\overset{\text{Def. 4.2.10}}{=} \delta_{T_i}\left(L\delta_i\left(\left[\lambda(\langle \psi_1 \rangle_{L_i}, \dots, \langle \psi_m \rangle_{L_i})\right]_{L_{i+1}}\right)\right) \\
&\overset{\text{Lemma 4.1.38}}{=} \delta_{T_i}\left(\left[\lambda(\mathrm{Ter}(\delta_i)(\langle \psi_1 \rangle_{L_i}), \dots, \mathrm{Ter}(\delta_i)(\langle \psi_m \rangle_{L_i}))\right]_{L\mathrm{P}T_i}\right) \\
&\overset{\text{Def. of } \delta}{=} d_{T_i}\left(\lambda(\mathrm{Ter}(\delta_i)(\langle \psi_1 \rangle_{L_i}), \dots, \mathrm{Ter}(\delta_i)(\langle \psi_m \rangle_{L_i}))\right) \\
&\overset{\text{Def. 4.3.1}}{=} \lambda_{T_i}\left(\left(\mathrm{Ter}(\delta_i)(\langle \psi_1 \rangle_{L_i})\right)^{\mathrm{P}T_i}, \dots, \left(\mathrm{Ter}(\delta_i)(\langle \psi_m \rangle_{L_i})\right)^{\mathrm{P}T_i}\right) \\
&\overset{\text{Equ. 4.1, page 77}}{=} \lambda_{T_i}\left(\delta_i(\langle \psi_1 \rangle_{L_i}), \dots, \delta_i(\langle \psi_m \rangle_{L_i})\right) \\
&\overset{\text{I.H.}}{=} \lambda_{T_i}(d_i(\psi_1), \dots, d_i(\psi_m)) \\
&\overset{\text{Def. 4.2.16}}{=} d_{i+1}(\lambda(\psi_1, \dots, \psi_m))
\end{aligned}$$

QED

We are ready to prove ω-step completeness of the logic:

4.3.17. Proposition. *If $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ is injective then we have for all $n \in \omega$ and all formulas $\phi, \psi \in \mathcal{L}_n(\Lambda)$ that*

$$d_n(\phi) = d_n(\psi) \quad \textit{implies} \quad \phi \equiv_n \psi.$$

Proof. Let n be a natural number and $\phi, \psi \in \mathcal{L}_n(\Lambda)$ formulas of modal depth $\leq n$ such that $d_n(\phi) = d_n(\psi)$. By Lemma 4.3.16 we have $\delta_n(\langle \phi \rangle_{L_n}) = d_n(\phi)$ and $\delta_n(\langle \psi \rangle_{L_n}) = d_n(\psi)$. Hence $\delta_n(\langle \phi \rangle_{L_n}) = \delta_n(\langle \psi \rangle_{L_n})$. From the fact that δ is injective it follows that δ_n is injective as well (cf. Prop. 4.3.15) and therefore we get $\langle \phi \rangle_{L_n} = \langle \psi \rangle_{L_n}$. Unraveling the definition of $\langle \_ \rangle_{L_n}$ (cf. Def. 4.2.10) this equation can be rewritten as

$$q_n\left([\phi]_{Qu\mathfrak{L}_n}\right) = q_n\left([\psi]_{Qu\mathfrak{L}_n}\right).$$

Because the map $q_n : Qu\mathfrak{L}_n \to L_n$ is an isomorphism (cf. Lemma 4.2.9) we obtain $[\phi]_{Qu\mathfrak{L}_n} = [\psi]_{Qu\mathfrak{L}_n}$. This, in turn, implies that $\phi \equiv_n \psi$, because we know from Lemma 4.2.8 that $\mathfrak{L}_n = \mathcal{L}_n(\Lambda)/\equiv_n$. QED

We proved both equation (4.4) and equation (4.5) of the introduction of this subsection. As mentioned above these equations enable us to prove completeness of $L(\Lambda, \mathtt{Ax})$ with respect to the coalgebraic semantics.

4.3.18. PROPOSITION. *If δ is injective, then $L(\Lambda, \mathtt{Ax})$ is sound and complete, i.e. we have* $\mathsf{Coalg}(T) \models \phi$ *iff* $\mathtt{Ax} \vdash \phi$ *for all* $\phi \in \mathcal{L}(\Lambda)$.

Proof. The direction from right to left is the soundness of the logic and it follows, as demonstrated in Proposition 4.3.6, from the existence of δ. For the other direction let ϕ be a formula such that $\mathsf{Coalg}(T) \models \phi$ and let n be the modal depth of ϕ, i.e. $\phi \in \mathcal{L}_n(\Lambda)$. Then in particular $C_n \models \phi$ and therefore by Proposition 4.3.13 we get $d_n(\phi) = T_n = d_n(\top)$. From Prop. 4.3.17 we know that the logic is ω-step complete because δ is injective. Therefore we can derive from $d_n(\phi) = d_n(\psi)$ that $\phi \equiv_n \top$. This implies according to Lemma 4.2.5 that $\mathtt{Ax} \vdash \phi$ and the proof is finished. QED

We demonstrated that the injectivity of δ implies soundness and completeness of the logic $\mathcal{L}\Lambda, \mathtt{Ax}$. In the case that we are working with the category of Stone spaces as our base category we proved that surjectivity of δ implies the expressiveness.

Before we summarize these results for $\mathsf{C} = \mathsf{Stone}$ in a theorem we introduce the notion of when the functor $L : \mathsf{BA} \to \mathsf{BA}$ is dual to T.

4.3.19. DEFINITION. We say that *L is dual to T* if the natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ from Lemma 4.3.3 is injective and surjective.

4.3.20. REMARK. If L is dual to T, i.e. if $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ is an isomorphism then we have indeed that L is isomorphic to the functor $T^{\partial} : \mathsf{BA} \to \mathsf{BA}$ that dually corresponds to $T : \mathsf{Stone} \to \mathsf{Stone}$ (cf. Def. 3.4.7 on page 60).

4.3.21. THEOREM. *Consider $T : \mathsf{Stone} \to \mathsf{Stone}$, a set of predicate liftings Λ for T and a set of axioms $\mathtt{Ax}$. Let $L : \mathsf{BA} \to \mathsf{BA}$ be the functor given by $\mathtt{Ax}$ (Definition 4.1.35). If L is dual to T, then $L(\Lambda, \mathtt{Ax})$ is sound, complete and expressive.*

Proof. The claim on soundness and completeness is contained in Propositions 4.3.6 and 4.3.18 above. Expressivity follows from 4.3.8. QED

4.4 A characterization of duality

In the previous section we have seen that the logic given by a set of predicate liftings Λ and a set $\mathtt{Ax}$ of axioms is sound, complete and expressive if the induced functor L is dual to T. In this section, we investigate conditions under which is the case. Our main result is Theorem 4.4.13, where we give a characterisation of this duality in terms of $\mathtt{Ax}$ and Λ. More specifically, we have that L is dual to T if the axioms are order-preserving and -reflecting (cf. Definition 4.2.18) and, additionally, the predicate liftings Λ allow to distinguish all elements of TX.

We discuss both aspects, the condition on the axioms and the condition of the predicate liftings, separately. First we will show that $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ exists and is injective iff the axioms are order-preserving and -reflecting. Then we will show for the case $\mathsf{C} = \mathsf{Stone}$ that δ is surjective iff the set of predicate liftings is "separating". Putting both results together yields the announced characterisation of duality.

4.4.1 Existence of δ and injectivity

We start by showing that the canonical natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ exists and is injective iff the set of axioms `Ax` is order-preserving and -reflecting. In the following we fix a set Λ of predicate liftings for T. The main result of this section can be formulated as follows.

4.4.1. Theorem. *Given a set of axioms* `Ax` *for* $\mathcal{L}(\Lambda)$*, then* `Ax` *is order-preserving and -reflecting iff the corresponding* δ *exists and is injective.*

We need some preparations in order to be able to prove the theorem, which we split into two parts. The first (and easy) part is the following lemma:

4.4.2. Lemma. *Given a set of axioms* `Ax` *for* $\mathcal{L}(\Lambda)$*, then*

(i) If `Ax` *is order-preserving, then* δ *exists.*

(ii) If `Ax` *is order-preserving and reflecting, then* δ *exists and is injective.*

Proof.

(i) Recall from Def. 4.3.1 that for $X \in \mathsf{C}$ the function d_X is defined as $\mathsf{Lift}((\_)^{\mathrm{P}X})$, i.e. the lifting of the term interpretation function

$$(\_)^{\mathrm{P}X} : \mathrm{Ter}(\mathrm{P}X) \to \mathrm{P}X.$$

According to our assumption $\mathsf{Lift}((\_)^{\mathrm{P}X})$ is preserving as $(\_)^{\mathrm{P}X}$ is preserving and therefore δ_X can be defined as described in 4.3.1. As this works for all $X \in \mathsf{C}$ we get a natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$.

(ii) Suppose now that `Ax` is order-preserving and -reflecting. Then as in the first case we can show that δ_X is a BA-morphism. As $d_X = \mathsf{Lift}((\_)^{\mathrm{P}X})$ is now also order-reflecting we obtain, using Lemma 4.1.21, that δ_X is injective.

QED

To prove the second half of the theorem we take a closer look at the definition of the lifting $\mathsf{Lift}(f)$ of some $f : (\mathbb{A}, \equiv) \to \mathrm{P}X$. First we need some notation.

4.4.3. Notation. *Let $\mathbb{A}$ be in $\mathsf{Alg}(\Sigma_{\mathsf{BA}})$, X an object in C and let $h : \mathbb{A} \to \mathrm{P}X$ be a Σ_{BA}-morphism. Then we denote by $h^* : \mathbb{A} \to \mathrm{Ter}(\mathrm{P}X)$ the Σ_{BA}-morphism mapping $a \in \mathbb{A}$ to $h(a) \in \mathrm{Ter}(\mathrm{P}X)$, i.e. we regard $h(a)$ as a term over $\mathrm{P}X$ (cf. Definition 4.1.10).*

The following observation is obvious.

4.4.4. Lemma. *Let $(\mathbb{A}, \equiv)$ be a pre-Boolean algebra, $X \in \mathsf{C}$ and let $h : \mathbb{A} \to \mathrm{P}X$ be a Σ_{BA}-morphism. Then $h : (\mathbb{A}, \equiv) \to \mathrm{P}X$ is a preserving map iff $h^* : (\mathbb{A}, \equiv) \to (\mathrm{Ter}(\mathrm{P}X), \mathrm{Diag}(\mathrm{P}X))$ is a preserving map and h is reflecting iff h^* is reflecting.*

4.4.5. Lemma. *Let $(\mathbb{A}, \equiv)$ be in PBA, X an object in C and let $h : (\mathbb{A}, \equiv) \to \mathrm{P}X$ be a preserving map. Then $\mathsf{Lift}(h) = d_X \circ \mathfrak{L}(h^*)$, where $\mathfrak{L}(h^*)$ is defined as in Definition 4.1.32.*

Proof. Let $\lambda(\psi_1, \ldots, \psi_m) \in \mathsf{Lift}(\mathbb{A})$. We calculate that

$$\begin{aligned}
\mathsf{Lift}(h)(\lambda(\psi_1, \ldots, \psi_m)) &= \lambda_X(h(\psi_1), \ldots, h(\psi_m)) \\
&= \lambda_X((h^*(\psi_1))^{\mathrm{P}X}, \ldots, (h^*(\psi_m))^{\mathrm{P}X}) \\
&= d_X\,(\lambda(h^*(\psi_1), \ldots, h^*(\psi_m))) \\
&= d_X\,(\mathfrak{L}(h^*)(\lambda(\psi_1, \ldots, \psi_m)))
\end{aligned}$$

QED

The second half of Theorem 4.4.1 is now an immediate consequence.

4.4.6. Lemma. *Let $h : \mathbb{A} \to \mathrm{P}X$ be a Σ_{BA}-morphism and let $(\mathbb{A}, \equiv)$ be a pre-Boolean algebra. Then*

1. *If δ exists and $h : (\mathbb{A}, \equiv) \to \mathrm{P}X$ is a preserving map , then*

$$\mathsf{Lift}(h) : (\mathsf{Lift}(\mathbb{A}), \mathsf{Lift}(\equiv)) \to \mathrm{P}TX$$

   is preserving.

2. *If in addition δ is injective and h is reflecting, then $\mathsf{Lift}(h)$ is also reflecting.*

Proof. Suppose δ exists and $h : (\mathbb{A}, \equiv) \to \mathrm{P}X$ is a preserving map. Then according to Lemma 4.4.5 we have $\mathsf{Lift}(h) = d_X \circ \mathfrak{L}(h^*)$. Because δ_X exists the function d_X must be a preserving map. Furthermore h^* is preserving and hence also $\mathfrak{L}(h^*)$ (cf. Lemma 4.1.33). Thus we can write $\mathsf{Lift}(h)$ as the composition of preserving maps and therefore $\mathsf{Lift}(h)$ itself is a preserving map. The proof of the second half of the lemma is completely analogous. QED

The proof of the Theorem 4.4.1 is now complete: Lemma 4.4.2 proves one direction and Lemma 4.4.6 the other direction.

4.4.2 Surjectivity

We now consider a logic for an endofunctor $T : \mathsf{Stone} \to \mathsf{Stone}$, i.e. $\mathsf{C} = \mathsf{Stone}$ and $\mathsf{P} = \mathbb{C}\mathrm{lp}$. Here we will see that requiring that the set of predicate liftings Λ for T is "separating" (a notion taken from [Pat04]) is equivalent to the fact that the canonical map δ is surjective.

4.4.7. Definition. Let $\mathbb{X} \in \mathsf{Stone}$.

1. A collection of clopens $C \subseteq \mathrm{Clp}(\mathbb{X})$ is called *separating* if the map
$$\begin{array}{rcl} s_C : \mathbb{X} & \to & \mathcal{P}(\mathrm{Clp}(\mathbb{X})) \\ x & \mapsto & \{U \in C \mid x \in U\} \end{array}$$
is injective.

2. A set of predicate liftings Λ for T is called *separating* if for all $\mathbb{X} \in \mathsf{Stone}$
$$\mathsf{Im}_\Lambda(\mathbb{X}) := \{\lambda_{\mathbb{X}}(U_1, \ldots, U_n) \mid \lambda \in \Lambda, U_1, \ldots, U_n \in \mathrm{Clp}(\mathbb{X})\}$$
is a separating set of clopens of $T\mathbb{X}$.

Intuitively a separating set of predicate liftings makes it possible to characterise points in $T\mathbb{X}$("successors") by lifted predicates over $\mathbb{X}$.

As it was shown in [Pat04] a coalgebraic modal language which has a separating set of predicate liftings is expressive. We will now see that provided we have a sound and complete logic for the functor T the fact that Λ is separating is equivalent to saying that the functor L defining the algebraic semantics of our logic is the dual of T.

Our main theorem states that δ is surjective if and only if the set Λ of predicate liftings is separating. Before we state (and prove) the theorem, we collect some facts on separating sets, which are necessary for the proof of the theorem.

4.4.8. Lemma. *Let $\mathbb{X} = (X, \tau) \in \mathsf{Stone}$ and let $\mathbb{A} \subseteq \mathrm{Clp}(\mathbb{X})$ be a subalgebra of $\mathrm{Clp}(\mathbb{X})$. Then $s_{\mathbb{A}}$ is injective iff $\mathbb{A} = \mathrm{Clp}(\mathbb{X})$.*

Proof. The implication from right to left is immediate. To prove the other direction suppose that $s_{\mathbb{A}}$ is injective. Then one can easily see that
$$\bigcap\{U \in \mathbb{A} \mid x \in U\} = \{x\} \tag{4.6}$$
for all $x \in X$. To prove $\mathbb{A} = \mathrm{Clp}(\mathbb{X})$ it suffices to show that $\mathbb{A}$ is a basis for the topology on X. Suppose that $W \subseteq X$ is open and let $x \in W$. We have to show that there is a clopen set $U \in \mathbb{A}$ such that $x \in U \subseteq W$. Because of (4.6) we know that for all $y \in -W$ there is some $U_y \in \mathbb{A}$ such that $x \notin U_y$ and $y \in U_y$. Hence $-W \subseteq \bigcup_{y \in -W} U_y$. Because of compactness of the topology there are $y_1, \ldots, y_n \in -W$ such that $-W \subseteq \bigcup_{i=1}^n U_{y_i}$. Define $V := -(\bigcup_{i=1}^n U_{y_i})$. Then $V \in \mathbb{A}$ and $x \in V \subseteq W$. Therefore $\mathbb{A}$ is a basis of the topology on X. QED

4.4.9. Lemma. *Let $\mathbb{X} = (X, \tau) \in \mathsf{Stone}$ and let $C \subseteq \mathrm{Clp}(\mathbb{X})$ be a clopen subbasis of the topology of $\mathbb{X}$. Then C is a separating set of clopens.*

Proof. Let $x, y \in X$ and $x \neq y$. Then there is a $U \in \mathrm{Clp}(\mathbb{X})$ such that $x \in U$ and $y \in -U$. As C is a subbasis of the topology there are $V_1, \ldots, V_n \in C$ such that

$$x \in \bigcap_{i=1}^{n} V_i \subseteq U.$$

But this means that there is at least one $V_j \in C$ such that $x \in V_j$ and $y \notin V_j$, and therefore $s_C(x) \neq s_C(y)$. QED

4.4.10. Lemma. *Let $\mathbb{X} = (X, \tau)$ be a Stone space, $C \subseteq \mathrm{Clp}(\mathbb{X})$ a subset of $\mathrm{Clp}(\mathbb{X})$ and define $-C := \{-U \mid U \in C\}$. Then*

$$s_C \textit{ injective} \;\Leftrightarrow\; s_{C\cup -C} \textit{ injective}$$

Proof. The direction from left to right is obvious. For the other direction, suppose $s_{C\cup -C}$ is injective and let $x, y \in X$, $x \neq y$. Then according to our assumption

$$s_{C\cup -C}(x) \neq s_{C\cup -C}(y).$$

Therefore we can assume that there is $V \in C \cup -C$ such that $x \in V$ and $y \in -V$. We distinguish the following cases:

Case $V \in C$. Then clearly $s_C(x) \neq s_C(y)$.

Case $V \in -C$. Then $-V \in C$ and hence $-V \in s_C(y)$ and $-V \notin s_C(x)$.

Since $V \in C \cup -C$, this finishes the proof. QED

Now we are ready to prove the main result of this section:

4.4.11. Theorem. *Let $T : \mathsf{Stone} \to \mathsf{Stone}$ be a functor, and suppose that $\mathcal{L}(\Lambda)$ is a logic for T that has a order-preserving set of axioms* `Ax`*. Then Λ is a separating set of predicate liftings iff the canonical $\delta : L\mathbb{C}\mathrm{lp} \Rightarrow \mathbb{C}\mathrm{lp}T$ is surjective.*

Proof. Given an order-preserving set of axioms we know that the map $d_{\mathbb{X}} := \mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}}) : \mathsf{Lift}(\mathrm{Ter}(\mathbb{C}\mathrm{lp}\mathbb{X})) \to \mathbb{C}\mathrm{lp}T\mathbb{X}$ factors through $\delta_{\mathbb{X}} : L\mathbb{C}\mathrm{lp}\mathbb{X} \to \mathbb{C}\mathrm{lp}T\mathbb{X} \in \mathsf{BA}$, because we proved in Theorem 4.4.1 that the fact that the set of axioms `Ax` is order-preserving implies that δ exists, i.e. that $d_{\mathbb{X}}$ factors through $\delta_{\mathbb{X}}$ for all $\mathbb{X} \in \mathsf{Stone}$. It is therefore obvious that we have the following equivalence:

$$\delta_{\mathbb{X}} \text{ is surjective for all } \mathbb{X} \in \mathsf{Stone} \quad \text{iff} \quad \mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}}) \text{ is surjective for all } \mathbb{X} \in \mathsf{Stone}.$$

We now show that the last property is equivalent to the fact that Λ is a separating set of liftings.

Suppose first that for an arbitrary $\mathbb{X} \in \mathsf{Stone}$ the map $\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})$ is surjective. As the domain of $\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})$ is closed under the boolean operations it can be easily seen that the image of $\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})$ is equal to $\langle \mathsf{Im}_\Lambda(\mathbb{X})\rangle_{\mathbb{C}\mathrm{lp}T\mathbb{X}}$, the subalgebra of $\mathbb{C}\mathrm{lp}T\mathbb{X}$ generated by $\mathsf{Im}_\Lambda(\mathbb{X})$. Hence we get

$$\mathbb{C}\mathrm{lp}T\mathbb{X} = \mathsf{Im}(\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})) = \langle \mathsf{Im}_\Lambda(\mathbb{X})\rangle_{\mathbb{C}\mathrm{lp}T\mathbb{X}}\,.$$

This implies that $\mathsf{Im}_\Lambda(\mathbb{X}) \cup -\mathsf{Im}_\Lambda(\mathbb{X})$ is a clopen subbasis of the topology of $T\mathbb{X}$, where again $-\mathsf{Im}_\Lambda(\mathbb{X}) := \{-U \mid U \in \mathsf{Im}_\Lambda(\mathbb{X})\}$. Using Lemma 4.4.9 and 4.4.10 we obtain that the map $s_{\mathsf{Im}_\Lambda(\mathbb{X})}$ is injective. As $\mathbb{X}$ was arbitrary we can conclude that Λ is separating. Now suppose that Λ is a separating set of liftings and s let $\mathbb{X} \in \mathsf{Stone}$. Then $s_{\mathsf{Im}_\lambda(\mathbb{X})}$ is injective which also implies the injectivity of $s_{\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})}$. As the image of $\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})$ is a subalgebra of $\mathbb{C}\mathrm{lp}T\mathbb{X}$ it follows by Lemma 4.4.8 that $\mathsf{Im}(\mathsf{Lift}((\_)^{\mathbb{C}\mathrm{lp}\mathbb{X}})) = \mathbb{C}\mathrm{lp}T\mathbb{X}$. QED

We note the following immediate consequence, which is the main result of this section:

4.4.12. Corollary. *Let $T : \mathsf{Stone} \to \mathsf{Stone}$ be a functor, and suppose that $\mathcal{L}(\Lambda)$ is a logic for T that has a order-preserving set of axioms* $\mathtt{Ax}$. *Then $\mathcal{L}(\Lambda)$ is expressive iff Λ is separating.*

Proof. Follows directly from the theorem and Proposition 4.3.8 QED

Summing up, we can now characterise duality between T and L in logical terms as follows:

4.4.13. Theorem. *Let $T : \mathsf{Stone} \to \mathsf{Stone}$, Λ a set of predicate liftings for T and let* $\mathtt{Ax}$ *be a set of axioms. The following are equivalent:*

(i) $\mathtt{Ax}$ *is order-preserving and reflecting, and Λ is separating*

(ii) L is dual to T.

Proof. From Theorem 4.4.1 we know that $\mathtt{Ax}$ is order-preserving and reflecting iff $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ exists and is injective. Theorem 4.4.11 tells us that, under the proviso that Ax is order-preserving, the language $\mathcal{L}(\Lambda)$ is expressive iff $\delta : L\mathrm{P} \Rightarrow \mathrm{P}T$ is surjective. Putting both statements together we arrive at the claim of the theorem. QED

Combining the above result with Theorem 4.3.21, both of the two equivalent conditions above provide us with a sound, complete and expressive logic for T-coalgebras.

4.5 Conclusion

General context

Our work stands in the broader context of employing Stone duality for providing a state-based semantics for logical calculi that are represented by algebras. Here we understand Stone duality in a more abstract sense, like e.g. in [Joh82]. In this way not only the already mentioned work by Jónsson and Tarski on Boolean algebras with operators ([JT51, JT52]) and the work by Goldblatt on descriptive general frames ([Gol76]) fit into this framework, but e.g. also the work by Abramsky on domain theory in logical form ([Abr91]).

The general pattern of the above listed approaches is as follows: one starts with a duality between a category A of algebras on the one hand, and a category of topological spaces X on the other hand. This duality is then extended, in the first example, to the duality between MA and DGF. In the work by Abramsky the duality is extended by applying dual constructions to the category $\mathsf{X} = \mathbf{SFP}$ of so-called **SFP**-domains and to the dual category A of the corresponding locales. In our setting we have as a basic duality the duality between $\mathsf{A} = \mathsf{BA}$ and $\mathsf{X} = \mathsf{Stone}$. In this chapter we showed how to lift this duality to a duality between functors $T : \mathsf{Stone} \to \mathsf{Stone}$ and $L : \mathsf{BA} \to \mathsf{BA}$, provided that we have an order-preserving and -reflecting set of axioms and a separating set of predicate liftings for T.

Our results

The presented work explains and generalizes earlier results by Pattinson on coalgebraic modal logic from [Pat03a] and [Pat04] (cf. introduction to Section 4.3). We defined an algebraic semantics of coalgebraic modal logic in terms of a category of algebras for a functor $L : \mathsf{BA} \to \mathsf{BA}$. Furthermore we connected the algebraic and the coalgebraic semantics via a natural transformation $\delta : L\mathrm{P} \Rightarrow \mathrm{P}L$ and showed that certain properties of δ correspond to properties of the logic. In the case $\mathsf{C} = \mathsf{Stone}$ we obtained what we called a logical characterization of duality: L is dual to T iff $L(\Lambda, \mathtt{Ax})$ is sound and expressive and Λ is expressive. Our results are based on the duality between Stone spaces and Boolean algebras, but our categorical formulation allows for further straightforward generalizations to other dualities, for example to the duality between the category of partially ordered sets and the category of spectral spaces.

Presenting functors

An interesting question concerning coalgebraic modal logics is, whether we can find for any functor $T : \mathsf{C} \to \mathsf{C}$ a set of predicate liftings Λ and a set of axioms $\mathtt{Ax}$ such that the language $\mathcal{L}(\Lambda)$ is expressive and the logic $L(\Lambda, \mathtt{Ax})$ is sound and complete with respect to the coalgebraic semantics. For $\mathsf{C} = \mathsf{Stone}$ the solution could be as follows: Given a functor $T : \mathsf{Stone} \to \mathsf{Stone}$ we can always look at the dual functor $T^{\partial} : \mathsf{BA} \to \mathsf{BA}$ of T. Given our results the question whether we can find an adequate logic for reasoning about T-coalgebras can be reformulated into the question whether we can represent the functor $T^\partial : \mathsf{BA} \to \mathsf{BA}$ with operations and equations, i.e. whether we can find a set of predicate liftings Λ and a set of axioms $\mathtt{Ax}$ such that the corresponding functor $L : \mathsf{BA} \to \mathsf{BA}$ is isomorphic to T^∂. This type of question is not new. A similar result has been proven already for Set-functors: Up to questions of size, any set-functor can be presented by operations and equations, see [Rei83, 1.5], [Ros81], and [AT90, Section III.3.2,III.4.3].

Completeness via Jónsson-Tarski

As mentioned in the introduction, our completeness criterion in terms of δ corresponds to the fact that we can lift the Stone representation embedding to the level of lifted predicates. The completeness proof, however, uses an argument involving the final sequence of the functor. From modal logic we know a much more direct completeness argument which uses the so-called Jónsson-Tarski theorem: for every modal algebra $\mathfrak{A} = (\mathbb{A}, f)$ we can define its canonical extension $\mathfrak{A}^* = (U_{\mathsf{Stone}}\mathbb{S}\mathrm{p}\mathbb{A}, f^*)$ such that the Stone representation map $j_{\mathbb{A}} : \mathbb{A} \to U_{\mathsf{Stone}}\mathbb{S}\mathrm{p}\mathbb{A}$ is a homomorphism between modal algebras. In [KKP05] we showed that this Jónsson-Tarski argument can be generalized to our coalgebraic setting if we make an additional assumption on the functor.

Chapter 5

Closure properties of coalgebra automata

There is a close connection between automata theory and the theory of coalgebras as has been pointed out in a number of papers, starting with [Rut98a]. It has to be stressed, however, that there is a fundamental difference between this work and the work on coalgebra automata which has been initiated by Venema in [Ven04]. Whereas the former work is concerned with the question of how to model automata as coalgebras the latter follows the slogan *automata are formulas* of some logic, namely formulas of coalgebraic fixed-point logic (cf. Section 2.3) in our case.

Like formulas of coalgebraic logic, which are either satisfied or refuted at some state in some given coalgebra, coalgebra automata accept or reject *rooted coalgebras*, i.e. coalgebras together with some designated state. Because coalgebras can be seen as abstract infinite objects coalgebra automata correspond to finite automata on infinite objects.

Such automata have already found important applications in areas of computer science where one investigates the ongoing behavior of nonterminating programs such as operating systems. As an example we mention the automata-based verification method of *model checking* [CGP00]. Work on automata on infinite objects also has a long and strong theoretical tradition and its results link the field to neighboring areas such as logic and game theory, see [GTW02] for an overview. We mention Rabin's decidability theorem [Rab69] for the monadic second order logic of trees as an outstanding example for a theoretical work in this area; to mention a more recent example, Janin & Walukiewicz [JW95] identified the modal μ-calculus as the bisimulation invariant fragment of the monadic second order logic of labeled transition systems.

An interesting phenomenon in work about automata on infinite objects has been that most key results hold for automata on infinite objects of different types alike, such as automata on words, trees or graphs. This naturally raises the question, whether these results can perhaps be formulated at the more general level of abstraction of coalgebra automata. In this chapter, which is in large parts based on the paper [KV05], we are going to answer this question in the positive by proving certain closure properties of coalgebra automata. In this way we obtain uniform proofs of existing results from automata on infinite words, trees and graphs. Furthermore we use the connection with coalgebraic fixed-point logic to obtain several logical corollaries of our automata-theoretic results.

This chapter is structured as follows: We will recall the definition of a T-(coalgebra) automaton for a standard, weak pullback preserving functor $T : \mathsf{Set} \to \mathsf{Set}$. Then we prove certain closure properties of T-automata and our main result, namely that for every alternating T-automaton we can find an equivalent non-deterministic one. Furthermore we show that the non-emptiness problem of T-automata is decidable provided that T maps finite sets to finite sets.

We then use the "automata are formulas" slogan to obtain corollaries of the above listed results for coalgebraic fixed-point logics: we show that coalgebraic fixed-point logics enjoy (a weak version of) the finite model property and, based on ideas from the proof that coalgebra automata are closed under alternation, we prove the soundness of a distributive law for the ∇-operator. In order to understand the relevance of this last result one should note that there is so far neither an axiomatisation for Moss's coalgebraic logic nor a conjecture of what such an axiomatisation could look like. Our distributive law can be added as a sound logical principle of coalgebraic logic to the list in [Mos99, Sec.6] and might help to find an axiomatisation either of that logic or of its finitary version.

The chapter is based on the earlier published paper [KV05] which is joint work with Yde Venema.

5.1 Coalgebra automata

Before we recall the definition of a coalgebra automaton from [Ven04] we first want to provide some intuition for how coalgebra automata naturally generalise automata running on infinite words, trees and graphs. This will be done by taking a closer look at the definition of an infinite graph automaton. All the automata in this chapter will be so-called *parity automata*. The acceptance condition of these automata is formulated in terms of *parity games*. The terminology that we are going to use, when talking about parity games, is listed in Appendix C.

5.1.1 Deterministic graph automata

5.1.1. Definition. Let C be a finite set. A *rooted graph* is a tuple (S, σ, s_I), where S is a set, $\sigma : S \to \mathcal{P}(S)$ is the successor function and $s_I \in S$ is the root. A *C-labeled rooted graph* is a tuple (S, σ, γ, s_I) such that (S, σ, s_I) is a rooted graph and $\gamma : S \to C$ is a (coloring) function assigning to each $s \in S$ its color $\gamma(s) \in C$.

The following definition of a deterministic graph automaton and its acceptance condition is essentially the same as the definition of a μ-automaton from [JW95], with the difference that we are only considering *deterministic* graph automata in this section.

5.1.2. Definition. A deterministic graph automaton is a tuple $\mathbb{A} = (A, a_I, C, \Delta, \Omega)$ where A is a finite set of states, $a_I \in A$ is the initial state, C is a finite set (the alphabet), $\Delta : C \times A \to \mathcal{P}A$ is the transition function and $\Omega : A \to \omega$ is a parity function, i.e. a function from A to the set of natural numbers that has finite range (cf. Definition C.0.9).

In order to be able to define when a graph automaton accepts a given graph we need the notion of a run of the automaton.

5.1.3. Definition. Let $\mathbb{A} := (A, a_I, C, \Delta, \Omega)$ be a deterministic graph automaton and $(\mathbb{S}, s_I) := (S, \sigma, \gamma, s_I)$ a C-labeled graph with labeling function $\gamma : S \to C$. A *run* of $\mathbb{A}$ on $(\mathbb{S}, s_I)$ is a rooted graph $(Y \subseteq S \times A, \rho : Y \to \mathcal{P}(Y), (s_I, a_I))$ such that for all $(s, a) \in Y$

- for all $a' \in \Delta(\gamma(s), a)$ there is an $s' \in \sigma(s)$ such that $(s', a') \in \rho(s, a)$
- for all $s' \in \sigma(s)$ there is an $a' \in \Delta(\gamma(s), a)$ such that $(s', a') \in \rho(s, a)$.

A run is called *accepting* if for all $(s, a) \in Y$ we have $\Delta(\gamma(s), a) \neq \emptyset$ and for all infinite sequences $\alpha = (s_0, a_0)(s_1, a_1)(s_2, a_2)\dots$ with $(s_0, a_0) = (s_I, a_I)$ and $(s_{i+1}, a_{i+1}) \in \rho(s_i, a_i)$ we have

$$\max\{\Omega(a) \mid (s, a) \in \mathrm{Inf}(\alpha)\} \quad \text{is even.},$$

where $\mathrm{Inf}(\alpha)$ is the set of states that occur infinitely often in α (cf. Def C.0.8). The automaton $\mathbb{A}$ *accepts* $(\mathbb{S}, s_I)$ if there is an accepting run.

5.1.4. Remark. The parity condition can be understood as follows: The acceptance condition should specify which runs of the automaton are accepting, i.e. an acceptance condition is a way of encoding subsets of all infinite runs of the automaton. The parity condition is particularly well-behaved as we will see when we move to the definition of an acceptance game: parity games are history-free determined.

The example of a deterministic graph automaton illustrates that an important part of constructing an accepting run of an automaton consists of constructing a bisimulation. This is the content of the next proposition.

5.1.5. Proposition. *Let C be equal to $\mathbf{1}$ (the "one-letter" alphabet), $\mathbb{A} = (A, a_I, C, \Delta, \Omega)$ a deterministic graph automaton and $(\mathbb{S}, s_I)$ a rooted graph. Then for $Y \subseteq S \times A$ we have*

Y is a bisimulation between (S, σ, s_I) and (A, Δ, a_I)

iff

there is $\rho : Y \to \mathcal{P}Y$ s.t. $(Y, \rho, (s_I, a_I))$ is a run of $\mathbb{A}$ on (S, σ, s_I).

where bisimulation refers to the usual notion of bisimulation between directed graphs or transition systems (cf. Example A.3.9).

Proof. Suppose first that there is a function $\rho : Y \to \mathcal{P}Y$ such that $(Y, \rho, (s_I, a_I))$ is a run of $\mathbb{A}$ on $(\mathbb{S}, s_I)$. Then spelling out the definition of a run we immediately get that Y is a bisimulation between (S, σ, s_I) and (A, Δ, a_I).

For the other direction let $Y \subseteq S \times A$ be a bisimulation between (S, σ, s_I) and (A, Δ, a_I). Then by definition $(s_I, a_I) \in Y$. Furthermore it is a well-known fact that bisimulations between directed graphs are exactly the $\mathcal{P}$-bisimulations between $\mathcal{P}$-coalgebras. Therefore there is a function $\rho : Y \to \mathcal{P}Y$ such that the following diagram commutes

$$\begin{array}{ccccc} S & \xleftarrow{\pi_A} & Y & \xrightarrow{\pi_S} & A \\ \downarrow{\sigma} & & \downarrow{\rho} & & \downarrow{\Delta} \\ \mathcal{P}S & \xleftarrow{\mathcal{P}\pi_A} & \mathcal{P}Y & \xrightarrow{\mathcal{P}\pi_S} & \mathcal{P}A \end{array}$$

It is easy to check that $(Y, \rho, (s_I, a_I))$ is a run of $\mathbb{A}$ on (S, σ, s_I). QED

5.1.6. Remark. We restrict our attention to the case that $C = \mathbf{1}$ in order to be able to underline the central role that bisimulations play in the acceptance condition of a graph automaton. If C contained more than one element, the connection between bisimulations and runs of the automaton would become unnecessarily more complicated.

Therefore the acceptance game for a deterministic graph automaton can be formulated using Baltag's bisimulation game from [Bal00], which we will present next.

5.1.2 The bisimulation game

We state the definition of the bisimulation game in full generality, i.e. not only for $\mathcal{P}$-bisimulations but for T-bisimulations for an arbitrary functor $T : \mathsf{Set} \to \mathsf{Set}$.

5.1.7. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ and $\mathbb{X} := (X, \gamma), \mathbb{Y} := (Y, \delta) \in \mathsf{Coalg}(T)$. Then the arena of the *T-bisimulation game* $\mathcal{G}(\mathbb{X}, \mathbb{Y})$ is given by the following table

Position: b	Player	Admissible moves: $E[b]$	$\Omega(b)$
$(x, y) \in X \times Y$	$\exists$	$\{Z \subseteq X \times Y \mid (\gamma(x), \delta(y)) \in \overline{T}Z\}$	0
$Z \in \mathcal{P}(X \times Y)$	$\forall$	Z	0

where the second column indicates whether a given position b belongs to player $\exists$ or $\forall$, i.e. whether $b \in B_\exists$ or $b \in B_\forall$ (cf. Def. C.0.9) and $\overline{T}Z$ is the relation lifting of Z (cf. Def. A.2.5).

5.1.8. Remark. Note that the parity function for the bisimulation game is just the constant function that assigns to every position the parity 0. This means that *all* infinite games are won by $\exists$.

Intuitively $\exists$ wants to show that two points are related by a bisimulation and $\forall$ tries to disprove $\exists$'s claim. This intuition is made explicit in the following proposition.

5.1.9. Proposition. *[Bal00] Let T : Set → Set be a weak pullback preserving functor and $(\mathbb{X}, x_I) := (X, \gamma, x_I)$, $(\mathbb{Y}, y_I) := (Y, \delta, y_I)$ rooted T-coalgebras. Then*

$$(\mathbb{X}, x_I) \leftrightarrows_T (\mathbb{Y}, y_I) \quad \text{iff} \quad \exists \text{ has a winning strategy in } \mathcal{G}(\mathbb{X}, \mathbb{Y}) \text{ from } (x_I, y_I).$$

Proof. We only sketch the proof. For the direction from left to right fix a T-bisimulation Z such that $(x_I, y_I) \in Z$. Then it is not difficult to check that the strategy of $\exists$ to move from any position $(x, y) \in X \times Y$ to Z is winning in $\mathcal{G}(\mathbb{X}, \mathbb{Y})$ from position (x_I, y_I). The other direction follows from the observation that the set $\mathrm{Win}_\exists(\mathcal{G}(\mathbb{X}, \mathbb{Y}))$ of winning positions of $\exists$ is a T-bisimulation. QED

The bisimulation game can be now used to reformulate the acceptance condition for deterministic graph automata in terms of a parity game. The goal of player $\exists$ will be to show that there is an accepting run of the automaton on a given rooted graph. In case we forget about the alphabet (we consider the trivial alphabet $C = \mathbf{1}$) we saw in Proposition 5.1.5 that this means that $\exists$ has to make sure that there is a bisimulation between the automaton and the rooted graph and this bisimulation fulfills additional properties specified by the parity function of the automaton. This leads us to the following reformulation of the acceptance condition of a graph automaton.

5.1.10. Definition. Let $\mathbb{A} = (A, a_I, C, \Delta, \Omega)$ be a deterministic graph automaton and $(\mathbb{S}, s_I) := (S, \sigma, \gamma, s_I)$ a rooted C-labeled graph with coloring function $\gamma : S \to C$. Then the *acceptance game* $\mathcal{G}(\mathbb{S}, \mathbb{A})$ is defined as the parity graph game (cf. Definition C.0.9) given by the following table

Position: b	Player	Admissible moves: $E[b]$	$\Omega'(b)$
$(s, a) \in S \times A$	$\exists$	$\{Y \subseteq S \times A \mid (\sigma(s), \Delta(\gamma(s), a)) \in \overline{\mathcal{P}}Y\}$	$\Omega(a)$
$Y \in \mathcal{P}(S \times A)$	$\forall$	Y	0

where $\Omega' : (S \times A) \cup \mathcal{P}(S \times A) \to \omega$ is the parity function of the acceptance game. We say that $\mathbb{A}$ *accepts* $(\mathbb{S}, s_I)$ if $\exists$ has a winning strategy in $\mathcal{G}(\mathbb{S}, \mathbb{A})$ starting from position (s_I, a_I).

It is not difficult to see that there is an accepting run of a given graph automaton $\mathbb{A}$ on a rooted graph $(\mathbb{S}, s_I)$ (cf. Def.5.1.3) iff $\mathbb{A}$ accepts $(\mathbb{S}, s_I)$ according to Definition 5.1.10.

The definition of the acceptance game of a deterministic graph automaton can now be generalised to automata working on T-coalgebras for an arbitrary standard weak pullback preserving functor T : Set → Set.

5.1.3 Coalgebra automata

Compared to the deterministic graph automata that we discussed in the last section T-coalgebra automata are a generalization in two directions: the first generalization is that T-coalgebra automata are operating on rooted T-coalgebras for some weak pullback preserving functor $T : \mathsf{Set} \to \mathsf{Set}$. The second generalization is a move from deterministic to alternating automata.

In the previous section we restricted our attention to deterministic automata because we wanted to focus on the important rôle of the bisimulation game in the acceptance game of an automaton. In a play of the acceptance game of a deterministic graph automaton $\exists$ has to ensure that for every position (s, a) the set of successors $\sigma(s)$ of s fulfill the conditions encoded in the set of successors $\Delta(a)$ of a. In a *non-deterministic* graph automaton, $\Delta(a)$ consists of several successor sets $\Phi_1, \Phi_2 \ldots \in \mathcal{P}(A)$ and $\exists$ has to show that *there is* a $\Phi \in \Delta(a)$ such that $\sigma(s)$ fulfills the conditions encoded in Φ. In an *alternating* graph automaton, $\Delta(a)$ consists of several sets of successor sets of a, i.e. $\Delta(a) = \{\Psi_1, \Psi_n, \ldots\}$ where $\Psi_i = \{\Phi^i_1, \Phi^i_2, \ldots\} \in \mathcal{P}(\mathcal{P}(A))$. It is now the task of $\exists$ to show that *there is* a $\Psi \in \Delta(a)$ such that *for all* $\Phi \in \Psi$ the set $\sigma(s)$ fulfills the conditions encoded in Φ. To sum it up: the move from deterministic to non-deterministic automata adds an existential quantifier and the move to alternating automata adds a universal quantifier to the acceptance condition.

Now we are prepared for the definition of a T-coalgebra automaton as it was introduced by Venema in [Ven04].

5.1.11. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a standard weak pullback preserving functor (cf. Appendix A.2.2). An (alternating) *T-automaton* is a quadruple $\mathbb{A} = (A, a_I, \Delta, \Omega)$ with A some finite set (of states), $a_I \in A$ the *root*, $\Delta : A \to \mathcal{P}(\mathcal{P}(TA))$ the *transition function* and $\Omega : A \to \omega$ a parity map. A T-automaton is called *non-deterministic* if all members of each $\Delta(a)$ are singleton sets. A T-automaton is called *deterministic* if $\Delta(a)$ has exactly one singleton set as its only element.

The acceptance condition for T-automata is formulated in terms of a parity game (cf. Definition C.0.9)

5.1.12. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a standard weak pullback preserving functor and $\mathbb{A} = (A, a_I, \Delta, \Omega)$ a T-automaton. Furthermore let $(\mathbb{S}, s_I) = (S, \sigma, s_I)$ be a rooted T-coalgebra. Then the acceptance game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ is given by the following table

Position: b	Player	Admissible moves: $E[b]$	$\Omega'(b)$
$(s, a) \in S \times A$	$\exists$	$\{(s, \Phi) \in S \times \mathcal{P}TA \mid \Phi \in \Delta(a)\}$	$\Omega(a)$
$(s, \Phi) \in S \times \mathcal{P}TA$	$\forall$	$\{(s, \phi) \in S \times TA \mid \phi \in \Phi\}$	0
$(s, \phi) \in S \times TA$	$\exists$	$\{Z \in \mathcal{P}(S \times A) \mid (\sigma(s), \phi) \in \overline{T}Z\}$	0
$Z \in \mathcal{P}S \times A$	$\forall$	Z	0

We say $\mathbb{A}$ accepts $(\mathbb{S}, s_I)$ if $\exists$ has a winning strategy from position (s_I, a_I) in $\mathcal{G}(\mathbb{S}, \mathbb{A})$. Positions of the form $(s, a) \in S \times A$ will be called *basic positions* of the game. A partial play of the game of the form

$$(s, a)\ (s, \Phi)\ (s, \phi)\ Z$$

with $(s, a) \in S \times A$, $(s, \Phi) \in S \times \mathcal{P}TA$, $(s, \phi) \in S \times TA$ and $Z \in \mathcal{P}(S \times A)$ will be called a *round* of the play.

A class of rooted T-coalgebras will be called *T-language* . A T-language L is *recognized* by some T-automaton $\mathbb{A}$ if a rooted T-coalgebra belongs to T iff it is accepted by $\mathbb{A}$. The T-language recognized by the T-automaton $\mathbb{A}$ will be denoted by L($\mathbb{A}$). A T-language is called *(non-deterministically) recognizable* if it is recognized by some (non-deterministic) T-automaton.

5.1.13. REMARK. It is clear from the definition of Ω' that only the *basic* positions of a play, i.e., positions of the form $(s, a) \in S \times A$, are relevant to determine the winner. Accordingly, in the sequel we will frequently represent a play of the game by the sequence of basic positions visited during the play.

The definition of the acceptance game of a T-automaton reflects that T-automata are generalising deterministic graph automata in two ways. The first two moves of the acceptance game are the part which reflects the generalization from deterministic to alternating automata: in a position (s, a) the two players are using their power to determine the "successor" of a. The second half of the game consists of the bisimulation game formulated for arbitrary functors $T : \mathsf{Set} \to \mathsf{Set}$ (cf. Definition 5.1.7). The fact that an essential part of the acceptance game of a T-automaton consists of the bisimulation game is also reflected in the following fact which was proven in [Ven04] stating that coalgebra automata work "modulo bisimulation".

5.1.14. FACT. [Ven04, Prop. 4.7] Let $T : \mathsf{Set} \to \mathsf{Set}$ be a standard weak pullback preserving functor and let $(\mathbb{S}_1, s^1_I), (\mathbb{S}_2, s^2_I)$ be rooted T-coalgebras such that $(\mathbb{S}_1, s^1_I) \leftrightarrows_T (\mathbb{S}_2, s^2_I)$. Then $\mathbb{A}$ accepts $(\mathbb{S}_1, s^1_I)$ iff $\mathbb{A}$ accepts $(\mathbb{S}_2, s^2_I)$.

But at first sight T-automata seem not to be a proper generalization of graph automata, as they work only on T-coalgebras without allowing any coloring. This leads to the definition of C-colored coalgebras and C-chromatic T-automata.

5.1.15. DEFINITION. Let $T : \mathsf{Set} \to \mathsf{Set}$, $C \in \mathsf{Set}$ and $\mathbb{S} = (S, \sigma)$ a T-coalgebra. Then a *C-coloring of* $\mathbb{S}$ is a function $\gamma : S \to C$. The *C-colored T-coalgebra* $\mathbb{S} \oplus \gamma := (S, \gamma, \sigma)$ can be identified with the $C \times T$-coalgebra $(S, \langle\gamma, \sigma\rangle)$.

5.1.16. DEFINITION. Let C be a finite set. A *C-chromatic T-automaton* is a quintuple $\mathbb{A} = (A, a_I, C, \Delta, \Omega)$ such that $\Delta : A \times C \to \mathcal{P}(\mathcal{P}(TA))$ (and A, a_I, and Ω as for T-automata).

Given such an automaton and a $C \times T$-coalgebra $\mathbb{S} = (S, \gamma, \sigma)$, the acceptance game $\mathcal{G}_C(\mathbb{A}, \mathbb{S})$ is defined as the acceptance game for T-automata with the only difference that $\exists$ has to move from a position (s, a) to a position (s, Φ) such that $\Phi \in \Delta(a, \gamma(s))$. See Table 5.1 for the details.

Position: b	Player	Admissible moves: $E[b]$	$\Omega'(b)$
$(s, a) \in S \times A$	$\exists$	$\{(s, \Phi) \in S \times \mathcal{P}TA \mid \Phi \in \Delta(a, \gamma(s))\}$	$\Omega(a)$
$(s, \Phi) \in S \times \mathcal{P}TA$	$\forall$	$\{(s, \phi) \in S \times TA \mid \phi \in \Phi\}$	0
$(s, \phi) \in S \times TA$	$\exists$	$\{Z \in \mathcal{P}(S \times A) \mid (\sigma(s), \phi) \in \overline{T}Z\}$	0
$Z \in \mathcal{P}(S \times A)$	$\forall$	Z	0

Table 5.1: Acceptance game for a chromatic T-automaton

5.1.17. Example. The well-known word, tree and graph automata are instantiations of this notion. They correspond to C-chromatic T-automata for the following functors T:

Automata on	Functor
infinite words	Id
infinite binary trees	$Id \times Id$
infinite graphs	$\mathcal{P}$

It was shown in [Ven04, Prop. 4.12] that C-chromatic T-automata and $C \times T$-automata have the same recognizing power. We will need the following fact about C-chromatic T-automata when proving that coalgebra automata are closed under projection (cf. Section 5.2.1 below).

5.1.18. Fact. To any $C \times T$-automaton $\mathbb{A}$ we can find a C-chromatic T-automaton $\mathbb{A}_C$, the *chromatic T-companion* of $\mathbb{A}$, such that $\mathbb{A}$ and $\mathbb{A}_C$ accept the same $C\times T$-coalgebras.

To be able to prove statements about T-automata we will have to reason about strategies of $\exists$ in the acceptance game. Parity games are known to enjoy a strong form of determinacy: in any position of the game board either $\exists$ or $\forall$ has a history-free winning strategy (cf. Theorem C.0.12). Therefore we can focus on $\exists$'s history-free strategies.

5.1.19. Definition. Given a T-coalgebra $\mathbb{S} = (S, \sigma)$ and a T-automaton $\mathbb{A} = (A, a_I, \Delta, \Omega)$ a *history-free strategy* of $\exists$ in $\mathcal{G}(\mathbb{S}, \mathbb{A})$ is a pair of functions

$$(\Phi : S \times A \to \mathcal{P}TA,\ Z : S \times TA \to \mathcal{P}(S \times A)).$$

A partial strategy of the kind $\Phi : S \times A \to \mathcal{P}TA$ will often be represented as a map $\Phi : S \to (\mathcal{P}TA)^A$; values of this map will be denoted as Φ_s, etc. Given a strategy (Φ, Z) of $\exists$ in $\mathcal{G}(\mathbb{S}, \mathbb{A})$ we let $\mathrm{Win}(\Phi, Z)$ denote the set of positions of the game at which (Φ, Z) is a history-free strategy of $\exists$.

5.2 Closure properties

We now come to the central part of this chapter - the discussion of closure properties of T-automata. This section is based on the paper [KV05]. Its main technical result can be formulated as follows.

5.2.1. THEOREM. *Let T be some standard set functor that preserves weak pullbacks. Then every T-automaton has a non-deterministic equivalent. Hence, a T-language is recognizable iff it is non-deterministically recognizable.*

When discussing closure properties we say that a class C of T-languages is closed under some operation on T-languages if whenever we apply this operation to a family of languages of C we obtain again a language in C. For example, one may easily prove that the class of recognizable T-languages is closed under taking intersection and union; with some more effort we will show that the class of non-deterministically recognizable T-languages is closed under projection. Theorem 5.2.1 allows us to strengthen the above list of closure properties as follows.

5.2.2. THEOREM. *Let T be some standard set functor that preserves weak pullbacks. Then the class of recognizable T-languages is closed under union, projections and intersection.*

Our *proofs* for these results are of course built on generalizations, to the coalgebraic level, of (well) known ideas from the theory of specific automata. This applies in particular to results on graph automata [JW95] and the abstract universal algebraic approach of [AN01].

Throughout this section we let $T : \mathsf{Set} \to \mathsf{Set}$ be some standard and weak pullback preserving functor.

5.2.1 Closure under union, intersection and projection

In this subsection we show that the class of non-deterministically recognizable languages is closed under taking union and projection, whereas the class of recognizable languages is shown to be closed under union and intersection. Combined with Theorem 5.2.1, this suffices to prove Theorem 5.2.2.

We first define the sum and product of two T-automata, and prove that they recognize, respectively, the union and the intersection of the languages associated with the original automata.

5.2.3. DEFINITION. Let $\mathbb{A}_1 = (A_1, a^1_I, \Delta_1, \Omega_1)$ and $\mathbb{A}_2 = (A_2, a^2_I, \Delta_2, \Omega_2)$ be two T-automata. We will define their *sum* $\mathbb{A}_\cup$ and *product* $\mathbb{A}_\cap$. Both of these automata will have the *disjoint union*

$$A_{12} := \{*\} \uplus A_1 \uplus A_2$$

as their collection of states. Also, the parity function Ω will be the same for both automata:

$$\Omega(a) := \begin{cases} 0 & \text{if } a = *, \\ \Omega_i(a) & \text{if } a \in A_i. \end{cases}$$

The only difference between the automata lies in the transition functions, which are defined as follows:

$$\Delta_{\cup}(a) := \begin{cases} \Delta_1(a_I^1) \cup \Delta_2(a_I^2) & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i, \end{cases}$$

$$\Delta_{\cap}(a) := \begin{cases} \{\Phi_1 \cup \Phi_2 \mid \Phi_i \in \Delta_i(a_I^i)\} & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i. \end{cases}$$

Finally, we put $\mathbb{A}_{\cup} := (A_{12}, a_I, \Delta_{\cup}, \Omega)$ and $\mathbb{A}_{\cap} := (A_{12}, a_I, \Delta_{\cap}, \Omega)$, where $a_I := *$.

The following proposition presents the announced closure properties.

5.2.4. Proposition. *Let $\mathbb{A}_1$ and $\mathbb{A}_2$ be two T-automata. Then for any pointed T-coalgebra $(\mathbb{S}, s_I)$ we have:*

1. *$\mathbb{A}_{\cup}$ accepts $(\mathbb{S}, s_I)$ iff $\mathbb{A}_1$ or $\mathbb{A}_2$ accepts $(\mathbb{S}, s_I)$,*
2. *$\mathbb{A}_{\cap}$ accepts $(\mathbb{S}, s_I)$ iff both $\mathbb{A}_1$ and $\mathbb{A}_2$ accept $(\mathbb{S}, s_I)$.*
3. *$\mathbb{A}_{\cup}$ is non-deterministic if $\mathbb{A}_1$ and $\mathbb{A}_2$ are so.*

Proof. First suppose that the automaton $\mathbb{A}_{\cup}$ accepts $(\mathbb{S}, s_I)$. Hence by definition, $\exists$ has a winning strategy (Φ, Z) in the acceptance game $\mathcal{G} := \mathcal{G}(\mathbb{S}, \mathbb{A}_{\cup})$ starting from position $(s_I, *)$. Let i be such that $\Phi(s_I, *) \in \Delta(a_I^i)$. It is then straightforward to verify that (Φ, Z), restricted to $\exists$'s positions in $\mathcal{G}(\mathbb{S}, \mathbb{A}_i)$, is a winning strategy for $\exists$ from position (s_I, a_I^i). From this it is immediate that $\mathbb{A}_i$ accepts $(\mathbb{S}, s_I)$. The other statements of the proof admit similarly straightforward proofs. QED

We now turn to the proof of the fact that T-automata are closed under projection. In the following all T-automata are assumed to be non-deterministic. To facilitate the presentation we will think of the transition function Δ as a map $A \to \mathcal{P}TA$ and the first component Φ of a strategy (Φ, Z) for $\exists$ in an acceptance game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ will be regarded as a function of type $A \times S \to TA$ (that is, we identify singleton sets with their unique elements).

5.2.5. Definition. Let C be a finite set, $\mathbb{A} = (A, a_I, \Delta, \Omega)$ be a $(C \times T)$-coalgebra automaton and $\mathbb{A}_C = (A, a_I, C, \Delta_C, \Omega)$ its C-chromatic T-companion, see Fact 5.1.18. Then we define the C-projection $\pi_C\mathbb{A} := (A, a_I, \Delta_\pi, \Omega)$ where $\Delta_\pi(a) := \bigcup_{c \in C} \Delta_C(c, a)$.

The C-projection of $\mathbb{A}$ accepts all the underlying rooted T-coalgebras of rooted $C \times T$-coalgebras that are accepted by $\mathbb{A}$.

5.2.6. Lemma. *If a C-chromatic T-automaton $\mathbb{A}$ accepts the $(C \times T)$-coalgebra $(\mathbb{S}, s_I) := (S, \gamma, \sigma, s_I)$ then $\pi_C\mathbb{A}$ accepts $(\mathbb{S}^\pi, s_I) := (S, \sigma, s_I)$.*

Proof. The proof is easy. One has to realize that a winning strategy (Φ, Z) of $\exists$ in the game for $\mathbb{A}_C$ is still a winning strategy of $\exists$ in the $\pi_C\mathbb{A}$ acceptance game. QED

The converse of this lemma however fails in general. Let $\mathbb{A}$ be some $C \times T$-automaton and let (S, σ, s_I) be a pointed T-coalgebra that is accepted by $\pi_C\mathbb{A}$. Then we know that $\exists$ has a winning strategy (Φ, Z) in $\mathcal{G}(\mathbb{S}, \pi_C\mathbb{A})$ from position (s_I, a_I). We would like to ensure that (Φ, Z) is also a winning strategy in $\mathcal{G}(\mathbb{S}, \mathbb{A}_C)$ by defining a coloring $\gamma : S \to C$ as follows: $\gamma(s) := c$ if there is a match of $\mathcal{G}(\mathbb{S}, \pi_C\mathbb{A})$, starting from position (s_I, a_I) and conform $\exists$'s strategy, in which a position (s, a) occurs and $\Phi_{s,a} \in \Delta_C(c, a)$. In general, however, there may be *distinct* positions (s, a_1) and (s, a_2) that $\forall$ may force the match to pass through, and it may not be possible to find a single $c \in C$ such that both $\Phi_{s,a_1} \in \Delta(c, a_1)$ and $\Phi_{s,a_2} \in \Delta(c, a_2)$. To avoid this problem we introduce now the notion of *strong* acceptance.

5.2.7. Definition. Let $\mathbb{A}$ be a T-automaton and $(\mathbb{S}, s_I)$ a rooted T-coalgebra. A history free strategy (Φ, Z) for $\exists$ in the game $\mathcal{G} = \mathcal{G}(\mathbb{S}, \mathbb{A})$ initialized at (s_I, a_I) is called *scattered* if the relation

$$\{(s_I, a_I)\} \cup \bigcup \{Z_{s,\phi} \subseteq S \times A \mid (s, \phi) \in \mathrm{Win}(\Phi, Z)\}$$

is the graph of some possibly partial function. Furthermore we say that $\mathbb{A}$ *strongly accepts* the rooted coalgebra $(\mathbb{S}, s_I)$ if $\exists$ has a scattered winning strategy in the game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ initialized at position (s_I, a_I).

Under the condition that $\mathbb{A}$ strongly accepts $(\mathbb{S}, s_I)$ we can now prove the converse of Lemma 5.2.6.

5.2.8. Lemma. *Let $\mathbb{A}$ be a $C \times T$-automaton, and let $(\mathbb{S}, s_I)$ be a rooted T-coalgebra that is strongly accepted by $\pi\mathbb{A}$. Then there is a C-coloring $\gamma : S \to C$ of $\mathbb{S}$ such that $\mathbb{A}$ accepts (S, γ, σ, s_I).*

Proof. Let (Φ, Z) be a scattered winning strategy for $\exists$ in $\mathcal{G}(\mathbb{S}, \pi\mathbb{A})$ that exists according to our assumption that $(\mathbb{S}, s_I)$ is strongly accepted by $\pi\mathbb{A}$. According to the definition of scatteredness we can assign to every $s \in S$ a state $a_s \in A$ such that $a_r = a_I$, and if $(s, a) \in Z_{s,\phi}$ for some winning position (s, ϕ) of $\exists$, then $a = a_s$. Then we define a function $\gamma : S \to C$ as follows. If there is a $c \in C$ such that $\Phi_{s,a_s} \in \Delta_C(c, a)$, then we pick such a c and put $\gamma(s) := c$; if there is no such c, then we define $\gamma(s) := d$ for some arbitrary $d \in C$. It follows from these definitions that (Φ, Z) is a winning strategy for $\exists$ in $\mathcal{G}_C(\mathbb{S} \oplus \gamma, \mathbb{A}_C)$ from position (s_I, a_I). From this it is immediate that $\mathbb{A}$ accepts (S, γ, σ, s_I). QED

At first sight the acceptance condition of a T-automaton seems to be strictly weaker than strong acceptance. One can easily think of an example of a T-coalgebra that is not strongly accepted by some T-automaton.

5.2.9. Example. Let $T = Id$ and $\mathbb{A} = (A, a_I, \Delta, \Omega)$ where $A := \{a_I, a\}$, $\Delta(a_I) := \{\{a_I\}, \{a\}\}$, $\Delta(a) := \{\{a\}\}$ and $\Omega(a_I) := 1, \Omega(a) := 0$. Furthermore let $(\mathbb{S}, s_I) := (\{*\}, id_{\{*\}}, *)$. Then $\mathbb{A}$ accepts $(\mathbb{S}, s_I)$ but does not strongly accept $(\mathbb{S}, s_I)$.

The next lemma, however, shows that if a rooted coalgebra is accepted by some automaton, but not strongly so, then we can always find a bisimilar pointed coalgebra that is strongly accepted.

5.2.10. Lemma. *Let $\mathbb{A}$ be a T-automaton, and let $(\mathbb{S}, s_I) = (S, \sigma, s_I)$ be a pointed T-coalgebra that is accepted by $\mathbb{A}$. Then there is a pointed T-coalgebra $(\overline{S}, \overline{\sigma}, \overline{s_I})$ such that*

1. $(\mathbb{S}, s_I) \leftrightarrows_T (\overline{S}, \overline{\sigma}, \overline{s_I})$ *and*
2. $\mathbb{A}$ *strongly accepts* $(\overline{S}, \overline{\sigma}, \overline{s_I})$.

Proof. The coalgebra $\overline{\mathbb{S}}$ will be based on the set $\overline{S} := S \times A$, and as the root $\overline{s}_I$ of $\overline{\mathbb{S}}$ we take the pair $(\overline{s}_I, a_I)$. For the definition of the coalgebra structure $\overline{\sigma}$, we need some auxiliary definitions.

First we define a coalgebra map $\tilde{\sigma} : S \times A \to T(S \times A)$ such that $(\overline{S}, \tilde{\sigma})$ is isomorphic to the A-fold coproduct of $\mathbb{S}$. Recall the well-known fact that the forgetful functor $U : \mathsf{Coalg}(T) \to \mathsf{Set}$ creates colimits (cf. [Bar93, Prop. 1.1]) and therefore the underlying set of a colimit and the canonical morphisms into the colimit can be computed as in Set. This means in particular that we can define $\tilde{\sigma} : S \times A \to T(S \times A)$ as follows: for $a \in A$ let $\kappa_a : S \to S \times A$ be the map that maps $s \in S$ to $(s, a) \in \overline{S}$. The set $S \times A$ together with the maps $\kappa_a : S \to S \times A$ is isomorphic to the A-fold coproduct of S in Set. Because U creates colimits this means that we can find a coalgebra map $\tilde{\sigma} : S \times A \to T(S \times A)$ such that $(\overline{S}, \tilde{\sigma}) \in \mathsf{Coalg}(T)$ together with the maps $\kappa_a : S \to \overline{S}$ is isomorphic to the the A-fold coproduct $\coprod_{a \in A} \mathbb{S}$ of $\mathbb{S}$ in $\mathsf{Coalg}(T)$. In particular this means that the κ_a's are $\mathsf{Coalg}(T)$-morphisms from $\mathbb{S}$ to $(\overline{S}, \tilde{\sigma})$. Let $\pi_S : S \times A \to S$ be the map that maps a pair $(s, a) \in S \times A$ to S. Then it is easy to show that the fact that the κ_a's are coalgebra morphisms implies that $\pi_S : (\overline{S}, \tilde{\sigma}) \to \mathbb{S}$ is a coalgebra morphism as well.

Second, given a relation $R \subseteq S \times A$, define the relation $\hat{R} \subseteq \overline{S} \times A$ by putting

$$\hat{R} := \{((s, a), a) \mid (s, a) \in R\}.$$

Then clearly we have that $R = \mathrm{Gr}(\pi_S)^{\smile} \circ \hat{R}$, and hence,

$$\overline{T}R = \mathrm{Gr}(T\pi_S)^{\smile} \circ \overline{T}\hat{R}. \tag{5.1}$$

Now let now (Φ, Z) be a winning strategy of $\exists$ in $\mathcal{G}(\mathbb{S}, \mathbb{A})$ from position (s_I, a_I) that exists according to our assumption. For the definition of $\overline{\sigma} : \overline{S} \to T\overline{S}$, consider an arbitrary element $(s, a) \in \overline{S}$, and distinguish cases. We first look at the case in which (Φ, Z) is a *winning* strategy of $\exists$ in the game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ from position (s, a), i.e.

$(s,a) \in \mathrm{Win}(\Phi, Z)$. Using (5.1), it follows from $(\sigma(s), \Phi_{s,a}) \in \overline{T}Z_{(s,\Phi_{s,a})}$, that we may define $\overline{\sigma}(s,a)$ to be some element in $T\overline{S}$ satisfying $(\sigma(s), \overline{\sigma}(s,a)) \in \mathrm{Gr}(T\pi_S)^{\smile}$ and $(\overline{\sigma}(s,a), \phi) \in \overline{T}\hat{Z}_{(s,\Phi_{s,a})}$. If, on the other hand, $(s,a) \notin \mathrm{Win}(\Phi, Z)$, then we simply put $\overline{\sigma}(s,a) := \tilde{\sigma}(s,a)$.

It is completely straightforward to check that the map π_S is in fact an T-coalgebraic homomorphism from $\overline{\mathbb{S}}$ onto $\mathbb{S}$. From this, the first statement of the proposition follows immediately.

For the second statement, define the strategy $(\overline{\Phi}, \overline{Z})$ with $\overline{\Phi} : \overline{S} \times A \to A$ and $\overline{Z} : \overline{S} \times TA \to \mathcal{P}(\overline{S} \times A)$ as follows:

$$\begin{array}{rcll} \overline{\Phi}: & ((s,a),b) & \mapsto & \Phi_{s,b} \\ \overline{Z}: & ((s,a),\phi) & \mapsto & \hat{Z}_{s,\phi}. \end{array}$$

Since all relations chosen by $\exists$ are of the form $\hat{R}$, and all elements of such relations are of the form $((s,a),b)$ with $a = b$, it is obvious that the set $\{((s,a_I),a_I)\} \cup \bigcup\{\hat{Z}_{s,\phi} \mid (s,\phi) \in S \times TA\}$ is functional. In other words, the strategy is scattered.

Thus it is left to prove that $(\overline{\Phi}, \overline{Z})$ guarantees $\exists$ to win any match of $\mathcal{G}(\overline{\mathbb{S}}, \mathbb{A})$ starting from $(\overline{s}_I, a_I)$. To see why this is the case, consider an arbitrary position $((s,a),a)$ with $(s,a) \in \mathrm{Win}(\Phi, Z)$, and abbreviate $\phi := \Phi_{s,a}$. Then by definition, $\overline{\Phi}((s,a),a) = \phi$ and $\overline{Z}((s,a),\phi) = \hat{Z}_{s,\phi} = \{((t,b),b) \mid (t,b) \in Z_{s,\phi}\}$. From this observation it is easy to derive that for any $\mathcal{G}(\overline{\mathbb{S}}, \mathbb{A})$ match $(\overline{s}_I, a_I)((s_1,a_1),a_1)((s_2,a_2),a_2)\ldots$ that is conform the strategy $(\overline{\Phi}, \overline{Z})$, the corresponding $\mathcal{G}(\mathbb{S}, \mathbb{A})$ match $(s_I, a_I)(s_1,a_1)(s_2,a_2)\ldots$ is conform (Φ, Z). And since this strategy was supposed to be winning for $\exists$ from (s_I, a_I), it follows that the $\mathcal{G}(\overline{\mathbb{S}}, \mathbb{A})$ match is, indeed, a win for $\exists$. This proves the second statement of the proposition. QED

5.2.11. Proposition. *Let $\mathbb{A}$ be some $(C \times T)$-automaton. Then the following are equivalent, for every rooted T-coalgebra $(\mathbb{S}, s_I)$:*

1. *$\pi_C\mathbb{A}$ accepts $(\mathbb{S}, s_I)$,*
2. *$\mathbb{A}$ accepts a $(C \times T)$-coalgebra $(S', \gamma, \sigma', s')$ such that (S', σ', s') and $(\mathbb{S}, s_I)$ are bisimilar.*

Proof. $1 \Rightarrow 2$: Suppose $\pi_C\mathbb{A}$ accepts $(\mathbb{S}, s_I)$. Then by Lemma 5.2.10 there is a pointed T-coalgebra (S', σ', s') that is bisimilar to $(\mathbb{S}, s)$ and that is strongly accepted by $\pi_C\mathbb{A}$. But then by Lemma 5.2.8 there is a coloring $\gamma : S' \to C$ such that $\mathbb{A}$ accepts $(S', \gamma, \sigma', s')$.
$2 \Rightarrow 1$: This follows from Lemma 5.2.6 and that T-automata do not distinguish between bisimilar T-coalgebras (cf. Fact 5.1.14). QED

5.2.2 From alternating automata to nondeterministic ones

In this section we prove the main technical result of this chapter, Theorem 5.2.1. That is, we will construct, for an arbitrary, alternating T-automaton an equivalent *non-*

deterministic T-automaton. Throughout this section we will be working with a fixed (but arbitrary) T-automaton $\mathbb{A} = (A, a_I, \Delta, \Omega)$.

Before going into the technical details of the construction, let us first provide some of the intuitions behind our approach. These intuitions ultimately go back to ideas of Muller and Schupp, see for instance [MS95], but in particular, our proof generalizes work by Janin and Walukiewicz [JW95], using the approach of Arnold and Niwiński [AN01].

The main idea is to bring the players' interaction pattern $\exists\forall\exists\forall$ in one round of the acceptance games for $\mathbb{A}$, into the 'strategic form' $\exists\forall$ (or more precisely: $\exists\exists\forall$). Concretely, consider a basic position $(s, a) \in S \times A$ in the acceptance game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ for some T-coalgebra $\mathbb{S}$. From this position a play proceeds as follows:

- $\exists$ picks $\Phi \in \Delta(a)$, moving to position (s, Φ);
- $\forall$ picks $\phi \in \Phi$, moving to position (s, ϕ);
- $\exists$ picks $Z_\phi \subseteq S \times A$ with $(\sigma(s), \phi) \in \overline{T} Z_\phi$ — this Z_ϕ is the new position;
- $\forall$ picks $(t, b) \in Z_\phi$ as the next basic position.

Now the crucial point is that $\exists$ may *gather* her family $\{Z_\phi \subseteq S \times A \mid \phi \in \Phi\}$ into *one single* relation $Z_\Phi \subseteq S \times \mathcal{P}A$, and that we may modify the game in such a way that this is an appropriate answer for $\exists$. This approach would suggest to take (representations of) subsets of A as the states of the new automaton $\mathbb{A}^d$.

However, as is well-known from the literature, such a straightforward subset construction may work for automata that operate on finite objects, in the case of automata for (possibly) infinite objects this approach fails to make some subtle but crucial distinctions. The remedy, which brings us to the second fundamental idea underlying our construction, is to use *binary relations on* A, rather than subsets of A, to bring the acceptance game into some kind of 'layered-strategic' form. Then, using the notion of a *trace* through a sequence of such relations, we have an established tool at our disposal for bringing the interaction pattern of the acceptance game into the required format. Our contribution here is to show that all of this can be done in the abstract context of coalgebras for an arbitrary standard, weak pullback preserving functor.

Now we are ready for the technical details of the construction.

5.2.12. Definition. Given a finite word $\rho = R_1 R_2 \ldots R_n$ over the set $Rel(A)$ of binary relations over A, a *trace* through ρ is an A-word $\alpha = a_0 a_1 a_2 \ldots a_k$ with $k \leq n$ such that $a_0 = a_I$ is the initial state a_I of the automaton, and $a_i R_{i+1} a_{i+1}$ for all $i < k$. Similar definitions apply to (finite or infinite) traces on infinite $Rel(A)$-words.

A trace α is a *trap for* $\exists$ if $\Delta(a_i) = \emptyset$ for some state a_i on α; a trace α is *bad* if it is a trap for $\exists$ or, in case α is infinite, if $\max\{\Omega(a_i) \mid i \in \mathrm{Inf}(\alpha)\}$ is odd.

As we will see, traces may be associated with matches of the acceptance game for $\mathbb{A}$, bad traces with the ones that are lost by $\exists$. Let us look at this in a bit more detail. As a consequence of the generality that we aim for, there are two different ways in which $\exists$ may loose a match. She may either get stuck at some finite stage of the match (either at a basic position or at a position of the form $(s, \phi) \in S \times TA$), or survive for infinitely many rounds but fail to establish the winning condition. Now the traces that are traps for $\exists$ will correspond to matches in which she gets stuck in a *basic* position, whereas the other kind of badness will turn out to be an encoding of $\exists$'s failing to win an infinite match. For finite matches that $\exists$ looses because of getting stuck in a non-basic position, we do not need a corresponding notion for traces.

The first proposition that we need is a variation on well-known results. It concerns the existence of a deterministic *word* automaton that accepts those words over $Rel(A)$ which contain no bad traces. Since there are two kinds of bad traces, this automaton needs to perform a double task: it needs to recognize traps for $\exists$, and it needs to take proper care of the infinite words. It will be convenient to have the automaton perform these two jobs more or less separately. That is, the automaton will have a special state $m_\forall$ signaling that $\exists$ has been trapped. In order to formulate the proposition we need some notation.

5.2.13. Notation. *Given a deterministic automaton $\mathbb{D} = (D, d_I, \Sigma, \delta, \Omega_{\mathbb{D}})$ operating on possibly infinite words with alphabet Σ and transition function δ, we let $\hat{\delta} : D \times \Sigma^* \to D$ denote the* iterated *transition function, inductively defined by $\hat{\delta}(d, \epsilon) = d$ and $\hat{\delta}(d, \alpha a) = \delta(\hat{\delta}(d, \alpha), a)$.*

5.2.14. Proposition. *There is a deterministic word automaton*

$$\mathbb{M}_0 = (M, m_I, Rel(A), \mu_0, \Omega_0),$$

operating on $Rel(A)$-words, and containing a special state $m_\forall$, such that:

1. *$\mu_0(m_\forall, R) = m_\forall$ for all $R \in Rel(A)$,*
2. *for any finite $Rel(A)$-word ρ: $\hat{\mu}(\rho) = m_\forall$ iff ρ contains a trap for $\exists$,*
3. *for any infinite $Rel(A)$-word ρ: $\mathbb{M}_0$ accepts ρ iff ρ contains no bad traces.*

Proof. We construct $\mathbb{M}_0$ in several steps.

Step 1: We define a non-deterministic word automaton $\mathbb{M}' = (M', m'_I, Rel(A), \mu', \Omega')$ by letting $M' := A \cup \{*\}$, $m'_I := a_I$ and

$$\mu'(m, R) := \begin{cases} R[a] & \text{if } a \neq * \text{ and } \Delta(a) \neq \emptyset \\ * & \text{otherwise,} \end{cases}$$

$$\Omega'(m) := \begin{cases} \Omega(m) + 1 & \text{if } m \in A \\ 0 & \text{if } m = *. \end{cases}$$

It is not difficult to check that $\mathbb{M}'$ accepts an infinite $Rel(A)$-word α iff α contains a bad trace.

Step 2: We use several standard constructions from automata theory: first we transform $\mathbb{M}'$ into an equivalent nondeterministic Büchi automaton (cf. [GTW02, Chapter 1]) , and transform the resulting automaton into an equivalent deterministic Muller automaton using the Safra construction (cf. [GTW02, Theorem 3.6]). The resulting deterministic Muller automaton can be easily transformed into a deterministic Muller automaton accepting the complement of the language of $\mathbb{M}'$. Finally we transform this deterministic Muller automaton into an equivalent deterministic parity automaton which we call $\overline{\mathbb{M}}'$. To sum it up we obtain a deterministic parity automaton $\overline{\mathbb{M}}' = (\overline{M}', \overline{m}'_I, Rel(A), \overline{\mu}', \overline{\Omega}')$ that accepts the complement of the language accepted by $\mathbb{M}'$, i.e. $\overline{\mathbb{M}}'$ accepts an infinite $Rel(A)$-word α iff it does not contain a bad trace.

Step 3: This is the last step of the construction in which we bring the automaton $\overline{\mathbb{M}}'$ in the special shape which is required by our proposition. To this aim we define $M := \overline{M}' \times \mathcal{P}A \cup \{m_\forall\}$, $m_I := \overline{m}'_I \times \{a_I\}$ and

$$\begin{aligned}
\mu_0((m, A'), R) &:= \begin{cases} (\overline{\mu}'(m, R), \bigcup_{a \in A'} R[a]) & \text{if} \quad \text{for all } a' \in A'.\Delta(a') \neq \emptyset \\ m_\forall & \phantom{\text{if} \quad} \text{otherwise} \end{cases} \\
\mu_0(m_\forall) &:= m_\forall. \\
\Omega((m, A')) &:= \overline{\Omega}'(m) \\
\Omega(m_\forall) &:= 1
\end{aligned}$$

QED

In the remainder of this section we *fix* the automaton $\mathbb{M}_0 = (M, m_I, \mu_0, \Omega_0)$ and state $m_\forall$ as given in Proposition 5.2.14. We turn now to the main construction of the proof. Below we define a non-deterministic automaton $\mathbb{M}_1$ which operates on $(\mathcal{P}TA)^A$-colored T-coalgebras, $\mathbb{S} \oplus \Phi$, that is, T-coalgebras $\mathbb{S} = (S, \sigma)$ that are colored by the map $\Phi : S \to (\mathcal{P}TA)^A$, i.e. by a (potential) *strategy* of $\exists$ in the game $\mathcal{G}(\mathbb{S}, \mathbb{A})$ that is *partial* in the sense of dealing with basic positions only. More precisely, for any position $(s, a) \in S \times A$, we let the value $\Phi_{s,a} \in \mathcal{P}TA$ encode the move $(s, \{\Phi_{s,a}\}) \in S \times \mathcal{P}TA$. Our aim with the automaton $\mathbb{M}_1$ is that it will recognize precisely those pointed $(\mathcal{P}TA)^A$-colored T-coalgebras (S, σ, Φ, s_I) of which Φ forms the basic part of a winning strategy in the game $\mathcal{G}(\mathbb{S}, \mathbb{A})$.Towards the end of this section we will see that this suffices to prove Theorem 5.2.1.

For the definition of $\mathbb{M}_1$ we need some preliminary definitions.

5.2.15. Definition. An object $\Xi \in T\mathcal{P}A$ is called an *T-redistribution* of a subset $\Phi \subseteq TA$ if $(\phi, \Xi) \in \overline{T}(\in_A)$ for all $\phi \in \Phi$.

An object $\Pi \in TRel(A)$ is called an *T-redistributive relational representation* of an element $\Phi \in (\mathcal{P}TA)^A$, or shortly: an *$T$-relation for* Φ, if $(Tev_a)(\Pi)$ is a redistribution of $\Phi(a)$ for all $a \in A$. Here $ev_a : Rel(A) \to \mathcal{P}A$ is the map given by $ev_a : R \mapsto R[a]$. The collection of T-relations for $\Phi \in (\mathcal{P}TA)^A$ is denoted as $\mathcal{R}_T(\Phi)$.

The intuitions on these notions are as follows. Concerning redistributions, the point is that for any T-coalgebra $\mathbb{S} = (S, s_I)$, any point $s \in S$ and any set $\Phi \in \mathcal{P}TA$, there is a 1-1 correspondence between:

- families $\{Z_\phi \subseteq S \times A \mid \phi \in \Phi\}$ of relations such that $(\sigma(s), \phi) \in \overline{T}Z_\phi$ for all $\phi \in \Phi$, and
- pairs (Z_Φ, Ξ) consisting of a relation $Z_\Phi \subseteq S \times \mathcal{P}A$, and an T-redistribution $\Xi \in T\mathcal{P}A$ of Φ, such that $(\sigma(s), \Xi) \in \overline{T}Z_\Phi$.

In brief, redistributions enable us to gather the information of a family $\{Z_\phi \subseteq S \times A \mid \phi \in \Phi\}$ of relation moves of $\exists$ into one single relation $Z_\Phi \subseteq S \times \mathcal{P}A$.

However, this regrouping of information on $\exists$'s strategy in terms of redistributions has one shortcoming: it is based on *subsets* of A whereas we already pointed out that such an encoding will not suffice to encode the full flow of information when transforming alternating automata into non-deterministic ones. This is where the notion of an T-relation for Φ comes in. The important observation is that any element Π of the set $TRel(A)$ has the right shape to represent a family $\{\Phi_a \in \mathcal{P}TA \mid a \in A\}$: the point is that we may use, for every $a \in A$, the map $Tev_a : TRel(A) \to T\mathcal{P}A$ to provide an element $(Tev_a)(\Pi)$ in the right set $T\mathcal{P}A$ of (potential) T-redistributions of Φ. Thus, the definition of a $\Pi \in TRel(A)$ being an T-redistributive relational representation of $\Phi \in (\mathcal{P}TA)^A$ forms, at least potentially, an adequate formalization of the requirement that Π and Φ 'fit well together'. As we will see below, it also forms the key to lead the flow of information in acceptance games for alternating automata into a non-deterministic channel.

5.2.16. Definition. Let $\mathbb{M}_1$ be the non-deterministic $(\mathcal{P}TA)^A$-chromatic T-automaton $(M, m_I, (\mathcal{P}TA)^A, \mu, \Omega_0)$, where $\mu : M \times (\mathcal{P}TA)^A \to \mathcal{P}\mathcal{P}TM$ is the map defined by

$$\mu(m, \Phi) := \begin{cases} \left\{ \{(T\mu_m)(\Pi)\} \mid \Pi \in \mathcal{R}_T(\Phi) \right\} & \text{if } m \neq m_\forall, \\ \emptyset & \text{if } m = m_\forall. \end{cases}$$

Here $\mu_m : Rel(A) \to M$ is given by $\mu_m(R) := \mu_0(m, R)$.

5.2.17. Remark. Let $\mathbb{M}_1$ be as above, and $\mathbb{S} = (S, \sigma, \Phi)$ some $(\mathcal{P}TA)^A$-colored T-coalgebra. Note that the acceptance game $\mathcal{G}(\mathbb{S}, \mathbb{M}_1)$ is summarized in Table 5.2.

Given the definition of μ, it is not hard to see that, from a position $(s, m) \in S \times M$, with subsequent moves of $\exists$, say, $(s, \{K\}) \in S \times \mathcal{P}TM$ and $Z \subseteq S \times M$, we may associate an element $\Pi \in TRel(A)$ and a relation $Y \subseteq S \times Rel(A)$ such that Π is an T-relation for Φ_s, $(T\mu_m)(\Pi) = K$ and $(\sigma(s), \Pi) \in \overline{T}Y$.

To start with, it is obvious from the definitions that there is some $\Pi \in \mathcal{R}_T(\Phi_s)$, such that $(T\mu_m)(\Pi) = K$. Now define the relation $Y := \{(t, R) \in S \times Rel(A) \mid (t, \mu_m(R)) \in Z\}$. Clearly, this relation is the composition of Z with the converse relation $\mathrm{Gr}(\mu_m)^{\breve{}}$ of the graph of the function μ_m. From this it follows that $\overline{T}Y = \overline{T}Z \circ \overline{T}(\mathrm{Gr}(\mu_m)^{\breve{}})$.

Position: b	Type	Player	Admissible moves:	$\Omega(b)$
(s, m)	$S \times M$	$\exists$	$\{(s, \{K\}) \in S \times \mathcal{P}TM \mid K \in \mu(m, \Phi(s))\}$	$\Omega_0(m)$
$(s, \{K\})$	$S \times \mathcal{P}TM$	$\forall$	$\{(s, K)\}$	0
(s, K)	$S \times TM$	$\exists$	$\{Z \in \mathcal{P}(S \times M) \mid (\sigma(s), K) \in \overline{T}Z\}$	0
Z	$\mathcal{P}(S \times M)$	$\forall$	Z	0

Table 5.2: Acceptance game for $\mathbb{M}_1$

Also, rewriting $(T\mu_m)(\Pi) = K$, we obtain that $(\Pi, K) \in \mathrm{Gr}(T\mu_m) = \overline{T}\mathrm{Gr}(\mu_m)$, so that $(K, \Pi) \in (\overline{T}\mathrm{Gr}(\mu_m))^{\smile} = \overline{T}(\mathrm{Gr}(\mu_m)^{\smile})$. Hence, from $(\sigma(s), K) \in \overline{T}Z$ it is immediate that $(\sigma(s), \Pi) \in \overline{T}Y$.

5.2.18. Proposition. *For any pointed T-coalgebra $(\mathbb{S}, s_I)$ and any $(\mathcal{P}TA)^A$-coloring Φ of S, the following are equivalent:*

1. *Φ is part of a winning strategy for $\exists$ in $\mathcal{G}(\mathbb{S}, \mathbb{A})$ form position (s_I, a_I);*
2. *$\mathbb{M}_1$ accepts $(\mathbb{S} \oplus \Phi, s_I)$.*

Proof. Recall that every infinite game may be represented as a tree, and that strategies of either player, limiting the possible course of actions, can be represented as *subtrees* of this game tree. Thus, both with a Φ-extending strategy of $\exists$ in $\mathcal{G} = \mathcal{G}(\mathbb{S}, \mathbb{A})$, and with a strategy of $\exists$ in the acceptance game $\mathcal{G}' = \mathcal{G}(\mathbb{S} \oplus \Phi, \mathbb{M}_1)$, we may associate such subtrees of the game trees of $\mathcal{G}$ and $\mathcal{G}'$, respectively. As it turns out, these two trees turn out to be rather similar, and in fact, may be coded up into one and the same structure. This observation forms the basis of our proof of the proposition.

More specifically, we will show the equivalence of both (1) and (2) to the statement (3) below.

3. There is a labeled tree
$$\mathbb{X} = (X, x_I, \xi, u, \Pi, Q),$$
where $x_I \in X$ and $\xi : X \to \mathcal{P}X$ denote, respectively, the root and the successor function of the tree, and $u : X \to S$, $\Pi : X \to T\mathit{Rel}(A)$, and $Q : X \to \mathit{Rel}(A)$ are labellings.

   This tree is supposed to satisfy the conditions 3a–3d below. Here, and in the sequel, we abbreviate Φ_{u_x} as Φ_x, and define $W_x := \{(u_y, Q_y) \mid y \in \xi(x)\}$. Branches of the tree start at the root, and thus induce (finite or infinite) words over $\mathit{Rel}(A)$.

   (a) $u_{x_I} = s_I$ and $Q_{x_I} = \{(a_I, a_I)\}$,

   (b) for all $x \in X$, Π_x is an T-relation for Φ_x,

   (c) for all $x \in X$, $(\sigma(u_x), \Pi_x) \in \overline{T}W_x$.

(d) $\mathbb{X}$ has no bad traces (that is, no branch of $\mathbb{X}$ induces a $Rel(A)$-word containing a bad trace).

As hinted at above, our intuition about $\mathbb{X}$ is that it represents a winning strategy for $\exists$ both in $\mathcal{G}$ and in $\mathcal{G}'$ (in the case of $\mathcal{G}$, of course, a strategy completing the partial strategy Φ). Counterstrategies of $\forall$ in $\mathcal{G}$ correspond to *branches* of $\mathbb{X}$, while his strategies in $\mathcal{G}$ appear as *traces* on $\mathbb{X}$. We show now that both statements (1) and (2) listed in the proposition are equivalent to the statement (3) concerning the existence of a labeled tree $\mathbb{X}$ satisfying the properties 3a-3d.

(1) $\Rightarrow$ (3)

Suppose that Φ, together with $Z : S \times TA \to \mathcal{P}(S \times A)$, is a history free winning strategy for $\exists$ in $\mathcal{G}$. We make some preparations for the definition of $\mathbb{X}$.

First, define the following sets $Y_{s,a} \subseteq S \times A$ and $Y_s \subseteq S \times Rel(A)$ by

$$\begin{aligned} Y_{s,a} &:= \bigcup_{\phi \in \Phi_{s,a}} Z_{s,\phi}, \\ Y_s &:= \{(t, R) \in S \times Rel(A) \mid \text{ for all } a \in A \,.\, R[a] = \{b \mid (t, b) \in Y_{s,a}\}\}. \end{aligned}$$

The set $Y_{s,a}$ contains exactly those positions that $\exists$, playing her strategy (Φ, Z), may expect as a possible next position after (s, a). The other set should be seen as a way to represent the entire family $\{Y_{s,a} \mid a \in A\}$. Recall that $ev_a : Rel(A) \to \mathcal{P}A$ is given by $ev_a : R \mapsto R[a]$. It is then easy to see that

$$Y_{s,a} = Y_s \circ \mathrm{Gr}(ev_a) \circ \ni_A \,. \tag{5.2}$$

But it will also be clear that the relation Y_s is functional; let $\zeta_s : S \to Rel(A)$ be the map such that $Y_s = \mathrm{Gr}(\zeta_s)$. It follows that $\overline{T}(Y_s) = \mathrm{Gr}(T\zeta_s)$, so if we define

$$\Pi_s := (T\zeta_s)(\sigma(s)),$$

then we have ensured that $(\sigma(s), \Pi_s) \in \overline{T}Y_s$.

The most important claim is that

$$\Pi_s \text{ is an } T\text{-relation for } \Phi_s. \tag{5.3}$$

To see why this holds, fix $a \in A$, and observe that it follows from (5.2), the definition of ζ_s, and the fact that $\overline{T}(\mathrm{Gr}(f)) = \mathrm{Gr}(Tf)$ for any function f, that

$$\overline{T}(Y_{s,a}) \;=\; \mathrm{Gr}(T\zeta_s) \circ \mathrm{Gr}(Tev_a) \circ \overline{T}(\ni_A). \tag{5.4}$$

Now take some $\phi \in \Phi_{s,a}$. By the assumption that Φ and Z form a winning strategy for $\exists$, it follows that $(\sigma(s), \phi) \in \overline{T}(Z_{s,\phi})$, so by monotonicity of $\overline{T}$ we obtain that $(\sigma(s), \phi) \in \overline{T}(Y_{s,a})$. Hence from (5.4) and the definition of Π_s we may infer that $(\Pi_s, \phi) \in \mathrm{Gr}(Tev_a) \circ \overline{T}(\ni_A)$. In other words, $((Tev_a)(\Pi_s), \phi) \in \overline{T}(\ni_A)$. Since ϕ was an arbitrary element of $\Phi_{s,a}$, this shows that $(Tev_a)(\Pi_s)$ is an T-redistribution of $\Phi_{s,a}$. Since this applies to all $a \in A$, we have finished the proof of (5.3).

Now we are ready for the definition of the labeled tree $\mathbb{X}$. The nodes of $\mathbb{X}$ will be taken from the set $S \times Rel(A) \times \omega$, and the labellings $u : X \to S$ and $Q : X \to Rel(A)$ are simply given as the first and second projection map. That is, $u_{(s,R,i)} := s$ and $Q_{(s,R,i)} := R$. The third projection function provides the height of the node. For the definition of the labeling $\Pi_x \in TRel(A)$, we simply look at the node u_x; that is, we put $\Pi_x := \Pi_{u_x}$. Hence, it remains to define the tree structure of $\mathbb{X}$, and this we will do by the following induction on the height of the nodes. For the *root* of $\mathbb{X}$ we take the triple $(s_I, \{(a_I, a_I)\}, 0)$, while the successor map ξ is given by

$$\xi(s, R, i) := Y_s \times \{i + 1\},$$

that is, $\xi(s, R, i)$ consists of those triples $(t, P, i+1)$ such that (t, P) belongs to the set Y_s.

Now that we have completely defined the structure $\mathbb{X}$, let us check that it satisfies the conditions (3a-3d). To start with, the condition (3a) has been directly cooked into the definition of $\mathbb{X}$, while (3b) immediately follows from (5.3). For condition (3c), it follows by a straightforward unraveling of the definitions that $W_x = Y_{u_x}$ for all $x \in X$. Since the definition of Π_s implies that $(\sigma(s), \Pi_s) \in \overline{T}Y_s$ for all $s \in S$, this directly gives (3c). Finally, for the last condition on $\mathbb{X}$, we may infer from the identity of W_x and Y_{u_x}, using the definitions of Y_s and $Y_{s,a}$, that any trace of $\mathbb{X}$ corresponds to a match of $\mathcal{G}$ in which $\exists$ plays her strategy (Φ, Z). Condition (3d) is then an immediate consequence of the fact that this strategy was supposed to be *winning* for $\exists$.

(3) $\Rightarrow$ (1)
Suppose that $\mathbb{X}$ is as described in (3). We have to prove that Φ can be extended to a winning strategy for $\exists$ in $\mathcal{G}$. Basically, what we will do is show that $\exists$ can 'keep the match on $\mathbb{X}$', in the following sense. Her strategy will ensure, for any match in which she plays this strategy, the existence of a branch $x_0 x_1 \dots$ (possibly finite) of $\mathbb{X}$ such that for each partial play $(s_0, a_0) \dots (s_k, a_k)$ of the match it holds that $s_i = u_{x_i}$ and $a_{i+1} \in Q_{x_{i+1}}[a_i]$ for all $i < k$. Since $\mathbb{X}$ contains no bad traces by (3d), this guarantees that she wins all infinite matches of the game. Hence, it suffices to prove that at any finite stage k of such a match, she either wins immediately, or else she can keep the above condition for one more round.

Suppose then that $\exists$ have been able to keep this condition for k steps, arriving at position $(s, a) = (s_k, a_k)$, with $s = u_x$. The first thing to note is that condition (3d) implies that $\Delta(a) \neq \emptyset$, since $a_0 \dots a_k$ is a trace of $\mathbb{X}$. But from $\Delta(a) \neq \emptyset$ it follows that $\exists$ may legitimately move $\Phi_{s,a} \in \Delta(a)$. If $\Phi_{s,a} = \emptyset$ then $\exists$ wins immediately, so suppose otherwise. Let $\forall$'s answer be $\phi \in \Phi_{s,a}$, then $\exists$ has to respond with a relation $Z \subseteq S \times A$ such that $(\sigma(s), \phi) \in \overline{T}Z$. Our suggestion to $\exists$ is to pick the relation $Z \subseteq S \times A$ given by

$$Z := W_x \circ \mathrm{Gr}(ev_a) \circ \ni_A .$$

If this is a legitimate move for $\exists$, then we are done. For, distinguish the following cases. If $Z = \emptyset$ then $\forall$ gets stuck so $\exists$ wins immediately. But if $Z \neq \emptyset$ then with any

$(t, b) \in Z$ that $\forall$ chooses as his next move we may associate, by definition of Z, a node $y \in \xi(x)$ such that $t = u_y$ and $(a, b) \in Q_y$. In other words, she either wins immediately or indeed manages to keep the match on $\mathbb{X}$ for one more round of the game.

Thus it is left to show that Z is a legal move for $\exists$ in $\mathcal{G}$; that is, we must show that

$$(\sigma(s), \phi) \in \overline{T}Z. \tag{5.5}$$

For this purpose, first observe that the definition of Z and the properties of T and $\overline{T}$ imply that

$$\overline{T}Z = \overline{T}W_x \circ \mathrm{Gr}(Tev_a) \circ \overline{T}(\ni_A). \tag{5.6}$$

Now it follows from property (3c) of $\mathbb{X}$ that $(\sigma(s), \Pi_x)$ belongs to $\overline{T}W_x$, and from property (3b), that $(Tev_a)(\Pi_x)$ is a redistribution of $\Phi_{s,a}$. In particular, it holds that $((Tev_a)(\Pi_x), \phi) \in \overline{T}(\ni_A)$, so that $(\Pi_x, \phi) \in \mathrm{Gr}(Tev_a) \circ \overline{T}(\ni_A)$. From this and (5.6) we may infer (5.5).

(2) $\Rightarrow$ (3)
Assume that f is a winning strategy for $\exists$ in the game $\mathcal{G}'$. We will define a labeled tree $\mathbb{X}$ that intuitively corresponds to the game tree of $\mathcal{G}'$, pruned according to this strategy f.

Formally, elements of X will be taken from the set $S \times M \times Rel(A) \times \omega$. This enables a very straightforward definition of the labeling functions u, Q and Π of $\mathbb{X}$. The state u_x in $\mathbb{S}$, and the relation $Q_x \in Rel(A)$ that we associate with a node x, are simply given by the first and third projection functions, respectively. That is, we put $u_{(s,m,R,i)} := s$ and $Q_{(s,m,R,i)} := R$.

We now turn to the definition of the tree structure of $\mathbb{X}$. As the root of the tree we take the element $(s_I, m_I, \{(a_I, a_I)\}, 0)$. For the definition of the successor map ξ, consider an arbitrary point $x = (s, m, R, i) \in X$. Inductively assume that (s, m) is a winning position for $\exists$ in the game $\mathcal{G}$. Suppose that $K \in TM$ and $Z \subseteq S \times M$ are locally the basis of $\exists$'s strategy f, that is, suppose that $f(s, m) = (s, \{K\})$ and $f(s, K) = Z$. Given our comments in Remark 5.2.17, we may assume the existence of an object $\Pi_{s,m} \in TRel(A)$ and a relation $Y_{s,m} \subseteq S \times Rel(A)$ such that $\Pi_{s,m} \in \mathcal{R}_T(\Phi_s)$ and $(\sigma(s), \Pi_{s,m}) \in \overline{T}Y_{s,m}$ and

$$\begin{aligned} K &= (T\mu_m)(\Pi_{s,m}) \\ Z &= \{(t, \mu_m(P)) \mid (t, P) \in Y_{s,m}\}. \end{aligned}$$

Now inductively we define

$$\begin{aligned} \xi(s, m, R, i) &:= \{(t, n, P, i + 1) \mid (t, P) \in Y_{s,m} \\ &\qquad\qquad \text{and } n = \mu_m(P)\}, \\ \Pi_{(s,m,R,i)} &:= \Pi_{s,m}. \end{aligned}$$

Clearly then, every element of $\xi(s)$ encodes a response of $\forall$ to $\exists$'s first two f-moves from position (s, m). This encoding is very direct: $(t, n, P, i + 1) \in \xi(s, m, R, i)$ corresponds to the move (t, n). It is then immediate that such (t, n) are again winning positions for $\exists$, so that the required inductive condition is kept. This finishes the definition of $\mathbb{X}$.

It is a straightforward exercise to check that $\mathbb{X}$, thus defined, satisfies the conditions (3a-3c). It should also require only little reflection to see that with each *branch* $x_0x_1\ldots$ of $\mathbb{X}$ we may associate a match of $\mathcal{G}'$ in which $\exists$ plays her winning strategy f. As we will see now, this is the key for proving that $\mathbb{X}$ satisfies (3d).

For the technical details, suppose, for contradiction, that $\mathbb{X}$ contains a bad trace α. We will only consider the case that α is finite, say, $\alpha = a_1\ldots a_k$, and bad because $\Delta(a_k) = \emptyset$. (The other kind of bad traces, which only applies to infinite traces, is treated in a very similar way.) Let $x_0\ldots x_k$ be the branch of $\mathbb{X}$ that contains α. Writing $x_i = (s_i, m_i, R_i, i)$ for all $i \leq k$, and unraveling the definitions, we find that this means that $a_{i+1} \in R_{i+1}[a_i]$ for all $i < k$, and that $m_0m_1\ldots m_k$ is the run of $\mathbb{M}_0$ on the finite $Rel(A)$-word $R_0R_1\ldots R_k$. Then it follows from Proposition 5.2.14 that $m_k = m_\forall$.

But, given the remark following the definition of ξ, we may equip $\forall$ with a strategy in $\mathcal{G}'$ which, when played against $\exists$'s f, ensures that the match passes through the basic positions $(s_0, m_0), \ldots, (s_k, m_k)$. By definition of $\mathbb{M}_1$, it follows from $m_k = m_\forall$ that $\mu(m_k, \Phi_{s_k}) = \emptyset$. Hence, $\exists$ gets stuck at (s_k, m_k) and immediately looses the game. Thus we arrive at the desired contradiction.

$(3) \Rightarrow (2)$
Let $\mathbb{X}$ be a labeled tree as specified in (3); we need to define a winning strategy for $\exists$ in $\mathcal{G}'$. The basic idea underlying this strategy will be that that the resulting match corresponds to a branch of $\mathbb{X}$.

More precisely, we will show that $\exists$ can maintain the condition that with every match (*) there exists a branch $x_0x_1\ldots$ (possibly finite) on $\mathbb{X}$ such that $u_{x_i} = s_i$ at every stage i of the match, and $m_{i+1} = \mu_0(m_i, Q_{x_{i+1}})$ (at least, in case the match reaches the stage $i+1$).

$\exists$ wins any match $(s_0, m_0)(s_1, m_1)\ldots$ in which she can maintain this condition infinitely long. To see why this is so, note that the infinite sequence $m_0m_1\ldots$ corresponds to the run of the automaton $\mathbb{M}_0$ on the infinite word $Q_{x_1}Q_{x_2}\ldots$. Now by condition (3d) on $\mathbb{X}$ this word contains no bad trace, so by Proposition 5.2.14 it is accepted by $\mathbb{M}_0$. But then the sequence $m_0m_1\ldots$ meets the acceptance condition Ω_0 of $\mathbb{M}_0$, which coincides with the acceptance condition of $\mathbb{M}_1$. It is then immediate from the definitions that the $\mathcal{G}'$-match $(s_0, m_0)(s_1, m_1)\ldots$ is won by $\exists$.

Hence it suffices to show that at any finite stage of a match satisfying the above condition, $\exists$ can either win the game in the current round, or else prolong the condition (*) for one more round. In order to show that this is the case, suppose that the match has reached position (s_k, m_k), and that play so far meets the above mentioned condition (*), relative to the branch $x_0x_1\ldots x_k$. In order to provide $\exists$ with a move to make at position (s_k, m_k), let us consider her set $\mu(m_k, \Phi_{s_k})$ of alternatives.

Since $\mathbb{X}$ contains no bad traces, it follows from Proposition 5.2.14 that $m_i \neq m_\forall$ for all $i \leq k$. Thus $\mu(m_k, \Phi_{s_k})$ is given as the set of all singletons $\{(T\nu)(\Pi)\}$ such that $\Pi \in TRel(A)$ is an T-relation for Φ_{s_k}, where $\nu : Rel(A) \to M$ is the map given by $\nu(R) := \mu_0(m_k, R)$. In particular, $\mu(m_k, \Phi_{s_k})$ is not empty (which would mean an immediate loss for $\exists$), since condition (3b), stating that $\Pi_{x_k} \in \mathcal{R}_T(\Phi_{s_k})$, guarantees that it contains $K := (T\nu)(\Pi_{s_k})$. Thus we may set $(s_k, \{K\})$ as a legitimate move for $\exists$ at position (s_k, m_k). After this, $\forall$ has no choice but to pick (s_k, K). Then let $\exists$ continue by moving $Z = \{(t, \nu(R)) \mid (t, R) \in W_{x_k}\}$, that is, $Z = W_{x_k} \circ \mathrm{Gr}(\nu)$. To verify the legitimacy of this move, note that $(\sigma(s_k), \Pi_{x_k}) \in \overline{T}W_{x_k}$ by (3c), so from $K = (T\nu)(\Pi_{x_k})$ it follows that $(\sigma(s_k), K) \in \overline{T}W_{x_k} \circ \mathrm{Gr}(T\nu) = \overline{T}Z$, as required.

Finally, suppose that $\forall$ responds by choosing some pair $(s_{k+1}, m_{k+1}) \in Z$ — the case that $Z = \emptyset$ is an immediate win for $\exists$. By definition of Z, there must be some $y \in \xi(x_k)$ such that $(y, Q_y) \in W_{x_k}$ and $m_{k+1} = \mu_0(m_k, Q_y) = \nu(Q_y)$. Hence by taking $x_{k+1} := y$ as the associated node of the new position (s_{k+1}, m_{k+1}) in the match, $\exists$ has maintained the condition (*) one more round, as required. QED

In the final step of the construction we have to transform $\mathbb{M}_1$ into a non-deterministic T-automaton $\mathbb{A}^d$ that is equivalent to $\mathbb{A}$. This last transformation is in fact easy — relatively that is: we need an application of the closure under projection of non-deterministically recognizable languages.

5.2.19. Definition. Let $\mathbb{A}^d$ be the T-automaton $(M, m_I, \mu^d, \Omega_0)$ where M, m_I and Ω_0 are as in Definition 5.2.16, while

$$\mu^d(m) := \bigcup_{e \in (\mathcal{P}TA)^A} \mu(m, e)$$

defines the transition map $\mu^d : M \to \mathcal{P}TM$.

It is easy to check that $\mathbb{A}^d$ is indeed non-deterministic, so clearly, the following proposition, which is a straightforward corollary of the Propositions 5.2.11 and 5.2.18, suffices to prove Theorem 5.2.1.

5.2.20. Proposition. *The automata $\mathbb{A}$ and $\mathbb{A}^d$ accept exactly the same rooted T-coalgebras.*

Proof. Let (S, σ, s) be a pointed T-coalgebra. We will show that $\mathbb{A}$ accepts (S, σ, s) if and only if $\mathbb{A}^d$ does. It is an immediate consequence of Proposition 5.2.11 that (S, σ, s) is accepted by $\mathbb{A}^d$ if and only if there is a pointed T-coalgebra (S', σ', s') which is bisimilar to (S, σ, s) and admits a $(\mathcal{P}TA)^A$-coloring Φ such that (S', σ', Φ, s') is accepted by $\mathbb{M}_1$. Hence, invoking Proposition 5.2.18 we see that (S, σ, s) is accepted by $\mathbb{A}^d$ if and only if it is bisimilar to a pointed T-coalgebra that is accepted by $\mathbb{A}$. The Proposition is then an immediate consequence of the fact that the class of coalgebras recognized by an T-automaton is closed under bisimilarity (cf. Fact 5.1.14). QED

A short remark concerning the complexity of our construction seems in order.

5.2.21. Remark. Although we do not go into the algorithmic details of our construction, we want to stress here that complexity theoretically, our results match known results in automata theory. If we define the size of an automaton as its number of states, the main observation is that the size of $\mathbb{A}^d$ is equal to the size of $\mathbb{M}_0$ and, in particular, does not depend on the functor T. In fact, combining well known results about word automata, one may show that basically, the size of $\mathbb{M}_0$ is exponential in the size of $\mathbb{A}$.

5.2.3 Non-emptiness of coalgebra automata

An important problem concerning automata on infinite objects is the so-called non-emptiness problem: Can we decide whether the language accepted by a given automaton is empty or not? Given the results from [Ven04] the non-emptiness problem for T-automata is equivalent to the question whether we can decide the satisfiability of a given formula of the coalgebraic fixed-point logic from Section 2.3. This will be made clear in Section 5.3 below.

Now we are going to give a partial solution for the non-emptiness problem of T-automata: If a T-automaton accepts some rooted T-coalgebra then it also accepts a finite one. For non-deterministic T-automata we can do better as the following theorem which is the main result of this section demonstrates.

5.2.22. Theorem. *Let $\mathbb{A} = (A, a_I, \Delta, \Omega)$ be a non-deterministic T-automaton. Then $\mathbb{A}$ accepts some rooted T-coalgebra $\mathbb{S} = (S, \sigma, s_I)$ iff $\mathbb{A}$ accepts a rooted T-coalgebra $\mathbb{S}' = (S', \sigma, s_I)$ with $|S'| \leq |A|$.*

Proof. Let $\mathbb{A}$ be a T-automaton and $(\mathbb{S}, s_I)$ a rooted T-coalgebra accepted by $\mathbb{A}$. Then w.l.o.g. $\exists$ has a scattered history-free winning strategy $(\Phi : S \times A \to TA, Z : S \times TA \to \mathcal{P}(S \times A))$ from position (s_I, a_I) in the acceptance game $\mathcal{G} := \mathcal{G}(\mathbb{S}, \mathbb{A})$ of the automaton.

The idea of the proof of the theorem is as follows: The coalgebra $\mathbb{S}'$ will be based on A - the carrier set of the automaton. Our aim is to show that for any a in A we can find some $\phi_a \in TA$ such that the T-coalgebra $(A, \phi : A \to TA, a_I)$ is accepted by $\mathbb{A}$ (again we use the notation of writing ϕ_a instead of $\phi(a)$). For the definition of ϕ we would like to use $\exists$'s winning strategy Φ by letting $\phi_a := \Phi_{(s,a)}$ for some $(s, a) \in \mathrm{Win}(\Phi, Z)$. The obvious problem with this approach lies in the fact that Φ depends on both s and a, i.e. there might be two different positions $(s_1, a), (s_2, a) \in \mathrm{Win}(\Phi, Z)$ such that $\Phi_{(s_1,a)} \neq \Phi_{(s_2,a)}$ and in this case it is unclear which of these values of Φ we should choose as the value for ϕ_a. This problem is overcome by making use once more of the history-free determinacy of parity games.

For $a \in \mathbb{A}$ we define $\mathbb{A}^a$ to be the automaton (A, a, Δ, Ω), in other words $\mathbb{A}^a$ is the automaton that is obtained from $\mathbb{A}$ by changing the initial state from a_I into a. Furthermore recall that $\mathrm{Win}(\Phi, Z)$ denote the set of positions of $\mathcal{G}$ for which (Φ, Z) is a winning strategy for $\exists$ and define

$$\mathrm{WRel}(\Phi, Z) := \{R \subseteq S \times A \mid \exists (s, \phi) \in \mathrm{Win}(\Phi, Z) \, . \, R = Z_{s,\phi}\},$$

i.e. $\mathrm{WRel}(\Phi, Z)$ is the set of potential "relational" moves of $\exists$ in a play of $\mathcal{G}$ according to her winning strategy. The *non-emptiness game* $\mathcal{G}'$ of $\mathbb{A}$ is given by the following table:

Position: b	Player	Admissible moves: $E[b]$	$\Omega'(b)$
$\{a \in A \mid \exists(s,a) \in \mathrm{Win}(\Phi, Z)\}$	$\exists$	$\Delta(a)$	$\Omega(a)$
$\phi \in TA$	$\exists$	$\mathrm{WRel}(\Phi, Z)$	0
$Z \in \mathcal{P}\mathrm{Win}(\Phi, Z)$	$\forall$	$\{a \in A \mid \exists s \,.\, (s,a) \in Z\}$	0

It is not difficult to see that (Φ, Z) can be used to construct a winning strategy for $\exists$ in $\mathcal{G}'$ from position a_I that is not history-free. To this aim we construct for each play of $\mathcal{G}'$ that starts in a_I a shadow play of $\mathcal{G}$ starting in (s_I, a_I) that is conform with (Φ, Z). During this construction we want to maintain the following condition $(*)$: if a position $a \in A$ is reached after n rounds of the play of $\mathcal{G}'$ then there exists a $s \in S$ such that the position $(s, a) \in S \times A$ is reached after n rounds of the corresponding shadow play of $\mathcal{G}$. Then infinite plays of $\mathcal{G}'$ are won by $\exists$ because the corresponding shadow plays of $\mathcal{G}$ are conform $\exists$'s winning strategy. All finite plays of $\mathcal{G}'$ are trivially won by $\exists$ because it is easy to see that $\exists$ cannot get stuck in such a play. As a result we get that $\exists$ wins all $\mathcal{G}'$-plays for which the condition $(*)$ is maintained.

Let us see now how $\exists$ is able to maintain condition $(*)$ during a $\mathcal{G}'$-play. Suppose that we are at position a in a play of $\mathcal{G}'$ that started in a_I and that we are in the shadow play of $\mathcal{G}$ in some position (s, a). Furthermore we assume that the (partial) play of $\mathcal{G}$ up to now was conform $\exists$'s winning strategy. Then $\exists$'s strategy in $\mathcal{G}$ would be to move to $\Phi_{s,a} \in TA$ and then to some relation $Z_{s,\Phi_{s,a}} \subseteq \mathrm{Win}(\Phi, Z)$ (as (Φ, Z) is a winning strategy we may assume that this relation only contains pairs $(s', a') \in \mathrm{Win}(\Phi, Z)$). We suggest $\exists$ to use this strategy also in $\mathcal{G}'$, i.e. $\exists$ moves in $\mathcal{G}'$ from a to $\Phi_{s,a}$ and then further to $Z_{s,\Phi_{s,a}}$. Then $\forall$ answers in $\mathcal{G}'$ by moving to some $a' \in A$ for which there is an $s' \in S$ with $(s', a') \in Z_{s,\Phi_{s,a}}$. This move of $\forall$ can be easily reflected in the shadow game: $\forall$'s corresponding choice in $\mathcal{G}$ would be to move to (s', a'). Hence $\exists$ can maintain condition $(*)$ and as a consequence win an arbitrary play of $\mathcal{G}'$ starting from a_I. Note, however, that the winning strategy of $\exists$, that we described, is not history-free, as she has to remember in every position of a play what her position in the shadow play would be.

But the history-free determinacy of parity games together with the fact that $\exists$ has a winning strategy in $\mathcal{G}'$ from position a_I implies that she also has a history-free winning strategy in $\mathcal{G}'$ from position a_I which can be encoded as a pair of functions

$$(\phi : A \to TA, Y : TA \to \mathrm{WRel}(\Phi, Z))\,.$$

We will show now that $\mathbb{A}$ accepts the rooted T-coalgebra given as $(\mathbb{A}(\phi), a_I) := (A, \phi, a_I)$. This will finish the proof of the theorem.

In order to prove that $(\mathbb{A}(\phi), a_I)$ is accepted by $\mathbb{A}$ we show that $\exists$ has a winning strategy in $\mathcal{G}_\phi := \mathcal{G}(\mathbb{A}(\phi), \mathbb{A})$ from position (a_I, a_I). More precisely we show that the pair of functions (Φ', Z') defined by

$$\begin{aligned} \Phi'_{a',a} &:= \phi_a \quad \text{for } (a', a) \in A \times A \\ Z'_{a,\phi} &:= Y_\phi^{\smile} \circ Y_\phi \quad \text{for } (a, \phi) \in A \times TA \end{aligned}$$

encodes a winning strategy for $\exists$ in $\mathcal{G}_\phi$ from position (a_I, a_I). This is done by showing that

1. all basic positions in a $\mathcal{G}_\phi$-play conform (Φ', Z') are of the form (a', a') for some $a' \in A$ and
2. for every partial $\mathcal{G}_\phi$ play $(a_I, a_I)(a_1, a_1)(a_2, a_2)\ldots(a_n, a_n)$ conform (Φ', Z') there is a partial $\mathcal{G}'$-play $a_I \ldots a_n$ that is conform (ϕ, Y).

Suppose we are at position (a, a) in an arbitrary play of $\mathcal{G}_\phi$ that started in position (a_I, a_I). Then $\exists$ moves to $\Phi'_{a,a} = \phi_a$ and further to $Z'_{a,\phi_a} = Y_\phi^{\smile} \circ Y_\phi$. In the corresponding $\mathcal{G}'$-play at position a she can move to ϕ and then further to Y_ϕ. These moves are conform her $\mathcal{G}'$-winning strategy (ϕ, Y). Now it is $\forall$'s turn to pick some element $(a_1, a_2) \in Y_\phi^{\smile} \circ Y_\phi$. As we assumed in the beginning of the proof that (Φ, Z) was a *scattered* winning strategy and $Y_\phi \in \mathrm{WRel}(\Phi, Z)$ it follows that $a_1 = a_2$. Obviously $\forall$ can also move in the corresponding $\mathcal{G}'$-game to a_1. Therefore the above conditions 1 and 2 can be maintained round-by-round. As (ϕ, Y) is winning for $\exists$ in $\mathcal{G}'$ from position a_I, we can conclude that (Φ', Z') is winning for $\exists$ in game $\mathcal{G}_\phi$ from position (a_I, a_I). QED

Using our previous result about alternation we get the following.

5.2.23. Corollary. *Let $T : \mathsf{Set} \to \mathsf{Set}$ be a standard weak pullback preserving functor and $\mathbb{A}$ a T-automaton. Then $\mathbb{A}$ accepts some rooted T-coalgebra $(\mathbb{S}, s_I)$ iff $\mathbb{A}$ accepts some rooted T-coalgebra with finite carrier set.*

Proof. Given a T-automaton $\mathbb{A}$ we can transform it to an equivalent non-deterministic T-automaton $\mathbb{A}^d$ (cf. Theorem 5.2.1) and then apply Theorem 5.2.22. QED

5.2.24. Remark. For concrete examples this is a well-known result: for example for automata working on infinite trees it means that every tree automaton that accepts some tree accepts at least one *regular* tree.

5.2.4 A remark about standardness

For T-automata operating on finite (algebraic) structures it is known that the condition that the functor $T : \mathsf{Set} \to \mathsf{Set}$ preserves weak pullbacks is necessary to be able to transform an alternating automaton into an equivalent non-deterministic one (cf. [AT90, Thm. VII.2.12]). This suggests that our assumption on the functor T being weak pullback preserving is necessary to prove the above closure properties. Our second assumption that the functor also should be standard, however, can safely be dropped as the following theorem, which is a special case of a theorem by Trnková, shows.

5.2.25. THEOREM. *Let S : Set → Set be a weak pullback preserving functor. Then there is a functor T : Set → Set such that*

1. *T is weak pullback preserving and standard and*
2. *there is a natural isomorphism $\tau : S \Rightarrow T$.*

Proof. A proof of this theorem can be found in [AT90, Theorem III.4.5]. As mentioned in the appendix our notion of standardness differs slightly from the one in *loc.cit.* which we call $\emptyset$-standardness (cf. Def. A.2.11). Because we are working with the additional assumption that the functors are weak pullback preserving, however, we know that the two notions of standardness coincide (cf. Lemma A.2.12) QED

5.2.26. LEMMA AND DEFINITION. *Let S : Set → Set and T : Set → Set be functors such that there is a natural isomorphism $\tau : S \Rightarrow T$ between them. Then the functor*

$$\begin{array}{rcl} G_\tau : \mathsf{Coalg}(S) & \to & \mathsf{Coalg}(T) \\ (S, \sigma) & \mapsto & (S, \tau_X \circ \sigma) \\ f & \mapsto & f \end{array}$$

is a category isomorphism.

Proof. Obviously G_τ is well-defined on objects. The well-definedness on morphism is a consequence of the naturaliy of τ. It is easy to check that this functor is a category isomorphism with the functor $G_{\tau^{-1}} : \mathsf{Coalg}(T) \to \mathsf{Coalg}(S)$ as its inverse. QED

Instead of defining S-automata for a weak pullback preserving but not necessarily standard functor we can use the automata for the standardized version of the functor for classifying S-coalgebras.

5.2.27. DEFINITION. Let S : Set → Set a weak pullback preserving functor, T : Set → Set a standardization of S, i.e. S is standard and there is a natural isomorphism $\tau : S \Rightarrow T$ between S and T. Furthermore let $\mathbb{A}$ be a T-automaton. Then we say $\mathbb{A}$ accepts a rooted S-coalgebra (X, γ, x_I) if $\mathbb{A}$ accepts $G_\tau(X, \gamma, x_I) := (X, \tau_X \circ \gamma, x_I)$.

5.2.28. EXAMPLE. An important example for a weak pullback preserving functor that is not standard is the filter functor $\mathcal{F}$: Set → Set which has been studied in [Gum01].

Recall that for an arbitrary set X a nonempty set $F \subseteq \mathcal{P}(X)$ is called a *filter over X* if

1. F is closed under finite intersections, i.e. $U \in F$ and $V \in F$ imply $U \cap V \in F$, and
2. F is upwards closed, i.e. $U \in F$ and $U \subseteq V$ imply $V \in F$.

The filter functor is defined as follows

$$\begin{array}{rcl} \mathcal{F} : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & \{F \subseteq \mathcal{P}(X) \mid F \text{ is a filter over } X\} \\ (f : X \to Y) & \mapsto & (\mathcal{F} f : \mathcal{F} X \to \mathcal{F} Y), \end{array}$$

where for a filter $F \subseteq \mathcal{P}(X)$ we have

$$(\mathcal{F} f)(X) := \{V \in \mathcal{P}(Y) \mid \exists U \in F \;.\; f[U] \subseteq V\}.$$

The definition of a standardization $\mathcal{F}'$ of $\mathcal{F}$ needs some preparation. We say that a filter F_1 over some set X_1 and a filter F_2 over some set X_2 are *equivalent* if

- for all $U \in F_1$ there is a $U' \in F_2$ such that $U \subseteq U'$, and vice versa
- for all $U' \in F_2$ there is a $U \in F_1$ such that $U' \subseteq U$.

In this case we write $F_1 \sim F_2$. Clearly $\sim$ is an equivalence relation on the class of all filters and we denote the equivalence class of a filter F over some set X by

$$[F] := \{F' \mid F' \text{ is a filter over some set } Y \text{ and } F' \sim F\}.$$

We are now ready to define the standard functor $\mathcal{F}'$:

$$\begin{array}{rcl} \mathcal{F}' : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & \{[F] \mid F \in \mathcal{F} X\} \\ (f : X \to Y) & \mapsto & (\mathcal{F}' f : \mathcal{F}' X \to \mathcal{F}' Y), \end{array}$$

where $(\mathcal{F}' f)([F]) := [(\mathcal{F} f)(F)]$ for $[F] \in \mathcal{F}' X$.

5.3 The connection with coalgebraic fixed-point logic

In the previous sections of this chapter we recalled the definition of coalgebra automata and investigated several of their properties. But how does this material fit into the framework of this thesis, i.e. where is the connection to finitary logics for coalgebra? The answer to this question was already given in detail by Venema in [Ven04], where the following two theorems are proven.

5.3.1. THEOREM. *[Ven04, Theorem 2] Let T be a standard, weak pullback preserving endofunctor on* Set. *Then every sentence $\phi \in \mu\mathcal{L}^T$ of coalgebraic fixed-point logic can be transformed into a T-automaton $\mathbb{A}_\phi$ such that for any rooted T-coalgebra $(\mathbb{S}, s_I)$:*

$$\mathbb{S}, s_I \models \phi \textit{ iff } \mathbb{A}_\phi \textit{ accepts } (\mathbb{S}, s_I).$$

5.3.2. THEOREM. *[Ven04, Theorem 3] Let T be a standard, weak pullback preserving endofunctor on* Set*. Then any T-automaton $\mathbb{A}$ can be transformed into a $\mu\mathcal{L}^T$-sentence $\phi_{\mathbb{A}}$ such that for any rooted T-coalgebra $(\mathbb{S}, s_I)$:*

$$\mathbb{A} \textit{ accepts } (\mathbb{S}, s_I) \textit{ iff } \mathbb{S}, s_I \models \phi_A.$$

In other words sentences of coalgebraic fixed-point logic and T-automata are essentially the same. This is a generalization of classical results from automata theory in the same way as T-automata generalise automata on infinite words, trees and graphs. The following table lists those classical correspondences between automata on certain structures on the one hand and formulas of a suitable logic on the other hand.

Automata on (possibly) infinite	formulas of
words	**S1S**
k-ary trees	**SkS**
graphs	the modal μ-calculus

5.3.1 Finite model property

An important application of the correspondence between automata and formulas is the reduction of the satisfiability problem of a logic to the non-emptiness problem of the corresponding automata.

For coalgebra automata this is done using Theorem 5.3.1: a given formula $\phi \in \mu\mathcal{L}^T$ is satisfiable in some T-coalgebra iff the automaton $\mathbb{A}_\phi$ accepts some rooted T-coalgebra.

In the previous section we saw that every T-automaton that accepts some rooted T-coalgebra, accepts at least one rooted T-coalgebra with finite carrier set. This gives us immediately the following corollary - a weak finite model property of coalgebraic fixed-point logic.

5.3.3. COROLLARY. *Let T : Set $\to$ Set be a standard, weak pullback preserving functor and $\phi \in \mu\mathcal{L}^T$ a sentence. Then ϕ is satisfiable in some rooted T-coalgebra $(\mathbb{S}, s_I)$ if ϕ is satisfiable in some rooted T-coalgebra $(\mathbb{S}', s_I')$ with a finite carrier set.*

Proof. This is a consequence of Theorem 5.3.1 and Corollary 5.2.23. QED

Decidability of coalgebraic fixed-point logic does, however, not follow from this finite model property. However, under the requirement that the functor T maps finite sets to finite sets we can easily obtain decidability.

5.3.4. COROLLARY. *Let T : Set $\to$ Set be a standard and weak pullback preserving functor that maps finite sets to finite sets. Then the problem whether a given sentence $\phi \in \mu\mathcal{L}^T$ is satisfiable in some rooted T-coalgebra $(\mathbb{S}, s_I)$ is decidable.*

Proof. Let $\phi \in \mu\mathcal{L}^T$ and $\mathbb{A}_\phi$ the corresponding T-automaton from Theorem 5.3.1. Then we first transform $\mathbb{A}_\phi$ into an equivalent non-deterministic automaton $\mathbb{A}_\phi^d = (A, a_I, \Delta, \Omega)$ as described in Section 5.2.2 above. A careful inspection of the proof of Theorem 5.2.22 shows that $\mathbb{A}_\phi^d$ accepts some rooted T-coalgebra $(\mathbb{S}, s_I)$ only if it accepts some rooted T-coalgebra (A, a_I, ϕ) with the property that $\phi(a) \in \Delta(a)$ for all $a \in A$. It is not difficult to see that there are only finitely many such T-coalgebras, because A and therefore also TA are finite sets. QED

5.3.2 A distributive law

The finite model property of coalgebraic fixed-point logic was an immediate consequence from our work on closure properties of T-automata. The material presented in this subsection is motivated by the work on T-automata, but we do not explicitly make use of the automata-theoretic framework.

Let us take a closer look at the precise connection between formulas in $\mu\mathcal{L}^T$ and T-automata: According to Theorem 5.3.2 we can assign to every automaton $\mathbb{A} = (A, a_I, \Delta, \Omega)$ an equivalent formula $\phi_\mathbb{A}$. For every $a \in A$ let us denote by $\mathbb{A}^a$ the automaton that we can obtain from $\mathbb{A}$ by changing the initial state a_I into a. Then we get for each $a \in A$ a corresponding formula $\phi_{\mathbb{A}^a}$. In this way we can define a function

$$\begin{array}{rcl} t : A & \to & \mu\mathcal{L}^T \\ a & \mapsto & \phi_{\mathbb{A}^a}. \end{array}$$

The following proposition is then not difficult to prove.

5.3.5. Proposition. *Let T : $\mathsf{Set} \to \mathsf{Set}$ be a standard, weak pullback preserving functor, $\mathbb{A} = (A, a_I, \Delta, \Omega)$ a T-automaton and $(\mathbb{S}, s_I)$ a rooted T-coalgebra. Then $\mathbb{A}$ accepts $(\mathbb{S}, s_I)$ iff*

$$\mathbb{S}, s_I \models \bigvee_{\Phi \in \Delta(a_I)} \bigwedge_{\phi \in \Phi} \nabla(Tt)(\phi).$$

This means that an alternating automaton accepts a rooted T-coalgebra $(\mathbb{S}, s_I)$ iff $(\mathbb{S}, s_I)$ satisfies a disjunction of conjunctions of ∇-formulas. In this perspective turning an alternating T-automaton into an equivalent non-deterministic one is equivalent to replacing a formula of the form $\bigvee \bigwedge \nabla\phi$ by an equivalent formula which is only a disjunction of ∇-formulas, i.e. into a formula of the form $\bigvee \nabla\psi$.

We know from Section 5.2.2 that we can always turn an alternating automaton into an equivalent nondeterministic one. This suggests that we can find to every formula of the form $\bigvee \bigwedge \nabla\phi$ an equivalent formula of type $\bigvee \nabla\psi$. In the remainder of the section we are going to prove that this is indeed the case. As a consequence we obtain a new validity of coalgebraic (fixed-point) logic. Again redistributions play an important role in the argument, but because we are working with finitary syntax we have to restrict ourselves to *finitary redistributions*. In the following we fix some (arbitrary) standard functor T : $\mathsf{Set} \to \mathsf{Set}$, that preserves weak pullbacks.

5.3.6. Definition. Let $\mathcal{X} \in \mathcal{P}TX$ be a subset of TX. Then the set $\mathcal{D}_\omega(\mathcal{X})$ of *finitary T-redistributions of $\mathcal{X}$* is defined as

$$\mathcal{D}_\omega(\mathcal{X}) := \{\Theta \in T_\omega(\mathcal{P}X) \mid \Theta \text{ is a } T\text{-redistribution of } \mathcal{X}\}.$$

where $T_\omega(\mathcal{P}X)$ consists of the elements of $T\mathcal{P}X$ that have finite T-base (cf. Def. 2.3.1 on page 26).

It will be convenient to view the conjunction of formulas as a relation in the following way.

5.3.7. Definition. Let $\mathcal{L} \subseteq \mu\mathcal{L}^T$ be a set of formulas. Then we define a relation $\sqcap_{\mathcal{L}} \subseteq \mathcal{P}\mathcal{L} \times \mathcal{L}$ by letting

$$(\Phi, \varphi) \in \sqcap_{\mathcal{L}} \quad \text{if} \quad \varphi = \bigwedge \Phi.$$

5.3.8. Remark. Obviously the relation $\sqcap_{\mathcal{L}}$ and therefore also its lifting $\overline{T}\sqcap_{\mathcal{L}}$ will be the graph of a possibly partial function. This allows in the following to identify the sets of successors $\sqcap_{\mathcal{L}}[\Phi]$ and $(\overline{T}\sqcap_{\mathcal{L}})[\Xi]$ with their only element.

In order to be able to state the main result of this section we need the following technical lemma. Expressed in words, it says that conjunction is well behaved with respect to the notion of finite T-base: consider a $\Xi \in T(\mathcal{P}_\omega(\mu\mathcal{L}^T))$ with finite T-base. Then its conjunction $\xi := (\overline{T}\sqcap_{\mathcal{L}})[\Xi]$ also has finite T-base and therefore $\nabla\xi$ is a formula in our finitary syntax.

5.3.9. Lemma. *Let $\Xi \in T_\omega(\mathcal{P}_\omega(\mu\mathcal{L}^T))$. Then $(\overline{T}\sqcap_{\mathcal{L}^T})[\Xi] \in T_\omega(\mu\mathcal{L}^T)$.*

Proof. Suppose $\Xi \in T_\omega(\mathcal{P}(\mu\mathcal{L}^T))$ and let $\xi := (\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]$. Furthermore define R to be the restriction of $\sqcap_{\mu\mathcal{L}^T}$ to $Base(\Xi)$. Then obviously $(\Xi, \xi) \in \overline{T}\sqcap_{\mu\mathcal{L}^T}$ and because of Lemma A.2.14 $(\Xi, \xi) \in \overline{T}R$. Hence $\xi \in F(\mathsf{rng}(R))$ by the definition of $\overline{T}$, where

$$\mathsf{rng}(R) := \{\varphi \in \mu\mathcal{L}^T \mid \exists \Phi \in Base(\Xi) \,.\, (\Phi, \varphi) \in R\}.$$

As $\mathsf{rng}(R)$ is obviously finite we finally get $\xi \in T_\omega(\mu\mathcal{L}^T)$. QED

We are now ready for stating the main theorem of this section.

5.3.10. Theorem. *Let $\mathcal{X} \subseteq_\omega T_\omega(\mu\mathcal{L}^T)$, $\mathbb{S} = (S, \sigma)$ a T-coalgebra and $s \in S$. Then*

$$\mathbb{S}, s \models \bigwedge_{\pi \in \mathcal{X}} \nabla\pi \quad \Leftrightarrow \quad \textit{there is } \Xi \in \mathcal{D}_\omega(\mathcal{X}) \textit{ s.t. } \mathbb{S}, s \models \nabla\left(\overline{T}\sqcap_{\mu\mathcal{L}^T}\right)[\Xi].$$

Proof. $\underline{\Rightarrow}$: Suppose $\mathbb{S}, s \models \bigwedge_{\pi \in \mathcal{X}} \nabla\pi$. We define $\mathcal{B} := \bigcup_{\pi \in \mathcal{X}} Base(\pi)$ and let

$$Z_\pi := \{(s', \varphi) \in S \times Base(\pi) \mid \mathbb{S}, s' \models \varphi\}.$$

According to our assumption we have for all $\pi \in \mathcal{X}$ that $(\sigma(s), \pi) \in \overline{T} Z_\pi$. Now let

$$Z := \{(s', \cup_{\pi \in \mathcal{X}} Z_\pi[s']) \mid s' \in S\} \subseteq S \times \mathcal{P}\mathcal{B}.$$

Note that by definition $Z \subseteq S \times \mathcal{P}\mathcal{B}$ is obviously the graph of a function and that therefore $\overline{T} Z$ is the graph of a function as well. Furthermore observe that $Z \circ \ni_{\mathcal{B}} \supseteq Z_\pi$ for all $\pi \in \mathcal{X}$. Hence for all $\pi \in \mathcal{X}$

$$(\sigma(s), \pi) \in \overline{T} Z_\pi \subseteq \overline{T}(Z \circ \ni_{\mathcal{B}}) = \overline{T} Z \circ \overline{T} \ni_{\mathcal{B}}.$$

Combining this with the fact that $\overline{T} Z$ is functional we get the existence of a unique $\Xi \in T\mathcal{P}\mathcal{B} \subseteq T\mathcal{P}\mu\mathcal{L}^T$ such that $(\sigma(s), \Xi) \in \overline{T} Z$ and $(\Xi, \pi) \in \overline{T} \ni_{\mathcal{B}}$ for all $\pi \in \mathcal{X}$. To show that $\Xi \in \mathcal{D}_\omega(\mathcal{X})$, it suffices to check that $\Xi \in T_\omega(\mathcal{P}\mu\mathcal{L}^T)$ but this is immediate, because $\Xi \in T\mathcal{P}\mathcal{B}$ and $\mathcal{B} \subseteq \mu\mathcal{L}^T$ is a finite set.
Note that Ξ fulfills the requirements of Lemma 5.3.9 and therefore $\nabla(\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]$ is a well-defined formula. We want to prove that in fact $\mathbb{S}, s \models \nabla(\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]$, i.e. that $(\sigma(s), (\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]) \in \overline{T} \models$. We know that

$$(\sigma(s), \Xi) \in \overline{T} Z \quad \text{and} \quad (\Xi, (\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]) \in \overline{T}\sqcap_{\mu\mathcal{L}^T}$$

so it suffices to show that $Z \circ \sqcap_{\mu\mathcal{L}^T} \subseteq \models$. Consider an arbitrary pair $(s, \psi) \in Z \circ \sqcap_{\mathcal{L}^T}$ and suppose $(s, \Psi) \in Z$ and $\psi = \bigwedge \Psi$. Then

$$\mathbb{S}, s \models \bigwedge \Psi \quad \Leftrightarrow \quad \forall \psi \in \Psi.\ \mathbb{S}, s \models \psi \quad \Leftarrow \quad Z \circ \ni \subseteq \models,$$

and as we know that the last statement is true we can conclude that $\mathbb{S}, s \models \bigwedge \Psi$.
$\underline{\Leftarrow}$: Suppose now that there is a $\Xi \in \mathcal{D}_\omega(\mathcal{X})$ such that $\mathbb{S}, s \models (\overline{T}\sqcap_{\mu\mathcal{L}^T})[\Xi]$. Then spelling out the definitions we get

$$(\sigma(s), \pi) \in \overline{T}(\models \circ \sqcap_{\mu\mathcal{L}^T}^{\smile} \circ \ni) \qquad \text{for all} \quad \pi \in \mathcal{X}$$

and it is easy to check that

$$\models \circ \sqcap_{\mu\mathcal{L}^T}^{\smile} \circ \ni \subseteq \models,$$

hence $(\sigma(s), \pi) \in \overline{T} \models$ for all $\pi \in \mathcal{X}$ and therefore $\mathbb{S}, s \vdash \bigwedge_{\pi \in \mathcal{X}} \nabla\pi$. QED

In case that T maps finite sets to finite sets we obtain a validity of coalgebraic logic.

5.3.11. Corollary. *Let T : Set $\to$ Set be a standard and weak pullback preserving functor that maps finite sets to finite sets. Then for all T-coalgebras $\mathbb{S}$ and all $\mathcal{X} \subseteq_\omega T_\omega(\mu\mathcal{L}^T)$ we get*

$$\mathbb{S} \models \bigwedge_{\pi \in \mathcal{X}} \nabla\pi \leftrightarrow \bigvee_{\Xi \in \mathcal{D}_\omega(\mathcal{X})} \nabla\left(\overline{T}\sqcap_{\mu\mathcal{L}^T}\right)[\Xi].$$

Proof. The claim follows immediately from the theorem and the fact that under the additional assumption on the functor the set $\mathcal{D}_\omega(\mathcal{X})$ is finite. QED

5.4 Concluding remarks

Relevance of Theorem 5.2.1

The fact that every T-automaton can be transformed into an equivalent non-deterministic one has many important consequences. This has been demonstrated first by Muller and Schupp in [MS95] for word and tree automata. In particular they show how to use this fact for proving Rabin's complementation lemma for tree automata: complementing an alternating tree automaton is trivial. Therefore the complementation of a nondeterministic automaton $\mathbb{A}$ can be reduced to the transformation of the alternating automaton, that accepts the complement of $\mathbb{A}$, into an equivalent nondeterministic one. This idea can, however, not easily carried over to arbitrary T-automata. The reason for this lies in the asymmetry of the bisimulation part of the acceptance game of an T-automaton. In the future we would like to investigate this question in more detail.

Another well known application of (an instance of) Theorem 5.2.1 is related to the modal μ-calculus. Janin and Walukiewicz establish in [JW95] a one-to-one correspondence between modal μ-formulas and μ-automata (graph automata) in general, and between *disjunctive* modal μ-formulas and non-deterministic μ-automata in particular. In this context the transformation of an arbitrary μ-automaton into an equivalent non-deterministic one corresponds to proving that every modal μ-formula is semantically equivalent to a disjunctive one. Disjunctive formulas play an important role in Walukiewicz's completeness proof of the modal μ-calculus in [Wal00]. A main step in his proof is that the semantic equivalence of an arbitrary formula to a disjunctive one is also a *provable equivalence* in Kozen's axiomatisation of the modal μ-calculus.

An application of Theorem 5.2.1 on the level of coalgebra automata, which we presented in this thesis, is the solution of the non-emptiness problem of T-automata and, as consequence, the proof of the finite model property of coalgebraic fixed-point logic. It will be interesting to compare our proofs to existing solutions to the non-emptiness problem of word, tree and graph automata. In particular we would like to compare our proof to existing proofs of the finite model property of the modal μ-calculus (cf. e.g. [Koz88, SE89]).

Other questions concerning coalgebra automata

The work on T-automata can be extended in many directions: as mentioned above in Section 5.2.4 our requirement that the functor T is standard can be dropped without problem. An immediate question is of course whether we can also define T-automata for functors that do not preserve weak pullbacks and that still have the closure properties.

In the main proofs of this chapter we heavily used the fact that parity games are history-free determined. Can we still prove similar results when we consider different acceptance conditions, such as the Büchi, Muller and Rabin condition? The main problem is that winning strategies cannot be assumed to be history-free anymore (cf.

[Zie98]).

Another variation on coalgebra automata would be to consider different base categories, i.e. to define T-automata for functors $T : \mathsf{C} \to \mathsf{C}$ for some arbitrary category C. One obvious choice here would be to take $\mathsf{C} = \mathsf{Stone}$. Acceptance by a T-automaton could then correspond to the fact, that the negation of a formula is not *provable*.

Finally we would like to point out that there seems to be a close connection to the work on coalgebraic trace semantics which was initiated by Jacobs in [Jac04] and carried further by Hasuo and Jacobs in [HJ05a, HJ05b].

Further consequences for coalgebraic (fixed-point) logic

We have shown that T-automata are closed under projection. It was observed by Venema that this has as an immediate corollary uniform interpolation of coalgebraic fixed-point logic following similar work by D'Agostino and Hollenberg in [DH00].

Adding negation to the language of coalgebraic (fixed-point) logic is is also an interesting issue which has been addressed already to some extent in [Mos99]. This is of course closely related to the above mentioned problem of complementation of T-automata: if we want to maintain the correspondence between formulas and automata while adding negation, we have to make sure that the recognizable languages are closed under complementation.

In case that the functor $T : \mathsf{Set} \to \mathsf{Set}$ maps finite sets to finite sets we showed that a certain distributive law is valid, i.e. it can be added as a sound axiom to a possible axiomatisation of coalgebraic (fixed-point) logic. The question of a complete axiomatisation of this logic remains of course open. A possible starting point could be to use Kozen's axiomatisation of the modal μ-calculus to obtain an axiomatisation of the T-logic for $T = \mathcal{P}$.

Appendix A

Category theory

A.1 Basic notions of category theory

For basic notions of category theory such as categories and functors we refer the reader to [Mac71]. The most important categories appearing in this thesis are listed in Table A.1.

A.1.1. Notation. *Let* C *be a category and* $A, B \in \mathsf{C}$. *Then we write* $\mathsf{C}(A, B)$ *for the collection of morphisms between A and B. If* $\mathsf{C} = \mathsf{Set}$ *we write* $[A, B]$ *for* $\mathsf{Set}(A, B)$.

A.1.2. Definition. Let C be a category and $T : \mathsf{C} \to \mathsf{C}$ a functor. Then C^{op} denotes the *opposite category* which has the same objects as C and where all the arrows are reversed, i.e. $\mathsf{C}(A, B) = \mathsf{C}^{op}(B, A)$ for all $A, B \in \mathsf{C}$. Given an arrow $f : A \to B \in \mathsf{C}$ we write $f^{op} : B \to A$ for the corresponding arrow in the opposite category. By $T^{op} : \mathsf{C}^{op} \to \mathsf{C}^{op}$ we denote the opposite of T, i.e. $T^{op}(A) = TA$ and $T^{op} f^{op} = (Tf)^{op}$.

A category that is important in this thesis is the category of algebras for a functor.

A.1.3. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor on some category C. Then a *T-algebra* is a pair (A, α) such that A is an object in C and $\alpha : TA \to A$ is a C-morphism. A *T-algebra morphism* between two T-algebras (A_1, α_1) and (A_1, α_2) is a C-morphism

Category	Objects	Morphisms
Set	sets	functions
Rel	sets	relations
BA	Boolean algebras	homomorphisms
Stone	Stone spaces	continuous functions
Cat	categories	natural transformations

Table A.1: Some important categories

$f : A_1 \to A_2$ such that $\alpha_2 \circ Tf = f \circ \alpha_1$. By $\mathsf{Alg}(T)$ we denote the category of T-algebras and T-algebra morphisms.

A.1.4. Definition. Let $T, S : \mathsf{C} \to \mathsf{D}$ be two functors. A *natural transformation* $\lambda : T \Rightarrow S$ is a family $(\lambda_X : TX \to SX)_{X \in \mathsf{C}}$ such that for all $X, Y \in \mathsf{C}$ and $f : X \to Y$ the following diagram commutes

$$\begin{array}{ccc} TX & \xrightarrow{\lambda_X} & SX \\ {\scriptstyle Tf}\downarrow & & \downarrow{\scriptstyle Sf} \\ TY & \xrightarrow{\lambda_Y} & SY \end{array}$$

A.1.5. Definition. Let C be a category. Then we define the following functors:

$$\begin{array}{rrcl} \mathsf{C}(\_, Y) : & \mathsf{C}^{op} & \to & \mathsf{Set} \\ & X & \mapsto & \mathsf{C}(X, Y) \\ & (f : X \to X') & \mapsto & \mathsf{C}(f, Y) \end{array} \qquad \begin{array}{rrcl} \mathsf{C}(X, \_) : & \mathsf{C} & \to & \mathsf{Set} \\ & Y & \mapsto & \mathsf{C}(X, Y) \\ & (g : Y \to Y') & \mapsto & \mathsf{C}(X, g) \end{array}$$

where

$$\begin{array}{rrcl} \mathsf{C}(f, Y) : & \mathsf{C}(X', Y) & \to & \mathsf{C}(X, Y) \\ & h & \mapsto & h \circ f \end{array} \qquad \begin{array}{rrcl} \mathsf{C}(X, g) : & \mathsf{C}(X, Y) & \to & \mathsf{C}(X, Y') \\ & h & \mapsto & g \circ h. \end{array}$$

A.1.6. Theorem (Yoneda lemma). *Let $S : \mathsf{C}^{op} \to \mathsf{Set}$ be a functor and let $X \in C$. Then there is a natural isomorphism*

$$\Theta_{S,X} : \mathsf{Cat}(\mathsf{C}(\_, X), S) \xrightarrow{\cong} SX.$$

A.2 Set-functors

A.2.1 Basic constructions

A.2.1. Definition. Given two functors $T_1, T_2 : \mathsf{Set} \to \mathsf{Set}$ we define their product

$$\begin{array}{rcl} T_1 \times T_2 : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & TX_1 \times TX_2 \\ f & \mapsto & Tf_1 \times Tf_2, \end{array}$$

and their sum

$$\begin{array}{rcl} T_1 + T_2 : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & TX_1 + TX_2 \\ f & \mapsto & Tf_1 + Tf_2. \end{array}$$

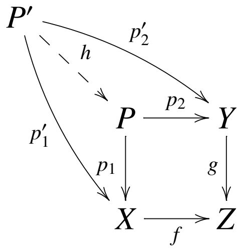

Figure A.1: Weak pullback

Here for two sets X_1, X_2 we denote by $X_1 \times X_2$ their cartesian product and by $X_1 + X_2$ their coproduct, i.e. their disjoint sum. Furthermore for two functions $f_1 : X_1 \times X_2$ and $f_2 : Y_1 \to Y_2$ we define

$$\begin{array}{rcl} f_1 \times f_2 : X_1 \times Y_1 & \to & X_2 \times Y_2 \\ (x, y) & \mapsto & (f_1(x), f_2(y)), \end{array}$$

and

$$\begin{array}{rcl} f_1 + f_2 : X_1 + Y_1 & \to & X_2 + Y_2 \\ x & \mapsto & \begin{cases} f_1(x) & \text{if} \quad x \in X_1 \\ f_2(x) & \text{if} \quad x \in Y_1 \end{cases} . \end{array}$$

A.2.2. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor and D a set. Then we define a functor

$$\begin{array}{rcl} T^D : \mathsf{Set} & \to & \mathsf{Set} \\ X & \mapsto & (TX)^D := [D, TX] \\ (f : X \to Y) & \mapsto & [D, Tf] : [D, X] \to [D, Y] \end{array}$$

If n is a natural number we write T^n for the functor $T^{\{0,\ldots,n-1\}}$.

A.2.2 Standard and weak pullback preserving functors

A.2.3. Definition. Given a category C and two morphisms $f : X \to Z$ and $g : Y \to Z$, we call the triple $(P, p_1 : P \to X, p_2 : P \to Y)$ a *weak pullback* of f and g if

(i) $f \circ p_1 = f \circ p_2$ and

(ii) for all triples $(P', p'_1 : P' \to X, p'_2 : P' \to Y)$ with $f \circ p'_1 = f \circ p'_2$ there is a morphism $h : P' \to P$ such that $p_1 \circ h = p'_1$ and $p_2 \circ h = p'_2$.

Figure A.1 illustrates the situation.

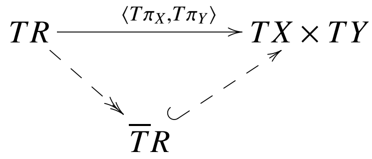

Figure A.2: Relation Lifting

A.2.4. Definition. A functor $T : \mathsf{C} \to \mathsf{C}$ is called *weak pullback preserving* if (P, p_1, p_2) being a weak pullback of f and g implies that (TP, Tp_1, Tp_2) is a weak pullback of Tf and Tg.

A central role in coalgebraic logic is played by the so-called relation lifting.

A.2.5. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor, $R \subseteq X \times Y$ a binary relation and $\pi_X : R \to X$, $\pi_Y : R \to Y$ the projection maps. Then we define the *lifted relation* $\overline{T}R \subseteq TX \times TY$ as the set

$$\overline{T}R := \{(x, y) \mid \exists z \in TR. T\pi_X(z) = x \ \& \ T\pi_Y(z) = y\}.$$

A.2.6. Remark. Expressed in more categorical terms $\overline{T}R$ is obtained by factoring the arrow $\langle T\pi_X, T\pi_Y \rangle : TR \to TX \times TY$ as shown in Figure A.2.

In case the functor T under consideration preserves weak pullbacks this relation lifting gives rise to a functor $T : \mathsf{Rel} \to \mathsf{Rel}$, where Rel is the category of sets and relations.

A.2.7. Fact. [Trn77] Let $T : \mathsf{Set} \to \mathsf{Set}$ be a weak pullback preserving functor. Then the following defines a functor $\overline{T} : \mathsf{Rel} \to \mathsf{Rel}$, the *unique extension of T to* Rel:

$$\begin{array}{rcl} \overline{T} : \mathsf{Rel} & \to & \mathsf{Rel} \\ X & \mapsto & TX \\ R \subseteq X \times Y & \mapsto & \overline{T}R. \end{array}$$

In particular this means that for a weak pullback preserving functor T we have $\overline{T}(R \circ S) = \overline{T}R \circ \overline{T}S$, i.e. $\overline{T}$ preserves the composition of relations.

A.2.8. Definition. A functor $T : \mathsf{Set} \to \mathsf{Set}$ is called *standard* if for all sets X, Y such that $X \subseteq Y$ we have $TX \subseteq TY$ and inclusion maps $i : X \hookrightarrow Y$ are mapped to inclusion maps $Ti : TX \hookrightarrow TY$.

Note that this definition of standardness, which has been used in various papers (cf. e.g. [Mos99, Ven04]) differs slightly from the one which was used by Adámek and Trnková e.g. in [AT90]. To state the latter definition we first have to introduce the following two Set-functors and the notion of a distinguished point of a Set-functor.

$$
\begin{array}{rcl}
C_1 : \mathsf{Set} & \to & \mathsf{Set} \\
X & \mapsto & 1 \\
\\
f & \mapsto & id_1
\end{array}
\qquad
\begin{array}{rcl}
C_{01} : \mathsf{Set} & \to & \mathsf{Set} \\
X & \mapsto & \begin{cases} \emptyset & \text{if} \quad X = \emptyset \\ 1 & \text{otherwise.} \end{cases} \\
(f : X \to Y) & \mapsto & \begin{cases} \emptyset & \text{if} \quad X = \emptyset \\ id_1 & \text{otherwise.} \end{cases}
\end{array}
$$

A.2.9. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor. Then a natural transformation $a : C_{01} \Rightarrow T$ is called a *distinguished point of T*.

A.2.10. Remark. A distinguished point of a functor T occurs "naturally" in TX for all nonempty sets X: If a is a distinguished point then $a_X : 1 \to TX$ can be identified with some element $\overline{a}_X \in TX$. The naturality of a means that for all non-empty sets X and Y and all functions $f : X \to Y$ we have $(Tf)(\overline{a}_X) = (Tf)(\overline{a}_Y)$.

Now we are ready to state the notion of standardness of a functor from [AT90]. In order to avoid confusion we call it $\emptyset$-standardness.

A.2.11. Definition. A functor $T : \mathsf{Set} \to \mathsf{Set}$ is called *$\emptyset$-standard* if T is standard and every distinguished point $a : C_{01} \Rightarrow T$ of T can be extended to a natural transformation $\overline{a} : C_1 \to T$.

In this thesis we are only considering standard functors that are also weak pullback preserving. For weak pullback preserving functors standardness and $\emptyset$-standardness coincide as the next lemma shows.

A.2.12. Lemma. *Let $T : \mathsf{Set} \to \mathsf{Set}$ be standard and weak pullback preserving. Then T is also $\emptyset$-standard.*

Proof. Let T be a standard and weak pullback preserving functor and let a be a distinguished point of T, i.e. $a : C_{01} \Rightarrow T$. In order to prove our claim we have to extend a to a natural transformation $\overline{a} : C_1 \Rightarrow T$. For all sets $X \neq \emptyset$ we put therefore $\overline{a}_X := a_X$.

Consider now two sets X and Y which are disjoint and nonempty and let $Z := X \cup Y$ their union. Then the following diagram is a pullback diagram in Set

$$
\begin{array}{ccc}
\emptyset & \xrightarrow{\emptyset_Y} & Y \\
{\scriptstyle \emptyset_X}\downarrow & & \downarrow{\scriptstyle i_Y} \\
X & \xrightarrow[i_X]{} & X \cup Z
\end{array}
$$

where the maps i_X, i_Y are the inclusion maps. Because T is weak pullback preserving the T-image of this diagram is a weak pullback diagram.

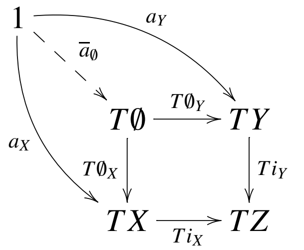

Because a is a natural transformation from C_{01} to T we get $Ti_X \circ a_X = Ti_Y \circ a_Y$ and therefore, by the fact that the square in the above diagram is a weak pullback, we get the existence of a function $\overline{a}_\emptyset : 1 \to T\emptyset$ such that

$$T\emptyset_X \circ \overline{a}_\emptyset = a_X. \tag{A.1}$$

In order to prove that with this definition $\overline{a}$ is indeed a natural transformation from C_1 to T it suffices to show that for an arbitrary nonempty set U the diagram below commutes.

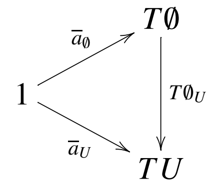

Take an arbitrary function $f : U \to X$. Then

$$\begin{aligned} T\emptyset_U \circ \overline{a}_\emptyset &= Tf \circ T\emptyset_X \circ \overline{a}_\emptyset && (\emptyset_U = f \circ \emptyset_X) \\ &= Tf \circ \overline{a}_X && \text{(Equation A.1)} \\ &= \overline{a}_U && \text{(naturality of } a) \end{aligned}$$

Hence the diagram commutes and $\overline{a}$ is a natural transformation of type $C_1 \Rightarrow T$ as required. QED

In particular this yields an important property of standard and weak pullback preserving Set-functors.

A.2.13. Fact. Let T : Set $\to$ Set be a standard and weak pullback preserving functor. Then T preserves finite intersections, i.e.

$$T(\bigcap_{i=1}^{n} U_i) = \bigcap_{i=1}^{n} TU_i.$$

Proof. The functor T is standard and weak pullback preserving and therefore according to Lemma A.2.12 also $\emptyset$-standard. For a proof of the fact that $\emptyset$-standard functors preserve finite intersections we refer the reader to [AT90, Prop. III.4.6]. QED

We will need the following properties of the lifting of standard Set-functors.

A.2.14. LEMMA. *Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor and $R \subseteq X \times Y$ a relation. Then*

1. $\overline{T}(R^{\smile}) = (\overline{T}R)^{\smile}$, *where $(\_)^{\smile}$ denotes the converse relation,*
2. $\overline{T}$ *is monotone, i.e. if $S \subseteq X \times Y$ is a relation such that $R \subseteq S$, then we have $\overline{T}R \subseteq \overline{T}S$, and*
3. *if T is standard, $\overline{T}$ commutes with restrictions, i.e. if $X' \subseteq X$, $Y' \subseteq Y$ and $R_{\restriction X' \times Y'} \neq \emptyset$ then*
$$\overline{T}(R_{\restriction X' \times Y'}) = \overline{T}R_{\restriction TX' \times TY'}.$$

Proof. Let $j : R \to R^{\smile}$ be the obvious isomorphism between R and $R^{\smile}$ mapping $(x, y) \in R$ to $(y, x) \in R^{\smile}$. Then (1) easily follows from the commutativity of the following chain of equivalences:

$$\begin{array}{rcl}
(y, x) \in (\overline{T}R)^{\smile} & \Leftrightarrow & \exists z \in TR \,.\, T\pi^R_Y(z) = y \text{ and } T\pi^R_X(z) = x \\
 & \Leftrightarrow & \exists z \in TR \,.\, T\pi^{R^{\smile}}_Y(Tj(z)) = y \text{ and } T\pi^{R^{\smile}}_X(Tj(z)) = x \\
 & \Leftrightarrow & \exists z' \in TR^{\smile} \,.\, T\pi^{R^{\smile}}_Y(z') = y \text{ and } T\pi^{R^{\smile}}_X(z') = x \\
 & \Leftrightarrow & (y, x) \in \overline{T}(R^{\smile}) \,.
\end{array}$$

The second claim of the lemma is an immediate consequence of the commutativity of the following diagram:

$$\begin{array}{ccccc}
TX & \xleftarrow{T\pi^S_X} & TS & \xrightarrow{T\pi^S_X} & TY \\
 & {\scriptstyle T\pi^R_X} \nwarrow & \cup\!\uparrow & \nearrow {\scriptstyle T\pi^R_Y} & \\
 & & TR & &
\end{array}$$

To prove the third claim note that because $R' := R_{\restriction X' \times Y'} \neq \emptyset$ we can define functions $i_X : X \to X'$, $i_Y : Y \to Y'$ and $i_R : R \to R'$ which are semi-inverse to the respective inclusion maps and which make the following diagram commute:

$$\begin{array}{ccccc}
X & \xleftarrow{\pi_X} & R & \xrightarrow{\pi_Y} & Y \\
{\scriptstyle i_X}\downarrow & & {\scriptstyle i_R}\downarrow & & {\scriptstyle i_Y}\downarrow \\
X' & \xleftarrow[\pi_{X'}]{} & R' & \xrightarrow[\pi_{Y'}]{} & Y'
\end{array}$$

The claim now follows from the following sequence of equivalences:

$$\begin{array}{rcl}
(x, y) \in (\overline{T}R)_{\restriction TX' \times TY'} & \Leftrightarrow & \exists z \in TR \,.\, T\pi_X(z) = x \in TX' \text{ and } T\pi_Y(z) = y \in TY' \\
 & \Leftrightarrow & \exists z \in TR \,.\, Ti_X(T\pi_X(z)) = x \text{ and } Ti_Y(T\pi_Y(z)) = y \\
 & \overset{\text{(diagram)}}{\Leftrightarrow} & \exists z \in TR \,.\, T\pi_{X'}(Ti_R(z)) = x \text{ and } T\pi_{Y'}(Ti_R(z)) = y \\
 & \Leftrightarrow & \exists z' \in TR' \,.\, T\pi_{X'}(z') = x \text{ and } T\pi_{Y'}(z') = y \\
 & \Leftrightarrow & (x, y) \in \overline{T}R'.
\end{array}$$

Here the second equivalence follows from standardness of T. QED

A.3 Coalgebras

In this section we introduce some basic notions from universal coalgebra.

A.3.1. Definition. Let C be a category and $T : \mathsf{C} \to \mathsf{C}$ be a functor. Then a *T-coalgebra* is a pair (X, γ) where $X \in \mathsf{C}$ and $\gamma : X \to TX \in \mathsf{C}$. Let (X, γ) and (Y, δ) be T-coalgebras. A *T-coalgebra morphism* $f : (X, \gamma) \to (Y, \delta)$ is a morphism $f : X \to Y \in \mathsf{C}$ such that the following diagram commutes:

$$\begin{array}{ccc} TX & \xrightarrow{Tf} & TY \\ \gamma \uparrow & & \uparrow \delta \\ X & \xrightarrow[f]{} & Y \end{array}$$

The category $\mathsf{Coalg}(T)$ has T-coalgebras as objects and T-coalgebra morphisms as arrows.

Rather than talking about T-coalgebras we will often talk about *rooted T-coalgebras*. A rooted T-coalgebra is a pair $(\mathbb{X}, x_I)$ where $\mathbb{X} = (X, \gamma)$ is a T-coalgebra and $x_I \in X$ is an element of X, the so-called "root" of $\mathbb{X}$.

A.3.2. Example. The category $\mathsf{Coalg}(\mathcal{P})$ corresponds to the category of Kripke frames and bounded morphisms (see e.g. [Rut95]).

One way of looking at coalgebras is that a coalgebra consists of some set of states X and the coalgebra map $\gamma : X \to TX$ allows us to *observe certain properties* of these states. Coalgebra morphisms preserve and reflect observable properties of objects, i.e. for a coalgebra morphism $f : (X, \gamma) \to (Y, \delta)$ we want that x and $f(x)$ are *observably* or *behaviourally equivalent*. It turns out that this notion of behavioural equivalence generalizes existing notions of bisimilarity and therefore we refer to it as *T-bisimilarity*.

A.3.3. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor, $\mathbb{X} = (X, \gamma), \mathbb{Y} = (Y, \delta) \in \mathsf{Coalg}(T)$, $x \in X$ and $y \in Y$. Then we say $(\mathbb{X}, x)$ and $(\mathbb{Y}, y)$ are *T-bisimilar* if there is a $(Z, \xi) \in \mathsf{Coalg}(T)$ and T-coalgebra morphisms $f : (X, \gamma) \to (Z, \xi), g : (Y, \delta) \to (Z, \xi)$ such that $f(x) = g(y)$. In this case we write $(\mathbb{X}, x) \leftrightarroweq_T (\mathbb{Y}, y)$.

A.3.4. Remark. Note that our notion of bisimilarity is often referred to as *behavioural equivalence*.

If T preserves weak pullbacks there is an alternative definition of bisimilarity in terms of T-bisimulations due to Aczel and Mendler (cf. [AM89]). We only state the definition for the case that our base category is Set. The generalization to an arbitrary category is straightforward, but not needed in this thesis.

A.3.5. Definition. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor, $\mathbb{X} = (X, \gamma), \mathbb{Y} = (Y, \delta) \in \mathsf{Coalg}(T)$ and $Z \subseteq X \times Y$ a binary relation. Then Z is called a T-bisimulation if there is a function $\rho : Z \to TZ$ such that the following diagram commutes

$$\begin{array}{ccccc} X & \xleftarrow{\pi_Y} & Z & \xrightarrow{\pi_X} & Y \\ \downarrow{\scriptstyle \gamma} & & \downarrow{\scriptstyle \rho} & & \downarrow{\scriptstyle \delta} \\ \mathcal{P}X & \xleftarrow{\mathcal{P}\pi_Y} & \mathcal{P}Z & \xrightarrow{\mathcal{P}\pi_X} & \mathcal{P}Y \end{array}$$

i.e. such that the projection maps π_X, π_Y are T-coalgebra morphisms.

Following the ideas of Hermida and Jacobs in [HJ98] the definition of a T-bisimulation can be reformulated using relation lifting.

A.3.6. Fact. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor, $\mathbb{X} = (X, \gamma), \mathbb{Y} = (Y, \delta) \in \mathsf{Coalg}(T)$ and $Z \subseteq X \times Y$. Then Z is a T-bisimulation iff for all $(x, y) \in Z$ we have $(\gamma(x), \delta(y)) \in \overline{T}Z$.

Proof. The claim, which is not difficult to prove, is an immediate corollary of [Rut98b, Thm. 1.2]. QED

If T preserves weak pullbacks, then T-bisimulations match with the notion of T-bisimilarity in the following sense.

A.3.7. Fact. Let $T : \mathsf{Set} \to \mathsf{Set}$ be a functor that preserves weak pullbacks and let $(\mathbb{X}, x_I) := (X, \gamma, x_I), (\mathbb{Y}, y_I) := (Y, \delta, y_I)$ be rooted T-coalgebras. Then $(\mathbb{X}, x_I)$ and $(\mathbb{Y}, y_I)$ are T-bisimilar iff there is a T-bisimulation $Z \subseteq X \times Y$ between $\mathbb{X}$ and $\mathbb{Y}$ with $(x_I, y_I) \in Z$.

A.3.8. Remark. A similar claim could be proven in the case that T does not preserve weak pullbacks if we worked with a more general definition of a T-bisimulation. We work with the less general definition of a T-bisimulation because we need the characterization in terms of relation lifting in the context of coalgebra automata.

A.3.9. Example. Let T be the power set functor $\mathcal{P}$. Then $\mathcal{P}$-bisimulations coincide with the standard notion for bisimulation for transition systems, i.e. given two $\mathcal{P}$-coalgebras (X, γ) and (Y, δ), a relation $Z \subseteq X \times Y$ is a bisimulation between (X, γ) and (Y, δ) iff $(x, y) \in Z$ implies

(i) for all $x' \in \gamma(x)$ there is a $y' \in \delta(y)$ with $(x', y') \in Z$ and

(ii) for all $y' \in \delta(y)$ there is an $x' \in \delta(x)$ with $(x', y') \in Z$.

Appendix B

Universal Algebra

In this part of the appendix we list definitions and facts from universal algebra that are essential for our presentation. For detailed references on the topic the reader is referred to [BS81, Wec92]. We use [Vic89] as a guideline for our account of presentations of algebras using generators and relations.

B.1 Algebras

B.1.1. Definition. A (finitary) *algebraic signature* is a set Σ of operation symbols together with a function $o : \Sigma \to \omega$ assigning to each operation $\sigma \in \Sigma$ its *arity* $o(\sigma) \in \omega$. If $o(\sigma) = n$ we call σ an n-ary operation.

B.1.2. Definition. The Boolean signature consists of the two constants ("0-ary operations") $\bot$ and $\top$, a unary operation $\neg$ and the binary operations $\vee$ and $\wedge$.

B.1.3. Definition. Let Σ be an algebraic signature. A Σ-*algebra* $\mathcal{A}$ consists of a set A together with operations $\sigma^{\mathcal{A}} : A^m \to A$ for every m and every m-ary operation symbol $\sigma \in \Sigma$. Given two Σ-algebras $\mathcal{A}_1, \mathcal{A}_2$ a function $h : \mathcal{A}_1 \to \mathcal{A}_2$ is a homomorphism if for all m-ary operations $\sigma \in \Sigma$ we have

$$h(\sigma^{\mathcal{A}_1}(a_1, \dots, a_m)) = \sigma^{\mathcal{A}_2}(h(a_1), \dots, h(a_m)).$$

Given a set X we denote by $\mathbf{T}_\Sigma(X)$ the term algebra generated by X, i.e.

$$\mathbf{T}_\Sigma(X) \ni t ::= x \in X \mid \sigma(t_1, \dots, t_n) \quad \text{for} \quad \sigma \in \Sigma \quad n\text{-ary},$$

and $\sigma^{\mathbf{T}_\Sigma(X)}(t_1, \dots, t_n) := \sigma(t_1, \dots, t_n)$. Furthermore we write $t(x_1, \dots, x_n)$ if the only generators $x \in X$ occurring in t are in $\{x_1, \dots, x_n\}$ and for terms $s_1, \dots, s_n$ the term $t(s_1, \dots, s_n)$ is the term which is obtained from $t(x_1, \dots, x_n)$ by replacing any occurrence of some $x_i, i \in \{1, \dots, n\}$ by the corresponding term s_i.

B.1.4. Definition. Let Σ be an algebraic signature and $\mathcal{A} = (A, \{\sigma^{\mathcal{A}}\}_{\sigma\in\Sigma})$ a Σ-algebra. An equivalence relation $R \subseteq A \times A$ is called a Σ-*congruence* if for all n-ary operation symbols $\sigma \in \Sigma$ and for $a_1, \ldots, a_n, a'_1, \ldots, a'_n$ we have

$$a_i R a'_i \quad \text{for all } i \in \{1, \ldots, n\} \quad \text{implies} \quad \sigma^{\mathcal{A}}(a_1, \ldots, a_n) R \sigma^{\mathcal{A}}(a'_1, \ldots, a'_n).$$

Given a congruence R on $\mathcal{A}$, the Σ-algebra $\mathcal{A}/R$ is the algebra which has the set A/R of R-equivalence classes as its carrier set and for each n-ary $\sigma \in \Sigma$ an operation

$$\begin{array}{rcl} \sigma^{\mathcal{A}/R} : (A/R)^n & \to & A/R \\ ([a_1]_{\mathcal{A}/R}, \ldots, [a_n]_{\mathcal{A}/R}) & \mapsto & \left[\sigma^{\mathcal{A}}(a_1, \ldots, a_n)\right]_{\mathcal{A}/R}. \end{array}$$

We also want to consider algebras for a signature that satisfy additional equations. We first give the formal definition of an equation and then introduced (Σ, E)-algebras, i.e. those Σ-algebras that satisfy the equations in E. In the following we fix some (sufficiently large) set $\mathcal{X}$ of equation variables.

B.1.5. Definition. Given a signature Σ and a set A we define the set of A-*equation terms (for* Σ*)* $\mathbf{EquT}_\Sigma(A)$ inductively as follows

$$\mathbf{EquT}_\Sigma(A) \ni e ::= x \in \mathcal{X} \mid a \in A \mid \sigma(e_1, \ldots, e_n) \quad \text{for} \quad \sigma \in \Sigma \quad n\text{-ary}.$$

A pair (e_1, e_2) of A-equation terms is called an A-equation and will be denoted by $e_1 \approx e_2$. Given a Σ-algebra $\mathcal{A}$, every equation term e with equation variables $x_1, \ldots, x_n$ gives rise to a term function $e^{\mathcal{A}} : A^n \to A$ where $e^{\mathcal{A}}(a_1, \ldots, a_n)$ is the term $e[x_1/a_1] \ldots [x_2/a_2]$, in which all the x_i's have been substituted by the a_i's, evaluated in $\mathcal{A}$. We say $\mathcal{A}$ satisfies an equation $e_1 \approx e_2$ if $e_1^{\mathcal{A}} = e_2^{\mathcal{A}}$ and denote this fact by $\mathcal{A} \models e_1 \approx e_2$.

B.1.6. Definition. An *algebraic theory* is a pair $\mathsf{T} = (\Sigma, E)$ such that Σ is an algebraic signature and $E \subseteq \mathbf{EquT}_\Sigma(\emptyset)$ is a set of equations.

B.1.7. Definition. Let $\mathsf{T} = (\Sigma, E)$ be an algebraic theory. Then $\mathcal{A}$ is an T-algebra if $\mathcal{A}$ is a Σ-algebra and we have $\mathcal{A} \models e_1 \approx e_2$ for all $e_1 \approx e_2 \in E$. We denote by $\mathsf{Alg}(\mathsf{T})$ the category of all T-algebras with homomorphisms as arrows between them, where a function $h : \mathcal{A}_1 \to \mathcal{A}_2$ is a homomorphism between T-algebras if h is a homomorphism between the underlying Σ-algebras.

B.1.8. Definition. Let T be an algebraic theory. Then a *presentation* $\mathsf{T}\langle G \mid R\rangle$ consists of a set G of generators and a set of relations

$$R \subseteq \{e_1 \approx e_2 \mid e_1, e_2 \in \mathbf{EquT}_\Sigma(G)\}.$$

B.1.9. Definition. Let $\mathsf{T} = (\Sigma, E)$ be an algebraic theory and $\mathsf{T}\langle G \mid R\rangle$ a presentation. A *model for* $\mathsf{T}\langle G \mid R\rangle$ is a T-algebra $\mathcal{B}$ such that there is a function $\mathsf{ingen} : G \to A$ and the algebra $\mathcal{A}^G$ is an $(\Sigma \cup G, E \cup R)$-algebra, where $\mathcal{A}^G$ is obtained by extending the algebra $\mathcal{A}$ with constants (0-ary operation symbols) g for every $g \in G$ and interpreting each g by $g^{\mathcal{A}^G} := \mathsf{ingen}(g)$. In particular $\mathcal{A}^G$ satisfies all the G-equations in R.

B.1.10. DEFINITION. Let T be an algebraic theory. Then a T-algebra $\mathcal{A}$ is presented by $\mathsf{T}\langle G \mid R\rangle$ if $\mathcal{A}$ is a model of $\mathsf{T}\langle G \mid R\rangle$ and if $\mathcal{B}$ is another model of $\mathsf{T}\langle G \mid R\rangle$ then there is a unique homomorphism $h : \mathcal{A} \to \mathcal{B}$ such that

$$h(\mathsf{ingen}(g)) = \mathsf{ingen}(g) \quad \text{for all } g \in G.$$

We are going to use the following facts about presentations.

B.1.11. FACT. Let T be an algebraic theory. Then the following holds[1]:

1. If $\mathcal{A}$ and $\mathcal{B}$ are presented by $\mathsf{T}\langle G \mid R\rangle$ then $\mathcal{A} \cong \mathcal{B}$.
2. For every presentation $\mathsf{T}\langle G \mid R\rangle$ there is an algebra $\mathcal{A}$ presented by $\mathsf{T}\langle G \mid R\rangle$.
3. Every T-algebra $\mathcal{A}$ is presented by some presentation $\mathsf{T}\langle G \mid R\rangle$.

B.2 Equational Logic

B.2.1. DEFINITION. Let Σ be an algebraic signature. The *consequence relation* $\vdash^{\Sigma}_{\mathbf{EL}}$ *of equational logic* relates sets of equations for the signature Σ with equations for Σ and is defined inductively by the following rules:

Axioms:
$$\frac{}{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2}\ (e_1 \approx e_2 \in E)$$

Reflexivity:
$$\frac{}{E \vdash^{\Sigma}_{\mathbf{EL}} e \approx e}$$

Symmetry:
$$\frac{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2}{E \vdash^{\Sigma}_{\mathbf{EL}} e_2 \approx e_1}$$

Transitivity:
$$\frac{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2 \qquad E \vdash^{\Sigma}_{\mathbf{EL}} e_2 \approx e_3}{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_3}$$

Congruence:
$$\frac{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx f_1 \qquad \dots \qquad E \vdash^{\Sigma}_{\mathbf{EL}} e_n \approx f_n}{E \vdash^{\Sigma}_{\mathbf{EL}} \sigma(e_1, \dots e_n) \approx \sigma(f_1, \dots, f_n)}\ (\sigma \in \Sigma)$$

Substitution:
$$\frac{E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2}{E \vdash^{\Sigma}_{\mathbf{EL}} e_1[x := f] \approx e_2[x := f]}\ (x \in \mathcal{X})$$

where $e[x := f]$ denotes the term obtained by replacing all occurrences in e of the variable x by the term f. A Σ-*derivation* of some equation $e_1 \approx e_2$ from a set of equations E is a finite tree such that

- the root is labeled by $E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$,

[1]Note that for 2. and 3. it is essential that we are working only with finitary signatures.

- if a node is labeled by $E \vdash^{\Sigma}_{\mathbf{EL}} e'_1 \approx e'_2$ and its children are labeled by $E \vdash^{\Sigma}_{\mathbf{EL}} f^1_1 \approx f^1_2, \ldots, E \vdash^{\Sigma}_{\mathbf{EL}} f^n_1 \approx f^n_2$ then

$$\frac{E \vdash^{\Sigma}_{\mathbf{EL}} f^1_1 \approx f^1_2 \qquad \ldots \qquad E \vdash^{\Sigma}_{\mathbf{EL}} f^n_1 \approx f^1_2}{E \vdash^{\Sigma}_{\mathbf{EL}} e'_1 \approx e'_2}$$

  is an instance of one of the above rules,

- and any unlabeled node of the tree is a leaf.

We write $\vdash_{\mathbf{EL}}$ instead of $\vdash^{\Sigma}_{\mathbf{EL}}$ if Σ is clear from the context. Furthermore for a Σ-derivation D of $E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$ we define $\mathrm{Occ}(D)$ to be the set of all equation terms occurring in D.

We need the following facts about equational logic.

B.2.2. Fact. Let Σ be an algebraic signature, $E \subset E'$ two sets of equations for Σ and suppose that $E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$. Then also $E' \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$.

B.2.3. Fact. Let G, G' be sets, $e_1 \approx e_2$ a G-equation term, E a set of G-equation terms and $f : G \to G'$ a function. Then

$$E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2 \quad \text{implies} \quad E[f] \vdash^{\Sigma}_{\mathbf{EL}} (e_1 \approx e_2)[f],$$

where $(e_1 \approx e_2)[f]$ is the G'-equation obtained from $e_1 \approx e_2$ by replacing all occurrences of parameters $g \in G$ by $f(g) \in G'$ and $E[f] := \{(e \approx e')[f] \mid e \approx e' \in E\}$.

B.2.4. Fact. Let Σ be an algebraic signature, let E be a set of G equations for Σ and suppose that there is a derivation of $E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$ for some G-equation $e_1 \approx e_2$. Then there is a derivation D of $E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$ such that $\mathrm{Occ}(D) \subseteq \mathbf{EquT}_{\Sigma}(G)$.

B.2.5. Theorem (Birkhoff). *Let (Σ, E) be an algebraic theory and $e_1, e_2 \in \mathbf{EquT}_{\Sigma}(\emptyset)$ be equation terms. Then the following are equivalent:*

1. *For all algebras $\mathcal{A} \in \mathsf{Alg}(\Sigma, E)$ we have $\mathcal{A} \models e_1 \approx e_2$, and*
2. *$E \vdash^{\Sigma}_{\mathbf{EL}} e_1 \approx e_2$.*

The definition of an algebra can be also formulated in category-theoretic terms. We use at several places the notion of an algebra for a functor, which is the exact dual of the definition of a coalgebra.

B.2.6. Definition. Let $T : \mathsf{C} \to \mathsf{C}$ be a functor. Then a T-algebra is a pair (A, α) where A is an object in C and $\alpha : TA \to A \in \mathsf{C}$ is a C-arrow. A T-homomorphism between two T-algebras (A_1, α_1) and (A_2, α_2) is a morphism $f : A_1 \to A_2 \in \mathsf{C}$ such that $\alpha_2 \circ T f = f \circ \alpha_1$. The category $\mathsf{Alg}(T)$ has T-algebras as objects and T-homomorphisms as arrows.

B.2.7. Remark. Categories of algebras for a signature correspond to categories of algebras for a functor. For the representation of algebras for an algebraic theory one usually has to consider algebras for a monad, i.e. algebras for a functor that fulfill certain extra conditions. For more details on the categorical treatment of algebras the reader is referred to [Man76].

Appendix C

Parity games

Here we introduce the terminology we need concerning parity games. For a detailed introduction to the subject we refer the reader to [GTW02].

When looking at graph games we will have to talk about infinite plays of the game. Infinite plays will correspond to infinite sequences of positions. Therefore we first have to fix our terminology regarding finite and infinite sequences.

C.0.8. Definition. Let B be a set. Then we define A^* to be the collection of finite sequence over A, A^ω to be the collection of infinite sequences over A and $A^\infty := A^* \cup A^\omega$ to be the collection of all sequences over A. Sequences, i.e. elements of A^∞, will be denoted by small greek letters $\alpha, \beta \ldots$. For some $\alpha \in A^\omega$ we let

$$\mathrm{Inf}(\alpha) := \{a \in A \mid a \text{ occurs infinitely often in } \alpha\}.$$

C.0.9. Definition. A *parity graph game* is a tuple $\mathcal{G} = (B_\exists, B_\forall, E, \Omega)$ where

- $B_\exists$, $B_\forall$ are disjoint sets of positions for the players $\exists$ (“Éloise”) and $\forall$ (“Abélard”) respectively,
- $E \subseteq (B_\exists \cup B_\forall) \times (B_\exists \cup B_\forall)$ is the edge relation and
- $\Omega : B_\exists \cup B_\forall \to \omega$ is a *parity function* , i.e. a function from $B_\exists \cup B_\forall$ to ω with finite range.

The set $B := B_\exists \cup B_\forall$ is called the *board* of the game $\mathcal{G}$, the triple $(B_\exists, B_\forall, E)$ is called the *arena* of $\mathcal{G}$.

A *play* or *match* consists of an initial state $b_I \in B$ and a sequence of moves of the players in that arena according to the following rules:

- In a position $b \in B_\exists$ player $\exists$ has to move to some position $b' \in E[b]$, where $E[b] := \{b' \in B_\exists \cup B_\forall \mid (b, b') \in E\}$.
- In a position $b \in B_\forall$ player $\forall$ has to move to some position $b' \in E[b]$.

We will therefore identify a play of $\mathcal{G}$ from starting from position $b_I \in B$ with a (possibly infinite) sequence of positions $b_0 b_1 b_2 \ldots$ such that $b_0 = b_I$ and $b_{i+1} \in E[b_i]$.

A play $\beta = b_0 b_1 b_2 \ldots$ from some position $b_I = b_0 \in B$ is said to be *complete* if either

- $\beta = b_0 \ldots b_n$ is finite and $E[b_n] = \emptyset$, or
- β is infinite.

If β is a complete play of $\mathcal{G}$ then we say $\exists$ *wins* β if either

- $\beta = b_0 \ldots b_n$ and $b_n \in B_\forall$ or
- $\beta \in B^\omega$ and $\max\{\Omega(b) \mid b \in \mathrm{Inf}(\beta)\}$ is even.

Otherwise $\forall$ *wins* β.

An important property of parity graph games is their history-free determinacy, i.e. the fact that starting from any position of the arena either of the players has a history-free winning strategy. We will now formally define the notion of such a winning strategy and then state the theorem.

C.0.10. Definition. Let $\mathcal{G} = (B_\exists, B_\forall, E, \Omega)$ be a parity graph game, $B = B_\exists \cup B_\forall$ the set of positions. A *strategy* for $\exists$ ($\forall$) is a function f mapping partial plays $b_0 \ldots b_n$ with $b_n \in B_\exists$ ($b_n \in B_\forall$) to some position b. We call f an *admissible strategy* for $\exists$ (for $\forall$) from position b if for all partial plays $\beta = b_0 \ldots b_n$ with $b_0 = b$ and $b_n \in B_\exists$ ($b_n \in B_\forall$) we have $f(\beta) \in E[b_n]$.

A strategy f for $\exists$ ($\forall$) is called *history-free* if for all partial plays β and β' that end on the same position $b \in B_\exists$ ($b \in B_\forall$) we get $f(\beta) = f(\beta')$, i.e. if f only depends only on the actual position of the play and not on its history.

Let f be a strategy for $\exists$ ($\forall$). Then a play β is called *conform with* f if for all proper prefixes $b_0 \ldots b_n$ of β ending on $b_n \in B_\exists$ ($b_n \in B_\forall$) we have that $b_0 \ldots b_n f(b_0 \ldots b_n)$ is a prefix of β.

A strategy f for $\exists$ ($\forall$) is called *winning strategy for* $\exists$ *(*$\forall$*)* in $\mathcal{G}$ starting from position $b \in B$ if f is an admissible strategy from position b and all complete plays β starting in b that are conform with f are won by $\exists$ ($\forall$).

A position $b \in B$ is called *winning position* for $\exists$ ($\forall$) if there is a winning strategy for $\exists$ ($\forall$) in $\mathcal{G}$ starting from position b.

C.0.11. Notation. *Let* $\mathcal{G} = (B_\exists, B_\forall, E, \Omega)$ *be a parity graph game. Then we define*

$$\begin{aligned} \mathrm{Win}_\exists(\mathcal{G}) &:= \{b \in B_\exists \cup B_\forall \mid b \textit{ is a winning position for } \exists\} \\ \mathrm{Win}_\forall(\mathcal{G}) &:= \{b \in B_\exists \cup B_\forall \mid b \textit{ is a winning position for } \forall\} \end{aligned}$$

Now we are ready to state the main result about parity games which is important for the results in Chapter 5.

C.0.12. Theorem. *Let $\mathcal{G} = (B_\exists, B_\forall, E, \Omega)$ be a parity graph game and $B = B_\exists \cup B_\forall$ the set of positions. Then $\mathcal{G}$ is history-free determined, i.e. $B = \mathrm{Win}_\exists(\mathcal{G}) \cup \mathrm{Win}_\forall(\mathcal{G})$ and a player that has a winning strategy from a position b in $\mathcal{G}$ has also a history-free winning strategy.*

Proof. The determinacy of parity graph games is a corollary of Borel determinacy ([Mar75]). For a transparent proof of the fact that parity games are history-free determined we refer the reader to [Zie98]. QED

Samenvatting

Het doel van deze dissertatie is het verbeteren van ons begrip van de hechte band tussen modale logica en co-algebra's. Deze band wordt niet alleen duidelijk uit het feit dat Kripke frames een speciaal type co-algebra's zijn, maar ook in meer algemene zin uit het feit dat de relatie tussen modale logica en co-algebra vanuit categorie-theoretisch perspectief gezien kan worden als de duale versie van de vruchtbare en bekende relatie tussen equationele logica en algebra.

In de literatuur zijn verschillende typen modale talen voorgesteld om over co-algebra's te redeneren. In deze dissertatie beschouwen wij de volgende drie typen: om te beginnen de inductief gedefinieerde talen voor Kripke polynomiale functoren, die werden ontwikkeld in successievelijke publicaties van Kurz, Rößiger en Jacobs; Pattinsons co-algebraische modale talen die voortkomen uit "predicate liftings"; en finitaire coalgebraïsche dekpuntstalen, die werden geintroduceerd door Venema als een aanpassing van Moss' infinitaire coalgebraïsche talen.

In deze dissertatie stellen wij ons op het standpunt dat voor een logische taal waarmee zinvol over co-algebra's geredeneerd kan worden, de syntax van eindige aard zou moeten zijn. Vandaar dat alle talen die wij bespreken eindig zijn. Talen met een eindige syntax missen echter in het algemeen de Hennessy-Milner eigenschap.

Om die reden is het een natuurlijke vraag of we een klasse van co-algebra's kunnen vinden die logica's met een eindige syntax toestaat, die desondanks de Hennessy-Milner eigenschap hebben. We stellen voor om dit vraagstuk op te lossen door een bekend concept uit de modale logica te generaliseren: descriptieve gegeneraliseerde frames. Deze descriptieve gegeneraliseerde frames kunnen gerepresenteerd worden als co-algebra's voor de Vietoris functor op de categorie van Stone topologieën. Vandaar dat Stone co-algebra's, dat wil zeggen co-algebra's voor functoren over de categorie der Stone topologieën, een natuurlijke generalisering van dit concept zijn.

Een manier om de expressiviteit van een modale taal te vergroten vormt het gebruik van de zogeheten dekpuntoperatoren. Venema's coalgebraïsche dekpuntslogica's hebben een finitaire syntax en bieden de mogelijkheid over oneindig, voortdurend gedrag te redeneren. Deze logica's kunnen beschouwd worden als een generalisatie van de modale mu-calculus, en zij kunnen, net als de modale mu-calculus, op een automaten-theoretische wijze geïnterpreteerd worden: er is een één-één correspondentie tussen formules van coalgebraïsche dekpuntslogica en de zogenoemde co-algebra automaten.

In deze dissertatie bewijzen we enkele afsluitingseigenschappen voor co-algebra automaten en tonen aan dat het 'non-emptiness probleem' van een co-algebra automaat in veel gevallen beslisbaar is. Onze resultaten kunnen vanuit twee perspectieven bekeken worden: In de eerste plaats generaliseren ze bekende resultaten aangaande automaten op oneindige objecten, zoals automaten op oneindige woorden, bomen en grafen. In de tweede plaats hebben onze resultaten logische gevolgen: we tonen aan dat alle coalgebraïsche dekpuntslogica's de eindige model eigenschap hebben. Hieruit volgt in het bijzonder een bewijs voor de eindige model eigenschap van de modale mu-calculus. Een ander gevolg is beslisbaarheid voor een grote klasse van coalgebraïsche dekpuntlogica's. Verder bewijzen we de correctheid van een bepaalde distributiewet voor de ∇-operator.

Deze dissertatie is als volgt ingedeeld: na de Introductie in Hoofdstuk 1, geven we een overzicht van de drie typen modale talen, die worden besproken in dit proefschrift.

Hoofdstuk 3 bevat een eerste toepassing van het idee om co-algebra's over Stone topologieën te beschouwen. We bekijken gedefinieerde logica's voor Kripke polynomiale functoren: voor elke Kripke polynomiale functoren definieren we een corresponderende functor op de categorie van Stone topologieën en verkrijgen we, in ons vocabulaire, de klasse der Vietoris polynomiale functoren. Voor elke dergelijke functor verkrijgen we de uiteindelijke co-algebra door middel van een aangepaste kanoniek model constructie. In het bijzonder impliceert deze constructie dat de talen die geassocieerd worden met Vietoris polynomiale functoren de Hennessy-Milner eigenschap hebben. Verder bewijzen we dat er voor elke Vietoris polynomiale functor T en de daarmee geassocieerde logica een adjunctie bestaat tussen de algebraïsche semantiek van de betreffende logica, gedefinieerd als een categorie van meersoortige algebra's, en de categorie van T-co-algebra's. Ten slotte karakteriseren we de meersoortige algebra's waarvoor de adjunctie in feite een categorie-theoretische equivalentie vormt.

In Hoofdstuk 4 richten we ons op co-algebraische modale logica's, die gegeven zijn in termen van een verzameling "predicate liftings" en een verzameling axioma's met een modale diepte van 1. Gegeven een endofunctor T op de categorie der verzamelingen of de categorie der Stone topologieën en een logica voor T ontwerpen we een functor L op de categorie der Boolse algebra's. De categorie der algebra's voor deze functor bepaalt de algebraische semantiek voor deze logica. Wij gebruiken deze algebraische semantiek om een categorie-theoretische analyse te geven van de condities waaronder de logica correct en volledig is met betrekking tot de co-algebraische semantiek. Dit doen we door de correctheid en volledigheid van de logica te relateren aan eigenschappen van een natuurlijke transformatie die de functoren L en T verbindt. In het geval dat T een functor is op Stone topologieën verkrijgen we het volgende resultaat: De logica is correct en volledig en heeft de Hennessy-Milner eigenschap als L duaal is aan T.

In Hoofdstuk 5 bewijzen we afsluitingseigenschappen van co-algebraische automaten en laten we zien hoe het "non-emptiness" probleem efficient opgelost kan worden voor een grote klasse van co-algebraische automaten. Het belangrijkste resultaat van dit hoofdstuk is het bewijs dat voor elke co-algebra automaat een equivalente non-deterministische automaat geconstrueerd kan worden. Het bewijs is uniform in alle typen co-algebraische automaten en in het speciale geval van boomautomaten impliceert het Rabins Complementerings-lemma.

Bibliography

[Abr88] S. Abramsky, *A Cook's tour of the finitary non-wellfounded sets*, Invited Lecture at BCTCS, 1988.

[Abr91] ———, *Domain theory in logical form*, Annals of Pure and Applied Logic **51** (1991), 1–77.

[Abr05] ———, *A Cook's tour of the finitary non-wellfounded sets*, We Will Show Them: Essays in honour of Dov Gabbay (S. Artemov, H. Barringer, A. d'Avila Garcez, L.C. Lamb, and John Woods, eds.), College Publications, vol. 1, 2005, pp. 1–18.

[Acz88] Peter Aczel, *Non-well-founded sets*, CSLI Lecture Notes, vol. 14, 1988.

[AK79] J. Adámek and V. Koubek, *Least Fixed Point of a Functor*, Journal of Computer and System Sciences **19** (1979), 163–178.

[AM89] P. Aczel and N.P. Mendler, *A Final Coalgebra Theorem*, Category Theory and Computer Science (London, UK), LNCS, vol. 389, Springer, 1989, pp. 357–365.

[AN01] A. Arnold and D. Niwiński, *Rudiments of μ-calculus*, Studies in Logic, vol. 146, North-Holland, 2001.

[AT90] J. Adámek and V. Trnková, *Automata and algebras in a category*, Kluwer Publishing Company, 1990.

[Bal00] A. Baltag, *A logic for coalgebraic simulation*, Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS) (H. Reichel, ed.), Electronic Notes in Theoretical Computer Science, vol. 33, 2000.

[Bar93] M. Barr, *Terminal coalgebras in well-founded set theory*, Theoretical Computer Science **114** (1993), no. 2, 299–315.

[BdV01] P. Blackburn, M. de Rijke, and Y. Venema, *Modal logic*, Cambridge Tracts in Theoretical Computer Science, vol. 53, Cambridge University Press, 2001.

[Ben76] J. van Benthem, *Modal Correspondence Theory*, Ph.D. thesis, Mathematisch Instituut & Instituut voor Grondslagenonderzoek, University of Amsterdam, 1976.

[BK05] M. Bonsangue and A. Kurz, *Duality for Logics of Transition Systems*, Proceedings of FoSSaCS '05, 2005.

[BM96] J. Barwise and L. Moss, *Vicious Circles*, CSLI Lecture Notes, vol. 60, 1996.

[BS81] S. N. Burris and H. P. Sankappanavar, *A Course in Universal Algebra*, Graduate Texts in Mathematics, no. 78, Springer, 1981.

[CGP00] E. Clarke, O. Grumberg, and D. Peled (eds.), *Model checking*, The MIT Press, 2000.

[Cîr04] C. Cîrstea, *A compositional approach to defining logics for coalgebras*, Theoretical Computer Science **327** (2004), 45–69.

[CP04] Corina Cîrstea and Dirk Pattinson, *Modular construction of modal logics.*, CONCUR, 2004, pp. 258–275.

[CZ97] A. Chagrov and M. Zakharyaschev, *Modal logic*, Clarendon Press, 1997.

[DH00] G. D'Agostino and M. Hollenberg, *Logical Questions Concerning the μ-Calculus: Interpolation, Lyndon and Łoś-Tarski.*, Journal of Symbolic Logic **65** (2000), no. 1, 310–332.

[Eng89] R. Engelking, *General topology*, Heldermann Verlag, 1989.

[Esa74] L.L. Esakia, *Topological Kripke models*, Soviet Mathematics Doklady **15** (1974), 147–151.

[Fin74] K. Fine, *An incomplete logic containing* **S4**, Theoria **40** (1974), 23–29.

[Ger96] J. Gerbrandy, *Characterizations of bisimulation and bounded bisimulations*, Workshop Proceedings on Observational Equivalence and Logical Equivalence, ESSLLI (M. de Rijke, ed.), Prague, 1996.

[Gol76] R. Goldblatt, *Metamathematics of modal logic, part I*, Reports on Mathematical Logic **6** (1976), 41–78.

[Gol01] ______, *What is the coalgebraic analogue of Birkhoff's variety theorem?*, Theoretical Computer Science **266** (2001), no. 1-2, 853–886.

[Gol03a] ———, *Enlargemants of a polynomial coalgebra*, Proceedings of the 7th and 8th Asian Logic Conferences (R. Downey et al., ed.), World Scientific, 2003, pp. 152–192.

[Gol03b] ———, *Observational ultraproducts of polynomial coalgebras*, Annals of Pure and Applied Logic **123** (2003), no. 1–3, 235–290.

[Gol05] ———, *Final Coalgebras and the Hennessy-Milner Property*, Annals of Pure and Applied Logic (2005), In Press.

[GTW02] E. Grädel, W. Thomas, and T. Wilke (eds.), *Automata, logic, and infinite games*, LNCS, vol. 2500, Springer, 2002.

[Gum01] H.P. Gumm, *Functors for Coalgebras*, Algebra Universalis **45** (2001), 135–147.

[HJ98] C. Hermida and B. Jacobs, *Structural induction and coinduction in a fibrational setting*, Information and Computation **145** (1998), no. 2, 107–152.

[HJ05a] I. Hasuo and B. Jacobs, *Coalgebraic trace semantics for probabilistic systems*, CALCO-jnr Workshop, 2005.

[HJ05b] ———, *Context-free languages via coalgebraic trace semantics*, CALCO 2005 (Springer, ed.), LNCS, vol. 3629, 2005, pp. 204–233.

[HK04] H. H. Hansen and C. Kupke, *A coalgebraic perspective on monotone modal logic*, Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS) (J. Adámek, ed.), Electronic Notes in Theoretical Computer Science, vol. 106, 2004.

[Jac96] B. Jacobs, *Objects and Classes, Co-Algebraically*, Object Orientation with Parallelism and Persistence, Kluwer Academic Publishers, 1996, pp. 83–103.

[Jac01] ———, *Many-Sorted Coalgebraic Modal Logic: a Model-theoretic Study*, Theoretical Informatics and Applications **35** (2001), no. 1, 31–59.

[Jac04] ———, *Trace Semantics for Coalgebras*, Electr. Notes Theor. Comput. Sci. **106** (2004), 167–184.

[Joh82] P.T. Johnstone, *Stone Spaces*, Cambridge Uniersity Press, Cambridge, 1982.

[JR97] B. Jacobs and J. Rutten, *A tutorial on (co)algebras and (co)induction*, EATCS Bulletin **62** (1997).

[JT51] B. Jónsson and A. Tarski, *Boolean algebras with operators I*, American Journal of Mathematics **73** (1951), 891–939.

[JT52] ———, *Boolean algebras with operators II*, American Journal of Mathematics **74** (1952), 127–162.

[JW95] D. Janin and I. Walukiewicz, *Automata for the modal μ-calculus and related results*, Proceedings 20th International Symposium on Mathematical Foundations of Computer Science (MFCS'95) (Berlin), LNCS, vol. 969, Springer, 1995, pp. 552–562.

[KKP04] C. Kupke, A. Kurz, and D. Pattinson, *Algebraic semantics for coalgebraic modal logic*, Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS) (J. Adámek, ed.), Electronic Notes in Theoretical Computer Science, vol. 106, 2004.

[KKP05] ———, *Ultrafilter extensions for coalgebras*, Algebra and Coalgebra in Computer Science, LNCS, vol. 3629, 2005.

[KKV04] C. Kupke, A. Kurz, and Y. Venema, *Stone coalgebras*, Theoretical Computer Science **327** (2004), 109–134.

[Koz88] D. Kozen, *A finite model theorem for the propositional μ-calculus*, Studia Logica **47** (1988), no. 3, 233–241.

[KP04] A. Kurz and A. Palmigiano, *Coalgebras and Modal Expansions of Logics*, Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS) (J. Adámek, ed.), Electronic Notes in Theoretical Computer Science, vol. 106, 2004.

[KP05] A. Kurz and D. Pattinson, *Coalgebraic Modal Logic of Finite Rank*, Mathematical Structures in Computer Science **15** (2005), no. 3, 453–473.

[KR02] Alexander Kurz and Jiří Rosický, *Modal predicates and coequations*, CMCS'02, ENTCS, vol. 65.1, Elsevier, 2002.

[Kra99] M. Kracht, *Tools and techniques in modal logic*, Elsevier, 1999.

[Kur00] A. Kurz, *Logics for Coalgebras and Applications to Computer Science*, Ph.D. thesis, Ludwig-Maximilians-Universität, 2000.

[Kur01a] A. Kurz, *A Co-Variety-Theorem for Modal Logic*, Advances in Modal Logic 2, CSLI, 2001, Selected Papers from AiML 2, Uppsala, 1998.

[Kur01b] A. Kurz, *Coalgebras and modal logic*, Course Notes for ESSLLI'2001, 2001, Available at `http://www.cs.le.ac.uk/people/akurz/works.html`.

[Kur01c] ———, *Specifying coalgebras with modal logic*, Theoretical Computer Science **260** (2001), 119–138.

[KV05] C. Kupke and Y. Venema, *Closure properties of coalgebra automata*, 20th Annual IEEE Symposium on Logic in Computer Science (LICS 2005), IEEE Press, 2005.

[Mac71] S. Mac Lane, *Category theory for the working mathematician*, Springer, 1971.

[Man76] E.G. Manes, *Algebraic Theories*, Graduate Texts in Mathematics, vol. 26, Springer, 1976.

[Mar75] D.A. Martin, *Borel Determinacy*, The Annals of Mathematics **102** (1975), no. 2, 363–371.

[Mic51] E. Michael, *Topologies on spaces of subsets*, Transactions of the American Mathematical Society **71** (1951), 152–182.

[Mos99] L. S. Moss, *Coalgebraic logic*, Annals of Pure and Applied Logic **96** (1999), 277–317.

[MS95] D.E. Muller and P.E. Schupp, *Simulating alternating tree automata by non-deterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra*, Theoretical Computer Science **141** (1995), no. 1–2, 69–107.

[MV04] L. S. Moss and I. D. Viglizzo, *Harsanyi type spaces and final coalgebras constructed from satisfied theories*, Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS) (J. Adámek, ed.), Electronic Notes in Theoretical Computer Science, vol. 106, 2004.

[Pal04] A. Palmigiano, *A coalgebraic view on positive modal logic*, Theoretical Computer Science **327** (2004), 175–195.

[Pat01] D. Pattinson, *Semantical principles in the modal logic of coalgebras*, Proceedings 18th International Symposium on Theoretical Aspects of Computer Science (STACS 2001) (Berlin), LNCS, vol. 2010, Springer, 2001.

[Pat03a] ———, *Coalgebraic modal logic: Soundness, completeness and decidability of local consequence*, Theoretical Computer Science **309** (2003), 177–193.

[Pat03b] ———, *An introduction to the theory of coalgebras*, North American Summer School in Logic, Language and Information (NASSLLI), 2003.

[Pat04] ———, *Expressive logics for coalgebras via terminal sequence induction*, Notre Dame Journal of Formal Logic **45** (2004), no. 1, 19–33.

[Rab69] M.O. Rabin, *Decidability of second-order theories and automata on infinite trees*, Transactions of the American Mathematical Society **141** (1969), 1–35.

[Rei83] J. Reiterman, *Algebraic theories and varieties of functor algebras*, Fund. Math. **118** (1983), 59–68.

[Rei95] H. Reichel, *An Approach to Object Semantics based on Terminal Co-Algebras*, Mathematical Structures in Computer Science **5** (1995), no. 2, 129–152.

[Ros81] J. Rosický, *On algebraic categories*, Universal Algebra (Proceedings Colloquium Esztergom 1977), Colloquium Mathematical Society J. Bolyai, vol. 29, 1981, pp. 662–690.

[Röß00] M. Rößiger, *Coalgebras and modal logic*, CMCS'00, ENTCS, vol. 33, 2000.

[Röß01] ______, *From modal logic to terminal coalgebras*, Theoretical Computer Science **260** (2001), 209–228.

[Rut95] J.J.M.M. Rutten, *A calculus of transition systems (towards universal coalgebra)*, Modal Logic and Process Algebra, CSLI Lecture Notes, vol. 53, CSLI, 1995.

[Rut98a] ______, *Automata and coinduction (an exercise in coalgebra)*, Proc. 9th International Conference on Concurrency Theory (CONCUR'98) (D. Sangiorigi and R. de Simone, eds.), LNCS, vol. 1466, Springer, 1998, pp. 194–218.

[Rut98b] ______, *Relators and Metric Bisimulation (Extended Abstract)*, Electronic Notes in Theoretical Computer Science **11** (1998), 1–7.

[Rut00] ______, *Universal coalgebra: A theory of systems*, Theoretical Computer Science **249** (2000), 3–80.

[Sch05] L. Schröder, *Expressivity of Coalgebraic Modal Logic: The Limits and Beyond*, Foundations of Software Science And Computation Structures (V. Sassone, ed.), Lecture Notes in Computer Science, vol. 3441, 2005, pp. 440–454.

[SE89] R.S. Street and E.A. Emerson, *An automata theoretic procedure for the propositional μ-calculus*, Information and Computation **81** (1989), 249–264.

[Sto36] M. H. Stone, *The Theory of Representation for Boolean Algebras*, Transactions of the American Mathematical Society **40** (1936), no. 1, 37–111.

[Sto37] ______, *Applications of the Theory of Boolean Rings to General Topology*, Transactions of the American Mathematical Society **41** (1937), no. 3, 375–481.

[Tho74] S. K. Thomason, *An Incompleteness Theorem in Modal Logic*, Theoria **40** (1974), 30–34.

[Trn77] V. Trnková, *Relational automata in a category and theory of languages*, Proc. FCT 1977, Springer, 1977, LNCS 56, pp. 340–355.

[Ven04] Y. Venema, *Automata and fixed point logics for coalgebras*, Electronic Notes in Theoretical Computer Science **106** (2004), 355–375.

[Ven06] ______, *Algebras and coalgebras*, Handbook of Modal Logic, 2006, To appear.

[Vic89] S. Vickers, *Topology Via Logic*, Cambridge University Press, 1989.

[Vie22] L. Vietoris, *Bereiche zweiter Ordnung*, Monatshefte für Mathematik und Physik **32** (1922), 258–280.

[Vig05] I. D. Viglizzo, *Final sequences and final coalgebras for measurable spaces*, Algebra and Coalgebra in Computer Science, LNCS, vol. 3629, 2005.

[Wal00] I. Walukiewicz, *Completeness of Kozen's Axiomatisation of the Propositional μ-Calculus*, Information and Computation **157** (2000), 142–182.

[Wec92] W. Wechler, *Universal algebra for computer scientists*, EATCS Monographs on Theoretical Computer Science, vol. 25, Springer, 1992.

[Wor05] J. Worrell, *On the final sequence of a finitary set functor*, Theoretical Computer Science **338** (2005), no. 1–3, 184–199.

[Zie98] W. Zielonka, *Infinite games on finitely coloured graphs with applications to automata on infinite trees*, Theoretical Computer Science **200** (1998), 135–183.

Index

Titles in the ILLC Dissertation Series:

ILLC DS-2001-01: **Maria Aloni**
Quantification under Conceptual Covers

ILLC DS-2001-02: **Alexander van den Bosch**
Rationality in Discovery - a study of Logic, Cognition, Computation and Neuropharmacology

ILLC DS-2001-03: **Erik de Haas**
Logics For OO Information Systems: a Semantic Study of Object Orientation from a Categorial Substructural Perspective

ILLC DS-2001-04: **Rosalie Iemhoff**
Provability Logic and Admissible Rules

ILLC DS-2001-05: **Eva Hoogland**
Definability and Interpolation: Model-theoretic investigations

ILLC DS-2001-06: **Ronald de Wolf**
Quantum Computing and Communication Complexity

ILLC DS-2001-07: **Katsumi Sasaki**
Logics and Provability

ILLC DS-2001-08: **Allard Tamminga**
Belief Dynamics. (Epistemo)logical Investigations

ILLC DS-2001-09: **Gwen Kerdiles**
Saying It with Pictures: a Logical Landscape of Conceptual Graphs

ILLC DS-2001-10: **Marc Pauly**
Logic for Social Software

ILLC DS-2002-01: **Nikos Massios**
Decision-Theoretic Robotic Surveillance

ILLC DS-2002-02: **Marco Aiello**
Spatial Reasoning: Theory and Practice

ILLC DS-2002-03: **Yuri Engelhardt**
The Language of Graphics

ILLC DS-2002-04: **Willem Klaas van Dam**
On Quantum Computation Theory

ILLC DS-2002-05: **Rosella Gennari**
Mapping Inferences: Constraint Propagation and Diamond Satisfaction

ILLC DS-2002-06: **Ivar Vermeulen**
A Logical Approach to Competition in Industries

ILLC DS-2003-01: **Barteld Kooi**
Knowledge, chance, and change

ILLC DS-2003-02: **Elisabeth Catherine Brouwer**
Imagining Metaphors: Cognitive Representation in Interpretation and Understanding

ILLC DS-2003-03: **Juan Heguiabehere**
Building Logic Toolboxes

ILLC DS-2003-04: **Christof Monz**
From Document Retrieval to Question Answering

ILLC DS-2004-01: **Hein Philipp Röhrig**
Quantum Query Complexity and Distributed Computing

ILLC DS-2004-02: **Sebastian Brand**
Rule-based Constraint Propagation: Theory and Applications

ILLC DS-2004-03: **Boudewijn de Bruin**
Explaining Games. On the Logic of Game Theoretic Explanations

ILLC DS-2005-01: **Balder David ten Cate**
Model theory for extended modal languages

ILLC DS-2005-02: **Willem-Jan van Hoeve**
Operations Research Techniques in Constraint Programming

ILLC DS-2005-03: **Rosja Mastop**
What can you do? Imperative mood in Semantic Theory

ILLC DS-2005-04: **Anna Pilatova**
A User's Guide to Proper names: Their Pragmatics and Semanics

ILLC DS-2006-01: **Troy Lee**
Kolmogorov complexity and formula size lower bounds

ILLC DS-2006-02: **Nick Bezhanishvili**
Lattices of intermediate and cylindric modal logics

ILLC DS-2006-03: **Clemens Kupke**
Finitary coalgebraic logics