What is the Coalgebraic Analogue of Birkhoff's Variety Theorem?

Robert Goldblatt*
School of Mathematical and Computing Sciences
Victoria University
P. O. Box 600, Wellington, New Zealand
Rob.Goldblatt@vuw.ac.nz

30 October 2000

Abstract

Logical definability is investigated for certain classes of coalgebras related to state-transition systems, hidden algebras and Kripke models. The *filter enlargement* of a coalgebra A is introduced as a new coalgebra A^+ whose states are special "observationally rich" filters on the state set of A. The *ultrafilter enlargement* is the subcoalgebra A^* of A^+ whose states are ultrafilters.

Boolean combinations of equations between terms of *observable* (or *output*) type are identified as a natural class of formulas for specifying properties of coalgebras. These *observable formulas* are permitted to have a single state variable, and form a language in which *modalities* describing the effects of state transitions are implicitly present. A^* and A^+ validate the same observable formulas.

It is shown that a class of coalgebras is definable by observable formulas iff the class is closed under disjoint unions, images of bisimulations, and (ultra)filter enlargements. (Closure under images of bisimulations is equivalent to closure under images and domains of coalgebraic morphisms.) Moreover, every set of observable formulas has the same models as some set of *conditional equations*.

Examples are constructed to show that the use of enlargements is essential in these characterisations, and that there are classes of coalgebras definable by conditional observable equations, but not by equations alone.

The main conclusion of the paper is that to structurally characterise classes of coalgebras that are logically definable by modal languages requires a new construction, of "Stone space" type, in addition to the coalgebraic duals of the three constructions (homomorphisms, subalgebras, direct products) that occur in Birkhoff's original variety theorem for algebras.

*I would like to thank the referees for their positive responses and helpful comments.

Keywords: Coalgebra, observationally rich filter, observable formula, bisimulation.

Contents

1 Introduction and Overview

Certain kinds of coalgebra have proven useful in the modelling of features of object-oriented programming languages [Rei95, Jac96]. Since coalgebras are dual to algebras (in the sense of category theory), this has provided impetus for the study of aspects of coalgebraic theory that correspond to parts of the general theory of algebras [Rut95, Rut00]. A notable case is Birkhoff's celebrated variety theorem [Bir35], stating that a class of algebras is definable by equations iff it is closed under homomorphic images, subalgebras, and direct products. These three constructions dualise to subcoalgebras, images of coalgebraic morphisms and disjoint unions, respectively, and hence a class of coalgebras closed under the latter operations has been dubbed a *covariety*. Several papers [GS98, Kur98a, Rut00] explore characterisations of covarieties that involve, variously, abstract dualisation of the notion of a variable; axiomatisations of the notions of formula and satisfaction; and the use of infinitary logic. Another approach [Roş98] characterises equationally specifiable classes of coalgebras in terms of constructions, one of which is itself defined in terms of preservation of equations and so is not fully "structural".

This paper provides a different perspective on questions of definability, and has something to say about both sides of the coin of results like Birkhoff's. On the one hand there is the question of the syntactic form of expressions used to define classes of structures, and on the other there is the nature of the constructions under which these classes are closed.

The kind of coalgebra to be considered here typically has a set X of *states* (which may be thought of as the possible realisations of some notion of object), and a set of operations which are divided into *methods* and *attributes*. Methods are of the type $m : X \times I \to X$ and attributes of the type $a : X \times I \to O$, where I is a set of inputs and O a set of outputs. A method can be thought of as a system of state transitions $x \mapsto m(x, b)$ parameterised by inputs b, while attributes assign outputs or "observable values" $a(x, b)$ to states x relative to an input. We may also suppress the input set (or regard it as having one element that can be ignored), to consider methods of type $X \to X$ and attributes of type $X \to O$. Thus this concept of coalgebra encompasses the idea of a transition system, such as an automaton or a Kripke frame for modal logic, as well as modelling the "hidden" state space of a specified class in an object-oriented language. Indeed the notion can be viewed as a *hidden algebra* [GM99, Roş98] having a single hidden sort.

On the syntactic side, examples of coalgebra specification in the literature [Rei95, Jac96] have a number of distinctive features that have influenced our choice of formalism:

(i) Terms and formulas typically contain a *single* state variable. We implement this by using a special symbol σ, which, rather than being a variable, may be thought of as a parameter denoting the "current state". Variables proper are constrained to take only *data elements*, i.e. inputs or outputs, as values.

(ii) An equation between terms is interpreted as meaning *equality* if those terms take data elements as values, and is interpreted as *bisimilarity* when they evaluate to states. To facilitate this distinction different formal symbols will be used for the two cases: *data equations* are written $t_1 \approx t_2$, while *state equations* are written $t_1 \simeq t_2$.

(iii) Significant properties of coalgebras sometimes require logical connectives as well as equations for their formulation. It is therefore appropriate to study *Boolean combinations* of equations, and our principal result is a characterisation of classes of coalgebras specifiable by such combinations.

Experience with other kinds of transition system modelling, such as in dynamic logic, has motivated the use of *modal logic* in the study of coalgebras [BM96, Kur98b]. In the present case it would be natural to employ modalities $[m.b]$ corresponding to state transitions, with a formula $[m.b]\varphi$ expressing the statement "after applying method m with parameter b, φ is true". But in the theory developed here, it turns out that this is expressed by the formula $\varphi(m(\sigma, b))$, obtained from φ by substituting the term $m(\sigma, b)$ in place of σ (Corollary 4.4 ff.). Thus by taking Boolean combinations of appropriate equations *we are already dealing with a modal language*, one in which the modalities are implicit and definable from other constructs.

Bisimilarity is intended to express the *behavioural indistinguishability* of states. A basic principle of the object-oriented philosophy is that states are not directly accessible. Computationally, all that we know about them is what can be learned by performing "experiments" and observing the results (outputs). An experiment consists of the application of a combination of methods, followed by an attribute (observation). This amounts to the evaluation of a term of output type [GM99, p. 292], and states prove to be bisimilar when they assign the same values to all output terms (Theorem 3.2). Hence satisfaction of state equations is reducible to satisfaction of certain output equations (Theorem 4.1) and suggests a prominent role for Boolean combinations of output equations – which we call *observable formulas.*

On the structural side, familiarity with modal model theory [GT75, Gol93] suggests that to characterise classes of coalgebras that are definable in a modal language will require some new construction, of "Stone space" type, in addition to the duals of Birkhoff's three constructions. For this purpose we introduce the notion of the *filter enlargement* of a coalgebra A. This is a coalgebra A^+ whose states are certain "observationally rich" filters on the state set of A. These filters are required to contain special sets of states defined by equations between terms of output type. Then the *ultrafilter enlargement* of A is defined to be the subcoalgebra A^* of A^+ consisting of those states that are *ultra*filters. We show that consideration of A^* and A^+ in specifying classes of coalgebras is unavoidable, because they preserve satisfaction of Boolean combinations of output equations (see Corollary 8.7). The proof of this fact requiries the establishment of relationships between A and A^+ (and A^*) reminiscent of the relationships between a structure and its reduced powers and ultrapowers that underlie Łoś's Theorem (see Theorems 8.3 and 8.4).

The principal result of the paper (Theorem 9.2) is that

> *a class of coalgebras is definable by a set of observable formulas if and only if it is closed under disjoint unions, images of bisimulation relations, and ultrafilter enlargements.*

(Closure under images of bisimulations is equivalent to closure under images and domains of coalgebraic morphisms.) Moreover, the proof of 9.2 shows that any class definable by a set of observable formulas is in fact also definable by a set of *conditional observable equations*

$$t \approx u \quad \textbf{if} \quad t_1 \approx u_1 \textbf{ and } \dots \textbf{ and } t_k \approx u_k.$$

These characterisations are analogous to a theorem from modal logic that appeared in [GT75], giving structural conditions under which a first-order axiomatisable class of Kripke frames is definable by modal propositional formulas. That involved the notion of the *ultrafilter extension* of a frame. Rutten [Rut95, p. 254] has asked about the extent to which the result generalises to arbitrary coalgebras, and this article provides a partial answer to his question. There are some significant differences between the present coalgebraic situation and the modal one. Ultrafilter enlargements of our coalgebras are instances of ultrafilter extensions of modal frames as far as their treatment of *methods* as concerned, but have additional structure produced by the *attributes* that is not found in standard modal model theory. Also, whereas a class of coalgebras defined by observable formulas is closed under ultrafilter enlargements and domains of coalgebraic morphisms, the corresponding properties do not always hold for the class of frames defined by a propositional modal formula. Whereas modal formulas are constructed from variables that take arbitrary sets of states as values, observable formulas constitute a language with a more specialised syntax and restricted expressive power. Further discussion of these comparisons is given at the end of Section 8 of this paper, and in Section 7 of [Gol00].

Section 10 concludes our work with a series of four examples:

1. It is shown that closure under (ultra)filter enlargements is independent of the other closure conditions discussed, by exhibiting a class of coalgebras that is closed under disjoint unions and images of bisimulations, hence under domains and images of morphisms, but not closed under ultrafilter or filter enlargements.

2. The second example shows that ultrafilter enlargements do not always preserve satisfaction of negations $\neg(t_1 \simeq t_2)$ of bisimilarity equations between state-valued terms (although they do preserve the equations themselves). Thus a characterisation of classes defined by such inequations, or other Boolean combinations of state equations, would have to use new closure concepts.

3. This example shows that the hypothesis of closure under domains of morphisms in our main results cannot be weakened to closure under subcoalgebras. A class of coalgebras is constructed that is closed under disjoint unions, images of morphisms, subcoalgebras, and filter and ultrafilter enlargements, but not closed under domains of morphisms.

4. The fourth example is a class of coalgebras that is definable by conditional observable equations, but not by equations alone.

The paper does not claim to fully resolve the question in its title. Rather it seeks to view the question in a new light, by proposing a different kind of answer to those considered so far. Birkhoff's theorem gives a *structural* characterisation of certain *logically definable* classes of algebras, and our main point is that simply dualising Birkhoff's three constructions will not provide a characterisation of classes of coalgebras definable by modal languages. The coalgebras we consider are extremely simple—they come from "monomial" endofunctors $\mathbf{Set} \to \mathbf{Set}$ on the category of sets that are products of functors of the form $X \mapsto X^I$ and $X \mapsto O^I$ where I and O are constant sets—but they are already complex enough to demonstrate our main point and to provide a new "Birkhoff theorem". However many natural specifications of object classes and abstract machines depend additionally on coproducts of functors for their definition [Jac96], and to extend our results in that direction will require the development of a more sophisticated syntax, pehaps along the lines of [Jac00]. Other problems to be addressed concern the issues of whether there is in fact a structural characterisation of classes defined just by equations ([Roş98] doubts this), or by Boolean combinations of *state* equations.

The proof of our principal result is model-theoretic and involves the theory of bisimulation relations. It is possible to take a more algebraic approach (like that of [GT75]) by using *Boolean algebras with operators* (BAO's) based on the powersets of the state sets of coalgebras. Then the proof can be transformed into an explicit application of Birkhoff's original theorem to these BAO's. This approach is fully worked out in a companion paper [Gol00].

2 Coalgebras and Terms

A *signature for coalgebras* is a sequence $(In, Out, Meth, Att)$ with the following properties.

- In is a collection of *input sets.*
- Out is a collection of *output sets.*
- $Meth$ is a collection of *method symbols.* Each symbol m in $Meth$ is assigned an input set I^m, called the *sort* of m.
- Att is a collection of *attribute symbols.* Each a in Att is assigned a pair (I^a, O^a) as its *sort*, with $I^a \in In$ and $O^a \in Out$.

Let $D = \bigcup(In \cup Out)$, the collection of all elements appearing in any input or output set. Members of D are *data elements.* We will assume that there is at least one output set containing *two distinct* elements (see Section 9 for the use of this assumption).

Fix a signature as described. A *coalgebra* for this signature is a structure

$$A = (X^A, \{m^A : m \in Meth\}, \{a^A : a \in Att\}),$$

consisting of a set X^A of *states*, a function $m^A : X^A \times I^m \to X^A$ for each method symbol m of sort I^m, and a function $a^A : X^A \times I^a \to O^a$ for each attribute symbol a of sort (I^a, O^a).

It is appropriate also to consider methods of the form $X^A \to X^A$ and attributes of the form $X^A \to O^a$, without any associated input sets. These can be subsumed under the present framework by taking I^m or I^a to be a one-element set $\{*\}$, and identifying X^A with $X^A \times \{*\}$.

The given definition of coalgebra can be equated to the categorical concept by associating with each set X the product set

$$T(X) = \prod_{m \in Meth} X^{I^m} \times \prod_{a \in Att} (O^a)^{I^a}.$$

Then a coalgebra A as above can be identified with the function $\chi : X^A \to T(X^A)$ having

$$\chi(x) = \big\langle \langle m_x^A : m \in Meth \rangle, \langle a_x^A : a \in Att \rangle \big\rangle,$$

where $m_x^A \in X^{I^m}$ is the function $b \mapsto m^A(x,b)$ and $a_x^A \in (O^a)^{I^a}$ is the function $c \mapsto a^A(x,c)$. The function T itself lifts naturally to an endofunctor $T : \mathbf{Set} \to \mathbf{Set}$ on the category of sets, and a coalgebra as a pair (X^A, χ) with $\chi : X^A \to T(X^A)$ is precisely a T-coalgebra as defined in category theory [BW85, p. 100].

In the absence of attribute functions, a coalgebra may be viewed as the collection of state-transition functions $x \mapsto m^A(x,b)$ for all $m \in Meth$ and $b \in I^m$. This corresponds to a Kripke frame for modal logic that consists of a collection of *functional* accessibility relations. That is a very particular and well understood case in modal semantics. But the presence of attributes immediately adds a new structural dimension going beyond the standard modal context, as is manifest in the new concepts of "observationally rich" ultrafilter and ultrafilter enlargement to be discussed in Section 8.

Terms

Let $\{V_P : P \in In \cup Out\}$ be a $In \cup Out$-indexed collection of sets of *variables*. Members of V_P are called *variables of type* P, and are intended to take data elements from P as values. A *term* is a finite string of symbols constructed from these variables, members of D as constants, and a special symbol σ, using method and attribute symbols. σ may be thought of as a parameter denoting the *current state*. If $\bar{v} = (v_1, \ldots, v_n)$ is a tuple of variables, we write $t(\bar{v})$ to indicate that all the variables occuring in term t are among $v_1, \ldots, v_n$. The syntax of terms is specified inductively by the rules given below. This syntax is also summarized in column 1 of Table 1.

Terms are expressions taking either data elements or states as values. They are therefore classified as *data terms* or *state terms*, with data terms further classified as *input terms* or *observable* (i.e. *output*) terms. A data term is *ground* if it has no variables. In any coalgebra each term has an associated *type*, which is either an input set, an output set, or X^A (see column 2 of Table 1). While it is possible for a given set to be both an input set and an output set (i.e. $In \cap Out \neq \emptyset$), this should not lead to confusion about the intended classification of a given term.

Here then are the rules for defining terms and their types:

- For each $P \in In \cup Out$, each v in V_P is a data term, called a *variable of type* P.
- For each $P \in In \cup Out$, each b in P is a data term, called a *constant of type* P.
- The symbol σ is a state term.
- If m is a method symbol of sort I^m, t_1 is a state term, and t_2 is an input term of type I^m, then $m(t_1, t_2)$ is a state term.
- If a is an attribute symbol of sort (I^a, O^a), t_1 is a state term, and t_2 is an input term of type I^a, then $a(t_1, t_2)$ is an observable term of type O^a.

Term $t(\bar{v})$	Type	Value $t^A[x,\bar{d}]$
variable v_i in V_P	P	d_i
constant b in I or O	I or O	b
σ	X^A	x
$m(t_1,t_2)$	X^A	$m^A(t_1^A[x,\bar{d}],t_2^A[x,\bar{d}])$
$a(t_1,t_2)$	O^a	$a^A(t_1^A[x,\bar{d}],t_2^A[x,\bar{d}])$

Table 1

A term takes a *value* belonging to its type once values have been given to its variables and a state has been specified as the denotation of σ. To define this we say that a tuple $\bar{d} = (d_1,\ldots,d_n)$ of data elements *matches* a tuple $\bar{v}$ of variables if $\bar{v}$ and $\bar{d}$ are of the same length and each v_i has the same type as the corresponding d_i. In that case $\bar{d}$ serves as a valuation, assigning values to the v_i's. The symbol

$$t^A[x,\bar{d}]$$

denotes the *value of term $t(\bar{v})$ in coalgebra A at state $x \in X^A$ under the assignment of $\bar{d}$ to $\bar{v}$*. This is defined inductively in column 3 of Table 1. If t is ground (no variables), then the value $t^A[x,\bar{d}]$ is independent of $\bar{d}$, and may be abbreviated to $t^A[x]$, the value of t *at state x*.

The term $t(\bar{d})$ defined in column 2 of Table 2 is the *ground* term obtained by replacing each occurrence of v_i in t by the constant d_i. It is readily seen that the value $t^A[x,\bar{d}]$ is the same as the value of $t(\bar{d})$ at x:

$$t^A[x,\bar{d}] = t(\bar{d})^A[x].$$

Thus discussion of the values of terms can in principle be reduced to discussion of the values of ground terms.

$t(\bar{v})$	$t(\bar{d})$	$t(u)$
variable v_i	d_i	v_i
constant b	b	b
σ	σ	u
$m(t_1,t_2)$	$m(t_1(\bar{d}),t_2(\bar{d}))$	$m(t_1(u),t_2(u))$
$a(t_1,t_2)$	$a(t_1(\bar{d}),t_2(\bar{d}))$	$a(t_1(u),t_2(u))$

Table 2

Now the value of a term t at a state of the form $m^A(x,b)$, resulting from the application of a method m^A, can be expressed as the value of an associated term at the initial state x. To explain this, we write $t(u)$ for the term obtained by replacing every occurrence of σ in t by a *state* term u. This notion is defined in column 3 of Table 2. Note that t and $t(u)$ are always of the same type.

The result we have in mind (Corollary 2.2 below) is that the value of t at state $m^A(x,b)$ is the same as the value at x of the term $t(m(\sigma,b))$ obtained by substituting $m(\sigma,b)$ for σ in t. More generally, evaluation of any state term u at x produces a new state $u^A[x]$, and the value of t at $u^A[x]$ is the same as the value of $t(u)$ at x. This is a relationship between syntactic substitution and semantic evaluation of terms that is familiar from first-order logic, where it is sometime called the Substitution Lemma (see for instance [EFT84, p. 52] or [End72, p. 126]). In the present context it takes on a more "dynamic" flavour:

Theorem 2.1 *Let u be a state term. Then for any term t, if $\bar{v}$ includes all the variables of $t(u)$, then*

$$t^A[u^A[x,\bar{d}\,],\bar{d}\,] = t(u)^A[x,\bar{d}\,],$$

for all $x \in X^A$, and all $\bar{d}$ matching $\bar{v}$,.

Proof. By induction on the length of t. If t is a variable v_i or constant b, then both sides of the equation evaluate to d_i, or to b, respectively. If t is σ, then $t(u)$ is u and both sides evaluate to $u^A[x,\bar{d}\,]$.

Now let t be $m(t_1,t_2)$ and make the induction hypothesis that the Theorem holds for state term t_1 and data term t_2. Then

$$\begin{array}{ll}
m(t_1,t_2)^A[u^A[x,\bar{d}\,],\bar{d}\,] & \\
= m^A(t_1^A[u^A[x,\bar{d}\,],\bar{d}\,], t_2^A[u^A[x,\bar{d}\,],\bar{d}\,]) & \text{Table 1} \\
= m^A(t_1(u)^A[x,\bar{d}\,], t_2(u)^A[x,\bar{d}\,]) & \text{hypothesis on } t_1,t_2 \\
= m(t_1(u),t_2(u))^A[x,\bar{d}\,] & \text{Table 1} \\
= m(t_1,t_2)(u)^A[x,\bar{d}\,] & \text{Table 2,}
\end{array}$$

showing that the result holds in this case.

Finally, if t is $a(t_1,t_2)$ and the result is assumed to hold for t_1 and t_2, then it holds for t by an argument that looks identical to the previous case, but with a in place of m. □

Corollary 2.2 $t^A[m^A(x,b),\bar{d}\,] = t(m(\sigma,b))^A[x,\bar{d}\,]$

Proof. $m^A(x,b) = m(\sigma,b)^A[x,\bar{d}\,]$. □

3 Bisimulation, Bisimilarity and Quotients

Bisimulations and bisimilarities are fundamental to what follows. We begin with a review of these cornerstone concepts from the theory of computational processes.

A binary relation $\rho \subseteq X^A \times X^B$ is called a *bisimulation from A to B* if methods preserve ρ-relatedness of states, and attributes assign the same observable values to ρ-related states, i.e.

$$x\rho y \quad \text{implies} \quad m^A(x,b)\,\rho\, m^B(y,b) \text{ and } a^A(x,c) = a^B(y,c)$$

for all $x \in X^A$, $y \in X^B$, $m \in Meth$, $b \in I^m$, $a \in Att$ and $c \in I^a$.

The union of any collection of such bisimulations is also a bisimulation from A to B. In particular, the union of *all* bisimulations from A to B is a bisimulation, to be denoted $\sim$, which is thereby the *largest* such bisimulation. If $x \sim y$ we say that x and y are *bisimilar*. Thus x and y are bisimilar when there is some bisimulation ρ with $x\rho y$.

Since bisimilarity can be defined between any coalgebras, it would be appropriate to use a notation like $\sim^{A,B}$ to denote its manifestation as a relation from A to B. In the case $A = B$, we may write $\sim^A$ for the restriction of this relation to the states of the coalgebra A. This elaboration of notation will only be used when needed for clarity.

Now the inverse ρ^{-1} of a bisimulation ρ from A to B is itself a bisimulation from B to A. If ρ_1 is a bisimulation from A_1 to A_2, and ρ_2 is a bisimulation from A_2 to A_3, then their composition $\rho_1 \circ \rho_2$ is a bisimulation from A_1 to A_3. In the case of a single coalgebra A, these facts can be used to show that $\sim_A$ is an *equivalence* relation on the state set X^A.

Theorem 3.1 *Let ρ be a bisimulation from A to B. Let $t(\bar{v})$ be a term, and suppose $\bar{d}$ matches $\bar{v}$. Then if $x\rho y$,*

(1) $t^A[x,\bar{d}]\,\rho\, t^B[y,\bar{d}]$ *if t is a state term; and*

(2) $t^A[x,\bar{d}] = t^B[y,\bar{d}]$ *if t is a data term.*

Proof. By induction on the length of t. Let $x\rho y$. If t is σ, then (1) just reasserts this assumption $x\rho y$. If t is a variable v_i or a constant b, then $t^A[x,\bar{d}]$ and $t^B[y,\bar{d}]$ are both either d_i or b, respectively, and so (2) holds for t.

Now let t be $m(t_1, t_2)$ and make the induction hypothesis that the Theorem holds for state term t_1 and data term t_2. This means that $t_1^A[x,\bar{d}]\,\rho\, t_1^B[y,\bar{d}]$ and $t_2^A[x,\bar{d}] = t_2^B[y,\bar{d}]$. Since ρ is a bisimulation, these last two equations lead to

$$m^A(t_1^A[x,\bar{d}], t_2^A[x,\bar{d}])\,\rho\, m^B(t_1^B[y,\bar{d}], t_2^B[y,\bar{d}]),$$

i.e. $m(t_1,t_2)^A[x,\bar{d}]\,\rho\, m(t_1,t_2)^B[y,\bar{d}]$, proving (1) in this case.

Similarly, if t is $a(t_1,t_2)$ and the Theorem holds for t_1 and t_2, the fact that ρ is a bisimulation leads to

$$a^A(t_1^A[x,\bar{d}], t_2^A[x,\bar{d}]) = a^B(t_1^B[y,\bar{d}], t_2^B[y,\bar{d}]),$$

giving (2) for this case of t. □

There are a number of logical characterisations of bisimilarity. As explained in Section 1, a ground observable term describes an "experiment" that can be performed upon a state to yield some observable value as output (see, e.g., [GM99,

p. 292], where these terms are called *contexts*). Two states are indistinguishable if they exhibit exactly the same behaviour, in the sense that all experiments on them produce the same results. This notion of indistinguishability is captured by bisimilarity, as the next theorem shows.

Theorem 3.2 *For any states $x \in X^A$ and $y \in X^B$ the following are equivalent.*

(1) $x \sim y$.

(2) $t^A[x, \bar{d}] = t^B[y, \bar{d}]$ *for all observable terms $t(\bar{v})$ and all $\bar{d}$ matching $\bar{v}$.*

(3) $t^A[x] = t^B[y]$ *for all ground observable terms t.*

(4) $t_1^A[x] = t_2^A[x]$ *iff* $t_1^B[y] = t_2^B[y]$ *for all ground observable terms t_1, t_2.*

Proof.
(1) implies (2): Suppose $x \sim y$. Then $x \rho y$ for some bisimulation ρ. Now (2) follows by Theorem 3.1(2).
(2) implies (3): immediate.
(3) implies (4): If (3) holds then $t_1^A[x] = t_1^B[y]$ and $t_2^A[x] = t_2^B[y]$, which directly implies (4).
(4) implies (1): Define a relation ρ by specifying that $z \rho w$ iff

$$t_1^A[z] = t_2^A[z] \text{ iff } t_1^B[w] = t_2^B[w] \text{ for all ground observable terms } t_1,\ t_2.$$

We show that this ρ is a bisimulation from A to B. Then if (4) holds for a particular x and y, we have $x \rho y$, and hence $x \sim y$ since $\sim$ is the largest bisimulation, so (1) holds.

Let $z \rho w$. First we show $m^A(z, b) \, \rho \, m^B(w, b)$. Noting that if t is a ground observable term, then so is $t(m(\sigma, b))$, we have for any ground observable t_1, t_2 that

$$\begin{array}{lll} & t_1^A[m^A(z,b)] = t_2^A[m^A(z,b)] & \\ \text{iff} & t_1(m(\sigma,b))^A[z] = t_2(m(\sigma,b))^A[z] & \text{by 2.2} \\ \text{iff} & t_1(m(\sigma,b))^B[w] = t_2(m(\sigma,b))^B[w] & \text{by 3.1(2), as } z\rho w \\ \text{iff} & t_1^B[m^B(w,b)] = t_2^B[m^B(w,b)] & \text{by 2.2.} \end{array}$$

This proves $m^A(z, b) \, \rho \, m^B(w, b)$.

Secondly we must show $a^A(z, c) = a^B(w, c)$. Let $a^A(z, c) = d$. Then $a(\sigma, c)^A[z] = d^A[z]$, so as $z \rho w$, $a(\sigma, c)^B[w] = d^B[w]$, i.e. $a^B(w, c) = d$ as required. □

This proof also shows that, in view of 3.1(2), replacing "observable" by "data" in each of (2), (3), and (4) gives further conditions equivalent to $x \sim y$. More will be said in the next section about the characterisation of bisimilarity.

Corollary 3.3 *Let ρ be a bisimulation from A to B, and suppose that $x \rho y$ and $x' \rho y'$. Then $x \sim x'$ in A iff $y \sim y'$ in B.*

Proof. If $x \sim^A x'$, then for any ground observable t, by repeated use of 3.1(2),

$$\begin{aligned} t^B[y] &= t^A[x] \quad \text{as } x\rho y \\ &= t^A[x'] \quad \text{as } x \sim^A x' \\ &= t^B[y'] \quad \text{as } x'\rho y', \end{aligned}$$

showing that $y \sim^B y'$ by 3.2(3). Likewise for the converse.

Alternatively, a set-theoretic proof not requiring 3.1 and 3.2 observes that if $x \sim^A x'$, then altogether $y\,\rho^{-1}\,x \sim^A x'\,\rho\,y'$. But the composition $\rho^{-1} \circ \sim^A \circ \rho$ is a bisimulation on B, so is contained in $\sim^B$, hence $y \sim^B y'$, etc. □

A new coalgebra can be formed from a given one by identifying bisimilar states. This is possible because $\sim^A$ is an *equivalence* relation on the state set X^A for which

$$x \sim^A y \quad \text{implies} \quad m^A(x,b) \sim^A m^A(y,b) \text{ and } a^A(x,c) = a^A(y,c). \tag{3.i}$$

These properties allow a coalgebra $A^\natural$ to be defined on on the set of $\sim$-equivalence classes of X^A. $A^\natural$ will be called the *quotient* of A. For each $x \in X^A$, put $|x| = \{y \in X^A : x \sim^A y\}$. Let

$$A^\natural = (X^{A^\natural}, \{m^{A^\natural} : m \in Meth\}, \{a^{A^\natural} : a \in Out\}),$$

where

$$\begin{aligned} X^{A^\natural} &= \{|x| : x \in X^A\}, \\ m^{A^\natural}(|x|, b) &= |m^A(x,b)|, \\ a^{A^\natural}(|x|, c) &= a^A(x,c). \end{aligned}$$

$m^{A^\natural}$ and $a^{A^\natural}$ are well-defined in view of (3.i). Their definitions ensure that the natural map $f^\natural : x \mapsto |x|$ is a bisimulation from A to $A^\natural$. (This map is an example of the important concept of *morphism*, to be studied in Section 5.) Corollary 3.3 can be applied, with $\rho = f^\natural$, to show that bisimilar states of $A^\natural$ are equal.

4 Formulas and their Satisfaction

A *data equation* is an expression of the form $t_1 \approx t_2$, where t_1 and t_2 are data terms. The equation is intended to express the *equality* of the values of t_1 and t_2. It is an *observable* equation if both of its terms are observable.

A *state equation* is an expression of the form $t_1 \simeq t_2$, where t_1 and t_2 are state terms. It is intended to express the *bisimilarity* of the values of t_1 and t_2.

A *formula* is any Boolean combination of equations, i.e. any expression built from equations by the usual logical connectives $\neg$, $\wedge$, $\rightarrow$, $\vee$, $\leftrightarrow$. An *observable/data/state* formula is one containing only observable/data/state terms (i.e.

built from observable/data/state equations only), and a *ground* formula is one without any variables. We will make particular use of formulas that are both ground and observable.

The notation

$$A, x \models \varphi[\bar{d}\,]$$

means that *formula* φ *is satisfied in coalgebra* A *by state* x *under assignment* $\bar{d}$, and is defined inductively by

$$\begin{array}{lll}
A, x \models t_1 \approx t_2[\bar{d}\,] & \text{iff} & t_1^A[x, \bar{d}\,] = t_2^A[x, \bar{d}\,] \\
A, x \models t_1 \simeq t_2[\bar{d}\,] & \text{iff} & t_1^A[x, \bar{d}\,] \sim^A t_2^A[x, \bar{d}\,] \\
A, x \models \neg\varphi[\bar{d}\,] & \text{iff} & \text{not } A, x \models \varphi[\bar{d}\,] \\
A, x \models \varphi_1 \wedge \varphi_2[\bar{d}\,] & \text{iff} & A, x \models \varphi_1[\bar{d}\,] \text{ and } A, x \models \varphi_2[\bar{d}\,] \\
A, x \models \varphi_1 \to \varphi_2[\bar{d}\,] & \text{iff} & A, x \models \varphi_1[\bar{d}\,] \text{ implies } A, x \models \varphi_2[\bar{d}\,],
\end{array}$$

and similarly for the other Boolean connectives. When $A, x \models \varphi[\bar{d}\,]$ we may also say that φ is *true at*, or *of*, x in A under $\bar{d}$.

As with terms, we write $\varphi(\bar{v}\,)$ to indicate that the variables occurring in φ are amongst those of $\bar{v}$, and use $\varphi(\bar{d}\,)$ to denote the ground formula obtained from φ by substituting d_i for v_i. $\varphi(\bar{v}\,)$ is *satisfied by* x per se, written $A, x \models \varphi$, if $A, x \models \varphi[\bar{d}\,]$ for all $\bar{d}$ matching $\bar{v}$.

In general it is true that

$$A, x \models \varphi_1 \wedge \varphi_2 \quad \text{iff} \quad A, x \models \varphi_1 \text{ and } A, x \models \varphi_2,$$

but it may be that neither $A, x \models \varphi$ nor $A, x \models \neg\varphi$, since φ may be satisfied at x by some $\bar{d}$'s and not by others. However if φ is a ground formula, then its satisfaction is independent of value assignment to variables, and in that case we get

$$A, x \models \neg\varphi \quad \text{iff} \quad \text{not } A, x \models \varphi.$$

It is straightforward to check, for any formula φ, that

$$A, x \models \varphi[\bar{d}\,] \text{ iff } A, x \models \varphi(\bar{d}\,), \tag{4.i}$$

and so satisfaction of φ by x reduces to satisfaction of the *set* of ground formulas

$$\{\varphi(\bar{d}\,) : \bar{d} \text{ matches } \bar{v}\,\}.$$

Coalgebra A is a *model* of formula φ, written $A \models \varphi$, if $A, x \models \varphi$ for all states $x \in X^A$. We may also say that A *models* φ, or that φ is *valid in* A, when this occurs. We write $Mod\,\varphi$ for the class of all models of φ, and

$$Mod\,\Phi = \{A : A \models \varphi \text{ for all } \varphi \in \Phi\}$$

for the class of all models of a set Φ of formulas.

In view of Theorem 3.2(4), we now see that states $x \in X^A$ and $y \in X^B$ of two coalgebras are bisimilar precisely when they satisfy the same *ground observable equations*, i.e.

$$A, x \models t_1 \approx t_2 \quad \text{iff} \quad B, y \models t_1 \approx t_2$$

for all ground observable terms t_1, t_2. If this holds, then x and y satisfy exactly the same ground observable formulas, and hence in view of (4.i) satisfy exactly the same observable formulas altogether.

Another observational characterisation of bisimilarity is given by the notion of *behavioural* satisfaction of state equations. The equation $u_1 \simeq u_2$ is *behaviourally satisfied* when $t(u_1) \approx t(u_2)$ is satisfied for all observable terms t [GM99, p. 292]. In this sense any state equation is semantically equivalent to an infinite conjunction of observable equations:

Theorem 4.1 *Let $u_1 \simeq u_2$ be a state equation with variables amongst $\bar{v}$. Then in any coalgebra A, if $x \in X^A$ and $\bar{d}$ matches $\bar{v}$, the following are equivalent:*

(1) $A, x \models u_1 \simeq u_2[\bar{d}]$.

(2) $A, x \models t(u_1) \approx t(u_2)[\bar{d}]$ for all ground observable terms t.

Proof. By definition, $A, x \models u_1 \simeq u_2[\bar{d}]$ iff state $u_1^A[x, \bar{d}]$ is bisimilar to state $u_2^A[x, \bar{d}]$. By Theorem 3.2(3), this bisimilarity holds iff $t^A[u_1^A[x, \bar{d}]] = t^A[u_2^A[x, \bar{d}]]$ for all ground observable terms t. But $t^A[u_i^A[x, \bar{d}]] = t(u_i)^A[x, \bar{d}]$ by Theorem 2.1, so

$$t^A[u_1^A[x, \bar{d}]] = t^A[u_2^A[x, \bar{d}]] \quad \text{iff} \quad A, x \models t(u_1) \approx t(u_2)[\bar{d}].$$

□

We now take up the question of preservation of satisfaction under bisimulation. A bisimulation ρ from A to B will be called *total* if its *domain*

$$\{x \in X^A : \exists y \in X^B (x\rho y)\}$$

is the whole of X^A. ρ is *surjective* if its *image*

$$\{y \in X^B : \exists x \in X^A (x\rho y)\}$$

is the whole of X^B.

Theorem 4.2 *Let ρ be a bisimulation from A to B and $\varphi(\bar{v})$ any formula.*

(1) *If $x\rho y$, then $A, x \models \varphi[\bar{d}]$ iff $B, y \models \varphi[\bar{d}]$, for all $\bar{d}$ matching $\bar{v}$.*

(2) *If $x\rho y$, then $A, x \models \varphi$ iff $B, y \models \varphi$.*

(3) *If ρ is surjective and $A \models \varphi$, then $B \models \varphi$.*

(4) *If ρ is total and $B \models \varphi$, then $A \models \varphi$.*

Proof.

For (1), let $x\rho y$. If φ is the data equation $t_1 \approx t_2$, then from 3.1(2), we have $t_i^A[x,\bar{d}\,] = t_i^B[y,\bar{d}\,]$ for $i = 1,2$. Hence

$$t_1^A[x,\bar{d}\,] = t_2^A[x,\bar{d}\,] \quad \text{iff} \quad t_1^B[y,\bar{d}\,] = t_2^B[y,\bar{d}\,],$$

giving $A, x \models \varphi[\bar{d}\,]$ iff $B, y \models \varphi[\bar{d}\,]$. If φ is the state equation $t_1 \simeq t_2$, then from 3.1(1), we have $t_i^A[x,\bar{d}\,]\,\rho\, t_i^B[y,\bar{d}\,]$ for $i = 1,2$. Hence by Corollary 3.3,

$$t_1^A[x,\bar{d}\,] \sim^A t_2^A[x,\bar{d}\,] \quad \text{iff} \quad t_1^B[y,\bar{d}\,] \sim^B t_2^B[y,\bar{d}\,],$$

which again gives the desired result. Thus x and y satisfy the same equations under $\bar{d}$. From this it follows readily that they satisfy the same Boolean combinations of equations, proving (1).

(2) follows directly from (1). For (3), suppose $A \models \varphi$ and ρ is surjective. Then for any $y \in X^B$ there is some $x \in X^A$ with $x\rho y$ and $A, x \models \varphi$, and so $B, y \models \varphi$ by (2). This shows $B \models \varphi$.

(4) is similar to (3), but using the fact that for each $x \in X^A$ there is some $y \in X^B$ with $x\rho y$. □

Application of a method causes a state transition $x \mapsto m^A(x,b)$ that may change the values of terms. Corollary 2.2 showed how the new values could be expressed in terms of values at the initial state x. This can now be extended to show how satisfaction of any formula by $m^A(x,b)$ can be characterised in terms of satisfaction by x. Again this works more generally for the transition $x \mapsto u^A[x]$ resulting from evaluation of any state term u. Writing $\varphi(u)$ for the formula obtained by substituting $u(\bar{v}\,)$ for σ in formula $\varphi(\bar{v}\,)$, we have:

Theorem 4.3

$$A, u^A[x,\bar{d}\,] \models \varphi[\bar{d}\,] \quad \textit{iff} \quad A, x \models \varphi(u)[\bar{d}\,]$$

for all $\bar{d}$ matching $\bar{v}$. Hence if u is ground,

$$A, u^A[x] \models \varphi \quad \textit{iff} \quad A, x \models \varphi(u).$$

Proof. If φ is the data equation $t_1 \approx t_2$, we want

$$\begin{array}{rcll} t_1^A[u^A[x,\bar{d}\,],\bar{d}\,] & = & t_2^A[u^A[x,\bar{d}\,],\bar{d}\,] & \quad \text{iff} \\ t_1(u)^A[x,\bar{d}\,] & = & t_2(u)^A[x,\bar{d}\,], & \end{array}$$

while if φ is the state equation $t_1 \simeq t_2$, we want the same condition but with $\sim$ in place of $=$. Both of these follow directly from Theorem 2.1, since this gives

$$t_i^A[u^A[x,\bar{d}\,],\bar{d}\,] = t_i(u)^A[x,\bar{d}\,]$$

for $i = 1,2$.

Thus the Theorem holds when φ is any equation. The inductive cases of the connectives are straightforward, using the fact that substitution respects these connectives, i.e. $(\neg\varphi)(u) = \neg(\varphi(u))$ etc. □

Corollary 4.4 $A, m^A(x,b) \models \varphi$ *iff* $A, x \models \varphi(m(\sigma,b))$. □

Associated with each formula φ is the set φ^A of all states in A that satisfy φ:

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

Relationships between some of these "truth sets" can be expressed by the standard Boolean set operations. For instance $(\varphi_1 \wedge \varphi_2)^A = \varphi_1^A \cap \varphi_2^A$ in general, and if φ is ground then $(\neg\varphi)^A = X^A - \varphi^A$.

Corollary 4.4 asserts that

$$m^A(x,b) \in \varphi^A \text{ iff } x \in \varphi(m(\sigma,b))^A. \tag{4.ii}$$

This fact can also be expressed by an operation on sets of states. For any method symbol m of sort I^m and any data element $b \in I^m$, define the function $[m.b]^A : \mathcal{P}(X^A) \to \mathcal{P}(X^A)$ (where $\mathcal{P}$ denotes powerset), by putting, for each $Y \subseteq X^A$,

$$[m.b]^A(Y) = \{x \in X^A : m^A(x,b) \in Y\}.$$

This operation preserves intersections, i.e.

$$[m.b]^A(Y_1 \cap Y_2) = [m.b]^A(Y_1) \cap [m.b]^A(Y_2),$$

and has $[m.b]^A(X^A) = X^A$. It thus has the properties needed to algebraically model a "box" modality in normal modal logic. In fact $[m.b]^A$ preserves set intersections and complements as well, becuase of the functionality (determinism) of m^A, and so models a rather special kind of modality with strong properties.

The operation $[m.b]^A$ allows result (4.ii) above to be given in the form

$$[m.b]^A(\varphi^A) = \varphi(m(\sigma,b))^A.$$

Thus the formula $\varphi(m(\sigma,b))$ expresses the modal assertion "after applying method m with parameter b, φ".

As a special case of this last equation we have the result

$$[m.b]^A(t \approx c)^A = (t(m(\sigma,b)) \approx c)^A \tag{4.iii}$$

for any data term t and constant c of the same type. The modal operators $[m.b]$ will be used in a new construction of coalgebras in Section 8, where equation (4.iii) plays an important role (see Theorem 8.1).

The article [Cor97] gives a complete deductive calculus for the valid observable equations over coalgebras of the kind considered here.

5 Morphisms and Subcoalgebras

A *morphism* from coalgebra A to coalgebra B is a function $f : X^A \to X^B$ between their state sets that preserves methods and attributes, meaning that for any state $x \in X^A$ the equation

$$f(m^A(x,b)) = m^B(f(x),b)$$

holds for all method symbols m and elements $b \in I^m$; while

$$a^A(x,c) = a^B(f(x),c)$$

holds for all attribute symbols a and $c \in I^a$. If f is surjective, then B is the *image* of A under this morphism. If f is bijective then it is an *isomorphism*, making A and B *isomorphic*.

When viewed as the relation $\{(x,y) : y = f(x)\}$, a morphism f is a bisimulation from A to B. Indeed a morphism is precisely a bisimulation that is *functional* and *total* (i.e. its domain is all of X^A). In particular the natural map $f^\natural(x) = |x|$ is a surjective morphism from A onto its quotient $A^\natural$, as defined at the end of Section 3.

If X^A is a subset of X^B, then A is called a *subcoalgebra* of B if the inclusion function $X^A \hookrightarrow X^B$ is a morphism. This holds precisely when

(a) each attribute a^A is the restriction of a^B to $X^A \times I^a$; and

(b) X^A is closed under all methods of B, i.e. $m^B(x,b) \in X^A$ whenever $x \in X^A$, and m^A is the restriction of m^B to $X^A \times I^m$.

Associated with any bisimulation ρ from A to B are two subcoalgebras: its domain and image. Whenever $x\rho y$ we have $m^A(x,b)\,\rho\, m^B(y,b)$, which puts $m^A(x,b)$ into the domain of ρ and $m^B(y,b)$ into the image. Thus the domain is closed under the methods of A, while the image is closed under the methods of B. Restricting the methods and attributes of A to the domain of ρ defines a subcoalgebra $\mathrm{Dom}\,\rho$ of A, while restricting the methods and attributes of B to the image of ρ defines a subcoalgebra $\rho(A)$ of B.

In the special case of a morphism $f : A \to B$ we get that the image $f(A)$ is a subcoalgebra of the codomain B. Moreover, if $f : A \to B$ is injective then it makes A isomorphic to the subcoalgebra $f(A)$.

Since a morphism is a bisimulation, the following results just reformulate Theorems 3.1 and 4.2 for morphisms.

Theorem 5.1 *Let $f : A \to B$ be a morphism. Then for any term $t(\bar{v})$, any formula $\varphi(\bar{v})$, any $\bar{d}$ matching $\bar{v}$, and any state $x \in X^A$:*

(1) $f(t^A[x,\bar{d}]) = t^B[f(x),\bar{d}]$ *if t is a state term;*

(2) $t^A[x,\bar{d}] = t^B[f(x),\bar{d}]$ *if t is a data term;*

(3) $A, x \models \varphi[\bar{d}]$ *iff* $B, f(x) \models \varphi[\bar{d}]$*;*

(4) $A, x \models \varphi$ *iff* $B, f(x) \models \varphi$*;*

(5) *If $A \models \varphi$, then $f(A) \models \varphi$;*

(6) *If $B \models \varphi$, then $A \models \varphi$.*

$\square$

An important special case of this is when A is a subcoalgebra of B. When f in Theorem 5.1 is the inclusion $X^A \hookrightarrow X^B$ having $f(x) = x$, we get the following.

Theorem 5.2 *Let A be a subcoalgebra of B. Then for any term $t(\bar{v})$, any formula $\varphi(\bar{v})$, any $\bar{d}$ matching $\bar{v}$, and any state $x \in X^A$:*

(1) $t^A[x, \bar{d}] = t^B[x, \bar{d}]$;

(2) $A, x \models \varphi[\bar{d}]$ *iff* $B, x \models \varphi[\bar{d}]$;

(3) $A, x \models \varphi$ *iff* $B, x \models \varphi$;

(4) *If* $B \models \varphi$, *then* $A \models \varphi$.

6 Disjoint Unions

If $\{A_j : j \in J\}$ is a collection of coalgebras, with

$$A_j = (X^{A_j}, \{m^{A_j} : m \in Meth\}, \{a^{A_j} : a \in Out\}),$$

then their *disjoint union* is the coalgebra

$$A = \coprod_J A_j = (\bigcup_J (X^{A_j} \times \{j\}), \{m^A : m \in Meth\}, \{a^A : a \in Out\})$$

where $m^A((x,j),b) = (m^{A_j}(x,b),j)$ and $a^A((x,j),c) = a^{A_j}(x,c)$.

Thus $\coprod_J A_j$ is the union of a set of pairwise disjoint copies $A_j \times \{j\}$ of the structures A_j. For each $k \in J$, the correspondence $x \mapsto (x,k)$ gives an injective morphism $A_k \to \coprod_J A_j$, whose image $A_k \times \{k\}$ is a subcoalgebra of $\coprod_J A_j$ isomorphic to A_k. In practice it is often convenient to identify this image with A_k, i.e. to regard the A_j's as being pairwise disjoint and $\coprod_J A_j$ as simply being their union. Then each A_j is itself a subcoalgebra of the disjoint union.

Theorem 6.1 *For any formula φ,*

$$\coprod_J A_j \models \varphi \text{ iff for all } j \in J,\ A_j \models \varphi.$$

Proof. Suppose every A_j is a model of φ. Then for any state (x,k) of the disjoint union we have $A_k, x \models \varphi$ and hence under the injective morphism $x \mapsto (x,k)$ we get $\coprod_J A_j, (x,k) \models \varphi$. This shows that $\coprod_J A_j \models \varphi$.

Conversely, if $\coprod_J A_j$ is a model of φ, then each A_j, being isomorphic to a subcoalgebra of the disjoint union, is also a model of φ (Theorem 5.2(4)). $\square$

7 Closure Properties

Let K be a class of coalgebras. The constructions we have defined provide several possible closure properties of K:

1. *Closure under images of bisimulations*: if $A \in K$ and there is a surjective bisimulation from A to B, then $B \in K$.

2. *Closure under domains of bisimulations*: if $B \in K$ and there is a total bisimulation from A to B, then $A \in K$.

3. *Closure under images of morphisms*: if $A \in K$ and there is a surjective morphism from A to B, then $B \in K$ (special case of 1).

4. *Closure under domains of morphisms*: if $B \in K$ and there is a morphism from A to B, then $A \in K$ (special case of 2).

5. *Closure under subcoalgebras*: if $B \in K$ and A is a subcoalgebra of B, then $A \in K$ (special case of 4, since the inclusion is a morphism from A to B).

6. *Closure under disjoint unions*: if $\{A_j : j \in J\} \subseteq K$ then $\coprod_J A_j \in K$.

The class $\mathit{Mod}\,\varphi = \{A : A \models \varphi\}$ of all models of *any* formula φ enjoys all of these closure properties, as we have seen from Theorems 4.2 and 6.1 (also 5.1 and 5.2 for the special cases). However these properties do not appear to be sufficient to ensure that a class K is of the form $\mathit{Mod}\,\varphi$, or even $\mathit{Mod}\,\Phi$ for some set Φ of formulas.

There are a number of interesting relationships between closure properties. The ones given next will be used in formulating the main result of this paper.

Theorem 7.1

(1) *K is closed under domains of bisimulations if and only if it is closed under images of bisimulations; and closed under domains of* surjective *bisimulations if and only if it is closed under images of* total *bisimulations.*

(2) *K is closed under images of bisimulations iff it is closed under both domains and images of morphisms.*

Proof.

(1) Coalgebra A is the domain of bisimulation ρ to B iff it is the image of bisimulation ρ^{-1} from B. Moreover ρ is surjective iff ρ^{-1} is total. (1) follows readily from these facts.

(2) If K is closed under images of bisimulations, then by (1) it is also closed under domains of bisimulations, and hence in particular closed under both domains and images of morphisms.

Conversely, assume that K is closed under domains and images of morphisms. Take $A \in K$ and let B be the image of a surjective bisimulation ρ from A. We want $B \in K$.

Now the subcoalgebra $\mathrm{Dom}\,\rho$ is the domain of the inclusion morphism to A, so belongs to K by closure under domains of morphisms. Let $f = \rho \circ f^{\natural}$ be the composition of ρ with the natural morphism $f^{\natural}(x) = |x|$ from B onto its quotient $B^\natural$. Then f is a bisimulation, being the composition of two bisimulations. Moreover f is functional, for if $x\rho y$ and $x\rho y'$, then $y \sim y'$ (Corollary 3.3), so $|y| = |y'|$. Thus f is a morphism from $\mathrm{Dom}\,\rho$ to $B^\natural$. Furthermore, f is surjective as ρ and $f^\natural$ are surjective. Closure of K under images of morphisms thus gives $B^\natural \in K$. But then $B \in K$ by closure under the domain of $f^\natural$.

□

8 Filter and Ultrafilter Enlargements

A *filter* on a coalgebra A is a non-empty collection F of subsets of the state set X^A that is closed under intersections and supersets of its members: if $Y, Z \in F$ then $Y \cap Z \in F$, and if $Y \subseteq Z \subseteq X^A$ and $Y \in F$ then $Z \in F$. Every filter contains X^A. F is *proper* if $\emptyset \notin F$.

An *ultrafilter* is a filter that is maximally proper, or equivalently that contains exactly one of Y and $X^A - Y$ for all $Y \subseteq X^A$. If an ultrafilter contains a union $Y \cup Z$, then it contains either Y or Z.

For every $G \subseteq \mathcal{P}(X^A)$ there is a smallest filter extending G, known as the filter *generated by* G. A set Z belongs to this generated filter iff there exist finitely many members $Y_0, \dots, Y_{k-1}$ of G with $Y_0 \cap \cdots \cap Y_{k-1} \subseteq Z$.

If $Y_0 \cap \cdots \cap Y_{k-1} \neq \emptyset$ whenever $Y_0, \dots, Y_{k-1} \in G$, then G has the *finite intersection property.* This is necessary and sufficient for the filter generated by G to be proper. Every G with the finite intersection property can be extended to an ultrafilter.

A proper filter F on A will be called *observationally rich*, or more briefly *rich*, if it satisfies the following condition:

> for any ground observable term t, of type O^t, there exists some $b \in O^t$ such that $(t \approx b)^A \in F$.

Recall that $(t \approx b)^A = \{x \in X^A : A, x \models t \approx b\} = \{x \in X^A : t^A[x] = b\}$.

The data element b corresponding to t in this condition is unique. For if $(t \approx b)^A$ and $(t \approx b')^A$ belong to F, then their intersection belongs to F and so is non-empty. Taking x as any member of this intersection gives $t^A[x] = b$ and $t^A[x] = b'$, hence $b = b'$.

It is notable that for a signature in which every output set is finite, all ultrafilters on A are rich. For if $O^t = \{b_1, \dots, b_n\}$, then

$$(t \approx b_1)^A \cup \cdots \cup (t \approx b_n)^A = X^A \in F,$$

so that if F is an ultrafilter then $(t \approx b_i)^A \in F$ for some $i \leq n$.

The *filter enlargement* of A is a new coalgebra A^+ whose state set $X^{A^+}$is the set of all rich proper filters of A. Its attributes are defined by

$$\begin{aligned} a^{A^+}(F, c) = b \quad &\text{iff} \quad (a(\sigma, c) \approx b)^A \in F \\ &\text{iff} \quad \{x \in X^A : a^A(x, c) = b\} \in F \end{aligned}$$

for all $F \in X^{A^+}$ and $c \in I^a$. Since $a(\sigma, c)$ is a ground observable term, the required data element $b \in O^a$ in this condition exists by the definition of "rich", and is unique as just explained.

The methods m^{A^+} of A^+ are defined with the help of the modal operators $[m.b]^A : \mathcal{P}(X^A) \to \mathcal{P}(X^A)$ introduced in Section 4, having

$$[m.b]^A(Y) = \{x \in X^A : m^A(x, b) \in Y\}.$$

For each $F \in X^{A^+}$ and $b \in I^m$, put

$$m^{A^+}(F, b) = \{Y \subseteq X^A : [m.b]^A(Y) \in F\}.$$

In other words,

$$Y \in m^{A^+}(F, b) \quad \text{iff} \quad \{x \in X^A : m^A(x, b) \in Y\} \in F.$$

Theorem 8.1 *$m^{A^+}(F, b)$ is a rich proper filter, and is an ultrafilter if F is one.*

Proof. $[m.b]^A$ has the following properties:

$$[m.b]^A(Y \cap Z) = [m.b]^A(Y) \cap [m.b]^A(Z),$$

$$Y \subseteq Z \text{ implies } [m.b]^A(Y) \subseteq [m.b]^A(Z),$$

$$[m.b]^A(X^A) = X^A,$$

$$[m.b]^A(\emptyset) = \emptyset.$$

These, together with the fact that F is a proper filter, ensure that $m^{A^+}(F, b)$ is a proper filter. The first gives closure under intersections; the second closure under supersets; the third guarantees that $m^{A^+}(F, b)$ is non-empty (as $X^A \in F$); and the fourth implies $m^{A^+}(F, b)$ is proper.

To show it is rich, take any ground observable term t of type O^t. Then $t(m(\sigma, b))$ is also a ground observable term of type O^t, so as F is rich there is some $c \in O^t$ such that $(t(m(\sigma, b)) \approx c)^A \in F$. But by equation (4.iii),

$$(t(m(\sigma, b)) \approx c)^A = [m.b]^A(t \approx c)^A,$$

so this implies that $(t \approx c)^A \in m^{A^+}(F, b)$, completing the proof that $m^{A^+}(F, b)$ is rich.

Finally, if F is an ultrafilter, the fact that

$$[m.b]^A(Y \cup Z) = [m.b]^A(Y) \cup [m.b]^A(Z)$$

can be used to show that if $Y \cup Z \in m^{A^+}(F, b)$, then either Y or Z is in $m^{A^+}(F, b)$, proving that $m^{A^+}(F, b)$ is an ultrafilter. □

Now let X^{A^*} be the set of rich *ultra*filters of A. The last part of Theorem 8.1 shows that X^{A^*} is closed under the methods m^{A^+}. Restricting all the methods and attributes of A^+ to X^{A^*} thus defines a subcoalgebra A^* of A^+ which will be called the *ultrafilter enlargement* of A.

Theorem 8.2

(1) *A is isomorphic to a subcoalgebra of A^*.*

(2) *There is a total surjective bisimulation from A^* to A^+.*

(3) *For any formula φ, $A^* \models \varphi$ iff $A^+ \models \varphi$.*

Proof.

(1) For each $x \in X^A$, the *principal* ultrafilter $F_x = \{Y \subseteq X^A : x \in Y\}$ is rich. Indeed if $b = t^A[x]$, then $(t \approx b)^A \in F_x$. The map $x \mapsto F_x$ is an injection of X^A into X^{A^*} that proves to be a morphism. The fact that it preserves methods follows because $Y \in m^{A^*}(F_x, b)$ iff $[m.b]^A(Y) \in F_x$ iff $x \in [m.b]^A(Y)$ iff $m^A(x,b) \in Y$ iff $Y \in F_{m^A(x,b)}$, which shows that $m^A(x,b) \mapsto m^{A^*}(F_x, b)$. Also if $b = a^A(x,c)$, then $(a(\sigma,c) \approx b)^A \in F_x$, so $a^{A^*}(F_x, c) = b$, showing that the injection preserves attributes.

(2) The claimed bisimulation is the *superset* relation. For $F \in X^{A^*}$ and $G \in X^{A^+}$, define $F\rho G$ iff $F \supseteq G$. Now suppose that $F\rho G$. Then for any $Y \in m^{A^+}(G,b)$, the set $[m.b]^A(Y)$ belongs to G and therefore to F, hence $Y \in m^{A^+}(F,b) = m^{A^*}(F,b)$. This shows that $m^{A^*}(F,b)\,\rho\, m^{A^+}(G,b)$. Moreover, if $a^{A^+}(G,c) = b$, then $(a(\sigma,c) \approx b)^A \in G \subseteq F$, so $a^{A^*}(F,c) = b = a^{A^+}(G,c)$. Altogether this establishes that ρ is a bisimulation. The fact that $F\rho F$ for any $F \in X^{A^*}$ makes it immediate that ρ is total. Finally, to show that ρ is surjective, take any $G \in X^{A^+}$: since G is a proper filter on X^A it can be extended to an ultrafilter F on X^A. This F will automatically be rich because G is rich, so $F \in X^{A^*}$ and $F\rho G$.

(3) If $A^+ \models \varphi$, then by Theorem 5.2(4) $A^* \models \varphi$, as A^* is subcoalgebra of A^+ by definition. Conversely, if $A^* \models \varphi$, then by Theorem 4.2(3) $A^+ \models \varphi$, as A^+ is the image of a bisimulation from A^* by (2).

□

Which model classes are closed under filter or ultrafilter enlargements? To explore that we need to investigate how the value of a term at state F in A^+ depends on the internal properties of the filter F.

Theorem 8.3 *Let $t(\bar{v})$ be any term, $F \in X^{A^+}$, and $\bar{d}$ match $\bar{v}$.*

(1) *If t is a state term, then for all $Y \subseteq X^A$,*

$$Y \in t^{A^+}[F, \bar{d}] \quad \textit{iff} \quad \{x \in X^A : t^A[x, \bar{d}] \in Y\} \in F.$$

(2) *If t is a data term, then*

$$\{x \in X^A : t^A[x, \bar{d}] = t^{A^+}[F, \bar{d}]\} \in F.$$

Proof. By induction on the length of t, taking the cases (1) and (2) together. When t is a data term (case (2)), we will let

$$Y_t = \{x \in X^A : t^A[x, \bar{d}\,] = t^{A^+}[F, \bar{d}\,]\,\}.$$

Now if t is σ, then $t^{A^+}[F, \bar{d}\,] = F$ and

$$\{x \in X^A : t^A[x, \bar{d}\,] \in Y\} = \{x : x \in Y\} = Y,$$

so (1) becomes the tautology $Y \in F$ iff $Y \in F$.

If t is a variable v_i or a constant b, then $t^A[x, \bar{d}\,]$ and $t^{A^+}[F, \bar{d}\,]$ are both either d_i or b, independently of state x, so $Y_t = X^A \in F$ and (2) holds.

Now let t be $m(t_1, t_2)$ and make the induction hypothesis that the Theorem holds for the state term t_1 and the data term t_2. For any $Y \subseteq X^A$, we have

$$Y \in m(t_1, t_2)^{A^+}[F, \bar{d}\,] = m^{A^+}(t_1^{A^+}[F, \bar{d}\,], t_2^{A^+}[F, \bar{d}\,])$$

iff (by definition of m^{A^+}) the set

$$Z = \{x \in X^A : m^A(x, t_2^{A^+}[F, \bar{d}\,]) \in Y\}$$

belongs to $t_1^{A^+}[F, \bar{d}\,]$. By the hypothesis (1) on t_1, this in turn holds iff

$$W_1 = \{x \in X^A : t_1^A[x, \bar{d}\,] \in Z\} \in F.$$

We have to show that this is equivalent to the statement

$$W_2 = \{x \in X^A : m(t_1, t_2)^A[x, \bar{d}\,] \in Y\} \in F.$$

Now

$$x \in W_1 \quad \text{iff} \quad m^A(t_1^A[x, \bar{d}\,], t_2^{A^+}[F, \bar{d}\,]) \in Y,$$

while

$$x \in W_2 \quad \text{iff} \quad m^A(t_1^A[x, \bar{d}\,], t_2^A[x, \bar{d}\,]) \in Y.$$

Since

$$x \in Y_{t_2} \quad \text{iff} \quad t_2^A[x, \bar{d}\,] = t_2^{A^+}[F, \bar{d}\,],$$

we have $W_1 \cap Y_{t_2} = W_2 \cap Y_{t_2}$. But $Y_{t_2} \in F$ by hypothesis (2) on t_2, and so as F is a filter, $W_1 \in F$ iff $W_2 \in F$, which completes the proof of (2) when t is $m(t_1, t_2)$.

The last case is where t is $a(t_1, t_2)$, with the induction hypothesis that the Theorem holds for t_1 and t_2. Let $b = a(t_1, t_2)^{A^+}[F, \bar{d}\,]$. Then

$$a^{A^+}(t_1^{A^+}[F, \bar{d}\,], t_2^{A^+}[F, \bar{d}\,]) = b,$$

so by definition of a^{A^+},

$$Z = \{x \in X^A : a^A(x, t_2^{A^+}[F, \bar{d}\,]) = b\} \in t_1^{A^+}[F, \bar{d}\,].$$

The hypothesis (1) on t_1 then gives

$$W = \{x \in X^A : t_1^A[x, \bar{d}\,] \in Z\} \in F.$$

But $x \in W$ iff $a^A(t_1^A[x, \bar{d}\,], t_2^{A^+}[F, \bar{d}\,]) = b$. Hence,

$$\begin{aligned} W \cap Y_{t_2} &\subseteq \{x \in X^A : a^A(t_1^A[x, \bar{d}\,], t_2^A[x, \bar{d}\,]) = b\} \\ &= \{x \in X^A : a(t_1, t_2)^A[x, \bar{d}\,] = a(t_1, t_2)^{A^+}[F, \bar{d}\,]\} \\ &= Y_{a(t_1,t_2)}. \end{aligned}$$

But $Y_{t_2} \in F$ by hypothesis (2) on t_2, so then $Y_{a(t_1,t_2)} \in F$, proving (2) when t is $a(t_1, t_2)$. □

Theorem 8.4 *Let $t_1(\bar{v}\,), t_2(\bar{v}\,)$ be terms of the same type, $F \in X^{A^+}$, and $\bar{d}$ match $\bar{v}$. Then*

$$\{x \in X^A : t_1^A[x, \bar{d}\,] = t_2^A[x, \bar{d}\,]\} \in F \quad \textit{implies} \quad t_1^{A^+}[F, \bar{d}\,] = t_2^{A^+}[F, \bar{d}\,].$$

Moreover, if t_1, t_2 are data terms, the converse is true.

Proof. Let $Z_{12} = \{x \in X^A : t_1^A[x, \bar{d}\,] = t_2^A[x, \bar{d}\,]\}$. If t_1, t_2 are of data type, then for $i = 1, 2$ put

$$Z_i = \{x \in X^A : t_i^A[x, \bar{d}\,] = t_i^{A^+}[F, \bar{d}\,]\}.$$

By Theorem 8.3(2), $Z_1, Z_2 \in F$. Thus if $Z_{12} \in F$, then the intersection $Z_1 \cap Z_{12} \cap Z_2$ belongs to F, so is non-empty. If x_0 is any member of this intersection,

$$t_1^{A^+}[F, \bar{d}\,] = t_1^A[x_0, \bar{d}\,] = t_2^A[x_0, \bar{d}\,] = t_2^{A^+}[F, \bar{d}\,],$$

giving the required result. Conversely, if $t_1^{A^+}[F, \bar{d}\,] = t_2^{A^+}[F, \bar{d}\,]$, then

$$Z_1 \cap Z_2 \subseteq Z_{12},$$

and so Z_{12} belongs to the filter F, because $Z_1, Z_2 \in F$.

Now for the case where t_1, t_2 are of state type. For each $Y \subseteq X^A$, put

$$Y_i = \{x \in X^A : t_i^A[x, \bar{d}\,] \in Y\}$$

for $i = 1, 2$. Then $Y_1 \cap Z_{12} = Y_2 \cap Z_{12}$. Thus if $Z_{12} \in F$, we get $Y_1 \in F$ iff $Y_2 \in F$, which by Theorem 8.3(1) means

$$Y \in t_1^{A^+}[F, \bar{d}\,] \quad \text{iff} \quad Y \in t_2^{A^+}[F, \bar{d}\,].$$

Since this holds for all $Y \subseteq X^A$, $t_1^{A^+}[F, \bar{d}\,] = t_2^{A^+}[F, \bar{d}\,]$. □

Corollary 8.5 *If u_1, u_2 are state terms, then*

$$\{x \in X^A : A, x \models u_1 \simeq u_2[\bar{d}]\} \in F \quad \textit{implies} \quad A^+, F \models u_1 \simeq u_2[\bar{d}].$$

Proof. Let $Y = \{x \in X^A : A, x \models u_1 \simeq u_2[\bar{d}]\}$. Assume $Y \in F$. If t is any ground observable term, then for all $x \in Y$, $A, x \models t(u_1) \approx t(u_2)[\bar{d}]$ by Theorem 4.1. So the set $\{x \in X^A : A, x \models t(u_1) \approx t(u_2)[\bar{d}]\}$ includes Y and therefore belongs to F. Theorem 8.4 then yields $A^+, F \models t(u_1) \approx t(u_2)[\bar{d}]$. Since this holds for all ground observable t, Theorem 4.1 again gives $A^+, F \models u_1 \simeq u_2[\bar{d}]$. □

The converse of this result is not true. In Section 10 there is an example of a coalgebra A that validates $\neg(m(\sigma) \simeq \sigma)$ while A^+, and indeed A^*, has states F for which $m^{A^*}(F) \sim F$. This shows that the model class *Mod* $\neg\varphi$ of a negated formula need not be closed under enlargements. We therefore look more closely at *positive* formulas, which by definition are those built from equations (of either type) using only conjunction $\wedge$ and disjunction $\vee$.

Theorem 8.6 *Let F be a rich ultrafilter of coalgebra A.*

(1) *If φ is any ground data formula, then*

$$\varphi^A \in F \quad \textit{iff} \quad A^*, F \models \varphi \quad \textit{iff} \quad A^+, F \models \varphi.$$

(2) *If φ is any ground positive formula, then*

$$\varphi^A \in F \quad \textit{implies} \quad A^*, F \models \varphi.$$

Proof.

(1) Since A^* is a subcoalgebra of A^+ and $F \in X^{A^*}$, $A^*, F \models \varphi$ iff $A^+, F \models \varphi$ in general by Theorem 5.2(2). Thus it suffices to show $\varphi^A \in F$ iff $A^+, F \models \varphi$.

If t_1, t_2 are ground data terms, then Theorem 8.4 states that $(t_1 \approx t_2)^A \in F$ iff $A^+, F \models (t_1 \approx t_2)$, so the desired result holds when φ is a data equation.

If the result holds for ground formulas φ_1, φ_2, then $(\varphi_1 \wedge \varphi_2)^A = \varphi_1^A \cap \varphi_2^A \in F$ iff $\varphi_1^A, \varphi_2^A \in F$ iff $A^+, F \models \varphi_1 \wedge \varphi_2$. Also $\neg\varphi_1^A = (X^A - \varphi_1^A) \in F$ iff $\varphi_1^A \notin F$ (as F is an ultrafilter) iff not $A^+, F \models \varphi_1$ iff $A^+, F \models \neg\varphi_1$ (as φ_1 is ground). Hence the result holds for $\varphi_1 \wedge \varphi_2$ and for $\neg\varphi_1$. This is enough to ensure that it holds for all Boolean combinations of ground data equations.

(2) If φ is a state equation, then (2) is given by Corollary 8.5. If φ is a data equation, (2) is a specialisation of (1).

The induction case of φ being a conjunction $(\varphi_1 \wedge \varphi_2)$ holds as for (1). Finally, suppose φ is a disjunction $(\varphi_1 \vee \varphi_2)$: if $(\varphi_1 \vee \varphi_2)^A = \varphi_1^A \cup \varphi_2^A \in F$, then as F is an ultrafilter one of φ_1^A, φ_2^A belongs to F, so $A^*, F \models \varphi_i$ for some i by induction hypothesis, hence $A^*, F \models \varphi_1 \vee \varphi_2$.

□

Corollary 8.7 *If φ is a data formula, or a positive formula, then*

$$A \models \varphi \quad \textit{iff} \quad A^* \models \varphi \quad \textit{iff} \quad A^+ \models \varphi.$$

Proof. First note that $A^* \models \varphi$ iff $A^+ \models \varphi$ by Theorem 8.2(3), and that $A^* \models \varphi$ implies $A \models \varphi$ in general, as A is isomorphic to a subcoalgebra of A^* (Theorem 8.2(1)). Thus it suffices to show that $A \models \varphi$ implies $A^* \models \varphi$.

Suppose that φ is a ground data or positive formula that is valid in A. Then $\varphi^A = X^A$, so for any $F \in X^{A^*}$, $\varphi^A \in F$, implying $A^*, F \models \varphi$ by Theorem 8.6. Hence A^* is a model of φ.

Thus the result holds for all ground data or positive formulas. It then holds for all such formulas by the reduction of satisfaction of formulas to satisfaction of ground formulas according to equation (4.i). □

Preservation of validity by enlargements extends beyond data formulas and positive formulas to a limited extent by the consideration of certain kinds of *conditional* formulas:

Theorem 8.8 *If φ is any formula of the form*

$$\varphi_0 \wedge \cdots \wedge \varphi_{k-1} \to \psi,$$

where the φ_i's are any data formulas and ψ is any positive formula, then,

$$A \models \varphi \quad \textit{iff} \quad A^* \models \varphi \quad \textit{iff} \quad A^+ \models \varphi.$$

Proof. As in the proof of Corollary 8.7, it is enough to consider *ground* φ, and to show that $A \models \varphi$ implies $A^* \models \varphi$. But for ground φ, if $A \models \varphi$, then

$$\varphi_0^A \cap \cdots \cap \varphi_{k-1}^A \subseteq \psi^A,$$

so if $A^*, F \models \varphi_i$ for all $i < k$, Theorem 8.6(1) gives $\varphi_i^A \in F$ for all $i < k$, implying $\psi^A \in F$ and so $A^*, F \models \psi$ by 8.6(2). This shows that $A^*, F \models \varphi$ for all states F in X^{A^*}. □

To conclude this section, here are some results that further illustrate the behaviour of enlargements of coalgebras.

Theorem 8.9 *For any bisimulation ρ from A to B there is a bisimulation ρ^+ from A^+ to B such that*

(1) *ρ^+ extends ρ via the injective morphism $x \mapsto F_x$ from A to A^*.*

(2) *If ρ is functional, then so is ρ^+.*

(3) *If B has finitely many states and ρ is total, then the domain of ρ^+ includes all states of A^*.*

Proof. For any $y \in X^B$, let $\rho \restriction y = \{x \in X^A : x\rho y\}$. For $F \in X^{A^+}$, put

$$F\,\rho^+\,y \quad \text{iff} \quad \rho \restriction y \in F.$$

First we show that methods preserve ρ^+-relatedness. Suppose $F\,\rho^+\,y$. Then $\rho \restriction y \in F$. Since methods preserve ρ-relatedness, for any $m \in Meth$ and $b \in I^m$,

$$\begin{aligned} \rho \restriction y &\subseteq \{x \in X^A : m^A(x,b)\,\rho\,m^B(y,b)\} \\ &= \{x \in X^A : m^A(x,b) \in \rho \restriction m^B(y,b)\}, \end{aligned}$$

so $\{x \in X^A : m^A(x,b) \in \rho \restriction m^B(y,b)\} \in F$. This implies $\rho \restriction m^B(y,b) \in m^{A^+}(F,b)$, and therefore $m^{A^+}(F,b)\,\rho^+\,m^A(y,b)$.

Next, ρ^+ preserve attributes: since ρ preserves attributes,

$$\{x : x\rho y\} \subseteq \{x : a^A(x,c) = a^B(y,c)\}.$$

Thus if $F\,\rho^+\,y$, $\{x : x\rho y\} \in F$, so for any $c \in I^a$, $\{x : a^A(x,c) = a^B(y,c)\} \in F$, which implies $a^{A^+}(F,c) = a^B(y,c)$. Hence ρ^+ is a bisimulation.

(1) If $x\rho y$, then $x \in \rho \restriction y$, so $\rho \restriction y \in F_x$, showing $F_x\,\rho^+\,y$. In this sense ρ^+ extends ρ via the injective morphism $x \mapsto F_x$ of Theorem 8.2(1).

(2) If $F\,\rho^+\,y$ and $F\,\rho^+\,z$, then the intersection $(\rho \restriction y) \cap (\rho \restriction z)$ belongs to F, so is nonempty. If x belongs to this intersection then $x\rho y$ and $x\rho z$, so if ρ is functional then $y = z$, showing ρ^+ is functional.

(3) If ρ is total, then for any filter F, $\bigcup_{y \in X^B}(\rho \restriction y) = X^A \in F$. If X^B is finite, this is a finite union, so then if F is an ultrafilter there is some $y \in X^B$ with $\rho \restriction y \in F$, hence $F\,\rho^+\,y$. So in this case the domain of ρ^+ includes all members of X^{A^*}.

□

Now let K_B be the class of all domains of bisimulations to coalgebra B. From the proof of Theorem 7.1(1), K_B can also be described as the class of all images of bisimulations from B.

Corollary 8.10 *K_B is closed under disjoint unions and images of bisimulations. If B is finite then K_B is closed under ultrafilter and filter enlargements.*

Proof. Suppose $\{A_j : j \in J\} \subseteq K_B$. Then for each $j \in J$ there is a total bisimulation ρ_j from A_j to B. Define $(x,j)\rho y$ iff $x\rho_j y$. Then ρ is a total bisimulation from $\coprod_J A_j$ to B, showing K_B is closed under disjoint unions.

Next suppose $A \in K_B$, with ρ a total bisimulation from A to B. Let coalgebra C be the image of a bisimulation τ from A. Then $\tau^{-1} \circ \rho$ is a bisimulation from C to B whose domain is C itself since τ is surjective. Hence $C \in K_B$, showing K_B is closed under images of bisimulations.

Now if B is finite, then any total bisimulation ρ from A to B lifts by Theorem 8.9 to give a total bisimulation from A^* to B. In other words, if $A \in K_B$ then

$A^* \in K_B$. But A^+ is the image of a bisimulation from A^* (Theorem 8.2(2)), so if $A^* \in K_B$ then $A^+ \in K_B$. $\square$

Ultrafilter enlargements are similar to the *canonical extensions* (also called *ultrafilter* extensions) of Kripke frames for propositional modal languages. Briefly, a Kripke frame for monomodal logic is a pair $A = (X^A, R^A)$ with R^A a binary relation on set X^A. For a language with several modal connectives, a frame has such a relation R^A for each connective. R^A may be identified with the function $x \mapsto \{y : xR^A y\}$ from X^A to the powerset $\mathcal{P}(X^A)$, so a frame as defined can be viewed as a coalgebra for the powerset endofunctor $\mathcal{P} : \mathbf{Set} \to \mathbf{Set}$. The *canonical extension* of frame A is the frame A^* whose points are the ultrafilters on X^A, with $FR^{A^*}G$ iff $\{Y : [R]^A(Y) \in F\} \subseteq G$, where $[R]^A(Y) = \{x \in X^A : \{y : xR^A y\} \subseteq Y\}$. Thus the ultrafilter enlargement of a coalgebra is essentially an instance of this modal construction as far as *methods* are concerned, as may be seen by identifying method m^A with the family of relations $R^A_b = \{\langle x, y\rangle : y = m^A(x, b)\}$ for all $b \in I^m$. But the presence of attributes adds additional complexity which requires the introduction of the new notion of *rich* ultrafilter, a notion that has its own intricate analysis as we have seen in Theorem 8.3 and its consequences. Even when all output sets of a signature are finite, so that all ultrafilters are rich, attributes constitute a new feature not found in the standard modal context.

The fact that ultrafilter enlargements of coalgebras preserve validity of observable formulas is also a point of distinction with the general modal propositional case. Languages for the latter have variables taking arbitrary subsets of X^A as values, and have formulas whose validity is not preserved by canonical extensions of Kripke frames. Further discussion of these comparisions with modal semantics is given in Section 7 of [Gol00].

9 Definability of Classes of Coalgebras

A *conditional equation* is any formula of the form

$$\varepsilon_0 \wedge \cdots \wedge \varepsilon_{k-1} \to \varepsilon,$$

where the ε_i's and ε are all equations. $\varepsilon_0, \cdots, \varepsilon_{k-1}$ are called the *premisses*.

Note that ε could be the equation $b \approx c$ where b, c are distinct elements of the same observable type (it was assumed at the outset that two such elements exist). In that case ε is never satisfied by any state, and the conditional equation is equivalent to the formula $\neg(\varepsilon_0 \wedge \cdots \wedge \varepsilon_{k-1})$.

Theorem 9.1 *Suppose that every ground observable conditional equation valid in coalgebra A is valid in coalgebra B. Then B is the image of a bisimulation from A^+.*

Proof. We show that the bisimilarity relation $\sim$ from A^+ to B is surjective (indeed if any bisimulation from A^+ to B is surjective, then $\sim$, as the largest such bisimulation, must be surjective).

Let y be any state of B. We must find a state F of A^+ with $F \sim y$. This is done by applying the characterisation of bisimilarity given in Theorem 3.2.

Define Φ_y to be the set of of all ground observable equations ε such that $B, y \models \varepsilon$. Let F be the filter on X^A generated by the collection

$$F_y = \{\varepsilon^A : \varepsilon \in \Phi_y\}.$$

Then for any ground observable equation ε we have

$$\varepsilon^A \in F \quad \text{iff} \quad \varepsilon \in \Phi_y. \tag{9.i}$$

From right to left this holds by definition. To see that the converse also holds, let $\varepsilon^A \in F$. Then since F is generated by F_y there exist $\varepsilon_0, \cdots, \varepsilon_{k-1} \in \Phi_y$ such that

$$\varepsilon_0^A \cap \cdots \cap \varepsilon_{k-1}^A \subseteq \varepsilon^A.$$

This means that the conditional equation $\varepsilon_0 \wedge \cdots \wedge \varepsilon_{k-1} \to \varepsilon$ is valid in A, and so by hypothesis is valid in B. But $B, y \models \varepsilon_i$ for all $i < k$, so $B, y \models \varepsilon$, giving $\varepsilon \in \Phi_y$ as desired.

(9.i) allows us to prove that F is proper and rich. First, to show it is proper let b, c be two distinct observable data elements of the same type. Then the equation $b \approx c$ does not belong to Φ_y, and hence $\emptyset = (b \approx c)^A \notin F$. Next, let t be any ground observable term of type O^t. Put $b = t^B[y]$. Then $(t \approx b) \in \Phi_y$, so $(t \approx b)^A \in F_y \subseteq F$. Hence F is rich.

Thus F is a state of A^+. (9.i) can be written as

$$A^+, F \models \varepsilon \quad \text{iff} \quad B, y \models \varepsilon$$

(using Theorem 8.4). Hence F and y satisfy the same ground observable equations, which implies, by Theorem 3.2(4), that $F \sim y$. $\square$

We are now ready to give the main result of this paper.

Theorem 9.2 *For any class K of coalgebras, the following are equivalent.*

(1) *K is the class of all models of some set of ground observable conditional equations.*

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

(3) *K is the class of all models of some set of observable formulas.*

(4) *K is closed under disjoint unions, domains and images of morphisms, and ultrafilter enlargements.*

(5) *K is closed under disjoint unions, images of bisimulations, and ultrafilter enlargements.*

(6) *K is closed under disjoint unions, images of bisimulations, and filter enlargements.*

Proof. By definition, (1) implies (2) and (2) implies (3). (3) implies (4) by Theorems 6.1 and 5.1 and Corollary 8.7. (4) implies (5) by Theorem 7.1(2). (5) implies (6) by Theorem 8.2(2), which shows that if $A^* \in K$ and K is closed under images of bisimulations, then $A^+ \in K$.

Now suppose (6) holds. To prove (1), let Φ be the set of all ground observable conditional equations that have all members of K as models. Then $K \subseteq Mod\,\Phi$ by definition. The heart of the matter is to prove the converse inclusion.

Let B be any member of $Mod\,\Phi$. Take Ψ to be the set of all ground observable conditional equations that are *not* valid in B. For each $\varphi \in \Psi$ there must be some coalgebra A_φ in K that is not a model of φ, or else φ would belong to Φ, implying $B \models \varphi$ as $B \in Mod\,\Phi$, contradicting $\varphi \in \Psi$. Let $A = \coprod_{\varphi \in \Psi} A_\varphi$ be the disjoint union of these coalgebras $A_\varphi \in K$.

Now $A \in K$ as K is closed under disjoint unions. Moreover each $\varphi \in \Psi$ is invalid in A, by Theorem 6.1, since it is invalid in A_φ. Hence any ground observable conditional equations valid in A is not in Ψ, and so is valid in B. Therefore by Theorem 9.1 there is a surjective bisimulation from A^+ onto B. But $A^+ \in K$ by closure of K under filter enlargements, and so $B \in K$ by closure under images of bisimulations.

This proves that $K = Mod\,\Phi$, and so (1) holds.

□

Recall that K_B is the class of all domains of bisimulations to B, or equivalently, images of bisimulations from B. For finite B, Theorem 8.10 leads to a logical characterisation of K_B. Let Φ_B be the set of all observable conditional equations valid in B.

Theorem 9.3 *If B is finite, then $K_B = Mod\,\Phi_B$.*

Proof. By 8.10, K_B is closed under disjoint unions, images of morphisms, and filter enlargements, so by the proof of Theorem 9.2, $K_B = Mod\,\Phi$, where Φ is the set of all observable conditional equations valid in all members of K_B. But $\Phi = \Phi_B$, because a formula is valid in B iff it is valid in the domain of all bisimulations to B. □

If φ is a conditional formula of the form

$$\varphi_0 \wedge \cdots \wedge \varphi_{k-1} \to \psi$$

described in Theorem 8.8, where the φ_i's are data formulas and ψ is positive, then $Mod\,\varphi$ is closed under disjoint unions, images of bisimulations and (ultra)filter enlargements. Hence by Theorem 9.2 φ is equivalent to some set Φ_φ of observable formulas, in the sense that φ and Φ_φ have the same models. But it is not necessary to appeal to 9.2 to reach this conclusion: an explicit syntactic definition of Φ_φ can be read off from the form of φ, giving also the stronger conclusion that

$$A, x \models \varphi[\bar{d}] \quad \text{iff} \quad \text{for all } \chi \in \Phi_\varphi,\ A, x \models \chi[\bar{d}].$$

First of all Φ_ψ is defined inductively for all positive ψ as in the following table, with the base case that ψ is a state equation being justified by Theorem 4.1.

ψ	Φ_ψ
$t_1 \approx t_2$	$\{t_1 \approx t_2\}$
$u_1 \simeq u_2$	$\{t(u_1) \approx t(u_2) : t \text{ is observable}\}$
$\psi_1 \wedge \psi_2$	$\Phi_{\psi_1} \cup \Phi_{\psi_2}$
$\psi_1 \vee \psi_2$	$\{\chi_1 \vee \chi_2 : \chi_1 \in \Phi_{\psi_1} \text{ and } \chi_2 \in \Phi_{\psi_2}\}$

Then for conditional φ as above, Φ_φ is defined to be the set

$$\{\varphi_0 \wedge \cdots \wedge \varphi_{k-1} \to \chi : \chi \in \Phi_\psi\}.$$

However, Theorem 9.2 allows the stronger conclusion that φ is equivalent to some set Φ'_φ of observable *conditional equations*. It would be interesting to know if a syntactic definition of such a set Φ'_φ can be given in similar fashion to that for Φ_φ.

10 Four Counter-Examples

The first example is a class K of coalgebras that is closed under disjoint unions and images of bisimulations, hence under domains and images of morphisms, but is not closed under ultrafilter enlargements. Consequently K is not closed under filter enlargements either, since A^* is the domain of the inclusion morphism to A^+.

The signature involved has two method symbols m and e, and one attribute symbol a. All three have no input sort, while a has output sort $\{0,1\}$. Thus in any coalgebra A, m^A and e^A are functions of the type $X^A \to X^A$, while a^A is of type $X^A \to \{0,1\}$.

Let $\omega = \{0,1,2,\ldots\}$ be the set of natural numbers. For any function f with the same domain and codomain, and any $k \in \omega$, f^k is the k-fold iteration of f, i.e. $f^0(x) = x$ and $f^{k+1} = f(f^k(x))$.

K is defined to be the class of all coalgebras in which any application of e^A is bisimilar to some finite number of iterations of m^A:

$$\forall x \in X^A \;\; \exists k \in \omega \;\; (m^A)^k(x) \sim^A e^A(x). \tag{10.i}$$

Closure of K under disjoint unions is straightforward: if A is the disjoint union of some coalgebras A_j from K, and x is a state of A, we may suppose that x is a state of some A_j, and A_j is itself a subcoalgebra of A. Then as $A_j \in K$, there exists $k \in \omega$ with $(m^{A_j})^k(x) \sim^{A_j} e^{A_j}(x)$. But A_j is closed under the methods of A, so $(m^{A_j})^k(x) = (m^A)^k(x)$ and $(e^{A_j})^k(x) = (e^A)^k(x)$, implying $(m^A)^k(x) \sim^A e^A(x)$.

For closure under images of bisimulations, let ρ be a surjective bisimulation from A to B, with A in K. Then if $y \in X^B$, there is some $x \in X^A$ with

$x\rho y$. By (10.i) there exists k with $(m^A)^k(x) \sim^A e^A(x)$. Since methods preserve ρ-relatedness, this implies $(m^A)^k(x)\,\rho\,(m^B)^k(y)$ and $e^A(x)\,\rho\,e^B(y)$ (see 3.1(1)). Hence $(m^B)^k(y) \sim^B e^B(y)$ by Corollary 3.3. This shows that (10.i) holds for B, and so $B \in K$.

To show K is not closed under ultrafilter enlargements, we construct a particular coalgebra A with $X^A = \omega$.

(i) $m^A : \omega \to \omega$ is the successor function $m^A(x) = x + 1$.

Hence in general $(m^A)^k(x) = x + k$.

(ii) $e^A : \omega \to \omega$ is any function having $e^A(x+1) > e^A(x) + x$ for all $x \in \omega$, and consequently $e^A(x) \geq x$.

For example, take $e^A(x) = \sum_{i=0}^{x} i$.

(iii) Write Y_e for the image $\{e^A(x) : x \in \omega\}$ of e^A. Then $a^A : \omega \to \{0,1\}$ is defined to be the characteristic function of Y_e: $a^A(y) = 1$ iff $y \in Y_e$. Thus

$$(a(\sigma) \approx 1)^A = \{x \in \omega : a^A(x) = 1\} = Y_e.$$

It is easy to see that $A \in K$: since $e^A(x) \geq x$ there exists k with $e^A(x) = x + k = (m^A)^k(x)$, hence $e^A(x) \sim (m^A)^k(x)$.

Now to show that $A^* \notin K$. Informally the idea is that e grows so fast on ω that eventually in A^* it takes a value that cannot be obtained by any finite number of iterations of m. For each $Y \subseteq \omega$, let

$$Y - k = \{y - k : k \leq y \in Y\} = \{x \in \omega : x + k \in Y\}.$$

Then for any $F \in X^{A^*}$, by definition of m^{A^*}

$$Y \in m^{A^*}(F) \quad \text{iff} \quad \{x \in X^A : x + 1 \in Y\} = Y - 1 \in F.$$

From this it follows that for $k \in \omega$,

$$(m^{A^*})^k(F) = \{Y \subseteq \omega : Y - k \in F\}. \tag{10.ii}$$

Now let $Y_k = \{x \in \omega : x + k \notin Y_e\}$, the complement of $Y_e - k$ in ω. Then the collection $\{Y_k : k \in \omega\}$ has the finite intersection property. To see this, take any $k_1, \dots, k_n \in \omega$. Choose any $x > k_1, \dots, k_n$. Then from the assumed properties of e^A we get that for each $j \leq n$,

$$e^A(x) < e^A(x) + 1 + k_j \leq e^A(x) + x < e^A(x+1).$$

Since e^A is a strictly increasing function, it follows that $e^A(x) + 1 + k_j$ is not an e^A-value, i.e. $e^A(x) + 1 + k_j \notin Y_e$, and therefore $e^A(x) + 1 \in Y_{k_j}$, for all $j \leq n$.

Since $\{Y_k : k \in \omega\}$ has the finite intersection property, it can be extended to an ultrafilter F on $X^A = \omega$. This ultrafilter is rich, as are all ultrafilters, because there are only finitely many observable data elements: for any term t of output type,

$$(t \approx 0)^A \cup (t \approx 1)^A = \omega \in F,$$

and so (exactly) one of $(t \approx 0)^A$ and $(t \approx 1)^A$ is in F.

Thus F is a state of A^*. But condition (10.i) fails for A^* at F. This follows because for each $k \in \omega$,

$$Y_e \in e^{A^*}(F) - (m^{A^*})^k(F).$$

To see this, first note that by definition of Y_e,

$$\{x \in X^A : e^A(x) \in Y_e\} = \omega \in F,$$

which makes $Y_e \in e^{A^*}(F)$. However $Y_e \notin (m^{A^*})^k(F)$, or else by (10.ii) $Y_e - k \in F$. But $Y_e - k$ is the complement of $Y_k \in F$, so this is impossible.

We thus see that the set $\{x \in \omega : a^A(x) = 1\} = Y_e$ belongs to $e^{A^*}(F)$, while its complement $\{x \in \omega : a^A(x) = 0\}$ belongs to $(m^{A^*})^k(F)$, showing that $a^{A^*}(e^{A^*}(F)) = 1$ while $a^{A^*}((m^{A^*})^k(F)) = 0$. Hence the term $a(\sigma)$ takes different values at these states, showing by Theorem 3.2 that $e^{A^*}(F)$ is not bisimilar to $(m^{A^*})^k(F)$ in A^* for any $k \in \omega$, and therefore that A^* cannot belong to K.

The second example is designed to show that the converse of Corollary 8.5 does not hold: $A^+, F \models u_1 \simeq u_2[\bar{d}]$ does not imply

$$\{x \in X^A : A, x \models u_1 \simeq u_2[\bar{d}]\,\} \in F.$$

Let A be a coalgebra with state set $X^A = \omega$, having

- a single method $m^A(x) = x + 1$ as in the previous example, and
- a single attribute a^A of sort $(\omega, \{0,1\})$, i.e. $a^A : \omega \times \omega \to \{0,1\}$, which is the characteristic function of the $\leq$ relation on ω. Thus $a^A(x,y) = 1$ iff $x \leq y$.

Take $u_1 = m(\sigma)$ and $u_2 = \sigma$. For any $x \in \omega$, let t be the ground observable term $a(\sigma, x)$. Then $t^A[\sigma^A[x]] = 1$ as $x \leq x$, but $t^A[m(\sigma)^A[x]] = 0$ as $x + 1 \not\leq x$. Therefore by Theorem 3.2, $\sigma^A[x]$ and $m(\sigma)^A[x]$ are not bisimilar, implying

$$\{x \in X^A : A, x \models m(\sigma) \simeq \sigma\} = \emptyset \notin F,$$

for any $F \in X^{A^+}$. This also shows that $A \models \neg(m(\sigma) \simeq \sigma)$.

Now take $F \in X^{A^*}$ to be any ultrafilter on X^A that is nonprincipal, i.e. $F \neq F_x$ for any $x \in X^A$. (As in the first example, all ultrafilters are rich.) For each $k \in \omega$,

$$\{x \in X^A : a^A(x,k) = 0\} = \{k+1, k+2, \dots\} \in F,$$

since a nonprincipal ultrafilter contains all cofinite sets. Thus $a^{A^*}(F,k) = 0$ for all inputs $k \in I^a$.

But $m^{A^*}(F)$ is also nonprincipal, because F is. For, by the analysis of the first example, if $m^{A^*}(F) = F_x$, then $x > 0$ and $F = F_{x-1}$. Hence also

$a^{A^*}(m^{A^*}(F), k) = 0$ for all inputs $k \in I^a$. This implies that $m^{A^*}(F)$ and F assign the same values to all ground observable terms, and so are bisimilar, by Theorem 3.2 again. Therefore $A^*, F \models u_1 \simeq u_2[\bar{d}]$, which also means $A^+, F \models u_1 \simeq u_2[\bar{d}]$.

This completes the counter-example to Corollary 8.5. It also shows that neither A^+ nor A^* is a model of $\neg(m(\sigma) \simeq \sigma)$, so neither filter nor ultrafilter enlargements preserve satisfaction of negations of state equations.

The third example shows that in part (4) of our main result 9.2 the hypothesis of closure under domains of morphisms cannot be weakened to closure under subcoalgebras. This is a class K that is closed under disjoint unions, images of morphisms, subcoalgebras, and filter enlargements, hence closed under ultrafilter enlargements, but not closed under domains of morphisms.

Let $Meth = \{m\}$ and $Att = \{a\}$, with m and a having no input sort and a having output sort $\{0, 1\}$. Let K be the class of coalgebras for this signature that satisfy $m^A(m^A(x)) = x$ for all $x \in X^A$. Closure of K under disjoint unions and subcoalgebras is straightforward to check. For closure under images of morphisms, let B be the image of $A \in K$ under morphism f. Then $m^B(m^B(f(x)) = f(m^A(m^A(x)) = f(x)$ in general. Since every member of X^B is of the form $f(x)$, this shows that $B \in K$.

For closure under filter enlargements, let $A \in K$ and take any $F \in X^{A^+}$. Then $Y \in m^{A^+}(m^{A^+}(F))$ iff $\{x \in X^A : m^A(m^A(x)) \in Y\} \in F$ iff $\{x : x \in Y\} = Y \in F$, showing that $m^{A^+}(m^{A^+}(F)) = F$. Hence $A^+ \in K$.

Now let A be the coalgebra having $X^A = \omega$, $m^A(x) = x + 1$, and $a^A(x) = 1$ iff x is even. Let B be the "quotient of A mod 2", defined by $X^B = \{0, 1\}$, $m^A(0) = 1$, $m^A(1) = 0$, $a^A(0) = 1$, and $a^A(1) = 0$. The characteristic function $f : \omega \to \{0, 1\}$ of the set of even numbers is a morphism $A \to B$. Here $f(x) = 1$ iff x is even, so f is in fact identical to a^A. But $B \in K$ while $A \notin K$, showing that K is not closed under domains of morphisms.

Notice that K is defined by a simple equation $(m(m(x)) \approx x)$ asserting *equality* of states. The characterisation of classes defined by (Boolean combinations of) such equations is a matter for further study, although there is less interest for coalgebraic theory in this type of equation than in those expressing bisimilarity $(t_1 \simeq t_2)$ of states.

The fourth example is a class of coalgebras that is definable by conditional equations, but not by equations alone. It is based on a pair of coalgebras presented in [Roş98] for a similar purpose, and is motivated by the specification of a class of objects called *flags* (e.g. [GM99, p. 291]).

There are three method symbols *up*, *down* and *rev* (for "reverse") and a single attribute symbol *up*?, all of which have empty input sort. The output sort of *up*? is $\{\mathtt{t}, \mathtt{f}, \mathtt{u}\}$, representing the truth-values *true*, *false* and *unknown*. Let B be the coalgebra with three states x_1, x_2, x_3 and methods and attribute as in the following table.

	up	$down$	rev	$up?$
x_1	x_1	x_2	x_2	$\mathtt{t}$
x_2	x_1	x_2	x_1	$\mathtt{f}$
x_3	x_1	x_2	x_3	$\mathtt{u}$

No two states are bisimilar, since they assign different values to the observable term $up?(\sigma)$. Thus state equations are interpreted in B by the *equality* relation.

The subset $\{x_1, x_2\}$ is closed under the three methods, so defines a subcoalgebra A of B. Note that all four operations in the table are either constant or injective, a property that is preserved by composition of operations. Hence the operation $x \mapsto t^B[x]$ defined by any ground term is either constant or injective. From this it follows that any two terms that agree on x_1 and x_2 will agree on x_3 as well, and therefore that *A and B validate exactly the same equations.*

Now let Φ be the set of all observable conditional equations valid in A and put $K = \mathit{Mod}\,\Phi$. Then K is closed under disjoint unions, images of bisimulations, and filter and ultrafilter enlargements. If K were definable by equations, then since A is in K and equationally indistinguishable from B, B would be in K too.

But $B \notin K$, because B is not a model of Φ, as shown by the formula

$$up?(rev(\sigma)) \approx up?(\sigma) \to up?(\sigma) \approx \mathtt{t}.$$

This conditional observable equation is not satisfied at x_3 in B, but is valid ("vacuously") in A since its antecedent is not satisfied at x_1 or x_2. An even simpler, though perhaps less natural, example is

$$rev(\sigma) \approx \sigma \to \mathtt{t} \approx \mathtt{f}.$$

References

[Bir35] Garret Birkhoff. On the Structure of Abstract Algebras. *Proceedings of the Cambridge Philosophical Society*, 31:433–454, 1935.

[BM96] Jon Barwise and Lawrence Moss. *Vicious Circles.* CSLI Lecture Notes No. 60. CSLI Publications, Stanford, California, 1996. Distributed by Cambridge University Press.

[BW85] Michael Barr and Charles Wells. *Toposes, Triples and Theories.* Springer-Verlag, 1985.

[Cor97] A. Corradini. A Complete Calculus for Equational Deduction in Coalgebraic Specification. Technical Report SEN-R9723, Centrum voor Wiskunde en Informatica (CWI), Amsterdam, 1997.

[EFT84] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic.* Springer-Verlag, 1984.

[End72] Herbert B. Enderton. *A Mathematical Introduction to Logic.* Academic Press, 1972.

[GM99] Joseph A. Goguen and Grant Malcolm. Hidden Coinduction: Behavioural Correctness Proofs for Objects. *Mathematical Structures in Computer Science*, 9:287–319, 1999.

[Gol93] Robert Goldblatt. *Mathematics of Modality.* CSLI Lecture Notes No. 43. CSLI Publications, Stanford, California, 1993. Distributed by Cambridge University Press.

[Gol00] Robert Goldblatt. Duality For Some Categories of Coalgebras. Research Report 00–2, School of Mathematical and Computing Sciences, Victoria University of Wellington, 2000. `http://www.mcs.vuw.ac.nz/~rob/papers/duality.ps` To appear in *Algebra Universalis.*

[GS98] H. Peter Gumm and Tobias Schröder. Covarieties and Complete Covarieties. *Electronic Notes in Theoretical Computer Science*, 11, 1998. `http://www.elsevier.nl/locate/entcs`

[GT75] R. I. Goldblatt and S. K. Thomason. Axiomatic Classes in Propositional Modal Logic. In J. N. Crossley, editor, *Algebra and Logic*, volume 450 of *Lecture Notes in Mathematics*, pages 163–173. Springer-Verlag, 1975.

[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`

[Kur98a] Alexander Kurz. Coalgebras and Modal Logic. In *Proceedings of Advances in Modal Logic '98, Uppsala*, pages 222–230, 1998.

[Kur98b] Alexander Kurz. Specifying Coalgebras with Modal Logic. *Electronic Notes in Theoretical Computer Science*, 11, 1998. `http://www.elsevier.nl/locate/entcs`

[Rei95] Horst Reichel. An Approach to Object Semantics Based on Terminal Co-Algebras. *Mathematical Structures in Computer Science*, 5:129–152, 1995.

[Roş98] Grigore Roşu. A Birkhoff-like Axiomatisability Result for Hidden Algebra and Coalgebra. *Electronic Notes in Theoretical Computer Science*, 11, 1998. `http://www.elsevier.nl/locate/entcs`

[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.