Sweedler Runners for Temporal Systems, I

From Graded Operational Syntax to Temporal Behaviour

Polynomial Moore Trees, Bounded Comodels, and Chu Evaluation

Companion research draft • 16 October 2026

Abstract

This paper asks how an ordered resource grading of an operational theory determines a category of bounded execution histories and a pairing between computations and observations. In the depth-filtered Moore theory the grading supplies budgets and sequencing, while the input operations label temporal paths and observations govern branching. Splitting its reader layer into observation sorts identifies graded Moore algebras with an admissible subcategory of algebras for a cartesian polynomial monad on the discrete sorts $\mathbb{N} \times O$. That monad is a composite of weakening and strictly graded update trees via an explicit distributive law. The opposite, comodel side is the slice topos $[\mathcal{J}_A, \mathbf{Set}]/\Omega$, with $\mathcal{J}_A$ a category of update words and unused budgets; in the unary case its opposite is the familiar interval category. Its terminal comodel has finite-horizon carriers $\Omega_n = O^{A^{\leq n}}$ and completes to Garner's O^{A^*}. We construct a temporal-valued balanced tensor, prove its tensor–Hom adjunction and second Chu transpose, and verify a graded interaction-law coherence theorem: evaluating a polynomial tree selects one leafwise temporal path by observation guards and induces precisely the tensor relation. A completion comparison need not be invertible, with explicit missing-lift and orbit-identification examples. Cost shifts and static framing illustrate applications; dynamic ownership, general resources, and sheaf conditions remain outside the main theorem. Throughout the scope is the ordinary, no-residue Moore theory. A companion paper adds cost-writer residues, amortised implementation certificates, and ownership-sensitive composition without identifying them with the no-residue semantics developed here.

Status. This document consolidates and audits a sequence of exploratory notes. Results proved below are separated from standard background, conditional extensions, and open problems. In particular, we do not claim an internal polynomial presentation of the original one-sorted monad, a general Sweedler measuring equivalence, or that every temporal presheaf is a sheaf.

Place in the two-paper study. This first paper develops the semantic foundation: graded operational syntax, bounded comodel behaviour, temporal Chu duality, and the comparison with completed futures. The companion *Sweedler Runners for Temporal Systems, II* [17] takes the next question—what quantitative cost and ownership guarantees a stateful implementation can carry—as its subject. Part II uses the finite-horizon execution setting developed here but introduces nontrivial writer residues and additional certificates; its results are not claims of a general measuring theorem for arbitrary effects.

Relation to the Temporal Interlude. A third, application-facing companion, *Getting to Grips with Sweedler Runners for Systems: Temporal Interlude—Execution, Synchronization, and Quantitative Verification* [16], starts instead from explicit stateful runners and unfolds their executions into interval sheaves. It proves pullback composition of one-read trajectories, treats intermediate execution phases and amortised durations, and checks deadlines using timed monitors and the clock object of a temporal topos. That is *physical realisation* rather than the bounded operational semantics established here: our presheaf category of input-labelled budgets is not identified with either of the Interlude's interval-sheaf toposes. The finite-input charts in Section 6 are a point of contact, not a general comparison theorem.

Contents

1 Why put operational syntax and time in the same picture? 3
1.1 The problem from the two directions . . . 3
1.2 The basic Hankel picture . . . 3
1.3 A synopsis of the main comparison . . . 4

2 The reader layer and its many-sorted splitting 4
2.1 The ungraded sorted theory . . . 5

3 Filtering the Moore monad and recognizing its algebras 6

4 The polynomial distributive law for graded syntax 7
4.1 An explicit two-step tree . . . 9

5 From polynomial operations to active and inert trees 9

6 Time budgets, bounded observations, and comodels 9
6.1 The graded comodel axioms . . . 10
6.2 The budgeted-word category . . . 11
6.3 Relation to Garner and polynomial cofree comonoids . . . 11
6.4 Finite input charts and representable clocks . . . 12
6.5 Complete finite history trees . . . 13

7 The one-observation baseline: an ordinary profunctor tensor 13

8 A temporal-valued algebra–comodel tensor 14
8.1 The first transpose: temporal handlers . . . 14
8.2 Free cancellation with a generator at an arbitrary grade . . . 15

9 Coherence of the graded interaction: trees, paths, and evaluation 16
9.1 Sorted execution, including reader evaluation . . . 16
9.2 A tree determines a family of labelled temporal arrows . . . 16
9.3 Guard cylinders select the executed leaf . . . 17

10 Cofree comodels and the two-sided Chu adjunction 19
10.1 Delay is right adjoint to shift . . . 19
10.2 Boolean tests separate free operational syntax . . . 20

11 Completion and the limits of finite tests 21
11.1 The opposed completion adjunctions . . . 21
11.2 Failure of surjectivity: compatible classes without states . . . 21
11.3 Failure of injectivity: completed orbit distinctions . . . 22
11.4 Bounded Hankel balancing . . . 22

12 Resource shifts and quantitative examples 23
12.1 The shift adjoint triple . . . 23
12.2 Weighted update costs . . . 23
12.3 A static tensor-frame theorem . . . 23
12.4 Fresh contexts and the limits of generalization . . . 24

13 Two worked finite examples **24**
13.1 A unary machine whose states become observable at different horizons 25
13.2 What the three-path tree forgets under ungraded completion 25

14 Conclusions and a controlled next programme **26**

A Elementary nerve reconstruction **27**

B Raw get–step trees and reader reductions **27**

C Audit ledger and scope of the claims **28**

1 Why put operational syntax and time in the same picture?

1.1 The problem from the two directions

In a clock-based approach to dynamics, an execution is a map *from* a clock into a system. Shortening or translating the clock restricts the execution, as in the temporal-sheaf approach of Schultz, Spivak, and Vasilakopoulou [14]. On the Sweedler side, a theory consists of operations and equations, and its comodels provide machines on which the operations can be executed, with complete futures organized by Garner's behaviour category [5]. An algebra is a way to interpret the operations, while the algebra–comodel tensor identifies a formal computation with its result when run against a state. At first these approaches organize different data: temporal subintervals versus algebraic terms of unbounded finite extent.

We test a sharper proposition. Suppose an ordered additive grading has been chosen, counting the updates of an operational theory. What temporal indexing category does the resulting graded comodel have, and how do the operations enrich that category by labelled paths? The grading supplies the resource order, the weakening maps and the additive composition of budgets; the theory supplies the update alphabet and branching. Time is thus not recovered from ungraded syntax alone. If so, complete observable futures should be a *completion* of finite-horizon comodels, not the primitive objects. The algebra–comodel evaluation should take values in a temporal category, retaining the horizon after a computation has executed.

The Moore theory is a useful first case. It has one family of state-changing operations, one pure observation operation, branching programs, and a simple equational presentation. It makes the two geometries especially visible: syntax branches over the possible *observations* O, while complete behaviour branches over possible *input histories* A^*. The graded pairing must connect these, rather than silently identify them.

1.2 The basic Hankel picture

For a Moore machine (S, o, δ) with a chosen point s_0, where $o : S \to O$ and $\delta : S \times A \to S$, the reachability and observability maps are

$$r_{s_0} : A^* \to S, \quad r_{s_0}(u) = \delta_u(s_0), \qquad \mathrm{beh} : S \to O^{A^*}, \quad \mathrm{beh}(s)(v) = o(\delta_v(s)).$$

Their composite transposes to the Hankel pairing

$$H_{s_0}(u, v) = o(\delta_{uv}(s_0)), \qquad H_{s_0}(uw, v) = H_{s_0}(u, wv). \tag{1}$$

The balancing law transfers a word from the executed history into the future experiment. At finite horizon N, this pairing is defined only for $|u| + |v| \le N$. The finite bound therefore has a genuine mathematical meaning: extending the history leaves less room to test the future.

In this paper the underlying theory has *no residual effect*: the comodel executes its own operations, as opposed to implementing them in a different effect theory. The extension to general residual Sweedler runners is a subsequent problem.

1.3 A synopsis of the main comparison

The four central constructions are

$$T_{\le n}(X) \quad \longleftrightarrow \quad (n,o)\text{-sorted polynomial trees},$$
$$\mathrm{Comod}_{\mathrm{gr}}(T) \simeq [\mathcal{J}_A, \mathbf{Set}]/\Omega, \qquad \Omega_n = O^{A^{\le n}},$$
$$B \mathbin{\widehat{\otimes}_T} S \in \mathcal{E} = \mathbf{Set}^{\mathbb{N}^{\mathrm{op}}}, \qquad \mathcal{A}(B,[S,K]) \cong \mathcal{E}(B \mathbin{\widehat{\otimes}_T} S, K) \cong \mathcal{C}(S, \langle B, K\rangle).$$

Beyond these descriptions, we check the coherence of the interaction law: interpreting a sorted polynomial tree in an algebra and balancing against a comodel agrees with executing the observation-guarded temporal path selected by the state. This is an explicit graded, sorted instance of the algebra–coalgebra interaction perspective of Katsumata–Rivas–Uustalu [11], rather than a wholly new kind of interaction law. The more distinctive results are the temporal-valued two-sided Chu duality and the noncommutation of evaluation with completion.

What would count as an answer? Our target is not another description of Moore machines in isolation. We want a single account in which a finite program has an algebraic meaning, a bounded machine state supplies the information needed to run it, and their encounter leaves a well-defined amount of future observation available. The three constructions must agree on execution, not merely resemble one another. This is why the evaluation pairing is part of the question from the start: recovering bounded observations without relating them to computations would leave the Sweedler half of the comparison unexplained.

How the argument progresses. The proof follows the order in which this target imposes requirements. We first need manageable *program shapes*. The equations for pure reads obscure those shapes in one-sorted syntax, so the next section separates observation sorts; the subsequent two sections add the depth grading and exhibit polynomial trees with active substitution and inert restriction. Having constructed the algebraic side, we turn to *bounded machines*: the same budget now measures what a state can reveal after successive updates. Their category produces the finite-horizon version of the behaviour picture, and its unary case gives the familiar temporal interval indexing.

At that point the two sides coexist but have not yet been paired. We therefore identify the ordinary profunctor case before defining the temporal-valued execution tensor, prove its first and second Chu transposes, and check that algebraic tree substitution and the selected temporal path give the same evaluation. Only then do we ask what is gained or lost by passing from all finite horizons to complete Garner futures. The final resource examples and finite machines test the scope of the constructions; they do not introduce further prerequisites. In particular, this paper is about *one graded theory with no residual effects*, not yet a compositional calculus for arbitrary Sweedler runners.

2 The reader layer and its many-sorted splitting

A first obstacle is syntactic rather than temporal. If repeated pure reads are represented as vertices of ordinary computation trees, the reader equations identify trees with visibly different branching shapes. This makes it hard to use the standard polynomial account of arities directly. Separating the possible observations into sorts removes precisely that obstacle: the choice of an observation becomes a coordinate of an algebra, while updates remain genuine tree vertices. We start without grades to see what this separation achieves before adding a time budget.

Fix an arbitrary set of update letters A and a *finite nonempty* set O of observations, unless stronger hypotheses are stated. The one-sorted Moore theory is generated by unary step_a for

$a \in A$ and an O-ary get, with the reader equations

$$\mathsf{get}(o \mapsto x) = x, \tag{2}$$

$$\mathsf{get}(o \mapsto \mathsf{get}(p \mapsto x_{o,p})) = \mathsf{get}(o \mapsto x_{o,o}). \tag{3}$$

There are no equations commuting a step with a read. The grade of a term counts its maximum number of step nodes along a branch; reading has grade zero.

Theorem 2.1 (Diagonal splitting)**.** *Let $R(X) = X^O$ denote the reader monad. The category $\mathrm{Alg}(R)$ is equivalent to the full subcategory of $\mathbf{Set}^O$ on families $(B_o)_{o\in O}$ which are either all nonempty or all empty. For a nonempty reader algebra, its carrier is $B \cong \prod_o B_o$ and its operation is*

$$\mathsf{get}\big((b^{(p)})_{p\in O}\big)_o = b_o^{(o)}.$$

The all-empty family corresponds to the empty reader algebra. The equivalence is natural on homomorphisms.

Proof. The diagonal identities are those of a Płonka diagonal algebra [1]; compare also the account of rectangular and diagonal modes in [2]. For a nonempty algebra choose $e \in B$. The coordinate retractions

$$\pi_o(x) = \mathsf{get}(p \mapsto [x \text{ when } p = o, \text{ and } e \text{ otherwise}])$$

are idempotent; their images B_o give a product decomposition, with inverse assembled using get. The diagonal law proves the coordinate identities and independence of the selected representatives. Homomorphisms preserve the induced coordinate congruences and therefore act coordinatewise. If $B = \varnothing$, nonemptiness of O makes $B^O = \varnothing$, so the empty algebra exists; there is no algebra whose underlying nonempty carrier is a product with a mixture of empty and inhabited coordinates. This is the classical diagonal-algebra decomposition; compare Płonka [1], as well as [2]. □

Remark 2.2 (What the splitting does and does not say)*.* Reader equations have been absorbed into product coordinates; a read is no longer a tree vertex. This does *not* imply that the original monad R on **Set** is cartesian: for $|O| > 1$, even its unit naturality square can fail to be a pullback. Polynomial arities will live on $\mathbf{Set}^O$ or $\mathbf{Set}^{\mathbb{N}\times O}$, not automatically on **Set**.

2.1 The ungraded sorted theory

Once the read operation has become coordinate selection, a step is simply a family of operations that takes one continuation for each possible subsequent observation. The following polynomial presentation records those continuations explicitly. Its value for the later argument is not merely a normal-form formula: it tells us which branching belongs to *syntax*, as opposed to the input histories that will index observable futures.

Let $P : \mathbf{Set}^O \to \mathbf{Set}^O$ be the polynomial endofunctor

$$(P\mathbf{B})_o = \coprod_{a\in A} \prod_{p\in O} B_p. \tag{4}$$

A P-algebra has one operation $\sigma_{a,o} : \prod_p B_p \to B_o$ for every a, o. Its free monad $M = P^*$ consists of finite O-coloured update trees [3, 4]. It includes algebras with mixed-empty sorts, so we retain the full subcategory of *admissible* M-algebras.

Proposition 2.3 (Ungraded Moore as sorted polynomial syntax)**.** *There is an equivalence $\mathrm{Alg}(T_{\mathrm{Moore}}) \simeq \mathrm{Alg}(M)_{\mathrm{adm}}$. With $\Delta X = (X)_{o\in O}$ and $\Pi\mathbf{B} = \prod_o B_o$, the underlying Moore monad is naturally isomorphic to $\Pi M \Delta$, with monad structure induced by the composite free–forgetful adjunction.*

Proof. A one-sorted Moore algebra is a reader algebra together with arbitrary unary operations $\sigma_a : B \to B$. Under Theorem 2.1, the coordinate maps of each σ_a are precisely the operations (4), with no further equations. This proves the algebra equivalence. The free sorted algebra on ΔX is admissible: an inhabitant $x \in X$ gives a leaf in every sort, while if $X = \varnothing$ and $O \neq \varnothing$ no finite tree has a leaf. Composing the left adjoints Δ and free M with their right adjoints gives the asserted monad $\Pi M \Delta$. □

Writing $Y_o = M(\Delta X)_o$, its trees satisfy the normal-form recursion

$$Y_o \cong X + \coprod_{a \in A} \prod_{p \in O} Y_p, \qquad T_{\mathrm{Moore}}(X) \cong \prod_o Y_o. \tag{5}$$

In particular, successive reads are not extra vertices in a normal tree. The apparent non-polynomial reader layer is a consequence of expressing a family of observation sorts one-sortedly.

Remark 2.4 (A simpler polynomial comparison: Mealy syntax). For deterministic Mealy machines, a standard free signature has one O-ary update operation for each $a \in A$, with no separate pure reader operation or reader equations. Its signature endofunctor on **Set** is already polynomial, $Q(X) = \coprod_{a \in A} X^O$, and its free monad has polynomial tree arities, so the sorted diagonal-algebra argument is unnecessary. Our graded polynomial and temporal interaction constructions have analogues there. We retain Moore theory as the principal example because the many-sorted splitting explains its observational equations and because its complete behaviours have the familiar O^{A^*} form; we do not claim the diagonal-splitting method is needed for Mealy syntax.

3 Filtering the Moore monad and recognizing its algebras

The sorted presentation gives us finite programs, but does not yet say how much future a program needs. The bridge to clocks begins by *choosing* a resource grading: here a budget bounds the number of state-changing steps on any branch. A read consumes no budget, so a program may ask a question without advancing time. The point of this section is to make the resource bookkeeping part of the algebraic laws, not an annotation added after a tree has been evaluated.

Let $T_n X$ be the subset of normal Moore terms in which each branch has at most n steps. We have inclusions $i_{m,n} : T_m \to T_n$ for $m \leq n$, with

$$\eta : \mathrm{Id} \to T_0, \qquad \mu_{m,n} : T_m T_n \to T_{m+n}. \tag{6}$$

Substitution adds upper bounds on depths; normalizing may lower them. Thus $T_\bullet$ is a lax monoidal functor from $(\mathbb{N}, \leq, +, 0)$ to endofunctors under composition: an ordered graded monad, or equivalently a filtered monad, in the sense of graded-monad semantics [6]. Note $T_0(X) = X^O$, not X.

Definition 3.1 (Graded algebra)**.** *A graded Moore algebra is an increasing sequence $B_0 \to B_1 \to \cdots$ together with maps*

$$\beta_{m,n} : T_m B_n \to B_{m+n}$$

satisfying the graded unit, associativity, and weakening laws. Its morphisms are levelwise natural transformations preserving these actions. We write $\mathcal{A} = \mathrm{Alg}_{\mathrm{gr}}(T)$.

Proposition 3.2 (Generators-and-relations description)**.** *A graded Moore algebra consists equivalently of reader algebras $g_n : B_n^O \to B_n$, reader-algebra homomorphisms $w_n : B_n \to B_{n+1}$, and arbitrary maps $\sigma_{a,n} : B_n \to B_{n+1}$ satisfying*

$$w_{n+1} \sigma_{a,n} = \sigma_{a,n+1} w_n. \tag{7}$$

There is no law commuting σ with g.

Proof. The grade-zero actions restrict to reader algebra structures; the grade-one generating steps induce σ. Weakening compatibility for steps is (7), and weakening compatibility for reads says that w preserves the reader structures. Conversely, evaluate a normal tree recursively, adding unused allowance with w. The reader identities and (7) ensure the result is independent of presentation and the graded monad axioms are exactly substitution compatibility. □

At each grade, diagonal splitting gives sort carriers $B_{n,o}$ and maps

$$w_{n,o} : B_{n,o} \to B_{n+1,o}, \qquad \sigma_{a,n,o} : \prod_{p \in O} B_{n,p} \to B_{n+1,o}, \tag{8}$$

with the coordinate equation

$$w_{n+1,o}\sigma_{a,n,o} = \sigma_{a,n+1,o}(\prod_{p} w_{n,p}). \tag{9}$$

4 The polynomial distributive law for graded syntax

Grading supplies an increasing sequence of sets of programs, but for the subsequent active–inert comparison we need more: an explicit description of what the operations *are* and how substitution changes their grades. There are two independent-looking ingredients, unused allowance and actual updates. The distributive law below is the compatibility that allows us to combine them without losing polynomial arities. It is also what ensures that weakening a whole computation agrees with weakening the continuations it will use.

The calculation in the preceding section is algebraic. We now make the tree arities canonical at the level of discrete grade–observation sorts. Let

$$C = \mathbb{N} \times O, \qquad \mathcal{D} = \mathbf{Set}^C,$$

where C is a *discrete set* of colours. This qualification prevents confusion with the diagram category of increasing grades.

Definition 4.1 (Two monads)**.** *Define a weakening monad W on $\mathcal{D}$ by*

$$(WX)_{n,o} = \coprod_{m \le n} X_{m,o}. \tag{10}$$

Its unit uses $m = n$, and its multiplication composes inequalities. Define a strictly budgeted update-tree monad G by

$$(GX)_{0,o} = X_{0,o}, \quad (GX)_{n+1,o} = X_{n+1,o} + \coprod_{a \in A} \prod_{p \in O} (GX)_{n,p}. \tag{11}$$

A G-tree rooted at (n, o) has an a-vertex with one input of every observation sort, each of budget $n - 1$; its leaves carry generators of precisely the remaining grades. Composition substitutes a tree into a matching leaf.

Both W and G are polynomial and cartesian monads. Although G is defined recursively, its trees have maximum depth n at root grade n, so no transfinite construction is needed for any fixed n.

Theorem 4.2 (Distributive law and polynomial normal forms)**.** *There is a distributive law*

$$\lambda : WG \Longrightarrow GW \tag{12}$$

that moves weakening down through update vertices. The composite $M_{\mathrm{gr}} = GW$ is a cartesian polynomial monad on $\mathcal{D}$ and its algebras are precisely the sorted data (8) satisfying (9). Its normal forms are

$$(M_{\mathrm{gr}}X)_{n,o} \cong \coprod_{\tau \in \mathsf{Tree}_{n,o}} \coprod_{\substack{m_\ell \le n - d(\ell) \\ \ell \in L(\tau)}} \prod_{\ell \in L(\tau)} X_{m_\ell, p(\ell)}, \tag{13}$$

where τ ranges over finite O-coloured, A-labelled update trees of height at most n, with leaf depths $d(\ell)$ and leaf observation sorts $p(\ell)$.

Proof. An element of $(WGX)_{n,o}$ is a pair consisting of a lower budget $m \le n$ and a G-tree t rooted at (m, o). On a leaf $x \in X_{m,o}$, λ produces the W-label $[m, x]$ at root budget n. When the tree starts with an update a, it moves the update to root budget n and recursively moves its p-children from grade $m - 1$ to grade $n - 1$. This recursively defines λ. The two Beck axioms involving weakening follow because successive changes of root budget leave precisely the same original grade on each leaf. The two axioms involving G follow by induction on the outer tree: moving weakening to the leaves commutes with substitution at every matching leaf. All maps are natural because the construction alters only shapes and grades, never variable labels.

The distributive law yields the composite monad GW. Its algebras are a W-algebra and a G-algebra on the same sorted family obeying the compatibility of pushing w through each generating update. This is exactly (9). Formula (13) follows by unfolding $G(WX)$ and choosing, for each leaf of residual grade $n - d(\ell)$, one original grade m_ℓ at most that residual grade. It is a coproduct of finite products of coordinate projections, hence a polynomial functor.

The unit is the one-leaf identity operation. Its naturality squares are pullbacks because the leaf's variable label is unchanged. Multiplication flattens a tree of trees, propagating outer weakening into inner trees before grafting. For a map of sorted generators $f : X \to Y$, specifying a flattened X-labelling and a compatible tree-of-trees decomposition over Y recovers a unique X-labelling of the original decomposition, leaf by leaf. This proves cartesianness of multiplication. The monad is therefore cartesian polynomial. □

Theorem 4.3 (Exact comparison with graded Moore algebras)**.** *For finite nonempty O there is an equivalence*

$$\boxed{\mathcal{A} \simeq \mathrm{Alg}(M_{\mathrm{gr}})_{\mathrm{adm}},}$$

where admissibility means that, at each grade n, either all $B_{n,o}$ are empty or all are nonempty. If $\underline{X}$ has $\underline{X}_{0,o} = X$ and $\underline{X}_{n,o} = \varnothing$ for $n > 0$, then

$$T_n(X) \cong \prod_{o \in O} (M_{\mathrm{gr}}\underline{X})_{n,o}. \tag{14}$$

The graded free algebra $F(X)$ generated at grade zero has $F(X)_n = T_n(X)$.

Proof. The category equivalence follows from stagewise diagonal splitting and Theorem 4.2. Reader-homomorphic weakenings act coordinatewise; arbitrary unary updates have precisely the coordinate maps (8). They obey exactly (9). For $\underline{X}$ every terminal leaf originates at grade zero, so the trees in (13) have depth at most n and leaves in X. Taking the product over initial observation coordinates reconstructs precisely the normal one-sorted computation trees of depth at most n. □

Remark 4.4 (The noncartesian obstruction is not being denied). For $O = \{0, 1\}$, the sorted family $B_{n,0} = 1$, $B_{n,1} = \varnothing$ at every grade admits a GW-algebra structure: all update-operation domains are empty and weakening is the evident map. It does not arise from a one-sorted reader algebra. Likewise the original one-sorted reader monad remains noncartesian. The polynomial result is on $\mathbf{Set}^C$, with an admissibility restriction on its algebras, not a new polynomial presentation of that original monad *on* **Set**.

4.1 An explicit two-step tree

Let $O = \{0, 1\}$, $A = \{a, b\}$. At root sort $(3, 0)$ an update a has children of sorts $(2, 0)$ and $(2, 1)$. Choose the sort-0 continuation to be a b-update with leaves at grade zero, and the sort-1 continuation to be a weakened return at grade one. The operation has three leaves: two labelled $(0, 0)$ and $(0, 1)$ reached by the input word ab, and one labelled $(1, 1)$ reached by a. This example will recur as the first nontrivial *guarded family of temporal paths*.

5 From polynomial operations to active and inert trees

Why retain the tree shapes after identifying the graded monad? The set of operations at each grade remembers how many steps a computation might take, but our intended temporal comparison must also recognize *parts* of an experiment and the ways programs are put together. Active and inert maps supply a controlled language for substitution and selecting a subcomputation. This is the syntactic counterpart of restricting a clock to a smaller temporal piece; the precise connection with temporal paths will be made only after the comodel side has been constructed.

The polynomial monad supplies not just a set of program trees, but a coloured operad $\mathcal{P}$ of operations, with colours $C = \mathbb{N} \times O$. An operation is a tree with root sort (n, o), leaves $((m_\ell, p_\ell))_\ell$, and inequality $m_\ell + d(\ell) \leq n$. Substitution at a leaf is operadic composition; permutation of an enumeration of leaves gives the symmetric action.

Let Ω_{den} denote the dendroidal category of finite rooted trees. The dendroidal nerve $N_d\mathcal{P}$ is a presheaf on Ω_{den}. Its category of elements

$$\Omega_{\mathcal{P}} = \int_{\Omega_{\mathrm{den}}} N_d\mathcal{P} \tag{15}$$

contains trees whose edges are colours and whose vertices are decorated by matching operations of $\mathcal{P}$.

Proposition 5.1 (Active–inert factorization survives decoration)**.** *The usual active–inert factorization of dendroidal maps lifts to $\Omega_{\mathcal{P}}$. Its inert maps select connected full subtrees preserving local incidence; active maps preserve outer boundaries and encode composition and units. The factorization is unique up to the standard isomorphism of the intermediate dendroidal tree.*

Proof. For a decorated map $(T, q) \to (U, r)$ with underlying dendroidal map $f = i \circ a$, form the intermediate decoration by pulling r back along i. Pulling it back along a recovers q. This constructs a decorated active–inert factorization, and uniqueness follows from uniqueness downstairs together with forced pullback of the intermediate decoration. □

A $\mathcal{P}$-algebra B determines a presheaf of edge labellings on $\Omega_{\mathcal{P}}$:

$$N_B(T, q) = \big\{(b_e) : b_e \in B_{\mathrm{col}(e)}, \quad b_{\mathrm{out}(v)} = q_v^B((b_{\mathrm{in}(v,\ell)})_\ell)\big\}.$$

An edge-labelling on a tree is equivalent to compatible labellings of its corollas and edges. The defining property of an algebra, as opposed to an arbitrary relation between inputs and outputs, is that each corolla's joint input-restriction map is a bijection. With these conditions, the elementary pieces recover the operations; active maps enforce substitution and units. The precise essential-image claim is recalled in Appendix A; it is restricted to this coloured decorated-tree setting and does not assert Raynor's graph nerve theorem [13, 9].

6 Time budgets, bounded observations, and comodels

Up to now the objects have been *programs*: a tree describes what may be asked and which continuations are available. A machine must answer those questions and choose a continuation as

it evolves. In ordinary Moore semantics the final comodel records every possible finite input word at once; that is an unbounded completed future. To compare with finite clocks we instead ask what a machine state can support with *only n updates available*. The resulting inverse sequence turns the same grading around: algebraic grades allow increasingly elaborate terms, whereas comodel horizons are restricted as time is consumed.

6.1 The graded comodel axioms

Let $\mathcal{E} = \mathbf{Set}^{\mathbb{N}^{\mathrm{op}}}$ be the topos of inverse sequences. A temporal object is

$$S_0 \xleftarrow{r_{1,0}} S_1 \xleftarrow{r_{2,1}} S_2 \longleftarrow \cdots .$$

The right-adjoint later modality from guarded temporal semantics [7] is

$$(\triangleright S)_0 = 1, \qquad (\triangleright S)_{n+1} = S_n.$$

Definition 6.1 (Budget-indexed comodel)**.** *A graded comodel has temporal carriers $S \in \mathcal{E}$ and natural cooperations*

$$\theta_{m,n,X} : T_m X \times S_{m+n} \longrightarrow X \times S_n \tag{16}$$

that satisfy the unit, substitution, grade-weakening and horizon-restriction equations. The category is $\mathcal{C} = \mathrm{Comod}_{\mathrm{gr}}(T)$.

Theorem 6.2 (Guarded presentation)**.** *There is an equivalence*

$$\boxed{\mathcal{C} \simeq \mathrm{Coalg}_{\mathcal{E}}(\Delta O \times (\triangleright -)^A).}$$

Equivalently, a graded comodel consists of observations $o_n : S_n \to O$ and update maps $\delta_{a,n} : S_{n+1} \to S_n$ satisfying

$$o_n r_{n+1,n} = o_{n+1}, \qquad r_{n+1,n}\delta_{a,n+1} = \delta_{a,n} r_{n+2,n+1}. \tag{17}$$

No condition forces an update to preserve observations.

Proof. The reader equations force the grade-zero cooperation to read a unique observation without changing the current state. The elementary grade-one updates give $\delta_{a,n}$. Naturality of these cooperations in the remaining horizon gives (17). Conversely, execute any bounded term recursively: inspect o_n on a read, apply δ_a on an update, and restrict if the allocated allowance exceeds the number of executed steps. The equations ensure this execution is independent of the term presentation, and the action of a substituted term is sequential execution. The displayed data are exactly the component form of a coalgebra for $\Delta O \times (\triangleright -)^A$. □

Theorem 6.3 (Terminal finite futures)**.** *The terminal graded comodel has components*

$$\boxed{\Omega_n = O^{A^{\leq n}}}, \tag{18}$$

with restrictions by truncation, $o_n(b) = b(\varepsilon)$, and

$$\delta_{a,n}(b)(w) = b(aw) \quad (|w| \leq n).$$

Its inverse limit is the ordinary final Moore comodel $\Omega_\infty \cong O^{A^*}$ with residual updates $\partial_a b(w) = b(aw)$.

Proof. A bounded future of horizon $n+1$ consists of a current observation and a horizon-n future for every possible input letter:

$$\Omega_{n+1} \cong O \times \Omega_n^A.$$

This proves terminality inductively, because the same recursion uniquely determines the behaviour maps from any guarded coalgebra. The limit of restrictions is a function assigning compatible observations to every finite input word, i.e. an element of O^{A^*}. The shift maps on the limit are precisely the residual updates. □

6.2 The budgeted-word category

The guarded-coalgebra equations describe transitions between horizons, but they have not yet produced a *category of temporal arrows*. Packaging restrictions and updates into one category shows exactly what the grading contributes and what the input alphabet adds: an arrow records an executed word together with allowance that may simply be discarded. This packaging is what will let us compare finite-horizon machines with both Garner's behaviour category and DOTS clock restrictions.

Define a category $\mathcal{J}_A$ with objects natural-number budgets and

$$\mathcal{J}_A(n,m) = \begin{cases} A^{\leq n-m}, & n \geq m, \\ \varnothing, & n < m. \end{cases} \tag{19}$$

An arrow labelled w executes w and discards the unused budget $n - m - |w|$. Composition concatenates labels in execution order. The generating arrows $n+1 \to n$ are the empty word (restriction) and the one-letter words (updates). Their commuting squares are the second equation in (17).

Theorem 6.4 (Temporal presheaf/slice description)**.** *There are natural equivalences*

$$\boxed{\mathcal{C} \simeq [\mathcal{J}_A, \mathbf{Set}]/\Omega \simeq [\textstyle\int_{\mathcal{J}_A} \Omega, \mathbf{Set}].} \tag{20}$$

Here $\Omega : \mathcal{J}_A \to \mathbf{Set}$ acts on $w : n \to m$ by $b \mapsto (v \mapsto b(wv))$, truncated to $|v| \leq m$. Thus $\mathcal{C}$ is a Grothendieck topos. The category of elements has bounded futures as objects and update/truncation arrows, providing a finite-horizon counterpart of Garner's behaviour category.

Proof. Data (17) give a functor $S : \mathcal{J}_A \to \mathbf{Set}$, since any interleaving of an update with unused-budget discards can be reordered using the commuting squares. The observation maps uniquely determine a natural transformation $\mathrm{beh} : S \to \Omega$ by

$$\mathrm{beh}_n(s)(w) = o_{n-|w|}(\delta_{w,n}(s)).$$

Conversely, a map of functors $S \to \Omega$ supplies o_n by evaluation at ε, and the generating arrows give restriction and update maps. Naturality forces the transformation to coincide with the displayed behaviour map, proving the first equivalence. The second is the category-of-elements equivalence: take fibres over each (n, b) and reconstruct S_n by disjoint union over the b's. □

For a unary input alphabet $A = \{a\}$, arrows $n \to m$ are determined by an offset $j \in \{0, \ldots, n-m\}$. Hence

$$\boxed{\mathcal{J}_A^{\mathrm{op}} \cong \mathsf{Int}_{\mathbb{N}} \cong \mathrm{Tw}(\mathbf{B}\mathbb{N}),} \tag{21}$$

with the convention that an interval embedding $m \to n$ is determined by its offset inside the larger interval. This is the interval-indexing structure used in temporal accounts of clock restrictions [14]; the representable-clock comparison with DOTS is made precise below. A coverage and sheaf gluing are additional choices, not consequences of (20).

6.3 Relation to Garner and polynomial cofree comonoids

For an ordinary Moore machine, the polynomial endofunctor

$$p(X) = O \times X^A \cong \coprod_{o \in O} X^A$$

has coalgebras consisting of observations and input-indexed transitions. Garner's behaviour category organizes their complete observable futures and transitions [5]. Niu–Spivak's polynomial

account identifies coalgebras for a polynomial interface with copresheaves on a category obtained from its cofree polynomial comonoid [12]. In this example its complete labelled A-ary trees have positions $b \in O^{A^*}$ and directions given by input histories. Our category of elements $\int_{\mathcal{J}_A} \Omega$ is a *bounded-horizon counterpart*, with additional truncation/discard arrows; it is not claimed to be literally the cofree-comonoid category at finite level. These comparisons locate the ordinary copresheaf representation in existing work. What is new here concerns its coupling to the *graded syntax* and the temporal-valued evaluation quotient.

6.4 Finite input charts and representable clocks

There is a further comparison that matters for the intended audience. The representables of the bounded behaviour category select a state with its *entire* branching future fixed. A DOTS clock, by contrast, can follow a specified input schedule and observe just one linear history. To avoid identifying these two different notions of clock, we next express a linear charted execution as a family of representable behaviours, indexed by the possible counterfactual extensions of its observed path.

Let $\mathcal{B}_{\mathrm{fin}} = \int_{\mathcal{J}_A} \Omega$. Under the equivalence (20), a comodel S corresponds to a covariant functor $F_S : \mathcal{B}_{\mathrm{fin}} \to \mathbf{Set}$ with

$$F_S(n, b) = \{s \in S_n : \mathrm{beh}_n(s) = b\}.$$

For any (n, b), the covariant representable $\mathcal{B}_{\mathrm{fin}}((n, b), -)$ selects precisely a state whose *entire bounded branching future* is b, by Yoneda. This is already close to Garner's representables, but it is not yet a linear DOTS clock: it retains the counterfactual futures after every alternative input word.

For a fixed input word $u = a_1 \cdots a_\ell$, $\ell \le n$, write $u_i = a_1 \cdots a_i$, $u_0 = \varepsilon$, and fix observations $\alpha_0, \ldots, \alpha_\ell \in O$. A *charted finite path* in S is a state $s \in S_n$ satisfying

$$o_{n-i}(\delta_{u_i,n}(s)) = \alpha_i \qquad (0 \le i \le \ell).$$

The indices record the remaining horizon; inputs are fixed as part of the chart, not quantified over by the linear trajectory. We obtain the following exact comparison.

Proposition 6.5 (Fixed-input path as a coproduct of representables)**.** *There is a natural bijection*

$$\boxed{\begin{aligned} E_{u,\alpha}(S) &:= \{s \in S_n : \ o_{n-i}(\delta_{u_i,n} s) = \alpha_i \ \forall i \le \ell\} \\ &\cong \coprod_{\substack{b \in \Omega_n \\ b(u_i) = \alpha_i \ \forall i \le \ell}} \mathrm{Nat}(\mathcal{B}_{\mathrm{fin}}((n, b), -), F_S). \end{aligned}} \tag{22}$$

Proof. Yoneda identifies the summand at b with $F_S(n, b) = \mathrm{beh}_n^{-1}(b)$. Every $s \in S_n$ has one bounded future $b = \mathrm{beh}_n(s)$. Its values along the chosen prefixes are precisely the observations in the definition of $E_{u,\alpha}(S)$. Taking the disjoint union over all possible counterfactual extensions b proves the assertion, naturally in comodel maps. □

The clock-based programme of Lynch–Myers–Rischel–Staton characterizes behaviours as morphisms *from* clock systems, yielding covariant representable functors and, for compositional systems theories, a Span-valued representable double functor [15]. Proposition 6.5 is a finite, fixed-chart comparison: its left side is the trajectory data of a linear input schedule, while its right side retains all the branching futures compatible with that schedule. To identify the *systems and clock morphisms* on both sides requires the input-charted morphisms of the clocks chapter; we do not claim that a branching representable is itself a DOTS linear clock. The known stochastic and nondeterministic clock constructions similarly should not be identified with our deterministic input branching.

6.5 Complete finite history trees

Let Hist_A consist of nonempty finite prefix-closed subsets $P \subseteq A^*$ with offset embeddings $P \to Q$ determined by words u such that $uP \subseteq Q$. For *finite nonempty* A, the complete trees $P_n = A^{\leq n}$ define a fully faithful embedding

$$\mathcal{J}_A^{\mathrm{op}} \hookrightarrow \mathsf{Hist}_A, \tag{23}$$

whose image is the full subcategory on these P_n. Indeed $uP_m \subseteq P_n$ iff $|u| \leq n - m$. The assumptions matter: $A = \varnothing$ makes all P_n identical, while for infinite A the complete trees are not finite. This locates our bounded time category inside a genuinely branching finite-history category, without identifying it with that entire category.

Remark 6.6 (The supplied grading, and what input operations add). The word-length map

$$\lambda : \mathcal{J}_A \longrightarrow \mathcal{J}_{\{*\}}, \qquad (n \xrightarrow{w} m) \longmapsto (n \xrightarrow{|w|} m)$$

is a functor, identity on budgets. The target records offset/duration; the source remembers the particular input word. Thus an ordered resource grading supplies the budgets and their composition, while A refines the temporal arrows. For positive integral update costs $c(a)$, the analogous category uses $|w|_c \leq n - m$; even for a unary letter of cost 2 it is no longer the ordinary discrete interval category. A continuous-time version would therefore have to start from an independently specified $\mathbb{R}_{\geq 0}$-graded flow theory, not from the ungraded Moore signature.

From input charts to physical trajectories. The paths classified above are indexed by finite input words and remaining budgets. In the *Temporal Interlude* [16], an explicit one-read runner is instead unfolded into a sheaf of full input, output, and state histories over contiguous intervals; its serial composition is a pullback of trajectory sheaves. Selecting a word from the present branching future supplies a linear input chart for such an execution, but the passage to a timed history also needs an implementation and a choice of physical clock. We therefore do not identify the two indexing categories or their corresponding toposes.

7 The one-observation baseline: an ordinary profunctor tensor

Before constructing a more elaborate pairing, it is useful to identify which part of the answer is already standard. If the observation set has one element, a computation has no read-dependent branches and the algebra and comodel sides reduce to a presheaf and copresheaf on the same budgeted-word category. Their evaluation is simply the familiar profunctor coend. The additional structure of the general Moore case can then be measured against this baseline, rather than mistaken for a new unary tensor construction.

Before the general interaction law, the case $O = 1$ should be isolated. There is then no genuine branching over observations. The graded Moore algebras are precisely functors

$$B : \mathcal{J}_A^{\mathrm{op}} \to \mathbf{Set},$$

while graded comodels are functors $S : \mathcal{J}_A \to \mathbf{Set}$, since $\Omega = 1$ and the slice (20) is trivial. The unary operad of graded polynomial operations is the category $\mathcal{J}_A^{\mathrm{op}}$: an operation from a leaf budget m to a root budget n is a word $w \in \mathcal{J}_A(n, m)$, including any unused allowance. This corrects a possible reversal of variance: it is $\mathcal{J}_A^{\mathrm{op}}$, not $\mathcal{J}_A$, on the algebra side.

For $k \geq 0$, the horizon shift $D_k S$ is again a covariant functor on $\mathcal{J}_A$. In the tensor defined below, balancing along weakenings and update operations is exactly balancing along all the arrows of $\mathcal{J}_A$. Thus

$$\boxed{(B \mathbin{\widehat{\otimes}_T} S)_k \cong \int^{n \in \mathcal{J}_A} B(n) \times (D_k S)(n) = B \otimes_{\mathcal{J}_A} D_k S.} \tag{24}$$

This is the ordinary coend pairing of a presheaf and a copresheaf. The first transpose is the standard tensor–Hom adjunction and cancellation against free graded generators is co-Yoneda. Accordingly, our contribution is *not* a new tensor formula for unary time. It is the extension to observation-branching sorted operations, its two-sided temporal Chu interpretation, and the analysis of completion. This baseline also clarifies that the finite-horizon object provides a temporal indexing category, while the observation fibres provide the extra data of the corresponding behaviour category.

8 A temporal-valued algebra–comodel tensor

We now have finite programs and finite-horizon implementations, but no single evaluation that retains both. The ordinary Sweedler-style execution quotient balances an algebraic term against a comodel state and returns a set; by itself it does not say how much temporal information remains after the computation. Introducing an extra horizon k makes that remainder visible. The resulting object is not simply a family of unrelated quotients: the restriction maps turn it into a temporal presheaf, and the balancing law will explain how an update may pass from the syntax side to the state side.

The ordinary algebra–comodel tensor is a Set-valued execution quotient. To retain finite-time observations, we keep an additional index recording the *remaining horizon*. For the rest of the main results let $B \in \mathcal{A}$ and $S \in \mathcal{C}$.

Definition 8.1 (Temporal balanced tensor)**.** *At horizon k, begin with*

$$\int^{n\in(\mathbb{N},\leq)} B_n \times S_{n+k},$$

where the coend uses the covariant algebra weakening and contravariant comodel restriction. Further identify

$$(\beta_{m,n}(t), s) \sim (b, s') \quad \text{if} \quad \theta_{m,n+k,B_n}(t, s) = (b, s'), \tag{25}$$

for $t \in T_m(B_n)$ and $s \in S_{m+n+k}$. The resulting set is $(B \mathbin{\widehat{\otimes}_T} S)_k$. Horizon restriction induces $(B \mathbin{\widehat{\otimes}_T} S)_{k+1} \to (B \mathbin{\widehat{\otimes}_T} S)_k$, and hence

$$\boxed{B \mathbin{\widehat{\otimes}_T} S \in \mathcal{E}.}$$

There are three elementary Moore execution relations, understood inside the component at horizon k:

$$(w_n b, s) \sim (b, r_{n+k+1,n+k} s), \qquad s \in S_{n+k+1}, \tag{26}$$

$$(g_n((b_o)_o), s) \sim (b_{o_{n+k}(s)}, s), \qquad s \in S_{n+k}, \tag{27}$$

$$(\sigma_{a,n}(b), s) \sim (b, \delta_{a,n+k}(s)), \qquad s \in S_{n+k+1}. \tag{28}$$

These generate the full execution relation because every normal term is obtained from reads and steps by finite substitution. The same relations are what make the tensor compatible with the algebra–comodel pairing in the ordinary ungraded setting.

8.1 The first transpose: temporal handlers

An execution quotient earns its place in a Sweedler comparison only if it supports the corresponding algebra–comodel evaluation principle. The first transpose answers this: a test of the tensor can instead be viewed as an algebraic interpretation of programs into temporal-valued functions on states. The degree of a program then appears as a shift of the state horizon. The subsequent free-algebra calculation checks this interpretation against ordinary execution of a program with a return label.

For $K \in \mathcal{E}$, set

$$\boxed{[S,K]_n = \mathcal{E}(D_nS, K), \qquad (D_nS)_k = S_{n+k}.} \tag{29}$$

The following formulas define a graded Moore algebra on these sets. Given $f \in [S,K]_n$, let

$$(w_nf)_k(s) = f_k(r_{n+k+1,n+k}s), \tag{30}$$

$$\big(g_n((f_o)_o)\big)_k(s) = (f_{o_{n+k}(s)})_k(s), \tag{31}$$

$$(\sigma_{a,n}f)_k(s) = f_k(\delta_{a,n+k}s). \tag{32}$$

Here the state s is taken at the unique stage for which the expression is typed. The read equations follow from the fact that the observation is single-valued and reading does not change the state. Compatibility of weakening and update is exactly (17).

Theorem 8.2 (Temporal tensor–Hom adjunction)**.** *For every $S \in \mathcal{C}$, $B \in \mathcal{A}$, and $K \in \mathcal{E}$ there is a natural bijection*

$$\boxed{\mathcal{E}(B \mathbin{\widehat{\otimes}_T} S, K) \cong \mathcal{A}(B, [S,K]).} \tag{33}$$

Thus $- \mathbin{\widehat{\otimes}_T} S : \mathcal{A} \to \mathcal{E}$ is a left adjoint and preserves colimits.

Proof. Given a natural map $h : B \mathbin{\widehat{\otimes}_T} S \to K$, transpose its representatives by

$$\overline{h}_n(b)_k(s) = h_k([b,s]), \qquad s \in S_{n+k}.$$

Horizon naturality makes $\overline{h}_n(b)$ an element of $\mathcal{E}(D_nS, K)$. The coend relation forces $\overline{h}$ to respect weakening; the reader and step relations (27)–(28) force it to preserve the operations (31)–(32). Hence $\overline{h}$ is a graded algebra morphism. Conversely a graded algebra map $f : B \to [S,K]$ defines $h_k([b,s]) = f_n(b)_k(s)$. Its compatibility with weakening and the generated algebra operations makes the formula independent of the representatives and the balancing relation. The two constructions are inverse and natural. □

8.2 Free cancellation with a generator at an arbitrary grade

Let $F^{(d)}X$ be the free graded algebra generated by a set X at grade d. Its stages are

$$(F^{(d)}X)_n = \begin{cases} \varnothing, & n < d, \\ T_{n-d}X, & n \geq d. \end{cases} \tag{34}$$

In grade d, an X-generator produces its grade-zero free reader term $\eta(x) \in T_0X = X^O$. This detail matters when evaluating against a state with a specific observation.

Lemma 8.3 (Free graded cancellation)**.** *There is a natural isomorphism of temporal objects*

$$\boxed{F^{(d)}X \mathbin{\widehat{\otimes}_T} S \cong X \times D_dS.} \tag{35}$$

Here $X \times -$ means product with the constant temporal object ΔX.

Proof. At residual horizon k, represent a class by a term $t \in T_{n-d}X$ and state $s \in S_{n+k}$, for $n \geq d$. Execute t against s, obtaining $(x, s') \in X \times S_{d+k}$. This is invariant under weakening and under execution relations, because evaluating a substituted program agrees with sequential execution. It defines the forward map. Conversely send (x, s') to $[\eta(x), s']$, with $\eta(x)$ at grade d. The graded unit gives one composite identity. For the other, the defining balancing relation identifies $[t,s]$ with the class of its executed return $[\eta(x), s']$. These maps respect restriction in k and are natural. □

In particular, for the free algebra $F(X) = F^{(0)}X$,

$$F(X) \mathbin{\widehat{\otimes}_T} \Omega \cong X \times \Omega, \qquad \varprojlim_k (F(X) \mathbin{\widehat{\otimes}_T} \Omega)_k \cong X \times O^{A^*}. \tag{36}$$

This verifies the cancellation that had previously appeared as an unproved dependency between the short notes.

9 Coherence of the graded interaction: trees, paths, and evaluation

Two descriptions of an update have now been developed independently. Polynomial syntax gives an operation with several potential return leaves; the bounded comodel determines which leaf an actual state reaches. The temporal tensor balances the algebraic and operational descriptions, but to claim these are parts of *one* construction we must verify that it identifies precisely the path the state executes, and that the identification respects substitution. This section supplies that check. The result is a coherence theorem, not a new source of branching choices.

The following comparison checks that the sorted polynomial operations, graded comodel cooperation, and balanced tensor implement the same interaction law. The proof is a coherence calculation, rather than an independent mechanism for selecting branches: comodel cooperation already performs that selection. It is nevertheless needed to identify the three categorical presentations precisely, and to control substitution.

9.1 Sorted execution, including reader evaluation

For a graded algebra B, split $B_n \cong \prod_o B_{n,o}$. For a comodel S, split each stage into observation fibres $S_{n,o} = o_n^{-1}(o)$, giving

$$S_n \cong \coprod_{o\in O} S_{n,o}, \quad r_{n,o} : S_{n+1,o} \to S_{n,o}, \quad \delta_{a,n,o} : S_{n+1,o} \to \coprod_{p\in O} S_{n,p}.$$

The comodel update chooses exactly one successor observation sort; the algebra update supplies a continuation for every sort. The sorted tensor at horizon k is the quotient of

$$\coprod_{n\geq 0} \coprod_{o\in O} B_{n,o} \times S_{n+k,o} \tag{37}$$

by weakening and update execution, where an update $\delta_{a,n+k,o}(s) \in S_{n+k,p}$ selects the coordinate b_p of the algebraic input tuple.

Proposition 9.1 (Sorting preserves the evaluation pairing)**.** *There is a natural isomorphism between the one-sorted temporal tensor of Definition 8.1 and the sorted execution quotient (37), compatible with the tensor–Hom transpose and grade-shifted free cancellation.*

Proof. Send $[b, s]$ with $b \in B_n \cong \prod_o B_{n,o}$ and $s \in S_{n+k,o}$ to $[b_o, s]$. The reader relation (27) ensures that only the chosen coordinate matters. Conversely, when B_n is inhabited, any coordinate b_o admits an extension to a tuple; the reader relation makes all choices equivalent. If B_n is empty, admissibility makes every sort empty, so no representatives are lost. Under this quotient identification, graded weakening is coordinatewise and step execution selects precisely one output fibre. The result is natural in horizon restriction and in algebra and comodel maps. □

9.2 A tree determines a family of labelled temporal arrows

Let $\mathrm{PathFam}(\mathcal{J}_A)$ be the coloured symmetric operad with colours $n \in \mathbb{N}$ and multimorphisms

$$\mathrm{PathFam}(\mathcal{J}_A)(m_1, \dots, m_r; n) = \prod_{i=1}^{r} \mathcal{J}_A(n, m_i).$$

An operation is a family of arrows sharing a source budget; substitution composes the arrows at the corresponding leaves. With this variance convention, $\mathrm{PathFam}(\mathcal{J}_A)$ is the cocartesian operad associated to $\mathcal{J}_A^{\mathrm{op}}$: its multimorphisms are products of unary hom-sets. Thus the genuinely new information in Π is the observation-branching polynomial shape from which a family of unary paths is extracted. For a sorted polynomial operation t with root (n, o) and leaf slots (m_ℓ, p_ℓ), let

$w_\ell \in A^*$ be the word of update labels along the root-to-leaf path. Its length equals the number of update vertices on that path; by the grade condition,

$$|w_\ell| + m_\ell \le n.$$

Thus w_ℓ specifies an arrow $n \to m_\ell$ of $\mathcal{J}_A$.

Theorem 9.2 (Leafwise temporal-path semantics)**.** *Forgetting the observation-sort labels extends to a morphism of coloured operads*

$$\boxed{\Pi : \mathcal{P} \longrightarrow \mathrm{PathFam}(\mathcal{J}_A), \qquad \Pi(t) = (n \xrightarrow{w_\ell} m_\ell)_{\ell \in L(t)}.} \tag{38}$$

It preserves units, symmetry and substitution exactly. If O has one element, the unary operations of $\mathcal{P}$ form $\mathcal{J}_A^{\mathrm{op}}$; for several observations, one cannot in general replace the family of paths by a single time arrow.

Proof. The grade inequality makes every component of (38) well-defined. The unit has the empty word. Substituting a tree u at a leaf ℓ of t replaces the leaf's word w_ℓ by concatenations $w_\ell v_j$, one for every leaf j of u. This is precisely composition of arrows in $\mathcal{J}_A$. The bijection between leaf sets on the two sides also verifies compatibility with symmetry. For $|O| = 1$, every tree is linear and a normal operation is a word with unused budget; viewed as an operation from input m to output n, it corresponds to an arrow $n \to m$ of $\mathcal{J}_A$, hence to $\mathcal{J}_A^{\mathrm{op}}$. □

Remark 9.3 (The weakness of numerical duration). Already with $O = \{0, 1\}$, a root update a whose two continuations respectively update by b and by c yields paths ab and ac. A map assigning a single temporal arrow to the entire tree would have to discard branch information. The correct comparison is leafwise. This is also why passing from trees to mere depth grades is strictly lossy.

9.3 Guard cylinders select the executed leaf

Let $\Omega_{n,o}$ be the bounded futures with initial observation o, and fix an operation t rooted at (n, o). Let k denote the desired horizon after execution. For a leaf ℓ define the guard

$$U_{t,\ell}^{(k)} \subseteq \Omega_{n+k,o}$$

to be those bounded futures whose answers to the observations encountered along the tree select exactly leaf ℓ. The guard is defined recursively: an update by a shifts the future to the subtree's horizon and then branches according to $b(a) \in O$. At a leaf it is the entire possible future. Define a residual map on each guard by

$$\rho_{t,\ell}^{(k)} : U_{t,\ell}^{(k)} \to \Omega_{m_\ell+k,p_\ell}, \qquad \rho_{t,\ell}^{(k)}(b)(v) = b(w_\ell v). \tag{39}$$

The inequality $|w_\ell| + m_\ell \le n$ gives enough coordinates to evaluate the right-hand side. Discarding unused budget is part of this arrow.

Proposition 9.4 (Guarded paths compose by pullback)**.** *For the terminal graded comodel, the sets $U_{t,\ell}^{(k)}$ form a finite disjoint partition of $\Omega_{n+k,o}$ by cylinder subsets (clopen for the product topology with O discrete). Each guard is nonempty. The spans*

$$\boxed{\Omega_{n+k,o} \longleftarrow U_{t,\ell}^{(k)} \xrightarrow{\rho_{t,\ell}^{(k)}} \Omega_{m_\ell+k,p_\ell}} \tag{40}$$

compose under sorted tree substitution by pullback; their temporal arrow labels compose by word concatenation. For an arbitrary graded Moore comodel, pulling these guards back along its bounded behaviour map can make some guards empty.

Proof. At an update node a, intersect the shifted guards of the p-children with the disjoint cylinders $\{b : b(a) = p\}$. Because different successor observations are unconstrained coordinates in the final future, every syntactic leaf has at least one compatible observation assignment; the guard is nonempty. Induction gives a finite cylinder partition, and restriction of futures gives the residual maps. For an operation u substituted at leaf ℓ, a new leaf (ℓ, j) is reached exactly when the original guard for ℓ is satisfied and the residual future satisfies the guard for j. Hence

$$U^{(k)}_{t\circ_\ell u,(\ell,j)} = U^{(k)}_{t,\ell} \cap (\rho^{(k)}_{t,\ell})^{-1}(U^{(k)}_{u,j}).$$

The residual shift is $\rho^{(k)}_{u,j}\rho^{(k)}_{t,\ell}$, and the associated word is $w_\ell w_j$. This is exactly pullback composition of the spans. The associated cooperation is the familiar map into the coproduct of successor-state fibres, as in ordinary polynomial-comodel semantics. □

Theorem 9.5 (Graded sorted interaction-law coherence)**.** *Let t be a graded sorted polynomial operation with root (n, o) and leaves (m_ℓ, p_ℓ); label its leaves by elements $b_\ell \in B_{m_\ell,p_\ell}$. For every $s \in S_{n+k,o}$, exactly one leaf ℓ is selected by the comodel's observations along t. Writing $s_\ell \in S_{m_\ell+k,p_\ell}$ for the state reached by its word w_ℓ followed by the allowed restriction, the following equality holds in the temporal Chu tensor:*

$$\boxed{[\![t]\!]_B((b_\ell)_\ell), s] = [b_\ell, s_\ell] \quad \text{in } (B \mathbin{\widehat{\otimes}_T} S)_k.} \tag{41}$$

The selected state is the action of the temporal arrow $\Pi(t)_\ell \in \mathcal{J}_A(n, m_\ell)$; for the terminal comodel its domain of selection is exactly $U^{(k)}_{t,\ell}$. Equality (41) is compatible with operadic substitution, and with the sorted/unsorted tensor isomorphism of Proposition 9.1.

Proof. Induct on the normal operation tree. At a leaf there is only weakening from its original grade to the root grade; (26) is precisely the asserted equality. At an update node a, the sorted algebra interprets an input tuple by $\sigma_{a,n-1,o}$ and the comodel performs $\delta_{a,n-1+k,o}$. This update lands in a uniquely determined observation fibre p, selecting the p-continuation. Relation (28), transported through Proposition 9.1, identifies the root interpretation with this selected continuation and new state. Apply the induction hypothesis to the selected subtree. The resulting word is a followed by its subtree's word, with any unused budget discarded by (26). This proves (41). It is independent of the order of substitutions because graded monad associativity identifies the two algebraic interpretations, while concatenation and naturality of the comodel identify their executed paths. For $S = \Omega$, the same induction is the partition-of-guards construction of Proposition 9.4. Hence the polynomial evaluation, guarded temporal path, and balanced tensor are three presentations of the same execution operation. □

Example 9.6 (Three paths in a branching computation). Take the tree described in Section 4.1: root budget 3, root observation 0, update a, then update b on the 0-branch and return immediately on the 1-branch. In a terminal bounded future $\Omega_{3,0}$ the leaves are selected by

$$\begin{aligned} U_1 &= \{b : b(a) = 0,\ b(ab) = 0\},\\ U_2 &= \{b : b(a) = 0,\ b(ab) = 1\},\\ U_3 &= \{b : b(a) = 1\}. \end{aligned}$$

The temporal arrows are $3 \xrightarrow{ab} 0$ for the first two leaves, and $3 \xrightarrow{a} 1$ for the third. At residual horizon k, the same paths act on $\Omega_{3+k,0}$ and return futures of horizons k, k, and $k+1$, respectively. Theorem 9.5 says that the same branch selection computes the balanced tensor class, regardless of which normal polynomial expression was used to construct the program.

10 Cofree comodels and the two-sided Chu adjunction

The interaction theorem connects the concrete program and execution semantics; it does not by itself make the pairing two-sided. So far the first transpose has transformed a *comodel* into an algebra of tests. The opposite direction asks whether a fixed *algebra* also determines a comodel representing all its evaluations. Cofree bounded systems and the shift–delay adjunction provide the required representation. Establishing it is what promotes the temporal tensor from a useful quotient to a genuine Chu *dual adjunction*, without claiming reflexivity.

The temporal tensor is already a categorical evaluation: its first transpose (33) turns a graded comodel into a graded algebra of temporal tests. The other transpose needs a comodel representing evaluations out of a fixed graded algebra. The existence and description of cofree graded Moore comodels provide it.

Theorem 10.1 (Cofree bounded Moore comodel)**.** *Let $U : \mathcal{C} \to \mathcal{E}$ forget the observations and updates. For any temporal object $K \in \mathcal{E}$, define*

$$\boxed{\mathrm{Cof}(K)_n = \prod_{w \in A^{\le n}} (O \times K_{n-|w|}).} \tag{42}$$

The root O-coordinate is the current observation, restrictions truncate histories and restrict K-coordinates, and updates re-root the array at an input letter. The root K-coordinate defines a counit $U\,\mathrm{Cof}(K) \to K$. There is a natural bijection

$$\boxed{\mathcal{C}(S, \mathrm{Cof}(K)) \cong \mathcal{E}(US, K).} \tag{43}$$

In particular $\mathrm{Cof}(1) = \Omega$ is terminal.

Proof. Given a natural transformation $p : US \to K$, map $s \in S_n$ to the labelled bounded tree

$$w \longmapsto \big(o_{n-|w|}(\delta_{w,n}s),\ p_{n-|w|}(\delta_{w,n}s)\big).$$

The comodel equations give naturality under restriction, and rerooting commutes with transitions. Conversely any comodel map $S \to \mathrm{Cof}(K)$ is determined by its projection to the root K-label; equivariance under update forces its remaining coordinates to be the values of the root label after executing the indicated word. These constructions are inverse and natural. □

10.1 Delay is right adjoint to shift

The shift D_dS has a right adjoint R_d given by

$$(R_dK)_k = \begin{cases} 1, & k < d, \\ K_{k-d}, & k \ge d. \end{cases} \tag{44}$$

In particular $R_d = \triangleright^d$. The isomorphism

$$[S, K]_d = \mathcal{E}(D_dS, K) \cong \mathcal{E}(S, R_dK) \tag{45}$$

expresses exactly how the resource shift transfers across evaluation: the state can be inspected d units further into the future, or the test output can instead be delayed by d units.

Theorem 10.2 (The second Chu transpose)**.** *For every $B \in \mathcal{A}$ and $K \in \mathcal{E}$ there is a graded Moore comodel $\langle B, K\rangle$, natural contravariantly in B, for which*

$$\boxed{\mathcal{A}(B, [S, K]) \ \cong\ \mathcal{E}(B \mathbin{\widehat{\otimes}_T} S, K) \ \cong\ \mathcal{C}(S, \langle B, K\rangle).} \tag{46}$$

For a free graded algebra with generators X at grade d,

$$\boxed{\langle F^{(d)}X, K\rangle \cong \mathrm{Cof}(R_d(K^X)).} \tag{47}$$

Here K^X is the pointwise power by the discrete set X. The two transposes give a contravariant adjunction between $\mathcal{A}$ and $\mathcal{C}$ for each fixed K; they are not generally inverse equivalences.

Proof. By Lemma 8.3, cartesian closedness of $\mathcal{E}$, and $D_d \dashv R_d$,

$$\begin{aligned}\mathcal{E}(F^{(d)}X \mathbin{\widehat{\otimes}_T} S, K) &\cong \mathcal{E}(X \times D_d S, K)\\ &\cong \mathcal{E}(D_d S, K^X)\\ &\cong \mathcal{E}(S, R_d(K^X))\\ &\cong \mathcal{C}(S, \mathrm{Cof}(R_d(K^X))).\end{aligned}$$

Every graded Moore algebra admits a coequalizer presentation by coproducts of grade-shifted free algebras: the operations form a finitary many-sorted algebraic theory after diagonal splitting, including the admissibility condition by restricting to its image. Equivalently, one can use the usual free-algebra presentation of the graded algebra category directly. The functor $B \mapsto B \mathbin{\widehat{\otimes}_T} S$ preserves colimits by Theorem 8.2. Hence a presentation of B expresses $S \mapsto \mathcal{E}(B \mathbin{\widehat{\otimes}_T} S, K)$ as the corresponding limit of the representable functors associated with free generators. Since $\mathcal{C}$ is complete by Theorem 6.4, the limit of these representing comodels exists, and represents the required functor. This constructs $\langle B, K\rangle$ naturally and proves the second bijection; the first was established in Theorem 8.2. □

Remark 10.3 (Scope of the adjunction). A categorical Chu pairing is not an assertion that all algebras are reflexive or that complete state information is uniquely reconstructible from tests. The canonical bidual maps

$$B \longrightarrow [\langle B, K\rangle, K], \qquad S \longrightarrow \langle [S, K], K\rangle$$

can fail to be invertible. Their possible kernels, images, and biextensional quotients are separate reachability–observability problems.

Proposition 10.4 (An elementary nonreflexive example). *Let $A = \{a\}$ and $O = \{0, 1\}$, and let B be the terminal graded Moore algebra, with $B_n = 1$ at every stage. For $K = \Delta 2$, the degree-zero component of the bidual map*

$$B \longrightarrow [\langle B, K\rangle, K]$$

is not surjective.

Proof. Write $S = \langle B, \Delta 2\rangle$. The terminal graded comodel Ω has nonempty stages. The constant-zero map $B \mathbin{\widehat{\otimes}_T} \Omega \to \Delta 2$ corresponds by Theorem 10.2 to a comodel map $\Omega \to S$; in particular S_0 is nonempty. Both constant-zero and constant-one maps $S \to \Delta 2$ are distinct natural transformations, hence distinct elements of $[S, \Delta 2]_0 = \mathcal{E}(S, \Delta 2)$. But B_0 has one element. Therefore its bidualization map cannot be surjective. No degenerate empty alphabet or singleton observation set is required. □

10.2 Boolean tests separate free operational syntax

Proposition 10.5 (Full abstraction and free syntax separation). *For every set X, distinct normal Moore terms have distinct evaluations*

$$T_{\mathrm{Moore}}(X) \to (A^* \times X)^{O^{A^*}}, \qquad t \mapsto (b \mapsto \mathrm{run}_t(b)). \tag{48}$$

Consequently the canonical map $F(X) \to [\langle F(X), \Delta 2\rangle, \Delta 2]$ is componentwise injective.

Proof. Use the coloured-tree normal forms (5). Distinct roots are distinguished by choosing the root observation; distinct return leaves are distinguished by their X labels, while a return and update are distinguished by their executed words. For two different update letters, the first executed letters differ. For identical update letters, choose a successor observation sort where the subtrees differ and recurse. Each finite assignment to a finite set of input histories extends to a complete future in O^{A^*}, so this induction yields an observation oracle separating any two distinct normal forms.

The cofree polar $\langle F(X), \Delta 2\rangle \cong \mathrm{Cof}((\Delta 2)^X)$ provides a Boolean test of return labels at every finite input history, alongside arbitrary observation futures. If two terms execute different words, choose Boolean labels that distinguish those words; if they execute the same word but return different labels, choose a Boolean function separating those labels. By the previous paragraph some oracle witnesses one of these alternatives, so a dual test distinguishes the terms. This proves the componentwise injectivity. (When X is empty, the assertion is vacuous.) □

11 Completion and the limits of finite tests

The most tempting conclusion at this point would be that the bounded theory is simply the ordinary theory viewed stage by stage. The final comodel does indeed complete to Garner's full future, but that does not settle what happens to the *paired evaluation*: forming equivalence classes of executions and taking inverse limits are different operations. This section tests whether they commute. The two counterexamples are included not as edge-case warnings but because they show two logically different failures, one of existence of a completed representative and the other of separation of completed executions.

Finite-horizon data are complete in one sense: every compatible system of observed coordinates defines a stream in O^{A^*}. But *evaluation quotients* need not commute with inverse limits. This distinction is essential if a Chu construction is to support reachability–observability minimization rather than just give an analogy.

11.1 The opposed completion adjunctions

Write $c_{\mathcal{A}} : \mathrm{Alg}(T) \to \mathcal{A}$ and $c_{\mathcal{C}} : \mathrm{Comod}(T) \to \mathcal{C}$ for the constant-grade functors. For finitary Moore theory, with finite O, there are adjunctions

$$\operatorname*{colim}_n : \mathcal{A} \rightleftarrows \mathrm{Alg}(T) : c_{\mathcal{A}}, \qquad c_{\mathcal{C}} : \mathrm{Comod}(T) \rightleftarrows \mathcal{C} : \lim_n . \tag{49}$$

The first follows because $T = \operatorname{colim}_n T_n$ and a compatible interpretation of all finite terms defines an ordinary algebra action. The second follows because compatible complete states can execute every finite program. The constant functors are fully faithful. Against an ordinary constant comodel, evaluation compares exactly:

$$B \mathbin{\widehat{\otimes}_T} c_{\mathcal{C}}(C) \cong \Delta\big((\operatorname{colim} B_n) \otimes_T C\big). \tag{50}$$

For general B, S there is only a comparison map

$$\boxed{\chi_{B,S} : (\operatorname*{colim}_n B_n) \otimes_T (\lim_n S_n) \longrightarrow \lim_k (B \mathbin{\widehat{\otimes}_T} S)_k.} \tag{51}$$

Its value at horizon k sends a class represented by $b \in B_n$ and a compatible state s to $[b, s_{n+k}]$. Lemma 8.3 makes χ an isomorphism when B is free graded (and hence for its retracts).

11.2 Failure of surjectivity: compatible classes without states

Example 11.1 (Missing complete representative). Take $A = \{a\}$, $O = 1$, and $B_n = 1$ constantly. Let

$$S_n = \{\infty\} \cup \{m \in \mathbb{N} : m \geq n\},$$

with restriction by inclusion, unique observation, and update $m \mapsto m-1$, $\infty \mapsto \infty$. These maps form a graded comodel. Its inverse limit is $\{\infty\}$. On the other hand, at every horizon k the single-step balancing relation identifies all finite states in S_k with one another, but never with ∞. Thus $(B \mathbin{\widehat{\otimes}_T} S)_k$ has two classes, and the restriction maps between these quotients are identities. The target of (51) has two elements, while the source has one. The comparison is not surjective.

11.3 Failure of injectivity: completed orbit distinctions

Example 11.2 (The 2-adic obstruction). Again take $A = \{a\}$, $O = 1$, and $B_n = 1$, but now let

$$S_n = \mathbb{Z}/2^{n+1}\mathbb{Z}$$

with restriction by modular reduction and update by addition of 1, with the evident source/target stages. All horizon sets are finite and all restriction maps are surjective. Every finite-stage balanced tensor quotient is a singleton, because the step relation identifies the cyclic $+1$ orbit. Consequently the target of (51) is one point. The inverse limit comodel is $\mathbb{Z}_2$ with update $x \mapsto x+1$, whose ordinary tensor with the terminal algebra is the orbit set $\mathbb{Z}_2/\mathbb{Z}$. This has many elements, since each orbit is countable. Therefore χ is not injective, despite finiteness and surjectivity at every stage.

Theorem 11.3 (Finite horizons rule out missing quotient classes)**.** *If B is a finite ordinary Moore algebra, regarded as constant graded, and all horizon sets S_n are finite, then the comparison map*

$$\chi_{c_{\mathcal{A}}(B),S} : B \otimes_T \varprojlim_n S_n \longrightarrow \varprojlim_k (c_{\mathcal{A}}(B) \mathbin{\widehat{\otimes}_T} S)_k$$

is surjective. Restriction maps need not be surjective.

Proof. At horizon k, constant algebra weakening makes the underlying coend canonically $B \times S_k$, and the balanced tensor is a quotient $q_k : B \times S_k \twoheadrightarrow Q_k$. Given a compatible family of quotient classes $(q_k)_k$, each fibre $F_k = q_k^{-1}(q_k)$ is nonempty and finite. The restriction maps take F_{k+1} into F_k. Build the finitely branching tree of finite compatible sequences of elements of these fibres. It has a node at every depth, so König's lemma supplies an infinite branch. This gives an element of $\varprojlim(B \times S_k) = B \times \varprojlim S_k$ whose evaluation class maps to the prescribed compatible family. Hence χ is surjective. □

Remark 11.4 (Two genuinely different failure mechanisms). The missing-lift example shows that a compatible family of *quotient classes* may have no compatible underlying state. The 2-adic example shows that two complete-state orbit classes can become indistinguishable under every finite quotient. The latter is precisely why finite-horizon reachability–observability should be formulated as a pairing, with its separation properties proved, rather than assuming completed evaluation is conservative.

11.4 Bounded Hankel balancing

For a pointed machine (S, o, δ, s_0) and $N \in \mathbb{N}$, define

$$H_N(u,v) = o(\delta_{uv}(s_0)), \qquad |u| + |v| \le N. \tag{52}$$

Whenever defined, $H_N(uw, v) = H_N(u, wv)$. The domains are triangular, not rectangular: executing a longer history consumes the time available for continuation tests. Taking compatible histories over all N recovers the ordinary Hankel matrix and its rows $u^{-1}L_{s_0} \in O^{A^*}$. The two Chu transposes encode the same balancing principle at the level of graded algebras and comodels, but do *not* assert that the history and test categories are involutively equivalent.

12 Resource shifts and quantitative examples

The completion results mark the end of the central graded Moore argument. What follows is deliberately a controlled change of perspective, not a new foundational layer: the same budget shifts that organize temporal evaluation also occur in cost-enriched and amortised semantics. We can therefore test which quantitative principles are already consequences of the tensor, and which require additional notions such as ownership and locality. This boundary matters if the construction is eventually to support compositional verification rather than only the description of behaviour.

The temporal direction of our pairing also reappears in earlier cost-enriched work and in amortised resource analysis. This section draws a boundary: the shift adjunctions and static frame strength are exact consequences of the preceding structure, but resource ownership and dynamic allocation are additional semantic requirements.

12.1 The shift adjoint triple

View $P = (\mathbb{N}, \geq, +)$ as the category indexing $\mathcal{E} = [P, \mathbf{Set}]$. Precomposition by translation $q \mapsto p + q$ defines

$$D_p X(q) = X(p + q).$$

Its left and right adjoints are

$$\boxed{L_p \dashv D_p \dashv R_p}, \qquad (L_p X)(q) = X(\max(0, q - p)), \qquad (R_p X)(q) = \begin{cases} 1, & q < p, \\ X(q - p), & q \geq p. \end{cases} \tag{53}$$

The left adjoint is Day convolution by the covariant representable $P(p, -)$, and R_p is delay by p stages. Equation (45) is the instance relevant to the tensor–Hom adjunction: an action on a future-shifted state corresponds to an action whose value is delayed.

For $p = 1$, $R_1 = \triangleright$ admits a natural map next : $\mathrm{Id} \to \triangleright$. In the topos of trees, equations of the form $f : \triangleright X \to X$ have a uniquely determined global fixed point, constructed inductively from the terminal stage at zero. But *shift adjoints alone do not imply guarded recursion*: for the unit resource $p = 0$, all three shifts are identities, and an arbitrary endomap need not have a unique fixed point. Also, reading has grade zero in Moore syntax, so a recursive definition guarded merely by get does not thereby advance time.

12.2 Weighted update costs

Assign each letter $a \in A$ a strictly positive integer cost $c(a)$. In the sorted tree monad replace an a-node at grade n by a node whose children have grade $n - c(a)$, whenever $n \geq c(a)$. The strictly graded update-tree functor becomes

$$(G^c X)_{n,o} = X_{n,o} + \coprod_{\substack{a \in A \\ c(a) \leq n}} \prod_{p \in O} (G^c X)_{n - c(a), p}.$$

Weakening still propagates to the leaves, yielding a distributive law $WG^c \Rightarrow G^c W$. Leafwise temporal arrows are now words with weighted length $|w|_c = \sum_i c(a_i)$ satisfying $|w|_c + m \leq n$. This keeps the polynomial calculation intact. It does *not* yet establish any separation-logic frame rule: it measures a global upper bound on execution cost.

12.3 A static tensor-frame theorem

Let F be a plain set of frame states, unaffected by execution and invisible to observation. The graded comodel $\mathrm{Fr}_F(S)$ has stages $S_n \times F$, restrictions $r \times \mathrm{Id}_F$, and updates $\delta_a \times \mathrm{Id}_F$. Since the

tensor's relations do not change the F-coordinate,

$$\boxed{B \mathbin{\widehat{\otimes}_T} \mathrm{Fr}_F(S) \cong (B \mathbin{\widehat{\otimes}_T} S) \times \Delta F.} \tag{54}$$

The proof passes product with F through the coends and quotient relations; it is valid because $- \times F$ preserves colimits in **Set**. The handler version is

$$[\mathrm{Fr}_F(S), K]_n \cong [S, K^{\Delta F}]_n.$$

This is a real strength of the temporal tensor but not full separation logic. In Atkey's work, resources live in an ordered commutative monoid and may be distributed over disjoint heap regions; a local command must not access the frame's heap or consume its credit [8]. That locality condition is more restrictive than keeping an arbitrary extra coordinate unchanged.

Example 12.1 (Owned credit is not a global time budget). For a unary Boolean machine whose update flips the bit and consumes one owned credit, put

$$S_n = \{(x, r) \mid x \in \{0, 1\},\ r \geq n\}, \quad o_n(x, r) = x, \quad \delta_{a,n}(x, r) = (1 - x, r - 1),$$

with restrictions given by inclusion. This is a graded comodel, and all S_n are inhabited. Nevertheless $\varprojlim_n S_n = \varnothing$: no state carries enough *finite* credit for every possible horizon. Appending a disjoint frame with separate credit preserves (54), whereas forgetting which credit is owned by which component does not preserve the right to spend it. This shows why resource budgeting and ownership must be kept distinct.

12.4 Fresh contexts and the limits of generalization

For a small symmetric monoidal category $\mathcal{R}$, shifts $D_c X(d) = X(c \otimes d)$ on $[\mathcal{R}, \mathbf{Set}]$ are precomposition functors and admit Kan-extension adjoints. With a coherent $\mathcal{R}$-graded monad and resource-sensitive cooperations, one can form the *candidate* coend

$$(B \widehat{\otimes}_T^{\mathcal{R}} S)(k) = \left(\int^{d \in \mathcal{R}} B(d) \times S(d \otimes k) \right) / \sim,$$

with covariant comodel and contravariant algebra carriers relative to $\mathcal{R}$. Under suitable grading and execution axioms the first tensor–Hom transpose is formal. Existence of cofree comodels and the second transpose requires additional hypotheses, and a guarded fixed-point theorem requires a genuinely decreasing, well-founded or locally contractive resource notion.

For finite name contexts, precomposition by adjoining one fresh position resembles the numerical shift. Its right adjoint, in a suitable finite-injection presheaf variance, collects data over all one-name deletions, $\prod_{s \in S} X(S \setminus \{s\})$. That phenomenon is interesting for syntax and deletion, but name abstraction and nominal freshness require support and permutation structure not captured by a numerical cost grade. We leave it outside the central result.

13 Two worked finite examples

After the abstract equivalences it is worth following actual states and programs through the constructions. The first example makes bounded observability, Hankel rows, and the execution tensor visible in one unary machine. The second returns to a branching tree and separates the roles of numerical budget, labelled temporal paths, and the observations that select one path. Together they show what each intermediate category preserves, and why none can simply be omitted.

These examples show how the same data look from the operational, temporal, and observational sides, rather than treating the categorical results as purely formal.

13.1 A unary machine whose states become observable at different horizons

Take $A = \{a\}$ and $O = \{0, 1\}$. Let the ordinary Moore machine have states $S = \{x, y, z\}$ and

$$o(x) = o(y) = 0, \quad o(z) = 1, \qquad \delta_a(x) = z, \quad \delta_a(y) = y, \quad \delta_a(z) = z.$$

Its constant-horizon graded comodel has stage $S_n = S$ for every n, with identity restriction maps. The unique map to the final graded comodel is the bounded behaviour map

$$\mathrm{beh}_n(s)(a^j) = o(\delta_a^j(s)), \qquad 0 \le j \le n.$$

At horizons zero and one these observations are

state	beh_0	beh_1	beh_∞
x	0	$(0, 1)$	$(0, 1, 1, 1, \ldots)$
y	0	$(0, 0)$	$(0, 0, 0, 0, \ldots)$
z	1	$(1, 1)$	$(1, 1, 1, 1, \ldots)$.

Thus x and y are indistinguishable at horizon zero but are separated by a one-step experiment. This makes the refinement of observable equivalence with additional budget concrete:

$$s \sim_n t \quad \Longleftrightarrow \quad \mathrm{beh}_n(s) = \mathrm{beh}_n(t), \qquad \sim_{n+1} \subseteq \sim_n .$$

With the initial point x, only x and z are reachable. The pointed infinite Hankel pairing has entries

$$H_x(a^i, a^j) = \begin{cases} 0, & i + j = 0, \\ 1, & i + j \ge 1, \end{cases}$$

and its two distinct rows give the minimal reachable observable realization. The third state y is observable as different but unreachable from the selected point. This is the exact asymmetry between the choice of an initial clock map and Garner's global assignment of a future to *every* state.

Now consider the grade-one program

$$t = \mathsf{step}_a\big(\mathsf{get}(0 \mapsto u, 1 \mapsto v)\big) \in T_1\{u, v\}.$$

Executing t from x chooses the return v and residual state z; from y it chooses u and residual state y. In the free graded algebra's temporal tensor, the defining relation therefore reads

$$[t, x] = [v, z], \qquad [t, y] = [u, y],$$

with the left-hand terms at grade one and the right-hand terms at grade zero, at matching remaining horizon. This is both the sorted-corolla rule and the leafwise temporal-path rule, now checked on a specific machine.

13.2 What the three-path tree forgets under ungraded completion

Take $A = \{a, b\}$ and $O = \{0, 1\}$ with the root-budget-three tree of Example 9.6. It tests the observation after a; on the 0-branch it then executes b and tests again, while on the 1-branch it returns. Its leafwise temporal paths are ab, ab, a. A plain duration upper bound records only "at most two updates" and cannot distinguish either the choices or the words. The path-family operad records all three words but still does not specify which one a given state executes. The finite cylinder guards partition the bounded-future object and supply exactly that missing choice. The balanced temporal tensor finally identifies the formal algebraic result of the entire tree with the selected return and its correctly truncated residual state.

The three steps of the comparison are therefore not redundant: *grade* bounds duration, *polynomial syntax* retains branching, and *comodel evaluation* selects one branch. Completion from Ω_n to Ω_∞ supplies unlimited consistent tests; it does not remove the need to keep the operational distinction among these three levels.

14 Conclusions and a controlled next programme

The route through the paper was determined by a single objective: to retain finite operational structure and finite observational structure *in one pairing*, rather than completing either side at the outset. Sorting and polynomial trees supplied the operational shapes; the chosen grading and the comodel equations supplied the bounded temporal category; the Chu tensor expressed their evaluation; and the comparison with complete futures exposed what completion preserves and what it can identify or fail to realize. This also pinpoints the limited but precise relationship with DOTS: the chosen grading supplies time, while input-labelled charts pick out linear histories inside a branching space of finite futures.

The Moore example supports a precise resting place for the syntax-first programme:

1. *A chosen ordered grading supplies bounded time.* The filtration by update depth is an ordered graded monad; the input signature refines its arrows by words, and the observation signature adds branching. Diagonal splitting converts the reader layer into observation sorts; the graded theory becomes an admissible part of a genuinely polynomial tree monad GW.

2. *Comodels carry temporal geometry.* Their inverse-stage data are precisely the coalgebras for $\Delta O \times (\triangleright -)^A$, and they form the presheaf slice $[\mathcal{J}_A, \mathbf{Set}]/\Omega$. For unary updates $\mathcal{J}_A^{\mathrm{op}}$ is the interval category; for branching input histories it embeds into complete finite history trees under stated finiteness hypotheses.

3. *The pairing retains temporal information.* The balanced tensor has a temporal tensor–Hom adjunction and, via cofree comodels, a second Chu transpose. This is a genuine dual adjunction, not an automatic reflexive equivalence.

4. *The graded interaction law is coherent across presentations.* A polynomial operation determines a family of labelled budget arrows; finite observations select one guard, and executing the selected path gives exactly the balanced tensor relation. Compatibility holds under substitution.

5. *Completion is mathematically nontrivial.* Garner's O^{A^*} is the inverse limit of the final bounded comodel, but quotienting by execution does not in general commute with that limit. The two counterexamples describe distinct obstructions.

This is already a worthwhile integrated result without claiming a universal theorem for arbitrary operational theories. The next stage should focus on three targeted problems: deriving general conditions for *injectivity* of the completion comparison χ (our 2-adic example shows why finiteness is insufficient); comparing fixed-input charted clocks with the full DOTS compositional framework; and constructing tensor products or multi-gradings of theories compatible with system wiring. As a first test for the last question, two independently clocked theories could carry an $\mathbb{N}^2$-grading, with a diagonal clock imposing synchronization. This is a proposed extension, not a theorem. General resource categories and dynamic ownership are further applications. A complete Raynor-style graph pattern for concurrency or a general residual Sweedler semantics would be a separate contribution, not a prerequisite for the present paper.

The step from Paper I to Paper II. The no-residue temporal Chu pairing developed here deliberately forgets implementation effects that have already been executed. Paper II [17] asks what changes when implementations emit resource costs and when those costs must be certified against a potential and assigned to particular owners. It begins from the same graded Moore resource coend, distinguishes it from residual cost writers, and studies certificate composition and restricted forms of framing and wiring. In this sense Paper II builds on the present semantics rather than supplying an alternative to it; the preservation of quantitative information requires additional structure that the current no-residue theory was not designed to retain.

Relation to the clocks chapter. The clocks chapter provides the systems-theoretic motivation: trajectories enter systems from clocks, and specifications or potentials test their behaviour. The present document explains how a chosen ordered resource grading, together with its labelled update operations, determines a temporal indexing category and compatible evaluation structure. It does not replace the clock-based account; it supplies its graded algebraic counterpart in a fully computed case.

Relation to physical temporal verification. The *Temporal Interlude* [16] addresses a different part of the same programme: it gives concrete realizations of runner operations as time-occupying trajectories, uses interval-sheaf pullbacks for synchronization, and proves specific deadline properties after attaching monitors or moving to the strict-interior temporal topos. Paper II [17] supplies the complementary theory of resource-certified implementations. A general functorial passage from the bounded comodel and Chu semantics of this paper, through certified residual runners, to those physical trajectory sheaves is an open comparison problem.

A Elementary nerve reconstruction

Here is the elementary form of the nerve argument underlying our active–inert description. The colours are $C = \mathbb{N} \times O$. Let E_c be a single edge of colour c, and let C_t be a corolla decorated by a graded polynomial operation $t : (c_1, \dots, c_r) \to c$. There are inclusions of each input edge and the output edge into C_t. For a $\mathcal{P}$-algebra B, $N_B(C_t)$ is the graph of its operation $t^B : \prod_i B_{c_i} \to B_c$.

Proposition A.1 (Decorated-tree Segal description)**.** *The nerve $B \mapsto N_B$ on the decorated tree category is fully faithful. Its essential image consists of presheaves whose tree values are limits of their vertex-corolla and edge values, and for which the joint input map $F(C_t) \to \prod_i F(E_{c_i})$ is bijective for every operation corolla. The original graded Moore algebras correspond to the additional stagewise admissibility condition on the edge values.*

Proof. For N_B, a compatible labelling of the vertices and their shared edges determines all edge labels, proving the Segal limit condition. Since each algebraic operation is a total function, its graph maps bijectively to the tuple of inputs. Conversely set $B_c = F(E_c)$; the corolla bijections and the output restriction define every operation t^B. Naturality for degeneracies supplies units, while naturality for active contractions gives compatibility with operadic substitution. The Segal limits reconstruct all tree values from corollas and edges. Maps of such presheaves are determined by their edge components and exactly preserve operations by naturality on corollas. The admissibility condition is the one characterized in Theorem 4.3. □

Remark A.2 (Raynor's setting is different). Raynor's circuit-algebra work handles graphs that may have loops and nontrivial unit/deletion phenomena and proves a Segal nerve theorem by a refined monadic approach. The present proof is a tree-algebra calculation for a cartesian polynomial monad after many-sorted splitting. Compare the general theory of monads with arities and their nerve theorems [10]. It does not establish her factorization claims for our reader-normalization quotients or extend automatically to wheeled props [9].

B Raw get–step trees and reader reductions

The many-sorted polynomial presentation removes explicit read nodes from normal forms, but raw one-sorted syntax can still be useful to locate local substitutions and observation tests. The two reader equations can be oriented as vacuous-read contraction and selective nested-read contraction:

$$\mathsf{get}(o \mapsto t) \longrightarrow t, \qquad \mathsf{get}(\dots, i \mapsto \mathsf{get}(p \mapsto t_p), \dots) \longrightarrow \mathsf{get}(\dots, i \mapsto t_i, \dots).$$

Selective contraction follows from the diagonal law and does not add an axiom. On finite trees, these reductions decrease the number of read nodes and yield the sorted normal forms, so they are terminating. The sorted semantic description establishes uniqueness of the normal form and therefore confluence of the terminating rewriting system.

Normalization does not preserve every raw subtree occurrence. For example, with $O = \{0, 1\}$, the raw term

$$\mathsf{get}\big(0 \mapsto \mathsf{get}(0 \mapsto x, 1 \mapsto y), \quad 1 \mapsto \mathsf{step}_a(z)\big)$$

reduces to $\mathsf{get}(0 \mapsto x, 1 \mapsto \mathsf{step}_a(z))$; the raw occurrence of y disappears because a second read without intervening step cannot change the current observation. If $b \in O^{A^*}$ is a complete observation oracle, the guard of an occurrence p is

$$U_{t,p} = \{b : \mathrm{run}_t(b) \text{ visits } p\}.$$

For the disappearing occurrence the guard is empty. Vacuous-read contraction instead coarsens the partition of guards. Executed update events and the final return can be indexed by event position, so equivalent terms have the same event space, even when their read-node occurrence sets differ. This explains why the older clopen occurrence-residual calculus remains useful *for comparison of presentations*, not as a replacement for coloured polynomial arities.

C Audit ledger and scope of the claims

The following ledger makes explicit which earlier dependencies are used in the integrated argument. It should be treated as a list of proof locations, not as a separate collection of conjectures.

Construction	Proof or boundary in this document
Reader splitting and empty sorts	Thm. 2.1; admissibility throughout.
Polynomial W, G and Beck law	Thm. 4.2; tree formula (13).
Comparison with graded Moore	Thm. 4.3; free formula (14).
Guarded and slice comodels	Thms. 6.2, 6.4.
Free tensor cancellation	Lemma 8.3, with explicit inverse.
Both Chu transposes	Thms. 8.2, 10.2.
Compatibility of all three semantics	Thm. 9.5; guards Prop. 9.4.
Limits of completion	Exs. 11.1, 11.2; Thm. 11.3.
Sheaf gluing and general residual effects	Not claimed.
Dynamic ownership and general resource Chu	Not claimed; only exact static frame result (54).

References

[1] J. Płonka, *Remarks on diagonal and generalized diagonal algebras*, Colloquium Mathematicum **15** (1966), 19–23. https://doi.org/10.4064/cm-15-1-19-23.

[2] A. Zamojska-Dzienio, *Medial modes and rectangular algebras*, Comment. Math. Univ. Carolin. **47**(1) (2006), 21–34.

[3] N. Gambino and J. Kock, *Polynomial functors and polynomial monads*, Math. Proc. Cambridge Philos. Soc. **154** (2013), 153–192. https://arxiv.org/abs/0906.4931.

[4] J. Kock, *Polynomial functors and trees*, Int. Math. Res. Not. (2011), 609–673. https://arxiv.org/abs/0807.2874.

[5] R. Garner, *The costructure–cosemantics adjunction for comodels for computational effects*, Math. Struct. Comput. Sci. **32** (2022), 374–419. https://arxiv.org/abs/2011.14520.

[6] S. Fujii, S. Katsumata, and P.-A. Melliès, *Towards a Formal Theory of Graded Monads*, FoSSaCS 2016, LNCS 9634, 513–530.

[7] L. Birkedal, R. E. Møgelberg, J. Schwinghammer, and K. Støvring, *First Steps in Synthetic Guarded Domain Theory: Step-Indexing in the Topos of Trees*, Logical Methods in Computer Science **8**(4), 2012.

[8] R. Atkey, *Amortised Resource Analysis with Separation Logic*, Logical Methods in Computer Science **7**(2), 2011. https://bentnib.org/amortised-sep-logic.html.

[9] S. Raynor, *Modular operads, iterated distributive laws and a nerve theorem for circuit algebras*, arXiv:2412.20262, version 4. https://arxiv.org/abs/2412.20262.

[10] C. Berger, P.-A. Melliès, and M. Weber, *Monads with arities and their associated theories*, Journal of Pure and Applied Algebra **216** (2012), 2029–2048. https://arxiv.org/abs/1101.3064.

[11] S.-y. Katsumata, E. Rivas, and T. Uustalu, *Interaction laws of monads and comonads*, LICS 2020. https://arxiv.org/abs/1912.13477.

[12] N. Niu and D. I. Spivak, *Polynomial Functors: A Mathematical Theory of Interaction*, Cambridge University Press, 2025; see especially Chapters 7–8. https://arxiv.org/abs/2312.00990.

[13] I. Moerdijk and I. Weiss, *Dendroidal Sets*, Algebraic & Geometric Topology **7** (2007), 1441–1470. https://arxiv.org/abs/math/0701293.

[14] P. Schultz, D. Spivak, and C. Vasilakopoulou, *Dynamical Systems and Sheaves*, Applied Categorical Structures **28** (2020), 1–57. https://arxiv.org/abs/1609.08086.

[15] O. Lynch, D. J. Myers, E. F. Rischel, and S. Staton, *Clock systems for stochastic and non-deterministic categorical systems theories*, arXiv:2603.29573, 2026.

[16] *Getting to Grips with Sweedler Runners for Systems: Temporal Interlude— Execution, Synchronization, and Quantitative Verification. From free observation theories to temporal sheaves and marked-event logic*, companion working manuscript (October 2026).

[17] *Sweedler Runners for Temporal Systems, II: Resource-Certified Implementations and Ownership*, companion research draft (October 2026).