From Behaviour Categories to Generalized Spaces

Sweedler homs, residual Moore machines, and nonrepresentable halos

Research interlude for the categorical systems tutorial

21 October 2026 – working note, with Spivak motivation

Abstract

We compare the generalized spaces extracted by Fairbanks–Carlson–Spivak from Sweedler comonads for deterministic Moore machines, Moore machines with writer residue, and Moore machines with finite nondeterministic residue. The first two comonads are polynomial, with an action category of residual behaviours and representable halos indexed by execution words. The third is the cofree comonad on the finitely branching Moore transition functor. A behaviour with one transition to a terminal successor has halo $X \times \mathcal{P}_{\mathrm{f}}^{+} X$, which is nonrepresentable: joint observations retain correlations erased by separate projections. More sharply, the halo calculation, together with FCS Proposition 4.1.16, proves that the entire FCS underlying category is discrete, whereas its topological reflection is the Alexandrov space of reachability, which is generally non-discrete. We include an indexed halo calculation, isolate the assumptions on the Moore theory, and distinguish the proved constructions from further applications to modal logic, probability, and concurrency.

1 Motivation and scope

In the preceding parts of the tutorial, a theory T specifies operations and equations while a residual monad R specifies what an implementation may leave unhandled. Its Sweedler hom $\{T, R\}$ is a comonad whose coalgebras are stateful R-residual runners of T: this is McDermott–Rivas–Uustalu, Theorem 1 [1], assuming a locally presentable base category and accessible monads T and R. For $R = \mathrm{id}$ these are ordinary comodels. Garner's behaviour category organizes comodels of a suitable accessible theory as a category of copresheaves. For polynomial comonads this is especially concrete: small categories correspond to polynomial comonads [3].

A complementary motivation comes from Spivak's 2022 seminar notes *The free monad, cofree comonad, and topological space associated to a polynomial* [4]. For a polynomial interface p, he considers the free monad $\mathfrak{m}_p$ of finite, well-founded p-trees and the cofree comonad $\mathfrak{c}_p$ of possibly infinite p-trees. The point set $\mathfrak{c}_p(1)$ carries a topology generated by tests specified by finite trees in $\mathfrak{m}_p(1)$; for finite p this is a Stone topology. Spivak contrasts the topos p-Coalg $\simeq \mathfrak{c}_p$-**Set**, associated with what he calls "singular experience", with sheaves on this finite-test space, associated with "counterfactual plans". A system may follow one particular history, while experiments specify finite information about many possible histories. This offers a motivation for studying not just behaviours but the relations among execution, experiments, and their logical interpretations.

There is a useful distinction of syntax here. For the Moore polynomial $p = Oy^A$, p-coalgebras are ordinary Moore machines, and their cofree comonad agrees with $\{T_{\mathrm{Moore}}(A, O), \mathrm{id}\}$: both are cofree on the same category of machines over **Set**. Yet Spivak's free polynomial monad $\mathfrak{m}_p$ is *not* the operational Moore theory $T_{\mathrm{Moore}}(A, O)$, whose getter satisfies the reader equations.

Consequently, a language of finite tree tests and a language of operations executed by a runner should not be identified simply because their behaviour comonads coincide.

Fairbanks, Carlson, and Spivak (hereafter FCS) ask what *space* an arbitrary comonad on **Set** determines. Their answer includes an underlying topological space, an underlying small category, and generalized "halos" around points [2]. We use the Sweedler theory as a supply of physically and computationally meaningful comonads, and ask exactly what those three constructions remember.

The resulting geometries should also be kept distinct. Spivak's finite-test topology records agreement under finite observations; for binary Moore streams it is the familiar prefix (Cantor) topology. The FCS topology computed below instead records closure under operational continuations, yielding an Alexandrov reachability topology. These generally differ even on the same point set. Our purpose here is to calculate the latter and the accompanying category and halos; the language of finite experiments, its possible Sweedler interaction, and reachability–observability pairings belong in a separate companion note.

This note is meant to follow the planned interlude on categories as comonoids in **Poly**, immediately after the tutorial's Part IX on behaviours. It does not purport to reconstruct FCS's general theory. All calculations below are in **Set**, for sets A of inputs and O of outputs, with $O \neq \varnothing$. The nonpolynomial and nondiscrete-topology examples require $A \neq \varnothing$, but the categorical discreteness theorem does not. We allow arbitrary input sets A for which the accessible cofree constructions exist; thinking of A as finite avoids branching terminology becoming misleading.

Provenance. The Sweedler–runner equivalence (under the hypotheses of [1, Theorem 1]), the correspondence between polynomial comonads and categories, and the FCS definitions are inherited results [1, 3, 2]. The explicit Moore calculations and the nonrepresentability and discreteness results are calculations presented here. Their newness relative to the wider literature has *not* been established.

2 The three generalized-space constructions

Let C be a comonad on **Set**. Its terminal C-coalgebra has carrier $C1$ (with structure $\delta_1 : C1 \to CC1$), so its set of *points* is

$$B_C = C1.$$

FCS equip B_C with the topology whose opens are the subsets underlying subcoalgebras of this terminal coalgebra (their Definition 1.1.9, following Proposition 1.1.7) [2].

2.1 The cofree runner behind CX

The definition of a halo becomes more transparent in the Sweedler case. Suppose $C = \{T, R\}$, and let $\mathrm{Run}_R(T)$ be the category of R-residual T-runners. The forgetful functor has a right adjoint, the *cofree-runner* functor:

$$U : \mathrm{Run}_R(T) \rightleftarrows \mathbf{Set} : G, \qquad U \dashv G, \qquad C = UG.$$

In particular, CX is the *underlying state set of the cofree residual runner* $G(X)$. Its counit

$$\varepsilon_X : CX \longrightarrow X$$

may be read as the auxiliary X-label of a state. This is not an additional operational effect: the X-labels are arbitrary observations of states, superimposed on the runner's T-operations and residual R-behaviour.

The adjunction supplies the key universal property. For any residual runner S and any function $\ell : US \to X$, there is a unique runner morphism

$$\widehat{\ell} : S \longrightarrow G(X), \qquad U\widehat{\ell} : US \longrightarrow CX, \qquad \varepsilon_X \circ U\widehat{\ell} = \ell.$$

Thus the cofree runner simultaneously records the operational behaviour of S and a freely chosen X-valued labelling of its states. We shall often write $\widehat{\ell}$ also for the underlying map $US \to CX$.

Because G is a right adjoint, $G(1)$ is terminal in $\mathrm{Run}_R(T)$. Its carrier $C1$ is the space of final behaviours. Every runner S has a unique behaviour morphism, whose underlying function we denote

$$\mathrm{beh}_S : US \longrightarrow C1.$$

Now apply G to the unique map $! : X \to 1$. Its underlying map

$$C(!) : CX \longrightarrow C1$$

forgets auxiliary X-labels while retaining the residual runner behaviour. Naturality and the universal property give

$$C(!) \circ U\widehat{\ell} = \mathrm{beh}_S\,. \tag{1}$$

In particular, the square

$$\begin{array}{ccc} CX & \xrightarrow{\varepsilon_X} & X \\ \downarrow{\scriptstyle C(!)} & & \downarrow{\scriptstyle !} \\ C1 & \longrightarrow & 1 \end{array}$$

expresses the two ways of forgetting information: discard the additional labels to obtain a behaviour, or discard both labels and behaviour.

2.2 Halos as fibres of the behaviour-forgetting map

For each $b \in B_C$, FCS's unlabelled *halo functor* is

$$H_b(X) = \{z \in CX \mid C(!)(z) = b\}. \tag{2}$$

For a Sweedler comonad this has a direct interpretation:

> $H_b(X)$ is the set of states of the cofree X-labelled runner whose underlying unlabelled behaviour is exactly b.

The behaviour b stays fixed while the set X of auxiliary labels varies. A map $f : X \to Y$ acts by relabelling: Cf sends a state of $G(X)$ over b to a state of $G(Y)$ over the same b, since $C(!) \circ Cf = C(!)$. Hence H_b is indeed a functor $\mathbf{Set} \to \mathbf{Set}$.

Equation (1) also explains the halo's relation to an *arbitrary* runner S. If $s \in US$ has behaviour b and we choose a labelling $\ell : US \to X$, then the cofree extension determines an element

$$\widehat{\ell}(s) \in H_b(X).$$

However, $H_b(X)$ is *not* the fibre of the behaviour map for one selected implementation S; it is the corresponding fibre in the *universal cofree* runner. It gathers every X-labelled refinement of b represented there, not merely the states of one particular runner.

Two examples preview the calculations below. In the deterministic Moore case

$$CX \cong O^{A^*} \times X^{A^*}, \qquad H_b(X) \cong X^{A^*}.$$

For a fixed behaviour $b : A^* \to O$, a halo element assigns an X-label to the state reached after each input word. In the finitely nondeterministic case, when b has one transition to a *terminal* successor up to bisimulation, we instead obtain

$$H_b(X) \cong X \times \mathcal{P}_{\mathrm{f}}^{+}(X).$$

There may be several differently X-labelled terminal successors even though they all collapse to one successor when the auxiliary labels are forgotten. This failure of a fixed, label-independent set of positions will be central in Section 6.

FCS regard the functor H_b contravariantly as a formal limit of sets (Definition 4.1.2). A representable halo has the form $H_b(X) \cong \mathbf{Set}(E_b, X) = X^{E_b}$ for a fixed position set E_b. The halo's *points* are natural transformations $H_b \Rightarrow \mathrm{id}_{\mathbf{Set}}$, equivalently maps from the formal halo to the singleton formal space.

The underlying category of C has objects $b \in B_C$ and arrows given by specializations between continuous maps $1 \to C$ (FCS Definition 3.1.12). FCS Proposition 4.1.16 identifies arrows *out of* b with points of H_b [2]. This is *not* the same as an arrow for each topological specialization, unless some further property ensures it.

FCS also retain a *labelled halo*, an indexed-set version of (2), defined from the associated comonad on $\mathbf{Set}^{B_C}$ (Definition 4.1.8). We shall calculate it in both polynomial and nondeterministic cases. Although the "unlabelled" halo $H_b(X)$ still accepts auxiliary X-labels, those labels all come from *one set* X. The labelled halo allows a different label set in each fibre over B_C. These indexed data should not be confused with a merely unlabelled halo carrying a map to the set of points (cf. their Warning 4.1.14).

3 The Moore presentation and residual runner equations

3.1 Theory and equations

The Moore theory $T_{\mathrm{Moore}}(A, O)$ has an O-ary operation get and free unary operations step_a, one for each $a \in A$. The getter satisfies the reader equations

$$\mathsf{get}(o \mapsto x) = x, \tag{G1}$$
$$\mathsf{get}\big(o \mapsto \mathsf{get}(p \mapsto x_{o,p})\big) = \mathsf{get}(o \mapsto x_{o,o}). \tag{G2}$$

There are no equations between the getter and updates, and no equations between distinct update terms. Thus observations may change after an update. These are the hypotheses for all formulas below; a more strongly quotiented theory would yield different runners.

An R-residual runner on a set S interprets get by a map

$$\gamma : S \longrightarrow R(O \times S)$$

and each step_a by $\rho_a : S \to RS$, satisfying the equations as interpreted in the state-passing monad $X \mapsto (R(X \times S))^S$. Morphisms are ordinary functions preserving these interpretations, as in the residual-runner category of [1].

3.2 Writer residue: get is forced to be pure

Let M be a monoid with identity e, and $R(X) = M \times X$. Initially write

$$\gamma(s) = (m(s), (o(s), g(s))), \qquad \rho_a(s) = (m_a(s), \delta_a(s)).$$

By (G1), evaluating a constant continuation x after γ must agree with returning x without an effect or state change:

$$(m(s), (x, g(s))) = (e, (x, s)).$$

Thus $m(s) = e$ and $g(s) = s$, and (G2) is automatic.

Proposition 3.1. *An $M \times -$-residual runner of $T_{\mathrm{Moore}}(A, O)$ is precisely a coalgebra*

$$S \longrightarrow O \times (M \times S)^A,$$

i.e. an output map $o : S \to O$ and arbitrary effectful updates $\rho_a : S \to M \times S$.

Proof. The argument above forces the getter to have the form $s \mapsto (e, (o(s), s))$. The free unary updates are unconstrained. Conversely these data interpret the two reader equations and extend uniquely to terms of the free-update theory, with writer effects multiplied in execution order. □

3.3 Finite-powerset residue: get is again deterministic

Let $R = \mathcal{P}_{\mathrm{f}}$, including the empty finite set, and initially take

$$\gamma : S \longrightarrow \mathcal{P}_{\mathrm{f}}(O \times S), \qquad \rho_a : S \longrightarrow \mathcal{P}_{\mathrm{f}} S.$$

Equation (G1) yields

$$\{s' \mid (o, s') \in \gamma(s)\} = \{s\}.$$

In particular $\gamma(s) = K_s \times \{s\}$ for a nonempty finite set $K_s \subseteq O$. Evaluate (G2) using the set of distinct variables $(o, p) \in O \times O$ as return values. Its two sides produce, respectively, $K_s \times K_s$ and $\{(o, o) \mid o \in K_s\}$. Equality requires K_s to be a singleton. Hence $\gamma(s) = \{(o(s), s)\}$.

Proposition 3.2. *A $\mathcal{P}_{\mathrm{f}}$-residual runner of $T_{\mathrm{Moore}}(A, O)$ is precisely a coalgebra*

$$S \longrightarrow O \times (\mathcal{P}_{\mathrm{f}} S)^A,$$

i.e. a deterministic Moore observation and one finite set of possible successors for each command.

Proof. The two getter equations give the form stated above. There are no constraints on the update interpretations. Conversely the stipulated structures satisfy (G1)–(G2), and the updates extend to terms through finite-powerset Kleisli composition. □

4 The polynomial cases: deterministic and writer Moore spaces

Write $W = A^*$, let $\epsilon \in W$ be the empty word, and use uv for first executing u, then v. For any $b : W \to O$, put $b_u(v) = b(uv)$.

Theorem 4.1. *There are isomorphisms of comonads*

$$C_0(X) := \{T_{\mathrm{Moore}}(A,O), \mathrm{id}\}(X) \cong O^W \times X^W, \tag{3}$$

$$C_M(X) := \{T_{\mathrm{Moore}}(A,O), M \times -\}(X) \cong O^W \times M^{W\times A} \times X^W. \tag{4}$$

Both are polynomial comonads, with representable unlabelled halos $H_q(X) \cong X^W$.

Proof. The respective runner endofunctors are $F_0(S) = O \times S^A$ and $F_M(S) = O \times (M \times S)^A$ by Proposition 3.1 and its special case $M = 1$. Their cofree coalgebras on X record the O-label and X-label at every word, and in the writer case the M-label emitted by each edge (w, a). This gives (3)–(4). The cofree comonads and the Sweedler homs induce the same categories of coalgebras *over* **Set** by [1], hence are isomorphic as comonads. Fixing the data $q \in C_M 1 = O^W \times M^{W\times A}$ leaves precisely X^W. □

For either case, write B for the point set (O^W or $O^W \times M^{W\times A}$). If q is a point, the residual q_w shifts every observation and every edge effect by the word w. Explicitly,

$$\varepsilon_X(q, f) = f(\epsilon), \qquad \Delta_X(q, f) = (q, w \mapsto (q_w, f_w)). \tag{5}$$

The associated category is the action category $B \rtimes W$:

$$\mathrm{Ob} = B, \qquad (q, w) : q \to q_w, \qquad (q, u); (q_u, v) = (q, uv).$$

For $R = \mathrm{id}$, this is the familiar concrete presentation of Garner's Moore behaviour category. Its copresheaves are Moore machines (as in Garner's description of comodels [5]): a machine S decomposes into fibres of its behaviour map $S \to B$, and the transitions give maps between those fibres. This is also the small category extracted by FCS from the polynomial comonad. Here the appropriate polynomial comonads are accessible, and the original theory $T_{\mathrm{Moore}}(A, O)$ is accessible since A and O are sets.

The opens of the FCS topology are the forward-invariant subsets:

$$U \subseteq B \text{ open} \quad \iff \quad q \in U \Longrightarrow q_w \in U \quad (\forall w \in W). \tag{6}$$

The topology is Alexandrov for the reachability preorder $q \preceq q'$ iff $q' = q_w$ for some w. Hence its interior on predicates is $\Box P(q) \iff \forall w\, P(q_w)$.

Remark 4.2 (Reachability, not the stream prefix topology). For $A = \{a\}$, points of the deterministic case are streams $O^{\mathbb{N}}$. The topology defined by (6) is the *shift-reachability Alexandrov topology*, not the product/prefix topology. Its minimal neighbourhood at a stream is its set of tails, while the halo still remembers *all words*, even if many words yield the same tail. For comparison, Spivak's 2022 Topos Institute talk [4] equips the point set of the cofree comonad on a polynomial p with a topology having the free-monad data $m_p(1)$ as a basis; when p is finite this is a Stone topology. The two constructions should not be conflated.

Our specialization convention is $x \preceq y$ iff *every* open containing x contains y. Neither of the reachability topologies need be T_0: for example the two periodic streams $0101\ldots$ and $1010\ldots$ are mutually reachable under shift and therefore topologically indistinguishable.

The labelled halo in these polynomial cases is equally simple. Given a B-indexed family $E = (E_q)_{q\in B}$,

$$\widehat{H}_q(E) \cong \prod_{w\in W} E_{q_w}. \tag{7}$$

The equality keeps the target point of each word, not only the underlying set of positions. Specialization along w evaluates the data at the w-position.

5 The nondeterministic Moore comonad

The runner endofunctor from Proposition 3.2 is

$$F(S) = O \times (\mathcal{P}_{\mathrm{f}} S)^A.$$

It is accessible, so the usual cofree-coalgebra construction applies.

Theorem 5.1. *The $\mathcal{P}_{\mathrm{f}}$-residual Moore Sweedler comonad is*

$$C_{\mathrm{nd}}(X) := \{T_{\mathrm{Moore}}(A, O), \mathcal{P}_{\mathrm{f}}\}(X) \cong \nu Z.(X \times O \times (\mathcal{P}_{\mathrm{f}} Z)^A). \tag{8}$$

Its points $B_{\mathrm{nd}} = C_{\mathrm{nd}}1$ are observable, finitely branching-per-input Moore behaviours modulo bisimulation.

Proof. For an accessible endofunctor F, the cofree comonad has underlying functor $X \mapsto \nu Z.(X \times FZ)$, with counit taking the root X-label and comultiplication labelling nodes by their rooted behaviours. Proposition 3.2 identifies its coalgebras *over* **Set** with residual Moore runners, so [1] identifies the cofree comonad with the Sweedler hom. □

Write a point as b with output o_b and finite a-successor sets $\mathrm{succ}_a(b) \subseteq B_{\mathrm{nd}}$. We regard behaviours as extensional final-coalgebra points, and labelled behaviours in (8) as bisimulation classes with labels in $O \times X$.

The opens are again forward-invariant:

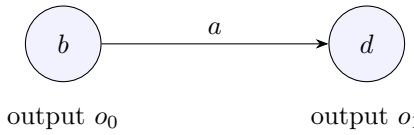
$$b \in U \implies \mathrm{succ}_a(b) \subseteq U \quad (\forall a \in A). \tag{9}$$

This follows because open subobjects of the final cofree-comonad coalgebra are the subsets supporting restricted F-coalgebra structures. Thus the topology is Alexandrov for the reflexive-transitive closure of successor reachability, using the convention $b \preceq c$ iff every open containing b contains c. In particular, the topology need not be T_0 when distinct behaviours are mutually reachable. This topological statement does *not* say that reachability steps are arrows of the FCS underlying category.

6 The one-successor halo: nonrepresentability

Fix $a \in A$ and $o_0, o_1 \in O$ (they need not be distinct). Let d be the behaviour with output o_1 and no successors under any input; let b have output o_0, a single a-successor d, and no successor for other inputs. Thus

$$\mathrm{succ}_a(b) = \{d\}, \qquad \mathrm{succ}_c(b) = \varnothing \quad (c \neq a).$$



Theorem 6.1 (Nonrepresentable halo)**.** *At the point b just defined, the unlabelled halo is naturally*

$$H_b(X) \cong X \times \mathcal{P}_{\mathrm{f}}^+(X), \tag{10}$$

where $\mathcal{P}_{\mathrm{f}}^+ X$ is the set of finite nonempty subsets of X. For $|A| \geq 1$, the Sweedler comonad C_{nd} is therefore not polynomial.

Proof. A labelled behaviour lies over b precisely when its root has output o_0 and some label $x \in X$, and it has a nonempty finite set of a-successors, all of which are deadlocks with output o_1. The successors' X-labels distinguish otherwise bisimilar deadlocks, so they form a finite nonempty *subset* $K \subseteq X$, giving (10).

To disprove representability, set $X = Y = \{0, 1\}$ and compare the different subsets of $X \times Y$

$$K_1 = \{(0,0),(1,1)\}, \qquad K_2 = \{(0,1),(1,0)\}.$$

They have identical images under both coordinate projections: $\pi_i(K_j) = \{0, 1\}$ for $i, j \in \{1, 2\}$. Hence the two distinct elements $((0,0), K_1)$ and $((0,0), K_2)$ of $H_b(X \times Y)$ have the same image in $H_b(X) \times H_b(Y)$. The canonical product comparison is not injective. A representable $X \mapsto \mathbf{Set}(E, X)$ preserves all products, so H_b is not representable. Every polynomial functor has representable fibres over its points, hence C_{nd} cannot be polynomial. □

The distinction is operational: separate observations of two channels record the possible values of each channel, but need not record *which pairs can occur together*. The two sets K_1 and K_2 have the same marginals and different joint possibilities. A representable halo would have fixed positions, independently of how a new label partitions the possible successors. Here additional labels can split a single unlabelled successor into several distinct labelled successors.

Proposition 6.2 (Indexed version)**.** *Let $E = (E_c)_{c \in B_{\mathrm{nd}}}$ be a family indexed by points of C_{nd}. The labelled halo at b, defined using the canonical sliced comonad on $\mathbf{Set}^{B_{\mathrm{nd}}}$, has the form*

$$\widehat{H}_b(E) \cong E_b \times \mathcal{P}_{\mathrm{f}}^{+}(E_d). \tag{11}$$

Restricting to constant families $E_c = X$ recovers (10).

Proof. The root has unlabelled behaviour b and hence requires an element of E_b. Every successor has unlabelled behaviour d and hence requires a label in E_d; an extensional finite collection of labelled deadlocks is a finite nonempty subset of E_d. Forgetting those labels identifies every successor with d. □

Equation (11) retains the source and destination types, useful for any subsequent semantics of state-dependent certificates. But the full notion of labelled halo also specifies how such data vary over *arbitrary* point-indexed families; it is not captured merely by adjoining a function to the unlabelled halo.

7 The categorical coreflection is discrete

The failure of representability is only the first obstruction. For finite-powerset residual Moore machines, *no* nonidentity arrows remain in the FCS underlying category. This is stronger than the one-successor calculation.

Lemma 7.1 (The one-successor case)**.** *Every natural transformation*

$$\theta_X : X \times \mathcal{P}_{\mathrm{f}}^{+}(X) \longrightarrow X$$

is the root projection, $\theta_X(x, K) = x$.

Proof. For $n \geq 2$, set $E_n = \{r, u_1, \ldots, u_n\}$. Every permutation of the u_i fixes the element $(r, \{u_1, \ldots, u_n\})$, so naturality forces θ_{E_n} of this element to be fixed by the full symmetric group $\mathfrak{S}_n$. Its only fixed element in E_n is r. Given any $(x, K) \in X \times \mathcal{P}_{\mathrm{f}}^{+}(X)$, take $n = \max(2, |K|)$ and a map $f : E_n \to X$ such that $f(r) = x$ and $\{f(u_i) : 1 \leq i \leq n\} = K$. Naturality yields

$$\theta_X(x, K) = f(\theta_{E_n}(r, \{u_1, \ldots, u_n\})) = f(r) = x.$$

The root projection is manifestly natural. □

For the general case it is important that the labelled behaviours may be infinite. We use a separate symmetric refinement of each given behaviour, not an allegedly universal finite-branching tree.

Lemma 7.2 (Symmetric duplication of a labelled behaviour)**.** *Let $b \in B_{\mathrm{nd}}$, let X be a set, and let $t \in H_b(X)$. There exist a set L, an element $\widetilde{t} \in H_b(L)$, and a function $f : L \to X$ such that:*

(i) $H_b(f)(\widetilde{t}) = t$*;*

(ii) the root of $\widetilde{t}$ has a label $r \in L$;

(iii) the only element of L fixed by every *permutation $\sigma : L \to L$ satisfying $H_b(\sigma)(\widetilde{t}) = \widetilde{t}$ is r.*

Proof. Unfold t into a rooted A-edge-labelled tree, with an O-output and an X-label at each node. Each node has finitely many successors for each fixed input a; the tree may have infinite depth. Its nodes form a set, indexed by finite branching paths.

At the *root only*, replace each immediate successor subtree (for every $a \in A$) by *two disjoint copies of that entire subtree*. Do not duplicate recursively inside either copy. Give *every node* in the resulting tree its own fresh label, distinct from the label of every other node, and let L be the set of all these labels; write $r \in L$ for the root label. Keep the original outputs and input-edge labels. Because every node has a distinct label, distinct rooted subtrees cannot be bisimilar as L-labelled behaviours. Thus this unfolding is already extensional, and its nodes really correspond to distinct labelled subbehaviours of $\widetilde{t} \in C_{\mathrm{nd}}(L)$.

Define $f : L \to X$ by returning the original X-label of the node from which each copy arose. Under f, each doubled pair of root successors becomes bisimilar and is identified by the *set*-valued successor operation $\mathcal{P}_{\mathrm{f}}$. The relabelled behaviour is therefore t, proving (i). Forgetting the labels also recovers b, so $\widetilde{t} \in H_b(L)$.

For each *immediate successor subtree of the root*, let σ interchange its two copies, including all their descendant labels, while fixing the root and all other labels. Relabelling by σ merely exchanges two members of an unordered successor set, and hence leaves $\widetilde{t}$ unchanged. Each nonroot label lies in one of these root-level copies and is moved by the corresponding swap; all the swaps fix r. Their common fixed-label set is exactly $\{r\}$, establishing (iii). Nothing in this argument requires finite depth: a swap exchanges whole subtrees, even when they contain infinite paths. □

Theorem 7.3 (Discreteness of the categorical coreflection)**.** *Let A be any set and $O \neq \varnothing$. For the comonad*

$$C_{\mathrm{nd}} = \{T_{\mathrm{Moore}}(A, O), \mathcal{P}_{\mathrm{f}}\},$$

every halo has exactly one point:

$$\boxed{\mathrm{Nat}(H_b, \mathrm{id}_{\mathbf{Set}}) \cong 1 \qquad (b \in C_{\mathrm{nd}}1).}$$

Consequently the category obtained by the FCS categorical coreflection of C_{nd} is the discrete category *on B_{nd}. In contrast, its FCS topological reflection is the Alexandrov topology for reflexive-transitive transition reachability, which is non-discrete if $A \neq \varnothing$ (and $O \neq \varnothing$).*

Proof. Root-label evaluation gives a natural transformation $H_b \Rightarrow \mathrm{id}$, so the displayed set is nonempty. Let $\theta : H_b \Rightarrow \mathrm{id}$ be any natural transformation and fix $t \in H_b(X)$. Take L, $\widetilde{t}$, and f from Lemma 7.2. For every permutation σ leaving $\widetilde{t}$ invariant, naturality gives

$$\sigma\big(\theta_L(\widetilde{t})\big) = \theta_L\big(H_b(\sigma)(\widetilde{t})\big) = \theta_L(\widetilde{t}).$$

Thus $\theta_L(\widetilde{t})$ is fixed by every such permutation, so by Lemma 7.2(iii) it is the root label r. A further use of naturality gives

$$\theta_X(t) = f\big(\theta_L(\widetilde{t})\big) = f(r),$$

the original root label of t. As t was arbitrary, θ is root evaluation.

FCS Proposition 4.1.16 [2] identifies specializations out of a point b with points of its halo. Thus each object of the underlying category has exactly one outgoing arrow, necessarily its identity. The assertion about topology is equation (9). For example the point b in Section 6 has a reachable but distinct deadlock successor d, so $b \preceq d$ topologically while there is no arrow $b \to d$ in the underlying category. □

Remark 7.4 (Why finite branching causes no problem at infinite depth)*.* The construction is chosen separately for each t. It doubles only the root's successor set for each input, so image-finiteness is preserved, independently of the depth of the original tree. The label set L may be infinite: halo naturality ranges over all sets, not merely finite ones. Extensional collapse of duplicate successor behaviours is essential to the argument.

Remark 7.5 (Finite multisets: a contrasting categorical coreflection)*.* The finite-multiset functor $\mathcal{B}_{\mathrm{f}}$ remembers how many copies of each successor occur, rather than merely their support. For the cofree comonad on $S \mapsto O \times (\mathcal{B}_{\mathrm{f}}S)^A$, let b_k be a root of output o_0 with exactly k occurrences of one terminal behaviour d of output o_1 under a fixed input a, and no other successors. Its halo is

$$H_{b_k}(X) \cong X \times \mathrm{Sym}^k(X), \qquad \mathrm{Sym}^k(X) = X^k/\mathfrak{S}_k,$$

where the second factor is the set of multisets of cardinality k. For $k = 1$, this is $X \times X$, and the second projection selects the unique successor naturally. By FCS Proposition 4.1.16 there is therefore an arrow $b_1 \to d$ in the underlying category. For $k \geq 2$ only root evaluation is natural: take $E = \{r, u_1, \ldots, u_k\}$ with pairwise distinct elements, and use the permutation action on the u_i to force any natural selector to return r; relabelling then covers every (x, K) with $|K| = k$ counted with multiplicity. Thus no individual successor is selected naturally when its multiplicity is at least two.

This shows that the underlying category of the multiset cofree comonad is *not* generally discrete, although it still forgets some branching. The stronger discreteness result for $\mathcal{P}_{\mathrm{f}}$ depends specifically on the extensional collapse of duplicate successors; it is not a consequence of nondeterministic branching alone.

This yields a clean separation between three structures:

$$\begin{gathered}\text{topology: existence of reachable successor behaviours,}\\ \text{underlying category: only identity arrows,}\\ \text{full halo: label-sensitive branching and joint correlations.}\end{gathered} \tag{12}$$

The categorical coreflection is discrete *despite* nontrivial reachability specializations. This is not a contradiction: FCS's topological reflection and categorical coreflection are distinct universal constructions.

8 Comparison and modal interpretation

	Deterministic Moore	Moore–writer	Moore–$\mathcal{P}_{\mathrm{f}}$
Runner functor	$O \times S^A$	$O \times (M \times S)^A$	$O \times (\mathcal{P}_{\mathrm{f}} S)^A$
Point set	O^{A^*}	$O^{A^*} \times M^{A^* \times A}$	$\nu Z.\, O \times (\mathcal{P}_{\mathrm{f}} Z)^A$
Underlying comonad	polynomial	polynomial	nonpolynomial
Example halo	X^{A^*}	X^{A^*}	$X \times \mathcal{P}_{\mathrm{f}}^{+} X$
Underlying topology	residual reachability	residual reachability	successor reachability
Underlying category	action category	action category	discrete

The topologies support a propositional S4 interior operator. Unlike one-step Hennessy–Milner logic with output atoms $[o]$ (which assert the current Moore output) and modalities $\langle a \rangle$, this interior only expresses $\Box$ along the reflexive-transitive reachability relation $\to^*$. For image-finite Moore transition systems, output atoms together with the one-step modalities characterize bisimilarity. The reachability topology is weaker: its possible failure of T_0 already shows that its opens need not separate distinct (hence non-bisimilar) points of the final coalgebra. For the nondeterministic space, if $\to^*$ is the reflexive-transitive closure of its transition relation, then

$$\Box P(b) \iff \forall c\, (b \to^* c \Longrightarrow P(c)).$$

But this modal predicate does not preserve information about *how* different observations correlate among nondeterministic successors. In (11), the relevant data live over different point-indexed fibres E_b and E_d. They could furnish an entry point to quantified modal semantics with state-dependent individuals or witnesses; the required logical doctrine and quantifier laws have not yet been constructed.

This also identifies a concrete limit of the analogy with differential cohesion. A nonrepresentable halo is a generalized neighbourhood, not by itself a Weil-algebra infinitesimal disk, a jet, or a synthetic tangent space. Those stronger geometric structures require extra assumptions on the ambient category or on the comonad.

9 Research agenda and editorial location

We recommend inserting this note after the planned *Poly interlude* following Part IX of the tutorial. The Poly interlude explains why categories arise from polynomial comonads; this note supplies a controlled nonpolynomial counterpoint. The next tasks are, in order:

1. **Indexed modal semantics.** Identify the logical structure induced by the labelled halo $E_b \times \mathcal{P}_{\mathrm{f}}^{+}(E_d)$, separating existential branching, universal branching, and joint-correlation predicates. Compare this with the witnessed logic developed in Part VI.A.
2. **More residual theories.** Repeat the finite branching calculation for probabilistic residuals, comparing ordinary marginalization against joint laws, and then for a true-concurrent residual theory; identify what is new beyond the powerset case.
3. **Behavioural functoriality.** Test whether changes of source and residual theories induce useful continuous maps of the generalized spaces and compatible maps of labelled halos.
4. **Relative and smooth settings.** Examine when the Sweedler construction can be performed in a slice or smooth topos and when its comonads preserve enough limits to support an ionadic interpretation. Do not infer this merely from the Set examples.

Status. Theorem 4.1, Theorem 5.1, the explicit nonrepresentability result, the indexed halo, and the categorical discreteness theorem are established here under the stated Moore presentation and standard cofree/Sweedler hypotheses. The claimed novelty relative to existing literature is not established; in particular, nonpolynomiality of powerset-like functors is unsurprising. The distinctive points worth comparing with the literature are the halo-level joint/marginal observation calculation and the topology-versus-discrete-category contrast. A limited search of the FCS appendix has not located this precise finite-powerset example, but no systematic priority audit has been carried out. The proposals concerning first-order modal logic, probability, concurrency, and differential cohesion are questions, not results of this note.

References

[1] D. McDermott, E. Rivas, and T. Uustalu, "Sweedler Theory of Monads," *FoSSaCS 2022*, LNCS 13242, pp. 428–448 (2022). doi:10.1007/978-3-030-99253-8_22.

[2] A. D. Fairbanks, K. Carlson, and D. I. Spivak, "Comonads as Spaces," arXiv:2607.15091v1 (2026), especially §1.1, Def. 3.1.12, and §4.1. https://arxiv.org/abs/2607.15091.

[3] D. Ahman and T. Uustalu, "Directed Containers as Categories," *EPTCS* 207, pp. 89–98 (2016). https://arxiv.org/abs/1604.01187.

[4] D. I. Spivak, "The free monad, cofree comonad, and topological space associated to a polynomial," Topos Institute, Berkeley Seminar, June 2022. https://dspivak.net/talks/pdfs/20220614-topos-scotttopology.pdf.

[5] R. Garner, "The costructure–cosemantics adjunction for comodels for computational effects," *Mathematical Structures in Computer Science* **32**(4), 374–419 (2022). doi:10.1017/S0960129521000219.