Sweedler Runners for Temporal Systems, II

Resource-Certified Implementations and Ownership

Cost, Temporal Evaluation, and Potentials

Companion research draft — 16 October 2026

Abstract

A stateful implementation of an operational specification must do more than preserve outputs: in resource-sensitive applications it should also preserve a quantitative guarantee, and this guarantee should survive composition with other implementations and systems. Building on the bounded Moore comodels and temporal Chu duality of the companion Paper I, this paper consolidates four investigations at the intersection of Sweedler runners, cost-indexed toposes, amortised resource analysis, and ownership semantics. We distinguish the opposite variances of graded algebras and bounded comodels, expressing finite-horizon evaluation as a resource coend followed by operational balancing. For deterministic runners into cost-writer residues, potential inequalities give pointwise colax coalgebra certificates which compose under residual substitution, with an exact law for unused allowance. The ordinary balanced Chu quotient need not remember actual cost. We distinguish a two-coordinate writer-twisted execution quotient, which retains actual cost and declared allowance, from the directional potential certificate, whose initial-state data need not descend. Restriction-stable graded potentials are shown to factor through the stage-zero carrier. Separating resource worlds then supply witnessed allocations, a bicategory of certified transition spans, and a canonical opfibration of allocations over resource extensions. A locality criterion gives canonical lifts of owned-state operations in the unadorned PCM case; proof-relevant ownership may still require extra lifts. Coherent operation lifts yield a frame theorem; finite-horizon resource coends satisfy an ownership base-change theorem in categories with the requisite pullback-stable colimits, including Grothendieck toposes. These results provide a reusable semantics for a restricted class of certified, controlled-wiring implementations. We explain their relation to automatic amortised resource analysis, Atkey's separation logic, Grodin–Harper's coalgebraic amortisation, and the cost-indexed models of λ-amor. We do not assert a general equivalence with internal Sweedler theory, an unrestricted DOTS wiring theorem, or higher measuring objects.

1 Why implementation needs a quantitative semantics

1.1 From an operational specification to a resource guarantee

An operational theory tells us what a system is allowed to do. A runner adds an implementation: it interprets operations using state and, potentially, a different residual effect theory. When the residual theory is the identity, the implementation looks like a comodel; When the residue is nontrivial, an abstract operation may itself perform a computation. This operational role of stateful runners originates in [3, 4]. This separation between specification and realisation is a major reason to use Sweedler theory in categorical systems theory.

For dependable systems, behavioural correctness is only part of the story. Two implementations may produce identical observations while using very different amounts of time, energy, memory, or communication. In a braking controller or a bounded-resource protocol, that difference can determine whether an implementation is acceptable. We would like a semantics that expresses

not only *what* is implemented, but *at what certified cost*, and then supports composition without redoing a global proof.

Amortised analysis offers a familiar mathematical pattern. Suppose a primitive abstract operation a changes a concrete implementation state s to $\delta_a s$, incurs actual cost $c_a(s)$, and has declared allowance b_a. A potential Φ certifies the implementation when

$$c_a(s) + \Phi(\delta_a s) \leq b_a + \Phi(s). \tag{1}$$

For a finite execution the intermediate potentials cancel, so the total cost is bounded by the sum of the declared allowances plus the initial stored potential. This is the principle behind the potential method, type-based automatic amortised resource analysis (AARA), and quantitative Hoare reasoning; see [1, 2, 10, 13, 11]. It is *not* new in this paper. The question we ask is different: can the certificate be fitted to the *runner and algebra–comodel pairing*, so that it survives implementation refinement, temporal observation, and suitably controlled system wiring?

The problem has three obstructions that are easy to miss. First, the *budget remaining to execute* and the *actual cost already incurred* are different measurements. Second, the ordinary algebra–comodel tensor deliberately equates a computation with the result of executing it; that equality can erase its cost. Third, composing two machines may inadvertently count the same stock of credit twice unless their ownership relationship is specified. None is solved simply by adding a cost number to a Moore-machine state.

1.2 The semantic foundation supplied by Paper I

This paper is the second of a pair. *Sweedler Runners for Temporal Systems, I: From Graded Operational Syntax to Temporal Behaviour* [15] starts from a depth-filtered Moore theory with *no residual effects*. It uses observation-sorted polynomial syntax to exhibit finite programs, identifies graded comodels with bounded temporal behaviours over the input-labelled budget category, and constructs a temporal-valued algebra–comodel tensor with both Chu transposes. Its interaction-coherence theorem explains how an observation selects the executed leaf of a finite program; its completion counterexamples show that finite evaluation cannot be replaced uncritically by evaluation on completed Garner futures.

These results provide the *semantic baseline* here. We keep the resource-indexed horizon and the distinction between formal computation and its executed outcome, but ask a question that Paper I intentionally leaves open: what if the execution of one abstract operation emits an actual, state-dependent cost, and this cost must remain certified through successive implementations and ownership-sensitive wiring? The answer requires a nontrivial cost-writer residue and quantitative potential inequalities. Neither construction is supplied by the no-residue Chu tensor alone. Nor do we assume the still-unproved identification of the external resource coend with an arbitrary internal Sweedler tensor.

1.3 Physical time and the Temporal Interlude

The application-facing *Temporal Interlude* [14] asks what a runner does *during* its implementation. It unfolds free-observation runners into sheaves of input, output, and internal-state histories, retaining intermediate phases when one source operation takes several target reads. It then uses amortised duration certificates to turn logical bounded response into a bound in physical time, composes monitored trajectories by pullbacks, and gives continuous-time and temporal-topos examples. This complements Paper I's abstract finite-horizon semantics and the present paper's accounting of resource-bearing residual effects.

The quantitative intersection is exact in the numerical, additive fragment: the Interlude's two-stage duration bound rescales an outer potential by an inner rate, yielding $q\Phi + \Psi$. The general writer-translation formula proved below is $f(\Phi) + \Psi$, which specializes to that form when $f(c) = qc$ (in a suitable numerical cost domain). This algebraic coincidence does *not* establish a general equivalence between our temporal Chu quotient and an interval trajectory sheaf; nor do the cost certificates alone supply non-Zeno behaviour or a temporal-logic satisfaction theorem.

1.4 The four steps of the argument

The route through the paper is dictated by those obstructions.

1. We first separate two categorical ways of representing cost: *availability at a grade* and *cost emitted by an effect*. The former is naturally indexed by an ordered resource category; the latter is carried by a writer residue. This gives the correct basis for comparing the Part VI cost semantics with our graded temporal Chu construction.
2. We then construct costed runners and prove that their potential certificates compose. This is the implementation theorem: it explains how an abstract allowance can remain meaningful even when one source operation is implemented by a state-dependent sequence of lower-level effects.
3. We place these certificates over finite-horizon evaluation. Here a sharp obstruction appears: actual costs do not generally descend to the ordinary Chu quotient. Resource-indexed execution must therefore be distinguished from the quotient which forgets execution history.
4. Finally, we add ownership. A partial resource combination and its witnessed allocations let us state when a frame remains untouched, when certificates compose across interfaces, and which part of the construction is stable under change of ambient topos.

The progression is from *resource geometry*, through *costed implementation*, to *compositional ownership*. Throughout we keep the general categorical machinery separate from policies specifying which operations, permissions, and transfers exist in a particular system.

1.5 Scope and status

The worked operational theory is deterministic Moore theory: a family of input-labelled transitions together with a pure observation. This is enough to test nontrivial writer residues and observation-controlled branching without importing the machinery for arbitrary effects. The resource calculations take place mainly in $(\mathbb{N}, +, 0, \leq)$, with statements for more general monoidal categories when the necessary hypotheses are explicit. The results about costed span composition and resource-coend base change do *not* require Moore syntax.

It is useful to compare this account with two lines of previous work. The λ-amor type system of Rajani–Gaboardi–Garg–Hoffmann [2] and its ordered-cost copresheaf interpretation in Binder–Corfield–Orchard–Rajani [9] supplies resource-indexed shifts, Day convolution and cost/potential adjunctions. Grodin–Harper [11] place amortisation in coalgebras of cost algebras, with colax maps and indexed composition; their theory is considerably broader than our deterministic writer fragment. Atkey's affine resource-BI semantics inside separation logic, with a Coq development and extracted verification-condition generator [10], explains what additional ownership reasoning an operational account must respect. Our aim is a *comparison and synthesis in the runner setting*, not a replacement for those established theories or a claim to automatic potential inference.

2 Two roles for cost: availability and residual output

2.1 Ordered budgets and their opposite variances

Let $\mathsf{P} = (\mathbb{N}, \geq)$ be the category with a unique arrow $n \to m$ when $n \geq m$, monoidal under addition. Define

$$\mathcal{E}_- = [\mathsf{P}, \mathbf{Set}], \qquad \mathcal{E}_+ = [\mathsf{P}^{\mathrm{op}}, \mathbf{Set}]. \tag{2}$$

An object $S \in \mathcal{E}_-$ is a descending diagram $\cdots \to S_2 \to S_1 \to S_0$. Its element $s \in S_n$ may be understood as a state carrying a certificate of at least n units of remaining execution capability. The maps forget the stronger certificate. An object $B \in \mathcal{E}_+$ is an increasing diagram $B_0 \to B_1 \to \cdots$; for a filtered operational algebra, allowing more budget admits more computations.

The model of cost in [9] takes copresheaves on the opposite cost order $(\mathbb{N}, \leq)$, so it lives in $\mathcal{E}_+$. Our graded Moore *comodel* carriers live in $\mathcal{E}_-$. One can describe both as $\mathcal{C}$-sets after choosing a cost category, but this is a change of variance, not a levelwise identification. Indeed an ascending diagram with $B_0 = \varnothing$, $B_1 = \{*\}$ has no corresponding descending diagram with the same level sets and maps.

This distinction is not a defect. The two variances are exactly what is needed for a coend pairing: a covariant action of increasing computation budgets is balanced against a contravariant action of remaining state capability. Keeping them separate will prevent us from claiming an internal Sweedler comparison before constructing one.

2.2 Shifts and Day convolution are reusable

More generally, let $(\mathcal{C}, \otimes, I)$ be a small symmetric monoidal category, and put $\mathcal{E}_{\mathcal{C}} = [\mathcal{C}, \mathbf{Set}]$ with Day convolution $\star$. Write $y_p = \mathcal{C}(p, -)$ for the covariant representable, and define the resource shift

$$D_p X(q) = X(p \otimes q).$$

The following familiar adjunction explains why the same shifts occur in diverse cost semantics.

Proposition 2.1 (Resource-shift adjunction)**.** *There is a natural adjunction and a Day-closure identification*

$$- \star y_p \;\dashv\; D_p, \qquad [y_p, X]_{\mathrm{Day}} \cong D_p X.$$

For $\mathcal{E}_-$ there is also a right adjoint R_p to D_p. In the natural-number instance, R_1 is the later or guarded-delay functor

$$(R_1 S)_0 = 1, \qquad (R_1 S)_{n+1} = S_n, \qquad R_p \cong R_1^p.$$

Proof. The Day formula and co-Yoneda give $(X \star y_p)(r) \cong \int^q X(q) \times \mathcal{C}(q \otimes p, r)$. The internal-hom formula therefore yields $[y_p, Z]_{\mathrm{Day}}(r) \cong \mathrm{Nat}(y_p, Z(r \otimes -)) \cong Z(r \otimes p)$. Day closure gives the stated adjunction. On the descending natural-number diagram category, the right Kan extension along translation has the displayed formula for $p = 1$; iteration gives the general case. The latter is the later modality of the topos-of-trees model of guarded recursion [8]. □

One reward is conceptual economy. D_p, its adjoints, and the associated weakening maps are not invented afresh for each operational signature. They follow from the resource category. What remains theory-dependent is how an operation interacts with these shifts and, in particular, which equations its algebra and comodel impose during execution.

2.3 Writer effects record a different kind of information

For a commutative monoid C, the cost-writer monad is

$$\mathsf{W}_C(X) = C \times X.$$

A computation emits an element of C in addition to its return value. Its algebras are C-actions: $\mathrm{Alg}(C \times -) \simeq [BC, \mathbf{Set}]$, where BC is the one-object category of the *discrete* monoid C. By contrast, the resource-indexed category $[\mathbf{C}_{\geq}, \mathbf{Set}]$ sees the inequalities between costs. These should not be conflated.

The two constructions answer different questions. In a resource-indexed presheaf, n tells us what a state or proof is available *at grade* n; in a writer computation, c is an *emitted cost*. No map from one to the other is automatic. A potential will relate them through an inequality, but a simple graded capability filtration is not itself an amortised potential: a graded step is only required to take S_{n+1} into S_n, and may preserve all physical credit rather than spend it.

This difference is the reason for introducing nontrivial Sweedler residues instead of trying to attach costs directly to ordinary behavioural quotients.

3 The resource coend and the operational quotient

3.1 Evaluation before executing an operation

Take an increasing diagram $B \in \mathcal{E}_+$, a descending diagram $S \in \mathcal{E}_-$, and an output horizon $k \geq 0$. The thin-cost pairing is

$$Q_k(B, S) = \int^{n \in \mathsf{P}} B_n \times S_{n+k}. \tag{3}$$

Concretely it is the quotient of $\coprod_n B_n \times S_{n+k}$ generated by

$$(w_{m,n} b, s) \sim (b, r_{n+k,m+k} s) \qquad (m \leq n). \tag{4}$$

This records precisely the fact that a computation can be described using a larger allowance while a bounded state can forget the corresponding surplus capability. The shift by k preserves the desired remaining horizon after evaluation. Nothing in (3) has yet imposed the law for executing an actual input-labelled operation.

Lemma 3.1 (Representable cancellation)**.** *For $d, k \geq 0$, with the contravariant representable $\mathsf{P}(-, d)$,*

$$Q_k(\mathsf{P}(-, d), S) \cong S_{d+k}.$$

Proof. This is co-Yoneda for the descending module $D_k S$. □

3.2 The unary baseline: ordinary profunctor composition

For a set of input labels A, let $\mathcal{J}_A$ have objects natural numbers and arrows

$$\mathcal{J}_A(n, m) = \{w \in A^* : |w| \leq n - m\} \quad (n \geq m).$$

Its arrows remember an input word as well as a permitted loss of budget; composition concatenates words. The thin cost category includes into $\mathcal{J}_A$ by sending every weakening to the empty word ε.

When the observation set is a singleton, graded Moore algebras are functors $B : \mathcal{J}_A^{\mathrm{op}} \to \mathbf{Set}$, and graded comodels are functors $S : \mathcal{J}_A \to \mathbf{Set}$. The resource pairing then becomes the familiar profunctor tensor once we impose one additional family of equations.

Theorem 3.2 (Unary execution comparison)**.** *For $|O| = 1$, the temporal algebra–comodel tensor is*

$$(B \mathbin{\widehat{\otimes}_T} S)_k \cong \int^{n \in \mathcal{J}_A} B_n \times S_{n+k}.$$

It is the quotient of $Q_k(B, S)$ by the step relations

$$[\sigma_{a,n}(b), s]_{n+1} \sim [b, \delta_{a,n+k}(s)]_n, \qquad s \in S_{n+k+1}. \tag{5}$$

The quotient is natural in the horizon and in algebra and comodel maps.

Proof. Balancing over $\mathcal{J}_A$ is balancing over all its arrows. The empty-word arrows give exactly (4); the single-letter arrows give (5). They generate $\mathcal{J}_A$, subject to the update–weakening compatibility, so the two displayed families generate all coend relations. Restricting the horizon commutes with these relations. □

This case is deliberately included as a baseline rather than claimed as novelty: it is ordinary profunctor composition. It tells us where the new questions begin. Multiple observations require branch selection, and residual writers require distinguishing costs that this quotient may forget.

3.3 Many observations and the precise comparison still missing

For finite nonempty O, diagonal splitting presents the graded Moore algebra by stagewise admissible sorts $B_{n,o}$, while a bounded Moore comodel has fibres $S_{n,o}$. Here the reader operation $\mathsf{get} : B_n^O \to B_n$ obeys the two reader equations selecting a constant branch and eliminating nested reads. For nonempty B_n, these diagonal-algebra equations split $B_n \cong \prod_{o \in O} B_{n,o}$, with $\mathsf{get}((b^{(p)})_p)_o = b_o^{(o)}$; the empty algebra corresponds to the all-empty family. Thus a family of sorts is *admissible* if all its carriers are empty or all are nonempty at each grade. Weakening acts separately on sorts and updates supply an O-tuple of continuations. We recall these facts here because the proposition below depends on the Moore-specific presentation, whose detailed tree-normal-form proof is given in Paper I [15], not repeated in this quantitative study. Put $\mathsf{P}_O = \mathsf{P} \times O_{\mathrm{disc}}$. The sorted resource coend is

$$Q_k^O(B, S) = \int^{(n,o) \in \mathsf{P}_O} B_{n,o} \times S_{n+k,o}.$$

An update corolla supplies one continuation of each observation sort, $\sigma_{a,n,o} : \prod_{p \in O} B_{n,p} \to B_{n+1,o}$, while a cooperation picks one successor, $\delta_{a,n+k,o} : S_{n+k+1,o} \to \coprod_{p \in O} S_{n+k,p}$.

Proposition 3.3 (Sorted execution comparison)**.** *For a stagewise admissible sorted Moore algebra, $(B \mathbin{\widehat{\otimes}_T} S)_k$ is the quotient of $Q_k^O(B, S)$ by*

$$[\sigma_{a,n,o}((b_p)_p), s]_{n+1,o} \sim [b_q, s']_{n,q} \quad \text{when } \delta_{a,n+k,o}(s) = (q, s'). \tag{6}$$

This identification is compatible with finite-horizon restriction.

Proof. Reader balancing first evaluates a pure read at the current observation sort: the component of a diagonal algebra relevant at o is $B_{n,o}$. The reader equations guarantee consistency of repeated reads. The residual cost weakenings yield Q_k^O, while executing an update gives exactly (6). These generate the same quotient as the usual graded algebra–comodel execution relations. □

It would be tempting to call this the Part VI internal Sweedler tensor, simply because both involve a cost-indexed category. That conclusion has *not* been proved. The construction above pairs two opposite variances externally and then imposes operational equations. An internal-Sweedler comparison must first construct an internal theory and identify its evaluation coequalizer with (6). This remains a precise open problem, not a change of notation.

The next section explains why that missing comparison is not an obstacle to a first quantitative implementation theorem: we can retain the cost itself in a writer residue and work with the explicit execution presentation already available.

4 Stateful writer residues and quantitative implementation

4.1 Why a residual theory matters

Suppose the source specification contains `increment`. It may be implemented by flipping bits in a binary counter. One increment at a number with many trailing one-bits performs many physical updates. Forgetting those internal operations yields the correct abstract transition, but it also forgets how much work has been done. A stateful runner into a cost-writer residue can retain the quantitative result while still treating the abstract and concrete operation theories separately.

Let T be the Moore theory with update labels A and pure observations in O, and let $\mathsf{W}_C(X) = C \times X$ for a commutative cost monoid C. A deterministic stateful implementation of T through W_C, carried by a set S, is described by observations $o : S \to O$, transitions $\delta_a : S \to S$, and emitted costs $c_a : S \to C$. Pure reads select their prescribed branches at zero cost. The universal property of the Moore operations then gives a monad morphism

$$T \longrightarrow \mathrm{St}_S\, \mathsf{W}_C, \qquad \mathrm{St}_S\, \mathsf{W}_C(X) = (C \times (X \times S))^S. \tag{7}$$

The reader laws hold because inspecting an unchanged state twice agrees with inspecting it once; compare the algebraic-effects runner framework of [3, 4]. This is the first instance of a *nontrivial residue*: output behaviour can stay the same while the writer component carries cost information.

An implementation may itself be implemented again. We therefore also need the elementary structure of runners *between writer theories*.

Proposition 4.1 (Writer runners are actions with cocycles)**.** *Let C, D be additive monoids and U a state set. A stateful runner*

$$\mathsf{W}_C \longrightarrow \mathrm{St}_U\, \mathsf{W}_D$$

is equivalent to maps $\nu : C \times U \to U$ and $h : C \times U \to D$ satisfying

$$\nu(0, u) = u, \qquad h(0, u) = 0, \tag{8}$$

$$\nu(c + d, u) = \nu(d, \nu(c, u)), \qquad h(c + d, u) = h(c, u) + h(d, \nu(c, u)). \tag{9}$$

The first map is a (right-convention) cost action; the second is an additive cost cocycle for that action.

Proof. Naturality in the return variable forces a writer computation (c, x), started from u, to return the same x, changing only the state and writer cost. The monad unit gives (8); the multiplication law for two consecutive writer costs gives (9). Conversely these equations make the indicated interpretation a monad morphism. □

The crucial feature is the state dependence of h: processing a writer effect labelled c may be cheap in one downstream state and expensive in another. This supplies exactly the flexibility needed for amortised implementation chains.

4.2 Potentials as colax cost certificates

For the rest of the principal numerical arguments take $C = \mathbb{N}$ with its usual order. Choose declared allowances $b_a \in \mathbb{N}$ and a potential $\Phi : S \to \mathbb{N}$. We call the source runner *certified* when it satisfies (1) for every a, s. The potential should be thought of as a quantitative assertion on the concrete state, rather than as an additional operation of the source syntax.

To compare this with the coalgebraic treatment of amortisation, introduce the free cost actions $\mathbb{N} \times S$ and $\mathbb{N}$ and their lifted transitions

$$\widehat{\delta}_a(d, s) = (d + c_a(s), \delta_a s), \qquad \widehat{\sigma}_a(d) = d + b_a, \qquad \widehat{\Phi}(d, s) = d + \Phi(s).$$

Equip these free actions with the pointwise cost order (holding the state coordinate fixed). The following calculation is the exact restricted bridge to the colax viewpoint of Grodin–Harper [11].

Theorem 4.2 (Potential–colax correspondence)**.** *The following statements are equivalent for a deterministic, writer-costed Moore runner:*

1. *the local inequalities $c_a(s) + \Phi(\delta_a s) \leq b_a + \Phi(s)$ hold;*
2. *the cost-action morphism $\widehat{\Phi}$ satisfies the pointwise colax inequalities*
$$\widehat{\Phi}\,\widehat{\delta}_a \leq \widehat{\sigma}_a\,\widehat{\Phi} \qquad (a \in A); \tag{10}$$
3. *each finite input word $w = a_1 \cdots a_m$, executed from s at total cost $c(w, s)$ with final state s_w, satisfies*
$$c(w, s) + \Phi(s_w) \leq \sum_{i=1}^{m} b_{a_i} + \Phi(s). \tag{11}$$

A Moore computation with pure reads obeys the same path inequality along the observation-selected branch.

Proof. Evaluating (10) at (d, s) gives the two numbers $d + c_a(s) + \Phi(\delta_a s)$ and $d + b_a + \Phi(s)$. Their comparison is exactly (1), and the converse follows by putting $d = 0$. Adding the local inequalities along a finite word cancels all intermediate potentials, proving (11); words of length one recover (1). A read changes neither state nor writer cost, so it only determines which inequalities along a term's execution path are used. □

The familiar telescope is thus also a cell expressing that the concrete costed coalgebra is simulated by an abstract allowance coalgebra. This is not a claim that *all* the colax coalgebra morphisms of [11] are Sweedler runners. Their endoprofunctor signatures, additional effects and indexed implementation composition require a substantially broader theory.

5 Composition of certified residual implementations

5.1 The question posed by multi-stage implementations

A certificate would be of limited use if changing the implementation of an underlying primitive invalidated the bound for the abstract operation. Suppose an initial runner realises a source

operation using writer cost in C, and a second runner interprets each such cost in terms of a more concrete cost monoid D. If a source update emits $c_a(s)$, the second implementation may charge $h(c_a(s), u)$, changing its own state from u to $\nu(c_a(s), u)$.

A numerical grade translation $f : C \to D$ need not be a pointwise bound on h. It is enough to certify it amortisedly. Assume C, D are ordered commutative monoids with monotone addition, $f : C \to D$ is a monotone monoid homomorphism, and there is a second potential $\Psi : U \to D$ satisfying

$$h(c, u) + \Psi(\nu(c, u)) \leq f(c) + \Psi(u) \qquad (c \in C, u \in U). \tag{12}$$

Theorem 5.1 (Certified residual composition)**.** *Let the Moore-to-writer runner* (7) *have a C-valued potential Φ and allowances b_a, satisfying* (1) *in C. Let the writer-to-writer runner (h, ν) carry (f, Ψ) satisfying* (12)*. The composite runner on $S \times U$ has primitive data*

$$\begin{aligned} \delta_a^{\mathrm{comp}}(s, u) &= \big(\delta_a s, \nu(c_a(s), u)\big), \\ c_a^{\mathrm{comp}}(s, u) &= h(c_a(s), u), & b_a^{\mathrm{comp}} &= f(b_a), \\ \Theta(s, u) &= f(\Phi(s)) + \Psi(u). \end{aligned}$$

It satisfies the composite certificate

$$c_a^{\mathrm{comp}}(s, u) + \Theta(\delta_a^{\mathrm{comp}}(s, u)) \leq f(b_a) + \Theta(s, u). \tag{13}$$

These certificates compose associatively under further certified writer translations, up to the canonical reassociation of state products.

Proof. Put $s' = \delta_a s$, $c = c_a(s)$, and $u' = \nu(c, u)$. The second certificate gives $h(c, u) + \Psi(u') \leq f(c) + \Psi(u)$. Monotonicity and additivity of f applied to the source certificate give $f(c) + f(\Phi(s')) \leq f(b_a) + f(\Phi(s))$. Adding yields (13). The underlying stateful interpretation is a runner by monad-morphism substitution, or directly by the cocycle equations (9). For a third translation with monoid map g and potential Ξ, either association gives $(g \circ f)(\Phi(s)) + g(\Psi(u)) + \Xi(v)$. The cost cocycles and state actions associate by the same substitution equations. □

The importance of this result is not the telescope alone, which belongs to ordinary amortised analysis, but its compatibility with *replacement of a residual implementation*. An abstract allowance continues to have a certified interpretation even after a lower-level operation acquires a stateful implementation of its own. Grodin–Harper's indexed coalgebraic composition [11] organizes related implementation changes; the present calculation is a restricted cost-writer comparison, not a proof that runner substitution is exactly reindexing in their indexed category.

Proposition 5.2 (Exact decomposition of slack)**.** *For $C = D = \mathbb{N}$ and additive f, let*

$$\begin{aligned} \Delta_a(s) &= b_a + \Phi(s) - c_a(s) - \Phi(\delta_a s), \\ \Delta_\eta(c, u) &= f(c) + \Psi(u) - h(c, u) - \Psi(\nu(c, u)). \end{aligned}$$

All slacks are nonnegative; every additive endomorphism $f : \mathbb{N} \to \mathbb{N}$ is multiplication by $f(1)$, and the composite transition has

$$\boxed{\Delta_a^{\mathrm{comp}}(s, u) = f(\Delta_a(s)) + \Delta_\eta(c_a(s), u).} \tag{14}$$

Proof. Expand the right-hand side and cancel $f(c_a(s))$. Additivity of f leaves the expression for the slack of (13) exactly. Subtraction is used only for natural-number costs. □

This gives a useful invariant for proofs and implementations: unused allowance is not an amorphous quantity which we merely know to be nonnegative. It decomposes into translated slack from the first stage and slack contributed by the second. That decomposition can be useful when tracking where a conservative bound originates.

5.2 A binary counter as a genuinely stateful residue

Example 5.3 (Incrementing a binary counter)**.** Let the abstract operation `increment` act on $n \in \mathbb{N}$ by $n \mapsto n+1$. If $t(n)$ is the number of trailing one-bits in the binary expansion of n, the number of bit flips is $c(n) = 1 + t(n)$. Put $\Phi(n) = \mathrm{popcount}(n)$. Then

$$\Phi(n+1) = \Phi(n) - t(n) + 1, \qquad c(n) + \Phi(n+1) = 2 + \Phi(n).$$

Thus an abstract sequence of m increments costs at most $2m + \Phi(n)$ bit flips, although no uniform constant worst-case bound exists for one increment over every initial state.

The same example can be moved to the *second* residue stage. Let an incoming writer cost $c \in \mathbb{N}$ instruct a downstream counter u to perform c increments. Set

$$\nu(c,u) = u + c, \qquad h(c,u) = \sum_{i=0}^{c-1} (1 + t(u+i)),$$
$$\Psi(u) = \mathrm{popcount}(u), \qquad f(c) = 2c.$$

The sum is a cost cocycle, and telescoping the bit-count identity gives

$$h(c,u) + \Psi(u+c) = 2c + \Psi(u).$$

So the downstream implementation is certified *with equality*. Composing it with any upstream certified writer runner automatically replaces its potential Φ by $2\Phi + \mathrm{popcount}$ and its allowance b_a by $2b_a$. This is a two-stage implementation theorem, not just the naming of a potential on a single counter.

The reader may now reasonably ask whether the certificate also survives passage to our temporal Chu evaluation. Answering that question is subtler because this evaluation contains a quotient; the next section isolates precisely which quotient is safe.

Comparison with timed runner composition. The duration-certificate theorem in the *Temporal Interlude* [14] has the same rescaling pattern for successive implementations: an outer rate r and an inner rate q give rate rq, and the potential becomes $q\Phi + \Psi$. Theorem 5.1 expresses the corresponding algebra for cost-writer residues and a monoid translation f. The Interlude additionally interprets completed source operations as events on a physical clock; its bounded-response and non-Zeno hypotheses are separate from the writer laws proved here.

6 Temporal evaluation and the cost-descent obstruction

6.1 The finite-horizon execution presentation

Let $T_{\le n}(X)$ denote bounded Moore terms with at most n updates along each execution branch, and let S_n be a bounded Moore comodel. Define

$$\mathrm{Ex}_k(X,S) := \int^{n \in \mathsf{P}} T_{\le n}(X) \times S_{n+k}. \tag{15}$$

This is the resource coend (3) for the free graded algebra carrier, *before* the additional relations identifying an update with its executed continuation. Its elements are equivalence classes of a computation, an initial bounded state, and a choice of permitted allowance.

Suppose actual costs $c_{a,n} : S_{n+1} \to \mathbb{N}$, potentials $\Phi_n : S_n \to \mathbb{N}$, and observations and updates commute with restriction, with

$$c_{a,n}(r_{n+2,n+1}s) = c_{a,n+1}(s), \qquad \Phi_n(r_{n+1,n}s) = \Phi_{n+1}(s). \tag{16}$$

Assume the local inequalities are graded:

$$c_{a,n}(s) + \Phi_n(\delta_{a,n}s) \leq b_a + \Phi_{n+1}(s), \qquad s \in S_{n+1}. \tag{17}$$

Cost at level n is *actual writer cost*; the index n measures the number of abstract updates still allowed. These need not coincide.

Remark 6.1 (Restriction stability makes potentials stage-zero data). The equalities (16) force

$$\Phi_n = \Phi_0 \circ r_{n,0}, \qquad c_{a,n} = c_{a,0} \circ r_{n+1,1}.$$

Thus a restriction-stable graded potential is a cocone $S \Rightarrow \Delta\mathbb{N}$, equivalently a function on S_0, since 0 is terminal in P and $\operatorname{colim}_{\mathsf{P}} S \cong S_0$. Theorem 6.2 is consequently a *descent theorem* for ungraded state potentials; its grading governs the evaluation, not the numerical potential itself. Horizon-dependent potentials can be obtained by allowing, for example, $\Phi_n(rs) \leq \Phi_{n+1}(s)$, but then numerical labels no longer descend as equalities to the ordinary Set-valued coend. For instance, on the singleton comodel $S_n = 1$, the choice $\Phi_n = n$, unit step cost and zero declared allowance satisfies the graded step inequality but not (16). An ordered/enriched quotient for such lax data is left open.

For an execution representative (t, s), let $w(t, s)$ be the selected sequence of update labels, let s' be the final bounded state, and let x be the returned variable. The final state belongs to S_k, after the executed updates and restrictions. Write $c(t, s)$ for the sum of actual costs and $b(t, s)$ for the sum of declared allowances along this path.

Theorem 6.2 (Temporal certificate over the resource presentation)**.** *Under* (16) *and* (17)*, the total actual cost, total allowance and initial/final potential define functions on each* $\mathrm{Ex}_k(X, S)$*, natural in the horizon. They satisfy*

$$\boxed{c(q) + \Phi_{\mathrm{out}}(q) \leq b(q) + \Phi_{\mathrm{in}}(q), \qquad q \in \mathrm{Ex}_k(X, S).} \tag{18}$$

These inequalities are compatible with substitution of Moore terms: writer costs and declared allowances concatenate and the intermediate potential telescopes.

Proof. Weakening a term and restricting its state change neither the selected execution path nor the values of the numerical labels, by the comodel restriction equations and (16). Hence all labels descend through the resource coend. The same equations imply horizon naturality. Applying (17) at each visited step and adding cancels intermediate potentials; pure reads only select a branch. Under term substitution the selected execution is the first path followed by a continuation path, with remaining allowance restricted as necessary. Writer multiplication adds their costs, while the common intermediate potential cancels. □

This is a pointwise ordered comparison on *cost-labelled execution presentations*. We have not constructed an abstract enriched measuring-object morphism for arbitrary effects. The theorem does provide the correct restricted semantics of the amortisation certificate at every temporal horizon.

6.2 Composition survives the presentation

If a second certified writer runner (h, ν, f, Ψ) has state set U independent of the horizon, then the composed graded state carrier is $S_n \times U$, with update

$$(s, u) \longmapsto (\delta_{a,n} s, \nu(c_{a,n}(s), u)).$$

The composite potential is $f(\Phi_n(s)) + \Psi(u)$. Restriction stability follows from (16). Moreover, since product with a fixed set preserves the coequalizer defining the coend,

$$\mathrm{Ex}_k(X, S \times U) \cong \mathrm{Ex}_k(X, S) \times U. \tag{19}$$

The cost cocycle makes second-stage execution of the total first-stage cost agree with successively interpreting its cost increments. Thus Theorem 5.1 acts directly on the temporal certificate (18), natural in k. A further writer implementation gives the same certificate after reassociating the state product.

6.3 Why the ordinary tensor may forget cost

The map from $\mathrm{Ex}_k(X, S)$ to the ordinary balanced Chu evaluation additionally identifies computations under operational step execution. Those equations preserve state and return value, but need not preserve how many writer units were emitted.

Proposition 6.3 (Non-descent of actual cost)**.** *There are certified deterministic costed runners whose total-cost function on* $\mathrm{Ex}_k(X, S)$ *does not factor through the balanced tensor* $T(X) \mathbin{\widehat{\otimes}_T} S$.

Proof. Use one observation, one input letter a, a singleton state at every horizon with identity restriction and update, and one return variable x. Give every a-step cost and allowance 1, with potential 0. The executions of x, $a(x)$, and $a(a(x))$ have costs $0, 1, 2$, so all satisfy (18) with equality. Operational step balancing, however, identifies them all with the same returned value and state in the tensor. A numerical function on that singleton cannot take three different values. □

This small counterexample is a structural guide for the whole paper. It tells us not to demand that actual cost become an invariant of ordinary behavioural equivalence. It does *not* prevent us from forming a better, writer-sensitive quotient. Later, ownership witnesses will likewise be retained before quotienting away their source information.

Proposition 6.4 (Writer-twisted execution and separate allowance)**.** *Assume the restriction-stable, deterministic free Moore setting of Theorem 6.2. Start with the free* $\mathbb{N}^2$*-set*

$$(\mathbb{N} \times \mathbb{N}) \times \mathrm{Ex}_k(X, S),$$

whose coordinates record already accumulated actual cost and already accumulated declared allowance. Impose the reader-execution relation at unchanged writer coordinates, and the twisted step *relation*

$$(d, e, [a(t), s]) \sim (d + c_a(s), e + b_a, [t, \delta_a s]), \tag{20}$$

with the grade indices implicit and all relations closed under addition in $\mathbb{N}^2$*. Let the resulting writer set be* $\mathrm{Ex}_k^{\mathrm{wr}}(X, S)$*. Then the total cost and allowance functions*

$$[d, e, t, s] \longmapsto d + c(t, s), \qquad [d, e, t, s] \longmapsto e + b(t, s)$$

are well defined. Executing the term induces a natural map

$$\mathrm{run}_k^{\mathrm{wr}} : \mathrm{Ex}_k^{\mathrm{wr}}(X, S) \longrightarrow \mathbb{N}^2 \times X \times S_k,$$

which is a bijection for the free Moore term algebra. Taking coinvariants for the $\mathbb{N}^2$*-action recovers the ordinary balanced execution tensor.* The initial potential is not asserted to descend through the twisted step relation.

Proof. At a step, the unexecuted continuation has cost $c(t, \delta_a s)$, so its total writer cost after shifting by $c_a(s)$ equals the old total $c(a(t), s)$; allowances satisfy the analogous equation when they are shifted by b_a, not by actual cost. Pure reader steps change neither quantity. Restriction stability verifies compatibility with the coend and the remaining horizon. Reducing an executable finite term to its uniquely selected leaf gives a representative $(d+c(t,s), e+b(t,s), x, s')$ with $s' \in S_k$. Conversely, leaf representatives are not identified with one another unless these writer coordinates, return value and final state agree, proving bijectivity. Coinvariants erase the two writer coordinates and impose precisely the untwisted read and update relations. □

Remark 6.5 (Why one writer coordinate is insufficient). Twisting (20) by actual cost alone does make total cost well defined. It does *not* make total allowance well defined unless $c_a(s) = b_a$: in general $d + b(a(t), s) \neq d + c_a(s) + b(t, \delta_a s)$. The two-coordinate writer repairs this. The *certificate* remains directional, since a step changes the state at which the potential is evaluated. An ordered coinserter or Pos-enriched coend which records that direction is a plausible next construction, but it is not supplied by the Set-quotient above.

7 Resource ownership is not just resource addition

7.1 Why separating resources changes the question

The last two sections concern one implementation state and its residual writer effect. Actual systems, however, have components that share an environment: two processes may use disjoint heap regions, a controller may access only the observations of another component, and credits may be transferred between owners. A total resource number is not enough to certify such situations. If two components each claim the same credit, adding their potential functions double-counts that credit.

Atkey's amortised separation logic [10] gives an existing principled response: describe the part of the world owned by each operation, combine assertions using a separating conjunction, and prove that operations do not disturb an unrelated frame. Our purpose here is not to reconstruct the logic or its inference engine. It is to identify the categorical *witnesses of ownership* needed to make our runner certificates compositional and to determine how they interact with finite-horizon evaluation.

Let $(W, \star, e)$ be a *partial commutative monoid* (PCM). We use the strong form of associativity appropriate to separation algebras: when a finite product is defined, its reparenthesizations and permutations are defined and equal. A typical world is $w = (h, r)$, with a finite heap h and natural-number credit r, and

$$(h, r) \star (h', r') = (h \uplus h', r + r') \tag{21}$$

when the heaps have disjoint domains. A decomposition $w = u \star v \star f$ records resources owned by two participating components and a further inert frame. In general there are many such decompositions of the same total world.

7.2 Day convolution and the loss of allocation witnesses

The resource-indexed topos of [9] uses Day convolution to combine grades. For $\mathcal{C} = (\mathbb{N}, \leq, +)$,

$$(P \star_{\mathrm{Day}} Q)(p) = \int^{u,v} \mathcal{C}(u+v, p) \times P(u) \times Q(v). \tag{22}$$

A raw representative includes a budget split $u+v \leq p$, and objects available at its two allocated grades. But the coend identifies representatives along weakening maps. It therefore represents an *extensional* combination, not necessarily a chosen allocation.

Proposition 7.1 (Day convolution need not retain ownership)**.** *Take the covariant representables $P = Q = y_1 = \mathcal{C}(1, -)$ on $(\mathbb{N}, \leq, +)$. Day–Yoneda gives $y_1 \star_{\mathrm{Day}} y_1 \cong y_2$. At total budget $p = 3$, the raw splits $(u, v) = (1, 1), (2, 1), (1, 2)$ all define one element: both owners' minimum requirements survive, but the coend forgets who receives the surplus credit.*

Proof. By the Day–Yoneda lemma, $y_1 \star y_1 \cong y_{1+1} = y_2$, whose value at 3 is a singleton. Each of the three displayed splits has a representative, and coend relations identify them. □

This proposition is deliberately elementary. Its significance is methodological: resource allocation is not recovered by running an inverse functor on the Day quotient. For proof-relevant certification we should retain the tuple (u, v, f) as additional dependent data. Later, when choosing to forget that data, we can recover an ordinary extensional semantics.

8 Certified transitions as a reusable compositional calculus

8.1 From graph runners to witnessed transitions

A stateful runner provides a deterministic transition for each primitive operation. But the elementary algebra of potential inequalities does not depend on determinism. A transition can carry an explicit witness z explaining *which* choice or branch was made. The natural categorical home for such witnesses is a span. This offers a reusable calculus before asking whether an arbitrary effect theory has a suitable Sweedler measuring object.

Definition 8.1 (Certified transition span)**.** An object is a set X equipped with a potential $\Phi : X \to \mathbb{N}$. A certified transition from (X, Φ) to (Y, Ψ) consists of

$$X \xleftarrow{s} E \xrightarrow{t} Y, \qquad c, b : E \to \mathbb{N},$$

with the pointwise inequality

$$c(z) + \Psi(tz) \leq b(z) + \Phi(sz), \qquad z \in E. \tag{23}$$

A 2-cell is a map of span apices preserving source, target, actual cost and declared allowance.

The apex E is itself proof-relevant. An input state may have several certified possible transitions, and we retain which one occurred. As every witness must satisfy the inequality, the interpretation is demonic or worst-case over the available branches, not an expected-cost semantics. This makes the construction a convenient ambient calculus for nondeterministic *execution witnesses*, even though we have not supplied a full nondeterministic residual Sweedler theory.

Theorem 8.2 (Bicategory of certified spans)**.** *Certified transition spans form a bicategory* **CertSpan**. *Horizontal composition is pullback of spans and addition of actual costs and allowances; the identity span has zero cost and zero allowance. If $z \in E$ and $z' \in F$ compose, their nonnegative slack satisfies*

$$\boxed{\Delta_{F\circ E}(z, z') = \Delta_E(z) + \Delta_F(z').} \tag{24}$$

Proof. If $tz = sz'$, add the inequalities for z and z': the intermediate potential occurs on opposite sides and cancels. This gives certification of the pullback composite with summed numerical labels. Zero-cost identity spans satisfy equality. The usual coherent pullback isomorphisms give bicategorical associativity and units; apex maps preserve the new labels under horizontal composition. Writing $\Delta_E(z) = b(z) + \Phi(sz) - c(z) - \Psi(tz)$ and similarly for F, cancellation proves (24). □

For an input-labelled deterministic writer runner, the graph of $\delta_a : X \to X$ gives

$$(X, \Phi) \xleftarrow{\mathrm{id}} X \xrightarrow{\delta_a} (X, \Phi), \qquad c = c_a, \quad b = b_a.$$

Its compositional certificate in **CertSpan** is exactly Theorem 4.2's path inequality. A Moore read restricts attention to the guard on which a particular output was observed; input substitution then composes the chosen guarded graph spans. This identifies a natural *span-level image* of our deterministic writer-runners without claiming that every certified span comes from a runner.

8.2 What the span calculus does and does not automate

The virtue of this construction is modest but important. We no longer need a new amortisation theorem for every operational signature: whenever its primitive transitions can be interpreted as certified spans, their finite sequential composition inherits a certificate. On the other hand, a span does not specify which observations are legitimate, which effects respect an algebraic theory's equations, or which resources a command is permitted to use. Those are additional questions, and they lead directly to ownership-dependent framing.

9 The canonical allocation opfibration

9.1 World extensions and allocation data

A frame rule starts by separating the part of a world an operation owns from a part that remains untouched. If such decompositions are arbitrary tuples, repeated re-framing becomes an ad hoc exercise. We therefore organize decompositions and their extensions into a category before referring to any operational theory.

Definition 9.1 (Resource-extension category)**.** Let $\mathrm{Ext}(W)$ have as objects the elements of a PCM W. A morphism $w \to w'$ is a *specified extension witness* g such that $w \star g = w'$. Composition is given by $g \star h$, and identities are labelled by e. The category remembers which extension was used, even when two extensions happen to have the same total result. For cancellative PCMs, including natural-number credit and disjoint heaps, an extension is determined by its endpoints, so $\mathrm{Ext}(W)$ is a preorder. The witness distinction only matters for noncancellative resource monoids: in $(\{0, 1\}, \max, 0)$, both extensions labelled 0 and 1 give different arrows $1 \to 1$.

For each total world w, define

$$\mathrm{Alloc}_2(w) = \{(u, v, f) \mid u \star v \star f = w\}.$$

If $g : w \to w \star g$ is an extension, leave both owned portions unchanged and add g to the frame:

$$(u, v, f) \longmapsto (u, v, f \star g). \tag{25}$$

Strong associativity makes this action well-defined and functorial.

Theorem 9.2 (Canonical allocation opfibration)**.** *The assignment $w \mapsto \mathrm{Alloc}_2(w)$, with (25), defines a functor*

$$\mathrm{Alloc}_2 : \mathrm{Ext}(W) \to \mathbf{Set}.$$

Its category of elements is a discrete opfibration

$$\boxed{\int_{\mathrm{Ext}(W)} \mathrm{Alloc}_2 \longrightarrow \mathrm{Ext}(W).} \tag{26}$$

An object is a witnessed decomposition into two owned portions and a frame. Its unique lift over a chosen resource extension adds that extension to the frame, leaving the owned portions fixed.

Proof. Associativity of $\star$ verifies identities and composition for $\mathrm{Ext}(W)$; it also gives the functor laws of (25). In the category of elements of a covariant set-valued functor, every arrow in the base has a unique lift with specified source, and that lift is opcartesian. □

This construction is canonical once the number and roles of owners and the distinguished frame are fixed. It is independent of the chosen program, cost model and Sweedler theory. What is *not* canonical is which reallocation of the same total world should be regarded as legal.

9.2 Local actions turn the allocation construction into operational structure

The opfibration just constructed describes *inert extension*, whereas the next section needs an *owned-state operation*. To connect them, impose the standard strong locality condition on a partial deterministic action $a : W \rightharpoonup W$: whenever $a(u)$ and $u \star g$ are defined, so are $a(u \star g)$ and $a(u) \star g$, and

$$a(u \star g) = a(u) \star g. \tag{27}$$

This is the separation-algebra notion of local action in a particularly concrete form; see [5]. Its definedness hypothesis is not automatic for partial programs.

Proposition 9.3 (Unique PCM lift of a local owned action)**.** *For each allocation $w = u \star v \star f$ with $a(u)$ defined, strong locality yields a* uniquely determined *lifted allocation*

$$(u, v, f) \longmapsto (a(u), v, f) \qquad \textit{over} \qquad a(w) = a(u) \star v \star f.$$

These lifts commute with inert frame extension. Consequently, composable local owned actions and frame extensions generate commuting squares on the allocation category of their admissible domains. No extra choice of ownership witness is needed for the plain PCM tuples.

Proof. Apply (27) with $g = v \star f$ to obtain the displayed decomposition of $a(w)$. The target tuple is forced by the stipulated changed owner and unchanged other owner and frame. Extending the frame by h gives on either path $(a(u), v, f \star h)$, using associativity and locality. Sequential local owned actions therefore compose compatibly with the framing functors on the subcategories where they are defined. □

Thus the allocation opfibration becomes operationally relevant on admissible local-action domains. It does *not* become an opfibration over a category of all world transitions: an action may fail to be enabled on a particular owned footprint even if it is enabled on the total world. Extending the construction to arbitrary effectful action spans or proof-relevant permissions requires specified lifting data, as in Section 10.

9.3 Permitted transfers replace sets by categories

Suppose for each w we have a category $\mathcal{A}_w$ whose objects are allocations in $\mathrm{Alloc}_2(w)$ and whose arrows represent permitted transfers of resources between owners. Assume each extension $g : w \to w \star g$ induces a framing functor

$$g_! : \mathcal{A}_w \to \mathcal{A}_{w\star g}, \qquad (u, v, f) \mapsto (u, v, f \star g),$$

coherently with composition of extensions.

Theorem 9.4 (Transfer-aware opfibration)**.** *Such functors give a (strict or pseudo-)functor* $\mathcal{A} : \mathrm{Ext}(W) \to \mathbf{Cat}$*, and its Grothendieck construction*

$$\int_{\mathrm{Ext}(W)} \mathcal{A} \longrightarrow \mathrm{Ext}(W)$$

is an opfibration. Its vertical arrows are permitted ownership transfers and its chosen opcartesian arrows are inert frame extensions. Conversely, a split opfibration with the specified allocation fibres recovers the framing functors.

Proof. The Grothendieck construction has a morphism over $g : w \to w'$ from $\alpha \in \mathcal{A}_w$ to $\beta \in \mathcal{A}_{w'}$ whenever there is a vertical arrow $g_!\alpha \to \beta$. The arrow represented by $\mathrm{id}_{g_!\alpha}$ is opcartesian. A cleavage that is split recovers strict transport, and arbitrary coherent cleavages recover the pseudofunctorial version. □

Example 9.5 (Credit transfers)**.** For $W = \mathbb{N}$, an allocation has (r_1, r_2, f) with $r_1 + r_2 + f = w$. A permitted arrow may transfer $d \le r_1$ credits from owner 1 to owner 2:

$$(r_1, r_2, f) \longrightarrow (r_1 - d, r_2 + d, f).$$

The arrow need not be invertible if the permission policy is directional. Extending the frame by g changes only f, so it commutes with this transfer. But the PCM itself does not choose such permissions, nor does it certify the cost of moving a physical heap cell.

10 Framing certified implementations

10.1 Coherent lifts are indispensable

An opfibration gives a systematic way to add an inert frame, but it does not say whether a command can be executed without interfering with the frame. Let $q : \Sigma \to X$ be the owned-state projection of a proof-relevant set of allocations, with a frame projection $\mathrm{fr} : \Sigma \to F$. Take

a certified transition span

$$X \xleftarrow{s} E \xrightarrow{t} X.$$

Pull back its input along q: $\widetilde{E} = E \times_{X,s,q} \Sigma$. To turn this into a transition on Σ, we also need a chosen target lift

$$\ell_E : \widetilde{E} \to \Sigma, \quad q\ell_E = t\,\mathrm{pr}_E, \quad \mathrm{fr}\ell_E = \mathrm{fr}\,\mathrm{pr}_\Sigma\,. \tag{28}$$

The first equation respects the underlying state transition; the second expresses frame preservation. Lifts must respect identities, compositions and the operational equations of the implementation.

Theorem 10.1 (Coherent witnessed frame theorem)**.** *Assume the lifts* (28) *exist and are coherent under composition. If E has a certificate for potential Φ, its lifted transition has a certificate for*

$$\widetilde{\Phi}(\zeta) = \Phi(q\zeta) + \Psi(\mathrm{fr}\zeta)$$

for any frame potential $\Psi : F \to \mathbb{N}$, with unchanged actual cost and allowance. The slack is unchanged. Furthermore, the lift preserves composition up to the canonical pullback isomorphisms:

$$\widetilde{(E' \circ E)} \cong \widetilde{E}' \circ \widetilde{E}.$$

Proof. The lifted transition keeps the frame coordinate fixed. Adding $\Psi(f)$ to the two sides of (23) proves the inequality and cancels in the slack. A two-step lifted execution consists of an initial allocation and two composable underlying execution witnesses. The intermediate allocation is supplied by the chosen first lift. Its coherence identifies the resulting pullback with the lift of the composite. Pullback associativity gives the required coherence. □

It is important that the lift appears as a *hypothesis*, not as a choice we claim to make for every operation. Proposition 9.3 shows that in the plain PCM case, a strong locality condition supplies a unique lift, commuting with the inert-frame arrows of the allocation opfibration. Thus the opfibration is used through its frame-extension functoriality, not merely displayed alongside the runner calculus. In relational ownership models a related stability condition may also yield a unique lift. With proof-relevant witnesses it may be nonunique, or impossible.

Proposition 10.2 (No automatic lift of proof-relevant ownership)**.** *There is no natural rule lifting arbitrary state transitions along arbitrary proof-relevant ownership projections, even when all target fibres are inhabited.*

Proof. Take a state transition $x \mapsto y$. Let the ownership fibre over x be a singleton and that over y a two-element set $\{v_0, v_1\}$. Any lift must choose one of v_0, v_1 for the unique starting witness. The automorphism interchanging v_0, v_1 preserves the projection but neither choice, proving that no such natural selection exists. □

This counterexample is useful for the programme's promise of generality: the framework can be uniform, but permissions and witnesses cannot in general be recovered from the coarse underlying state. It is appropriate that system designers or verification tools supply, infer or check this additional information.

10.2 Reallocation requires its own quantitative certificate

Let $\tau : \mathsf{A} \to \mathsf{A}$ be a permitted transfer of allocation witnesses (possibly leaving the total physical world fixed), with potential $\Theta : \mathsf{A} \to \mathbb{N}$, actual cost c_τ, and allowance b_τ. Its certificate is

$$c_\tau(z) + \Theta(\tau z) \leq b_\tau(z) + \Theta(z). \tag{29}$$

A protocol alternating certified component updates and certified transfers inherits a global amortised bound by adding these inequalities. In particular, a zero-cost, zero-allowance transfer must *not increase* total potential. Moving a token between two accounts counted equally in Θ can be free; moving an object whose potential depends on its owner need not be.

One should distinguish this transfer certificate from the categorical fact that transfers compose. The allocation opfibration says *which transfers are admitted*; (29) states what they *cost*. This separation between permissions and quantitative guarantees is the main reason for retaining both categorical and numeric structures.

11 Temporal base change and controlled wiring

11.1 The most portable construction

The ownership fibration concerns decompositions of resource worlds. Independently, the resource coend of Section 3 has a strong stability property: ownership witnesses can be pulled back across *any* finite-horizon resource coend, even in another ambient topos. This is the cleanest change-of-topos result in our present development, and it does not depend on diagonal splitting, Moore equations or the existence of a Sweedler measuring object.

Let $I = (\mathbb{N}, \leq)$, $\mathcal{E}$ be a locally cartesian closed category with the countable coproducts and coequalisers forming the coend, preserved by pullback (in particular a Grothendieck topos), $B : I \to \mathcal{E}$ an increasing resource diagram and $S : I^{\mathrm{op}} \to \mathcal{E}$ a decreasing one. The maps $S_j \to S_0$ make every $B_n \times S_{n+k}$ an object over S_0. Therefore

$$\mathrm{Ex}_k(B, S) = \int^{n \in I} B_n \times S_{n+k} \tag{30}$$

is itself naturally an object of $\mathcal{E}/S_0$, with structure map induced by the initial state. Importantly, step execution has *not* been imposed: the map to the initial state need not survive the full operational balancing relations, since an executed update changes that state.

Theorem 11.1 (Topos-internal base change for resource evaluation)**.** *For any arrow $q : Q \to S_0$, define a pulled-back diagram $S_n^Q = S_n \times_{S_0} Q$. There is a canonical isomorphism in $\mathcal{E}/Q$, natural in all the data,*

$$\boxed{\mathrm{Ex}_k(B, S^Q) \cong \mathrm{Ex}_k(B, S) \times_{S_0} Q.} \tag{31}$$

For any geometric morphism $f : \mathcal{F} \to \mathcal{E}$, the inverse-image functor preserves the resource coend and this comparison:

$$f^* \,\mathrm{Ex}_k(B, S) \cong \mathrm{Ex}_k(f^*B, f^*S).$$

No algebraic theory or runner is needed for these two assertions.

Proof. The coend (30) is a colimit in the slice $\mathcal{E}/S_0$. In a locally cartesian closed category, the pullback functor $q^* : \mathcal{E}/S_0 \to \mathcal{E}/Q$ has a right adjoint Π_q, hence preserves colimits. It also preserves the relevant products over S_0:

$$q^*(B_n \times S_{n+k}) \cong B_n \times (S_{n+k} \times_{S_0} Q).$$

Pulling back the coend diagram gives precisely the diagram defining the left-hand side of (31). A geometric inverse-image functor preserves finite limits and colimits, and therefore the coend, its structural map and the pullback comparison. □

The theorem is deliberately narrower than an assertion that *all* Sweedler constructions transport along geometric morphisms. Internal polynomial functors may use dependent products, and measuring objects may involve further internal homs or adjoints; arbitrary inverse-image functors need not preserve those.

Proposition 11.2 (Finite-signature Moore execution commutes with inverse image)**.** *Let $f : \mathcal{F} \to \mathcal{E}$ be a geometric morphism of Grothendieck toposes. Suppose the internal bounded Moore theory has a* finite constant *observation set and a constant input signature, and its operations and execution relations are presented by finite products, coproducts and coequalisers, with finite-horizon terms formed by the usual depth-bounded inductive syntax. For internal graded algebras B and bounded comodels S of that signature, there is a natural comparison isomorphism*

$$f^*(B \mathbin{\widehat{\otimes}_T} S)_k \cong (f^*B \mathbin{\widehat{\otimes}_{f^*T}} f^*S)_k.$$

This is a statement about the explicitly presented Moore tensor*, not preservation of arbitrary internal Sweedler measuring objects.*

Proof. For a finite constant observation set, each branch-selection operation is a finite-product/finite-coproduct construction. Induction on a depth bound constructs the term objects using such operations and colimits; inverse image preserves them because it is left exact and cocontinuous. Both the resource coend and the remaining reader/update balancing are coequalisers of coproducts of these diagrams. Applying f^* therefore yields the same coequaliser as forming the balanced tensor in $\mathcal{F}$. □

For *guarded execution* the comparison also depends on a coherent operation lift to ownership witnesses. At a finite constant observation object, selecting a branch is a finite disjoint coproduct decomposition, preserved by inverse image; the associated pullback guards and finite composites of cooperations are preserved as well. This gives the pathwise comparison claimed below, under the same explicit finite-signature hypotheses. It does not assert descent to the original-state slice after operational balancing: updates change the state.

11.2 From base change to executed paths

Suppose now that S is a graded comodel and the ownership projection $q : Q \to S_0$ comes with coherent lifts of its operations. A guarded execution tree partitions the relevant bounded states according to the observed branch. Its guards are constructed by pullbacks and finite coproducts; its selected paths are constructed by composition of the cooperations. The chosen lifts make those constructions commute with the pullback along q. Thus the resource-coend comparison also respects selected paths and substitution of polynomial terms. Without these lifts, the base-change theorem remains true about the underlying resource diagrams but there is no induced ownership comodel to execute.

This is the correct categorical interpretation of our earlier frame results. The generic pullback is a fact about toposes; its application to *operational evaluation* depends on the locally specified action of the operation on ownership witnesses. It also explains why we cannot simply move the ownership witness to the final Chu quotient: that quotient may no longer retain the initial state over which the pullback was taken.

11.3 Observation-dependent feedback with separated resources

Consider two graded Moore components S^1, S^2 with update alphabets A_1, A_2, observations O_1, O_2, and restriction-stable actual costs and potentials Φ^1, Φ^2. Let F be a frame. For each horizon, suppose an admissible ownership relation

$$\Sigma_n \subseteq S_n^1 \times S_n^2 \times F$$

is preserved by restrictions and by every permitted component update: an A_1-update changes only the first owned state and restricts the second to the shorter horizon, and vice versa. Define the global potential

$$\Theta(s_1, s_2, f) = \Phi^1(s_1) + \Phi^2(s_2) + \Psi(f).$$

The controller may inspect the observations of both participating components while leaving F untouched.

Theorem 11.3 (Controlled ownership-sensitive wiring)**.** *Let*

$$w : \Gamma \times O_1 \times O_2 \to A_1 \sqcup A_2$$

be a restriction-compatible, read-only selection wire. Assume each chosen component operation preserves Σ_n, its nonselected component is only restricted, the frame remains unchanged, and each external input $\gamma \in \Gamma$ has an allowance B_γ at least as large as every selected internal allowance $b^i_{w(\gamma,o_1,o_2)}$. Then the wired graded system is certified by Θ with allowance B_γ. At each selected step its slack equals the selected component's slack plus the nonnegative enlargement of allowance. These inequalities telescope over finite wired executions and are compatible with temporal restriction.

Proof. For an A_1-step, add the component certificate $c_a^1(s_1) + \Phi^1(\delta_a^1 s_1) \le b_a^1 + \Phi^1(s_1)$ to the equalities expressing restriction stability of Φ^2 and invariance of $\Psi(f)$. The same argument works for an A_2-step. Preservation of the admissible ownership relation keeps the transition inside the allowed global carriers. Observation and restriction compatibility ensure the choice of internal letter commutes with shortening the horizon. Enlarging the allowance from b_a^i to B_γ adds the stated nonnegative slack. Repeated application telescopes the potential, as before. □

This is a genuine feedback case: the observation of one component may determine which component will execute the next operation. It is nevertheless *not* unrestricted DOTS wiring. A wire which modifies shared memory, spends the other component's credits, or adds its own unaccounted cost is outside the theorem. A synchronous step updating both independent owners is covered after adding both costs and allowances and requiring joint preservation of the ownership witness. A changing ownership allocation is covered only when the transfer itself has a certificate (29) and an admissible operational lift.

12 A running example with implementation, framing, and transfer

The binary-counter calculation demonstrates all the stages with minimal machinery. Let two counters have values m_1, m_2, stored in disjoint heap footprints. Each abstract `increment` has

one source tick; the number of residual bit flips in the chosen counter is $c_i(m_i) = 1 + t(m_i)$. Their sum potential

$$\Theta_0(m_1, m_2) = \mathrm{popcount}(m_1) + \mathrm{popcount}(m_2)$$

satisfies, for each chosen counter,

$$c_i(m_i) + \Theta_0(m_1', m_2') = 2 + \Theta_0(m_1, m_2).$$

Thus a controller that consults the two parities and decides which counter to increment has allowance 2 bit flips per external increment, in the amortised sense. It does not matter how often the controller switches between them: the potential telescopes along the actual chosen sequence. The controller reads their observations but does not change the disjointness of their footprints.

Adjoin a frame f with independent potential $\Psi(f)$. If the selected update leaves f fixed, then

$$\Theta(m_1, m_2, f) = \Theta_0(m_1, m_2) + \Psi(f)$$

certifies the same operations with exactly the same slack. Here the ownership witness is the chosen decomposition of the global heap into two disjoint footprints and a frame. If we forget that decomposition and remember only total credit, we may no longer know which counter is allowed to spend which resource, even though the numerical inequality remains superficially meaningful.

Suppose in addition that the two owners hold *separately owned* credit tokens r_1, r_2. An owner may perform an increment with *zero newly supplied allowance* only if it can spend two of its *own* tokens. The permitted transition on owner i is then

$$(m_i, r_i) \longmapsto (m_i + 1, r_i - 2), \qquad r_i \geq 2.$$

With

$$\Theta' = \mathrm{popcount}(m_1) + \mathrm{popcount}(m_2) + r_1 + r_2 + \Psi(f),$$

the binary-counter identity yields for that transition

$$c_i(m_i) + \Theta'(\text{successor}) = \Theta'(\text{initial}).$$

An increment of counter 1 is not permitted when $r_1 < 2$, even if $r_1 + r_2 \geq 2$. The ownership witness is now essential: the same total credit supports some actions but not others.

A separately permitted transfer $(r_1, r_2 + 2) \mapsto (r_1 + 2, r_2)$ leaves both counters unchanged and preserves Θ'. It therefore has a zero-cost, zero-allowance certificate, and *enables* an otherwise unavailable increment by owner 1. Thus frame extension, local-action lifting (Proposition 9.3), and transfer permissions perform genuinely different jobs in this one example. Transfers of physical heap cells with owner-dependent potential need further certification; they are not free merely because the physical heap is unchanged.

There are two boundaries to the illustration. First, the resource units are abstract bit flips; a bound on processor time or energy requires a further soundness argument relating bit flips to an actual implementation cost model. Second, the numerical sum does not itself confer permission to spend a token held by the other owner. The local-action lift and the transfer policy, compatible with the allocation opfibration, enforce that distinction.

13 How the construction relates to existing resource analysis

13.1 What we borrow, and what remains to be supplied

It is important to claim novelty at the correct level. The potential method, efficient proof-rule design, automatic resource inference, and separating resource logics have extensive prior development. The AARA survey [13] emphasizes syntax-directed resource annotations and reduction of many inference tasks to linear optimization. Atkey [10] combines affine BI resource reasoning with consumable credit inside separation logic, and related verified methods using time credits appear in [6, 7]. Atkey also implements machine-checked soundness and verification-condition generation for heap-manipulating programs. Grodin–Harper [11, 12] explain amortisation coalgebraically, including mixed-variance operations and compositional implementations in an indexed setting. The cost-indexed model of [9] relates cost and potential by adjoint graded functors and supplies a Day-convolution semantics.

The present results contribute a restricted *runner-oriented juxtaposition* of these constructions. The resource coend of Section 3 is an external module pairing; Theorems 4.2 and 5.1 show how a potential certificate follows a deterministic Sweedler implementation with a nontrivial writer residue; Theorem 6.2 places the inequality over temporal execution; and Sections 9–11 supply a general setting for ownership and coherent framing. The allocation opfibration is a conventional Grothendieck construction applied to the appropriate resource functor, not an unfamiliar new category-theoretic principle.

The division of labour suggests a route toward usable verification. An AARA-style solver could infer a candidate potential from a finite basis of quantitative observables. A local proof checker could verify (1) for the relevant runner cooperations. The theorems here would then transport these already-certified bounds through writer-residual substitutions and ownership-stable wiring. This is a design objective, *not* a claim that the present paper implements an inference algorithm.

13.2 The higher setting, and why we stop here

There is a natural way to combine the cost-action and resource-order pictures one categorical dimension higher. Let $\mathbf{C}_{\geq}$ be an ordered monoidal category of costs. Its delooping $B\mathbf{C}_{\geq}$ is a one-object bicategory whose 1-cells are costs and whose 2-cells are inequalities between costs. A pseudofunctor

$$B\mathbf{C}_{\geq} \longrightarrow \mathbf{Cat} \tag{32}$$

is a categorical action of costs, or an actegory. For additive natural-number credit, adding p credits to the frame defines a strict action on the *total* allocation category, not on one fibre: the total world changes. If $p \leq p'$, the canonical extension by $p' - p$ supplies the components of a comparison natural transformation, given by opcartesian frame-extension arrows. Thus ownership allocations give concrete categorical objects on which cost actions act.

This makes contact with the higher-runner programme of Appendix D: a categorified runner might mediate statefully between two cost actions, with potential inequalities as colax 2-cells. But there is no higher Sweedler measuring object or comparison 2-functor constructed here. The action (32) is an entry point and a test case, not proof of the desired measuring theorem. In particular, a general PCM has only partial resource addition and need not produce a global action of a monoidal cost category on all allocations.

Likewise the construction here is not an equivalence between Grodin–Harper's indexed coalgebra category and the double category of Sweedler runners. Their coalgebraic composition does suggest a promising route to such a comparison, but it must preserve the cost order, witness

transport and the laws of the residual effects; those properties are not automatic from the results above.

14 Conclusion: a resting point and a precise agenda

The two papers address complementary questions. Paper I [15] constructs the bounded operational and observational semantics, establishes the temporal Chu duality, and explains its relationship to completed behaviours. Here we have retained that finite-horizon baseline while adding the quantitative information needed to reason about an implementation: writer-residual costs, composable potential certificates, and witnessed ownership. The second paper does not re-prove the semantic results of the first; it studies additional data and compatibility laws that the no-residue theory cannot be expected to encode on its own.

We began with the practical demand that an operational implementation should admit a compositional quantitative certificate. The order of constructions was forced by what a naive approach forgets. Resource indices provide weakening and finite-horizon availability, but not actual cost. Writer residues preserve accumulated cost, and potential inequalities make them compositional. The ordinary Chu quotient may erase those costs. A writer-twisted quotient preserves costs and allowances as residual effects, while potential certificates still require source-state information or a genuinely ordered refinement. Ownership cannot be recovered from a total resource value, so proof-relevant resource splits are retained and organized into an opfibration. A locality condition makes owned-operation lifts canonical for plain PCM allocations; richer proof-relevant permissions require coherent operation lifts to make framing and controlled feedback compatible with certification. Finally, the resource base-change theorem, and a stronger finite-signature Moore instance, identify precisely which temporal evaluation constructions transport robustly between toposes.

For a new deterministic costed operational theory, the reusable tasks are now clear. One specifies its operations and equations, a cost-writer implementation, a potential or inferred potential template, and the operation's legal action on ownership witnesses. Writer composition, certified-span composition, framing, and resource-coend base change no longer require new proofs. The operadic presentation of the theory and its precise execution-balancing relations may still require analysis. The current results do not promise that arbitrary theories are automatically gradable or that general Sweedler measurement commutes with change of ambient topos.

The following four questions form a natural next programme, but *none is a prerequisite for treating the present restricted theory as complete*:

1. *Internal comparison.* Under which hypotheses is the external resource coend with operational balancing the evaluation of a Sweedler theory internal to a cost-indexed topos?
2. *General residues.* Can the quantitative certificate calculus be extended from deterministic cost writers to branching, probabilistic or other effects, with a suitable inequality-valued interaction law?
3. *General wiring.* Can the ownership opfibration and certified spans be incorporated in a double category of interfaces and runners, compatible with genuinely shared-state DOTS wiring?
4. *Higher actions.* Do cost actegories and colax potential cells admit a higher measuring construction that formalizes resource-certified stateful translations?

Our result is therefore neither a complete resource analysis tool nor a universal theorem about Sweedler theory. It is a *coherent, reusable fragment* linking graded execution, nontrivial

residues, amortised certification and witnessed ownership. The fragment is sufficiently explicit to support implementation examples and sufficiently constrained to show what must still be added before the methods extend to unrestricted systems.

From resource certificates to physical deadlines. The *Temporal Interlude* [14] supplies an application-level continuation to these calculations: a stateful runner is unfolded into an interval-indexed machine, a duration potential bounds the time between logical completion events, and a timed monitor turns a response deadline into a state-safety assertion. In contrast, this paper records writer effects, quantitative slack, and ownership witnesses before (and sometimes beyond) behavioural identification. Connecting the constructions by a general realization functor compatible with the full temporal Chu tensor and ownership-sensitive wiring remains open.

A A compact map of the four source notes

This integration draws on four working documents, in their development order. The map records their distinct mathematical roles so that the consolidation can be checked against the detailed proofs rather than obscuring their provenance.

1. *Two Cost Worlds and Their Interaction*. Resource-indexed variance, the external coend comparison, deterministic writer-runner/colax cells, and categorical cost actions. Its source distinctions are used in Sections 2–4.
2. *Certified Residual Runners and Temporal Chu Evaluation*. Writer actions and cocycles, associative certificate composition, exact slack, temporal cost labels and independent feedback. They form Sections 5–6 and part of Section 11.
3. *Ownership-Indexed Certified Runners*. Witnessed disjointness, guarded-frame evaluation, controlled ownership-sensitive wiring, separately certified transfers and the Day-convolution loss-of-witness example. They motivate Sections 7, 10 and 11.
4. *Ownership Fibrations and Certified Residual Execution*. Allocation and transfer opfibrations, the certified-span bicategory, coherent frame lifts, the topos-internal coend theorem and categorical cost actions. They supply Sections 8–11 and the higher-categorical outlook.

The central limitations have been retained rather than smoothed over: the ordinary Chu quotient can lose numerical costs; the Part VI internal Sweedler tensor has not been identified with the external coend; proof-relevant ownership lifts are not canonical; and general shared-resource or effectful wiring remains open.

References

[1] R. E. Tarjan, *Amortized Computational Complexity*, SIAM Journal on Algebraic Discrete Methods **6**(2), 306–318 (1985), https://doi.org/10.1137/0606031.

[2] V. Rajani, M. Gaboardi, D. Garg, and J. Hoffmann, *A Unifying Type-Theory for Higher-Order (Amortized) Cost Analysis*, Proceedings of the ACM on Programming Languages **5**, POPL, article 27 (2021), https://doi.org/10.1145/3434308.

[3] T. Uustalu, *Stateful Runners of Effectful Computations*, Electronic Notes in Theoretical Computer Science **319**, 403–421 (2015), https://doi.org/10.1016/j.entcs.2015.12.024.

[4] D. Ahman and A. Bauer, *Runners in Action*, ESOP 2020, LNCS **12075**, https://doi.org/10.1007/978-3-030-44914-8_2.

[5] C. Calcagno, P. W. O'Hearn, and H. Yang, *Local Action and Abstract Separation Logic*, LICS 2007, https://doi.org/10.1109/LICS.2007.30.

[6] A. Charguéraud and F. Pottier, *Verifying the Correctness and Amortized Complexity of a Union-Find Implementation in Separation Logic with Time Credits*, Journal of Automated Reasoning **62**, 331–365 (2019), https://doi.org/10.1007/s10817-017-9431-7.

[7] G. Mével, J.-H. Jourdan, and F. Pottier, *Time Credits and Time Receipts in Iris*, ESOP 2019, https://doi.org/10.1007/978-3-030-17184-1_1.

[8] 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, https://doi.org/10.2168/LMCS-8(4:1)2012.

[9] D. Binder, D. Corfield, D. Orchard, and V. Rajani, *Categorical Models of Amortized Cost: An Adjoint Relationship between Cost and Potential*, arXiv:2608.09635 (2026), https://arxiv.org/abs/2608.09635.

[10] R. Atkey, *Amortised Resource Analysis with Separation Logic*, Logical Methods in Computer Science **7**(2), article 17 (2011), https://arxiv.org/abs/1104.1998.

[11] H. Grodin and R. Harper, *Amortized Analysis via Coalgebra*, Electronic Notes in Theoretical Informatics and Computer Science **4**, MFPS XL (2024), https://doi.org/10.46298/entics.14797.

[12] H. Grodin and R. Harper, *Amortized Analysis via Coinduction*, LIPIcs CALCO (2023), https://doi.org/10.4230/LIPIcs.CALCO.2023.23.

[13] J. Hoffmann and S. Jost, *Two Decades of Automatic Amortized Resource Analysis*, Mathematical Structures in Computer Science (2022), https://doi.org/10.1017/S0960129521000487.

[14] *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).

[15] *Sweedler Runners for Temporal Systems, I: From Graded Operational Syntax to Temporal Behaviour*, companion research draft (October 2026).