

A Modal Proof Theory for Final Polynomial Coalgebras

David Friggens and Robert Goldblatt *

Centre for Logic, Language and Computation

Victoria University, P. O. Box 600, Wellington, New Zealand

29 November 2005

Abstract

An infinitary proof theory is developed for modal logics whose models are coalgebras of polynomial functors on the category of sets. The canonical model method from modal logic is adapted to construct a final coalgebra for any polynomial functor. The states of this final coalgebra are certain “maximal” sets of formulas that have natural syntactic closure properties.

The syntax of these logics extends that of previously developed modal languages for polynomial coalgebras by adding formulas that express the “termination” of certain functions induced by transition paths. A completeness theorem is proven for the logic of functors which have the *Lindenbaum* property that every consistent set of formulas has a maximal extension. This property is shown to hold if the deducibility relation is generated by countably many inference rules.

A counter-example to completeness is also given. This is a polynomial functor that is not Lindenbaum: it has an uncountable set of formulas that is deductively consistent but has no maximal extension and is unsatisfiable, even though all of its countable subsets are satisfiable.

1 Introduction

If $T : \mathbf{Set} \rightarrow \mathbf{Set}$ is a functor on the category of sets, then a T -coalgebra is a pair $\langle A, \alpha \rangle$ with A being a set and α a function of the form $A \rightarrow TA$. This concept has proven useful in modelling various computational structures and systems, including data structures (infinite lists, streams, trees), state-based systems (automata, labelled transition systems, process algebras) and classes in object-oriented programming languages [34, 18, 38, 39, 23, 22]. Typically A is thought of as a set of *states*, and α as a *transition structure*. The T -coalgebras form a category under a natural notion of coalgebraic morphism $f : \langle A, \alpha \rangle \rightarrow \langle B, \beta \rangle$ given by a function $f : A \rightarrow B$ that preserves the transition structures in a suitable sense.

Particular importance attaches to the notion of a *final* (or *terminal*) coalgebra, which is a coalgebra $\langle C, \gamma \rangle$ such that for each coalgebra $\langle A, \alpha \rangle$ there is exactly one coalgebraic morphism from $\langle A, \alpha \rangle$ to $\langle C, \gamma \rangle$. In the context of process algebra, the states of a final coalgebra can be thought of as representing all possible “observable behaviours” of processes, because states x of $\langle A, \alpha \rangle$ and y of $\langle B, \beta \rangle$ are “observationally indistinguishable” precisely when they are identified by the unique morphisms from $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ to $\langle C, \gamma \rangle$.

A final T -coalgebra is unique up to isomorphism, if it exists. Conditions under which it exists, and representations of it when it does, have been extensively studied [3, 4, 5, 18, 39, 25, 37]. The principal aim of this paper is to show how the proof theory of modal logic can give an elegant construction of the final coalgebra for any *polynomial* functor. A functor is polynomial if it can be inductively constructed from the identity functor and constant-valued functors by forming products $T_1A \times T_2A$, coproducts (disjoint unions) $T_1A + T_2A$, and exponentials $(TA)^I$ with fixed exponent I . Polynomial functors can be thought of as being constructed by these operations from some fixed sets given in advance, and a polynomial coalgebra is a very general kind of deterministic transition system for which the value of a constant functor is a set of “observable outputs” and an exponent I constitutes a set of “inputs”.

*Corresponding author: rob.goldblatt@vuw.ac.nz. The authors acknowledge support from a Logic and Computation programme funded by the New Zealand Institute of Mathematics and its Applications, and thank the referees for their helpful suggestions.

There have been a number of proposals of formal languages and logics for characterising properties of coalgebras. A desideratum of such logics is that they have a semantic satisfaction relation $A, \alpha, x \models \varphi$, expressing “formula φ is true/satisfied at state x in coalgebra $\langle A, \alpha \rangle$ ”, that provides a logical characterisation of observational indistinguishability in the following form:

$$\begin{aligned} &x \text{ is observationally indistinguishable from } y \\ &\quad \text{if, and only if,} \\ &\text{for all formulas } \varphi, A, \alpha, x \models \varphi \text{ iff } B, \beta, y \models \varphi. \end{aligned}$$

In other words, observational indistinguishability is identical to logical indistinguishability. If such a logic exists we say that it, or the functor T , has the *Hennessey-Milner property*, after the pioneers of this idea for process algebra [16, 17].

The first explicit coalgebraic logic with the Hennessy-Milner property was introduced by Moss [32] for a broadly defined class of functors T that have final coalgebras. The language involved was *infinitary*, allowing formation of conjunctions of infinite sets of formulas, and was motivated by ideas from modal logic. Finitary modal languages with the Hennessy-Milner property were subsequently developed for more specific types of functor, beginning with the work of Kurz [27, 28], Rößiger [35, 37] and Jacobs [20] on polynomial functors. The fundamental *canonical model* construction was adapted in [35, 28] to build polynomial coalgebras. This construction originated in the method introduced by Henkin [15] for proving completeness of first-order logic, and was adapted to modal logic by Lemmon and Scott [30] and others [8, 31]. The essence of the method is to define a model whose states are certain “maximal” sets of formulas with special closure properties determined by the proof theory of the logic, and to show that a formula φ is satisfied in this model at a maximal x precisely when $\varphi \in x$. The technique was used in [35, 28] to construct final polynomial coalgebras as canonical models, under the restriction that any constant output set involved in the formation of T is *finite*.

In this paper we show that the canonical model method can be used to construct a final T -coalgebra for *any* polynomial T , including those that have infinite constant output sets, such as the set of natural numbers or even uncountable sets. This is done by developing an *infinitary* proof theory for a *finitary* modal language determined by T . The proof theory is infinitary in the sense that it involves a deducibility relation $\Gamma \vdash \varphi$, interpreted as “formula φ is deducible from the set of formulas Γ ”, which may hold concurrently with $\Gamma' \vdash \varphi$ failing for every finite $\Gamma' \subseteq \Gamma$. So the deduction of φ may depend on infinitely many “premisses”. In particular, it may be that Γ is deductively inconsistent, in the sense that $\Gamma \vdash \perp$ where \perp is a constant false formula, while at the same time each finite subset of Γ is consistent. Proof theories of this kind were developed in [9] for standard modal logics, and are adapted here to the coalgebraic setting. While the general framework of [9] carries over, there are a many novel features involved, including the axioms and rules of inference used, and the canonical model construction itself, which is distinctively coalgebraic.

Our modal language is finitary in the sense that its formulas are finite sequences of symbols. Its syntax extends that of Rößiger’s polynomial language by a new construct expressing assertions about the existence of transition functions induced by certain *path* expressions. This path notion was developed in [35, 37] and then in [19, 20] and provides a way of representing the internal structure and formation of a polynomial functor T . We write $T \xrightarrow{p} S$ to indicate that p is a path from T to functor S : such paths exist whenever S is a component of the formation of T . Paths induce a partial function $p_A : TA \multimap SA$ for each set A , and this composes with a T -transition $\alpha : A \rightarrow TA$ to give a partial function $p_A \circ \alpha : A \multimap SA$. p is a *state path* if S is the identity functor, so $SA = A$, and is an *observation path* if S is a constant functor, so $SA = D$ for some fixed set D of observable values. The language of [35, 37] has formulas $(p)c$ for each such observation path p and each $c \in D$. The formula $(p)c$ is true at state x in T -coalgebra $\langle A, \alpha \rangle$ when $p_A(\alpha(x))$ is defined and equal to c . There are also modalities $[p]$ for each state path p , with a formula $[p]\varphi$ being true at state x when φ is true at $p_A(\alpha(x))$, provided that the latter is defined. Thus $[p]\varphi$ expresses the assertion “after the state-transition $p_A \circ \alpha$, φ will be true”.

Now if the the output set D associated with an observation path p is infinite, then the language of [35, 37] is unable to express the condition that the transition function $p_A \circ \alpha$ is defined, i.e. that there is a “terminating”, or “halting”, transition induced by p . If $p_A(\alpha(x))$ exists, then the formula $(p)c$ is true at x for some (indeed for one) $c \in D$. Thus the requirement that $p_A(\alpha(x))$ be defined is expressed by the disjunction of the infinite set of formulas $\{(p)c : c \in D\}$. Hence the condition that $p_A(\alpha(x))$ be *undefined* is expressed by the *conjunction* of $\{\neg(p)c : c \in D\}$. But conjunctions and disjunctions of infinite sets do not exist in a finitary language. This is

the essential reason why the canonical model constructions of [35, 28] were restricted to polynomial functors formed from finite output sets.

Here this restriction is overcome by extending the syntax to add a new *atomic* formula $(p)\downarrow$ for each path p , with the semantics

$$A, \alpha, x \models (p)\downarrow \quad \text{iff} \quad \alpha(x) \text{ belongs to the domain of } p_A.$$

There is a price for this solution: the language remains finitary, but its proof theory becomes infinitary and its semantics exhibits failures of compactness. This is inevitable and unavoidable as soon as output sets of observation paths are allowed to be infinite. To see why, consider the set of formulas

$$\{\neg(p)c : c \in D\} \cup \{(p)\downarrow\}.$$

This set is unsatisfiable, in the sense that there is no state at which all of its members can be simultaneously true, but each of its finite subsets may be satisfiable when D is infinite. Correspondingly, our proof theory should make this set inconsistent while allowing all of its finite subsets to be consistent. By the same token, the deducibility relation should have $\{\neg(p)c : c \in D\} \vdash \neg(p)\downarrow$ while allowing that $\{\neg(p)c : c \in D'\} \not\vdash \neg(p)\downarrow$ for all finite $D' \subseteq D$.

The proof theory we develop will indeed fulfil $\{\neg(p)c : c \in D\} \vdash \neg(p)\downarrow$ for all observation paths p , as well as satisfying other “inference rules” built from these by the modalities and the implication connective, such as

$$\{\psi \rightarrow [q]\neg(p)c : c \in D\} \vdash \psi \rightarrow [q]\neg(p)\downarrow.$$

Our definition of a “maximal” set of formulas will include the requirement of closure under such inference rules. A canonical T -coalgebra $\langle A_T, \alpha_T \rangle$ will be constructed with A_T as the set of all these maximal sets and a “Truth Lemma” proven, showing that

$$A_T, \alpha_T, x \models \varphi \quad \text{iff} \quad \varphi \in x.$$

From this it will follow that $\langle A_T, \alpha_T \rangle$ is a final T -coalgebra. The explanation for this reveals the naturalness of using the canonical model construction here. Each state b in a coalgebra $\langle B, \beta \rangle$ determines the “truth set”

$$\{\varphi \mid B, \beta, b \models \varphi\},$$

consisting of all formulas that are true at b . This truth set proves to be maximal, and hence is itself a member of A_T , i.e. is state of the coalgebra $\langle A_T, \alpha_T \rangle$. This defines a map from B to A_T which proves to be the unique morphism between the coalgebras.

The Truth Lemma says that the truth set $\{\varphi \mid A_T, \alpha_T, x \models \varphi\}$ of a state x in A_T is just x itself. So the states of the canonical coalgebra are precisely all the truth sets of all states of all coalgebras, and in this sense the final coalgebra represents “all possible situations”.

A *completeness* theorem also follows from the Truth Lemma, stating that if φ is a semantic consequence of Γ (i.e. φ is satisfied by any state at which Γ is satisfied), then $\Gamma \vdash \varphi$. Equivalently, if Γ is deductively consistent ($\Gamma \not\vdash \perp$), then Γ is satisfiable at some coalgebraic state. But these completeness results require the *Lindenbaum property* that every consistent set of formulas has a maximal extension. We show that this property does hold under a countably proviso on the set of infinitary inference rules of the kind exemplified above. We characterise this proviso in terms of the number of paths that T has, and give examples in Section 7 illustrating the range of possibilities for this.

Experience with infinitary logic indicates that some such cardinality constraint on completeness is to be expected (we also give examples of that experience in Section 7). Indeed we show that there are cases of incompleteness here, by exhibiting a simple polynomial functor for which the Lindenbaum property fails. In this example a T -coalgebra is any function $\alpha : A \longrightarrow \omega^{\mathbb{R}}$, with \mathbb{R} the set of real numbers and ω the set of natural numbers. The associated logic has a set Γ of formulas that is deductively consistent but has no maximal extension and is not satisfiable at any state of any T -coalgebra. This unsatisfiable Γ is uncountable, and all of its countable subsets *are* satisfiable.

Here is an outline of the paper. In the next Section we review the basic coalgebraic theory that will be used, including the notion of *bisimilarity* that gives a mathematical formulation of the concept of observational

indistinguishability, and a characterisation of bisimilarity in terms of the behaviour of path-transitions. Section 3 sets out the formal syntax and semantics of our logic for a polynomial functor T , including the basic semantic consequence relation $\Gamma \models_{\overline{T}} \varphi$, and confirms that the logic has the Hennessy-Milner property for T -coalgebras. Section 4 begins the study of proof theory, introducing axioms and certain inferentially closed sets of formulas called *theories*. Section 5 uses theories to define the deducibility relation $\vdash_{\overline{T}}$ determined by T , and establishes its main properties, as well as introducing maximal sets and developing their relationships. Section 6 constructs the canonical coalgebra and shows that it is final. Section 7 studies the Lindenbaum property and completeness theorems. Section 8 gives the above-mentioned incompleteness example. Section 9 closes the paper with a discussion of possible generalisations and questions for further study.

2 Coalgebras and paths of polynomial functors

We begin by establishing some notation concerning sets and functions. The identity function on a set A is denoted id_A . The symbol $\circ \rightarrow$ will be used for partial functions. Thus $f : A \circ \rightarrow B$ means that f is a function with codomain B whose domain, $\text{Dom } f$, is a subset of A .

The cartesian product $A_1 \times A_2$ of two sets has associated *projections* $\pi_j : A_1 \times A_2 \rightarrow A_j$ for $j \in \{1, 2\}$. The *coproduct* $A_1 + A_2$ of A_1, A_2 is their disjoint union, with injective *insertion* functions $\iota_j : A_j \rightarrow A_1 + A_2$ for $j \in \{1, 2\}$. Each element of $A_1 + A_2$ is equal to $\iota_j(x)$ for a unique j and a unique $x \in A_j$. Associated with each insertion ι_j is its partial inverse, the *extraction* function $\varepsilon_j : A_1 + A_2 \circ \rightarrow A_j$ having $\varepsilon_j(y) = x$ iff $\iota_j(x) = y$. Thus $\text{Dom } \varepsilon_j = \iota_j A_j$, i.e. $y \in \text{Dom } \varepsilon_j$ iff $y = \iota_j(x)$ for some $x \in A_j$.

The *D-th exponential* of a set A is the set A^D of all functions from set D to A . For each $d \in D$ there is the *evaluation-at-d* function $ev_d : A^D \rightarrow A$ having $ev_d(f) = f(d)$.

DEFINITION 2.1 (POLYNOMIAL FUNCTORS) A functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ assigns a set TA to each set A , and a function $Tf : TA \rightarrow TB$ to each function $f : A \rightarrow B$ in such a way that $T\text{id}_A = \text{id}_{TA}$ and $T(g \circ f) = Tg \circ Tf$. The *identity functor* Id has $\text{Id}A = A$ and $\text{Id}f = f$. For each set D , the *constant functor* \overline{D} has $\overline{D}A = D$ and $\overline{D}f = \text{id}_D$.

A functor T is *polynomial* if it can be obtained in finitely many steps from Id and/or constant functors \overline{D} with $D \neq \emptyset$ by forming product, coproduct, and exponential functors with constant exponents. These operations on functors are as follows.

- product functors: $T_1 \times T_2$ acts on sets by $A \mapsto T_1A \times T_2A$, with $(T_1 \times T_2)f$ being the function $\langle x_1, x_2 \rangle \mapsto \langle T_1(f)(x_1), T_2(f)(x_2) \rangle$.
- coproduct functors: $T_1 + T_2$ has $A \mapsto T_1A + T_2A$ on sets; while $(T_1 + T_2)f$ acts by $\iota_j(x) \mapsto \iota_j(T_j(f)(x))$.
- exponential functors: T^D has $A \mapsto (TA)^D$, while $T^D(f)$ acts by $g \mapsto T(f) \circ g$.

DEFINITION 2.2 (COMPONENTS) Any functor involved in the formation of T is a *component* of T . Formally, the set $\text{comp}(T)$ of T -components is inductively defined by putting $\text{comp}(T) = \{T\}$ if $T = \text{Id}$ or $T = \overline{D}$; $\text{comp}(T) = \{T\} \cup \text{comp}(T_1) \cup \text{comp}(T_2)$ if $T = T_1 \times T_2$ or $T = T_1 + T_2$; and $\text{comp}(T^D) = \{T^D\} \cup \text{comp}(T)$.

It is evident that $\text{comp}(T)$ is finite and always contains at least Id or some constant functor \overline{D} .

DEFINITION 2.3 (COALGEBRAS) A *T-coalgebra* is a pair $\langle A, \alpha \rangle$, where A is a set and α is a function of the form $A \rightarrow TA$. A is called the *state set* and α is called the *transition structure*.

Since A can be recovered as $\text{Dom } \alpha$ we often refer to $\langle A, \alpha \rangle$ simply as α .

DEFINITION 2.4 (COALGEBRAIC MORPHISMS) Let $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ be T -coalgebras. A function $f : A \rightarrow B$ is a *(T-)morphism* from $\langle A, \alpha \rangle$ to $\langle B, \beta \rangle$ if $\beta \circ f = Tf \circ \alpha$, i.e. the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \alpha \downarrow & & \downarrow \beta \\ TA & \xrightarrow{Tf} & TB \end{array}$$

The identity function on A is a T -morphism $\text{id}_A : \langle A, \alpha \rangle \longrightarrow \langle A, \alpha \rangle$. The T -coalgebras and their morphisms form a category under functional composition of morphisms. An isomorphism in this category is a bijective T -morphism.

There are many illustrations in the literature (e.g. [34, 18, 38, 39, 23, 22]) showing how data structures and state-based systems can be presented as coalgebras.

DEFINITION 2.5 (FINAL COALGEBRA) A T -coalgebra $\langle A, \alpha \rangle$ is *final* if, for any T -coalgebra $\langle B, \beta \rangle$, there exists a unique T -morphism from $\langle B, \beta \rangle$ to $\langle A, \alpha \rangle$.

Thus a final T -coalgebra is a terminal object in the category of T -coalgebras, so any two final coalgebras are isomorphic. There have been several studies of conditions on T that ensure there is a final T -coalgebra [3, 4, 39, 25]. In particular, a final coalgebra exists for all polynomial functors.

DEFINITION 2.6 (BISIMULATION [3] & BISIMILARITY) For T -coalgebras $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ a relation $R \subseteq A \times B$ is a (T) -*bisimulation* from α to β if there exists a transition structure $\rho : R \longrightarrow TR$ on R such that the projections are morphisms from $\langle R, \rho \rangle$ to α and β , i.e. the following diagram commutes:

$$\begin{array}{ccccc} A & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & B \\ \alpha \downarrow & & \downarrow \rho & & \downarrow \beta \\ TA & \xleftarrow{T\pi_1} & TR & \xrightarrow{T\pi_2} & TB \end{array}$$

The union of any collection of bisimulations from α to β is a bisimulation [39, Section 5]. Hence there is a largest bisimulation from α to β , called *bisimilarity*, which we denote by \sim . Two states $x \in A$ and $y \in B$ are *bisimilar*, $x \sim y$, if there exists a bisimulation $R \subseteq A \times B$ with $\langle x, y \rangle \in R$. Bisimilarity is a mathematical formulation of the notion of observational/behavioural indistinguishability.

DEFINITION 2.7 (PATHS) A *path* is a finite list (possibly empty) of symbols of the kinds $\pi_1, \pi_2, \varepsilon_1, \varepsilon_2, \text{ev}_d$. We write $p.q$ for the associative operation of concatenation of lists p and q . The notation $T \xrightarrow{p} S$ means that p is a path from functor T to functor S , defined inductively on the formation of T as follows:

- $T \xrightarrow{\langle \rangle} T$, where $\langle \rangle$ is the empty path
- $T_1 \times T_2 \xrightarrow{\pi_j.q} S$, whenever $j \in \{1, 2\}$ and $T_j \xrightarrow{q} S$
- $T_1 + T_2 \xrightarrow{\varepsilon_j.q} S$, whenever $j \in \{1, 2\}$ and $T_j \xrightarrow{q} S$
- $T^D \xrightarrow{\text{ev}_d.q} S$ for every $d \in D$, whenever $T \xrightarrow{q} S$.

It is evident that if there is a path $T \xrightarrow{p} S$, then S is one of the components of T . Conversely, if $S \in \text{comp}(T)$ then there exists a path from T to S . Paths can be composed by concatenating lists: if $T_1 \xrightarrow{p} T_2$ and $T_2 \xrightarrow{q} T_3$, then $T_1 \xrightarrow{p.q} T_3$. A path $T \xrightarrow{p} S$ is a *state path* if $S = \text{Id}$, and an *observation path* if $S = \overline{D}$ for some set D .¹ Every path can be extended either to a state path or to an observation path.

DEFINITION 2.8 (PATH FUNCTIONS) A path $T \xrightarrow{p} S$ induces a partial function $p_A : TA \circ \longrightarrow SA$ for every set A , defined by induction on the length of p as follows:

- $\langle \rangle_A : TA \longrightarrow TA$ is the identity function id_{TA} and is total.
- $(\pi_j.p)_A = p_A \circ \pi_j$, the composition of $T_1A \times T_2A \xrightarrow{\pi_j} T_jA \xrightarrow{p_A} SA$. Hence $x \in \text{Dom}(\pi_j.p)_A$ iff $\pi_j(x) \in \text{Dom} p_A$.
- $(\varepsilon_j.p)_A = p_A \circ \varepsilon_j$, the composition of $T_1A + T_2A \xrightarrow{\varepsilon_j} T_jA \xrightarrow{p_A} SA$. Hence $x \in \text{Dom}(\varepsilon_j.p)_A$ iff $x \in \text{Dom} \varepsilon_j$ and $\varepsilon_j(x) \in \text{Dom} p_A$.

¹Observation paths and state paths are called “positions” in [37].

- $(ev_d.p)_A = p_A \circ ev_d$, the composition of $(TA)^D \xrightarrow{ev_d} TA \xrightarrow{p_A} SA$. Hence $f \in \text{Dom}(ev_d.p)_A$ iff $f(d) \in \text{Dom } p_A$.

Concatenation of paths corresponds to composition of functions, in the sense that $(p.q)_A = q_A \circ p_A$. Note that if no extraction symbol ε_j occurs in p , then p_A is always a total function. It is the presence of coproducts that introduces partiality into this theory. ■

An important role played by paths is to provide a characterization of a T -bisimulation as a relation that is “preserved” by the partial functions induced by state and observation paths from T . To explain this we adopt the convention that whenever we write “ $f(x) Q g(y)$ ” for some relation Q and some partial functions f and g we mean that $x \in \text{Dom } f$ iff $y \in \text{Dom } g$ and if both $f(x)$ and $g(y)$ are defined then $\langle f(x), g(y) \rangle \in Q$. In particular, we use this convention for the relation “ $=$ ”. Proofs of the following results can be found in [12, Section 5].

THEOREM 2.9 *Let $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ be T -coalgebras.*

(1) $R \subseteq A \times B$ is a T -bisimulation iff $x R y$ implies

- $p_A(\alpha(x)) = p_B(\beta(y))$ for every observation path $T \xrightarrow{p} \overline{D}$, and
- $p_A(\alpha(x)) R p_B(\beta(y))$ for every state path $T \xrightarrow{p} \text{Id}$.

(2) $f : A \rightarrow B$ is a morphism from α to β iff

- $p_A(\alpha(x)) = p_B(\beta(f(x)))$ for every observation path p , and
- $f(p_A(\alpha(x))) = p_B(\beta(f(x)))$ for every state path p .

(3) If $f : A \rightarrow B$ is a morphism from α to β , then for any path $T \xrightarrow{p} S$,

- $\alpha(x) \in \text{Dom } p_A$ iff $\beta(f(x)) \in \text{Dom } p_B$.

■

3 Syntax and Semantics of Formulas

We now define a Hennessy-Milner style modal language for a polynomial functor T which will remain fixed throughout the rest of the paper. The language consists of propositional formulas that are “constant”, i.e. there are no propositional variables.

DEFINITION 3.1 The set of *well-formed formulas* (wff) Φ_T is defined inductively to consist of the following:

- \perp
- $(p)\downarrow$, for every path $T \xrightarrow{p} S$
- $(p)c$, for every observation path $T \xrightarrow{p} \overline{D}$ and $c \in D$
- $\varphi \rightarrow \psi$, for every $\varphi, \psi \in \Phi_T$
- $[p]\varphi$, for every state path $T \xrightarrow{p} \text{Id}$ and $\varphi \in \Phi_T$.

The connectives $\neg, \top, \vee, \wedge, \leftrightarrow$ are defined in the usual way from \perp and \rightarrow . In particular $\neg\varphi$ is $\varphi \rightarrow \perp$. We also write $(p)\uparrow$ for the wff $\neg(p)\downarrow$.

Φ_T includes the formulas of the language of [37], which are generated essentially as above but without the formation of $(p)\downarrow$. Note that Φ_T may be uncountable, since the sets D may be uncountable, and/or because there may be uncountably many paths p .

DEFINITION 3.2 (TRUTH AND CONSEQUENCE) The truth relation $\alpha, x \models \varphi$ is defined inductively on the formation of φ , as follows for all T -coalgebras $\langle A, \alpha \rangle$, with $x \in A$:

- $\alpha, x \not\models \perp$
- $\alpha, x \models (p)\downarrow$ iff $\alpha(x) \in \text{Dom } p_A$
- $\alpha, x \models (p)c$ iff $\alpha, x \models (p)\downarrow$ and $p_A(\alpha(x)) = c$
- $\alpha, x \models [p]\varphi$ iff $\alpha, x \models (p)\downarrow$ implies $\alpha, p_A(\alpha(x)) \models \varphi$
- $\alpha, x \models \varphi \rightarrow \psi$ iff $\alpha, x \models \varphi$ implies $\alpha, x \models \psi$

Thus $\alpha, x \models \neg\varphi$ iff $\alpha, x \not\models \varphi$, and similarly the other standard connectives have their usual semantics.

We say that φ is *true* at x in α , or x *satisfies* φ , if $\alpha, x \models \varphi$, and that φ is *valid* in α , $\alpha \models \varphi$, if it is true at all states in A . The set $\{\psi \in \Phi_T \mid \alpha, x \models \psi\}$ is called the *truth set* of x in α . A set $\Gamma \subseteq \Phi_T$ is *true/satisfied* at x in α , $\alpha, x \models \Gamma$, if $\alpha, x \models \varphi$ for all $\varphi \in \Gamma$.

Semantic consequence relations are defined by

$$\begin{aligned} \Gamma \models^\alpha \varphi &\text{ iff } (\forall x \in A) \alpha, x \models \Gamma \text{ implies } \alpha, x \models \varphi \\ \Gamma \models_{\overline{T}} \varphi &\text{ iff } \Gamma \models^\alpha \varphi \text{ for all } T\text{-coalgebras } \alpha. \end{aligned}$$

Satisfaction of formulas is invariant under the action of morphisms:

LEMMA 3.3 *Let $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ be T -coalgebras and $f : A \rightarrow B$ be a morphism from α to β . Then for every $\varphi \in \Phi_T$, $\alpha, x \models \varphi$ iff $\beta, f(x) \models \varphi$.*

Proof This is proven for all $x \in A$ by induction on the construction of φ .

For any path p from T we have $\alpha(x) \in \text{Dom } p_A$ iff $\beta(f(x)) \in \text{Dom } p_B$, by Theorem 2.9(3), hence $\alpha, x \models (p)\downarrow$ iff $\beta, f(x) \models (p)\downarrow$.

If p is an observation path, then $p_A(\alpha(x)) = p_B(\beta(f(x)))$ by 2.9(2), so $p_A(\alpha(x)) = c$ iff $p_B(\beta(f(x))) = c$, and hence $\alpha, x \models (p)c$ iff $\beta, f(x) \models (p)c$.

If p is a state path, then $f(p_A(\alpha(x))) = p_B(\beta(f(x)))$, so assuming the result for φ gives $\alpha, p_A(\alpha(x)) \models \varphi$ iff $\beta, f(p_A(\alpha(x))) \models \varphi$ iff $\beta, p_B(\beta(f(x))) \models \varphi$, which leads to $\alpha, x \models [p]\varphi$ iff $\beta, f(x) \models [p]\varphi$.

The cases of the propositional connectives are straightforward. ■

It is a pertinent question as to what makes a formal language appropriate for a given class of coalgebras. We have stressed the Hennessy-Milner property that logical equivalence of states should coincide with bisimilarity. That property is already possessed by the language of [37]. But it is desirable also that the language be powerful enough to allow effective model-building, and that is why we needed to add the formulas of type $(p)\downarrow$. Of course adding more formulas preserves the property of logically distinguishing states that are not bisimilar, but then we need to check that it does not also cause some states to be distinguished that *are* bisimilar.

THEOREM 3.4 (HENNESSY-MILNER PROPERTY) *Let $\langle A, \alpha \rangle$ and $\langle B, \beta \rangle$ be T -coalgebras with $x \in A$ and $y \in B$. The following are equivalent:*

- (1) $x \sim y$
- (2) $\alpha, x \models \varphi$ iff $\beta, y \models \varphi$ for every $\varphi \in \Phi_T$.

Proof (1) \Rightarrow (2): Let $\langle R, \rho \rangle$ be a bisimulation from α to β with $x R y$. The projections $\pi_1 : R \rightarrow A$ and $\pi_2 : R \rightarrow B$ are morphisms so by Lemma 3.3, $\alpha, x \models \varphi$ iff $\rho, \langle x, y \rangle \models \varphi$ iff $\beta, y \models \varphi$.

(1) \Leftarrow (2): If (2) holds, then in particular $\alpha, x \models \varphi$ iff $\beta, y \models \varphi$ for every φ that has no occurrence of a ‘‘halting’’ formula $(p)\downarrow$, and [37, Proposition 2.8] shows that this is sufficient to prove $x \sim y$. But we sketch a proof anyway.

Let $R = \{\langle x, y \rangle \in A \times B \mid (2) \text{ holds}\}$. It is enough to show that R is a bisimulation relation, for then $\langle x, y \rangle \in R$ implies $x \sim y$. By Theorem 2.9(1) we need to show that

- (i) $p_A(\alpha(x)) = p_B(\beta(y))$, for every observation path p , and
- (ii) $p_A(\alpha(x)) R p_B(\beta(y))$, for every state path p .

The property (i) is captured by the fact that $\alpha, x \models (p)c$ iff $\beta, y \models (p)c$, and the property (ii) is captured by the fact that $\alpha, x \models [p]\varphi$ iff $\beta, y \models [p]\varphi$. ■

4 Axioms, Inference Rules and Theories

By “proof theory” we mean the study of a binary relation $\Gamma \vdash \varphi$, from sets Γ of formulas to formulas φ , that is intended to capture the notion that φ is deducible/derivable/provable from members of Γ by using certain axioms and rules of inference. The definition of \vdash will depend on the syntactic shape of the formulas involved, together with basic set-theoretic properties of sets of formulas. We then seek to obtain *soundness* and *completeness* results to show that \vdash is identical to some semantically defined relation, such as the consequence relation \models_T of the previous section.

There is more than one approach to defining deducibility relations. Classically, $\Gamma \vdash \varphi$ was often taken to mean that there is a *proof-sequence* from Γ to φ , i.e. a sequence of formulas ending at φ , with each member of the sequence being either a member of Γ , an axiom, or derivable from previous members of the sequence by a rule of inference. This approach works well when all inference rules have finitely many premisses. Then proof-sequences can be constrained to be finite, and the relation \vdash is *finitary* in the sense that whenever $\Gamma \vdash \varphi$ then $\Gamma' \vdash \varphi$ for some finite set $\Gamma' \subseteq \Gamma$. But if there are rules with infinitely many premisses, then proof-sequences may be transfinite in length, and their analysis requires the arithmetic of infinite ordinals [24].

Working with concatenations of transfinite sequences can be cumbersome. Consequently a more “axiomatic” approach to \vdash was developed, using the general theory of *inductive definitions*, in which an inductively defined set is given as the *least fixed point* of a monotonic operator on sets [2]. The operator in question takes each set of formulas to its closure under the relevant axioms and rules of inference. A fixed point of this operator, i.e. a set of formulas that is closed under the axioms and rules, is called a *theory*, and $\Gamma \vdash \varphi$ is defined to hold when φ belongs to every theory extending Γ . This results in the set $\{\varphi \mid \Gamma \vdash \varphi\}$ of formulas deducible from Γ being inductively characterised as the least theory extending Γ . See for example [1, 6] for extensive use of this kind of formulation of infinitary proof theory.

We will take this approach to \vdash here, adapting a framework for infinitary modal logic developed in [9], but will also make some use of classical proof-sequences in Section 8. We first discuss axioms, rules and theories that are particular to our coalgebraic language, and then in the next Section introduce deducibility relations and “maximal” sets of formulas.

DEFINITION 4.1 (AXIOMS) The set of *T-axioms*, Ax_T , consists of the following formulas.

1. All instances of propositional tautologies.
2. $(\langle \rangle) \downarrow$

For each path $T \xrightarrow{p} S_1 \times S_2$ and $j \in \{1, 2\}$:

3. $(p) \downarrow \leftrightarrow (p.\pi_j) \downarrow$

For each path $T \xrightarrow{p} S^D$ and all $d \in D$:

4. $(p) \downarrow \leftrightarrow (p.ev_d) \downarrow$

For each path $T \xrightarrow{p} S_1 + S_2$:

5. $[(p) \downarrow \leftrightarrow (p.\varepsilon_1) \downarrow \vee (p.\varepsilon_2) \downarrow] \wedge \neg[(p.\varepsilon_1) \downarrow \wedge (p.\varepsilon_2) \downarrow]$

For each observation path $T \xrightarrow{p} \overline{D}$ and all $c, d \in D$ such that $c \neq d$:

6. $(p)c \rightarrow \neg(p)d$
7. $(p)d \rightarrow (p) \downarrow$

For each state path $T \xrightarrow{p} \text{Id}$ and all $\varphi, \psi \in \Phi_T$:

8. $\neg[p]\varphi \rightarrow [p]\neg\varphi$
9. $(p) \downarrow \rightarrow \neg[p] \perp$

10. $(p)\uparrow \rightarrow [p]\perp$

11. $[p](\varphi \rightarrow \psi) \rightarrow ([p]\varphi \rightarrow [p]\psi)$.

These axioms express natural properties of the structure of coalgebras and their path functions. Axiom 3 expresses the fact that the path function $(p.\pi_j)_A = \pi_j \circ p_A$ is defined precisely when p_A is defined, since the projection function π_j is total. Likewise for axiom 4, as evaluation functions are total. Axiom 5 expresses the fact that the domain of p_A is the disjoint union of the domains of $(p.\varepsilon_1)_A$ and $(p.\varepsilon_2)_A$. Axiom 11 is the well known axiom K (for “Kripke”) from classical modal logic. Axioms 9 and 10 together express the fact that $(p)\downarrow$ is true precisely when $[p]\perp$ is not, and could have been presented as the biconditional $(p)\downarrow \leftrightarrow \neg [p]\perp$. But each has its own role to play (in Lemmas 4.11(4) and 6.5 respectively), so it is convenient to separate them.

With similar observations about the other axioms, we are led to the conclusion that

THEOREM 4.2 *All T-axioms are valid in all T-coalgebras.* ■

We now begin the study of syntactic closure properties of sets of formulas.

DEFINITION 4.3 (MODAL CLOSURE) A set Δ of formulas is *modally closed* if $[p]\varphi \in \Delta$ whenever $\varphi \in \Delta$ and p is any state path from T . The *modal closure* Γ^* of any set of formulas Γ is the smallest modally closed set that extends Γ . This Γ^* consists of all formulas of the form $[p_0]\dots[p_{n-1}]\varphi$ where $\varphi \in \Gamma$ and p_0, \dots, p_{n-1} is any finite sequence (possibly empty) of state paths.

Many systems of modal logic have a set of theorems that is modally closed: if φ is a theorem then so is $\Box\varphi$, where \Box is any modality of “box” type, like our modalities $[p]$. Indeed the presentation of such logics usually has the inference rule *from φ infer $\Box\varphi$* , known as the rule of Necessitation, or the “Box rule”. Alternatively, for finitary logics this rule can be stipulated just for the case that φ is an axiom, and then derived for theorems in general by using the appropriate version of axiom K (our axiom 11). Here we will achieve this effect by taking the modal closure Ax_T^* of the set of T-axioms and building it in to the notion of a *theory*, which we define next. This will allow certain versions of the Box rule to be derived later (see Lemma 5.4).

DEFINITION 4.4 A set $\Gamma \subseteq \Phi_T$ is:

- *closed under Detachment* if $\varphi, \varphi \rightarrow \psi \in \Delta$ implies $\psi \in \Delta$;
- a *theory* if it includes the modal closure Ax_T^* of the set of axioms and is closed under Detachment;
- *negation complete* if for every $\varphi \in \Phi_T$, either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$;
- *\perp -free* if $\perp \notin \Gamma$.

LEMMA 4.5

(1) *If Γ is a negation complete theory, then for every $\varphi, \psi \in \Phi_T$,*

$$\varphi \rightarrow \psi \in \Gamma \quad \text{iff} \quad \varphi \in \Gamma \text{ implies } \psi \in \Gamma.$$

(2) *If Γ is a \perp -free negation complete theory, then:*

$$\begin{aligned} \neg\varphi \in \Gamma & \quad \text{iff} \quad \varphi \notin \Gamma \\ \varphi \wedge \psi \in \Gamma & \quad \text{iff} \quad \varphi \in \Gamma \text{ and } \psi \in \Gamma \\ \varphi \vee \psi \in \Gamma & \quad \text{iff} \quad \varphi \in \Gamma \text{ or } \psi \in \Gamma \\ \varphi \leftrightarrow \psi \in \Gamma & \quad \text{iff} \quad (\varphi \in \Gamma \text{ iff } \psi \in \Gamma). \end{aligned}$$

(3) *Every truth set is a negation complete \perp -free theory.*

Proof (1) and (2) follow by standard arguments, using the fact that all tautologies are in Γ by axiom 1.

For (3), observe first that the set $\Gamma^\alpha = \{\psi \mid \alpha \models \psi\}$ of formulas valid in coalgebra α contains all axioms by Theorem 4.2, and is modal closed because if ψ is true at all states of α then so is $[p]\psi$ for any state path p . Hence $Ax_T^* \subseteq \Gamma^\alpha$. Then each truth set $\{\psi \mid \alpha, x \models \psi\}$ includes Γ^α and hence includes Ax_T^* ; is closed under Detachment by the semantics of $\varphi \rightarrow \psi$; is negation complete by the semantics of $\neg\varphi$; and is \perp -free by the semantics of \perp . ■

DEFINITION 4.6 Let 2^{Φ_T} be the powerset of Φ_T , with $\Gamma \in 2^{\Phi_T}$ and $\mathcal{R} \subseteq 2^{\Phi_T} \times \Phi_T$.

- An *inference rule*, or just *rule*, is a pair $\langle \Sigma, \varphi \rangle \in 2^{\Phi_T} \times \Phi_T$. Here Σ may be thought of as a set of *premisses*, and φ as a *conclusion*.
- Γ is *closed under the rule* $\langle \Sigma, \varphi \rangle$ if $\Sigma \subseteq \Gamma$ implies $\varphi \in \Gamma$, i.e. if $\Sigma \not\subseteq \Gamma$ or $\varphi \in \Gamma$.
- Γ is *\mathcal{R} -closed* if it is closed under every rule belonging to \mathcal{R} .
- Γ is an *\mathcal{R} -theory* if it is \mathcal{R} -closed and is a theory (i.e. is closed under Detachment and $Ax_T^* \subseteq \Gamma$). In particular, an \emptyset -theory is just a theory as in Definition 4.4. ■

The functor T determines a special relation $\mathcal{R}_T \subseteq 2^{\Phi_T} \times \Phi_T$ that is central to our proof theory and is defined as follows:

DEFINITION 4.7

- For each observation path $T \xrightarrow{p} \overline{D}$, \mathcal{I}_p is the inference rule $\langle \{\neg(p)d \mid d \in D\}, (p)\uparrow \rangle$.
- For each state path q , $[q]\Sigma = \{[q]\psi \mid \psi \in \Sigma\}$.
- For $\Sigma \cup \{\psi\} \subseteq \Phi_T$, $\psi \rightarrow \Sigma = \{\psi \rightarrow \theta \mid \theta \in \Sigma\}$.
- \mathcal{R}_T is the smallest relation (i.e. intersection of all relations) satisfying
 - $\mathcal{I}_p \in \mathcal{R}_T$ for all observation paths p ;
 - if $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$, then $\langle [q]\Sigma, [q]\varphi \rangle \in \mathcal{R}_T$ for every state path q ;
 - if $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$, then $\langle \psi \rightarrow \Sigma, \psi \rightarrow \varphi \rangle \in \mathcal{R}_T$ for every $\psi \in \Phi_T$.

THEOREM 4.8 $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$ implies $\Sigma \models_{\overline{T}} \varphi$.

Proof Each T -coalgebra α has $\{\neg(p)d \mid d \in D\} \models^\alpha (p)\uparrow$ for any observation path p . Also, from $\Sigma \models^\alpha \varphi$ it follows that $[q]\Sigma \models^\alpha [q]\varphi$ for all state paths q and $\psi \rightarrow \Sigma \models^\alpha \psi \rightarrow \varphi$ for all $\psi \in \Phi_T$. Thus the relation $\{\langle \Sigma, \varphi \rangle \mid \Sigma \models^\alpha \varphi\}$ satisfies the three closure properties defining \mathcal{R}_T , and so includes \mathcal{R}_T as the smallest such relation. Hence $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$ implies $\Sigma \models^\alpha \varphi$ for all T -coalgebras α . ■

COROLLARY 4.9 Every truth set is an \mathcal{R}_T -theory.

Proof If Γ is the truth set of x in α , and $\Sigma \subseteq \Gamma$ with $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$, then $\alpha, x \models \Sigma$ and $\Sigma \models_{\overline{T}} \varphi$, so $\alpha, x \models \varphi$ and $\varphi \in \Gamma$. Thus Γ is \mathcal{R}_T -closed. But Γ is a theory by Lemma 4.5(3). ■

DEFINITION 4.10 For each state path p and $\Delta \subseteq \Phi_T$, let $\Delta_p = \{\varphi \mid [p]\varphi \in \Delta\}$.

This operation is crucial both to the proof theory of the modalities $[p]$ and to the construction of a canonical/final coalgebra.

LEMMA 4.11 For any $\Delta \subseteq \Phi_T$, any $\mathcal{R} \subseteq 2^{\Phi_T} \times \Phi_T$, and any state path p :

- (1) If Δ is a theory, then so is Δ_p .
- (2) If $\langle \Sigma, \varphi \rangle \in \mathcal{R}$ implies $\langle [p]\Sigma, [p]\varphi \rangle \in \mathcal{R}$, then if Δ is \mathcal{R} -closed so is Δ_p .
- (3) If Δ is negation complete, then so is Δ_p .
- (4) If $(p)\downarrow \in \Delta$, then if Δ is \perp -free so is Δ_p .

Proof

- (1) If $\varphi \in Ax_T^*$, then $[p]\varphi \in Ax_T^* \subseteq \Delta$, and so $\varphi \in \Delta_p$. Hence $Ax_T^* \subseteq \Delta_p$. If $\varphi \rightarrow \psi, \varphi \in \Delta_p$, then $[p](\varphi \rightarrow \psi), [p]\varphi \in \Delta$. But every instance of axiom 11 is in Δ , so closure of Δ under Detachment gives $[p]\psi \in \Delta$, hence $\psi \in \Delta_p$. Thus Δ_p is closed under Detachment.

- (2) Let Δ be \mathcal{R} -closed. Then if $\langle \Sigma, \varphi \rangle \in \mathcal{R}$ and $\Sigma \subseteq \Delta_p$, we get $[p]\Sigma \subseteq \Delta$ and $\langle [p]\Sigma, [p]\varphi \rangle \in \mathcal{R}$, hence $[p]\varphi \in \Delta$. This shows that Δ_p is \mathcal{R} -closed.
- (3) If Δ is negation complete, then for every φ either $[p]\varphi \in \Delta$ or $\neg [p]\varphi \in \Delta$. But $\neg [p]\varphi \rightarrow [p]\neg\varphi \in \Delta$ by axiom 8, so either $[p]\varphi \in \Delta$ or $[p]\neg\varphi \in \Delta$, by Detachment. Thus either $\varphi \in \Delta_p$ or $\neg\varphi \in \Delta_p$, showing that Δ_p is negation complete.
- (4) If $(p)\downarrow \in \Delta$, then as $(p)\downarrow \rightarrow \neg [p]\perp \in \Delta$ by axiom 9, $\neg [p]\perp \in \Delta$ by Detachment, i.e. $[p]\perp \rightarrow \perp \in \Delta$. Thus if $\perp \notin \Delta$, then $[p]\perp \notin \Delta$ by Detachment, and so $\perp \notin \Delta_p$. ■

5 Deducibility and Maximality

We now use \mathcal{R}_T -theories to define a deducibility relation $\vdash_{\bar{T}}$ that will eventually be seen to be identical to the semantic consequence relation $\models_{\bar{T}}$ for many T .

DEFINITION 5.1 Let $\Gamma \vdash_{\bar{T}} \varphi$ mean that $\varphi \in \bigcap \{ \Delta \mid \Gamma \subseteq \Delta \text{ and } \Delta \text{ is an } \mathcal{R}_T\text{-theory} \}$ i.e. that φ belongs to every \mathcal{R}_T -theory extending Γ .

Since the class of \mathcal{R}_T -theories is closed under intersection, it follows that the set $\{ \varphi \in \Phi_T \mid \Gamma \vdash_{\bar{T}} \varphi \}$ is the *smallest \mathcal{R}_T -theory extending Γ* .

A number of properties follow directly from this definition, and are left to the reader to check:

LEMMA 5.2

- (1) If $\varphi \in \Gamma \cup Ax_T$, then $\Gamma \vdash_{\bar{T}} \varphi$.
- (2) \mathcal{R}_T -theories are deductively closed: if $\Gamma \vdash_{\bar{T}} \varphi$ and Γ is itself an \mathcal{R}_T -theory, then $\varphi \in \Gamma$.
- (3) $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$ implies $\Sigma \vdash_{\bar{T}} \varphi$. ■

THEOREM 5.3 (SOUNDNESS) If $\Gamma \vdash_{\bar{T}} \varphi$ then $\Gamma \models_{\bar{T}} \varphi$.

Proof Suppose $\Gamma \vdash_{\bar{T}} \varphi$ and $\alpha, x \models \Gamma$ for some T -coalgebra $\langle A, \alpha \rangle$ and $x \in A$. We need to show $\alpha, x \models \varphi$, so let $\Delta = \{ \psi \mid \alpha, x \models \psi \}$. Then Δ is an \mathcal{R}_T -theory by Corollary 4.9. Therefore, since $\Gamma \subseteq \Delta$ and $\Gamma \vdash_{\bar{T}} \varphi$, $\varphi \in \Delta$. Hence $\alpha, x \models \varphi$. ■

LEMMA 5.4

- (1) Cut rule (CT): If $\Gamma \vdash_{\bar{T}} \psi$ for all $\psi \in \Delta$ and $\Delta \vdash_{\bar{T}} \varphi$, then $\Gamma \vdash_{\bar{T}} \varphi$.
- (2) Deduction theorem (DT): $\Gamma \cup \{ \varphi \} \vdash_{\bar{T}} \psi$ implies $\Gamma \vdash_{\bar{T}} \varphi \rightarrow \psi$.
- (3) Monotonicity: If $\Gamma \vdash_{\bar{T}} \varphi$ and $\Gamma \subseteq \Delta$, then $\Delta \vdash_{\bar{T}} \varphi$.
- (4) Detachment: If $\Gamma \vdash_{\bar{T}} \varphi$ and $\Gamma \vdash_{\bar{T}} \varphi \rightarrow \psi$, then $\Gamma \vdash_{\bar{T}} \psi$.
- (5) If $\Gamma \vdash_{\bar{T}} \varphi$ and $\Gamma \cup \{ \varphi \} \vdash_{\bar{T}} \perp$, then $\Gamma \vdash_{\bar{T}} \perp$.
- (6) If $\Gamma \cup \{ \neg\varphi \} \vdash_{\bar{T}} \perp$, then $\Gamma \vdash_{\bar{T}} \varphi$.
- (7) Box rule: If $\Gamma \vdash_{\bar{T}} \varphi$, then $[p]\Gamma \vdash_{\bar{T}} [p]\varphi$ for all state paths p .
- (8) Implication rule: If $\Gamma \vdash_{\bar{T}} \varphi$, then $\psi \rightarrow \Gamma \vdash_{\bar{T}} \psi \rightarrow \varphi$.
- (9) If $\emptyset \vdash_{\bar{T}} \varphi$, then $\Gamma \vdash_{\bar{T}} [p]\varphi$ for all $\Gamma \subseteq \Phi_T$.

Proof (1)–(8) can be proven as in [9], but we give the proof for (7) since our use of the set Ax_T^* is slightly different to the setup in [9]. Suppose then that $\Gamma \vdash_{\bar{T}} \varphi$. To show $[p]\Gamma \vdash_{\bar{T}} [p]\varphi$, let Δ be an \mathcal{R}_T -theory with $[p]\Gamma \subseteq \Delta$. Then $\Gamma \subseteq \Delta_p$. But Δ_p is an \mathcal{R}_T -theory by Lemma 4.11, so from $\Gamma \vdash_{\bar{T}} \varphi$ we get $\varphi \in \Delta_p$, hence $[p]\varphi \in \Delta$.

For (9), if $\emptyset \vdash_{\bar{T}} \varphi$ then $[p]\emptyset \vdash_{\bar{T}} [p]\varphi$ by the Box rule (7). But $[p]\emptyset = \emptyset \subseteq \Gamma$, so then $\Gamma \vdash_{\bar{T}} [p]\varphi$ by Monotonicity (3). ■

A deducibility relation gives rise to various notions of deductive consistency:

DEFINITION 5.5 A set Γ of wffs is

- \vdash_{\neg} -inconsistent if $\Gamma \vdash_{\neg} \perp$, and \vdash_{\neg} -consistent otherwise;
- *finitely \vdash_{\neg} -consistent* if all finite subsets of Γ are \vdash_{\neg} -consistent;
- *maximally finitely \vdash_{\neg} -consistent* if it is finitely \vdash_{\neg} -consistent but has no proper extension that is finitely \vdash_{\neg} -consistent;
- *maximal* if it is a negation complete and \vdash_{\neg} -consistent \mathcal{R}_{\neg} -theory.

THEOREM 5.6 *Every truth-set is maximal.*

Proof Any truth set is \vdash_{\neg} -consistent: if $\Gamma = \{\psi \mid \alpha, x \models \psi\}$, then $\alpha, x \models \Gamma$ so $\Gamma \not\vdash_{\neg} \perp$, and therefore $\Gamma \not\vdash_{\neg} \perp$ by Soundness (5.3). It follows from Lemma 4.5(3) and Corollary 4.9 that every truth-set is maximal. ■

The most obvious example of a \vdash_{\neg} -inconsistent set is one containing \perp . Then next most obvious is one including $\{\varphi, \neg\varphi\} = \{\varphi, \varphi \rightarrow \perp\}$ for some φ , since any theory containing $\varphi, \neg\varphi$ will contain \perp by closure under Detachment.

Here are the main relationships between the various notions of consistency described in Definition 5.5:

LEMMA 5.7

- (1) *If Γ is finitely \vdash_{\neg} -consistent, then so is one of $\Gamma \cup \{\varphi\}$ and $\Gamma \cup \{\neg\varphi\}$ for each φ .*
- (2) *If Γ is negation complete and finitely \vdash_{\neg} -consistent, then it is a theory.*
- (3) *Γ is maximally finitely \vdash_{\neg} -consistent iff it is negation complete and finitely \vdash_{\neg} -consistent.*
- (4) *If Γ is maximal, then it is maximally finitely \vdash_{\neg} -consistent and hence is maximally \vdash_{\neg} -consistent.*
- (5) *An \mathcal{R}_{\neg} -theory is \vdash_{\neg} -consistent iff it is \perp -free.*

Proof (1)–(4) can be proven as in [9]. For (5), observe that an \mathcal{R}_{\neg} -theory Γ has $\Gamma \vdash_{\neg} \varphi$ iff $\varphi \in \Gamma$ in general by Lemma 5.2. In particular $\Gamma \not\vdash_{\neg} \perp$ iff $\perp \notin \Gamma$. ■

The following result will be needed in our construction of a final coalgebra.

LEMMA 5.8 *Let Γ be a maximal set of wffs. For each observation path $T \xrightarrow{p} \overline{D}$, if $(p)\downarrow \in \Gamma$ then $(p)d \in \Gamma$ for a unique $d \in D$.*

Proof Let $(p)\downarrow \in \Gamma$. Then $(p)\uparrow \notin \Gamma$, as Γ is \vdash_{\neg} -consistent. Since Γ is \mathcal{R}_{\neg} -closed it is closed under the rule $\mathcal{I}_p = \langle \{\neg(p)d \mid d \in D\}, (p)\uparrow \rangle$, so then $\neg(p)d \notin \Gamma$ for at least one d . Hence $(p)d \in \Gamma$ by negation completeness. In fact d is unique, for if $d \neq c \in D$, then $(p)c \rightarrow \neg(p)d \in \Gamma$ by axiom 6, so $(p)c \notin \Gamma$ by closure under Detachment. ■

The companion result for state paths is

LEMMA 5.9 *Let Δ be a maximal set of wffs. For each state path $T \xrightarrow{p} \text{Id}$, if $(p)\downarrow \in \Gamma$, then Δ_p is maximal.*

Proof By Lemma 4.11, as Δ is a negation complete \mathcal{R}_{\neg} -theory, so too is Δ_p , and as $(p)\downarrow \in \Gamma$ and $\perp \notin \Delta$, $\perp \notin \Delta_p$. But then by Lemma 5.7(5), Δ_p is \vdash_{\neg} -consistent. ■

A simpler characterisation of maximal sets will now be obtained. Inference rules of the form $\langle \psi \rightarrow \Sigma, \psi \rightarrow \varphi \rangle$ were built into the definition of \mathcal{R}_T , and hence of \vdash_T , because they are needed to establish the crucial deduction property DT (see [9, 9.3.3(4)]). It transpires that for negation complete sets, closure under such implicational rules can be derived from the other properties of \mathcal{R}_T , as we will now see.

DEFINITION 5.10 \mathcal{R}_T^- is the smallest set of inference rules satisfying

- $\mathcal{I}_p \in \mathcal{R}_T^-$ for all observation paths p ;
- if $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T^-$, then $\langle [p]\Sigma, [p]\varphi \rangle \in \mathcal{R}_T^-$ for every state path p .

LEMMA 5.11 Every negation complete \mathcal{R}_T^- -theory is an \mathcal{R}_T -theory.

Proof Let Θ be the set of all negation complete \mathcal{R}_T^- -theories. Define \mathcal{R} to be the set of all rules $\langle \Sigma, \varphi \rangle$ such that every member of Θ is closed under $\langle \Sigma, \varphi \rangle$. We will show that \mathcal{R} has the three properties defining \mathcal{R}_T . Hence $\mathcal{R}_T \subseteq \mathcal{R}$ because \mathcal{R}_T is the smallest set of rules having these properties. But each member of Θ is \mathcal{R} -closed by definition of \mathcal{R} , so then is \mathcal{R}_T -closed, and hence is an \mathcal{R}_T -theory.

First, $\mathcal{I}_p \in \mathcal{R}$ because $\mathcal{I}_p \in \mathcal{R}_T^-$ and every member of Θ is \mathcal{R}_T^- -closed. Secondly, suppose $\langle \Sigma, \varphi \rangle \in \mathcal{R}$. Then for any state path p , to show that $\langle [p]\Sigma, [p]\varphi \rangle \in \mathcal{R}$, take any $\Delta \in \Theta$ and suppose $[p]\Sigma \subseteq \Delta$. Then $\Sigma \subseteq \Delta_p$. But by Lemma 4.11, Δ_p is a negation complete \mathcal{R}_T^- -theory because Δ is, so $\Delta_p \in \Theta$ and hence Δ_p is closed under $\langle \Sigma, \varphi \rangle$, giving $\varphi \in \Delta_p$, whence $[p]\varphi \in \Delta$. This proves that any $\Delta \in \Theta$ is closed under $\langle [p]\Sigma, [p]\varphi \rangle$, so that rule is in \mathcal{R} .

Thirdly, given $\langle \Sigma, \varphi \rangle \in \mathcal{R}$ consider the rule $\langle \psi \rightarrow \Sigma, \psi \rightarrow \varphi \rangle$. If $\Delta \in \Theta$ and $\psi \rightarrow \varphi \notin \Delta$, then Lemma 4.5(1), $\psi \in \Delta$ and $\varphi \notin \Delta$. Since Δ is closed under $\langle \Sigma, \varphi \rangle$, $\Sigma \not\subseteq \Delta$. Taking a $\theta \in \Sigma$ with $\theta \notin \Delta$ we get $\psi \rightarrow \theta \notin \Delta$ as Δ is a theory, so $(\psi \rightarrow \Sigma) \not\subseteq \Delta$. Hence Δ is closed under the rule $\langle \psi \rightarrow \Sigma, \psi \rightarrow \varphi \rangle$, which thus belongs to \mathcal{R} . That completes the proof of the Lemma. ■

COROLLARY 5.12 Γ is maximal if, and only if, it is maximally finitely \vdash_T -consistent and \mathcal{R}_T^- -closed.

Proof From left to right holds by Lemma 5.7(4) and the fact that $\mathcal{R}_T^- \subseteq \mathcal{R}_T$ so \mathcal{R}_T -closure implies \mathcal{R}_T^- -closure.

Conversely, if Γ is maximally finitely \vdash_T -consistent and \mathcal{R}_T^- -closed, then by (2) and (3) of Lemma 5.7 it is a negation complete theory, hence an \mathcal{R}_T^- -theory, so is an \mathcal{R}_T -theory by Lemma 5.11. Moreover $\perp \notin \Gamma$ as Γ is finitely \vdash_T -consistent, so then Γ is \vdash_T -consistent by 5.7(5). ■

By similar arguments we can also characterise maximality without reference to the deducibility relation \vdash_T thus: Γ is maximal iff it is a negation complete \perp -free \mathcal{R}_T^- -theory. The formulation of 5.12 is however more convenient for the completeness theorem to come later.

6 The Canonical T -Coalgebra

In this Section we will see that every maximal set is a truth set. This is established by constructing a single T -coalgebra $\langle A_T, \alpha_T \rangle$ whose states are the maximal sets, and showing that each maximal set is the truth set of itself as a state in A_T . Since all truth sets are maximal (Theorem 5.6), this implies that the truth set of any state in any coalgebra is equal to the truth set of some member of A_T . It will turn out that $\langle A_T, \alpha_T \rangle$ is a final T -coalgebra.

DEFINITION 6.1 $A_T = \{x \subseteq \Phi_T \mid x \text{ is maximal}\}$.

To define the transition structure α_T we use the proof theory to associate with each path $T \xrightarrow{q} S$ a partial function $\alpha_q : A_T \rightarrow SA_T$ which will prove to be equal to the function $q_{A_T} \circ \alpha_T$ (see Corollary 6.4 below). α_T itself arises when the construction is applied to the empty path. More precisely, when defining the α_q 's in Lemma 6.2 we show that $\alpha_{p.q} = q_{A_T} \circ \alpha_p$ for all concatenated paths $p.q$, and then put $p = \langle \rangle$.

The reader may find it helpful to read the statements of Lemma 6.2 and Definition 6.3, and then Corollary 6.4, before coming to the proof of 6.2. A comparison with the canonical coalgebra constructions of [36] and [21] may be found in Section 9 at the end.

LEMMA 6.2 For each T -component S and each path $T \xrightarrow{p} S$ there exists a partial function $\alpha_p : A_T \circ \rightarrow SA_T$ with $\text{Dom } \alpha_p = \{x \in A_T \mid (p)\downarrow \in x\}$, such that for every path $S \xrightarrow{q} U$,

$$\text{Dom } \alpha_{p,q} = \text{Dom } (q_{A_T} \circ \alpha_p) = \{x \in \text{Dom } \alpha_p \mid \alpha_p(x) \in \text{Dom } q_{A_T}\}$$

and $\alpha_{p,q} = q_{A_T} \circ \alpha_p$:

$$\begin{array}{ccc} A_T & \xrightarrow{\alpha_p} & SA_T \\ & \searrow \alpha_{p,q} & \downarrow q_{A_T} \\ & & UA_T \end{array}$$

Proof Define $\text{Dom } \alpha_p = \{x \in A_T \mid (p)\downarrow \in x\}$ for all paths $T \xrightarrow{p} S$. Then define $\alpha_p(x)$ for all $x \in \text{Dom } \alpha_p$ by induction on the formation of the component functor S , as follows.

CASE $T \xrightarrow{p} \bar{D}$: By Lemma 5.8 there is a unique $d \in D$ such that $(p)d \in x$. Define $\alpha_p(x) = d$.

CASE $T \xrightarrow{p} \text{Id}$: Define $\alpha_p(x) = x_p = \{\varphi \mid [p]\varphi \in x\}$ (see Definition 4.10). By Lemma 5.9, $x_p \in A_T$.

Now for the inductive cases. These follow a similar pattern, in which it is also shown in each case that there is some path symbol s ($= \pi_j$ or ev_d or ε_j) such that

$$\alpha_{p,s} = (s)_{A_T} \circ \alpha_p. \quad (6.i)$$

CASE $T \xrightarrow{p} S_1 \times S_2$: For $j \in \{1, 2\}$, make the induction hypothesis on S_j that for the path $T \xrightarrow{p.\pi_j} S_j$, the function $\alpha_{p.\pi_j} : A_T \circ \rightarrow S_j A_T$ has been defined. By axiom 3, $(p)\downarrow \in x$ iff $(p.\pi_j)\downarrow \in x$, so $\text{Dom } \alpha_p = \text{Dom } \alpha_{p.\pi_j}$. Define $\alpha_p(x) = \langle \alpha_{p.\pi_1}(x), \alpha_{p.\pi_2}(x) \rangle$. This makes

$$\alpha_{p.\pi_j} = (\pi_j)_{A_T} \circ \alpha_p,$$

which is equation (6.i) in this case, where $(\pi_j)_{A_T}$ is the projection $S_1 A_T \times S_2 A_T \rightarrow S_j A_T$.

CASE $T \xrightarrow{p} S^D$: By axiom 4, $\text{Dom } \alpha_p = \text{Dom } \alpha_{p.ev_d}$ for all $d \in D$, where $\alpha_{p.ev_d} : A_T \circ \rightarrow SA_T$. Define $\alpha_p(x)(d) = \alpha_{p.ev_d}(x)$ for all d , to obtain $\alpha_p(x) \in (SA_T)^D = (S^D)A_T$.

This makes $\alpha_{p.ev_d} = (ev_d)_{A_T} \circ \alpha_p$, where $(ev_d)_{A_T}$ is the evaluation function $(SA_T)^D \rightarrow SA_T$.

CASE $T \xrightarrow{p} S_1 + S_2$: By axiom 5, $\text{Dom } \alpha_p$ is the disjoint union of $\text{Dom } \alpha_{p.\varepsilon_1}$ and $\text{Dom } \alpha_{p.\varepsilon_2}$. Define $\alpha_p(x) = \iota_j \alpha_{p.\varepsilon_j}(x)$ for the unique $j \in \{1, 2\}$ with $x \in \text{Dom } \alpha_{p.\varepsilon_j}$.

This makes $\alpha_{p.\varepsilon_j} = (\varepsilon_j)_{A_T} \circ \alpha_p$, where $(\varepsilon_j)_{A_T}$ is the extraction $S_1 A_T + S_2 A_T \circ \rightarrow S_j A_T$.

This completes the definition of α_p for any $T \xrightarrow{p} S$, with equation (6.i) satisfied appropriately in all the inductive cases. Then for each $S \xrightarrow{q} U$ we prove the rest of the Lemma by induction on the formation of S again. In each case, if $q = \langle \rangle$ then $U = S$ with $p.q = p$ and $q_{A_T} = \text{id}_{SA_T}$, hence $\alpha_{p,q}$ and $q_{A_T} \circ \alpha_p$ are identical as required. In particular, if $S = \text{Id}$ or \bar{D} , then we must have $q = \langle \rangle$, so the result holds as just stated.

The inductive cases $S = S_1 \times S_2$ or S_1^D or $S_1 + S_2$ now all follow the same pattern using (6.i): if $q \neq \langle \rangle$, then $q = s.r$ for some symbol s and some paths $S \xrightarrow{s} S' \xrightarrow{r} U$. By induction hypothesis on S' , applied to the paths $T \xrightarrow{p.s} S' \xrightarrow{r} U$, we get

$$\text{Dom } \alpha_{p,q} = \text{Dom } \alpha_{p.s.r} = \text{Dom } (r_{A_T} \circ \alpha_{p.s})$$

and $\alpha_{p,q} = \alpha_{p.s.r} = r_{A_T} \circ \alpha_{p.s}$. But then since $\alpha_{p.s} = (s)_{A_T} \circ \alpha_p$ by (6.i), we have

$$\text{Dom } \alpha_{p,q} = \text{Dom } (r_{A_T} \circ (s)_{A_T} \circ \alpha_p) = \text{Dom } ((s.r)_{A_T} \circ \alpha_p),$$


and $\alpha_{p,q} = (s.r)_{A_T} \circ \alpha_p$ as required. ■

DEFINITION 6.3 Let $\alpha_T = \alpha_{\langle \rangle}$, i.e. α_p for $p = T \xrightarrow{\langle \rangle} T$. $\langle A_T, \alpha_T \rangle$ is called the *canonical T-coalgebra*.

COROLLARY 6.4

(1) $\text{Dom } \alpha_T = A_T$.

(2) For every path $T \xrightarrow{q} U$, $\text{Dom } \alpha_q = \{x \in A_T \mid \alpha_T(x) \in \text{Dom } q_{A_T}\}$ and $\alpha_q = q_{A_T} \circ \alpha_T$:



Proof

(1) Every maximal set contains axiom 2, $\langle \rangle \downarrow$, so $\text{Dom } \alpha_{\langle \rangle} = A_T$.

(2) Let $p = \langle \rangle$ in Lemma 6.2. ■

We can now prove the fundamental result showing that each maximal set x is a truth set, viz. $x = \{\varphi \mid \alpha_T, x \models \varphi\}$:

LEMMA 6.5 (TRUTH LEMMA) $\alpha_T, x \models \varphi$ iff $\varphi \in x$.

Proof By induction on the structure of φ , we show that for all $x \in A_T$, $\alpha_T, x \models \varphi$ iff $\varphi \in x$:

CASE $\varphi = \perp$: $\alpha_T, x \not\models \perp$ and x is \perp -free.

CASE $\varphi = (p) \downarrow$: $\alpha_T, x \models (p) \downarrow$
iff $\alpha_T(x) \in \text{Dom } p_{A_T}$ by the semantics of $(p) \downarrow$
iff $x \in \text{Dom } \alpha_p$ by Corollary 6.4
iff $(p) \downarrow \in x$ by the definition of α_p .

CASE $\varphi = (p)d$: (\Rightarrow) If $\alpha_T, x \models (p)d$ then $\alpha_T, x \models (p) \downarrow$ and $p_{A_T}(\alpha_T(x)) = d$. Therefore $(p) \downarrow \in x$, by the previous CASE, and $\alpha_p(x) = d$, by Corollary 6.4. Hence, by Lemma 5.8 and the definition of α_p (Lemma 6.2), $(p)d \in x$.

(\Leftarrow) If $(p)d \in x$ then $(p) \downarrow \in x$ by axiom 7. By its definition (Lemma 6.2), $\alpha_p(x) = d$ and thus $p_{A_T}(\alpha_T(x)) = d$, by Corollary 6.4. Since $\alpha_T, x \models (p) \downarrow$ by the previous CASE it follows that $\alpha_T, x \models (p)d$.

Now assume the lemma holds for θ and ψ .

CASE $\varphi = \theta \rightarrow \psi$:

$\alpha_T, x \models \theta \rightarrow \psi$
iff $\alpha_T, x \models \theta$ implies $\alpha_T, x \models \psi$ by the semantics of \rightarrow
iff $\theta \in x$ implies $\psi \in x$ by the induction hypothesis
iff $\theta \rightarrow \psi \in x$ by Lemma 4.5(1).

CASE $\varphi = [p]\psi$: (\Rightarrow) Suppose $\alpha_T, x \models [p]\psi$. If $\alpha_T, x \not\models (p) \downarrow$, then $(p) \uparrow \in x$, and so using axiom 10, $[p]\perp \in x$. Also $\perp \rightarrow \psi$ is a tautology, so $[p](\perp \rightarrow \psi) \in Ax_T^* \subseteq x$. But by axiom 11

$$[p](\perp \rightarrow \psi) \rightarrow ([p]\perp \rightarrow [p]\psi) \in x,$$

and so by Detachment it follows that $[p]\psi \in x$.

Otherwise, $\alpha_T, x \models (p) \downarrow$ and $\alpha_T, p_{A_T}(\alpha_T(x)) \models \psi$. Therefore $(p) \downarrow \in x$ and $\alpha_T, x_p \models \psi$, as $p_{A_T}(\alpha_T(x)) = \alpha_p(x) = x_p$ by Lemma 6.2 and Corollary 6.4. Hence $\psi \in x_p$ by the induction hypothesis, and so $[p]\psi \in x$.

(\Leftarrow) Suppose $[p]\psi \in x$. By negation completeness either $(p) \uparrow \in x$ or $(p) \downarrow \in x$. For the former it immediately follows that $\alpha_T, x \not\models (p) \downarrow$ and therefore $\alpha_T, x \models [p]\psi$. For the latter, $\alpha_T, x \models (p) \downarrow$ and $\psi \in x_p \in A_T$ so, by the induction hypothesis, $\alpha_T, x_p \models \psi$. But $x_p = p_{A_T}(\alpha_T(x))$; hence $\alpha_T, x \models [p]\psi$. ■

Now we are able to show that the canonical T -coalgebra is a final object in the category of T -coalgebras. The reason is natural and conceptually appealing: states can be represented by their truth sets. Any state b in a coalgebra $\langle B, \beta \rangle$ has a truth set in this coalgebra that is maximal and therefore a member of A_T . This defines a map $B \rightarrow A_T$ which proves to be the unique morphism between the coalgebras. The proof of that uses the characterisation of morphisms from Theorem 2.9(2), which is specific to *polynomial* functors.

THEOREM 6.6 $\langle A_T, \alpha_T \rangle$ is a final T -coalgebra.

Proof Let $\langle B, \beta \rangle$ be any T -coalgebra. Define

$$!_\beta : B \rightarrow A_T : b \mapsto \{\varphi \mid \beta, b \models \varphi\}.$$

Now we show that $!_\beta$ is the unique morphism from β to α_T .

- $!_\beta(b) \in A_T$: $!_\beta(b)$ is a truth set, so is in A_T by Lemma 5.6.
- $!_\beta$ is a morphism: by Theorem 2.9(2) we need to show $p_B(\beta(b)) = p_{A_T}(\alpha_T(!_\beta(b)))$ for observation paths and $!_\beta(p_B(\beta(b))) = p_{A_T}(\alpha_T(!_\beta(b)))$ for state paths.

Firstly,

$$\begin{aligned} & \beta(b) \in \text{Dom } p_B \\ \text{iff } & \beta, b \models (p)\downarrow && \text{by the semantics of } (p)\downarrow \\ \text{iff } & (p)\downarrow \in !_\beta(b) && \text{by the definition of } !_\beta \\ \text{iff } & \alpha_T, !_\beta(b) \models (p)\downarrow && \text{by the Truth Lemma (6.5)} \\ \text{iff } & \alpha_T(!_\beta(b)) \in \text{Dom } p_{A_T} && \text{by the semantics of } (p)\downarrow. \end{aligned}$$

For observation paths, when $\beta, b \models (p)\downarrow$:

$$\begin{aligned} & p_B(\beta(b)) = c \\ \text{iff } & \beta, b \models (p)c && \text{by the semantics of } (p)c \\ \text{iff } & (p)c \in !_\beta(b) && \text{by the definition of } !_\beta \\ \text{iff } & \alpha_T, !_\beta(b) \models (p)c && \text{by the Truth Lemma (6.5)} \\ \text{iff } & p_{A_T}(\alpha_T(!_\beta(b))) = c && \text{by the semantics of } (p)c, \end{aligned}$$

so $p_B(\beta(b)) = p_{A_T}(\alpha_T(!_\beta(b)))$.

For state paths, when $\beta, b \models (p)\downarrow$:

$$\begin{aligned} & \varphi \in !_\beta(p_B(\beta(b))) \\ \text{iff } & \beta, p_B(\beta(b)) \models \varphi && \text{by the definition of } !_\beta \\ \text{iff } & \beta, b \models [p]\varphi && \text{by the semantics of } [p]\varphi \\ \text{iff } & [p]\varphi \in !_\beta(b) && \text{by the definition of } !_\beta \\ \text{iff } & \varphi \in \alpha_p(!_\beta(b)) && \text{by the definition of } \alpha_p \\ \text{iff } & \varphi \in p_{A_T}(\alpha_T(!_\beta(b))) && \text{by Corollary 6.4,} \end{aligned}$$

so $!_\beta(p_B(\beta(b))) = p_{A_T}(\alpha_T(!_\beta(b)))$.

- uniqueness of $!_\beta$: Let $f : B \rightarrow A_T$ be a morphism from β to α_T . Now we need to show $\varphi \in f(b)$ iff $\varphi \in !_\beta(b)$:

$$\begin{aligned} & \varphi \in !_\beta(b) \\ \text{iff } & \beta, b \models \varphi && \text{by the definition of } !_\beta \\ \text{iff } & \alpha_T, f(b) \models \varphi && \text{by Lemma 3.3} \\ \text{iff } & \varphi \in f(b) && \text{by the Truth Lemma (6.5).} \quad \blacksquare \end{aligned}$$

7 Completeness

“Completeness” of the deducibility relation \vdash_T with respect to our semantics would state that $\Gamma \vDash_T \varphi$ implies $\Gamma \vdash_T \varphi$. An equivalent formulation of this assertion is that every \vdash_T -consistent set is satisfiable at some coalgebraic state (using Lemma 5.4(6)).

DEFINITION 7.1 The functor T is called *Lindenbaum* if every \vdash_T -consistent set of formulas is a subset of some maximal set.

THEOREM 7.2 (COMPLETENESS) *If T is Lindenbaum, then following are equivalent:*

- (1) $\Gamma \vdash_T \varphi$
- (2) $\Gamma \vDash_T \varphi$
- (3) $\Gamma \vDash^{\alpha_T} \varphi$.

Proof

(1) \Rightarrow (2): Theorem 5.3.

(2) \Rightarrow (3): Follows as α_T is a T -coalgebra.

(3) \Rightarrow (1): If $\Gamma \not\vdash_T \varphi$ then $\Gamma \cup \{\neg\varphi\}$ is \vdash_T -consistent, by Lemma 5.4(6). Therefore, by the Lindenbaum property, there exists a maximal $x \supseteq \Gamma \cup \{\neg\varphi\}$. By the Truth Lemma (6.5), $\alpha_T, x \vDash \Gamma$ and $\alpha_T, x \not\vDash \varphi$, hence $\Gamma \not\vDash^{\alpha_T} \varphi$. ■

The Lindenbaum property is in fact *necessary*, as well as sufficient, for completeness to hold. To see this, suppose $\Gamma \vDash_T \varphi$ implies $\Gamma \vdash_T \varphi$. Then if Γ is \vdash_T -consistent, we have $\Gamma \not\vdash_T \perp$ and hence $\Gamma \not\vDash_T \perp$, so $\alpha, x \vDash \Gamma$ for some α, x . Then the truth set $\{\psi \mid \alpha, x \vDash \psi\}$ extends Γ and is maximal by Theorem 5.6.

Infinitary proof relations often satisfy Lindenbaum and completeness properties only under some cardinality condition. For instance the predicate logic $\mathcal{L}_{\omega_1\omega}$, whose formulas admit denumerably long conjunctions and disjunctions but only finite strings of quantifiers, has a standard proof relation that satisfies $\Gamma \vdash \varphi$ iff $\Gamma \vDash \varphi$ for countable Γ , but not in general. Indeed [40] cites an example of a negation complete and consistent set of $\mathcal{L}_{\omega_1\omega}$ -formulas that is unsatisfiable, and gives another that is consistent but has no negation complete and consistent extensions.

Another well-known example, in the case of languages with finite-length formulas, is ω -logic. This concerns first-order languages that include constants for all members of the set ω of natural numbers, and whose proof theory has the ω -rule

$$\text{from } \{\varphi(n) \mid n \in \omega\} \text{ infer } \forall x\varphi(x).$$

The completeness theorem for ω -logic in [7] works for countable languages only. Similarly, for our coalgebraic logic we can derive the Lindenbaum property under a cardinality constraint:

THEOREM 7.3 *If \mathcal{R}_T^- is countable, then every \vdash_T -consistent subset of Φ_T can be extended to a maximal set.*

Proof Fix an enumeration $\langle \Sigma_0, \varphi_0 \rangle, \langle \Sigma_1, \varphi_1 \rangle, \dots, \langle \Sigma_n, \varphi_n \rangle, \dots$ of the countable set \mathcal{R}_T^- . Since $\mathcal{R}_T^- \subseteq \mathcal{R}_T$, we have $\Sigma_n \vdash_T \varphi_n$ for all n .

Suppose Γ is \vdash_T -consistent. Let $\Delta_0 = \Gamma$. Now assume inductively that Δ_n is defined and \vdash_T -consistent. If $\Delta_n \vdash_T \varphi_n$, then let

$$\Delta_{n+1} = \Delta_n \cup \{\varphi_n\},$$

which is \vdash_T -consistent because Δ_n is, by Lemma 5.4(5). Alternatively, $\Delta_n \not\vdash_T \varphi_n$. But $\Sigma_n \vdash_T \varphi_n$, so in that case by the cut rule CT there must exist a $\psi \in \Sigma_n$ such that $\Delta_n \not\vdash_T \psi$. Let

$$\Delta_{n+1} = \Delta_n \cup \{\neg\psi\},$$

which is \vdash_T -consistent by Lemma 5.4(6).

Next, let $\Delta = \bigcup_{n \geq 0} \Delta_n$. By construction, for all n , if $\varphi_n \notin \Delta$ then $\neg\psi \in \Delta$ for some $\psi \in \Sigma_n$.

Δ is finitely \vdash_T -consistent — any finite subset of Δ is a subset of some Δ_n , which is \vdash_T -consistent. More generally, the union of any chain of finitely \vdash_T -consistent sets is finitely \vdash_T -consistent, so Zorn's Lemma applies to the \subseteq -ordered set

$$\{\Delta' \subseteq \Phi_T \mid \Delta \subseteq \Delta' \text{ and } \Delta' \text{ is finitely } \vdash_T\text{-consistent}\}$$

to provide an extension Δ' of Δ that is maximally finitely \vdash_T -consistent. (Alternatively, by a standard argument, we could enumerate Φ_T and use Lemma 5.7(1) to proceed inductively along this (possibly transfinite) enumeration to build Δ' as a negation complete finitely \vdash_T -consistent extension of Δ .)

Now for each n , if $\varphi_n \notin \Delta'$ then $\varphi_n \notin \Delta$, so by construction there exists $\psi \in \Sigma_n$ with $\neg\psi \in \Delta \subseteq \Delta'$, hence $\psi \notin \Delta'$ or else $\{\neg\psi, \psi\}$ would contradict Δ' being finitely \vdash_T -consistent. This shows that Δ' is closed under the rule $\langle \Sigma_n, \varphi_n \rangle$ for all n , so is \mathcal{R}_T^- -closed. By Corollary 5.12 it follows that the extension Δ' of Γ is maximal. ■

The status of the countability of \mathcal{R}_T^- is clarified by the following results.

THEOREM 7.4

- (1) \mathcal{R}_T^- is countable if, and only if, either
 - (i) T has no observation paths; or
 - (ii) T has countably many paths.
- (2) T has no observation paths precisely when it has no constant components.
- (3) T has countably many paths precisely when every exponential T -component S^E has a countable exponent set E .

Proof

- (1) Suppose \mathcal{R}_T^- is countable. If T does have an observation path then there is a rule of the form \mathcal{I}_p , so $\mathcal{R}_T^- \neq \emptyset$. But for any rule $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T^-$ we have $\langle [q]\Sigma, [q]\varphi \rangle \in \mathcal{R}_T^-$ for all state paths q , so there can only be countably many state paths. Also, there are only countably many rules of the form \mathcal{I}_p , so there can only be countably many observation paths. Since any path can be extended to a state or observation path, there are then only countably many paths altogether.

Conversely, if (i) holds then there are no \mathcal{I}_p rules, so \mathcal{R}_T^- is the countable set \emptyset . On the other hand if (ii) holds put $\mathcal{R}_0 = \{\mathcal{I}_p \mid p \text{ is an observation path}\}$, and inductively

$$\mathcal{R}_{n+1} = \{\langle [q]\Sigma, [q]\varphi \rangle : \langle \Sigma, \varphi \rangle \in \mathcal{R}_n \text{ and } q \text{ is a state path}\}.$$

Then, inductively, each \mathcal{R}_n is countable, hence so is $\bigcup_{n \geq 0} \mathcal{R}_n = \mathcal{R}_T^-$.

- (2) T has paths to all its components, so there is an observation path iff there is a constant component.
- (3) A path is a finite list of symbols of the form $\pi_1, \pi_2, \varepsilon_1, \varepsilon_2, ev_d$. Since T has finitely many components, if all exponents occurring in T are countable, then there are countably many symbols of the form ev_d in the list-alphabet, hence countably many finite lists.

On the other hand, if T has a component S^E with E uncountable, then since there exists a path $T \xrightarrow{p} S^E$, the paths include the uncountable collection $\{p.ev_d \mid d \in E\}$. ■

LEMMA 7.5 *If T has no constant components with **infinite** constant set, then \vdash_T is finitary.*

Proof If every constant component \bar{D} of T is determined by a *finite* set D , then every rule $\langle \Sigma, \varphi \rangle$ of the form \mathcal{I}_p has a finite premiss-set Σ (this holds vacuously when T has no constant components, since then $\mathcal{R}_T^- = \emptyset$). But then every rule in \mathcal{R}_T^- has a finite premiss set. By standard arguments this implies that the deducibility relation \vdash_T is finitary, i.e. if $\Gamma \vdash_T \varphi$ then $\Gamma' \vdash_T \varphi$ for some finite $\Gamma' \subseteq \Gamma$. Indeed $\Gamma \vdash_T \varphi$ iff there exists a finite sequence $\varphi_0, \dots, \varphi_n = \varphi$ such that each φ_i is an axiom, or a member of Γ , or can be inferred from previous members of the sequence either by Detachment or by some rule from \mathcal{R}_T^- (see the proof of Lemma 8.1 below for a similar result). ■

COROLLARY 7.6 *If \vdash_T is infinitary, then \mathcal{R}_T^- is countable if, and only if, T has countably many paths.*

Proof If T has no observation paths, then it has no constant components at all, hence by Lemma 7.5, \vdash_T is finitary. Contrapositively, if \vdash_T is infinitary then it must have observation paths, so the Corollary now follows from Theorem 7.4(1). ■

In the light of these facts, a range of possibilities can be observed:

- If \vdash_T is finitary, then T is Lindenbaum and so \vdash_T is complete. For if \vdash_T is finitary, then a set of formulas is \vdash_T -consistent iff it is finitely \vdash_T -consistent, and hence the union of any chain of \vdash_T -consistent sets is \vdash_T -consistent. This means that the inductive argument of the proof of Theorem 7.3 still works for transfinite inductions, and so can work with an enumeration of \mathcal{R}_T^- of any length. Thus any \vdash_T -consistent set can be shown to have a \vdash_T -maximal extension even when \mathcal{R}_T^- is uncountable.
- It is possible to have $\mathcal{R}_T^- = \emptyset$, hence T is Lindenbaum, while Φ_T is *uncountable*. If T is constructed from the identity functor Id by any of the polynomial operations, then T has no constant components \overline{D} , hence no observation paths and no rules of the form \mathcal{I}_p , so $\mathcal{R}_T^- = \emptyset$. For example, if $T = \text{Id}^{\mathbb{R}}$, where \mathbb{R} is the set of real numbers, then $\mathcal{R}_T^- = \emptyset$ while Φ_T includes the uncountably many formulas $(ev_r)\downarrow$ for all $r \in \mathbb{R}$. Note that if $\mathcal{R}_T^- = \emptyset$, then also $\mathcal{R}_T = \emptyset$. In this case in fact every theory is an \mathcal{R}_T -theory, which implies that \vdash_T is finitary.
- It is possible to have \mathcal{R}_T^- countable (and even finite), hence T is Lindenbaum, while \mathcal{R}_T is uncountable. For instance, let $T = \overline{\mathbb{R}}$ and put $p = \langle \rangle$ (the only path), with $\mathcal{I}_p = \langle \Sigma, \varphi \rangle$. Then $\mathcal{R}_T^- = \{\mathcal{I}_p\}$, since there are no state paths, while \mathcal{R}_T has the uncountably many rules $\langle (p)r \rightarrow \Sigma, (p)r \rightarrow \varphi \rangle$ for all $r \in \mathbb{R}$.
- T may still be Lindenbaum when \mathcal{R}_T^- is uncountable. This holds if $T = \overline{D}^{\mathbb{R}}$ with D any finite set (even a one-element set). Since the only constant component of T has a finite constant set, \vdash_T is finitary by Lemma 7.5, so T is Lindenbaum as above. But Theorem 7.4 implies that \mathcal{R}_T^- is uncountable, since T has a constant component and an uncountable exponent \mathbb{R} .

8 Incompleteness

We now present an example to show that the Lindenbaum property and completeness of \vdash_T can fail when \mathcal{R}_T^- is uncountable. In a sense to be explained, this is the simplest possible example of incompleteness.

We have just seen that if T has no infinite constant component, then \vdash_T is finitary so T is Lindenbaum regardless of the size of \mathcal{R}_T^- . But even if T has infinite constant sets, then it will still be Lindenbaum provided that any exponent E of a component S^E is countable, for then \mathcal{R}_T^- will be countable by parts (1) and (3) of Theorem 7.4. Thus any potential counter-example to completeness will have to contain at least one infinite constant component, and at least one *uncountable* exponent. The simplest such case is to take T to be the exponential functor $\overline{\omega}^{\mathbb{R}}$ (any uncountable set would do for the exponent here). In that case T has the uncountably many observation paths $T \xrightarrow{ev_r} \overline{\omega}$ for $r \in \mathbb{R}$, and these are all the non-empty paths there are. The only non-trivial path functions are the (total) evaluation functions $ev_r : \omega^{\mathbb{R}} \rightarrow \omega$.

Now for each set $X \subseteq \mathbb{R}$, define a set Γ_X of wffs by putting

$$\Gamma_X = \{(ev_r)n \rightarrow \neg(ev_s)n \mid r, s \in X, r \neq s, \text{ and } n \in \omega\}.$$

Then $\Gamma_{\mathbb{R}}$ is itself unsatisfiable. To see this, suppose on the contrary that $\alpha, x \models \Gamma_{\mathbb{R}}$ for some α and x . Define $f : \mathbb{R} \rightarrow \omega$ by putting $f(r) = \alpha(x)(r) = ev_r(\alpha(x))$. Then $f(r) = n$ iff $\alpha, x \models (ev_r)n$, and so from the truth of all members of $\Gamma_{\mathbb{R}}$ at x we get $f(r) \neq f(s)$ for all $r \neq s \in \mathbb{R}$. But this impossible, as there is no injective function from \mathbb{R} to ω . Hence $\Gamma_{\mathbb{R}} \not\models \perp$.

On the other hand, if X is *countable*, then Γ_X is satisfiable: take any injective function $f : X \rightarrow \omega$, put $A = \{x\}$, and let $\alpha(x)$ be any function belonging to $\omega^{\mathbb{R}}$ that agrees with f on X . Then $\alpha, x \models (ev_r)n \rightarrow \neg(ev_s)n$ for all $r \neq s \in X$. Hence $\Gamma_X \not\models \perp$ for all countable X .

It turns out that $\Gamma_{\mathbb{R}}$ is \vdash_T -consistent. To prove this we observe that there is an infinite analogue of the principle used in the proof of Lemma 7.5 that finite premiss sets lead to a finitary deducibility relation. If all premiss sets of the rules from \mathcal{R}_T have fewer than κ members, where κ is a regular cardinal, then any instance

of the relation $\Gamma \vdash_T \varphi$ is witnessed by a proof-sequence of length less than κ , so $\Gamma \vdash_T \varphi$ implies $\Gamma' \vdash_T \varphi$ for some subset Γ' of Γ with fewer than κ members (see [2, 1.3]). In particular, we will sketch below a proof of

LEMMA 8.1 *For $T = \overline{\omega}^{\mathbb{R}}$, if $\Gamma \vdash_T \varphi$ then $\Gamma' \vdash_T \varphi$ for some countable $\Gamma' \subseteq \Gamma$. ■*

This Lemma shows that if $\Gamma_{\mathbb{R}}$ were \vdash_T -inconsistent, then $\Gamma' \vdash_T \perp$ for some countable $\Gamma' \subseteq \Gamma_{\mathbb{R}}$. Then $\Gamma' \subseteq \Gamma_X$ for some countable $X \subseteq \mathbb{R}$, so $\Gamma_X \vdash_T \perp$. But then $\Gamma_X \Vdash_T \perp$ by Soundness, contradicting the satisfiability of Γ_X as above.

Thus we see that $\Gamma_{\mathbb{R}} \Vdash_T \perp$ but $\Gamma_{\mathbb{R}} \not\vdash_T \perp$, so completeness fails. This also shows directly why the Lindenbaum property fails: $\Gamma_{\mathbb{R}}$ is \vdash_T -consistent but has no maximal extension, or else it would be satisfied at such an extension in the canonical coalgebra $\langle A_T, \alpha_T \rangle$, by the Truth Lemma.

Proof of Lemma 8.1:

Note first that the \mathcal{I}_p -rules determined by observation paths all have the form $\langle \{\neg(ev_r)n \mid n \in \omega\}, (ev_r)\uparrow \rangle$ for some $r \in \mathbb{R}$, and so have a countable premiss-set. From this it can be seen that if $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$, then Σ is countable.

Now define a Γ -sequence to be any sequence of wffs of the form $\langle \varphi_\mu \mid \mu \leq \nu \rangle$, such that ν is a countable ordinal and for all $\mu \leq \nu$:

- $\varphi_\mu \in \Gamma$; or
- φ_μ is an axiom; or
- φ_μ follows from previous members of the sequence by Detachment, i.e. there exist $\kappa, \lambda < \mu$ with $\varphi_\lambda = (\varphi_\kappa \rightarrow \varphi_\mu)$; or
- φ_μ follows from previous members of the sequence by a rule from \mathcal{R}_T , i.e. there exists $\langle \Sigma, \varphi \rangle \in \mathcal{R}_T$ such that $\Sigma \subseteq \{\varphi_\kappa \mid \kappa < \mu\}$ and $\varphi = \varphi_\mu$.

Next define a relation $\vdash_{\overline{\omega}}$ by putting $\Gamma \vdash_{\overline{\omega}} \varphi$ iff there exists a Γ -sequence as above with $\varphi_\nu = \varphi$. If Δ is any maximal set with $\Gamma \subseteq \Delta$, then the closure properties of Δ ensure that every member of every Γ -sequence belongs to Δ . From this it follows that $\Gamma \vdash_{\overline{\omega}} \varphi$ implies $\Gamma \vdash_T \varphi$. The converse is also true, and is shown by proving that $\Delta_\Gamma = \{\varphi \mid \Gamma \vdash_{\overline{\omega}} \varphi\}$ is an \mathcal{R}_T -theory with $\Gamma \subseteq \Delta_\Gamma$. The proof, whose details are left to the reader, requires the construction of certain Γ -sequences by concatenation of other such sequences. The fact that the premiss-set of each rule in \mathcal{R}_T is countable is crucial here in allowing all the required sequences to be indexed by countable ordinals. Then if $\Gamma \vdash_T \varphi$, the \mathcal{R}_T -theory Δ_Γ must contain φ , so $\Gamma \vdash_{\overline{\omega}} \varphi$.

To complete the proof of Lemma 8.1, suppose $\Gamma \vdash_T \varphi$. Then there is a Γ -sequence $\langle \varphi_\mu \mid \mu \leq \nu \rangle$ with $\varphi_\nu = \varphi$. Put $\Gamma' = \Gamma \cap \{\varphi_\mu \mid \mu \leq \nu\}$. Then $\langle \varphi_\mu \mid \mu \leq \nu \rangle$ is a Γ' -sequence showing $\Gamma' \vdash_T \varphi$, and Γ' is countable. ■

9 Comparisons and Questions

As far as we know this is the first paper to develop a systematic infinitary proof theory for finitary formulas in coalgebraic logic. We have used it to show that the canonical model method can be extended to give a natural logical construction of final coalgebras for all polynomial functors. The essential new features allowing this were the ‘‘halting formulas’’ $(p)\downarrow$ and their associated inference rules. It may be asked whether the technique can be adapted to other kinds of functor. Relevant to this is the result of [14] that on an abstract level the existence of a final T -coalgebra is equivalent to the existence of a logical system, with a relation of satisfaction of formulas by coalgebraic states, that has the Hennessy-Milner property and a set of formulas (rather than a proper class of formulas, as can happen if enough infinite conjunctions and disjunctions are permitted).

A different approach to polynomial coalgebraic logic, closer to the classical equational logic of universal algebra, was introduced in [10, 12]. In this approach the atomic formulas are equations between terms for algebraic expressions like $p(\alpha(x))$ where p is a path expression, α a symbol for transition structures, and x a state-valued variable. Boolean combinations of such equations provide a set of formulas whose semantics fulfills the Hennessy-Milner property. This is a more expressive language than that of the present paper, since it includes formulas with the same semantics as the constructs $(p)\downarrow$, $(p)c$ and $[p]\varphi$ but also provides syntax for many other polynomial operations: projections, pairings, insertions, case-analyses, evaluations, lambda

abstractions, functional applications. It would be of interest to investigate whether there is a suitable proof theory and canonical model construction for this richer language.

As mentioned in the Introduction, a polynomial coalgebra can be viewed as a very general kind of deterministic automaton. Non-deterministic transition systems can be modelled by operations involving the powerset functor \mathcal{P} , where $\mathcal{P}A$ is the set of all subsets of A . The “Kripke polynomial” functors are those constructible by the polynomial operations and \mathcal{P} . This class of functors was introduced in [36] and studied further [21]. The full use of the powerset functor prevents there being a final coalgebra, but there is still interest in canonical models and questions of completeness. The canonical coalgebra constructions in [36] and [21] differ from the one given here. Both of those papers consider only finitary deducibility relations for functors that have finite constant sets. They also take a many-sorted approach to syntax and proof theory, defining a class of formulas of sort S , and relation \vdash_S over the set of S -formulas, for each component S of the main functor T . Rößiger works with the set M_S of maximally \vdash_S -consistent sets of S -formulas, and defines certain functions $\alpha_S : M_S \rightarrow S(M_T)$ for all T -components S . In particular this gives a function $\alpha_T : M_T \rightarrow T(M_T)$ which is taken as the canonical T -coalgebra. Thus this coalgebra is built out of T -formulas, whereas our A_T 's correspond to M_{Id} and are built out of Id-formulas.

Jacobs works algebraically with indexed families of Boolean algebras with operators, each index corresponding to a component of T . From these, coalgebras are built using the representation theory of Boolean algebras. In this approach ultrafilters play the role of algebraic analogues of maximal sets of formulas. Interpreted syntactically, it could be said that the approach builds functions of the form $r_S : M_S \rightarrow S(M_{\text{Id}})$. Then the function $r_T : M_T \rightarrow T(M_{\text{Id}})$ is composed with a certain map $M_{\text{Id}} \rightarrow M_T$ to give a function $M_{\text{Id}} \rightarrow T(M_{\text{Id}})$ that serves as a canonical T -coalgebra.

In these terms, our method in Section 6 was to use the internal structure of T as given by paths $T \xrightarrow{p} S$ to build a partial function $\alpha_p : M_{\text{Id}} \circ \rightarrow S M_{\text{Id}}$, which turn out to be $p_{A_T} \circ \alpha_T$, and to derive $\alpha_T : M_{\text{Id}} \rightarrow T M_{\text{Id}}$ from the case that p is the empty path. It may be possible to take this approach with nondeterministic coalgebras, modifying our path functions to set-valued functions $p_A : TA \rightarrow \mathcal{P}SA$, or equivalently to binary relations $R_p \subseteq TA \times SA$, to obtain a structural analysis reminiscent of the Kripke relational semantics for classical modal logic.

It would also be of interest to develop an algebraic analogue of this approach. The construction would be more complex than Jacob's, in that it would not be possible to use all ultrafilters of a Boolean algebra when the formation of T involves infinite constant sets. Only those ultrafilters that have an infinitary “richness” property analogous to the \mathcal{R}_T -closure of maximal sets would be admissible. A theory of rich ultrafilters for polynomial coalgebras has been extensively developed in [10, 11, 13]).

Another related study is the work on Stone coalgebras in [26]. This involves coalgebras whose state-set carries a Stone-space topology, with the Vietoris functor on Stone spaces being used in place of \mathcal{P} . The Boolean representations of Jacobs [21] are adapted to give a construction of a final coalgebra for each “Vietoris polynomial functor”. Here it would appear that *compactness* replaces *finiteness* as the constraint on the constant sets occurring in functors. A natural line of enquiry then would be to see if our canonical coalgebra construction could be lifted to this topological setting to give an alternative construction of these final coalgebras. This may involve topological conditions on path relations, such as the *point-image-closure* requirement that sets of the form $\{y \mid x R_p y\}$ are closed.

Finally, a comment on *finitely branching* non-determinism, which can be modelled by replacing \mathcal{P} by the finitary powerset functor \mathcal{P}_ω , where $\mathcal{P}_\omega A$ is the set of all finite subsets of A . This imposes the *image-finiteness* property that point-image sets $\{y \mid x R_A^p y\}$ are finite. It is known that a final coalgebra exists for any functor built from polynomial operations and \mathcal{P}_ω . But as far as we are aware there is no known construction, either in classical modal logic or in coalgebraic logic more generally, that produces canonical models that are image-finite. Further investigation of this may be worthwhile.

References

- [1] Peter Aczel. Infinitary logic and the Barwise compactness theorem. In John Bell, Julian Cole, Graham Priest, and Alan Slomson, editors, *Proceedings of the Bertrand Russell Memorial Logic Conference (Uldum, Denmark, 1971)*, pages 234–277, Leeds, 1973.

- [2] Peter Aczel. An introduction to inductive definitions. In Jon Barwise, editor, *Handbook of Mathematical Logic*, pages 739–782. North-Holland, 1977.
- [3] Peter Aczel and Nax Mendler. A final coalgebra theorem. In D. H. Pitt et al., editors, *Category Theory and Computer Science. Proceedings 1989*, volume 389 of *Lecture Notes in Computer Science*, pages 357–365. Springer-Verlag, 1989.
- [4] Michael Barr. Terminal coalgebras in well-founded set theory. *Theoretical Computer Science*, 114:299–315, 1993. Corrigendum in [5].
- [5] Michael Barr. Additions and corrections to “Terminal Coalgebras in Well-founded Set Theory”. *Theoretical Computer Science*, 124:189–192, 1994.
- [6] Jon Barwise. *Admissible Sets and Structures*. Springer-Verlag, Berlin, Heidelberg, 1975.
- [7] C. C. Chang and H. J. Keisler. *Model Theory*. North-Holland, Amsterdam, 1973.
- [8] M. J. Cresswell. A Henkin completeness theorem for T. *Notre Dame Journal of Formal Logic*, 8:186–190, 1967.
- [9] Robert Goldblatt. A framework for infinitary modal logic. In *Mathematics of Modality*, number 43 in CSLI Lecture Notes, chapter 9, pages 213–299. CSLI Publications, Stanford, California, 1993.
- [10] Robert Goldblatt. What is the coalgebraic analogue of Birkhoff’s variety theorem? *Theoretical Computer Science*, 266:853–886, 2001.
- [11] Robert Goldblatt. Enlargements of polynomial coalgebras. In Rod Downey et al., editors, *Proceedings of the 7th and 8th Asian Logic Conferences*, pages 152–192. World Scientific, 2003.
- [12] Robert Goldblatt. Equational logic of polynomial coalgebras. In Philippe Balbiani, Nobu-Yuki Suzuki, Frank Wolter, and Michael Zakharyashev, editors, *Advances in Modal Logic, Volume 4*, pages 149–184. King’s College Publications, King’s College London, 2003. www.aiml.net.
- [13] Robert Goldblatt. Observational ultraproducts of polynomial coalgebras. *Annals of Pure and Applied Logic*, 123:235–290, 2003.
- [14] Robert Goldblatt. Final coalgebras and the Hennessy-Milner property. *Annals of Pure and Applied Logic*, 138(1–3):77–93, 2005.
- [15] Leon Henkin. The completeness of the first-order functional calculus. *Journal of Symbolic Logic*, 14:159–166, 1949.
- [16] Matthew Hennessy and Robin Milner. On observing nondeterminism and concurrency. In J. W. de Bakker and J. van Leeuwen, editors, *Automata, Languages and Programming. Proceedings 1980*, volume 85 of *Lecture Notes in Computer Science*, pages 299–309. Springer-Verlag, 1980.
- [17] Matthew Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the Association for Computing Machinery*, 32:137–161, 1985.
- [18] Bart Jacobs. Objects and classes, co-algebraically. In B. Freitag, C.B. Jones, C. Lengauer, and H.-J. Schek, editors, *Object-Oriented Programming with Parallelism and Persistence*, pages 83–103. Kluwer Acad. Publ., 1996.
- [19] Bart Jacobs. The temporal logic of coalgebras via Galois algebras. Technical Report CSR-R9906, Computing Science Institute, University of Nijmegen, April 1999.
- [20] Bart Jacobs. Towards a duality result in coalgebraic modal logic. *Electronic Notes in Theoretical Computer Science*, 33, 2000.
- [21] Bart Jacobs. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theoretical Informatics and Applications*, 35:31–59, 2001.

- [22] Bart Jacobs. Exercises in coalgebraic specification. In R. Backhouse, R. Crole, and J. Gibbons, editors, *Algebraic and Coalgebraic Methods in the Mathematics of Program Construction*, volume 2297 of *Lecture Notes in Computer Science*, pages 237–280. Springer, 2002.
- [23] Bart Jacobs and Jan Rutten. A tutorial on (co)algebras and (co)induction. *Bulletin of the European Association for Theoretical Computer Science*, 62:222–259, 1997.
- [24] Carol R. Karp. *Languages With Expressions of Infinite Length*. North-Holland, Amsterdam, 1964.
- [25] Yasuo Kawahara and Masao Mori. A small final coalgebra theorem. *Theoretical Computer Science*, 233:129–145, 2000.
- [26] Clemens Kupke, Alexander Kurz, and Yde Venema. Stone coalgebras. *Theoretical Computer Science*, 327:109–134, 2004.
- [27] Alexander Kurz. Specifying coalgebras with modal logic. *Electronic Notes in Theoretical Computer Science*, 11, 1998.
- [28] Alexander Kurz. Specifying coalgebras with modal logic. *Theoretical Computer Science*, 260:119–138, 2001.
- [29] E. J. Lemmon. *An Introduction to Modal Logic*, volume 11 of *American Philosophical Quarterly Monograph Series*. Basil Blackwell, Oxford, 1977. Written in collaboration with Dana Scott. Edited by Krister Segerberg.
- [30] E. J. Lemmon and Dana Scott. Intensional logic. Preliminary draft of initial chapters by E. J. Lemmon, Stanford University (later published as [29]), July 1966.
- [31] D. C. Makinson. On some completeness theorems in modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 12:379–384, 1966.
- [32] Lawrence S. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. Erratum in [33].
- [33] Lawrence S. Moss. Erratum to “Coalgebraic Logic”. *Annals of Pure and Applied Logic*, 99:241–259, 1999.
- [34] Horst Reichel. An approach to object semantics based on terminal co-algebras. *Mathematical Structures in Computer Science*, 5:129–152, 1995.
- [35] Martin Rößiger. From modal logic to terminal coalgebras. Technical Report MATH-AL-3-1998, Technische Universität Dresden, 1998.
- [36] Martin Rößiger. Coalgebras and modal logic. *Electronic Notes in Theoretical Computer Science*, 33, 2000.
- [37] Martin Rößiger. From modal logic to terminal coalgebras. *Theoretical Computer Science*, 260:209–228, 2001.
- [38] J. J. M. M. Rutten. A calculus of transition systems (towards universal coalgebra). In Alban Ponse, Maarten de Rijke, and Yde Venema, editors, *Modal Logic and Process Algebra*, CSLI Lecture Notes No. 53, pages 231–256. CSLI Publications, Stanford, California, 1995.
- [39] J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3–80, 2000.
- [40] Dana Scott. Logic with denumerably long formulas and finite strings of quantifiers. In J. W. Addison, L. Henkin, and A. Tarski, editors, *The Theory of Models*, pages 329–341. North-Holland, 1965.