Connections of Coalgebra and Semantic Modeling

Lawrence S. Moss
Department of Mathematics
Indiana University
Bloomington, IN 47405 USA
lsm@cs.indiana.edu

ABSTRACT

The aim of this tutorial is to present the area of coalgebra to people interested in the kinds of semantic modeling that is prominent at TARK. Coalgebra is a general study of a great many kinds of models, and these include type spaces and Kripke models, and many others. But the theory is not overly general, it is not a theory of absolutely everything. The tutorial is designed to be a short introduction to a substantial technical field, bearing in mind that this is nearly impossible. It is also intended to bring together literatures from theoretical computer science and game theory.

Categories and Subject Descriptors

F.3.2 [**Semantics of Programming Languages**]: Algebraic Approaches to Semantics; F.4.1 [**Mathematical Logic**]: Modal Logic

General Terms

Theory of computation, category theory, Harsanyi type space, modal logic, probabilistic logic

Keywords

coalgebra, type space, modal logic

1. INTRODUCTION

This paper is a tutorial introduction to topics in *coalgebra* that I think could be of interest to participants in the TARK 2011 meeting, and to the larger community of people interested in a set of very basic semantic issues related to probabilistic systems in the theoretical parts of computer science and economics. It is partly a "survey" and partly a "tutorial" as such, but I intend my lecture to have more of a conversational flavor and thus be more of a tutorial.

Here is one way to the issues in this paper. In the literature on game theory and economics, there exists a long discussion of *universal type spaces*, starting with Harsanyi [H].

TARK 2011, July 12-14, 2011, Groningen, The Netherlands.

They are mathematical structures used in modeling imperfect information. The idea is that agents are described by their *types*, and these types give us "beliefs about the world", "beliefs about each other's beliefs about the world", "beliefs about each other's beliefs about each other's beliefs about the world", etc. That is, the formal concept of a type is intended to capture in one object an unfolding infinite hierarchy related to *interactive, higher-order belief*. Since types could be defined in terms of beliefs about types of the other agents, it seems that there is something circular about this very concept. Finally, there is a question about what the ultimate, most abstract type space could possibly be.

Harsanyi did not really formalize type spaces in his original paper; this was left to later researchers starting with Böge and Eisele's paper [BE]. To make matters concrete, I want to quote a precise definition, taken from Heifetz and Samet [HS].

Type spaces Fix a measurable space S the elements of which are called *states of nature*.

Definition 3.1. A *type space* on S is a pair

$$((T_i)_{i\in I_0}, (m_i)_{i\in I}),$$

or (T, m) for short, where

1. $T_0 = S$, and T_i , for $i \in I$, is a measurable space.
2. For each $i \in I$, m_i is a measurable function $m_i : T_i \to \Delta(T)$.
3. For each $i \in I$ and $t_i \in T_i$, the marginal of $m_i(t_i)$ on T_i is t_i.

Type morphisms Let (T, m) and (T', m') be type spaces on S. ...

Definition 3.2. Let $(\phi_i)_{i\in I_0}$ be an I_0-tuple, of measurable functions $\phi.i : T_i \to T'_i$. The induced function $\phi : T \to T'$ is called a *type morphism* if,

1. ϕ_0 is the identity on S;
2. for each $i \in I$, $m'_i\phi_i = \hat{\phi}m_i$.

The morphism is a type isomorphism if ϕ is an isomorphism (or equivalently, if ϕ_i is an isomorphism for each $i \in I_0$).

Definition 3.3. A type space T^* on S is *universal* if for every type space T on S there is a unique type morphism from T to T^*.

Again, we have quoted Heifetz and Samet [HS], omitting some comments between the definitions. We should explain the notation a bit. Here I is a fixed set of *players*, I_0 =

$I \cup \{0\}$, with 0 standing for "nature." For a measurable space M, $\Delta(M)$ is the set of probability measures on M, itself made into a measurable space using the smallest σ-algebra containing the sets

$$\beta^p(E) = \{\mu : \mu(E) \geq p\},$$

where E is measurable in M. The function $\hat{\phi}$ in Definition 3.2 is of particular importance in our discussion, and for our purposes it would be better written as

$$\Delta\phi : \Delta T \to \Delta T'.$$

Here a general definition of this Δ. Let (M, Σ) and (N, Σ') be measurable spaces. If $f : M \to N$ is a measurable function, then for $\mu \in \Delta(M)$ and $A \in \Sigma'$, $(\Delta f)(\mu)(A) = \mu(f^{-1}(A))$. That is, $(\Delta f)(\mu) = \mu \circ f^{-1}$.

The main point of [HS] and several other papers on this topic (such as [BE, BD, MZ, Me, Va], to name just a few) is to construct and study the universal type space. By now there is a modest literature on the topic, and certainly this paper is a contribution to that literature. But are not only interested in the construction of the universal space, we also have in mind a different set of questions.

1. Is there an underlying mathematical construction that all of the work in the area exploits? Or are all of the different papers in it completely disjoint pieces of work?

2. Is there a way to understand the role of *topology* in constructions of universal type spaces?

3. Suppose one is not happy with type spaces in the first place and then proposes a different sort of mathematical definition for use in connection with game theory. Would there be an "easy" way to get the theory going? That is, would we expect an easy way to formulate the basic notions of equivalence of models and ask about universal structures? Or would we expect an extended literature?

Question (i) is the kind of question that category theory can help with, since it includes studies of very general constructions that have arisen in different areas of mathematics.

Here is an explanation of question (ii): The constructions of universal type spaces in that paper and most of the succeeding literature worked on categories which combined measure-theoretic and topological structure. The topological work is to some extent unfortunate, since it does not appear to be close to the original motivations for type spaces. Heifetz and Samet's paper [HS] was the first to avoid the topological setting, and it seems to have been written explicitly for the purpose of getting away from the extra topological assumptions that had been common in the area. Another paper which has the same goal is Pintér [Pi]. We shall have more to say about both papers later.

The final question in our list is one that should interest people in the area, but I have not seen any discussion of it. The semantic models used in game theory and applied logic are always controversial, and much of the discussion on the foundational side concerns the adequacy of various mathematical models and the appropriateness of various assumptions that we have about them. At some point people may propose different models. Certainly the computer science literature that this paper also addresses is rife with different kinds of model, more so than the economics/game theory literature. The work that I am reporting on is of potential application to new models. One cannot be sure that a completely new kind of model will fall under the purview of coalgebra, of course. But I feel confident in saying that a mathematical model addressing the same kind of issues (uncertainty in multi-agent settings, higher-order belief, actions which might change the model in addition to the players' beliefs) is likely to be a coalgebra of some functor or other, or to be an example of something related. I say this based on the experience and usefullness of coalgebras in many settings.

2. BACKGROUND: CATEGORY THEORY

When we first mentioned type spaces in Section 2.1, we noted that they came with a notion of *morphism*, and also that the definition involved applying the probability measure operation Δ not only to a space M but to a function ϕ. The notion of a *category* provides the setting where one can do all of this. A category (say $\mathcal{A}$) comes with *objects* (A, B, X, ...) and *morphisms* (f, g, ϕ, ...); these are undefined terms and so can be instantiated in a number of ways. The morphisms are required to have specific objects as their *domain* and *codomain*; we indicate these by writing $f : A \to B$. A category is required to have *identity morphisms* on its objects, and it also must have a partially-defined operation of *composition* which must be associative respect the identity morphisms in the obvious sense. (We are being terse here for lack of space, and the reader who is not familiar with any of this will no doubt have to consult other sources.) Given categories $\mathcal{A}$ and $\mathcal{B}$, a *functor* $F : \mathcal{A} \to \mathcal{B}$ is a mapping of objects of $\mathcal{B}$ to objects of $\mathcal{A}$, and similarly for morphisms, and this mapping must preserve the identities and compositions.

We are going to be interested in several categories $\mathcal{A}$ and *endofunctors* $F : \mathcal{A} \to \mathcal{A}$. Chief among the categories are the category Set of all sets, with morphisms the functions, and for our purposes we also need the category Meas of all measurable sets M and measurable functions. As examples of $F : \mathsf{Set} \to \mathsf{Set}$, we mention the *power set* functor $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$ defined by: $\mathcal{P}(X)$ is the set of all subsets of X, and for $f : X \to Y$, $\mathcal{P}f : \mathcal{P}X \to \mathcal{P}Y$ is defined by taking forward images: $A \subseteq X \mapsto \{f(x) : x \in A\}$. We have also seen the endofunctor $\Delta : \mathsf{Meas} \to \mathsf{Meas}$.

An *algebra for F* is a pair (A, a) with A an object of $\mathcal{A}$ and $a : FA \to A$. Dually, a *coalgebra for F* is a pair (A, a) with A an object of $\mathcal{A}$ and $a : A \to FA$.

Given two algebras, (A, a) and (B, b), a *morphism* from the first to the second is a morphism $f : A \to B$ so that $fa = bFf$. An *initial algebra* is an algebra with the additional property that for every algebra, there is a unique morphism from the first to the second:

$$\text{initial algebra} \qquad \begin{array}{ccc} FA & \xrightarrow{a} & A \\ {\scriptstyle Ff}\downarrow & & \downarrow{\scriptstyle f} \\ FB & \xrightarrow[b]{} & B \end{array}$$

We obtain the dual concepts of a *coalgebra morphism* and *final coalgebra* by turning the arrows around.

$$\begin{array}{ccc} A & \xrightarrow{a} & FA \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle Ff} \\ B & \xrightarrow[b]{} & FB \end{array} \qquad \text{final coalgebra}$$

We are interested in coalgebra in this paper because type spaces (see Section 2.1) are coalgebras, and formulating things this way gets the notion of a morphism correct "automatically." Moreover, we have some other examples of interest.

Examples of coalgebras of functors on Set.

A coalgebra for $\mathcal{P}$ is a set X together with a function from X to its power set. This is a "repackaged" form of an undirected graph: think of the function as specifying, for each point $x \in X$, the successors of x in the graph.

A Kripke model would be a coalgebra for $FX = \mathcal{P}X \times A$, where A is the power set of some underlying set of atomic propositions. The good news is that the notion of a p-morphism from modal logic falls out as a special case of a morphism of coalgebras.

Similarly, various types of automata are (suitably repackaged as) coalgebras of other functors on Set (see Rutten [R] for many examples).

Let N be a set of *players* and let $S_1, \ldots, S_N$ be non-empty *strategy sets* for the players. A *game frame over* N is a function of the form

$$C \to \{((S_1, \ldots, S_N), f) : f : \prod S_N \to C\}.$$

It associates to each state $c \in C$ a strategic game with strategy sets $S_1, \ldots S_N$ and outcome function f.

Second, a *conditional frame* is a map of the form

$$C \to \{f : \mathcal{P}(C) \to \mathcal{P}(C) \mid f \text{ is a function}\}.$$

Every state yields a *selection function* that assigns *properties* (subsets of the frame) to *conditions* (subsets of the frame again).

All of these are coalgebras of different endofunctors on the category Set. The theory would then study notions of equivalence of models, generalizations of modal logic, principles of (co-)recursion, and the like.

Types spaces as coalgebras of functors on Meas *and* Meas^I.

Harsanyi type spaces are close to coalgebras on the category Meas of the endofunctor $\Delta(M \times Id)$ for a fixed space M. (Another way to write this functor is $FX = \Delta(M \times X)$.) But type spaces are about several "players" or agents, so they really are coalgebras of a related functor on a category Meas^I for a fixed set I. And then one imposes some conditions on the functor related to the intuitions of self-knowledge by the players.

We'll make all of this precise in a moment, but the reader not interested in these details may omit this passage. For most purposes related to the construction, it is easier to think of functors such as $FX = \Delta X \times S$ or $FX = \Delta(X \times S)$ than to work with the category Meas^I which we now introduce. This discussion is from [MV1].

Fix a finite set I of *players*. We assume that $0 \notin I$ and as above we let $I_0 = I \cup \{0\}$. The objects of Meas^I are families $X = (X_i)_{i \in I}$ of measurable spaces, and the morphisms are also families of measurable maps. Fix a measurable space M to represent the "states of nature", and write X_0 for M. Each player should have beliefs about nature and about the beliefs of the other players. Moreover, players should know their own types. After some work, both conceptual and measure-theoretic, we arrive at the following formulation.

Let $\mathsf{C} : \mathsf{Meas}^I \to \mathsf{Meas}$ be the functor given by

$$\mathsf{C}X \quad = \quad \prod_{i \in I_0} X_i$$

At first glance, it might look like what we want is to consider for each $x_i \in X_i$, a probability measure on $\mathsf{C}X$. However, this is *not* what we want because it misses an important intuition concerning type spaces. This is that . In other words, each player i should only have beliefs about the (joint distribution on) other players' beliefs; i's own beliefs should not even enter in.

Define functors $U_i : \mathsf{Meas}^I \to \mathsf{Meas}$ given by

$$U_i(X) \quad = \quad \Pi\{X_j \mid j \in I_0, j \neq i\}.$$

U_i acts the obvious way on morphisms. Note that U_i depends on the space X_0 of states of nature, even though our notation does not mention this.

Let $F : \mathsf{Meas}^I \to \mathsf{Meas}^I$ be defined by

$$\mathsf{F}(X) \quad = \quad (\Delta U_i(X))_{i \in I}$$

As before Δ is the probability measure space functor, and once again, our notation elides the underlying space X_0 of states of nature.

This way, instead of having a family of functions in Meas, each one of them with a condition on one of its marginals, any morphism in Meas^I works as a coalgebra structure. The particular functor we use automatically takes care of the condition on marginals.

With these preliminaries done, we can now mention a reformulation of the notion of a type space, due to Viglizzo [Vi]:

A *Harsanyi type space (over M)* is a coalgebra for the functor F in the category Meas^I. A *universal type space* is a final coalgebra for F in Meas^I.

This paper has several aims, and one of them is to discuss the matter of *finding a final coalgebra* if one exists. Before we start in on this topic, we should explain why we are interested in it.

Lemma 2.1 (Lambek [L]). *Let $\mathcal{A}$ be any category, let $F : \mathcal{A} \to \mathcal{A}$ be any endofunctor, and let (A, a) be an initial algebra of F. Then a has an inverse b. (That is, $ba = id_{FA}$, and $ab = id_A$.)*

The same goes for final coalgebras.

By Lambek's Lemma, if one wants, for example, an object X isomorphic to FX, one might as well try to find an initial algebra or final coalgebra of F. Even more, if one wants a solution to $X \cong FX$ which is as "big as possible," then one should try for a final coalgebra: not only will $X \cong FX$ (by Lambek's Lemma), but every Y satisfying $Y \cong FY$ will (be a coalgebra and hence) automatically map uniquely into X.

One last point before we turn to a construction method, summarizing what we have done so far: the definition in Section 2.1 of a *type space* is essentially that of a coalgebra of a certain functor on Meas, and with this definition, the notion of a *type space morphism* comes automatically, as does the definition of a a universal type space.

2.1 Final colagebras obtained by taking limits

We turn to the general matter of constructing final coalgebras. There are many constructions; all have special conditions so the matter is rather technical. The most common methods are generalizations of the following well-known result.

THEOREM 2.2 (KLEENE). *Let $\mathcal{A}$ be a complete partial order with a least element $\perp$. Then every continuous endofunction $F : \mathcal{A} \to \mathcal{A}$ has a least fixed point μF given concretely by*

$$\mu F = \sup_{n\in\omega} F^n(\perp).$$

The reason we are interested in this result is that posets are catgories (the objects are the points of the poset, and the morphisms $f : p \to q$ are just the statements that $p \leq q$). Then an endofunctor on this category is the same as a *montone* (or *order-preserving* function on the poset. Kleene's Theorem not only gives a least fixed point, it also gives an initial algebra of the functor. Indeed, we are not really interested in Kleene's Theorem but in generalizations of it, and in dualizations of those generalizations, etc. Figure 1 shows the "categorification" of the order-theoretic concepts in Kleene's Theorem. In each line, the order-theoretic concept on the left is a special case of the category-theoretic concept to its right. Of special interest is the generalization of the completeness condition on the poset and the continuity condition on the function. We assume that every ω-chain, that is, a functor from ω to $\mathcal{A}$ has a colimit, and in particular we consider the *initial ω-chain* shown in (1). On needs this in Kleene's Theorem: for example, the natural numbers $(N, \leq)$ with $fn = n + 1$ is monotone and yet has no fixed points.

DEFINITION 2.3. *An* initial object *in a category is an object* 0 *with the property that for every object A, there is exactly one morphism from* 0 *to A, denoted* $!_A : 0 \to A$.

By the initial ω-chain *of an endofunctor* $F : \mathcal{A} \to \mathcal{A}$ *is meant the ω-chain*

$$0 \xrightarrow{\ !\ } F0 \xrightarrow{F!} F^2 0 \xrightarrow{F^2!} \cdots \qquad (1)$$

For an example of an initial object, in Set it is the empty set. In the category of groups and group homomorphisms, it is the trivial one-element group; this group is also *final*: there is exactly one homomorphism into it from every group. Returning to Set, the final objects there are exactly the singleton (one-element) sets.

Our main result in this section is Theorem 2.5, the categorification of Kleene's Theorem, due to Jiří Adámek. We need a few definitions.

A *cocone* over the initial ω-chain is an object A of $\mathcal{A}$ together with a family of morphisms $a_n : F^n 0 \to A$ such that $a_n = a_{n+1} F^n!$ for all $n < \omega$. A *colimit* of the initial ω-chain is a cocone $(C, c_n : F^n 0 \to C)$ over it with the universal property that if $(A, \alpha_n : F^n 1 \to A)$ is any cocone over the initial ω-chain, then there is a unique *factorizing morphism* $f : C \to A$ such that for all $n < \omega$, $a_n = f c_n$.

CONSTRUCTION 2.4. *Every F-algebra (A, α) induces a family of maps $\alpha_n : F^n 0 \to A$ over the initial ω-chain as follows: $\alpha_0 : 0 \to A$ is unique (since 0 is initial) and*

$$\alpha_{n+1} = \alpha \cdot F\alpha_n : F(F^n 0) \to A. \qquad (2)$$

order	category
preorder $(A, \leq)$	category $\mathcal{A}$
$x \leq y$ and $y \leq x$	A and B are isomorphic
least element 0	initial object 0
monotone $F : A \to A$	functor $F : \mathcal{A} \to \mathcal{A}$
pre-fixed point: $Fx \leq x$	F-algebra: $f : FA \to A$
ω-chain	functor from $(\omega, \leq)$ to $\mathcal{A}$
F is continuous	F preserves ω-colimits
least pre-fixed point: $Fx \leq x$	initial F-algebra

Figure 1: Generalizing Kleene's Theorem to categories

This is a cocone, and we call it the cocone (A, α_n) the cocone induced by *A*.

Let $c_n : F^n 0 \to \mu F$ be the colimit of the initial ω-chain. Applying F to each object and morphism in (1) yields another ω-chain

$$F0 \xrightarrow{F!} F^2 0 \xrightarrow{F^2!} \cdots \qquad (3)$$

which obviously has the same colimit as (1).

This leads to the following result:

THEOREM 2.5. *[A74] Let $\mathcal{A}$ be a category with initial object* 0 *and with colimits of ω-chains. If $F : \mathcal{A} \to \mathcal{A}$ preserves ω-colimits, then it has an initial algebra.*

We are going to give the proof in full so that the reader familiar with Kleene's Theorem can see the ways in which various steps are generalized.

PROOF. Since F preserves colimits, Fc_n is a colimit cocone of (3). We have another cocone of (3), namely (c_{n+1}). Hence there is a unique morphism $\varphi : F(\mu F) \to \mu F$ be uniqueley defined by the condition that that

$$\varphi \cdot Fc_n = c_{n+1} : F(F^n 0) \to \mu F.$$

for all n. We claim that the F-algebra $(\mu F, \phi)$ is initial. To check this, let (A, α) be any F-algebra. Consider the cocone $\alpha_n : F^n 0 \to A$ from Construction 2.4. It is easy to check that $(F\alpha_n)$ is a cocone from (3) to $F(\mu F)$. Let $h : \mu F \to A$ be the unique morphism with $h \cdot c_n = \alpha_n$ for all n. To see that $h \cdot \varphi = \alpha \cdot Fh$, we check that both are mediating morphisms for the cocone (α_{n+1}) of (3). That is, we check that for all n,

$$\begin{array}{rcl} (h\phi)Fc_n & = & \alpha_{n+1} \\ (\alpha Fh)Fc_n & = & \alpha_{n+1} \end{array}$$

For the first assertion, $h(\phi Fc_n) = hc_{n+1} = \alpha_{n+1}$. For the second, consider the diagram below:

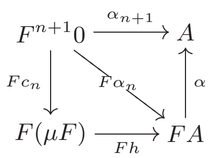

The upper triangle commutes by (2), and the lower by definition of h. So the square commutes, showing our that $(\alpha Fh)Fc_n = \alpha_{n+1}$. This verifies that h is an algebra morphism.

For the uniqueness, suppose that $k: \mu F \to A$ is also an F-algebra morphism: $k \cdot \varphi = \alpha \cdot Fk$. We show by induction on n that $k \cdot c_n = \alpha_n$; then the uniqueness of h shows that $k = h$. For $n = 0$, kc_n and α_n are both morphisms with domain 0, so they are the same. Assuming that $kc_n = \alpha_n$, we see that

$$\begin{array}{rcll} kc_{n+1} & = & k\phi Fc_n & \text{by definition of } \phi \\ & = & \alpha F(kc_n) & k \text{ is an } F\text{-algebra morphism} \\ & = & \alpha \cdot F\alpha_n & \text{by induction hypothesis} \\ & = & \alpha_{n+1} & \text{by (2)} \end{array}$$

This concludes the induction, and hence the overall proof. □

To obtain an initial algebra of an endofunctor $F\mathcal{A} \to \mathcal{A}$, it is not really necessary that $\mathcal{A}$ have colimits of all ω-chains or that F preserve all ω-colimits. It is sufficient to assume that the colimit of the initial ω-chain exists and that F preserve this colimit. That is, these are the only facts about $\mathcal{A}$ and F that are used in the proof of Theorem 2.5. In many cases, it is just as easy to verify the stronger requirements that we stated in Theorem 2.5 than it is to verify the special cases used in the proof. But this is not always the case, and in topological and measure-theoretic settings (such as those needed in work on type spaces) the sufficient conditions hold but the stronger ones do not.

The details in the proof of Theorem 2.5 are based on the exposition in Adámek et al. [AMM].

Here is an example of an initial algebra. Consider the category Set of sets, and the functor $\mathcal{D} : \mathsf{Set} \to \mathsf{Set}$ taking a set X to the set of discrete probability measures on X. This is a function $\pi : X \to [0, 1]$ which is 0 on all but finitely many points, and with the property that $\sum_{x \in X} \pi x = 1$. The action of $\mathcal{D}$ on functions is by summing inverse images, as in its cousin Δ. Let A be a set of *observations*, and let $FX = \mathcal{D}X + A$. Then a coalgebra for F is basically a Markov chain: it would be pair (S, f) where S a set of states, and each state $s \in S$ is either an end-state (when $f(s)$ is an observation in the set A), or else $f(s) \in \mathcal{D}(S)$, and we have a probability distribution on the same set of states, giving a probabilistic "next-state" function. The initial algebra of F may be found using Theorem 2.5. However, it would take some work to calculate the colimit of the initial ω-chain. To make a long story short, we can repackage the initial algebra as the set T of finite "probability trees". These would be trees with the property that their leaves are labeled in the observation set A, each non-leaf has children with edges labeled by numbers that sum to 1, and with the extra property that the tree is "reduced" in an appropriate sense. Coalgebra overall has the general language needed to say what "reduced" means, and to give the parallel to similar notions (for example in automata theory). We invite you to try to formulate the appropriate definition in this case to see what is involved on the specific level related to this functor. In addition, it would be good to ferret out the initial algebra structure $FT \to T$.

2.2 Dualization

One of the features of category theory is that concepts and results *dualize* by turning the arrows around on morphisms. We have already seen this with the concepts of *algebra* and *coalgebra*. Recall from Section 2.1 that a *coalgebra* for an endofunctor F is an object A together with a morphism $a: A \to FA$. One dualizes initial objects 0 to final objects 1, ω-chains to ω^{op}-chains (that is, functors from ω^{op} to $\mathcal{A}$), colimits to limits, and the initial ω-chain of Definition 2.3 to the *final* ω^{op}*-chain* given by

$$1 \xleftarrow{\;!\;} F1 \xleftarrow{\;F!\;} F^2 1 \xleftarrow{\;F^2!\;} \cdots \tag{4}$$

With this in mind, the dualization of Theorem 2.5 is straightforward.

THEOREM 2.6 (BARR [BA]). *Let $\mathcal{A}$ be a category with a final object 1 and with limits of ω^{op}-chains. If $F : \mathcal{A} \to \mathcal{A}$ preserves limits of ω^{op}-chains, then it has a final coalgebra νF obtained by taking a limit:*

$$\nu F = \lim_{n \in \omega^{op}} F^n 1.$$

NOTATION 2.7. *We denote by*

$$\ell_n : \lim_{i \in \omega^{op}} F^i 1 \to F^n 1$$

the projections of the limit of the final ω^{op}-chain of F. If F has a final coalgebra, we denote it by

$$\nu F \quad or \quad \nu X.F(X).$$

Every coalgebra $\alpha : A \to FA$ induces a canonical cone over the final ω^{op}-chain of F by induction: $\alpha_0 : A \to 1$ is uniquely determined and $\alpha_{n+1} = F\alpha_n\alpha : A \to F^{n+1}1$.

2.3 Applications

Theorem 2.6 accounts for most of the final coalgebra results in the literature. For example, it applies to every functor on Set which is built from constant sets, products, coproducts (disjoint unions), and functions from fixed sets. As an example, and to point out some of the directions that coalgebra pursues, I will mention the final coalgebra of functors such as $FX = A \times X$, where A is a fixed set. This is the set A^∞ of *streams over* A; concretely, if can be taken to be functions from the natural numbers into A. The important data is the final coalgebra *structure* $A^\infty \to A \times A^\infty$. In the concrete case that I mentioned, this is $f \mapsto (f(0, g)$, where $g(n) = f(n + 1)$. However, once one has A^∞ and its structure, one can for many purposes forget the details of the construction and instead simply use the fact that we have a final coalgebra.

As a special case, consider the case when A is the set $\mathbb{R}$ of real numbers. In this case, we have an especially nice representation of the final coalgebra: it is the set An of *real analytic functions*, see [PE]. Here one considers $f : \mathbb{R} \to \mathbb{R}$ such that for every n there is an open interval around 0 such that $f'(x)$ is defined at each point in the interval, and f agrees with its Taylor series there. The coalgebra structure $\alpha : An \to \mathbb{R} \times An$ is given by

$$\phi(x) \mapsto (\phi(0), \phi'(x))$$

Here is an example of a coalgebra: $(\{a, b, c, d\}, f)$, where

$$f(a) = (0, b) \quad f(b) = (1, c) \quad f(c) = (0, d) \quad f(d) = (-1, a)$$

By "corecursion", that is, by finality, there is a unique coalgebra morphism

$$\phi : (\{a, b, c, d\}, f) \to (An, \alpha).$$

It is easy to check that

$$\phi(a) = \sin x \quad \phi(b) = \cos x \quad \phi(c) = -\sin x \quad \phi(c) = -\cos x.$$

That is, the assignment ϕ satisfies the condition of a coalgebra morphism, and by finality it is the unique such.

2.4 Measurable spaces

Let Meas be the category of measurable spaces and measurable functions. We consider this category and also the subcategory SB of *standard Borel spaces.* These are the measurable spaces (M, Σ) such that for some Polish topology (i.e., a topology generated by a complete separable metric) $\mathcal{T}$ on M, Σ is the family of Borel sets generated by $\mathcal{T}$. (For background here, see Kechris [Ke].

A related functor is the *subprobability measure functor* S : Meas $\to$ Meas. A *subprobability measure* on a space M is a σ-additive function μ from the Borel subsets of M to $[0, 1]$. We do not require that $\mu(M) = 1$. The subprobability measures on a space themselves form a measurable space with the σ-algebra defined as for Δ. It is known that there is a metric on $S(M)$ making $S(M)$ into a standard Borel space; see e.g., Doberkat [Dob],

The category SB is closed under countable coproducts and countable limits in Meas (see [Sch]). Moreover, it is closed under the functors S and Δ.

PROPOSITION 2.8. *S preserves the surjectivity of morphisms in* SB. *(See Kechris [Ke], Proposition 1.101.)*

THEOREM 2.9 (VIGLIZZO [Vi]). Δ : Meas $\to$ Meas *does not preserve limits of ω^{op}-chains.*

THEOREM 2.10. *The functor* Δ : SB $\to$ SB *preserves the limit of its final ω^{op}-chain.*

This result follows from the celebrated Kolmogorov Extension Theorem, stated below:

THEOREM 2.11 (KOLMOGOROV). *Let*

$$X_0 \xleftarrow{f_0} X_1 \xleftarrow{f_1} X_2 \qquad \cdots$$

be an ω^{op}-chain in SB, *and assume in addition that each f_n is surjective. Let $X = \lim X_n$, and let $\pi_n : X \to X_n$ be the projection. Let $\mu_n \in \Delta X_n$ be Borel measures such that $\Delta f_n(\mu_{n+1}) = \mu_n$ for all n. Then there is a unique $\mu \in \Delta X$ so that for all n, $\Delta\pi_n(\mu) = \mu_n$.*

In the form we have stated it, Theorem 2.10 was used in Schubert [Sch] to prove that a large collection of functors on SB have final coalgebras. These include the constant functors, Δ and S, and functors built from these by composition, product and coproduct. At the heart of things are the general category theoretic result (Theorem 2.6, the dual of Theorem 2.5) and some specific results on standard Borel spaces, such as Kolmogorov's Theorem (and a few other special facts).

Heifetz [He] carries out a parallel development, this time for regular measures on a Hausdorff space.

Similarly, den Hartog and de Vink [HV] study complete ultrametric spaces and non-expanding maps. They show that the functor $Meas$ is locally non-expanding. where the functor takes a space to its compactly-supported Borel probability measures, using the following metric:

$$d(\mu, \nu) = \inf\{\varepsilon : (\forall x)\mu(B(x, \varepsilon) = \nu(B(x, \varepsilon))\}$$

The functor works on morphisms in the obvious way.

Situations where Theorem 2.6 does not apply.

Having pointed out the uses of Theorem 2.6, we must mention that there are important settings in which it does not apply. Most notable of these is the case of the finite power set functor $\mathcal{P}_f$ on Set (cf. [AK79, W]) and the case of functors such as Δ on measurable spaces. Concerning the latter example, the relevant result is that all functors on Meas built from constants, Δ, products, and coproducts have final coalgebras. The first proof of this is in [MV2], building on work in Heifetz and Samet [HS]. Both of these used a variant of coalgebraic modal logic as we shall see it just below. A different approach to the problem of universal Harsanyi spaces may be found in Pintér [Pi]. Here the work goes via a new extension result along the lines of Kolmogorov's Theorem, but more applicable to the purely measure-theoretic setting. It is not yet known whether Pintér's Theorem shows that all functors on Meas built from constants, Δ, products, and coproducts have final coalgebras.

3. COALGEBRAIC GENERALIZATION OF MODAL LOGICS

Up until now in this paper, the focus has been on the construction of final coalgebras as a way to understand the literature on universal Harsanyi type spaces. At this point, I want to shift the focus to a different matter, that of *coalgebraic generalizations of the Kripke semantics of modal logic.* There are several reasons for doing this, including another connection to universal type spaces. Before this, let me mention two reconstructions of modal logic.

Let (A, α) be a coalgebra for F. We have maps $\alpha_n : A \to F^n 1$ as follows: α_0 is by finality, and $\alpha_{n+1} = F\alpha_n\alpha$. This gives a cone over the final sequence: an easy induction shows the required fact that for all n, $\alpha_n = F^n!\alpha_{n+1}$. We have a unique $\alpha^\dagger : A \to L$ such that for all n, $\alpha_n = l_n\alpha^\dagger$. We call the family α_n *the cone over the final sequence determined by* (A, α), and $\alpha^\dagger$ is the *connecting morphism.*

LEMMA 3.1 (MOSS [Mo$_2$]). $\alpha^\dagger$ *is the unique* coalgebra-to-algebra *morphism, the unique morphism such that the diagram below commutes:*

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & FA \\ {\scriptstyle \alpha^\dagger}\downarrow & & \downarrow{\scriptstyle F\alpha^\dagger} \\ L & \xleftarrow[m]{} & FL \end{array}$$

We shall use an easy consequence of Lemma 3.1.

LEMMA 3.2. *Let (A, α) and (B, β) be coalgebras for F and let $\chi : A \to B$ be a coalgebra morphism. Let $\alpha^\dagger$ and $\beta^\dagger$ be the connecting morphisms for the cones over the final sequence from these coalgebras. Then $\beta^\dagger\chi = \alpha^\dagger$.*

Let L^* be the subset of L consisting of the functions $g \in L$ such that for some coalgebra (A, α) and some $a \in A$, $g = \alpha^\dagger(a)$. In words, we take the union of all images of all of the connecting maps from coalgebras to the limit L. Then we have an inclusion $i : L^* \to L$. Also, the connecting map $\alpha^\dagger : A \to L$ factors as ij_A for some (unique) $j_A : A \to L^*$. That is, each point in the image of j_A does belong to L^*. Our main result in this section, Theorem 3.5, shows that L^* is the carrier of a final coalgebra structure. The main point is to define the structure map on L^*, and this is where the injectivity of m is critical. The details are in Lemma 3.3.

LEMMA 3.3. *Assume that $m : FL \to L$ is injective.*

1. *For each $g \in L^*$, there is a unique $x \in FL^*$ such that $(mFi)x = i(g)$.*
2. *There is a function $l^* : L^* \to FL^*$ so that the square on the right below commutes:*

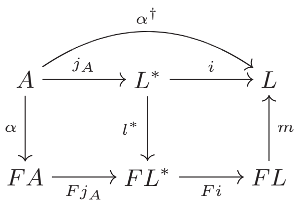

(5)

3. *Finally, for any coalgebra (A, α) the square on the left commutes, as does the top (by definition).*

PROOF. Let $g \in L^*$. Let (B, β) and $b \in B$ be such that $i(g) = \beta^\dagger(b)$. Let $j_B : B \to L^*$ be such that $ij_B = \beta^\dagger$. Let $x = (Fj_B\beta)b$. Then

$$\begin{array}{rcll} (mFi)x & = & (mFiFj_B\beta)b & \text{by definition of } x \\ & = & (mF\beta^\dagger\beta)b & \text{since } \beta^\dagger = ij_B \\ & = & \beta^\dagger(b) & \text{by Lemma 3.1} \\ & = & i(g) & \text{by definition of } b \end{array}$$

For the uniqueness of x, note that we may assume that L^* is non-empty. (For if L^* were empty, then there would be no coalgebras for F whatsoever, except for the one with empty carrier. Then F would be the constant functor with value $\emptyset$. But then there would be just one coalgebra for F, and this would be the final coalgebra.) And now, recall that every functor on Set preserves injectivity of maps with non-empty domain. So we see that Fi is injective. So is m, by the assumption in this result. We thus see that mFi is injective; hence x is unique.

The uniqueness defines l^*. (That is, for each $g \in L^*$ we used some coalgebra (B, β) or other to define $l^*(g)$, but the definition was independent of the coalgebra used.) The commutativity of the right square follows. For the one on the left, let $a \in A$. We show that $l^*(j_A(a)) = (Fj_A\alpha)a$ by showing that the latter object satisfies the condition which uniquely defines the former. The verification is:

$$(mFiFj_A\alpha)a \quad = \quad (mF\alpha^\dagger\alpha)a \quad = \quad \alpha^\dagger(a) \quad = \quad i(j_A).$$

(We are using Lemma 3.1.) $\square$

Once again, by Lemma 3.3 we have a map $l^* : L^* \to FL^*$. Being a coalgebra structure for F, l^* gives rise to a connecting map $(l^*)^\dagger : L^* \to L$.

LEMMA 3.4. $(l^*)^\dagger = i$.

PROOF. By Lemma 3.3, i is a coalgebra-to-algebra morphism: $mFil^* = i$. Hence we have $(l^*)^\dagger = i$ by the uniqueness part of Lemma 3.1. $\square$

The previous lemma is the key property of l^*. For readers interested in modal logic, it calls to mind the parallel lemma for the theory map from canonical model into itself. The surjectivity is tantamount to the completeness theorem for the logic.

THEOREM 3.5. *(L^*, l^*) is a final coalgebra for F: for all coalgebras A, $j_A : A \to L^*$ is the unique coalgebra morphism.*

PROOF. Let (A, α) be a coalgebra, and let $\alpha^\dagger : A \to L$ be its connecting map. Let $j_A : A \to L^*$ be as above, so that $ij_A = \alpha^\dagger$. By Lemma 3.3, $j_A : A \to L^*$ is a morphism of coalgebras.

It remains to prove the uniqueness of j_A. Let $k : A \to L^*$ be a coalgebra morphism. By Lemma 3.2, $\alpha^\dagger = (l^*)^\dagger k$, and by the last lemma, this is ik. Hence $ik = \alpha^\dagger = ij_A$. Since i is injective, $k = j_A$. $\square$

There are proof systems for this logical system, and currently much is known about it.

Why the logic is a generalization of standard modal logic.

the connection of maximal consistent sets in K to the elements of $\mathcal{P}^\omega 1$. This connection is made in the rest of this section.

Let $\mathcal{L}$ be the set of modal sentences. There are maps $c_n : \mathcal{P}_f^n 1 \to \mathcal{L}$ given by $c_0(*) = \top$ (some tautology), and for $S \in \mathcal{P}_f(\mathcal{P}_f^n 1)$,

$$c_{n+1}(S) \quad = \quad \bigwedge_{x \in S} \Diamond c_n(x) \ \wedge \ \Box \bigvee_{x \in S} c_n(x)$$

(All conjunctions shown are finite, of course. Empty conjunctions count as $\top$ and empty disjunctions as $\bot$.) Incidentally, $c_{n+1}(S)$ is usually written $\nabla(S)$ in the literature on coalgebraic modal logic.

Let C_n be the image of c_n. This set is sometimes called the set of *canonical sentences of height* n, or the set of *Fine normal forms of height* n.

LEMMA 3.6 (FINE [F]). *For every pointed graph (G, g), there is exactly one canonical sentence ϕ of height n such that $g \models \phi$ in G. Moreover, if ψ is any modal sentence of height $\leq n$, then either $\vdash \phi \to \psi$ or $\vdash \phi \to \neg\psi$ in K.*

This lemma gives a representation of canonical cones from graphs to the the final sequence of $\mathcal{P}_f$. Given $(A, \to)$, the map $\alpha_n : A \to \mathcal{P}^n 1$ takes a to the unique canonical sentence of height n satisfied by a in A.

Applying this to the canonical graph $(C, \to)$ itself, and using the Truth Lemma, we see that the canonical map $\gamma_n : C \to \mathcal{P}^n 1$ is given by

$$\gamma_n(u) \quad = \quad \text{the unique canonical sentence of height } n \text{ in } u. \tag{6}$$

THEOREM 3.7. *The cone $\gamma_n : C \to \mathcal{P}^n 1$ is a limit cone.*

PROOF. Let ϕ_n be a sequence of modal formulas with the property that for all n, $\mathcal{P}^n!(\phi_{n+1}) = \phi_n$. We need to show that there is a unique modal theory u containing all ϕ_n. For the existence, we need only check that $\{\phi_n : n \in \omega\}$ is consistent. (Once this is established, then by Zorn's Lemma, this set will have a maximal consistent superset.) An easy induction on n shows that

$$\vdash \psi \to \mathcal{P}^n!(\psi)$$

for all $\psi \in \mathcal{P}^{n+1}1$ (see [Mo$_1$]). It follows from this general observation and the finiteness of proofs that $\{\phi_n : n \in \omega\}$ is consistent.

For the uniqueness, suppose that u and v contain ϕ_n for all n. Then we claim that $u = v$. To see this, let ψ be any modal formula. Let n be the modal height of ψ. Then by Lemma 3.6, either $\vdash \phi_n \to \psi$ or $\vdash \phi_n \to \neg\psi$. Without loss of generality, assume that the first alternative holds. Then since u and v are closed under deduction (this follows from maximal consistency), they both contain ψ. In this way, $u = v$. $\square$

COROLLARY 3.8. *$\mathcal{P}^\omega 1$ may be described as the set C of maximal consistent modal theories, with maps $\gamma_n C \to \mathcal{P}^n 1$ as in (6).*

Conclusion.

The upshot of this discussion is that $\mathcal{P}_f^\omega(1)$ is isomorphic to the *canonical model* and $\nu\mathcal{P}_f$ to the submodel consisting of *theories of pointed finitely branching models*.

For more on the normal forms, including the proof of Lemma 3.6 and applications to modal completeness for standard modal logics, see Fine [F] and also Moss [Mo_1].

3.1 Addendum: generalizing modal logic via predicate liftings

In Section 3, we saw one way to generalize the Kripke semantics of modal logic to functors on Set. The generalization was not a perfect one in the sense that one obtains a fragment of standard modal logic and not the whole thing. From the standpoint of using the logic, there is much to be desired. So for this reason, it is interesting to study different generalizations of modal logic. This is the topic of *coalgebraic modal logic*. One has a generalization of many modal logics, including standard modal logic K, conditional logic, probabilistic modal logic, and various combinations of these. Moreover, one has complexity results for the satisfiability problem which are *generic* (that is, independent of the particular logic). This area is quite developed, and in fact has given new results on existing logics, many of which could be of interest to the TARK community. Some references here include C. Cîrstea et al [CKPSV], Schröder and Pattinson [SP], and Kurz [K].

4. FOR FURTHER READING

Introductory material on coalgebra may be found in papers such as Gumm [GU], Jacobs and Rutten [JR], Pattinson [P_2], and Rutten [R]. An extensive treatment of final coalgebra results, including those relevant to type spaces, is promised in Adámek et al [AMM].

5. ACKNOWLEDGMENTS

My thanks to Eric Pacuit and the TARK Program Committee for their invitation to present this tutorial.

6. REFERENCES

[A74] J. Adámek, Free algebras and automata realizations in the language of categories, *Comment.Math. Univ. Carolinae* 14 (1974), 589–602.

[AK79] J. Adámek and V. Koubek, On the least fixed point of a set functor, *J. Comput. Syst. Sci.* 19 (1979), 163–178.

[AMM] J. Adámek, S. Milius, and L. S. Moss, Initial Algebras and terminal coalgebras: a survey. Ms., 2011.

[Ba] M. Barr, Terminal coalgebras in well-founded set theory, *Theoret. Comput. Sci.* 114 (1993) no. 2, 299–315. Additions and corrections in *Theoret. Comput. Sci.* 124 (1994), no. 1, 189–192.

[BE] W. Böge, T. Eisele, On solutions of Bayesian games, *Internat. J. Game Theory* 8 (4) (1979) 193–215.

[BD] A. Brandenburger and E. Dekel, Hierarchies of beliefs and common knowledge. *J. Economic Theory*, 59 (1993), 189–198.

[CKPSV] C. Cîrstea, A. Kurz, D. Pattinson, L. Schröder, Y. Venema (2011), Modal logics are coalgebraic. *Computer Journal*, 54 (1), 31-41.

[Dob] E.-E. Doberkat, *Stochastic Relations: Foundations for Markov Transition Systems.* Chapman Hall & CRC Press, Boca Raton, 2007.

[F] K. Fine, Normal forms in modal logic. *Notre Dame J. Formal Logic* 16 (1975), 229–237.

[GU] H. P. Gumm, Universelle Coalgebra, an appendix of T. Ihringer, *Allgemeine Algebra*, Helderman Verlag, Berlin 2003 (in German)

[H] J. C. Harsányi, Games with incomplete information played by "Bayesian" players. I. The basic model, *Management Sci.* 14 (1967), 159–182.

[HV] J. den Hartog and de Vink, Building metric strucutres with the *Meas*-functor, Liber Americorum Jaco de Bakker, F. de Boer, M. van der Heijden, P. Klint and J. Rutten (eds.), CWI, Amsterdam, 2002, 93–108.

[He] A. Heifetz, The bayesian formulation of incomplete information–the non-compact case. *International J. Game Theory*, 21 (1993), 329–338.

[HS] A. Heifetz and D. Samet, Topology-free typology of beliefs. *J. Economic Theory*, 82 (1998), 324–341.

[JR] B. Jacobs and J. Rutten, A tutorial on (Co)algebras and (co)induction. *EATCS Bulletin* 62, 1997, p.222-259.

[Ke] A. Kechris, *Classical Descriptive Set Theory*, Graduate Texts in Mathematics 156, Springer-Verlag, New York, 1995.

[K] A. Kurz, Coalgebras and their logics. Logic Column of the *SIGACT News* 37 (2), pp. 57-77, 2006.

[L] J. Lambek, A Fixpoint Theorem for Complete Categories, *Math. Z.* 103 (1968), 151–161.

[MZ] J. F. Mertens and S. Zamir, Formulations of bayesian analysis for games with incomplete information, *International J.Game Theory*, 14 (1985), 1–29.

[Me] M. Meier, Universal knowledge-belief structures. *Games and Economic Behavior* 62 (2008) 53-66.

[Mo_1] L. S. Moss, Finite models constructed from canonical formulas, *J. Philos. Logic* 36 (2007), no. 6, 605–640.

[Mo_2] L. Moss, A note on expressive coalgebraic logics for finitary set functors, *Journal of Logic and Computation*, to appear, 12 pp.

[MV1] L. S. Moss and I. D. Viglizzo, Harsanyi type spaces and final coalgebras constructed from satisfied theories. *Electronic Notes in Theoretical Computer Science* 106 (2004) 279–295.

[MV2] L. S. Moss and I. D. Viglizzo, Final coalgebras for measurable spaces, *Inform. and Comput.* 204 (2006), no. 4, 610–636.

[P_1] D. Pattinson, Expressive logics for coalgebras via terminal sequence induction. *Notre Dame Journal of Formal Logic*, 45(1):19-33, 2004.

[P_2] D. Pattinson, An Introduction to the Theory of Coalgebras, course notes from NASSLLI 2003, available from `http://www.indiana.edu/~nasslli/program.html`.

[PE] D. Pavlović, M. H. Escardó, Calculus in coinductive form, In: Proc. Thirteenth Annual IEEE Symposium on Logic in Computer Science (Indianapolis, IN, 1998), 408–417, IEEE Computer Soc., Los Alamitos, CA, 1998.

[Pi] M. Pintér, The existence of an inverse limit of an inverse system of measure spaces - A purely measurable case, *Acta Mathematica Hungarica* 2010, 126, 65-77.

[R] J. Rutten, Universal coalgebra: a theory of systems, *Theoret. Comput. Sci.* 249(1), 2000, 3–80.

[SP] L. Schröder and D. Pattinson, Rank-1 modal logics are coalgebraic. *Journal of Logic and Computation*, 20 (5), 2010, 1113-1147.

[Sch] C. Schubert, Final coalgebras for measure-polynomial functors. In: Proc. Conference on Theory and Applications of Models of Computation (TAMC'2009), ShangSha, P.R. China, to appear in *Lecture Notes Comput. Sci.*

[Va] S. Vassilakis, Economic Data Types, Working Paper 252, Department of Economics, University of Pittsburgh, 1989.

[Vi] I. Viglizzo, *Coalgebras on Measurable Spaces.* Ph.D. thesis, Indiana University, Bloomington, 2005.

[W] J. Worrell, On the final sequence of a finitary set functor, *Theoret. Comput. Sci.* 338 (2005), no. 1–3, 184–199.