

Getting to Grips with Sweedler Runners for Systems

A conversational tutorial, Part IX: behaviours, reachability, and observability
Final comodels, Garner’s behaviour category, minimal realisation, and temporal properties

October 2026

Aim. The preceding parts concentrated on how systems are specified, implemented, and composed. We now ask a different question: *what can actually be observed of a state by running the system?* For a theory or monad, T , Garner’s cosemantics gives a final comodel and, more richly, a small behaviour category $\mathbb{B}_T$ from which the whole category of T -comodels can be recovered as a presheaf-like functor category. This gives a natural language for observability. Adding an initial state gives reachability. Putting the two together leads to minimal realisation. We end by explaining why the Manna–Pnueli hierarchy of temporal properties belongs nearby: safety is only the beginning once specifications are predicates on whole behaviours rather than on instantaneous states.

Contents

1	From implementation to behaviour	2
2	A first example: one unary action	2
3	The final comodel	3
4	The input theory: final behaviour is a stream	4
5	Moore machines: final behaviour as an input/output function	4
6	Garner’s behaviour category $\mathbb{B}_T$	5
6.1	Back to the unary example	6
6.2	For Moore machines	6
7	Reachability from a representable	6
8	Observability as subterminality	7
9	Minimal realisation: the Yoneda-to-terminal image	7
10	Where reachability–observability duality enters	8
11	Residual runners already induce profunctorial behaviour	8
12	Why the predicate layer sees safety first	9

13 Linear time and branching behaviour must be separated	10
14 The Manna–Pnueli hierarchy and the need for a clock	10
14.1 Topology on final behaviours	11
15 A useful picture of the whole chapter	11
16 A test case: hidden Markov models and backward duality	12
17 Where this leaves the Run programme	16

1 From implementation to behaviour

A runner tells us how a theory is realised. In the simplest case,

$$T \rightsquigarrow \text{Id},$$

a runner is just a T -comodel. Its carrier S is a set of internal states, and every operation of the theory has a cooperation telling the environment how to respond and how the state changes.

So far we have mostly looked inward: what structure must S possess in order to implement the theory? Now we look outward:

Which distinctions between states can an external observer actually detect?

Two different states may give exactly the same answers to every possible experiment. Conversely, some states may never be reached from the chosen initial condition. These are the two classical directions:

- reachability : which internal states can actually occur?
- observability : which internal states can be distinguished externally?

The remarkable point is that the comodel story already contains a canonical space of behaviours against which every state can be compared.

2 A first example: one unary action

Return to the theory T freely generated by one unary operation u . A T -comodel is simply a set S with an endomap

$$f : S \longrightarrow S.$$

In Part I this was our simplest dynamical system.

What is its externally visible behaviour if the theory contains no observation operation at all? There is only one possible observable behaviour. Whatever state we start in, the system can perform another u -step, and another, but it never emits data. The final comodel therefore has one state:

$$1 \xrightarrow{\text{id}} 1.$$

Every dynamical system has the unique map

$$S \longrightarrow 1.$$

If we looked only at the underlying set of the final comodel, we would conclude that all states are observationally identical.

Yet the theory certainly has dynamics: we may execute u , u^2 , u^3 , and so on. Where did that information go?

It reappears when we pass from the *set of behaviours* to Garner's *category of behaviours*. In this example the behaviour category has one object but many endomorphisms, generated by u : morally it is the one-object category associated to the monoid

$$(\mathbb{N}, +, 0).$$

Thus a functor

$$\mathbb{N} \longrightarrow \mathbf{Set}$$

is precisely a set equipped with an endomap.

This tiny example already separates two notions that we will keep distinct:

final comodel = possible behaviours as states

while

behaviour category = behaviours together with residual commands between them.

The second can retain dynamics even when the first has only one point.

3 The final comodel

Let T be an accessible monad on $\mathbf{Set}$, presented if we like by an algebraic theory. The category of T -comodels is comonadic over $\mathbf{Set}$, and in particular has a final object. Let us denote its underlying set by

$$B_T^0.$$

An element of B_T^0 is a possible complete T -behaviour. The set B_T^0 may be empty: for example, a theory with a constant operation has no nonempty comodels. In that case the behaviour category below is empty and $\mathbf{Comod}(T) \simeq 1$.

For every T -comodel S , finality gives a unique comodel morphism

$$\text{beh}_S : S \longrightarrow B_T^0.$$

We will call this the *behaviour map*.

Conceptually,

$$\text{beh}_S(s)$$

is everything an external user of the theory could ever learn by interacting with state s . If

$$\text{beh}_S(s) = \text{beh}_S(s'),$$

then no experiment expressible in the theory can distinguish s from s' .

Garner gives a particularly general description of the elements of the final comodel as *admissible behaviours*: coherent families

$$\beta_A : T(A) \longrightarrow A$$

which say, uniformly in A , how the state runs any T -computation far enough to return its value [1]. We will not need the axioms here. The important idea is that a behaviour is already a complete recipe for answering all possible computations.

4 The input theory: final behaviour is a stream

Take the free theory with one A -ary operation **read**. A comodel has a cooperation

$$\gamma : S \longrightarrow A \times S.$$

Write

$$\gamma(s) = (o(s), \sigma(s)).$$

Starting from s , repeated reads reveal

$$o(s), \quad o(\sigma s), \quad o(\sigma^2 s), \quad \dots$$

so every state determines an infinite stream in A :

$$\text{beh}_S(s) = (o(s), o(\sigma s), o(\sigma^2 s), \dots).$$

The final comodel is therefore the familiar stream coalgebra

$$A^{\mathbb{N}} \xrightarrow{\langle \text{head}, \text{tail} \rangle} A \times A^{\mathbb{N}}.$$

This is perhaps the cleanest example of the slogan

a state is observable only through the behaviour it generates.

Two internal states may be represented very differently and yet determine exactly the same stream.

5 Moore machines: final behaviour as an input/output function

Now return to the Moore theory $T_M(A, O)$. A comodel is a Moore machine

$$S, \quad \delta : A \times S \longrightarrow S, \quad o : S \longrightarrow O.$$

For a word $w = a_1 \cdots a_n \in A^*$, write

$$\delta_w(s)$$

for the state reached after processing the word. Then a state s determines the function

$$\beta_s : A^* \longrightarrow O, \quad \beta_s(w) = o(\delta_w(s)).$$

Thus the final behaviour space is

$$B_{T_M(A, O)}^0 \cong O^{A^*}.$$

The empty word records the current output:

$$\beta_s(\varepsilon) = o(s),$$

and input a moves to the left derivative

$$(a^{-1}\beta)(w) = \beta(aw).$$

So the final Moore machine has state space O^{A^*} , output

$$\beta \mapsto \beta(\varepsilon),$$

and transition

$$(a, \beta) \mapsto a^{-1}\beta.$$

The behaviour map is exactly

$$s \mapsto \beta_s.$$

For automata theorists this is completely familiar: a state is identified extensionally by all future observations after all possible input words. The point for us is that this standard construction appears canonically as the map into the final $T_M(A, O)$ -comodel.

6 Garner's behaviour category $\mathbb{B}_T$

The final comodel gives the *objects* of Garner's richer construction. The behaviour category

$$\mathbb{B}_T$$

has objects the admissible behaviours $\beta \in B_T^0$. For a behaviour β , arrows out of β are represented by commands in $T(1)$, modulo Garner's behaviour-dependent equivalence $\sim_\beta$; the class $[m]$ has target the residual behaviour β_m . Thus $\mathbb{B}_T$ does not refine the observational equivalence encoded by B_T^0 : its objects are still exactly the complete behaviours. What it adds is dynamics between residual behaviours.

Garner's theorem is strikingly strong:

$$\boxed{\mathbf{Comod}(T) \simeq [\mathbb{B}_T, \mathbf{Set}]}$$

for every accessible monad T on **Set**, compatibly with the forgetful functors to **Set** [1]. If a comodel corresponds to a functor

$$F : \mathbb{B}_T \rightarrow \mathbf{Set},$$

then its underlying state set is

$$US \cong \coprod_{\beta \in B_T^0} F(\beta),$$

and the behaviour map is simply the projection sending an element of the fibre $F(\beta)$ to β .

This description will do most of the work in the rest of the chapter. A comodel is a family of internal states indexed by behaviours, with residual commands transporting states between the fibres.

6.1 Back to the unary example

For one unary generator there is one behaviour object and a monoid of commands. The theorem becomes

$$\mathbf{Comod}(T) \simeq [\mathbb{N}, \mathbf{Set}],$$

exactly the category of sets with an endomap. The single object of $\mathbb{B}_T$ records that there is only one observable behaviour; its endomorphisms retain the dynamics that the final behaviour set alone cannot see.

6.2 For Moore machines

For the Moore theory the description is exact. The objects are functions

$$\beta : A^* \rightarrow O,$$

and A^* acts on them by left derivatives

$$\beta \mapsto w^{-1}\beta, \quad (w^{-1}\beta)(v) = \beta(wv).$$

Hence $\mathbb{B}_{T_M(A,O)}$ is the translation category for this action:

$$\mathbb{B}_{T_M}(\beta, \gamma) = \{w \in A^* \mid w^{-1}\beta = \gamma\}.$$

A functor $\mathbb{B}_{T_M} \rightarrow \mathbf{Set}$ is therefore a Moore machine sorted by its extensional behaviours. Notice that distinct words leading to the same residual behaviour are not thereby identified: the hom-set remembers the word witnessing the residual move.

7 Reachability from a representable

Choose an initial state $s_0 \in S$, and write

$$\beta = \text{beh}_S(s_0).$$

Under Garner's equivalence, the pointed comodel is an element $s_0 \in F(\beta)$. By Yoneda this is equivalently a natural transformation

$$y(\beta) = \mathbb{B}_T(\beta, -) \longrightarrow F.$$

The representable $y(\beta)$ is therefore the universal pointed comodel having initial behaviour β . Its states are the residual commands out of β , namely the classes in $T(1)/\sim_\beta$.

The reachable part is exactly the image of this Yoneda map:

$$\text{Reach}(S, s_0) = \text{im}(y(\beta) \rightarrow F).$$

On underlying sets this is the familiar formula

$$\text{Reach}(S, s_0) = \{\text{run}_m(s_0) \mid m \in T(1)\}.$$

No extra “closure” clause is required: this set is already closed under the cooperations, since further residual commands are represented by composition in $\mathbb{B}_T$.

For a Moore machine, $y(\beta)$ is the unfolding from β , with one state for each input word, and

$$\text{Reach}(S, s_0) = \{\delta_w(s_0) \mid w \in A^*\}.$$

Thus reachability is controlled by a representable, and the choice of a point is essential.

8 Observability as subterminality

The behaviour map is injective exactly when every fibre $F(\beta)$ contains at most one element. Hence

$$S \text{ observable} \iff F \rightarrow 1 \text{ is monic.}$$

Observable comodels are precisely the subterminal objects of $[\mathbb{B}_T, \mathbf{Set}]$. Equivalently, they correspond to sets of behaviours closed under residual transitions: cosieves in $\mathbb{B}_T$.

For a Moore machine, observability says

$$o(\delta_w(s)) = o(\delta_w(s')) \text{ for every } w \in A^* \implies s = s'.$$

The terminal functor 1 is the final comodel. Thus observability is controlled by the terminal object, whereas reachability is controlled by a representable. This is the first precise form of the asymmetry that will matter below.

9 Minimal realisation: the Yoneda-to-terminal image

Fix a behaviour $\beta \in B_T^0$. There is a unique map

$$y(\beta) \longrightarrow 1.$$

Because $[\mathbb{B}_T, \mathbf{Set}]$ is a presheaf topos, it has the usual epi-mono image factorisation. Define

$$\text{Min}(\beta) = \text{im}(y(\beta) \rightarrow 1).$$

For any pointed realisation (S, s_0) of β , the Yoneda map factors as

$$y(\beta) \longrightarrow \text{Reach}(S, s_0) \longrightarrow \text{Min}(\beta) \hookrightarrow 1.$$

The first image discards unreachable states; the second identifies reachable states with the same behaviour. In particular, every reachable realisation of β subjects onto the same minimal observable realisation.

This is a general Myhill–Nerode pattern for every accessible monad on $\mathbf{Set}$. Reachable pointed realisations of β are quotients of the representable $y(\beta)$; equivalently, they amount to congruences on the residual-command set $T(1)/\sim_\beta$ compatible with its $\mathbb{B}_T$ -action. The minimal realisation is the coarsest such quotient that retains exactly the residual behaviours of β .

For Moore machines,

$$\text{Min}(\beta) = \{w^{-1}\beta \mid w \in A^*\},$$

with duplicate derivatives identified. This is precisely the usual Nerode-style minimal Moore realisation.

The categorical spine. For a pointed comodel of behaviour β , reachability and observability are organised by

$$y(\beta) \twoheadrightarrow \text{Reach}(S, s_0) \twoheadrightarrow \text{Min}(\beta) \hookrightarrow 1.$$

Reachability starts from the representable generated by a point; observability ends at the terminal comodel. Garner’s theorem turns the informal two-stage reduction into an ordinary image-factorisation argument in a topos.

10 Where reachability–observability duality enters

So far reachability and observability are two different reductions:

$$\begin{array}{c} \text{subobject generated from initial data} \\ \text{versus} \\ \text{quotient by observational indistinguishability.} \end{array}$$

In classical linear systems these two constructions become dual under linear duality. Reachable subspaces of a system correspond to annihilators of unobservable subspaces of the dual system; controllability and observability exchange roles.

This suggests an important goal for $\mathbf{Run}(\mathcal{D})$, but also an important warning. There is no canonical operation

$$(T \rightsquigarrow R) \not\Rightarrow (R \rightsquigarrow T).$$

Sweedler hom is not symmetric in its two arguments. If a genuine reachability–observability duality is to exist, it should come from *additional duality in the ambient category*: for example a contravariant involution, linear duals on a suitable subcategory, or another doctrine transporting theories, handlers, and runners together.

The presheaf description makes the obstruction particularly sharp. Reachability is generated by a representable $y(\beta)$, whereas observability is expressed by subterminality inside 1. The topos structure itself does not exchange representables and subterminals. Thus the tutorial moral is:

Reachability and observability exist generally; an exchange between them needs extra structure.

In special settings that extra structure can come from linear duality or from a powerset-style dualisation; it is not supplied by Garner’s behaviour category alone. We will return to this when smooth and linear systems enter the story.

11 Residual runners already induce profunctorial behaviour

The preceding sections used $T \rightsquigarrow \text{Id}$, but Garner’s theorem also settles the first Set-based step towards relative behaviour. Let

$$Y : T \rightsquigarrow R$$

be a runner, with structure

$$\theta_X : T(X) \times Y \longrightarrow R(X \times Y).$$

Every R -comodel S closes the residual effect and produces a T -comodel with carrier $Y \times S$. Thus the runner determines a functor

$$\Phi_Y : \mathbf{Comod}(R) \longrightarrow \mathbf{Comod}(T), \quad S \longmapsto Y \times S.$$

This is exactly the closure-by-an- R -comodel construction, but performed uniformly in the downstream comodel.

The functor Φ_Y is cocontinuous. The forgetful functor from comodels creates colimits, cartesian product with the fixed set Y preserves colimits in $\mathbf{Set}$, and the corresponding comparison map is therefore an isomorphism. Transporting across Garner’s equivalences gives a cocontinuous functor

$$[\mathbb{B}_R, \mathbf{Set}] \longrightarrow [\mathbb{B}_T, \mathbf{Set}].$$

By the standard representation of cocontinuous functors between presheaf categories, this is determined by a profunctor

$$P_Y : \mathbb{B}_R^{\text{op}} \times \mathbb{B}_T \longrightarrow \mathbf{Set}.$$

Explicitly, if $y_R(\gamma) = \mathbb{B}_R(\gamma, -)$, then

$$P_Y(\gamma, \beta) = \Phi_Y(y_R(\gamma))(\beta).$$

In words, it is the set of states in $Y \times y_R(\gamma)$ whose induced T -behaviour is β .

This construction has the expected composition law. Sequentially composing runners corresponds, up to the canonical coherence isomorphism, to composing the functors Φ ; under the presheaf–profunctor correspondence this becomes coend composition of the profunctors P_Y . A tight monad map is the residual-state-one special case and recovers Garner’s contravariant cofunctoriality on behaviour categories [1].

The profunctor is therefore not a separate conjectural semantics added to the comodel story. In the $\mathbf{Set}$ -based case it is the representing object for the operation “close this runner against every downstream comodel”, expressed through Garner’s behaviour categories.

What remains open is the genuinely double-categorical and base-change story: the treatment of runner 2-cells, faithfulness and essential image, and above all the extension from $\mathbf{Set}$ to the other duoidal bases used elsewhere in this tutorial. Those are the points at which the later behavioural functor to $\mathbf{CofProf}$ has real work to do.

12 Why the predicate layer sees safety first

Garner’s equivalence also clarifies exactly what Part V’s state-predicate layer captures. Let a comodel correspond to $F : \mathbb{B}_T \rightarrow \mathbf{Set}$. A subobject

$$G \hookrightarrow F$$

is precisely a family of subsets of the state fibres closed under every residual transition. On underlying states, it is an invariant subset. Hence the subobject lattice $\text{Sub}(F)$ is the logic of invariants of the comodel.

The inclusion of invariant subsets into all subsets of US ,

$$\text{Sub}(F) \hookrightarrow \mathcal{P}(US),$$

has two familiar adjoints. The left adjoint sends a set of states to the smallest invariant containing it: reachability closure, a least-fixpoint or μ -construction. The right adjoint sends a predicate $p \subseteq US$ to the largest invariant contained in p : a greatest-fixpoint or ν -construction. For an initial state s_0 ,

every state reachable from s_0 satisfies p

exactly when s_0 lies in this greatest invariant. In temporal notation this is the familiar safety form AGp .

Thus Part V’s certification layer is not merely *similar* to safety reasoning: at the Boolean $\mathbf{Set}$ -valued level it is the internal subobject logic of the comodel topos. To express liveness one must enrich the semantic object beyond invariant subobjects.

13 Linear time and branching behaviour must be separated

There is one important level distinction. For the input-only stream theory the final behaviour object is already linear:

$$B_T^0 = A^{\mathbb{N}}.$$

A temporal property can therefore be a language of infinite streams.

For a Moore theory, however,

$$B_{T_M}^0 = O^{A^*}$$

is a complete input-indexed *tree* of future observations. A subset of O^{A^*} is consequently a branching-time property of Moore behaviours, not automatically a linear-time Manna–Pnueli property. To compare with linear time one first extracts a trace along an infinite input sequence. For example, a linear trace property

$$L \subseteq (A \times O)^\omega$$

induces the universal branching property

$$\hat{L} = \{\beta \in O^{A^*} \mid \text{every } a_0 a_1 \cdots \in A^\omega \text{ induces a trace in } L\}.$$

Not every branching property is obtained this way.

The same distinction appears in the fixpoint logic of state predicates. The negation of $AG \neg p$ is $EF p$, not $AF p$. Thus guarantee is not simply the Boolean complement of safety once the system branches over inputs. In input-free theories, such as unary dynamics or streams generated by a single successor, this branching/linear distinction collapses.

14 The Manna–Pnueli hierarchy and the need for a clock

Manna and Pnueli organise linear-time ω -regular properties into a hierarchy that becomes visible once a theory supplies a notion of infinite stepwise execution. Typical forms are

$$Gp, \quad Fp, \quad GFp, \quad FGP,$$

corresponding respectively to safety, guarantee, recurrence, and persistence; obligation is generated by Boolean combinations at the safety/guarantee level, and reactivity gives the full level generated by recurrence and persistence [4, 5].

There is a useful fixpoint mnemonic:

$$\begin{array}{ll} \text{safety} & : \nu, \\ \text{guarantee} & : \mu, \\ \text{recurrence} & : \nu\mu, \\ \text{persistence} & : \mu\nu. \end{array}$$

This makes the connection with the preceding section suggestive: reachability closure already supplies a μ -construction, while invariant interior supplies a ν -construction. We should not yet identify this formal alternation with the whole Manna–Pnueli hierarchy, but it indicates the right kind of structure a richer temporal layer must expose.

There is also a genuine obstruction that is easy to miss. Safety in the form $AG p$ only asks for closure under all commands and is therefore invariant under presentation of the monad. By contrast, $AF p$, $GF p$, and $FG p$ require a notion of successive atomic steps or of infinite runs. An

arbitrary accessible monad need not contain a canonical “next” operation: equations may identify composites, and state-like theories can collapse the prefix structure almost completely.

Two possible remedies are worth keeping in view. One can make time explicit by equipping the theory with a grading, in the spirit of graded monads, so that one-step and n -step computations remain distinguished. Alternatively, one can try to define infinite runs intrinsically from the prefix preorder on the residual-command set $T(1)/\sim_\beta$, for example by suitable ideals. The latter is presentation-independent but can degenerate badly for theories such as state. We will therefore treat general liveness as requiring additional temporal structure, rather than pretending that every monad already carries a canonical clock.

Automata theory reflects the same hierarchy. Deterministic Büchi conditions capture recurrence, deterministic co-Büchi conditions capture persistence, while weak deterministic automata sit at the obligation level rather than being interchangeable with co-Büchi recognition [5].

Three levels should not be conflated. Invariant subobjects of a comodel give state-based safety. Predicates on O^{A^*} give branching properties of complete Moore behaviours. Manna–Pnueli properties live on linear infinite traces and require, in addition, a suitable notion of step or run.

14.1 Topology on final behaviours

For free algebraic theories there is also a useful topological formulation. Final comodels carry a natural topology generated by finite observations, and continuous maps between them recover the expected stream-processor semantics [6]. In the familiar trace spaces this places the temporal hierarchy alongside the classical topological one: safety properties are closed, guarantee properties open, recurrence properties are of G_δ -type, and persistence properties of F_σ -type. This is a promising intrinsic bridge between finite observation and temporal classification. We will not claim the same topological characterisation for arbitrary accessible monads without further hypotheses.

15 A useful picture of the whole chapter

The central construction is no longer best pictured as a chain ending in a Boolean-valued map, because a temporal predicate on the underlying set of final behaviours is not a comodel morphism. The categorical part is instead

$$y(\beta) \longrightarrow \text{Reach}(S, s_0) \longrightarrow \text{Min}(\beta) \longleftarrow 1,$$

all inside $[\mathbb{B}_T, \mathbf{Set}]$. Separately, a specification tests the appropriate semantic object: an invariant subobject for state safety, a subset of B_T^0 for extensional branching behaviour, or a language of infinite traces once a linear-time semantics has been chosen.

This gives a useful four-level distinction:

internal state	: $s \in S$,
reachable state	: $s \in \text{Reach}(S, s_0)$,
extensional behaviour	: $\text{beh}_S(s) \in B_T^0$,
temporal trace	: an infinite path extracted using additional step structure.

The behaviour category $\mathbb{B}_T$ links the middle levels by recording residual commands. Keeping these levels separate prevents both the reachability/observability and the branching/linear-time distinctions from being hidden by notation.

16 A test case: hidden Markov models and backward duality

The reachability–observability duality becomes substantially less formal once probability and partial observation are introduced. A useful test case is provided by the work of Kim and Mehta on nonlinear filtering [2, 3]. They consider a hidden Markov model with hidden state X_t and noisy observation process Z_t , and recover a genuine analogue of the classical Kalman controllability–observability duality. Their striking result is that the dual control system is not another forward stochastic system but a *backward stochastic differential equation*.

This is precisely the sort of phenomenon for which the runner perspective ought to provide a structural explanation. In the linear theory, duality between controllability and observability can be expressed by matrix transposition. For a nonlinear stochastic system there is no comparable finite-dimensional transpose available. Nevertheless, there remains a much more elementary variance reversal: states or distributions are propagated forward, whereas observables are pulled backwards.

To see this, begin with the familiar pairing between a measure μ on a state space S and an observable $f : S \rightarrow \mathbb{R}$,

$$\langle \mu, f \rangle = \mu(f).$$

If $P_{s,t}$ denotes forward Markov propagation of states or distributions, then there is a corresponding pullback of observables. Fixing a terminal observable f at time T , one obtains schematically

$$Y_t = P_{t,T}^* f.$$

Thus the observable naturally evolves from the terminal time back towards the initial time. In the deterministic Markov case this is simply the backward Kolmogorov equation

$$-\frac{\partial Y_t}{\partial t} = \mathcal{A}Y_t, \quad Y_T = f,$$

where $\mathcal{A}$ is the generator of the Markov process. The reversal of temporal direction is therefore not peculiar to stochastic control. It expresses the elementary fact that forward propagation of states induces contravariant transport of tests.

This is already suggestive from the Sweedler perspective. A runner may be read covariantly on its state or comodel side and contravariantly on its handler or observable side. The distinction we have repeatedly encountered between running states forward and transporting handlers against a runner now appears as a temporal distinction:

$$\boxed{\text{forward evolution of states} \quad \longleftrightarrow \quad \text{backward evolution of observables.}}$$

The familiar transpose of linear control theory should then be regarded as the finite-dimensional linear shadow of this more primitive model–comodel reversal.

For a hidden Markov model the situation is richer because the observable propagated backwards must respect the information actually available through the observation process. Kim and Mehta make this precise by first passing to the unnormalised nonlinear filter. The Zakai equation is linear in the initial measure, and therefore defines a linear operator

$$\mathcal{L}^\dagger : \mathcal{M}(S) \longrightarrow L^2_{\mathcal{Z}}([0, T]; \mathbb{R}^m) \times \mathbb{R},$$

which records the exposed observation behaviour generated by an initial signed measure. Stochastic observability is equivalent to

$$\ker(\mathcal{L}^\dagger) = 0.$$

They then construct the adjoint operator

$$\mathcal{L} : L^2_{\mathbb{Z}}([0, T]; \mathbb{R}^m) \times \mathbb{R} \longrightarrow \mathcal{Y},$$

for a suitable function space $\mathcal{Y}$, by means of a backward stochastic differential equation. In their notation the dual system takes the form

$$-dY_t = (\mathcal{A}Y_t + h^\top(U_t + V_t))dt - V_t^\top dZ_t, \quad Y_T = c\mathbf{1},$$

where U_t is the control input and the pair (Y_t, V_t) is adapted to the observation filtration. The operator is

$$\mathcal{L}(U, c) = Y_0.$$

Its range

$$\mathcal{C}_T = \text{ran}(\mathcal{L})$$

is the controllable space of the dual BSDE, and Kim and Mehta prove

$$\mathcal{L}^\dagger = \mathcal{L}^*.$$

Consequently,

$$\boxed{\text{HMM observable} \iff \text{dual BSDE controllable.}}$$

For infinite-dimensional state spaces, controllability is naturally interpreted as density of $\text{ran}(\mathcal{L})$ in the chosen observable space rather than literal equality.

There are really two distinct pieces of structure in this construction, and it is useful not to conflate them. The *backward* direction comes from variance. States are pushed forward; observables are pulled back. The specifically *stochastic* part comes from the requirement that this backward-transported observable remain adapted to the forward observation filtration. A terminal condition cannot in general be propagated backwards as an adapted process using only a drift term. The additional process V_t supplies the martingale component,

$$-V_t^\top dZ_t,$$

required by stochastic adaptation. Schematically,

$$\begin{array}{c} \text{state-observable duality} \\ \Downarrow \\ \text{states forward, observables backward} \\ + \text{adaptation to stochastic observations} \\ \Downarrow \\ \text{BSDE.} \end{array}$$

This suggests that the appearance of a BSDE is much less mysterious from the runner viewpoint than it might initially seem. The Sweedler structure predicts the backward transport; stochastic enrichment supplies the martingale correction.

The control term is equally significant. The same observation map h which exposes hidden state in the forward HMM reappears as the channel through which control enters the backward system. Thus an *output interface* on the forward system becomes, after dualisation, an *input interface* on the dual system. This is exactly the structural pattern one expects from controllability–observability duality. In the linear Gaussian case, restricting to deterministic controls gives $V = 0$, and the BSDE reduces on the finite-dimensional space of linear observables to

$$-\dot{y}_t = Ay_t + Hu_t, \quad y_T = 0,$$

recovering the classical dual control system. The usual matrix transpose picture is therefore recovered as a special finite-dimensional case, rather than serving as the definition of duality.

There is a useful warning here. The Kim–Mehta BSDE should not be identified simply with the classical backward Zakai equation. Because the forward Zakai equation is linear, an adjoint equation had long been known:

$$-d\eta_t = \mathcal{A}\eta_t dt + h\eta_t \overleftarrow{dZ}_t.$$

This equation is important, for example, in smoothing. But Kim and Mehta stress that it is not the required controllability dual: it has no control input, and its stochastic integral is backward rather than being adapted to the ordinary forward observation filtration. Their BSDE is constructed so that the adjoint of the observable–behaviour map becomes an actual controlled system. From our point of view, this distinction is revealing. Merely reversing the state dynamics produces an adjoint evolution; incorporating the exposed observation interface is what turns that evolution into a controllability problem.

This fits naturally with the behavioural viewpoint developed in this part. An HMM exposes observation behaviour while retaining hidden residual state. One may picture this schematically as

$$\text{hidden initial state} \rightsquigarrow \text{observation trajectory}.$$

Observability asks whether distinct initial states or distributions can be distinguished by these exposed behaviours. It is therefore a forward, outgoing question. The dual problem instead fixes admissible interventions at the exposed boundary and asks which functions of the hidden initial state can thereby be generated. It is an incoming, contravariant question:

$$\text{controls at the observation boundary} \rightsquigarrow \text{functions on hidden state}.$$

This is very close to the two variances already present in a profunctor. If the behavioural construction

$$\text{Beh} : \mathbf{Run}(\mathcal{D}) \longrightarrow \mathbf{CofProf}$$

can be extended to an appropriate stochastic or probabilistically enriched setting, we should expect forward propagation of hidden states and backward propagation of tests to arise as the two readings of the same behavioural object rather than as separately postulated constructions.

The example therefore suggests a sharper test for the Sweedler programme. One would like to find an ambient probabilistic duoidal category in which a hidden Markov model is represented as a runner, derive from that runner both its action on measure-valued states and its contravariant action on function-valued handlers, and ask whether the latter, once adaptation to the observation filtration is taken into account, produces the Kim–Mehta equation

$$-dY_t = (\mathcal{A}Y_t + h^\top(U_t + V_t)) dt - V_t^\top dZ_t.$$

The first probabilistic step is already available in the Sweedler framework. Let $\mathbf{Dist}$ denote the finite-distribution monad on $\mathbf{Set}$, or equivalently use the residual probability theory $\mathbf{rand}$ introduced in Part II. A deterministic autonomous Mealy system with observation set O is a comodel of $T_{\text{Mealy}}(1, O)$, hence has structure

$$S \longrightarrow O \times S.$$

A runner

$$Y : T_{\text{Mealy}}(1, O) \rightsquigarrow \mathbf{Dist}$$

replaces this deterministic cooperation by a stochastic one,

$$\kappa : Y \longrightarrow \mathbf{Dist}(O \times Y).$$

Thus a discrete hidden Markov model is already a probabilistic Mealy runner: its latent state is the runner state Y , while one abstract observation-and-step operation is implemented by a residual probabilistic computation.

In the usual presentation an HMM is specified by a transition kernel

$$q : Y \longrightarrow \mathbf{Dist}(Y)$$

and an emission kernel

$$e : Y \longrightarrow \mathbf{Dist}(O).$$

The commutative strength of $\mathbf{Dist}$ combines these into the joint kernel

$$Y \longrightarrow \mathbf{Dist}(O \times Y)$$

required by the runner. Different conventions concerning whether emission occurs before or after transition correspond simply to the two familiar Kleisli factorizations of this joint kernel.

The difficult probabilistic structure therefore enters one stage later. A runner generates the joint stochastic dynamics, but filtering requires conditioning: after an observation one wishes to replace a prior distribution on Y by the corresponding posterior distribution. Markov categories with conditionals provide an abstract setting in which hidden Markov models, Bayesian filtering and smoothing can all be formulated. The question for the present programme is consequently not whether HMMs can be represented as runners, but how conditioning interacts with the Sweedler operations, behaviour, and the reachability–observability pairing.

For finite probability this investigation can remain inside

$$[\mathbf{Set}, \mathbf{Set}]_{\text{acc}}.$$

For continuous probability a natural next base is the category $\mathbf{Qbs}$ of quasi-Borel spaces. It is cartesian closed and locally presentable, and carries a commutative probability monad. This suggests the candidate Sweedler universe

$$[\mathbf{Qbs}, \mathbf{Qbs}]_{\text{acc}},$$

with the probability monad as a residual theory. Establishing the required duoidal closure and the corresponding measuring objects would turn the discrete construction above into a measure-theoretic one.

This substantially narrows the gap to the filtering–control duality of Kim and Mehta. The runner supplies the hidden stochastic dynamics; conditioning supplies Bayesian state estimation.

What remains is to identify the control-side construction and determine whether the duality between observability and controllability, and ultimately between filtering and optimal control, is induced by the same Sweedler pairing that relates behaviour and reachability elsewhere in this tutorial.

At present this should be regarded as a conjectural programme rather than a theorem of the runner formalism. But it is a strong test. If the BSDE emerges canonically from the mate or handler action associated to the stochastic runner, rather than being inserted by hand, then the framework would explain a genuinely nontrivial extension of Kalman duality.

We shall return to this question in Part XA, where hidden Markov models are developed as probabilistic runners and conditioning is added to obtain Bayesian filtering. The present section should therefore be read as identifying the reachability–observability structure that the later probabilistic theory will have to explain. From what we have already seen, the guiding slogan is simple:

Filtering pushes beliefs forward; duality pulls tests backward.

For hidden Markov models, pulling a test backwards while keeping it adapted to the observations is precisely the sort of operation that leads from an ordinary backward evolution equation to a BSDE.

17 Where this leaves the **Run** programme

This chapter has concentrated on the boundary case

$$T \rightsquigarrow \text{Id},$$

because there Garner’s theory is already available and very concrete. It gives us three firm pieces of mathematics:

1. every accessible T has a final comodel of admissible behaviours;
2. there is a behaviour category $\mathbb{B}_T$ whose arrows encode residual commands;
3. the whole category of T -comodels is equivalent to $[\mathbb{B}_T, \mathbf{Set}]$.

These facts make reachability, observability, and minimal realisation natural rather than imported notions.

For the wider **Run**($\mathcal{D}$) programme, the main open questions are now sharper:

- How does the **Set**-based profunctor P_Y extend to runner 2-cells and to the full double functorial behaviour semantics?
- Which other duoidal bases admit an analogue of Garner’s presheaf representation and hence of P_Y ?
- Under what additional duality on $\mathcal{D}$ do representable-generated reachability and subterminal observability become genuinely dual?
- What extra grading, path, or trace structure is needed before temporal specification semantics can support the full Manna–Pnueli hierarchy and compose along runners?

The next part changes the geometry once more. In smooth systems, the fundamental transition is no longer a next-state map

$$A \times S \rightarrow S,$$

but an infinitesimal change

$$A \times S \rightarrow TS.$$

That will let us compare the Run viewpoint with DOTS-style changes of tangency and with parameterised ODEs.

Part IX in one sentence. Garner’s equivalence $\mathbf{Comod}(T) \simeq [\mathbb{B}_T, \mathbf{Set}]$ makes reachability the image of a representable, observability subterminality, and minimal realisation the image of $y(\beta) \rightarrow 1$; runners then induce Set-based behavioural profunctors, while liveness requires us to distinguish branching behaviours from linear traces and to supply a genuine notion of temporal step.

References

- [1] R. Garner, *The costructure–cosemantics adjunction for comodels for computational effects*, Mathematical Structures in Computer Science 32(4), 374–419, 2022; published online 2021. DOI: 10.1017/S0960129521000219.
- [2] J. W. Kim and P. G. Mehta, *Duality for Nonlinear Filtering I: Observability*, arXiv:2208.06586, 2022.
- [3] J. W. Kim and P. G. Mehta, *Duality for Nonlinear Filtering II: Optimal Control*, arXiv:2208.06587, 2022.
- [4] Z. Manna and A. Pnueli, *A hierarchy of temporal properties*, Proceedings of PODC, 1990.
- [5] A. Duret-Lutz et al., *Exploring the temporal hierarchy of Manna & Pnueli*, Spot documentation, <https://spot.lre.epita.fr/hierarchy.html>.
- [6] T. Yoshida, *Continuous Functions on Final Comodels of Free Algebraic Theories*, arXiv:2211.11897, 2022.