

ALGEBRAS AND COALGEBRAS

Yde Venema

1	Introduction	2
2	Basics of modal logic	3
3	Modal logic in algebraic form	7
4	Varieties of expanded Boolean algebras	12
5	Frames and algebras	17
5.1	Introduction	17
5.2	Complex duality	18
5.3	Ultrafilter frames	21
5.4	Topological duality	24
5.5	Simplicity and Subdirect irreducibility	26
5.6	Class operations	28
6	Logics and varieties	31
6.1	Completeness	32
6.2	Canonicity	35
6.3	Interpolation	39
7	Case study: canonical equations	41
7.1	Introduction	41
7.2	Canonical extensions of Boolean algebras	43
7.3	Extending maps to the canonical extension	45
7.4	Composite maps	47
7.5	Canonical equations	50
7.6	Further remarks	52
8	Special algebraic topics	54
8.1	Tense logic	54
8.2	Global modality & discriminator varieties	56
9	Coalgebras: an introduction	59
10	Final coalgebras	64
11	Bisimulation & behavioral equivalence	67
12	Covarieties	73
13	Modal logic and coalgebras	76
14	Co-Birkhoff theorems and cofree coalgebras	82
15	Duality of algebra and coalgebra	86
A	Basics of universal algebra and category theory	88

1 INTRODUCTION

Modal logic is not an isolated field. When studied from a mathematical perspective, it has evident connections with many other areas in logic, mathematics and theoretical computer science. Other chapters of this handbook point out some of the links between modal logic and areas like (finite) model theory or automata theory. Here we will outline the *algebraic* and *coalgebraic* environments of the theory of modal logic.

First we approach modal logic with the methodology of *algebraic logic*, a discipline which aims at studying all kinds of logics using tools and techniques from universal algebra — in fact, much of the theory of universal algebra was developed in tandem with that of algebraic logic. The idea is to associate, with any logic L , a class $\text{Alg}(L)$ of algebras, in such a way that (natural) logical properties of L correspond to (natural) algebraic properties of $\text{Alg}(L)$. Carrying out this program for modal logic, we find that normal modal logics have algebraic counterparts in varieties of *Boolean algebras with operators* (BAOs). In the simplest case of monomodal logics, the algebras that we are dealing with are simply *modal algebras*, that is, expansions of Boolean algebras with a single, unary operation that preserves finite joins (disjunctions). One advantage of the algebraic semantics over the relational one is that it allows a general *completeness* result, but the algebraic approach may also serve to prove many significant results concerning properties of modal logics such as completeness, canonicity, and interpolation. As we will see, a crucial observation in the algebraic theory of modal logic is that standard algebraic constructions correspond to well-known operations on Kripke frames. These correspondences can be made precise in the form of categorical *dualities*, which may serve to explain much of the interaction between modal logic and universal algebra. Our discussion of the algebraic approach towards modal logics takes up the sections 3 to 8.

The *coalgebraic* perspective on modal logic is much more recent (see section 9 for references). Coalgebras are simple but fundamental mathematical structures that capture the essence of dynamic or evolving systems. The theory of universal coalgebra seeks to provide a general framework for the study of notions related to (possibly infinite) behavior such as invariance, and observational indistinguishability. When it comes to modal logic, an important difference with the algebraic perspective is that coalgebras *generalize* rather than *dualize* the model theory of modal logic. Many familiar notions and constructions, such as bisimulations and bounded morphisms, have analogues in other fields, and find their natural place at the level of coalgebra. Perhaps even more important is the realization that one may generalize the concept of modal logic from Kripke frames to arbitrary coalgebras. In fact, the link between (these generalizations of) modal logic and coalgebra is so tight, that one may even claim that modal logic is the natural logic for coalgebras — just like equational logic is that for algebra. The second and last part of this chapter, starting from section 9, is devoted to coalgebra.

What is the point of taking such an abstract perspective on modal logic, be it algebraic or coalgebraic? Obviously, making the above kind of mathematical generalizations, one should not aim at solving all concrete problems for specific modal logics. Rather, the approach may serve to isolate those aspects of a problem that are easy in the sense of being solvable by general means; it thus enables us to focus on the remaining aspects that are specific to the problem at hand. To give an example, it is certainly not the case that all modal formulas are canonical, but Sahlqvist's theorem considerably simplifies completeness proofs by taking care of the canonical part of the axiomatization. A second

benefit of embedding modal logic in its mathematical context is that it may lead to a better understanding of notions from modal logic. Taking an example from coalgebra, the notion of a bounded morphism between Kripke models (or frames), becomes much more natural once we understand that it coincides with the natural coalgebraic notion of a homomorphism.

Our main aim with this chapter is to give the reader an impression of both the algebraic and the coalgebraic perspective on modal logic. Our focus will be on concepts and ideas, but we will also mention important techniques and landmark results; proofs, or rather proof sketches, are given as much as possible. Despite its over-average length, a text of this size cannot come close to being comprehensive; our main selection criterion has been to focus on *generality* of methods and results. Unfortunately, even some important topics have fallen prey to this, most particularly, the *algebras of relations*, even though they played and continue to play a crucial role in the history of algebraic logic. Fortunately, these kinds of BAOS are well documented elsewhere, see for instance HENKIN, MONK & TARSKI [57] for cylindric algebras, or HIRSCH & HODKINSON [58] for relation algebras. A second topic receiving only fragmented attention is *historical context*. While we do attribute results as much as possible, readers with an interest in the (fascinating!) history of modal logic, will not find much to suit their taste here. Rather, they should consult GOLDBLATT [44], or perhaps the historical notes of BLACKBURN, DE RIJKE & VENEMA [13]. Finally, a warning: in this chapter we assume familiarity with basic notions from category theory (such as functors, duality), universal algebra (such as congruences, free algebras), and more specifically, Boolean algebras. Readers encountering unfamiliar concepts in this chapter are advised to consult some text book in universal algebra or category theory. For convenience, in an appendix we have summed up all the material that we consider to be background knowledge.

Acknowledgments Many people have contributed to this chapter by commenting on earlier versions of it. I am grateful to Nick Bezhanishvili, Mai Gehrke, Çiğdem Gencer, Rob Goldblatt, H. Peter Gumm, Gaelle Fontaine, Helle Hansen, Ian Hodkinson, Ramon Jansana, Clemens Kupke, Alessandra Palmigiano, Jan Rutten, and Mark Theunissen for their help. Great is my gratitude to Alexander Kurz, who provided a large number of suggestions to improve the chapter.

2 BASICS OF MODAL LOGIC

In this section we briefly review the basic definitions of modal logic. Starting with syntax, we take a fairly general approach towards modal languages and allow modal connectives of arbitrary finite rank. A *modal similarity type* is a set τ of modal connectives, together with an arity function $ar : \tau \rightarrow \omega$ assigning to each symbol $\nabla \in \tau$ a *rank* or *arity* $ar(\nabla)$. Given a modal similarity type τ and a set X of variables we inductively define the set $Fma_\tau(X)$ of *modal τ -formulas in X* by the following rule:

$$\varphi ::= x \in X \mid \top \mid \perp \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \nabla(\varphi_1, \dots, \varphi_n)$$

with $\nabla \in \tau$ and $n = ar(\nabla)$. We will use standard abbreviations such as \rightarrow and \leftrightarrow ; we also define the *dual* operator Δ of $\nabla \in \tau$ as $\Delta(\varphi_1, \dots, \varphi_n) := \neg\nabla(\neg\varphi_1, \dots, \neg\varphi_n)$. Unary modalities are usually called *diamonds*, and their duals, *boxes*; to denote these modalities we reserve (possibly indexed) symbols of the shape \diamond and \square , respectively.

Throughout this chapter we will work with an arbitrary but fixed modal similarity type τ . Often, we will provide proofs only for the *basic modal similarity type* which consists of a single diamond that will always simply be denoted as \diamond (its dual as \square). Unless explicitly stated otherwise, we are always dealing with a fixed, countably infinite set X of variables; in order not to clutter up notation we will suppress explicit references to X as much as possible.

It will be convenient to have names and notation for some special formulas that behave just like ordinary diamond formulas of the form $\diamond v$. Fix a special dummy variable v . In the basic modal language, we may define a *compound diamond* as any disjunction of formulas of the form $\diamond^n v$ (here $\diamond^0 \varphi := \varphi$ and $\diamond^{n+1} \varphi := \diamond \diamond^n \varphi$). In a language with diamonds only, the compound diamonds may be defined as follows:

$$\blacklozenge ::= v \mid \diamond_i \blacklozenge \mid \blacklozenge_1 \vee \blacklozenge_2.$$

An example is $\diamond_1 \diamond_0 (\diamond_1 v \vee v) \vee \diamond_1 \diamond_1 v$. We will write $\blacklozenge \varphi$ for the formula in which every occurrence of v is substituted by φ (note that v is the *unique* variable occurring in a compound diamond). Induced and compound *boxes* are defined in the obvious, dual, way.

The general case, which readers may safely choose to skip, is a bit more involved. For any modality ∇ of arity $n > 1$, and any $1 \leq i \leq n$, the formula $\diamond_{(\nabla, i)} v := \nabla(\top, \dots, \top, v, \top, \dots, \top)$ (i.e., all arguments are \top except for the i -th one which is v) is called the i -th *induced diamond* of ∇ . The collection $CD(\tau)$ of *compound diamonds* of τ is defined via:

$$\blacklozenge ::= v \mid \diamond_i \blacklozenge \mid \diamond_{(\nabla, i)} \blacklozenge \mid \blacklozenge_1 \vee \blacklozenge_2.$$

Modal logic can be approached from a semantic or from a purely syntactic/axiomatic angle. In this chapter we follow both approaches, starting with the semantic one.

DEFINITION 1.1. A τ -*frame* is a structure $\mathbb{S} = \langle S, R \rangle$ where S is a non-empty set of objects called *states*, *points*, or *worlds*, and R is an interpretation assigning an $n + 1$ -ary relation R_∇ on S to every n -ary modal connective ∇ . A *valuation* on \mathbb{S} is a map $V : X \rightarrow \mathcal{P}(S)$ assigning a subset of S to each variable x . A τ -*model* is a structure $\mathbb{M} = \langle S, R, V \rangle$ such that $\langle S, R \rangle$ is a τ -frame, on which V is a valuation; the frame $\langle S, R \rangle$ is called the *underlying frame* of \mathbb{M} .

The notion of *truth* is defined by formula induction. The set of points where φ is true will always be denoted as $\llbracket \varphi \rrbracket$.

DEFINITION 1.2. Given a τ -model \mathbb{M} , we define by induction when a formula φ is *true* at a state s of \mathbb{M} , notation: $\mathbb{M}, s \Vdash \varphi$:

$$\begin{array}{ll} \mathbb{M}, s \Vdash x & \text{if } s \in V(x), \\ \mathbb{M}, s \Vdash \top & \text{always,} \\ \mathbb{M}, s \Vdash \perp & \text{never,} \\ \mathbb{M}, s \Vdash \neg \varphi & \text{if } \mathbb{M}, s \not\Vdash \varphi, \\ \mathbb{M}, s \Vdash \varphi \wedge \psi & \text{if } \mathbb{M}, s \Vdash \varphi \text{ and } \mathbb{M}, s \Vdash \psi, \\ \mathbb{M}, s \Vdash \varphi \vee \psi & \text{if } \mathbb{M}, s \Vdash \varphi \text{ or } \mathbb{M}, s \Vdash \psi, \\ \mathbb{M}, s \Vdash \nabla(\varphi_1, \dots, \varphi_n) & \text{if } R_\nabla s s_1 \dots s_n \text{ for some } s_1, \dots, s_n \text{ such that } \mathbb{M}, s_i \Vdash \varphi_i \text{ for each } i. \end{array}$$

We write $\mathbb{M} \Vdash \varphi$ if φ is true *throughout* \mathbb{M} , that is, true at every state of \mathbb{M} .

DEFINITION 1.3. Given a τ -frame \mathbb{S} , we say that a modal formula φ is *valid* in \mathbb{S} , notation: $\mathbb{S} \Vdash \varphi$ if φ is true throughout any model based on \mathbb{S} . Similarly standard definitions apply to sets of formulas and classes of frames.

Using the notation $Q[s] := \{t \mid Qst\}$ for any binary relation Q , we define the relation $R_{\mathbb{S}}$ such that $R_{\mathbb{S}}[s]$ consists of those points that can be reached from s in one step using any of the accessibility relations, and R^ω as the reflexive and transitive closure of $R_{\mathbb{S}}$.

We may extend the interpretation R of a τ -frame $\mathbb{S} = \langle S, R \rangle$ to the compound diamonds by putting

$$\begin{aligned} R_v &:= Id (= \{(s, s) \mid s \in S\}) \\ R_{\diamond_{(\nabla, i)} \blacklozenge} &:= \{(s, s_i) \mid R_{\nabla} s s_1 \cdots s_n \text{ for some } s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n \in S\} \circ R_{\blacklozenge}, \\ R_{\diamond \blacklozenge} &:= R_{\diamond} \circ R_{\blacklozenge}, \\ R_{\blacklozenge_1 \vee \blacklozenge_2} &:= R_{\blacklozenge_1} \cup R_{\blacklozenge_2}. \end{aligned}$$

It is then straightforward to verify, in any frame \mathbb{S} , that $R^\omega = \bigcup_{\blacklozenge \in CD(\tau)} R_{\blacklozenge}$, and that for any valuation V it holds that

$$\mathbb{S}, V, s \Vdash \blacklozenge \varphi \text{ iff } \mathbb{S}, V, t \Vdash \varphi \text{ for some } t \text{ with } R_{\blacklozenge} s t.$$

That is, compound diamonds indeed behave like diamonds.

Frames and models do not exist in isolation. Given two τ -frames \mathbb{S} and \mathbb{S}' , a map $\theta : S \rightarrow S'$ is called a *bounded morphism* from \mathbb{S} to \mathbb{S}' , notation: $\theta : \mathbb{S} \rightarrow \mathbb{S}'$, if θ satisfies the following conditions for all $\nabla \in \tau$:

(forth) $R_{\nabla} s s_1 \dots s_n$ only if $R'_{\nabla} \theta(s) \theta(s_1) \dots \theta(s_n)$, and

(back) $R'_{\nabla} \theta(s) s'_1 \dots s'_n$ only if there are s_1, \dots, s_n such that $R_{\nabla} s s_1 \dots s_n$ and $\theta(s_i) = s'_i$ for each i .

We let Fr_τ denote the category with τ -frames as objects and bounded morphisms as arrows.

If such a bounded morphism θ is surjective, we call \mathbb{S}' a *bounded morphic image* of \mathbb{S} , notation: $\mathbb{S} \twoheadrightarrow \mathbb{S}'$; if θ is injective we write $\mathbb{S} \hookrightarrow \mathbb{S}'$ and call the subframe of \mathbb{S}' based on the image $\theta[S]$ a *generated subframe* of \mathbb{S}' . We leave it for the reader to verify that the structure $\langle T, R \upharpoonright T \rangle$ (where $R \upharpoonright T$ maps each $\nabla \in \tau$ to the restriction of R_{∇} to T) is a generated subframe of \mathbb{S} if and only if T is a *hereditary* subset of S , that is, if $t \in T$ then $R_{\nabla} t t_1 \dots t_n$ implies that all the t_i belong to T . Given a point r in \mathbb{S} , we denote with \mathbb{S}_r the least generated subframe containing r ; the domain of this subframe is thus the set $R^\omega[r]$. If $\mathbb{S} = \mathbb{S}_r$ we call r a *root* of \mathbb{S} , and say that \mathbb{S} is *rooted*. Finally, given a *family* $\{\mathbb{S}_i \mid i \in I\}$ of τ -frames, we define its *disjoint union* $\sum_{i \in I} \mathbb{S}_i$ as the structure $\langle \sum_{i \in I} S_i, \{R_{\nabla} \mid \nabla \in \tau\} \rangle$, where the domain $\sum_{i \in I} S_i = \bigcup_{i \in I} \{i\} \times S_i$ is the disjoint union of the domains S_i , and the relation R_{∇} is given by $R_{\nabla}(i, s)(i_1, s_1) \dots (i_n, s_n) : \iff i = i_1 = \dots = i_n$ and $(R_i)_{\nabla} s s_1 \dots s_n$.

REMARK 1.4. More general than Kripke frames are the neighborhood frames, which we now review very briefly, and for the basic modal similarity type only. The reader can find more details on these structures in Chapter 1 of this volume. A *neighborhood frame* is a structure $\mathbb{S} = \langle S, \sigma \rangle$ with $\sigma : S \rightarrow \mathcal{P}\mathcal{P}(S)$; such a structure is called *monotone* if $\sigma(s)$

is upwards closed for all $s \in S$, that is, $X \in \sigma(s)$ and $X \subseteq Y$ imply $Y \in \sigma(s)$. Elements of $\sigma(s)$ are called *neighborhoods* of s , and the semantics of the modality ∇ (we will not use \diamond and \square in this context) in a *neighborhood model* $\mathbb{M} = \langle S, \sigma, V \rangle$ with $V : X \rightarrow \mathcal{P}(S)$ a valuation is given by

$$\mathbb{M}, s \Vdash \nabla \varphi \text{ if } \llbracket \varphi \rrbracket \in \sigma(s), \quad (1.1)$$

that is, $\nabla \varphi$ holds at s iff s has a neighborhood of φ -points. Both the box and the diamond interpretation in Kripke models follow the pattern of (1.1): take $\sigma_\diamond(s) = \{A \subseteq S \mid A \cap R[s] \neq \emptyset\}$, and $\sigma_\square(s) = \{A \subseteq S \mid R[s] \subseteq A\}$, respectively.

A map $f : S \rightarrow S'$ is a *neighborhood morphism* between two neighborhood frames $\langle S, \sigma \rangle$ and $\langle S', \sigma' \rangle$ if for all $s \in S$ and all $X' \subseteq S'$ it holds that $X' \in \sigma'(fs)$ iff $f^{-1}[X'] \in \sigma(s)$.

Now we turn to the more syntactic approach towards modal logic. We identify logics with sets of theorems — the more general approach based on *consequence relations* will be discussed in Chapter 8 of this book. A *modal τ -logic* is then a set $L \subseteq Fma_\tau$ which (i) contains all classical propositional tautologies, and (ii) is closed under the derivation rules (MP) of *Modus Ponens* (if both φ and $\varphi \rightarrow \psi$ belong to L then so does ψ), and (US) of *uniform substitution* (if φ belongs to L then so do all of its substitution instances). If a formula φ belongs to a modal logic L then we say that φ is a *theorem* of L , notation: $\vdash_L \varphi$.

A modal logic is called *classical* if it is closed under the following rule: $\vdash_L \varphi_i \leftrightarrow \psi_i \Rightarrow \vdash_L \nabla(\varphi_1, \dots, \varphi_n) \leftrightarrow \nabla(\psi_1, \dots, \psi_n)$; *monotone* if it is closed under $\vdash_L \varphi_i \rightarrow \psi_i \Rightarrow \vdash_L \nabla(\varphi_1, \dots, \varphi_n) \rightarrow \nabla(\psi_1, \dots, \psi_n)$; and *normal* if it contains in addition, for each $\nabla \in \tau$, the axioms $\neg \nabla^i \perp$ and $\nabla(\bar{p}, q \vee q', \bar{r}) \rightarrow \nabla(\bar{p}, q, \bar{r}) \vee \nabla(\bar{p}, q', \bar{r})$ where \bar{p} and \bar{r} denote arbitrary sequences of propositional variables of combined length $ar(\nabla) - 1$. We leave it as an exercise for the reader to verify that this definition coincides with the standard one in the case of basic modal logic.

The *minimal* classical, monotone and normal modal logics for a similarity type τ are denoted as \mathbf{C}_τ , \mathbf{M}_τ and \mathbf{K}_τ , respectively. Here we use the convention that \mathbf{C} , \mathbf{M} and \mathbf{K} denote the minimal logics for the basic modal similarity type. It is easy to see that the collection of normal modal logics is closed under taking arbitrary intersections and therefore forms a complete lattice under the inclusion ordering. Hence, with any set Γ of τ -formulas we may associate the *least* normal modal τ -logic extending \mathbf{K} and containing all formulas in Γ ; this logic is denoted as $\mathbf{K}_\tau.\Gamma$. We say that this logic is *axiomatized* by Γ , since any theorem in $\mathbf{K}_\tau.\Gamma$ can be obtained as the result of a *derivation* from the axioms of the logic (including formulas in Γ) using its derivation rules. Similar definitions and notation apply to extensions of \mathbf{C}_τ and \mathbf{M}_τ .

The validity relation \Vdash between frames and formulas induces a Galois connection consisting of two maps, *Log* and *Fr*, defined as follows. Given a class \mathcal{C} of frames, $Log(\mathcal{C})$, the *logic of \mathcal{C}* , is the set of modal formulas that are valid in \mathcal{C} . Conversely, given a set Γ of formulas, let $Fr(\Gamma)$ denote the class of frames on which Γ is valid. (We call this a Galois connection because we always have $\mathcal{C} \subseteq Fr(\Gamma)$ iff $\Gamma \subseteq Log(\mathcal{C})$.) The *stable* sets of formulas of this connection, that is, the sets Γ such that $\Gamma = Log(Fr(\Gamma))$ are called (*Kripke*) *complete logics* — we leave it for the reader to verify that such sets are indeed normal modal logics. On the other side, the *stable* frame classes, that is, the ones that are closed under the composition $Fr \circ Log$, are called (*modally*) *definable*. Not all modal logics are Kripke complete (see Chapter 7 of this volume) and not all frame classes are modally definable (see Chapter 1 of this volume).

3 MODAL LOGIC IN ALGEBRAIC FORM

As indicated in the introduction, it is the aim of algebraic logic to study logic by algebraic means. Nowadays, most people will associate modal logic primarily with relational structures, but, as with other branches of logic, the 19th century infancy of modern symbolic modal logic was completely algebraic, see MACCOLL [82]). Somehow during the 20th century however, the traditions of algebraic logic and of modal logic got separated, and for decades proceeded without any interaction whatsoever. In particular, while JÓNSSON & TARSKI [70] introduced not only Boolean algebras with operators and their representation over relational structures, but also the rudiments of canonicity and correspondence theory, this seminal work did not mention modal logic, and it was completely overlooked by modal logicians for many years. This is not to say that algebras were to remain absent from the modal logic tradition — they were introduced by LEMMON [80]. But only in the 1970s, probably with the discovery of the fundamental incompleteness of the relational semantics by THOMASON [102], did universal algebraic (and topological) methods regain importance — as examples we mention BLOK [14], ESAKIA [23], GOLDBLATT [37, 38], and RAUTENBERG [91]. And it would even have to wait until the 1990s before the algebraic and modal traditions would be completely rejoined, with collaborations between modal and algebraic logicians (leading to, for instance, the introduction of the guarded fragment in ANDRÉKA, VAN BENTHEM & NÉMETI [7]), with modal logicians investigating algebras of relations from a modal perspective (MARX & VENEMA [84]), or with algebraic logicians responding to the modal tradition (JÓNSSON [69]). It is from this perspective that the algebraic part of this chapter has been written.

Before we explain how to *algebraize* modal logic using the key structures of Boolean algebras with operators (BAOs), let us first briefly introduce the algebraic perspective on (propositional) logic itself. Think of proposition letters as atomic objects referring to entities called propositions, and of connectives as function symbols to be interpreted as operations on propositions. Then notice the complete analogy between the definitions of formulas and terms, respectively, and already we have worked our way towards one of the key ideas underlying the algebraic approach towards (propositional) logic: *propositional formulas can be seen as algebraic terms denoting propositions*.

DEFINITION 1.5. Given a modal similarity type τ , we define its corresponding algebraic similarity type $Bool_\tau$ simply as the union of τ with the Boolean similarity type $Bool = \{\top, \perp, \neg, \wedge, \vee\}$.

We will use \approx as the equality symbol of this algebraic language; as abbreviations we use $\not\approx$ and \preceq in their standard meaning. Since the standard Boolean symbols are function symbols in this algebraic language, we will not use them to denote Boolean combination of equations. For that purpose we let the symbols $\&$ and \Rightarrow denote conjunction and implication, respectively.

The set $Fma_\tau(X)$ of formulas over a set of variables X can then be identified with the set $Ter_{Bool_\tau}(X)$ of algebraic $Bool_\tau$ -terms over X . More importantly, we may impose *algebraic structure* on formulas.

DEFINITION 1.6. The τ -formula algebra is the structure $\mathbb{F}ma_\tau := \langle Fma_\tau, \{\heartsuit^{\mathbb{F}ma_\tau} \mid \heartsuit \in Bool_\tau\} \rangle$, where for each (Boolean or modal) connective \heartsuit , its interpretation

$$\heartsuit^{\mathbb{F}ma_\tau} : (\varphi_1, \dots, \varphi_n) \mapsto \heartsuit(\varphi_1, \dots, \varphi_n)$$

defines a map of arity $n = ar(\heartsuit)$ on Fma_τ .

As a first advantage of this algebraic point of view, recall that substitutions are completely determined by their values on the variables. Putting this algebraically, for any function σ assigning formulas to variables, the substitution induced by σ is the unique extension $\tilde{\sigma}$ of σ to an *endomorphism* on the formula algebra. More generally, it is easy to see that given an arbitrary algebra \mathbb{A} of type $Bool_\tau$, any assignment mapping variables to elements of the carrier of \mathbb{A} has a unique extension $\tilde{\alpha}$ which is a *homomorphism* from Fma_τ to \mathbb{A} . That is, we have the following result.

PROPOSITION 1.7. *Fma_τ is the ω -generated absolutely free algebra of the similarity type $Bool_\tau$.*

Logical languages may now be interpreted in many different kinds of algebras; but of course, we are only interested in structures that can plausibly be viewed as algebras of propositions.

EXAMPLE 1.8. Consider the *truth value* algebra $\mathbb{2}$ of the Boolean similarity type. Its carrier is given as the set $2 = \{0, 1\}$ where 0 ('false') and 1 ('true') are the classical truth values, while its interpretation of the Boolean connectives/function symbols is given by the standard *truth tables*. Given a *valuation* $V : X \rightarrow 2$ of truth values to propositional variables, we can simply *compute* the truth value $\tilde{V}(\varphi)$ of any propositional formula φ , using the unique homomorphism $\tilde{V} : Fma_\tau \rightarrow \mathbb{2}$ extending the assignment V . That is, we see another manifestation of the absolute freeness of the formula algebra.

The algebras arising from the relational semantics of modal languages are the so-called complex algebras. (This terminology dates back to the times when subsets of groups were referred to as *complexes* of the group.)

DEFINITION 1.9. Given an $n + 1$ -ary relation R on a set S , define the n -ary map $\langle R \rangle$ on the power set of S by

$$\langle R \rangle(a_1, \dots, a_n) := \{s \in S \mid Rss_1 \dots s_n \text{ for some } s_1, \dots, s_n \text{ with } s_i \in a_i \text{ for all } i\}.$$

The *complex algebra* \mathbb{S}^+ of a τ -frame \mathbb{S} is obtained by expanding the power set algebra $\mathbb{P}(S)$ with operations $\langle R_\nabla \rangle$ for each modal connective ∇ ; that is,

$$\mathbb{S}^+ := \langle \mathbb{P}(S), S, \emptyset, \sim_S, \cap, \cup, \{\langle R_\nabla \rangle \mid \nabla \in \tau\} \rangle. \quad (1.2)$$

Given a frame class \mathbf{C} , we let $\mathbf{Cm}(\mathbf{C})$ denote the class of complex algebras of frames in \mathbf{C} ; conversely, for a class \mathbf{K} of algebras, $\mathbf{Str}(\mathbf{K})$ denotes the class of frames whose complex algebras belong to \mathbf{K} .

REMARK 1.10. More generally, given a neighborhood frame $\mathbb{S} = \langle S, \sigma \rangle$, define the map $\sigma^+ : \mathbb{P}(S) \rightarrow \mathbb{P}(S)$ by $\sigma^+(A) := \{s \in S \mid A \in \sigma(s)\}$, and define \mathbb{S}^+ as the expansion of $\mathbb{P}(S)$ with the operation σ^+ .

From the perspective of complex algebras, a *valuation* is nothing but an *assignment* of variables to elements of the complex algebra of \mathbb{S}^+ . Furthermore, and much more importantly, given a valuation V on a frame \mathbb{S} , a straightforward induction proves that

$$\mathbb{S}, V, s \Vdash \varphi \text{ iff } s \in \tilde{V}(\varphi), \quad (1.3)$$

where $\tilde{V} : Fma_\tau \rightarrow \mathbb{S}^+$ is the unique homomorphism extending V . With the *meaning function* $\llbracket \cdot \rrbracket^{\mathbb{S}, V}$ defined as the function that maps a formula φ to its *extension* $\llbracket \varphi \rrbracket := \{s \in S \mid \mathbb{S}, V, s \Vdash \varphi\}$, what (1.3) reveals is that, in a slogan, *meaning is a homomorphism*:

PROPOSITION 1.11. *Let V be some valuation on a τ -frame \mathbb{S} . Then the meaning function $\llbracket \cdot \rrbracket^{\mathbb{S}, V}$ is the unique homomorphism $\tilde{V} : \mathbb{Fma}_\tau \rightarrow \mathbb{S}^+$ that extends V .*

As a corollary of this, let $\varphi \approx$ denote the equation $\varphi \approx \top$, then we find that for any τ -frame \mathbb{S} , and any τ -formulas φ, ψ :

$$\mathbb{S} \Vdash \varphi \text{ iff } \mathbb{S}^+ \models \varphi \approx \quad \text{and} \quad \mathbb{S} \Vdash \varphi \leftrightarrow \psi \text{ iff } \mathbb{S}^+ \models \varphi \approx \psi, \quad (1.4)$$

that is, the validity of a *formula* in the frame \mathbb{S} corresponds to that of an *equation*¹ in the complex algebra of \mathbb{S} , and vice versa. We have arrived at one of the most fundamental notions of algebraic logic, namely, that of a class of algebras *algebraizing* a logic.

DEFINITION 1.12. Let L be a modal τ -logic, and \mathbb{K} a class of $Bool_\tau$ -algebras. We say that \mathbb{K} *algebraizes* L , if we have

$$\vdash_L \varphi \quad \text{iff} \quad \mathbb{K} \models \varphi \approx, \quad (1.5)$$

$$\mathbb{K} \models \varphi \approx \psi \quad \text{iff} \quad \vdash_L \varphi \leftrightarrow \psi, \quad (1.6)$$

for all formulas/terms φ and ψ .

The point of this definition is to alert the reader that algebraizations constitute stronger links between logics and classes of algebras than the mere existence of a completeness result, as would be expressed by (1.5) on its own. If the class \mathbb{K} algebraizes the modal logic L , then it is not just the case that \mathbb{K} contains all the information of L through the translation $(\cdot)^\approx$, but also, L encodes the full equational theory of \mathbb{K} through the translation mapping an equation $\varphi \approx \psi$ to the formula $\varphi \leftrightarrow \psi$. Furthermore, the second translation is an *inverse* to the first one in the sense that if we translate the formula φ back and forth, the result $\varphi \leftrightarrow \top$ is L -equivalent to φ . Given the Boolean backbone of modal logics, this property holds vacuously, so there is no need to formulate this as an additional clause of the definition.

Also, observe that it immediately follows from the definition that if \mathbb{K} algebraizes L , then so does the variety generated by \mathbb{K} .

REMARK 1.13. The above definition is a specific instantiation of a much wider notion, which is due to BLOK & PIGOZZI [16]. The basic idea of a class of algebras *algebraizing* a logic always involves uniform translations from formulas to equations, and from equations to formulas, that are, modulo equivalence, inverse to each other. But the general case is of course not limited to modal logics, or to logics extending classical propositional logic; also, the translations may be from formulas to sets of equations, and from equations to sets of formulas.

The most important point is however that the natural habitat of the concept is that of *consequence relations* rather than of logics (in our sense of the word, that is, of logics as sets of sentences). In this more general setting, the requirement that the translations are each other's inverse, is expressed on the logical side by means of the consequence relation, and can equivalently be described on the algebraic side using (infinitary versions of) quasi-equations. For more details on modal consequence relations and the way to algebraize them, the reader is referred to Chapter 8 of this volume. For the general theory of algebraizing logics, see CZELAKOWSKI [21] or FONT & JANSANA [26].

¹In the sequel, we will be sloppy about the distinction between a formula and its equational translation, writing for instance $\mathbb{A} \models \varphi$ if we mean $\mathbb{A} \models \varphi \approx$.

In any case, it will be clear that we can already state our first algebraization result, the proof of which is immediate from (1.4):

THEOREM 1.14. *Let \mathbf{C} be a class of τ -frames. Then $\text{Cm}(\mathbf{C})$ algebraizes $\text{Log}(\mathbf{C})$.*

Turning to the algebraization of arbitrary modal logics, we now introduce the key players: *Boolean algebras with operators*, together with some related concepts.

DEFINITION 1.15. Given two Boolean algebras \mathbb{B} and \mathbb{B}' , it is often convenient to call a function $f : B \rightarrow B'$ a *map from \mathbb{B} to \mathbb{B}'* . Such a map is called *monotone* if $a \leq b$ in \mathbb{B} implies $f(a) \leq' f(b)$ in \mathbb{B}' , *normal* if $f(\perp) = \perp'$, and *additive* if² $f(a \vee b) = f(a) \vee' f(b)$, and *multiplicative* if $f(a \wedge b) = f(a) \wedge' f(b)$. We will call an operation $f : B^n \rightarrow B$ an *operator* if it is normal and additive in each of its coordinates.

BAE_τ denotes the class of τ -expanded Boolean algebra, (shortly, τ -BAES), that is, of algebras

$$\mathbb{A} = \langle A, \top, \perp, -, \wedge, \vee, \{\nabla^{\mathbb{A}} \mid \nabla \in \tau\} \rangle$$

with a Boolean *reduct* $\langle A, \top, \perp, -, \wedge, \vee \rangle$ that is indeed a Boolean algebra. \mathbb{A} is called a *monotone τ -expanded Boolean algebra*, or a τ -BAM, if each $\nabla^{\mathbb{A}}$ is a monotone operation, and a *Boolean algebra with τ -operators*, or τ -BAO, if each $\nabla^{\mathbb{A}}$ is an operator. The classes of these algebras are denoted as, respectively, BAM_τ and BAO_τ . In the case of the basic modal similarity type, we speak of *modal algebras* rather than of τ -BAOs; MA denotes the class of these algebras. Given a set Γ of modal τ -formulas, and a class \mathbf{K} of τ -expanded Boolean algebras, we define $\mathbf{K}(\Gamma)$ as the class of algebras in \mathbf{K} that validate the set of equations $\Gamma^\approx := \{\gamma \approx \top \mid \gamma \in \Gamma\}$.

Given two τ -BAES \mathbb{A} and \mathbb{A}' , we call a map $\eta : A \rightarrow A'$ a *Boolean homomorphism* if it is a homomorphism from the Boolean reduct of \mathbb{A} to that of \mathbb{A}' , and a *modal homomorphism* if it is a homomorphism with respect to the modal operations. Thus a homomorphism between two τ -BAES is a map that is both a Boolean and a modal homomorphism. We let BAE_τ , BAM_τ , etc. also denote the category with the τ -BAES, \dots , as objects and the homomorphisms as arrows.

EXAMPLE 1.16. Algebras of the form \mathbb{S}^+ , with \mathbb{S} some τ -frame, are the prime specimens of Boolean algebras with operators. These algebras are sometimes referred to as *concrete* BAOS.

More generally, the complex algebra of a neighborhood frame (see Remark 1.10) is an example of a BAE for the basic modal similarity type; it is easy to see that such an \mathbb{S}^+ belongs to BAM iff \mathbb{S} is a monotone neighborhood frame.

Our terminological convention will be that properties of and notions pertaining to Boolean algebras (such as atomicity, completeness, filters, \dots) apply to an expanded Boolean algebra as they apply to its underlying Boolean algebras.

All of the properties defined in Definition 1.15 can be given in equational form, so all of the classes defined there are in fact *varieties*. In the next section we discuss the algebraic properties of these varieties; let us first see why they are so important from a logical perspective. This can be formulated very concisely.

THEOREM 1.17. *Let Γ be a set of modal τ -formulas. Then $\text{BAE}_\tau(\Gamma)$ algebraizes $\mathbf{C}_\tau.\Gamma$, $\text{BAM}_\tau(\Gamma)$ algebraizes $\mathbf{M}_\tau.\Gamma$, and $\text{BAO}_\tau(\Gamma)$ algebraizes $\mathbf{K}_\tau.\Gamma$. In particular, $\text{MA}(\Gamma)$ algebraizes $\mathbf{K}.\Gamma$.*

²Observe that we write \vee and \vee' rather than $\vee^{\mathbb{A}}$ and $\vee^{\mathbb{A}'}$, respectively; this convention will always apply to the interpretations of the Boolean symbols, and sometimes to the modal connectives as well.

Note that this theorem implies a general, algebraic, *completeness* result: for instance, concerning modal logics in the basic modal similarity type, it states that

$$\vdash_{\mathbf{K},\Gamma} \varphi \text{ iff } \mathbf{MA}(\Gamma) \models \varphi^{\approx}. \quad (1.7)$$

That is to say, φ is a theorem of the logic *axiomatized* by Γ if and only if φ is valid in the class of algebras *defined* by Γ .

The key tool in the *proof* of Theorem 1.17 is played by the so-called *Lindenbaum-Tarski algebra* of a logic. The introduction of this fundamental tool is based on the observation that for all classical modal logics, the notion of logical equivalence is a congruence on the formula algebra.

DEFINITION 1.18. Let L be a modal τ -logic. The relation \equiv_L between formulas is defined by putting $\varphi \equiv_L \psi$ if $\varphi \leftrightarrow \psi$ is an L -theorem.

PROPOSITION 1.19. For any classical modal τ -logic L , the relation \equiv_L is a congruence on the formula algebra $\mathbb{F}\mathbf{ma}_\tau$.

DEFINITION 1.20. Given a modal τ -logic L , we denote the congruence class of the formula χ under the relation \equiv_L by $[\chi]_L$; for a set of formulas Φ , we let $[\Phi]_L$ denote the set $\{[\varphi]_L \mid \varphi \in \Phi\}$. The quotient algebra $\mathbb{F}\mathbf{ma}_\tau/\equiv_L$ is called the *Lindenbaum-Tarski algebra* of L , notation: \mathbb{F}_L .

Note that the elements of the Lindenbaum-Tarski algebra \mathbb{F}_L are the equivalence classes of the relation \equiv_L of the set $F\mathbf{ma}_\tau$. The algebraic operations are defined as follows: $\top^{\mathbb{F}_L} = [\top]_L$, $\perp^{\mathbb{F}_L} = [\perp]_L$, $[\varphi]_L \wedge^{\mathbb{F}_L} [\psi]_L = [\varphi \wedge \psi]_L$, etc. We briefly remind the reader that all of these definitions could be parameterized by making the set X of variables explicit.

It is hard to overestimate the importance of Lindenbaum-Tarski algebras. For a start, the algebra \mathbb{F}_L contains all the information of its logic L , in the following sense.

THEOREM 1.21. Let L be a modal logic for some similarity type τ . Then for any two τ -formulas φ and ψ , we have

$$\mathbb{F}_L \models \varphi \approx \psi \text{ iff } \varphi \equiv_L \psi.$$

Proof. For the direction from left to right, consider the natural assignment $\nu : x \mapsto [x]_L$. It follows from the validity of $\varphi \approx \psi$ in \mathbb{F}_L that $\tilde{\nu}(\varphi) = \tilde{\nu}(\psi)$. But an easy formula induction shows that $\tilde{\nu}(\chi) = [\chi]_L$, for all formulas χ . Hence we obtain that $[\varphi]_L = [\psi]_L$, that is, $\varphi \equiv_L \psi$.

For the reverse direction, let α be some assignment on the Lindenbaum-Tarski algebra. Choose for each variable x a representative $\sigma(x)$ of the equivalence class $\alpha(x)$; that is, for each variable x we have that $\alpha(x) = [\sigma(x)]_L$. Note that this map σ is nothing but a substitution; recall that $\tilde{\sigma}$ is the extension of σ to all formulas. It is not hard to prove that all formulas χ satisfy $\tilde{\alpha}(\chi) = [\tilde{\sigma}(\chi)]_L$. But it follows from $\varphi \equiv_L \psi$ that $\tilde{\sigma}(\varphi) \equiv_L \tilde{\sigma}(\psi)$, since L is closed under uniform substitution. Hence we find that $\tilde{\alpha}(\varphi) = \tilde{\alpha}(\psi)$. And since α was arbitrary, this shows that $\mathbb{F}_L \models \varphi \approx \psi$, as required. ■

On the other hand, Lindenbaum-Tarski algebras play an important algebraic role as well, as is concisely formulated in the following Theorem.

THEOREM 1.22. For any classical modal τ -logic L , \mathbb{F}_L is the ω -generated free algebra for the variety $\mathbf{BAE}_\tau(L)$.

Proof. Let \mathbb{A} be an algebra in $\text{BAE}_\tau(L)$, and consider an arbitrary map $\alpha : [X]_L \rightarrow A$ (recall that X denotes the set of variables, and that $[X]_L = \{[x]_L \mid x \in X\}$). We will prove that α can be extended to a homomorphism from \mathbb{F}_L to \mathbb{A} .

To this aim, consider the composition $\alpha \circ \nu : X \rightarrow A$ of α with the natural map $\nu : x \mapsto [x]_L$. It follows from the universal mapping property of $\mathbb{F}\mathbf{ma}_\tau$ over X that this map can be extended to a homomorphism $\widetilde{\alpha \circ \nu} : \mathbb{F}\mathbf{ma}_\tau \rightarrow \mathbb{A}$.

We claim that $\ker(\widetilde{\nu}) \subseteq \ker(\widetilde{\alpha \circ \nu})$. To see this, consider formulas φ and ψ such that $(\varphi, \psi) \in \ker(\nu)$; then $[\varphi]_L = [\psi]_L$, and so $\varphi \equiv_L \psi$. It follows from \mathbb{A} being in $\text{BAE}_\tau(L)$ that $\mathbb{A} \models \varphi \approx \psi$, so $\varphi \approx \psi$ certainly holds in \mathbb{A} under the assignment $\alpha \circ \nu$. But that is just another way of saying that $(\varphi, \psi) \in \ker(\widetilde{\alpha \circ \nu})$.

But then from this claim it follows that the map $\widetilde{\alpha} : \mathbb{F}\mathbf{ma}_\tau / \equiv_L \rightarrow A$, given by

$$\widetilde{\alpha}([\varphi]_L) := \widetilde{\alpha \circ \nu}(\varphi)$$

is well-defined. It is not hard to show that $\widetilde{\alpha}$ is in fact a homomorphism from \mathbb{F}_L to \mathbb{A} , and since it clearly extends α , we have established the universal mapping property of \mathbb{F}_L for $\text{BAE}_\tau(L)$ over $[X]_L$. \blacksquare

Finally, in order to prove the Algebraization Theorem 1.17 from these two theorems, we need one additional result concerning varieties of the form $\text{BAE}_\tau(L)$ if L is a modal logic axiomatized by a set Γ of formulas. We leave the rather tedious but straightforward proof of this proposition as an exercise for the reader.

PROPOSITION 1.23. *Let Γ be a set of τ -formulas. Then $\text{BAE}_\tau(\mathbf{C}_\tau.\Gamma) = \text{BAE}_\tau(\Gamma)$, $\text{BAE}_\tau(\mathbf{M}_\tau.\Gamma) = \text{BAM}_\tau(\Gamma)$, and $\text{BAE}_\tau(\mathbf{K}_\tau.\Gamma) = \text{BAO}_\tau(\Gamma)$.*

This finishes our introduction to the algebraization of modal logics. In section 6 we will have a lot more to say about the link between normal modal logics and varieties of BAOS.

4 VARIETIES OF EXPANDED BOOLEAN ALGEBRAS

In this section we discuss what the theory of universal algebra has to say about Boolean algebras with operators and their siblings.

Lattices of congruences

A very important theme in universal algebra has been to relate the properties of a variety to the shape of the congruence lattices of its algebras. In the case of Boolean algebras and their expansions, this has turned out to be particularly fruitful.

DEFINITION 1.24. An algebra \mathbb{A} has *permuting congruences* if $\Theta_1 \circ \Theta_2 = \Theta_2 \circ \Theta_1$ for all congruences Θ_1, Θ_2 over \mathbb{A} ; \mathbb{A} is *congruence distributive* if $\text{Cg}(\mathbb{A})$, its lattice of congruences, is distributive.

These properties hold of a variety if they hold of each of its members; and a variety is called *arithmetical* if it is both congruence permutable and congruence distributive.

It is a rather strong property for an algebra to have permuting congruences, or to be congruence distributive, and both notions have important applications. Concerning the second notion, we will see an important property of congruence distributive varieties

in Theorem 1.35. In order to motivate the first concept here we just mention that it allows a considerable simplification in the computation of joins in congruence lattices: whereas in general the join $\Theta_1 \vee \Theta_2$ of two congruences Θ_1 and Θ_2 is given as $\Theta_1 \vee \Theta_2 = \Theta_1 \cup (\Theta_1 \circ \Theta_2) \cup (\Theta_1 \circ \Theta_2 \circ \Theta_1) \cup \dots$, in the case of permuting congruence this rearranges itself as $\Theta_1 \vee \Theta_2 = \Theta_1 \circ \Theta_2$.

THEOREM 1.25. *Varieties of expanded Boolean algebras are arithmetical.*

Proof. This proof can be seen as a consequence of a result by A. Pixley, who proved that a variety is arithmetical if and only if it admits the definition of so-called *Mal'cev* and $\frac{2}{3}$ -majority terms. For some detail, let \mathbf{V} be a variety of expanded Boolean algebras. First consider the ternary (Boolean) term $p(x, y, z)$ given by

$$p(x, y, z) := (x \wedge z) \vee (x \wedge \neg y \wedge \neg z) \vee (\neg x \wedge \neg y \wedge z).$$

We leave it for the reader to verify that this is a *Mal'cev* term, that is,

$$\mathbf{V} \models p(x, x, z) \approx z \text{ and } \mathbf{V} \models p(x, z, z) \approx x. \quad (1.8)$$

From this it follows that \mathbf{V} is congruence permutable: let \mathbb{A} be some algebra in the variety and let $a, b \in A$ be elements such that $(a, b) \in \Theta_1 \circ \Theta_2$ for some congruences Θ_1 and Θ_2 . Then there is some $c \in A$ with $(a, c) \in \Theta_1$ and $(c, b) \in \Theta_2$. From this it follows that $(a, b) \in \Theta_2 \circ \Theta_1$, because

$$a = p^{\mathbb{A}}(a, b, b)\Theta_2 p^{\mathbb{A}}(a, c, b)\Theta_1 p^{\mathbb{A}}(c, c, b) = b.$$

This proves that $\Theta_1 \circ \Theta_2 \subseteq \Theta_2 \circ \Theta_1$ which means that \mathbb{A} has permuting congruences. Congruence distributivity can be proven in a similar way: consider the term M given by

$$M(x, y, z) := (x \vee y) \wedge (y \vee z) \wedge (z \vee x).$$

The reader will have little trouble in showing that

$$\mathbf{V} \models M(x, x, y) \approx M(x, y, x) \approx M(y, x, x) \approx x, \quad (1.9)$$

i.e., M is a $\frac{2}{3}$ -majority term. In a similar way as above we can then use (1.9) to show \mathbf{V} is congruence distributive. ■

Congruences and filters

One of the nicest features of BAOs is that their congruences can be characterized by certain *subsets* of the algebra.

DEFINITION 1.26. Let \mathbb{B} be a Boolean algebra. A subset $F \subseteq B$ is called a (*Boolean*) *filter* of \mathbb{B} if it (i) contains the top element of \mathbb{B} , (ii) is closed under taking meets (that is, if $a, b \in F$ then $a \wedge b \in F$), and (iii) is an *up-set* (that is, $a \in F$ and $a \leq b$ imply $b \in F$). A filter F is *proper* if it does not contain the bottom element of \mathbb{B} , or equivalently, if $F \neq B$. We let $Fi(\mathbb{B})$ denote the collection of filters of \mathbb{B} .

EXAMPLE 1.27. It is not difficult to see that the set $Fi(\mathbb{B})$ is closed under taking intersections; hence, we may speak of the *smallest* filter F_D containing a given set $D \subseteq B$; this filter can also be defined as the following set

$$F_D = \{\top\} \cup \{b \in B \mid \text{there are } d_1, \dots, d_n \in D \text{ such that } d_1 \wedge \dots \wedge d_n \leq b\},$$

which explains why we also refer to this set as the filter *generated by* D . In case that D is a singleton $\{a\}$, we write $a\uparrow$ for $F_{\{a\}}$; this set is called the *principal filter generated by* a . Clearly we have $a\uparrow := \{b \in \mathbb{B} \mid a \leq b\}$.

The filter F_D is proper iff D has the so-called *finite meet property* (that is, $\bigwedge D_0 > \perp$ for all finite subsets $D_0 \subseteq D$).

DEFINITION 1.28. Let \mathbb{A} be a BAO; a subset $F \subseteq A$ is a *modal* or *open filter* of \mathbb{A} if F is a filter of (the underlying BA of) \mathbb{A} which is closed under the application of boxes; that is, $a \in F$ implies $\Box_i a \in F$ for all boxes \Box_i . (If the language has modalities of arity higher > 1 , we need to strengthen this to requiring that F is closed under the application of induced boxes.) The collection of modal filters of \mathbb{A} is denoted as $MF\mathfrak{i}(\mathbb{A})$.

In any BAO \mathbb{A} , the sets $\{\top^{\mathbb{A}}\}$ and A are modal filters; the singleton $\{\top^{\mathbb{A}}\}$ is called the *trivial* (modal) filter of \mathbb{A} , and any filter different from A is called *proper*.

The following theorem will prove to be extremely useful.

THEOREM 1.29. *Let \mathbb{A} be a Boolean algebra with operators. Then*

1. *the collection $MF\mathfrak{i}(\mathbb{A})$ is closed under taking arbitrary intersections and hence forms a complete lattice with respect to the subset ordering;*
2. *this lattice is isomorphic to the congruence lattice of \mathbb{A} through the isomorphism $\Pi : MF\mathfrak{i}(\mathbb{A}) \rightarrow Cg(\mathbb{A})$ given by*

$$\Pi_M := \{(a, b) \in A \times A \mid a \leftrightarrow b \in M\},$$

and its inverse $N : Cg(\mathbb{A}) \rightarrow MF\mathfrak{i}(\mathbb{A})$ by

$$N_{\Theta} := \{a \in A \mid (a, \top) \in \Theta\}.$$

It follows from the completeness of the lattice of modal filters of a BAO \mathbb{A} , that with each set $D \subseteq A$ we may associate the *smallest* modal filter M_D including D . The following proposition explains why we also refer to M_D as the modal filter *generated by* D :

PROPOSITION 1.30. *Let \mathbb{A} be a Boolean algebra with τ -operators, and D a subset of A . Then*

$$M_D = \{a \in A \mid \blacksquare_1 d_1 \wedge \dots \wedge \blacksquare_n d_n \leq a \text{ for some } \blacksquare_1, \dots, \blacksquare_n \in CD(\tau), d_1, \dots, d_n \in D\}.$$

In particular, when D is a singleton, say, $D = \{d\}$, we find

$$M_d = \{a \in A \mid \blacksquare d \leq a \text{ for some } \blacksquare \in CD(\tau), d \in D\}.$$

Subdirect irreducibility

We now turn to the algebraic notion of subdirect irreducibility, which plays an important role in the analysis of varieties. The motivation for introducing this concept is the search for the universal algebraic analogon of the prime numbers, as it were. That is, we want to isolate a class of algebraic building blocks that are basic in the sense that (i) every algebra may be decomposed into basic ones, while (ii) the basic ones themselves only

allow trivial decompositions. Now there are various interpretations of the words ‘basic’ and ‘decomposition’.

DEFINITION 1.31. An algebra \mathbb{A} is *simple* if its only homomorphic images are \mathbb{A} itself and the trivial algebra of its similarity type, and *directly indecomposable* if in any decomposition $\mathbb{A} \cong \prod \mathbb{A}_i$, \mathbb{A} is isomorphic to one of the \mathbb{A}_i .

Both of these notions are important and interesting, but neither one is exactly what we want. The notion of simplicity is too restrictive since not every variety is generated by its simple members. And, whereas every finite algebra is isomorphic to a direct product of directly indecomposable algebras, this does not hold for all infinite algebras. For instance, it is not hard to see that the algebra $\mathfrak{2}$ of Example 1.8 is the only nontrivial directly indecomposable Boolean algebra, while a straightforward cardinality argument shows that no countably infinite algebra can be isomorphic to a direct power of $\mathfrak{2}$.

Hence, in order to meet our criteria, we arrive at a notion which at first sight may seem somewhat involved. In words, an algebra is *subdirectly irreducible* iff it does not allow a proper *subdirect* decomposition.

DEFINITION 1.32. Let \mathbb{A} be an algebra, and $\{\mathbb{A}_i \mid i \in I\}$ a family of algebras of the same type. An embedding η of \mathbb{A} into $\prod_{i \in I} \mathbb{A}_i$ is called *subdirect* if $\pi_i \circ \eta$ is surjective for each projection function π_i . If \mathbb{A} is a subalgebra of $\prod_{i \in I} \mathbb{A}_i$, then we say that \mathbb{A} is a *subdirect product* of the family $\{\mathbb{A}_i \mid i \in I\}$, or that the family forms a *subdirect decomposition* of \mathbb{A} , if the inclusion map is a subdirect embedding.

\mathbb{A} is called *subdirectly irreducible*, or, briefly, *s.i.*, if for every subdirect embedding $\eta : \mathbb{A} \rightarrow \prod_{i \in I} \mathbb{A}_i$ there is an $i \in I$ such that $\pi_i \circ \eta : \mathbb{A} \rightarrow \mathbb{A}_i$ is an isomorphism.

In practice, one always uses a nice characterization of subdirect irreducibility in terms of the congruence lattice of the algebra, and similarly for simple and directly indecomposable algebras. For the proof of this proposition we refer to any standard textbook on universal algebra. For a proper understanding of its formulation, recall that any algebra \mathbb{A} always has at least two congruences: the *diagonal* relation $\Delta_A = \{(a, a) \mid a \in A\}$, and the *global* relation $\Upsilon_A = A \times A$.

PROPOSITION 1.33. *Let \mathbb{A} be an algebra. Then*

1. \mathbb{A} is simple iff $Cg(\mathbb{A}) = \{\Delta, \Upsilon\}$;
2. \mathbb{A} is directly indecomposable iff there are no two congruences Θ_1 and Θ_2 such that $\Theta_1 \wedge \Theta_2 = \Delta$ and $\Theta_1 \circ \Theta_2 = \Upsilon$;
3. \mathbb{A} is subdirectly irreducible iff it has a monolith, that is, a smallest non-diagonal congruence.

The following theorem can be read as stating that, indeed, subdirect irreducibility is the proper concept when it comes to finding the basic building blocks of varieties.

THEOREM 1.34 (Birkhoff). *Every algebra can be subdirectly decomposed into subdirectly irreducible algebras. As a corollary, every variety is generated by its subdirectly irreducible members.*

As a corollary of this theorem, we see that the study of the lattice of subvarieties of a given variety can be conducted by way of inspecting the s.i. members of the variety. In the case of expanded Boolean algebras, the logical meaning of this is that it gives us a

tool for the study of extensions of a given modal logic. For, as we will see in section 6 that the subvarieties of the variety determined by a modal logic, correspond to the extensions of that logic. Also, because expanded Boolean algebras are congruence distributive, we may apply *Jónsson's Lemma*. This result involves the class operations H , S and Pu , which are defined in the appendix.

THEOREM 1.35 (Jónsson). *Let K be a class of algebras such that $\text{Var}(K)$ is congruence distributive. Then all subdirectly irreducible members of $\text{Var}(K)$ belong to $\text{HSPu}(K)$.*

The use of this theorem lies in the fact that if K generates a congruence distributive variety V , then the s.i. members of V still resemble the algebras in K in many ways. For instance, if K is a *finite* set of *finite* members, then $Pu(K) = K$; hence we obtain the following result for finitely generated varieties of expanded Boolean algebras.

COROLLARY 1.36. *Let K be a finite set of finite τ -expanded Boolean algebras. Then $\text{Var}(K)$ only has finitely many subvarieties, each of which is determined by a subset of $\text{HS}(K)$.*

Finally, restricting our attention to Boolean algebras with operators, we encounter yet another nice property, namely that we can characterize subdirect irreducibility of an algebra by the existence of one *single* element — one with rather special properties, that is.

DEFINITION 1.37. An element e of a BAO \mathbb{A} is called *essential* or an *opremum* if $e < \top$, while for all $b < \top$ there is a compound modality \blacklozenge such that $\blacklozenge b \leq e$. Dually, we say that an element ρ is *radical* in \mathbb{A} , or a *radix* of \mathbb{A} , if $\rho > \perp$, while for all $a > \perp$ there is a compound modality \blacklozenge such that $\rho \leq \blacklozenge a$.

Clearly, an element e of a BAO is essential iff its complement $\neg e$ is radical. In the sequel this fact will be used implicitly, context deciding which formulation is the most convenient.

EXAMPLE 1.38. Let \mathbb{S} be a *rooted* frame with root r . It is easy to see that the singleton $\{r\}$ is radical in \mathbb{S}^+ : let $a \subseteq S$ be a nonempty element of \mathbb{S}^+ . Take an element s from a ; since r is a root of \mathbb{S} , there must be some compound modality \blacklozenge such that $R_{\blacklozenge}rs$; from this it is immediate that $\{r\} \subseteq \langle R_{\blacklozenge}a \rangle$.

The following theorem (or at least, the more important statement concerning subdirect irreducibility) is due to RAUTENBERG, see for instance [91].

THEOREM 1.39. *Let \mathbb{A} be a nontrivial Boolean algebra with τ -operators. Then \mathbb{A} is simple iff every non-top element of A is essential, and subdirectly irreducible iff it has an essential element.*

Proof. It follows immediately from Theorem 1.29 that \mathbb{A} is s.i. iff it has a smallest non-trivial modal filter, and it is not hard to see that any such filter is of the form M_e for some element e of \mathbb{A} . The proof of the statement on subdirect irreducibility is thus complete if we can show that for an arbitrary element $e \in A$:

$$M_e \text{ is a smallest nontrivial modal filter iff } e \text{ is essential.} \quad (1.10)$$

First suppose that M_e is a smallest nontrivial modal filter. Since M_e is nontrivial, it follows immediately that $e \neq \top$. In order to show that e is essential, consider an arbitrary element $a < \top \in A$, and consider the filter M_a generated by a . It follows

from our assumption on M_e that $M_e \subseteq M_a$, so that $e \in M_a$. Hence we may deduce from Proposition 1.30 that there is some compound modality \blacklozenge such that $\blacksquare a \leq e$. This suffices to prove that e is essential.

For the converse direction, suppose that e is essential, and let M be an arbitrary nontrivial modal filter on \mathbb{A} . That is, $M \neq \{\top\}$, so M contains an element $a \neq \top$. but then it follows from the essentiality of e that there is some compound modality \blacklozenge such that $\blacksquare a \leq e$; this shows that $e \in M$, whence $M_e \subseteq M$. In other words, M_e is the smallest modal filter on \mathbb{A} .

The proof concerning simplicity is completely similar and therefore left as an exercise. \blacksquare

5 FRAMES AND ALGEBRAS

5.1 Introduction

The algebraic study of modal logic was started in section 3. Its main result, Theorem 1.17, links normal modal logics to varieties of Boolean algebras with operators by stating a general algebraization result. But no matter how well-behaved these algebras are, most modal logicians will still prefer the relational semantics, either because they find it more intuitive, or because frames simply happen to be the structures in which they take an (application driven) interest. Hence there is an obvious need to understand the precise relation between the worlds of frames and algebras, respectively. As we will discuss in this section, much of this relation can be understood within the framework of two dualities, both of which relate algebras to (topological) frames, and one forgetful functor. In order to explain why *two* dualities are needed, it is best to consider *finite* structures first. For the sake of a smooth presentation we confine ourselves to the basic modal language.

Let $\mathbf{FinFram}$ and \mathbf{FinMA} denote the respective categories of finite frames with bounded morphisms, and of finite modal algebras with homomorphisms. Recall that in Definition 1.9 we coded up a frame $\mathbb{S} = (S, R)$ by means of its *complex algebra* \mathbb{S}^+ . Conversely, if $\mathbb{A} = \langle A, \perp, \top, -, \wedge, \vee, \diamond \rangle$ is a finite modal algebra, then we can base a frame on the set $At(\mathbb{A})$ of *atoms* (see Definition 1.40) of \mathbb{A} by putting

$$R_{\diamond}pq : \iff p \leq \diamond q.$$

It is then easy to see that

$$\mathbb{S} \cong (\mathbb{S}^+)_+ \text{ and } \mathbb{A} \cong (\mathbb{A}_+)^+$$

for an arbitrary finite frame \mathbb{S} and an arbitrary finite modal algebra \mathbb{A} . And, with the appropriate extension of the constructions $(\cdot)^+$ and $(\cdot)_+$ to functors, we can in fact establish that

$$(\cdot)^+ \text{ and } (\cdot)_+ \text{ form a dual equivalence between } \mathbf{FinFram} \text{ and } \mathbf{FinMA}. \quad (1.11)$$

Unfortunately, there is no way to remove the restriction to finite structures in (1.11) and obtain a dual equivalence between the categories \mathbf{Fr} and \mathbf{MA} . In fact, since the category \mathbf{MA} has an initial object (the free modal algebra over zero generators), while \mathbf{Fr} does not have a final object (cf. section 10 for details), *no* duality whatsoever can be established between these two categories. However, there *is* a natural way to associate a frame with an arbitrary modal algebra \mathbb{A} , if we let *ultrafilters* generalize the notion of

an atom. That is, we can simply base the *ultrafilter frame* \mathbb{A}_\bullet of \mathbb{A} on the collection of ultrafilters of (the Boolean reduct of) \mathbb{A} by putting

$$R_{\diamond uv} : \iff \diamond a \in u \text{ for all } a \in v.$$

Again, this construction can be extended to a functor $(\cdot)_\bullet$ from \mathbf{MA} to \mathbf{Fr} .

We will see that there is interesting interaction between the functors $(\cdot)^+$ and $(\cdot)_\bullet$. The most important result is the Jónsson-Tarski representation theorem stating that every modal algebra \mathbb{A} can be embedded in its ‘double dual’ $\mathbb{A}^\sigma := (\mathbb{A}_\bullet)^+$. As we will see in the next section, this result lies at the root of the application of algebra in modal completeness results.

While there is no duality between the categories \mathbf{Fr} and \mathbf{MA} , with some modifications, both functors $(\cdot)^+$ and $(\cdot)_\bullet$ do provide interesting dualities. Here there are two basic observations. First, the complex duality functor $(\cdot)^+$ is injective on objects; that is, any frame may be recovered (modulo isomorphism) from its complex algebra. Second, although the functor $(\cdot)_\bullet$ does not have this property (see Example 1.53), there is a simple remedy for this problem, namely, to *add* the missing information, topologically encoded, to the frame \mathbb{A}_\bullet of an algebra \mathbb{A} . Thus we see that two fairly nice dualities can be found if we remove the finiteness constraint on *either* side of the duality (1.11):

- a ‘complex’ or ‘discrete’ duality obtains (see Theorem 1.47) if we consider the entire category on the frame side, and a *subcategory* of *perfect* algebras with *complete* homomorphisms on the other side;
- a ‘topological’ duality obtains (see Theorem 1.67) if, conversely, we keep the category on the algebra side intact, but *add topological structure* on the frame side.

Both dualities restrict to (1.11) in the finite case, and the topological and the complex duality are linked by the functor that *forgets* the topological structure on the frame side. Furthermore, similar results can be proved connecting (monotone) neighborhood frames and (monotone) expanded Boolean algebras. In fact, the picture sketched above applies to far wider contexts [68].

For a brief overview of this section, below we first introduce the above mentioned functors and dualities, in some detail. We then see how the algebraic notions of subdirect irreducibility and simplicity turn up on the other side of this duality. We finish the section with a brief discussion of the interaction of the functors $(\cdot)^+$ and $(\cdot)_\bullet$ with more ‘intrinsic’ constructions on algebras and frames such as products and disjoint unions.

5.2 Complex duality

We have already seen how to transform frames into algebras; we now consider these complex algebras from a more abstract perspective. In order to characterize them among the class of all Boolean algebras with operators, we need some terminology.

DEFINITION 1.40. A Boolean algebra \mathbb{B} is called *complete* if it is complete as a lattice, that is, if every subset X of B has both a meet (or greatest lower bound) $\bigwedge X$ and a join (or least upper bound) $\bigvee X$. \mathbb{B} is called *atomic* if below every non-bottom element of \mathbb{B} there is an *atom*, (i.e., an element p satisfying $\perp < p$ while there is no a such that $\perp < a < p$).

Now let \mathbb{B} and \mathbb{B}' be two Boolean algebras; a map $f : B \rightarrow B'$ is called *completely additive* if it preserves all non-empty joins, that is, if for all non-empty subsets X of B for which $\bigvee X$ exists, it holds that

$$f(\bigvee X) = \bigvee' f[X].$$

An n -ary operation f on a Boolean algebra \mathbb{B} is called a *complete operator* if it preserves all joins in each coordinate (or, equivalently, if it is normal and completely additive in each of its coordinates). Finally, a Boolean algebra with operators is called *perfect* if it is complete and atomic, and all its operators are complete.

The reader can easily verify that all complex algebras are perfect. It is equally easy to see that every finite BAO is perfect, since such an algebra has no infinite joins, and a straightforward induction proves that operators preserve finite joins in each of their arguments. For an example of an operator that is not complete, let S be an infinite set, and define $f : \mathcal{P}(S) \rightarrow \mathcal{P}(S)$ by putting $f(X) = X$ if X is finite while $f(X) = S$ otherwise.

In the very same way as we defined above for finite structures, given a perfect BAO we can define a frame based on the set of atoms of \mathbb{A} . In fact, for the definition to make sense, we only need the BAO to be atomic.

DEFINITION 1.41. Let \mathbb{B} be an atomic Boolean algebra, and f an n -ary operator on \mathbb{B} . Define the $n + 1$ -ary relation Q_f on $At(\mathbb{B})$ by

$$Q_f p_0 p_1 \dots p_n : \iff p_0 \leq f(p_1, \dots, p_n).$$

Given an atomic τ -BAO \mathbb{A} , define its *atom structure* \mathbb{A}_+ as the τ -frame $\mathbb{A}_+ = \langle At\mathbb{A}, \{Q_{\nabla^{\mathbb{A}}} \mid \nabla \in \tau\} \rangle$.

Now that we have ways to turn frames into atomic algebras and vice versa, the natural question is how these constructions interact. The following proposition seems to be folklore.

PROPOSITION 1.42. *Let, for a modal similarity type τ , \mathbb{S} be a τ -frame, and \mathbb{A} an atomic τ -BAO. Then*

1. $\mathbb{S} \cong (\mathbb{S}^+)_+$;
2. $\mathbb{A} \cong (\mathbb{A}_+)^+$ iff \mathbb{A} is perfect.

Proof. Concerning the first part, it is straightforward to verify that the map $\iota : x \mapsto \{x\}$ is the required isomorphism. For the second item, let the map $\epsilon : A \rightarrow \mathcal{P}(At(\mathbb{A}))$ be given by $\epsilon(a) := \{p \in At\mathbb{A} \mid p \leq a\}$. The crucial observation in the proof is that

$$\epsilon \text{ embeds } \mathbb{A} \text{ into } (\mathbb{A}_+)^+ \text{ iff all operations of } \mathbb{A} \text{ are complete.} \quad (1.12)$$

This map is then an isomorphism iff \mathbb{A} is perfect. ■

As we will see now, the link between frames and algebras is not restricted to objects. With the natural definition for morphisms between perfect BAOS, we will see how to turn bounded morphisms between frames into these *complete* BAO homomorphisms, and vice versa.

DEFINITION 1.43. Let \mathbb{A} and \mathbb{A}' be two perfect τ -BAOs. A *complete homomorphism* from \mathbb{A} to \mathbb{A}' is a homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ which preserves *all* meets and joins. That is, for every subset $X \subseteq A$ we have that $\eta(\bigvee X) = \bigvee' \eta[X]$ and $\eta(\bigwedge X) = \bigwedge' \eta[X]$. We let BAO_τ^+ denote the category of perfect Boolean algebras with τ -operators as objects, and complete homomorphisms as arrows.

DEFINITION 1.44. Let \mathbb{S} and \mathbb{S}' two τ -frames. Given a bounded morphism $\theta : \mathbb{S} \rightarrow \mathbb{S}'$, define the map $\theta^+ : \mathcal{P}(S') \rightarrow \mathcal{P}(S)$ by

$$\theta^+(X') := \{s \in S \mid \theta(s) \in X'\}.$$

Conversely, given perfect τ -BAOs \mathbb{A} and \mathbb{A}' and a complete homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$, define the map $\eta_+ : \text{At}(\mathbb{A}') \rightarrow A$, which can be shown to map atoms to atoms, by

$$\eta_+(p') := \bigwedge \{a \in A \mid p' \leq \eta(a)\}.$$

It is our aim to prove that $(\cdot)^+$ and $(\cdot)_+$ form a duality between the categories Fr_τ and BAO_τ^+ . We first show functoriality:

PROPOSITION 1.45. $(\cdot)^+$ is a contravariant functor from Fr_τ to BAO_τ^+ .

Proof. The important issue here is that for any bounded morphism $\theta : \mathbb{S} \rightarrow \mathbb{S}'$, the map θ^+ is a complete homomorphism from \mathbb{S}'^+ to \mathbb{S}^+ . It is easy to see that θ^+ is a complete Boolean homomorphism between the respective power set algebras; in order to prove that it is also a modal homomorphism, it suffices to show that for an $n + 1$ -ary relation R we have

$$\langle R \rangle(\theta^+(X_1), \dots, \theta^+(X_n)) = \theta^+(\langle R' \rangle(X_1, \dots, X_n)) \quad (1.13)$$

in case θ is a bounded morphism with respect to R and R' . Here it is interesting to note that in fact the inclusion \subseteq is equivalent to the *forth property*, and the converse inclusion \supseteq , to the *back property* of θ . In a way, (1.13) can be seen as a piece of evidence that bounded morphisms provide in fact the right kind of morphism between frames. ■

PROPOSITION 1.46. $(\cdot)_+$ is a contravariant functor from BAO_τ^+ to Fr_τ .

Proof. Here the first point is to prove that if $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ is a complete Boolean homomorphism between the perfect τ -BAOs \mathbb{A} and \mathbb{A}' , then η_+ maps atoms of \mathbb{A}' to atoms of \mathbb{A} . To see this, let p' be an atom of \mathbb{A}' ; it suffices to show that $\eta_+(p')$ is *join prime* in \mathbb{A} . That is, we assume that $\eta_+(p') \leq \bigvee X$ for some $X \subseteq A$, and have to show that $\eta_+(p') \leq x$ for some $x \in X$. From the assumption we may derive that

$$p' \leq' \eta \eta_+(p') \leq' \eta(\bigvee X) = \bigvee' \eta[X].$$

Here the first inequality directly follows from the definition of $\eta_+(p')$. But since p' is an atom of \mathbb{A}' , the fact that $p' \leq' \bigvee' \eta[X]$ implies that $p' \leq' \eta(x)$ for some $x \in X$. The definition of $\eta_+(p')$ then immediately gives that $\eta_+(p') \leq' x$.

Unfortunately, we do not have the space here to prove that if η is in addition a modal homomorphism, then η_+ is a bounded morphism, or that the operation $(\cdot)_+$ commutes with function composition, i.e., that $(\theta \circ \eta)_+ = \eta_+ \circ \theta_+$ if $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ and $\theta : \mathbb{A}' \rightarrow \mathbb{A}''$ are complete homomorphisms. ■

The following result, that we will refer to as the *complex duality* for BAOs, is due to THOMASON [103] (for the basic modal logic case).

THEOREM 1.47. *The functors $(\cdot)^+$ and $(\cdot)_+$ constitute a dual equivalence between the categories Fr_τ and BAO_τ^+ .*

Proof. Given the results already established, it suffices to prove that the isomorphisms $\iota_{\mathbb{S}} : \mathbb{S} \rightarrow (\mathbb{S}^+)_+$ and $\epsilon_{\mathbb{A}} : \mathbb{A} \rightarrow (\mathbb{A}_+)^+$, defined in the proof of Proposition 1.42, are *natural*. For instance, concerning ϵ , we have to prove that $\epsilon_{\mathbb{A}'} \circ \eta = (\eta_+)^+ \circ \epsilon_{\mathbb{A}}$ for an arbitrary complete homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$. The reader can easily verify this by a direct calculation. ■

5.3 Ultrafilter frames

Now let us see how to remove the restriction to finite structures on the algebra side of (1.11); our first goal is to represent arbitrary (that is, not necessarily finite or even atomic) algebras by frames. But, given a BAO \mathbb{A} , what to take as the points of a frame representing \mathbb{A} ? This problem of course already appears on the Boolean level, and its solution is provided by Stone's representation theorem. This celebrated piece of mathematics states that every Boolean algebra can be embedded in the set algebra over its *ultrafilters*; let us briefly review the basic facts concerning ultrafilters.

DEFINITION 1.48. Let \mathbb{B} be a Boolean algebra. An *ultrafilter* of \mathbb{B} is a proper filter u such that either a or $\neg a$ belongs to u , for all $a \in B$. The collection of ultrafilters of \mathbb{B} is denoted as $Uf(\mathbb{B})$. Given a set S , we sometimes refer to ultrafilters of the power set algebra of S as *ultrafilters over S* .

EXAMPLE 1.49. Given a set S , and an element $s \in S$, define the *principal* ultrafilter π_s as the set $\{X \subseteq S \mid s \in X\}$. It is straightforward to verify that this set is indeed an ultrafilter over S . More generally, if p is an atom of the BA \mathbb{B} , then the principal filter $p\uparrow = \{a \in B \mid p \leq a\}$ is in fact an ultrafilter; it is in this sense that ultrafilters form a generalization of atoms.

For an example of a non-principal ultrafilter, consider the Boolean algebra of finite and cofinite sets of some infinite set T ; the collection of cofinite subsets of T forms an ultrafilter of this algebra.

As a last example, ultrafilters can be seen to generalize the notion of a *maximal consistent set*. Consider the Lindenbaum-Tarski algebra \mathbb{F}_L of a modal logic L ; it is easy to verify that Φ is a maximal L -consistent set of formulas if and only if the set $\{[\varphi]_L \mid \varphi \in \Phi\}$ is an ultrafilter of \mathbb{F}_L .

Ultrafilters can be characterized as the proper filters that are maximal with respect to the inclusion ordering; this identification provides the key tool for establishing the existence of ultrafilters, as the proof sketch of the following Theorem reveals.

THEOREM 1.50 (Ultrafilter Theorem). *Any proper filter of a Boolean algebra \mathbb{B} can be extended to an ultrafilter of \mathbb{B} .*

Proof. Given a proper filter F , apply Zorn's Lemma to the collection C of proper filters that extend F , and obtain a proper filter u that is maximal in C . It is not hard to prove that u is in fact a maximal proper filter, and from this it easily follows that u is an ultrafilter. ■

Stone's representation theorem suggests to take the collection $Uf(\mathbb{A})$ of a BAO \mathbb{A} as the *domain* of a frame that will represent \mathbb{A} ; for the *accessibility relation* on this ultrafilter frame we will (in the case of the basic modal similarity type) make the ultrafilter v visible from u if there is no explicit information preventing this; that is, if there is no $a \in v$ with $\diamond a \notin u$. For an arbitrary similarity type we have the following definition.

DEFINITION 1.51. Given an n -ary operator f on the Boolean algebra \mathbb{B} , define its *dual relation* R_f as the $n + 1$ -ary relation on $Uf(\mathbb{B})$ given by:

$$R_f u_1 \cdots u_n : \iff f(a_1, \dots, a_n) \in u \text{ for all } a_1 \in u_1, \dots, a_n \in u_n.$$

Now let \mathbb{A} be a Boolean algebra with τ -operators; then we define the *ultrafilter frame* or *canonical structure* of \mathbb{A} as the structure

$$\mathbb{A}_\bullet := \langle Uf(\mathbb{A}), \{R_{\nabla^{\mathbb{A}}} \mid \nabla \in \tau\} \rangle.$$

Given a class \mathbf{K} of algebras, we let $\text{Cst}(\mathbf{K})$ denote the class of ultrafilter frames of algebras in \mathbf{K} .

EXAMPLE 1.52. Recall from Chapter 2 of this volume that the *canonical frame* of a normal modal logic L is the structure $\mathbb{C}_L = \langle C, R \rangle$ where C is the set of maximally L -consistent set of formulas, and (we confine ourselves to the basic modal similarity type) R is the canonical accessibility relation given by $Ruv : \iff \diamond \varphi \in u$ for all φ in v . Using the identification that we made in Example 1.49 of maximal L -consistent sets with ultrafilters of the Lindenbaum-Tarski algebra \mathbb{F}_L , it is fairly obvious that the canonical frame for L is isomorphic to the ultrafilter frame of \mathbb{F}_L .

As a second example of the ultrafilter frame construction we mention that the *ultrafilter extension* \mathbf{ueS} of a frame \mathbb{S} (as defined in Chapter 5 of this volume) is nothing but the 'double dual' $(\mathbb{S}^+)_\bullet$ of \mathbb{S} . Verifying this is simply a matter of unraveling the definitions.

Unlike the complex algebra functor, the ultrafilter frame construction is not injective.

EXAMPLE 1.53. Let A be the collection of finite and cofinite subsets of \mathbb{N} , and let B contain in addition those sets of natural numbers that differ in at most finitely many elements from either the set E of evens or the set O of odds. Both A and B are closed under the Boolean operations, and it is easy to see that A has exactly one non-principal ultrafilter, and B , exactly two: one containing the set E , and one the set O . Now suppose that we create algebras \mathbb{A} and \mathbb{B} by endowing A and B with some dummy operator, say, the identity map. Then we find that the respective ultrafilter frames \mathbb{A}_\bullet and \mathbb{B}_\bullet are isomorphic: both have countably many points, and in both cases, the accessibility relation is simply the diagonal. But the algebras \mathbb{A} and \mathbb{B} are clearly not isomorphic.

As we will see further on, the following theorem from JÓNSSON & TARSKI [70] is not only vital when it comes to applications of the algebraic approach in modal completeness theory. It is also a manifestation of a fundamental mathematical concept, namely that of a *representation* theorem stating that every abstract structure in an axiomatically defined class is in fact isomorphic to a concrete, 'intended' structure of the kind that the axioms try to capture.

THEOREM 1.54 (Jónsson-Tarski Representation Theorem). *Let \mathbb{A} be a Boolean algebra with τ -operators. Then the Stone representation map $\hat{\cdot} : \mathbb{A} \rightarrow \mathcal{P}Uf(\mathbb{A})$ given by*

$$\hat{a} := \{u \in Uf(\mathbb{A}) \mid a \in u\} \tag{1.14}$$

is an embedding of \mathbb{A} into $(\mathbb{A}_\bullet)^+$.

Proof. We omit details concerning the Boolean part of this theorem, which is of course nothing but Stone's representation theorem for Boolean algebras.

Concerning the additional operations, we restrict ourselves to the basic modal similarity type. So we consider a modal algebra $\mathbb{A} = \langle A, \top, \perp, -, \wedge, \vee, \diamond \rangle$ and show that

$$\widehat{\diamond}a = \langle R_\diamond \rangle \widehat{a}. \quad (1.15)$$

First we consider an ultrafilter $u \in \langle R_\diamond \rangle \widehat{a}$. It follows by the definition of $\langle R_\diamond \rangle$ (see 1.9) that there is an ultrafilter v such that $R_\diamond uv$ and $v \in \widehat{a}$, that is, $a \in v$. Then by definition of R_\diamond it follows that $\diamond a \in u$, and, hence, that $u \in \widehat{\diamond}a$. This proves that $\widehat{\diamond}a \subseteq \langle R_\diamond \rangle \widehat{a}$.

For the converse direction, take an arbitrary ultrafilter $u \in \widehat{\diamond}a$; that is, $\diamond a \in u$. We have to come up with an ultrafilter v such that (i) $R_\diamond uv$ and (ii) $v \in \widehat{a}$, or, equivalently, $a \in v$. We first reformulate the first condition:

$$R_\diamond uv \text{ iff } a \in v \text{ for all } a \text{ with } -\diamond -a \in u. \quad (1.16)$$

Hence, by the Ultrafilter Theorem 1.50 it suffices to show that the set $\{x \in A \mid -\diamond -x \in u\} \cup \{a\}$ has the finite meet property, see Example 1.27. In order to prove this, first observe that the set $\{x \in A \mid -\diamond -x \in u\}$ is closed under taking meets — this easily follows from the additivity of \diamond and the fact that u is a filter.

But then it is left to show that $x \wedge a > \perp$ for any $x \in A$ with $-\diamond -x \in u$. Suppose for contradiction that $x \wedge a = \perp$. We obtain $a \leq -x$ so $\diamond a \leq \diamond -x$ by monotonicity of \diamond , and so we find $\diamond -x$ in u because $\diamond a \in u$. This gives the desired contradiction since we already had $-\diamond -x$ in u . ■

DEFINITION 1.55. Given a Boolean algebra with τ -operators \mathbb{A} , the ‘double dual’ algebra $(\mathbb{A}_\bullet)^+$ is known as the *canonical embedding algebra of \mathbb{A}* , the *canonical extension of \mathbb{A}* and the *perfect extension of \mathbb{A}* ; we will mainly use the second term, and usually denote the structure as \mathbb{A}^σ .

The Jónsson-Tarski theorem thus states that the constructions $(\cdot)^+$ and $(\cdot)_\bullet$ interact well if we start with algebras: $\mathbb{A} \mapsto (\mathbb{A}_\bullet)^+$ for every BAO \mathbb{A} . Unfortunately, if we start with frames, then the return is less safe: for a τ -frame \mathbb{S} , the map $s \mapsto \pi_s$ (assigning to points of \mathbb{S} their associated principal ultrafilters) is an embedding of \mathbb{S} into $(\mathbb{S}^+)_\bullet$ only if \mathbb{S} is *image finite*. (In fact, the condition of image-finiteness is also sufficient.) And if \mathbb{S} contains a point from which paths of arbitrary finite length emanate, but no infinite path, then there is *no* bounded morphism from \mathbb{S} to $(\mathbb{S}^+)_\bullet$ whatsoever. From this it follows that there is no way to extend the ultrafilter frame construction to a functor that is *adjoint* to that of taking complex algebras. This is a notable divergence of the case of Boolean algebras per se (that is, without operators) — the formation of the canonical extension \mathbb{B}^σ of a Boolean algebra \mathbb{B} is a *free* construction, see [68] for more information on these matters.

Nevertheless, the operation of taking ultrafilter frames can be extended to a functor, as follows.

DEFINITION 1.56. Let \mathbb{A} and \mathbb{A}' be two Boolean algebras with τ -operators. Given a homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$, we define the map $\eta_\bullet : Uf\mathbb{A}' \rightarrow \mathcal{P}(A)$, which can be shown to map ultrafilters to ultrafilters, by putting

$$\eta_\bullet(u') := \{a \in A \mid \eta(a) \in u'\}. \quad (1.17)$$

PROPOSITION 1.57. $(\cdot)_\bullet$ is a contravariant functor from BAO_τ to Fr_τ .

Proof. If $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ is a Boolean homomorphism, then it follows almost immediately that η_\bullet maps ultrafilters to ultrafilters, while it is not too hard either to prove that, for any modality ∇ of rank, say, n :

$$\begin{aligned} \eta_\bullet \text{ has the forth property for } R_\nabla & \quad \text{if} \quad \nabla'(\eta a_1, \dots, \eta a_n) \leq \eta(\nabla(a_1, \dots, a_n)), \\ \eta_\bullet \text{ has the back property for } R_\nabla & \quad \text{if} \quad \nabla'(\eta a_1, \dots, \eta a_n) \geq \eta(\nabla(a_1, \dots, a_n)). \end{aligned}$$

This shows that η_\bullet is a bounded morphism from \mathbb{A}'_\bullet to \mathbb{A}_\bullet if $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ is a homomorphism. It is then left to show that $(\cdot)_\bullet$ is functorial, and in particular, that $(\eta \circ \theta)_\bullet = \theta_\bullet \circ \eta_\bullet$ for homomorphisms $\theta : \mathbb{A} \rightarrow \mathbb{A}'$ and $\eta : \mathbb{A}' \rightarrow \mathbb{A}''$. This can be checked by a straightforward calculation which we leave for the reader. \blacksquare

5.4 Topological duality

In the previous subsection we encountered a problem of the functor $(\cdot)_\bullet$: in general, algebras cannot be retrieved from their ultrafilter frames. A very simple remedy is then to *add* this information to the frame by melting algebra and frame into one structure. Since this issue already pertains at the level of Boolean algebras (without additional operations), that is where we start the discussion.

DEFINITION 1.58. A *field of sets* is a pair $\langle S, A \rangle$ with $A \subseteq \mathcal{P}(S)$ being closed under all Boolean set-theoretic operations, or equivalently, with A such that $\langle A, S, \emptyset, \sim_S, \cap, \cup \rangle$ is a subalgebra of $\mathbb{P}S$. The elements of A are called the *admissible* subsets of S .

Given a Boolean algebra $\mathbb{A} = \langle A, \top, \perp, -, \wedge, \vee \rangle$, put $\widehat{A} := \{\widehat{a} \subseteq \text{Uf}(\mathbb{A}) \mid a \in A\}$, with $\widehat{\cdot}$ as in (1.15), and define $\mathbb{A}_* := \langle \text{Uf}(\mathbb{A}), \widehat{A} \rangle$ as the *associated field of sets* of \mathbb{A} . Conversely, the *associated Boolean algebra* of a field of sets $\mathbb{S} = \langle S, A \rangle$ is the structure $\mathbb{S}^* := \langle A, S, \emptyset, \sim_S, \cap, \cup \rangle$.

It will be clear that the Boolean algebras \mathbb{A} and $(\mathbb{A}_*)^*$ will always be isomorphic; however, we will only have that $\mathbb{S} \cong (\mathbb{S}^*)^*$ if \mathbb{S} has some special properties.

DEFINITION 1.59. A field of sets $\mathbb{S} = \langle S, A \rangle$ is *discrete* if A contains all singletons of S , *differentiated* if for any two distinct points $s \neq t$ of S there is a set $a \in A$ such that $s \in a$ and $t \notin a$, and *full* if $A = \mathcal{P}(S)$. \mathbb{S} is *compact* if every subset of A with the finite intersection property has a non-empty intersection, and *descriptive* if it is both differentiated and compact.

In a descriptive field of sets, the points and the admissible sets are in *balance*: there are sufficiently many admissible sets to separate distinct points, while there are enough points to witness all the ultrafilters of the algebra. More precisely, one can prove that for any field of sets $\mathbb{S} = \langle S, A \rangle$, the map

$$s \mapsto \{a \in A \mid s \in a\} \tag{1.18}$$

provides an bijection between \mathbb{S} and the collection of ultrafilters of \mathbb{S}^* iff \mathbb{S} is descriptive.

REMARK 1.60. Our terminology strongly suggests a topological connection. In order to make this explicit, note that the collection of admissible sets of a field of sets $\mathbb{S} = \langle S, A \rangle$ forms a basis for a topology σ_A ; and that, conversely, we may take the set $\text{Clp}(\mathbb{X})$ of *clopen* (that is, closed and open) elements of a topology $\mathbb{X} = (X, \tau)$ as a collection of

admissible sets. In accordance with this, we define a subset $X \subseteq S$ to be *open* if it is a union of admissible sets, and *closed* if it is an intersection of admissible sets. Thus the study of fields of sets takes us into a rather specific branch of set-theoretic topology in which all spaces are *zero-dimensional*, that is, have a basis of clopens.

One may prove for any field of sets $\mathbb{S} = \langle S, A \rangle$ that \mathbb{S} is descriptive iff $\langle S, \sigma_A \rangle$ is a *Stone space*, that is, σ_A is a compact, Hausdorff and zero-dimensional topology. Basically then, descriptive fields of sets and Stone spaces are two ways of formulating the same mathematical objects; the difference is no more than a matter of focus, be it on the topology itself, or rather on its sets of clopens.

The topological nature also comes out clearly when we discuss *morphisms*.

DEFINITION 1.61. Given two fields of sets $\mathbb{S} = \langle S, A \rangle$ and $\mathbb{S}' = \langle S', A' \rangle$, we call a map $\theta : S \rightarrow S'$ *continuous* if the set

$$\theta^*(a') := \{s \in S \mid \theta(s) \in a'\} \quad (1.19)$$

belongs to A for all $a' \in A'$.

We define the dual $\eta_* : Uf\mathbb{A}' \rightarrow Uf\mathbb{A}$ of a morphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ between two Boolean algebras as the map $\eta_*(u') := \{a \in A \mid \eta(a) \in u'\}$.

Without further proof we mention (our reformulation of) the following seminal result from STONE [101] (see JOHNSTONE [68] for an extensive discussion of its impact).

THEOREM 1.62 (Stone duality). *The functors $(\cdot)^*$ and $(\cdot)_*$ form a dual equivalence between the category of Boolean algebras with homomorphism, and that of descriptive fields of sets with continuous maps.*

The duality for BAOs can now be developed by incorporating the ultrafilter functor $(\cdot)_\bullet$ into the Stone duality: the dual object representing a Boolean algebra with operators will combine the BAO and its dual Kripke frame in one structure.

DEFINITION 1.63. A *general τ -frame* is a structure $\mathbb{G} = \langle G, R, A \rangle$, where $R = \{R_\nabla \mid \nabla \in \tau\}$ is a family of relations on G , such that (i) $\langle G, R \rangle$ is a τ -frame and (ii) $\langle G, A \rangle$ is a field of sets such that (iii) A is closed under the operation $\langle R_\nabla \rangle$ for each operation symbol $\nabla \in \tau$. The structure $\langle G, R \rangle$ is called the *underlying Kripke frame* of \mathbb{G} .

Given a general frame $\mathbb{G} = \langle G, R, A \rangle$, define \mathbb{G}^* as the subalgebra of $\langle G, R \rangle^+$ with carrier A . Conversely, given a τ -BAO \mathbb{A} , define its dual general frame \mathbb{A}_* as the structure $\langle Uf(\mathbb{A}), \{R_{\nabla^*} \mid \nabla \in \tau\}, \hat{A} \rangle$.

As in the case of the duals of Boolean algebras, general frames of the form \mathbb{A}_* are rather special, also with respect to the interaction between their relational and the topological side. We let notions like differentiatedness apply to a general frame $\langle G, R, A \rangle$ as it applies to the underlying field of sets $\langle G, A \rangle$.

DEFINITION 1.64. A general frame $\mathbb{G} = \langle G, R, A \rangle$ is *tight* if every tuple (s, s_1, \dots, s_n) which is *not* in the relation R_∇ (with ∇ an arbitrary relation symbol of arity n) is witnessed by admissible sets a_1, \dots, a_n such that $s_i \in a_i$ for each i , while $s \notin \langle R_\nabla \rangle(a_1, \dots, a_n)$. \mathbb{G} is *refined* if it is both differentiated and tight, and *descriptive* if it is both refined and compact.

REMARK 1.65. An easy proof shows that we may reformulate the property of tightness equivalently by requiring that (restricting to the basic modal language here) $R[s] = \bigcap \{a \in A \mid s \in [R]a\}$ for each point s in $\mathbb{G} = \langle G, R, A \rangle$. In other words, the relation R

is *point-closed*, since each point of \mathbb{G} has a *closed* successor set — closed in the induced topology σ_A , that is. Thus from a topological perspective, descriptive general frames can be identified with point-closed relational Stone spaces.

In order to turn the constructions $(\cdot)^*$ and $(\cdot)_*$ into functors we have to introduce morphisms between (descriptive) general frames as well. Again, we combine modal and topological aspects in the natural way.

DEFINITION 1.66. Given two general frames $\mathbb{G} = \langle G, R, A \rangle$ and $\mathbb{G}' = \langle G', R', A' \rangle$, a map $\theta : G \rightarrow G'$ is called a *continuous bounded morphism* if it is both a bounded morphism from $\langle G, R \rangle$ to $\langle G', R' \rangle$ and a continuous map from $\langle G, A \rangle$ to $\langle G', A' \rangle$. The category of descriptive general τ -frames with continuous bounded morphisms is denoted as DGF_τ .

Now let us see how $(\cdot)^*$ and $(\cdot)_*$ operate on morphisms. For the definition of θ^* for θ a continuous bounded morphism we refer to (1.19); conversely, given a homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ between two τ -BAOs, define η_* as in Definition 1.61, that is: $\eta_*(u') := \{a \in A \mid \eta(a) \in u'\}$. We have now arrived at the main result of this subsection, Theorem 1.67 below, which is due to GOLDBLATT [37, 39]. Independently, ESAKIA [23] came up with such a duality for a more specific variety of algebras.

THEOREM 1.67. *The functors $(\cdot)^*$ and $(\cdot)_*$ constitute a dual equivalence between the categories BAO_τ and DGF_τ .*

Proof. It is rather straightforward to verify that $(\cdot)^*$ and $(\cdot)_*$ are functors which form a dual adjunction between the categories DGF_τ and BAO_τ . It is then left to show that $\mathbb{G} \cong (\mathbb{G}^*)_*$ for any descriptive general frame \mathbb{G} , and that $\mathbb{A} \cong (\mathbb{A}_*)^*$ for any Boolean algebra with τ -operators \mathbb{A} . But both of these claims are easy to establish: for the first isomorphism, take the map of (1.18); and for the second isomorphism, simply take the Stone embedding $\hat{\cdot}$ of (1.14). The proof details are left to the reader. ■

It is straightforward to derive from this duality that for any class \mathbb{C} of general frames, the class of dual algebras algebraizes \mathbb{C} (once we have properly defined all notions involved), but we leave the details for the reader.

5.5 Simplicity and Subdirect irreducibility

As an application of these dualities, let us look at the frame counterparts of the notions of simplicity and subdirect irreducibility. In the complex duality of section 5.2, this question has a satisfactory answer, at least for subdirect irreducibility:

THEOREM 1.68. *Let \mathbb{S} be a τ -frame. Then*

1. \mathbb{S}^+ is simple only if each point is a root of \mathbb{S} ;
2. \mathbb{S}^+ is subdirectly irreducible iff \mathbb{S} is rooted.

Proof. Concerning subdirect irreducibility, the direction from right to left, first mentioned in GOLDBLATT [39], was already treated in Example 1.38. The proof of the converse implication appeared first in SAMBIN [99]. For its details, suppose that ρ is a radix of the algebra \mathbb{S}^+ , and consider an arbitrary point $s \in S$. Then by definition of radicality we find that $\rho \subseteq \langle R_\blacklozenge \rangle \{s\}$ for some compound modality \blacklozenge . It is easy to see

that this implies $R_\blacklozenge rs$ for each $r \in \rho$, so that each element of ρ is in fact a root of \mathbb{S} . Hence, if \mathbb{S}^+ is simple, then every point is a root of \mathbb{S} , since every non-empty subset of S is a radix of \mathbb{S}^+ by Theorem 1.39. If \mathbb{S}^+ is s.i., then by the same theorem it has at least one radix; rootedness of \mathbb{S} thus follows from the fact that radical elements are non-empty by definition. ■

Perhaps contrary to the reader's expectation, the converse of Theorem 1.68(1) is not true.

EXAMPLE 1.69. Consider the frame $\mathbb{Z} = (Z, R)$ for the basic modal similarity type, with Z as the set of integers and Rxy iff $|x - y| = 1$. Then clearly every integer is a root of \mathbb{Z} , while on the other hand, \mathbb{Z}^+ is not simple. An easy way to see this is by proving that the only radical elements of \mathbb{Z}^+ are the finite subsets of \mathbb{Z} .

In the topological duality of section 5.4, the correspondence between subdirect irreducibility and rootedness is not so nice either. In general, subdirect irreducibility of \mathbb{A} neither implies rootedness of \mathbb{A}_* , nor is it implied by it, as the following examples from SAMBIN [99] witness.

EXAMPLE 1.70. For an example of the first kind, take the subalgebra $\mathbb{A} \rightarrow \langle N, > \rangle^+$ based on the collection of finite and cofinite subsets of the set N of natural numbers. As we will see later on, \mathbb{A} is not subdirectly irreducible. However, the frame \mathbb{A}_\bullet is rooted, since it adds one reflexive point ω (corresponding to the ultrafilter of the cofinite sets) to $\langle N, > \rangle$, in such a way that ω sees all other points.

Conversely, consider the frame \mathbb{Z} of the previous example, and take its subalgebra \mathbb{B} based on the finite and cofinite sets. It is easy to see that \mathbb{B} is s.i.: simply note that every singleton is radical. However, the one reflexive point ∞ that \mathbb{B}_* adds to \mathbb{Z} is not related to any other point in \mathbb{B}_* . Hence, \mathbb{B}_* provides an example of an s.i. algebra of which the dual general frame has no roots at all.

These examples indicate that if we are looking for a characterization of the notion of subdirect irreducibility, it does not suffice to look at the dual Kripke frame alone: we have to take the topology into account. Our characterization will be in terms of so-called topological roots or, briefly, topo-roots. Recall that a *root* of a τ -frame $\mathbb{S} = \langle S, R \rangle$ is a point r of \mathbb{S} such that $R^\omega[r] = S$, where the relation R^ω is given as the union of the accessibility relations of the compound diamonds. It is straightforward to verify that in a frame of the form \mathbb{A}_* this boils down to

$$R^\omega uv \text{ iff there is a compound diamond } \blacklozenge \text{ with } \blacklozenge a \in u \text{ for all } a \in v. \quad (1.20)$$

Our definition of the *topo-reachability* relation is obtained by swapping the universal and the existential quantifier in (1.20).

DEFINITION 1.71. Given a Boolean algebra with operators \mathbb{A} , define the *topo-reachability* relation $R^* \subseteq Uf\mathbb{A} \times Uf\mathbb{A}$ as follows:

$$R^* uv \text{ iff for all } a \in v \text{ there is a compound diamond } \blacklozenge \text{ with } \blacklozenge a \in u. \quad (1.21)$$

We let $T_{\mathbb{A}}$ denote the set of *topo-roots* of \mathbb{A}_* ; that is, the collection of those ultrafilters u such that $R^*[u] = Uf\mathbb{A}$.

The topological terminology will be clarified by the following alternative characterization of R^* .

PROPOSITION 1.72. *Let \mathbb{A} be some Boolean algebra with τ -operators, and u some ultrafilter of \mathbb{A} . Then $R^*[u] = \overline{R^\omega[u]}$; that is, $R^*[u]$ is the topological closure of $R^\omega[u]$ in the Stone topology of \mathbb{A}_* .*

As the following theorem from VENEMA [108] witnesses, topo-roots provide the right tool for the characterization of the notions of simplicity and subdirect irreducibility.

THEOREM 1.73. *Let \mathbb{A} be a Boolean algebra with τ -operators. Then*

1. \mathbb{A} is simple iff $T_{\mathbb{A}} = Uf\mathbb{A}$;
2. \mathbb{A} is subdirectly irreducible iff $T_{\mathbb{A}}$ is open and non-empty.

Unfortunately, we do not have the space for a proof or even a proof sketch. We confine ourselves to noting that the proof makes use of the correspondence between modal filters of \mathbb{A} and closed, hereditary subsets of \mathbb{A}_* .

EXAMPLE 1.74. It is now obvious why the algebra \mathbb{A} of Example 1.70 is not s.i.: its dual frame does have a (single) root ω but the set $\{\omega\}$ of roots is not open in the topology of \mathbb{A}_* . The algebra \mathbb{B} of the same example on the other hand *is* s.i. Whereas its dual frame \mathbb{B}_\bullet has no roots at all, almost every point of \mathbb{B}_* is a topo-root.

As corollaries of the last theorem we obtain some (well-)known results showing that in many cases, nicer characterizations are indeed possible. We call a Boolean algebra with operators ω -transitive if it has a *master modality*, that is, a compound diamond \diamond such that $\blacklozenge a \leq \diamond a$ for all compound diamonds \blacklozenge and all a in \mathbb{A} . (With some authors, this property goes under the name of *weak transitivity*). The following result is due to SAMBIN [99] (whereas in the closely related field of intuitionistic logic, similar characterizations of s.i. Heyting algebras in terms of their dual structures had been known for some time, cf. ESAKIA [24]).

COROLLARY 1.75. *Let \mathbb{A} be an ω -transitive Boolean algebra with operators. Then \mathbb{A} is subdirectly irreducible iff the collection of roots of \mathbb{A}_* is non-empty and open.*

Proof. This follows from Theorem 1.73 by the observation that if \mathbb{A} is ω -transitive, then $R^* = R^\omega = R_\diamond$ (where \diamond is the master modality of \mathbb{A}), whence the notions of root and topo-root coincide. ■

Results concerning the duals of *finite* BAOs are already covered by Theorem 1.68, since for finite BAOs the complex and the topological dualities coincide.

5.6 Class operations

While the functors $(\cdot)^+$ and $(\cdot)_\bullet$ do not form a duality, they do provide an interesting link between the categories Fr_τ and BAO_τ . We already discussed the role of the ‘double duals’, that is, the *canonical embedding algebra* $\mathbb{A}^\sigma = (\mathbb{A}_\bullet)^+$ of a BAO \mathbb{A} , and the *ultrafilter extension* $(\mathbb{S}^+)_\bullet$ of a frame \mathbb{S} . But there is also a wealth of results concerning the direct interaction of the mentioned functors with the more ‘intrinsic’ constructions on algebras and frames. We confine ourselves here to the algebraic operations of taking homomorphic images, subalgebras and products, and their frame counterparts of taking generated subframes, bounded morphic images, and disjoint unions. The results listed

in Theorem 1.76 are more or less direct consequences of the dualities established earlier on; therefore, we leave the proofs to the reader.

THEOREM 1.76. *Let \mathbb{S}, \mathbb{S}' and all \mathbb{S}_i with $i \in I$ be τ -frames, and let \mathbb{A}, \mathbb{A}' and all \mathbb{A}_j with $j \in J$ be Boolean algebras with τ -operators. Then*

1. $\theta : \mathbb{S} \rightarrow \mathbb{S}'$ only if $\theta^+ : \mathbb{S}'^+ \rightarrow \mathbb{S}^+$;
2. $\theta : \mathbb{S} \rightarrow \mathbb{S}'$ only if $\theta^+ : \mathbb{S}'^+ \rightarrow \mathbb{S}^+$;
3. $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ only if $\eta_\bullet : \mathbb{A}'_\bullet \rightarrow \mathbb{A}_\bullet$;
4. $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ only if $\eta_\bullet : \mathbb{A}'_\bullet \rightarrow \mathbb{A}_\bullet$;
5. $(\sum_{i \in I} \mathbb{S}_i)^+ \cong \prod_{i \in I} \mathbb{S}_i^+$;
6. $(\prod_{j \in J} \mathbb{A}_j)_\bullet \cong \sum_{j \in J} (\mathbb{A}_j)_\bullet$ if J is finite.

In general it is not true that the ultrafilter frame $(\prod_{j \in J} \mathbb{A}_j)_\bullet$ is isomorphic to the disjoint union $\sum_{j \in J} (\mathbb{A}_j)_\bullet$: the problem is that for infinite J , not every ultrafilter of the product can be linked to an ultrafilter of one of the factors. Fortunately, we do have the following ‘second best’ connection, essentially due to GEHRKE [27], which states that the ultrafilter frame of the product is isomorphic to the disjoint union of the ultrafilter frames of all *ultraproducts* of the original algebras over the index set.

THEOREM 1.77. *Let $\{\mathbb{A}_i \mid i \in I\}$ be a family of Boolean algebras with τ -operators. Then*

$$\left(\prod_{i \in I} \mathbb{A}_i \right)_\bullet \cong \sum_{D \in \text{Uf}(I)} \left(\prod_D \mathbb{A}_i \right)_\bullet.$$

Proof. Given an element a of $\mathbb{A} := \prod_I \mathbb{A}_i$, let $d(a) := \{i \in I \mid a(i) \neq \perp\}$ be the *support set* of a . Then it is not hard to prove that $d[u] := \{d(a) \mid a \in u\}$ is an ultrafilter over I for every $u \in \text{Uf}(\mathbb{A})$.

Now given an ultrafilter D over I , the natural homomorphism $\nu^D : a \mapsto a/D$ is a surjective homomorphism from \mathbb{A} onto $\mathbb{A}_D := \prod_D \mathbb{A}_i$. So by Theorem 1.76(4), its dual $\nu^D_\bullet : (\mathbb{A}_D)_\bullet \rightarrow \mathbb{A}_\bullet$ is a frame embedding. We now claim that

$$\text{Range}(\nu^D_\bullet) = \{u \in \text{Uf} \mathbb{A} \mid d[u] = D\}. \quad (1.22)$$

For the inclusion \subseteq , take an arbitrary ultrafilter z of \mathbb{A}_D . For any $a \in \nu^D_\bullet(z)$, it holds by definition that $\nu^D(a) = a/D$ belongs to z ; but then a/D must be distinct from the bottom element of \mathbb{A}_D . Hence $d(a) \in D$ by definition of d . Since this applies to arbitrary $a \in \nu^D_\bullet(z)$ it follows that $d[\nu^D_\bullet(z)] \subseteq D$. But then we must have equality because both $d[\nu^D_\bullet(z)]$ and D are ultrafilters over I . For the converse inclusion, if $u \in \text{Uf} \mathbb{A}$ satisfies $d[u] = D$, then the set $u_D := \{a/D \mid a \in u\}$ is easily seen to be an ultrafilter of \mathbb{A}_D which satisfies $\nu^D_\bullet(u_D) = u$. This proves (1.22).

Clearly for each ultrafilter D over I , $\text{Range}(\nu^D_\bullet)$ is (the domain of) a generated subframe of \mathbb{A}_\bullet ; it now follows from the fact that $d[u] \in \text{Uf}(I)$ and (1.22) that these subframes are mutually disjoint, but jointly cover the full domain $\text{Uf} \mathbb{A}$ of \mathbb{A}_\bullet . From this the theorem is immediate. ■

On the basis of the Theorems 1.76 and 1.77 we may develop a ‘calculus of class operations’. For instance, letting \mathbf{S}_f denote the operation of taking generated subframes, Theorem 1.76(1) can be read as stating ‘ $\mathbf{CmS}_f \leq \mathbf{HCm}$ ’, meaning that $\mathbf{CmS}_f(\mathbf{C}) \subseteq \mathbf{HCm}(\mathbf{C})$ for every frame class \mathbf{C} . There are many constructions of either frames or algebras that have been investigated, and many results, similar to the Theorems 1.76 and 1.77, have been obtained. The interested reader is referred to work by GOLDBLATT, for instance [40, 41].

Unfortunately, we have only space here for one further example (which will be used in the next section).

PROPOSITION 1.78. *For any class \mathbf{C} of frames, $\mathbf{PuCm}(\mathbf{C}) \subseteq \mathbf{SCmPu}(\mathbf{C})$.*

Proof. Let $\{\mathbb{S}_i \mid i \in I\}$ be a family of τ -frames, and let D be an ultrafilter over I . Define the map $\eta : \prod_I \mathcal{P}(\mathbb{S}_i)/D \rightarrow \mathcal{P}(\prod_I \mathbb{S}_i/D)$ by putting, for s/D in $\prod_I \mathbb{S}_i/D$:

$$s/D \in \eta(a/D) \iff \{i \in I \mid s(i) \in a(i)\} \in D.$$

We leave it for the reader to verify that this is a well-defined embedding of $\prod_I \mathbb{S}_i^+/D$ into $(\prod_I \mathbb{S}_i/D)^+$. ■

We will give one application of the Theorems 1.76 and 1.77 here, more use of these results will be made in the next sections. Theorem 1.79 below, due to GOLDBLATT & THOMASON [47], can be read as a modal dual of Birkhoff’s theorem identifying varieties with equational classes. For a definition of Birkhoff’s theorem from a *coalgebraic* perspective, the reader is referred to section 14.

THEOREM 1.79 (Goldblatt-Thomason Theorem). *Let \mathbf{C} be a class of τ -frames. Then*

1. *if \mathbf{C} is modally definable then it reflects ultrafilter extensions, and is closed under taking bounded morphic images, generated subframes and disjoint unions;*
2. *the converse of (1) holds if \mathbf{C} is closed under taking ultrapowers, (for instance, if \mathbf{C} is elementary).*

Proof. First assume that \mathbf{C} is modally definable; that is, $\mathbf{C} = \mathbf{Fr}(\Gamma)$ for some set Γ of modal τ -formulas (in fact, we may take Γ to be the logic of \mathbf{C} , but this is not relevant now). Now suppose that the frame \mathbb{S}' is the bounded morphic image of some \mathbb{S} in \mathbf{C} . From \mathbb{S} in \mathbf{C} it follows that $\mathbb{S} \Vdash \Gamma$ whence $\mathbb{S}^+ \models \Gamma^\approx$; but at the same time we see that by Theorem 1.76(2), \mathbb{S}'^+ is a subalgebra of \mathbb{S}^+ . Hence also $\mathbb{S}'^+ \models \Gamma^\approx$, so $\mathbb{S}' \models \Gamma$ which immediately implies that \mathbb{S}' belongs to \mathbf{C} . This shows that \mathbf{C} is closed under taking bounded morphic images; the case of generated subframes and disjoint unions is proved similarly.

Now suppose that the ultrafilter extension $\mathbf{ueS} = (\mathbb{S}^+)_\bullet$ belongs to \mathbf{C} . Then $((\mathbb{S}^+)_\bullet)^+ \models \Gamma^\approx$, and so $\mathbb{S}^+ \Vdash \Gamma^\approx$ since \mathbb{S}^+ is a subalgebra of $((\mathbb{S}^+)_\bullet)^+$ by the Jónsson-Tarski Theorem 1.54. But from $\mathbb{S}^+ \models \Gamma^\approx$ it follows that $\mathbb{S} \Vdash \Gamma$ whence \mathbb{S} belongs to \mathbf{C} . This shows that \mathbf{C} reflects ultrafilter extensions, and thus proves part (1).

For the second part, assume that \mathbf{C} enjoys all of the listed closure properties. In order to prove that $\mathbf{C} = \mathbf{Fr}(\mathbf{Log}(\mathbf{C}))$, take an arbitrary frame \mathbb{S} such that $\mathbb{S} \Vdash \mathbf{Log}(\mathbf{C})$. It suffices to show that \mathbb{S} actually belongs to \mathbf{C} .

It follows from $\mathbb{S} \Vdash \text{Log}(\mathbb{C})$ that \mathbb{S}^+ validates the equational theory of the class $\text{Cm}(\mathbb{C})$, and so by Birkhoff's variety theorem \mathbb{S}^+ belongs to the variety $\text{VarCm}(\mathbb{C})$ generated by the class of complex algebras over \mathbb{C} . Then by Tarski's HSP-theorem, \mathbb{S}^+ belongs to $\text{HSPCm}(\mathbb{C})$. That is, for some family $\{\mathbb{F}_i \mid i \in I\}$ of frames in \mathbb{C} , and some algebra \mathbb{A} we have that

$$\mathbb{S}^+ \leftarrow \mathbb{A} \rightarrow \prod_I \mathbb{F}_i^+.$$

Note that $\prod_I \mathbb{F}_i^+ \cong (\sum_I \mathbb{F}_i)^+$ by Theorem 1.76(5), and that $\mathbb{F} := \sum_I \mathbb{F}_i$ belongs to \mathbb{C} . Then using Theorem 1.76(3) and (4) we find that

$$(\mathbb{S}^+)_{\bullet} \rightarrow \mathbb{A}_{\bullet} \leftarrow (\mathbb{F}^+)_{\bullet}.$$

Now it follows by Theorem 90 in Chapter 5 of this volume that $(\mathbb{F}^+)_{\bullet}$ is a bounded morphic image of some *ultrapower* \mathbb{F}^J/D of \mathbb{F} . Then by the various listed closure properties of \mathbb{C} , we show that subsequently, each of the frames \mathbb{F}^J/D , $(\mathbb{F}^+)_{\bullet}$, \mathbb{A}_{\bullet} and $(\mathbb{S}^+)_{\bullet}$ belong to \mathbb{C} . Finally then, also \mathbb{S} belongs to \mathbb{C} since its ultrafilter extension $(\mathbb{S}^+)_{\bullet}$ does so. ■

6 LOGICS AND VARIETIES

This section, which forms the heart of the algebra part of this chapter, discusses the connection between normal modal logics (NMLs) and varieties of BAOS. The main part of the section consists in showing how standard properties of a logic turn up on the algebraic side of the picture, but we start with showing how the *lattice* of normal modal logics is dually isomorphic to that of the varieties of BAOS.

DEFINITION 1.80. Given a normal modal logic L , we say that a normal modal logic L' is a *normal extension* of L simply if $L \subseteq L'$. The lattice of normal extensions of L is denoted as $\text{NExt}(L)$.

We have already seen that with every normal modal τ -logic we may associate a variety $\text{BAO}_{\tau}(L)$ of τ -BAOs. Conversely, every class of these algebras gives rise to a normal modal logic.

DEFINITION 1.81. Given a class \mathbb{K} of Boolean algebras with τ -operators, we define $\text{Log}(\mathbb{K}) := \{\varphi \in \text{Fma}_{\tau} \mid \mathbb{K} \models \varphi\}$.

The following theorem then describes the intimate connection between normal modal logics and varieties of BAOS. Similar results can be proved about arbitrary modal logics and varieties of BAES, and about monotone modal logics and varieties of BAMS.

THEOREM 1.82.

1. The maps $\text{BAO}_{\tau}(\cdot)$ and $\text{Log}(\cdot)$ form a Galois connection, in the sense that for every set Γ of τ -formulas, and every class \mathbb{K} of Boolean algebras with τ -operators, $\Gamma \subseteq \text{Log}(\mathbb{K})$ iff $\mathbb{K} \subseteq \text{BAO}_{\tau}(\Gamma)$.
2. The stable formula sets of this connection are precisely the normal modal τ -logics, while the stable classes of algebras are precisely the varieties of Boolean algebras with τ -operators.
3. Hence, Log is a dual isomorphism between the lattice of subvarieties of BAO_{τ} and the lattice $\text{NExt}(\mathbb{K}_{\tau})$ of normal modal τ -logics

Proof. It is not hard to see the Galois connection, since we have $\Gamma \subseteq \text{Log}(\mathbf{K})$ iff $\mathbb{A} \models \gamma \approx$ for all \mathbb{A} in \mathbf{K} and all $\gamma \in \Gamma$ iff $\mathbf{K} \subseteq \text{BAO}_\tau(\Gamma)$.

Now let Γ be a stable set of formulas of this connection, that is, suppose that $\Gamma = \text{Log}(\text{BAO}_\tau(\Gamma))$; one easily infers that such a Γ must be a normal modal logic. Conversely, if L is a normal modal logic, then $L = \text{Log}(\text{BAO}_\tau(L))$ by the Algebraization Theorem 1.17.

At the other side of the connection, it is immediate from the definition that every class $\text{BAO}_\tau(\Gamma)$ is a variety. Conversely, assume that \mathbf{V} is a variety of τ -BAOs. Then clearly $\mathbf{V} \subseteq \text{BAO}_\tau(\text{Log}(\mathbf{V}))$ since this holds for any class; for the opposite inclusion, by Birkhoff's variety theorem it suffices to show that $\text{BAO}_\tau(\text{Log}(\mathbf{V}))$ validates every equation of \mathbf{V} . So suppose that $\mathbf{V} \models \varphi \approx \psi$; then $\mathbf{V} \models (\varphi \leftrightarrow \psi) \approx \top$ since \mathbf{V} has a Boolean basis; from this it follows that $\varphi \leftrightarrow \psi \in \text{Log}(\mathbf{V})$, whence $\text{BAO}_\tau(\text{Log}(\mathbf{V}))$ validates the equation $(\varphi \leftrightarrow \psi) \approx \top$, by definition. But $\text{BAO}_\tau(\text{Log}(\mathbf{V}))$ also has a Boolean basis, so we find that $\text{BAO}_\tau(\text{Log}(\mathbf{V})) \models \varphi \approx \psi$, as required.

The last part of the theorem is then immediate by the general theory of Galois connections. ■

The dual isomorphism given by Theorem 1.82, linking the lattice of normal modal logics to that of varieties of BAOs, has yielded a wealth of information on modal logics. For instance, universal algebraic theory on *splitting algebras* led algebraically minded modal logicians to strong results on the *degree of Kripke incompleteness* of a modal logic, see for instance BLOK [15]. We will not discuss the lattice of modal logics any further in this chapter, referring the reader to the Chapters 7 and 8 of this volume.

Instead we turn to the question, how standard logical phenomena fit in the algebraic framework presented so far. The answer to this question depends on the issue at stake, so let us consider a number of examples:

completeness is a property not so much of a single logic but rather of a pair of logics. For instance, Kripke completeness of a logic L means that L coincides with the logic of its frame class \mathbf{C} . Algebraically, this corresponds to the fact that the variety $\text{BAO}_\tau(L)$ is generated by the class of complex algebras $\text{Cm}(\mathbf{C})$. More details will be provided in subsection 6.1.

canonicity of a modal logic L has, as we will see in subsection 6.2, an algebraic counterpart in the property of a class of algebras being closed under taking canonical extensions.

correspondence is more about formulas, or equations, than about logics, or varieties of algebras. Nevertheless, it has a clear algebraic meaning: We can say that an equation $s \approx t$ *corresponds, over a frame class \mathbf{C}* to a first-order formula α in the language of frames, if, for all frames \mathbb{S} in \mathbf{C} , we have that $\mathbb{S}^+ \models s \approx t$ iff $\mathbb{S} \models \alpha$.

interpolation is a property of a normal modal logic. In subsection 6.3 we will see that it corresponds to an amalgamation property on the algebraic side.

Let us now move to a more detailed discussion of some of these issues.

6.1 Completeness

As we mentioned already, Theorem 1.17 can be read as a *general algebraic completeness* result. So in this respect the algebraic semantics behaves much better than the relational

one: Classes of Kripke frames are generally not adequate for revealing all distinctions between normal modal logics, see Chapter 7 of this volume for the details. It clearly *means* something for a modal logic to be Kripke complete, so what about the associated algebraic variety? For an answer, recall the notion of a *perfect* BAO from Definition 1.40.

THEOREM 1.83. *A normal modal τ -logic L is (Kripke) complete iff $\text{BAO}_\tau(L)$ is generated by its perfect members.*

Proof. Straightforward by the observation that any variety \mathbf{V} of BAOs is generated by its perfect members iff its equational theory coincides with that of the class $\text{CmStr}(\mathbf{V})$. ■

This inspires the following definition.

DEFINITION 1.84. A variety \mathbf{V} of Boolean algebras with τ -operators is called (*Kripke*) *complete* if \mathbf{V} is generated by its perfect members.

The phenomenon of Kripke incompleteness of normal modal logics is thus algebraically reflected by the fact that many different varieties of BAOs may share the same class of perfect members.

The formulation of Theorem 1.83 strongly suggests that Kripke completeness is only one of a *family* of properties pertaining to normal modal logics. In fact, one may wonder whether varieties of Boolean algebras with operators are generated by those of their members that meet any given constraint. For instance, we might consider varieties that are generated by their *finite* members. Since every finite BAO is perfect this gives a strong version of Kripke completeness that is known on the logical side as the *finite model property* of the logic.

In this respect it is also interesting to see what happens if we consider *weakenings* or *variations* of the notion of perfection. For instance, recall that perfection of a BAO is the conjunction of three properties: atomicity and completeness of the underlying Boolean algebra, and complete additivity of the operators. Hence, we may naturally ask which varieties of BAOs are generated by their atomic members, their complete and completely additive members, etc. Recent investigations have provided answers to some of these questions. First however, we mention a result of BUSZKOWSKI [18] which has been around for almost twenty years already, but which seems to have received little attention. Call a first-order formula or equation in the language of Boolean algebras with operators *modally guarded* if every variable occurs within the scope of a modality.

THEOREM 1.85. *Let \mathbf{V} be a variety of expanded Boolean algebras which is axiomatized by modally guarded equations. Then \mathbf{V} is generated by its atomic members.*

Proof. Given two BAEs \mathbb{A} and \mathbb{A}' , call an embedding $\eta : \mathbb{A} \rightarrow \mathbb{A}'$ *guarded* if for all guarded formulas $\varphi(x_1, \dots, x_k)$, and all $a_1, \dots, a_k \in \mathbb{A}$, it holds that $\mathbb{A} \models \varphi[a_1, \dots, a_k]$ iff $\mathbb{A}' \models \varphi[\eta a_1, \dots, \eta a_k]$. Then

$$\text{every BAE } \mathbb{A} \text{ has a guarded embedding into an atomic BAE.} \quad (1.23)$$

It is straightforward to prove the theorem from (1.23): Any algebra \mathbb{A} in \mathbf{V} can be embedded into an atomic BAE \mathbb{B} that satisfies the same guarded sentences as \mathbb{A} , and thus in particular, also belongs to \mathbf{V} .

For a proof of (1.23), let \mathbb{A} be some τ -expanded Boolean algebra. By the Stone representation theorem, we may assume that for some set X , \mathbb{A} is of the form $\langle A, X, \emptyset, \sim_X, \cup, \cap, \{\nabla^{\mathbb{A}} \mid$

$\nabla \in \tau\}$). In fact, we may assume that every non-empty $a \in A$ is an *infinite* subset of X . (Otherwise, replace X with the set $X \times \omega$ and, using the natural embedding $P \mapsto P \times \omega$ of the power set algebra of X into that of $X \times \omega$, continue with the image of \mathbb{A} under this map.) Now let B be the collection of those subsets b of X that differ in at most finitely many elements from some element of A ; that is,

$$B := \{b \subseteq X \mid (a \cap \sim_X b) \cup (b \cap \sim_X a) \text{ is finite, for some } a \in A\}.$$

It is not hard to see that for every $b \in B$ there is in fact a *unique* element $a \in A$ such that the symmetric difference $(a \cap \sim_X b) \cup (b \cap \sim_X a)$ is finite; this element will be denoted as b^* .

One then easily proves that the structure $\langle B, X, \emptyset, \sim_X, \cup, \cap \rangle$ is an atomic Boolean algebra, so if we define, for $\nabla \in \tau$:

$$\nabla^{\mathbb{B}}(b_1, \dots, b_n) := \nabla^{\mathbb{A}}(b_1^*, \dots, b_n^*),$$

we obtain a τ -expanded Boolean algebra \mathbb{B} . Finally, a straightforward induction on the complexity of guarded formulas shows that the identity map is the required guarded embedding of \mathbb{A} into \mathbb{B} . This proves (1.23). ■

However, the restriction to *guarded* axioms in Theorem 1.85 is essential, as the following result of VENEMA [106] implies that there are varieties of BAOs that have *no* atomic members.

THEOREM 1.86. *There are nontrivial varieties of Boolean algebras with operators of which all members are atomless.*

Proof. The basic idea underlying this proof is straightforward: construct a particular, nontrivial, BAO \mathbb{A} , and a unary term $\pi(x)$ such that the formula $\alpha \equiv \forall x(\perp \prec x \Rightarrow \perp \prec \pi(x) \prec x)$ holds in \mathbb{A} . This shows not only that \mathbb{A} is atomless, but that this atomlessness is witnessed by a *term function*.

Lacking the space for further details concerning the construction of \mathbb{A} , we briefly sketch how to prove the theorem from here. Let \mathbf{K} be the class of BAOs satisfying α . Without loss of generality, assume that \mathbf{K} has a global modality (see section 8.2). It then follows that the class $\mathbf{SP}(\mathbf{K})$ is a variety, and thus, that the formula α , being a universal Horn sentence, holds in every member of this variety. But then every such algebra is atomless, so the theorem follows if we can prove that \mathbf{K} is nontrivial. But this is an immediate consequence of the existence of the algebra \mathbb{A} . ■

Regarding the order/lattice theoretic property of completeness, a similar result obtains, due to LITAK [81].

THEOREM 1.87. *There are nontrivial varieties of Boolean algebras with operators without complete members.*

Proof. Consider the similarity type of *tense logic*, as in section 8.1. Let $\mathbb{S} = \langle \mathbb{N}, \prec \rangle$ be the bidirectional frame of the natural numbers with the standard ordering. That is, we interpret the diamonds \diamond_F and \diamond_P via the relations \prec and \succ , respectively. Furthermore, let \mathbb{A} be the subalgebra of \mathbb{S}^+ based on the collection of finite and cofinite subsets of \mathbb{N} . We claim that $\mathbf{Var}(\mathbb{A})$, the variety generated by \mathbb{A} , has no complete members. Suppose for contradiction that \mathbb{C} is a complete member of $\mathbf{Var}(\mathbb{A})$.

Each natural number n is, inside \mathbb{S} , the unique point satisfying the variable free formula $\varphi_n := \diamond_P^n \top \wedge \square_P^{n+1} \perp$. Observe that the inequalities $\varphi_n \wedge \varphi_m \preceq \perp$ (for $m \neq n$), and $\varphi_n \preceq \diamond_F \varphi_{n+1}$ hold in \mathbb{A} , hence in $\text{Var}(\mathbb{A})$, and therefore, in \mathbb{C} . Define $a_n := \varphi_{2n}^{\mathbb{C}}$ and $b_n := \varphi_{2n+1}^{\mathbb{C}}$. It is then immediate that $a_n \leq \diamond_F b_n$, $b_n \leq \diamond_F a_{n+1}$, and $a_n \wedge b_m = \perp$, for all m, n (we write \diamond_F rather than $\diamond_F^{\mathbb{C}}$). But \mathbb{C} is complete, so it contains elements $a = \bigvee_n a_n$ and $b = \bigvee_n b_n$, for which we easily derive that $a \leq \diamond_F b$, $b \leq \diamond_F a$, and $a \wedge b = \perp$. Hence, from the fact that $\mathbb{C} \models \diamond_F \diamond_F x \preceq \diamond_F x$ it follows that $a \leq \diamond_F a \wedge \diamond_F a$, whence $a \wedge (\square_F a \vee \square_F \neg a) = \perp$. Thus \mathbb{C} refutes the inequality $\diamond_F x \preceq \diamond_F (x \wedge \square_F x \vee \square_F \neg x)$, while a straightforward proof shows this inequality to hold in \mathbb{A} , and hence, in $\text{Var}(\mathbb{A})$. This provides the required contradiction. ■

For more information on such notions of incompleteness that are weaker than Kripke incompleteness, the reader is referred to LITAK [81]. To mention one open problem: it is not known whether an analogue of the previous two results can be proved for the notion of completely additivity.

6.2 Canonicity

In Chapter 2 of this volume, a normal modal logic L is defined to be *canonical* if $\mathbb{C}_L \Vdash L$, where \mathbb{C}_L is the canonical frame for the logic L . In order to put this in an algebraic perspective, first note that $\mathbb{C}_L \Vdash L$ is equivalent to the requirement that $\mathbb{C}_L^+ \models L^{\approx}$. Also, recall from Example 1.52 that the canonical frame for L is isomorphic to the ultrafilter frame of the Lindenbaum-Tarski algebra \mathbb{F}_L . Hence, we see that the issue is whether $(\mathbb{F}_L)^\sigma = ((\mathbb{F}_L)_\bullet)^+ \models L^{\approx}$, whereas we know that $\mathbb{F}_L \models L^{\approx}$, cf. Theorem 1.21. This inspires the following definition.

DEFINITION 1.88. A class of Boolean algebras with τ -operators is *canonical* if it is closed under taking canonical embedding algebras. Accordingly, an equation η is called *canonical* if the variety $\text{BAO}_\tau(\eta)$ is canonical, that is, if $\mathbb{A} \models \eta$ only if $\mathbb{A}^\sigma \models \eta$, for all BAOS \mathbb{A} .

From the definition it is obvious that any normal modal logic is canonical if the variety $\text{BAO}_\tau(L)$ is canonical, but what about the converse implication? Here we need to be a bit more precise about the definition of the canonical frame; in particular, about the size of the set of variables. For, observe that the notion of *maximality* of an L -consistent set of formulas depends on the surrounding set of formulas, and hence, on the set X of variables. Thus the shape of the canonical frame \mathbb{C}_L depends on the *size* of the set X of variables; in order to make this dependence explicit, we will write $\mathbb{C}_L(X)$ for the canonical frame in which the points are maximal L -consistent subsets of $\text{Fma}(X)$. A similar convention applies to Lindenbaum-Tarski algebras. Taking this cardinal subtlety into account, we arrive at a sharpened definition of the logical concept of canonicity.

DEFINITION 1.89. A normal modal logic L is *canonical* if $\mathbb{C}_L(X) \Vdash L$ for *all* sets X . A formula φ is called *canonical* if $\mathbb{C}_L(X) \Vdash \varphi$ for all normal modal logics L containing φ .

Fortunately, we can prove that the logical and the algebraic notion of canonicity coincide.

THEOREM 1.90. *For any normal modal τ -logic L , L is canonical iff $\text{BAO}_\tau(L)$ is a canonical variety.*

Proof. Let \mathbb{A} be an arbitrary algebra in $\text{BAO}_\tau(L)$, and let X be a set containing a separate variable x_a for each $a \in A$. Then \mathbb{A} is a homomorphic image of $\mathbb{F}_L(X)$ by the fact that $\mathbb{F}_L(X)$ is the free algebra for $\text{BAO}_\tau(L)$ over the set $[X]_L$, see Theorem 1.22 for the case of countable X . Now two applications of Theorem 1.76 show that $(\mathbb{F}_L(X))^\sigma \twoheadrightarrow \mathbb{A}^\sigma$. But $(\mathbb{F}_L(X))^\sigma$ belongs to $\text{BAO}_\tau(L)$ by canonicity of L , and so \mathbb{A}^σ is in $\text{BAO}_\tau(L)$ because varieties are closed under taking homomorphic images. ■

It is not known whether, for the variety $\text{BAO}_\tau(L)$ to be canonical, it suffices that the canonical frames for countable variable sets validate L . Leaving this question as an open problem, we turn to the logical *motivation* of the concept of canonicity. This lies in its applications in modal completeness theory, see Chapter 2 of this volume for details. Algebraically, these applications are connected to the following result.

THEOREM 1.91. *Let \mathcal{V} be a variety of Boolean algebras with τ -operators. If \mathcal{V} is canonical, then \mathcal{V} is complete.*

Proof. If \mathcal{V} is canonical then $\mathcal{V} \subseteq \text{SCmCst}(\mathcal{V})$ so clearly \mathcal{V} is generated by its perfect members. ■

So where do we find canonical varieties? In general there seem to be two roads here, a syntactic and a model-theoretic one. The syntactic approach is the most important one for applications. Basically, the idea is to find out whether a logic is canonical on the basis of the syntactic shape of the axioms. Now in general it is *undecidable* whether a given formula φ is canonical (see KRACHT [72] for a proof). Fortunately, however, there are fairly large classes of canonical formulas that occur frequently in practice, and are easily recognized. We confine our attention here to *Sahlqvist formulas* — these are also discussed in the Chapters 1, 5 and 7 of this volume.

In the sequel it will be convenient to assume that the primitive symbols of our language are, besides the Boolean connectives $\top, \perp, \neg, \wedge$ and \vee , and the modalities $\{\nabla \mid \nabla \in \tau\}$, also the implication symbol \rightarrow , and the dual modalities $\{\Delta \mid \nabla \in \tau\}$. Also, recall that boxes are the duals of diamonds, that is, of *unary* modal operators.

DEFINITION 1.92. Given a modal similarity type τ , we define the following classes of terms/formulas. A *boxed atom* is a variable, possibly preceded by a string of boxes. A formula π is *positive* (*negative*) if all of its variables are in the scope of an even (odd, respectively) number of negation symbols. A *Sahlqvist formula* is a formula of the form $\varphi \rightarrow \psi$, where φ is built up from negative formulas, boxed atoms, and constants, using only modalities, \wedge and \vee , while ψ is a positive formula.

The following results are some of the most celebrated general results in modal logic. Theorem 1.93 below, from SAHLQVIST [98], put the crown on the work of many contemporary modal logicians.

THEOREM 1.93 (Sahlqvist Canonicity). *Every Sahlqvist formula is canonical.*

For the *proof*, the reader is referred to section 7. As a corollary of this theorem and the *correspondence* result for Sahlqvist formulas (see Chapter 1 of this volume), we obtain the following.

COROLLARY 1.94. *Let $L = \mathbf{K}_\tau.\Sigma$ be a normal modal logic axiomatized by a collection Σ of Sahlqvist axioms. Then L is sound and complete with respect to the class of frames defined by the first-order correspondents of the formulas of Σ .*

REMARK 1.95. Although the Sahlqvist canonicity theorem takes care of most of the canonical formulas that one encounters in practice, it certainly does not cover the concept completely. For instance, GORANKO & VAKARELOV [49] widen the class to that of so-called *inductive* formulas, see Chapter 5 of this volume for some discussion. JÓNSSON [69] generalizes an example of FINE [25] to the result that for every positive formula $\varphi(x)$, the equation $\varphi(x \vee y) \approx \varphi(x) \vee \varphi(y)$ is canonical. And of course, there are individual examples of canonical formulas, such as the conjunction of the transitivity axiom 4 and the McKinsey axiom $\Box \Diamond x \leq \Diamond \Box x$, cf. [69] for an algebraic proof.

As we mentioned, a second way to arrive at canonical varieties of BAOs proceeds via a model-theoretic road. The basic idea here is that varieties are canonical if they can be *generated* in a certain way. A first and seminal result in this direction was the following.

THEOREM 1.96 (Fine). *If \mathbf{K} is an elementary class of frames, then $\text{Log}(\mathbf{K})$ is a canonical normal modal logic.*

Algebraically, Theorem 1.96 reads that elementary frame classes generate canonical varieties. This result points at an intriguing connection between elementary frame classes and canonical varieties. In particular, it has been an open problem for a long time whether the converse of Fine's theorem would hold as well, that is, whether every canonical variety would be generated by some elementary frame class. Recently however, this issue has been settled negatively in GOLDBLATT, HODKINSON & VENEMA [46].

THEOREM 1.97. *There is a canonical variety that is not generated by any elementary frame class.*

Proof. The example that we give here is based on a famous graph-theoretic result due to Erdős. Here a *graph* is a pair $\mathbb{G} = (G, E)$ with E an irreflexive, symmetric relation on G . A *k-coloring* of \mathbb{G} is a partition of G into k *independent* sets, i.e., sets containing no pair of neighboring vertices. The *chromatic number* $\chi(\mathbb{G})$ of \mathbb{G} is the smallest number k for which it has a k -coloring, and ∞ if it has no finite coloring. A *cycle* in \mathbb{G} is a path $x_1 E x_2 E \dots E x_n E x_1$ such that $n \geq 3$ and x_1, \dots, x_n are all distinct vertices; the length of this cycle is n .

Now intuitively, a lack of short cycles, indicating a certain 'looseness' of the graph, should make it easy to color a graph with few colors, but ERDŐS [22] reveals the existence of a sequence of finite graphs whose n -th member \mathbb{G}_n has chromaticity bigger than n while \mathbb{G}_n has *no* cycles of length $\leq n$. Fix such a sequence $\{\mathbb{G}_n \mid n \geq 2\}$, under the additional assumption that $|\mathbb{G}_n| > |\mathbb{G}_m|$ if $n > m$. (Here $|\mathbb{G}|$ denotes the number of vertices in \mathbb{G} .)

The modal similarity type ϵ of our variety EG will have two diamonds, \Diamond and \mathbf{E} . On a graph \mathbb{G} , the first of these will be interpreted through the edge relation, and the second, through the *global relation* $\Upsilon_G = G \times G$. That is, \mathbf{E} is a *global modality*, cf. section 8.2. In the sequel we will blur the distinction between the structures $\langle G, E, \Upsilon_G \rangle$ and $\langle G, E \rangle$, for instance calling $\langle G, E, \Upsilon_G \rangle^+$ the complex algebra of \mathbb{G} , and denoting it, accordingly, as \mathbb{G}^+ .

For the definition of EG we extend the notion of chromaticity to arbitrary algebras. An element a of an ϵ -BAO \mathbb{A} is called *independent* if $a \wedge \Diamond a = \perp$; write $\chi(\mathbb{A})$ for the chromatic number of \mathbb{A} , that is, for the least k such that there are independent a_1, \dots, a_k with $a_1 \vee \dots \vee a_k = \top$ and $a_i \wedge a_j = \perp$ for $i \neq j$, putting $\chi(\mathbb{A}) = \infty$ if there is no finite such k . Note that this definition generalizes the one given earlier, in the sense that for any graph \mathbb{G} , $\chi(\mathbb{G}) = \chi(\mathbb{G}^+)$.

Now let $\psi_{n,m}$ be the first order formula in this algebraic language stating that if \mathbb{A} has at least 2^n elements, then $\chi(\mathbb{A}) > m$, and define

$$\begin{aligned}\Psi &:= \{\psi_{1,2}\} \cup \{\psi_{|\mathbb{G}_n|,n} \mid n \geq 2\}, \\ \Gamma &:= \{x \preceq \mathbf{E}x, \mathbf{E}\mathbf{E}x \preceq \mathbf{E}x, \mathbf{E}\neg\mathbf{E}\neg x \preceq x, \diamond x \preceq \mathbf{E}x\}.\end{aligned}$$

Note that Γ is the set of equations defining \mathbf{E} to be a global modality, cf. Definition 1.135 for the logical incarnation of Γ . Let \mathbf{C} denote the class of algebras satisfying the formulas $\Psi \cup \Gamma$, and let \mathbf{EG} denote the variety generated by \mathbf{C} . It follows from Theorem 1.139 that $\mathbf{EG} = \mathbf{SP}(\mathbf{C})$.

We first show that \mathbf{EG} is canonical. Note that since \mathbf{C} is an elementary class, it suffices by Theorem 1.98 below to prove that \mathbf{C} itself is canonical. Take an arbitrary algebra \mathbb{A} in \mathbf{C} . If \mathbb{A} is finite, then $\mathbb{A}^\sigma \cong \mathbb{A}$ is in \mathbf{C} by assumption. If \mathbb{A} is infinite, then $|\mathbb{A}| > 2^{|\mathbb{G}_n|}$ for all $n \geq 2$, so by $\mathbb{A} \models \psi_{|\mathbb{G}_n|,n}$ we obtain that $\chi(\mathbb{A}) > n$ for all $n \geq 2$. Clearly then $\chi(\mathbb{A}) = \infty$; from this we may derive that the ultrafilter frame \mathbb{A}_\bullet has a *reflexive* point, which implies that $(\mathbb{A}_\bullet)^+$, being the complex algebra of \mathbb{A}_\bullet , has infinite chromaticity as well. But then we see that $\mathbb{A}^\sigma \models \psi_{m,n}$ for all m, n , so we certainly have $\mathbb{A}^\sigma \models \Psi$. It is easily seen that the formulas Γ are canonical, so that we have proved that \mathbb{A}^σ belongs to \mathbf{C} .

It is left to prove that \mathbf{EG} is not elementarily generated. Theorem 4.12 of GOLDBLATT [40] states that any variety \mathbf{V} of BAOs which is elementarily generated, is generated by an elementary frame class \mathbf{K} such that $\mathbf{Cst}(\mathbf{V}) \subseteq \mathbf{K} \subseteq \mathbf{Str}(\mathbf{V})$. Hence, for our purpose it suffices to come up with a family of frames in $\mathbf{Cst}(\mathbf{EG})$ that provide an ultraproduct outside $\mathbf{Str}(\mathbf{EG})$, and the obvious candidates for this are the Erdős frames $\{\mathbb{G}_n \mid n \geq 2\}$. It is easy to check that $\mathbb{G}_n^+ \models \Psi$ for each $n \geq 2$, so each \mathbb{G}_n^+ belongs to \mathbf{C} . But then all Erdős frames belong to $\mathbf{Cst}(\mathbf{C})$, because each \mathbb{G}_n , being finite, is isomorphic to $(\mathbb{G}_n^+)_\bullet$. Now take a non-principal ultrafilter D over the set $\omega \setminus \{0, 1\}$. Observe that for each k , only finitely many of the \mathbb{G}_n have any cycles of length k ; hence, by Łoś' theorem, the ultraproduct $\prod_D \mathbb{G}_n$ has no cycles at all, and hence, it is *2-colorable*.

This shows that $\prod_D \mathbb{G}_n$ does not belong to \mathbf{C} , since it follows from $\mathbf{C} \models \psi_{1,2}$ that every nontrivial algebra in \mathbf{C} has chromaticity at least three. But fairly direct proofs show that $\chi(\prod_I \mathbb{A}_i) \geq \chi(\mathbb{A}_i)$ for all i , and that $\chi(\mathbb{A}) \geq \chi(\mathbb{A}')$ if $\mathbb{A} \twoheadrightarrow \mathbb{A}'$. This implies that $\chi(\mathbb{A}) > 2$ for all \mathbb{A} in $\mathbf{SP}(\mathbf{C})$, so by the fact that $\mathbf{SP}(\mathbf{C}) = \mathbf{EG}$ it follows that $(\prod_D \mathbb{G}_n)^+$ does not belong to \mathbf{EG} . ■

Nevertheless, the converse of Fine's theorem may fail be true in general, in many interesting cases it does hold — we refer to GOLDBLATT, HODKINSON & VENEMA [46] for a state of the art survey. Note that it is still an open problem whether every *finitely axiomatizable* canonical variety is elementarily generated.

Finally, recent work has put Fine's result in a wider algebraic context. We formulate the following theorem for Boolean algebras with operators, but in fact, it holds in a much wider setting, see for instance GEHRKE & HARDING [28].

THEOREM 1.98. *Let \mathbf{K} be a class of Boolean algebras with τ -operators which is closed under taking ultraproducts and canonical extensions. Then the variety generated by \mathbf{K} is canonical.*

Proof. Let \mathbb{A} be in the variety generated by \mathbf{K} ; we will show \mathbb{A}^σ to belong to $\mathbf{Var}(\mathbf{K})$ as well. By Tarski's 'HSP'-theorem, there is a family $\{\mathbb{B}_i \mid i \in I\} \subseteq \mathbf{K}$, and an algebra

\mathbb{B} such that $\mathbb{A} \leftarrow \mathbb{B} \rightarrow \prod_I \mathbb{B}_i$. Then it follows from two times two applications of Theorem 1.76 that $\mathbb{A}^\sigma \leftarrow \mathbb{B}^\sigma \rightarrow (\prod_I \mathbb{B}_i)^\sigma$, so it suffices to show that $(\prod_I \mathbb{B}_i)^\sigma$ belongs to $\text{Var}(\mathbb{K})$. However, we may infer from Theorem 1.77 and Theorem 1.76(5) that

$$\left(\prod_I \mathbb{B}_i\right)^\sigma \cong \prod_{D \in \text{Uf}(I)} \left(\prod_D \mathbb{A}_i\right)^\sigma. \quad (1.24)$$

But by the assumptions on \mathbb{K} , each algebra $(\prod_D \mathbb{A}_i)^\sigma$ belongs to \mathbb{K} , and so the product (1.24) is in $\text{P}(\mathbb{K}) \subseteq \text{Var}(\mathbb{K})$, as required. ■

From the above result we can derive Fine's Theorem as follows. Suppose that \mathbb{C} is a frame class, closed under taking ultraproducts; for instance, let \mathbb{C} be elementary. Then consider the class $\text{SCm}(\mathbb{C})$ of sub-complex algebras over \mathbb{C} . This class can be shown to be closed under taking ultraproducts as a corollary of Proposition 1.78, and closed under taking canonical extensions as a corollary of Theorem 1.76 and Theorem 90 in Chapter 5 of this volume. Application of Theorem 1.98 then yields the desired result.

6.3 Interpolation


In the last part of this section we discuss another fundamental property of logics: interpolation. Interpolation is important for applications because it allows reasoning systems to be set up in a modular way. Since we have confined our attention to logics in the form of sets of theorems, the version of interpolation that we will consider here is the following.

DEFINITION 1.99. A modal logic L has the *local* or *Craig interpolation property* if for every two formulas φ and ψ such that $\vdash_L \varphi \rightarrow \psi$ there is an *interpolant*, that is, a formula χ with $\vdash_L \varphi \rightarrow \chi$ and $\vdash_L \chi \rightarrow \psi$ and such that each variable of χ occurs both in φ and in ψ .

The algebraic counterpart of interpolation involves the notion of *amalgamation*.

DEFINITION 1.100. Let \mathbb{K} be a class of algebras.

A *V-formation* in \mathbb{K} is a quintuple, presented as $\mathbb{B}_1 \xleftarrow{e_1} \mathbb{B}_0 \xrightarrow{e_2} \mathbb{B}_2$, and consisting of three algebras \mathbb{B}_0 , \mathbb{B}_1 and \mathbb{B}_2 in \mathbb{K} , linked by two embeddings e_0 and e_1 . An *amalgam* of this V-formation is a formation $\mathbb{B}_1 \xrightarrow{f_1} \mathbb{B}_{12} \xleftarrow{f_2} \mathbb{B}_2$ such that $f_1 \circ e_1 = f_2 \circ e_2$. Such a amalgam is a *superamalgam* if for all distinct i and j , and all $b_i \in B_i$ and $b_j \in B_j$: $f_i(b_i) \leq_{12} f_j(b_j)$ only if there is some $b_0 \in B_0$ with $b_i \leq_i e_i(b_0)$ and $e_j(b_0) \leq_j b_j$.



\mathbb{K} is said to have the *(super)amalgamation property* if every V-formation in \mathbb{K} has a (super)amalgam in \mathbb{K} .

In words, an amalgam is a superamalgam if whenever a B_i -element is smaller (in B_{12}) than a B_j -element, then this is *witnessed* by a B_0 -element. The basic result connecting interpolation and amalgamation is from MAKSIMOVA [83].


THEOREM 1.101. *Let L be a normal modal τ -logic. Then L has the local interpolation property if and only if $\text{BAO}_\tau(L)$ has superamalgamation.*

Proof. Fix L . In the proof of this theorem we will frequently consider Lindenbaum-Tarski algebras for L over various distinct sets of variables. Our notational convention will be that these sets of variables will always be called X_0, X_1, X_2 and X_{12} , with $X_0 = X_1 \cap X_2$ and $X_{12} = X_1 \cup X_2$; that \mathbb{F}_i denotes the Lindenbaum-Tarski algebra over X_i ; that $[\varphi]_i$ denotes the equivalence class of the formula φ under the L -equivalence relation \equiv_L within the set $Fma(X_i)$; and, finally, if $X_i \subseteq X_j$, that $\iota_{i,j}$ denotes the map given by $[\varphi]_i \mapsto [\varphi]_j$. We leave it for the reader to verify that $\iota_{i,j}$ is an *embedding* of \mathbb{F}_i into \mathbb{F}_j .

It is not hard to prove that L has local interpolation iff for all sets X_1 and X_2 of variables, the formation $\mathbb{F}_1 \xrightarrow{\iota_{1,12}} \mathbb{F}_{12} \xleftarrow{\iota_{2,12}} \mathbb{F}_2$ is a superamalgam of the V-formation $\mathbb{F}_1 \xleftarrow{\iota_{0,1}} \mathbb{F}_0 \xrightarrow{\iota_{0,2}} \mathbb{F}_2$. This observation already takes care of the direction from right to left of the theorem.

For the other direction we have to work harder. Consider a V-formation $\mathbb{B}_1 \xleftarrow{e_1} \mathbb{B}_0 \xrightarrow{e_2} \mathbb{B}_2$ in $\text{BAO}_\tau(L)$. Without loss of generality we may assume that $B_0 = B_1 \cap B_2$. Wanting to use local interpolation of L to find a superamalgam of this V-formation, we translate the V-formation into syntax.

With $X_i := \{x_b \mid b \in B_i\}$ for each $i \in \{0, 1, 2\}$, let $\beta_i : \mathbb{F}_i \rightarrow \mathbb{B}_i$ be the unique homomorphism determined by the map $[x_b] \mapsto b$, cf. the picture. Clearly each β_i is surjective, whence by universal algebra, each \mathbb{B}_i is isomorphic to the algebra $\mathbb{F}_i / \ker(\beta_i)$. Let M_i be the modal filter of \mathbb{F}_i associated with the congruence $\ker(\beta_i)$ (as in Theorem 1.29), and let M be the modal filter of \mathbb{F}_{12} generated by the union of M_1 and M_2 , or, to be more precise, by the set $\iota_{1,12}[M_1] \cup \iota_{2,12}[M_2]$. We claim that the algebra \mathbb{F}_{12}/Π_M is the required superamalgam, with Π_M the congruence associated with M , again, as in Theorem 1.29.



Proving this, the crucial observation is that $[\varphi]_{12}$ belongs to M iff there are formulas $\varphi_1 \in Fma(X_1)$ and $\varphi_2 \in Fma(X_2)$ such that $\vdash_L (\varphi_1 \wedge \varphi_2) \rightarrow \varphi$, and $[\varphi_i]_i \in M_i$ for $i = 1, 2$. From this, using local interpolation, it may be derived that for formulas $\psi_1 \in Fma(X_1)$ and $\psi_2 \in Fma(X_2)$, we have $[\psi_1 \rightarrow \psi_2]_{12} \in M$ iff there is a $\chi \in Fma(X_0)$ such that $[\psi_1 \rightarrow \chi]_1 \in M_1$ and $[\chi \rightarrow \psi_2]_2 \in M_2$. And from this the desired properties of \mathbb{F}_{12}/Π_M follow almost immediately. ■

This theorem can be applied to obtain a fairly general interpolation result for *canonical* modal logics that define nice frame classes. We need the following definition.

DEFINITION 1.102. Let \mathbb{S}_1 and \mathbb{S}_2 be two τ -frames. The direct product $\mathbb{S}_1 \times \mathbb{S}_2$ of these frames is the frame based on the Cartesian product $S_1 \times S_2$, with the relations defined coordinate-wise (for instance, in the case of a binary relation R , we put $R(s_1, s_2)(t_1, t_2)$ if $R_1 s_1 t_1$ and $R_2 s_2 t_2$). A subframe \mathbb{Z} of $\mathbb{S}_1 \times \mathbb{S}_2$ is called a *zigzag product* of \mathbb{S}_1 and \mathbb{S}_2 if \mathbb{Z} is a hereditary subset of the product frame on which the projection maps are surjective.

Clearly then zigzag products are substructures of direct products. A different perspective is that zigzag products of \mathbb{S}_1 and \mathbb{S}_2 are given by those bisimulations Z between \mathbb{S}_1 and \mathbb{S}_2 that are *full*, i.e., have domain S_1 and range S_2 .

As an example of a zigzag product, consider two surjective bounded morphisms θ_1, θ_2 with $\theta_i : \mathbb{S}_i \rightarrow \mathbb{S}_0$. Then the frame $\mathbb{E}(\theta_1, \theta_2)$ based on the set $\{(s_1, s_2) \in S_1 \times S_2 \mid \theta_1(s_1) = \theta_2(s_2)\}$ is a zigzag product of \mathbb{S}_1 and \mathbb{S}_2 . We call this the zigzag product *induced* by θ_1 and θ_2 .

The following theorem, which is a generalization from MARX [84] of a result by NÉMETI [87], is useful for proving that a canonical logic has interpolation.

THEOREM 1.103. *Let \mathbf{K} be a class of Boolean algebras with τ -frames, and \mathbf{C} a class of τ -frames such that $\text{Cst}(\mathbf{K}) \subseteq \mathbf{C}$, $\text{Cm}(\mathbf{C}) \subseteq \mathbf{K}$, and \mathbf{C} is closed under taking zigzag products. Then \mathbf{K} has the superamalgamation property.*

Proof. Suppose that \mathbf{K} and \mathbf{C} have the listed properties, and consider a V-formation

$$\mathbb{B} \xleftarrow{\alpha} \mathbb{A} \xrightarrow{\alpha'} \mathbb{B}'. \quad (1.25)$$

It follows from Theorem 1.76(3) that $\mathbb{B}_\bullet \xrightarrow{\alpha_\bullet} \mathbb{A}_\bullet \xleftarrow{\alpha'_\bullet} \mathbb{B}'_\bullet$. Now let \mathbb{E} be the zigzag product of \mathbb{B}_\bullet and \mathbb{B}'_\bullet induced by the bounded morphisms α_\bullet and α'_\bullet . Note that \mathbb{E} belongs to \mathbf{C} by the listed closure properties. Letting π and π' be the (surjective!) bounded morphisms from \mathbb{E} onto \mathbb{B}_\bullet and \mathbb{B}'_\bullet , respectively, we see that $\mathbb{B}_\bullet \xleftarrow{\pi} \mathbb{E} \xrightarrow{\pi'} \mathbb{B}'_\bullet$. It then follows from Theorem 1.76(2) and Theorem 1.54 that

$$\mathbb{B} \xrightarrow{\hat{}} \mathbb{B}^\sigma \xrightarrow{\pi^+} \mathbb{E}^+ \xleftarrow{\pi'^+} \mathbb{B}'^\sigma \xleftarrow{\hat{}} \mathbb{B}'. \quad (1.26)$$

We claim that in fact, (1.26) is a superamalgam of (1.25), but leave further proof details for the reader. ■

As a corollary of this theorem, suppose that Γ is a set of canonical formulas defining an elementary frame class that is closed under taking direct products and substructures — for instance, Γ corresponds to a set of universal Horn sentences. Then $\mathbf{K}_\tau.\Gamma$ has Craig interpolation.

Chapter 8 of this volume contains more information on interpolation. Related properties, such as Beth definability, also have algebraic characterizations; for details we refer to HOOGLAND [59].

7 CASE STUDY: CANONICAL EQUATIONS

7.1 Introduction

In this section we address the question, which equations are *canonical*, that is, remain valid when we move from a BAO \mathbb{A} to its canonical embedding algebra \mathbb{A}^σ . In other words, we are interested in properties that move to certain *superalgebras*.

Earlier on we defined \mathbb{A}^σ via a concrete construction, namely, as the ‘double dual’ $(\mathbb{A}_\bullet)^+$: the complex algebra of the ultrafilter frame of \mathbb{A} . In this section we will take a rather more abstract approach in which we first consider the canonical extension \mathbb{B}^σ of the Boolean reduct \mathbb{B} of \mathbb{A} ; this \mathbb{B}^σ is not constructed but axiomatically characterized as the (modulo isomorphism) unique completion of \mathbb{B} in which \mathbb{B} is *dense* and *compact*. Then the property of density suggests a canonical way to extend the interpretation of the operators on \mathbb{B} to operations on \mathbb{B}^σ , thus providing the canonical extension \mathbb{A}^σ of \mathbb{A} .

This algebraic method originates with the original BAO paper JÓNSSON AND TARSKI [70], but it differs from the duality-based approach of for instance SAMBIN & VACCARO [100] that modal logicians usually take. In order to compare the two approaches, consider the following picture, introducing the four main characters of this story:

$$\begin{array}{|c|c|} \hline \mathbb{A} & \mathbb{A}_* \\ \hline \mathbb{A}^\sigma & \mathbb{A}_\bullet \\ \hline \end{array} \quad (1.27)$$

In the duality-based approach, one compares the frame (frame-based) structures on the right hand side of the picture, cf. the discussion on the notion of *persistence* in Chapter 5 of this volume, while the algebraic method stays purely on the left hand side, basically by encoding the relevant topological concepts into the algebraic framework. An advantage of the duality-based method is that it allows a treatment of canonicity in tandem with correspondence; on the other hand, the more abstract and ‘duality-free’ nature of the other approach enables its transportation to a much wider setting than that of canonical extensions of Boolean algebras with operators. In recent years, the algebraic approach has proven its use for *lattices* expanded with *arbitrary* operations, and has been applied to other kinds of completions than the perfect extension of Jónsson and Tarski.

Our exposition of this algebraic approach in the sections 7.2 to 7.5 is based on work by JÓNSSON [69], GEHRKE & JÓNSSON [30, 31, 32] and GEHRKE & HARDING [28], while the very similar approach by GHILARDI & MELONI [34] should also be mentioned here. In our presentation we try to be as general as possible while keeping the section self-contained, and staying within the framework of Boolean algebras. Almost all our formulations apply to lattice-ordered algebras as well, however; we will come back to this issue towards the end of the section when we discuss further generalizations of the theory presented here.

For an outline, recall that the validity of equations can be formulated using term functions:

$$\mathbb{A} \models s \approx t \text{ iff } s^{\mathbb{A}} = t^{\mathbb{A}}. \quad (1.28)$$

Hence, for the canonical extension of \mathbb{A} , we find that

$$\mathbb{A}^\sigma \models s \approx t \text{ iff } s^{\mathbb{A}^\sigma} = t^{\mathbb{A}^\sigma}. \quad (1.29)$$

Now suppose that we have developed a canonical way to extend an n -ary map $f : A^n \rightarrow A$ to an n -ary map $f^\sigma : (A^\sigma)^n \rightarrow A^\sigma$; it then immediately follows from (1.28) that

$$\mathbb{A} \models s \approx t \text{ only if } (s^{\mathbb{A}})^\sigma = (t^{\mathbb{A}})^\sigma. \quad (1.30)$$

Hence, in case s and t are *stable* on \mathbb{A} , that is, if $(s^{\mathbb{A}})^\sigma = s^{\mathbb{A}^\sigma}$ and $(t^{\mathbb{A}})^\sigma = t^{\mathbb{A}^\sigma}$, then we may infer from $\mathbb{A} \models s \approx t$ that $\mathbb{A}^\sigma \models s \approx t$. This motivates a careful analysis of the relation between the functions $s^{\mathbb{A}^\sigma}$ (the term function of s in \mathbb{A}^σ) and $(s^{\mathbb{A}})^\sigma$ (the extension to \mathbb{A}^σ of the term function $s^{\mathbb{A}}$). This analysis crucially involves the question, which f and g satisfy $(f \circ g)^\sigma = f^\sigma \circ g^\sigma$. We will see that such cases of $(\cdot)^\sigma$ distributing over function composition admit a satisfactory explanation in terms of ‘matching continuity properties’ of the maps f^σ and g^σ . For this purpose we will endow canonical extensions of Boolean algebras with topological structure.

7.2 Canonical extensions of Boolean algebras

In this section we define the canonical extension of a Boolean algebra \mathbb{B} as the unique *completion* of \mathbb{B} in which \mathbb{B} is *dense* and *compact*. We introduce these notions one by one.

A Boolean algebra \mathbb{C} is a *completion* of a Boolean algebra \mathbb{B} if \mathbb{C} is complete and \mathbb{B} is a subalgebra of \mathbb{C} . If \mathbb{C} agrees with \mathbb{B} on *all* meets and joins, then we call \mathbb{C} a *regular completion* of \mathbb{B} , but in general we do not require completions to be regular. Thus the notation \bigvee for finite joins is unambiguous, but not so for infinite joins. Our convention will be that $\bigvee X$ always denotes $\bigvee^{\mathbb{C}} X$, that is, the join taken in the completion.

For an example of a completion, consider a field of sets $\mathbb{S} = \langle S, A \rangle$ and note that the power set algebra $\mathbb{P}S$ is a completion of \mathbb{S}^* .

Before we define the concept of density, we introduce some preliminary notions. Given a completion \mathbb{C} of the Boolean algebra \mathbb{B} , we call an element $c \in C$ *closed* (*open*) if c is the meet (join, respectively) in \mathbb{C} of elements in B . We let $K_{\mathbb{C}}(B)$ and $O_{\mathbb{C}}(B)$ denote the collections of closed and open elements, respectively. Objects (such as the elements of B) that are both closed and open are called *clopen*. This terminology is in accordance with the topological perspective on fields of sets as in Remark 1.60. In the sequel, we may write $K_{\mathbb{C}}$, $K(B)$, or even K , instead of $K_{\mathbb{C}}(B)$, if the suppressed details are clear from context; and similarly for the set $O_{\mathbb{C}}(B)$.

We say that \mathbb{B} is *meet-dense* in \mathbb{C} if $K_{\mathbb{C}}(B) = C$, *join-dense* if $O_{\mathbb{C}}(B) = C$, and *dense* if $K_{\mathbb{C}}(O_{\mathbb{C}}(B)) = O_{\mathbb{C}}(K_{\mathbb{C}}(B)) = C$. In words, A is dense in \mathbb{C} if every element of C is both a meet of open elements, and a join of closed elements. As a simple example of join-density, note that a Boolean algebra is atomic iff the collection of atoms forms a join-dense set. Building on this, we leave it as an exercise for the reader to verify that a field of sets $\mathbb{S} = \langle S, A \rangle$ is differentiated iff \mathbb{S}^* is dense in $\mathbb{P}S$.

Now we turn to the notion of compactness. Given a completion \mathbb{C} of the Boolean algebra \mathbb{B} , we say that \mathbb{B} is *compact in* \mathbb{C} if for all sets X and Y of closed and open elements, respectively, $\bigwedge X \leq \bigvee Y$ implies the existence of finite subsets $X_0 \subseteq X$, $Y_0 \subseteq Y$ such that $\bigwedge X_0 \leq \bigvee Y_0$. An alternative (but equivalent) characterization of compactness is that, for any closed p and open u ,

$$p \leq u \text{ only if } p \leq b \leq u \text{ for some } b \in B,$$

as can easily be verified. Also note that, again, our definition of compactness coincides with standard topological terminology; this easily follows from the observation that for any pair C, U of collections of subsets of a set S , we have $\bigcap C \subseteq \bigcup U$ iff $S \subseteq \bigcup U \cup \{\sim_{S C} \mid c \in C\}$.

We are now ready to define canonical extensions.

DEFINITION 1.104. A completion \mathbb{C} of the Boolean algebra \mathbb{B} is called a *canonical extension* of \mathbb{B} if \mathbb{B} is both compact and dense in \mathbb{C} .

It is in fact a rather strong property for one Boolean algebra to be the canonical extension of another. To start with, every Boolean algebra has a *unique* canonical extension.

THEOREM 1.105. *Let \mathbb{B} be some Boolean algebra. Then*

1. (*existence*) \mathbb{B} has a canonical extension;
2. (*unicity*) Any two canonical extensions of \mathbb{B} are isomorphic via a unique isomorphism that restricts to the identity on B .

Proof. Recall from the topological duality that $\mathbb{B}_* = \langle \text{Uf}\mathbb{B}, \widehat{B} \rangle$ is a differentiated and compact field of sets. By the comments made above it should be clear that $\mathbb{P}(\text{Uf}\mathbb{B})$ is a canonical extension of \mathbb{B} .

For unicity, suppose that \mathbb{C} is a canonical extension of \mathbb{B} . We leave it as an exercise for the reader to verify that, by compactness, the map $F \mapsto \bigwedge F$ forms a dual (that is, order-reversing) isomorphism between the lattice $\langle \text{Fi}(\mathbb{B}), \subseteq \rangle$ and the induced ordering on the set $K(B)$ of closed elements. Its inverse is given by the map $p \mapsto \{a \in B \mid a \geq p\}$. Similarly, there is a dual isomorphism between the lattice of ideals of \mathbb{B} , and the induced ordering of the open elements. Also, we have for p closed and u open, that $p \leq u$ iff there is an $a \in B$ with $p \leq a \leq u$, and that $u \leq p$ iff $a \leq b$ for all a and b in A with $a \leq u$ and $p \leq b$. In other words, by *compactness* the induced poset on the set $K \cup O$ of closed or open elements is completely determined by the ordering of \mathbb{B} . This suffices to prove the theorem, since by *density*, the elements of \mathbb{C} can be identified with the pairs (L, U) of subsets of C such that L is the collection of closed lower bounds of U , and U is the collection of open upper bounds of L . Summarizing, we see that together, compactness and density completely fix the order relation of the canonical extension. ■

The above theorem justifies our speaking of ‘the’ canonical extension of a Boolean algebra \mathbb{B} ; this algebra will be denoted as \mathbb{B}^σ . Furthermore, we need the following facts.

PROPOSITION 1.106. *Let \mathbb{C} be a canonical extension of the Boolean algebra \mathbb{B} . Then*

1. $B = K(B) \cap O(B)$; that is, B coincides with the set of clopen elements of \mathbb{C} ;
2. the set $K(B)$ forms a sublattice of \mathbb{C} which is closed under taking infinitary meets;
3. \mathbb{C} is atomic and $\text{At}\mathbb{C} \subseteq K(B)$; that is, all atoms are closed.

We leave the proof of this proposition to the reader; note that by Theorem 1.105, it suffices to restrict attention to the double dual $\mathbb{P}(\text{Uf}\mathbb{B})$ of \mathbb{B} . For instance, part (3) follows almost immediately from the identification of atoms of $\mathbb{P}(\text{Uf}\mathbb{B})$ with ultrafilters of \mathbb{B} .

As a last introductory remark, we note that canonical extensions interact well with finite products and order duals. Concerning the latter notion, recall that the *order dual* of a Boolean algebra $\mathbb{B} = \langle B, \top, \perp, -, \wedge, \vee \rangle$ is the structure $\mathbb{B}^\partial = \langle B, \perp, \top, -, \vee, \wedge \rangle$. The fact, that \mathbb{B}^∂ is a Boolean algebra as well, enables us to shorten quite a lot of definitions and proofs by referring to the *principle of order duality*: Every fact concerning Boolean algebras remains valid after swapping \top with \perp , \wedge with \vee , etc.

PROPOSITION 1.107. *Let $\mathbb{B}_1, \dots, \mathbb{B}_n$ be Boolean algebras. Then*

1. $(\mathbb{B}_1 \times \dots \times \mathbb{B}_n)^\sigma \cong \mathbb{B}_1^\sigma \times \dots \times \mathbb{B}_n^\sigma$;
2. $(\mathbb{B}^\partial)^\sigma \cong (\mathbb{B}^\sigma)^\partial$;

Proof. Both statements can be proved on the basis of Theorem 1.76. As intermediate steps, one can prove facts like $K(\mathbb{B}_1) \times \dots \times \mathbb{B}_n = K(\mathbb{B}_1) \times \dots \times K(\mathbb{B}_n)$ and $K(\mathbb{B}^\partial) = O(\mathbb{B})$. ■

7.3 Extending maps to the canonical extension

In the introduction to this section we saw that in order to investigate the canonicity of an equation $s \approx t$, it is useful to define extensions of the term functions on a BAO to maps on the canonical extension of the BAO. But in fact, there are canonical ways to extend an *arbitrary* map between two Boolean algebras \mathbb{A} and \mathbb{B} , to a map between \mathbb{A}^σ and \mathbb{B}^σ . This general definition will be discussed at the end of this section — for the time being we will confine ourselves to extensions of *monotone* maps.

The easiest way to understand these definitions is to break them down in two steps. For a start, the definition of closed and open elements suggests the following extension of $f : \mathbb{A} \rightarrow \mathbb{B}$ to a map \bar{f} defined on $K(A) \cup O(A)$:

$$\begin{aligned} \bar{f}(p) &:= \bigwedge \{f(a) \mid p \leq a \in A\} & \text{for } p \in K(A), \\ \bar{f}(u) &:= \bigvee \{f(a) \mid u \geq a \in A\} & \text{for } u \in O(A). \end{aligned} \quad (1.31)$$

Note that this is a correct definition because $K \cap O = A$ by Proposition 1.106(1), that $\bar{f}(a) = f(a)$ for $a \in A$ by monotonicity of f , and that \bar{f} itself is also order preserving.

Now for the second step of the construction. The fact that every element is both the join of the closed elements below it, and the meet of the opens above it, suggests *two* ways to proceed:

$$\begin{aligned} f^\sigma(x) &:= \bigvee \{\bar{f}(p) \mid x \geq p \in K(A)\}, \\ f^\pi(x) &:= \bigwedge \{\bar{f}(u) \mid x \leq u \in O(A)\}. \end{aligned} \quad (1.32)$$

The maps f^σ and f^π are called the *lower* and *upper extension* of f , respectively.

Let us first gather some basic facts concerning these definitions. The following proposition says that the names ‘lower’, ‘upper’, and ‘extension’ are well chosen.

PROPOSITION 1.108. *Let $f : \mathbb{A} \rightarrow \mathbb{B}$ be a monotone map between Boolean algebras. Then*

1. both f^σ and f^π extend f ;
2. $f^\sigma \leq f^\pi$, with equality holding on the closed and on the open elements.

Proof. The first statement is immediate by the definitions and the monotonicity of \bar{f} . For the second statement, take, for $x \in A^\sigma$, a closed $p \leq x$ and an open $u \geq x$. By compactness there is an $a \in [p, u] \cap A$. This element satisfies $\bar{f}(p) \leq f(a) \leq \bar{f}(u)$ by definition of \bar{f} ; hence $f^\sigma(x) \leq f^\pi(x)$ by definition of f^σ and f^π . Finally, for closed p we may derive from the first part of the proposition that $f^\pi(p) \leq \bar{f}(p)$, and from the monotonicity of \bar{f} that $\bar{f}(p) = f^\sigma(p)$. Thus we obtain the desired equality $f^\sigma = f^\pi$ on K . The result for opens follows by order duality. ■

Maps for which the lower and upper extension coincide are obviously of interest.

DEFINITION 1.109. A monotone map f between Boolean algebras is called *smooth* if $f^\sigma = f^\pi$.

EXAMPLE 1.110. As a first example of a smooth operation, consider the *global modality* g on a Boolean algebra \mathbb{B} , given by $g(\perp) = \perp$ while $g(b) = \top$ for $b > \perp$, see Definition 1.131. It is easy to see that \bar{g} satisfies these conditions as well, whence it is equally easy to infer that both g^σ and g^π coincide with the global modality of \mathbb{B}^σ ; smoothness

is then immediate. Similarly, one can prove that the *meet* and *join* operations of \mathbb{B} are smooth, and that their extensions coincide with the meet and the join of \mathbb{B}^σ , respectively.

For an operation that is not smooth, consider the composition of the global modality with the meet operation, i.e., the map $f : B^2 \rightarrow B$ given by $f(a, b) = \perp$ if $a \wedge b = \perp$, while $f(a, b) = \top$ otherwise. Now if \mathbb{B} is infinite, then \mathbb{B}^σ must contain some element c which is closed but not open; a straightforward verification shows that for such a c , we have that $f^\sigma(c, -c) = \perp$, while $f^\pi(c, -c) = \top$. This shows that not even operators are smooth.

While it may not be the case that the lower and the upper extension agree in all cases, both kinds of extensions generally display good behavior; often they even improve on the original map. For the definitions of the notions mentioned in the theorem below, see Definition 1.15 and 1.40.

PROPOSITION 1.111. *Let $f : \mathbb{A} \rightarrow \mathbb{B}$ be a map between Boolean algebras. Then*

1. *if f is monotone then so is f^σ ;*
2. *if f is an operator then f^σ is a complete operator;*
3. *if f is additive or multiplicative then f is smooth.*

Proof. The proof of the first statement is easy and hence omitted, while we postpone the proof of the last statement (it is in fact a rather straightforward consequence of the Propositions 1.116 and 1.117). For the remaining part, we need to show that if f is normal and additive in each coordinate, then f^σ is normal and completely additive in each coordinate. Leaving the easy proof for normality as an exercise for the reader, concerning additivity, we will prove that if $f : \mathbb{A}_0 \times \mathbb{A}_1 \rightarrow \mathbb{B}$ is additive in its first coordinate and monotone in its second, then f^σ preserves all non-empty joins in its first coordinate.

Fix elements $x_0 \in A_0^\sigma$ and $x_1 \in A_1^\sigma$. By atomicity of \mathbb{B}^σ , and monotonicity of f^σ , it suffices to prove, for an arbitrary atom p of \mathbb{B}^σ :

$$p \leq f^\sigma(x_0, x_1) \text{ only if there is a } q \in At_0 \text{ with } p \leq f^\sigma(q, x_1), \quad (1.33)$$

where At_0 denotes the set of atoms in A_0^σ below x_0 . Note that since $f^\sigma(x_0, x_1) = \bigvee \{f^\sigma(c_0, c_1) \mid x_i \geq c_i \in K(A_i)\}$ we may safely assume that both x_0 and x_1 are closed.

Now suppose for contradiction that (1.33) fails. Then for some atom p of \mathbb{B}^σ we have $p \leq f^\sigma(x_0, x_1)$ while for each $q \in At_0$ there are, by definition of f^σ , elements $a_{q,0} \in A_0$ above q and $a_{q,1} \in A_1$ above x_1 , such that $p \not\leq f^\sigma(a_{q,0}, a_{q,1})$. It follows that $x_0 = \bigvee At_0 \leq \bigvee \{a_{q,0} \mid q \in At_0\}$, whence by compactness $x_0 \leq \bigvee \{a_{q,0} \mid q \in F\}$ for some finite set $F \subseteq At_0$.

Now observe that the join $a_0 = \bigvee \{a_{q,0} \mid q \in F\}$ is in A_0 , and the meet $a_1 = \bigwedge \{a_{q,1} \mid q \in F\}$ is in A_1 . Clearly $p \not\leq f^\sigma(a_{q,0}, a_{q,1})$ for each $q \in F$; since p is an atom this means $p \not\leq \bigvee \{f^\sigma(a_{q,0}, a_{q,1}) \mid q \in F\} = f(a_0, a_1)$, where in the last identity we use the additivity of f in its first coordinate.

On the other hand, from $x_0 \leq a_0$ and $x_1 \leq a_1$ it follows that $f^\sigma(x_0, x_1) \leq f(a_0, a_1)$ which gives the desired contradiction. \blacksquare

In the proof above we already used the fact that complete additivity of f^σ means that it is completely determined by its values on the *atoms* of \mathbb{B}^σ . Now recall that (in the concrete representation of) \mathbb{B}^σ , the atoms are nothing but the *ultrafilters* of \mathbb{B} . From this the following proposition is immediate.

PROPOSITION 1.112. *Let \mathbb{A} be some Boolean algebra with τ -operators with underlying Boolean algebra \mathbb{B} . Then $\mathbb{A}^\sigma := (\mathbb{A}_\bullet)^+$ is isomorphic to the algebra \mathbb{B}^σ expanded with the family $\{(\nabla^{\mathbb{A}})^\sigma \mid \nabla \in \tau\}$ of complete operators.*

This proposition, which can be summarized as ‘ $\nabla^{\mathbb{A}^\sigma} = (\nabla^{\mathbb{A}})^\sigma$ ’, will be used throughout the sequel, but always implicitly.

7.4 Composite maps

We now investigate the interaction between composing maps between Boolean algebras and taking their canonical extensions. That is, we will take a look at the relation between the maps $(gf)^\sigma$ and $g^\sigma f^\sigma$ for maps $f : \mathbb{A} \rightarrow \mathbb{A}'$, and $g : \mathbb{A}' \rightarrow \mathbb{A}''$. We are obviously eager to find cases in which we have $(gf)^\sigma = g^\sigma f^\sigma$, but also conditions under which one of the inequalities (\leq or \geq) apply will turn out to be of interest. As we will see shortly, many of these conditions can naturally be described in *topological* terms.

For this purpose, we will introduce no less than six topologies on each set A^σ . Fortunately, these topologies can be neatly organized in two families, each consisting of an upper, a lower and a join topology. As a terminological convention, let us call a map between the algebras \mathbb{A}^σ and \mathbb{B}^σ (ρ, ρ') -*continuous*, if it is a continuous function between the topological spaces $\langle A^\sigma, \rho \rangle$ and $\langle B^\sigma, \rho' \rangle$.

The first family is that of the Scott topologies. Although these can already be defined on arbitrary partial orders, here we will only consider topologies on canonical extensions of BAOs. Recall that a subset D of a partial order is called *up-directed*, if every pair of elements of D has an upper bound in D .

DEFINITION 1.113. Given a Boolean algebra \mathbb{B} , call a subset U of B^σ *Scott open* if U is an up-set such that $U \cap D \neq \emptyset$ for every up-directed set D with $\bigvee D \in U$. The *Scott topology* is defined as the collection γ^\uparrow of Scott open sets; the topology γ^\downarrow is given by the principle of order duality, and we define $\gamma := \{U \cap V \mid U \in \gamma^\uparrow, V \in \gamma^\downarrow\}$ as the join of γ^\uparrow and γ^\downarrow in the lattice of topologies over B .

In practice it is sometimes easier to work with the *closed* sets in the Scott topology; these are precisely the down-sets of \mathbb{C} that are closed under taking up-directed unions. From this observation one easily derives the (well-known) fact that a map between partial orders is Scott continuous (that is, $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous) iff it preserves up-directed joins. But this implies that a map is completely additive iff it is both additive and Scott continuous, which may help to explain the relevance of the Scott topologies for our purposes.

We now turn to the second family of topologies. Recall from Example 1.27 that for an arbitrary element b of a Boolean algebra \mathbb{B} , the sets b^\uparrow and b^\downarrow are defined as $b^\uparrow = \{a \in B \mid b \leq a\}$ and $b^\downarrow = \{a \in B \mid a \leq b\}$.

PROPOSITION 1.114. *For any Boolean algebra \mathbb{B} , the sets $\sigma^\uparrow := \{p^\uparrow \mid p \in K\}$ and $\sigma^\downarrow := \{u^\downarrow \mid u \in O\}$ both form a topology on A^σ ; and so does the set $\sigma := \{p^\uparrow \cap u^\downarrow \mid K \ni p \leq u \in O\}$, which is in fact identical to the join $\sigma^\uparrow \vee \sigma^\downarrow$ in the lattice of topologies on A^σ .*

In the sequel, we will write $[p, u]$ for the *interval* between p and u , that is, $[p, u] = p\uparrow \cap u\downarrow$.

Proof. The fact that σ^\uparrow is a topology follows from the fact that the set $K(A)$ is closed under finitary joins and arbitrary meets of \mathbb{A}^σ , see Proposition 1.106(2). ■

REMARK 1.115. As suggested by notation, the topology σ is closely connected to the kind of inclusion of \mathbb{B} in \mathbb{B}^σ . Let us just mention a couple of salient facts here. First, it is easy to see that the set $\{[a, b] \mid a, b \in B\}$ is a *basis* for σ . This reveals that the set B is topologically *dense* in σ , in the sense that every σ -open set contains an element of B . But also, B constitutes the collection of *isolated points* of σ — recall that a point x is isolated in a topology if the singleton $\{x\}$ is open. It is the latter two properties that make it possible to extend arbitrary maps between Boolean algebras to their extensions; we will come back to this at the end of this section.

The following proposition, which links the two topological families, will be crucial when it comes to finding the ‘matching continuities’ mentioned in the introduction.

PROPOSITION 1.116. *Let \mathbb{A} be a Boolean algebra. Then $\gamma^\uparrow \subseteq \sigma^\uparrow$, $\gamma^\downarrow \subseteq \sigma^\downarrow$ and $\gamma \subseteq \sigma$.*

Proof. Confining ourselves to the first claim, it suffices to prove that $U = \bigcup\{p\uparrow \mid p \in U \cap K\}$ for an arbitrary Scott open set $U \subseteq A^\sigma$. The crucial observation here is that every $u \in U$ is the *up-directed* join of the closed elements below it. Further proof details are left to the reader. ■

The following proposition is a first sign that these topologies can be useful.

PROPOSITION 1.117. *Let $f : \mathbb{A} \rightarrow \mathbb{B}$ be a monotone map between the Boolean algebras \mathbb{A} and \mathbb{B} . Then*

1. f^σ is the largest monotone $(\sigma, \gamma^\uparrow)$ -continuous extension of f ;
2. f is smooth iff f^σ is (σ, γ) -continuous;
3. if f is an operator then f^σ is $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous;
4. if f is additive then f^σ is $(\sigma^\downarrow, \sigma^\downarrow)$ -continuous.
5. if f is multiplicative then f^σ is $(\sigma^\uparrow, \sigma^\uparrow)$ -continuous.

Proof. Concerning the first part of the proposition, we already know from Proposition 1.108 that f^σ is an extension of f . Now for $x \in A^\sigma$ take an arbitrary Scott open set $V \subseteq B^\sigma$ with $f^\sigma(x) \in V$. That is, $\bigvee\{f^\sigma(p) \mid x \geq p \in K(A)\} \in V$. Now it is easy to see that the collection $Q := \{f^\sigma(p) \mid x \geq p \in K(A)\}$ is up-directed, so $Q \cap V \neq \emptyset$. In other words, there is a closed $p \leq x$ with $f^\sigma(p) \in V$. But then by monotonicity of f^σ we have that $f^\sigma[p\uparrow] \subseteq V$. Since $x \in p\uparrow \in \sigma$ this suffices to prove that f^σ is $(\sigma, \gamma^\uparrow)$ -continuous, while by Proposition 1.111 it is monotone.

In order to show that f is the largest such map, take a monotone $(\sigma, \gamma^\uparrow)$ -continuous extension $g : A^\sigma \rightarrow B^\sigma$ of f , and suppose for contradiction that $g(x) \not\leq f^\sigma(x)$ for some $x \in A^\sigma$. By atomicity of \mathbb{B}^σ there must be an *atom* p of \mathbb{B}^σ which lies below $g(x)$, but not below $f^\sigma(x)$. Because $g(x) \in p\uparrow \in \gamma^\uparrow$, the continuity of g provides us with a $c \in K$ such that $c \leq x$ and $g[c\uparrow] \subseteq p\uparrow$. In other words, we find that $p \leq g(c)$ whence by monotonicity

it follows that $p \leq g(a)$ for all $a \in A$ above c . But then by the fact that g extends f , and the definition of f^σ , we may infer that $p \leq f^\sigma(c)$. From this we obtain, as the required contradiction, that $p \leq f^\sigma(x)$.

For part (2), it follows from part (1) by order duality that f^π is the smallest monotone $(\sigma, \gamma^\downarrow)$ -continuous extension of f . Hence if f is smooth, then $f^\sigma = f^\pi$ is both $(\sigma, \gamma^\uparrow)$ - and $(\sigma, \gamma^\downarrow)$ -continuous, and hence, (σ, γ) -continuous. Conversely, if f^σ is (σ, γ) -continuous, then it is, a fortiori, $(\sigma, \gamma^\downarrow)$ -continuous. This implies, again by the order dual of part (1), that $f^\pi \leq f^\sigma$; but then we have equality because of Proposition 1.108(2).

Concerning part (3), if $f : \mathbb{A}^n \rightarrow \mathbb{A}$ is an operator then by Proposition 1.111(2), $f^\sigma : (\mathbb{A}^\sigma)^n \rightarrow \mathbb{A}^\sigma$ is additive in each coordinate. From this it is straightforward to derive that f^σ preserves up-directed joins.

For part (4), suppose that $f : \mathbb{A} \rightarrow \mathbb{B}$ is additive, and take an arbitrary σ^\downarrow -open subset $u\downarrow$ of B^σ , that is, $u \in O(B)$. It follows by Proposition 1.111(2) that f^σ preserves all non-empty joins. From this one may derive that the set $(f^\sigma)^{-1}[u\downarrow]$ is either empty, in which case it certainly belongs to σ , or else it is of the form $v\downarrow$, where $v = \bigvee (f^\sigma)^{-1}[u\downarrow]$ satisfies $f^\sigma(v) \leq u$. In order to show that $v\downarrow$ is open in σ , it suffices to prove that v is an open element of \mathbb{A}^σ .

Consider an arbitrary closed element $p \leq v$; then $\bigwedge f[p\uparrow \cap A] = f^\sigma(p) \leq f^\sigma(v) \leq u$. Hence by compactness there is a finite set $F \subseteq p\uparrow \cap A$ such that $\bigwedge f[F] \leq u$. Putting $a_p := \bigwedge F$ we find that $a_p \in A$, $p \leq a_p$ and $a_p \leq v$ since $f(a_p) \leq \bigwedge f[F]$. Clearly then $v = \bigvee \{p \mid v \geq p \in K\} \leq \bigvee \{a_p \mid v \geq p \in K\} \leq v$ which shows that v is *identical* to the second join, and hence, open.

Finally, part (5) follows from part (4) by order duality. ■

As we announced already in the introduction to this section, the following properties will be crucial in proving canonicity results further on. The reason for this lies in the observation that for some terms t , we may apply Proposition 1.118(2) by the fact that the term function $t^{\mathbb{A}^\sigma}$ in the canonical extension \mathbb{A}^σ can be decomposed as $t^{\mathbb{A}^\sigma} = g^\sigma \circ f^\sigma$ where g^σ is (τ, γ^\uparrow) -continuous and f^σ is (σ, τ) -continuous, for some ‘intermediate’ topology τ . This is the principle of *matching continuities* that we mentioned in the introduction.

PROPOSITION 1.118. *Let $f : \mathbb{A} \rightarrow \mathbb{B}$ and $g : \mathbb{B} \rightarrow \mathbb{C}$ be monotone maps between the Boolean algebras \mathbb{A} , \mathbb{B} and \mathbb{C} . Then*

1. $(gf)^\sigma \leq g^\sigma f^\sigma$;
2. $(gf)^\sigma \geq g^\sigma f^\sigma$ whenever $g^\sigma f^\sigma$ is $(\sigma, \gamma^\uparrow)$ -continuous.

Proof. Part (2) of the proposition is an immediate consequence of Proposition 1.117(1) since $g^\sigma f^\sigma$ is an extension of gf (and gf is monotone). Concerning part (1), we first show that $(gf)^\sigma(p) \leq g^\sigma f^\sigma(p)$ for closed p . Note that

$$\begin{aligned} (gf)^\sigma(p) &= \bigwedge \{gf(a) \mid p \leq a \in A\}, \\ g^\sigma f^\sigma(p) &= \bigwedge \{g(b) \mid f^\sigma(p) \leq b \in B\}. \end{aligned}$$

where the latter identity holds because $f^\sigma(p)$ is closed in \mathbb{A}^σ . Now take a $b \in B$ with $f^\sigma(p) \leq b$. As $f^\sigma(p) = \bigwedge \{f(a) \mid p \leq a \in A\}$ is a down-directed meet, compactness

provides some $a \in A$ with $p \leq a$ and $f(a) \leq b$. Then $(gf)^\sigma(p) \leq gf(a) \leq gb$; and hence, $(gf)^\sigma(p) \leq g^\sigma f^\sigma(p)$.

Now we turn to arbitrary $x \in A^\sigma$. Note that

$$\begin{aligned} (gf)^\sigma(x) &= \bigvee \{(gf)^\sigma(p) \mid x \geq p \in K(A)\}, \\ g^\sigma f^\sigma(x) &= \bigvee \{g^\sigma(q) \mid f^\sigma(x) \geq q \in K(B)\}. \end{aligned}$$

Take an arbitrary $p \in K(A)$ with $p \leq x$; then $(gf)^\sigma(p) \leq g^\sigma f^\sigma(p)$, as we just saw. Since $f^\sigma(x) \geq f^\sigma(p) \in K(B)$, this shows that every joinand $(gf)^\sigma(p)$ of $(gf)^\sigma(x)$ is below some joinand $g^\sigma(q)$ of $g^\sigma f^\sigma(x)$. This suffices to prove the desired inequality. ■

7.5 Canonical equations

Time to harvest. The key idea for proving canonicity results for an equation $s \approx t$ will be to use properties of the *term functions* $s^\mathbb{A}$ and $t^\mathbb{A}$. Recall that for a term $t(x_1, \dots, x_n)$, the term function $t^\mathbb{A} : A^n \rightarrow A$ is inductively defined as follows:

$$\begin{aligned} x_i^\mathbb{A} &:= \pi_i^n, \\ (\heartsuit(t_1, \dots, t_n))^\mathbb{A} &:= \heartsuit^\mathbb{A} \circ \langle t_1^\mathbb{A}, \dots, t_n^\mathbb{A} \rangle. \end{aligned}$$

where $\pi_i^n : (a_1, \dots, a_n) \mapsto a_i$ is the i -th projection function, and, for maps $f_1, \dots, f_n : X \rightarrow Y$, the map $\langle f_1, \dots, f_n \rangle : X \rightarrow Y^n$ is given by $\langle f_1, \dots, f_n \rangle(x) = (f_1(x), \dots, f_n(x))$.

In the context of canonical extensions the following definitions are crucial.

DEFINITION 1.119. A term t is *expanding* on an expanded Boolean algebra \mathbb{A} if $(t^\mathbb{A})^\sigma \leq t^{\mathbb{A}^\sigma}$, *contracting* if $(t^\mathbb{A})^\sigma \geq t^{\mathbb{A}^\sigma}$, and *stable* if $(t^\mathbb{A})^\sigma = t^{\mathbb{A}^\sigma}$. We let these properties apply to classes of algebras in case they apply to all members of the class.

PROPOSITION 1.120. *Let s and t be two τ -terms, and \mathbb{K} a class of τ -expanded Boolean algebras. If s is contracting and t is expanding on \mathbb{K} , then the inequality $s \preceq t$ is canonical on \mathbb{K} .*

Proof. Consider an algebra \mathbb{A} in \mathbb{K} such that $\mathbb{A} \models s \preceq t$. In other words, we have $s^\mathbb{A} \leq t^\mathbb{A}$, so that $(s^\mathbb{A})^\sigma \leq (t^\mathbb{A})^\sigma$. But then by the assumptions on s and t it follows that $s^{\mathbb{A}^\sigma} \leq (s^\mathbb{A})^\sigma \leq (t^\mathbb{A})^\sigma \leq t^{\mathbb{A}^\sigma}$, which shows that $\mathbb{A}^\sigma \models s \preceq t$. ■

So which terms are contracting, and which ones are expanding? Here the topologies prove their value. Before moving on to these results, we need to get one technicality out of the way. Basically, the following proposition states that the product map $\langle f_1, \dots, f_n \rangle$ behaves as well as one could hope for.

PROPOSITION 1.121. *Let f_1, \dots, f_n be monotone maps between the Boolean algebras \mathbb{A} and \mathbb{B} . Then*

$$\langle f_1, \dots, f_n \rangle^\sigma = \langle f_1^\sigma, \dots, f_n^\sigma \rangle,$$

and for all $\rho, \rho' \in \{\gamma^\downarrow, \gamma^\uparrow, \gamma, \sigma^\downarrow, \sigma^\uparrow, \sigma\}$ it holds that

$$\langle f_1, \dots, f_n \rangle^\sigma \text{ is } (\rho, \rho')\text{-continuous iff each } f_i^\sigma \text{ is } (\rho, \rho')\text{-continuous.}$$

We leave the rather tedious but not very difficult proof of this proposition to the reader, and move on to more interesting facts. First we associate topological properties with term functions.

PROPOSITION 1.122. *Let \mathbb{A} be a τ -expanded Boolean algebra, and t a τ -term. Then*

1. *If \mathbb{A} interprets all connectives in t as operators, then $t^{\mathbb{A}\sigma}$ is $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous.*
2. *If \mathbb{A} interprets all connectives in t as additive maps, then $t^{\mathbb{A}\sigma}$ is $(\sigma^\downarrow, \sigma^\downarrow)$ -continuous.*
3. *If \mathbb{A} interprets all connectives in t as multiplicative maps, then $t^{\mathbb{A}\sigma}$ is $(\sigma^\uparrow, \sigma^\uparrow)$ -continuous.*

Proof. All three statements can be proved by a straightforward term induction, using the Propositions 1.117 and 1.121 for the induction step. For the induction base, note that the projection maps are both join- and meet preserving, and hence, their canonical extensions have all the continuity properties mentioned in the statements of this proposition. ■

Here we arrive at the core of the algebraic approach towards the canonicity of equations. On the basis of the syntactic shape of some terms we can see whether it is expanding or stable. In Theorem 1.123 we give some sample results; observe that the key idea in the proof of part (3) is the principle of ‘matching continuities’ as described before Proposition 1.118.

THEOREM 1.123. *Let \mathbb{A} be a τ -expanded Boolean algebra, and t a τ -term. Then*

1. *If \mathbb{A} interprets all connectives in t as monotone maps, then t is expanding.*
2. *If \mathbb{A} interprets all connectives in t as operators or dual operators, then t is stable.*
3. *If t is of the form $s(u_1, \dots, u_n)$ such that \mathbb{A} interprets all connectives in s as operators, and all connectives in each of the u_i as meet-preserving operations, then t is stable.*

Proof. Part (1) is proved by term induction. The base case is immediate from the definitions. For the inductive step, suppose that $t \equiv \nabla(t_1, \dots, t_n)$, then

$$\begin{aligned}
 (t^{\mathbb{A}})^\sigma &= (\nabla^{\mathbb{A}} \circ \langle t_1^{\mathbb{A}}, \dots, t_n^{\mathbb{A}} \rangle)^\sigma \\
 &\leq (\nabla^{\mathbb{A}})^\sigma \circ \langle t_1^{\mathbb{A}}, \dots, t_n^{\mathbb{A}} \rangle^\sigma \\
 &= \nabla^{\mathbb{A}\sigma} \circ \langle (t_1^{\mathbb{A}})^\sigma, \dots, (t_n^{\mathbb{A}})^\sigma \rangle \\
 &\leq \nabla^{\mathbb{A}\sigma} \circ \langle t_1^{\mathbb{A}\sigma}, \dots, t_n^{\mathbb{A}\sigma} \rangle \\
 &= t^{\mathbb{A}\sigma}.
 \end{aligned}$$

Here the first and last step are by definition, the second step is by Proposition 1.118(1) and monotonicity, the third step is by definition of $\nabla^{\mathbb{A}\sigma} = (\nabla^{\mathbb{A}})^\sigma$ and by Proposition 1.121, and the fourth step is by the inductive hypothesis and the monotonicity of $\nabla^{\mathbb{A}\sigma}$.

For part (2) and (3) it suffices to prove that $t^{\mathbb{A}\sigma} \leq (t^{\mathbb{A}})^\sigma$, since the opposite inequality holds by part (1). In the case of part (2) this follows from a straightforward induction, whereas for part (3) we need the principle of matching topologies.

Let t be as described in part (3), then $t^{\mathbb{A}\sigma} = s^{\mathbb{A}\sigma} \circ \langle u_1^{\mathbb{A}\sigma}, \dots, u_n^{\mathbb{A}\sigma} \rangle = (s^{\mathbb{A}})^\sigma \circ \langle (u_1^{\mathbb{A}})^\sigma, \dots, (u_n^{\mathbb{A}})^\sigma \rangle$ with the second identity holding by part (2). Also, note that by Proposition 1.117, the term function $s^{\mathbb{A}\sigma}$ is $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous, and each $u_i^{\mathbb{A}\sigma}$ is $(\sigma^\uparrow, \sigma^\uparrow)$ -continuous. From this we infer by Proposition 1.121 that the map $\langle (u_1^{\mathbb{A}})^\sigma, \dots, (u_n^{\mathbb{A}})^\sigma \rangle = \langle u_1^{\mathbb{A}}, \dots, u_n^{\mathbb{A}} \rangle^\sigma$

is $(\sigma^\uparrow, \sigma^\uparrow)$ -continuous as well, whence by $\gamma^\uparrow \subseteq \sigma^\uparrow$ it is $(\sigma^\uparrow, \gamma^\uparrow)$ -continuous. Thus the $(\gamma^\uparrow, \gamma^\uparrow)$ -continuity of $s^{\mathbb{A}^\sigma}$ matches with the $(\sigma^\uparrow, \gamma^\uparrow)$ -continuity of $\langle u_1^\mathbb{A}, \dots, u_n^\mathbb{A} \rangle^\sigma$. Hence, we may apply Proposition 1.118(2), and find that $t^{\mathbb{A}^\sigma} = (s^\mathbb{A})^\sigma \circ \langle u_1^\mathbb{A}, \dots, u_n^\mathbb{A} \rangle^\sigma \leq (s^\mathbb{A} \circ \langle u_1^\mathbb{A}, \dots, u_n^\mathbb{A} \rangle)^\sigma = (t^\mathbb{A})^\sigma$, as desired. ■

As a sample application, we show how Sahlqvist canonicity is an easy consequence of the previous theorem.

COROLLARY 1.124. *Sahlqvist equations are canonical over the class of all Boolean algebras with τ -operators.*

Proof. First we treat inequalities of the form $\varphi(\beta_1, \dots, \beta_n) \leq \psi$, where φ only uses \wedge , \vee and modalities, all β_i are boxed atoms, and ψ is positive. But then it is immediate by the previous proposition that $\varphi(\beta_1, \dots, \beta_n)$ is stable, while ψ is expanding. Hence the result follows from Proposition 1.120.

Now consider an arbitrary Sahlqvist inequality. Without loss of generality we may assume that it is in fact an equation of the form

$$\varphi(\beta_1, \dots, \beta_n, \neg\psi_1, \dots, \neg\psi_k) \approx \perp, \quad (1.34)$$

where φ and the β 's are as before, while all ψ_j are positive formulas. It is easy to see that this equation is equivalent to the quasi-equation

$$\left(\bigwedge_{1 \leq i \leq n} x_i \leq \neg\psi_i \right) \Rightarrow \varphi(\beta_1, \dots, \beta_n, x_1, \dots, x_k) \approx \perp,$$

which in its turn is equivalent to

$$\left(\bigwedge_{1 \leq i \leq n} x_i \wedge \psi_i \approx \perp \right) \Rightarrow \varphi(\beta_1, \dots, \beta_n, x_1, \dots, x_k) \approx \perp. \quad (1.35)$$

Now suppose that we *add* a diamond \mathbf{E} to the language, and interpret this diamond as the global modality on every algebra (see section 8.2). Then clearly the quasi-equation (1.35) is equivalent to the formula

$$\varphi(\beta_1, \dots, \beta_n, x_1, \dots, x_k) \preceq \bigvee_{1 \leq i \leq n} \mathbf{E}(x_i \wedge \psi_i). \quad (1.36)$$

(Note that this reduction of a quasi-equation to an equivalent equation is a specific example of Proposition 1.138.)

The result then follows by the observation that (1.36) is a Sahlqvist inequality of the kind already treated, together with the fact that the canonical extension of the global modality is again the global modality (see Remark 1.110). ■

7.6 Further remarks

The ideas described in this section allow for variations and generalizations in at least two directions.

To start with, the algebraic approach has already been put to work for a far wider class of structures than just Boolean algebras with operators. In particular, nothing in the

theory crucially depends on the *Boolean* nature of the underlying order of the algebras. The notion of a canonical extension, with all the results in section 7.2 pertaining to them, has been extended to (first distributive and then) arbitrary lattices, with work on partial orders under way.

Furthermore, the restriction to monotone operations is not necessary either; *arbitrary* maps between lattices can be extended to maps between their canonical extensions. First suppose that we are dealing with a dense set X' in a topology $\langle X, \rho \rangle$, and let $f : X' \rightarrow C$ be a map from X' to the carrier C of a complete lattice \mathbb{C} . Then define

$$\begin{aligned} f^\sigma(x) &:= \bigvee \{ \bigwedge f[U \cap X'] \mid x \in U \in \rho \}, \\ f^\pi(x) &:= \bigwedge \{ \bigvee f[U \cap X'] \mid x \in U \in \rho \}. \end{aligned} \quad (1.37)$$

In order to apply this definition for the canonical extension of a map f between two lattices \mathbb{L} and \mathbb{M} , note that (just like in the case for Boolean algebras, see Remark 1.115) the carrier L of \mathbb{L} forms a dense subset of the σ -topology over the carrier \mathbb{L}^σ . Also observe that f^σ and f^π are *extensions* of f because all elements of L are *isolated* points of f , and that for *monotone* f , (1.37) agrees with (1.32).

Finally, it is not just the definitions that translate to the more general setting of lattice expansions (that is, lattices with additional operations), the same holds for the theory. To mention just one example: one may prove that any equation $s \approx t$ is canonical provided that all the primitive symbols (including the join operation \wedge) occurring in s and t are interpreted as operators. Details can be found in for instance GEHRKE & HARDING [28].

The second generalization that we want to mention involves other ways of completing lattices and lattice expansions, such as the *MacNeille completion*, which generalizes Dedekind's construction of the reals from the rationals to arbitrary partial orders. For a characterization in the style of this section, one may start by proving that any lattice \mathbb{L} has a (modulo isomorphism) unique completion \mathbb{L}^μ , its *MacNeille completion*, in which \mathbb{L} is both join- and meet dense. This way of extending lattices is obviously similar to that of the canonical extension, but a substantial difference is that the MacNeille completion agrees with the original lattices on *all* meets and joins, whereas the canonical extension only agrees on the *finite* ones.

In any case, it follows from join- and meet density, that any map between two lattices can be extended to a map between their MacNeille completions, in two ways. In the case of a *monotone* operation f between two lattices \mathbb{L} and \mathbb{M} , we define the *lower extension* \check{f} and the *upper extension* \hat{f} by

$$\begin{aligned} \check{f}(x) &:= \bigvee \{ f(a) \mid x \geq a \in L \} \\ \hat{f}(x) &:= \bigwedge \{ f(a) \mid x \leq a \in L \} \end{aligned}$$

Clearly then, almost all questions concerning canonical extensions have an obvious counterpart for MacNeille completions. Generally speaking, MacNeille completions are less well-behaved than canonical extensions; for instances, unary operators (diamonds) are no longer smooth, and the variety of modal algebras is not closed under taking lower MacNeille completions. Probably for this reason, MONK [85] introduced the notion of the MacNeille completion of a BAO only for Boolean algebras with *complete* operators. On the other hand, in case the primitive operations are *residuated* (see Proposition 1.129),

the situation improves; for instance, GIVANT & VENEMA [36] show that the validity of all Sahlqvist equations is preserved under taking MacNeille completions of *tense* algebras. As a final remark, there are interesting connections between the MacNeille completion and the canonical extension of a lattice expansion: for instance, GEHRKE, HARDING & VENEMA [29] prove that the canonical extension of lattice expansion \mathbb{A} can be embedded in the MacNeille completion of some ultrapower of \mathbb{A} . As a consequence, every variety of lattice expansions that is closed under taking MacNeille completions, is also canonical in the sense of canonical extensions.

8 SPECIAL ALGEBRAIC TOPICS

In this final section on algebra we discuss the algebraic perspective on two further issues in modal logic.

8.1 *Tense logic*

Our first example concerns *tense logic*; as its name already indicates, this branch of modal logic originates in the formal semantics of natural language, cf. Chapter 19 of this volume.

DEFINITION 1.125. The modal similarity type ϑ of *tense logic* is fixed by its two diamonds, \diamond_F and \diamond_P .

The letters \diamond_F and \diamond_P are mnemonic of **f**uture and **p**ast, respectively. This already indicates that the standard interpretation of this language is in frames representing a flow of time, such that \diamond_F obtains the meaning ‘sometime in the future’, and dually \diamond_P means ‘sometime in the past’. Tense logic thus forms a rather simple example of *temporal* logic, cf. Chapter 11 of this volume. Here we abstract from the temporal interpretations of tense logics; what is then left is that in the intended frames for this language, the two diamonds of the language are interpreted along the two directions of a *single* binary relation.

DEFINITION 1.126. A ϑ -frame $\mathbb{S} = \langle S, R_F, R_P \rangle$ is called *bidirectional* if R_F and R_P are each other’s converse.

This definition explains why a ϑ -frame is often represented simply as the pair $\langle S, R_F \rangle$. Turning to logic, we define the following.

DEFINITION 1.127. A modal ϑ -logic L is a *tense logic* if both formulas $p \rightarrow \Box_F \diamond_P p$ and $p \rightarrow \Box_P \diamond_F p$ are theorems of L ; the minimal tense logic is denoted as \mathbf{K}_t . Algebraically, a *tense algebra* is a Boolean algebra expanded with monotone ϑ -operations satisfying the corresponding equations $x \preceq \Box_F \diamond_P x$ and $x \preceq \Box_P \diamond_F x$.

It is easy to see that \mathbb{S}^+ is a tense algebra if and only if \mathbb{S} is a bidirectional frame. In the other direction, it is not a priori clear whether we can extract a useful frame from an arbitrary tense algebra: First we must show that tense algebras are Boolean algebras with operators. In fact, already JÓNSSON & TARSKI [70] show something better.

THEOREM 1.128. *Let $\mathbb{A} = \langle A, \top, \perp, -, \wedge, \vee, \diamond_F, \diamond_P \rangle$ be a tense algebra. Then*

1. *the operations \diamond_F and \diamond_P are complete operators;*

2. the structure \mathbb{A}_\bullet is a bidirectional frame, and the algebra \mathbb{A}^σ is again a tense algebra.

Proof. For part 1 of the Theorem, let $a \in A$ be the least upper bound of some subset X of A . Then by monotonicity, $\diamond_F a$ is an upper bound of the set $\diamond_F[X]$. Now suppose that b is also an upper bound of this set, that is, $\diamond_F x \leq b$ for all $x \in X$. From this it follows, for each $x \in X$, that $x \leq \square_P \diamond_F x \leq \square_P b$ (here we use monotonicity of \square_P , which is easily proven). Thus we see that $a \leq \square_P b$ by our assumption on a . But then by monotonicity of \diamond_F we obtain that $\diamond_F a \leq \diamond_F \square_P b \leq b$. This proves that $\diamond_F a$ is in fact the least upper bound of the set $\diamond_F[X]$.

Concerning the second part of the theorem, that \mathbb{A}^σ is a tense algebra is a special of the Sahlqvist Canonicity Theorem 1.93; the bidirectionality of \mathbb{A}_\bullet is then immediate since $\mathbb{A}^\sigma = (\mathbb{A}_\bullet)^+$. ■

There is a lot more to say about the complete additivity of the diamonds in tense algebras. To start with, the definition of tense algebras can be reformulated using either of the algebraically more familiar notions of *conjugation* or *residuation*.

PROPOSITION 1.129. *Let $\mathbb{A} = \langle A, \top, \perp, -, \wedge, \vee, \diamond_F, \diamond_P \rangle$ be a monotone ϑ -expanded Boolean algebra. Then the following are equivalent:*

1. \mathbb{A} is a tense algebra,
2. \diamond_F and \diamond_P are conjugated operations, that is, they satisfy the following:

$$\mathbb{A} \models \forall xy (x \wedge \diamond_F y \approx \perp \Leftrightarrow y \wedge \diamond_P x \approx \perp), \quad (1.38)$$

3. \diamond_F and \square_P form a residual pair, that is,

$$\mathbb{A} \models \forall xy (\diamond_F x \preceq y \Leftrightarrow x \preceq \square_P y). \quad (1.39)$$

This connection with residuation shows that from a general mathematical perspective, tense logic is not just *any* bimodal logic: It provides the modal logic manifestation of the fundamental category theoretic concept of *adjoint functors*. Theorem 1.128(1) is thus a rather special case of the category theoretic fact that left adjoint functors preserve all (existing) colimits.

Another nice property of tense logic that should be mentioned here is that somehow, tense algebras are *richer* than ordinary Boolean algebras with operators. For instance, consider an atomic modal algebra \mathbb{A} , and suppose that \mathbb{A} satisfies some Sahlqvist equation η . Then it is *not* guaranteed that the atom structure \mathbb{A}_\bullet (see Definition 1.41) satisfies the first-order correspondent c_η of η , not even if the diamond of \mathbb{A} is completely additive. However, in case \mathbb{A} is a *tense algebra*, it contains sufficient information to enforce this.

THEOREM 1.130. *Let \mathbb{A} be an atomic tense algebra. Then for every Sahlqvist equation η : $\mathbb{A} \models \eta$ iff $\mathbb{A}_+ \models c_\eta$ iff $(\mathbb{A}_\bullet)^+ \models \eta$.*

Proof. Clearly, the equivalence of the last two statements follows from Sahlqvist correspondence theory. For the implication from right to left, it suffices to observe that \mathbb{A} is a subalgebra of $(\mathbb{A}_\bullet)^+$ because of the complete additivity of the operators. This follows from (1.12) in the proof of Proposition 1.42.

The remaining implication is a special case of the preservation of Sahlqvist equations under taking (lower) MacNeille completions of tense algebras, see the end of section 7 for some discussion, and GIVANT & VENEMA [36] for proofs. ■

Finally, tense algebras play a role in other part of universal algebra as well. For instance, any *lattice* can be represented as the sublattice of a tense algebra that has the solution set of the equation $x \approx \Box_P \Diamond_F x$ as its carrier. This idea basically goes back to BIRKHOFF [12]; for more details, the reader is referred to HARDING [56].

Nevertheless, despite their rather special characteristics, just like all bimodal logics, tense logics can be *simulated* by monomodal ones; for details we refer to Chapter 8 of this volume.

8.2 Global modality \mathcal{E} discriminator varieties

Recent years have witnessed an increasing interest in formalisms that enhance the expressive power of standard modal languages, see for instance Chapter 14 of this volume. In such a pursuit, one naturally arrives at the *global* or *universal* modality \mathbf{E} which has the global relation $S \times S$ of a frame \mathbb{S} as its (intended) accessibility relation, see GORANKO & PASSY [48]. But also, a large number of standard logics come with an intended semantics in which the global relation interprets some more complex term of the language: as an example we mention the compound modality $\Diamond_F \Diamond_P$ in the tense logic over any linear flow of time.

DEFINITION 1.131. Algebraically, we define the *global modality* or *unary discriminator* over a Boolean algebra (with operators) \mathbb{B} as the function given by

$$b \mapsto \begin{cases} \perp & \text{if } b = \perp, \\ \top & \text{if } b > \perp. \end{cases}$$

The term $\gamma(x)$ is called a *global modality* or *unary discriminator term* over an expanded Boolean algebra \mathbb{A} if it is interpreted as the global modality on \mathbb{A} .

This notion can be seen as the BAO manifestation of the well-known algebraic concept of a *discriminator*, see JIPSEN [67] for a first explicit discussion of the connections.

DEFINITION 1.132. We call a ternary term d a *discriminator term* over an algebra \mathbb{A} if it is interpreted as the discriminator function on A , that is, if $d^{\mathbb{A}}(a, b, c) = a$ if $a \neq b$, and $d^{\mathbb{A}}(a, b, c) = c$ if $a = b$. Any variety \mathbf{V} generated by a class of algebras with a common discriminator term, is called a *discriminator variety*.

PROPOSITION 1.133. *Let \mathbb{A} be a τ -expanded Boolean algebra.*

1. *If γ is a global modality for \mathbb{A} , then the term $(\gamma(\neg(x \leftrightarrow y)) \wedge x) \vee (\gamma(\neg(x \leftrightarrow y)) \wedge z)$ is a discriminator term for \mathbb{A} .*
2. *If $d(x, y, z)$ is a discriminator term for \mathbb{A} , then the term $\neg d(\perp, x, \top)$ is a global modality for \mathbb{A} .*

Before going into further detail of the connection with the global modality, let us, for future reference, list some of the many nice properties that discriminator varieties have.

THEOREM 1.134. *Let \mathbf{K} be a class of algebras with a discriminator term d . Then*

1. *all algebras in \mathbf{K} are simple;*
2. *$\text{Var}(\mathbf{K})$ is congruence-distributive and congruence-permutable;*

3. all subdirectly irreducible algebras in $\text{Var}(\mathbf{K})$ are simple, and vice versa;
4. $\text{Var}(\mathbf{K})$ is semi-simple; that is, every algebra in $\text{Var}(\mathbf{K})$ is a subdirect product of simple algebras.
5. d is a discriminator term for every simple algebra in $\text{Var}(\mathbf{K})$.

Proof. For the first statement of the theorem, define the term $s(x, y, u, v) := d(d(x, y, u), d(x, y, v), v)$. It is easy to see that s is a so-called *switching term* for \mathbf{K} ; that is, for every \mathbb{A} in \mathbf{K} , and for all a, b, c and d in \mathbb{A} :

$$s^{\mathbb{A}}(a, b, c, d) = \begin{cases} c & \text{if } a = b, \\ d & \text{if } a \neq b. \end{cases}$$

Now let $\Theta \neq \Delta_{\mathbb{A}}$ be a congruence of \mathbb{A} ; then there are two elements $a \neq b$ with $(a, b) \in \Theta$. But then we find $(c, d) = (s^{\mathbb{A}}(a, a, c, d), s^{\mathbb{A}}(a, b, c, d)) \in \Theta$ for every c and d in Θ . In other words, such a Θ must be the trivial congruence $A \times A$. But this clearly means that \mathbb{A} is simple. Details of the proof of the second statement, which is similar to that of Theorem 1.25, are left to the reader.

For the third part of the theorem, it is not hard to verify that d is a discriminator term for $\text{SPu}(\mathbf{K})$ as well, whence $\text{SPu}(\mathbf{K})$ consists of simple algebras by part (1). So by definition of simplicity, we find that $\text{HSPu}(\mathbf{K}) = \text{SPu}(\mathbf{K})$; hence, all algebras in $\text{HSPu}(\mathbf{K})$ are simple. However, by part 2 we may apply Jónsson's Lemma, which states that all s.i. members of $\text{Var}(\mathbf{K})$ belong to $\text{HSPu}(\mathbf{K})$. Thus every s.i. algebra in $\text{Var}(\mathbf{K})$ is simple.

Part (4) is immediate from part (3) by Birkhoff's subdirect indecomposability theorem, while the final statement follows from the fact that every simple algebra belongs to $\text{SPu}(\mathbf{K})$, and thus shares the discriminator term of \mathbf{K} . ■

In particular, since the notions of simplicity and subdirect irreducibility coincide in a discriminator variety, its subvarieties are completely determined by its simple members. Let us now see how these issues are axiomatized in normal modal logics.

DEFINITION 1.135. A τ -formula $\gamma(x)$ is a *global modality* for a normal modal τ -logic L if the formulas Γ

- $\nabla(x_1, \dots, x_n) \rightarrow \gamma(x_i)$ for every $\nabla \in \tau$, and every $i \in \{1, \dots, n\}$;
- $x \rightarrow \gamma(x)$, $\gamma(\gamma(x)) \rightarrow \gamma(x)$ and $\gamma(\neg\gamma(\neg x)) \rightarrow x$;

are theorems of L .

That is, L defines a global modality iff there is a term $\gamma(x)$ that satisfies the **S5** axioms, plus the inclusion axiom $\nabla^i x \rightarrow \gamma(x)$ for every induced diamond ∇^i . It is not hard to derive that such an axiomatically defined global modality $\gamma(x)$ also has $\gamma(\neg\gamma(x)) \rightarrow \neg\gamma(x)$, and $\vdash_L \blacklozenge x \rightarrow \gamma(x)$ for all compound diamonds \blacklozenge .

The terminology of Definition 1.135 is justified by the following Proposition, which is essentially taken from JIPSEN [67].

PROPOSITION 1.136. *Let L be a normal modal τ -logic, and $\gamma(x)$ a τ -formula. Then $\gamma(x)$ is a global modality for L if and only if $\text{BAO}_{\tau}(L) = \text{Var}(\mathbf{K})$ for some class \mathbf{K} of algebras sharing γ as a global modality.*

Proof. The direction from right to left is immediate by the fact that any unary discriminator term satisfies all the formulas listed in Definition 1.135.

For the other direction, by Theorem 1.34 it suffices to show that γ is a unary discriminator term on subdirectly irreducible algebras in $\text{BAO}_\tau(L)$. In order to prove this, suppose for contradiction that \mathbb{A} has a radical element ρ , while $\gamma^\mathbb{A}$ is not the global modality on \mathbb{A} . That is, some $a \in A$ satisfies $a \neq \perp$ while $\gamma^\mathbb{A}(a) \neq \top$, whence $-\gamma^\mathbb{A}(a) \neq \perp$. Since ρ is radical in \mathbb{A} there are compound diamonds \blacklozenge_1 and \blacklozenge_2 such that $\rho \leq \blacklozenge_1 a$ and $\rho \leq \blacklozenge_2 -\gamma^\mathbb{A}(a)$. However, from $\rho \leq \blacklozenge_1 a$ we obtain $\rho \leq \gamma^\mathbb{A}(a)$, while from $\rho \leq \blacklozenge_2 -\gamma^\mathbb{A}(a)$ we may infer that $\rho \leq \gamma^\mathbb{A}(-\gamma^\mathbb{A}(a)) \leq -\gamma^\mathbb{A}(a)$. This contradicts the fact that $\rho > \perp$, and so we may conclude that γ is the global modality on \mathbb{A} . ■

A very useful property of discriminators is that they allow the effective replacement of universal sentences with equations. In the case of BAOs, this works out as follows.

DEFINITION 1.137. Suppose that $\gamma(x)$ is a global modality term for \mathbf{K} . Inductively we define a function λ mapping quantifier-free formulas (in the first order language of BAOs) to τ -terms:

$$\begin{aligned} s \approx t &\mapsto (s \wedge \neg t) \vee (\neg s \wedge t), \\ \sim P &\mapsto \neg \gamma(\lambda_P), \\ P \& Q &\mapsto \lambda_P \vee \lambda_Q. \end{aligned}$$

THEOREM 1.138. *Let \mathbf{K} be a class of Boolean algebras with τ -operators with a discriminator term γ . Then any universal formula P is equivalent over \mathbf{K} to the equation $\lambda_{P'} \approx \perp$, where P' is the quantifier-free part of P .*

Proof. A straightforward induction shows that for any algebra \mathbb{A} in \mathbf{K} , any assignment α on \mathbb{A} and any quantifier-free formula P it holds that

$$\mathbb{A} \models_\alpha P \text{ iff } \mathbb{A} \models_\alpha \lambda_P \approx \perp.$$

From this, the statement of the theorem is immediate. ■

Working with discriminator classes has many advantages. For instances, if \mathbf{K} is a discriminator class, then we may generate $\text{Var}(\mathbf{K})$ from \mathbf{K} just by taking products and subalgebras (that is, homomorphic images are not needed). The result in this generality is due to GIVANT [35].

THEOREM 1.139. *Let \mathbf{K} be a class of Boolean algebras with a common global modality term $\gamma(x)$.*

1. *If $\text{Pu}(\mathbf{K}) \subseteq \mathbf{S}(\mathbf{K})$, then $\text{SP}(\mathbf{K})$ is a variety and $\mathbf{S}(\mathbf{K})$ is the universal class of simple algebras in $\text{SP}(\mathbf{K})$.*
2. *If \mathbf{K} is axiomatized by a set Φ of universal formulas, then $\text{SP}(\mathbf{K})$ is axiomatized by the set $\{\lambda_P \approx \perp \mid P \in \Phi\}$, together with the set Γ of Definition 1.135.*

Proof. Assume that $\text{Pu}(\mathbf{K}) \subseteq \mathbf{S}(\mathbf{K})$, then it is easy to see that the class $\mathbf{S}(\mathbf{K})$ is closed under taking ultraproducts and subalgebras. It then follows by standard universal algebra, see [17, Theorem 2.20], that $\mathbf{S}(\mathbf{K})$ is a universal class, that is, an elementary class axiomatized by universal formulas.

By assumption, the algebras in \mathbf{K} have a common discriminator term, and, hence, we find, reasoning as in the proof of Theorem 1.134(3), that $\text{SirVar}(\mathbf{K}) = \text{SPu}(\mathbf{K})$, where $\text{SirVar}(\mathbf{K})$ denotes the class of s.i. members in $\text{Var}(\mathbf{K})$. Thus by the assumption we find that $\text{SirVar}(\mathbf{K}) = \mathbf{S}(\mathbf{K})$ and therefore, $\mathbf{S}(\mathbf{K})$ is the class of simple algebras in $\text{Var}(\mathbf{K})$, since the notions of simplicity and subdirect irreducibility coincide. Finally then, by Birkhoff's and Jónsson's theorems, the variety $\text{Var}(\mathbf{K})$ is the class of subdirect products of algebras in $\text{HSPu}(\mathbf{K}) = \mathbf{S}(\mathbf{K})$; a straightforward calculation then will show that $\text{Var}(\mathbf{K}) = \text{SP}(\mathbf{K})$.

Part two of the theorem is a straightforward consequence of Proposition 1.136 and Theorem 1.138. ■

Finally, for more information on the global modality, the reader is referred to Chapter 8 of this volume.

9 COALGEBRAS: AN INTRODUCTION

This section forms a brief introduction to the field of *Coalgebra*. While certain kinds of coalgebras had already been studied in the sixties, the field really took off after it was realized that coalgebra can be conceived as a general and uniform theory of dynamic systems, taken in a broad sense.

Many structures in mathematics and theoretical computer science can naturally be represented as coalgebras. Probably the first example was provided by ACZEL [2], who models transition systems and non-well-founded sets as coalgebras. On the basis of Aczel's work, BARWISE & MOSS [11] discuss a wide range of phenomena involving the notions of circularity and self-reference, with applications ranging from theoretical economics to the semantics of natural language. A second paradigmatic specimen of coalgebras in computer science is given by (deterministic) automata, see RUTTEN [96]. Further important examples include the representation of infinite data structures, and the formal modeling of objects and classes in object oriented programming, see REICHEL [92] or JACOBS [61]. But for modal logicians, it will be Kripke frames and models that provide the prime examples of coalgebras; this link goes back to at least ABRAMSKY [1]. In fact, the model theory of modal logic is coalgebraic in nature, so modal logicians entering the field will have much the same experience as group theorists learning about universal algebra, in that they will recognize many familiar notions and results, lifted to a higher level of generality and abstraction.

For readers that want to learn more about coalgebras, the literature harbors some well written introductions and surveys (although at the time of writing there is no text book or monograph available). We refer the reader to JACOBS & RUTTEN [65] for a very accessible introduction, and to RUTTEN [97] or GUMM [50] for comprehensive surveys. IHRINGER [60] has an appendix on coalgebras by Gumm. For more details on the connection between coalgebra and modal logic, the reader may consult KURZ [75] or PATTINSON [90].

What then are coalgebras? The most concrete, state-based specimens, called *systems*, simply consist of a set S endowed with some kind of transition, formally modeled as some map σ from S to another set ΩS . Here Ω is some functor constituting the *type* or *signature* of the coalgebra at stake. The transition map provides some kind of structure on S , but whereas *algebraic* operations are ways to *construct* complex objects out of simple ones, coalgebraic operations, going *out of* the carrier set, should be seen as ways

to *unfold* or *observe* objects. This explains the central role of the notion of *behavior* in the theory of coalgebras.

More generally, given an endofunctor Ω on some base category \mathbb{C} , an Ω -coalgebra is a pair $\mathbb{C} = \langle C, \gamma \rangle$, with C an arbitrary object in \mathbb{C} , and γ a \mathbb{C} -arrow from C to ΩC . The full functorial power of Ω comes in when we turn Ω -coalgebras into a category $\text{Coalg}(\Omega)$ by introducing morphisms: A homomorphism from $\langle C, \gamma \rangle$ to $\langle C', \gamma' \rangle$ is an arrow $f : C \rightarrow C'$ such that $\gamma' \circ f = (\Omega f) \circ \gamma$. This set-up enables the canonical definition of two notions of equivalence between coalgebras, namely, bisimulation and behavioral equivalence. As we will see as well, the definitions make the concept of a coalgebra very similar to that of an algebra. However, if one makes this connection mathematically precise, it turns out that coalgebras over the base category \mathbb{C} are *dual* to algebras over the *opposite* category \mathbb{C}^{op} . This explains not only the name ‘coalgebra’, but, as we will see, also many of the peculiarities of *universal coalgebra*, that is, the general coalgebraic theory of systems.

Given the nature of coalgebra as a very general model of state-based dynamics, there is a natural place for *modal logic* as a formalism for reasoning about behavior. It was MOSS [11, 86] who realized that one may generalize the concept of modal logic from Kripke frames and models to coalgebras over arbitrary set functors. Over subsequent years, the development and study of modal languages for the specification of properties of coalgebras has been actively pursued and studied by various authors, including JACOBS [62, 64], KURZ [77, 76], PATTINSON [88, 89], and RÖSSIGER [95]. In fact, as we will see, the link between modal logics and coalgebra is so tight, that one may even claim that modal logic is the natural logic for coalgebras — just like equational logic is that for algebra.

We now turn to the technical development of the topic, starting with the definition of a *coalgebra*.

DEFINITION 1.140. Given an endofunctor Ω on a category \mathbb{C} , an Ω -coalgebra is a pair $\mathbb{A} = (A, \alpha)$, where A is an object of \mathbb{C} called the *carrier* of \mathbb{A} , and $\alpha : A \rightarrow \Omega A$ is an arrow in \mathbb{C} , called the *transition map* of \mathbb{A} . In case Ω is an endofunctor on Set , Ω -coalgebras may also be called Ω -systems; a *pointed Ω -system* is a triple $\langle A, \alpha, a \rangle$ such that $\langle A, \alpha \rangle$ is an Ω -system, and a is a state in \mathbb{A} , that is, an element of A .

As we mentioned already, the action of the functor Ω on the *arrows* of the category \mathbb{C} will be needed when we introduce, in Definition 1.148 below, homomorphisms between Ω -coalgebras. First we consider some examples of systems.

EXAMPLE 1.141. Probably the simplest example of a system is that of an C -colored set, that is, a pair $\langle S, \gamma : S \rightarrow C \rangle$. No matter where we start, this system can only display the color of the current state, and halt after doing so.

A slightly more interesting example is provided by a black box machine which may be prompted to display a value, or color, from C , and to move on to a next state. These states are internal to the machine, that is, invisible to an outside observer. Such a machine can abstractly be modeled as a coalgebra $\mu : M \rightarrow C \times M$, with $\pi_0(\mu(s)) \in C$ denoting the current value of the machine, and $\pi_1(\mu(s)) \in M$ representing the machine’s next internal state. (Here $\pi_0 : C \times M \rightarrow C$ and $\pi_1 : C \times M \rightarrow M$ are the projection functions.)

EXAMPLE 1.142. For our second example, we turn to automata theory. Recall that *deterministic automata* are usually modeled as quintuples $\mathbb{A} = \langle A, a_I, C, \delta, F \rangle$ such that A is the state space of the automaton \mathbb{A} , $a_I \in A$ is its initial state, C its *alphabet*,

$\delta : A \times C \rightarrow A$ its *transition function* and finally, $F \subseteq A$ its collection of *accepting states*.

Now observe that we may represent F by its characteristic map $\chi_F : A \rightarrow 2$ (with 2 denoting the set $\{0, 1\}$) which maps $a \in A$ to 1 if $a \in F$, and to 0 if $a \notin F$. Furthermore, we can and will view δ as a map from $A \rightarrow A^C$, where A^C denotes the collection of maps from C to A . Thus we see that we may represent a deterministic automaton over the alphabet C as a pointed system over the functor $S \mapsto 2 \times S^C$.

EXAMPLE 1.143. Our third example provides the crown witness when it comes to the connection between coalgebra and modal logics: We will now see that *frames* and *models* are in fact coalgebras in disguise. The crucial observation is here that a binary relation $R \subseteq S \times S$ can be represented by the function $R[\cdot] : S \rightarrow \mathcal{P}(S)$ mapping a point s to the collection $R[s]$ of its successors. Thus frames for the basic modal similarity type correspond to coalgebras over the covariant power set functor \mathcal{P} . (This functor maps a set S to its power set $\mathcal{P}(S)$ and a function $f : S \rightarrow S'$ to the image map $\mathcal{P}f$ given by $(\mathcal{P}f)(X) := f[X](= \{f(x) \mid x \in X\})$.)

Similarly, a ternary relation $T \subseteq S^3$ can be modeled as the function $T[\cdot] : S \rightarrow \mathcal{P}(S^2)$ given by $T[s] = \{(t_1, t_2) \in S^2 \mid Tst_1t_2\}$. Thus for any modal similarity type τ , we can represent τ -frames as coalgebras for the functor $S \mapsto \prod_{\nabla \in \tau} \mathcal{P}(S^{ar(\nabla)})$. Also note that *image finite frames*, that is, frames in which $R[s]$ is a finite set for all points s , correspond to coalgebras over the *finitary* power set functor \mathcal{P}_ω .

Concerning models, in this section we let \mathbf{Prop} denote the set of propositional variables. It is easy to see that a valuation $V : \mathbf{Prop} \rightarrow \mathcal{P}(S)$ on a frame $\mathbb{S} = \langle S, R \rangle$ could equivalently have been defined as a $\mathcal{P}(\mathbf{Prop})$ -coloring of S , that is, as the map sending a state s to the collection $V^{-1}[s] = \{p \in \mathbf{Prop} \mid s \in V(p)\}$ of proposition letters holding at s . Thus models for the basic modal similarity type can be identified with coalgebras of the functor Ω given by $X \mapsto \mathcal{P}(\mathbf{Prop}) \times \mathcal{P}(X)$.

EXAMPLE 1.144. For our last example, let $\check{\mathcal{P}}$ denote the *contravariant* power set functor. This functor agrees with the covariant power set functor on objects, while on arrows $\check{\mathcal{P}}$ takes *inverse* images. That is, for $f : A \rightarrow A'$, the function $\check{\mathcal{P}}f : \mathcal{P}A' \rightarrow \mathcal{P}A$ is given by $(\check{\mathcal{P}}f)(X') := f^{-1}[X'] (= \{x \in A \mid f(x) \in X'\})$. Note that $\check{\mathcal{P}}$ is not a functor from \mathbf{Set} to \mathbf{Set} , and thus does not produce coalgebras. Its composition with itself, however, *is* an endofunctor on \mathbf{Set} , so that we may consider $\check{\mathcal{P}} \circ \check{\mathcal{P}}$ -coalgebras. Because the transition function σ of such a coalgebra $\langle S, \sigma \rangle$ is a function $\sigma : S \rightarrow \mathcal{P}\mathcal{P}S$, the structure $\langle S, \sigma \rangle$ may also be seen as a *neighborhood frame*, as discussed in Chapter 1 of this volume.

Some variants of the functor $\check{\mathcal{P}} \circ \check{\mathcal{P}}$ are of interest as well — we discuss the examples $\mathcal{U}_{\check{\mathcal{P}}}$ and $\mathcal{F}_{\check{\mathcal{P}}}$. Recall that $\check{\mathcal{P}} \circ \check{\mathcal{P}}(S) = \mathcal{P}\mathcal{P}(S)$ is the set of *all* collections of subsets of S . $\mathcal{U}_{\check{\mathcal{P}}}(S)$ denotes the set of all *upward closed* collections of subsets of S , while $\mathcal{F}_{\check{\mathcal{P}}}(S)$ denotes the set of all *filters* of S . On arrows, these functors coincide with $\check{\mathcal{P}} \circ \check{\mathcal{P}}$; more precisely, for $f : S \rightarrow S'$, we set $\mathcal{U}_{\check{\mathcal{P}}}f$ and $\mathcal{F}_{\check{\mathcal{P}}}f$ as the restrictions of $(\check{\mathcal{P}} \circ \check{\mathcal{P}})f$ to $\mathcal{U}_{\check{\mathcal{P}}}S$ and $\mathcal{F}_{\check{\mathcal{P}}}S$, respectively.

It is not hard to show that $\mathcal{U}_{\check{\mathcal{P}}}$ and $\mathcal{F}_{\check{\mathcal{P}}}$ are indeed functors $\mathbf{Set} \rightarrow \mathbf{Set}$. The reader may in fact be familiar with (some) coalgebras for these functors. It can easily be verified that the $\mathcal{U}_{\check{\mathcal{P}}}$ -coalgebras correspond exactly to the *monotonic neighborhood frames* that were mentioned in Chapter 1 as the superset closed neighborhood frames. Prime examples of $\mathcal{F}_{\check{\mathcal{P}}}$ -coalgebras are the topological spaces (that were also mentioned in Chapter 1, be it rather implicitly under the name of *topological semantics*). To see this, represent the topology σ on the set S by the function mapping a point $s \in S$ to the collection

$\{U \in \sigma \mid s \in U\}$ of its neighborhoods.

EXAMPLE 1.145. For each set functor Ω , the empty set \emptyset , with the unique map from \emptyset to $\Omega\emptyset$, provides an Ω -coalgebra.

The functors mentioned in the Examples 1.141, 1.142 and 1.143, are examples of so-called *Kripke polynomial functors* which share some pleasant properties as we will see further on.

DEFINITION 1.146. The collection of *polynomial functors* is inductively defined as follows:

$$K ::= \mathcal{I} \mid C \mid K_0 + K_1 \mid K_0 \times K_1 \mid K^D. \quad (1.40)$$

Here \mathcal{I} denotes the identity functor on the category **Set**; C the constant functor $X \mapsto C$; $K_0 + K_1$ the coproduct functor $X \mapsto K_0(X) + K_1(X)$; $K_0 \times K_1$ the product functor; and K^D denotes the exponent functor $X \mapsto K(X)^D$.

Similarly, the collection of *Kripke polynomial functors* is given by

$$K ::= \mathcal{I} \mid C \mid K_0 + K_1 \mid K_0 \times K_1 \mid K^D \mid \mathcal{P}K, \quad (1.41)$$

where $\mathcal{P}K$ is the composition of K with the power set functor \mathcal{P} . Replacing \mathcal{P} with the *finite power set functor* \mathcal{P}_ω , and demanding the exponent D in K^D to be finite, we obtain the collection of *finitary Kripke polynomial functors*.

In each of these cases, the set $\text{Ing}K$ of *ingredient functors* of a (Kripke) polynomial functor K is defined by an obvious induction, with clauses $\text{Ing}(\mathcal{I}) := \{\mathcal{I}\}$, $\text{Ing}(\mathcal{P}K) := \{\mathcal{P}K\} \cup \text{Ing}(K)$, etc.

With the notation of this definition, Example 1.141 provides examples of coalgebras for the functors C and $\mathcal{I} \times C$. Deterministic automata over the alphabet C are $2 \times \mathcal{I}^C$ -coalgebras. Kripke frames are $\mathcal{P}\mathcal{I}$ -coalgebras, and Kripke models are coalgebras for the functor $\mathcal{P}\text{Prop} \times \mathcal{P}\mathcal{I}$. (Note that in the format (1.41), the power set functor as such is not a Kripke polynomial functor: It has to be represented as the functor $\mathcal{P}\mathcal{I}$. In the sequel, we will keep working with Kripke frames as \mathcal{P} -coalgebras, unless explicitly mentioned otherwise.)

After **Set**, the base category for coalgebras that carries most interest to modal logicians, is probably that of *Stone spaces*.

EXAMPLE 1.147. Recall from Remark 1.60 that a Stone space is pair $\mathbb{S} = (S, \sigma)$ such that σ is a compact Hausdorff space with a basis of clopens. Let **Stone** denote the category with Stone spaces as objects, and continuous maps as arrows. We will show that *descriptive general frames* can be viewed as **Stone**-coalgebras for the so-called *Vietoris functor* \mathcal{V} — for details on this observation, which is due to ABRAMSKY [1], see KUPKE, KURZ & VENEMA [74].

This functor, which forms the topological counterpart of the power set functor, is defined as follows. Given a topological space $\mathbb{S} = \langle S, \sigma \rangle$, let $K(\mathbb{S})$ denote the collection of closed subsets of S , and let $\ni \subseteq K(\mathbb{S}) \times S$ denote the converse membership relation. Then (in accordance with our earlier notation), we define, for any subset $U \subseteq S$, the sets $\langle \ni \rangle U = \{F \in K(\mathbb{S}) \mid F \cap U \neq \emptyset\}$ and $[\ni]U = \{F \in K(\mathbb{S}) \mid F \subseteq U\}$. The topology on $K(\mathbb{S})$, generated by taking the collection $\{\langle \ni \rangle U, [\ni]U \mid U \in \sigma\}$ as a subbasis, is called the *Vietoris topology* of σ , and the resulting space, the *Vietoris space* $\mathcal{V}(\mathbb{S})$ associated with \mathbb{S} .

The Vietoris construction preserves several properties of topological spaces; in particular, if \mathbb{S} is a Stone space, then so is $\mathcal{V}(\mathbb{S})$. Also, we may extend it to a functor, by defining, for a continuous map $f : \mathbb{S} \rightarrow \mathbb{S}'$, the function $\mathcal{V}f$ as the image map given by $(\mathcal{V}f)(X) := f[X]$. Here we omit the proof that $\mathcal{V}f$ is indeed an arrow in the category **Stone**, i.e., that it is a *continuous* map from $\mathcal{V}(\mathbb{S})$ to $\mathcal{V}(\mathbb{S}')$.

Now let $\mathbb{G} = \langle G, R, A \rangle$ be a descriptive general frame (cf. Definition 1.64), with associated Stone space σ_A . Recall from Remark 1.65 that the map $R[\cdot]$ mapping a point in G to the collection of its successors, is a function from G to $K(\langle G, \sigma_A \rangle)$. It is not too hard to prove that this is in fact a *continuous* map from $\langle G, \sigma_A \rangle$ to its Vietoris space. Thus we may represent \mathbb{G} as the Stone coalgebra $\langle \langle G, \sigma_A \rangle, R[\cdot] \rangle$.

Obviously, coalgebras are not studied in isolation; the following definition provides a natural notion of a map between coalgebras that preserves the transition structure.

DEFINITION 1.148. Let $\mathbb{A} = \langle A, \alpha \rangle$ and $\mathbb{A}' = \langle A', \alpha' \rangle$ be two coalgebras for the functor $\Omega : \mathbf{C} \rightarrow \mathbf{C}$. Then a *homomorphism* from \mathbb{A} to \mathbb{A}' is an arrow $f : A \rightarrow A'$ for which the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{f} & A' \\ \alpha \downarrow & & \downarrow \alpha' \\ \Omega A & \xrightarrow{\Omega f} & \Omega A' \end{array}$$

EXAMPLE 1.149. The homomorphisms for \mathcal{P} -coalgebras coincide with the bounded morphisms between Kripke frames. To see this, let $\mathbb{S} = \langle S, R \rangle$ and $\mathbb{S}' = \langle S', R' \rangle$ be two frames (for the basic modal similarity type), and consider their respective coalgebraic representations $\langle S, \sigma \rangle$ and $\langle S', \sigma' \rangle$, as in Example 1.143.

Now consider a map $f : S \rightarrow S'$. It is straightforward to show that

$$\begin{aligned} f \text{ satisfies the forth condition} & \quad \text{iff} \quad (\mathcal{P}f) \circ \sigma(s) \subseteq \sigma' \circ f(s) \text{ for all } s \in S, \\ f \text{ satisfies the back condition} & \quad \text{iff} \quad (\mathcal{P}f) \circ \sigma(s) \supseteq \sigma' \circ f(s) \text{ for all } s \in S. \end{aligned}$$

This shows that f is a bounded morphism from \mathbb{S} to \mathbb{S}' if and only if it is a coalgebra homomorphism from $\langle S, \sigma \rangle$ to $\langle S', \sigma' \rangle$, and provides perhaps the most convincing argument that the notion of a bounded morphism is a natural one.

EXAMPLE 1.150. Let \mathbb{X} and \mathbb{X}' be two topological spaces, represented as coalgebras $\mathbb{X} = \langle X, \xi \rangle$ and $\mathbb{X}' = \langle X', \xi' \rangle$ for the filter functor $\mathcal{F}_{\mathcal{P}}$ of Example 1.144. We leave it for the reader to check that a map $f : S \rightarrow S'$ is an $\mathcal{F}_{\mathcal{P}}$ -coalgebra homomorphism iff f is continuous and open (i.e., not only do we require $f^{-1}[U']$ to be open in \mathbb{X} if U' is open in \mathbb{X}' , but also $f[U]$ must be open in \mathbb{X}' for all \mathbb{X} -open U).

Likewise, one can prove that the coalgebraic notion of a homomorphism between monotone neighborhood frames, represented as coalgebras for the functor $\mathcal{U}_{\mathcal{P}}$, corresponds to that of a bounded morphism for these structures as defined in section 2.

It is easy to check that the collection of coalgebra homomorphisms contains all identity arrows and is closed under arrow composition. Hence, the Ω -coalgebras with their homomorphisms form a category.

DEFINITION 1.151. For any functor $\Omega : \mathbf{C} \rightarrow \mathbf{C}$, we let $\text{Coalg}(\Omega)$ denote the category with Ω -coalgebras as objects and the corresponding homomorphisms as arrows. The

category \mathbf{C} is called the *base category* of $\mathbf{Coalg}(\Omega)$.

The reader will already be familiar with a number of (isomorphic copies of) these categories. For instance, Example 1.149 shows in fact that the category \mathbf{Fr} (of frames with bounded morphisms) is *isomorphic* to the category $\mathbf{Coalg}(\mathcal{P})$ of \mathcal{P} -coalgebras. Likewise, elaborating Example 1.147, one can prove that the category \mathbf{DGF} (of descriptive general frames with continuous bounded morphisms, see Definition 1.66) is isomorphic to the category of Stone coalgebras for the Vietoris functor. Of course, it is these isomorphisms that justify our classification of modal structures as coalgebras, not so much the simple fact that the objects in isolation can be presented in coalgebraic format.

REMARK 1.152. Recall that an *algebra* over a signature Ω is a set A with an Ω -indexed collection $\{f^A \mid A^{ar(f)} \rightarrow A\}$ of operations. These operations may be combined into a single map $\alpha : \sum_{f \in \Omega} A^{ar(f)} \rightarrow A$, where $\sum_{f \in \Omega} A^{ar(f)}$ denotes the *coproduct* (or *sum*, or *disjoint union*) of the sets $\{A^{ar(f)} \mid f \in \Omega\}$. It is not hard to verify that a map $g : A \rightarrow A'$ is an algebraic homomorphism between the algebras $\mathbb{A} = \langle A, \alpha \rangle$ and $\mathbb{A}' = \langle A', \alpha' \rangle$ iff the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{f} & A' \\ \alpha \uparrow & & \uparrow \alpha' \\ \Omega A & \xrightarrow{\Omega f} & \Omega A' \end{array}$$

where we now view the signature Ω as the polynomial set *functor* $\sum_{f \in \Omega} \mathcal{I}^{ar(f)}$. That is, Ω operates as well on functions between sets. This naturally suggests the following generalization.

Given an endofunctor Ω on a category \mathbf{C} , an Ω -algebra is a pair $\mathbb{A} = \langle A, \alpha \rangle$ where $\alpha : \Omega A \rightarrow A$ is an arrow in \mathbf{C} . A homomorphism from an Ω -algebra \mathbb{A} to an Ω -algebra \mathbb{A}' is an arrow $f : A \rightarrow A'$ such that $f \circ \alpha = \alpha' \circ (\Omega f)$. The induced category is denoted as $\mathbf{Alg}(\Omega)$.

Now the obvious similarities between the notions of algebra and coalgebra can be made very precise. The basic observation, which also explains the name ‘coalgebra’, is that a coalgebra $\mathbb{C} = \langle C, \gamma : C \rightarrow \Omega C \rangle$ over a base category \mathbf{C} can also be seen as an algebra in the *opposite* category \mathbf{C}^{op} — we will come back to this issue in section 15. Note however, that universal coalgebra, dealing with arbitrary set functors, is more general than (what is usually called) universal algebra, which involves only polynomial functors.

10 FINAL COALGEBRAS

DEFINITION 1.153. A functor $\Omega : \mathbf{C} \rightarrow \mathbf{C}$ is said to *admit a final* or *terminal coalgebra* if the category $\mathbf{Coalg}(\Omega)$ has a final object, that is, a coalgebra \mathbb{Z} such that from every coalgebra \mathbb{A} in $\mathbf{Coalg}(\Omega)$ there is a unique homomorphism $!_{\mathbb{A}} : \mathbb{A} \rightarrow \mathbb{Z}$.

Functors admitting a final coalgebra are of special interest. In the case of state-based coalgebras, one reason for this is that final coalgebras often provide an intuitive encoding of the notion of *behavior*. And in fact, many interesting and well-known mathematical objects can be naturally associated with the final coalgebra of some functor.

EXAMPLE 1.154. Consider a black box machine $\mathbb{M} = \langle M, \mu \rangle$ as in Example 1.141.

Starting from, say, state x_0 , the machine makes a transition $\mu(x_0) = (c_0, x_1)$ and continues with $\mu(x_1) = (c_1, x_2)$, $\mu(x_2) = (c_2, x_3)$, etc. Since the states x_0, x_1, \dots are internal to the machine, the only observable part of this dynamics is the infinite sequence or *stream* $\text{beh}(x_0) = (c_0, c_1, c_2, \dots) \in C^\omega$ of values in the data set C .

The collection C^ω of all infinite words over C forms itself a system for the functor $C \times \mathcal{I}$. Simply endow the set C^ω with the transition structure γ splitting an infinite stream $u = c_0c_1c_2\dots$ into its *head* $h(u) = c_0$ and its *tail* $t(u) = c_1c_2c_3\dots$. Putting $\gamma(u) = (h(u), t(u))$, one easily proves that the behavior map $x \mapsto \text{beh}(x)$ is the unique homomorphism from \mathbb{M} to this coalgebra $\langle C^\omega, \gamma \rangle$. This shows that $\langle C^\omega, \gamma \rangle$ is the final object in the category $\text{Coalg}(C \times \mathcal{I})$.

EXAMPLE 1.155. For a second example, consider again the coalgebraic representation of a deterministic automaton over the alphabet C as a $2 \times \mathcal{I}^C$ -coalgebra. Now we will see that the collection $\mathcal{P}(C^*)$ of all *languages* over C provides (the carrier of) the final coalgebra. We can turn this set $\mathcal{P}(C^*)$ into a coalgebra by imposing on it the following transition function $\lambda : \mathcal{P}(C^*) \rightarrow 2 \times \mathcal{P}(C^*)^C$. Writing $\lambda(L) = (\lambda_0(L), \lambda_1(L))$, we define $\lambda_0(L) := 1$ iff the empty string belongs to L , and $\lambda_1(L)(c) := \{w \in C^* \mid cw \in L\}$. (The latter set is sometimes called the *c-derivative* of L .)

We leave it for the reader to verify that with this definition, the structure $\langle \mathcal{P}(C^*), \lambda \rangle$ forms the final object in $\text{Coalg}(2 \times \mathcal{I}^C)$. Given a $2 \times \mathcal{I}^C$ -coalgebra \mathbb{A} , the unique homomorphism $!_{\mathbb{A}} : \mathbb{A} \rightarrow \langle \mathcal{P}(C^*), \lambda \rangle$ maps a state $a \in A$ to the language that is accepted by the automaton that we obtain by taking a as initial state of \mathbb{A} .

EXAMPLE 1.156. An interesting example in modal logic is provided by the final coalgebra for the Vietoris functor \mathcal{V} of Example 1.147. The existence of a final \mathcal{V} -coalgebra is in fact an immediate consequence of the isomorphism $\text{Coalg}(\mathcal{V}) \cong \text{DGF}$, and the duality between DGF and MA (the category of modal algebras with homomorphisms). MA has an initial object (namely, the Lindenbaum-Tarski algebra generated by the empty set of variables, or, equivalently, the free modal algebra over zero generators), and so by duality, $\text{Coalg}(\mathcal{V})$ must have a final object. In fact, the *canonical* descriptive general frame, based on the set of maximal consistent closed formulas, fulfills this role — a nice and perhaps quite unexpected application of this construction.

An important application of final coalgebras is provided by the principle of *coinduction*, which is one of the fundamental coalgebraic notions. There are two sides to this principle: it serves both as an important proof tool and as an elegant means of providing definitions. As a definition principle, coinduction is based on the existence of unique homomorphisms into the final Ω -system $\mathbb{Z} = \langle Z, \zeta \rangle$. For, suppose that we can endow a set S with an Ω -coalgebra map $\sigma : S \rightarrow \Omega S$, thus obtaining the Ω -system \mathbb{S} . Then there is a unique function $f_\sigma = !_{\mathbb{S}} : S \rightarrow Z$ which is consistent with the coalgebra specification σ , in the sense that it is a coalgebraic homomorphism from $\langle S, \sigma \rangle$ to \mathbb{Z} . Thus the function f_σ is *defined* by coinduction from (the specification) σ .

EXAMPLE 1.157. For instance, take the function that merges two streams by taking elements from either stream in turn. For a coinductive definition of this map, define the transition map $\text{zip} : C^\omega \times C^\omega \rightarrow C \times (C^\omega \times C^\omega)$ as follows:

$$\text{zip}(u, v) := (h(u), (v, t(u))),$$

where h and t are the head and tail maps of Example 1.154. Then by finality there is a unique homomorphism $f_{\text{zip}} : C^\omega \times C^\omega \rightarrow C^\omega$. One may verify that this indeed defines

the map that zips two streams together.

The previous example is fairly typical in that it uses coinduction to define a function from a product of the final system to itself. It should also be noted that coinduction works particularly well for structures that combine algebraic and coalgebraic features, such as streams of data objects which are subject themselves to algebraic operations.

Unfortunately, final coalgebras do not exist for every functor Ω . For instance, **Set**-endofunctors involving the power set functor in a nontrivial way, will generally not admit a final coalgebra; in particular, there is no final Kripke frame or model. By Cantor's theorem, these results are immediate consequence of the following proposition, which is due to Lambek [79].

PROPOSITION 1.158. *Let $\Omega : \mathbf{C} \rightarrow \mathbf{C}$ be some functor admitting a final system $\mathbb{Z} = \langle Z, \zeta \rangle$. Then ζ is an isomorphism (in \mathbf{C}) between Z and ΩZ .*

Proof. Suppose that $\mathbb{Z} = \langle Z, \zeta \rangle$ is the final object of $\mathbf{Coalg}(\Omega)$. It can easily be verified that ζ is in fact a coalgebra homomorphism from \mathbb{Z} to $\mathbb{Z}_2 := \langle \Omega Z, \Omega \zeta \rangle$. But then the composition $!_{\mathbb{Z}_2} \circ \zeta$ is a coalgebra homomorphism from \mathbb{Z} to itself, just like the identity arrow id_Z on Z . Thus by uniqueness it follows that $!_{\mathbb{Z}_2} \circ \zeta = id_Z$. For the reverse composition $\zeta \circ !_{\mathbb{Z}_2}$ we have, by the fact that $!_{\mathbb{Z}_2}$ is a homomorphism, that $\zeta \circ !_{\mathbb{Z}_2} = \Omega !_{\mathbb{Z}_2} \circ \Omega \zeta = \Omega(!_{\mathbb{Z}_2} \circ \zeta) = \Omega(id_Z) = id_{\Omega Z}$. From this the result is immediate. ■

So which functors admit final coalgebras? Some good sufficient conditions are known.

DEFINITION 1.159. Let Ω be some set functor, and κ some cardinal. Call Ω κ -small if

$$\Omega(S) = \bigcup \{ (\Omega \iota)[\Omega(A)] \mid \iota : A \hookrightarrow S, |A| < \kappa \},$$

for all sets $S \neq \emptyset$. Ω is *small* if it is small for some cardinal κ .

In words, the definition requires every element of $\Omega(S)$ to be in the range of $\Omega \iota$ for an appropriate inclusion map $\iota : A \hookrightarrow S$. In case Ω is a *standard* functor (meaning that Ω maps inclusions $\iota : A \hookrightarrow B$ to inclusions $(\Omega \iota) : \Omega A \hookrightarrow \Omega B$), the definition boils down to the requirement that $\Omega(S) = \bigcup \{ \Omega(A) \mid A \subseteq S, |A| < \kappa \}$. The notion of smallness is easily seen to be equivalent to the instantiation in **Set** of the more general notion of *accessibility*, and it is also equivalent to the concept of *boundedness*, cf. ADÁMEK & PORST [6] for details.

Examples of small functors abound; for instance, whenever we replace, in a Kripke polynomial functor, the power set functor by a bounded variant such as the finite power set functor, the result is a small functor. For instance, the finite power set functor \mathcal{P}_ω is ω -small. The following result, due to ACZEL & MENDLER [3] and BARR [9], witnesses the importance of the notion.

FACT 1.160. Every small set functor admits a final coalgebra.

As one of the immediate corollaries of this fact, the categories of image finite frames and image finite models, which can be represented as coalgebras for the functor \mathcal{P}_ω , and $\mathcal{P}\mathbf{Prop} \times \mathcal{P}_\omega$, respectively, have final objects.

REMARK 1.161. For **Set**-based functors that do not admit a final coalgebra, one may *create* a final coalgebra — at least, if one is willing to allow coalgebras with a class rather than a set as their carrier. Let **SET** be the category that has classes as objects,

and set-continuous functions as arrows. These are functions $f : C \rightarrow C'$ between classes with the property that $f(C) = \bigcup\{f(S) \mid S \subseteq C \text{ and } S \text{ is a set}\}$. An endofunctor on **SET** is *set-based* if for each class C and each $c \in \Omega(C)$ there is a set $S \subseteq C$ such that $c \in (\Omega\iota)[\Omega(S)]$, where $\iota : S \rightarrow C$ is the inclusion map. (If the set functor is standard, this boils down to requiring that Ω is a set-continuous map on objects.) Now ACZEL & MENDLER [3] proved that every set-based endofunctor $\Xi : \mathbf{SET} \rightarrow \mathbf{SET}$ admits a final coalgebra. The similarity to Fact 1.160 is no coincidence: BARR [9] showed that the result of Aczel & Mendler can in fact be reformulated as Fact 1.160.

This fact can be used as follows. Given an endofunctor Ω on **Set**, there is a *unique* way to extend Ω to a set based endofunctor Ω^+ on **SET**. (On objects, simply put $\Omega^+(C) := \bigcup\{(\Omega\iota)[\Omega(S) \mid \iota : S \hookrightarrow C, S \text{ a set}]\}$.) The theorem of Aczel & Mendler then guarantees the existence of a final object Z in $\mathbf{Coalg}(\Omega^+)$. This coalgebra will be class-based if Ω does not admit a final coalgebra, but it will be final, not only with respect to the set-based coalgebras in $\mathbf{Coalg}(\Omega^+)$, but also with respect to the class-based ones. As an important instance of this idea, ACZEL [2] showed that the class of non-well-founded sets provides the final coalgebra for (the **SET**-based extension of) the power set functor.

REMARK 1.162. Whether the functor admits a final coalgebra or not, one may always (try to) approximate it. The *final* or *terminal sequence* associated with a given set functor Ω , is an ordinal indexed sequence of objects $\langle Z_\alpha \rangle$ with maps $p_\beta^\alpha : Z_\alpha \rightarrow Z_\beta$ for $\beta \leq \alpha$, such that (i) $Z_{\alpha+1} = \Omega Z_\alpha$ and $p_{\beta+1}^{\alpha+1} = \Omega p_\beta^\alpha$, (ii) $p_\alpha^\alpha = id_{Z_\alpha}$ and $p_\gamma^\beta \circ p_\beta^\alpha = p_\gamma^\alpha$, (iii) if λ is a limit ordinal, then Z_λ with $\{p_\alpha^\lambda \mid \alpha < \lambda\}$ is a limit of the diagram with objects $\{Z_\alpha \mid \alpha < \lambda\}$ and arrows $\{p_\beta^\alpha \mid \alpha, \beta < \lambda\}$. (In particular, taking 0 to be a limit ordinal, we find that $Z_0 = 1$ is some initial object 1 of the category **Set**.) It is not hard to prove that, modulo isomorphism, the final sequence is uniquely determined by these conditions.

Intuitively, it can be seen as an approximation of the final coalgebra for Ω . That is, where elements of the final coalgebra represent ‘complete’ behavior, elements of Z_α represent behavior that can be performed in α steps. To make this precise and formal, observe that for any Ω -coalgebra \mathbb{S} there is a *unique* ordinal-indexed class of functions $!_\alpha : S \rightarrow Z_\alpha$ such that $!_0$ is fixed by the finality of Z_0 in **Set**, $!_{\alpha+1} = (\Omega!_\alpha) \circ \sigma$, and for limit λ , $!_\lambda$ is given as the unique map $!_\lambda : S \rightarrow Z_\lambda$ such that $!_\lambda = p_\alpha^\lambda \circ !_\alpha$ for all $\alpha < \lambda$. It is not hard to prove that, for instance, $\mathbb{S}, s \equiv_\Omega \mathbb{S}', s'$ implies that $!_\alpha(s) = !_\alpha(s')$ for all α .

The relation with final coalgebras can be made precise, as follows. On the one hand, if the final sequence converges, in the sense that some arrow $p_\alpha^{\alpha+1}$ is a bijection, then the coalgebra $\langle Z_\alpha, (p_\alpha^{\alpha+1})^{-1} \rangle$ is a final coalgebra for Ω . And conversely, under some constraints on Ω , ADÁMEK & KOUBEK [5] proved that if Ω admits a final coalgebra, then the final sequence converges to it. More information on the final sequence of set functors can be found in WORRELL [109].

11 BISIMULATION & BEHAVIORAL EQUIVALENCE

In this section we discuss the most important notions of equivalence between systems: behavioral equivalence and bisimulation. Both of these generalize the concept of a bisimulation between two Kripke models.

Probably the most intuitive notion of equivalence between systems is that of *behavioral*, or *observational*, equivalence. The idea here is to consider two states to be similar if we

cannot distinguish them by observations, because they display the same behavior. For instance, we call two deterministic automata (pointed $2 \times \mathcal{I}^C$ -coalgebras) equivalent if they recognize the same language. In case the functor Ω admits a final coalgebra \mathbb{Z} , this idea is easily formalized by making state s_0 in coalgebra \mathbb{S}_0 equivalent to state s_1 in coalgebra \mathbb{S}_1 if $!_{\mathbb{S}_0}(s_0) = !_{\mathbb{S}_1}(s_1)$. In case the functor does not admit a final coalgebra, we generalize this demand as follows.

DEFINITION 1.163. Let $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{S}' = \langle S', \sigma' \rangle$ be two systems for the set functor Ω . Then $s \in S$ and $s' \in S'$ are *behaviorally equivalent*, notation: $\mathbb{S}, s \equiv_{\Omega} \mathbb{S}', s'$ if there is an Ω -system $\mathbb{X} = \langle X, \xi \rangle$ and homomorphisms $f : \mathbb{S} \rightarrow \mathbb{X}$ and $f' : \mathbb{S}' \rightarrow \mathbb{X}$ such that $f(s) = f'(s')$.

REMARK 1.164. It is easily checked that in case Ω admits a final coalgebra, then indeed $\mathbb{S}, s \equiv_{\Omega} \mathbb{S}', s'$ iff $!_{\mathbb{S}}(s) = !_{\mathbb{S}'}(s')$. In the case that Ω does not admit a final coalgebra, then one may show that behavioral equivalence is captured in the same way by the final coalgebra of the extension Ω^+ of Ω to the category SET, see Remark 1.161.

REMARK 1.165. As a variation of behavioral equivalence, the final sequence can be used to study behavior, in a way that is not unlike modal logic. For instance, call two pointed Ω -systems (\mathbb{S}, s) and (\mathbb{S}', s') α -*equivalent* if $!_{\alpha}(s) = !_{\alpha}(s')$. In the case of Kripke models, this notion coincides with that of bounded bisimilarity, see Chapter 5 of this volume. One may prove that behavioral equivalence itself coincide with the intersection of α -equivalence for all ordinals α .

In almost all cases of interest, behavioral equivalence can be characterized via the equally fundamental concept of *bisimilarity*, which is due to ACZEL & MENDLER [3]. The definition of bisimilarity and bisimulations may not be so intuitive at first sight, but, as we will see, these notions have some rather elegant mathematical properties.

DEFINITION 1.166. Let $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{S}' = \langle S', \sigma' \rangle$ be two systems for the set functor Ω . A relation $B \subseteq S \times S'$ is called a *bisimulation between \mathbb{S} and \mathbb{S}'* , if we can endow it with a coalgebra map $\beta : B \rightarrow \Omega B$, in such a way that the two projections $\pi : B \rightarrow S$ and $\pi' : B \rightarrow S'$ are homomorphisms from $\langle B, \beta \rangle$ to \mathbb{S} and \mathbb{S}' , respectively:

$$\begin{array}{ccccc} S & \xleftarrow{\pi} & B & \xrightarrow{\pi'} & S' \\ \sigma \downarrow & & \beta \downarrow & & \sigma' \downarrow \\ \Omega S & \xleftarrow{\Omega\pi} & \Omega B & \xrightarrow{\Omega\pi'} & \Omega S' \end{array}$$

If there exists a bisimulation B with $(s, s') \in B$, we say that s and s' are *bisimilar*, notation: $\mathbb{S}, s \leftrightarrow \mathbb{S}', s'$ (or $B : \mathbb{S}, s \leftrightarrow \mathbb{S}', s'$ in case we want to make the bisimulation B explicit).

Finally, if $\mathbb{S} = \mathbb{S}'$ we say that B is a bisimulation *on \mathbb{S}* ; if this B happens to be an equivalence relation, we call it a *bisimulation equivalence* on \mathbb{S} .

REMARK 1.167. Intuitively, bisimulation equivalences correspond to *congruences* in universal algebra. To make this analogy somewhat more precise, call a relation $R \subseteq A_0 \times A_1$, linking the carrier sets of two Ω -algebras \mathbb{A}_0 and \mathbb{A}_1 , *substitutive* if there exists an *algebraic structure* $\rho : \Omega R \rightarrow R$, such that the two projections $\pi_i : R \rightarrow A_i$ are (algebraic) homomorphisms. This is clearly an algebraic *analogue* (rather than a dual version) of a bisimulation, so that the correspondence between congruences and bisimulation equivalences obtains through the observation that a congruence is nothing

but a substitutive equivalence relation.

EXAMPLE 1.168. Let $\mathbb{S}_0 = \langle S_0, \sigma_0 \rangle$ and $\mathbb{S}_1 = \langle S_1, \sigma_1 \rangle$ be two coalgebras over the functor $\mathcal{P}\text{Prop} \times \mathcal{P}$. That is, \mathbb{S}_0 and \mathbb{S}_1 are Kripke models in coalgebraic shape; write $\sigma_i(s) = (\lambda_i(s), R_i[s])$, where $\lambda_i(s)$ is the collection of proposition letters true at s in \mathbb{S}_i , and $R_i[s]$ is the successor set of s in \mathbb{S}_i , as in the examples 1.143 and 1.149. Now consider an arbitrary relation $B \subseteq S_0 \times S_1$. It is a very instructive exercise to check that B is a bisimulation in the coalgebraic sense if and only if it is a bisimulation in the sense of Kripke models. Recall that the latter property means that for any pair $(s_0, s_1) \in B$:

(*atom*) $p \in \lambda_0(s)$ iff $p \in \lambda_1(s)$, for all $p \in \text{Prop}$;

(*forth*) for all $t_0 \in R_0[s_0]$ there is some $t_1 \in R_1[s_1]$ with $(t_0, t_1) \in B$;

(*back*) for all $t_1 \in R_1[s_1]$ there is some $t_0 \in R_0[s_0]$ with $(t_1, t_0) \in B$.

One way to prove this equivalence uses the fact that bounded morphisms coincide with coalgebra morphisms, cf. Example 1.149. Details are left to the reader.

EXAMPLE 1.169. Recall from Example 1.142 that deterministic automata over an alphabet C can be represented as $2 \times \mathcal{I}^C$ -coalgebras. Now let $\mathbb{A} = (A, o, \nu)$ and $\mathbb{A}' = (A', o', \nu')$ be two such automata. We leave it for the reader to verify that $B \subseteq A \times A'$ is a bisimulation between \mathbb{A} and \mathbb{A}' iff every pair $(s, s') \in B$ satisfies (i) $o(s) = o'(s')$ and (ii) $(\nu(s)(c), \nu'(s')(c)) \in B$ for every $c \in C$. In this case it is easy to see that bisimilar states are also behaviorally equivalent.

EXAMPLE 1.170. For an arbitrary set functor Ω , it is easy to see that for any coalgebra \mathbb{S} , the diagonal relation $\Delta_{\mathbb{S}}$ is a bisimulation equivalence on \mathbb{S} . Furthermore, the converse of a bisimulation is again a bisimulation. However, the collection of bisimulations is not in general closed under taking relational composition.

Finally, homomorphisms can be seen as functional bisimulations. To be more precise, let $f : S_0 \rightarrow S_1$ be a function between the carriers of two Ω -coalgebras \mathbb{S}_0 and \mathbb{S}_1 . Recall that the *graph* of f is the relation $G_f := \{(s, f(s)) \mid s \in S_0\}$. Then it holds that

$$f \text{ is a coalgebraic homomorphism iff its graph is a bisimulation.} \quad (1.42)$$

In order to see why this is so, first suppose that $G_f : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1$. Since the projection map $\pi_0 : G_f \rightarrow S_0$ is a bijective homomorphism, its inverse π_0^{-1} is also a homomorphism. But then $f = \pi_1 \circ \pi_0^{-1}$, as the composition of two homomorphisms, is also a homomorphism. For the other direction, suppose that f is a homomorphism; then it is straightforward to verify that the map $(\Omega\pi_0)^{-1} \circ \sigma \circ \pi_0$ equips the set G_f with the required coalgebraic structure.

Bisimulations admit an elegant alternative characterization which involves the notion of *relation lifting*. As an example, consider the power set functor \mathcal{P} . Recall that $B \subseteq S_0 \times S_1$ is a bisimulation between $\mathbb{S}_0 = \langle S_0, R_0[\cdot] \rangle$ and $\mathbb{S}_1 = \langle S_1, R_1[\cdot] \rangle$ iff B satisfies the conditions (*back*) and (*forth*) of Example 1.168. Now suppose that we define, for an arbitrary relation $R \subseteq S_0 \times S_1$, the relation $\overline{\mathcal{P}}(R) \subseteq \mathcal{P}(S_0) \times \mathcal{P}(S_1)$ by putting

$$\overline{\mathcal{P}}(R) := \{(Q_0, Q_1) \mid \forall q_0 \in Q_0 \exists q_1 \in Q_1. (q_0, q_1) \in R \text{ and } \forall q_1 \in Q_1 \exists q_0 \in Q_0. (q_0, q_1) \in R\}. \quad (1.43)$$

In other words, we *lift* the relation R to the level of the power sets of S_0 and S_1 . The definition of a bisimulation between \mathcal{P} -coalgebras can now be nicely characterized as follows:

$$B : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1 \text{ iff } (R_0[s_0], R_1[s_1]) \in \overline{\mathcal{P}}(B) \text{ for all } (s_0, s_1) \in B.$$

This nice way of characterizing bisimulation via relation lifting is not limited to the power set functor — it applies in fact to *every* set functor.

DEFINITION 1.171. Let \mathbb{S}_0 and \mathbb{S}_1 be two coalgebras for some set functor Ω . Given a relation $R \subseteq S_0 \times S_1$, consider the following diagram, where $\pi_i : R \rightarrow S_i$ and $p_i : \Omega S_0 \times \Omega S_1 \rightarrow \Omega S_i$ denote the projection maps.

$$\begin{array}{ccc} S_0 & \xleftarrow{\pi_0} R & \xrightarrow{\pi_1} S_1 \\ \sigma_0 \downarrow & & \downarrow \sigma_1 \\ \Omega S_0 & \xleftarrow{\Omega\pi_0} \Omega R & \xrightarrow{\Omega\pi_1} \Omega S_1 \\ & \swarrow p_0 \quad \rho_R \quad \searrow p_1 & \\ & \Omega S_0 \times \Omega S_1 & \end{array}$$

It follows from the category theoretic properties of the product $\Omega S_0 \times \Omega S_1$ that there is a unique map $\rho_R = \langle \Omega\pi_0, \Omega\pi_1 \rangle$ from ΩR to $\Omega S_0 \times \Omega S_1$ such that $p_i \circ \rho_R = \Omega\pi_i$ for $i = 0, 1$. We define the *relation lifting* of R as the relation

$$\overline{\Omega}R := \{((\Omega\pi_0)(u), (\Omega\pi_1)(u)) \mid u \in \Omega R\}, \quad (1.44)$$

that is, $\overline{\Omega}R$ is the image of ΩR under ρ_R .

The results listed in the following theorem, which summarize the most important properties of bisimulations, basically date back to ACZEL & MENDLER [3].

THEOREM 1.172. *Let \mathbb{S}_0 and \mathbb{S}_1 be two coalgebras for some set functor Ω .*

1. $B : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1$ iff $(\sigma_0(s_0), \sigma_1(s_1)) \in \overline{\Omega}(B)$ for all $(s_0, s_1) \in B$.
2. The collection of bisimulations between \mathbb{S}_0 and \mathbb{S}_1 forms a complete lattice under the inclusion order, with joins given by unions.
3. The bisimilarity relation \rightleftharpoons is the largest bisimulation between \mathbb{S}_0 and \mathbb{S}_1 .

Proof. The first part of the theorem is an almost immediate consequence of the definitions, so we leave the details to the reader.

The crucial observation in the proof of the other two parts is that

$$\overline{\Omega} : \mathcal{P}(S_0 \times S_1) \rightarrow \mathcal{P}(\Omega S_0 \times \Omega S_1) \text{ is a monotone operation.} \quad (1.45)$$

For a proof, let $R \subseteq R'$ be two relations between S_0 and S_1 , with $\iota : R \rightarrow R'$ denoting the inclusion map. By definition of $\overline{\Omega}$, we may without loss of generality represent an arbitrary element of $\overline{\Omega}(R)$ as a pair $\rho_R(u) = ((\Omega\pi_0)(u), (\Omega\pi_1)(u))$ for some $u \in \Omega R$. Define $u' := (\Omega\iota)(u)$, then u' belongs to $\Omega R'$, and for each i we find that $(\Omega\pi'_i)(u') = (\Omega\pi'_i \circ \Omega\iota)(u) = (\Omega(\pi'_i \circ \iota))(u) = (\Omega\pi_i)(u)$. That is, $\rho_R(u) = \rho_{R'}(u')$, which shows that $\rho_R(u)$ belongs to $\overline{\Omega}(R')$. This proves (1.45).

Now for the proof of part 2, recall that a partial order is a complete lattice if it closed under arbitrary joins. Hence, it suffices to prove that the union B of a collection $\{B_j \mid j \in J\}$ of bisimulations is again a bisimulation. Take an arbitrary pair $(s_0, s_1) \in B$. Then (s_0, s_1) belongs to B_j for some $j \in J$. Hence, by part 1, we find (s_0, s_1) in $\overline{\Omega}(B_j)$, so $(s_0, s_1) \in \overline{\Omega}(B)$ by the monotonicity of $\overline{\Omega}$. But then B is a bisimulation by part 1.

Finally, for part 3, note that it is an immediate consequence of part 2 that \Leftrightarrow , being the union of all bisimulations between \mathbb{S}_0 and \mathbb{S}_1 , is a bisimulation itself. Hence, by definition, it is the greatest bisimulation between \mathbb{S}_0 and \mathbb{S}_1 . In fact, it follows by the Knaster-Tarski theorem (on fixed points of monotone operations on complete lattices), that \Leftrightarrow is in fact the greatest fixed point of the map $\Lambda : R \mapsto \{(s_0, s_1) \mid (\sigma_0(s_0), \sigma_1(s_1)) \in \overline{\Omega}(R)\}$. ■

In the case of Kripke polynomial functors, relation lifting can be characterized using *induction* on the construction of the functor, cf. JACOBS [63].

PROPOSITION 1.173. *Let S and S' be two sets, and $R \subseteq S \times S'$ a binary relation between S and S' . Then the following induction defines the relation lifting $\overline{K}(R) \subseteq KS \times KS'$, for each Kripke polynomial functor K :*

$$\begin{aligned} \overline{I}(R) &:= R, \\ \overline{C}(R) &:= \Delta_C, \\ \overline{K_0 \times K_1}(R) &:= \{(x_0, x_1), (x'_0, x'_1) \mid (x_0, x'_0) \in \overline{K_0}(R) \text{ and } (x_1, x'_1) \in \overline{K_1}(R)\}, \\ \overline{K_0 + K_1}(R) &:= \{(\kappa_0 x_0, \kappa_0 x'_0) \mid (x_0, x'_0) \in \overline{K_0}(R)\} \cup \{(\kappa_1 x_1, \kappa_1 x'_1) \mid (x_1, x'_1) \in \overline{K_1}(R)\}, \\ \overline{K^D}(R) &:= \{(f, f') \mid (f(d), f'(d)) \in \overline{K}(R) \text{ for all } d \in D\}, \\ \overline{PK}(R) &:= \{(Q, Q') \mid \forall q \in Q \exists q' \in Q'. (q, q') \in \overline{K}(R) \text{ and } \forall q' \in Q' \exists q \in Q. (q, q') \in \overline{K}(R)\}. \end{aligned}$$

Now that we have defined these two notions of equivalence between coalgebras, the obvious question is how they relate to each other. One direction is clear: bisimilarity is a sufficient condition for behavioral equivalence.

PROPOSITION 1.174. *Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ be some functor, and let s_0 and s_1 be states of the Ω -coalgebras \mathbb{S}_0 and \mathbb{S}_1 , respectively. Then $\mathbb{S}_0, s_0 \Leftrightarrow \mathbb{S}_1, s_1$ implies $\mathbb{S}_0, s_0 \equiv_{\Omega} \mathbb{S}_1, s_1$.*


Proof. The proof of this proposition is, in the general case, similar to the one of Theorem 1.177 below (with an application of *pushouts* instead of pullbacks), so we omit details.

In the special case that Ω admits a final coalgebra, a very simple proof obtains. Assume that $B : \mathbb{S}_0 \Leftrightarrow \mathbb{S}_1$, and let $\beta : B \rightarrow \Omega B$ be a coalgebra map witnessing this. It follows from the definitions that both $!_{\mathbb{S}_0} \circ \pi_0$ and $!_{\mathbb{S}_1} \circ \pi_1$ are coalgebraic homomorphisms from $\langle B, \beta \rangle$ to the final coalgebra, so from finality it follows that $!_{\mathbb{S}_0} \circ \pi_0 = !_{\mathbb{S}_1} \circ \pi_1$. From this it is immediate that $B \subseteq \equiv_{\Omega}$. Hence in particular, since \Leftrightarrow is itself a bisimulation, we see that $\Leftrightarrow \subseteq \equiv_{\Omega}$. ■

In general however, bisimilarity is a strictly stronger notion than behavioral equivalence. For instance, the two notions do not coincide in the case of monotone neighborhood frames (coalgebras for the functor $\mathcal{U}_{\mathcal{P}}$ of Example 1.144). The reader is referred to HANSEN & KUPKE [55] for details. Here we just mention that behavioral equivalence, which for monotone neighborhood frames is formulated exactly like the *topobisimilarity* defined in Chapter 1 of this volume, seems to be the more natural notion.

For a constraint on the functor that guarantees the two notions to coincide, consider the following.

DEFINITION 1.175. A *weak pullback* of two arrows $f_0 : A_0 \rightarrow B$, $f_1 : A_1 \rightarrow B$ in a category \mathbf{C} is a pair of arrows $p_0 : W \rightarrow A_0$, $p_1 : W \rightarrow A_1$ such that (i) $f_0 \circ p_0 = f_1 \circ p_1$, while (ii) for every pair $p'_0 : W' \rightarrow A_0$, $p'_1 : W' \rightarrow A_1$ that also satisfies $f_0 \circ p'_0 = f_1 \circ p'_1$, there is a mediating arrow $w' : W' \rightarrow W$ such that $p_0 \circ w' = p'_0$ and $p_1 \circ w' = p'_1$. A functor $\Omega : \mathbf{C} \rightarrow \mathbf{C}'$ *preserves weak pullbacks* if for any weak pullback (p_0, p_1) of any (f_0, f_1) in \mathbf{C} , the pair $(\Omega p_0, \Omega p_1)$ is a weak pullback of $(\Omega f_0, \Omega f_1)$ in \mathbf{C}' .



Note that the mediating arrow w' need not be unique: adding this requirement to the definition would give the more familiar, and stronger, notion of a *pullback*. The category **Set** has pullbacks: for $f_0 : A_0 \rightarrow B$ and $f_1 : A_1 \rightarrow B$, we can take the projections to A_0 and A_1 from the set $pb(f_0, f_1) := \{(a_0, a_1) \in A_0 \times A_1 \mid f_0(a_0) = f_1(a_1)\}$.

Many but not all endofunctors on **Set** in fact preserve weak pullbacks.

PROPOSITION 1.176. *All polynomial functors preserve pullbacks, and all Kripke polynomial functors preserve weak pullbacks.*

This *prima facie* rather exotic property is of great importance in the theory of universal coalgebra. The main reason for this is that Ω preserving weak pullbacks is equivalent to $\overline{\Omega}$ commuting with relational composition, that is, satisfying $\overline{\Omega}(R \circ R') = \overline{\Omega}(R) \circ \overline{\Omega}(R')$. In fact, one may show that any set functor Ω preserves weak pullbacks if and only if $\overline{\Omega}$ is an endofunctor on the category with sets as objects and binary relations as arrows. This result is often attributed to CARBONI, KELLY & WOOD [19], but it already follows from earlier work by Trnková [104, 105] and Barr [10]. In any case, the importance of the notion in the theory of coalgebras lies in the results from RUTTEN [97] that are given in the next theorem.

THEOREM 1.177. *Assume that the functor $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ preserves weak pullbacks. Then the collection of bisimulations is closed under taking relational composition, and the notions of bisimilarity and behavioral equivalence coincide.*


Proof. We leave the proof of the first statement as an exercise for the reader, and concentrate on the second statement. Let s_0 and s_1 be states of the Ω -coalgebras \mathbb{S}_0 and \mathbb{S}_1 , respectively. We need to prove that $\mathbb{S}_0, s_0 \rightleftharpoons \mathbb{S}_1, s_1$ iff $\mathbb{S}_0, s_0 \equiv_{\Omega} \mathbb{S}_1, s_1$. Because of Proposition 1.174 it suffices to prove the direction from right to left.

Let $f_0 : \mathbb{S}_0 \rightarrow \mathbb{X}$ and $f_1 : \mathbb{S}_1 \rightarrow \mathbb{X}$ be two homomorphisms such that $f_0(s_0) = f_1(s_1)$. Then in **Set**, the set $B := \{(s_0, s_1) \in \mathbb{S}_0 \times \mathbb{S}_1 \mid f_0(s_0) = f_1(s_1)\}$, together with the projection functions $\pi_0 : B \rightarrow \mathbb{S}_0$ and $\pi_1 : B \rightarrow \mathbb{S}_1$ constitutes a pullback of f_0 and f_1 , cf. the square in the foreground of the picture. Because Ω preserves weak pullbacks, the diagram in the background of the picture is a weak pullback diagram in **Set**.

Now consider the two arrows $\sigma_i \circ \pi_i : B \rightarrow \Omega(S_i)$. First observe that $\Omega f_i \circ \sigma_i = \xi \circ f_i$ for each i , because each f_i is a coalgebra homomorphism. Hence, chasing the diagram we find that

$$\Omega f_0 \circ \sigma_0 \circ \pi_0 = \xi \circ f_0 \circ \pi_0 = \xi \circ f_1 \circ \pi_1 = \Omega f_1 \circ \sigma_1 \circ \pi_1.$$

Since $\Omega\pi_0$ and $\Omega\pi_1$ form a weak pullback of Ωf_0 and Ωf_1 , this implies the existence of a mediating function $\beta : B \rightarrow \Omega B$ such that $\Omega\pi_i \circ \beta = \sigma_i \circ \pi_i$. In other words, $\mathbb{B} := \langle B, \beta \rangle$ is an Ω -coalgebra, and the projection maps π_0 and π_1 are homomorphisms from \mathbb{B} to \mathbb{S}_0 and \mathbb{S}_1 , respectively.



We finish the section with a brief discussion of *coinduction* as a coalgebraic proof principle. This principle states, for a system \mathbb{S} , that $\simeq \subseteq \Delta_{\mathbb{S}}$; or equivalently, that every bisimulation is a subset of the diagonal $\Delta_{\mathbb{S}}$. The importance of this principle is that, when applicable to \mathbb{S} , in order to prove the identity of two states in \mathbb{S} , it suffices to show that they are linked by some bisimulation. It is not hard to prove that *final coalgebras*, if existing, satisfy the principle of coinduction. This principle has surprisingly powerful applications. For instance, since the class of non-well-founded sets is (in $\text{Coalg}(\mathcal{P}^+)$, cf. Remark 1.161) the final coalgebra of the power set functor, bisimilarity may serve as a notion of identity between sets, see ACZEL [2]. As a second example, RUTTEN [96] is a presentation of the theory of deterministic automata and (regular) languages in which coinduction on the final coalgebra of Example 1.155 is the basic proof principle.

12 COVARIETIES

What is the coalgebraic analog of a variety? In other words, what are natural closure operations on classes of coalgebras? We start with homomorphic images.

DEFINITION 1.178. Let Ω be some endofunctor on Set . If $\varphi : \mathbb{A} \rightarrow \mathbb{B}$ is a surjective homomorphism between the Ω -coalgebras \mathbb{A} and \mathbb{B} , then we say that \mathbb{B} is a *homomorphic image* of \mathbb{A} .

In universal algebra, one finds a one-one correspondence between homomorphic images and congruences. Something similar applies here, but the analogy is perfect only in the case of functors that preserve weak pullbacks.

PROPOSITION 1.179. Let $\mathbb{S} = \langle S, \sigma \rangle$ be an Ω -coalgebra for some set functor Ω . Then

1. Given a bisimulation equivalence E on \mathbb{S} , there is a unique coalgebra structure σ' on S/E such that the quotient map $\nu : S \rightarrow S/E$ is a homomorphism.
2. If Ω preserves weak pullbacks, then $\ker(\varphi)$ is a bisimulation equivalence for any homomorphism $\varphi : \mathbb{S} \rightarrow \mathbb{S}'$.

Proof. For part 1, the coalgebra map σ' can be defined by putting $\sigma'([s]_E) := (\Omega\nu) \circ \sigma(s)$. Further proof details can be found in RUTTEN [97]. For the second part of the proposition,

observe that $\ker(\varphi)$ is the relational composition of the graph of φ with its converse. The result then follows from Theorem 1.177. ■

The next class operation that we consider is that of taking subcoalgebras.

DEFINITION 1.180. Let $\mathbb{A} = \langle A, \alpha \rangle$ and $\mathbb{S} = \langle S, \sigma \rangle$ be two Ω -coalgebras, such that S is a subset of A . If the inclusion map $\iota : S \rightarrow A$ is a homomorphism from $\langle S, \sigma \rangle$ to $\langle A, \alpha \rangle$, then we say that S is *open* with respect to \mathbb{A} , and we call the structure $\langle S, \alpha \upharpoonright_S \rangle$ a *subcoalgebra* of \mathbb{A} .

Interestingly enough, the transition map of a subcoalgebra is completely determined by the underlying open set:

PROPOSITION 1.181. Let $\mathbb{S}_0 = \langle S, \sigma_0 \rangle$ and $\mathbb{S}_1 = \langle S, \sigma_1 \rangle$ be two subcoalgebras of the coalgebra \mathbb{A} . Then $\sigma_0 = \sigma_1$.

Proof. The case of S being empty is trivial, so suppose otherwise. Then from the assumption that \mathbb{S}_0 and \mathbb{S}_1 are subcoalgebras of \mathbb{A} , we may infer that $(\Omega\iota) \circ \sigma_0 = \alpha \circ \iota = (\Omega\iota) \circ \sigma_1$, where ι is the inclusion map of S into A . It follows from the functoriality of Ω that $\Omega\iota$ is an injection, so that we may conclude that $\sigma_0 = \sigma_1$. ■

Some further observations concerning subcoalgebras are in order. First of all, the topological terminology is justified by the following proposition.

PROPOSITION 1.182. Given a coalgebra \mathbb{A} for some set functor Ω , the collection $\tau_{\mathbb{A}}$ of \mathbb{A} -open sets forms a topology.

Proof. Closure of $\tau_{\mathbb{A}}$ under taking (arbitrary) unions follows from Theorem 1.172, together with the observation that

$$S \subseteq A \text{ is open with respect to } \mathbb{A} \text{ iff } \Delta_S \text{ is a bisimulation on } \mathbb{A}, \quad (1.46)$$

which in its turn is an immediate consequence of (1.42). We skip the proof of the fact that the *intersection* of two opens is open, since it requires a little more work. We refer the reader to GUMM & SCHRÖDER [54] for the details. ■

It follows from the Proposition above that, given a subset S of (the carrier of) a coalgebra \mathbb{A} , there is a largest subcoalgebra of \mathbb{A} (of which the carrier is) contained in S : Its universe is given as the union of all open subsets of S . It also follows from Proposition 1.182 that the collection $\tau_{\mathbb{A}}$ of open subsets of A forms a *complete* lattice under set inclusion. Hence, given a subset S of A , there is an open set $U \subseteq A$ which is the *meet* of the collection $\{Q \in \tau_{\mathbb{A}} \mid S \subseteq Q\}$. However, there is no guarantee that U is also the *intersection* of this collection, or, indeed, that S is actually a subset of U . Thus we may not in general speak of the smallest subcoalgebra containing a given subset, as the following example from GUMM [50] witnesses.

EXAMPLE 1.183. Consider the standard Euclidean topology on the real numbers, seen as a coalgebra for the filter functor $\mathcal{F}_{\mathcal{P}}$, cf. Example 1.144. One can show, that a set S of reals is open in the topological sense iff it is open in the sense of Definition 1.180 — in fact, this holds for any topology. Now take an arbitrary point r in \mathbb{R} . Obviously, we have that the *meet* of all open neighborhoods containing r is the empty set.

Before we turn to further coalgebraic constructions, consider the following natural link between homomorphic images and subcoalgebras.

PROPOSITION 1.184. *Given a coalgebraic homomorphism $\varphi : \mathbb{A} \rightarrow \mathbb{B}$, there is a (unique) subcoalgebra $\varphi[\mathbb{A}]$ of \mathbb{B} such that $\varphi : \mathbb{A} \rightarrow \varphi[\mathbb{A}]$ is a surjective homomorphism.*

Proof. For a proof of this proposition, let $S := \varphi[A]$ be the (set-theoretic) image of A under φ , and let $f : S \rightarrow A$ be a right inverse of φ , that is, $\varphi(f(s)) = s$ for all $s \in S$. Now define $\sigma : S \rightarrow \Omega S$ by $\sigma := \Omega\varphi \circ \alpha \circ f$. It can be shown that the resulting structure \mathbb{S} is always a subcoalgebra of \mathbb{B} , and that $\varphi : \mathbb{A} \rightarrow \mathbb{S}$ is a surjective homomorphism; for details the reader is referred to RUTTEN [97]. ■

Our last example of a coalgebraic construction concerns the straightforward generalization of the disjoint union of Kripke models and frames. The idea is as follows. Recall that in **Set**, a concrete representation of the coproduct of a collection $\{A_i \mid i \in I\}$ of sets is given by the disjoint union $\biguplus_I A_i$, together with the inclusions/embeddings $e_i : A_i \rightarrow \biguplus_I A_i$. Hence, the defining property of coproducts provides the key ingredient of the coalgebraic notion of a coproduct, or sum of a family of coalgebras.

DEFINITION 1.185. The *sum* $\coprod_I \mathbb{A}_i$ of a family $\{\mathbb{A}_i \mid i \in I\}$ of coalgebras for some set functor Ω , is defined by endowing the disjoint union $A := \biguplus_I A_i$ with the unique map $\alpha : A \rightarrow \Omega A$ which turns all embeddings $e_i : A_i \rightarrow A$ into homomorphisms.

We have now gathered all the basic class operations needed to define the notion of a covariety, which was introduced in RUTTEN [97] as the natural dual of a variety in universal algebra.

DEFINITION 1.186. Let Ω be some endofunctor on **Set**. A class of Ω -coalgebras is a *covariety* if it is closed under taking homomorphic images, subcoalgebras and sums. The smallest covariety containing a class \mathbb{K} of Ω -coalgebras is called the covariety *generated* by \mathbb{K} , notation: $\text{Covar}(\mathbb{K})$.

As in the case of universal algebra, in order to obtain a more succinct characterization of the covariety generated by a class of coalgebras, one may develop a calculus of class operations.

DEFINITION 1.187. Let \mathbb{H} , \mathbb{S} and $\mathbb{\Sigma}$ denote the class operations of taking (isomorphic copies of) homomorphic images, subcoalgebras, and sums, respectively.

On the basis of these (and other) operations one may investigate the validity of ‘inequalities’ like $\mathbb{HS} \leq \mathbb{SH}$ (meaning that $\mathbb{HS}(\mathbb{K}) \subseteq \mathbb{SH}(\mathbb{K})$ for all classes \mathbb{K} of coalgebras). Results of these kind lead to the following coalgebraic analog of Tarski’s HSP-theorem in universal algebra, due to GUMM & SCHRÖDER [53].

THEOREM 1.188. *Let \mathbb{K} be a class of Ω -coalgebras for some set functor Ω . Then $\text{Covar}(\mathbb{K}) = \mathbb{SH}\mathbb{\Sigma}(\mathbb{K})$.*

Proof. It is straightforward to prove the theorem on the basis of the idempotency of the class operations \mathbb{H} , \mathbb{S} and $\mathbb{\Sigma}$, together with the following three ‘inequalities’: $\mathbb{HS} \leq \mathbb{SH}$, $\mathbb{\Sigma}\mathbb{S} \leq \mathbb{S}\mathbb{\Sigma}$, and $\mathbb{\Sigma}\mathbb{H} \leq \mathbb{H}\mathbb{\Sigma}$. For proofs of these (and more) inequalities, the reader is referred to GUMM & SCHRÖDER [53]. ■

As in the case of varieties, one may wonder about the basic building blocks of varieties. Dualizing the notion of subdirect irreducibility, we arrive at the following definition. It

uses the notion of a conjunct sum, which is known, in the case of Kripke frames, under the name of *bounded union*.

DEFINITION 1.189. Let \mathbb{A} be some Ω -coalgebra for some set functor Ω . A *conjunct representation* \mathbb{A} by a family $\{\mathbb{A}_i \mid i \in I\}$ of coalgebras is a family of embeddings $\{e_i : \mathbb{A}_i \rightarrow \mathbb{A} \mid i \in I\}$ such that $A = \bigcup_{i \in I} e_i[A_i]$. In this case we call \mathbb{A} a *conjunct sum* of the \mathbb{A}_i . A coalgebra \mathbb{A} is called *conjunctly irreducible* if each of its conjunct representations is trivial in the sense that one of the embeddings is an isomorphism.

Covarieties are easily seen to be closed under taking conjunct sums — we will use this fact without further notice.

Given the results on dualizing the notion of subdirect irreducibility in section 5, in particular, Theorem 1.68, one would expect that conjunct irreducibility can be explained in terms of roots. Call a state s of a system \mathbb{S} a *root* of \mathbb{S} if \mathbb{S} itself is the only subcoalgebra of \mathbb{S} that contains s . It is then fairly easy to prove that a coalgebra is conjunctly irreducible if and only if it has a root. However, GUMM [50] proves that there is no analog of Birkhoff's s.i. theorem here, at least not for an arbitrary functor. For instance, expanding Example 1.183, one easily shows that a topological coalgebra will generally not be a conjunct sum of rooted coalgebras.

13 MODAL LOGIC AND COALGEBRAS

If coalgebras are mathematical structures that represent the essence of dynamics, then there is an obvious need for logics to represent and reason about properties of such structures. This is of particular importance for computer scientists who are interested in the formal specification and verification of the behavior of a system. The *kind* of properties that one wants to describe formally may differ from one application to another, but it seems natural to restrict attention to properties that are invariant under behavioral equivalence. MOSS [11, 86] was the first to realize that such properties can be conveniently formalized in a version of *modal logic*, properly generalized from Kripke structures to systems for an arbitrary set functor. This connection between modal logic and coalgebra has provided a quite active research area. At the time of writing, quite a few proposals for coalgebraic modal logics are around; most of them are roughly based on one of the approaches to be discussed in this section.

We start with Moss' original approach, which is also the most general. In order to introduce his formalism, we first put ordinary modal logic in a slightly different perspective by introducing a new connective ∇ . The *meaning* of this modality, which takes a *set* of formulas as its argument, can be summarized by presenting the formula $\nabla\Phi$, with Φ a set of formulas, as the following abbreviation:

$$\nabla\Phi := \square \bigvee \Phi \wedge \bigwedge \diamond\Phi, \quad (1.47)$$

where $\diamond\Phi$ denotes the set $\{\diamond\varphi \mid \varphi \in \Phi\}$, and \bigvee and \bigwedge denote disjunction and conjunction. We do not want to exclude the possibility that Φ is an infinite set — coalgebraic logic is generally of an infinitary nature. The operator ∇ pops up in a number of areas in modal logic, cf. for instance the characteristic formulas of Chapter 5. We may also decide to treat this ∇ as a *primitive* connective. As long as we keep \vee and \top in our language, both the standard diamond and box connective are definable in terms of ∇ ,

since we have the following equivalences:

$$\begin{aligned}\diamond\varphi &\equiv \nabla\{\varphi, \top\}, \\ \square\varphi &\equiv \nabla\emptyset \vee \nabla\{\varphi\},\end{aligned}$$

so that we may in fact *replace* the diamond and box with this new modality.

Spelling out the truth definition of $\nabla\Phi$, we see that it can in fact be expressed in terms of the *relation lifting* that we defined in section 11. For, let $\mathbb{S} = \langle S, \lambda, R[\cdot] \rangle$ be a modal model in coalgebraic shape. Then it is straightforward to verify that $\mathbb{S}, s \Vdash \nabla\Phi$ if and only if the pair $(R[s], \Phi)$ belongs to the relation lifting $\overline{\mathcal{P}}(\Vdash_{\mathbb{S}})$ of the satisfaction relation $\Vdash_{\mathbb{S}} \subseteq S \times \Phi$: Every $\varphi \in \Phi$ must hold at some successor $t \in R[s]$, and at every successor t of s some $\varphi \in \Phi$ must hold, see (1.43). This fundamental insight paves the way for Moss' development of *coalgebraic logic*, in which the same principle is applied to an arbitrary (but fixed) set functor Ω . Basically, the idea is to have

$$\mathbb{S}, s \Vdash_{\mathbb{S}} \nabla P \text{ iff } (P, \sigma(s)) \in \overline{\Omega}(\Vdash_{\mathbb{S}}).$$

Note that in this perspective, the satisfaction relation is much like a bisimulation between a language and a coalgebra; this observation was first made and exploited in BALTAG [8].

In order to provide a more precise definition, recall from Remark 1.161 that we may uniquely extend Ω to a set based endofunctor Ω^+ on the category SET that has classes as objects, and set-continuous functions as arrows. For convenience, we follow MOSS [86] in that we confine our attention to *standard* set functors, that is, functors that map inclusions to inclusions.

DEFINITION 1.190. Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ be a standard set functor that preserves weak pullbacks. Then \mathcal{L}_{Ω} , the language of *coalgebraic formulas for Ω* , is defined as the least class \mathcal{C} such that (i) $\bigwedge \Phi \in \mathcal{L}_{\Omega}$ if $\Phi \subseteq \mathcal{L}_{\Omega}$ is a *set* of formulas, and (ii) $\nabla P \in \mathcal{L}_{\Omega}$ for any $P \in \Omega^+(\mathcal{L}_{\Omega})$.

Categorically, $\langle \mathcal{L}_{\Omega}, \bigwedge, \nabla \rangle$ can be characterized as the *initial algebra* of the functor $(\mathcal{P} + \Omega)^+$. This explains our move to the category SET: if we want to guarantee the *existence* of such a structure, for reasons similar as given in the discussion following Proposition 1.158, we need to allow class-based algebras.

DEFINITION 1.191. Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ preserve weak pullbacks. Given an Ω -coalgebra $\mathbb{S} = \langle S, \sigma \rangle$, define $\Vdash_{\mathbb{S}} \subseteq S \times \mathcal{L}_{\Omega}$ as the least relation satisfying

$$\begin{aligned}s \Vdash_{\mathbb{S}} \bigwedge \Phi &\text{ if } s \Vdash_{\mathbb{S}} \varphi \text{ for all } \varphi \in \Phi, \\ s \Vdash_{\mathbb{S}} \nabla P &\text{ if } (P, \sigma(s)) \in \overline{\Omega}(W) \text{ for some set } W \subseteq \Vdash_{\mathbb{S}}.\end{aligned}$$

EXAMPLE 1.192. Consider the functor $\mathcal{P}\mathbf{Prop} \times \mathcal{P}$ of Kripke models. Unraveling the definitions, we find that an arbitrary element of $\Omega^+(\mathcal{L}_{\Omega})$ must be of the form (A, Φ) with $A \subseteq \mathbf{Prop}$ a set of proposition letters, and $\Phi \subseteq \mathcal{L}_{\Omega}$ a set of formulas. It is not hard to verify that

$$\mathbb{S}, s \Vdash \nabla(A, \Phi) \text{ iff } \mathbb{S}, s \Vdash \bigwedge A \wedge \bigwedge \neg(\mathbf{Prop} \setminus A) \wedge \square \bigvee \Phi \wedge \bigwedge \diamond\Phi,$$

where $\bigwedge \neg(\mathbf{Prop} \setminus A)$ denotes the formula $\bigwedge \{\neg p \mid p \in P \setminus A\}$. It is instructive to observe the difference between this and (1.47) which displays an arbitrary ∇ -formula for the functor \mathcal{P} of Kripke frames as opposed to models.

EXAMPLE 1.193. For another example, an arbitrary element of the class $\Omega^+(\mathcal{L}_\Omega)$, where Ω is now the functor $\mathcal{I} \times \mathcal{I}$, must be a *pair* of formulas, say, (φ_0, φ_1) . Clearly then we have

$$\mathbb{S}, s \Vdash \nabla(\varphi_0, \varphi_1) \text{ iff } \mathbb{S}, \pi_0(\sigma(s)) \Vdash \varphi_0 \text{ and } \mathbb{S}, \pi_1(\sigma(s)) \Vdash \varphi_1.$$

This in fact implies that *all* formulas are true at *all* states of *all* coalgebras; in other words, in the absence of propositions, the language \mathcal{L}_Ω may be rather uninteresting.

Obviously, many variations of this language exist, or may be defined. For instance, it is easy to develop finitary versions of the language, while independently of this, one may add Boolean connectives like negation or (infinitary) disjunction. Interestingly, \mathcal{L}_Ω on its own is already powerful enough to characterize behavior. Theorem 1.194 below shows that it has the *Hennessy-Milner property* (cf. Chapter 5 of this volume): non-bisimilarity of two points is witnessed by some formula in the language.

THEOREM 1.194. *Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ preserve weak pullbacks, and let \mathbb{S} and \mathbb{S}' be two Ω -coalgebras. Then for any pair of states $s \in S, s' \in S'$:*

$$\mathbb{S}, s \rightleftharpoons \mathbb{S}', s' \text{ iff } s \text{ and } s' \text{ satisfy the same } \mathcal{L}_F\text{-formulas.}$$

Proof. The direction from left to right is proved by induction on the complexity of formulas. That is, we define Θ to be the class of formulas on which all bisimilar points in \mathbb{S} and \mathbb{S}' agree. Then we prove that $\Theta = \mathcal{L}_\Omega$ by showing that Θ is closed under \bigwedge and ∇ (in the sense that $\bigwedge \Phi \in \Theta$ for all subsets $\Phi \subseteq \Theta$, and that $\nabla P \in \Theta$ for all $P \in \Omega(\Theta)$). We leave the fairly straightforward details as an exercise for the reader.

The proof for the other direction is analogous to that of Karp's Theorem for modal logic (see Chapter 5 of this volume), so we confine ourselves to a brief sketch here. Given an Ω -system \mathbb{S} , by ordinal induction we define a family $\varphi_\alpha^{\mathbb{S}} : S \rightarrow \mathcal{L}_\Omega$ as follows (we omit the superscript):

$$\begin{aligned} \varphi_0(s) &:= \top \\ \varphi_{\alpha+1}(s) &:= \nabla(\Omega\varphi_\alpha)(\sigma(s)), \\ \varphi_\lambda(s) &:= \bigwedge \{\varphi_\alpha(s) \mid \alpha < \lambda\}. \end{aligned}$$

One approach to the proof would then be to show that the relation \equiv_φ , defined via $s \equiv_\varphi t$ if $\varphi_\alpha(s) = \varphi_\alpha(t)$ for all α , is itself a bisimulation. \blacksquare

Moss' definition provides powerful languages, of which syntax and semantics uniformly depend on the coalgebraic signature, but his systems are not very welcoming to our intuitions on modal languages as extensions of propositional logic with diamonds and boxes that are interpreted via accessibility relations. BALTAG [8] introduces variants of Moss' language in which the connectives $\square \bigvee \Phi$ and $\bigwedge \diamond \Phi$ of (1.47) are (separately) generalized from Kripke frames to arbitrary functors, but also his formalism is far too abstract for practical purposes. It therefore seems worthwhile to develop more 'concrete' and practical alternatives to \mathcal{L}_Ω .

In the case of Kripke polynomial functors, the concrete, inductive definition of the functor allows for more down to earth modal languages, as was first observed by KURZ [77]. Here we present a formalism that was introduced in ROESSIGER [95], and studied by

JACOBS [62]. From the perspective of modal logic, its only non-standard feature is that both its syntax and semantics are *sorted* by the set $IngK$ of ingredient functors of K .

DEFINITION 1.195. Fix a Kripke polynomial functor K . We define the language $Fma_K = \bigcup_{\Lambda \in Ing(K)} Fma_K(\Lambda)$ of K -sorted modal formulas, by the following induction. (All functors appearing in the definition below are supposed to be ingredient functors of K .)

- $\perp \in Fma_K(\Lambda)$ for every $\Lambda \in Ing(K)$;
- if $\varphi, \psi \in Fma_K(\Lambda)$ then $\neg\varphi, \varphi \vee \psi \in Fma_K(\Lambda)$;
- if $c \in C$ then $c \in Fma_K(C)$;
- if $\varphi \in Fma_K(\Lambda_i)$ then $\diamond_{\kappa_i}\varphi \in Fma_K(\Lambda_0 + \Lambda_1)$;
- if $\varphi \in Fma_K(\Lambda_i)$ then $\diamond_{\pi_i}\varphi \in Fma_K(\Lambda_0 \times \Lambda_1)$;
- if $\varphi \in Fma_K(\Lambda)$ then $\diamond_d\varphi \in Fma_K(\Lambda^D)$ for all $d \in D$;
- if $\varphi \in Fma_K(\Lambda)$ then $\diamond_{\supset}\varphi \in Fma_K(\mathcal{P}\Lambda)$;
- if $\varphi \in Fma_K(K)$ then $\odot\varphi \in Fma_K(\mathcal{I})$.

We say that φ is of sort Λ if $\varphi \in Fma_K(\Lambda)$ — note that this sort need not be unique.

How do we interpret these formulas in coalgebras? Intuitively, with each K -coalgebra \mathbb{S} , we associate a multi-sorted frame based on the set $\bigcup_{\Lambda \in Ing(K)} \Lambda(S)$. The accessibility relations of this frame (which we will not make explicit) are completely determined by the shape of the functor. For instance, to link the set $(\Lambda_0 + \Lambda_1)(S)$ to $\Lambda_0(S)$, we lay down the relation $R_{\kappa_0} = \{(\kappa_0 s_0, s_0) \mid s_0 \in \Lambda_0(S)\}$. Likewise, the converse membership relation \supset provides the accessibility relation from $\mathcal{P}\Lambda(S)$ to $\Lambda(S)$.

DEFINITION 1.196. Let $\mathbb{S} = \langle S, \sigma \rangle$ be a K -coalgebra for some Kripke polynomial functor K . By formula induction we define a sorted satisfaction relation $\Vdash = \bigcup_{\Lambda \in Ing(K)} \Vdash_{\Lambda}$, with $\Vdash_{\Lambda} \subseteq \Lambda(S) \times Fma_K(\Lambda)$:

$$\begin{array}{ll}
s \Vdash_{\Lambda} \perp & : \text{ never,} \\
s \Vdash_{\Lambda} \neg\varphi & \text{if } s \not\Vdash_{\Lambda} \varphi \text{ (but } s \in \Lambda(S)\text{),} \\
s \Vdash_{\Lambda} \varphi \vee \psi & \text{if } s \Vdash_{\Lambda} \varphi \text{ or } s \Vdash_{\Lambda} \psi, \\
s \Vdash_C c & \text{if } s = c, \\
s \Vdash_{\Lambda_0 + \Lambda_1} \diamond_{\kappa_i}\varphi & \text{if } s = \kappa_i(t) \text{ for some } t \in \Lambda_i(S) \text{ with } t \Vdash_{\Lambda_i} \varphi, \\
s \Vdash_{\Lambda_0 \times \Lambda_1} \diamond_{\pi_i}\varphi & \text{if } s = (s_0, s_1) \text{ and } s_i \Vdash_{\Lambda_i} \varphi, \\
s \Vdash_{\Lambda^D} \diamond_d\varphi & \text{if } s(d) \Vdash_{\Lambda} \varphi, \\
s \Vdash_{\mathcal{P}\Lambda} \diamond_{\supset}\varphi & \text{if there is some } t \in s \text{ with } t \Vdash_{\Lambda} \varphi, \\
s \Vdash_{\mathcal{I}} \odot\varphi & \text{if } \sigma(s) \Vdash_K \varphi.
\end{array}$$

Furthermore we employ the usual terminology concerning validity, etc.

EXAMPLE 1.197. Consider the functor $\Omega = \mathcal{P}\mathbf{Prop} \times \mathcal{P}(\mathcal{I} \times \mathcal{I})$ corresponding to Kripke models based on frames with a ternary accessibility relation T . In the standard modal

language for such models, we would be working with a binary modality \diamond , whereas here, we are dealing with four unary modalities: \odot , \diamond_{\exists} , \diamond_{π_1} and \diamond_{π_2} . We leave it for the reader to verify that the modal formula $\varphi_1 \diamond \varphi_2$ in the first language can be rendered as $\odot \diamond_{\exists} (\diamond_{\pi_0} \varphi_1 \wedge \diamond_{\pi_1} \varphi_2)$ in the second. That is, we have

$$\mathbb{S}, s \models \odot \diamond_{\exists} (\diamond_{\pi_0} \varphi_1 \wedge \diamond_{\pi_1} \varphi_2) \text{ iff there are } t_1, t_2 \text{ with } Tst_1t_2 \text{ and } \mathbb{S}, t_i \Vdash \varphi_i.$$

Bisimulation invariance of this language is easily proved:

PROPOSITION 1.198. *Assume that K is some Kripke polynomial functor, and let \mathbb{S} and \mathbb{S}' be two K -coalgebras. Then for any pair of states $s \in S$, $s' \in S'$:*

$$\mathbb{S}, s \Leftrightarrow \mathbb{S}', s' \text{ only if } s \text{ and } s' \text{ satisfy the same formulas in } Fma_K.$$

Proof. Fix a bisimulation B between \mathbb{S} and \mathbb{S}' . We claim that for any formula φ of type $\Lambda \in \text{Ing}(K)$, it holds for any pair $(s, s') \in \Lambda(S) \times \Lambda(S')$ that

$$\mathbb{S}, s \Vdash_{\Lambda} \varphi \text{ iff } \mathbb{S}', s' \Vdash_{\Lambda} \varphi,$$

provided that (s, s') belong to the relation lifting $\bar{\Lambda}(B)$ of B . The proof is by a straightforward formula induction. ■

The basic modal theory of this formalism has been developed. For instance, analogous to Theorem 38 in Chapter 5 of this volume, one may prove that if K is a *finitary* Kripke polynomial functor, then the language Fma_K has the Hennessy-Milner property. Also, results concerning completeness and decidability are known. The interested reader is referred to RÖSSIGER [95] and JACOBS [62].

We now move to the third approach towards coalgebraic modal logic. PATTINSON [89] combines the generality of the first formalism with the concreteness of the second. That is, the approach applies to arbitrary set functors, but provides languages with standard diamonds and boxes. First we present a simplified version, which is based on the idea to *extract diamonds out of the natural transformations* from the coalgebra functor Ω to the power set functor \mathcal{P} . Recall that a natural transformation $\lambda : \Omega \rightarrow \mathcal{P}$ provides an arrow $\lambda_S : \Omega(S) \rightarrow \mathcal{P}(S)$ for each set S , in such a way that for each function $f : S \rightarrow S'$, the following diagram commutes:

$$\begin{array}{ccccc} S & & \Omega S & \xrightarrow{\lambda_S} & \mathcal{P} S \\ f \downarrow & & \Omega f \downarrow & & \mathcal{P} f \downarrow \\ S' & & \Omega S' & \xrightarrow{\lambda_{S'}} & \mathcal{P} S' \end{array}$$

Thus if we have an Ω -coalgebra $\mathbb{S} = \langle S, \sigma \rangle$, we may define a relation $R_{\lambda} \subseteq S \times S$ for such a λ by putting $R_{\lambda}st$ if $t \in \lambda_S(\sigma(s))$. We may then introduce a diamond \diamond_{λ} which takes this R_{λ} as its accessibility relation. Natural transformations $\lambda : \Omega \rightarrow \mathcal{P}$ thus literally transform Ω -coalgebras into \mathcal{P} -coalgebras, that is, Kripke frames.

Similarly, if we want to have *atomic propositions* in our language, consider any natural transformation ν from Ω to the constant functor \mathbf{Prop} . We then make $p \in \mathbf{Prop}$ true at s depending on whether p is an element of the set $\nu_S(\sigma(s))$ or not. It is as if we add the valuation V_{ν} to \mathbb{S} given by $V_{\nu}(p) := \{s \in S \mid p \in \nu_S(\sigma(s))\}$.

DEFINITION 1.199. Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ be some functor, $\nu : \Omega \rightarrow \mathbf{Prop}$ some natural transformation, and Λ some collection of natural transformations $\Omega \rightarrow \mathcal{P}$. Then $\mathcal{L}_{\nu, \Lambda}$ is the standard modal language we obtain by taking \mathbf{Prop} as the collection of propositional variables, and $\tau_{\Lambda} := \{\diamond_{\lambda} \mid \lambda \in \Lambda\}$ as the modal similarity type.

It will now be obvious how these formulas are interpreted in Ω -coalgebras. We confine ourselves to the following clauses of the inductive truth definition:

$$\begin{aligned} \mathbb{S}, s \Vdash p & \quad \text{if} \quad p \in \nu(\sigma(s)), \\ \mathbb{S}, s \Vdash \diamond_{\lambda} \varphi & \quad \text{if} \quad \mathbb{S}, t \Vdash \varphi \text{ for some } t \in \lambda_S(\sigma(s)). \end{aligned}$$

In other words, an Ω -coalgebra \mathbb{S} is treated as the Kripke model $\langle S, \{R_{\lambda} \mid \lambda \in \Lambda\}, V_{\nu} \rangle$. The reason to require the transformations to be *natural* is to guarantee invariance under behavioral equivalence.

PROPOSITION 1.200. *Let Ω , ν and Λ be as in Definition 1.199. Then for any pair \mathbb{S}, \mathbb{S}' of Ω -coalgebras, and any pair of states $s \in S, s' \in S'$:*

$$\mathbb{S}, s \equiv_{\Omega} \mathbb{S}', s' \text{ only if } s \text{ and } s' \text{ satisfy the same } \mathcal{L}_{\nu, \Lambda}\text{-formulas.}$$

Proof. It suffices to prove that for any coalgebraic homomorphism $f : \mathbb{S} \rightarrow \mathbb{S}'$, each state s in \mathbb{S} satisfies the same $\mathcal{L}_{\nu, \Lambda}$ -formulas as $f(s)$ in \mathbb{S}' . This inductive proof is in fact straightforward, the crucial observation being that the naturality of the transformations guarantees that f is a bounded morphism between the Kripke models associated with \mathbb{S} and \mathbb{S}' . ■

For the more general picture, Pattinson uses *predicate liftings* (from $\mathcal{P}S$ to $\mathcal{P}\Omega S$) to obtain modal operators. In order to introduce these, note that the semantics of the modal operator \diamond_{λ} could have been expressed as follows:

$$\mathbb{S}, s \Vdash \diamond_{\lambda} \varphi \text{ iff } \sigma(s) \in \mu_S^{\lambda}(\llbracket \varphi \rrbracket),$$

where $\mu_S^{\lambda} : \mathcal{P}S \rightarrow \mathcal{P}\Omega S$ is given by $A \mapsto \{\Gamma \in \Omega S \mid \lambda_S(\Gamma) \cap A \neq \emptyset\}$, and $\llbracket \varphi \rrbracket$ denotes the extension of φ in \mathbb{S} . In fact, it can be shown that μ^{λ} is a natural transformation from the contravariant power set functor $\check{\mathcal{P}}$ to the functor $\check{\mathcal{P}} \circ \Omega$. Generalizing this, we arrive at the following definition.

DEFINITION 1.201. A *predicate lifting* for a set functor Ω is a natural transformation $\mu : \check{\mathcal{P}} \rightarrow \check{\mathcal{P}} \circ \Omega$. With each predicate lifting we can associate a modal operator \diamond_{μ} , with the following semantics:

$$\mathbb{S}, s \Vdash \diamond_{\mu} \varphi \text{ iff } \sigma(s) \in \mu_S(\llbracket \varphi \rrbracket).$$

And as before, it is the naturality of the transformation that ensures that this language is invariant under behavioral equivalence.

In order to finish this section, a number of remarks are in order. First, the above mentioned versions of coalgebraic logic are open for the standard expressive enhancements that we know from extended modal logic. As examples we mention JACOBS [64], who adds past operators (as in section 8.1) to a variant of the formalism defined in the Definitions 1.195 and 1.196, and VENEMA [107], who develops a finitary fixed point version of Moss' logic.

Second, it should be mentioned that for certain *polynomial* functors, coalgebraic specification languages have been developed of an *equational* rather than modal nature. Very roughly, the idea is that coalgebras for such a polynomial functor K can be represented by a structured collection of partial functions on the carrier of the coalgebra. From the perspective of Definition 1.196, this can be explained by the observation that in the absence of the power set functor, each and every accessibility relation of the multi-sorted frame is in fact (the graph of) a partial function. Lacking the space for an appropriate survey of this more equational perspective, we only mention one interesting idea which adds some *modal flavor* to equational logic. In coalgebraic approaches towards specification theory, such as that of *hidden algebra*, a *state equation* $t_1 \simeq t_2$ holds of a state s in a coalgebra \mathbb{S} if $t_1^{\mathbb{S}}(s)$ and $t_2^{\mathbb{S}}(s)$ evaluate to *bisimilar* (rather than identical) states in \mathbb{S} . We refer the reader to GOLDBLATT [42, 43] and ROŞU [94] for more details; in particular, GOLDBLATT [43] contains a clear discussion of this overlap area between modal and equational logic.

Third, KURZ & PATTINSON [78] establish a link between coalgebraic predicates and the *final sequence*, see Remark 1.162: they argue that finitary predicates correspond to subsets of some set Z_n (n finite) occurring in the final sequence. This work is in fact closely related to that of GHILARDI [33], even though the word ‘coalgebra’ is not mentioned in the latter work.

Finally, there is an interesting connection between Hennessy-Milner results and final coalgebras: Goldblatt [45] proves that a set functor Ω admits a final coalgebra iff there is a coalgebraic modal language for Ω , which has the Hennessy-Milner property and is based on a *set* (rather than a proper class) of formulas.

14 CO-BIRKHOFF THEOREMS AND COFREE COALGEBRAS

In order to give the reader some impression of universal coalgebra at work, we discuss one result, or better, one cluster of results, in some detail. The topic that we have chosen concerns the coalgebraic version of Birkhoff’s variety theorem; recall that this result in universal algebra states that a class C of algebras is a *variety*, (that is, closed under the class operations H , S and P), if and only if it is equationally definable. Thus in essence, Birkhoff established a link between two different ways of characterizing algebraic classes: a logical one, in terms of the validity of certain formulas, and a structural one, in terms of certain class operations.

If we are after a co-Birkhoff result, two roads seem open to us. Since we have already developed the concept of a covariety, the most obvious thing to do would be to try and find out what corresponds to it, logically. An alternative approach would be to investigate the structural counterpart of the logical languages developed in the previous section. Here we follow the first road, but interestingly, it leads us to (very natural generalizations of) modal languages! This provides justification for our earlier claim that modal logic is dual to equational logic.

In the proof of Birkhoff’s theorem, *free* algebras play a key role; thus it will come as no surprise that we will be looking at *cofree* coalgebras here. However, these structures do not serve as proof tools only, they have a quite intuitive meaning as well. To explain this, first note that many set functors provide coalgebraic structures that come with a notion of *output*. For instance, the black box machines of Example 1.141 may be prompted to display some value, the states of the automata of Example 1.142 output 0 or 1 depending


on whether they are final or not, and the states of a Kripke model satisfy some set of propositional variables. For a general functor $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$, such a notion of output may not be available. However, nothing prevents us from *adding* an extra output feature to the functor.

DEFINITION 1.202. Let Ω be some set functor, and C a set of objects that we will call *colors*. A C -coloring of an Ω -coalgebra $\mathbb{A} = \langle A, \alpha \rangle$ is a map $\gamma : A \rightarrow C$; the structure $\langle A, \alpha, \gamma \rangle$ will be called the coalgebra \mathbb{A} colored by γ .

As a prime example, Kripke models can be seen as $\mathcal{P}\mathbf{Prop}$ -colored Kripke frames. In general, C -colored Ω -coalgebras may be identified with Ω_C -coalgebras, where Ω_C is the functor $C \times \Omega$; this provides us with a category of C -colored Ω -coalgebras. Spelling it out, $f : S \rightarrow S'$ is a morphism from $\langle S, \sigma, \gamma \rangle$ to $\langle S', \sigma', \gamma' \rangle$ if f is an Ω -coalgebra homomorphism from $\langle S, \sigma \rangle$ to $\langle S', \sigma' \rangle$ such that $\gamma(s) = \gamma'(fs)$ for all $s \in S$.

Colors can be seen as the coalgebraic duals of variables, colorings as the duals of assignments. This brings us to the definition of a cofree coalgebra, which is the formal dual of the notion of a free algebra. We recall the latter notion, for the purposes of the present context, as follows. Let $\Omega : \mathbf{Set} \rightarrow \mathbf{Set}$ be some set functor, X a set of variables, and $T = \langle T, \tau : \Omega T \rightarrow T \rangle$ some Ω -algebra such that $e : X \rightarrow T$ is some kind of injection. (Here we deviate from the more standard presentation, where e is taken to be an inclusion map.) Then T , with e , is called *free over X* if for every Ω -algebra $\mathbb{A} = \langle A, \alpha \rangle$ and every assignment $f : X \rightarrow A$, there is a unique homomorphism $\tilde{f} : T \rightarrow \mathbb{A}$ such that $f = \tilde{f} \circ e$.

DEFINITION 1.203. Let Ω be a set functor, C a set of colors, and \mathbb{Z} some Ω -coalgebra with a coloring $\gamma : Z \rightarrow C$. Then \mathbb{Z} (with γ) is called (*absolutely*) *cofree over C* if for every Ω -coalgebra $\mathbb{A} = \langle A, \alpha \rangle$ and every coloring $g : A \rightarrow C$ of \mathbb{A} , there is a unique homomorphism $\tilde{g} : \mathbb{A} \rightarrow \mathbb{Z}$ such that $g = \gamma \circ \tilde{g}$.



Observe that the diagram above is not properly typed (it mixes arrows from different categories). A more proper formulation of the notion of cofreeness would involve the right adjoint to the forgetful functor from $\mathbf{Coalg}(\Omega)$ to \mathbf{Set} .

It is immediate from the definitions that an Ω -coalgebra with coloring $\gamma : T \rightarrow C$ is cofree over C iff the structure $\langle T, \tau, \gamma \rangle$ is a *final* coalgebra for the functor $\Omega_C = C \times \Omega$. This explains that we may view the carrier Z of such a cofree coalgebra as the collection of all *behavior patterns* expressible in the output set C . And this perspective paves the way for a dual version of Birkhoff's variety theorem, by providing a natural means for characterizing classes of coalgebras in terms of permitted, or forbidden, behaviors.

DEFINITION 1.204. Let Ω be some set functor, and let \mathbb{Z} , with coloring $\gamma : Z \rightarrow C$, be the cofree coalgebra over some set C of colors. Given a set Q in \mathbb{Z} , let $\mathbf{Cov}(Q)$ be the class of Ω -coalgebras \mathbb{A} such that $\eta[A] \subseteq Q$ for all homomorphisms $\eta : \mathbb{A} \rightarrow \mathbb{Z}$.

And conversely, given a class \mathbf{K} of Ω -coalgebras, define $\mathbf{Bhv}(\mathbf{K}) \subseteq Z$ to be the union of all images $\tilde{g}[A]$ in which \tilde{g} arises from some C -coloring g of some coalgebra \mathbb{A} in \mathbf{K} .

There are all kinds of interesting facts concerning these two maps. For instance, it is fairly obvious from the definitions that \mathbf{Bhv} and \mathbf{Cov} form a (dual) Galois connection: For any class \mathbf{K} of Ω -coalgebras, and any set Q of behavior patterns, we have

$$\mathbf{Bhv}(\mathbf{K}) \subseteq Q \text{ iff } \mathbf{K} \subseteq \mathbf{Cov}(Q). \tag{1.48}$$

We will have use for this fact in the proof of a first co-Birkhoff result, which is basically

due to RUTTEN [97]. In the remainder of this section we restrict our attention to *small* functors, in order to ensure the existence of final and cofree coalgebras.

THEOREM 1.205. *Let Ω be some endofunctor on \mathbf{Set} which is κ -small for some cardinal κ . Then for any set C of size κ , the cofree coalgebra over C exists, and a class \mathbf{K} of Ω -coalgebras is a covariety iff $\mathbf{K} = \mathbf{Cov}(Q)$ for some set Q of behavior patterns.*

Proof. It follows from the assumption on Ω that the functor $\Omega_C = C \times \Omega$ has a final coalgebra. However, we already observed that this structure may be represented as a triple $\langle Z, \zeta, \gamma \rangle$ such that $\mathbb{Z} = \langle Z, \zeta \rangle$, with coloring γ , is the cofree Ω -coalgebra over C . We fix this \mathbb{Z} and γ for the remainder of the proof.

In order to show that $\mathbf{Cov}(Q)$ is a covariety, one needs to subsequently prove closure under taking homomorphic images, subcoalgebras, and sums. Here we restrict our attention to the proof for subcoalgebras, because that is the only part where the cofreeness of \mathbb{Z} is used.

Suppose that \mathbb{A} is a subcoalgebra of \mathbb{B} , with inclusion ι , while \mathbb{B} belongs to $\mathbf{Cov}(Q)$; we need to show that \mathbb{A} also belongs to this class. For that purpose, consider a homomorphism $\eta : \mathbb{A} \rightarrow \mathbb{Z}$, and observe that $\gamma \circ \eta : A \rightarrow C$ is a coloring of \mathbb{A} . Clearly this coloring can be extended to a coloring $g : B \rightarrow C$ of \mathbb{B} . Let $\tilde{g} : \mathbb{B} \rightarrow \mathbb{Z}$ be the unique homomorphism such that $g = \gamma \circ \tilde{g}$ — such a map exists by the cofreeness of \mathbb{Z} .

$$\begin{array}{ccc} A & \xrightarrow{\eta} & Z \\ \downarrow \iota & \nearrow \tilde{g} & \downarrow \gamma \\ B & \xrightarrow{g} & C \end{array}$$

Now $g = \gamma \circ \tilde{g}$, so that $\gamma \circ \tilde{g} \circ \iota = g \circ \iota$. But g was chosen so that $g \circ \iota = \gamma \circ \eta$. Hence we find that $\gamma \circ \tilde{g} \circ \iota = \gamma \circ \eta$, so by the cofreeness of \mathbb{Z} with respect to colorings of \mathbb{A} , we find that $\tilde{g} \circ \iota = \eta$, that is, \tilde{g} extends η . From this it is immediate that $\eta[A] = \tilde{g}|_A[A] \subseteq \tilde{g}[B]$, so that $\eta[A] \subseteq Q$ by the assumption that \mathbb{B} belongs to $\mathbf{Cov}(Q)$.

For the other direction of the theorem, suppose that \mathbf{K} is a covariety; we claim that

$$\mathbf{K} = \mathbf{Cov}(Bhv(\mathbf{K})). \quad (1.49)$$

The inclusion \subseteq is immediate from (1.48). For the opposite inclusion, it easily follows from the definitions that $Bhv(\mathbf{K})$ is \mathbb{Z} -open. Let \mathbb{B}_κ be the (unique) subcoalgebra of \mathbb{Z} with carrier set $Bhv(\mathbf{K})$. It is not hard to prove that \mathbb{B}_κ is a conjunct sum of algebras in \mathbf{K} , which implies that \mathbb{B}_κ actually belongs to \mathbf{K} since covarieties are closed under taking conjunct sums. Hence, in order to prove the remaining inclusion \supseteq of (1.49), it suffices to show that

$$\text{every coalgebra in } \mathbf{Cov}(Bhv(\mathbf{K})) \text{ is a conjunct sum of subcoalgebras of } \mathbb{B}_\kappa. \quad (1.50)$$

Take an arbitrary coalgebra \mathbb{A} in $\mathbf{Cov}(Bhv(\mathbf{K}))$. From the κ -smallness of Ω it may be derived that \mathbb{A} is the conjunct sum of coalgebras \mathbb{A}_i , each of size at most κ . Clearly then it suffices to prove that each \mathbb{A}_i belongs to \mathbf{K} , since covarieties are closed under taking conjunct sums.

Fix some $i \in I$; clearly $\mathbf{Cov}(Bhv(\mathbf{K}))$, being closed under taking subcoalgebras, contains \mathbb{A}_i . Since $|A_i| \leq \kappa = |C|$, there is an injective coloring $e_i : A_i \rightarrow C$. Hence by cofreeness of \mathbb{Z} there is a unique homomorphism $\tilde{e}_i : \mathbb{A}_i \rightarrow \mathbb{Z}$ such that $e_i = \gamma \circ \tilde{e}_i$. This \tilde{e}_i must also be injective, which implies that \mathbb{A}_i is isomorphic to its image $\tilde{e}_i[A_i]$. But, since \mathbb{A}_i belongs to $\mathbf{Cov}(Bhv(\mathbf{K}))$, the structure $\tilde{e}_i[A_i]$ is a subcoalgebra of \mathbb{B}_κ , and thus, belongs to \mathbf{K} . From this it is immediate that each \mathbb{A}_i belongs to \mathbf{K} , and thus, so does the conjunct sum \mathbb{A} . \blacksquare

Clearly, not only the statement, but also the proof of Theorem 1.205 is dual to that of Birkhoff's variety theorem. For instance, the coalgebra \mathbb{B}_κ clearly fulfills the role of the *cofree* coalgebra for the class \mathbf{K} over the color set C . What seems to be missing from Theorem 1.205, however, is some notion of logic, involving *syntax*. (It should be noted that also in the algebraic case, the straightforward characterization of varieties in terms of equations only obtains in the case of relatively simple functors.) Since we are discussing a dual of Birkhoff's theorem, the question this raises is: what are *co-equations*?

Given the nature of systems as state-based models of dynamics, it seems natural to require that formulas describe *behavior*. This would provide natural constraints on possible coequational languages, namely, that formulas are evaluated at states, in such a way that truth is invariant under behavioral equivalence. Furthermore, we allow the use of colors in order to obtain sufficient expressive power. It was an insight of KURZ [76] that these requirements may also be read as a natural *definition* of coalgebraic modal logic.

DEFINITION 1.206. Let Ω be some set functor. A *coalgebraic modal language* for Ω consists of a set C of colors, a class \mathcal{L}_C of formulas, and, for each C -colored Ω -coalgebra $\langle \mathbb{S}, g \rangle$, a truth or satisfaction relation $\Vdash^{\mathbb{S}, g} \subseteq S \times \mathcal{L}_C$ such that \Vdash is invariant under behavioral equivalence. That is, if $\langle \mathbb{S}, g \rangle, s \equiv_{\Omega_C} \langle \mathbb{T}, h \rangle, t$, then $\mathbb{S}, s \equiv_{\mathcal{L}_C} \mathbb{T}, t$, where the latter notation indicates that s in $\langle \mathbb{S}, g \rangle$ and t in $\langle \mathbb{T}, h \rangle$ satisfy exactly the same \mathcal{L}_C -formulas φ .

In the sequel we will use notation and terminology from modal logic. For instance, we write $\langle \mathbb{S}, g \rangle, s \Vdash \varphi$ instead of $s \Vdash^{\mathbb{S}, g} \varphi$, and we define $\mathbb{S}, g \Vdash \varphi$ and $\mathbb{S} \Vdash \varphi$ by quantifying over all elements and all valuations, respectively.

How can we link such modal languages to the cofree coalgebra? The idea here is that modal formulas correspond to subcoalgebras: if \mathbb{Z} , with C -coloring γ is a cofree coalgebra over C , then define

$$\llbracket \varphi \rrbracket^{\mathbb{Z}, \gamma} := \{z \in Z \mid \mathbb{Z}, \gamma, z \Vdash \varphi\}.$$

Using the behavioral invariance of the logic, it is not hard to see that $\llbracket \varphi \rrbracket$ (we usually omit superscripts) is always \mathbb{Z} -open. Now one way to obtain nice co-Birkhoff results is to require the modal language to be expressive enough for the converse to hold as well.

DEFINITION 1.207. Let Ω be some κ -small set functor, and let $(C, \mathcal{L}_C, \Vdash)$, with $|C| = \kappa$ constitute a coalgebraic modal logic for Ω . This modal logic is called *expressive* if every open set of the C -cofree coalgebra \mathbb{Z} is of the form $\llbracket \varphi \rrbracket$ for some formula φ .

This may seem a strong requirement on a language, but expressive languages are not hard to come by.

EXAMPLE 1.208. Under some mild additional assumptions on Ω , one may show that Moss' logic of Definition 1.190 and 1.191, extended with infinite *disjunctions*, is expressive. For a proof sketch: strengthen Theorem 1.194 by proving that for any pointed Ω_C -system (\mathbb{S}, s) , there is a formula $\varphi^{\mathbb{S}, s}$ such that for all pointed Ω_C -system (\mathbb{S}', s') one has that $\mathbb{S}', s' \Vdash \varphi^{\mathbb{S}, s}$ iff $\mathbb{S}', s' \equiv_{\Omega_C} \mathbb{S}, s$. Then, given an open set U of the cofree Ω_C -coalgebra \mathbb{Z} , one may define $\varphi^U := \bigvee \{\varphi^{\mathbb{Z}, u} \mid u \in U\}$.

Now the next theorem bears witness to the tight link between modal logic and coalgebras. It is due to KURZ [76], while a very similar result was proved in GUMM & SCHRÖDER [53].

THEOREM 1.209. *Let Ω be some C -small set functor, and let $\langle C, \mathcal{L}_C, \Vdash \rangle$ constitute an expressive coalgebraic modal logic for Ω . Then a class \mathbf{K} of Ω -coalgebras is a covariety iff for some formula φ , \mathbf{K} is the class of all Ω -coalgebras \mathbb{S} such that $\mathbb{S} \Vdash \varphi$.*

Proof. Let \mathbb{Z} , with coloring $\gamma : Z \rightarrow C$, be the cofree Ω -coalgebra over C . Given a formula φ , it is a direct consequence of cofreeness and truth invariance, that for any Ω -coalgebra \mathbb{S} with C -coloring g , and for any state s in \mathbb{S} , we have

$$\mathbb{S}, g, s \Vdash \varphi \text{ iff } \tilde{g}(s) \in \llbracket \varphi \rrbracket, \quad (1.51)$$

from which one easily derives that for any \mathcal{L}_C -formula φ :

$$\text{Cov}(\llbracket \varphi \rrbracket) \text{ is the class of all } \Omega\text{-coalgebras } \mathbb{S} \text{ such that } \mathbb{S} \Vdash \varphi. \quad (1.52)$$

From (1.52) the direction ‘ \Leftarrow ’ of the Theorem is immediate. For the other direction, suppose that \mathbf{K} is a covariety. Then by expressiveness, $\text{Bhv}(\mathbf{K}) = \llbracket \varphi \rrbracket$ for some formula φ , so by (1.49) and (1.52) it follows that $\mathbf{K} = \text{Cov}(\text{Bhv}(\mathbf{K})) = \text{Cov}(\llbracket \varphi \rrbracket)$, as required. ■

Although this theorem, being formulated in terms of a fairly general notion of modal logic, may still seem to be rather abstract, it does provide a useful tool to provide more concrete results. For instance, given Example 1.208, as a corollary to Theorem 1.209 one may obtain very general modal co-Birkhoff results for Moss’ coalgebraic logic. Or, to give an even more concrete corollary of Theorem 1.209, call an (ordinary) modal frame κ -bounded for some cardinal κ if every point has less than κ successors.

COROLLARY 1.210. *A class \mathbf{K} of κ -bounded frame is (within the class of all κ -bounded frames) definable by means of infinitary modal formulas, if and only if \mathbf{K} is closed under taking generated subframes, homomorphic images and disjoint unions.*

The reader who compares the above two result to the Goldblatt-Thomasson Theorem 1.79, may be puzzled by the absence of ultrafilter extensions here. The explanation for this absence is of course that such Stone-type completions are not relevant in the presence of infinite disjunctions and conjunctions. If one takes the alternative road to co-Birkhoff theorems and starts, not from the notion of a covariety, but from a *finitary* coalgebraic logical formalism, one will find that notions like ultrafilter extensions or ultraproducts are needed in the characterization of definable classes of coalgebras. Results in this direction can be found in for instance GOLDBLATT [42, 43] or ROŞU [94].

Finally, the search for coalgebraic versions of Birkhoff’s variety theorem has received considerable attention in the coalgebraic literature, as is witnessed by many contributions in [66, 93, 20, 52, 4]. Perhaps GUMM [51] should get some special mentioning for developing an alternative coequational syntax based on equivalence classes of infinite labeled trees.

15 DUALITY OF ALGEBRA AND COALGEBRA

Various other coalgebraic topics may be of interest to modal logicians, but here we confine ourselves to a brief discussion of the duality between algebra and coalgebra.

In remark 1.152 we already observed that some of the similarities between algebra and coalgebra are based on the fact that a coalgebra $\mathbb{C} = \langle C, \gamma : C \rightarrow \Omega C \rangle$ over an

endofunctor $\Omega : \mathbf{C} \rightarrow \mathbf{C}$ can also be seen as an algebra in the *opposite* category \mathbf{C}^{op} . In fact, it is a trivial exercise to show that

$$\mathbf{Coalg}(\Omega) = (\mathbf{Alg}(\Omega^{op}))^{op}. \quad (1.53)$$

That is, the category of Ω -coalgebras is dually isomorphic to the category of algebras over the functor Ω^{op} (which acts on objects and arrows just like Ω does, the difference being that Ω^{op} is an endofunctor on \mathbf{C}^{op}).

This duality between algebras and coalgebras has been a major guideline in the development of universal coalgebra, see RUTTEN [97]. To mention just one example (many more can be found in the text): whereas *initial* algebras play an important role in universal algebra, it is the *final* objects that are relevant in coalgebra. For instance, whereas the principle of induction is based on the fact that initial algebras have no proper subalgebras, the dual *coinduction* principle boils down to the fact that final coalgebras have no proper quotients. However, it is important to realize that in (1.53) the *base* category has been dualized. This means, for instance, that systems, or **Set**-coalgebras, correspond, not so much to algebras over **Set**, as to algebras over the opposite category \mathbf{Set}^{op} (which happens to be equivalent to the category of complete and atomic Boolean algebras with complete homomorphisms). As a consequence, a general theory of systems cannot be obtained by a straightforward dualization of universal (**Set**-based) algebra. On the other hand, the fact that systems are, just like standard algebras, ‘sets with structure’, indicates that many universal algebraic concepts may apply to coalgebra by *analogy* rather than by duality — see for instance Proposition 1.179. Thus, the universal coalgebraic theory of systems is an interesting mix of dualized and non-dualized universal algebra, with, of course, some characteristics of its own.

In case that there is an informative *duality* for the base category \mathbf{C} , more can be said of (1.53). This applies for instance to the just mentioned duality of the category **Set**, but for the present purpose we prefer to focus on the category **Stone** of Stone spaces. The point is, that since **Stone** is dually equivalent to the well-known category **BA** of Boolean algebras, every endofunctor Ω on **Stone** induces an endofunctor $\Omega^* := (\cdot)_* \circ \Omega \circ (\cdot)^*$ on **BA**. It is then an immediate consequence of (1.53) that the categories $\mathbf{Coalg}(\Omega)$ and $\mathbf{Alg}(\Omega^*)$ are dually equivalent:

$$\mathbf{Coalg}(\Omega) \rightleftharpoons \mathbf{Alg}(\Omega^*). \quad (1.54)$$

For an example of this, consider the Vietoris functor \mathcal{V} of Example 1.147. Concretely, the behavior of its dual functor $\mathcal{V}^* : \mathbf{BA} \rightarrow \mathbf{BA}$ on objects is as follows. To a Boolean algebra \mathbb{B} it assigns the Boolean algebra $\mathcal{V}^*(\mathbb{B})$ freely generated by the set $\{\diamond b \mid b \in B\}$, subject to the axioms $\diamond \perp = \perp$ and $\diamond a \vee \diamond b = \diamond(a \vee b)$. Since the category $\mathbf{Coalg}(\mathcal{V})$ is dually equivalent to that of modal algebras, we thus see that the latter category, **MA**, may be represented as an algebraic category $\mathbf{Alg}(\mathcal{V}^*)$. This insight in fact provided the very first connection between modal logic and coalgebra, see ABRAMSKY [1]. Recently, the duality that (1.54) provides between algebra and coalgebra has been used to prove results on coalgebraic modal *logics*, where we now use the word ‘logic’ in the technical sense. For instance, JACOBS [62] and KUPKE, KURZ & VENEMA [74] use dualities in the style of (1.54) to prove completeness results for the multi-sorted modal logic of Definition 1.195 and 1.196. KUPKE, KURZ & PATTINSON [73] apply the above framework in order to characterize properties of arbitrary coalgebraic modal logics.

Let us finish the chapter with the observation that *both* of the fundamental dualities underlying the mathematical theory of modal logic are nontrivial instances of an

algebra/coalgebra duality. This means that the algebraic and the coalgebraic approach towards modal logic may be fruitfully operated in tandem. We believe that a thorough study of the interaction of algebra and coalgebra will provide a better understanding, not only of modal logic itself, but also of its mathematical surroundings.

A BASICS OF UNIVERSAL ALGEBRA AND CATEGORY THEORY

This section provides some technical preliminaries to this Chapter; we briefly review notation and terminology on universal algebra and category theory.

If we equip a set with a collection of finitary operations, we call the resulting structure an *algebra*; two such structures are called *similar* if their operations correspond in number and rank. In order to formalize this notion we introduce the notion of a *similarity type* as a set Σ of function symbols each of which comes with a nonnegative integer to be called its *rank* or *arity*. The arity of a function symbol f is denoted as $ar(f)$. Function symbols of rank zero are called *constants*.

The similarity type of (bounded) lattices is the set $Latt = \{\top, \perp, \wedge, \vee\}$ where \top ('top') and \perp ('bottom') are constants, and \wedge ('meet') and \vee ('join') are binary symbols. As the similarity type for Boolean algebras we take the set $Bool = \{\top, \perp, \neg, \wedge, \vee\}$ where \top , \perp , \wedge and \vee are as before, and \neg ('complementation') is a unary symbol.

A Σ -algebra is then a pair $\mathbb{A} = (A, I)$, in which the *interpretation* I assigns to each function symbol $f \in \Sigma$ an operation of arity $ar(f)$ on the *carrier* A of the algebra. Usually we write $f^{\mathbb{A}}$ rather than $I(f)$, and denote the algebra $\mathbb{A} = \langle A, I \rangle$ by $\mathbb{A} = \langle A, \{f^{\mathbb{A}} \mid f \in \Sigma\} \rangle$. As an example, let, for a set S , $\mathbb{P}(S) = \langle \mathcal{P}(S), S, \emptyset, \sim_S, \cap, \cup \rangle$ be the power set algebra, where \sim_S denotes the unary operation of complementation with respect to S . An algebra is called *trivial* if it has just one element; this completely determines the behavior of the operations.

A *homomorphism* from a Σ -algebra \mathbb{A} to a similar algebra \mathbb{B} is a map $\theta : A \rightarrow B$ that preserves Σ -structure, in the sense that, for all $f \in \Sigma$, and all a_1, \dots, a_n in A (where $n = ar(f)$):

$$\theta(f^{\mathbb{A}}(a_1, \dots, a_n)) = f^{\mathbb{B}}(\theta a_1, \dots, \theta a_n). \quad (55)$$

An injective homomorphism is called an *embedding* and a surjective one, an *epimorphism*; an *isomorphism* is a bijective homomorphism. A homomorphism with the same source as target algebra is called an *endomorphism* in general, and an *automorphism* if it is bijective.

Homomorphisms are closely related to special equivalence relations: a *congruence* on \mathbb{A} is an equivalence relation \sim satisfying, for all $f \in \Sigma$:

$$\text{if } a_1 \sim b_1 \ \& \ \dots \ \& \ a_n \sim b_n, \text{ then } f^{\mathbb{A}}(a_1, \dots, a_n) \sim f^{\mathbb{A}}(b_1, \dots, b_n), \quad (56)$$

where n is the rank of f . Given a congruence \sim on \mathbb{A} , the *quotient algebra of \mathbb{A} by \sim* is the algebra \mathbb{A}/\sim whose carrier is the set $A/\sim := \{[a] \mid a \in A\}$ of equivalence classes of A under \sim , and whose operations are defined by

$$f_{\mathbb{A}/\sim}([a_1], \dots, [a_n]) = [f_{\mathbb{A}}(a_1, \dots, a_n)].$$

(This is well-defined by (56).) The close connection between homomorphisms and congruences is formed by the fact that if $\theta : \mathbb{A} \rightarrow \mathbb{B}$ is a homomorphism, its *kernel*

$\ker(\theta) := \{(a, b) \in A \times A \mid \theta(a) = \theta(b)\}$ is a congruence on \mathbb{A} , while, on the other hand, for any congruence \sim on \mathbb{A} , the associated *natural map* ν_{\sim} taking an element $a \in A$ to its equivalence class $[a]$ is a surjective homomorphism from \mathbb{A} onto \mathbb{A}/\sim .

The set of congruences $Cg\mathbb{A}$ of an algebra \mathbb{A} forms in fact a complete lattice under the subset ordering; this lattice is denoted as $\mathbb{C}g(\mathbb{A})$; the meet operation of this lattice is simply their intersection, while the join of two congruences is given by $\Theta_1 \vee \Theta_2 = \Theta_1 \cup (\Theta_1 \circ \Theta_2) \cup (\Theta_1 \circ \Theta_2 \circ \Theta_1) \cup \dots$.

A Σ -algebra \mathbb{A} is a *subalgebra* of a Σ -algebra \mathbb{B} if $A \subseteq B$ and for all $f \in \Sigma$, the operation $f^{\mathbb{A}}$ coincides with the restriction of $f^{\mathbb{B}}$ to A . The *direct product* $\mathbb{A} = \prod_{i \in I} \mathbb{A}_i$ of a family of Σ -algebras is an algebra with carrier $\prod_{i \in I} A_i$ and such that for $f \in \Sigma$ and $a_1, \dots, a_n \in \prod_{i \in I} A_i$:

$$f^{\mathbb{A}}(a_1, \dots, a_n)(i) := f^{\mathbb{A}_i}(a_1(i), \dots, a_n(i))$$

We assume familiarity with the notions of ultraproduct and ultrapower.

Given a class \mathbf{K} of algebras, we let $\mathbf{H}(\mathbf{K})$ denote the class of homomorphic images of algebras in \mathbf{K} ; $\mathbf{S}(\mathbf{K})$ is the class of isomorphic copies of subalgebras of algebras in \mathbf{K} , and likewise definitions applies for the class operations \mathbf{P} (products), \mathbf{Pu} (ultraproducts) and \mathbf{Pw} (ultrapowers).

A class of algebras is called a *variety* if it is closed under taking subalgebras, homomorphisms, and products; the smallest variety containing a class \mathbf{K} is called the *variety generated by \mathbf{K}* , notation: $\mathbf{Var}(\mathbf{K})$. Using inequalities like $\mathbf{SH} \leq \mathbf{HS}$ (meaning that, for any class of algebras \mathbf{K} , $\mathbf{SH}(\mathbf{K})$ is a subclass of $\mathbf{HS}(\mathbf{K})$), together with the idempotence of the class operations \mathbf{S} , \mathbf{H} and \mathbf{P} , one can prove Tarski's Theorem stating that

$$\mathbf{Var}(\mathbf{K}) = \mathbf{HSP}(\mathbf{K}) \tag{57}$$

for any class of algebras \mathbf{K} .

Given a similarity type Σ and a set of variables X , we define the set $Ter_{\Sigma}(X)$ of Σ -terms over X by a straightforward induction: it is the smallest including X which contains $f(t_1, \dots, t_n)$ whenever it contains t_1, \dots, t_n and $f \in \Sigma$ is a function symbol of rank n . (In particular, $Ter_{\Sigma}(X)$ contains all constants in Σ .) In this chapter we adopt the convention that unless explicitly indicated otherwise, X denotes a countably infinite set of variables; we often omit explicit reference to X , writing for instance Ter_{Σ} rather than $Ter_{\Sigma}(X)$, etc. Also, writing $s(x_1, \dots, x_n)$ for a term s , we indicate that the variables occurring in s are among x_1, \dots, x_n .

Given an assignment α of a set X of variables to (the carrier A of) an algebra \mathbb{A} , we inductively define the *meaning* $\tilde{\alpha}(s)$ of a term s as follows:

$$\begin{aligned} \tilde{\alpha}(x) &= \alpha(x) \\ \tilde{\alpha}(f(t_1, \dots, t_n)) &= f^{\mathbb{A}}(\tilde{\alpha}(t_1), \dots, \tilde{\alpha}(t_n)). \end{aligned} \tag{58}$$

Thus any term $s(x_1, \dots, x_n)$ induces an n -ary *term function* $s^{\mathbb{A}}$ on \mathbb{A} , given by $s^{\mathbb{A}}(a_1, \dots, a_n) = \tilde{\alpha}(s)$, where α is any assignment mapping each x_i to a_i . (Of course, $s^{\mathbb{A}}$ can also be given an inductive definition.)

Using the close resemblance between the second clause of (58) and (55), we can turn the meaning function into a real homomorphism by imposing Σ -algebra structure on the set $Ter_{\Sigma}(X)$, obtaining the *term algebra* $\mathbf{Ter}_{\Sigma}(X)$. The idea is to interpret the function

symbol $f \in \Sigma$ as follows:

$$f^{\text{Ter}_\Sigma(X)} : (t_1, \dots, t_n) \mapsto f(t_1, \dots, t_n).$$

Elaborating on this perspective, let \mathbf{K} be a class of Σ -algebras, and \mathbb{F} a Σ -algebra generated by a set $X \subseteq F$. Suppose that for every \mathbb{A} in \mathbf{K} and every map $\alpha : X \rightarrow A$ there is a homomorphism $\tilde{\alpha} : \mathbb{F} \rightarrow \mathbb{A}$ extending α . Then we say that \mathbb{F} has the *universal mapping property for \mathbf{K} over X* , or that \mathbb{F} is *free for \mathbf{K} over X* . The identities of (58) thus reveal that $\text{Ter}_\Sigma(X)$ is free over X for the class of *all* Σ -algebras; for this reason it is often referred to as the *absolutely free algebra over X* .

Free algebras have a number of important properties of which we mention the following:

- every algebra in \mathbf{K} is a homomorphic image of a free algebra over an appropriately large set of generators;
- all free algebras for \mathbf{K} belongs to the class $\text{SP}(\mathbf{K})$;
- if \mathbb{F} and \mathbb{F}' are free for \mathbf{K} over the generator sets X and X' , respectively, and X and X' have the same cardinality, then \mathbb{F} and \mathbb{F}' are isomorphic.

Universal algebra may on the one hand be seen as generalizing the study of individual classes of algebras such as groups, fields, or lattices. On the other hand we may consider it as a rather special branch of model theory in which one is interested in structures for a language without relation symbols. The standard language for talking about such structures is *equational*.

An *equation* is nothing but a pair (s, t) of terms, always denoted as $s \approx t$. The equation $s \approx t$ (with $s, t \in \text{Ter}_\Sigma(X)$) is *true* or *holds* in the algebra \mathbb{A} *under the assignment* $\alpha : X \rightarrow A$, notation: $\mathbb{A} \models_\alpha s \approx t$ if s and t obtain the same meaning in \mathbb{A} under α , that is, if $\tilde{\alpha}(s) = \tilde{\alpha}(t)$. An equation $s \approx t$ *holds in* the algebra \mathbb{A} , or, equivalently, the algebra \mathbb{A} *satisfies* the equation $s \approx t$, notation: $\mathbb{A} \models s \approx t$, if $\mathbb{A} \models_\alpha s \approx t$ for every assignment α .

The relation \models induces a Galois connection between sets of formulas and classes of algebras; the polarities of this connection are given as the maps *Equ* and *Mod*, where $\text{Equ}(\mathbf{K})$ is the set of all equations that hold in \mathbf{K} , and $\text{Mod}(E)$ denotes the class of algebras that satisfy every equation in E . The classes of algebras that are stable under this connection, that is, the classes \mathbf{K} of the form $\text{Mod}(E)$ for some set E of equations, are called *equational classes*. An important result by Birkhoff states that this notion coincides with that of a variety, and that for any class \mathbf{K} of algebras it holds that

$$\text{Mod}(\text{Equ}(\mathbf{K})) = \text{Var}(\mathbf{K}). \quad (59)$$

The relation

$$s \equiv_{\mathbf{K}} t : \iff \mathbf{K} \models s \approx t$$

corresponding to the set $\text{Equ}(\mathbf{K})$ is in fact a *congruence* on the term algebra Ter_Σ . The algebra $\text{Ter}_\Sigma(X) / \equiv_{\mathbf{K}}$ has the universal mapping property for \mathbf{K} over $[X]$ (the set of equivalence classes of X under $\equiv_{\mathbf{K}}$), which, together with the third fact on free algebras listed above, explains why we call it the *free algebra for \mathbf{K} over $[X]$* .

A *category* \mathbf{C} consists of a class $\text{Ob}(\mathbf{C})$ of *objects*, and for each pair of objects A, B , a family $\mathbf{C}(A, B)$ of *arrows*. If f belongs to the latter set, we write $f : A \rightarrow B$, and call A

the *domain* and B the *codomain* of the arrow. The collection of arrows is endowed with some algebraic structure: for every object A of \mathbf{C} there is an arrow $id_A : A \rightarrow A$, and every pair $f : A \rightarrow B$, $g : B \rightarrow C$ can be uniquely *composed* to an arrow $g \circ f : A \rightarrow C$. These operations are supposed to satisfy the associative law for composition, while the appropriate identity arrows are left- and right neutral elements. An arrow $f : A \rightarrow B$ is an *iso* if it has an *inverse*, that is, an arrow $g : B \rightarrow A$ such that $f \circ g = id_B$ and $g \circ f = id_A$. Examples of categories are \mathbf{Set} , the class of sets with functions, and, for every similarity type Σ , the class $\mathbf{Alg}(\Sigma)$ of Σ -algebras, with homomorphisms as arrows. The *opposite* category \mathbf{C}^{op} of a given category \mathbf{C} has the same objects as \mathbf{C} , while $\mathbf{C}^{op}(A, B) = \mathbf{C}(B, A)$ for all objects A, B from \mathbf{C} , and the operations on arrows are defined in the obvious way.

An object X is *initial* in a category \mathbf{C} if for every object A in \mathbf{C} there is a unique arrow $\alpha : X \rightarrow A$, and *final* if for all A there is a unique $\alpha : A \rightarrow X$. In \mathbf{Set} , the empty set is initial, and the final objects are precisely the singletons. A *product* of two objects A_0 and A_1 in a category \mathbf{C} consists of a triple $(A, \alpha_0 : A \rightarrow A_0, \alpha_1 : A \rightarrow A_1)$, such that for every triple $(A', \alpha'_0 : A' \rightarrow A_0, \alpha'_1 : A' \rightarrow A_1)$ there is a unique arrow $f : A' \rightarrow A$ such that $\alpha_i \circ f = \alpha'_i$ for both i . *Coproducts* of A_0 and A_1 are defined dually as triples $(A, \alpha_0 : A_0 \rightarrow A, \alpha_1 : A_1 \rightarrow A)$, such that for every triple $(A', \alpha'_0 : A_0 \rightarrow A', \alpha'_1 : A_1 \rightarrow A')$ there is a unique arrow $f : A \rightarrow A'$ such that $f \circ \alpha_i = \alpha'_i$ for each i . The category \mathbf{Set} has both products and coproducts — that is, every pair (S_0, S_1) of sets has both a product (for which we may take the cartesian product $S_0 \times S_1$ together with the two projection functions $\pi_i : S_0 \times S_1 \rightarrow S_i$), and a coproduct (for which we may take the disjoint union $S_0 \uplus S_1 = S_0 \times \{0\} \cup S_1 \times \{1\}$ together with the coproduct maps κ_0 and κ_1 given by $\kappa_i(s) = (s, i)$).

A *functor* $\Omega : \mathbf{C} \rightarrow \mathbf{D}$ from a category \mathbf{C} to a category \mathbf{D} consists of an operation mapping objects and arrows of \mathbf{C} to objects and arrows of \mathbf{D} , respectively, in such a way that $\Omega f : \Omega A \rightarrow \Omega B$ if $f : A \rightarrow B$, $\Omega(id_A) = id_{\Omega A}$ and $\Omega(g \circ f) = (\Omega g) \circ (\Omega f)$ for all objects and arrows involved. A functor $\Omega : \mathbf{C} \rightarrow \mathbf{D}^{op}$ is sometimes called a *contravariant* functor from \mathbf{C} to \mathbf{D} . An *endofunctor* on \mathbf{C} is a functor $\Omega : \mathbf{C} \rightarrow \mathbf{C}$.

As examples we consider the following *set functors* (that is, endofunctors on \mathbf{Set}): (i) for a fixed set C , the *constant* functor mapping all sets to C and all arrows to id_C ; this functor is denoted as C , (ii) the *power set functor* \mathcal{P} , which maps any set S to its power set $\mathcal{P}S$, and any map $f : S \rightarrow S'$ to the map $\mathcal{P}f : \mathcal{P}S \rightarrow \mathcal{P}S'$ given by $\mathcal{P}f : X \mapsto \{fx \mid x \in X\}$, and (iii) for every cardinal κ , the variant \mathcal{P}_κ of the power set functor, which maps any set S to the collection $\mathcal{P}_\kappa S := \{X \subseteq S \mid \kappa > |X|\}$, and agrees with \mathcal{P} on the arrows for which is defined. Furthermore, given two functors Ω_0 and Ω_1 , their *product functor* $\Omega_0 \times \Omega_1$ is given (on objects) by $(\Omega_0 \times \Omega_1)S := \Omega_0 S \times \Omega_1 S$, while for $f : S \rightarrow S'$, the map $(\Omega_0 \times \Omega_1)f$ is given as $((\Omega_0 \times \Omega_1)f)(\sigma_0, \sigma_1) := ((\Omega_0 f)(\sigma_0), (\Omega_1 f)(\sigma_1))$. The *coproduct functor* is defined similarly. Finally, every category \mathbf{C} admits the *identity functor* $\mathcal{I}_\mathbf{C} : \mathbf{C} \rightarrow \mathbf{C}$ which is the identity on both objects and arrows of \mathbf{C} .

Let \mathbf{C} and \mathbf{D} be two categories, and let Ω and Ψ be two functors from \mathbf{C} to \mathbf{D} . A *natural transformation* τ from Ω to Ψ , notation $\tau : \Omega \Rightarrow \Psi$, consists of \mathbf{D} -arrows $\tau_A : \Omega A \rightarrow \Psi A$ such that $\tau_B \circ \Omega f = \Psi f \circ \tau_A$ for each $f : A \rightarrow B$ in \mathbf{C} .

Finally, let $\Omega : \mathbf{C} \rightarrow \mathbf{D}$ and $\Psi : \mathbf{D} \rightarrow \mathbf{C}$ be two functors linking the categories \mathbf{C} and \mathbf{D} . Ω and Ψ constitute an *equivalence* between \mathbf{C} and \mathbf{D} if their compositions are naturally isomorphic to the identity functors, that is, if there are natural transformations $\sigma : \mathcal{I}_\mathbf{C} \Rightarrow \Psi\Omega$ and $\tau : \mathcal{I}_\mathbf{D} \Rightarrow \Omega\Psi$ such that all arrows $\sigma_A : A \rightarrow \Psi\Omega A$ and $\tau_B : B \rightarrow \Omega\Psi B$ are isos. If such Ω and Ψ exist, then the categories \mathbf{C} and \mathbf{D} are called *equivalent*; if Ω

and Ψ are in fact each other's inverse (both on maps and on arrows) then \mathcal{C} and \mathcal{D} are *isomorphic*. If Ω and Ψ form a *dual equivalence* between the categories \mathcal{C} and \mathcal{D} , that is, an equivalence between the categories \mathcal{C} and \mathcal{D}^{op} , then we say that the categories are *dual* or *dually equivalent* to each other.

INDEX

- additive, 11
 - completely, 20
- algebraizes, 10
- atom, 20
- atom structure, 21
- atomic, 20

- BAO, 11
- base category, 68
- behavioral equivalence, 74
- bisimilar, 74
- bisimulation
 - between coalgebras, 74
 - equivalence, 74
- Boolean algebra
 - expanded, 11
 - with operators, 11
- box, 4

- canonical
 - class, 38
 - equation, 38
 - extension, 25, 47
 - formula, 38
 - normal modal logic, 38
- Cm, 9
- coalgebra, 65
- cofree, 91
- color, 90
- coloring, 90
- compact, 46
- complete
 - Boolean algebra, 20
 - homomorphism, 21
 - variety, 35
- completion, 46
 - MacNeille, 56
- complex algebra, 9
- congruence distributive, 14
- covariety, 81
- Cst, 24

- dense, 46
- diamond, 4
 - compound, 4
 - induced, 4
- discriminator
 - unary, 60
- discriminator term, 60
- duality
 - complex, 22

- essential, 17
- expanding, 53

- field of sets, 26
- filter, 15
 - modal/open, 15
 - principal, 15
 - proper, 15
- final coalgebra, 70
- final sequence, 72
- formula algebra, 8

- general frame, 27
 - descriptive, 27
- global modality, 60, 61

- homomorphic image
 - coalgebraic, 80
- homomorphism
 - between coalgebras, 67

- interpolation, local, 42

- Lindenbaum-Tarski algebra, 12

- modal logic
 - coalgebraic, 93
- monotone, 11
- multiplicative, 11

- neighborhood frame, 6
- normal, 11

normal extension, 34

open set
of coalgebra, 80

operator, 11
complete, 20

perfect, 20

pointed, 65

polynomial functor, 66
finitary Kripke, 67
Kripke, 67

pullback, 78
weak, 78

radical, 17

relation lifting, 76

root, 6

Sahlqvist formula, 39

similarity type, basic modal, 4

similarity type, modal, 4

simple, 16

small, 71

smooth, 49

stable, 53

standard, 71

Stone space, 26

Str, 9

subcoalgebra, 80

subdirectly irreducible, 16

sum, 81
conjunct, 82

superamalgamation, 42

system, 65

tense algebra, 58

tense logic, 58

ultrafilter, 23

ultrafilter frame, 24

up-directed, 50

up-set, 15

Vietoris construction, 67

BIBLIOGRAPHY

- [1] S. Abramsky. A Cook's Tour of the Finitary Non-Well-Founded Sets, 1988. Invited Lecture at BCTCS.
- [2] P. Aczel. *Non-well-founded sets*, volume 14 of *CSLI Lecture Notes*. CSLI, 1988.
- [3] P. Aczel and N. Mendler. A final coalgebra theorem. In D.H. Pitt, A. Poigné, and D.E. Rydeheard, editors, *Category Theory and Computer Science*, volume 389 of *Lecture Notes in Computer Science*. Springer, 1989.
- [4] J. Adámek, editor. *Coalgebraic Methods in Computer Science (CMCS'04)*, volume ?? of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2004.
- [5] J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theoretical Computer Science*, 150:57–75, 1995.
- [6] J. Adámek and H.-E. Porst. From varieties of algebras to covarieties of coalgebras. In Corradini et al. [20].
- [7] H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
- [8] A. Baltag. A logic for coalgebraic simulation. In Reichel [93].
- [9] M. Barr. Terminal coalgebras in well-founded set theory. *Theoretical Computer Science*, 114(2):299–315, 1993. (Corrigendum in *Theor.Comp.Sci.* 124:189–192,1994).
- [10] M. Barr. Relational algebras. In S. Mac Lane, editor, *Reports of the Midwest Category Seminar IV*, volume 137 of *Lecture Notes in Mathematics*. Springer Verlag, 1997.
- [11] J. Barwise and L. Moss. *Vicious Circles*, volume 60 of *CSLI Lecture Notes*. CSLI Publications, 1996.
- [12] G. Birkhoff. *Lattice Theory*. Number 25 in Colloquium Publications. American Mathematical Society, Providence, 1967. Third edition.
- [13] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [14] W. Blok. *Varieties of Interior Algebras*. PhD thesis, Mathematisch Instituut, University of Amsterdam, 1976.
- [15] W. Blok. The lattice of modal algebras: An algebraic investigation. *Journal of Symbolic Logic*, 45:221–236, 1980.
- [16] W. Blok and D. Pigozzi. Algebraizable logics. *Memoirs of the American Mathematical Society*, 77(396), 1989.
- [17] S. Burris and H.P. Sankappanavar. *A Course in Universal Algebra*. Graduate Texts in Mathematics. Springer, 1981.
- [18] W. Buszkowski. Embedding Boolean structures into atomic Boolean structures. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 32:227–228, 1986.
- [19] A. Carboni, G. Kelly, and R. Wood. A 2-categorical approach to change of base and geometric morphisms I. Technical Report 90-1, Department of Pure Mathematics, University of Sydney, 1990.
- [20] A. Corradini, M. Lenisa, and U. Montanari, editors. *Coalgebraic Methods in Computer Science (CMCS'01)*, volume 44 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2001.
- [21] J. Czelakowski. *Protoalgebraic Logic*. Trends in Logic. Kluwer Academic Press, 2001.
- [22] Paul Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.
- [23] L. Esakia. Topological Kripke models. *Soviet Mathematics Doklady*, 15:147–151, 1974.
- [24] L. Esakia. To the theory of modal and superintuitionistic systems. In V. Smirnov, editor, *Proceedings of the USSR Symposium on Logical Inference*, pages 147–172. Nauka, Moscow, 1979. (In Russian).
- [25] K. Fine. Some connections between elementary and modal logic. In Kanger [71].
- [26] J.M. Font and R. Jansana. *A general algebraic semantics for sentential logics*. Lecture Notes in Logic. Springer, 1996.
- [27] M. Gehrke. The order structure of Stone spaces and the T_D -separation axiom. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 37:5–15, 1991.
- [28] M. Gehrke and J. Harding. Bounded lattice expansions. *Journal of Algebra*, 238:345–371, 2001.
- [29] M. Gehrke, J. Harding, and Y. Venema. MacNeille completions and canonical extensions. *Transactions of the American Mathematical Society*, to appear.
- [30] M. Gehrke and B. Jónsson. Bounded distributive lattices with operators. *Mathematica Japonica*, 40:207–215, 1994.
- [31] M. Gehrke and B. Jónsson. Monotone bounded distributive lattices expansions. *Mathematica Japonica*, 52:197–213, 2000.
- [32] M. Gehrke and B. Jónsson. Bounded distributive lattice expansions. *Mathematica Scandinavica*, 94:3–45, 2004.
- [33] S. Ghilardi. An algebraic theory of normal forms. *Annals of Pure and Applied Logic*, 71:189–245, 1995.
- [34] S. Ghilardi and G. Meloni. Constructive canonicity in non-classical logics. *Annals of Pure and Applied Logic*, 86:1–32, 1997.

- [35] S. Givant. Universal classes of simple relation algebras. *The Journal of Symbolic Logic*, 64:575–589, 1999.
- [36] S. Givant and Y. Venema. The preservations of Sahlqvist equations in completions of Boolean algebras with operators. *Algebra Universalis*, 41:47–84, 1999.
- [37] R. Goldblatt. Metamathematics of modal logic I. *Reports on Mathematical Logic*, 6:41–78, 1976.
- [38] R. Goldblatt. Metamathematics of modal logic II. *Reports on Mathematical Logic*, 7:21–52, 1976.
- [39] R. Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, 38:173–241, 1989.
- [40] R. Goldblatt. Elementary generation and canonicity for varieties of Boolean algebras with operators. *Algebra Universalis*, 34:551–607, 1995.
- [41] R. Goldblatt. Algebraic polymodal logic: a survey. *Logic Journal of the IGPL*, 8:393–450, 2000.
- [42] R. Goldblatt. A calculus of terms for coalgebras of polynomial functors. In Corradini et al. [20].
- [43] R. Goldblatt. What is the coalgebraic analogue of Birkhoff’s variety theorem. *Theoretical Computer Science*, 266:853–886, 2001.
- [44] R. Goldblatt. Mathematical modal logic: a view of its evolution. *Journal of Applied Logic*, 1:309–392, 2003. An updated version is available from the web site of the author.
- [45] R. Goldblatt. Final coalgebras and the hennessy-milner property. *Annals of Pure and Applied Logic*, to appear.
- [46] R. Goldblatt, I. Hodkinson, and Y. Venema. Erdős graphs resolve Fine’s canonicity problem. *Bulletin of Symbolic Logic*, 10:186–208, 2004.
- [47] R. Goldblatt and S. Thomason. Axiomatic classes in propositional modal logic. In J. Crossley, editor, *Algebra and Logic*, pages 163–173. Springer, 1974.
- [48] V. Goranko and S. Passy. Using the universal modality: Gains and questions. *Journal of Logic and Computation*, 2:5–30, 1992.
- [49] V. Goranko and D. Vakarelov. Sahlqvist formulas unleashed in polyadic modal languages. In F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev, editors, *Advances in Modal Logic, Volume 3*, pages 221–240. World Scientific, 2002.
- [50] H. Peter Gumm. Elements of the general theory of coalgebras, 1999. Lecture Notes for LUATCS’99, Rand Africaans University, Johannesburg; available from <http://www.mathematik.uni-marburg.de/~gumm/>.
- [51] H. Peter Gumm. Birkhoff’s variety theorem for coalgebras. *Contributions to General Algebra*, 13:159–173, 2000.
- [52] H. Peter Gumm, editor. *Coalgebraic Methods in Computer Science (CMCS’03)*, volume 82 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2003.
- [53] H. Peter Gumm and T. Schröder. Covarieties and complete covarieties. *Theoretical Computer Science*, 260:71–86, 2001.
- [54] H. Peter Gumm and T. Schröder. Coalgebras of bounded type. *Mathematical Structures in Computer Science*, 12:565–578, 2002.
- [55] H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. In Adámek [4].
- [56] J. Harding. Canonical completions of lattices and ortholattices. *Tatra Mountains Mathematical Publications*, 15:85–96, 1998.
- [57] L. Henkin, J.D. Monk, and A. Tarski. *Cylindric Algebras. Part 1. Part 2*. North-Holland Publishing Company, Amsterdam, 1971, 1985.
- [58] R. Hirsch and I.M. Hodkinson. *Relation Algebras by Games*. Number 147 in Studies in Logic. Elsevier, Amsterdam, 2002.
- [59] E. Hoogland. *Definability and Interpolation. Model-theoretic investigations*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2001.
- [60] T. Ihringer. *Allgemeine Algebra. Mit einem Anhang über Universelle Coalgebra von H. P. Gumm*, volume 10 of *Berliner Studienreihe zur Mathematik*. Heldermann Verlag, 2003.
- [61] B. Jacobs. Objects and classes, coalgebraically. In B. Freitag, C. Jones, C. Lengauer, and H.-J. Schek, editors, *Object-Oriented with Parallelism and Persistence*, volume 44, pages 83–103. Kluwer Academic Publishers, 1997.
- [62] B. Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoretical Informatics and Applications*, 35(1):31–59, 2001.
- [63] B. Jacobs. Exercises in coalgebraic specification. In R. Backhouse, R. Crole, and J. Gibbons, editors, *Algebraic and Coalgebraic Methods in the Mathematics of Program Construction*, volume 44 of *Lecture Notes in Computer Science*, pages 237–280. Springer, 2002.
- [64] B. Jacobs. The temporal logic of coalgebras via Galois algebras. *Mathematical Structure in Computer Science*, 12:875–903, 2002.
- [65] B. Jacobs and J. Rutten. A tutorial on (co)algebras and (co)induction. *Bulletin of the European Association for Theoretical Computer Science*, 62:222–259, 1997.
- [66] B. Jacobs and J. Rutten, editors. *Coalgebraic Methods in Computer Science (CMCS’00)*, volume 19 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1999.

- [67] P. Jipsen. Discriminator varieties of Boolean algebras with residuated operators. In C. Rauszer, editor, *Algebraic Methods in Logic and Computer Science*, volume 28 of *Banach Center Publications*, pages 239–252. Polish Academy of Sciences, 1993.
- [68] P.J. Johnstone. *Stone Spaces*, volume 3 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1982.
- [69] B. Jónsson. On the canonicity of Sahlqvist identities. *Studia Logica*, 4:473–491, 1994.
- [70] B. Jónsson and A. Tarski. Boolean algebras with operators, Part I. *American Journal of Mathematics*, 73:891–939, 1952.
- [71] S. Kanger, editor. *Proceedings of the Third Scandinavian Logic Symposium. Uppsala 1973*. North-Holland Publishing Company, 1975.
- [72] M. Kracht. *Tools and Techniques in Modal Logic*. Number 142 in *Studies in Logic*. Elsevier, Amsterdam, 1999.
- [73] C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. In Adámek [4].
- [74] C. Kupke, A. Kurz, and Y. Venema. Stone coalgebras. *Theoretical Computer Science*, 327:109–134, 2004.
- [75] A. Kurz. Coalgebras and modal logic, 1999. Lecture Notes for ESSLLI'01; available from <http://www.folli.uva.nl/CD/2001/courses/readers/K40.pdf>.
- [76] A. Kurz. A co-variety-theorem for modal logic. In M. Zakharyashev, K. Segerberg, M. de Rijke, and H. Wansing, editors, *Advances in Modal Logic, Volume 2*. CSLI Publications, 2000.
- [77] A. Kurz. Specifying coalgebras with modal logic. *Theoretical Computer Science*, 260:119–138, 2001.
- [78] A. Kurz and D. Pattinson. Coalgebraic modal logic of finite rank. *Mathematical Structures in Computer Science*, to appear. Also available as CWI Technical Report SEN-R0222, 2002.
- [79] J. Lambek. A fixpoint theorem for complete categories. *Mathematische Zeitschrift*, 103:151–161, 1968.
- [80] E. Lemmon. Algebraic semantics for modal logics, Parts I & II. *Journal of Symbolic Logic*, pages 46–65 & 191–218, 1966.
- [81] T. Litak. On notions of completeness weaker than Kripke incompleteness. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing, editors, *Advances in Modal Logic, Volume 5*. King's College London Publications, to appear.
- [82] H. MacColl. *Symbolic Logic and its Applications*. Longmans, Green, and Co., London, 1906.
- [83] L. Maksimova. Interpolation theorems in modal logic and amalgable varieties of topological Boolean algebras. *Algebra and Logic*, 18:348–370, 1979.
- [84] M. Marx and Y. Venema. *Multidimensional Modal Logic*, volume 4 of *Applied Logic Series*. Kluwer Academic Publishers, 1997.
- [85] J. Monk. Completions of Boolean algebras with operators. *Mathematische Nachrichten*, 46:47–55, 1970.
- [86] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *Ann.P.Appl.Log.* 99:241–259, 1999).
- [87] I. Németi. Cylindric-relativized set algebras have strong amalgamation. *Journal of Symbolic Logic*, 50:689–700, 1985.
- [88] D. Pattinson. Semantical principles in the modal logic of coalgebras. In H. Reichel and A. Ferreira, editors, *Proc. 18th International Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, volume 2010 of *LNCS*, Berlin, 2001. Springer.
- [89] D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309:177–193, 2003.
- [90] D. Pattinson. An introduction to the theory of coalgebras, 2003. Notes accompanying a course at NASSLLI 2003; available from <http://www.indiana.edu/~nasslli/>.
- [91] W. Rautenberg. *Klassische und nichtklassische Aussagenlogik*. Vieweg & Sohn, 1979.
- [92] H. Reichel. An approach to object semantics based on terminal coalgebras. *Mathematical Structures in Computer Science*, 5:129–152, 1995.
- [93] H. Reichel, editor. *Coalgebraic Methods in Computer Science (CMCS'00)*, volume 33 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2000.
- [94] G. Roşu. Equational axiomatizability for coalgebra. *Theoretical Computer Science*, 260:229–247, 2001.
- [95] M. Röbiger. Coalgebras and modal logic. In Reichel [93], pages 299–320.
- [96] J. Rutten. Automata and coinduction (an exercise in coalgebra). In D. Sangiorigi and R. de Simone, editors, *Proceedings of CONCUR '98*, volume 1466 of *Lecture Notes in Computer Science*, pages 194–218. Springer, 1998.
- [97] J. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
- [98] H. Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. In Kanger [71], pages 110–143.
- [99] G. Sambin. Subdirectly irreducible modal algebras and initial frames. *Studia Logica*, 62:269–282, 1999.

- [100] G. Sambin and V. Vaccaro. A topological proof of Sahlqvist's theorem. *Journal of Symbolic Logic*, 54:992–999, 1989.
- [101] M. Stone. The theory of representations for Boolean algebras. *Transactions of the American Mathematical Society*, 40:37–111, 1936.
- [102] S. Thomason. Semantic analysis of tense logics. *Journal of Symbolic Logic*, 37:150–158, 1972.
- [103] S. Thomason. Categories of frames for modal logics. *Journal of Symbolic Logic*, 40:439–442, 1975.
- [104] V. Trnková. Relational automata in a category and theory of languages. In M. Karpinski, editor, *Fundamentals of computation theory: proceedings of the 1977 international FCT-conference*, volume 56 of *Lecture Notes in Computer Science*, pages 340–355. Springer, 1977.
- [105] V. Trnková. General theory of relational automata. *Fundamenta Informaticae*, 3:189–234, 1979.
- [106] Y. Venema. Atomless varieties. *Journal of Symbolic Logic*, 68:607–614, 2003.
- [107] Y. Venema. Automata and fixed point logics for coalgebras. In Adámek [4].
- [108] Y. Venema. A dual characterization of subdirectly irreducible BAOS. *Studia Logica*, 77:105–115, 2004.
- [109] J. Worrell. Terminal sequences for accessible endofunctors. In Jacobs and Rutten [66], pages 39–53.