Modal Logic and the Vietoris Functor

Venema, Y.; Vosmaer, J.

DOI
10.1007/978-94-017-8860-1_6

Publication date
2014

Document Version
Submitted manuscript

Published in
Leo Esakia on Duality in Modal and Intuitionistic Logics

Modal logic and the Vietoris functor

Yde Venema and Jacob Vosmaer

Dedicated to the memory of Leo Esakia, who was and will remain a great source of inspiration, both as a logician and as a person

Abstract In [17], Esakia uses the Vietoris topology to give a coalgebra-flavored definition of topological Kripke frames, thus relating the *Vietoris topology*, *modal logic* and *coalgebra*. In this chapter, we sketch some of the thematically related mathematical developments that followed. Specifically, we look at Stone duality for the Vietoris hyperspace and the Vietoris powerlocale, and at recent work combining coalgebraic modal logic and the Vietoris functor.

1 Introduction

The *Vietoris hyperspace* is a topological construction on compact Hausdorff spaces, which was introduced in 1922 by Leopold Vietoris [42] as a generalization of the Hausdorff metric. Given a compact Hausdorff space X, one can obtain the Vietoris topology on $\mathrm{K}X$, the set of compact subsets of X, by generating a topology from a basis, consisting of all sets of the form

$$\nabla\{U_1,\dots,U_n\} := \{F \in \mathrm{K}X \mid F \subseteq \bigcup_{i=1}^{n} U_i \text{ and } \forall i \leq n, F \cap U_i \neq \emptyset\},$$

where $\{U_1,\dots,U_n\}$ ranges over the collection of finite sets of opens in X. Alternatively, one can generate the Vietoris topology from a subbasis, consisting of open sets

$$[U] := \{F \in \mathrm{K}X \mid F \subseteq U\} \text{ and } \langle U\rangle := \{F \in \mathrm{K}X \mid F \cap U \neq \emptyset\},$$

Yde Venema · Jacob Vosmaer
Institute for Logic, Language and Computation, University of Amsterdam
e-mail: Y.Venema@uva.nl, contact@jacobvosmaer.nl

where U ranges over the open subsets of X. This construction can be seen as a functor on the category of compact Hausdorff spaces and continuous functions: if $f\colon X \to Y$ is continuous, then so is $\mathrm{V}f\colon \mathrm{V}X \to \mathrm{V}Y$, where $\mathrm{V}f\colon F \mapsto f[F]$ is taking forward images.

With his 1974 paper [17], Leo Esakia was the first to point out that there is a connection between the Vietoris topology and *modal logic*: he defines his *topological Kripke frames* using the Vietoris topology, and links these structures to modal algebras via a Stone-type duality. In fact, from a modern viewpoint, Esakia's topological Kripke frames are *coalgebras*, and his duality is a key example of a nontrivial algebra/coalgebra duality. This chapter will explore some of the further connections within this picture – comprising the Vietoris topology, modal logic and coalgebra – that have since been discovered in the mathematical landscape. In particular, we will look at how modal logic can help one to understand the Vietoris construction.

Generally, modal logicians think of topological structures and Stone-type dualities as *tools* for understanding modal logics; tools that are of interest primarily or at least partly because the standard Kripke semantics is too coarse a tool for bringing out subtle differences between modal logics. In this paper, we take an opposite view, namely of modal logic, and *coalgebraic logic*, as a tool for understanding the Vietoris topology. In §2 we consider the basic case. We discuss the use of Boolean modal logic for describing the Stone dual of the Vietoris functor on Stone spaces, and the relation of this idea to coalgebra and Esakia's work [17]. In §3 we see that the relation between the Vietoris construction and modal logic generalizes from Stone spaces to compact Hausdorff spaces. This takes us into locale theory, where the modal logic approach has been used to generalize the Vietoris construction even beyond compact Hausdorff spaces, to stably locally compact spaces. As examples of situations where we find spaces which are not compact Hausdorff, we consider distributive lattices and algebraic domains. Finally, in §4, we investigate a recent perspective on the Vietoris construction, namely, via the *nabla* modality and *Moss' coalgebraic logic*. This leads to a new presentation of the Vietoris construction in locale theory, as well as a new direction of generalization.

This chapter can serve as a first guide through the mathematical landscape that we just sketched, by providing a tour along some well-known results, and relating these to new work. Throughout, we have assumed that the reader has at least some basic familiarity with the following subjects: propositional modal logic and its Kripke semantics, basic general topology and category theory, Stone duality for Boolean algebras, and frames and locales as used in point-free topology. At the end of each (sub)section we provide some historical notes and pointers to the literature (in particular we provide references for facts that are mentioned without proof in the main text).

2 The main ideas in the Boolean case

In his 1974 paper [17], Esakia presented duality results for *topological Kripke frames* and modal algebras by building on Stone duality for Boolean algebras. Topological Kripke frames, more commonly known as descriptive general frames, play an important role in the model theory of modal logic, because unlike "ordinary" (discrete) Kripke frames, they provide a complete semantics for modal logic. In his definition of a topological Kripke frame, Esakia interestingly uses the Vietoris topology and the idea that Kripke frames can be seen as what we nowadays call coalgebras. These choices together foreshadow two influential ideas, which can be seen as red threads running through the research we discuss in this chapter:

1. Modal logic can be used to present the *Stone dual of the Vietoris functor*;
2. Certain "modal variants" of Stone duality can be *categorically separated* into dualities for their base logics and their modalities by stating them as *algebra/coalgebra duality* results.

In this section we will discuss the above two ideas in the "basic" case of Boolean algebras and Stone spaces. Our givens are the contravariant functors $\mathrm{K\Omega}\colon \mathbf{Stone} \to \mathbf{BA}$ and $\mathrm{spec}\colon \mathbf{BA} \to \mathbf{Stone}$, which constitute the dual equivalence $\mathbf{BA} \simeq \mathbf{Stone}^{op}$, and the covariant endofunctor $\mathrm{V}\colon \mathbf{Stone} \to \mathbf{Stone}$. We can present these three functors in one picture as follows:

$$\begin{tikzcd} \mathbf{BA} \arrow[r, bend right, "\mathrm{spec}"'] \arrow[r, phantom, "\simeq"] & \mathbf{Stone} \arrow[l, bend right, "\mathrm{K\Omega}"'] \arrow[loop right, "\mathrm{V}"] \end{tikzcd} \tag{1}$$

Can we do something about the asymmetry in this picture? Can we define a functor on Boolean algebras, in "algebraic" terms, which is dual to V? In §2.1, we will see that this is indeed the case. Specifically, we can use modal logic to describe a functor $\mathrm{M_f}\colon \mathbf{BA} \to \mathbf{BA}$, which can be seen as the Stone dual of $\mathrm{V}\colon \mathbf{Stone} \to \mathbf{Stone}$.

$$\begin{tikzcd} \mathbf{BA} \arrow[loop left, "\mathrm{M_f}"] \arrow[r, bend right, "\mathrm{spec}"'] \arrow[r, phantom, "\simeq"] & \mathbf{Stone} \arrow[l, bend right, "\mathrm{K\Omega}"'] \arrow[loop right, "\mathrm{V}"] \end{tikzcd} \tag{2}$$

We can do two things with the resulting picture: we can use it to frame Esakia's duality as an algebra/coalgebra duality, which is what we will do in §2.2, but we can also view it as an archetype, and ask ourselves: can we generalize this picture? In §3, we will see that $\mathrm{M_f}$ is essentially a restriction of the Vietoris powerlocale, a more general construction on locales, and that one can also prove various duality results for $\mathrm{M_f}$.

2.1 The Stone dual of the Vietoris functor

Our goal in this subsection is to present the fact that the functor $\mathrm{M_f}$, which is presented using *modal logic*, is the Stone dual of $\mathrm{V}\colon \mathbf{Stone} \to \mathbf{Stone}$. To visualize this we pull apart diagram (2), which gives us the following:

$$\begin{tikzcd}
\mathbf{BA} \arrow[r, bend right, "\mathrm{spec}"'] \arrow[d, "\mathrm{M_f}"'] \arrow[r, phantom, "\simeq"] & \mathbf{Stone} \arrow[l, bend right, "\mathrm{K}\Omega"'] \arrow[d, "\mathrm{V}"] \\
\mathbf{BA} \arrow[r, bend right, "\mathrm{spec}"'] \arrow[r, phantom, "\simeq"] & \mathbf{Stone} \arrow[l, bend right, "\mathrm{K}\Omega"']
\end{tikzcd} \tag{3}$$

What we mean by saying that $\mathrm{M_f}$ is the Stone dual of V is that the above diagram commutes up to isomorphism. We will make this claim more precise shortly. The subscript “f” on $\mathrm{M_f}$ denotes that this is a construction which constructs *finitary* algebras; we will see an infinitary version of $\mathrm{M_f}$ in §3.

Starting from a Boolean algebra $\mathbb{A} = \langle A; \wedge, \vee, \neg, 0, 1\rangle$, we can define a new Boolean algebra “based on” $\mathbb{A}$ using the following presentation by generators and relations:

Definition 1. Let $\mathbb{A}$ be a Boolean algebra. We define $\mathrm{M_f}\,\mathbb{A}$ to be the Boolean algebra generated by the set $\{\Box a \mid a \in A\} \cup \{\Diamond a \mid a \in A\}$, subject to the following relations:

$$\begin{aligned}
\Box 1 &= 1; & \Diamond 0 &= 0; \\
\Box(a \wedge b) &= \Box a \wedge \Box b; & \Diamond(a \vee b) &= \Diamond a \vee \Diamond b; \\
\Box(a \vee b) &\leq \Box a \vee \Diamond b; & \Box a \wedge \Diamond b &\leq \Diamond(a \wedge b).
\end{aligned}$$

One may obtain $\mathrm{M_f}\,\mathbb{A}$ by taking the quotient, over the relations listed, of the free Boolean algebra generated by the set $\{\Box a \mid a \in A\}$ and $\{\Diamond a \mid a \in A\}$. In this definition, the sets $\{\Box a \mid a \in A\}$ and $\{\Diamond a \mid a \in A\}$ represent two distinct copies of A; we use boxes and diamonds to denote the respective elements of these sets in order to underline the connection with modal logic. Observe that the relations are nothing more than an algebraic axiomatization of the Boolean modal logic **K**; the last two relations (the *interaction axioms*) imply that $\Diamond \neg a$ is the Boolean complement of $\Box a$: simply substitute $\neg a$ for b (also see Remark 1).

The action of $\mathrm{M_f}$ on Boolean algebra homomorphisms is defined as follows: given a Boolean algebra homomorphism $f\colon \mathbb{A} \to \mathbb{B}$, we can map the generators of $\mathrm{M_f}\,\mathbb{A}$ into $\mathrm{M_f}\,\mathbb{B}$ in the straightforward way, namely by sending

$$\Box a \mapsto \Box f(a) \text{ and } \Diamond a \mapsto \Diamond f(a).$$

Since this mapping respects the relations on $\mathrm{M_f}\,\mathbb{A}$, we obtain a unique Boolean algebra homomorphism $\mathrm{M_f}\,f\colon\ \mathrm{M_f}\,\mathbb{A} \to \mathrm{M_f}\,\mathbb{B}$. This completes our description of the functor $\mathrm{M_f}\colon \mathbf{BA} \to \mathbf{BA}$.

We will now state more precisely what we mean by saying that diagram (3) commutes.

Fact 1. *There exist natural isomorphisms such that for any Boolean algebra $\mathbb{A}$ and any Stone space X, we have*

1. $\mathrm{M_f}(\mathrm{K}\Omega X) \simeq \mathrm{K}\Omega(\mathrm{V}X)$*, and*
2. $\mathrm{spec}(\mathrm{M_f}\,\mathbb{A}) \simeq \mathrm{V}(\mathrm{spec}\,\mathbb{A})$.

Proof sketch. It follows from the fact that $\mathrm{K}\Omega$ and spec form a dual equivalence of categories, that statements (1) and (2) are in fact equivalent. Below we will sketch a proof of the fact that $\mathrm{spec}(\mathrm{M_f}\,\mathbb{A}) \simeq \mathrm{V}(\mathrm{spec}\,\mathbb{A})$ for any Boolean algebra $\mathbb{A}$. We leave the proof of the naturality of this isomorphism to the reader, and use without warning the fact that in this setting, the compact sets coincide with the closed ones.

1. The elements of $\mathrm{V}(\mathrm{spec}\,\mathbb{A})$, i.e., the closed subsets of $\mathrm{spec}\,\mathbb{A}$, are in a 1-1 correspondence with $\mathrm{Filt}\,\mathbb{A}$, the filters of $\mathbb{A}$. We can topologize $\mathrm{Filt}\,\mathbb{A}$ by generating a topology from
$$[a] := \{F \in \mathrm{Filt}\,\mathbb{A} \mid a \in F\}, \text{ and}$$
$$\langle a \rangle := \{F \in \mathrm{Filt}\,\mathbb{A} \mid \forall b \in F, a \wedge b > 0\},$$
where a ranges over the elements of $\mathbb{A}$. Using this topology, $\mathrm{Filt}\,\mathbb{A}$ is homeomorphic to $\mathrm{V}(\mathrm{spec}\,\mathbb{A})$.
2. We view the elements of $\mathrm{spec}(\mathrm{M_f}\,\mathbb{A})$, i.e., the ultrafilters of $\mathrm{M_f}\,\mathbb{A}$, as Boolean homomorphisms $p\colon \mathrm{M_f}\,\mathbb{A} \to \mathbf{2}$, where $\mathbf{2}$ is the two-element Boolean algebra. Given a homomorphism $p\colon \mathrm{M_f}\,\mathbb{A} \to \mathbf{2}$, we define
$$F_p := \{a \in A \mid p(\Box a) = 1\}.$$
This gives us a map from $\mathrm{spec}(\mathrm{M_f}\,\mathbb{A})$ to $\mathrm{Filt}\,\mathbb{A}$.
3. Conversely, given a filter $F \in \mathrm{Filt}\,\mathbb{A}$, we define a map p_F from the generators of $\mathrm{M_f}\,\mathbb{A}$ to $\mathbf{2}$ by specifying
$$p_F(\Box a) = \begin{cases} 1 & \text{if } a \in F; \\ 0 & \text{otherwise,} \end{cases}$$
for the $\Box$-generators, and
$$p_F(\Diamond a) = \begin{cases} 1 & \text{if } \forall b \in F, a \wedge b > 0; \\ 0 & \text{otherwise,} \end{cases}$$
for the $\Diamond$-generators. One can verify that this mapping extends to a Boolean homomorphism $p_F\colon \mathrm{M_f}\,\mathbb{A} \to \mathbf{2}$ by checking the relations from Definition 1. Thus, we have defined a map from $\mathrm{Filt}\,\mathbb{A}$ to $\mathrm{spec}(\mathrm{M_f}\,\mathbb{A})$.

4. Finally, we must show that the assignments $p \mapsto F_p$ and $F \mapsto p_F$ are both continuous, and that for all $F \in \mathrm{Filt}\,\mathbb{A}$, $F = F_{p_F}$ and for all $p\colon \mathrm{M_f}\,\mathbb{A} \to \mathbf{2}$, $p = p_{F_p}$. $\square$

Remark 1. In Boolean modal logic, the modalities $\Box$ and $\Diamond$ are interdefinable. For the functor $\mathrm{M_f}\colon \mathbf{BA} \to \mathbf{BA}$, this is reflected by the following fact. Given a Boolean algebra $\mathbb{A}$, we define $\mathrm{M}_\Box\,\mathbb{A}$ to be the Boolean algebra generated by the set $\{\Box a \mid a \in A\}$, subject to the relations $\Box 1 = 1$ and $\Box(a \wedge b) = \Box a \wedge \Box b$. One can easily show that $\mathrm{M}_\Box$ is a functor on the category of Boolean algebras; moreover, there exists a natural isomorphism such that for any Boolean algebra $\mathbb{A}$, $\mathrm{M_f}\,\mathbb{A} \simeq \mathrm{M}_\Box\,\mathbb{A}$; this isomorphism can be obtained by sending each $\Box a$-generator of $\mathrm{M_f}\,\mathbb{A}$ to the corresponding $\Box a$ in $\mathrm{M}_\Box\,\mathbb{A}$, and each $\Diamond a$ of $\mathrm{M_f}\,\mathbb{A}$ to $\neg\Box\neg a$ in $\mathrm{M}_\Box\,\mathbb{A}$. Indeed, all of the narrative in §2.1 above could have been stated in terms of the functor $\mathrm{M}_\Box$ rather than $\mathrm{M_f}$.

Notes

The Boolean case of Stone duality for the Vietoris functor, as discussed above, is discussed in more detail by Kupke et al. in [27]. See the notes for §3.1 for more sources.

2.2 *Algebra/coalgebra duality*

In this subsection we will use our new knowledge of the functor $\mathrm{M_f}\colon \mathbf{BA} \to \mathbf{BA}$ to state an archetypical algebra/coalgebra duality result: the duality between *Vietoris coalgebras over Stone spaces* and $\mathrm{M_f}$-*algebras over Boolean algebras*. We then discuss the relation of this duality with the original results of Esakia, and its impact on the completeness theory of modal logic.

2.2.1 Algebras and coalgebras

First, we recall the categorical notions of F-algebras and coalgebras. Let $\mathrm{F}\colon \mathbf{C} \to \mathbf{C}$ be an endofunctor on a category $\mathbf{C}$. The category $\mathrm{Alg}_{\mathbf{C}}(\mathrm{F})$, of F-*algebras over* $\mathbf{C}$ has as its *objects* all $\mathbf{C}$-morphisms of the shape $h\colon \mathrm{F}X \to X$, where X, the 'carrier set' of the algebra, ranges over the objects of $\mathbf{C}$. A *morphism* between F-algebras $h\colon \mathrm{F}X \to X$ and $h'\colon \mathrm{F}X' \to X'$ is a $\mathbf{C}$-morphism $f\colon X \to X'$ such that $f \circ h = h' \circ \mathrm{F}f$, i.e., such that the following square commutes:

$$\begin{tikzcd}
\mathrm{F}X \arrow[r, "h"] \arrow[d, "\mathrm{F}f"'] & X \arrow[d, "f"] \\
\mathrm{F}X' \arrow[r, "h'"'] & X'
\end{tikzcd}$$

Below we will see that the category of modal algebras and modal algebra homomorphisms can be presented as a category of F-algebras over **BA**, the category of Boolean algebras.

The category $\mathrm{Coalg}_{\mathbf{C}}(F)$, of F-*coalgebras over* **C**, is defined dually: F-coalgebras are morphisms $h\colon X \to \mathrm{F}X$, and morphisms of F-coalgebras must make a similar square commute:

$$\begin{tikzcd} X \arrow[r, "h"] \arrow[d, "f"'] & \mathrm{F}X \arrow[d, "\mathrm{F}f"] \\ X' \arrow[r, "h'"'] & \mathrm{F}X' \end{tikzcd}$$

An important example of F-coalgebras is given by *Kripke frames*. If $\mathrm{P}\colon \mathbf{Set} \to \mathbf{Set}$ is the covariant powerset functor, then the category of Kripke frames and bounded morphisms can be presented as $\mathrm{Coalg}_{\mathbf{Set}}(\mathrm{P})$, the category of P-coalgebras over **Set**. If $\langle X, R\rangle$ is a Kripke frame, then we can equivalently present the accessibility relation $R \subseteq X \times X$ as the *successor map* $\rho_R\colon X \to \mathrm{P}X$, where $\rho_R\colon x \mapsto \{y \in X \mid Rxy\}$. Moreover, one can easily verify that coalgebra morphisms between P-coalgebras are precisely bounded morphisms.

In [17], Esakia defined *topological* Kripke frames in a similar way: a topological Kripke frame consists of a Stone space X and a binary relation $R \subseteq X \times X$ such that $\rho_R\colon X \to \mathrm{V}X$ is continuous as a map into the Vietoris hyperspace of X. This is noteworthy because the idea to view Kripke frames as P-coalgebras only started to gain popularity through the work of Aczel in the late 1980s [5].

2.2.2 Duality for Vietoris coalgebras

Using Stone duality and Fact 1, it is now an elementary exercise in category theory to see that the category of $\mathrm{M_f}$-algebras over **BA** is dually equivalent to the category of V-coalgebras over **Stone**.

Fact 2. $\mathrm{Alg}_{\mathbf{BA}}(\mathrm{M_f}) \simeq \left(\mathrm{Coalg}_{\mathbf{Stone}}(\mathrm{V})\right)^{op}$.

In order to relate this fact to Esakia's results, we need to do a little more work. Particularly, on the algebraic side, Esakia is not working with algebras for the functor $\mathrm{M_f}$, but with the category **MA** of *modal algebras* and modal algebra homomorphisms. Interestingly, the categories $\mathrm{Alg}_{\mathbf{BA}}(\mathrm{M_f})$ and **MA** are *isomorphic*:

Fact 3. $\mathrm{Alg}_{\mathbf{BA}}(\mathrm{M_f}) \cong \mathbf{MA}$.

Proof sketch. Let $\mathbb{A}$ be a Boolean algebra with underlying set A, and let (f,g) be a pair of functions $f,g\colon A \to A$. We call (f,g) a *modal expansion of* $\mathbb{A}$ if the algebraic structure $\langle \mathbb{A}, f, g\rangle$ is a modal algebra, i.e. f preserves $\wedge$ and 1, g preserves $\vee$ and 0, and $\neg \circ f = g \circ \neg$. The key insight underlying the proof of Fact 3 concerns the existence, for a given Boolean algebra $\mathbb{A}$, of a 1-1 correspondence between the modal expansions of $\mathbb{A}$ and the set $\mathrm{Hom}_{\mathbf{BA}}(\mathrm{M_f}\,\mathbb{A}, \mathbb{A})$ of Boolean algebra homomorphisms from $\mathrm{M_f}\,\mathbb{A}$ to $\mathbb{A}$: if (f,g) is a modal expansion of $\mathbb{A}$, then the assignment

$\Box a \mapsto f(a)$ and $\Diamond a \mapsto g(a)$ uniquely determines a Boolean homomorphism from $\mathrm{M_f}\,\mathbb{A}$ to $\mathbb{A}$, and conversely, if $h\colon \mathrm{M_f}\,\mathbb{A} \to \mathbb{A}$ is a Boolean homomorphism, then the maps $a \mapsto h(\Box a)$ and $a \mapsto h(\Diamond a)$ define a modal expansion of $\mathbb{A}$. $\Box$

From Facts 2 and 3, we can now deduce the following modern version of Esakia's duality result, which states that modal algebras are dually equivalent to Vietoris coalgebras over Stone spaces:

Fact 4. $\mathbf{MA} \simeq \big(\mathrm{Coalg}_{\mathbf{Stone}}(\mathrm{V})\big)^{op}$.

To conclude this subsection, we briefly indicate how the duality between modal algebras and Vietoris coalgebras is used in the completeness theory of modal logic. Again, the key insight here is that Vietoris coalgebras can be seen as topological Kripke frames; in particular, by forgetting the topology of this structure, we obtain an ordinary Kripke frame. This 'forgetting' can be formalized as a functor U from the category $\mathrm{Coalg}_{\mathbf{Stone}}(\mathrm{V})$ to the category $\mathrm{Coalg}_{\mathbf{Set}}(\mathrm{P})$ of P-coalgebras over **Set**, which as we know is isomorphic to the category of Kripke frames and bounded morphisms. The completeness of modal logic can then be proved by showing that every modal algebra $\mathbb{A}$ can be embedded into the full complex algebra of the underlying Kripke frame of the dual Vietoris coalgebra of $\mathbb{A}$. We will briefly revisit the relation between modal logic and coalgebra in §4.2.1.

Notes

There are many good introductions to Stone duality; our notation stems from [23]. More detailed discussions of duality for modal algebras and Vietoris coalgebras can be found in the work of Abramsky [3] and Kupke et al. [27].

Regarding Esakia's duality for topological Kripke frames, it should be noted that in his paper [17], Esakia is mainly interested in the duality between closure algebras and reflexive, transitive topological Kripke frames, and the duality between Heyting algebras and (what are now called) *Esakia spaces*: reflexive, transitive and anti-symmetric topological Kripke frames. The coalgebraic view of Esakia spaces, already present in Esakia's original paper, has also been discussed by Davey and Galati in [16].

3 Varying the base categories

In §2.1, we have seen that the functor $\mathrm{V}\colon \mathbf{Stone} \to \mathbf{Stone}$, the Vietoris hyperspace construction restricted to Stone spaces, is dual to the functor $\mathrm{M_f}\colon \mathbf{BA} \to \mathbf{BA}$, which is presented using Boolean modal logic. In this section we will see that in the compact Hausdorff case, the Vietoris hyperspace is dual to a construction on locales which uses geometric modal logic: the *Vietoris powerlocale*.

In §3.1, we will see how the duality from §2.1 can be extended to compact regular locales and compact Hausdorff spaces, and how this locale-theoretic approach suggests a generalization of the Vietoris hyperspace from compact Hausdorff spaces to a hyperspace construction on *stable locally compact spaces*. In §3.2, we look at an important example which is not covered by the compact Hausdorff case, namely Stone duality for distributive lattices and both coherent and Priestley spaces. Finally, in §3.3, we will see how the locale-theoretic Vietoris construction also is the Stone dual of the *Plotkin powerdomain* construction on algebraic domains.

3.1 The Vietoris powerlocale

Vietoris introduced his hyperspace construction to topologize the set of all closed subsets of a compact Hausdorff space. To extend the duality result from §2.1 beyond Stone spaces, we can use Stone duality as it is used in locale theory: as the categorical *equivalence* between spatial locales and sober spaces.

3.1.1 Compact Hausdorff spaces and compact regular locales

Using the Axiom of Choice, the equivalence between spatial locales and sober spaces restricts to an equivalence between *compact regular locales* and compact Hausdorff spaces. Recall that if $\mathbb{A}$ is a locale and $a,b \in \mathbb{A}$, we say a is *well inside* b ($a \lessdot b$) if there is a c such that $a \wedge c = 0$ and $b \vee c = 1$. Equivalently, $a \lessdot b$ iff $a^* \vee b = 1$, where a^* is the pseudo-complement of a. If $U,V \in \Omega X$ are open subsets of a topological space X, then $U \lessdot V$ iff $\mathrm{cl}(U) \subseteq V$. We say $\mathbb{A}$ is *regular* if for every $a \in \mathbb{A}$, $a = \bigvee\{b \mid b \lessdot a\}$. Furthermore, $\mathbb{A}$ is *compact* if for every non-empty directed set S, $1 \leq \bigvee S$ implies $1 \in S$.

Knowing that Ω, the functor sending a space to its locale of opens, and pt, the functor sending a locale to its space of points, constitute an equivalence between **KRegLoc**, the category of compact regular locales, and **KHaus**, the category of compact Hausdorff spaces and continuous maps, we can now draw the following picture:

$$\begin{tikzcd} \mathbf{KRegLoc} \arrow[r, bend right, "\mathrm{pt}"'] \arrow[r, phantom, "\simeq"] & \mathbf{KHaus} \arrow[l, bend right, "\Omega"'] \arrow[loop right, "\mathrm{V}"] \end{tikzcd} \tag{4}$$

Again, the question is: can we find an endofunctor on **KRegLoc**, defined in "algebraic" terms, corresponding to $\mathrm{V}\colon$ **KHaus** $\to$ **KHaus**? Indeed we can, using the following modification of Definition 1:

Definition 2. Let $\mathbb{A}$ be a locale. We define $\mathrm{M}\,\mathbb{A}$, the *Vietoris powerlocale of* $\mathbb{A}$, to be the locale generated by the set $\{\Box a \mid a \in A\} \cup \{\Diamond a \mid a \in A\}$, subject to the following relations:

$$\begin{array}{cc} \Box 1 = 1; & \Diamond 0 = 0; \\ \Box(a \wedge b) = \Box a \wedge \Box b; & \Diamond(a \vee b) = \Diamond a \vee \Diamond b; \\ \Box \text{preserves directed joins;} & \Diamond \text{preserves directed joins;} \\ \Box(a \vee b) \leq \Box a \vee \Diamond b; & \Box a \wedge \Diamond b \leq \Diamond(a \wedge b). \end{array}$$

The action of M on frame homomorphisms is defined as in the case of **BA**.

Readers who raise their eyebrows at the above definition, worrying about the fact that we are using generators and relations to define an algebra with an *infinitary* signature, can rest assured: for locales, this is not a problem; see [23, §II.1] or [39]. Observe that the only difference between Definitions 1 and 2, apart from the shift from Boolean algebras to locales, is the additional stipulation that $\Box$ and $\Diamond$ preserve directed joins. From a logical viewpoint, this amounts to a shift from Boolean propositional logic to *geometric* propositional logic, i.e., the logic of finite conjunctions and infinite disjunctions which is preeminent in locale theory and topos theory. Also note that although we are currently interested in the restriction of M to compact regular locales, Definition 2 is stated for arbitrary locales.

We can draw the following diagram now that we have our functor M on locales; as before, we will see that the diagram commutes up to natural isomorphism.

$$\begin{tikzcd} \mathbf{KRegLoc} \arrow[r, bend right, "\mathrm{pt}"'] \arrow[d, "\mathrm{M}"'] \arrow[r, phantom, "\simeq"] & \mathbf{KHaus} \arrow[l, bend right, "\Omega"'] \arrow[d, "\mathrm{V}"] \\ \mathbf{KRegLoc} \arrow[r, bend right, "\mathrm{pt}"'] \arrow[r, phantom, "\simeq"] & \mathbf{KHaus} \arrow[l, bend right, "\Omega"'] \end{tikzcd} \tag{5}$$

In other words, M restricted to compact regular locales is the Stone dual of V on compact Hausdorff spaces:

Fact 5. *The functor* $\mathrm{M}\colon \mathbf{Loc} \to \mathbf{Loc}$ *restricts to an endofunctor on compact regular locales. Moreover, there exist natural isomorphisms such that for any compact Hausdorff space X and for any compact regular locale $\mathbb{A}$, we have*

1. $\mathrm{M}(\Omega X) \simeq \Omega(\mathrm{V}X)$*, and*
2. $\mathrm{pt}(\mathrm{M}\mathbb{A}) \simeq \mathrm{V}(\mathrm{pt}\,\mathbb{A})$.

3.1.2 Beyond compact Hausdorff: stably locally compact spaces

Recall that in §3.1.1 we asked ourselves what the Stone dual of the Vietoris hyperspace construction on compact Hausdorff spaces is. Now that we have defined the Vietoris powerlocale construction for arbitrary locales, we can ask ourselves: what is the Stone dual of the Vietoris powerlocale, beyond the compact Hausdorff case? In other words, if $\mathbb{A}$ is a spatial locale and $X = \mathrm{pt}\,\mathbb{A}$ is its equivalent sober space of points, can we define a hyperspace $\mathrm{V}X$ based on X, such that $\mathrm{M}\mathbb{A}$ is equivalent to $\mathrm{V}X$?

In its full generality, this question is ill posed. For example, if we take $\mathbb{A}$ to be the open-set lattice of $\mathbb{Q}$, the set of rational numbers equipped with their usual topology, then $\mathrm{M}\,\mathbb{A}$ does not have a Stone dual because it is not *spatial*, i.e., $\mathrm{M}\,\mathbb{A} \not\simeq \Omega \circ \mathrm{pt}(\mathrm{M}\,\mathbb{A})$ (see p. 177 of [24]). Below we will see, however, that we can ask *and* affirmatively answer this question in the case of stably locally compact spaces.

Recall that a topological space is *sober* if it is T_0 and if every irreducibly closed set is the closure of a singleton. A subset U of a topological space X is *saturated* if it is an intersection of opens; equivalently, U is saturated if it is an upper set in the specialization order of X. A topological space is *stably locally compact* if X is sober, locally compact, and binary intersections of compact saturated sets are compact.

Definition 3. Let X be a stably locally compact space. A *lens* is an intersection of a saturated set with a closed set. We define $\mathrm{V}X$, the Vietoris hyperspace of X, to be the collection of compact lenses of X with the topology generated by the usual subbasic opens,

$$[U] = \{L \in \mathrm{V}X \mid L \subseteq U\} \text{ and}$$
$$\langle U \rangle = \{L \in \mathrm{V}X \mid L \cap U \neq \emptyset\},$$

where U ranges over the opens of X.

The choice of compact *lenses*, rather than arbitrary compact subsets of X, is dictated by the desideratum that $\mathrm{V}X$ is again T_0: the original space X may have too many compact subsets.

What are the localic analogs of stably locally compact spaces? Recall that the *way-below* relation on a dcpo (directed complete partial order) $\mathbb{D}$ is defined as follows: we say that a is way below b ($a \ll b$) if for every directed set S with $b \leq \bigvee S$, there is a $c \in S$ such that $a \leq c$. A dcpo $\mathbb{D}$ is *continuous* if for every $a \in \mathbb{D}$, the set $\{b \in \mathbb{D} \mid b \ll a\}$ is directed and $a = \bigvee\{b \in L \mid b \ll a\}$. Now let $\mathbb{A}$ be a locale. We say $\mathbb{A}$ is *stably locally compact* if the dcpo reduct of $\mathbb{A}$ is continuous and for all $a, b, c \in \mathbb{A}$, if $a \ll b$ and $a \ll c$ then $a \ll b \wedge c$.

Fact 6. *1. Both* M *and* V *preserve stable local compactness;*
2. If $\mathbb{A}$ is a stably locally compact locale and X is a stably locally compact space, then both $\mathrm{M}(\Omega X) \simeq \Omega(\mathrm{V}X)$ *and* $\mathrm{V}(\mathrm{pt}\,\mathbb{A}) \simeq \mathrm{pt}(\mathrm{M}\,\mathbb{A})$.

Notes

The equivalence between the categories **KRegLoc** and **KHaus** was established by Isbell [21], see also [7]. The Vietoris powerlocale was first introduced by Johnstone in [23, Ch. III §4], where he also proves the results contained in Fact 5. We also recommend [23] as an introduction to locale theory and the duality between compact regular locales and compact Hausdorff spaces. For an introduction to stably locally compact spaces, we refer to Gierz. et al. [19].

The results contained in Fact 6 are also due to Johnstone [24]. For a discussion of the equivalence between stably locally compact locales and stably locally compact spaces, we suggest reading [26, §1.2 and 1.3]. An alternative account of the Vietoris construction in both localic and spatial form is given by Simmons in [36].

Finally, we would like to point out two alternative approaches to the question “What is the Stone dual of the Vietoris powerlocale?”. Firstly, this question has often been approached by (more) *constructive* means, diverging from the “classical” perspective we take in this chapter. This is the case in the work of Johnstone [24] we referred to in §3.1 and of Vickers [40]. Secondly we would like to point out the work of Palmigiano and Venema [33], who use *Chu spaces* to find the Stone dual of the Vietoris powerlocale, taking inspiration from the success of *relation lifting* (see §4) in coalgebraic logic. Yet another approach uses so-called de Vries algebras [10].

3.2 Distributive lattices and the Vietoris construction

We will now look at the Vietoris functor in relation to an important example of stably locally compact spaces which are not necessarily Hausdorff, namely, the Stone duals of distributive lattices: coherent spaces and Priestley spaces. In this subsection we will look at four different versions of the Vietoris functor, each of which acts on a category (dually) equivalent to **DL**, the category of bounded distributive lattices and (bounded) distributive lattice homomorphisms (throughout this chapter, lattices are assumed to be bounded). The final aim is to show that the three squares in diagram (6) commute up to isomorphism.

$$\begin{tikzcd}
\mathbf{DL} \arrow[r, shift right, bend right, "\mathrm{Idl}"'] \arrow[d, "\mathrm{M_f}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohLoc} \arrow[l, shift right, bend right, "\mathrm{K}"'] \arrow[r, shift right, bend right, "\mathrm{pt}"'] \arrow[d, "\mathrm{M}"] \arrow[r, phantom, "\simeq"] & \mathbf{CohSp} \arrow[l, shift right, bend right, "\Omega"'] \arrow[r, shift right, bend right, "\mathrm{Patch}"'] \arrow[d, "\mathrm{V}"'] \arrow[r, phantom, "\cong"] & \mathbf{Priestley} \arrow[l, shift right, bend right, "\mathrm{OpenUpper}"'] \arrow[d, "\mathrm{V_c}"] \\
\mathbf{DL} \arrow[r, shift right, bend right, "\mathrm{Idl}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohLoc} \arrow[l, shift right, bend right, "\mathrm{K}"'] \arrow[r, shift right, bend right, "\mathrm{pt}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohSp} \arrow[l, shift right, bend right, "\Omega"'] \arrow[r, shift right, bend right, "\mathrm{Patch}"'] \arrow[r, phantom, "\cong"] & \mathbf{Priestley} \arrow[l, shift right, bend right, "\mathrm{OpenUpper}"']
\end{tikzcd} \tag{6}$$

In §3.2.1, we look at a distributive lattice version of the functor $\mathrm{M_f}$ and its relation to M. In §3.2.2, we will see how M restricted to coherent locales corresponds to the compact lens hyperspace of Definition 3. Finally, in §3.2.3, we will see how to construct the *convex* Vietoris hyperspace of a Priestley space.

3.2.1 Distributive lattices and coherent locales

We start by looking closer at the left square in diagram (6).

$$
\begin{tikzcd}
\mathbf{DL} \arrow[r, bend right, "\mathrm{Idl}"'] \arrow[r, phantom, "\simeq"] \arrow[d, "\mathrm{M_f}"'] & \mathbf{CohLoc} \arrow[l, bend right, "\mathrm{K}"'] \arrow[d, "\mathrm{M}"] \\
\mathbf{DL} \arrow[r, bend right, "\mathrm{Idl}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohLoc} \arrow[l, bend right, "\mathrm{K}"']
\end{tikzcd}
$$

By **CohLoc** we denote the category of coherent locales and coherent maps. A locale $\mathbb{A}$ is *coherent* if $\mathbb{A}$ is algebraic, meaning that every $a \in \mathbb{A}$ is a directed join of finite (also called compact) elements, and if additionally $\mathrm{K}\,\mathbb{A}$, the poset of finite elements of $\mathbb{A}$, forms a (distributive) lattice. Equivalently, $\mathbb{A}$ has to be the ideal completion of a distributive lattice. In fact, the ideal completion functor Idl is one half of a dual equivalence between the category **DL** of distributive lattices and lattice homomorphisms, and **CohLoc** of coherent locales and coherent maps; the other half is the functor K which sends a coherent locale to its distributive lattice of finite elements.

To understand the vertical arrows in the left square of diagram (6) we need to introduce the functor $\mathrm{M_f}$ on distributive lattices.

Definition 4. Let $\mathbb{A}$ be a distributive lattice. We define $\mathrm{M_f}\,\mathbb{A}$ to be the distributive lattice generated by the set $\{\Box a \mid a \in A\} \cup \{\Diamond a \mid a \in A\}$, subject to the following relations:

$$
\begin{aligned}
\Box 1 &= 1; & \Diamond 0 &= 0;\\
\Box(a \wedge b) &= \Box a \wedge \Box b; & \Diamond(a \vee b) &= \Diamond a \vee \Diamond b;\\
\Box(a \vee b) &\leq \Box a \vee \Diamond b; & \Box a \wedge \Diamond b &\leq \Diamond(a \wedge b).
\end{aligned}
$$

The action of M on lattice homomorphisms is defined as before: given $f\colon \mathbb{A} \to \mathbb{B}$, we let $\mathrm{M}\,f$ be the extension of $\Box a \mapsto \Box f(a)$ and $\Diamond a \mapsto \Diamond f(a)$.

Note that Definition 4 differs from Definition 1 only because we are generating a distributive lattice rather than a Boolean algebra. This difference is quite subtle due to the following fact.

Fact 7. *Let* $\mathrm{U}\colon \mathbf{BA} \to \mathbf{DL}$ *denote the forgetful functor that sends a Boolean algebra to its underlying (distributive) lattice. Then there exists a natural isomorphism such that for any Boolean algebra* $\mathbb{A}$*, we have* $\mathrm{U}(\mathrm{M_f}\,\mathbb{A}) \simeq \mathrm{M_f}(\mathrm{U}\,\mathbb{A})$.

The above fact corresponds to the well-known fact in Boolean modal logic that any modal formula containing arbitrary negations is equivalent to a modal formula in which negations are only applied to proposition letters — this observation can also be used in a proof of Fact 7.

We can now explicitly state the content of the first square of diagram (6), namely, that the functor $\mathrm{M_f}$ on distributive lattices is equivalent to the Vietoris powerlocale M restricted to coherent locales.

Fact 8. *The Vietoris powerlocale functor* $\mathrm{M}\colon \mathbf{Loc} \to \mathbf{Loc}$ *restricts to an endofunctor on coherent locales. Moreover, there exist natural isomorphisms such that for any coherent locale* $\mathbb{A}$ *and for any distributive lattice* $\mathbb{L}$*, we have*

1. $\mathrm{M}(\mathrm{Idl}\,\mathbb{L}) \simeq \mathrm{Idl}(\mathrm{M_f}\,\mathbb{L})$, *and*
2. $\mathrm{M_f}(\mathrm{K}\,\mathbb{A}) \simeq \mathrm{K}(\mathrm{M}\,\mathbb{A})$.

3.2.2 Coherent locales and coherent spaces

We move on to the middle square of diagram (6), in which we encounter the Vietoris functor on coherent spaces.

$$\begin{tikzcd}
\mathbf{CohLoc} \arrow[d, "\mathrm{M}"'] \arrow[r, bend right, "\mathrm{pt}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohSp} \arrow[l, bend right, "\Omega"'] \arrow[d, "\mathrm{V}"] \\
\mathbf{CohLoc} \arrow[r, bend right, "\mathrm{pt}"'] \arrow[r, phantom, "\simeq"] & \mathbf{CohSp} \arrow[l, bend right, "\Omega"']
\end{tikzcd}$$

By **CohSp** we denote the category of *coherent spaces* and coherent maps. Recall that a coherent space is a (compact) sober space with a basis of compact opens, with the additional property that any finite intersection of compact opens is compact. (Coherent spaces/maps are also known as *spectral* spaces/maps.) A continuous map between coherent spaces is called coherent if the inverse image of a compact open set is compact.

Definition 5. Let X be a coherent space. We define $\mathrm{V}X$ to be the Vietoris hyperspace of compact lenses introduced in Definition 3. Moreover, if $f\colon X \to Y$ is a coherent map between coherent spaces, we define $\mathrm{V}f\colon\ \mathrm{V}X \to \mathrm{V}Y$ as follows:

$$\mathrm{V}f\colon L \mapsto {\uparrow} f[L] \cap \mathrm{cl}(f[L]),$$

where ${\uparrow} f[L]$ is the saturation of $f[L]$, i.e., its upward closure in the specialization order, and $\mathrm{cl}(f[L])$ is the closure of $f[L]$.

The reason we need to take a "lens closure" in the definition of $\mathrm{V}f$ above is that unlike compactness, the property of being a lens is not stable under forward images of continuous functions.

Fact 9. *The construction* V *described above is well-defined, and it is an endofunctor on the category of coherent spaces and coherent maps. Moreover, there exist natural isomorphisms such that for any coherent locale* $\mathbb{A}$ *and for any coherent space* X*, we have*

1. $\mathrm{V}(\mathrm{pt}\,\mathbb{A}) \simeq \mathrm{pt}(\mathrm{M}\,\mathbb{A})$, *and*
2. $\Omega(\mathrm{V}X) \simeq \mathrm{M}(\Omega X)$.

We can refine Definition 5 by exploiting the special role of compact open sets in coherent spaces:

Fact 10. *Let X be a coherent space. If $U \subseteq X$ is compact open in X, then so are $[U]$ and $\langle U\rangle$ in $\mathrm{V}X$. In fact, the sets of the form $[U]$ and $\langle U\rangle$, with U ranging over the compact opens of X, form a sub-base for the topology on $\mathrm{V}X$.*

Here we are essentially using the fact that $\mathrm{M_f}\colon \mathbf{DL} \to \mathbf{DL}$ is the Stone dual of $\mathrm{V}\colon \mathbf{CohSp} \to \mathbf{CohSp}$.

3.2.3 Coherent spaces and Priestley spaces

We will now discuss the final square of diagram (6), and learn about the Vietoris construction for Priestley spaces.

$$\begin{tikzcd}
\mathbf{CohSp} \arrow[r, bend right, "\mathrm{Patch}"'] \arrow[d, "\mathrm{V}"'] \arrow[r, phantom, "\cong"] & \mathbf{Priestley} \arrow[l, bend right, "\mathrm{OpenUpper}"'] \arrow[d, "\mathrm{V_c}"] \\
\mathbf{CohSp} \arrow[r, bend right, "\mathrm{Patch}"'] \arrow[r, phantom, "\cong"] & \mathbf{Priestley} \arrow[l, bend right, "\mathrm{OpenUpper}"']
\end{tikzcd} \tag{7}$$

By **Priestley** we denote the category of *Priestley spaces* and order-preserving continuous maps. A Priestley space is a partially ordered compact space $\langle X, \leq, \tau\rangle$, with the additional property that if $x, y \in X$ such that $x \not\leq y$, then there exists a clopen upper set $U \subseteq X$ such that $x \in U \not\ni y$. As a consequence, Priestley spaces are Hausdorff. The categories **CohSp** and **Priestley** are *isomorphic*: we can transform coherent spaces into Priestley spaces and vice versa, and these transformations are mutually inverse. If $\langle X, \tau\rangle$ is a coherent space, then $\langle X, \leq_\tau, \mathrm{patch}(\tau)\rangle$ is a Priestley space, where $\leq_\tau$ is the specialization order of τ and $\mathrm{patch}(\tau)$ is the *patch topology* of τ, i.e., the topology generated by the open sets of τ and the complements of compact saturated sets. This allows one to define a functor $\mathrm{Patch}\colon \mathbf{CohSp} \to \mathbf{Priestley}$, which leaves the set-theoretic functions underlying coherent maps unchanged. We can also go from Priestley spaces to coherent spaces: if $\langle X, \leq, \sigma\rangle$ is a Priestley space, then $\langle X, \sigma^{\uparrow}\rangle$ is a coherent space, where $\sigma^{\uparrow}$ is the collection of open upper sets of $\langle X, \leq, \sigma\rangle$. This gives us a functor $\mathrm{OpenUpper}\colon \mathbf{Priestley} \to \mathbf{CohSp}$, which again leaves the functions underlying the morphisms unchanged. The functors OpenUpper and Patch form an isomorphism of categories: if X is a coherent space and if Y is a Priestley space, then

$$\mathrm{OpenUpper}\left(\mathrm{Patch}\, X\right) = X \text{ and } \mathrm{Patch}\left(\mathrm{OpenUpper}\, Y\right) = Y.$$

For a detailed account of this connection, see Cornish [15].

Before we introduce the Vietoris construction on Priestley spaces, we will take a closer look at the patch topology, and in particular the patch topology of $\mathrm{V}X$ when X is a coherent space.

Fact 11. *Let $\langle X,\tau\rangle$ be a coherent space. The patch topology* $\mathrm{patch}(\tau)$ *of τ is generated by the following base:*

$$\{U \setminus V \mid U, V \text{ compact open in } \tau\}.$$

Topological properties with respect to τ often correspond to *order*-topological properties with respect to $\mathrm{patch}(\tau)$.

Fact 12. *Let $\langle X,\tau\rangle$ be a coherent space, and let $\leq$ be its specialization order.*

1. *The* open *subsets of X are precisely the* patch-open upper *subsets of X.*
2. *The* closed *subsets of X are precisely the* patch-closed lower *subsets of X.*
3. *The* compact saturated *subsets of X are precisely the* patch-closed upper *subsets of X.*
4. *The* compact open *subsets of X are precisely the* patch-clopen upper *subsets of X.*

Lemma 1. *Let $\langle X,\tau\rangle$ be a coherent space, let $\leq$ be its specialization order, and let L be a compact lens. Then (1) L is patch-compact; and (2) $\downarrow L$ is closed.*

Proof. Let L be a compact lens. Since L is a lens, $L = \uparrow L \cap \mathrm{cl}(L)$. Because all opens are upper sets, a subset $\mathscr{C} \subseteq \tau$ covers L iff it covers $\uparrow L$; it follows that L is compact iff $\uparrow L$ is compact. Since we assumed that L is compact, so is $\uparrow L$, whence by Fact 12(3) above, $\uparrow L$ must be patch-closed. By Fact 12(2), $\mathrm{cl}(L)$ is also patch-closed. It follows that $L = \uparrow L \cap \mathrm{cl}(L)$ is patch-closed, and because $\mathrm{patch}(\tau)$ is a compact Hausdorff topology, L is also patch-compact. This proves statement (1); as for the second statement, since $\langle X, \leq, \mathrm{patch}(\tau)\rangle$ is a Priestley space, it follows from e.g. [23, Ch. 7, §1] that $\downarrow L$ is patch-closed. By Fact 12(2), $\downarrow L$ is also closed w.r.t. τ. □

A subset U of a poset $\mathbb{P}$ is called *convex* if $U = \uparrow U \cap \downarrow U$. If U,V are subsets of $\mathbb{P}$, we say that U is below V in the *Egli-Milner order* ($U \leq_{EM} V$) if both $U \subseteq \downarrow V$ and $\uparrow U \subseteq V$. In other words, $U \leq_{EM} V$ iff

$$\forall x \in U, \exists y \in V \text{ such that } x \leq y, \text{ and } \forall y \in V, \exists x \in U \text{ such that } x \leq y. \tag{8}$$

Proposition 1. *Let $\langle X,\tau\rangle$ be a coherent space and let $\leq$ be its specialization order.*

1. *The compact lenses of X are precisely the patch-compact convex subsets of X.*
2. *The specialization order of $\mathbb{V}X$ is $\leq_{EM}$.*
3. *The patch topology of $\mathbb{V}X$ is generated by sets of the form*

$$[U], \langle U\rangle, [X \setminus U], \langle X \setminus U\rangle,$$

where U ranges over the compact opens of X.

Proof. 1. Suppose $L \subseteq X$ is a compact lens. Then by Lemma 1(1), L is patch-compact. Since $L = \uparrow L \cap \mathrm{cl}(L)$ and $\mathrm{cl}(L)$ is always a lower set, it is easy to see that L is convex. For the converse, suppose that L is a patch-compact convex set. Because

$\tau \subseteq \text{patch}(\tau)$, L must also be compact w.r.t. τ. Moreover, since $L = \uparrow L \cap \downarrow L$, and $\downarrow L$ is closed by Lemma 1(2), we see that L is a lens.

2. Let L and M be points of $\mathrm{V}X$, i.e., compact lenses of X. Observe that L is below M in the specialization order of $\mathrm{V}X$ iff

$$\forall U \in \tau, L \in [U] \Rightarrow M \in [U], \text{ and } \forall U \in \tau, L \in \langle U \rangle \Rightarrow M \in \langle U \rangle. \tag{9}$$

Suppose that (9) holds for L and M. Then since $\uparrow L = \bigcap\{U \in \tau \mid L \subseteq U\}$, it follows from the left half of (9) that $M \subseteq \uparrow L$. Moreover, if we take $U = X \setminus \downarrow M$, then by Lemma 1(2), U is open. Now $M \notin \langle U \rangle$, so by the right side of (9), $L \notin \langle U \rangle$, i.e., $L \cap (X \setminus \downarrow M) = \emptyset$, so that $L \subseteq \downarrow M$. We conclude that $L \leq_{EM} M$.

Conversely, suppose that $L \leq_{EM} M$, so that $M \subseteq \uparrow L$ and $L \subseteq \downarrow M$. If U is an open set such that $L \in [U]$, i.e., such that $L \subseteq U$, then $\uparrow L \subseteq U$ so since we assumed $M \subseteq \uparrow L$, $M \in [U]$. And if U is an open set such that $M \notin \langle U \rangle$, i.e., such that $M \cap U = \emptyset$, then since U is an upper set, it is also the case that $\downarrow M \cap U = \emptyset$. But then since we assumed that $L \subseteq \downarrow M$, we see that $L \cap U = \emptyset$, so that $L \notin \langle U \rangle$. It follows that (9) holds.

3. Observe that if U is a compact open set, then since

$$[X \setminus U] = \mathrm{V}X \setminus \langle U \rangle \text{ and } \langle X \setminus U \rangle = \mathrm{V}X \setminus [U], \tag{10}$$

it follows from Fact 11 that $[X \setminus U]$ and $\langle X \setminus U \rangle$ are patch-open sets in $\mathrm{V}X$.

It follows from Fact 10 that every compact open of $\mathrm{V}X$ can be expressed as a finite union of finite intersections of sets of the form $[U]$ and $\langle U \rangle$, where U ranges over compact opens of X. Using De Morgan's laws and the distributive laws, one can see that the *complement* of a compact open set in $\mathrm{V}X$ can therefore be expressed as a finite union of finite intersections of sets $\mathrm{V}X \setminus [U]$ and $\mathrm{V}X \setminus \langle U \rangle$, with U still ranging over compact opens. Using (10), we see therefore that the complements of compact opens of $\mathrm{V}X$ can be obtained as finite unions of finite intersections of sets $[X \setminus U]$ and $\langle X \setminus U \rangle$. It now follows by Fact 11 that the patch topology of $\mathrm{V}X$ is generated by sets of the form $[U], \langle U \rangle, [X \setminus U], \langle X \setminus U \rangle$, with U ranging over the compact opens of X. $\square$

We will now define the Vietoris construction on Priestley spaces.

Definition 6. Let X be a Priestley space. We define $\mathrm{V_c}X$, the *Vietoris convex hyperspace* of X, to be the collection of compact convex subsets of X, ordered by the Egli-Milner order $\leq_{EM}$ and topologized by the usual subbasic opens $[U]$ and $\langle U \rangle$, with U ranging over the clopen upper and clopen lower sets. If $f\colon X \to Y$ is a morphism of Priestley spaces, i.e., if f is a continuous order-preserving map, then we define

$$\mathrm{V_c} f\colon F \mapsto \uparrow f[F] \cap \downarrow f[F].$$

In other words, $\mathrm{V_c} f$ sends each compact set F to the "convex closure" of its forward image $f[F]$.

In light of Proposition 1, the following should come as no surprise:

Theorem 13. *The construction* $\mathrm{V_c}$ *described above is an endofunctor on the category of Priestley spaces and continuous order-preserving maps.*

In fact, the Vietoris convex hyperspace on Priestley spaces coincides with the Vietoris hyperspace of compact lenses on coherent spaces, i.e., diagram (7) *commutes:*

$$\mathrm{V_c} \circ \mathrm{Patch} = \mathrm{Patch} \circ \mathrm{V} \ \textit{and} \ \mathrm{V} \circ \mathrm{OpenUpper} = \mathrm{OpenUpper} \circ \mathrm{V_c}\,.$$

Notes

Facts 8 and 9 can be found (implicitly) in Johnstone's [24]; we do not know a reference for Fact 7, which corresponds to a well-known fact in modal logic.

The origins of Definition 6 and Theorem 13 are not entirely clear to us. Definition 6 is mentioned by Palmigiano in a paper [32] which focuses on a different kind of Vietoris construction for Priestley spaces. Theorem 13 is stated by Bezhanishvili and Kurz in [11], who then refer to [24] and [32]. None of these sources spells out a proof however, so we decided to include one here.

A detailed discussion of Facts 11 and 12, and the isomorphism between the categories of coherent spaces and Priestley spaces, both in relation to *bitopological spaces*, can be found in [9]. An earlier discussion of the patch topology can be found in [20].

3.3 Algebraic domains and the Plotkin powerdomain

In this final subsection of §3, we will look at Stone duality for the Vietoris powerlocale from an opposite perspective. Namely, we will look at algebraic domains and the Plotkin powerdomain, and we will see that the Stone dual of the Plotkin powerdomain is the Vietoris powerlocale.

Domains, the structures which are studied in domain theory for applications such as semantics for programming languages, are ordered structures which one can simultaneously regard as topological spaces. Crucially, the topology of a domain is uniquely determined by its order (namely, it is the *Scott topology*), and conversely, the order on a domain is uniquely determined by its topology (namely, it is the *specialization order*). From a topological viewpoint, one could say that domains are classes of T_0 spaces which are defined using order-theoretic properties of their specialization orders.

In several important cases, the natural topology of a domain (the Scott topology) can be understood via Stone duality. In this subsection we will consider algebraic domains, a class of directed complete partial orders (dcpo's) which happen to have the property that they are sober in their Scott topologies. Consequently, algebraic domains can be understood in *three* different ways: (1) as dcpo's, (2) as topological spaces, and (3) dually, as locales.

First, we recall the definition of algebraic domains. An element p of a dcpo $\mathbb{D}$ is called *finite* if for all directed S such that $p \leq \bigvee S$, there is a $c \in S$ such that $p \leq c$. We denote the poset of finite elements of $\mathbb{D}$ by $\mathrm{K}\mathbb{D}$. We say $\mathbb{D}$ is an *algebraic domain* if $\mathbb{D}$ is a dcpo such that for all $a \in D$, the set $\{b \in \mathrm{K}\mathbb{D} \mid b \leq a\}$ is directed and

$a = \bigvee\{b \in \mathrm{K}\mathbb{D} \mid b \leq a\}$. Every algebraic domain $\mathbb{D}$ is completely determined by its finite elements; specifically, $\mathbb{D} \simeq \mathrm{Idl}(\mathrm{K}\mathbb{D})$, where Idl stands for taking the ideal completion. (Note that $\mathrm{K}\mathbb{D}$ is a join semilattice, so that ideals can be defined as usual.)

The *Scott topology* on a domain $\mathbb{D}$ is defined as the collection of all upper sets which are inaccessible by directed joins; we denote this topology (and also the locale it induces) by $\Sigma\mathbb{D}$. This allows us to transform domains into locales. Moreover, if we convert the locale $\Sigma\mathbb{D}$ back into a space of points using Stone duality, we find that $\mathrm{pt}(\Sigma\mathbb{D})$, viewed as a dcpo, is isomorphic to $\mathbb{D}$, *assuming* $\mathbb{D}$ *is algebraic*. (The order on $\mathrm{pt}(\Sigma\mathbb{D})$ is the specialization order.)

Powerdomain constructions were introduced in domain theory to model *branching* of computational processes. One particular powerdomain construction is the so-called *Plotkin powerdomain*, which is defined as a free dcpo semi-lattice construction. For algebraic domains, the following surprising characterization of the Plotkin powerdomains is known: if $\mathbb{D}$ is an algebraic domain, then its Plotkin powerdomain can be presented as the ideal completion of the convex subsets of $\mathrm{K}\mathbb{D}$, ordered by the Egli-Milner order (see (8)).

Given the Plotkin powerdomain construction on algebraic domains, and the fact that algebraic domains can be seen as the dual spaces of locales, we can now ask ourselves the question: what is the Stone dual of the Plotkin powerdomain? The answer is that the formation of the Plotkin powerdomains corresponds exactly to the formation of the Vietoris powerlocale.

Fact 14. *Let* $\mathbb{D}$ *be an algebraic domain and let* $\mathrm{Pl}\,\mathbb{D}$ *be its Plotkin powerdomain. Then* $\mathrm{M}(\Sigma\mathbb{D}) \simeq \Sigma(\mathrm{Pl}\,\mathbb{D})$.

Notes

For a general introduction to domain theory we refer to [19], or to [4] in connection with power constructions. Fact 14 is due to Robinson [34]. A natural generalization of it would be to consider continuous rather than algebraic domains. Vickers [41] discusses powerdomains and powerlocales in the context of continuous lattices, but he does not address the specific problem of generalizing Fact 14.

Above, we have left out a discussion of Abramsky's *Domain theory in logical form* [2], for lack of space. In a nutshell, Abramsky exploits Stone duality for the intersection of algebraic domains (§3.3) and coherent spaces. Within this context, the Plotkin powerdomain is Stone dual to the functor $\mathrm{M_f}$ on distributive lattices, a fact which is used to study *bisimulation* in [1]. For an introduction to the very powerful framework of "Domain theory in logical form" we refer the reader to [2] or [4].

4 The Vietoris construction and the nabla modality

If we look at our discussion of Stone duality for the Vietoris functor in §2.1, we see an asymmetry in the presentations of the hyperspace topology on the one hand and the logical/algebraic powerlocale constructions on the other hand. The hyperspace of a compact Hausdorff space X can be topologized in two equivalent ways, namely using *basic opens* of the shape

$$\nabla\{U_1,\ldots,U_n\} := \{F \in \mathrm{K}X \mid F \subseteq \bigcup_{i=1}^{n} U_i \text{ and } \forall i \leq n, F \cap U_i \neq \emptyset\},$$

(∇ is pronounced "nabla") versus using *subbasic opens* of the shape

$$[U] := \{F \in \mathrm{K}X \mid F \subseteq U\} \text{ and}$$
$$\langle U\rangle := \{F \in \mathrm{K}X \mid F \cap U \neq \emptyset\},$$

where U and the U_i range over the opens of X. The powerlocales $\mathrm{M}\,\mathbb{A}$ and $\mathrm{M_f}\,\mathbb{A}$, on the other hand, we only presented using box ($\Box$) and diamond ($\Diamond$) in combination with positive modal logic.

The co-existence of these two distinct definitions of the Vietoris construction on topological spaces naturally raises the question, how to give a presentation of the Vietoris powerlocale directly in terms of nabla (∇); similarly, it is an interesting problem how to axiomatize modal logic in terms of the nabla modality. In this section we will see how ideas from the theory of *coalgebra*, and more specifically, *coalgebraic modal logic* may be used to address and solve these problems. As a by-product of this coalgebraic approach, we will see that the Vietoris construction V can be seen as an *instance* of a more general construction which is parametrized by a 'coalgebra functor' on the category **Set**: Given such a functor T we will define the notion of a T-powerlocale functor on the category of locales, in such a way that the Vietoris construction corresponds to the case where T is the power set functor P.

We will first have a brief look at the nabla modality as a derived connective in §4.1. In §4.2, we will introduce the syntax and semantics of Moss' coalgebraic modal logic. In §4.3, we introduce the Carioca axiom system, which is sound and complete with respect to Moss' coalgebraic logic. In §4.4 we will then show how these axioms can be applied to the Vietoris powerlocale, and how they even lead to a notion of generalized powerlocale.

4.1 Nabla-expressions

In this subsection, we will look at the nabla modality as a derived connective. From this point of view, the nabla modality is simply an *expression* consisting of $\Box$ and $\Diamond$ modalities. The main result we discuss is the fact that every element of a powerlocale can be expressed as a disjunction of nabla expressions.

Definition 7. Let $\mathbb{A}$ be a locale, distributive lattice or Boolean algebra. A *nabla-expression over* $\mathbb{A}$ is a term of the shape

$$\Box\left(\bigvee\alpha\right) \wedge \bigwedge_{a\in\alpha} \Diamond a,$$

where $\alpha \subseteq \mathbb{A}$ is a finite subset of $\mathbb{A}$.

It is not hard to see that if $\mathbb{A} = \Omega X$, then the nabla-expressions over $\mathbb{A}$ correspond precisely to the basic open subsets $\nabla\{U_1,\ldots,U_n\}$ of the Vietoris hyperspace $\mathrm{V}X$. The fact that the sets $\nabla\{U_1,\ldots,U_n\}$ form a basis for the Vietoris topology, rather than a subbasis, can also be expressed algebraically:

Fact 15. *1. If $\mathbb{A}$ is a locale then every element of* $\mathrm{M}\,\mathbb{A}$ *can be expressed as a join of nabla-expressions over $\mathbb{A}$;*
2. If $\mathbb{A}$ is a distributive lattice or a Boolean algebra, then every element of $\mathrm{M_f}\,\mathbb{A}$ *can be expressed as a* finite *join of nabla-expressions over $\mathbb{A}$.*

Proof sketch. We will briefly discuss the case where $\mathbb{A}$ is a distributive lattice. Suppose $x \in \mathrm{M_f}\,\mathbb{A}$. Because $\mathrm{M_f}\,\mathbb{A}$ is generated by (equivalence classes of) elements of the shape $\Box a$ and $\Diamond b$, we may assume that x is a disjunction of terms of the shape

$$\textstyle\bigwedge_I \Box a_i \wedge \bigwedge_J \Diamond b_j,$$

where I, J are finite index sets and the a_i, b_j come from $\mathbb{A}$. It will suffice to show that such conjunctions can be obtained as disjunctions of nabla-expressions.

Because $\Box$ preserves finite meets, we will assume we have a single $\Box$-conjunct $\Box a$ (if $I = \emptyset$, this will be the term $\Box 1$). We will now show that the following term can be obtained as a disjunction of at most two nabla-expressions:

$$\textstyle\Box a \wedge \bigwedge_J \Diamond b_j.$$

For the case that $|J| = 0$, we leave it as an exercise for the reader to show that

$$\Box a = (\Box a \wedge \Diamond a) \vee \big(\Box(\textstyle\bigvee\emptyset) \wedge \textstyle\bigwedge\emptyset\big),$$

which is a binary disjunction of nabla-expressions.

We will now assume that $|J| > 0$, and we will show that in this case we get just one nabla-expression. To do this, we will use the following equations, which can be easily derived from the axioms in Definition 4:

$$\Box c \wedge \Diamond d = \Box c \wedge \Diamond(c \wedge d); \tag{11}$$
$$\Box c \wedge \Diamond d = \Box c \wedge \Diamond c \wedge \Diamond d. \tag{12}$$

We now see that

$$\begin{array}{ll}
\Box a \wedge \bigwedge_J \Diamond b_j & \\
= \Box a \wedge \bigwedge_J \Diamond(b_j \wedge a) & \text{by (11) } (|J| \text{ times}),\\
= \Box a \wedge \Diamond a \wedge \bigwedge_J \Diamond(b_j \wedge a) & \text{by (12) since } |J| > 0,\\
= \Box\big(a \vee \bigvee_J (b_j \wedge a)\big) \wedge \Diamond a \wedge \bigwedge_J \Diamond(b_j \wedge a) & \text{by order theory.}
\end{array}$$

The final expression above is now indeed a nabla-expression, for

$$\alpha = \{a\} \cup \{b_j \wedge a \mid j \in J\}.$$

Since we assumed $x \in \mathrm{M_f}\,\mathbb{A}$ to be a finite disjunction of conjunctions of $\Box a$'s and $\Diamond b_j$'s, and since each such conjunction is the disjunction of at most two nabla-expressions, it follows that x itself is also a finite disjunction of nabla-expressions. The same argument can be applied in the locale case. $\Box$

Notes

What we call "nabla expressions" above have been used, in one form or another, both in modal logic and in locale/domain-theoretic investigations of the powerlocale. For modal logic, see e.g. the normal forms used by Fine [18]; for locale theory, see e.g. Johnstone [23, 24] and Robinson [34]. None of these sources, however, explicitly state or prove Fact 15.

4.2 Moss' coalgebraic logic

In this subsection we introduce the syntax and semantics of Moss' coalgebraic logic. We start with an observation about the semantics of nabla-expressions in Kripke frames. We will then very briefly review some of the background of coalgebra and coalgebraic logic in §4.2.1. In §4.2.2, we introduce relation lifting, a technique which sits at the heart of Moss' coalgebraic logic. In §4.2.3, we then introduce the syntax and semantics of Moss' coalgebraic logic.

Suppose that $\mathfrak{F} = \langle X, R\rangle$ is a Kripke frame, where $R \subseteq X \times X$, and suppose we have a nabla-expression

$$\Box(\bigvee_{i=1}^{n} \varphi_i) \wedge \bigwedge_{i=1}^{n} \Diamond \varphi_i.$$

For simplicity, we assume $\varphi_1, \dots, \varphi_n$ are closed formulas, i.e., they contain no proposition letters. What is the semantics of our nabla-expression? If $x \in X$, then

$$x \Vdash_{\mathfrak{F}} \Box(\bigvee_{i=1}^{n} \varphi_i) \wedge \bigwedge_{i=1}^{n} \Diamond \varphi_i$$

if and only if

$$\forall y \in R[x], \exists i \leq n, y \Vdash_{\mathfrak{F}} \varphi_i \text{ and } \forall i \leq n, \exists y \in R[x], y \Vdash_{\mathfrak{F}} \varphi_i, \tag{13}$$

where $R[x]$ is the set of R-successors of x. If we view $\Vdash_{\mathfrak{F}}$ as a binary relation between X and the set of all closed modal formulas, then we can abbreviate (13) as follows:

$$R[x]\ (\Vdash_{\mathfrak{F}})_{EM}\ \{\varphi_1, \dots, \varphi_n\},$$

where $(\cdot)_{EM}$ stands for taking the *Egli-Milner* lifting of a binary relation (see (8) in §3.3).

Guided by this observation, we now consider a variant of the standard modal language in which we take the ∇ modality to be a *primitive* modality, with the following semantics on a given Kripke frame $\mathfrak{F}$:

$$x \Vdash_{\mathfrak{F}} \nabla\{\varphi_1, \ldots, \varphi_n\} \text{ iff } R[x]\ (\Vdash_{\mathfrak{F}})_{EM}\ \{\varphi_1, \ldots, \varphi_n\}. \tag{14}$$

It is not hard to verify that using (14),

$$x \Vdash_{\mathfrak{F}} \Box\varphi \text{ iff } x \Vdash_{\mathfrak{F}} \nabla\{\varphi\} \vee \nabla\emptyset,$$

and that

$$x \Vdash_{\mathfrak{F}} \Diamond\varphi \text{ iff } x \Vdash_{\mathfrak{F}} \nabla\{\varphi, \top\}.$$

What makes the reformulation interesting is that the semantics (14) allows for *coalgebraic generalizations*. As we will see, the key for turning the above observation about Kripke frames into a logical language and semantics for more general coalgebras is to use *relation lifting*.

4.2.1 Coalgebra and coalgebraic modal logic

The theory of *coalgebra* aims to provide a general mathematical framework for the study of state-based evolving systems. Given an endofunctor T on the category **Set** of sets with functions, we already saw the definition of a coalgebra of type T, or briefly: a T-coalgebra, as a pair (S, σ) where S is some set and $\sigma\colon S \to \mathrm{T}S$. The set S is called the *carrier* of the coalgebra, elements of which are called *states*; σ is called the *transition map* of the coalgebra. A T-coalgebra morphism between coalgebras $\sigma\colon S \to \mathrm{T}S$ and $\sigma'\colon S' \to \mathrm{T}S'$ is simply a function $f\colon S \to S'$ such that $\mathrm{T}f \circ \sigma = \sigma' \circ f$.

$$\begin{tikzcd} S \arrow[r, "\sigma"] \arrow[d, "f"'] & \mathrm{T}S \arrow[d, "\mathrm{T}f"] \\ S' \arrow[r, "\sigma'"'] & \mathrm{T}S' \end{tikzcd}$$

The coalgebraic approach to state-based systems combines mathematical simplicity with wide applicability: many features of computation, such as input, output, non-determinism, probability or interaction between agents, can be encoded in the functor T. Examples of coalgebras are Kripke frames, Kripke models, deterministic automata, topologies (with continuous open maps), and Markov chains.

The key notion of equivalence in coalgebra is that of two states s and s' in coalgebras (S, σ) and (S', σ') being *behaviorally equivalent*, notation: $(S, \sigma), s \simeq (S', \sigma'), s'$; this relation holds if there are coalgebra morphisms f, f' with a common codomain such that $f(s) = f'(s')$. As the name suggests, behaviorally equivalent states are considered to display the same behavior, and hence, to be essentially the same.

Coalgebraic logics are designed and studied in order to reason formally about coalgebras and their behavior; one of the main applications of this approach is the design of specification and verification languages for coalgebras. An (abstract) coalgebraic logic is a pair $(\mathscr{L}, \Vdash^{\mathscr{L}})$ such that $\mathscr{L}$ is a set of *formulas* and $\Vdash$ is a collection of relations associating with each T-coalgebra (S,σ) a binary relation $\Vdash^{\mathscr{L}}_{(S,\sigma)} \subseteq S \times \mathscr{L}$. If $s \Vdash^{\mathscr{L}}_{(S,\sigma)} \varphi$ we say that the formula φ is *true* or *satisfied* at s in (S,σ), and we will often write $(S,\sigma), s \Vdash \varphi$.

A natural criterion for a coalgebraic logic is that it cannot make a distinction between behaviorally equivalent states. A formula φ is *behaviorally invariant* if for all pairs of behaviorally equivalent pointed coalgebras $(S,\sigma), s \simeq (S',\sigma'), s'$ it holds that $(S,\sigma), s \Vdash \varphi \iff (S',\sigma'), s' \Vdash \varphi$. A coalgebraic language is *adequate* if all of its formulas are behaviorally invariant. An example of an adequate language is classical modal logic interpreted on P-coalgebras, i.e., on Kripke frames.

Given the prominence of Kripke frames and models as examples of coalgebras, it is not surprising to see that standard modal logic can be suitably generalized to provide adequate coalgebraic logics for coalgebras of arbitrary type. There are in fact distinct ways to do this; here we will focus on the approach based on the notion of relation lifting.

4.2.2 Relation lifting

Relation lifting is nothing more than a particular way of extending a coalgebra type functor $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ to a functor $\overline{\mathrm{T}}\colon \mathbf{Rel} \to \mathbf{Rel}$ on the category of sets and binary relations. For our purposes, we restrict attention to transition types that preserve *weak pullbacks*.

A *weak pullback* of two morphisms $f\colon X \to Z$ and $g\colon Y \to Z$ with a shared codomain Z is a pair of morphisms $p_X\colon P \to X$ and $p_Y\colon P \to Y$ with a shared domain P, such that (1) $f \circ p_X = g \circ p_Y$, and (2) for any other pair of morphisms $q_X\colon Q \to X$ and $q_Y\colon Q \to Y$ with $f \circ q_X = g \circ q_Y$, there is a morphism $q\colon Q \to P$ such that $p_X \circ q = q_X$ and $p_Y \circ q = q_Y$. This pullback is "weak" because we are not requiring q to be unique.

$$
\begin{tikzcd}
Q \arrow[dr, dotted, "q"] \arrow[drr, bend left, "q_Y"] \arrow[ddr, bend right, "q_X"'] & & \\
& P \arrow[r, "p_Y"] \arrow[d, "p_X"'] & Y \arrow[d, "g"] \\
& X \arrow[r, "f"'] & Z
\end{tikzcd}
$$

Saying that $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ preserves weak pullbacks means that if $p_X\colon P \to X$ and $p_Y\colon P \to Y$ form a weak pullback of $f\colon X \to Z$ and $g\colon Y \to Z$, then $\mathrm{T}p_X\colon \mathrm{T}P \to \mathrm{T}X$ and $\mathrm{T}p_Y\colon \mathrm{T}P \to \mathrm{T}Y$ form a weak pullback of $\mathrm{T}f\colon \mathrm{T}X \to \mathrm{T}Z$ and $\mathrm{T}g\colon \mathrm{T}Y \to \mathrm{T}Z$. Examples of weak pullback-preserving endofunctors on the category of sets include the identity functor, constant functors, the covariant powerset functor, the multiset functor, the distribution functor, and finite products and sums of such functors.

We will now define the notion of relation lifting.

Definition 8. Let $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ be a weak pullback-preserving functor, and let $R \subseteq X \times Y$ be a binary relation between sets X and Y. We denote the left and right projections of R as $\pi_X\colon R \to X$ and $\pi_Y\colon R \to Y$, respectively. Let $\alpha \in \mathrm{T}X$ and $\beta \in \mathrm{T}Y$; we now define

$$\alpha\,\overline{\mathrm{T}}R\,\beta :\Leftrightarrow \exists \delta \in \mathrm{T}R, \mathrm{T}\pi_X(\delta) = \alpha \text{ and } \mathrm{T}\pi_Y(\delta) = \beta.$$

We call $\overline{\mathrm{T}}R$ the T-*lifting of* R.

Observe that $\overline{\mathrm{T}}R$ is simply the binary relation between $\mathrm{T}X$ and $\mathrm{T}Y$ induced by the span

$$\mathrm{T}X \xleftarrow{\mathrm{T}\pi_X} \mathrm{T}R \xrightarrow{\mathrm{T}\pi_Y} \mathrm{T}Y.$$

Example 1. Recall that the covariant powerset functor is an example of a weak pullback-preserving functor. Now for any binary relation $R \subseteq X \times Y$, the P-lifting of R is precisely the Egli-Milner lifting $R_{EM} \subseteq \mathrm{P}X \times \mathrm{P}Y$. In other words, if $\alpha \in \mathrm{P}X$ and $\beta \in \mathrm{P}Y$, then $\alpha\,\overline{\mathrm{P}}R\,\beta$ iff

$$\forall x \in \alpha, \exists y \in \beta \text{ s.t. } x\,R\,y \text{ and } \forall y \in \beta, \exists x \in X \text{ s.t. } x\,R\,y.$$

Recall that **Set** can be embedded in the category **Rel** of sets and binary relations, using the functor $\mathrm{Graph}\colon \mathbf{Set} \to \mathbf{Rel}$, defined as $\mathrm{Graph}\colon X \mapsto X$ and

$$\mathrm{Graph}\colon X \xrightarrow{f} Y \mapsto \{(x, f(x)) \mid x \in X\},$$

where we view the right-hand-side above as a binary relation between X and Y. The desired property that turns $\overline{\mathrm{T}}$ into a *lifting* is that it makes the following diagram commute:

$$\begin{tikzcd}
\mathbf{Rel} \arrow[r, "\overline{\mathrm{T}}"] & \mathbf{Rel} \\
\mathbf{Set} \arrow[u, "\mathrm{Graph}"] \arrow[r, "\mathrm{T}"'] & \mathbf{Set} \arrow[u, "\mathrm{Graph}"']
\end{tikzcd}$$

The condition that the functor T preserves weak pullbacks is needed to ensure that $\overline{\mathrm{T}}$ is indeed a *functor*.

Fact 16. *Let* $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ *be a functor. Then* $\overline{\mathrm{T}}$ *is a lifting in the sense described above, that is:*

$$\mathrm{Graph} \circ \mathrm{T} = \overline{\mathrm{T}} \circ \mathrm{Graph}.$$

Moreover, $\overline{\mathrm{T}}$ *is a functor on* **Rel***, the category of sets and binary relations, iff* T *preserves weak pullbacks.*

4.2.3 Syntax and semantics of Moss' coalgebraic logic

We will now present the syntax and semantics of Moss' coalgebraic logic for an arbitrary weak pullback-preserving functor $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$. We will make additional assumptions about T. Firstly, we assume that T is *standard*; in the case that T preserves weak pullbacks we can take this to mean that T preserves *inclusions* (that is, if $\iota : X \hookrightarrow Y$ is an inclusion map, then $\mathrm{T}\iota : \mathrm{T}X \hookrightarrow \mathrm{T}Y$ is the inclusion map witnessing that $\mathrm{T}X$ is a subset of $\mathrm{T}Y$). This assumption is innocuous from the viewpoint of **Set**-coalgebras, because for any $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ there is a standard $\mathrm{T}'\colon \mathbf{Set} \to \mathbf{Set}$ such that the category of T-coalgebras is equivalent to the category of T'-coalgebras.

If we would leave it at this, only assuming $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ is standard and weak pullback-preserving, we could already define Moss' language, and indeed this is what he does in [31]. A downside of this approach is, however, that one might obtain formulas with *infinitely* many subformulas. This can be avoided by requiring that T satisfies the following condition for all sets X:

$$\mathrm{T}X = \bigcup\{\mathrm{T}X' \mid X' \subseteq X,\ X' \text{ finite}\}. \tag{15}$$

We say T is *finitary* if it satisfies (15).

If the coalgebra functor $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ one happens to be interested in is not finitary, this can be remedied. For each set X, we can define

$$\mathrm{T}_\omega X := \bigcup\{\mathrm{T}X' \mid X' \subseteq X,\ X' \text{ finite}\}.$$

Using the assumption that T is standard, this gives us a functor $\mathrm{T}_\omega\colon \mathbf{Set} \to \mathbf{Set}$. For the covariant powerset functor $\mathrm{P}\colon \mathbf{Set} \to \mathbf{Set}$, the above definition of T_ω yields precisely the finite powerset functor $\mathrm{P}_\omega\colon \mathbf{Set} \to \mathbf{Set}$.

From the general viewpoint of coalgebraic logic, one would want to consider both T and T_ω when understanding Moss' logic. Our current viewpoint, however, is focused on the Carioca derivation system, and there we only really need T_ω. To simplify our presentation and notation, we will therefore assume from here on that $\mathrm{T} = \mathrm{T}_\omega$, i.e., that T is finitary.

We will now define the finitary, Boolean version of Moss' coalgebraic language. Note that again for simplicity, we are working with the closed fragment.

Definition 9. Let $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ be a finitary, standard, weak pullback-preserving functor. We define $\mathscr{L}_\mathrm{T}$, the *closed (0-variable) Moss language for* T, to be the smallest set such that (1) $\top, \bot \in \mathscr{L}_\mathrm{T}$, (2) if $\varphi \in \mathscr{L}_\mathrm{T}$ then also $\neg\varphi \in \mathscr{L}_\mathrm{T}$, (3) if $\varphi, \psi \in \mathscr{L}_\mathrm{T}$ then also $\varphi \wedge \psi \in \mathscr{L}_\mathrm{T}$ and $\varphi \vee \psi \in \mathscr{L}_\mathrm{T}$, and (4) if $\alpha \in \mathrm{T}\mathscr{L}_\mathrm{T}$, then $\nabla\alpha \in \mathscr{L}_\mathrm{T}$.

The *coalgebraic semantics of* $\mathscr{L}_\mathrm{T}$ is defined as follows. Suppose we have a T-coalgebra $\sigma\colon S \to \mathrm{T}S$; we will define a satisfaction relation $\Vdash_\sigma$ between S (the set of states of our coalgebra) and $\mathscr{L}_\mathrm{T}$. Let $x \in S$; then we inductively define

1. $x \Vdash_\sigma \top$ and $x \nVdash_\sigma \bot$;
2. For all $\varphi, \psi \in \mathscr{L}_\mathrm{T}$, $x \Vdash_\sigma \varphi \wedge \psi$ iff $x \Vdash_\sigma \varphi$ and $x \Vdash_\sigma \psi$ (and similarly for $\varphi \vee \psi$ and $\neg\varphi$);
3. For all $\alpha \in \mathrm{T}\left(\mathscr{L}_\mathrm{T}\right)$, $x \Vdash_\sigma \nabla\alpha$ iff $\sigma(x)\ \overline{\mathrm{T}}\left(\Vdash_\sigma\right)\ \alpha$.

Note that if we choose $\mathrm{T} = \mathrm{P}_\omega$, the finite powerset functor, then the semantics in Definition 9 gives us precisely the syntax and semantics for the nabla we saw above in (14), since $\mathrm{P}_\omega\left(\mathscr{L}_{\mathrm{P}_\omega}\right)$ is the collection of finite sets of $\mathscr{L}_{\mathrm{P}_\omega}$-formulas, and the P_ω-lifting of $\Vdash_\sigma$ is precisely the Egli-Milner lifting of $\Vdash_\sigma$.

Notes

A classic reference for the theory of coalgebras is Rutten [35]. For a recent overview to the area of coalgebraic logic, with pointers to introductory literature, we suggest [13] or [30]. The idea to use nabla as a primitive modality plays an important role in the work of both Barwise and Moss [8] and Janin and Walukiewicz [22]. The idea to use nabla as a coalgebraic modality is due to Moss [31]. For a more detailed discussion of the material in this subsection, including detailed proofs of the technical results, we refer to [29]. The observation in Fact 16, that $\overline{\mathrm{T}}$ is a functor on **Rel** iff T preserves weak pullbacks, goes back to Trnková [37].

4.3 The Carioca derivation system

We will now introduce the Carioca derivation system. The aim of this derivation system is to enable us to derive exactly those *inequalities* of formulas in Moss' language, that are valid on all T-coalgebras. In order to state the axioms and rules of the Carioca system, we will first have to introduce two new concepts: lifted conjunctions and disjunctions, and slim redistributions.

The inequalities we are considering are those of the form $\varphi \preccurlyeq \psi$, for $\varphi, \psi \in \mathscr{L}_{\mathrm{T}}$. We say $\varphi \preccurlyeq \psi$ is valid on a coalgebra $\sigma\colon S \to \mathrm{T}S$ if for all $x \in S$ such that $x \Vdash_\sigma \varphi$, it is also the case that $x \Vdash_\sigma \psi$. If $\varphi \preccurlyeq \psi$ is valid on *all* T-coalgebras, we write $\varphi \Vdash_{\mathrm{T}} \psi$.

When writing the Carioca axioms, we think of the formation of disjunctions and conjunctions as *functions* from $\mathrm{P}_\omega \mathscr{L}_{\mathrm{T}}$ to $\mathscr{L}_{\mathrm{T}}$, i.e., one can consider the maps $\bigvee\colon \mathrm{P}_\omega \mathscr{L}_{\mathrm{T}} \to \mathscr{L}_{\mathrm{T}}$ and $\bigwedge\colon \mathrm{P}_\omega \mathscr{L}_{\mathrm{T}} \to \mathscr{L}_{\mathrm{T}}$ as maps in **Set**. Consequently, we can also apply T to $\bigvee$ and $\bigwedge$, which gives us maps

$$\mathrm{T}\bigvee\colon \mathrm{T}\mathrm{P}_\omega \mathscr{L}_{\mathrm{T}} \to \mathrm{T}\mathscr{L}_{\mathrm{T}} \text{ and } \mathrm{T}\bigwedge\colon \mathrm{T}\mathrm{P}_\omega \mathscr{L}_{\mathrm{T}} \to \mathrm{T}\mathscr{L}_{\mathrm{T}}.$$

If $\Phi \in \mathrm{T}\mathrm{P}_\omega \mathscr{L}_{\mathrm{T}}$, we call $\mathrm{T}\bigvee(\Phi)$ and $\mathrm{T}\bigwedge(\Phi)$ a T-*lifted* disjunction and conjunction, respectively.

Example 2. If $\mathrm{T} = \mathrm{P}_\omega$ and we apply the P_ω-lifted disjunction operation $\mathrm{P}_\omega\bigvee$ to an element $\Phi = \{S_1, \ldots, S_n\} \in \mathrm{P}_\omega \mathrm{P}_\omega(\mathscr{L}_{\mathrm{P}_\omega})$, we obtain a forward image:

$$\mathrm{P}_\omega \bigvee\big(\{S_1, \ldots, S_n\}\big) = \{\bigvee S_1, \ldots, \bigvee S_n\}.$$

The final concept we will now introduce is that of *slim redistributions*.

Definition 10. Let T: **Set** $\to$ **Set** be a finitary, standard, weak pullback-preserving functor and let X be a set. If $\alpha \in \mathrm{T}X$, then we define the *base* of α to be the following intersection:

$$\mathrm{Base}(\alpha) := \bigcap\big\{X' \subseteq X \mid \alpha \in \mathrm{T}X'\big\}.$$

Now if $C \in \mathrm{P}_\omega \mathrm{T}X$ is a finite collection of elements of $\mathrm{T}X$, then we define a *slim redistribution of* C to be an element Ψ such that

$$\Psi \in \mathrm{T}\mathrm{P}_\omega\left(\textstyle\bigcup_{\alpha\in C} \mathrm{Base}(\alpha)\right) \text{ and for all } \alpha \in C,\ \alpha\, \overline{\mathrm{T}} \in \Psi.$$

We denote the set of all slim redistributions of C by $\mathrm{SRD}_{\mathrm{T}}(C)$.

Intuitively, the idea is that $\mathrm{Base}(\alpha)$ is the smallest set $X' \subseteq X$ such that $\alpha \in \mathrm{T}X'$.

Fact 17. *Let* $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ *be a finitary, standard, weak pullback-preserving functor and let* X *be a set. Then for all* $\alpha \in \mathrm{T}X$ *and all sets* Y *it holds that* $\alpha \in \mathrm{T}Y$ *iff* $\mathrm{Base}(\alpha) \subseteq Y$. *In fact,* Base *is a natural transformation from* T *to* P_ω.

Example 3. In the case that $\mathrm{T} = \mathrm{P}_\omega$, Definition 10 can be simplified as follows. Firstly, if $\mathrm{T} = \mathrm{P}_\omega$ and $\alpha \in \mathrm{P}_\omega X$ is simply a finite subset of X, then the smallest subset $X' \subseteq X$ such that $\alpha \in \mathrm{T}X'$ is α itself; in other words, Base is the identity if $\mathrm{T} = \mathrm{P}_\omega$.

Secondly, if $C \in \mathrm{P}_\omega \mathrm{P}_\omega X$ is a finite collection of finite subsets of X, then

$$\mathrm{SRD}_{\mathrm{P}_\omega}(C) = \left\{\Psi \in \mathrm{P}_\omega \mathrm{P}_\omega(\bigcup C) \mid \forall \alpha \in C,\, \alpha\, \overline{\mathrm{P}}{\in}\, \Psi\right\},$$

where $\overline{\mathrm{P}}{\in}$ is the Egli-Milner lifting of the element relation, viewed as a binary relation $\in \subseteq X \times \mathrm{P}_\omega X$. It is now not hard to see that if $\Psi = \{S_1, \ldots, S_m\}$ then

$$\alpha\, \overline{\mathrm{P}}{\in}\, \{S_1, \ldots, S_m\} \text{ iff } \alpha \subseteq \textstyle\bigcup_{i=1}^m S_i \text{ and } \forall i \leq m, \alpha \cap S_i \neq \emptyset.$$

Thus, for $\mathrm{T} = \mathrm{P}_\omega$, we see that Ψ is a slim redistribution of $C \in \mathrm{P}_\omega \mathrm{P}_\omega X$ iff

$$\Psi \in \mathrm{P}_\omega \mathrm{P}_\omega X \text{ such that } \bigcup C = \bigcup \Psi \text{ and } \forall \alpha \in C, \forall S \in \Psi,\, \alpha \cap S \neq \emptyset. \tag{16}$$

We are now ready to define the Carioca derivation system. For the sake of simplicity, we will present a simplified version in which all disjunctions and conjunctions are finite. The simplification we use to achieve this, is to assume that T maps finite sets to finite sets.

Definition 11. Let $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ be a standard, finitary, weak pullback-preserving functor. Additionally, we assume that $\mathrm{T}X$ is finite whenever X is finite. The *Carioca derivation system* $\vdash_{\mathrm{T}}$ consists of a complete set of axioms and rules for all Boolean inequalities, combined with the following rule and axioms:

$$\begin{array}{lll}
(\nabla 1) & \alpha\, \overline{\mathrm{T}}{\preccurlyeq}\, \beta \Rightarrow\, \vdash_{\mathrm{T}} \nabla\alpha \preccurlyeq \nabla\beta & (\alpha, \beta \in \mathrm{T}\mathscr{L}_{\mathrm{T}}); \\
(\nabla 2) & \vdash_{\mathrm{T}} \bigwedge_{\alpha \in C} \nabla\alpha \preccurlyeq \bigvee\left\{\nabla \mathrm{T}\bigwedge(\Psi) \mid \Psi \in \mathrm{SRD}_{\mathrm{T}}(C)\right\} & (C \in \mathrm{P}_\omega \mathrm{T}\mathscr{L}_{\mathrm{T}}); \\
(\nabla 3.f) & \vdash_{\mathrm{T}} \nabla \mathrm{T}\bigvee(\Phi) \preccurlyeq \bigvee\left\{\nabla\beta \mid \beta\, \overline{\mathrm{T}}{\in}\, \Phi\right\} & (\Phi \in \mathrm{T}\mathrm{P}_\omega \mathscr{L}_{\mathrm{T}}).
\end{array}$$

(Note that it is provable in $\vdash_{\mathrm{T}}$ that $(\nabla 2)$ and $(\nabla 3)$ are in fact equations rather than inequalities.)

Example 4. We will make the rule and axioms above more concrete for the case $\mathrm{T} = \mathrm{P}_\omega$.

Starting with $(\nabla 1)$, suppose that $\alpha, \beta \in \mathrm{P}_\omega \mathscr{L}_{\mathrm{P}_\omega}$ are finite sets of formulas. Rule $(\nabla 1)$ says that in case that $\alpha\, \overline{\mathrm{P}}{\preccurlyeq}\, \beta$, i.e., in case that $\alpha \preccurlyeq_{EM} \beta$, then $\vdash_{\mathrm{P}_\omega} \nabla\alpha \preccurlyeq \nabla\beta$. In other words, if

$$\forall \varphi \in \alpha, \exists \psi \in \beta \text{ s.t. } \vdash_{\mathrm{P}_\omega} \varphi \preccurlyeq \psi \text{ and } \forall \psi \in \beta, \exists \varphi \in \alpha \text{ s.t. } \vdash_{\mathrm{P}_\omega} \varphi \preccurlyeq \psi, \tag{17}$$

then $\vdash_{\mathrm{P}_\omega} \nabla\alpha \preccurlyeq \nabla\beta$. Intuitively, this means that we can derive that $\nabla\beta$ is $\preccurlyeq$-related to (follows from) $\nabla\alpha$, provided that the elements of β are $\preccurlyeq$-related to those of α in an "Egli-Milner" way.

Moving on to $(\nabla 2)$, we see that if $C \in \mathrm{P}_\omega \mathrm{P}_\omega \mathscr{L}_{\mathrm{P}_\omega}$ is a finite collection of finite sets of formulas, then

$$\vdash_{\mathrm{P}_\omega} \bigwedge_{\alpha \in C} \nabla\alpha \preccurlyeq \bigvee \left\{ \nabla\{\bigwedge S_1, \ldots, \bigwedge S_n\} \mid \{S_1, \ldots, S_n\} \in \mathrm{SRD}_{\mathrm{P}_\omega}(C) \right\}, \tag{18}$$

where we refer the reader to (16) for a description of $\mathrm{SRD}_{\mathrm{P}_\omega}(C)$. Intuitively, this means that any conjunction of ∇-formulas is equivalent to a disjunction of nablas of lifted conjunctions.

Finally, $(\nabla 3.f)$ can be simplified as follows. Suppose we have $\Phi = \{S_1, \ldots, S_n\} \in \mathrm{P}_\omega \mathrm{P}_\omega(\mathscr{L}_{\mathrm{P}_\omega})$; then $(\nabla 3.f)$ boils down to the axiom

$$\vdash_{\mathrm{P}_\omega} \nabla\{\bigvee S_1, \ldots, \bigvee S_n\} \preccurlyeq \bigvee \left\{ \nabla\beta \mid \beta \subseteq \bigcup_{i=1}^n S_i \text{ and } \forall i \leq n, \beta \cap S_i \neq \emptyset \right\}. \tag{19}$$

As a further simplification, one could also write the following:

$$\vdash_{\mathrm{P}_\omega} \nabla(\alpha \cup \{\bigvee S\}) \preccurlyeq \bigvee \left\{ \nabla(\alpha \cup \beta) \mid \beta \subseteq S \text{ and } \beta \neq \emptyset \right\}. \tag{20}$$

One can inductively derive (19) from (20). Regardless of how we look at $(\nabla 3.f)$, the intuitive content of this axiom is that finite disjunctions "under" nablas distribute to disjunctions of nablas.

Fact 18. *Let* $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ *be a standard, finitary, weak pullback-preserving functor, with the added property that* $\mathrm{T}X$ *is finite whenever* X *is finite.*

The Carioca derivation system for T *is sound and complete with respect to* T*-validity: for all* $\varphi, \psi \in \mathscr{L}_{\mathrm{T}}$,

$$\vdash_{\mathrm{T}} \varphi \preccurlyeq \psi \textit{ iff } \varphi \Vdash_{\mathrm{T}} \psi.$$

Notes

A first axiomatization of the nabla modality (in the power set case) was given by Palmigiano and Venema [33]; this calculus was streamlined by Bílková et al. [12] into a formulation admitting a generalization to the arbitrary case in the Carioca system. (The name 'Carioca' refers to the fact that this version of the axiomatization was formulated in Rio de Janeiro.) Fact 18, the completeness of the Carioca system, was proved by Kupke et al. [28, 29]; the latter work also contains a discussion (with proof) of Fact 17.

4.4 The T-powerlocale

Having acquainted ourselves with Moss' coalgebraic logic and the Carioca derivation system, we now introduce the T-powerlocale construction. This is a generalization of the Vietoris powerlocale construction, using techniques from coalgebraic logic. Because the Carioca axioms are parametric in their coalgebra type functor T, so is the T-powerlocale construction. We will see that certain properties of the Vietoris functor can be proved at the more general level of the T-powerlocale, and as a corollary, we show how the Vietoris powerlocale can be presented using nablas as generators, rather than boxes and diamonds. Recall that locales have finite meets and arbitrary joins; in a locale $\mathbb{A}$ we represent these maps as $\bigwedge : \mathrm{P}_\omega A \to A$ and $\bigvee : \mathrm{P}A \to A$, respectively.

Definition 12. Let $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ be a standard, finitary, weak pullback-preserving functor and let $\mathbb{A}$ be a locale with an underlying set of opens A. We define $\mathrm{V}_\mathrm{T}\,\mathbb{A}$, the T-*powerlocale of* $\mathbb{A}$, to be the locale generated by the set $\{\nabla\alpha \mid \alpha \in \mathrm{T}A\}$, subject to the following relations:

$$
\begin{array}{lll}
(\nabla 1) & \nabla\alpha \leq \nabla\beta \text{ if } \alpha\,\overline{\mathrm{T}}{\leq}\,\beta & (\alpha,\beta \in \mathrm{T}A);\\
(\nabla 2) & \bigwedge_{\alpha\in C}\nabla\alpha \leq \bigvee\left\{\nabla\mathrm{T}\bigwedge(\Psi) \mid \Psi \in \mathrm{SRD}_\mathrm{T}(C)\right\} & (C \in \mathrm{P}_\omega\mathrm{T}A);\\
(\nabla 3) & \nabla\mathrm{T}\bigvee(\Phi) \leq \bigvee\left\{\nabla\beta \mid \beta\,\overline{\mathrm{T}}{\in}\,\Phi\right\} & (\Phi \in \mathrm{TP}A).
\end{array}
$$

Note that the only real difference between the Carioca axioms in Definition 11 and the relations in Definition 12 above is the difference between $(\nabla 3.f)$ and $(\nabla 3)$. We will later see how this corresponds to the difference between finite disjunctions (as found in Boolean algebras and distributive lattices) and infinite disjunctions (as found in locales).

Fact 19. *The construction described in Definition 12 defines a functor* $\mathrm{V}_\mathrm{T}\colon \mathbf{Loc} \to \mathbf{Loc}$.

The functor V_T we have just introduced has several additional properties which can be proved at an abstract level. As an example, note the following fact.

Fact 20. *Let* $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ *be a standard, finitary, weak pullback-preserving functor.*

1. *The functor* $\mathrm{V}_\mathrm{T}\colon \mathbf{Loc} \to \mathbf{Loc}$ *preserves regularity.*
2. *If we further assume that* $\mathrm{T}X$ *is finite for every finite set* X*, then* V_T *preserves the combination of compactness and zero-dimensionality.*

In the introduction of §4, we motivated our discussion of nabla expressions and Moss' coalgebraic logic with the question whether we could describe the Vietoris powerlocale using nabla. The following fact asserts that the Carioca axioms indeed allow us to do this.

Fact 21. *The* P_ω*-powerlocale is the Vietoris powerlocale.*

Proof sketch. Suppose that $\mathbb{A}$ is a locale; we must show that $\mathrm{V}\,\mathbb{A} \simeq \mathrm{V}_{\mathrm{P}_\omega}\,\mathbb{A}$. This is achieved by defining frame morphisms in both directions, and showing that these morphisms are mutually inverse. From $\mathrm{V}_{\mathrm{P}_\omega}\,\mathbb{A}$ to $\mathrm{V}\,\mathbb{A}$, we send

$$\nabla\alpha \mapsto \Box\left(\bigvee\alpha\right) \wedge \bigwedge_{a\in\alpha} \Diamond a.$$

From $\mathrm{V}\,\mathbb{A}$ to $\mathrm{V}_{\mathrm{P}_\omega}\,\mathbb{A}$, we use the following assignments:

$$\Diamond a \mapsto \nabla\{a,1\} \text{ and } \Box a \mapsto \nabla\{a\} \vee \nabla\emptyset.$$

For details about the rest of the proof, we refer the reader to [33] or [38]. $\Box$

Below in Fact 23 we will look in more detail at nabla presentations of the Vietoris powerlocale. Before we do so, however, we will introduce an alternative presentation of V_{T}, in which we exploit the fact that in the language of locales we can use *infinite* disjunctions.

Fact 22. *The relation* $(\nabla 2)$ *in Definition 12 can equivalently be replaced by the following pair of relations:*

$$(\nabla 2.0) \quad 1 \leq \bigvee\left\{\nabla\alpha \mid \alpha \in \mathrm{T}A\right\};$$
$$(\nabla 2.2) \quad \nabla\alpha \wedge \nabla\beta \leq \bigvee\left\{\nabla\gamma \mid \gamma\,\overline{\mathrm{T}}\!\leq \alpha \textit{ and } \gamma\,\overline{\mathrm{T}}\!\leq \beta\right\}.$$

Note that the suffixes ".0" and ".2" indicate nullary and binary conjunctions, respectively. Returning to the case $\mathrm{T} = \mathrm{P}_\omega$, we will now give a concrete nabla-presentation of V.

Fact 23. *Let* $\mathbb{A}$ *be a locale. We can present* $\mathrm{V}\,\mathbb{A}$*, the Vietoris powerlocale of* $\mathbb{A}$*, as the locale generated by the set* $\{\nabla\alpha \mid \alpha \in \mathrm{P}_\omega A\}$*, subject to the following relations:*

$$(\nabla 1) \qquad \nabla\alpha \leq \nabla\beta \quad (\textit{if } \alpha \leq_{EM} \beta);$$
$$(\nabla 2) \qquad \bigwedge_{\alpha\in C}\nabla\alpha \leq \bigvee\left\{\nabla\{\bigwedge S_1,\ldots,\bigwedge S_n\} \mid \{S_1,\ldots,S_n\} \in \mathrm{SRD}_{\mathrm{P}_\omega}(C)\right\},$$

where C ranges over the finite subsets of $\mathrm{P}_\omega A$*, also see* (16)*; and*

$$(\nabla 3) \qquad \nabla\{\bigvee S_1,\ldots,\bigvee S_n\} \leq \bigvee\left\{\nabla\beta \mid \beta \subseteq \bigcup_{i\leq n} S_i \textit{ and } \forall i \leq n,\, \beta \cap S_i \neq \emptyset\right\},$$

where the S_i *range over (possibly infinite) subsets of A. Moreover, the* $(\nabla 2)$ *relation can be replaced by the following pair of relations:*

$$(\nabla 2.0) \qquad 1 \leq \bigvee\left\{\nabla\alpha \mid \alpha \in \mathrm{P}_\omega A\right\};$$
$$(\nabla 2.2) \qquad \nabla\alpha \wedge \nabla\beta \leq \bigvee\left\{\nabla\gamma \mid \gamma \leq_{EM} \alpha \textit{ and } \gamma \leq_{EM} \beta\right\},$$

and the $(\nabla 3)$ *relation can be replaced by the following inductive version:*

$$(\nabla 3.ind) \quad \nabla(\alpha \cup \{\bigvee S\}) \leq \bigvee\left\{\nabla(\alpha\cup\beta) \mid \beta \in \mathrm{P}_\omega S \textit{ and } \beta \neq \emptyset\right\} \quad (S \in \mathrm{P}A).$$

We have now seen how to present the Vietoris powerlocale, and more generally the T-powerlocale, using nablas. We can improve on this still, by showing that "every element of $\mathrm{V}_{\mathrm{T}}\,\mathbb{A}$ is a disjunction of nablas" in a rather strong sense. Recall that a *suplattice* is a complete join-semilattice, and that any locale has an underlying suplattice.

Definition 13. Let $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ be a standard, finitary, weak pullback-preserving functor and let $\mathbb{L}$ be a suplattice. We define $\mathrm{W}_{\mathrm{T}}\,\mathbb{L}$, the T-*powerlattice of* $\mathbb{L}$, to be the suplattice generated by the $\{\nabla\alpha \mid \alpha \in \mathrm{T}L\}$, subject to the following relations:

$$
\begin{array}{lll}
(\nabla 1) & \nabla\alpha \leq \nabla\beta \text{ if } \alpha \, \overline{\mathrm{T}}\!\leq \beta & (\alpha, \beta \in \mathrm{T}L); \\
(\nabla 3) & \nabla \mathrm{T}\bigvee(\Phi) \leq \bigvee\left\{\nabla\beta \mid \beta \, \overline{\mathrm{T}}\!\in \Phi\right\} & (\Phi \in \mathrm{TP}L).
\end{array}
$$

If we now let U denote the (contravariant) forgetful functor from **Loc** to **SupLat**, the category of suplattices and suplattice morphisms, we can draw the following picture:

$$
\begin{tikzcd}
\mathbf{Loc} \arrow[r, "\mathrm{V}_{\mathrm{T}}"] \arrow[d, "\mathrm{U}"'] & \mathbf{Loc} \arrow[d, "\mathrm{U}"] \\
\mathbf{SupLat} \arrow[r, "\mathrm{W}_{\mathrm{T}}"'] & \mathbf{SupLat}
\end{tikzcd}
$$

We would like to emphasize that the following result, like Facts 19, 20 and 22, holds not only for the Vietoris powerlocale but for the T-powerlocale in general.

Fact 24. *Let* $\mathrm{T}\colon \mathbf{Set} \to \mathbf{Set}$ *be a standard, finitary, weak pullback-preserving functor. Then there exists a natural transformation such that for all locales* $\mathbb{A}$, $\mathrm{U}(\mathrm{V}_{\mathrm{T}}\,\mathbb{A}) \simeq \mathrm{W}_{\mathrm{T}}(\mathrm{U}\,\mathbb{A})$.

The proof of Fact 24 uses flat sites, a technique from formal topology [14], which is meant to capture the notion of a *basis* of a topological space. From a logical viewpoint, Fact 24 tells us that (1) any $(\wedge, \bigvee)$-formula in Moss' coalgebraic language for T is $((\nabla 1), (\nabla 2), (\nabla 3))$-equivalent to a $\bigvee$-formula, and that (2) for any inequality between $\bigvee$-formulas derived using $((\nabla 1), (\nabla 2), (\nabla 3))$, there is a $((\nabla 1), (\nabla 3))$-derivation which proves that inequality.

Notes

The T-powerlocale was introduced by Venema, Vosmaer and Vickers in [38]; this is also where one can find the above results. (An early version can be found in [43, Ch. 5].) For more information on the method of using sup-lattices to obtain results like our Fact 24 the reader is referred to [25].

Acknowledgements We would like to thank Liang-Ting Chen and Steve Vickers for providing helpful pointers to the literature. Additionally, we would like to thank the anonymous referee, whose thorough criticism contributed significantly to the quality of this paper, and in particular to the presentation and choice of contents in §3.2.3. Finally, we are grateful to Guram Bezhanishvili, who provided many suggestions for improving the presentation of the chapter.

References

1. S. Abramsky. A domain equation for bisimulation. *Inform. and Comput.*, 92(2):161–218, 1991.
2. S. Abramsky. Domain theory in logical form. *Ann. Pure Appl. Logic*, 51(1-2):1–77, 1991. Second Annual IEEE Symposium on Logic in Computer Science (Ithaca, NY, 1987).
3. S. Abramsky. A cook's tour of the finitary non-well-founded sets. In Sergei Artemov, Howard Barringer, Artur d'Avila Garcez, Luis C. Lamb, and John Woods, editors, *We Will Show Them: Essays in honour of Dov Gabbay*, volume 1, pages 1–18. College Publications, 2005.
4. S. Abramsky and A. Jung. Domain theory. In *Handbook of logic in computer science*, volume 3, pages 1–168. Clarendon Press, Oxford, 1994.
5. P. Aczel. *Non-well-founded sets*, volume 14 of *CSLI Lecture Notes*. Stanford University Center for the Study of Language and Information, Stanford, CA, 1988.
6. C. Areces and R. I. Goldblatt, editors. *Advances in Modal Logic 7, Nancy 2008*. College Publications, 2008.
7. B. Banaschewksi and C.J. Mulvey. Stone-Cech compactification of locales. I. *Houston Journal of Mathematics*, 6:301–312, 1980.
8. J. Barwise and L. Moss. *Vicious circles*, volume 60 of *CSLI Lecture Notes*. CSLI Publications, Stanford, CA, 1996.
9. G. Bezhanishvili, N. Bezhanishvili, D. Gabelaia, and A. Kurz. Bitopological duality for distributive lattices and Heyting algebras. *Math. Structures Comput. Sci.*, 20(3):359–393, 2010.
10. G. Bezhanishvili, N. Bezhanishvili, and J. Harding. Modal compact Hausdorff spaces. *Journal of Logic and Computation*, 2012.
11. N. Bezhanishvili and A. Kurz. Free modal algebras: A coalgebraic perspective. In T. Mossakowski, U. Montanari, and M. Haveraaen, editors, *Algebra and Coalgebra in Computer Science*, volume 4624 of *Lecture Notes in Computer Science*, pages 143–157. Springer Berlin / Heidelberg, 2007.
12. M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality. In Areces and Goldblatt [6], pages 1–21.
13. C. Cîrstea, A. Kurz, D. Pattinson, L. Schröder, and Y. Venema. Modal logics are coalgebraic. *The Computer Journal*, 54:524–538, 2011.
14. T. Coquand, G. Sambin, J. Smith, and S. Valentini. Inductively generated formal topologies. *Ann. Pure Appl. Logic*, 124(1-3):71–106, 2003.
15. W. H. Cornish. On H. Priestley's dual of the category of bounded distributive lattices. *Mat. Vesnik*, 12(27)(4):329–332, 1975.
16. B. A. Davey and J. C. Galati. A coalgebraic view of Heyting duality. *Studia Logica*, 75(3):259–270, 2003.
17. L. L. Esakia. Topological Kripke models. *Dokl. Akad. Nauk SSSR*, 214:298–301, 1974.
18. K. Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2):229–237, 1975.
19. G. Gierz, K.H. Hofmann, K. Keimel, J.D. Lawson, M. Mislove, and D.S. Scott. *Continuous Lattices and Domains*. Cambridge University Press, 2003.
20. M. Hochster. Prime ideal structure in commutative rings. *Transactions of the American Mathematical Society*, 142:43–60, 1969.
21. J.R. Isbell. Atomless parts of spaces. *Mathematica Scandinavica*, 31:5–32, 1972.
22. D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Mathematical foundations of computer science 1995 (Prague)*, volume 969 of *Lecture Notes in Comput. Sci.*, pages 552–562. Springer, Berlin, 1995.
23. P. T. Johnstone. *Stone spaces*, volume 3 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1982.
24. P. T. Johnstone. Vietoris locales and localic semilattices. In *Continuous lattices and their applications (Bremen, 1982)*, volume 101 of *Lecture Notes in Pure and Appl. Math.*, pages 155–180. Dekker, New York, 1985.

25. A. Jung, M. A. Moshier, and S. J. Vickers. Presenting dcpos and dcpo algebras. In *Proceedings of the 24th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXIV)*, volume 218 of *Electronic Notes in Theoretical Computer Science*, pages 209–229, 2008.
26. M. Kegelmann. Continuous domains in logical form. *Electronic Notes in Theoretical Computer Science*, 49:1–166, 2002.
27. C. Kupke, A. Kurz, and Y. Venema. Stone coalgebras. *Theoret. Comput. Sci.*, 327(1-2):109–134, 2004.
28. C. Kupke, A. Kurz, and Y. Venema. Completeness of the finitary Moss logic. In Areces and Goldblatt [6], pages 193–217.
29. C. Kupke, K. Kurz, and Y. Venema. Completeness for the coalgebraic cover modality. *Logical Methods in Computer Science*, 8(3), 2012.
30. C. Kupke and D. Pattinson. Coalgebraic semantics of modal logics: an overview. *Theoretical Computer Science*, 412:5070–5094, 2011.
31. L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *Ann.P.Appl.Log.* 99:241–259, 1999).
32. A. Palmigiano. A coalgebraic view on positive modal logic. *Theoret. Comput. Sci.*, 327(1-2):175–195, 2004.
33. A. Palmigiano and Y. Venema. Nabla algebras and Chu spaces. In T. Mossakowski, U. Montanari, and M. Haveraaen, editors, *Algebra and Coalgebra in Computer Science. Second International Conference, CALCO 2007, Bergen, Norway, August 20-24, 2007*, volume 4624 of *LNCS*, pages 394–408. Springer Verlag, 2007.
34. E. Robinson. Power-domains, modalities and the Vietoris monad. Technical Report UCAM-CL-TR-98, University of Cambridge, Computer Laboratory, October 1986.
35. J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoret. Comput. Sci.*, 249(1):3–80, 2000. Modern algebra and its applications (Nashville, TN, 1996).
36. H. Simmons. The Vietoris modifications of a frame. Unpublished manuscript, 79pp., available online at `http://www.cs.man.ac.uk/~hsimmons/`, 2004.
37. V. Trnková. General theory of relational automata. *Fundamenta Informaticae*, 3(2):189–234, 1980.
38. Y. Venema, S. J. Vickers, and J. Vosmaer. Generalized powerlocales via relation lifting. *Mathematical Structures in Computer Science*, 23:142–199, 2013.
39. S. J. Vickers. *Topology via logic*, volume 5 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 1989.
40. S. J. Vickers. Constructive points of powerlocales. *Mathematical Proceedings of the Cambridge Philosophical Society*, 122:207–222, 1997.
41. S.J. Vickers. Information systems for continuous posets. *Theoretical Computer Science*, 114(2):201 – 229, 1993.
42. L. Vietoris. Bereiche zweiter Ordnung. *Monatsh. Math. Phys.*, 32(1):258–280, 1922.
43. J. Vosmaer. *Logic, Algebra and Topology. Investigations into canonical extensions, duality theory and point-free topology*. PhD thesis, Institute for Logic, Language and Computation, December 2010. Promotores: M. Gehrke and Y. Venema.