ENLARGEMENTS OF POLYNOMIAL COALGEBRAS*

ROBERT GOLDBLATT
Centre for Logic, Language and Computation
Victoria University of Wellington, New Zealand
E-mail: `Rob.Goldblatt@vuw.ac.nz`

We continue a programme of study of the model theory of coalgebras of polynomial functors on the category of sets. Each such coalgebra α is shown to have an "enlargement" to a new coalgebra $E\alpha$ whose states are certain ultrafilters on the state-set of α.

This construction is used to give a new characterization, in terms of structural closure properties, of classes of coalgebras that are defined by "observable" formulas, these being Boolean combinations of equations between terms that take observable values.

It is shown that the E-construction can be replaced by a modification that restricts to ultrafilters whose members are definable in α. Both constructions are examined from the category-theoretic perspective, and shown to generate monads on the category of coalgebras concerned.

1. Introduction and Overview

This paper continues a series [Gol01c,a,b] of articles on the equational logic and model theory of coalgebras for certain functors $T : \mathbf{Set} \to \mathbf{Set}$ on the category of sets. A *T-coalgebra* is a pair (A, α) comprising a set A, thought of as a set of "states", and a function $\alpha : A \to TA$ called the *transition structure*. We study the case of functors T that are *polynomial*, i.e. constructed from constant-valued functors and the identity functor by forming products, coproducts, and exponential functors with constant exponent. Many data structures and systems of interest to computer science – such as lists, streams, trees, automata, and classes in object-oriented programming languages – can be modelled as coalgebras for polynomial functors [Rei95, Jac96, Rut95, Rut00]. This has motivated the development of a theory of "universal coalgebra" [Rut95, Rut00], by analogy with, and categorically dual to, the study of abstract algebras.

---

*Prepared using Paul Taylor's `diagrams` package.

The article [Gol01a] developed a type-theoretic calculus of *terms* for operations on polynomial coalgebras and explored its semantics. A special role is played by terms that take "observable" values. A polynomial coalgebra can be thought of as being constructed, using products and coproducts etc, from some fixed sets of observable elements given in advance. Computationally, the states of a coalgebra are regarded as not being directly accessible, but can only be investigated by performing certain "experiments" in the form of coalgebraic operations that yield observable values. Hence the emphasis on observable terms. It was shown in [Gol01a] that Boolean combinations of equations between observable terms form a natural language of *observable formulas* for specifying coalgebraic properties. In particular such formulas give a logical characterization of the fundamental relation of *bisimilarity*, or observational indistinguishability, between states of coalgebras: two states are bisimilar precisely when they satisfy the same observable formulas [Gol01a, Theorem 7.2].

The subsequent article [Gol01b] adapted the theory of ultrapowers to the context of polynomial coalgebras. Given a T-coalgebra α and an ultrafilter U, the standard theory of ultrapowers produces a structure α^U that is not a T-coalgebra. It was shown how to modify α^U to overcome this problem, by removing certain states. The result was the notion of the *observational ultrapower* α^+ of α with respect to U. This notion was then used to give a structural characterization of classes of coalgebras definable by observable formulas: a class K is the class of all models of a set of such formulas iff K is closed under disjoint unions, images of bisimilarity relations, and observational ultrapowers [Gol01b, Theorem 7.1].

The purpose of the present paper is to replace the α^+ construction in this last result by a "Stone space like" construction that is intrinsic to α. We define the *ultrafilter enlargement* $E\alpha$ of α as a new coalgebra whose states are certain ultrafilters on the state set of α. $E\alpha$ is a homomorphic image of any observational ultrapower α^+ that is sufficiently saturated, and we make use of this fact to transfer the analysis of α^+ to $E\alpha$. In particular we use the version of Łoś's Theorem developed for α^+ in [Gol01b] to study the conditions for truth in $E\alpha$ at a state (ultrafilter) F. If φ^α is the set of states in α at which formula φ is true, then it transpires that

$$E\alpha, F \models \varphi \quad \text{iff} \quad \varphi^\alpha \in F,$$

which may be interpreted as saying that φ is true in $E\alpha$ at state F iff it is true in α at a set of states that is "large in the sense of F" (see the Truth Lemma 3.3 below).

State-sets of the form φ^α may be called *definable* in α. The collection $\boldsymbol{Def}_\alpha$ of all such definable sets is a Boolean algebra, and an alternative construction to $E\alpha$ results by taking states to be ultrafilters of this algebra $\boldsymbol{Def}_\alpha$, rather than of the algebra of *all* subsets of α. The result is the *definable enlargement* $\varDelta\alpha$ of α. $\varDelta\alpha$ is a homomorphic image of $E\alpha$, and can be realized as the quotient of $E\alpha$ by the bisimilarity relation.

The category-theoretic nature of the $E\alpha$ and $\varDelta\alpha$ constructions are investigated: in the final section of the paper we show that each gives rise to a *monad* structure on the category of coalgebras.

A polynomial functor is *monomial* if it is constructed without the use of coproducts. A notion of ultrafilter enlargement for monomial coalgebras was developed in [Gold01c]. Its theory is much simpler than the one described here: the presence of coproducts introduces considerable complexity associated with the *partiality* of certain "path functions" expressing the dynamics of the transition structure α.

2. Essential Background

A substantial conceptual framework and notational system is developed in [Gol01a] and [Gol01b], all of which is essential to understanding the constructions and results given here. We now review this material, in order to make the present paper reasonably self-contained and accessible.

2.1. *Polynomial Functors and Coalgebras*

First, here is the notation for products, powers and coproducts of sets. For $j = 1$ and 2, $\pi_j : A_1 \times A_2 \to A_j$ is the *projection* function from the *product* set $A_1 \times A_2$ onto A_j, i.e. $\pi_j(a_1, a_2) = a_j$. The *pairing* of two functions of the form $f_1 : A \to B_1$ and $f_2 : A \to B_2$ is the function $\langle f_1, f_2 \rangle : A \to B_1 \times B_2$ given by $f(a) = (f_1(a), f_2(a))$. The *product* of two functions of the form $f_1 : A_1 \to B_1$ and $f_2 : A_2 \to B_2$ is the function $f_1 \times f_2 : A_1 \times A_2 \to B_1 \times B_2$ that maps (a_1, a_2) to $(f_1(a_1), f_2(a_2))$. Thus $\pi_j((f_1 \times f_2)(x)) = f_j(\pi_j(x))$.

The *coproduct* $A_1 + A_2$ of sets A_1, A_2 is their disjoint union, with injective *insertion* functions $\iota_j : A_j \to A_1 + A_2$ for $j = 1$ and 2. Each element of $A_1 + A_2$ is equal to $\iota_j(a)$ for a unique j and a unique $a \in A_j$. The *coproduct* of two functions of the form $f_1 : A_1 \to B_1$ and $f_2 : A_2 \to B_2$ is the function $f_1 + f_2 : A_1 + A_2 \to B_1 + B_2$ that maps $\iota_j(a)$ to $\iota_j(f_j(a))$.

The *D-th power* of set A is the set A^D of all functions from set D to A. The D-th power of a function $f : A \to B$ is the function $f^D : A^D \to B^D$ having $f^D(g) = f \circ g$ for all $g : D \to A$. The *evaluation* function

$eval : A^D \times D \to A$ has $eval(f, d) = f(d)$. For each $d \in D$ there is the *evaluation-at-d* function $ev_d : A^D \to A$ having $ev_d(f) = eval(f, d) = f(d)$.

The symbol $\circ\!\!\longrightarrow$ will be used for partial functions. Thus $f : A \circ\!\!\longrightarrow B$ means that f is a function with codomain B and domain $\mathrm{Dom}\, f \subseteq A$. Associated with each insertion $\iota_j : A_j \to A_1 + A_2$ is its partial inverse, the *extraction* function $\varepsilon_j : A_1 + A_2 \circ\!\!\longrightarrow A_j$ having $\varepsilon_j(y) = x$ iff $\iota_j(x) = y$. Thus $\mathrm{Dom}\, \varepsilon_j = \iota_j A_j$, i.e. $y \in \mathrm{Dom}\, \varepsilon_j$ iff $y = \iota_j(x)$ for some $x \in A_j$. These extraction functions play a vital role in the analysis of coalgebras built out of coproducts. Observe that the coproduct $f_1 + f_2$ of two functions has $(f_1 + f_2)(x) = \iota_j(f_j(\varepsilon_j(x)))$ for some j.

The identity function on a set A is denoted $\mathrm{id}_A : A \to A$. If A is a subset of B, then $A \hookrightarrow B$ is the inclusion function from A to B.

Polynomial functors are formed from the following constructions of endofunctors $T : \mathbf{Set} \to \mathbf{Set}$.

- For a fixed set $D \neq \emptyset$, the *constant functor* $\bar{D}$ has $\bar{D}(A) = D$ on sets A and $\bar{D}(f) = \mathrm{id}_D$ on functions f.
- The *identity functor* Id has $\mathrm{Id}A = A$ and $\mathrm{Id}f = f$.
- The product $T_1 \times T_2$ of two functors has $T_1 \times T_2(A) = T_1 A \times T_2 A$, and, for a function $f : A \to B$, has $T_1 \times T_2(f)$ being the product function
$$T_1(f) \times T_2(f) : T_1 A \times T_2 A \to T_1 B \times T_2 B.$$
- The coproduct $T_1 + T_2$ of two functors has $T_1 + T_2(A) = T_1 A + T_2 A$, and for $f : A \to B$, has $T_1 + T_2(f)$ being the coproduct function
$$T_1(f) + T_2(f) : T_1 A + T_2 A \to T_1 B + T_2 B.$$
- The D-th power functor T^D of a functor T has $T^D A = (TA)^D$, and for $f : A \to B$, has $T^D(f)$ being the function $(T(f))^D : (TA)^D \to (TB)^D$ that acts by $g \mapsto T(f) \circ g$. Thus $T^D(f)(g)(d) = T(f)(g(d))$.

A functor T is *polynomial* if it is constructed from constant functors and Id by finitely many applications of products, coproducts and powers. Any functor formed as part of the construction of T is a *component* of T. A polynomial functor that does not have Id as a component must be constant.

A T*-coalgebra* is a pair (A, α) comprising a set A and a function $A \xrightarrow{\alpha} TA$. A is the set of *states* and α is the *transition structure* of the coalgebra. Note that A is determined as the domain $\mathrm{Dom}\, \alpha$ of α, so we can identify the coalgebra with its transition structure, i.e. a T-coalgebra is any function of the form $\alpha : \mathrm{Dom}\, \alpha \to T(\mathrm{Dom}\, \alpha)$. A *morphism* from

T-coalgebra α to T-coalgebra β is a function $f : \mathrm{Dom}\,\alpha \to \mathrm{Dom}\,\beta$ between their state sets which commutes with their transition structures in the sense that $\beta \circ f = Tf \circ \alpha$, i.e. the following diagram commutes:

$$\begin{tikzcd}
\mathrm{Dom}\,\alpha \arrow[r, "f"] \arrow[d, "\alpha"'] & \mathrm{Dom}\,\beta \arrow[d, "\beta"] \\
T(\mathrm{Dom}\,\alpha) \arrow[r, "Tf"] & T(\mathrm{Dom}\,\beta)
\end{tikzcd}$$

f is an *isomorphism* if it has an inverse that is also a coalgebraic morphism (or equivalently, if it is bijective).

If $\mathrm{Dom}\,\alpha \subseteq \mathrm{Dom}\,\beta$, then α is a *subcoalgebra* of β iff the inclusion function $\mathrm{Dom}\,\alpha \hookrightarrow \mathrm{Dom}\,\beta$ is a morphism from α to β. More generally, if X is a subset of B, then there exists *at most one* transition $\beta' : X \to TX$ for which the inclusion $X \hookrightarrow B$ is a T-morphism from β' to β [Rut00, Proposition 6.1]. Thus it makes sense to talk about the set X being a subcoalgebra of β.

For any morphism $f : (A, \alpha) \to (B, \beta)$, the image $f(A)$ is a subcoalgebra of β, and if f is injective then coalgebra α is isomorphic to this image coalgebra [Rut00, Theorem 6.3].

Every set $\{\alpha_i : i \in I\}$ of T-coalgebras has a *disjoint union* $\sum_I \alpha_i$, which is a T-coalgebra whose domain is the disjoint union of the $\mathrm{Dom}\,\alpha_i$'s and whose transition structure acts as α_j on the summand $\iota_j \mathrm{Dom}\,\alpha_j$ of $\mathrm{Dom} \sum_I \alpha_i$. More precisely, this transition is given by $\iota_j(a) \mapsto T(\iota_j)(\alpha_j(a))$, with the insertion $\iota_j : \mathrm{Dom}\,\alpha_j \to \mathrm{Dom} \sum_I \alpha_i$ being an injective morphism making α_j isomorphic to a subcoalgebra of the disjoint union (see [Rut00, Section 4]).

2.2. *Paths and Bisimulations*

If (A, α) and (B, β) are T-coalgebras, then a relation $R \subseteq A \times B$ is a T-*bisimulation* from α to β if there exists a transition structure $\rho : R \to TR$ on R such that the projections from R to A and B are coalgebraic morphisms from ρ to α and β, i.e. the following diagram commutes:

$$\begin{tikzcd}
A \arrow[d, "\alpha"'] & R \arrow[l, "\pi_1"'] \arrow[r, "\pi_2"] \arrow[d, "\rho"] & B \arrow[d, "\beta"] \\
TA & TR \arrow[l, "T\pi_1"] \arrow[r, "T\pi_2"'] & TB
\end{tikzcd}$$

We may say that coalgebra β is the *image* of the bisimulation, or is the image of α under the bisimulation, if R is surjective, i.e. every member of B is in the image of R. Dually, α is the *domain* of the bisimulation if R is a total relation, i.e. $\operatorname{Dom} R = A$.

A function $f : A \to B$ is a morphism from α to β iff its graph $\{(a, f(a)) : a \in A\}$ is a bisimulation from α to β [Rutt00, Theorem 2.5]: a morphism is essentially a functional bisimulation. When $\operatorname{Dom}\alpha \subseteq \operatorname{Dom}\beta$, α is a subcoalgebra of β iff the identity relation on $\operatorname{Dom}\alpha$ is a bisimulation from α to β.

The above categorical definition of bisimulation appeared in [AM89]. It has a characterization [Her93, HJ98] in terms of "liftings" of relations $R \subseteq A \times B$ to relations $R^T \subseteq TA \times TB$. This in turn was transformed in [Gol01a] to another characterization of bisimulations that uses the idea of "paths" between functors, an idea introduced in [Jac00, Section 6]. Informally, the construction of polynomial T can be parsed into a tree of component functors. The paths we use are just the paths through this tree.

Formally, a *path* is a finite list of symbols of the kinds π_j, ε_j, ev_d. Write $p.q$ for the concatenation of lists p and q. The notation $T \overset{p}{\rightsquigarrow} S$ means that p is a path from functor T to functor S, and is defined by the following conditions

- $T \overset{\langle\rangle}{\rightsquigarrow} T$, where $\langle\rangle$ is the empty path.
- $T_1 \times T_2 \overset{\pi_j.p}{\rightsquigarrow} S$ whenever $T_j \overset{p}{\rightsquigarrow} S$, for $j = 1, 2$.
- $T_1 + T_2 \overset{\varepsilon_j.p}{\rightsquigarrow} S$ whenever $T_j \overset{p}{\rightsquigarrow} S$, for $j = 1, 2$.
- $T^D \overset{ev_d.p}{\rightsquigarrow} S$ for all $d \in D$ whenever $T \overset{p}{\rightsquigarrow} S$.

It is evident that for any path $T \rightsquigarrow S$, S is one of the components of T. Paths can be composed by concatenating lists: if $T_1 \overset{p}{\rightsquigarrow} T_2$ and $T_2 \overset{q}{\rightsquigarrow} T_3$, then $T_1 \overset{p.q}{\rightsquigarrow} T_3$.

A path $T \overset{p}{\rightsquigarrow} S$ induces a partial function $p_A : TA \circ\!\!\longrightarrow SA$ for each set A, defined by induction on the length of p as follows.

- $\langle\rangle_A : TA \circ\!\!\longrightarrow TA$ is the identity function id_{TA}, so is totally defined.
- $(\pi_j.p)_A = p_A \circ \pi_j$, the composition of $T_1A \times T_2A \xrightarrow{\pi_j} T_jA \overset{p_A}{\circ\!\!\longrightarrow} SA$. Thus $x \in \mathrm{Dom}\,(\pi_j.p)_A$ iff $\pi_j(x) \in \mathrm{Dom}\,p_A$.
- $(\varepsilon_j.p)_A = p_A \circ \varepsilon_j$, the composition of $T_1A + T_2A \overset{\varepsilon_j}{\circ\!\!\longrightarrow} T_jA \overset{p_A}{\circ\!\!\longrightarrow} SA$. Thus $x \in \mathrm{Dom}\,(\varepsilon_j.p)_A$ iff $x \in \mathrm{Dom}\,\varepsilon_j$ and $\varepsilon_j(x) \in \mathrm{Dom}\,p_A$.
- $(ev_d.p)_A = p_A \circ ev_d$, the composition of $(TA)^D \xrightarrow{ev_d} TA \overset{p_A}{\circ\!\!\longrightarrow} SA$. Thus $f \in \mathrm{Dom}\,(ev_d.p)_A$ iff $f(d) \in \mathrm{Dom}\,p_A$.

Concatenation of paths corresponds to composition of functions, in the sense that $(p.q)_A = q_A \circ p_A$.

A path $T \rightsquigarrow S$ is a *state path* if $S = \mathrm{Id}$, an *observation path* if $S = \bar{D}$ for some set D, and a *basic* path if it is either. A T-bisimulation can be characterized 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 $f(x)$ is defined iff $g(y)$ is defined, and $(f(x), g(y)) \in Q$ when they are both defined.

Theorem 2.1. [Gol01a, Theorem 5.7]
If $A \xrightarrow{\alpha} TA$ and $B \xrightarrow{\beta} TB$ are coalgebras for a polynomial functor T, then a relation $R \subseteq A \times B$ is a T-bisimulation if, and only if, xRy implies

(1) *for all state paths $T \overset{p}{\rightsquigarrow} \mathrm{Id}$, $p_A(\alpha(x))\,R\,p_B(\beta(y))$; and*
(2) *for all observation paths $T \overset{p}{\rightsquigarrow} \bar{D}$, $p_A(\alpha(x)) = p_B(\beta(y))$.* □

The inverse of a bisimulation is a bisimulation, and the union of any collection of bisimulations from α to β is a bisimulation [Rutt00, Section 5]. Hence there is a largest bisimulation from α to β, which is called *bisimilarity*. We denote this by $\sim$. States x and y are *bisimilar*, $x \sim y$, when xRy for some bisimulation R between α and β. This is intended to capture the notion that x and y are observationally indistinguishable.

2.3. *Types, Terms, and Formulas*

Fix a set $\mathbb{O}$ of symbols called *observable types*, and a collection $\{[\![o]\!] : o \in \mathbb{O}\}$ of non-empty sets indexed by $\mathbb{O}$. $[\![o]\!]$ is the *denotation* of o, and its members are called *observable elements*, or *constants* of type o. The set of *types over* $\mathbb{O}$, or $\mathbb{O}$-*types*, is the smallest set $\mathbb{T}$ such that $\mathbb{O} \subseteq \mathbb{T}$, $\mathsf{St} \in \mathbb{T}$ and

(1) if $\sigma_1, \sigma_2 \in \mathbb{T}$ then $\sigma_1 \times \sigma_2,\ \sigma_1 + \sigma_2 \in \mathbb{T}$;
(2) if $\sigma \in \mathbb{T}$ and $o \in \mathbb{O}$, then $o \Rightarrow \sigma \in \mathbb{T}$.

A *subtype* of an $\mathbb{O}$-type τ is any type that occurs in the formation of τ. St is a type symbol that will denote the state set of a given coalgebra. The symbol "o" will always be reserved for members of $\mathbb{O}$. $o \Rightarrow \sigma$ is a power type: such types will always have an observable exponent. A type is *rigid* if it does not have St as a subtype. The set of rigid types is thus the smallest set that includes $\mathbb{O}$ and satisfies (1) and (2).

Each $\mathbb{O}$-type σ determines a polynomial functor $|\sigma| : \mathbf{Set} \to \mathbf{Set}$. For $o \in \mathbb{O}$, $|o|$ is the constant functor $\bar{D}$ where $D = [\![o]\!]$; $|\mathsf{St}|$ is the identity functor Id; and inductively

$$|\sigma_1 \times \sigma_2| = |\sigma_1| \times |\sigma_2|, \quad |\sigma_1 + \sigma_2| = |\sigma_1| + |\sigma_2|, \quad |o \Rightarrow \sigma| = |\sigma|^{[\![o]\!]}.$$

For denotations of types, we write $[\![\sigma]\!]_A$ for the set $|\sigma|A$. Thus we have $[\![o]\!]_A = [\![o]\!]$, $[\![\mathsf{St}]\!]_A = A$,

$$\begin{aligned} [\![\sigma_1 \times \sigma_2]\!]_A &= [\![\sigma_1]\!]_A \times [\![\sigma_2]\!]_A \\ [\![\sigma_1 + \sigma_2]\!]_A &= [\![\sigma_1]\!]_A + [\![\sigma_2]\!]_A \\ [\![o \Rightarrow \sigma]\!]_A &= [\![\sigma]\!]_A{}^{[\![o]\!]}. \end{aligned}$$

If σ is a rigid type then $|\sigma|$ is a constant functor, so $[\![\sigma]\!]_A$ is a fixed set whose definition does not depend on A and may be written $[\![\sigma]\!]$. A τ-*coalgebra* is a coalgebra (A, α) for the functor $|\tau|$, i.e. α is a function of the form $A \to [\![\tau]\!]_A$.

To define *terms* we fix a denumerable set *Var* of *variables* and define a *context* to be a finite (possible empty) list

$$\Gamma = (v_1 : \sigma_1, \ldots, v_n : \sigma_n)$$

of assignments of $\mathbb{O}$-types σ_i to variables v_i, with the proviso that $v_1, \ldots, v_n$ are all distinct. Γ is a *rigid* context if all of the σ_i's are rigid types. Concatenation of lists Γ and Γ' with disjoint sets of variables is written Γ, Γ'. A *term-in-context* is an expression of the form

$$\Gamma \rhd M : \sigma,$$

which signifies that M is a "raw" term of type σ in context Γ. This may be abbreviated to $\Gamma \rhd M$ if the type of the term is understood. If $\sigma \in \mathbb{O}$, then the term is *observable*.

Figure 1 gives axioms that legislate certain *base terms* into existence, and rules for generating new terms from given ones. Axiom (Con) states that an observable element is a constant term of its type, while the raw term s in axiom (St) is a special parameter which will be interpreted as the "current" state in a coalgebra. The rules for products, coproducts and powers are the standard ones for introduction and transformation of terms of those types. The raw term in the consequent of rule (Case) is sometimes abbreviated to $\mathsf{case}(N, M_1, M_2)$.

Axioms

$$\text{(Var)}\quad \frac{v \in Var}{v : \sigma \rhd v : \sigma} \qquad \text{(Con)}\quad \frac{c \in [\![o]\!]}{\emptyset \rhd c : o} \qquad \text{(St)}\quad \frac{}{\emptyset \rhd \mathsf{s} : \mathsf{St}}$$

Weakening

$$\text{(Weak)}\quad \frac{\Gamma, \Gamma' \rhd M : \sigma}{\Gamma, v : \sigma', \Gamma' \rhd M : \sigma} \qquad \text{where } v \text{ does not occur in } \Gamma \text{ or } \Gamma'.$$

Product Types

$$\text{(Proj}_1\text{)}\quad \frac{\Gamma \rhd M : \sigma_1 \times \sigma_2}{\Gamma \rhd \pi_1 M : \sigma_1} \qquad \text{(Proj}_2\text{)}\quad \frac{\Gamma \rhd M : \sigma_1 \times \sigma_2}{\Gamma \rhd \pi_2 M : \sigma_2}$$

$$\text{(Pair)}\quad \frac{\Gamma \rhd M_1 : \sigma_1 \qquad \Gamma \rhd M_2 : \sigma_2}{\Gamma \rhd \langle M_1, M_2 \rangle : \sigma_1 \times \sigma_2}$$

Coproduct Types

$$\text{(In}_1\text{)}\quad \frac{\Gamma \rhd M : \sigma_1}{\Gamma \rhd \iota_1 M : \sigma_1 + \sigma_2} \qquad \text{(In}_2\text{)}\quad \frac{\Gamma \rhd M : \sigma_2}{\Gamma \rhd \iota_2 M : \sigma_1 + \sigma_2}$$

$$\text{(Case)}\quad \frac{\Gamma \rhd N : \sigma_1 + \sigma_2 \qquad \Gamma, v_1 : \sigma_1 \rhd M_1 : \sigma \qquad \Gamma, v_2 : \sigma_2 \rhd M_2 : \sigma}{\Gamma \rhd \mathsf{case}\ N\ \mathsf{of}\ [\iota_1 v_1 \mapsto M_1 \mid \iota_2 v_2 \mapsto M_2] : \sigma}$$

Power Types

$$\text{(App)}\quad \frac{\Gamma \rhd M : o \Rightarrow \sigma \qquad \Gamma \rhd N : o}{\Gamma \rhd M \cdot N : \sigma} \qquad \text{(Abs)}\quad \frac{\Gamma, v : o \rhd M : \sigma}{\Gamma \rhd (\lambda v.M) : o \Rightarrow \sigma}$$

Figure 1. Axioms and Rules for Generating Terms

Bindings of variables in raw terms occur in lambda-abstractions and case terms: the v in the consequent of rule (Abs) and the v_j's in the consequent of (Case) are bound in those terms. It is readily shown that in any term $\Gamma \rhd M$, all free variables of M appear in the list Γ. A *ground* term is one of the form $\emptyset \rhd M : \sigma$, which may be abbreviated to $M : \sigma$, or just to the raw term M. Thus a ground term has no free variables. Note that a ground term may contain the state parameter s, which behaves nonetheless like a variable in that it can denote any member of $\mathrm{Dom}\,\alpha$, as will be seen in the semantics presented in Section 2.4. There exist ground terms of every type, as may be seen by induction on type formation.

A term is defined to be *rigid* if its context is rigid. This entails that any free variable of the term is assigned a rigid type by Γ, so its type is formed without use of St. Of course all ground terms are rigid.

τ-*Terms*

For a given $\mathbb{O}$-type τ, a τ-*term* is any term that can be generated by the axioms and rules of Figure 1 together with the additional rule

$$(\tau\text{-Tr}) \quad \frac{\Gamma \rhd M : \mathsf{St}}{\Gamma \rhd \mathsf{tr}(M) : \tau} \; .$$

Note that from this rule and the axiom (St) we can derive the ground τ-term

$$\emptyset \rhd \mathsf{tr}(\mathsf{s}) : \tau.$$

The symbol tr will denote the transition structure of a τ-coalgebra $A \xrightarrow{\alpha} [\![\tau]\!]_A$. If M is interpreted as the state x of α, then $\mathsf{tr}(M)$ is interpreted as $\alpha(x)$.

τ-*Formulas*

An *equation-in-context* has the form $\Gamma \rhd M_1 \approx M_2$ where $\Gamma \rhd M_1$ and $\Gamma \rhd M_2$ are terms *of the same type.* A *formula-in-context* has the form $\Gamma \rhd \varphi$, with the expression φ being constructed from equations $M_1 \approx M_2$ by propositional connectives. Formation rules for formulas are given in Figure 2, using the connectives $\neg$ and $\wedge$. The other standard connectives $\vee$, $\rightarrow$, and $\leftrightarrow$ can be introduced as definitional abbreviations in the usual way. A formula $\emptyset \rhd \varphi$ with empty context is *ground*, and may be abbreviated to φ. A *rigid* formula is one whose context is rigid.

Equations

$$\text{(Eq)}\quad \frac{\Gamma \rhd M_1 : \sigma \qquad \Gamma \rhd M_2 : \sigma}{\Gamma \rhd M_1 \approx M_2}$$

Weakening

$$\text{(Weak)}\quad \frac{\Gamma, \Gamma' \rhd \varphi}{\Gamma, v : \sigma', \Gamma' \rhd \varphi} \qquad \text{where } v \text{ does not occur in } \Gamma \text{ or } \Gamma'.$$

Connectives

$$\text{(Neg)}\quad \frac{\Gamma \rhd \varphi}{\Gamma \rhd \neg\varphi} \qquad\qquad \text{(Con)}\quad \frac{\Gamma \rhd \varphi_1 \qquad \Gamma \rhd \varphi_2}{\Gamma \rhd \varphi_1 \wedge \varphi_2}$$

Figure 2. Formation Rules for Formulas

A τ-*formula* is one that is generated by using only τ-terms as premisses in the rule (Eq). An *observable* formula is one that uses only terms of observable type in forming its constituent equations.

2.4. *Semantics of Terms and Formulas*

A τ-coalgebra $\alpha : A \to |\tau|A$ interprets types σ and contexts $\Gamma = (v_1 : \sigma_1, \ldots, v_n : \sigma_n)$ by putting $[\![\,\sigma\,]\!]_\alpha = |\sigma|(\mathrm{Dom}\,\alpha) = [\![\,\sigma\,]\!]_A$, and

$$[\![\,\Gamma\,]\!]_\alpha = [\![\,\sigma_1\,]\!]_\alpha \times \cdots \times [\![\,\sigma_n\,]\!]_\alpha$$

(so $[\![\,\emptyset\,]\!]_\alpha$ is the empty product 1). The *denotation* of each τ-term $\Gamma \rhd M : \sigma$, relative to the coalgebra α, is a function

$$[\![\,\Gamma \rhd M : \sigma\,]\!]_\alpha : A \times [\![\,\Gamma\,]\!]_\alpha \longrightarrow [\![\,\sigma\,]\!]_\alpha,$$

defined by induction on the formation of terms. For empty contexts,

$$A \times [\![\,\emptyset\,]\!]_\alpha = A \times 1 \cong A,$$

so we replace $A \times [\![\,\emptyset\,]\!]_\alpha$ by A itself and interpret a ground term $\emptyset \rhd M : \sigma$ as a function $A \to [\![\,\sigma\,]\!]_\alpha$. The definition of denotations for terms given by the axioms of Figure 1, and the rule (τ-Tr), are as follows.

Var:

$[\![\,v : \sigma \rhd v : \sigma\,]\!]_\alpha : A \times [\![\,\sigma\,]\!]_\alpha \to [\![\,\sigma\,]\!]_\alpha$ is the right projection function.

Con:

$[\![\,\emptyset \rhd c : o\,]\!]_\alpha : A \to [\![\,o\,]\!]$ is the constant function with value c.

St:

$[\![\emptyset \rhd \mathsf{s} : \mathsf{St}]\!]_\alpha : A \to [\![\mathsf{St}]\!]_\alpha$ is the identity function $A \to A$.

τ-Tr:

$[\![\Gamma \rhd \mathsf{tr}(M) : \tau]\!]_\alpha : A \times [\![\Gamma]\!]_\alpha \to [\![\tau]\!]_\alpha$ is the composition of the functions

$$\begin{tikzcd} A \times [\![\Gamma]\!]_\alpha \arrow[r, "{[\![\Gamma \rhd M : \mathsf{St}]\!]_\alpha}"] & A \arrow[r, "\alpha"] & {[\![\tau]\!]_\alpha}. \end{tikzcd}$$

The denotations of terms generated inductively by the rules of Figure 1 are given by definitions that are standard in categorical logic (see [Gold01a, Section 4] or [Pit00]).

Substitution of Terms

The term $N[M/v]$ is the result of substituting the raw term M for free occurrences of the variable v in N. The following rule is derivable:

$$\text{(Subst)} \quad \frac{\Gamma \rhd M : \sigma \qquad \Gamma, v : \sigma \rhd N : \sigma'}{\Gamma \rhd N[M/v] : \sigma'}$$

The semantics of terms obeys the basic principle that substitution is interpreted as *composition* of denotations [Pit00, 2.2]. Because of the special role of the state set A, this takes the form

$$[\![\Gamma \rhd N[M/v]]\!]_\alpha = [\![\Gamma, v : \sigma \rhd N]\!]_\alpha \circ \langle \pi_1, \pi_2, [\![\Gamma \rhd M]\!]_\alpha \rangle,$$

so that the following diagram commutes:

$$\begin{tikzcd} A \times [\![\Gamma]\!]_\alpha \arrow[r, "{\langle \pi_1, \pi_2, [\![M]\!]_\alpha \rangle}"] \arrow[dr, "{[\![N[M/v]]\!]_\alpha}"'] & A \times [\![\Gamma]\!]_\alpha \times [\![\sigma]\!]_\alpha \arrow[d, "{[\![N]\!]_\alpha}"] \\ & {[\![\sigma']\!]_\alpha} \end{tikzcd}$$

Substitution for the State Parameter

The term $N[M/\mathsf{s}]$ is the result of substituting M for the state parameter s in N, according to the derivable rule

$$\text{(s-Subst)} \quad \frac{\Gamma \rhd M : \mathsf{St} \qquad \Gamma \rhd N : \sigma'}{\Gamma \rhd N[M/\mathsf{s}] : \sigma'}$$

with the semantics $[\![\,\Gamma \rhd N[M/\mathsf{s}]\,]\!]_\alpha = [\![\,\Gamma \rhd N\,]\!]_\alpha \circ \langle\, [\![\,\Gamma \rhd M\,]\!]_\alpha, \pi_2 \rangle$:

$$\begin{tikzcd}
A \times [\![\,\Gamma\,]\!]_\alpha \arrow[r, "{\langle\, [\![\,M\,]\!]_\alpha, \pi_2\rangle}"] \arrow[dr, "{[\![\,N[M/\mathsf{s}]\,]\!]_\alpha}"'] & A \times [\![\,\Gamma\,]\!]_\alpha \arrow[d, "{[\![\,N\,]\!]_\alpha}"] \\
 & [\![\,\sigma'\,]\!]_\alpha
\end{tikzcd}$$

For ground terms ($\Gamma = \emptyset$), this takes the simple form

$$[\![\,N[M/\mathsf{s}]\,]\!]_\alpha = [\![\,N\,]\!]_\alpha \circ [\![\,M\,]\!]_\alpha.$$

Semantics of Formulas

A τ-equation $\Gamma \rhd M_1 \approx M_2$ is said to be *valid* in coalgebra α if the α-denotations $[\![\,\Gamma \rhd M_1\,]\!]_\alpha$ and $[\![\,\Gamma \rhd M_2\,]\!]_\alpha$ of the terms $\Gamma \rhd M_j$ are identical. More generally we introduce a satisfaction relation

$$\alpha, x, \gamma \models \Gamma \rhd \varphi,$$

for τ-formulas in τ-coalgebras, which expresses that $\Gamma \rhd \varphi$ is *satisfied*, or *true*, in α at state x under the value-assigment $\gamma \in [\![\,\Gamma\,]\!]_\alpha$ to the variables of context Γ. This is defined inductively by

$$\begin{array}{lll}
\alpha, x, \gamma \models \Gamma \rhd M_1 \approx M_2 & \text{iff} & [\![\,\Gamma \rhd M_1\,]\!]_\alpha(x,\gamma) = [\![\,\Gamma \rhd M_2\,]\!]_\alpha(x,\gamma), \\
\alpha, x, \gamma \models \Gamma \rhd \neg\varphi & \text{iff} & \text{not } \ \alpha, x, \gamma \models \Gamma \rhd \varphi, \\
\alpha, x, \gamma \models \Gamma \rhd \varphi_1 \wedge \varphi_2 & \text{iff} & \alpha, x, \gamma \models \Gamma \rhd \varphi_1 \text{ and } \alpha, x, \gamma \models \Gamma \rhd \varphi_2.
\end{array}$$

$\Gamma \rhd \varphi$ is *true at* x, written $\alpha, x \models \Gamma \rhd \varphi$, if $\alpha, x, \gamma \models \Gamma \rhd \varphi$ for all $\gamma \in [\![\,\Gamma\,]\!]_\alpha$. α is a *model of* $\Gamma \rhd \varphi$, written $\alpha \models \Gamma \rhd \varphi$, if $\alpha, x, \models \Gamma \rhd \varphi$ for all states $x \in \operatorname{Dom} \alpha$. In that case we also say that $\Gamma \rhd \varphi$ is *valid in* the coalgebra α.

The following result is proven in [Gol01a, Section 5].

Theorem 2.2. *The class $\{\alpha : \alpha \models \Gamma \rhd \varphi\}$ of all models of an observable formula is closed under domains and images of bisimulations, including domains and images of morphisms as well as subcoalgebras. If $\Gamma \rhd \varphi$ is rigid and observable, then its class of models is also closed under disjoint unions.* □

By substituting a ground term M for the state parameter s in given formulas we can produce formulas $\varphi[M/\mathsf{s}]$ that express the modal assertion that φ

will be true after execution of the state transition $x \mapsto [\![M]\!]_\alpha(x)$ defined by M. This is the content of the following result.

Theorem 2.3. [Gol01a, Theorem 6.5]
If M is any ground term of type St, and φ any ground formula, then in any τ-coalgebra (A, α),

$$\alpha, [\![M]\!]_\alpha(x) \models \varphi \quad \textit{iff} \quad \alpha, x \models \varphi[M/\mathsf{s}].$$

2.5. *The Role of Observable Formulas*

Observable terms and formulas (and especially ground ones) play a role in the theory of polynomial coalgebras comparable to that played by standard terms and equations in the theory of abstract algebras. We record here some results that will be needed from [Gol01a], concerning ways in which observable terms and formulas characterize structural aspects of coalgebras.

Theorem 2.4. [Gol01a, Corollary 5.3]
Let $\Gamma \rhd \varphi$ be any rigid observable τ-formula. If R is a $|\tau|$-bisimulation from α to β and xRy, then

$$\alpha, x \models \Gamma \rhd \varphi \quad \textit{iff} \quad \beta, y \models \Gamma \rhd \varphi.$$

In particular, if $f : A \to B$ is a morphism from α to β, then for any $x \in A$,

$$\alpha, x \models \Gamma \rhd \varphi \quad \textit{iff} \quad \beta, f(x) \models \Gamma \rhd \varphi.$$

Consequently, if f is a surjective morphism,

$$\alpha \models \Gamma \rhd \varphi \quad \textit{iff} \quad \beta \models \Gamma \rhd \varphi.$$

□

Theorem 2.5. [Gol01a, Theorem 6.7]
A function $f : A \to B$ between τ-coalgebras (A, α) and (B, β) is a $|\tau|$-morphism if, and only if, for all $x \in A$,

(1) $[\![M]\!]_\alpha(x) = [\![M]\!]_\beta(f(x))$ *for all ground τ-terms M of observable type; and*
(2) $f([\![M]\!]_\alpha(x)) = [\![M]\!]_\beta(f(x))$ *for all ground τ-terms M of type St.*

□

2.6. *Defining Path Action and Bisimilarity*

The action of a path function is definable by a (ground) term, in the following sense.

Lemma 2.6. (Path Lemma) [Gol01a, Theorem 6.1]
For any path $|\tau| \overset{p}{\rightsquigarrow} |\sigma|$ *and variable* v *there exists a* tr*-free* τ*-term of the form*

$$v : \tau \rhd \bar{p} : \sigma$$

such that for any τ*-coalgebra* (A, α) *and any* $x \in A$*, if* $\alpha(x) \in \operatorname{Dom} p_A$ *then*

$$p_A(\alpha(x)) = [\![\, \bar{p}[\mathsf{tr}(\mathsf{s})/v] \,]\!]_\alpha(x).$$

$\square$

Note that by the substitution rule (Subst), $\bar{p}[\mathsf{tr}(\mathsf{s})/v]$ is a ground term of type σ, since $\mathsf{tr}(\mathsf{s})$ is a ground term of type τ. The term function $[\![\, \bar{p}[\mathsf{tr}(\mathsf{s})/v] \,]\!]_\alpha$ has domain A, and so may not be identical to $p_A \circ \alpha$ if p_A is partial. This is only an issue when the path p includes an extraction symbol ε_j (for otherwise p_A is total), but use of case allows the construction of observable terms that "discriminate" between the two summands of a coproduct $[\![\, \tau_1 \,]\!]_A + [\![\, \tau_2 \,]\!]_A$ and determine whether $p_A(\alpha(x))$ is defined [Gol01a, Section 6]. For this to work it is necessary to assume that there is available at least one observable type μ that is *non-trivial* in the sense that $[\![\, \mu \,]\!]$ has at least two distinct members. This is a plausible assumption in dealing with notions that are to be discriminated by observable behaviour. Define a relation $\equiv_{\alpha\beta}$ between the state sets of two τ-coalgebras by putting

$$x \equiv_{\alpha\beta} y \;\text{ iff every } \textit{ground observable} \text{ term } M \text{ has } [\![\, M \,]\!]_\alpha(x) = [\![\, M \,]\!]_\beta(y).$$

If τ has at least one non-trivial observable subtype, then $\equiv_{\alpha\beta}$ is a bisimulation from α to β [Gol01a, Lemma 7.1]. Moreover it proves to be the largest such bisimulation, giving a logical definition of bisimilarity. The precise situation is as follows.

Theorem 2.7. [Gol01a, Theorem 7.2]
Let (A, α) *and* (B, β) *be* τ*-coalgebras, where* τ *has at least one non-trivial observable subtype. Then for any* $x \in A$ *and* $y \in B$*, the following are equivalent:*

(1) x *and* y *are bisimilar:* $x \sim y$.
(2) $\alpha, x \models \Gamma \rhd \varphi$ *iff* $\beta, y \models \Gamma \rhd \varphi$ *for all rigid observable formulas* $\Gamma \rhd \varphi$.

(3) $\alpha, x \models \varphi$ *iff* $\beta, y \models \varphi$ *for all ground observable formulas* φ.
(4) $\alpha, x \models M \approx N$ *implies* $\beta, y \models M \approx N$ *for all ground observable terms* M *and* N.
(5) $[\![M]\!]_\alpha(x) = [\![M]\!]_\beta(y)$ *for all ground observable terms* M, *i.e.* $x \equiv_{\alpha\beta} y$.

$\square$

2.7. *Observational Ultrapowers*

Here we review the ultrapower construction of Section 4 of [Gol01b]. Let U be an ultrafilter on a set I. For each set A, there is an equivalence relation $=_U$ on the I-th power A^I of A defined by

$$f =_U g \text{ iff } \{i \in I : f(i) = g(i)\} \in U.$$

Each $f \in A^I$ has the equivalence class $f^U = \{g \in A^I : f =_U g\}$. The quotient set

$$A^U = \{f^U : f \in A^I\}$$

is called the *ultrapower of* A *with respect to* U. A notation that will be useful below is to write $f \in_U X$, for $X \subseteq A$, when $\{i \in I : f(i) \in X\} \in U$. We may also safely write $f^U \in_U X$ in this case, since in general $f \in_U X$ iff $g \in_U X$ whenever $f =_U g$.

There is a natural injection $e_A : A \rightarrowtail A^U$ given by $e_A(a) = \bar{a}^U$, where $\bar{a} \in A^I$ is the constant function on I with value a. The distinction between a and $\bar{a}^U$ is sometimes blurred, allowing A to be identified with the subset $e_A(A)$ of A^U.

A map $\theta : A \to B$ has a *U-lifting* to $\theta^U : A^U \to B^U$ where $\theta^U(f^U) = (\theta \circ f)^U$. This works also for a partial $\theta : A \circ\!\!\longrightarrow B$, providing a U-lifting $\theta^U : A^U \circ\!\!\longrightarrow B^U$ in the same way, with the proviso that $f^U \in \operatorname{Dom}\theta^U$ precisely when $f \in_U \operatorname{Dom}\theta$, i.e. when $\{i \in I : f(i) \in \operatorname{Dom}\theta\} \in U$. Moreover, U-lifting commutes with functional composition: given also $\eta : B \circ\!\!\longrightarrow C$ we have $(\eta \circ \theta)^U = \eta^U \circ \theta^U : A^U \circ\!\!\longrightarrow C^U$.

Now let $\alpha : A \to [\![\tau]\!]_A$ be a τ-coalgebra. The transition structure α lifts to a function $\alpha^U : A^U \to [\![\tau]\!]_A^U$, but this α^U is not a τ-coalgebra on A^U since its codomain is $[\![\tau]\!]_A^U = (|\tau|(A))^U$ rather than $[\![\tau]\!]_{A^U} = |\tau|(A^U)$. To overcome this obstruction it is necessary to remove some points from A^U. The key to understanding which ones are to be retained is provided by considering the U-lifting of the α-denotation of a ground observable term $M : o$. This is the function $[\![M]\!]_\alpha^U : A^U \to [\![o]\!]^U$. To act as a denotation for M it should assign values in $[\![o]\!]$, viewed as a subset of $[\![o]\!]^U$. In other words we should have

$$\llbracket M \rrbracket^U_\alpha(x) \in e\llbracket o \rrbracket = \{\bar{c}^U : c \in \llbracket o \rrbracket\} \subseteq \llbracket o \rrbracket^U.$$

We are thus led to define an element x of A^U to be *observable* if $\llbracket M \rrbracket^U_\alpha(x) \in e\llbracket o \rrbracket$ for every ground observable τ-term $M : o$. If $x = f^U$, this means that for each such M there exists an observable element $c_M \in \llbracket o \rrbracket$ such that $\llbracket M \rrbracket^U_\alpha(x) = \bar{c}_M{}^U$ and so

$$\{i \in I : \llbracket M \rrbracket_\alpha(f(i)) = c_M\} \in U.$$

Put $A^+ = \{x \in A^U : x \text{ is observable}\}$. For each $a \in A$ and any ground $M : o$,

$$\llbracket M \rrbracket^U_\alpha(e_A(a)) = \llbracket M \rrbracket^U_\alpha(\bar{a}^U) = (\llbracket M \rrbracket_\alpha \circ \bar{a})^U = \left(\overline{\llbracket M \rrbracket_\alpha(a)}\right)^U \in e\llbracket o \rrbracket,$$

so $e_A(a)$ is observable. Thus e_A embeds A into A^+, allowing us to view A^+ as an extension of A.

The definition of a $|\tau|$-transition structure α^+ on A^+ depends on the nature of the functor $|\sigma|$, which can be analysed in terms of paths $|\tau| \rightsquigarrow |\sigma|$. The definition of α^+ is founded on the following technical result, whose proof proceeds by induction on the length of σ.

Theorem 2.8. [Go01b, Theorem 4.1]
For any path $|\tau| \overset{p}{\rightsquigarrow} |\sigma|$ beginning at $|\tau|$ there exists a partial function $(p_A \circ \alpha)^+ : A^+ \circ\!\!\longrightarrow \llbracket \sigma \rrbracket_{A^+}$ with domain $A^+ \cap \mathrm{Dom}\,(p_A \circ \alpha)^U$ such that the diagram

$$\begin{tikzcd}
A \arrow[r, tail, "e_A"] \arrow[d, o->, "p_A \circ \alpha"'] & A^+ \arrow[d, o->, "(p_A \circ \alpha)^+"] \\
\llbracket \sigma \rrbracket_A \arrow[r, "|\sigma|e_A"] & \llbracket \sigma \rrbracket_{A^+}
\end{tikzcd}$$

commutes wherever defined: if $a \in \mathrm{Dom}\,(p_A \circ \alpha)$ then $e_A(a) \in \mathrm{Dom}\,(p_A \circ \alpha)^+$ and

$$|\sigma|e_A((p_A \circ \alpha)(a)) = (p_A \circ \alpha)^+(e_A(a)).$$

□

Now when $\sigma = \tau$ and p is the empty path, so that $p_A = \mathrm{id}_A$, Theorem 2.8 gives a function $\alpha^+ : A^+ \circ\!\!\longrightarrow \llbracket \tau \rrbracket_{A^+}$ whose domain is $A^+ \cap \mathrm{Dom}\,\alpha^U = A^+$, hence α^+ is total, such that the following diagram commutes.

$$\begin{tikzcd}
A \arrow[r, tail, "e_A"] \arrow[d, "\alpha"'] & A^+ \arrow[d, "\alpha^+"] \\
\llbracket \tau \rrbracket_A \arrow[r, "|\tau| e_A"] & \llbracket \tau \rrbracket_{A^+}
\end{tikzcd}$$

α^+ is thus a τ-coalgebra, which we call the *observational ultrapower* of α over U.

A feature of this construction is that the class of models of an observable formula is closed under observational ultrapowers. More strongly, if $\Gamma \rhd \varphi$ is observable, then

$$\alpha \models \Gamma \rhd \varphi \quad \text{if, and only if,} \quad \alpha^+ \models \Gamma \rhd \varphi$$

[Gol01b, Corollary 5.3]. This follows from a coalgebraic version of Łoś's Theorem, which for our present purposes takes the following form.

Theorem 2.9. [Gol01b, Theorem 5.2]
If $\Gamma \rhd \varphi$ *is an observable* τ*-formula, and* $f^U \in A^+$*, then*
$\alpha^+, f^U \models \Gamma \rhd \varphi$ *if, and only if,* $\{i \in I : \alpha, f(i) \models \Gamma \rhd \varphi\} \in U$. □

Ultrapowers are used in first-order model theory to build extensions of structures that are "saturated" (full of elements). We make use of a kind of saturation notion, which is expressed by saying that α^+ is *enlarging* if the following property holds:

> *any collection* $\mathcal{S}$ *of subsets of* A *with the finite intersection property has a "nonstandard element in its intersection". This element is an* $x \in A^U$ *such that for each* $X \in \mathcal{S}$, $x \in_U X$.

Enlarging observational ultrapowers can be obtained by choosing a suitable ultrafilter U [Go01b, Section 6]. A structural characterization of logically definable classes of polynomial coalgebras can now be stated:

Theorem 2.10. [Go01b, Theorem 7.1]
If τ *has at least one non-trivial observable subtype, then for any class* K *of* τ*-coalgebras, the following are equivalent.*

(1) K *is the class of all models of some set of ground observable formulas.*
(2) K *is the class of all models of some set of rigid observable formulas.*

(3) *K is closed under disjoint unions, images of bisimulations, and observational ultrapowers.*
(4) *K is closed under disjoint unions, images of bisimilarity relations, and enlarging observational ultrapowers.*

3. Ultrafilter Enlargements

One of the main purposes of the present paper is to show that observational ultrapowers can be replaced in Theorem 2.10 by a "Stone space like" construction of the *ultrafilter enlargement* $E\alpha$, an object that is intrinsically determined by the coalgebra α itself.

Assume from now that τ is a type that has at least one non-trivial observable subtype. In any τ-coalgebra $A \xrightarrow{\alpha} [\![\tau]\!]_A$, each formula $\Gamma \rhd \varphi$ defines in A the "truth set"

$$(\Gamma \rhd \varphi)^\alpha = \{x \in A : \alpha, x \models \Gamma \rhd \varphi\}$$

of all states at which the formula is true. Notice that for any morphism $f : (A, \alpha) \to (B, \beta)$, Theorem 2.4 states that if $\Gamma \rhd \varphi$ is rigid and observable, then in general $x \in (\Gamma \rhd \varphi)^\alpha$ iff $f(x) \in (\Gamma \rhd \varphi)^\beta$, and so $(\Gamma \rhd \varphi)^\alpha = f^{-1}(\Gamma \rhd \varphi)^\beta$.

An ultrafilter F on A will be called *observationally α-rich*, or more briefly just *rich* when α is understood, if it satisfies the following condition:

> *for any ground observable term $M : o$ there exists some observable element $c_M \in [\![o]\!]$ such that the truth set*
>
> $$(M \approx c_M)^\alpha = \{x \in A : [\![M]\!]_\alpha(x) = c_M\}$$
>
> *belongs to F.*

The element c_M corresponding to M in this condition is unique, for if $(M \approx c)^\alpha$ and $(M \approx d)^\alpha$ belong to F then their intersection does too, hence is non-empty. But if $x \in (M \approx c)^\alpha \cap (M \approx d)^\alpha$, then $c = [\![M]\!]_\alpha(x) = d$.

The set of α-rich ultrafilters on A will be denoted EA. Each subset X of A determines the subset X^{EA} of EA defined by

$$X^{EA} = \{F \in EA : X \in F\},$$

and the map $X \mapsto X^{EA}$ preserves the Boolean set operations $\cap$, $\cup$ and $-$.

Members of EA can be constructed from the states of any observational ultrapower $\alpha^+ : A^+ \to [\![\tau]\!]_{A^+}$ of α over an ultrafilter U on some set I. For $x \in A^+$ define

$$\Phi_U(x) = \{X \subseteq A : x \in_U X\}.$$

Note that if $x = f^U$, then for any $X \subseteq A$,

$$X \in \Phi_U(f^U) \quad \text{iff} \quad \{i \in I : f(i) \in X\} \in U.$$

It is standard theory that $\Phi_U(f^U)$ is an ultrafilter. But since f^U is an *observable* element of A^U, for each ground term $M : o$ there is some $c_M \in [\![o]\!]$ such that the set

$$\{i \in I : [\![M]\!]_\alpha(f(i)) = c_M\} = \{i : f(i) \in (M \approx c_M)^\alpha\}$$

belongs to U, which implies $(M \approx c_M)^\alpha \in \Phi_U(f^U)$. Thus $\Phi_U(f^U)$ is observationally α-rich, and we have a function $\Phi_U : A^+ \to EA$.

Lemma 3.1. *If α^+ is enlarging, then $\Phi_U : A^+ \to EA$ is surjective.*

Proof. Any $F \in EA$ has the finite intersection property, so by the enlarging property there is some $x \in A^U$ with $x \in_U X$ for all $X \in F$. Since F is rich, for each ground $M : o$ there is some $c_M \in [\![o]\!]$ with $x \in_U (M \approx c_M)^\alpha$. This shows that x is observable, i.e. $x \in A^+$. But $F \subseteq \Phi_U(x)$, so the maximality of F as a filter ensures that $F = \Phi_U(x)$. □

The definition of a τ-coalgebra structure $E\alpha$ on EA is a matter of similar complexity to the definition of α^+ and requires an induction on paths from $|\tau|$ to its component functors, similar to the proof of Theorem 2.8. We formulate the construction in a way that will enable us to transfer structure from any observational ultrapower of α to $E\alpha$ by the maps Φ_U.

Theorem 3.2. *For any path $|\tau| \overset{p}{\rightsquigarrow} |\sigma|$ beginning at $|\tau|$ there exists a partial function $E(p_A \circ \alpha) : EA \circ\!\!\longrightarrow [\![\sigma]\!]_{EA}$ with domain equal to*

$$(\mathrm{Dom}\, p_A \circ \alpha)^{EA} = \{F \in EA : \mathrm{Dom}\,(p_A \circ \alpha) \in F\}$$

such that for any observational ultrapower A^+ of A over an ultrafilter U the following diagram commutes wherever defined

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p_A \circ \alpha)^+"', "\circ" {at start}] & EA \arrow[d, "E(p_A \circ \alpha)"] \\
{[\![\sigma]\!]_{A^+}} \arrow[r, "|\sigma|\Phi_U"] & {[\![\sigma]\!]_{EA}}
\end{tikzcd}$$

i.e., if $x \in \mathrm{Dom}\,(p_A \circ \alpha)^+$, then $\Phi_U(x) \in (\mathrm{Dom}\, p_A \circ \alpha)^{EA}$ and

$$E(p_A \circ \alpha)(\Phi_U(x)) = |\sigma|\Phi_U((p_A \circ \alpha)^+(x)).$$

□

The lengthy proof of this theorem is deferred to the next section. We proceed here to explore its consequences. In particular, when $\sigma = \tau$ and p is the empty path with $p_A = \mathrm{id}_A$, we get a commuting diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "\alpha^+"'] & EA \arrow[d, "E\alpha"] \\
\llbracket \tau \rrbracket_{A^+} \arrow[r, "|\tau|\Phi_U"] & \llbracket \tau \rrbracket_{EA}
\end{tikzcd}$$

with the domain of $E\alpha$ being $(\mathrm{Dom}\,\alpha)^{EA}$. Thus $\mathrm{Dom}\,E\alpha = EA$, since $\mathrm{Dom}\,\alpha = A \in F$ for all $F \in EA$, so $E\alpha$ is a total function. This gives the definition of $E\alpha$ as a τ-coalgebra and the diagram shows that Φ_U is a morphism from α^+ to $E\alpha$. $E\alpha$ is called the *ultrafilter enlargement of* the coalgebra α.

If the diagram of Theorem 3.2 is composed with the square

$$\begin{tikzcd}
A \arrow[r, tail, "e_A"] \arrow[d, "p_A \circ \alpha"', {Circle[open]}->] & A^+ \arrow[d, "(p_A \circ \alpha)^+", {Circle[open]}->] \\
\llbracket \sigma \rrbracket_A \arrow[r, "|\sigma| e_A"] & \llbracket \sigma \rrbracket_{A^+}
\end{tikzcd}$$

of Theorem 2.8, the result is a commuting square

$$\begin{tikzcd}
A \arrow[r, tail, "\eta_A"] \arrow[d, "p_A \circ \alpha"', {Circle[open]}->] & EA \arrow[d, "E(p_A \circ \alpha)", {Circle[open]}->] \\
\llbracket \sigma \rrbracket_A \arrow[r, "|\sigma| \eta_A"] & \llbracket \sigma \rrbracket_{EA}
\end{tikzcd}$$

where $\eta_A = \Phi_U \circ e_A$ is the injection $a \mapsto \{X \subseteq A : a \in X\}$. In the case that p is the empty path, this becomes

$$\begin{tikzcd}
A \arrow[r, tail, "\eta_A"] \arrow[d, "\alpha"'] & EA \arrow[d, "E\alpha"] \\
\llbracket \tau \rrbracket_A \arrow[r, "|\tau|\eta_A"] & \llbracket \tau \rrbracket_{EA}
\end{tikzcd}$$

which shows that η_A is an injective *morphism* $\alpha \to E\alpha$ of τ-coalgebras that makes α isomorphic to a subcoalgebra of $E\alpha$.

One of the benefits of Theorem 3.2 is that it enables the deep analysis of Łoś's Theorem (2.9) to be transferred to give information about truth conditions in the coalgebra $E\alpha$:

Lemma 3.3. (Truth Lemma)
For any rigid observable formula $\Gamma \rhd \varphi$*, and any state* $F \in EA$*,*

$$E\alpha, F \models \Gamma \rhd \varphi \quad \textit{iff} \quad (\Gamma \rhd \varphi)^{\alpha} \in F.$$

Proof. Let α^+ be an enlarging observational ultrapower of α with the associated map $\Phi_U : \alpha^+ \to EA$ being a surjection (Lemma 3.1). Given $F \in EA$, choose $f^U \in A^+$ such that $F = \Phi_U(f^U)$.

Then by the invariance under a morphism of truth of a rigid observable formula at a state (Theorem 2.4), we have that

$$E\alpha, \Phi_U(f^U) \models \Gamma \rhd \varphi \quad \text{iff} \quad \alpha^+, f^U \models \Gamma \rhd \varphi.$$

By Łoś's Theorem 2.9, the latter condition is equivalent to

$$\{i : \alpha, f(i) \models \Gamma \rhd \varphi\} \in U,$$

i.e. to $f^U \in_U (\Gamma \rhd \varphi)^{\alpha}$, and hence is equivalent to to $(\Gamma \rhd \varphi)^{\alpha} \in \Phi_U(f^U)$ by definition of Φ_U. $\square$

We can now show that the class of all models of a rigid observable formula is closed under ultrafilter enlargements, and indeed is invariant under this construction:

Corollary 3.4. *For any rigid observable* $\Gamma \rhd \varphi$*,*

$$\alpha \models \Gamma \rhd \varphi \quad \textit{iff} \quad E\alpha \models \Gamma \rhd \varphi.$$

Proof. If $\alpha \models \Gamma \rhd \varphi$, then for each state $F \in EA$, $(\Gamma \rhd \varphi)^\alpha = A \in F$, so $E\alpha, F \models \Gamma \rhd \varphi$ by Lemma 3.3. This shows that $E\alpha \models \Gamma \rhd \varphi$.

Conversely, if $E\alpha \models \Gamma \rhd \varphi$, then $\alpha \models \Gamma \rhd \varphi$ follows by Theorem 2.2, as α is isomorphic to a subalgebra of $E\alpha$ by the morphism η_A.

Alternatively, a more streamlined proof is that for any *enlarging* α^+, $\alpha \models \Gamma \rhd \varphi$ iff $\alpha^+ \models \Gamma \rhd \varphi$ by a Corollary to Łoś's Theorem mentioned earlier [Gol01b, Corollary 5.3], while $\alpha^+ \models \Gamma \rhd \varphi$ iff $E\alpha \models \Gamma \rhd \varphi$ by the last part of Theorem 2.4, which yields that validity in a coalgebra is invariant under the *surjective* morphism $\Phi_U : \alpha^+ \to E\alpha$. □

Now if M is a ground term of type St, then for any $F \in EA$, the denotation value $[\![M]\!]_{E\alpha}(F)$ is a state of $E\alpha$, i.e. an observationally rich ultrafilter on A. The following is a useful characterization of members of this ultrafilter in terms of F.

Lemma 3.5. (State-Term Lemma)
Let M be ground term of type St. Then for any $F \in EA$, and any $X \subseteq A$,

$$X \in [\![M]\!]_{E\alpha}(F) \quad \textit{iff} \quad [\![M]\!]_\alpha^{-1}(X) \in F.$$

Proof. Given $F \in EA$, let $F = \Phi_U(f^U)$, where $\Phi_U : A^+ \to EA$ is the surjection given by an enlarging ultrapower of α.

Now it was shown in [Gol01b, Theorem 5.1(4)] that for ground $M : \mathsf{St}$, the denotation $[\![M]\!]_{\alpha^+}$ of M in α^+ is just the restriction to A^+ of the U-lifting $[\![M]\!]_\alpha^U$ of the denotation $[\![M]\!]_\alpha : A \to A$. Thus $[\![M]\!]_{\alpha^+}(f^U) = [\![M]\!]_\alpha^U(f^U)$. Hence as the morphism Φ_U preserves denotation values (Theorem 2.5(2)), $[\![M]\!]_{E\alpha}(\Phi_U(f^U)) = \Phi_U([\![M]\!]_{\alpha^+}(f^U)) = \Phi_U([\![M]\!]_\alpha^U(f^U))$. Then

$$\begin{array}{ll} X \in [\![M]\!]_{E\alpha}(F) & \\ \text{iff } X \in \Phi_U([\![M]\!]_\alpha^U(f^U)) & \text{from above} \\ \text{iff } [\![M]\!]_\alpha^U(f^U) \in_U X & \text{definition of } \Phi_U \\ \text{iff } ([\![M]\!]_\alpha \circ f)^U \in_U X & \text{definition of } [\![M]\!]_\alpha^U \\ \text{iff } \{i \in I : [\![M]\!]_\alpha(f(i)) \in X\} \in U & \text{definition of } \in_U \\ \text{iff } \{i \in I : f(i) \in [\![M]\!]_\alpha^{-1}(X)\} \in U & \text{definition of } [\![M]\!]_\alpha^{-1} \\ \text{iff } [\![M]\!]_\alpha^{-1}(X) \in \Phi_U(f^U) = F & \text{definition of } \Phi_U. \end{array}$$

□

Corollary 3.6. *For any $X \subseteq A$,*

$$[\![M]\!]_{E\alpha}^{-1}(X^{EA}) = ([\![M]\!]_\alpha^{-1}(X))^{EA}.$$

Proof. For any $F \in EA$, from Lemma 3.5 and the definition of X^{EA} we get that

$$[\![M]\!]_{E\alpha}(F) \in X^{EA} \quad \text{iff} \quad F \in ([\![M]\!]_{\alpha}^{-1}(X))^{EA}. \qquad \square$$

4. The Proof of Theorem 3.2

This section could be skipped over on a first reading of the paper, if the reader wishes to continue at this point with the conceptual development.

The proof of Theorem 3.2 proceeds by induction on the formation of the end-type σ of the path p. In each case we *define* $\mathrm{Dom}\, E(p_A \circ \alpha)$ to be the set

$$(\mathrm{Dom}\, p_A \circ \alpha)^{EA} = \{F \in EA : \mathrm{Dom}\,(p_A \circ \alpha) \in F\}.$$

But then if $x \in \mathrm{Dom}\,(p_A \circ \alpha)^+$, Theorem 2.8 implies that $x \in \mathrm{Dom}\,(p_A \circ \alpha)^U$, which means that $x \in_U \mathrm{Dom}\,(p_A \circ \alpha)$, and therefore $\mathrm{Dom}\,(p_A \circ \alpha) \in \Phi_U(x)$, i.e. $\Phi_U(x) \in (\mathrm{Dom}\, p_A \circ \alpha)^{EA}$ as required.

Thus the burden of the proof in each case is to define $E(p_A \circ \alpha)$ itself in such a way that $E(p_A \circ \alpha)(\Phi_U(x)) = |\sigma|\Phi_U((p_A \circ \alpha)^+(x))$:

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p_A \circ \alpha)^+"'] & EA \arrow[d, "E(p_A \circ \alpha)"] \\
{[\![\sigma]\!]_{A^+}} \arrow[r, "|\sigma|\Phi_U"] & {[\![\sigma]\!]_{EA}}
\end{tikzcd}$$

To understand the proof it is necessary to know the definition of the function $(p_A \circ \alpha)^+$. This will be stated in each case, but it might be beneficial if the reader had access to the proof of Theorem 2.8 given in [Gol01b, Theorem 4.1].

The induction begins with the base cases of observable types and the type St.

$\boxed{\text{Case } \sigma \in \mathbb{O}}$ Let $D = [\![\sigma]\!]$. Then the above diagram in this case is

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p_A \circ \alpha)^+"'] & EA \arrow[d, "E(p_A \circ \alpha)"] \\
D \arrow[r, "\mathrm{id}"] & D
\end{tikzcd}$$

Let $M_p = \bar{p}[\mathsf{tr}(\mathsf{s})/v]$ be the ground term of type σ given by the Path Lemma 2.6. For $F \in \mathrm{Dom}\, E(p_A \circ \alpha)$, let $E(p_A \circ \alpha)(F)$ be the unique $d \in D$ such that the truth set $(M_p \approx d)^\alpha$ belongs to the rich ultrafilter F.

Now if $x = f^U \in \mathrm{Dom}\,(p_A \circ \alpha)^+$, then as x is observable, there is some $c \in D$ such that $\llbracket M_p \rrbracket^U_\alpha(x) = \bar{c}^U$. *In the proof of Theorem 2.8,* $(p_A \circ \alpha)^+(x)$ *is defined to be this* c. But now

$$\{i \in I : f(i) \in (M_p \approx c)^\alpha\} = \{i \in I : \llbracket M_p \rrbracket_\alpha(f(i)) = c\} \in U,$$

so $x \in_U (M_p \approx c)^\alpha$. Hence $(M_p \approx c)^\alpha \in \Phi_U(x)$, implying that c is the value of $E(p_A \circ \alpha)$ at $\Phi_U(x)$, i.e. $(p_A \circ \alpha)^+(x) = E(p_A \circ \alpha)(\Phi_U(x))$, as required for the diagram to commute.

$\boxed{\text{Case } \sigma = \mathsf{St}}$ Here $|\sigma|$ is the identity functor Id, so the diagram becomes

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, o->, "(p_A \circ \alpha)^+"'] & EA \arrow[d, o->, "E(p_A \circ \alpha)"] \\
A^+ \arrow[r, "\Phi_U"] & EA
\end{tikzcd}$$

Define a unary operation $[p_A] : \mathcal{P}A \to \mathcal{P}A$ on the powerset $\mathcal{P}A$ of A by putting

$$\begin{aligned}
[p_A]X &= \{a \in A : a \in \mathrm{Dom}\,(p_A \circ \alpha) \text{ implies } p_A(\alpha(a)) \in X\} \\
&= -\mathrm{Dom}\,(p_A \circ \alpha) \cup \{a \in \mathrm{Dom}\,(p_A \circ \alpha) : p_A(\alpha(a)) \in X\}.
\end{aligned}$$

It is straightforward to verify that

$$[p_A](X \cap Y) = [p_A]X \cap [p_A]Y \tag{i}$$

$$[p_A](X \cup Y) = [p_A]X \cup [p_A]Y \tag{ii}$$

$$[p_A]\emptyset = -\mathrm{Dom}\,(p_A \circ \alpha). \tag{iii}$$

For any ultrafilter F on A, let F_p be the inverse image of F under $[p_A]$:

$$F_p = \{X \subseteq A : [p_A]X \in F\}.$$

Since F is a filter, it follows from (i) that $X \cap Y \in F_p$ iff $X, Y \in F_p$, which means that F_p is a filter. Then if $F \in (\mathrm{Dom}\, p_A \circ \alpha)^{EA}$, since $\mathrm{Dom}\,(p_A \circ \alpha) \in F$ we get $[p_A]\emptyset \notin F$ by (iii), so $\emptyset \notin F_p$, and therefore F_p is proper. But (ii) implies that F_p is prime ($X \cup Y \in F_p$ only if $X \in F_p$ or $Y \in F_p$), so altogether F_p is an ultrafilter on A in this case.

Moreover, if F is rich, then so is F_p. For, given a ground observable term $N : o$, consider the term $N[M_p/\mathsf{s}] : o$, where $M_p = \bar{p}[\mathsf{tr}(\mathsf{s})/v] : \mathsf{St}$ as in the Path Lemma 2.6. By richness of F there is some element $c \in \llbracket o \rrbracket$ such that the truth set $(N[M_p/\mathsf{s}] \approx c)^\alpha$ belongs to F. But

$$(N[M_p/\mathsf{s}] \approx c)^\alpha \subseteq [p_A](N \approx c)^\alpha,$$

since if $(N[M_p/\mathsf{s}] \approx c)$ is true in α at a, then if $a \in \mathrm{Dom}\,(p_A \circ \alpha)$ it follows by Theorem 2.3 that $(N \approx c)$ is true at $\llbracket M_p \rrbracket_\alpha(a) = p_A(\alpha(a))$, so $p_A(\alpha(a)) \in (N \approx c)^\alpha$, showing $a \in [p_A](N \approx c)^\alpha$. As F is closed under supersets, we then get $[p_A](N \approx c)^\alpha \in F$, hence $(N \approx c)^\alpha \in F_p$, establishing the richness of F_p.

Altogether now we have shown that F_p is a rich ultrafilter whenever $F \in (\mathrm{Dom}\, p_A \circ \alpha)^{EA}$, so we can define $E(p_A \circ \alpha)(F)$ to be F_p. Thus in general,

$$X \in E(p_A \circ \alpha)(F) \quad \text{iff} \quad [p_A]X \in F.$$

But in this case of $\sigma = \mathsf{St}$, $(p_A \circ \alpha)^+$ is defined to be the restriction of $(p_A \circ \alpha)^U$ to A^+, so if $x = f^U \in \mathrm{Dom}\,(p_A \circ \alpha)^+$, then $f \in_U \mathrm{Dom}\,(p_A \circ \alpha)$ and $(p_A \circ \alpha)^+(x) = (p_A \circ \alpha)^U(f^U) = g^U$, say. Thus the set

$$J = \{i : f(i) \in \mathrm{Dom}\,(p_A \circ \alpha) \text{ and } p_A(\alpha(f(i))) = g(i)\}$$

belongs to U. But for any $X \subseteq A$,

$$J \cap \{i : g(i) \in X\} = J \cap \{i : f(i) \in [p_A]X\},$$

so as $J \in U$ we have $\{i : g(i) \in X\} \in U$ iff $\{i : f(i) \in [p_A]X\} \in U$. This means that $X \in \Phi_U(g^U)$ iff $[p_A]X \in \Phi_U(f^U)$. But by definition, $[p_A]X \in \Phi_U(f^U)$ iff $X \in E(p_A \circ \alpha)(\Phi_U(f^U))$. This establishes that

$$E(p_A \circ \alpha)(\Phi_U(x)) = \Phi_U(g^U) = \Phi_U((p_A \circ \alpha)^+(x)),$$

making the last diagram commute as required.

$\boxed{\text{Case } \sigma = \sigma_1 \times \sigma_2}$ In this first inductive case we make the hypothesis that the statement of Theorem 3.2 holds for σ_1 and σ_2. From the path $|\tau| \overset{p}{\rightsquigarrow} |\sigma|$ we obtain, for $j = 1$ and 2, the path $p^j = |\tau| \overset{p.\pi_j}{\rightsquigarrow} |\sigma_j|$ and, by the induction hypothesis, a partial function $E(p^j_A \circ \alpha)$ such that the diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p^j_A \circ \alpha)^+"'] & EA \arrow[d, "E(p^j_A \circ \alpha)"] \\
\llbracket \sigma_j \rrbracket_{A^+} \arrow[r, "|\sigma_j|\Phi_U"] & \llbracket \sigma_j \rrbracket_{EA}
\end{tikzcd}$$

fulfils Theorem 3.2 for any ultrafilter U.

Now $p^j_A = \pi_j \circ p_A$, where π_j projects $\llbracket \sigma_1 \rrbracket_A \times \llbracket \sigma_2 \rrbracket_A$ onto $\llbracket \sigma_j \rrbracket_A$, so as π_j and α are total, $\mathrm{Dom}\,(p^j_A \circ \alpha) = \mathrm{Dom}\,(p_A \circ \alpha)$. Thus if $\mathrm{Dom}\,(p_A \circ \alpha) \in F \in EA$, then by induction hypothesis $E(p^j_A \circ \alpha)(F)$ is defined for $j = 1, 2$, so we can define

$$E(p_A \circ \alpha)(F) = \langle E(p^1_A \circ \alpha)(F), E(p^2_A \circ \alpha)(F) \rangle.$$

This yields the diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p_A \circ \alpha)^+"'] & EA \arrow[d, "E(p_A \circ \alpha)"] \\
\llbracket \sigma_1 \rrbracket_{A^+} \times \llbracket \sigma_2 \rrbracket_{A^+} \arrow[r, "|\sigma_1 \times \sigma_2|\Phi_U"] & \llbracket \sigma_1 \rrbracket_{EA} \times \llbracket \sigma_2 \rrbracket_{EA}
\end{tikzcd}$$

In this case of $\sigma = \sigma_1 \times \sigma_2$, $(p_A \circ \alpha)^+$ is defined to be the pairing function $\langle (p^1_A \circ \alpha)^+, (p^2_A \circ \alpha)^+ \rangle$, so if $x \in \mathrm{Dom}\,(p_A \circ \alpha)^+$, then for $j = 1, 2$,

$$\begin{aligned}
&\pi_j[E(p_A \circ \alpha)(\Phi_U(x))] \\
&= E(p^j_A \circ \alpha)(\Phi_U(x)) && \text{definition of } E(p_A \circ \alpha) \\
&= |\sigma_j|\Phi_U((p^j_A \circ \alpha)^+(x)) && \text{second-to-last diagram} \\
&= |\sigma_j|\Phi_U(\pi_j((p_A \circ \alpha)^+(x))) && \text{definition of } (p_A \circ \alpha)^+ \\
&= \pi_j[|\sigma_1 \times \sigma_2|\Phi_U((p_A \circ \alpha)^+(x))] && \text{definition of } |\sigma_1 \times \sigma_2|.
\end{aligned}$$

Hence

$$E(p_A \circ \alpha)(\Phi_U(x)) = |\sigma_1 \times \sigma_2|\Phi_U((p_A \circ \alpha)^+(x)),$$

making the last diagram commute as required.

$\boxed{\text{Case } \sigma = \sigma_1 + \sigma_2}$ Assume Theorem 3.2 holds for σ_1 and σ_2. This time we define p^j to be the path $|\tau| \overset{p.\varepsilon_j}{\rightsquigarrow} |\sigma_j|$ and, by the induction hypothesis, have a partial function $E(p^j_A \circ \alpha)$ such that the same diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, "(p^j_A \circ \alpha)^+"'] & EA \arrow[d, "E(p^j_A \circ \alpha)"] \\
\llbracket \sigma_j \rrbracket_{A^+} \arrow[r, "|\sigma_j|\Phi_U"] & \llbracket \sigma_j \rrbracket_{EA}
\end{tikzcd}$$

fulfils Theorem 3.2. But now $p_A^j = \varepsilon_j \circ p_A$, where ε_j is the (partial) extraction from $[\![\,\sigma_1\,]\!]_A + [\![\,\sigma_2\,]\!]_A$ to $[\![\,\sigma_j\,]\!]_A$, and Dom $(p_A \circ \alpha)$ is the disjoint union of Dom $(p_A^1 \circ \alpha)$ and Dom $(p_A^2 \circ \alpha)$. Thus if Dom $(p_A \circ \alpha) \in F \in EA$, then Dom $(p_A^j \circ \alpha) \in F$ for exactly one j, and we define

$$E(p_A \circ \alpha)(F) = \iota_j(E(p_A^j \circ \alpha)(F))$$

for this j, where ι_j is the insertion of $[\![\,\sigma_j\,]\!]_{EA}$ into $[\![\,\sigma_1\,]\!]_{EA} + [\![\,\sigma_2\,]\!]_{EA}$. This yields the diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, o->, "(p_A \circ \alpha)^+"'] & EA \arrow[d, o->, "E(p_A \circ \alpha)"] \\
{[\![\,\sigma_1\,]\!]_{A^+} + [\![\,\sigma_2\,]\!]_{A^+}} \arrow[r, "|\sigma_1 + \sigma_2|\Phi_U"] & {[\![\,\sigma_1\,]\!]_{EA} + [\![\,\sigma_2\,]\!]_{EA}}
\end{tikzcd}$$

In this coproduct case, Dom $(p_A \circ \alpha)^+$ is the disjoint union of Dom $(p_A^1 \circ \alpha)^+$ and Dom $(p_A^2 \circ \alpha)^+$, and $(p_A \circ \alpha)^+(x) = \iota_j((p_A^j \circ \alpha)^+(x))$ for the unique j such that $(p_A^j \circ \alpha)^+(x)$ is defined. Then

$$\begin{aligned}
&E(p_A \circ \alpha)(\Phi_U(x)) \\
&= \iota_j(E(p_A^j \circ \alpha)(\Phi_U(x))) && \text{definition of } E(p_A \circ \alpha) \\
&= \iota_j(|\sigma_j|\Phi_U((p_A^j \circ \alpha)^+(x))) && \text{second-to-last diagram} \\
&= |\sigma_1 + \sigma_2|\Phi_U(\iota_j(p_A^j \circ \alpha)^+(x)) && \text{definition of } |\sigma_1 + \sigma_2|. \\
&= |\sigma_1 + \sigma_2|\Phi_U((p_A \circ \alpha)^+(x)) && \text{definition of } (p_A \circ \alpha)^+,
\end{aligned}$$

making the last diagram commute as required.

$\boxed{\text{Case of } o \Rightarrow \sigma}$ Assume the Theorem holds for σ. Let $D = [\![\,o\,]\!]$. Then from the path $|\tau| \overset{p}{\rightsquigarrow} |o \Rightarrow \sigma|$ we obtain, for each $d \in D$, the path $p^d = |\tau| \overset{p.ev_d}{\rightsquigarrow} |\sigma|$ and, by hypothesis on σ, a partial function $E(p_A^d \circ \alpha)$ such that the diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, o->, "(p_A^d \circ \alpha)^+"'] & EA \arrow[d, o->, "E(p_A^d \circ \alpha)"] \\
{[\![\,\sigma\,]\!]_{A^+}} \arrow[r, "|\sigma|\Phi_U"] & {[\![\,\sigma\,]\!]_{EA}}
\end{tikzcd}$$

fulfils Theorem 3.2. Here $p_A^d = ev_d \circ p_A$, with $ev_d : [\![\sigma]\!]_A^D \to [\![\sigma]\!]_A$, so as ev_d and α are total, $\mathrm{Dom}\,(p_A^d \circ \alpha) = \mathrm{Dom}\,(p_A \circ \alpha)$. Thus if $\mathrm{Dom}\,(p_A \circ \alpha) \in F \in EA$, then by induction hypothesis on σ, $E(p_A^d \circ \alpha)(F)$ is defined for all $d \in D$, so we can define $E(p_A \circ \alpha)(F)$ as a function of type $D \to [\![\sigma]\!]_{EA}$ by putting

$$E(p_A \circ \alpha)(F)(d) = E(p_A^d \circ \alpha)(F).$$

This yields the diagram

$$\begin{tikzcd}
A^+ \arrow[r, "\Phi_U"] \arrow[d, o->, "(p_A \circ \alpha)^+"'] & EA \arrow[d, o->, "E(p_A \circ \alpha)"] \\
{[\![\sigma]\!]_{A^+}^D} \arrow[r, "|o \Rightarrow \sigma|\Phi_U"] & {[\![\sigma]\!]_{EA}^D}
\end{tikzcd}$$

In this power-type case, $\mathrm{Dom}\,(p_A \circ \alpha)^+ = \mathrm{Dom}\,(p_A^d \circ \alpha)^+$ for all $d \in D$, with

$$(p_A \circ \alpha)^+(x)(d) = (p_A^d \circ \alpha)^+(x) \in [\![\sigma]\!]_{A^+}.$$

Then for all $d \in D$,

$$\begin{array}{ll}
E(p_A \circ \alpha)(\Phi_U(x))(d) & \\
= E(p_A^d \circ \alpha)(\Phi_U(x)) & \text{definition of } E(p_A \circ \alpha) \\
= |\sigma|\Phi_U((p_A^d \circ \alpha)^+(x)) & \text{second-to-last diagram} \\
= |\sigma|\Phi_U((p_A \circ \alpha)^+(x)(d)) & \text{definition of } (p_A \circ \alpha)^+. \\
= |o \Rightarrow \sigma|\Phi_U((p_A \circ \alpha)^+(x))(d) & \text{definition of } |o \Rightarrow \sigma|.
\end{array}$$

Hence

$$E(p_A \circ \alpha)(\Phi_U(x)) = |o \Rightarrow \sigma|\Phi_U((p_A \circ \alpha)^+(x)),$$

making the last diagram commute as required.

This completes the inductive proof of Theorem 3.2. □

5. Definable Enlargements

We now consider a modification of the ultrafilter enlargement construction. This will produce a natural quotient of the coalgebra $E\alpha$ by focusing on the truth sets

$$\varphi^\alpha = \{x \in A : \alpha, x \models \varphi\}$$

of *ground observable* formulas φ. Such sets may be called *definable*, and the collection

$$\boldsymbol{Def}^{\alpha} = \{\varphi^{\alpha} : \varphi \text{ is ground and observable}\}$$

is a Boolean algebra of subsets of A. This follows because in general $\varphi_1^{\alpha} \cap \varphi_2^{\alpha} = (\varphi_1^{\alpha} \wedge \varphi_2^{\alpha})$ and $A - \varphi^{\alpha} = (\neg\varphi)^{\alpha}$, so $\boldsymbol{Def}^{\alpha}$ is a subalgebra of the powerset Boolean algebra $\mathcal{P}A$.

Now let $\varDelta A$ be the set of all observationally rich ultrafilters of the Boolean algebra $\boldsymbol{Def}^{\alpha}$. Hence a member of $\varDelta A$ is a collection of definable sets. Note that the sets $(M \approx c_M)^{\alpha}$ required for the definition of "observationally rich" are all defined by ground observable formulas, so such sets belong to $\boldsymbol{Def}^{\alpha}$.

Let $\theta_{\alpha} : EA \to \varDelta A$ be the restriction map taking each $F \in EA$ to

$$\theta_{\alpha}(F) = F \cap \boldsymbol{Def}^{\alpha} = \{\varphi^{\alpha} : \varphi^{\alpha} \in F\}.$$

It is readily checked that $\theta_{\alpha}(F)$ belongs to $\varDelta A$ when $F \in EA$. Moreover, θ_{α} is surjective: for any $H \in \varDelta A$, H has the finite intersection property so extends to an ultrafilter F on A which is rich because H is rich. Then $F \in EA$ and $H \subseteq \theta_{\alpha}(F)$, so $H = \theta_{\alpha}(F)$ as H is a maximal filter in $\boldsymbol{Def}^{\alpha}$.

Theorem 5.1. *$\theta_{\alpha}(F) = \theta_{\alpha}(G)$ if, and only if, F and G are bisimilar states in the coalgebra $E\alpha$.*

Proof. $\theta_{\alpha}(F) = \theta_{\alpha}(G)$ iff F and G contain the same sets of the form φ^{α} with φ ground and observable. By the Truth Lemma 3.3 this means precisely that $E\alpha, F \models \varphi$ iff $E\alpha, G \models \varphi$ for all such φ. But by Theorem 2.7(3), this holds iff F and G are bisimilar. □

Now let $R = \{\langle F, G\rangle : F, G \in EA \text{ and } F \text{ and } G \text{ are bisimilar}\}$. Since R is a $|\tau|$-bisimulation (the largest one), there exists a transition $\rho : R \to [\![\tau]\!]_R$ such that the diagram

$$\begin{tikzcd}
R \arrow[r, "\pi_j"] \arrow[d, "\rho"'] & EA \arrow[d, "E\alpha"] \\
{[\![\tau]\!]_R} \arrow[r, "|\tau|\pi_j"] & {[\![\tau]\!]_{EA}}
\end{tikzcd}$$

commutes for $j = 1$ and $j = 2$.

Lemma 5.2. *$\theta_\alpha(F) = \theta_\alpha(G)$ implies $|\tau|\theta_\alpha(E\alpha(F)) = |\tau|\theta_\alpha(E\alpha(G))$, where $|\tau|\theta_\alpha : [\![\tau]\!]_{EA} \to [\![\tau]\!]_{\Delta A}$ is the result of applying the functor $|\tau|$ to $\theta_\alpha : EA \to \Delta A$.*

Proof. From the last diagram, for $j = 1, 2$,

$$|\tau|\theta_\alpha \circ |\tau|\pi_j \circ \rho = |\tau|\theta_\alpha \circ E\alpha \circ \pi_j,$$

so as $|\tau|$ is a functor,

$$|\tau|(\theta_\alpha \circ \pi_j) \circ \rho = |\tau|\theta_\alpha \circ E\alpha \circ \pi_j.$$

But Theorem 5.1 states that $\langle F, G\rangle \in R$ iff $\theta_\alpha(F) = \theta_\alpha(G)$, so the functions $\theta_\alpha \circ \pi_1$and $\theta_\alpha \circ \pi_2$ from R to ΔA are identical, hence by the last displayed equation,

$$|\tau|\theta_\alpha \circ E\alpha \circ \pi_1 = |\tau|\theta_\alpha \circ E\alpha \circ \pi_2.$$

Thus if $\theta_\alpha(F) = \theta_\alpha(G)$, then $\langle F, G\rangle \in R$ with

$$|\tau|\theta_\alpha \circ E\alpha \circ \pi_1\langle F, G\rangle = |\tau|\theta_\alpha \circ E\alpha \circ \pi_2\langle F, G\rangle,$$

i.e. $|\tau|\theta_\alpha(E\alpha(F)) = |\tau|\theta_\alpha(E\alpha(G))$ as desired. □

Theorem 5.3. *There is a unique function $\Delta\alpha : \Delta A \to [\![\tau]\!]_{\Delta A}$ making the following diagram commute.*

$$\begin{tikzcd}
EA \arrow[r, "\theta_\alpha"] \arrow[d, "E\alpha"'] & \Delta A \arrow[d, dotted, "\Delta\alpha"] \\
{[\![\tau]\!]_{EA}} \arrow[r, "|\tau|\theta_\alpha"] & {[\![\tau]\!]_{\Delta A}}
\end{tikzcd}$$

Proof. Define $\Delta\alpha$ by putting $\Delta\alpha(\theta_\alpha(F)) = |\tau|\theta_\alpha(E\alpha(F))$. Lemma 5.2 ensures that this is *well-defined.* Since θ_α is surjective, the domain of $\Delta\alpha$ is ΔA. The definition of $\Delta\alpha$ makes the diagram commute and is the only definition that can do so. □

This result defines $\Delta\alpha$ as a τ-coalgebra and, importantly, makes θ_α a surjective *morphism* from $E\alpha$ to $\Delta\alpha$. $\Delta\alpha$ is the *definable enlargement of* α. Theorem 5.1 states that the kernel of θ_α is the bisimilarity relation on $E\alpha$, so $\Delta\alpha$ is isomorphic to the quotient of $E\alpha$ by bisimilarity. Hence $\Delta\alpha$ is a *simple* coalgebra, i.e. itself has no proper quotients [Rut00, Proposition

8.2]. In $\Delta\alpha$ itself, bisimilar states are equal. That also follows from Theorem 5.1, since bisimilarity is invariant under morphisms, so $\theta_\alpha(F)$ and $\theta_\alpha(G)$ are bisimilar in $\Delta\alpha$ precisely when F and G are bisimilar in $E\alpha$, i.e. precisely when $\theta_\alpha(F) = \theta_\alpha(G)$.

The morphism θ_α can be used to transfer the Truth Lemma 3.3 for $E\alpha$, and its Corollary 3.4, to the corresponding results for $\Delta\alpha$:

Theorem 5.4. *Let $\Gamma \rhd \varphi$ be a rigid observable formula.*

(1) *For any $G \in \Delta A$, $\Delta\alpha, G \models \Gamma \rhd \varphi$ iff $(\Gamma \rhd \varphi)^\alpha \in G$.*
(2) *$\alpha \models \Gamma \rhd \varphi$ iff $\Delta\alpha \models \Gamma \rhd \varphi$.*

Proof.

(1) Given G, choose $F \in EA$ with $G = \theta_\alpha(F)$. Then as θ_α is a morphism, Theorem 2.4 yields that $\Delta\alpha, \theta_\alpha(F) \models \Gamma \rhd \varphi$ iff $E\alpha, F \models \Gamma \rhd \varphi$, which in turn holds iff $(\Gamma \rhd \varphi)^\alpha \in G$ by the Truth Lemma 3.3.
(2) From Corollary 3.4 we already know that $\alpha \models \Gamma \rhd \varphi$ iff $E\alpha \models \Gamma \rhd \varphi$. But as the morphism θ_α is surjective, Theorem 2.4 yields that $E\alpha \models \Gamma \rhd \varphi$ iff $\Delta\alpha \models \Gamma \rhd \varphi$. □

The morphism θ_α can also be used to transfer the State-Term Lemma 3.5 and its Corollary 3.6 to $\Delta\alpha$:

Lemma 5.5. *Let M be ground term of type St and let $X \in \boldsymbol{Def}_\alpha$. Then for any $G \in EA$,*

$$X \in [\![M]\!]_{\Delta\alpha}(G) \quad \textit{iff} \quad [\![M]\!]_\alpha^{-1}(X) \in G.$$

Consequently,

$$[\![M]\!]_{\Delta\alpha}^{-1}(X^{\Delta A}) = ([\![M]\!]_\alpha^{-1}(X))^{\Delta A},$$

where in general $Y^{\Delta A} = \{G \in \Delta A : Y \in G\}$.

Proof. Note first that if $X = \varphi^\alpha$, then Theorem 2.3 states that $[\![M]\!]_\alpha(x) \in X$ iff $x \in \varphi[M/\mathsf{s}]^\alpha$, so $[\![M]\!]_\alpha^{-1}(X) = \varphi[M/\mathsf{s}]^\alpha$, showing that $[\![M]\!]_\alpha^{-1}(X)$ is also definable.

Now let $G = \theta_\alpha(F)$ with $F \in EA$. Then as θ_a is a morphism,

$$[\![M]\!]_{\Delta\alpha}(G) = \theta_\alpha([\![M]\!]_{E\alpha}(F)) = [\![M]\!]_{E\alpha}(F) \cap \boldsymbol{Def}_\alpha,$$

so as X is definable, $X \in [\![M]\!]_{\Delta\alpha}(G)$ iff $X \in [\![M]\!]_{E\alpha}(F)$, which holds iff $[\![M]\!]_\alpha^{-1}(X) \in F$ by Lemma 3.5. But $[\![M]\!]_\alpha^{-1}(X) \in F$ iff $[\![M]\!]_\alpha^{-1}(X) \in G$, since $[\![M]\!]_\alpha^{-1}(X)$ is definable as we just saw.

The rest of the Lemma then follows straightforwardly. □

It follows from Theorem 5.4(2) that the class of all models of a rigid observable formula is closed under definable enlargements. In fact, in the structural characterization of such model classes set out in Theorem 2.10, observational ultrapowers can be replaced by ultrafilter enlargements, or by definable enlargements. To see this, first consider a class K of coalgebras that is closed under images of bisimulations. Then in particular it is closed under domains and images of coalgebraic morphisms, which means that for any surjective morphism $f : \alpha \twoheadrightarrow \beta$ we have $\alpha \in K$ iff $\beta \in K$. This follows because the image of f is the image of the bisimulation R_f (the graph of f), while the domain of f is the image of the inverse relation R_f^{-1}, which is also a bisimulation.

Now for any τ-coalgebra α, if α^+ is an enlarging observational ultrapower of α we have surjective morphisms

$$\begin{tikzcd} \alpha^+ \arrow[r, two heads, "\alpha^+"] & E\alpha \arrow[r, two heads, "\theta_\alpha"] & \Delta\alpha. \end{tikzcd}$$

Thus if K is closed under images of bisimulations, and contains one of these three coalgebras, then it contains the other two as well. This observation, together with the equivalences of Theorem 2.10, yields the following extension of that Theorem.

Theorem 5.6. *If τ has at least one non-trivial observable subtype, then for any class K of τ-coalgebras, the following are equivalent.*

(1) *K is the class of all models of some set of rigid observable formulas.*
(2) *K is closed under disjoint unions, images of bisimulations, and ultrafilter enlargements.*
(3) *K is closed under disjoint unions, images of bisimulations, and definable enlargements.* □

6. Monads From Enlargements

In this section a category-theoretic perspective on coalgebraic enlargements is developed. The operation of assigning to each set A the collection of all ultrafilters on A gives rise to a categorical structure on the category **Set** of sets and functions that is known as a *monad* or *triple* (see [ML71, Chapter VI] or [Man76]). In a similar way, the $E\alpha$ construction gives rise to a monad on the category τ-**Coalg** of τ-coalgebras and their morphisms.

For any morphism $f : (A, \alpha) \to (B, \beta)$ of τ-coalgebras, define a function Ef on EA by putting

$$Ef(F) = \{Y \subseteq B : f^{-1}Y \in F\}.$$

Lemma 6.1. *Ef is a morphism $(EA, E\alpha) \to (EB, E\beta)$.*

Proof. It is standard theory that $Ef(F)$ is an ultrafilter on B whenever F is an ultrafilter on A. To show it is observationally rich we use the fact, from the second sentence of Theorem 2.4, that for any ground observable formula φ, we have $x \in \varphi^\alpha$ iff $f(x) \in \varphi^\beta$ in general, and so $\varphi^\alpha = f^{-1}\varphi^\beta$.

For any ground observable term $M : o$, α-richness of F implies that $(M \approx c)^\alpha \in F$ for some $c \in \llbracket o \rrbracket$. Thus $f^{-1}(M \approx c)^\beta \in F$ by the last observation, and so $(M \approx c)^\beta \in Ef(F)$. This shows that $Ef(F)$ is a β-rich ultrafilter, so that Ef is indeed a function from EA to EB.

Then to show Ef is a morphism it suffices, by Theorem 2.5, to show that for any $F \in EA$ and any ground term M,

(1) $\llbracket M \rrbracket_{E\alpha}(F) = \llbracket M \rrbracket_{E\beta}(Ef(F))$ if M is observable;
and
(2) $Ef(\llbracket M \rrbracket_{E\alpha}(F)) = \llbracket M \rrbracket_{E\beta}(Ef(F))$ if M is of type St.

For (1), let $\llbracket M \rrbracket_{E\alpha}(F) = c$. Then $E\alpha, F \models M \approx c$, so $(M \approx c)^\alpha \in F$ by the Truth Lemma 3.3. It follows as above that $(M \approx c)^\beta \in Ef(F)$, hence $E\beta, Ef(F) \models M \approx c$ by 3.3 again. Thus

$$\llbracket M \rrbracket_{E\beta}(Ef(F)) = c = \llbracket M \rrbracket_{E\alpha}(F).$$

For (2), as f is a morphism Theorem 2.5(2) states that the diagram

$$\begin{tikzcd}
A \arrow[r, "f"] \arrow[d, "\llbracket M \rrbracket_\alpha"'] & B \arrow[d, "\llbracket M \rrbracket_\beta"] \\
A \arrow[r, "f"] & B
\end{tikzcd}$$

commutes. Hence for any $Y \subseteq B$,

$$\llbracket M \rrbracket_\alpha^{-1}(f^{-1}Y) = f^{-1}(\llbracket M \rrbracket_\beta^{-1}Y).$$

Then

$$\begin{aligned}
& Y \in Ef([\![M]\!]_{E\alpha}(F)) && \\
& \text{iff } f^{-1}Y \in [\![M]\!]_{E\alpha}(F) && \text{definition of } Ef \\
& \text{iff } [\![M]\!]_{\alpha}^{-1}(f^{-1}Y) \in F && \text{State-Term Lemma 3.5} \\
& \text{iff } f^{-1}([\![M]\!]_{\beta}^{-1}Y) \in F && \text{from above} \\
& \text{iff } [\![M]\!]_{\beta}^{-1}Y \in Ef(F) && \text{definition of } Ef \\
& \text{iff } Y \in [\![M]\!]_{E\beta}(Ef(F)) && \text{State-Term Lemma 3.5.}
\end{aligned}$$

Since this holds for all $Y \subseteq B$, (2) follows. □

It is now readily seen that the assignments $\alpha \mapsto E\alpha$ and $f \mapsto Ef$ provide a functor $E : \tau\text{-}\mathbf{Coalg} \to \tau\text{-}\mathbf{Coalg}$ on the category of τ-coalgebras.

Theorem 6.2. *The morphisms $\eta_A : \alpha \to E\alpha$ are the components of a natural transformation η from the identity functor on τ-**Coalg** to the functor E.*

Proof. This amounts to the claim that for any morphism f the diagram

$$\begin{tikzcd}
A \arrow[r, "\eta_A"] \arrow[d, "f"'] & EA \arrow[d, "Ef"] \\
B \arrow[r, "\eta_B"] & EB
\end{tikzcd}$$

commutes in **Set**. But it is a simple set-theoretic calculation to show that $Ef(\eta_A(x)) = \eta_B(f(x))$ for all $x \in A$. □

Composing the functor E with itself gives the functor EE on τ-**Coalg** that assigns to each coalgebra (A, α) a coalgebra $(EEA, EE\alpha)$ whose states are the $E\alpha$-rich ultrafilters on EA. A function μ_α is defined on EEA by putting

$$\mu_\alpha(p) = \{X \subseteq A : X^{EA} \in p\},$$

where $X^{EA} = \{F \in EA : X \in F\}$, as in Section 3. Note that the notation μ_α is preferable to μ_A, since the definition depends on EA and hence on α. By contrast, the definition of η_A depends only on the set A.

Theorem 6.3. *The functions μ_α are the components of a natural transformation μ from EE to E.*

Proof. First it must be shown that μ_α is an arrow in τ-**Coalg** (a morphism) from $EE\alpha$ to $E\alpha$. The fact that the map $X \mapsto X^{EA}$ preserves the Boolean set operations ensures that for each ultrafilter $p \in EEA$, $\mu_\alpha(p)$ is an ultrafilter on A. Moreover $\mu_\alpha(p)$ is α-rich: for any ground term $M : o$ there is some $c \in [\![o]\!]$ with $(M \approx c)^{E\alpha} \in p$, and then by the Truth Lemma 3.3,

$$\{F \in EA : (M \approx c)^\alpha \in F\} = \{F : E\alpha, F \models M \approx c\} = (M \approx c)^{E\alpha} \in p,$$

so $(M \approx c)^\alpha \in \mu_\alpha(p)$ by definition of μ_α. This shows that $\mu_\alpha(p)$ is rich, so μ_α is a function from EEA to EA.

To show that μ_α is a morphism we apply Theorem 2.5, as in the proof of Lemma 6.1, this time showing that for any $p \in EA$ and any ground term M,

(1) $[\![M]\!]_{EE\alpha}(p) = [\![M]\!]_{E\alpha}(\mu_\alpha(p))$ if M is observable;
and
(2) $\mu_\alpha([\![M]\!]_{EE\alpha}(p)) = [\![M]\!]_{E\alpha}(\mu_\alpha(p))$ if M is of type St.

For (1), there exists an element c such that $(M \approx c)^{E\alpha} \in p$ and $(M \approx c)^\alpha \in \mu_\alpha(p)$ as in the previous paragraph. By the Truth Lemma these imply that $EE\alpha, p \models M \approx c$ and $E\alpha, \mu_\alpha(p) \models M \approx c$, so that

$$[\![M]\!]_{EE\alpha}(p) = [\![M]\!]_{E\alpha}(\mu_\alpha(p)) = c.$$

For (2), we reason that for any $X \subseteq A$,

$$\begin{array}{ll} X \in \mu_\alpha([\![M]\!]_{EE\alpha}(p)) & \\ \text{iff } X^{EA} \in [\![M]\!]_{EE\alpha}(p) & \text{definition of } \mu_\alpha \\ \text{iff } [\![M]\!]^{-1}_{E\alpha}(X^{EA}) \in p & \text{State-Term Lemma 3.5} \\ \text{iff } ([\![M]\!]^{-1}_{\alpha}(X))^{EA} \in p & \text{Corollary 3.6} \\ \text{iff } [\![M]\!]^{-1}_{\alpha} X \in \mu_\alpha(p) & \text{definition of } \mu_\alpha \\ \text{iff } X \in [\![M]\!]_{E\alpha}(\mu_\alpha(p)) & \text{State-Term Lemma 3.5.} \end{array}$$

Thus $\mu_\alpha([\![M]\!]_{EE\alpha}(p)) = [\![M]\!]_{E\alpha}(\mu_\alpha(p))$, completing the proof that μ_α is a morphism in τ-**Coalg**.

Finally, to show μ is natural it must be shown that the diagram

$$\begin{tikzcd}
EE\alpha \arrow[r, "\mu_\alpha"] \arrow[d, "EEf"'] & E\alpha \arrow[d, "Ef"] \\
EE\beta \arrow[r, "\mu_\beta"] & E\beta
\end{tikzcd}$$

commutes in τ-**Coalg** whenever f is a morphism from α to β. This requires that

$$\begin{tikzcd}
EEA \arrow[r, "\mu_\alpha"] \arrow[d, "EEf"'] & EA \arrow[d, "Ef"] \\
EEB \arrow[r, "\mu_\beta"] & EB
\end{tikzcd}$$

commutes in **Set**, where A and B are the state sets of α and β. The proof of this is set-theoretic, requiring no further coalgebraic analysis, and is essentially part of the standard theory of ultrafilters [Man76, Section 1.3]. The details are left to the interested reader, who would find it useful to first show that for any $Y \subseteq B$,

$$(f^{-1}Y)^{EA} = (Ef)^{-1}(Y^{EB}).$$

□

The triple $\langle E, \eta, \mu \rangle$ forms a *monad* on the category τ-**Coalg**. In addition to the naturality of η and μ (Theorems 6.2 and 6.3), this means that for any τ-coalgebra (A, α) the following diagrams commute.

$$\begin{tikzcd}
EEE\alpha \arrow[r, "E\mu_\alpha"] \arrow[d, "\mu_{E\alpha}"'] & EE\alpha \arrow[d, "\mu_\alpha"] \\
EE\alpha \arrow[r, "\mu_\alpha"] & E\alpha
\end{tikzcd}
\qquad
\begin{tikzcd}
E\alpha \arrow[r, "E\eta_A"] \arrow[dr, "\mathrm{id}"'] & EE\alpha \arrow[d, "\mu_\alpha"] & E\alpha \arrow[l, "\eta_{E\alpha}"'] \arrow[dl, "\mathrm{id}"] \\
 & E\alpha &
\end{tikzcd}$$

Demonstration of this reduces to showing commutativity of the corresponding diagrams in **Set** that result from replacing $E\alpha$ by EA. Again these are standard ultrafilter calculations that need not be reproduced here. The reader who is interested to check the details would find it useful, in the case of the left diagram, to first show that for any $X \subseteq A$,

$$\mu_\alpha^{-1}(X^{EA}) = (X^{EA})^{EEA}.$$

The Definable Case

The construction $\alpha \mapsto \Delta\alpha$ also gives rise to a monad on τ-**Coalg**. First of all, Δ extends to a functor on τ-**Coalg** that assigns to each morphism $f : (A, \alpha) \to (B, \beta)$ the function $\Delta f : \Delta A \to \Delta B$ having

$$\Delta f(G) = \{Y \in \boldsymbol{Def}_\beta : f^{-1}Y \in G\}.$$

The proof that Δf is a morphism from $\Delta\alpha$ to $\Delta\beta$ is similar to the proof that Ef is a morphism, using results 5.4 and 5.5 in place of 3.3 and 3.5. It is readily seen that the diagram

$$\begin{tikzcd}
E\alpha \arrow[r, "\theta_\alpha"] \arrow[d, "Ef"'] & \Delta\alpha \arrow[d, "\Delta f"] \\
E\beta \arrow[r, "\theta_\beta"] & \Delta\beta
\end{tikzcd}$$

commutes, so the morphisms θ_α are the components of a natural transformation θ from E to Δ.

A function $\eta^\Delta_\alpha : A \to \Delta A$ is defined by

$$\eta^\Delta_\alpha(x) = \{X \in \boldsymbol{Def}_\alpha : x \in X\} = \theta_\alpha(\eta_A(x)).$$

Then η^Δ_α is a morphism from α to $\Delta\alpha$, being the composition of the morphisms η_A and θ_α. The η^Δ_α's are the components of a natural transformation from the identity functor on τ-**Coalg** to Δ, the composition of η and θ. Note that, unlike η_A, η^Δ_α need not be injective: in general $\eta^\Delta_\alpha(x) = \eta^\Delta_\alpha(y)$ iff x and y satisfy the same ground observable formulas in α, which holds iff x and y are *bisimilar* (Theorem 2.7). Thus η^Δ_α is injective precisely when bisimilar states in α are equal.

A natural transformation $\mu^\Delta : \Delta\Delta \to \Delta$ is given by defining

$$\mu^\Delta_\alpha(p) = \{X \in \boldsymbol{Def}_\alpha : X^{\Delta A} \in p\},$$

where $X^{\Delta A} = \{G \in \Delta A : X \in G\} = \theta_\alpha(X^{EA})$ (see Lemma 5.5). The proof that μ^Δ_α is a morphism $\Delta\Delta\alpha \to \Delta\alpha$ is analogous to the proof that $\mu_\alpha : EE\alpha \to E\alpha$ is a morphism.

The triple $\langle \Delta, \eta^\Delta, \mu^\Delta \rangle$ forms a monad on τ-**Coalg**, but one of a special kind, as the functor Δ is "idempotent up to isomorphism", in the sense that $\Delta\alpha$ and $\Delta\Delta\alpha$ are isomorphic. A "logical" explanation of this is that if φ and ψ are ground observable formulas then by Theorem 5.4(2),

$$\alpha \models \varphi \leftrightarrow \psi \quad \text{iff} \quad \Delta\alpha \models \varphi \leftrightarrow \psi,$$

and so $\varphi^{\alpha} = \psi^{\alpha}$ iff $\varphi^{\Delta\alpha} = \psi^{\Delta\alpha}$. Thus the map $\varphi^{\alpha} \mapsto \varphi^{\Delta\alpha}$ is a well-defined bijection between the Boolean algebras $\boldsymbol{Def}_{\alpha}$ and $\boldsymbol{Def}_{\Delta\alpha}$ of definable subsets of α and $\Delta\alpha$, respectively. Moreover this map is a Boolean isomorphism and gives a bijection between the sets of α-rich ultrafilters of $\boldsymbol{Def}_{\alpha}$ and $\Delta\alpha$-rich ultrafilters of $\boldsymbol{Def}_{\Delta\alpha}$, taking $G \in \Delta A$ to $\{\varphi^{\Delta\alpha} : \varphi^{\alpha} \in G\} \in \Delta\Delta A$. This gives the isomorphism $\Delta\alpha \cong \Delta\Delta\alpha$.

But the version of the Truth Lemma for Δ given in Theorem 5.4(1) shows that

$$\varphi^{\alpha} \in G \quad \text{iff} \quad G \in \varphi^{\Delta\alpha},$$

so the isomorphism is the map $G \mapsto \{\varphi^{\Delta\alpha} : G \in \varphi^{\Delta\alpha}\} = \eta^{\Delta}_{\Delta\alpha}(G)$.

In other words, this isomorphism is just the component

$$\eta^{\Delta}_{\Delta\alpha} : \Delta\alpha \to \Delta\Delta\alpha$$

of the natural transformation η^{Δ}. Another proof that this component is a bijection follows from the observations firstly that $\eta^{\Delta}_{\Delta\alpha}$ is injective because $\Delta\alpha$ is a simple coalgebra in which bisimilar states are equal, and secondly that $\eta^{\Delta}_{\Delta\alpha}$ is surjective because for any $p \in \Delta\Delta A$ the set $G = \{\varphi^{\alpha} : \varphi^{\Delta\alpha} \in p\}$ is a rich ultrafilter of $\boldsymbol{Def}_{\alpha}$ with $\eta^{\Delta}_{\Delta\alpha}(G) = p$.

It is noteworthy that part of this monad structure on Δ is the property that the diagram

$$\begin{tikzcd}
\Delta\alpha \arrow[r, "\eta^{\Delta}_{\Delta\alpha}"] \arrow[dr, "\mathrm{id}"'] & \Delta\Delta\alpha \arrow[d, "\mu^{\Delta}_{\alpha}"] \\
 & \Delta\alpha
\end{tikzcd}$$

commutes, so in fact the component μ^{Δ}_{α} of the natural transformation μ^{Δ} is itself the inverse of the isomorphism $\eta^{\Delta}_{\Delta\alpha}$, and hence is also an isomorphism.

A monad on a category has an associated category of *algebras*. In the case of the Δ-monad, an algebra is a pair (α, f) with $f : \Delta\alpha \to \alpha$ a morphism for which the following commute:

$$\begin{tikzcd}
\Delta\Delta\alpha \arrow[r, "\mu^{\Delta}_{\alpha}"] \arrow[d, "\Delta f"'] & \Delta\alpha \arrow[d, "f"] \\
\Delta\alpha \arrow[r, "f"] & \alpha
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
\alpha \arrow[r, "\eta^{\Delta}_{\alpha}"] \arrow[dr, "\mathrm{id}"'] & \Delta\alpha \arrow[d, "f"] \\
 & \alpha
\end{tikzcd}$$

But for an idempotent monad like Δ, in which the components μ^{Δ}_{α} are all isomorphisms, any such algebra (α, f) has f an isomorphism [Bor94, Proposition 4.2.3].

For the ultrafilter monad on the category **Set**, the associated category of algebras is isomorphic to the category of compact Hausdorff topological spaces and continuous functions – this is Manes' Theorem, see [Man76] or [Joh82, Section III 2]. It would be of interest to know whether this situation lifts from **Set** to τ-**Coalg**, replacing the ultrafilter monad by the monad of E. Is there some topology that can be imposed on polynomial coalgebras that identifies a natural class of topological coalgebras isomorphic to the category of E-algebras $f : E\alpha \to \alpha$?

References

AM89. 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.

Bor94. Francis Borceux. *Handbook of Categorical Algebra 2. Categories and Structures.* Cambridge University Press, 1994.

Gol01a. Robert Goldblatt. Equational logic of polynomial coalgebras. In *Advances in Modal Logic, volume 4*, World Scientific (to appear). Manuscript available at `http://www.mcs.vuw.ac.nz/~rob`

Gol01b. Robert Goldblatt. Observational ultrapowers of polynomial coalgebras. Manuscript available at `http://www.mcs.vuw.ac.nz/~rob`

Gol01c. Robert Goldblatt. What is the coalgebraic analogue of Birkhoff's variety theorem? *Theoretical Computer Science*, 266:853–886, 2001.

Her93. Claudio Hermida. *Fibrations, Logical Predicates and Indeterminates.* PhD thesis, University of Edinburgh, 1993. Techn. rep. LFCS-93-277. Also available as Aarhus Univ. DAIMI Techn. rep. PB-462.

HJ98. Claudio Hermida and Bart Jacobs. Structural induction and coinduction in a fibrational setting. *Information and Computation*, 145:107–152, 1998.

Jac96. Bart Jacobs. Objects and classes, coalgebraically. In B. Freitag, C. B. Jones, C. Lengauer, and H.-J. Schek, editors, *Object-Orientation with Parallelism and Persistence*, pages 83–103. Kluwer Academic Publishers, 1996.

Jac00. Bart Jacobs. Towards a duality result in coalgebraic modal logic. *Electronic Notes in Theoretical Computer Science*, 33, 2000. `http://www.elsevier.nl/locate/entcs`.

Joh82. P. T. Johnstone. *Stone Spaces.* Cambridge University Press, 1982.

Man76. Ernest G. Manes. *Algebraic Theories.* Springer-Verlag, 1976.

ML71. Saunders Mac Lane. *Categories for the Working Mathematician.* Springer-Verlag, 1971.

Pit00. Andrew M. Pitts. Categorical logic. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science, Volume 5: Algebraic and Logical Structures*, chapter 2. Oxford University Press, 2000.

Rei95. Horst Reichel. An approach to object semantics based on terminal co-algebras. *Mathematical Structures in Computer Science*, 5:129–152, 1995.

Rut95. 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.

Rut00. J.J.M.M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3–80, 2000.