A Note on Expressive Coalgebraic Logics for Finitary Set Functors

Lawrence S. Moss
Department of Mathematics
Indiana University
Bloomington, IN 47405 USA

September 17, 2008

Abstract

This paper has two purposes. The first is to present a final coalgebra construction for finitary endofunctors on *Set* that uses a certain subset L^* of the limit L of the first ω terms in the final sequence. L^* is the set of points in L which arise from all coalgebras using their canonical morphisms into L, and it was used earlier for different purposes in Kurz and Pattinson [5]. Viglizzo in [11] showed that the same set L^* carried a final coalgebra structure for functors in a certain inductively defined family. Our first goal is to generalize this to all finitary endofunctors; the result is implicit in Worrell [12]. The second goal is to use the final coalgebra construction to propose coalgebraic logics similar to those in [6] but for all finitary endofunctors F on *Set*. This time one can dispense with all conditions on F, construct a logical language $\mathcal{L}_F$ directly from it, and prove that two points in a coalgebra satisfy the same sentences of $\mathcal{L}_F$ iff they are identified by the final coalgebra morphism. The language $\mathcal{L}_F$ is very spare, having no boolean connectives. This work on $\mathcal{L}_F$ is thus a re-working of coalgebraic logic for finitary functors on sets.

1 Introduction

The introduction of coalgebraic logics had two sources. The earliest was an observation by Jan Rutten in his now-classic paper [9]: Kripke frames and models as they appear in modal logic are coalgebras of certain endofunctors on *Set*. At the same time, logical systems which had a "modal feel" such as the Hennessy-Milner logic coming from concurrency theory were also interpreted on coalgebras. Jon Barwise and I were interested in connections to hypersets, and the second source of coalgebraic logics was our result in *Vicious Circles* [3] that worlds in Kripke models may be characterized up to bisimulation by single sentences in *infinitary* modal logic. A version of this strong Hennessy-Milner-type property had actually been obtained earlier for modal logic, in van Benthem and Bergstra [4]. It was then natural to try to see the sentences in various modal logics as "observations" in some general sense, and then to try to use the concepts in coalgebra as a way to define logics from functors. The first work in this direction was another chapter in [3]. But that book shunned category theory and so worked with operators on sets instead of functors.

Having categorical concepts improved matters. The goal of our paper [6] was to show the following: To each endofunctor F on Set satisfying some weak conditions, one should obtain a logical language $\mathcal{L}_F$ which would be interpreted on the coalgebras of F. The logics should be appropriate for coalgebraic reasoning in the sense that for all coalgebras (A, α) for F and all $a, b \in A$,

$$a \text{ and } b \text{ are bisimilar in } A \text{ iff they satisfy the same sentences } \phi \in \mathcal{L}_F \text{ in } (A, \alpha). \tag{1}$$

This condition is called *expressivity* for the language $\mathcal{L}_F$. Moreover, taking F to be a functor associated with semantic models that already existed should result in a language $\mathcal{L}_F$ which had already been independently studied, or at least looked natural. Also, because the results in [3] for the power set functor pointed to *infinitary* modal logic, it was also of interest to look at the strong version of (1), namely *strong expressivity*: for all (A, α) and all $a \in A$ there is a sentence $\phi \in \mathcal{L}_F$ such that for all $b \in A$:

$$a \text{ and } b \text{ are bisimilar in } A \text{ iff } b \text{ satisfies } \phi_a. \tag{2}$$

I'll call the approach of [6] the the *direct* approach to generalizing modal logics to coalgebra. The results coming from this approach were mixed. On the one hand, for functors F meeting three conditions, one could formulate a language $\mathcal{L}_F$ and obtain the strong expressivity result formulated in (2). We'll spell out those conditions in due course. But on the other, the languages coming from the direct approach were not generalizations of standard systems such as classical modal logic, rather they generalized *fragments* of them. In fact, most of the papers that came after [6] complained that the syntax of $\mathcal{L}_F$ was not much of a syntax at all. This led to the main body of work in coalgebraic modal logic, where one begins not just with a functor but with a family of predicate liftings for it. (See Pattinson [7] for the first paper and, e.g., Schröder [10] for a more recent on issues related to that of the current paper.) From these, one then defines $\mathcal{L}_F$; this time the syntax looks much more familiar in the concrete cases of interest. I'll call this the *predicate lifting* approach. Then one of the main issues in the resulting large and impressive literature has again been to establish expressivity in the form of (1). The functors again had some requirements, different from the ones in [6]. Of special note here was the requirement that F be *finitary*; this made it possible to have $\mathcal{L}_F$ be a countable set.

So at this point, it will be useful to mention the actual conditions on the functors F in the direct approach. They were: (a) F is standard (it preserves inclusions of sets); (b) F preserves weak pullbacks; and (c) F is uniform. It would take us too far afield to define uniformity. It certainly seemed from the presentation in [6] that conditions like uniformity would be needed to show the expressivity of $\mathcal{L}_F$, and the weak pullback preservation was even more important: the basic lemmas in [6] seemed to require it at every step.

In addition, when one compared the expressivity results for the direct approach to those in the predicate lifting approach, the following question was natural: if one took the direct logics and only looked at some "small" fragment of them, say the fragment obtained by restricting the infinitary conjunctions to be finite, and if F were finitary to begin with, would $\mathcal{L}_F$ be expressive?

The main result in Section 4 answers this question affirmatively. Actually, the work here points to a different way to execute the direct approach in the first place. If one works with a finitary endofunctor on *Set*, it is possible to define $\mathcal{L}_F$ and give its semantics in a very straightforward and simple way, and then to prove the expressivity in a few steps. It turns out that conjunction is *not needed* in the language. Further, one does not need *any* requirements on F except that it be finitary. However, when dealing with functors that need not preserve weak pullbacks, the appropriate statement of expressivity would refer not to bisimilarity but rather to *behavioral equivalence*: a and b have the same image in a final coalgebra for F iff they satisfy the same sentences $\phi \in \mathcal{L}_F$ in (A, α). We define $\mathcal{L}_F$ and prove expressivity in this sense in Section 4.

This modest paper has two points. We have been discussing the contribution to the study of coalgebraic versions of modal logic. The second point is a representation theorem for the final coalgebra of certain endofunctors on *Set*, namely the finitary ones. It turns out that this second point is used in establishing the expressivity result for $\mathcal{L}_F$. It may well be of independent interest, and we discuss it in the next two sections.

2 Final coalgebras via the final sequence

Let F be an endofunctor on a category C. Perhaps the most common way to obtain a final coalgebra for F is to consider the *final sequence*, the ω^{op}-chain

$$1 \xleftarrow{!} F1 \xleftarrow{F!} F^2 1 \qquad \cdots \qquad F^n 1 \xleftarrow{F^n!} F^{n+1} 1 \qquad \cdots \tag{3}$$

Here and below, 1 is *final object* in the category, and $! : F1 \to 1$ is by finality.

A *cone over the final sequence* is an object A together with a family of morphisms $a_n : A \to F^n 1$ such that $a_n = F^n! \cdot a_{n+1}$ for all n. A *limit of the final sequence* is a cone over it $(L, l_n : L \to F^n 1)$ with the universal property that if $(A, a_n : A \to F^n 1)$ is any cone over the final sequence, then there is a unique *connecting morphism* $f : A \to L$ such that for all n, $a_n = l_n \cdot f$.

Assume that the limit above exists. There always is a morphism $m : FL \to L$, the connecting morphism from the cone $(FL, a_n : FL \to F^n 1)$, where a_0 is by finality, and $a_{n+1} = Fl_n$. (To check that this is indeed a cone, note that our limit cone has $l_{n-1} = F^{n-1}! \cdot l_n$, and so $Fl_{n-1} = F^n! \cdot Fl_n$.) So we have a unique $m : FL \to L$ such that for all n, $l_n \cdot m = a_n$. That is, $l_0 \cdot m$ is the map from FL to 1, and for all n,

$$l_{n+1} \cdot m \quad = \quad Fl_n. \tag{4}$$

Assuming that the limit L of the chain (3) exists and that F preserves it, then as shown in Adámek [1], m is an isomorphism, and as a result, (L, m^{-1}) is a final coalgebra of F. This result is well-known. We shall not need details of the proof. Instead, we are concerned with the situation where m is *not* an isomorphism but is only injective.

The "combinatory" fact in Lemma 2.1 below is due to James Worrell and is crucial for our work. Actually, Worrell weakens the hypothesis from "finitary" to "κ-accessible" functors, where κ is some infinite cardinal.

Lemma 2.1 (Worrell [12]) *Let $F : Set \to Set$ be finitary. Then $m : FL \to L$ is injective.*

We wish to prove this, not only to make this paper self-contained but also because it is the only place where we explicitly call on the notion of a finitary functor. Recall that in Set, L may be taken to be the set of functions g with domain ω such that for all n, $g(n) \in F^n 1$ and such that for all n $F^n!(g(n+1)) = g(n)$. Moreover, we take $l_n(g) = g(n)$ for all n.

We now sketch the proof of Lemma 2.1. (We are actually proving a special case of the result from [12].) Let $x, y \in FL$, and assume that $mx = my$. By the assumption that F is finitary, there is a finite set S, a map $f : S \to L$ and $x', y' \in FS$ such that $x = (Ff)x'$ and $y = (Ff)y'$.

The image set $f[S]$ of f is also finite. We claim that for some k, $l_k \cdot f$ is injective. Here is the reasoning. Let a and b be different members of $f[S]$. We think of a and b as functions on the natural numbers, and each l_n works by applying these to the number n. Since $a \neq b$, we must have some $n = n(a, b)$ such that $l_n(a) = a(n) \neq b(n) = l_n(b)$. And for all $m > n$, we also have $l_m(a) \neq l_m(b)$. The set

$$\mathcal{S} \quad = \quad \{(a, b) \in f[S] \times f[S] : a \neq b\}$$

is also finite, and so $\max\{n(a, b) : (a, b) \in \mathcal{S}\}$ is a natural number, say k. For this k, $l_k \cdot f$ is injective.

Functors on Set preserve injectivity of maps with non-empty domain, and so $F(l_k \cdot f)$ is injective. By (4), this is the same as $l_{k+1} \cdot m \cdot Ff$. Hence $m \cdot Ff$ is injective. We have $(m \cdot Ff)x' = mx = my = (m \cdot Ff)y'$. So $x' = y'$, and thus $x = y$.

The following is perhaps the main point in [12]: Assuming only that $m : FL \to L$ is injective, the limit of the next ω steps in the final sequence

$$L \xleftarrow{m} FL \xleftarrow{F^2 m} F^2 L \qquad \cdots \qquad F^n L \xleftarrow{F^{n+1} m} F^{n+1} L \qquad \cdots \tag{5}$$

is the carrier of a final F-coalgebra structure.

2.1 Cones derived over the final sequences via coalgebras

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

Lemma 2.2 *$\alpha^\dagger$ is the unique* coalgebra-to-algebra *morphism, the unique morphism such that the diagram below commutes:*

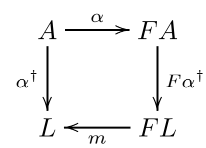

Proof We check first that the diagram commutes. For this, we show that $m \cdot F\alpha^\dagger \cdot \alpha$ satisfies the condition uniquely defining $\alpha^\dagger$. For all n,

$$l_{n+1} \cdot m \cdot F\alpha^\dagger \cdot \alpha \quad = \quad Fl_n \cdot F\alpha^\dagger \cdot \alpha \quad = \quad F\alpha_n \cdot \alpha \quad = \quad \alpha_{n+1}.$$

Also, $l_0 \cdot m \cdot F\alpha^\dagger \cdot \alpha = \alpha_0$ by finality.

For the uniqueness, suppose that $\beta = m \cdot F\beta \cdot \alpha$. We check that for all n, $\alpha_n = l_n \cdot \beta$. For $n = 0$, this is clear. Assuming that $\alpha_n = l_n \cdot \beta$, note that

$$\begin{array}{rcll} \alpha_{n+1} & = & F\alpha_n \cdot \alpha & \text{by definition of } \alpha_n \\ & = & F(l_n \cdot \beta) \cdot \alpha & \text{by induction hypothesis} \\ & = & l_{n+1} \cdot m \cdot F\beta \cdot \alpha & \text{by (4)} \\ & = & l_{n+1} \cdot \beta & \text{by assumption} \end{array}$$

This completes the proof. ⊣

We shall use an easy consequence of Lemma 2.2.

Lemma 2.3 *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 $\beta^\dagger \cdot \chi = \alpha^\dagger$.*

3 Final coalgebras for finitary functors

We are still concerned with finitary endofunctors on *Set*. We define a special subset L^* of the limit L, and show that L^* is the carrier of a final coalgebra structure.

The carrier set 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 $i \cdot j_A$ 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 3.3, 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 3.1.

Before we turn to that result, we should mention the origin of the definition of L^*. As far as I know, its first use was in Kurz and Pattinson [5], a paper connected with the second topic of this one, but not with a final coalgebra construction. Viglizzo [11] used what we are calling L^* to construct the final coalgebra of the *probabilistic Kripke functors* on *Set*; these are the smallest collection containing the identity, constants, the finite power set, and the discrete distribution functor and closed under product, coproduct, and function space from a fixed set. Now all of the functors in this class are finitary, and so in a sense our observation in this section is merely a confirmation that the work in [11] generalizes to the setting of finitary functors.

Lemma 3.1 *Assume that $m : FL \to L$ is injective.*

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

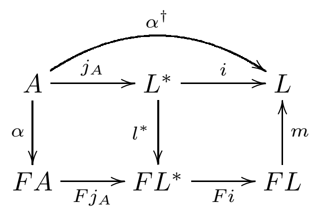

(6)

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 $i \cdot j_B = \beta^\dagger$. Let $x = (Fj_B \cdot \beta)b$. Then

$$\begin{array}{rcll} (m \cdot Fi)x & = & (m \cdot Fi \cdot Fj_B \cdot \beta)b & \text{by definition of } x \\ & = & (m \cdot F\beta^\dagger \cdot \beta)b & \text{since } \beta^\dagger = i \cdot j_B \\ & = & \beta^\dagger(b) & \text{by Lemma 2.2} \\ & = & i(g) & \text{by definition of } b \end{array}$$

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 $m \cdot Fi$ 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 $l^*(j_A(a)) = (Fj_A \cdot \alpha)a$ by showing that the latter object satisfies the condition which uniquely defines the former. The verification is:

$$(m \cdot Fi \cdot Fj_A \cdot \alpha)a \quad = \quad (m \cdot F\alpha^\dagger \cdot \alpha)a \quad = \quad \alpha^\dagger(a) \quad = \quad i(j_A).$$

(We are using Lemma 2.2.) ⊣

Once again, by Lemma 3.1 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 3.2 $(l^*)^\dagger = i$.

Proof By Lemma 3.1, i is a coalgebra-to-algebra morphism: $m \cdot Fi \cdot l^* = i$. Hence we have $(l^*)^\dagger = i$ by the uniqueness part of Lemma 2.2. ⊣

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 3.3 *(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 $i \cdot j_A = \alpha^\dagger$. By Lemma 3.1, $j_A : A \to L^*$ 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 2.3, $\alpha^\dagger = (l^*)^\dagger \cdot k$, and by the last lemma, this is $i \cdot k$. Hence $i \cdot k = \alpha^\dagger = i \cdot j_A$. Since i is injective, $k = j_A$. ⊣

3.1 The original development

Referees suggested that the result of the last section it is implicit in Worrell [12]: one may take for the carrier of a final coalgebra of a finitary functor the union of all images of all coalgebras in the limit of the the first ω terms of the final sequence. To substantiate this, we shall re-prove this result in a different way, by calling on further results in Worrell's paper. To present things in a way which shortens the original argument, we also use Lemmas 2.2 and 2.3.

Lemma 3.4 *Assume again that $m : FL \to L$ is injective, and let l be such that $l \cdot m = \mathsf{id}_{FL}$. For every F-coalgebra (A, α), $\alpha^\dagger$ is a coalgebra morphism from (A, α) to (L, l). As a result, (L, l) is a* weakly final *coalgebra for F.*

Proof We use Lemma 2.2: $F\alpha^\dagger \cdot a = l \cdot m \cdot F\alpha^\dagger \cdot a = l \cdot \alpha^\dagger$. ⊣

This is Proposition 7 in Worrell [12]. Proposition 8 is that $l^\dagger$ is idempotent. We can prove this as follows: Apply Lemma 3.4 to $l : L \to FL$ to see that $l^\dagger$ is a coalgebra morphism from (L, l) to itself. Then by Lemma 2.3, $l^\dagger \cdot l^\dagger = l^\dagger$.

Define L^* to be the image of $l^\dagger$, and let $i : L^* \to L$ be the inclusion. So $i \cdot j_L = l^\dagger$, and by idempotence $j_L \cdot i = \mathsf{id}_{L^*}$. Note that $(m \cdot Fi) \cdot (Fj_L \cdot l) = m \cdot Fl^\dagger \cdot l = l^\dagger$ (we used Lemma 2.2), and $(Fj_L \cdot l) \cdot (m \cdot Fi) = \mathsf{id}_{FL^*}$. Let $l^* : L^* \to FL^*$ be $Fj_L \cdot l \cdot i$. The inverse of l^* is $j_L \cdot m \cdot Fi$. We look back at (6), or rather the special case of this when $a : A \to FA$ is $l : L \to FL$. Then with this choice of l^*, the two squares commute. We omit the calculations. The point is that we have *defined* L^* and l^* in such a way that l^* is a coalgebra morphism, and also so that i is a coalgebra-to-algebra morphism. Then the proof that l^* is a final coalgebra is just as before (see Lemma 3.2 and Theorem 3.3).

Recall that L^* was defined using idempotence of $l^\dagger$. The one remaining point in this development would be to show that L^* is just what we said it was in the previous section, the union of the images of all coalgebras inside L. This would mean that for all coalgebras (A, a), the final coalgebra map from A to L^* (it is j_A) satisfies $i \cdot j_A = \alpha^\dagger$. But j_A is a coalgebra morphism, and $i = (l^*)^\dagger$, and so this follows from Lemma 2.3.

A final remark It might also be interesting to view Theorem 3.3 from a different angle. Let F be any endofunctor on *Set* which has a final coalgebra. It follows from a theorem in Adámek and Koubek [2] that for every coalgebra (A, α) and all $a, b \in A$, a and b are identified by the final coalgebra morphism iff all of their images in the (*possibly transfinite*) final sequence of F agree. In this paper, we are only concerned with finitary endofunctors, and Theorem 3.3 gives a refinement: a and b are identified by the final coalgebra morphism iff all of their images in the *first* ω *steps* of final sequence of F agree. This point was made earlier by Dirk Pattinson, in Theorem 4.1 of [7], and (with a proof) in Theorem 3.1.11 of [8].

4 An expressive coalgebraic logic in the finitary setting

At this point, we return to the discussion of coalgebraic generalizations of modal logic from our Introduction.

Let F be a finitary functor on sets. Let $\mathcal{L} = \mathcal{L}_F$ be $\coprod_{n=0}^{\infty} F^n 1$, where we understand $F^0 1 = 1$, and take 1 to be a singleton set which we call $\{true\}$.

Here is how we interpret $\mathcal{L}$ on a coalgebra for F, say (A, α). We define *satisfaction relations* $S_n \subseteq A \times F^n 1$ by recursion on n. (We should write $S_n(A)$ or $S_{n,A}$ to indicate the dependence on the coalgebra involved. We do use this notation in Lemma 4.3 below, since there are two coalgebras involved. But outside of that result, we suppress the mention of the underlying structure.)

If $\phi \in F^n 1$ we sometimes say that a *satisfies* ϕ if $(a, \phi) \in S_n$, and write $a \models \phi$.

We start with $S_0 = A \times 1$, so that all $a \in A$ satisfy *true*.
Given $S_n \subseteq A \times F^n 1$, we have $\pi_1 : S_n \to A$ and $\pi_2 : S_n \to F^n 1$. Let

$$S_{n+1} \quad = \quad \{(a, \phi) : (\exists w \in F(S_n))\ (F\pi_1)w = \alpha(a) \text{ and } (F\pi_2)w = \phi\}.$$

Recall also that we have $\alpha_n : A \to F^n 1$; α_0 is by finality, and $\alpha_{n+1} = F\alpha_n \cdot \alpha$.

Lemma 4.1 *Consider each α_n as a subset of $A \times F^n 1$. For all n, $\alpha_n \subseteq S_n$.*

Proof By induction on n. For $n = 0$, we have $\alpha_0 = A \times 1 = S_0$. Assume that $\alpha_n \subseteq S_n$. Let $f : A \to S_n$ be given by $f(a) = (a, \alpha_n(a))$. Then $\pi_1 \cdot f = \mathsf{id}_A$, and $\pi_2 \cdot f = \alpha_n$. For each $a \in A$, let $w_a = (Ff \cdot \alpha)a$. Then

$$\begin{array}{lllll} (F\pi_1)w_a & = & (F\mathsf{id}_A \cdot \alpha)a & = & \alpha(a) \\ (F\pi_2)w_a & = & (F\alpha_n \cdot \alpha)a & = & \alpha_{n+1}(a) \end{array}$$

This completes the proof. ⊣

Lemma 4.2 *For all n, $\alpha_n = S_n$.*

Proof By induction on n. We again have the case $n = 0$ by definition. Assume that $\alpha_n = S_n$. Suppose that $(a, \phi) \in S_{n+1}$, and let $v \in F(S_n)$ be such that $(F\pi_1)v = \alpha(a)$ and $(F\pi_2)v = \psi$. By our last result, let $w \in F(S_n)$ be such that $(F\pi_1)w = \alpha(a)$ and $(F\pi_2)w = \alpha_{n+1}(a)$. Recall that π_1 and π_2 are the projections from S_n to A and to $F^n 1$. But $S_n = \alpha_n$ by our induction hypothesis. And thus $\pi_2 = \alpha_n \cdot \pi_1$. (This is true for projections from any function whatsoever, when considered as a relation between its domain and codomain.) It follows that

$$\psi = (F\pi_2)v = F\alpha_n(F\pi_1 v) = F\alpha_n(\alpha(a)) = F\alpha_n(F\pi_1 w) = (F\pi_2)w = \alpha_{n+1}(a).$$

In other words, the only element of $F^{n+1}1$ satisfied by a is $\alpha_{n+1}(a)$. ⊣

One can also prove Lemma 4.2 directly by induction on n, using the relation lifting $\overline{F}$ of F. Without assumptions on F, $\overline{F}$ need not be an endofunctor on the category of relations, but this is not an obstacle to its use here. Observe that S_{n+1} is by definition the relational composition $Gr(\alpha) \circ \overline{F}(S_n)$. (We are using $Gr(f)$ for the graph relation of a function, and $\circ$ denotes relational composition.) It is a general fact about relation lifting that $\overline{F}(Gr(f)) = Gr(Ff)$. An easy induction then shows that $S_n = Gr(\alpha_n)$. This is Lemma 4.2.

Lemma 4.3 *Coalgebra morphisms preserve the semantics. That is, if $f : (A, \alpha) \to (B, \beta)$ is a coalgebra morphism and $\phi \in \mathcal{L}$, then $a \models \phi$ iff $fa \models \phi$.*

Proof An easy induction shows that for each n, the diagram below commutes:

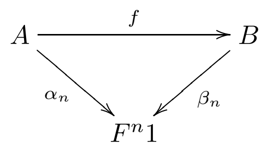

In more detail: the commutativity is clear for $n = 0$. Assuming it for n, we have

$$\alpha_{n+1} \quad = \quad F\alpha_n \cdot \alpha \quad = \quad F\beta_n \cdot Ff \cdot \alpha \quad = \quad F\beta_n \cdot \beta \cdot f \quad = \quad \beta_{n+1} \cdot f \ .$$

Let $\phi \in F^n 1$ and $a \in A$. Then we have the following chain of equivalences:

$a \models \phi$	iff	$(a, \phi) \in S_n$	by definition of S_n
	iff	$(a, \phi) \in \alpha_n$	by Lemma 4.2
	iff	$\alpha_n(a) = \phi$	
	iff	$\beta_n(f(a)) = \phi$	by the diagram above
	iff	$f(a) \models \phi$	as above

This completes the proof. ⊣

Theorem 4.4 *Let F be finitary, and let (C, γ) be a final coalgebra for F. Let (A, α) be a coalgebra for F, and let $\psi : A \to C$ be a final coalgebra morphism. Then $a, b \in A$ satisfy the same sentences of $\mathcal{L}$ iff $\psi(a) = \psi(b)$.*

Proof We work with the representation of the final coalgebra for F given in Section 3. In our earlier notation, we may assume that $(C, \gamma) = (L^*, l^*)$ and that $\psi = j_A$. Suppose that a and b satisfy the same sentences of $\mathcal{L}$. Due to Lemma 4.2, for all n, $\alpha_n(a) = \alpha_n(b)$. Thus $\alpha^\dagger(a) = \alpha^\dagger(b)$. Since $i \cdot j_A = \alpha^\dagger$ and i is injective, we see that $j_A(a) = j_A(b)$. Conversely, if $j_A(a) = j_A(b)$, then by two applications of Lemma 4.3, the three points a, b, and $j_A(a) = j_A(b)$ all satisfy exactly the same sentences of $\mathcal{L}$. ⊣

5 Discussion

We close with some examples of the logical systems $\mathcal{L}_F$ and with comments on the general construction.

First, let F be the finite power set functor $\mathcal{P}_{fin}$. In this case, $\mathcal{L}_F$ is the (disjoint) union of $\{true\}$, $\{\emptyset, \{true\}\}$, …, $\mathcal{P}^n_{fin}\{true\}$, …. The semantics of *true* is trivial. Every other sentence ϕ is a set of sentences. In a coalgebra (A, α),

$$a \models \phi \quad \text{iff} \quad (\forall b \in \alpha(a))(\exists \psi \in \phi) b \models \phi \text{ and } (\forall \psi \in \phi)(\exists b \in \alpha(a)) b \models \phi$$

We may view a coalgebra for F as a Kripke model for modal logic, except where there are no atomic propositions. The sentences in $\mathcal{L}_F$ are then the *characterizing sentences* of modal logic (also called the *Fine normal forms*).

Second, let $FX = \{a,b\} + (X \times X)$, where $a \neq b$ are arbitrary sets. Then $\mathcal{L}_F$ contains *true*, a, b, $(\mathit{true}, \mathit{true})$, (a,b), $((a,b),(b,b))$, etc. Here is an example: let $X = \{1,2,3,4\}$, let $f : X \to FX$ be $f(1) = a$, $f(2) = b$, $f(3) = (2,1)$, and $f(4) = (3,2)$. We have $3 \models (b,a)$, and also $4 \models ((\mathit{true}, \mathit{true}), b)$. For this functor F, it is easy to directly verify that $\mathcal{L}_F$ is expressive.

These examples point to a general defect in the logics $\mathcal{L}_F$: not only do they lack the boolean connectives, in addition, the logics are in general not even closed under the functor. That is, $F(\mathcal{L}_F)$ is in general not a subset of $\mathcal{L}_F$. In fact, there are sentences of $F(\mathcal{L}_F)$ which are not semantically equivalent to sentences of $\mathcal{L}_F$. (For a concrete example coming from the functor just above, consider (a, true).)

For this reason, I would think that anyone unhappy with the syntax of coalgebraic logic as in [6] would be even less pleased with the proposal here. And so the main point of having the language $\mathcal{L}_F$ would be that it appears to be a minimal language for expressivity: to show any language L to be expressive, even one obtained using predicate liftings or other devices, one need only show that L includes sentences which are logically equivalent to those of $\mathcal{L}_F$.

We conclude with a few remarks going back to our point in the Introduction about which conditions on finitary functors are needed in the study of coalgebraic logics. The condition of uniformity in [6] mainly came into play for the sake of the full power set. Indeed, the main application of uniformity was to express the coalgebraic logic of a uniform functor in terms of infinitary modal logic, a language for which we already had the strong expressivity. If one restricts attention to accessible functors, my suspicion is that uniformity will not be needed in expressivity results. But getting back to the finitary case, the only case considered in this paper, it appears now that extra conditions, particularly standardness and preservation of weak pullbacks, might well be needed to obtain languages stronger than the language $\mathcal{L}_F$ from this paper. For example, to get a language $\mathcal{L}_1$ closed under the functor, we would take $\mathcal{L}_1$ to be the least fixed point (initial algebra) of $1 + F$. So we have

$$\mathcal{L}_1 \quad = \quad 1 + F\mathcal{L}_1.$$

(Recall that our $\mathcal{L}_F$ is $\coprod_n F^n 1$.) It is convenient here to assume that F is standard so that we may write $\mathcal{L}_1$ as $\bigcup_n \mathcal{L}_1(n)$, where $\mathcal{L}_1(0) = \emptyset$, and $\mathcal{L}_1(n{+}1) = 1 + F\mathcal{L}_1(n)$. And again assuming the standardness, it is not so hard to re-prove the results of the previous section to show the expressivity for all finitary F. There are a few more details, but not significantly more than we have seen in Section 4.

To have the additional closure under binary conjunction, one would take a language like $\mathcal{L}_2$, the least solution of

$$\mathcal{L}_2 \quad = \quad 1 + F\mathcal{L}_2 + (\mathcal{L}_2 \times \mathcal{L}_2).$$

Alternatively, one could just take the finite power set and consider

$$\mathcal{L}_3 \quad = \quad F\mathcal{L}_3 + \mathcal{P}_{fin}\mathcal{L}_3.$$

This is exactly the finite version of what was done in earlier in [6]. In fact, one would like to go even further and have a language $\mathcal{L}_4$ with a syntactic negation, and for this one only needs to add another summand.

But at this point, even to prove the basic facts on semantics for languages of this type seems to require that F preserve weak pullbacks, or perhaps a weaker condition such as preservation of weak monos. This matter is the main open question coming from this paper: given a finitary functor F are any conditions needed to guarantee the existence of a language for coalgebras of F which is expressive, closed under the functor itself, and under the finite boolean operations? If so, what are the weakest conditions on F?

Acknowledgment

I thank the referees for several important suggestions, including the observations in Section 3.1, and also for the short proof of Lemma 4.3.

References

[1] Jiří Adámek, Free algebras and automata realizations in the language of categories. *Commentationes Mathematicae Universitatis Carolinae* 15 (1974), 589–602.

[2] Jiří Adámek and Vaćlav Koubek, On the greatest fixed point of a set functor, *Theoretical Computer Science* 150 (1995), 57–75.

[3] Jon Barwise and Lawrence Moss. *Vicious Circles*. CSLI Lecture Notes Number 60, CSLI Publications, Stanford, 1996.

[4] Johan van Benthem and Jan Bergstra, Logic of transition systems. *Journal of Logic, Language, and Information* 3 (1995), no. 4, 247–283.

[5] Alexander Kurz and Dirk Pattinson, Coalgebraic modal logic of finite rank. *Mathematical Structures in Computer Science*, 15(3) (2005):453–473.

[6] Lawrence S. Moss, Coalgebraic logic, *Annals of Pure and Applied Logic* 96 (1999), no. 1-3, 277–317; also vol. 99 (1999) no. 1–3, 241–259.

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

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

[9] J.J.M.M. Rutten, Universal coalgebra: a theory of systems. *Theoretical Computer Science* (2000), 249, no. 1, 3–80.

[10] Lutz Schröder, Expressivity of coalgebraic modal logic: the limits and beyond. To appear in *Theoretical Computer Science*.

[11] Ignacio Viglizzo, **Coalgebras on Measurable Spaces**. Ph.D. dissertation, Indiana University, Bloomington, 2005.

[12] James Worrell, On the final sequence of a finitary set functor. *Theoretical Computer Science* 338 (2005), no. 1–3, 184–199.