ON MINIMAL COALGEBRAS

H. PETER GUMM

Abstract. We define an out-degree for F-coalgebras and show that the coalgebras of outdegree at most κ form a covariety. As a subcategory of all F-coalgebras, this class has a terminal object, which for many problems can stand in for the terminal F-coalgebra, which need not exist in general. As examples, we derive structure theoretic results about minimal coalgebras, showing that, for instance minimization of coalgebras is functorial, that products of finitely many minimal coalgebras exist and are given by their largest common subcoalgebra, that minimal subcoalgebras have no inner endomorphisms and show how minimal subcoalgebras can be constructed from Moore-automata. Since the elements of minimal subcoalgebras must correspond uniquely to the formulae of any logic characterizing observational equivalence, we give in the last section a straightforward and self-contained account of the coalgebraic logic of D. Pattinson and L. Schröder, which we believe is simpler and more direct than the original exposition.

For every automaton $\mathcal{A}$ there exists a minimal automaton $\nabla(\mathcal{A})$, which displays the same behavior as $\mathcal{A}$. In the case of acceptors, this means that $\mathcal{A}$ and $\nabla(\mathcal{A})$ recognize the same language, and in the more general case of Moore-Automata it means that equal inputs generate the same outputs. Minimality, of course, refers to the cardinality of (the state set of) any automaton displaying the same behavior.

Turning to coalgebras, there are two possible notions of state equivalence, to begin with. Given two coalgebras $\mathcal{A}$ and $\mathcal{B}$ one may consider states $a \in A$ and $b \in B$ equivalent if they are *bisimilar*, or, alternatively, if they are *observationally equivalent*. Whereas these two notions agree for automata, they may differ for general coalgebras, unless the type functor F weakly preserves kernels, see [7]. In general, bisimilar states are observationally equivalent, but the converse need not hold. Observational equivalence, restricted to a single coalgebra, is a congruence relation, i.e. the kernel of some homomorphism, while bisimilarity may fail to be transitive.

For these reasons we choose observational equivalence as our notion of equivalence. It follows that for every coalgebra $\mathcal{A}$ there exists an equivalent minimal coalgebra $\nabla(\mathcal{A})$. Now $\mathcal{A}$ and $\mathcal{B}$ are observationally equivalent just in case $\nabla(\mathcal{A})$ and $\nabla(\mathcal{B})$ are isomorphic. Generalizing a result of Kianpi and Jugnia [10], we show that the class of all minimal coalgebras forms a full subcategory of the category of all F-coalgebras and ∇ is a functor, which is left adjoint to the inclusion functor.

This and similar results are most easily obtained when a terminal F-coalgebra $\mathcal{T}$ exists. In this case $\nabla(\mathcal{A})$ is isomorphic to the image of $\mathcal{A}$ under the unique homomorphism $\tau_{\mathcal{A}} : \mathcal{A} \to \mathcal{T}$ and arbitrary coalgebras $\mathcal{A}$ and $\mathcal{B}$ are equivalent iff their terminal images are identical, i.e. $\tau_{\mathcal{A}}[A] = \tau_{\mathcal{B}}[B]$. From this many results follow easily that, for instance,

- products of finitely many minimal coalgebras exist,
- the structure maps of minimal coalgebras must be injective,

- minimal coalgebras have no internal endomorphisms, etc.

Unfortunately, though, not every functor F admits a terminal coalgebra. This is due to Lambek's Lemma (c.f.[17]), which implies that the structure map $\tau : T \to F(T)$ of a terminal coalgebra $\mathcal{T}$ must be bijective. The powerset functor $\mathbb{P}$, for instance, does not admit a terminal coalgebra, as the equation $T = \mathbb{P}(T)$ has no solution in T.

In this case the natural rescue is to replace $\mathbb{P}$ by a bounded powerset functor $\mathbb{P}_\kappa$ for some sufficiently large κ and to realize that terminal coalgebras $\mathcal{T}_\kappa$ do exist for $\mathbb{P}_\kappa$. Such a replacement can be made for any $\mathcal{S}et$-functor by introducing the notion of out-degree for arbitrary coalgebras and we show how one obtains any arbitrary coalgebra from an appropriate Moore-automaton. In particular, any set of coalgebras is contained in the subcategory of all coalgebras with out-degree(κ). This class forms a covariety and it has, as a subcategory of F-coalgebras a terminal object. For applications involving at most a set of coalgebras and requiring the existence of a terminal coalgebra, we often may pretend that this terminal coalgebra does exists. All that is required is to choose κ large enough so that from the perspective of the application there is no difference between the κ-bounded terminal coalgebra and a "real" terminal coalgebra. As example application, we gather some results on minimal coalgebras, which are easy to prove by this method.

Given any logic, which is sound and expressive with respect to observational equivalence, the elements of the minimal coalgebras $\nabla\mathcal{A}$ correspond to the formulas inequivalent modulo $\mathcal{A}$. Therefore, we take this perspective in the last section to give an account of coalgebraic modal logic due to D. Pattinson[13] and L. Schröder[16], which we believe is more direct and straightforward in its presentation than can be found in the literature.

1. Basics

An important property true in the category of sets is that every epi-mono-square has a (necessarily unique) diagonal. That is, given a square $m \circ f = g \circ e$, where e is epi and m mono, there is a unique d such that $d \circ e = f$ and $m \circ d = g$. This is sometimes called the "diagonal fill-in property", and it is easily verified by checking that the set of all pairs $(e(x), f(x))$ defines the graph of the required function d.

$$\begin{tikzcd} \circ \arrow[r, two heads, "e"] \arrow[d, "f"'] & \circ \arrow[d, "g"] \arrow[dl, dashed, "d"'] \\ \circ \arrow[r, tail, "m"'] & \circ \end{tikzcd}$$

We shall need this lemma in a slightly more general form in 4.7:

Lemma 1.1. *Let $(X_i)_{i\in I}$ be a family of sets and $(e_i : X_i \to Y)_{i\in I}$ and $(f_i : X_i \to Z)_{i\in I}$ maps, so that the following diagram commutes for all $i \in I$:*

$$\begin{tikzcd} X_i \arrow[r, "e_i"] \arrow[d, "f_i"'] & Y \arrow[d, "g"] \\ Z \arrow[r, tail, "m"'] & W \end{tikzcd}$$

If m is mono and the e_i are jointly epi, then there exists a unique $d : Y \to Z$ making all arising triangles commutative.

1.1. **Coalgebras and homomorphisms.** Let $F : \mathcal{S}et \to \mathcal{S}et$ be any functor. By an *F-coalgebra* we understand a pair $\mathcal{A} = (A, \alpha_{\mathcal{A}})$ consisting of a set A and a map $\alpha : A \to F(A)$. A *homomorphism* φ to another F-coalgebra $\mathcal{B} = (B, \alpha_{\mathcal{B}})$ is just a map making the obvious diagram commute:

$$\begin{tikzcd} A \arrow[r, "\varphi"] \arrow[d, "\alpha_{\mathcal{A}}"'] & B \arrow[d, "\alpha_{\mathcal{B}}"] \\ F(A) \arrow[r, "F\varphi"] & F(B) \end{tikzcd}$$

If φ is the natural inclusion $\subseteq_A^B$, then $\mathcal{A}$ is called a *subcoalgebra* of $\mathcal{B}$. We denote this by writing $\mathcal{A} \le \mathcal{B}$. The class of all F-coalgebras with homomorphisms as defined above forms a category $\mathcal{S}et_F$, in which all colimits exist. In fact, the forgetful functor, associating with a coalgebra $\mathcal{A} = (A, \alpha_{\mathcal{A}})$ its base set A, creates and reflects colimits, see [15]. Some limits always exist, such as arbitrary equalizers, see [7], others, such as e.g. products, and the terminal object in particular, may fail to exist.

A *congruence* θ on a coalgebra $\mathcal{A} = (A, \alpha_{\mathcal{A}})$ is defined to be the kernel (in $\mathcal{S}et$) of any homomorphism with domain $\mathcal{A}$. It is well known that the set of all congruences on $\mathcal{A}$ forms a lattice with smallest element the identity relation id_A, and a largest congruence, which we shall call $\nabla_{\mathcal{A}}$. The latter may in general be properly contained in the universal relation $A \times A$.

1.2. ***Bisimilarity and observational equivalence.*** *Bisimulations* are compatible relations between coalgebras. By definition, a bisimulation R between coalgebras $\mathcal{A}$ and $\mathcal{B}$ is a binary relation $R \subseteq A \times B$ which can be equipped with some coalgebra structure $\rho : R \to F(R)$ so that the projections $\pi_1 : R \to A$ and $\pi_2 : R \to B$ become coalgebra homomorphisms:

$$\begin{tikzcd} A \arrow[d, "\alpha_{\mathcal{A}}"'] & R \arrow[l, "\pi_1"'] \arrow[r, "\pi_2"] \arrow[d, dashed, "\exists\rho"] & B \arrow[d, "\alpha_{\mathcal{B}}"] \\ F(A) & F(R) \arrow[l, "F\pi_1"'] \arrow[r, "F\pi_2"] & F(B) \end{tikzcd}$$

Note that a coalgebra structure $\rho : R \to FR$ witnessing that R is a bisimulation is not necessarily unique, see [5]. Nevertheless, there always exists a largest bisimulation $\sim_{\mathcal{A},\mathcal{B}}$ between any two coalgebras $\mathcal{A}$ and $\mathcal{B}$, written $\sim_{\mathcal{A}}$ when $\mathcal{A} = \mathcal{B}$. This largest bisimulation can be conveniently characterized as follows [3]:

Proposition 1.2. *Given $a \in A$ and $b \in B$, then $a \sim_{\mathcal{A},\mathcal{B}} b$ if and only if there exists a coalgebra $\mathcal{P} = (P, \alpha_P)$, homomorphism $\varphi_1 : \mathcal{P} \to \mathcal{A}$ and $\varphi_2 : \mathcal{P} \to \mathcal{B}$ and an element $p \in P$ so that $\varphi_1(p) = a$ and $\varphi_2(p) = b$.*

We say that $a \in \mathcal{A}$ and $b \in \mathcal{B}$ are *bisimilar*, and we write $a \sim b$, when $(a, b) \in \sim_{\mathcal{A},\mathcal{B}}$. Unfortunately, $\sim_{\mathcal{A}}$ need not be transitive, unless the type functor F weakly preserves kernel pairs, [7].

Dually to this characterization of bisimilarity, we can call two states $a \in \mathcal{A}$ and $b \in \mathcal{B}$ *observationally equivalent*, provided that there is a coalgebra $\mathcal{Q}$ and homomorphisms $\psi_1 : \mathcal{A} \to \mathcal{Q}$, $\psi_2 : \mathcal{B} \to \mathcal{Q}$ with $\psi_1(a) = \psi_2(b)$.

If a and b are states of the same coalgebra $\mathcal{A}$, it follows that (a, b) is in the kernel of the coequalizer of ψ_1 and ψ_2, which means that observational equivalence is given by the largest congruence relation $\nabla_{\mathcal{A}}$ of $\mathcal{A}$. Similarly, $a \in \mathcal{A}$ and $b \in \mathcal{B}$ are observationally equivalent iff $a \nabla_{\mathcal{A}+\mathcal{B}} b$ where $\mathcal{A} + \mathcal{B}$ is the sum of the coalgebras $\mathcal{A}$ and $\mathcal{B}$.

2. Minimal coalgebras

A coalgebra is called *minimal*, if $\nabla_{\mathcal{A}} = id$, i.e. if the only congruence relation is the identity. We define $\nabla(\mathcal{A}) := \mathcal{A}/\nabla_{\mathcal{A}}$, so $\mathcal{A}$ is minimal iff $\mathcal{A} = \nabla(\mathcal{A})$. (In earlier work we had called minimal coalgebras *strongly simple*, see [3].) We shall need to make use of the following observation :

Proposition 2.1. *A coalgebra $\mathcal{A}$ is minimal if and only if every homomorphism with domain $\mathcal{A}$ is regular mono.*

This follows from the facts that $\pi_{\nabla} : \mathcal{A} \to \nabla(\mathcal{A})$ can be obtained as the limit of all homomorphisms with domain $\mathcal{A}$ and that any homomorphism has an epi-regular mono factorization. Regular monos are precisely the injective homomorphisms, see [7].

Minimization of coalgebras is in fact functorial. Kianpi and Jugnia show this in [10] under the additional hypothesis that the type functor F weakly preserves kernel pairs. In that case ∇ is left adjoint to the inclusion of minimal coalgebras in the category of all coalgebras. Here we prove this result without any hypothesis on the type functor, and, not surprisingly, the proof becomes much more straightforward. Actually, it turns out that the minimal coalgebras form an epi-reflective implicational subcategory of $\mathcal{S}et_F$. The reflection for each coalgebra $\mathcal{A}$ is given by the unique homomorphism $\pi_{\nabla_{\mathcal{A}}} : \mathcal{A} \to \nabla(\mathcal{A})$. This follows from the following lemma:

Lemma 2.2. *For each homomorphism $\varphi : \mathcal{A} \to \mathcal{B}$ there exists a unique homomorphism $\nabla\varphi : \nabla(\mathcal{A}) \to \nabla(\mathcal{B})$ such that the following diagram commutes:*

$$\begin{tikzcd} \mathcal{A} \arrow[r, two heads, "\pi_{\nabla_{\mathcal{A}}}"] \arrow[d, "\varphi"'] & \nabla(\mathcal{A}) \arrow[d, dashed, "\nabla\varphi"] \\ \mathcal{B} \arrow[r, two heads, "\pi_{\nabla_{\mathcal{B}}}"'] & \nabla(\mathcal{B}) \end{tikzcd}$$

$\nabla\varphi$ is always regular mono, and it is iso whenever φ is epi.

Proof. Form the pushout in $\mathcal{S}et_F$ of $\pi_{\nabla_{\mathcal{A}}}$ and $\pi_{\nabla_{\mathcal{B}}} \circ \varphi$. As $\pi_{\nabla_{\mathcal{A}}}$ is epi and $\nabla(\mathcal{B})$ minimal, the resulting morphism from $\nabla(\mathcal{B})$ to the pushout must be epi and regular mono, hence iso. Thus $\nabla(\mathcal{B})$ itself is the pushout object and we can define $\nabla\varphi$ as the pushout map from $\nabla(\mathcal{A})$ to $\nabla(\mathcal{B})$. $\square$

From the lemma one obtains immediately that ∇ defines a functor from $\mathcal{S}et_F$ to the full subcategory of minimal coalgebras. Moreover, since $\nabla(\nabla(\mathcal{C})) = \nabla(\mathcal{C})$ for each coalgebra $\mathcal{C}$, the same lemma with $\nabla(\mathcal{C})$ in place of $\mathcal{B}$ shows that for each homomorphism $\varphi : \mathcal{A} \to \nabla(\mathcal{C})$ there exists a unique homomorphism $\nabla\varphi : \nabla(\mathcal{A}) \to \nabla(\mathcal{C})$ with $\varphi = \nabla\varphi \circ \pi_{\nabla_{\mathcal{A}}}$. This proves the following result:

Theorem 2.3. *∇ is a functor from $\mathcal{S}et_F$ to the full subcategory of minimal coalgebras which is left adjoint to the inclusion functor and has unit $\pi_{\nabla_{\mathcal{A}}} : \mathcal{A} \to \nabla(\mathcal{A})$.*

The subcategory of minimal coalgebras can be defined by implications, where an implication in the sense of [2] is just a regular epi $e : A \twoheadrightarrow B$, and an object C *satisfies* e, provided that for every morphism $f : A \to C$ there exists a (necessarily unique) $\tilde{f} : B \to C$ with $f = \tilde{f} \circ e$. Thus the lemma yields:

Corollary 2.4. *The subcategory of minimal coalgebras forms an implicational class, defined by the class of all implications* $\pi_{\nabla_{\mathcal{A}}} : \mathcal{A} \twoheadrightarrow \nabla(\mathcal{A})$.

2.1. **Products of minimal coalgebras.** Products of coalgebras need not exist. In [6] one finds examples of a small Kripke-Structures whose binary product is shown not to exist. For minimal coalgebras the situation is different, even for arbitrary functors F:

Theorem 2.5. *Let $\mathcal{A}$ and $\mathcal{B}$ be minimal coalgebras, then their product exists, and is isomorphic to their largest common subcoalgebra.*

Proof. $\sim_{\mathcal{A},\mathcal{B}}$ carries a coalgebra structure so that the projections $\pi_1 : \sim_{\mathcal{A},\mathcal{B}} \to \mathcal{A}$ and $\pi_2 : \sim_{\mathcal{A},\mathcal{B}} \to \mathcal{B}$ are homomorphisms. We form their pushout $\mathcal{P}$ with homomorphisms $\psi_1 : \mathcal{A} \to \mathcal{P}$ and $\psi_2 : \mathcal{B} \to \mathcal{P}$. These must be regular mono (injective) since $\mathcal{A}$ and $\mathcal{B}$ are minimal, so we can identify the latter with their images $\psi_1[A]$ and $\psi_2[B]$ as subcoalgebras of $\mathcal{P}$. Henceforth we will assume that ψ_1 and ψ_2 are the canonical injections into $\mathcal{P}$. According to [6], subcoalgebras are closed under finite intersections, so we have the subcoalgebra $\mathcal{A} \cap \mathcal{B}$ of $\mathcal{P}$, which we now claim to be the product of $\mathcal{A}$ with $\mathcal{B}$.

To verify this, let a competitor $\mathcal{Q}$ with homomorphisms $\varphi_1 : \mathcal{Q} \to \mathcal{A}$ and $\varphi_2 : \mathcal{Q} \to \mathcal{B}$ be given, then Proposition 1.2 yields a map (not necessarily a homomorphism) $f : Q \to \sim_{\mathcal{A},\mathcal{B}}$ with $\pi_1 \circ f = \varphi_1$ and $\pi_2 \circ f = \varphi_2$. Therefore, $\psi_1 \circ \varphi_1 = \psi_1 \circ \pi_1 \circ f = \psi_2 \circ \pi_2 \circ f = \psi_2 \circ \varphi_2$. This guarantees the required mediating homomorphism from $\mathcal{Q}$ to $\mathcal{A} \cap \mathcal{B}$, whose uniqueness is obvious.

$$
\begin{tikzcd}
 & & & \mathcal{A} \arrow[dr, hook, "\psi_1"] & \\
\mathcal{Q} \arrow[r, dashed, "f"] \arrow[urrr, bend left, "\varphi_1"] \arrow[drrr, bend right, "\varphi_2"'] & \sim_{\mathcal{A},\mathcal{B}} \arrow[urr, bend left, "\pi_1"] \arrow[drr, bend right, "\pi_2"'] \arrow[r, shift left, "\chi"] & \mathcal{A}\cap\mathcal{B} \arrow[l, dashed, shift left, "d"] \arrow[ur, hook, "\subseteq^{\mathcal{A}}"'] \arrow[dr, hook, "\subseteq^{\mathcal{B}}"] & & \mathcal{P} \\
 & & & \mathcal{B} \arrow[ur, hook, "\psi_2"'] &
\end{tikzcd}
$$

Finally, we show that $\sim_{\mathcal{A},\mathcal{B}}$ is isomorphic to $\mathcal{A} \cap \mathcal{B}$. For any other subcoalgebras $\mathcal{U} \leq \mathcal{A}$ and $\mathcal{V} \leq \mathcal{B}$ with $\mathcal{U} \cong \mathcal{V}$, the graph of the isomorphism would have to be a bisimulation and hence contained in $\sim_{\mathcal{A},\mathcal{B}}$.

A set-map $d : A \cap B \to \sim_{\mathcal{A},\mathcal{B}}$ with $\pi_1 \circ d = \subseteq^{\mathcal{A}}$ and $\pi_2 \circ d = \subseteq^{\mathcal{B}}$ can be obtained from proposition 1.2. From the equation $\psi_1 \circ \pi_1 = \psi_2 \circ \pi_2$ we obtain a homomorphism χ in the other direction for which we check that $\pi_1 = \subseteq^{\mathcal{A}} \circ \chi$ and likewise $\pi_2 = \subseteq^{\mathcal{B}} \circ \chi$. Since the π_i are jointly mono and so are $\subseteq^{\mathcal{A}}$ and $\subseteq^{\mathcal{B}}$, one gets that χ and d are mutually inverse, so they are in fact isomorphisms. □

2.2. **Terminal Coalgebras.** An F-coalgebra $\mathcal{W}$ is called *weakly terminal*, if for every F-coalgebra $\mathcal{A}$ there exists a homomorphism $\varphi : \mathcal{A} \to \mathcal{W}$. A coalgebra $\mathcal{T}$ is called *terminal* if there is always a unique such homomorphism, i.e. $\mathcal{T}$ is a terminal object in the category $\mathcal{S}et_F$. The next proposition is well known ([3]):

Proposition 2.6. *Let $\mathcal{W}$ be weakly terminal, then $\nabla(\mathcal{W})$ is terminal.*

Lambek's Lemma implies that the structure map $\alpha_{\mathcal{T}} : T \to FT$ of a terminal coalgebra $\mathcal{T}$ must be bijective. Consequently, terminal coalgebras need not exist, for instance if F is the powerset functor $\mathbb{P}$. If the terminal coalgebra $\mathcal{T}$ does exist, then for each $\mathcal{A}$ the terminal morphism $\tau : \nabla(\mathcal{A}) \to \mathcal{T}$ must be injective. Conversely, no subcoalgebra $\mathcal{U} \leq \mathcal{T}$ can be the source of a non-injective homomorphism $\varphi : \mathcal{U} \to \mathcal{B}$, as the unique homomorphism $\tau_{\mathcal{B}} : \mathcal{B} \to \mathcal{T}$ would have to compose yielding $\tau_{\mathcal{B}} \circ \varphi = \subseteq^{\mathcal{T}}_{\mathcal{U}}$. Therefore, we have:

Proposition 2.7. *If the terminal F-coalgebra exists, then the minimal F-coalgebras are precisely the subcoalgebras of $\mathcal{T}$.*

Of course, this has many consequences regarding the structure of minimal coalgebras, which are more or less trivially checked :

Proposition 2.8. *If the terminal coalgebra exists then*

- *subcoalgebras of minimal coalgebras are minimal,*
- *each minimal coalgebra has an injective structure map,*
- *minimal coalgebras have no nontrivial inner endomorphisms,*
- *products of finitely many minimal coalgebras exist.*

As it turns out, all the statement of the above proposition can be proved without assuming the existence of the terminal coalgebra. The last item, without that hypothesis, was in fact already proven in Theorem 2.5. It is repeated it here with that additional assumption, in order to demonstrate how obvious the proof would have been in the presence of a terminal coalgebra. In that case, both $\mathcal{A}$ and $\mathcal{B}$ would be subcoalgebras of the terminal coalgebra $\mathcal{T}$ and it is clear that the intersection of two subobjects of a terminal object is the same as their product.

With the ease that such results are obtained in the presence of a terminal coalgebra, the question arises whether they hold true, and how one proves them in its absence.

Recall that the non-existence of a terminal coalgebra, in the case, e.g. of the powerset functor $\mathbb{P}$, is caused by limitations of set theory. Allowing class-based coalgebras, one can indeed show that a terminal coalgebra does exist, but its carrier may be a proper class, see [1]. Using such results forces one to change the *definition of coalgebra*, to worry about the axiomatics of sets and classes and to re-inspect the theory as to how much of it remains valid and can be used for reasoning in the extended setting of class-based coalgebras.

If at all possible, we would like to avoid this, yet derive results as the previous one with equal ease. We claim that this is often possible, indeed, and we shall see that we can always assume the existence of a set-based terminal coalgebra in contexts where only a set of coalgebras is involved, or more precisely where we can put a bound on the *out-degree* (a term defined later) of the class of coalgebras under consideration.

For the first three examples in the above list, such a strategy works well, yet admittedly, it fails for the last example. The reason is that checking whether something is a product of F-coalgebras, requires considering as competitors all coalgebras $\mathcal{Q}$ from the whole *class* Set_F, rather than being able to limit consideration to a *set* of competitors or at least to a class with *bounded out-degree* (a term to be defined shortly), from which $\mathcal{Q}$ can be chosen. This justifies the separate statement and proof of Theorem 2.5.

3. Moore-Automata

Before continuing the general theory, we recall from the work of Rutten[14, 15] the basic theory of automata from the coalgebraic standpoint. Such automata are not only our prime examples of coalgebras, furnishing us with good intuitions and useful concepts, but we also will see that every coalgebra $\mathcal{A} = (A, \alpha)$ arises from some automaton (A, E, γ, δ) with the same base set A and an appropriate alphabet E, see Theorem 4.4.

3.1. **Moore automata over alphabet E.** A *Moore-Automaton* with input alphabet E and output alphabet D is given by a transition map $\delta : A \times E \to A$ and an output map $\gamma : A \to D$. Let $E^\star$ be the set of words (finite lists) of elements of E, then one usually extends δ to a map $\delta^\star : A \times E^\star \to A$ by

(1) $\delta^\star(a, \varepsilon) := a$, and
(2) $\delta^\star(a, e.w) := \delta^\star(\delta(a, e), w)$,

where ε denotes the empty word and $e.w$ is the word obtained by prefixing e to w.

If D is a two-element set $2 = \{0, 1\}$, then the automaton is called an *acceptor*. Fixing any starting state $a_0 \in A$, a word $w \in E^\star$ is *accepted* if $\delta^\star(a_0, w) = 1$.

3.2. **Automata as coalgebras.** The maps γ and δ can be combined into a single map $\alpha : A \to D \times A^E$, sending an element $a \in A$ to the pair $(\gamma(a), \delta_a)$, where the second component denotes the map δ with first argument fixed at a. Thus, the automaton is indeed a coalgebra for the functor with object map $F(X) = D \times X^E$, which sends a map $f : X \to Y$ to $Ff : D \times X^E \to D \times Y^E$ where $Ff(d, \sigma) := (d, f \circ \sigma)$.

The homomorphism condition, as defined for arbitrary coalgebras, translates straightforwardly to the usual notion of homomorphism between Moore-automata. That is, a map $\varphi : A \to B$ is a homomorphism between $D \times (-)^E$-coalgebras $\mathcal{A} = (A, \alpha_\mathcal{A})$ and $\mathcal{B} = (B, \alpha_B)$ iff with respect to δ and γ it satisfies:

(1) $\varphi(\delta_\mathcal{A}(a, e)) = \delta_\mathcal{B}(\varphi(a), e)$, and
(2) $\gamma_\mathcal{A}(a) = \gamma_\mathcal{B}(\varphi(a))$.

Intuitively, the *behavior* of a state a can be determined by feeding it a sequence w of inputs and finally observing the generated output $\gamma(\delta^\star(a, w))$. Thus, two states a, a' show the same behavior, iff $\gamma\delta^\star(a, w) = \gamma\delta^\star(a', w)$ for each $w \in E^\star$. Since the terminal coalgebra ought to contain one representative for each possible behavior, this suggests the following construction.

3.3. **The terminal Moore-Automaton.** Choose as underlying set $T = D^{E^\star}$, the set of all maps $t : E^\star \to D$. The structure on $D^{E^\star}$ is defined as

(1) $\delta_\mathcal{T}(t, e)(w) = t(e.w)$, and
(2) $\gamma_\mathcal{T}(t) := t(\varepsilon)$.

One may visualize the elements of $D^{E^\star}$ as E-branching infinite trees t whose nodes are labeled with elements from D. Then $\gamma_\mathcal{T}(t)$ is the label at the root of t and $\delta_\mathcal{T}(t, e) = t(e)$ is the subtree whose root is the e-th child of t. More general, each word $w \in E^\star$ determines a path from the root of t to some inner node whose label is precisely $\gamma\delta^\star(t, w) = t(w)(\varepsilon)$.

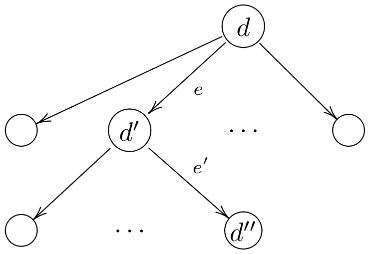

$d = \gamma(t)$

$d' = \gamma(\delta(t, e))$

$d'' = \gamma(\delta^\star(t, ee'))$

To see that $\mathcal{T}$ is indeed terminal, let $\mathcal{A} = (A, \delta_{\mathcal{A}}, \gamma_{\mathcal{A}})$ be any Moore-automaton. Assuming that $\tau : \mathcal{A} \to \mathcal{T}$ is a homomorphism, we obtain, simultaneously, for all $a \in A$, a recursive definition of $\tau(a)$:

(1) $\tau(a)(\varepsilon) \quad = \gamma_{\mathcal{T}}(\tau(a)) = \gamma_{\mathcal{A}}(a)$, and
(2) $\tau(a)(e.w) = \delta_{\mathcal{T}}(\tau(a), e)(w) = \tau(\delta_{\mathcal{A}}(a, e))(w)$.

These equations establish existence and uniqueness of τ. Using the *-notation, we can give a concise definition of τ which shows that it is just the map that unfolds the automaton at each state a into the tree $\tau(a)$:

$$\tau(a)(w) = \gamma_{\mathcal{A}} \delta^\star_{\mathcal{A}}(a, w).$$

This equation is easily verified using the recursive definitions of $\tau(a)$ and of $\delta^\star$. It also implies that two states have the same behavior iff they are observational equivalent.

4. The out-degree of a coalgebra

The *out-degree* of a state in a Moore-automaton is the number of immediate successors of a given state, so this is obviously bounded by the cardinality of E. On the other hand, *Kripke structures*, i.e. coalgebras of type $\mathbb{P}$, may have unlimited out-degree, even though the out-degree of any element in a fixed Kripke structure $\mathcal{K}$ is bounded by the cardinality of the underlying set of $\mathcal{K}$.

Here we shall associate an out-degree with any state in any arbitrary coalgebra. Again, the out-degree of a single state turns out to be bounded by the cardinality of the coalgebra. The coalgebras of a given functor F turn out to be uniformly bounded by some cardinality κ if and only if the functor is $\kappa-$bounded, that is if it is $\kappa-$accessible.

4.1. Bounded parts of $\mathcal{S}et$-functors.

We start with the well known approximation of $\mathcal{S}et$-endofunctors by κ-accessible subfunctors, see for instance [11, 18]. Let $F : \mathcal{S}et \to \mathcal{S}et$ be any set-endofunctor and κ a cardinal, then a natural transformation $\eta^\kappa : F(\kappa) \times (-)^\kappa \stackrel{\cdot}{\to} F(-)$ is given by $\eta^\kappa_X(u, \sigma) = (F\sigma)(u)$ for any set X. Its image-factorization yields a $\mathcal{S}et$-endofunctor which we shall call F_κ. Thus the following is a commuting diagram of natural transformations:

$$\begin{tikzcd}
F(\kappa) \times (-)^\kappa \arrow[rr, "\eta^\kappa"] \arrow[dr] & & F(-) \\
& F_\kappa(-) \arrow[ur, hook, "\subseteq"'] &
\end{tikzcd}$$

F_κ(X) can be described on objects as $F_\kappa(X) = \{(F\sigma)(u) \mid \sigma : \kappa \to X,\, u \in F(\kappa)\}$ and on maps $f : X \to Y$ as $(F_\kappa f)(F\sigma(u)) := (F(f \circ \sigma))(u)$. Clearly, F_κ is a functor, which agrees with F on the subcategory of sets of cardinality at most κ. Obviously, F is the directed union of the F_κ, i.e.

$$F(X) = \bigcup_{\kappa \in Card} F_\kappa(X).$$

F is called κ-*bounded*, if F is equal to F_κ for some cardinal κ. One easily checks:

Lemma 4.1. *Each F_κ is a subfunctor of F and for $\kappa \le \kappa'$, each pair of maps $\pi : \kappa' \to \kappa$ and $\iota : \kappa \to \kappa'$ with $\pi \circ \iota = id_\kappa$ gives rise to a natural inclusion $\mu : F_\kappa \subseteq F_{\kappa'}$.*

4.2. **Coalgebras of out-degree κ.** Since $F(X) = \bigcup_{\kappa \in Card} F_\kappa(X)$, there exists for each element $u \in F(X)$ a smallest κ with $u \in F_\kappa(X)$. We shall call this the *degree* of $u \in F(X)$.

Definition 4.2. Let $\mathcal{A} = (A, \alpha_{\mathcal{A}})$ be a coalgebra. For each $a \in A$, the *out-degree* of a will be the degree of $\alpha_{\mathcal{A}}(a)$. The out-degree of $\mathcal{A}$ is the supremum of the out-degrees of all $a \in A$.

In other words, the out-degree of a coalgebra $\mathcal{A} = (A, \alpha_A)$ is the smallest κ so that $\alpha_{\mathcal{A}} : A \to F(A)$ factors through $F_\kappa(A)$. Obviously, the out-degree of any $a \in A$, hence also the out-degree of $\mathcal{A}$ is bounded by $|A|$, the cardinality of A, that is

$$\textit{out-degree}(\mathcal{A}) \le |A|.$$

Example 4.3. In a Moore-automaton $\mathcal{A} = (A, \delta, \gamma)$ with input alphabet E and output set D, we recall that $F(X) = D \times X^E$, and $\alpha(a) = (\gamma(a), \delta_a)$ where $\delta_a = \lambda x : E.\delta(a, x)$. For any state $a \in A$, we calculate:

$$\begin{aligned}
\textit{out-degree}(a) &= min\{\kappa \mid \exists \sigma : \kappa \to X.\alpha(a) \in F\sigma[F(\kappa)\} \\
&= min\{\kappa \mid \exists \sigma : \kappa \to X.(\gamma(a), \delta_a) \in (id_D, \sigma \circ -)[D \times \kappa^E]\} \\
&= min\{\kappa \mid \exists \sigma : \kappa \to X.\exists \tau : E \to \kappa.\delta_a = \sigma \circ \tau\} \\
&= min\{\kappa \mid |\delta_a[E]| \le \kappa\} \\
&= |\{\delta(a, e) \mid e \in E\}|.
\end{aligned}$$

Thus, the out-degree of any state a is indeed just the number of immediate successors of a.

Given any coalgebra $\mathcal{A}$, let κ be its out-degree, then its structure map $\alpha_{\mathcal{A}}$ factors through $F_\kappa(A)$. Utilizing the surjective transformation $\eta^\kappa : F(\kappa) \times (-)^\kappa \twoheadrightarrow F_\kappa(-)$ and the axiom of choice, we can decompose $\alpha_{\mathcal{A}}$ as an automaton structure, followed by a component of the natural transformation η^κ:

$$\begin{tikzcd}
A \arrow[r, "\alpha_{\mathcal{A}}"] \arrow[d, dashed, "\exists\,(\gamma,\delta_-)"'] \arrow[dr, dotted] & F(A) \\
F(\kappa) \times A^\kappa \arrow[r, two heads, "\eta^\kappa"] & F_\kappa(A) \arrow[u, hook, "\subseteq"']
\end{tikzcd}$$

To be precise, the automaton has input set κ and output set $F(\kappa)$. After choosing a right-inverse r to η^κ_A, the automata structure is given by $\gamma(a) := \pi_1(r(\alpha(a)))$

and $\delta(a,e) := (\pi_2(r(\alpha(a))))(e)$. With the notation $\delta_a(e) := \delta(a,e)$, we obtain $\alpha(a) = \eta^\kappa_A(r(\alpha(a))) = F\pi_2(r(\alpha(a)))(\pi_1(r(\alpha(a)))) = (F\delta_a)(\gamma(a))$, thus recovering the coalgebra structure from the automaton. This yields:

Theorem 4.4. *For every coalgebra $\mathcal{A} = (A, \alpha_\mathcal{A})$ there exists a Moore-automaton $\mathcal{M} = (A, E, \gamma, \delta)$ with $|E| \leq$ out-degree$(\mathcal{A})$ such that $\alpha_\mathcal{A}(a) = (F\delta_a)(\gamma(a))$.*

The main result of this section will be that for any Set-functor F and any cardinality κ, the class of all F-coalgebras of out-degree at most κ forms a covariety, i.e. is closed under sums, homomorphic images and subcoalgebras. In order to show this, we have to identify and utilize appropriate features of the natural transformations η^κ.

4.3. **Natural transformations.** Let F and G be set-endofunctors and $\eta : F \stackrel{\cdot}{\to} G$ a natural transformation. η is called *epi-transformation* (resp. *mono-transformation*) if all component maps η_X are epi (resp. mono). η called *cartesian* if for each $f : X \to Y$ the naturality square

$$\begin{tikzcd} FX \arrow[r, "\eta_X"] \arrow[d, "Ff"'] & GX \arrow[d, "Gf"] \\ FY \arrow[r, "\eta_Y"] & GY \end{tikzcd}$$

is a pullback. If this holds only for f mono, then η called *sub-cartesian* (see [4]) or *taut* (see [12]). We call η *essentially subcartesian*, if the diagram is subcartesian unless, perhaps, for $X = \emptyset$.

Given any transformation $\eta : F \to G$, then each F-coalgebra $\mathcal{A} = (A, \alpha_\mathcal{A})$ gives rise to a G-coalgebra $\eta\mathcal{A} := (A, \eta_A \circ \alpha_\mathcal{A})$. This correspondence establishes a functor from $\mathcal{S}et_F$ to $\mathcal{S}et_G$, which we shall also denote by η. For any subclass $\mathfrak{K}$ of Set_F denote by $\eta\mathfrak{K}$ the image of $\mathfrak{K}$ under the functor $\eta : \mathcal{S}et_F \to \mathcal{S}et_G$, then the main result of this section is prepared by the following lemma:

Lemma 4.5. *Let $\eta : F \to G$ be a natural transformation, then*

(1) ηSet_F is closed under sums,
(2) if η is mono, then ηSet_F is closed under homomorphic images,
(3) η is essentially subcartesian if and only if ηSet_F is closed under taking subcoalgebras.

Proof. The first statement follows from the fact that the forgetful functor from coalgebras to their base set creates colimits. Thus the sum of the $\mathcal{A}_i$ in $\mathcal{S}et_F$ as well as the sum of the $\eta\mathcal{A}_i$ in $\mathcal{S}et_G$ have the same base set $A := \Sigma_{i\in I} A_i$ and the canonical embeddings $e_i : A_i \to A$ are the same as in $\mathcal{S}et$. Let $\alpha : A \to F(A)$ and $\alpha' : A \to G(A)$ be their corresponding structure maps. We need to show that $\alpha' = \eta_A \circ \alpha$, but this is easily obtained by precomposing with the e_i and using the fact that these are jointly epi in $\mathcal{S}et$.

$$\begin{tikzcd}
A_i \arrow[r, "e_i"] \arrow[d, "\alpha_i"] \arrow[dd, bend right=60, "\eta\circ\alpha_{A_i}"'] & A \arrow[d, "\alpha"'] \arrow[dd, bend left=60, "\alpha'"] \\
F(A_i) \arrow[r, "Fe_i"] \arrow[d, "\eta_{\mathcal{A}_i}"] & F(A) \arrow[d, "\eta_A"] \\
G(A_i) \arrow[r, "Ge_i"] & G(A)
\end{tikzcd}$$

For the proof of the second statement start with some $\mathcal{A} = (A, \eta_{\mathcal{A}} \circ \alpha_{\mathcal{A}})$ from ηSet_F, and let $\varphi : \mathcal{A} \twoheadrightarrow \mathcal{B}$ be an epimorphism to some $\mathcal{B}$ from Set_G. Then φ is epi in Set and the following diagram indicates how an F-structure on B can be obtained as a diagonal fill-in of an appropriate epi-mono-square. This yields an F-coalgebra which renders $\mathcal{B}$ in ηSet_F.

$$\begin{tikzcd}
A \arrow[r, two heads, "\varphi"] \arrow[d, "\alpha_{\mathcal{A}}"'] & B \arrow[d, dotted] \arrow[dd, bend left=60, "\alpha_{\mathcal{B}}"] \\
F(A) \arrow[r, "F\varphi"] \arrow[d, tail, "\eta_{\mathcal{A}}"'] & F(B) \arrow[d, tail, "\eta_{\mathcal{B}}"'] \\
G(A) \arrow[r, "G\varphi"] & G(B)
\end{tikzcd}$$

For the third claim, one easily checks that ηSet_F is closed under subcoalgebras and obviously contains the empty coalgebra, whenever η is essentially subcartesian.

For the converse, consider the naturality diagram with monomorphic f and elements $u \in GX$ and $v \in FY$ so that $(Gf)(u) = \eta_Y(v)$. We must produce an element $w \in FX$ with $\eta_X(w) = u$ and $(Ff)(w) = v$.

$$\begin{tikzcd}
 & & FX \arrow[r, tail, "Ff"] \arrow[d, "\eta_X"] & FY \arrow[d, "\eta_Y"] & \ni & v \\
u & \in & GX \arrow[r, tail, "Gf"] & GY & &
\end{tikzcd}$$

Consider the map $c_u^X : X \to GX$ with constant value u and $c_v^Y : Y \to FY$ with constant value v. Clearly, $\mathcal{X} = (X, c_u^X)$ is a subcoalgebra of $\mathcal{Y} = (Y, \eta_Y \circ c_u^Y)$ and the latter is in ηSet_F. By hypothesis then $\mathcal{X}$ is in ηSet_F, whence c_u^X must factor as $\eta_X \circ \alpha$ through $F(X)$. Choosing an arbitrary element $x_0 \in X$, an element with the desired properties is obtained as $\alpha(x_0)$. $\square$

Theorem 4.6. *Given any covariety $\mathfrak{V}$ of F-coalgebras, then the class $\mathfrak{V}_\kappa$ of all coalgebras in $\mathfrak{V}$ of out-degree at most κ forms a subcovariety.*

Proof. As intersections of covarieties are covarieties, it is enough to show the result for $\mathfrak{V} = Set_F$. Since F_κ is a subfunctor of F, parts 1 and 2 of the the above lemma apply. It remains to check that the inclusion $F_\kappa \subseteq F$ is essentially subcartesian.

Given $u \in F(X)$ and $v \in F_\kappa(Y)$ with $(F \subseteq)(u) = v$, there exists some map $\sigma : \kappa \to Y$ and some element $a \in F(\kappa)$ with $F(\sigma)(a) = v$. As we are allowed to assume that X is nonempty, we can choose any map $\tau : Y \to X$ which is a left inverse to the inclusion $\subseteq : X \to Y$ and calculate

$$\begin{aligned}(F(\tau \circ \sigma))(a) &= (F\tau)(F\sigma)(a) \\ &= (F\tau)(F \subseteq)(u) \\ &= F_{id_X}(u) \\ &= u.\end{aligned}$$

This proves that $u \in F_\kappa(X)$. Automatically $(F_\kappa \subseteq)(u) = v$. □

4.4. **Terminal coalgebras.** For any functor F, the bounded functors F_k approximate F in the sense that $F_\kappa \le F_{\kappa'} \le F$ for each $\kappa \le \kappa'$, and F is the union of all F_κ. Similarly, the class Set_F of F-coalgebras is the union of an increasing sequence of the classes Set_{F_κ}, each of which is subcovariety of Set_F.

Theorem 4.7. *Let F and G be Set-endofunctors and $\eta : F \stackrel{\cdot}{\to} G$ a mono-transformation.*

1. If $\mathcal{A}$ is a minimal F-coalgebra then $\eta\mathcal{A}$ is a minimal G-coalgebra.
2. If a terminal G-coalgebra $\mathcal{T}_G$ exists, then a terminal F-coalgebra $\mathcal{T}_F$ exists too, and $\eta\mathcal{T}_F \le \mathcal{T}_G$.

Proof. (1) Given a G-homomorphism $\varphi : \eta\mathcal{A} \to \mathcal{B}$, we may in fact assume that φ is surjective, for otherwise we could replace $\mathcal{B}$ with the image of $\mathcal{A}$ under φ. We need to show that φ is injective. By Lemma 4.5(2), $\mathcal{B}$ is in ηSet_F, from which it follows that φ is an F-homomorphism, hence injective.

(2) Let $\mathcal{T} = (T, \alpha_\mathcal{T})$ be the terminal G-coalgebra. The subset $U := \bigcup\{\varphi[\eta\mathcal{A}] \mid \mathcal{A} \in \mathcal{S}et_F,\ \varphi : \eta\mathcal{A} \to \mathcal{T}\}$ is a union of subcoalgebras of $\mathcal{T}$, so it is itself a G-subcoalgebra $\mathcal{U} = (U, \alpha_\mathcal{U})$ of $\mathcal{T}$. We can choose a set $(\mathcal{A}_i)_{i\in I}$ of F-coalgebras and homomorphisms $\varphi_i : \eta\mathcal{A}_i \to \mathcal{U}$ which are jointly epi. Utilizing Lemma 1.1, the following diagram indicates how an F-coalgebra structure β on U can be obtained as a diagonal fill-in, so that $\alpha_\mathcal{U} = \eta_\mathcal{U} \circ \beta$, that is $\mathcal{U} = \eta(U, \beta)$.

$$\begin{tikzcd}
A_i \arrow[r, "\varphi"] \arrow[d, "\alpha_{\mathcal{A}_i}"'] & U \arrow[r, hook] \arrow[d, dotted, "\beta"] \arrow[dd, bend left, "\alpha_\mathcal{U}"] & T \arrow[dd, bend left, "\alpha_\mathcal{T}"] \\
F(A_i) \arrow[r, "F\varphi"] \arrow[d, tail, "\eta_{\mathcal{A}_i}"'] & F(U) \arrow[d, tail, "\eta_\mathcal{U}"'] & \\
G(A_i) \arrow[r, "G\varphi"] & G(U) \arrow[r, hook] & G(T)
\end{tikzcd}$$

We claim that (U, β) is the terminal F-coalgebra. Indeed, given any F-coalgebra $\mathcal{A}$, there is a unique G-homomorphism $\varphi : \eta\mathcal{A} \to \mathcal{U}$ with $\alpha_\mathcal{U} \circ \varphi = G\varphi \circ \eta_A \circ \alpha_\mathcal{A}$. It follows that $\eta_U \circ \beta \circ \varphi = \eta_U \circ F\varphi \circ \alpha_\mathcal{A}$, whence $\beta \circ \varphi = F\varphi \circ \alpha_\mathcal{A}$. Therefore, φ is an F-homomorphism. Uniqueness of φ is obvious, so (U, β) is terminal in $\mathcal{S}et_F$. □

Since we have $F_\kappa \leq F_{\kappa'} \leq F$ whenever $\kappa \leq \kappa'$ for each functor F, it follows that $\eta^\kappa \mathcal{T}_\kappa \leq \eta^{\kappa'} \mathcal{T}_{\kappa'} \leq \mathcal{T}$ for the corresponding terminal coalgebras, provided $\mathcal{T}$ exists. Thanks to the epi-transformation $\eta^\kappa : F(\kappa) \times (-)^\kappa \stackrel{\cdot}{\twoheadrightarrow} F_\kappa(-)$, each F_κ-coalgebra is of the form $\eta\mathcal{A}$ for some Moore-Automaton with input alphabet κ and output set $F(\kappa)$ as described in 4.4. Terminal Moore-Automata exist, so the following result from [9] guarantees the existence of all $\mathcal{T}_\kappa$ and shows us how they are built:

Theorem 4.8. *([9], sect. 4.3) If $\mathcal{W} = (W, \alpha_{\mathcal{W}})$ is a (weakly) terminal F-coalgebra and $\eta : F \stackrel{\cdot}{\twoheadrightarrow} G$ an epi-transformation, then $\nabla(\eta\mathcal{W})$ is the terminal G-coalgebra.*

The following diagram summarizes the construction of the terminal coalgebra with out-degree κ. Starting with the terminal Moore-Automaton $\mathcal{M} = (M, \kappa, \gamma, \delta)$, we obtain the structure map $\alpha : M \to F_\kappa(M)$ defined by $\alpha(m) = (F\delta_m)(\gamma(m))$, then we factor this F_κ-coalgebra by its largest congruence relation ∇.

$$\begin{tikzcd}
F(\kappa) \times M^\kappa \arrow[rr] \arrow[dr, two heads, "\eta"] & & F(M) \arrow[r, two heads] & F(M/\nabla) \\
& F_\kappa(M) \arrow[ur, hook] \arrow[r, two heads] & F_\kappa(M/\nabla) \arrow[ur, hook] & \\
M \arrow[uu, "(\gamma,\delta)"] \arrow[ur, "\alpha"] \arrow[r, two heads] & M/\nabla \arrow[ur] & &
\end{tikzcd}$$

To show, for instance, that a minimal coalgebra $\mathcal{A}$ does have an injective structure map, we consider $\mathcal{A}$ as subcoalgebra of $\mathcal{T}_\kappa$ where $\kappa \geq$ *out-degree*$(\mathcal{A})$ and use the fact that the structure map of a terminal coalgebra is always bijective. All other items of proposition 2.8 are proved similarly without assuming the precondition of existence of the terminal coalgebra. An exception, as we have mentioned is the last item, for which we have given a separate proof in Theorem 2.5.

5. Coalgebraic Modal Logic

A logic for coalgebras should provide a language and a semantics that can describe the states in a coalgebra up to observational equivalence. Thus logical expressions should be able to tell apart the elements of minimal coalgebras, and only these. Therefore, logical formulae correspond, up to logical equivalence, exactly to the elements of minimal coalgebras.

There is, in fact, a modal logic, developed by D. Pattinson[13] and L. Schröder [16], which under some mild conditions captures exactly observational equivalence. Since it complements well our study of minimal coalgebras, we give here an account, which we believe is quite a bit simpler and more straightforward than the one found in the original literature.

5.1. **Coalgebraic logic in general.** Any coalgebraic logic for a functor F must consist of a logical language $\mathcal{L}$ and a validity relation $\models_{\mathcal{A}} \subseteq A \times \mathcal{L}$ for each coalgebra $\mathcal{A} = (A, \alpha_{\mathcal{A}})$. We write

$$a \models_{\mathcal{A}} \phi$$

if $(a, \phi) \in \models_{\mathcal{A}}$ for $a \in A$ and $\phi \in \mathcal{L}$. We shall drop the lower index, if it is clear from the context.

Given a formula ϕ, we write $[\![\phi]\!]_{\mathcal{A}}$ for the characteristic function of the set of elements of A *defined by* ϕ, that is $[\![\phi]\!]_A : A \to 2$ is given as

$$[\![\phi]\!]_{\mathcal{A}}(a) = \begin{cases} 1 & \text{if}\, a \models_{\mathcal{A}} \phi \\ 0 & \text{else.} \end{cases}$$

Given coalgebras $\mathcal{A} = (A, \alpha_{\mathcal{A}})$ and $\mathcal{B} = (B, \alpha_B)$ then elements $a \in A$ and $b \in B$ are called *logically equivalent*, and we write $a \approx b$ if they satisfy the same formulae, i.e, if for all $\phi \in \mathcal{L}$ we have $a \models_{\mathcal{A}} \phi \iff b \models_{\mathcal{B}} \phi$. We write $a/\approx$ for the $\approx$-class containing a, and $A/\approx$ for the factor set of A by $\approx$.

A logic $(\mathcal{L}, \models)$ is called *adequate* or *admissible*, if observationally equivalent elements are logically equivalent, i.e. if $\nabla \subseteq \approx$, and *expressive*, if logically equivalent elements are observationally equivalent, i.e. $\approx \subseteq \nabla$. Thus a logic is adequate and expressive just in case it can distinguish two elements iff they are mapped to different elements in the minimal factor, that is iff $\approx = \nabla$.

A logic is called κ-ary, if it has negation $\neg$ and conjunctions $\bigwedge_{i \in I}$ for index sets I with $|I| < \kappa$ defined by the obvious semantics: $a \models \neg\phi \iff a \not\models \phi$ and $a \models \bigwedge_{i\in I} \phi_i \iff \forall i \in I. a \models \phi_i$.

The following lemma asserts that on κ-small subsets $U \subseteq A$ families of logically inequivalent elements can be separated by some logical formula. To be precise:

Lemma 5.1. *Given $U \subseteq A$ with $|U| < \kappa$ and a map $f : A/\approx \to 2$, then there is a formula ϕ_f such for each $u \in U$*

$$u \models \phi_f \iff f(u/\approx) = 1.$$

Proof. For any two elements $u, v \in U$ with $u \not\approx v$, there must be a formula $\phi_{u,v}$ holding in u but not in v. Then $\bigvee_{u \not\approx v \in U} \phi_{u,v}$ defines u among all other (non-equivalent) elements of U. We obtain the required formula as

$$\phi_f := \bigvee_{f(u/\approx)=1} \bigwedge_{u \not\approx v} \phi_{u,v}.$$

□

5.2. **Modal logic.** For *image-finite* Kripke-Structures, i.e. coalgebras of type $\mathbb{P}_\omega(-)$, a logic $\mathcal{L}$ satisfying the requirements of being both adequate and expressive, is Hennessy-Milner logic, defined by the grammar

$$\phi ::= \top \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \Box\phi$$

where the only interesting semantical clause is

$$a \models \Box\phi :\iff \forall a' \in \alpha(a).\, a' \models \phi.$$

Completeness of this logic for image finite Kripke structures was established by Hennessy and Milner [8].

5.3. **Pattinson-Schröder logic.** In an attempt to mimic this for arbitrary coalgebras, Dirk Pattinson ([13]) studied so called “predicate liftings” λ, intended to translate predicates on the base set A to predicates on $F(A)$. Each such predicate lifting λ gives rise to a modality $[\lambda]$ and Pattinson gave conditions for such predicate liftings to yield adequate and expressive logics. Lutz Schröder ([16]) then observed that the relevant liftings can all be obtained from subsets of $F(2)$:

A semantic map $[\![\phi]\!] : A \to 2$ lifts to a map $F\,[\![\phi]\!] : F(A) \to F(2)$, so each subset $W \subseteq F(2)$ defines a modality $[W]$ specified by

$$a \models_{\mathcal{A}} [W]\phi \;:\Longleftrightarrow\; F\,[\![\phi]\!](\alpha_{\mathcal{A}}(a)) \in W.$$

The syntax for κ-ary coalgebraic modal logic is therefore given by the grammar

$$\phi ::= \top \mid \bigwedge_{i \in I} \phi_i \mid \neg\phi \mid [W]\phi$$

where W runs through the subsets of $F(2)$ and I may be any index set with $|I| < \kappa$. The semantics of $\top$, $\bigwedge$, and $\neg$ is as usual, and the semantics of $[W]\phi$ is the one defined in the previous paragraph.

Intuitively, we like to think of $F(A)$ as a set of *patterns* (e.g. tuples, lists, trees or equivalence classes of trees) with certain positions (coordinates, leaves or nodes) occupied by elements of A. A map $F\,[\![\phi]\!] : A \to 2$ transforms such a pattern by relabeling the positions satisfying ϕ by 1 and those not satisfying ϕ by 0. The result is a 0-1-pattern (an element of $F(2)$) which might or might not be a member of $W \subseteq F(2)$. Thus $a \models [W]\phi$ holds iff the 0-1-pattern associated by ϕ to $\alpha_{\mathcal{A}}(a)$, i.e. the successor pattern of a with respect to ϕ, is a member of W.

5.4. **Admissibility.**

Lemma. [16]*If $\varphi : \mathcal{A} \to \mathcal{B}$ is a coalgebra homomorphism, ϕ a formula of modal logic and $a \in A$, then $a \models \phi \iff \varphi(a) \models \phi$.*

Proof. The claim is obviously true for $\phi = \top$ and is easily carried inductively from ϕ to $\neg\phi$ and from a collection $(\phi_i)_{i \in I}$ to $\bigwedge_{i \in I} \phi_i$. To handle the last syntactic clause, assume that the result is true for ϕ, that is $[\![\phi]\!]_{\mathcal{B}} \circ \varphi = [\![\phi]\!]_A$. Applying F to this equation and drawing the structure maps of $\mathcal{A}$ and $\mathcal{B}$ we obtain the following diagram, from which one can immediately read off that $F\,[\![\phi]\!]_{\mathcal{A}} \circ \alpha_{\mathcal{A}} = F\,[\![\phi]\!]_{\mathcal{B}} \circ \alpha_{\mathcal{B}} \circ \varphi$, so in particular, $a \models_{\mathcal{A}} [W]\phi \iff \varphi(a) \models_{\mathcal{B}} [W]\phi$ for any $W \subseteq F(2)$.

$$\begin{tikzcd}
A \arrow[rr, "\varphi"] \arrow[d, "\alpha_{\mathcal{A}}"] & & B \arrow[d, "\alpha_{\mathcal{B}}"] \\
F(A) \arrow[rr, "F\varphi"] \arrow[dr, "F\,[\![\phi]\!]_{\mathcal{A}}"'] & & F(B) \arrow[dl, "F\,[\![\phi]\!]_{B}"] \\
& F(2) &
\end{tikzcd}$$

□

5.5. **Expressivity.**

To show expressivity, the functor is assumed to be *separating,* in the sense that for any pair $u, v \in F(X)$ with $u \neq v$ there exists some map $f : X \to 2$ with $Ff(u) \neq Ff(v)$.

Theorem 5.2. *(L. Schröder* [16]*) Let the functor F be separating and κ-accessible for a regular cardinal κ. Then the κ-ary logic for F is expressive.*

Proof. Consider the equivalence relation $\approx$ on $\mathcal{A} = (A, \alpha)$ given by logical equivalence and the canonical projection $\pi_{\approx} : A \to A/\approx$. Admissibility of the logic, as established in the previous lemma, amounts to the inclusion $\nabla \subseteq \approx$. For expressiveness it remains to show $\approx \subseteq \nabla$.

We attempt to put a coalgebra structure on $A/\approx$ that turns $\pi_\approx$ into a homomorphism, for then $\approx$ will be a congruence, whence $\approx \subseteq \nabla$. To obtain such a structure map, it is (necessary and) sufficient to show the following

Claim 5.3. $ker\pi_\approx \subseteq ker(F\pi_\approx \circ \alpha)$.

$$\begin{tikzcd}
 & {} \arrow[d, dotted, shift right, "a"'] \arrow[d, dotted, shift left, "b"] & & \\
U \arrow[r, hook] & A \arrow[r, two heads, "\pi_\approx"] \arrow[rr, bend left, "\llbracket\phi\rrbracket_{f\mathcal{A}}"] \arrow[d, "\alpha"] \arrow[dl, dotted] & A/\approx \arrow[r, "f"] \arrow[d, dashed, tail] & 2 \\
F(U) \arrow[r, hook] & F(A) \arrow[r, two heads, "F\pi_\approx"] \arrow[rr, bend right, "F\llbracket\phi\rrbracket_{f\mathcal{A}}"'] & F(A/\approx) \arrow[r, "Ff"] & F(2) \\
 & & & {} \arrow[u, dotted, shift right, "w_a"] \arrow[u, dotted, shift left, "w_b"']
\end{tikzcd}$$

By way of contradiction, assume that $a, b \in A$ be given with $\pi_\approx(a) = \pi_\approx(b)$, i.e. $a \approx b$, but $(F\pi_\approx \circ \alpha)(a) \neq (F\pi_\approx \circ \alpha)(b)$. The separability assumption applied to $X = A/\approx$ provides a map $f : A/\approx \to 2$ with

$$w_a := (Ff \circ F\pi_\approx \circ \alpha)(a) \neq (Ff \circ F\pi_\approx \circ \alpha)(b) =: w_b.$$

Since F is κ-bounded and κ regular, there exists a subset $U \subseteq A$ with $|U| < \kappa$ and $\alpha(a), \alpha(b) \in F(U)$. By Lemma 5.1 there is a formula ϕ_f so that $\pi_\approx \circ f$ and $\llbracket\phi_f\rrbracket$ agree on U. Applying the functor F, it follows that $Ff \circ F\pi_\approx$ and $F\llbracket\phi_f\rrbracket$ agree on $F(U)$. Therefore, a satisfies $[\{w_a\}]\phi_f$, but b does not, contradicting $a \approx b$. □

For coalgebraic logic to be expressive, the theorem requires that for each set X the family $(Ff)_{f\in 2^X}$ is a mono-source. While this is true for commonly used coalgebraic type functors, it is also easy to construct examples to the contrary. For any equational theory Σ that cannot be specified solely by two-variable equations, the free-algebra functor $F_\Sigma(X)$ provides such a counterexample.

Separating functors can in fact be easily identified as being precisely subfunctors of the functor $Q(X) = F(2)^{2^X}$, which is the composition of the contravariant hom-functors $Q_1 = hom(-, F(2))$ and $Q_2 = hom(-, 2)$. To be precise:

Proposition 5.4. *A $\mathcal{S}et$-functor F is separating iff the natural transformation* $\mu : F \to F(2)^{2^-}$ *given by* $\mu_X(u) = ((Ff)(u))_{f:X\to 2}$ *is injective.*

This association of functors with subfunctors of $F(2)^{2^-}$ has an easy interpretation: Each element of $F(X)$ is being associated with the set of all possible $0-1-$patterns that can arise from it by substituting the elements of X with either 0 or 1. F is separable iff any two elements of $F(X)$ can be distinguished by their $0-1-$patterns.

If this is not the case, Schröder shows that one can always obtain a complete logic for κ-accessible functors, if polyadic modalities are allowed. In our language, 0-1-patterns must be replaced by $\{0,1\}^\kappa$-patterns, i.e. subsets of $F(2^\kappa)$. The proof essentially remains intact, see [16].

References

[1] P. Aczel and N. Mendler. A final coalgebra theorem. In D.H. Pitt et al., editors, *Proceedings category theory and computer science*, Lecture Notes in Computer Science, pages 357–365. Springer, 1989.

[2] J. Adámek, H. Herrlich, and G.E. Strecker. *Abstract and Concrete Categories*. John Wiley & Sons, 1990.

[3] H.P. Gumm. Elements of the general theory of coalgebras. LUATCS 99, Rand Afrikaans University, Johannesburg, South Africa, 1999.

[4] H.P. Gumm. From T-coalgebras to filter structures and transition systems. In D.H. Fiadeiro et al., editors, *Algebra and Coalgebra in Computer Science*, volume 3629 of *Lecture Notes in Computer Science*, pages 194–212. Springer, 2005.

[5] H.P. Gumm and T. Schröder. Products of coalgebras. *Algebra Universalis*, 46:163–185, 2001.

[6] H.P. Gumm and T. Schröder. Coalgebras of bounded type. *Math. Struct. in Comp. Science*, 12:565–578, 2002.

[7] H.P. Gumm and T. Schröder. Types and coalgebraic structure. *Algebra Universalis*, 53:229–252, 2005.

[8] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *J. Assoc. Comput. Mach.*, 32:137–161, 1985.

[9] Th. Ihringer and H.P.Gumm. *Allgemeine Algebra*. Heldermann Verlag, 2003.

[10] M. Kianpi and C.N. Jugnia. A simplification functor for coalgebras. *International Journal of Mathematics and Mathematical Sciences*, 2006.

[11] V. Koubek. Set functors. *Comm. Math. Univ. Carolinae*, 12(1):175–195, 1971.

[12] E.G. Manes. Implementing collection classes with monads. *Math. Struct. in Comp. Science*, 8:231–276, 1998.

[13] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic*, 45:19–33, 2004.

[14] J.J.M.M. Rutten. Automata and coinduction (an exercise in coalgebra). In D. Sangiorigi and R. de Simone, editors, *Proceedings of CONCUR '98*, number 1466 in LNCS, pages 194–218, 1998.

[15] J.J.M.M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3–80, 2000.

[16] L. Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. In V. Sassone, editor, *Foundations of Software Science and Computation Structures (FOSSACS)*, Lecture Notes in Computer Science, pages 440–454. Springer, 2005.

[17] M.B. Smyth and G.D. Plotkin. The category-theoretic solution of recursive domain equations. *SIAM Journal of Computation*, 11(4):761–783, 1982.

[18] V. Trnková. On descriptive classification of set-functors I. *Comm. Math. Univ. Carolinae*, 12(1):323–352, 1971.