Self-Similarity along Compact Seams

Metric Completion and Final Coalgebras in a Condensed ∞-Topos

Anonymous submission

Abstract

Finite point gluing produces final coalgebras by completing an initial algebra. We extend that construction to a finite directed system of sorts, compact marked boundary spaces, and compact overlap quotients. An exact formula for the metric across the overlap reduces admissibility to the boundary seed. It gives an isometric algebra on the initial chain, and completion yields a compact final coalgebra for sets, closed-marked topological spaces, and short maps with isometric markings. Graph-directed branch lengths may vary, and individual edges may have weight one if products along long paths converge to zero. For a compatible geometric system we construct a canonical continuous quotient onto the attractor and prove that this is a homeomorphism exactly when the attractor realizes the prescribed overlap relation. In the condensed ∞-topos the same compact object represents the final coalgebra, including in slices over higher bases and after descent of the specifications. The single-sort, finite-boundary specialization is naturally identical to the quotient functor, completed metric, and coalgebra of Ratnayake and Manokaran, and recovers their Hutchinson comparison with its bi-Lipschitz conclusion. A Lean development checks the algebraic specialization.

Keywords. Self-similarity, final coalgebra, compact gluing, graph-directed iterated function system, condensed mathematics, descent.
MSC 2020. 18N60, 18C50, 54E50, 28A80.

Contents

1 Introduction **3**
1.1 Main results . . . 3
1.2 Relation to the literature . . . 5
1.3 Exactness of the specialization . . . 5
1.4 Conventions . . . 6

2 Marked gluing and its initial chain **6**
2.1 The data and the pushout . . . 6
2.2 Closed quotients and the finite-depth pieces . . . 7

3 The seam metric **9**
3.1 Path pseudometrics and the boundary test . . . 9
3.2 Distances through the seam . . . 9
3.3 Completeness and dense embeddings . . . 11
3.4 A finite shortest-path description . . . 12

4 The compact completion of the chain **13**
4.1 Path products . . . 13
4.2 The stage metrics and the shift . . . 13
4.3 Total boundedness and completion . . . 14
4.4 Algebraic independence of the weights . . . 15

5 The final coalgebra and its mediating maps **15**
5.1 Finality in sets . . . 15
5.2 Finality for continuous coalgebras . . . 16
5.3 Shortness and quantitative stability . . . 17
5.4 Changing the scales . . . 17

6 Graph-directed attractors and exact overlaps **18**
6.1 Compatible geometric realizations . . . 18
6.2 Initial stages and orbit injectivity . . . 19
6.3 Address spaces and their equivalence relations . . . 19
6.4 The completed overlap criterion . . . 20
6.5 Bounds and failure of finite orbit detection . . . 21

7 Final coalgebras in a condensed ∞-topos **22**
7.1 The site and the marked endofunctor . . . 22
7.2 Evaluation at the singleton . . . 23
7.3 The completed space represents the higher fixed point . . . 24
7.4 Coinduction on compact plots . . . 25
7.5 An effective presentation by addresses . . . 26
7.6 Descent and monodromy . . . 27
7.7 Why the cover condition matters . . . 28

8 Equivalence with finite point gluing **28**
8.1 An isomorphism of gluing functors . . . 28
8.2 Finite boundary estimates and admissibility . . . 29
8.3 The module presentation . . . 30
8.4 Result-by-result comparison . . . 31

9 Classical and compact-seam examples **32**
9.1 The interval and its dyadic algebra . . . 32
9.2 Binary streams, the Cantor set, and p-adic integers . . . 32
9.3 A geometric lemma for intrinsic metrics . . . 33
9.4 The Sierpiński gasket . . . 34
9.5 An interval system with two sorts and a unit weight . . . 35
9.6 A square with shared edges . . . 35
9.7 The carpet with its full overlap relation . . . 35
9.8 Nested radicals as a continuous coalgebra . . . 37

10 Changing the overlap relation **37**
10.1 Role of the hypotheses . . . 38
10.2 Conclusion . . . 38

A Lean verification of the specialization **39**

B Quantitative comparison formulas **40**
B.1 Hölder bound for a coalgebra map 40
B.2 Iterated orbit estimates 40
B.3 Forward bounds when the scales differ 40

1 Introduction

Completing a sequence of increasingly fine descriptions is a familiar way to construct a fractal. A categorical construction has to explain two independent facts. The finite descriptions must form an initial algebra, and the complete object must admit a unique map from any coalgebra of the same gluing functor. Ratnayake and Manokaran prove both for a finite point-gluing specification with a metrically admissible scale [18]. They also compare the result with the classical attractor of an iterated function system under injectivity and bi-Lipschitz hypotheses. Their example of the Sierpiński carpet exhibits a different problem from choosing the correct scale. Two first-level tiles can share a whole interval. Recording only coincidences among finitely many distinguished points cannot record that relation.

We allow a branch to have a sort different from that of the space it helps construct. Each sort has a compact metric space of marks. For each output sort, the disjoint union of the boundaries of its incoming pieces maps onto a compact space of seams. This single quotient records simultaneous identifications, including identifications within a branch. Every seam point has a boundary representative. Output marks are selected by an embedding into the seam space. Boundary distances are part of the specification. The criterion for their compatibility with the construction is an isometry at the first step of the initial chain.

The distinction between marks and seams matters. The output marks must be carried through each substitution, whereas a seam may include points that have no output label. For a square subdivided into four squares, it is convenient to mark the entire perimeter. The seams include the inside grid lines as well. In the initial chain, the successive stages are compact skeletons. The stages embed isometrically and their completion fills the square. There is an analogous construction of the Sierpiński carpet with a path metric on the surviving grid. In both examples the branch quotient identifies shared edges, which do not fit the point-gluing relation in [18].

A second distinction concerns the topology of the answer. The set of points of a complete solution is final for the associated functor on marked sets. The metric and the topology are extra structure with their own universal properties. An ∞-topos of sheaves on compact Hausdorff probes retains this topology and tests the construction against families over arbitrary spaces of parameters. Its compact covers can resolve which branch contains the values of a plot even when that plot crosses a seam. This property underlies the higher finality theorem. It also explains the role of the site. For example, the identity plot near the midpoint of an interval cannot be locally lifted into one closed half for the open-cover topology.

1.1 Main results

A specification in this paper consists of finite sorts s, finite nonempty branch sets E_s, target sorts $t(e)$, nonempty compact marked spaces B_s, seam quotients q_s, output embeddings b_s, and

weights $r_e \in (0, 1]$. The gluing functor is a pushout under the output marks. Its n-fold iterate has a branch for each directed word w of length n and weight r_w. Write ϵ_n for the maximum of those weights. We assume $\epsilon_n \to 0$. On a finite graph, this condition is equivalent to requiring each directed cycle to contain an edge of weight strictly less than one.

Theorem 1.1 (Metric completion and finality)**.** *Suppose that the path pseudometric on the first-step seam spaces is a metric inducing their compact topology, and that the output maps b_s are isometric for these metrics. Then the initial algebra G of the marked-set gluing functor is the union of an isometric initial chain, starting at B. Its componentwise Cauchy completion $\widehat{G}$ is compact. The algebra extends to an isometric isomorphism $a\colon \mathcal{F}\widehat{G} \to \widehat{G}$. Its inverse is final in marked sets with injective markings, in Hausdorff spaces with closed embedded markings and continuous maps, and in bounded metric spaces with isometric markings and short maps.*

The proof first computes distances in a single substitution. The only endpoints of a nontrivial zero-cost jump lie over the seam. The subspace metric on the seam therefore depends only on the metric of the boundary seed. The resulting two-endpoint formula controls positivity, topology, dense inclusions, and completion. The coalgebra proof in sets uses contraction of a mapping operator with ϵ_n as its n-step bound. The topological proof uses the finite closed cover by depth-n pieces. It remains valid when the domain of a coalgebra is noncompact and has no continuous retraction onto its marked boundary.

Theorem 1.2 (Attractor comparison)**.** *Suppose that a graph-directed system on complete bounded metric spaces is compatible with the marks and seams and has shrinking products of branch Lipschitz constants. Let A be its compact attractor and O its marked orbit. There is a canonical continuous surjection $h\colon \widehat{G} \to A$ extending the initial-algebra map $G \to O$. It is a homeomorphism if and only if the induced algebra $\mathcal{F}A \to A$ is injective. If $G \to O$ is bi-Lipschitz and the orbit algebra is injective, then h is bi-Lipschitz and gives the final coalgebra on A in the marked set and continuous categories.*

The forward comparison is established by factoring the geometric address evaluation through the canonical one. The proof follows two addresses up to their first seam identification. All addresses of an output mark have the same geometric value because the marks satisfy the branch equations at every depth. This makes the map to the attractor available before any injectivity or inverse Lipschitz estimate is known. Orbit injectivity and exactness of the attractor overlaps are separate hypotheses. We give an example in which all finite orbit stages are injective and the completed map identifies pairs.

Theorem 1.3 (Higher finality)**.** *Under the hypotheses of Theorem 1.1, let y send a compact Hausdorff space to its representable object in the ∞-topos of sheaves on compact Hausdorff spaces for finite jointly surjective covers. Form the gluing functor with homotopy pushouts and monic markings yB_s. Then $y\widehat{G}$, with the inverse of the induced algebra, is final. In the slice over any base object T, the corresponding final coalgebra is $T \times y\widehat{G} \to T$. Specifications obtained from these by descent have final coalgebras that descend along the same covers. Evaluation at the singleton recovers the marked-set final coalgebra on discrete inputs.*

Here finality means that the mapping *space* of coalgebra morphisms is contractible. The input can carry homotopy in its sections and in its base. In particular a finite group action on a specification gives a descended final object over its classifying object. A reflection of the interval yields a family with nontrivial monodromy. The higher proof uses the metric to turn

a unique pointwise map into a map on every compact plot. It uses the sheaf condition to lift finite observations, which is weaker than lifting a whole infinite address.

1.2 Relation to the literature

Hutchinson's attractor theorem supplies the geometric comparison in the single-sort contractive case [9]. Graph-directed fractal geometry has long allowed different types and branch scales [15]. The role of the graph here is to organize categorical substitution and to determine a bound on the diameter of the images of all depth-n branches. The graph-directed attractor is used with its full overlap relation, including when an overlap is a compact subset of a boundary. We do not use an open set condition or a dimension formula.

Leinster's universal solutions provide a general theory of simultaneous gluing equations based on modules [11]. His results include a universality theorem for systems satisfying a nondegeneracy condition and comparisons with IFS attractors. Hasuo, Jacobs, and Niqui study the resulting symbolic descriptions and the injectivity of their interpretations [8]. Our one-step pushout is a concrete marked presentation. Its metric uses a common interface that can be an infinite compact space, and we prove an explicit criterion on the seed metric that ensures the completion description. The proof of the address comparison uses exactness of a seam quotient and gives an attractor criterion without invoking a module classification. In the finite pointed case we identify the associated two-object module to fix the relation with the module construction.

The use of a mapping-space contraction has antecedents in metric semantics of recursive domain equations [2, 19]. Barr and Adámek treat related relationships between initial algebras, final coalgebras and completions for classes of set functors [3, 1]. Their ultrametric and limit-based constructions address different categories and hypotheses. Our metric can be Euclidean, can arise from path lengths, and need not be an ultrametric. The category of sets with multiple injective markings need not admit the usual terminal object of the unrestricted undercategory. We use an initial chain with an isometric seed and prove the final property directly.

The compact-cover site is the one used in condensed mathematics [5]. We use it to define a concrete ∞-topos, and use standard higher-topos facts about universality of colimits, truncation and descent [12, 13]. The higher proof supplies a terminal mapping-space computation for this specific gluing functor. It does not require a general identification of arbitrary metric completions with higher limits. For example, taking a homotopy pushout without the compact-cover argument would leave unresolved whether a plot into the complete geometric object has the required local lifts.

1.3 Exactness of the specialization

Our comparison with [18] is functorial. For one sort and a discrete finite boundary I, take $K = (M \times I)/\overline{R}$, q the class map, $b(j) = [D_j]$, and $r_e = \rho$. We construct a natural isomorphism of pointed functors, show that it is an isometry for every non-degenerate pointed metric input, and identify both the initial chain and its completion. The final coalgebra equivalence is over the same set and metric categories, so it identifies the unique mediating maps as well as their carriers. The orbit algebra and the completed comparison have the same branch formulas, which proves the corresponding identification for the categorical Hutchinson theorem. Section 8 gives a result-by-result crosswalk. Appendix A specifies the machine-checked algebraic part of this identification.

1.4 Conventions

All metric spaces in the main construction have distances at most one. A map is *short* if it has Lipschitz constant at most one. A closed marking of X by B is an embedding onto a closed subspace. We work with classical sets in the metric and topological arguments. A fixed universe is understood for small sets, and a larger one for the category of sheaves on compact Hausdorff probes of bounded universe size. All spaces used as probes in the proofs, including the finite pullbacks and compact subsets that refine covers, lie in the probe universe. In the higher sections, Map denotes a space, pushouts are homotopy pushouts, and $\tau_{\leq 0}$ denotes the 0-truncation. A mono in an ∞-topos is a (-1)-truncated map.

2 Marked gluing and its initial chain

2.1 The data and the pushout

Let S be a finite nonempty set of sorts. For $s \in S$, let E_s be a finite nonempty set. We draw an edge $e \in E_s$ from s to $t(e)$ because a point of sort s is obtained by substituting a point of sort $t(e)$ into the e-th branch. Let B_s and K_s be nonempty compact Hausdorff spaces. The specification has maps

$$J_s = \coprod_{e \in E_s} B_{t(e)} \xrightarrow{q_s} K_s, \qquad B_s \xrightarrow{b_s} K_s. \tag{2.1}$$

We require q_s to be a continuous surjection and b_s to be an embedding. Since J_s is compact and K_s is Hausdorff, q_s is a closed quotient map. These data can first be considered in sets. No metric or rate enters the algebraic construction.

Write $\mathsf{Set}_B^{\mathrm{inj}}$ for the category of S-tuples of injections $j_s \colon B_s \to X_s$ with all maps commuting with the j_s. A morphism need not be injective on unmarked points and may send them to marks. In the topological version $\mathsf{Top}_B^{\mathrm{cl}}$, all X_s are Hausdorff and the markings are closed embeddings. An object is specified by the underlying tuple and its marking. The initial object is B, with identity markings. We denote its unique map to X by j_X.

Define the underlying component of $\mathcal{F}X$ by the pushout

$$\begin{array}{ccc} \coprod_{e \in E_s} B_{t(e)} & \xrightarrow{\quad q_s \quad} & K_s \\ {\scriptstyle \coprod j_{t(e)}}\downarrow & & \downarrow \\ \coprod_{e \in E_s} X_{t(e)} & \xrightarrow[\quad p_{s,X} \quad]{} & (\mathcal{F}X)_s. \end{array} \tag{2.2}$$

The output mark is $B_s \xrightarrow{b_s} K_s \to (\mathcal{F}X)_s$. The branch $p_e^X \colon X_{t(e)} \to (\mathcal{F}X)_s$ is the restriction of $p_{s,X}$ to its e-th summand. We write $[e, x]$ for $p_e^X(x)$ in sets. Surjectivity of q_s means that every point of $(\mathcal{F}X)_s$ is of the form $[e, x]$.

Lemma 2.1 (Normal form)**.** *There is a canonical bijection, natural for maps that carry unmarked points to unmarked points,*

$$(\mathcal{F}X)_s \;\cong\; K_s \;\sqcup\; \coprod_{e \in E_s} \bigl(X_{t(e)} \setminus j_{t(e)} B_{t(e)}\bigr).$$

For arbitrary marked maps the induced map fixes K_s. On an unmarked representative it takes $[e, x]$ to $[e, f(x)]$, which may belong to K_s. Two branch representatives are equal if and only if they are the same representative, or both are marked and have the same q_s-image.

Proof. Send a marked representative $(e, j_{t(e)}b)$ to $q_s(e, b)$. Send every other representative to its own element of the displayed disjoint union. There is a well-defined function because each $j_{t(e)}$ is injective. It agrees with the identity on the seam in the pushout square. Its inverse uses the inclusion of K_s and the unmarked branch representatives. This gives the equality criterion and determines the action on maps. □

Consequently, $\mathcal{F}$ preserves componentwise injective maps of marked objects. Indeed, an injective marked map reflects membership in the marked subset, and Lemma 2.1 then shows injectivity on the unmarked summands and on the seam. $\mathcal{F}$ also preserves colimits of increasing sequences of injective maps in the marked-set category. For such a sequence, a representative in the union lies at some stage. Equality of two representatives in the gluing of the union is detected in a common stage, because membership in B and equality of marked representatives are both detected there. The canonical comparison is thus surjective and injective. One could obtain preservation of filtered colimits directly from the pushout and coproduct, but the representative description also fixes the maps used in the metric argument.

Proposition 2.2 (The initial algebra)**.** *Let $A_0 = B$, $A_{n+1} = \mathcal{F}A_n$, $i_0 = j_{\mathcal{F}B}$, and $i_n = \mathcal{F}^n i_0 \colon A_n \to A_{n+1}$. In $\mathsf{Set}_B^{\mathrm{inj}}$, the colimit G of this chain has an initial algebra structure $a_G \colon \mathcal{F}G \to G$. If $\iota_n \colon A_n \to G$ is the inclusion, this structure is the unique map satisfying*

$$a_G \mathcal{F}(\iota_n) = \iota_{n+1}. \tag{2.3}$$

It is an isomorphism.

Proof. The first inclusion identifies B_s with $b_s B_s \subseteq K_s$ and is injective. All subsequent inclusions are injective by Lemma 2.1. Since $\mathcal{F}$ preserves this colimit, the shifted chain and the original chain give the same colimit and the shift comparison is an isomorphism. Let $\alpha \colon \mathcal{F}X \to X$ be any marked algebra. Define $u_0 = j_X$ and $u_{n+1} = \alpha \mathcal{F} u_n$. At stage zero, $u_1 i_0 = u_0$ because all maps are marked. Applying $\mathcal{F}$ and α proves compatibility at every stage. The induced $u \colon G \to X$ is a homomorphism by (2.3). Every homomorphism has these restrictions, which proves uniqueness. □

The same initiality assertion holds on the full marked-set undercategory when $\mathcal{F}$ is defined there by the pushout. Indeed the initial chain and its carrier colimit are unchanged. For a target algebra with possibly noninjective markings, the same compatible maps u_n are still defined, and the branch representatives and seam equations give (2.3) and the same uniqueness argument. Non-degeneracy enters the finality theorem, where all marked maps into the completed carrier must exist.

This construction allows nontrivial identifications inside one marked boundary. Injectivity of the chain maps is a property of applying $\mathcal{F}$ to an injective marked map. It does not assert injectivity of each branch p_e^X. That branch is injective precisely when q_s is injective on the e-th boundary summand. Keeping these two kinds of injectivity distinct is necessary for the attractor comparison.

2.2 Closed quotients and the finite-depth pieces

Proposition 2.3 (Topological gluing)**.** *For $X \in \mathsf{Top}_B^{\mathrm{cl}}$, the topological pushout in (2.2) is Hausdorff. The map $p_{s,X}$ is a closed surjection with compact fibers, and the inclusion of K_s is a closed embedding. Thus $\mathcal{F}$ is an endofunctor of $\mathsf{Top}_B^{\mathrm{cl}}$.*

Proof. Put $W = \coprod_e X_{t(e)}$, and regard J_s as its closed compact marked subset. Give the set quotient of W the quotient topology. For a closed subset $D \subseteq W$ its saturation is

$$D \ \cup \ q_s^{-1}(q_s(D \cap J_s)).$$

The added set is compact and is closed in the Hausdorff space W. The quotient map is closed. Its fibers are singletons away from the seam and closed subsets of the compact set J_s on the seam. Two disjoint compact subsets of a Hausdorff space can be separated by disjoint open neighborhoods. Given two distinct fibers, choose such neighborhoods U, V and remove the closed images of $W \setminus U$ and $W \setminus V$, respectively. This gives disjoint open neighborhoods in the quotient. The map K_s into the quotient is continuous since q_s is a quotient. It is an embedding by compactness and Hausdorffness. Its image is closed. The same quotient topology gives the topological pushout, as a compatible continuous map on W descends and agrees with the proposed map on K_s by surjectivity of q_s. □

A word of length n at sort s is a composable sequence of n edges starting at s. Let $W_{s,n}$ be the finite set of such words, $t(w)$ their terminal sort, and take $W_{s,0} = \{\varnothing_s\}$ with $t(\varnothing_s) = s$. Iteration gives maps

$$p_w^X \colon X_{t(w)} \longrightarrow (\mathcal{F}^n X)_s, \qquad p_{s,n,X} \colon \coprod_{w \in W_{s,n}} X_{t(w)} \longrightarrow (\mathcal{F}^n X)_s.$$

Each $p_{s,n,X}$ is a surjection. In the closed-marked topological category it is closed. This follows by induction, since the length-$(n+1)$ quotient is the composite of a finite coproduct of closed length-n quotients with $p_{s,\mathcal{F}^n X}$. In particular every piece $p_w^X(X_{t(w)})$ is closed. Compact fibers are also preserved under this composition. To see this for any closed map with compact fibers, an open cover of the preimage of a compact set admits a finite subcover on each fiber. A neighborhood of each fiber in the range has its full preimage covered by those finitely many opens, since the map is closed. A finite subcover in the range completes the argument.

Lemma 2.4 (Tracking a word)**.** *For every marked $f \colon X \to Y$,*

$$(\mathcal{F}^n f)_s p_w^X = p_w^Y f_{t(w)}.$$

For an algebra $a \colon \mathcal{F}Y \to Y$ write $a^{[0]} = \mathrm{id}$ and $a^{[n+1]} = a\mathcal{F}(a^{[n]})$, and for a coalgebra $c \colon X \to \mathcal{F}X$ write $c^{[0]} = \mathrm{id}$ and $c^{[n+1]} = \mathcal{F}^n(c)c^{[n]}$. A marked map f satisfies $f = a\mathcal{F}(f)c$ if and only if it satisfies this identity for all iterates,

$$f = a^{[n]}\mathcal{F}^n(f)c^{[n]} \quad (n \geq 1).$$

Proof. The first identity holds at depth one by the pushout and iterates. Inducting on n gives $(u \mapsto a\mathcal{F}(u)c)^n(f) = a^{[n]}\mathcal{F}^n(f)c^{[n]}$, using functoriality and associativity of composition. The second assertion follows. □

In a multitype system, $W_{s,n}$ is more informative than the set of sequences in the union of all branch alphabets. For instance a branch can have weight one, then lead to a type for which every outgoing branch strictly shrinks. Typing is needed to measure the products that can actually occur and to identify the correct boundary when the last sort of a word changes.

3 The seam metric

3.1 Path pseudometrics and the boundary test

Give B_s a metric of diameter at most one inducing its topology, and choose $r_e \in (0,1]$. For the moment let X be a family of one-bounded metric spaces with isometrically embedded markings. Equip the coproduct $W_s(X) = \coprod_e X_{t(e)}$ with the metric

$$\mathrm{d}_W((e,x),(f,y)) = \begin{cases} r_e \mathrm{d}_{X_{t(e)}}(x,y), & e = f, \\ 1, & e \neq f. \end{cases} \tag{3.1}$$

The disjoint-union topology is the metric topology. For $z, z' \in (\mathcal{F}X)_s$, take the infimum of the total cost of finite chains that start over z, end over z', and alternate paid d_W steps and zero-cost jumps between representatives of the same point. Denote the result by $\mathrm{d}_{\mathcal{F}X}(z,z')$. A one-step chain gives the upper bound one. Inserting or deleting a zero step allows concatenation, so this is a pseudometric. It is also the largest pseudometric on the quotient making $p_{s,X}$ short. Indeed, any such pseudometric is bounded above by the cost of each chain by the triangle inequality.

Apply this construction to $X = B$. The set quotient is K_s. Denote the resulting pseudometric by δ_s. For $e \in E_s$,

$$\delta_s(q_s(e,b), q_s(e,c)) \leq r_e \mathrm{d}_{B_{t(e)}}(b,c). \tag{3.2}$$

If δ_s is separated, (3.2) implies continuity of q_s into (K_s, δ_s). The identity from the originally specified K_s to (K_s, δ_s) is continuous by the quotient property, hence is a homeomorphism by compactness. In particular (K_s, δ_s) is compact.

Definition 3.1 (Metric admissibility)**.** A weighted specification is *metrically admissible* if every δ_s is separated and $b_s \colon B_s \to (K_s, \delta_s)$ is an isometry.

The separation clause is automatic for finite J_s. It is a genuine metric condition for compact J_s. For example, let the boundary of a single branch be $[0,1]$ with its ordinary metric and let q be the Cantor function onto $[0,1]$. Its kernel is a closed equivalence relation. The quotient pseudodistance between its endpoint values is zero. To see this, pass through the complementary intervals removed at the first n stages of the middle-thirds construction at zero cost. The total length of the other steps is bounded by $(2/3)^n$ before multiplying by the branch weight. The same argument between any two boundary representatives gives pseudodistance zero. Taking b to be the identity into the topological quotient meets the topological hypotheses, but fails Definition 3.1. This example isolates the need for the boundary test even when every equivalence class is compact.

3.2 Distances through the seam

For representatives $z = p_e^X(x)$ and $z' = p_f^X(y)$ define

$$D_{e,f}(x,y) = \begin{cases} r_e \mathrm{d}_{X_{t(e)}}(x,y), & e = f, \\ +\infty, & e \neq f. \end{cases}$$

Write j_t for the marking of sort t. The cost of entering and leaving the seam is

$$\begin{aligned} V_{e,f}(x,y) = \inf_{b \in B_{t(e)},\ c \in B_{t(f)}} \{ & r_e \mathrm{d}_{X_{t(e)}}(x, j_{t(e)} b) \\ & + \delta_s(q_s(e,b), q_s(f,c)) + r_f \mathrm{d}_{X_{t(f)}}(j_{t(f)} c, y) \}. \end{aligned} \tag{3.3}$$

Theorem 3.2 (Two-endpoint formula)**.** *The restriction of the quotient pseudometric to K_s is δ_s. For all choices of representatives,*

$$\mathrm{d}_{\mathcal{F}X}(z, z') = \min\{1, D_{e,f}(x, y), V_{e,f}(x, y)\}. \tag{3.4}$$

If δ_s is separated, the infimum in (3.3) is a minimum. The quotient pseudometric is then separated, and its topology agrees with the closed topological quotient of Proposition 2.3.

Proof. A chain of cost less than one has no paid step between different branches. Consecutive paid steps in one branch with identical meeting representatives can be combined by the triangle inequality. At a nontrivial zero-cost jump, both representatives are marked, by Lemma 2.1. For a chain between seam points, all remaining paid steps after this simplification can be taken to join marked points. Indeed choose marked representatives of the endpoints, and group consecutive steps between the endpoint and the next seam visit, or between two seam visits, into one step in their branch. Its endpoints are marked. The marking is isometric, so the resulting chain is a boundary chain of the same or lower cost. Conversely a boundary chain is a chain in X. Truncating at one establishes equality of the two pseudometrics on the seam.

For general endpoints, precompose and postcompose any chain with zero-cost jumps from and to the specified representatives. If the chain stays in a single copy without visiting the seam, the triangle inequality gives the direct term. Otherwise take its first and last seam visits. The portion before the first visit takes place in the branch e, and has length at least its corresponding entry term in (3.3). If z is already on the seam, take b to be its specified marked representative, so that this term is zero. The same argument applies after the last visit in branch f. The middle portion costs at least the boundary pseudodistance just proved. This gives the lower bound in (3.4), including the possibility of chains of cost at least one. Conversely, take the direct paid step when it is available, or join x to jb, take a boundary chain approximating δ_s, and join jc to y. Together with the universal bound one these give the other inequality.

Suppose now that δ_s is separated. The objective in (3.3) is continuous on the compact product of boundaries, by (3.2) and continuity of the metric δ_s. Thus it attains its minimum. Put

$$R_e(x) = r_e \min_{b \in B_{t(e)}} \mathrm{d}_{X_{t(e)}}(x, j_{t(e)} b). \tag{3.5}$$

For an unmarked point, $R_e(x) > 0$. The quotient distance from $p_e^X(x)$ to K_s is exactly $R_e(x)$. A closest boundary representative supplies the upper bound. For the lower bound, the first seam visit of any chain of cost less than one has cost at least $R_e(x)$; and $R_e(x) \leq 1$. It follows that a point off the seam has positive distance to every seam point. Two points off the seam have positive distance if they are distinct, since the via-seam cost is bounded below by the sum of the two positive radii and the direct term is positive or infinite. Distinct seam points are separated by δ_s.

The map from the topological quotient to the metric quotient is continuous by the shortness of $p_{s,X}$ and the quotient property. To prove that it is a homeomorphism, let $z_n \to z$ in the metric. If z is off the seam, then for all sufficiently large n the metric distance is less than $\min(1, R_e(x))$. It must be realized by the direct term and z_n has an unmarked representative in the same branch, converging there to x. If z is on the seam, choose for each z_n a seam point k_n minimizing its distance to the seam. For z_n on the seam take $k_n = z_n$. Otherwise choose $k_n = q_s(e_n, b_n)$ with b_n minimizing (3.5). Then $k_n \to z$ in the seam metric and hence in the compact seam topology. For any subsequence of off-seam terms, there is a further subsequence with fixed

$e_n = e$ and $b_n \to b$. The representatives x_n on this further subsequence converge to $j_{t(e)}b$ because $R_e(x_n) \to 0$ and $r_e > 0$. Their images converge in the quotient topology to $q_s(e, b) = z$. A subsequence already on the seam has the same conclusion. Every subsequence has a further subsequence converging to z in the quotient topology. If a quotient-open neighborhood of z failed eventually to contain the original sequence, the subsequence outside it would contradict this property. Thus every quotient-open set is sequentially open for the metric topology, which is first-countable. □

Several features of (3.4) are important. The underlying set quotient has not been changed by identifying distance-zero classes. The proof establishes separation for exactly the specified equivalence relation. A branch itself may have shortcuts through the seam. Its map is r_e-Lipschitz, but its subspace metric need not be $r_e \mathrm{d}_X$. By contrast, the comparison between two stages coming from a marked isometry is itself an isometry, because the formula retains all boundary choices in both stages.

Proposition 3.3 (Functorial estimates)**.** *Under metric admissibility, $\mathcal{F}$ acts on one-bounded metric spaces with isometric markings, with either continuous or short maps. If a marked map is componentwise L-Lipschitz for $L \geq 1$, then so is $\mathcal{F}f$. Isometric embeddings are preserved. For componentwise uniform continuity, $\mathcal{F}f$ is uniformly continuous. The metric formula and these conclusions also apply to two families with different metrics on the same marked spaces, when both induced boundary seam pseudometrics are separated and the markings induce the specified topologies.*

Proof. The output markings are isometric by Theorem 3.2 and Definition 3.1. Continuity for continuous f follows from the identification with the topological pushout. On a paid within-branch step an L-Lipschitz map changes the cost by a factor at most L; the cross-branch cost is at most one, and zero jumps remain zero jumps. This proves the Lipschitz estimate by taking infima. An isometric embedding of marked inputs leaves the boundary distances, the direct term and the objective in (3.3) unchanged. By injectivity in Lemma 2.1 it is therefore an isometric embedding on the output.

For the last assertion first use the same argument with the actual metric pulled back from each input onto its marked spaces. If its seam metric is separated, Theorem 3.2 applies with that induced metric in place of the original seed metric. To see uniform continuity in this generality, take a common modulus for the finite family of f_t. For a sufficiently small output-domain distance, (3.4) gives either a direct paid step of small length or a via-seam representation with each of its three nonnegative terms small. The finitely many positive r_e and the common modulus control the entry and exit terms in the target. The two seam metrics induce the same compact topology, so the identity between them is uniformly continuous. It controls the middle term. This gives uniform continuity even when the two marked metrics differ. When the boundary is isometric on both sides the middle term is unchanged. □

3.3 Completeness and dense embeddings

Proposition 3.4 (Completion commutes with gluing)**.** *If the induced seam metrics are separated, $\mathcal{F}$ takes componentwise complete marked spaces to complete spaces. It takes dense isometric marked embeddings to dense isometric embeddings. For an admissible specification and isometrically marked X there is a unique natural isometric isomorphism*

$$C(\mathcal{F}X) \;\simeq\; \mathcal{F}(CX) \tag{3.6}$$

extending $\mathcal{F}(X \to CX)$. It is natural for short maps, and for uniformly continuous maps whenever both sides are defined with the relevant separated boundary metrics.

Proof. Let z_n be Cauchy in $(\mathcal{F}X)_s$. The numbers $\mathrm{d}_{\mathcal{F}X}(z_n, K_s)$ converge to a nonnegative number r because distance to a nonempty set is a short function. If $r = 0$, choose $k_n \in K_s$ with $\mathrm{d}(z_n, k_n) = \mathrm{d}(z_n, K_s)$, using (3.5) or $k_n = z_n$. The triangle inequality makes (k_n) Cauchy. Compactness of the seam gives its limit, which is also the limit of z_n.

If $r > 0$, for large m, n the quotient distance $\mathrm{d}(z_n, z_m)$ is smaller than both 1 and $r/2$, whereas each seam radius is larger than $r/2$. The via-seam term cannot give this distance. The direct term in (3.4) forces all terms in a single branch and shows that their representatives are Cauchy in that complete input component. Its limit maps to the desired limit. This proves completeness. For density, each point of the target has a representative $[e, x]$. Approximate x by points of the dense domain in its branch and use the branch Lipschitz bound. The embedding is isometric by Proposition 3.3. The characterization of Cauchy completion by a dense isometry gives (3.6). Naturality follows by uniqueness of continuous extension from a dense subset, using Proposition 3.3 for uniform continuity. □

The completion assertion depends on keeping the marked subset fixed as a compact set. In an extension where marks are merely a dense subset of a boundary, the completion could add new marked points and hence new identifications. Our square and carpet examples take the entire compact outer boundary at stage zero. This supplies the same seam at every later completion step.

3.4 A finite shortest-path description

When all boundaries are finite, (3.4) may be written as a shortest-path formula for the entire quotient metric. Let z, z' be fixed. Introduce vertices s_0, t_0 for their representative sets and a vertex for each seam class. For any two distinct vertices, put an edge of length equal to the minimum paid cost among representatives in a common branch. Omit that edge if there are no such representatives. Write $L(z, z')$ for the shortest such path length, with $+\infty$ for disconnected endpoints. Then

$$\mathrm{d}_{\mathcal{F}X}(z, z') = \min\{1, L(z, z')\}. \tag{3.7}$$

For a chain of cost less than one, grouping adjacent steps with identical meeting representatives gives a walk in this graph. Conversely each edge of a walk has representatives that attain its weight. At a seam vertex they may be joined for zero cost. Since cycles can be deleted and the graph is finite, the shortest finite path is attained. If the metric distance is one, a single paid step attains it. For unequal endpoints any path between them must have a positive edge because an edge of weight zero occurs only between duplicate vertices representing the same quotient point. This also proves separation for arbitrary bounded input metrics with finite injective markings.

When every branch has weight ρ and the seed has the discrete metric, distinct seam classes joined in one branch have weight exactly ρ. The boundary test is thus equivalent to saying that every pair of output classes is distinct and has graph distance at least $1/\rho$ when graph edges count as one. Disconnected classes have infinite graph distance. If $\rho \geq 1/2$, this reduces to the absence of an edge joining two distinct output classes. For smaller ρ, a path of two or more edges may already violate admissibility. The numerical test separates the choices of branch incidence, equality of labels, and scale.

4 The compact completion of the chain

4.1 Path products

Let $r_w = \prod_{e \in w} r_e$, with $r_{\varnothing_s} = 1$, and put

$$\epsilon_n = \max_{s \in S,\ w \in W_{s,n}} r_w.$$

The numbers are positive, at most one, and submultiplicative. The metric on every iterate satisfies

$$\mathrm{d}_{(\mathcal{F}^n X)_s}(p_w^X(x), p_w^X(y)) \ \leq \ r_w \mathrm{d}_{X_{t(w)}}(x, y). \tag{4.1}$$

This is proved by composing the branch estimates at each depth. It holds for the entire iterated gluing metric even when some branch has internal identifications or when the graph admits several paths with the same terminal sort.

Lemma 4.1 (Eventual strict contraction)**.** *For a finite graph with $0 < r_e \leq 1$, the following conditions are equivalent.*

1. *The sequence ϵ_n tends to zero.*
2. *There exists N with $\epsilon_N < 1$.*
3. *Every directed cycle has product strictly less than one.*

If these conditions hold, $\epsilon_n \leq \epsilon_N^{\lfloor n/N \rfloor}$ for such an N.

Proof. Submultiplicativity proves the equivalence of the first two statements and the displayed bound. If there is a cycle with product one, all of its weights are one, and arbitrarily long words following that cycle have weight one. Conversely, a word of length $|S|$ consisting only of weight-one edges must repeat a vertex, hence contains a directed cycle with product one. If no such cycle exists, $\epsilon_{|S|} < 1$. □

This condition permits a finite string of isometrically scaled transitions between sorts before any quantitative contraction occurs. For instance a type with a single edge of weight one to a second type is admissible for the contraction argument when every return to the first type goes through a strictly weighted edge. A one-step estimate using the maximum edge weight would give only nonexpansiveness in that case. Passing to N steps retains the scale of the input components and gives a contraction.

4.2 The stage metrics and the shift

Give $A_n = \mathcal{F}^n B$ the iterative path metric. By Definition 3.1 and Proposition 3.3, i_0 and all $i_n = \mathcal{F}^n i_0$ are isometric embeddings. Each component of A_n is compact, by induction using its presentation as a continuous image of a finite coproduct of compact spaces. There is one metric on the set colimit G for which every ι_n is an isometry. For two points, choose a common stage and use the distance there. Isometric inclusions make this independent of the stage, and a common stage for three points proves the triangle inequality.

Proposition 4.2 (Isometric shift)**.** *The algebra a_G of Proposition 2.2 is an isometric isomorphism for the path metric on $\mathcal{F}G$ and the chain metric on G. Moreover G with this algebra is initial among algebras with short structure maps on isometrically marked, one-bounded metric families.*

Proof. Two elements of $\mathcal{F}G$ have representatives in $\mathcal{F}A_n$ for a common n. The inclusion $\mathcal{F}A_n \to \mathcal{F}G$ is an isometry, and a_G restricts there to ι_{n+1} by (2.3). This proves preservation of all distances. Its inverse already exists on sets. For a short algebra structure, the initial map $B \to X$ is isometric, and all maps u_n in the proof of Proposition 2.2 are short by induction. Their union is short because every pair of points belongs to a common stage. The marked-set uniqueness implies uniqueness in the short category. □

The colimit statement uses the metric from isometric stages, which is appropriate for short maps. Continuous maps that are merely continuous on every stage need not be continuous for the metric colimit. For example an increasing union of finite discrete subsets can have a nondiscrete metric topology. We use the initial algebra in sets and the short category here. The continuous finality result follows separately from closed covers.

4.3 Total boundedness and completion

Fix a point $b^*_s \in B_s$ for each sort. For any $z \in G_s$, since $a_G^{[n]}$ is an isometric isomorphism and $p_{s,n,G}$ is onto, there are $w \in W_{s,n}$ and $x \in G_{t(w)}$ with

$$z = a_G^{[n]} p_w^G(x), \qquad \mathrm{d}_G(z, a_G^{[n]} p_w^G(j_{G,t(w)} b^*_{t(w)})) \le \epsilon_n.$$

The second term belongs to the image of A_n. There are only finitely many centers $a_G^{[n]} p_w^G(jb^*)$ for fixed s, n. This gives total boundedness whenever $\epsilon_n \to 0$. The argument uses one sample per word even if A_n is an infinite compact space. Completion gives a compact metric space $\widehat{G}_s = C(G_s)$, with the same isometric marked boundary.

By Proposition 3.4, the completion of the isometric a_G identifies with a unique isometric isomorphism $a\colon \mathcal{F}\widehat{G} \to \widehat{G}$. Let $\kappa\colon G \to \widehat{G}$ be the dense isometry. The commuting square

$$\begin{array}{ccc} \mathcal{F}G & \xrightarrow{a_G} & G \\ {\scriptstyle \mathcal{F}\kappa}\downarrow & & \downarrow{\scriptstyle \kappa} \\ \mathcal{F}\widehat{G} & \xrightarrow[a]{} & \widehat{G} \end{array} \tag{4.2}$$

characterizes a as a continuous map, since $\mathcal{F}\kappa$ is dense. In the one-sort finite setting this is the same characterization as the completed algebra in [18]. We use $\gamma = a^{-1}$ as the coalgebra structure.

Lemma 4.3 (Complete fixed points)**.** *Suppose P is a complete isometrically marked, one-bounded family and $\alpha\colon \mathcal{F}P \to P$ an isometric isomorphism. Under eventual contraction, the unique algebra map $G \to P$ is a dense isometry, and its completion is an isometric algebra isomorphism $\widehat{G} \to P$.*

Proof. The stage maps $A_n \to P$ are isometries by induction starting with $B \to P$. Thus their union is isometric. For any $p \in P_s$, write $p = \alpha^{[n]} p_w^P(x)$. Replacing x by the mark $b^*_{t(w)}$ gives a

point in the image of A_n within ϵ_n of p. Hence the image is dense. Completion of an isometry with dense image into a complete metric space gives an isometric isomorphism. The algebra identity extends along the dense inclusion $\mathcal{F}G \to \mathcal{F}\widehat{G}$ by continuity. □

The lemma provides a practical identification principle. Given a geometric family with the same metrics on the boundaries, it suffices to establish that the branch quotient map onto it is an isometry. One need not identify the metric on every stage by a separate calculation. In Section 9 we use this for a square and for a finite Markov graph. We also give an explicit intrinsic metric for the carpet for which this approach works.

4.4 Algebraic independence of the weights

The set functor and the topological functor in Section 2 use only (2.1). Thus two admissible weight systems with eventual contraction have the same initial set algebra and the same marked set final coalgebra up to unique isomorphism, once finality is proved. Topological finality further gives a canonical homeomorphism between their completions. This homeomorphism restricts to the identity on G. It is generally sensitive to the chosen metrics when regarded as a Lipschitz map. In the unglued binary example, the distances with weights r and r' are r^ℓ and $(r')^\ell$ on a pair with common prefix length ℓ. The quotient descriptions and finality agree for all $r, r' \in (0,1)$, while the ratio of distances is unbounded in one direction unless $r = r'$.

5 The final coalgebra and its mediating maps

5.1 Finality in sets

For a marked-set coalgebra $c\colon X \to \mathcal{F}X$ let H_X be the set of marked maps $u\colon X \to \widehat{G}$. A map in H_X takes $j_{X,s}(b)$ to $j_{\widehat{G},s}(b)$ and can take every other point of X_s to $j_{\widehat{G},s}(b_s^*)$. This shows $H_X \neq \varnothing$ without assuming any topology on X. Give H_X the metric

$$D(u,v) = \max_{s\in S} \sup_{x\in X_s} \mathrm{d}_{\widehat{G}_s}(u_s(x), v_s(x)). \tag{5.1}$$

It is one-bounded and complete. A Cauchy sequence converges at every x by completeness of $\widehat{G}_s$. Taking limits in the uniform Cauchy estimate shows convergence for D. The prescribed values on marks persist under pointwise limits.

Consider the operator $T_c(u) = a\mathcal{F}(u)c$ on H_X. This uses the set action of $\mathcal{F}$ on u, even though a is an isometric map. For $x \in X_s$, take any representative

$$c_s^{[n]}(x) = p_w^X(y).$$

By Lemma 2.4,

$$(T_c^n u)_s(x) = a_s^{[n]} p_w^{\widehat{G}}(u_{t(w)}(y)).$$

The same representative may be used for v. Combining the isometry $a^{[n]}$ and (4.1) yields

$$D(T_c^n u, T_c^n v) \le \epsilon_n D(u,v). \tag{5.2}$$

The estimate holds without choosing representatives consistently as x varies, since it bounds the same value for every representative. For an N with $\epsilon_N < 1$, Banach's theorem gives a unique

fixed point f of T_c^N. Then $T_c(f)$ is a fixed point of T_c^N because the two operators commute. Thus $T_c(f) = f$, and any fixed point of T_c equals f. The fixed-point equation is equivalent to

$$\gamma f = \mathcal{F}(f)c. \tag{5.3}$$

This proves the marked-set finality in Theorem 1.1. The use of injective markings is essential to the nonemptiness of H_X. If the source collapses two labels whose target labels differ, there is no marked map at all.

Corollary 5.1 (Finite-depth evaluation)**.** *Fix any depth n and write $c_s^{[n]}(x) = p_w^X(y)$. Then $f_s(x)$ belongs to the compact subset*

$$P_w = a_s^{[n]} p_w^{\widehat{G}}(\widehat{G}_{t(w)}) \subseteq \widehat{G}_s \tag{5.4}$$

and its distance to $a_s^{[n]} p_w^{\widehat{G}}(j_{\widehat{G},t(w)} b)$ is at most r_w for every $b \in B_{t(w)}$. In particular such approximations converge uniformly as $n \to \infty$ for arbitrary choices of w, y, b.

Proof. Apply Lemma 2.4 to the fixed point, and use (4.1) and the bound $\mathrm{d}_{\widehat{G}} \leq 1$. Since P_w has diameter at most r_w, the same argument bounds the distance between any two points obtained from the same representative word. A second choice of representative still gives the same $f_s(x)$ and the same ϵ_n estimate. □

An estimate for iteration is $D(T_c^n u, f) \leq \epsilon_n D(u, f)$. Thus the iterative computation can start at any marked map, not just one with constant values on the complement. This also explains why extending a coalgebra map from a small subobject, or from a choice of representatives, does not add choices to the final map.

5.2 Finality for continuous coalgebras

Let c now be a continuous marked coalgebra on a Hausdorff family with closed embedded marks. All $c^{[n]}$ are continuous by Proposition 2.3. For $w \in W_{s,n}$, define

$$U_w = (c_s^{[n]})^{-1}(p_w^X(X_{t(w)})) \subseteq X_s. \tag{5.5}$$

For fixed s, n, these are a finite closed cover by the closedness of the branch images. For every $x \in U_w$, Corollary 5.1 gives $f_s(x) \in P_w$. Given $x \in X_s$, remove the union of those U_w which do not contain x. The complement is an open neighborhood of x. Every y in it is in at least one U_w also containing x. Consequently $\mathrm{d}(f_s(x), f_s(y)) \leq \epsilon_n$. For each $\varepsilon > 0$, choosing n with $\epsilon_n < \varepsilon$ proves continuity. The uniqueness from sets gives finality in $\mathsf{Top}_B^{\mathrm{cl}}$.

This proof uses the existence of a closed cover by possible words at each finite depth. A word chosen at one point does not have to extend to a continuous address near that point. It suffices that a nearby point cannot lie exclusively in pieces not containing the original point. This observation will be used in Section 7, where the finite covers are instead obtained by lifting a probe through an effective epimorphism.

In the metric categories with continuous arrows, Theorem 3.2 identifies the metric functor on its domain with the topological functor after forgetting distances. Hence a continuous metric coalgebra gives a closed-marked topological coalgebra and its unique map is continuous. This applies both to the isometric-boundary category and to the original finite non-degenerate metric category, in which the marked distances need not equal the seed distances.

5.3 Shortness and quantitative stability

Assume that X has isometric markings and c is short. The subset of H_X consisting of short maps is closed for D. Nonemptiness of that subset has not been assumed. To see that the final map is short, compare representatives of $c^{[n]}(x)$ and $c^{[n]}(y)$ in $\mathcal{F}^n X$. In each chosen word replace the tail by $j_X(b^*)$. This moves each point by at most ϵ_n. The resulting pair belongs to the isometrically embedded $\mathcal{F}^n B$. Its corresponding images in $\widehat{G}$ have the same distance, because the inclusions of the initial chain and its completion are isometric. By Corollary 5.1, those images are each within ϵ_n of the final values. Hence

$$\mathrm{d}_{\widehat{G}}(f(x), f(y)) \le 4\epsilon_n + \mathrm{d}_{\mathcal{F}^n X}(c^{[n]}(x), c^{[n]}(y)) \le 4\epsilon_n + \mathrm{d}_X(x, y). \tag{5.6}$$

Letting $n \to \infty$ proves shortness and the remaining assertion of Theorem 1.1. This argument does not require any extension theorem for short maps from a boundary to the whole domain. It is important, for example, for a boundary embedded in a metric space that admits no short retraction onto it.

More generally the same reasoning compares an induced boundary metric with the canonical one. Write B_X for B with the metric pulled back from X. If the iterated metrics are separated and the identity $B_X \to B$ is componentwise L-Lipschitz for $L \ge 1$, then $\mathcal{F}^n B_X \to \mathcal{F}^n B$ is L-Lipschitz and $\mathcal{F}^n B_X \to \mathcal{F}^n X$ is isometric. Thus

$$\mathrm{d}_{\widehat{G}}(f(x), f(y)) \;\le\; 2(1+L)\epsilon_n + L\mathrm{d}_{\mathcal{F}^n X}(c^{[n]}(x), c^{[n]}(y)). \tag{5.7}$$

The two contributions $2\epsilon_n$ and $2L\epsilon_n$ come respectively from replacing the tails at the target and from replacing them at the domain. For arbitrary finite markings the separation hypothesis is automatic at all depths. For compact boundaries, a bi-Lipschitz comparison between B_X and B suffices, since the corresponding comparisons of chain costs remain bi-Lipschitz with the same bounds at every iteration. This follows before passing to the quotient by applying the Lipschitz inequalities to every paid step.

If c is short, (5.7) gives the bound L on f. If c is K-Lipschitz for $K > 1$, then $c^{[n]}$ is K^n-Lipschitz by Proposition 3.3. Together with the exponential bound of Lemma 4.1, this gives a Hölder estimate with exponent $\log(1/\theta)/\log(K/\theta)$ whenever $\epsilon_n \le C\theta^n$ with $0 < \theta < 1$. A precise constant is recorded in Appendix B. In the original specialization $\epsilon_n = \rho^n$ and $L = 1/\Delta(X)$, so (5.7) is the source estimate and its short-map consequence.

5.4 Changing the scales

Suppose the same topological data admit two choices of weights or of metrics on the boundaries, and both satisfy admissibility and eventual contraction. The topological gluing functor is the same in both constructions, so their completed coalgebras admit a unique isomorphism in $\mathsf{Top}^{\mathrm{cl}}_B$. Its underlying marked map agrees on G with the initial-algebra inclusions. Indeed, as an isomorphism of coalgebras with invertible structure it is also an algebra map. Its restriction to G is then the unique initial-algebra map into the second completion. This proves the independence assertion of Section 4 without identifying different path metrics.

6 Graph-directed attractors and exact overlaps

6.1 Compatible geometric realizations

For each sort take a nonempty complete one-bounded metric space Z_s and a continuous embedding $\alpha_s\colon B_s \to Z_s$. For $e \in E_s$, take a Lipschitz map $\phi_e\colon Z_{t(e)} \to Z_s$ with bound $0 < \lambda_e \le 1$. Write $\lambda_w = \prod_{e\in w} \lambda_e$ and $\eta_n = \max_{s,w\in W_{s,n}} \lambda_w$. We assume $\eta_n \to 0$. The weights r_e controlling the chain metric may differ from λ_e. Compatibility means there are maps $\beta_s\colon K_s \to Z_s$ satisfying

$$\beta_s q_s(e,b) = \phi_e \alpha_{t(e)}(b), \qquad \beta_s b_s = \alpha_s. \tag{6.1}$$

The first equation gives uniqueness and continuity of β_s because q_s is a quotient. It allows β_s to identify more seam points than the specification does. In the finite fixed-point specification of [18], K_s is already the image of the boundary and β_s is injective.

For a family $C = (C_s)$ of nonempty compact subsets of Z_s, set

$$(\mathcal{H}C)_s = \bigcup_{e\in E_s} \phi_e(C_{t(e)}).$$

On the product of nonempty compact hyperspaces use the maximum of the Hausdorff distances. Each hyperspace of a complete metric space is complete. One way to see this is to select a subsequence with summably decreasing Hausdorff gaps, approximate a point of one compact set by a sequence of points in the later sets, and use completeness of Z_s. The set of limits is closed, and has finite nets because it stays within arbitrarily small Hausdorff distance of one compact set in the subsequence. This yields compactness and convergence of the original Cauchy sequence. The finite product is also complete. Since

$$(\mathcal{H}^n C)_s = \bigcup_{w\in W_{s,n}} \phi_w(C_{t(w)}), \qquad d_H(\mathcal{H}^n C, \mathcal{H}^n C') \le \eta_n d_H(C, C'),$$

Banach's theorem for a strict iterate gives a unique fixed family A of $\mathcal{H}$. As in Section 5, the fixed point of the iterate is fixed under $\mathcal{H}$. This is the graph-directed compact attractor.

Let $O_0 = (\alpha_s B_s)$, $O_n = \mathcal{H}^n O_0$, and $O_s = \bigcup_n O_{n,s}$. Compatibility gives $O_0 \subseteq \mathcal{H}O_0$ since q_s is surjective. The stages are therefore nested. They converge to A in Hausdorff distance by the contraction estimate. Each O_n lies in A, because it lies in all later stages and their distance from A tends to zero. Conversely every point of A_s is a limit of points in the $O_{n,s}$. It follows that O is a dense marked orbit. When O_0 consists of the individual branch fixed points this is exactly the orbit used in the classical single-sort argument.

For $Q = O_n, O, A$, the branch equations define the corresponding algebra or stage map by

$$\Phi([e,x]) = \phi_e(x), \qquad \Phi(k) = \beta_s(k) \quad (k \in K_s). \tag{6.2}$$

On O_n the map has codomain O_{n+1}, and on O, A it has codomain O, A, respectively. It is always continuous in the topological gluing presentation and surjective. All the marks are embeddings since they are induced by α_s. In metric categories where the orbit and attractor gluing path metrics are separated, Φ is L_Φ-Lipschitz for

$$L_\Phi = \max(\{1\} \cup \{\lambda_e/r_e : e \in \coprod_s E_s\}). \tag{6.3}$$

Indeed a step within branch e has scaled cost r_ed and image distance at most λ_ed. A cross-branch step costs one and the ambient distances are bounded by one. Zero-cost jumps preserve

the image. Apply the triangle inequality and take the infimum. Thus Φ is short when every $\lambda_e \le r_e$. Orbit boundary metrics may be different from the canonical ones. In Theorem 1.2 the comparison of the orbit and attractor as topological algebras does not require separation of their own gluing path metrics.

6.2 Initial stages and orbit injectivity

Let $h_0 \colon G \to O$ be the unique marked-set algebra map. Its restriction to $A_n = \mathcal{F}^n B$ maps onto O_n by the stage recursion in Proposition 2.2. We keep the notation A without a subscript for the geometric attractor. The use of A_n for the canonical stages is local to the chain and independent of the geometric space.

Proposition 6.1 (Injectivity at the orbit)**.** *The following conditions are equivalent.*

1. *Every stage map $\mathcal{F}O_n \to O_{n+1}$ of (6.2) is injective.*
2. *The orbit algebra $\Phi_O \colon \mathcal{F}O \to O$ is injective.*
3. *$h_0 \colon G \to O$ is bijective.*

When they hold, each canonical stage is identified with its geometric orbit stage by a marked bijection compatible with the next stage map.

Proof. The stage inclusions and their images under $\mathcal{F}$ are injective. Any two representatives in $\mathcal{F}O$ are represented in a common stage, and the same Φ formula is used there. This proves equivalence of the first two conditions. Assuming the first, the stage comparisons are bijections by induction, starting with $\alpha_s \colon B_s \to O_{0,s}$. Their union is h_0. Conversely, for bijective h_0, the algebra identity expresses $\Phi_O = h_0 a_G (\mathcal{F}h_0)^{-1}$, which is bijective. □

This proposition tests the relation that is generated on the orbit. In general the attractor can contain coincidences that are limits of orbit points but are not witnessed by equality between two orbit points. The next subsection gives a criterion for the completed comparison.

6.3 Address spaces and their equivalence relations

For a sort s, the space $\mathcal{W}_s$ of infinite composable words is the inverse limit of the finite sets $W_{s,n}$ under prefix restriction. Give it the compact subspace topology from the product of edge alphabets. It is nonempty because every sort has an outgoing branch. For $w \in W_{s,n}$, the compact subset P_w of (5.4) is the image of the n-fold composition of the branch maps $p_e = a_s p_e^{\widehat{G}}$ into $\widehat{G}_s$. These images are nested along prefix extensions and their diameters are at most r_w. Thus a word ω determines exactly one point in their intersection. This gives a map

$$\pi_s \colon \mathcal{W}_s \to \widehat{G}_s.$$

It is continuous because the diameter of any prefix piece of depth n is at most ϵ_n. It is surjective because at each step the finitely many pieces cover the parent, and an address for a point is obtained by continuing through a branch preimage. This argument allows a branch to be noninjective. The resulting address map is a quotient map by compactness and Hausdorffness. Geometric branch maps and bounds η_n give in exactly the same way a continuous surjection

$$\pi_s^A \colon \mathcal{W}_s \to A_s.$$

The formulas respect prefixing by any edge for which the type matches. In particular the two address maps have the same domain, regardless of the ambient metrics or the behavior of their inverse branches.

Lemma 6.2 (Addresses of marks)**.** *If $\pi_s(\omega) = j_{\widehat{G},s}(b)$ for a mark b, then $\pi_s^A(\omega) = \alpha_s(b)$.*

Proof. Write $\omega = e\omega'$. A preimage of $j_{\widehat{G},s}(b)$ under the branch p_e must be marked, by Lemma 2.1 applied to $\gamma = a^{-1}$ and the fact that $j_{\widehat{G},s}(b)$ corresponds to $b_s(b) \in K_s$. Hence $\pi_{t(e)}(\omega') = j_{\widehat{G},t(e)}(b')$ for some b' with $q_s(e, b') = b_s(b)$. By (6.1), $\alpha_s(b) = \phi_e \alpha_{t(e)}(b')$. Repeat along the address. For every prefix w of length n, this expresses $\alpha_s(b)$ as ϕ_w of a geometric mark. Thus $\alpha_s(b)$ belongs to $\phi_w(A_{t(w)})$ for every such prefix. Their intersection is the single point $\pi_s^A(\omega)$. □

Proposition 6.3 (Canonical quotient to the attractor)**.** *The implication*

$$\pi_s(\omega) = \pi_s(\omega') \quad \Longrightarrow \quad \pi_s^A(\omega) = \pi_s^A(\omega')$$

holds for each sort. Consequently $\pi^A = h\pi$ for a unique continuous surjection $h\colon \widehat{G} \to A$. This is a marked algebra map and extends $h_0\colon G \to O$ after the inclusion $O \subseteq A$.

Proof. If the addresses agree, the implication is immediate. Otherwise compare their first edges and the canonical values of their tails. If the edges and the canonical tail values agree, discard that edge and continue. Because the addresses differ at a finite position, this process reaches a depth where the edges differ or the canonical tail values differ. At that depth there are two distinct representatives of the same point in the pushout. By Lemma 2.1, both tails are marked and the two marked representatives have the same q-value. By Lemma 6.2, their geometric tail values are the corresponding α-marks. By (6.1) the geometric branch images coincide. Applying the common preceding prefix proves the implication.

Since π is a quotient map and π^A is continuous and surjective, the factor h is continuous and surjective. Marks are preserved by Lemma 6.2. Every point has an address, so prefixing it proves that $hp_e = \phi_e h$ for each branch. These identities determine the algebra equation by surjectivity of the branch quotient. Its restriction to G is the unique initial-algebra map into A, which is the inclusion of O composed with h_0. □

The argument identifies the two kinds of equality of addresses. The canonical kernel is generated in the following explicit sense. When distinct addresses evaluate to the same point, they are obtained by a common prefix followed by two different representatives from the same prescribed boundary class, where each tail is an address of its specified mark. Conversely every such pair evaluates to the same point. An address of a mark need not be unique, and the possible equalities at the marked tails are built into this description. The same statement is valid for infinite compact boundary spaces.

6.4 The completed overlap criterion

Theorem 6.4 (Exact geometric overlaps)**.** *The following are equivalent for a compatible geometric realization.*

1. *The comparison $h\colon \widehat{G} \to A$ is a homeomorphism.*
2. *The geometric gluing map $\Phi_A\colon \mathcal{F}A \to A$ is injective.*

3. *For every sort s, the equivalence relations* $\ker(\pi_s)$ *and* $\ker(\pi_s^A)$ *coincide.*

4. *The orbit algebra is injective and* $h_0^{-1}\colon O \to G$ *is uniformly continuous for the inherited metrics.*

When these conditions hold, Φ_A^{-1} is a final coalgebra in the marked set and closed-marked topological categories. If h_0 is bi-Lipschitz, the homeomorphism h is bi-Lipschitz with the same comparison constants.

Proof. The first condition implies the second by the algebra equation and the isomorphism a. If Φ_A is injective, it is a continuous bijection from a compact space to a Hausdorff space in each sort, and hence is a homeomorphism. Its inverse defines a coalgebra. Let $g\colon A \to \widehat{G}$ be its unique map to the final coalgebra. The algebra identity for h makes h a coalgebra map since both structure maps are invertible. Therefore $gh = \mathrm{id}$, by finality on $\widehat{G}$. Since h is onto it is an isomorphism. This proves equivalence of the first two conditions and the transported finality. The third condition is equivalent to injectivity of h because π is onto and $\pi^A = h\pi$.

Assume the first condition. Then h_0 is a bijection by restriction to G and $O = h_0(G)$, and the inverse of h is uniformly continuous on the compact attractor. This gives the fourth condition using Proposition 6.1. Conversely, under the fourth condition, the inverse of the orbit comparison extends uniquely to $g\colon A \to \widehat{G}$ by uniform continuity and completeness. Its composite with h equals the identity on the dense subset G, hence equals the identity everywhere. This gives the first condition. If the two inequalities on the orbit comparison are Lipschitz inequalities, they pass to limits along sequences from G and O and give the same inequalities on the completions. □

In a metric setting where $C(\mathcal{F}O) \cong \mathcal{F}C(O)$, the completed orbit algebra of [18] is the algebra Φ_A here. To see this, apply (6.2) on the dense image of $\mathcal{F}O$ in $\mathcal{F}A$ and use continuity. In the finite point case the necessary preservation of dense isometries follows from Proposition 3.4. This identifies the actual completed maps, in addition to identifying their universal properties.

6.5 Bounds and failure of finite orbit detection

If the boundary comparison $B \to O_0$ is L_0-Lipschitz and $\lambda_e \le r_e$ for all e, then each stage map $\mathcal{F}^n B \to O_n$ is $L = \max(1, L_0)$-Lipschitz. This follows from (6.3) and preservation of the Lipschitz bound by gluing. Passing to the chain and then the completion gives the same bound for h_0 and h. A lower bound on h_0 follows from a uniform estimate for the inverses of the iterated orbit maps, provided the inverse boundary comparison is Lipschitz. We give the formula in Appendix B and recover the constants of the source paper in Section 8.

The topology of the orbit comparison also has a one-step criterion. Under orbit injectivity, if Φ_O^{-1} is continuous then O is a continuous coalgebra and its map to $\widehat{G}$ is the inclusion of G composed with h_0^{-1}. Since that inclusion is a topological embedding, h_0^{-1} is continuous. Conversely, h_0 is continuous by Proposition 6.3, and if h_0^{-1} is continuous then $\Phi_O^{-1} = \mathcal{F}(h_0)a_G^{-1}h_0^{-1}$ is continuous. This is an assertion about the orbit topology. Uniform continuity of the inverse is the condition for extending it to the completion.

For a strict example of the difference, take a single mark 0, two branches with no gluing, and weight $1/2$ on each branch. The distinguished output is the first branch applied to the mark. Realize the branches on the interval by $\phi_0(x) = x/2$ and $\phi_1(x) = (1 + x)/2$. The orbit is the dyadic rationals in $[0, 1)$. The two orbit images lie on opposite sides of $1/2$, with only the second

containing $1/2$. Thus the orbit algebra is injective. The attractor is the full interval, and both branches now contain $1/2$. Its two codes $1000\ldots$ and $0111\ldots$ lie in disjoint first-branch pieces of the canonical unglued binary space. Thus the comparison collapses them. This example uses a compatible marked seed, with only one branch fixed point chosen. It shows why a bijective comparison of initial algebras alone cannot establish a homeomorphism of completions.

7 Final coalgebras in a condensed ∞-topos

7.1 The site and the marked endofunctor

Let $\mathsf{CompHaus}$ be the category of compact Hausdorff spaces in the chosen universe. Equip it with the Grothendieck topology generated by finite jointly surjective families. The empty family covers the empty space. These families are stable under pullback and composition. We write

$$\mathsf{Cond}_\infty = \mathrm{Shv}_{\mathcal{S}}(\mathsf{CompHaus}), \qquad y(L)(U) = \mathrm{Cont}(U, L)$$

where the displayed section sets are viewed as discrete spaces. The topology is subcanonical. Indeed, a continuous function defined on a finite coproduct of compact spaces and constant on the fibers of a surjection descends continuously to its Hausdorff target. This gives the equalizer condition for the representable, and the higher descent condition for a set-valued presheaf reduces to that condition. In particular y is fully faithful and preserves finite limits. Finite disjoint unions of compact spaces become coproducts of their representables as sheaves. To see this, the summands give a disjoint covering family, and a plot into their topological disjoint union is locally a plot into one summand. A surjection of compact Hausdorff spaces induces an effective epimorphism of representables, by pulling back its cover along each plot. This is the compact-cover presentation of condensed objects [5].

For an object $T \in \mathsf{Cond}_\infty$, take all the coefficients yB_s, yK_s and maps from (2.1) in the slice $\mathsf{Cond}_{\infty/T}$ by taking their products with T. Let $\mathcal{C}_T$ denote the full subcategory of tuples under $T \times yB_s$ in the slice whose markings are monomorphisms. In this section $\mathcal{F}_T$ is the homotopy pushout functor with these coefficients:

$$(\mathcal{F}_T X)_s = (T \times yK_s) \sqcup_{\coprod_{e \in E_s} (T \times yB_{t(e)})} \coprod_{e \in E_s} X_{t(e)}. \tag{7.1}$$

All coproducts and the pushout are in the slice. The output marking is induced by b_s. In an ∞-topos, a pushout of a monomorphism is a monomorphism. For this fact one can compute first in presheaves, where a mono of spaces is the inclusion of a union of components up to equivalence, and then apply left-exact sheafification. The maps from $T \times yK_s$ into (7.1) are therefore mono. The output markings are also mono. Products with a fixed base commute with the colimits used here, by universality of colimits in an ∞-topos [12, Sec. 6.1].

The branch family $\coprod_e X_{t(e)} \to (\mathcal{F}_T X)_s$ is an effective epimorphism. It is obtained by pushing out the effective epimorphism of boundary coproduct onto $T \times yK_s$, and effective epimorphisms are stable under pushout. For instance their (-1)-truncated image factorization, or the Čech nerve characterization together with universality, gives this stability. Since finite coproducts and composites preserve effective epimorphisms, the length-n branch family

$$\coprod_{w \in W_{s,n}} X_{t(w)} \longrightarrow (\mathcal{F}_T^n X)_s \tag{7.2}$$

is again an effective epimorphism. Its maps are the iterates of the branch maps of (7.1). For any morphism f, the counterpart of Lemma 2.4 is a commuting diagram up to its specified natural homotopy.

Lemma 7.1 (Local branch lifts)**.** *Let $x\colon y(U) \to (\mathcal{F}_T^n X)_s$ be a plot over some map $y(U) \to T$. There is a finite jointly surjective family $U_i \to U$, a word w_i for each i, and a lift $y(U_i) \to X_{t(w_i)}$ whose image along the branch agrees with $x|_{U_i}$ up to homotopy in the slice.*

Proof. An effective epimorphism of sheaves of spaces is locally surjective on components. To pass from this assertion to a lift of the particular plot, first use epimorphy on 0-truncated sheaves. A section of the 0-truncation is locally represented by a section of the original sheaf, because truncation of a sheaf is obtained by sheafifying presheaf truncation. Equality of the resulting components is also witnessed locally. Applying this to (7.2) gives local lifts to the coproduct with paths over the original section. Refine once more to obtain a section of one summand, since a coproduct sheaf is locally represented by sections of its summands. Equivalently, apply local surjectivity to the jointly epimorphic family of the individual branch maps. All refinements can be taken in the generating pretopology, whose finite covering families compose. The lifts take their maps to T from the pulled-back plot. This is the slice formulation of the local lifting property. □

This lemma does not require the ∞-topos to have enough points. It constructs finite observations of each plot. The distinction is important because the singleton probes used below need not detect arbitrary morphisms of condensed objects. They will detect maps with the specific representable targets of this paper.

7.2 Evaluation at the singleton

Write 1 for the singleton compact Hausdorff space. Define

$$\mathrm{ev}_1\colon \mathsf{Cond}_\infty \longrightarrow \mathcal{S}, \qquad X \longmapsto X(1).$$

As evaluation, it preserves limits. It also preserves colimits. Here is a useful proof of the latter fact that applies to higher sheaves. For a space V, the presheaf

$$D(V)(U) = \mathrm{Map}_{\mathcal{S}}\big((|U|)^{\mathrm{disc}}, V\big)$$

is the right Kan extension from the full singleton subcategory of probes. It is a sheaf. The underlying-set map of a jointly surjective finite family is surjective, and the realization of its Čech nerve in spaces is the discrete set $|U|$. Applying $\mathrm{Map}(-, V)$ turns this into the descent limit. Thus D is right adjoint to ev_1 on sheaves, proving the colimit assertion. In particular $\mathrm{ev}_1(yL) = |L|^{\mathrm{disc}}$, and ev_1 carries $\mathcal{F}_T$ to the analogous pushout functor over $T(1)$ in spaces.

Lemma 7.2 (The set comparison inside spaces)**.** *For discrete coefficients B_s, K_s as above, the discrete space on each $\widehat{G}_s$ is the underlying object of the final coalgebra for the homotopy pushout gluing functor on monically marked spaces. In the slice over any space V, the constant family with these fibers is final for the fiberwise functor.*

Proof. For a mono $B_s^{\mathrm{disc}} \to X_s$, the map $B_s \to \pi_0 X_s$ is injective. The 0-truncation of the homotopy pushout is the pushout of the sets of components, because 0-truncation is left adjoint

to the discrete-space inclusion. Conversely the homotopy pushout of the discrete marked-set span is the ordinary pushout as a discrete space. Indeed the leg from the boundary coproduct to the branch coproduct is injective. The homotopy pushout thus splits off the unmarked discrete summands and replaces the marked ones by the seam. This identifies the action and the markings on all discrete inputs.

Put $Y = \widehat{G}^{\mathrm{disc}}$. The mapping spaces $\mathrm{Map}_{\mathcal{C}}(X, Y)$ and $\mathrm{Map}_{\mathcal{C}}(X, \mathcal{F}Y)$ in the marked category are 0-truncated. Their points are the marked functions on $\pi_0 X_s$ into the relevant sets. For a coalgebra c on X, the mapping space to the candidate coalgebra is the homotopy equalizer of the two maps

$$\mathrm{Map}_{\mathcal{C}}(X, Y) \;\rightrightarrows\; \mathrm{Map}_{\mathcal{C}}(X, \mathcal{F}Y), \qquad f \longmapsto \gamma f, \quad f \longmapsto \mathcal{F}(f)c. \tag{7.3}$$

This description is the mapping-space formula in the lax equalizer defining $\mathsf{Coalg}(\mathcal{F})$: its points are marked maps with a path between the two composites. The two spaces in (7.3) are discrete, so the homotopy equalizer is the set of coalgebra maps on components as a discrete space. By set finality this space is contractible. In a slice $\mathcal{S}_{/V} \simeq \mathrm{Fun}(V, \mathcal{S})$, coefficients are constant and the operations are pointwise. A coalgebra is thus a V-indexed family of coalgebras with coherent transports. Mapping to the constant terminal coalgebra is pointwise terminal, hence terminal in the functor category. □

Markings at different points of a base space can be related by transport. The slice assertion uses monomorphisms in the slice, which are checked on the fibers of a map of spaces over the base. It does not choose representatives of components of V or assign arbitrary transport maps.

7.3 The completed space represents the higher fixed point

Proposition 7.3 (Representable fixed point)**.** *Write $Y_s = y(\widehat{G}_s)$. In the absolute case $T = 1$, the maps from the seam and branches of the compact algebra induce an equivalence*

$$\mathcal{F}_1 Y \;\simeq\; Y.$$

After multiplying by any T, this identifies $\mathcal{F}_T(T \times Y)$ with $T \times Y$ over T.

Proof. All the objects in the span for Y are 0-truncated and one leg is mono. Its homotopy pushout is 0-truncated. One can first compute the pushout of sets objectwise in presheaves. There the homotopy pushout is a set because one of its legs is injective, and sheafification preserves truncated objects by left exactness. Let $\xi \colon \mathcal{F}_1 Y \to Y$ be the resulting map of sheaves of sets.

The branch maps to $\widehat{G}_s$ define a jointly surjective finite family of continuous maps of compact Hausdorff spaces. Given a continuous map $u \colon U \to \widehat{G}_s$, pull the branch maps back along u. Each pullback is compact Hausdorff, and the resulting maps to U jointly cover it. On each of these probes u factors through a branch. This proves local surjectivity of ξ_s.

For injectivity, suppose two sections of $(\mathcal{F}_1 Y)_s(U)$ have the same image. Refine by a finite cover so both are represented by branch plots, using Lemma 7.1. On one such probe U', write their representatives as (e, x) and (f, x'), where x, x' are continuous maps into the appropriate compact components. By Lemma 2.1, at each point of U' the representatives are either equal, or both are marked and have the same q_s-image. The first condition defines a closed subset (empty when $e \neq f$). The second also defines a closed subset, since the markings are closed embeddings

and equality in K_s is closed. Together these subsets are a finite compact cover of U'. On the first, the representatives agree as plots. On the second, both plots factor continuously through their marked boundaries, and their q_s-images agree as plots into the seam. They represent the same section on each part. By descent the original sections agree. Thus ξ_s is injective and locally surjective, so is an isomorphism of sheaves of sets. Taking products with T proves the slice statement. □

The proof checks the pushout itself. For compact spaces a pushout along a compact embedded boundary is still a closed quotient, but continuity of the quotient map alone would not establish the sheaf pushout property. The finite closed cover in this proof is what turns the pointwise equality criterion into equality of sections.

7.4 Coinduction on compact plots

Theorem 7.4 (Relative higher finality)**.** *Let T be any object of Cond_∞. With the inverse of the algebra in Proposition 7.3, the tuple $T \times y\widehat{G}$ is a final $\mathcal{F}_T$-coalgebra in $\mathcal{C}_T$.*

Proof. Write $Y_T = T \times y\widehat{G}$ and let $c\colon X \to \mathcal{F}_T X$ be a coalgebra. Evaluate at the singleton. By colimit and limit preservation, this is a coalgebra in monically marked spaces over $T(1)$ for the functor of Lemma 7.2. There is a coalgebra map $f_*\colon X(1) \to Y_T(1)$ over $T(1)$, unique up to a contractible choice. The second projection of $f_{*,s}$ is a map into the discrete space $|\widehat{G}_s|^{\mathrm{disc}}$. In particular it assigns a value to each singleton plot, constant on paths.

Given a plot $x \in X_s(U)$, define the function $f_x\colon |U| \to |\widehat{G}_s|$ by evaluating $f_{*,s}$ at the restrictions of x to singletons and taking its second projection. This assignment respects paths of plots and is natural in U. We show that f_x is continuous. Fix n and apply Lemma 7.1 to $c_s^{[n]}x$. Obtain maps $p_i\colon U_i \to U$ and lifts along words $w_i \in W_{s,n}$. On evaluation at any point of U_i the coalgebra equation for f_*, iterated n times, implies that

$$f_x(p_i(u)) \;\in\; P_{w_i} \qquad (u \in U_i). \tag{7.4}$$

Indeed after evaluation the branch representation maps under the pointwise solution into the corresponding iterated branch of $\widehat{G}$. A path witnessing the lift has no effect on its value in the discrete target. This applies over each point of the base $T(1)$, with the same bound $\operatorname{diam} P_{w_i} \le \epsilon_n$.

Each $p_i(U_i)$ is closed, and these sets cover U. Around any $u \in U$, exclude the finite union of those not containing u. Every point v of the remaining open neighborhood belongs to an image $p_i(U_i)$ that also contains u. By (7.4) its f_x-value is within ϵ_n of $f_x(u)$. Taking n arbitrarily large proves continuity. We obtain a natural transformation $X_s \to y\widehat{G}_s$ by $x \mapsto f_x$. Since the target is 0-truncated, the map on components of section spaces suffices to define this transformation. Combining with the given $X_s \to T$ defines $f\colon X \to Y_T$ in the slice.

For any compact L, two maps into yL agreeing on singleton evaluation agree on all plots. In fact the two images of a plot $x \in X_s(U)$ are continuous functions $U \to L$ and agree at every point by naturality. Thus maps into Y_T, and into $\mathcal{F}_T Y_T \simeq Y_T$, are determined in the slice by their evaluations at 1. The same holds with marks, since those conditions are equations in spaces of maps into relatively 0-truncated targets. It follows that f preserves the marks and satisfies the coalgebra equation, since f_* does. Conversely every coalgebra morphism must evaluate to f_* and is therefore f. Both mapping spaces in the slice version of (7.3) are 0-truncated, since maps over T into $T \times y\widehat{G}_s$ are maps into the 0-truncated sheaf $y\widehat{G}_s$. The same applies to $\mathcal{F}_T Y_T$ by Proposition 7.3. The homotopy equalizer with exactly one point is contractible. □

This proves the relative part of Theorem 1.3. At $T = 1$, evaluation on discrete marked inputs identifies the structure and every mediating map with the ones in Section 5. On an arbitrary probe, the same map is required to vary continuously in the probe variable. This is stronger than putting the discrete marked-set final object into a constant presheaf. Even when $\widehat{G}_s = [0, 1]$, $y\widehat{G}_s$ contains the identity plot of the interval, which would be locally constant only on a discrete domain.

7.5 An effective presentation by addresses

The address maps themselves give a second presentation in the condensed ∞-topos. This presentation records the overlap relation as a groupoid object and also relates the canonical and geometric quotients.

Proposition 7.5 (Address groupoids)**.** *For each sort s, put $R_s = \mathcal{W}_s \times_{\widehat{G}_s} \mathcal{W}_s$, the compact kernel pair of π_s, with the subspace topology. The groupoid object with objects $y\mathcal{W}_s$ and arrows yR_s is effective in Cond_∞. Its realization is $y\widehat{G}_s$. For a compatible geometric attractor, $R_s^A = \mathcal{W}_s \times_{A_s} \mathcal{W}_s$ similarly presents yA_s, and the inclusion $R_s \to R_s^A$ induces the map yh_s on realizations. Both constructions and the comparison commute with product base change and with the descent of equivariant specifications.*

Proof. The address map π_s is a surjection of compact Hausdorff spaces. Therefore $y\pi_s$ is an effective epimorphism by the compact-cover topology. Its Čech nerve in degree n is the $(n + 1)$-fold fiber product of $y\mathcal{W}_s$ over $y\widehat{G}_s$. It is represented by the corresponding compact fiber product because Yoneda preserves limits. The arrow space in degree one is yR_s, and its composition is composition of pairs with a matching intermediate address. The realization is $y\widehat{G}_s$ by effectivity. Apply the same argument to the surjection π_s^A. By Proposition 6.3, the first fiber products map into the second ones by the identity on address coordinates. The induced map of Čech realizations is yh_s, since the realization maps commute with the epimorphisms from $y\mathcal{W}_s$. Pullback preserves colimits and finite limits in an ∞-topos, so the same description holds after base change. Equivariant comparison maps commute with the descent equivalences at every degree, and slice descent identifies their realizations with the descended quotients. □

The Čech nerve keeps each equality of address values as a relation between the corresponding lifts, with all its composable strings. Although its realization is 0-truncated in the absolute case, it is a homotopy colimit. A comparison between two such realizations is an equivalence exactly when the kernel pairs agree. One direction follows from the equality criterion of Theorem 6.4. For the other direction, an equivalence of the realizations induces equivalences on the kernel pairs by pullback. Since the kernel pairs here are representables of closed subsets of the same compact space, this is equivalent to equality of their subsets. Thus the exact-overlap criterion can be stated equivalently at the level of groupoid objects, where it requires no metric estimate on the inverse.

For a direct example, in the glued interval R contains the two directions of the relation between the binary streams $0111\ldots$ and $1000\ldots$, and all the pairs formed from these by prepending a common finite prefix. The diagonal is also present. The relation is closed. In particular a sequence of distinct pairs with increasingly long common prefixes can converge to a diagonal pair, which is already in the relation. The individual quotient map need not admit a continuous section. A continuous section of the binary evaluation would embed the interval into

the totally disconnected address space, which is impossible. It is the effective quotient and its local compact lifts, rather than the choice of a global lift, that give the presentation.

For a second comparison, the compact-seam gluing construction on arbitrary compact marked Hausdorff inputs X is preserved by y:

$$y\big((\mathcal{F}X)_s\big) \simeq yK_s \sqcup_{\coprod_{e\in E_s} yB_{t(e)}} \coprod_{e\in E_s} yX_{t(e)}.$$

Indeed $\mathcal{F}X$ is compact by Proposition 2.3, and the proof of Proposition 7.3 only uses the finite compact branch cover and the representative equality criterion, both valid for X. The same proof is natural in marked maps. By full faithfulness of y, this identifies the compact Hausdorff coalgebras with the full subcategory of condensed coalgebras whose carriers are compact representables with the designated markings. In particular the closed topology of a branch overlap is retained in this embedding. Evaluation at the singleton then forgets that topology and returns precisely the marked-set construction. These two comparisons explain why the same compact completion can be tested either on ordinary topological inputs or on higher sheaves with additional section and descent data.

7.6 Descent and monodromy

Let $U \to T$ be an effective epimorphism, with Čech nerve $U_\bullet$. Suppose the compact-seam coefficients are given in $\mathsf{Cond}_{\infty/U}$, are constant compact data there (or are so on a further effective cover), and have descent equivalences in the Čech nerve. We allow relabeling of sorts and branches by coherent permutations. We require the descent equivalences to identify the markings, seam maps, and substitutions, so that they induce equivalences of the marked categories and their gluing functors. Isometric automorphisms of a compact-seam specification give such equivalences. The rate condition is checked on each local specification. The coefficients can be viewed globally as a finite family indexed by the descended sort and branch objects.

By descent in an ∞-topos, the slice over T is the limit of the slices over the Čech nerve. The same conclusion holds for the categories of marked families with the stated descent data. One applies slice descent to the carriers and markings, and tests monicity after the effective cover. Homotopy pushouts commute with these pullbacks, so the coalgebra categories with this data are also obtained by descent. More explicitly, both the carriers and their structure arrows are objects of diagram categories in the slices, and the condition that a structure arrow has codomain the specified pushout is stable under pullback. Thus maps of coalgebras, including the homotopies in (7.3), are recovered as limits of their restrictions.

On each U_n the base changes of the local completed coalgebras are final by Theorem 7.4. The spaces of equivalences between two final objects in each such category are contractible. Their restrictions along the nerve therefore give coherent descent data: the spaces of choices of compatible equivalences and higher compatibilities are limits of contractible spaces. The descended object is final because its mapping spaces from any descended coalgebra are limits of contractible mapping spaces. This proves the descent assertion of Theorem 1.3. It also gives compatibility with further base change for these locally constant specifications, since the same argument applies to the pulled-back cover.

For a concrete instance, let a finite group H act by automorphisms of the specification. Take the action on the initial chain and on its metric completion. The equivariant coalgebra descends to the classifying object BH via the cover $1 \to BH$. For the interval, take $H = C_2$, acting by

reflection, interchanging endpoints and the two branches and fixing the central seam. The total object over BC_2 is the homotopy quotient of the representable interval by that action, with the final coalgebra in the corresponding twisted marked category. Its fiber along the singleton cover is $y[0,1]$. On singleton evaluation, its action groupoid has a component with stabilizer C_2 over the midpoint. A two-element orbit, such as that of the two endpoints, has trivial stabilizer. This example exhibits nontrivial monodromy and isotropy in the total higher object even though the absolute representing sheaf is 0-truncated.

7.7 Why the cover condition matters

There is a corresponding comparison to sheaves for open covers of ordinary spaces. In that setting the map from the gluing sheaf for two interval halves to the sheaf of continuous interval-valued functions need not be locally surjective. Take the identity plot on an open interval around $1/2$. Every open neighborhood of $1/2$ contains points in both halves. Hence the identity does not locally lift into either single closed half, and it does not factor through the seam. Every section of the pushout sheaf is locally represented by a branch or a seam section, so the identity cannot lie locally in its image. The compact-cover site admits the two closed half preimages as a cover of a compact plot. This accounts for the extra locality used in Proposition 7.3 and Theorem 7.4.

There is also a reason to retain monic markings in the higher category. For the interval specification, if both marks of the terminal one-point set are identified, gluing its two copies identifies them and the result is again a singleton. The singleton is then a final coalgebra in the unrestricted bipointed set category. The monic marking hypothesis is therefore necessary to recover the interval as a final coalgebra in the particular comparison made here.

8 Equivalence with finite point gluing

8.1 An isomorphism of gluing functors

Fix the notation of Ratnayake and Manokaran [18]. A single sort has the finite label set $I = [k]$, branch set M, output labels $D_j = (m_j, i_j)$, and a reflexive symmetric relation R on $M \times I$ with equivalence closure $\overline{R}$. Put

$$B = I, \qquad K = (M \times I)/\overline{R}, \qquad q(m,i) = [(m,i)], \qquad b(j) = [D_j]. \tag{8.1}$$

The sets I, K have their finite discrete topologies. Give I its discrete metric and put $r_m = \rho$ for all m. The condition that b be injective is exactly the non-degeneracy condition on the output classes. An object of $\mathsf{Set}_B^{\mathrm{inj}}$ is exactly a non-degenerate k-pointed set. A topologically closed marking by this finite discrete B into a Hausdorff space is the same as a choice of k distinct points.

Theorem 8.1 (Exact finite-point specialization)**.** *For (8.1) the marked-set functor $\mathcal{F}$ is naturally isomorphic to $\Sigma \otimes -$ of [18]. For any one-bounded k-pointed metric space with distinct labels, the corresponding comparison is an isometric isomorphism for its two quotient metrics. Under this comparison the seed test, initial chain, initial algebra, completed algebra, final coalgebra, and canonical mediating maps agree. For a compatible IFS with the source paper's choice of fixed-point marks, the orbit algebra and completed Hutchinson comparison agree as well.*

Proof. Write $j(i) = x_i$. The set pushout maps each (m, x_i) to $q(m,i)$, and nothing else is identified except what these equations imply. There is a map from the pushout to the original

quotient sending $[m, x]$ to the class $m \otimes x$ and $q(m, i)$ to $m \otimes x_i$. It is well defined because pairs related by $\overline{R}$ give equal classes after replacing labels by marked points. In the other direction, send $m \otimes x$ to $[m, x]$. Each generating R-identification holds in the pushout, so the map descends. The composites are the identity on every representative, including seam representatives since q is onto. This proves the isomorphism. It carries $b(j)$ to $m_j \otimes x_{i_j}$, and its branch formula commutes with every pointed map. The same proof works for arbitrary pointed sets if one uses the actual pushout and does not require monic markings.

Before quotienting, the two constructions have the same set $M \times X$, the same cost $\rho \mathrm{d}_X$ inside each copy, and the same cost one between copies. The comparison identifies their quotient fibers, as just proved. Thus for every pair of endpoints there is equality of the families of admissible chains with their costs, and the two infima are equal. This is an equality for all input metrics under the same representatives, even when the distances between marked points are smaller than one. For finite marks the path metrics separate points by (3.7). The initial map $I \to \mathcal{F}I$ is exactly the map $j \mapsto [D_j]$ into the seam with its quotient metric. Thus its being isometric is the metric admissibility test in both papers.

At depth zero the chain comparison is the identity on I. At depth $n + 1$ use the natural functor comparison and apply the functor to the depth-n comparison. Naturality makes these stage isometries commute with all chain inclusions. Therefore the induced map of colimits is isometric. The initial algebra maps agree by the shift equation (2.3). The completion is the unique extension of that isometry from a dense subset. The completed algebras agree by (4.2), and hence so do their inverses. The transport of coalgebra diagrams along the natural isomorphism is the identity on underlying marked maps, so it identifies the universal mediating maps in sets and in continuous metric spaces.

For the IFS, identify $q(m, i)$ with $\phi_m(\alpha_i)$ using the fixed-point kernel relation. The output $q(j, j)$ maps to α_j, so (6.1) holds. The orbit stages are identical, and (6.2) is the original branch formula. The initial comparison is unique and therefore the same map. The identification $C(O) \cong A$ is the isometry taking a Cauchy sequence to its ambient limit. Under it the completed orbit algebra restricts on branch representatives of O to the same maps as Φ_A. The inclusion $\mathcal{F}O \to \mathcal{F}A$ is dense, by Proposition 3.4. The two continuous maps agree on that dense image, so they coincide. This identifies the IFS comparison and its transported coalgebra. □

The categories in this theorem should be read with their precise maps. In the metric category of [18], a morphism is continuous and pointed, and the input is only required to have distinct marks. This is broader than the category of short maps with isometric marks in Theorem 1.1. Our formula for finite boundaries separates for all distinct marks, and Theorem 3.2 identifies the resulting topologies with $\mathsf{Top}_B^{\mathrm{cl}}$. It follows that continuous finality specializes to the full category of the source. For short maps, restricting to discrete marked distances gives exactly the normalized short-map finality there.

8.2 Finite boundary estimates and admissibility

Let $\Delta(X)$ be the minimum positive distance between distinct marked points, with convention $\Delta(X) = 1$ for one mark. We make explicit three finite consequences of the shortest-path proof.

First, for a reduced minimum chain of cost less than one with n paid steps, every paid step is within one copy and all $n - 2$ interior paid steps have distinct marked endpoints. Hence $(n - 2)\rho\Delta(X) < 1$. If there is a one-step minimum, it already satisfies the bound $n < 2 + (\rho\Delta(X))^{-1}$. Deleting redundant zero jumps and steps does not increase the score, and

at unequal endpoints it leaves at least one paid step. The set of possible reduced chains under this bound is finite because the endpoint fibers and all boundary representative sets are finite. This recovers both attainment and the bounded reduced-chain form of the finite metric formula.

Second, distinct output classes have copy-separated supports if there is no edge between them in the unit-edge boundary graph on the discrete seed. They must be copy-separated for admissibility, since otherwise their distance is at most $\rho < 1$. If $\rho \geq 1/2$ and they are copy-separated, any connecting path has at least two positive edges, so its cost is at least one. This is the equivalence of copy separation and admissibility in that scale range. More generally in an arbitrary finite marked metric input, each nonzero seam edge has weight at least $\rho\Delta(X)$. If the output supports are disjoint and $\rho \geq 1/2$, their distance is therefore at least $\Delta(X)$. The one-mark case follows by convention. At smaller scales the precise shortest-path test of Section 3 replaces the two-edge lower bound.

Third, (5.7) specializes with $L = 1/\Delta(X)$ and $\epsilon_n = \rho^n$. This recovers the finite-stage estimate for the unique mediating map. For a short coalgebra it gives the Lipschitz bound $1/\Delta(X)$, and the output distances of a normalized object with $\Delta(X) = 1$ are again normalized by the seed test. In the original IFS specialization let Δ be the separation of the distinct fixed points and λ a common branch Lipschitz bound. Each fixed point α_j occurs in its own output class $q(j, j)$. The support of that class contains m precisely when $\alpha_j \in \phi_m(O_0)$. Thus pairwise disjoint output supports are equivalent to

$$\phi_m(O_0) \cap O_0 = \{\alpha_m\} \qquad (m \in M).$$

If $\phi_m(\alpha_i) = \alpha_j$ with $m \neq j$, then $\Delta \leq \mathrm{d}(\alpha_m, \alpha_j) \leq \lambda \mathrm{d}(\alpha_m, \alpha_i) \leq \lambda$. This proves the source implication from $\lambda < \Delta$. Its numerical hypotheses concern the first-step boundary only, whereas orbit injectivity concerns all depths.

8.3 The module presentation

The algebraic comparison also identifies the two-object module used in the source. Let $\mathbb{A}_k$ have objects $0, 1$ and arrows $u_i \colon 0 \to 1$ for $i \in I$, together with identities. For a marked set X, let $D_X \colon \mathbb{A}_k \to \mathsf{Set}$ take 0 to the singleton, 1 to X, and u_i to the i-th mark. In the convention $L \colon \mathbb{A}_k^{\mathrm{op}} \times \mathbb{A}_k \to \mathsf{Set}$, the module has values

$$L(0,0) = \{*\}, \quad L(1,0) = \varnothing, \quad L(0,1) = K, \quad L(1,1) = M.$$

The nonempty nonidentity actions are $M \to K$ by $m \mapsto q(m, i)$ for the arrow u_i in the first variable, and $\{*\} \to K$ by $* \mapsto b(i)$ for u_i in the second variable. The coend $\int^p L(p, 1) \times D_X(p)$ has coproduct of terms $K \sqcup (M \times X)$ and imposes the equations $q(m, i) = (m, j(i))$. The coend at 0 is a singleton. Its arrows out of 0 select the marks $b(i)$. This gives the same pushout, including on degenerate inputs.

For this category, a covariant set-valued row is non-degenerate in the sense used by Leinster exactly when its arrows out of the 0-component are injective and have disjoint images [11, Thm. 4.11]. There are no further requirements from compositions of nonidentity arrows, since there are none. The row with 0-component empty satisfies the condition. The row with 0-component a singleton satisfies it if and only if the $b(i)$ are distinct. This recovers the equivalence between module non-degeneracy, distinct output classes, and preservation of the non-degenerate marked category.

8.4 Result-by-result comparison

The table uses the numbering of arXiv:2611.07584v1. Each row identifies the objects and arrows of the source result before making its comparison. The proof of Theorem 8.1 shows that the entries describe the same constructions, so none of the rows relies on similarity of statements alone.

Source result	Specialization and identification
Propositions 4, 7 and Lemmas 5–8	The seam K is the finite class set of $\overline{R}$. Lemma 2.1 and Theorem 8.1 identify the fibers and preserve injective maps, and the marking is mono exactly for distinct output classes. Copy-separated supports imply that those classes are distinct.
Proposition 9 and Lemmas 10–12	Equation (3.7) is the finite boundary graph with the same endpoint and seam vertices. The bound on reduced chains above and (3.5) give the finite local estimates. Proposition 3.4 gives completeness with no normalization of input marked distances.
Propositions 13–14 and 18	Propositions 3.3 and 3.4 give continuous, Lipschitz, and isometric preservation and the natural completion comparison. The arrows and the dense inclusion are identified by Theorem 8.1.
Propositions 15–17	Propositions 2.2 and 4.2 give the same set initial algebra and isometric structure. The discrete seed distance between output classes is the clipped shortest path length, proving the source admissibility equivalence for $\rho \geq 1/2$.
Theorem 1 and Lemma 19	The set part of Theorem 1.1 is over the identical non-degenerate marked category. The completed algebra inverse is the same structure map. Corollary 5.1 is the estimate on representatives with $\epsilon_n = \rho^n$.
Theorem 2, Lemma 20 and Proposition 21	The topological final map is continuous in the original metric topology. Equation (5.7) gives the source finite-stage inequality and the $1/\Delta(X)$ short-coalgebra estimate. $\Delta(X) = 1$ identifies the normalized short categories.
Lemmas 22–23	The fixed-point kernel relation identifies output classes precisely when their fixed points coincide. Output isolation equals copy separation by the support formula above, and $\lambda < \Delta$ is sufficient.
Propositions 24–28	Equation (6.3) gives the source bound $\max(1, \lambda/\rho)$. Proposition 6.1 gives the same orbit-injectivity comparison. The last subsection of Section 6 gives the topological inverse criterion, and Appendix B gives the iterated Lipschitz criterion.
Theorem 3	With the source bi-Lipschitz and orbit-injectivity hypotheses, Theorem 6.4 produces the unique bi-Lipschitz extension and identifies the transported structure with the completed orbit algebra by density. The continuous and set finalities follow in the source categories.
Propositions 29–30	The two-object module yields the same pushout by its coend. Non-degeneracy of its rows is exactly injectivity of the output marking.

The hypotheses of the source theorem remain sufficient when read in this comparison. The theorem on exact completed overlaps adds a separate characterization. It can give a topological isomorphism even when the initial comparison is not bi-Lipschitz, as happens for a mismatch of

the symbolic and Euclidean scales of the Cantor set. The initial algebra, the orbit algebra and the inverse completed structure in that special case are still those of the original construction.

9 Classical and compact-seam examples

9.1 The interval and its dyadic algebra

Let $B = \{0, 1\}$ with the discrete metric, and let the two copies have weights $1/2$. The only nontrivial boundary identification is the right endpoint of the first copy with the left endpoint of the second. The seam K has three points. The output marks are the two outer endpoints. The graph on the discrete seed has a path of two edges between them, so Definition 3.1 holds.

Let $P = [0, 1]$ with the usual metric, and use the branches $x \mapsto x/2$ and $x \mapsto (1 + x)/2$. The induced branch map onto P is an isometry from the gluing metric. Within each half, the direct path has cost $|x - y|$ in output coordinates, and when the endpoints lie in different halves the path through $1/2$ has this same cost. The ambient distance is always a lower bound on the quotient costs by the triangle inequality. The same calculation at stage n identifies A_n with $\{j/2^n : 0 \le j \le 2^n\}$. Thus the initial algebra consists of the dyadic points and the completion is isometric to P. The inverse algebra sends x to $[0, 2x]$ for $0 \le x \le 1/2$ and to $[1, 2x - 1]$ for $1/2 \le x \le 1$, using branch indices $0, 1$. The representatives agree at $1/2$. This is the source's interval coalgebra and the familiar interval equation of Freyd [7, 11, 4].

The binary address $d_1 d_2 \cdots$ evaluates to $\sum_{j \ge 1} d_j 2^{-j}$. Equal addresses either coincide or are obtained by a common prefix and the two boundary addresses $0111 \cdots$ and $1000 \cdots$. Thus in this example the boundary classification in Proposition 6.3 is precisely the two binary expansions of an interior dyadic point. This is why the glued coalgebra is an interval while the unglued one is the space of binary streams. The two quotients differ already in the fibers of the address evaluation.

9.2 Binary streams, the Cantor set, and p-adic integers

For two branches with no gluing, take the same two-element boundary and place its output marks in the distinct first branches, using the same label in each. The initial chain consists of finite words followed by a constant tail, with extension of a word by one copy of the repeated label. Its union is the eventually constant binary streams. The output metrics give the ultrametric r^ℓ for two unequal streams with longest common prefix length ℓ, for any $0 < r < 1$. Completion gives all streams. The inverse algebra takes the first digit and the remaining stream.

With geometric branches $x/3$ and $(x + 2)/3$, the attractor is the middle-thirds Cantor set, and the address map is

$$\omega \longmapsto \sum_{j \ge 1} 2\omega_j 3^{-j}.$$

There are no additional overlaps. If two streams have common prefix length ℓ and then differ, the difference between their images is bounded between $3^{-\ell-1}$ and $3^{-\ell}$. The first bound comes from the gap between the two first divergent subintervals, and the second comes from the length of their common containing interval. At $r = 1/3$ the comparison is bi-Lipschitz. Conversely the images of 0^ω and $0^n 1^\omega$ are at distance 3^{-n}, while the symbolic distance is r^n. If $r \ne 1/3$, one of the comparison constants diverges. Theorem 6.4 gives a homeomorphism at every admissible scale because the geometric pieces remain disjoint. This recovers the original example and makes the distinction between topology and the quantitative comparison explicit.

For a prime p, take p unglued branches, digits in $\{0, \dots, p-1\}$, and weight p^{-1}. With a single mark, choose the output corresponding to digit zero applied to the mark. The stages give streams eventually zero and evaluation by $\sum_{j \geq 0} a_j p^j$ identifies them with the nonnegative integers in $\mathbb{Z}_p$. With p marks choose the a-th output to be digit a applied to the a-th mark. The initial algebra is then the eventually constant streams. The geometric mark for a is the fixed point $a/(1-p)$ of $\phi_a(z) = a + pz$. A finite prefix $a_0 \dots a_{n-1}$ followed by the tail b^ω evaluates to

$$\sum_{j=0}^{n-1} a_j p^j + \frac{bp^n}{1-p}.$$

The two presentations have different sets of marks and the same complete metric carrier. In either case, if the first differing digit is at position ℓ, the difference of the evaluations is p^ℓ times a p-adic unit. This proves that evaluation is an isometry for the prefix metric $p^{-\ell}$. Completeness and uniqueness of p-adic expansions identify the completion with $\mathbb{Z}_p$. The algebra is $(a, z) \mapsto a + pz$ and the inverse sends z to its residue a modulo p and the tail $(z-a)/p$. This preserves the designated fixed points in the p-marked version. It recovers both p-adic constructions of [18] and the earlier one-marked metric completion in [14].

9.3 A geometric lemma for intrinsic metrics

We use an intrinsic metric in the gasket and carpet calculations. If L is a subset of a normed plane, write d_L^{int} for the infimum of lengths in the norm of rectifiable paths in L connecting the endpoints. Suppose L has a finite closed cover by sets L_i, each of which admits a retraction from L that does not increase the length of any path. Then

$$d_L^{\mathrm{int}}(x, y) = \inf \left\{ \sum_{j=1}^{m} d_{L_{i_j}}^{\mathrm{int}}(x_{j-1}, x_j) : x_0 = x,\ x_m = y,\ x_{j-1}, x_j \in L_{i_j} \right\}. \tag{9.1}$$

The inequality from the right to the left follows by concatenating paths of length arbitrarily close to the terms on the right. To prove the other direction, take a path $\gamma\colon [0,1] \to L$ of finite length. At time $t < 1$, among the tiles containing $\gamma(t)$ choose one whose last visit in $[t, 1]$ is latest. That last visit is strictly later than t. Otherwise choose times strictly decreasing to t; a single tile contains infinitely many of their images, and by closedness it also contains $\gamma(t)$, a contradiction. Jump to the last visit time of the selected tile. Each selected tile has then been exhausted and cannot be selected again with a strictly later visit. The procedure terminates after at most the number of tiles and produces a finite partition of the path with both endpoints of each subpath in a common tile. Retraction of each subpath to that tile gives a chain of total intrinsic cost at most the original length. Taking the infimum proves (9.1). This argument permits an arbitrary number of seam crossings by the original path.

Here is a second useful observation. Consider a compact residual set in a convex polygon after removing countably many open polygonal holes whose boundaries remain in the residual set. Suppose any two boundary points of one hole can be joined along its boundary by a path of length at most C times their norm distance. For any segment between residual points, its intersections with hole interiors are disjoint open subintervals. Replace each open subinterval by one of the boundary paths, traversed at constant speed. For finitely many replacements the resulting path has speed bounded by C times that of the segment, using $C \geq 1$. As more replacements are made the paths converge uniformly: the displacement on a remaining interval

is bounded by a constant times its length, and only finitely many intervals have length larger than any given threshold. The limit remains in the polygon and avoids every hole. Its speed is bounded by the same constant in the Lipschitz sense, and its length is at most C times the length of the segment. It follows that the intrinsic metric is bi-Lipschitz equivalent to the restriction of the ambient norm, with lower bound the norm distance and upper bound C times that distance. In particular the intrinsic metric is compact on the residual set.

9.4 The Sierpiński gasket

Let L be the gasket in a Euclidean equilateral triangle of side length one with vertices v_0, v_1, v_2. Its branch maps are $x \mapsto (x + v_i)/2$. Their images meet pairwise at the midpoints $(v_i + v_j)/2$, which are exactly the coincidences of first-level marked points. Use these three vertices as discrete marks. The output classes lie in distinct branches, so $\rho = 1/2$ is admissible. This recovers the point-glued functor of [17].

For completeness, its metric comparison can be proved directly. In an equilateral triangular hole, two points on adjacent sides can be joined via their common vertex. If their distances to the vertex are a, b, then the square of their chord length is a^2+b^2-ab, and $(a+b)^2 \le 4(a^2+b^2-ab)$. Points on the same side are joined along it at their Euclidean distance. Thus the residual-set observation gives

$$\|x - y\|_2 \le d_L^{\mathrm{int}}(x, y) \le 2\|x - y\|_2. \tag{9.2}$$

Nearest-point projection onto the convex first-level triangle of a tile is Euclidean short and fixes that triangle. A point outside the triangle projects to its boundary, which is in the tile's gasket. A point of the gasket already inside the triangle belongs to that tile, since other first-level tiles can meet it only at marked vertices. Consequently the projection restricts to a retraction of L onto the tile and does not increase path lengths. The intrinsic metric restricted to the tile is therefore equal to the metric computed using paths confined to the tile, which is $1/2$ times d_L^{int} under the branch similarity.

To see the one-boundedness of d_L^{int}, let R be the supremum of its distances from an outer vertex to any point, over the three choices of vertex. It is finite by (9.2). In the first-level tile containing the vertex the distance is at most $R/2$. In another tile, travel first along an outer half-edge to the shared vertex and then inside the other tile. The distance is at most $1/2 + R/2$. Therefore $R \le 1$. Two points in different tiles can go through their common marked vertex at total cost at most one. Two points in the same tile are at distance at most half the diameter of L. Taking the supremum of all distances shows that the diameter is at most one. The three outer vertices are at distance one from each other by the ambient lower bound and the presence of the outer sides.

The set gluing map is bijective onto L by the first-level intersection description. The within-branch costs are the intrinsic metrics of the tiles. Formula (9.1), together with one-boundedness, identifies the quotient metric with d_L^{int}. More explicitly, the ambient intrinsic distance bounds the cost of every paid step and every quotient chain, and (9.1) produces chains of arbitrarily close cost. Lemma 4.3 then identifies the canonical completion isometrically with the intrinsic metric and bi-Lipschitz with the Euclidean metric. The finite stages are the images of the three vertices under words of length n, so the set and orbit comparisons coincide with the source's gasket. This gives the same continuous final coalgebra and a direct quantitative comparison in this presentation.

9.5 An interval system with two sorts and a unit weight

Let the sorts be s, t, each with two discrete marks. Sort s has a single branch to sort t with weight one and no identifications. Sort t has two branches to sort s with weights $u, 1-u$, where $0 < u < 1$, glued end to end as in the interval. The output labels are the two outer endpoints. The first-step seed maps are isometric at both sorts, since at sort t the path between its outputs has cost $u + (1-u) = 1$.

Take $P_s = P_t = [0,1]$. Use the identity for the s-branch, and use the two similarities $x \mapsto ux$ and $x \mapsto u + (1-u)x$ for the t-branches. The gluing quotient onto P_t is isometric by the path through u, while the quotient onto P_s is the identity. Every path of length n encounters at least $\lfloor n/2 \rfloor$ of the strict weights. Hence $\epsilon_n = \max(u, 1-u)^{\lfloor n/2 \rfloor}$ and Lemma 4.3 identifies this tuple with the completion. At every depth, the intervals in sort t keep track of the cuts in the variable-ratio binary subdivision. The case $u \neq 1/2$ has distinct branch scales and the edge of weight one is essential to the displayed graph, so a common one-step contraction constant for all branches would not establish its finality. The compact-cover theorem applies to the same two-type family.

9.6 A square with shared edges

Take $P = [0,1]^2$ with metric $\mathrm{d}_P(x,y) = \frac{1}{2}\|x - y\|_1$, and let $B = \partial P$ have the induced metric. Use the four translations of the half-size square, with weight $1/2$, indexed by $\ell \in \{0,1\}^2$ and branch maps $\phi_\ell(x) = (x + \ell)/2$. Put

$$K = \bigcup_{\ell \in \{0,1\}^2} \phi_\ell(B), \qquad q(\ell, b) = \phi_\ell(b), \qquad b \colon B \hookrightarrow K.$$

The square gluing is a bijection from the pushout because distinct tile interiors do not meet and every tile boundary overlap is represented by the same element of K.

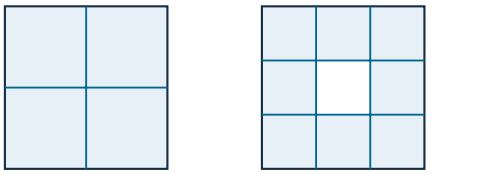
A straight segment in P can be subdivided into finitely many segments inside the four tiles by cutting at the grid lines. Segments contained in a grid line are also contained in a tile. The successive lengths in the scaled ℓ^1 metric sum to the original length. Thus the quotient distance is at most d_P; it is at least d_P because each branch step has its ambient cost and the cross-branch cost one is a global upper bound on ambient distances. The quotient is an isometry. Since B has the induced metric, the first part of Theorem 3.2 identifies the metric of K with its induced subspace metric. Therefore the specification is admissible. Each stage is the union of the boundaries of the depth-n dyadic squares and the completion is P. This exhibits a genuinely infinite compact boundary: any one of the open shared edges contains points that would be unmarked in a finite point specification and hence could not be identified by its normal form.

9.7 The carpet with its full overlap relation

Let $P \subseteq [0,1]^2$ be the usual Sierpiński carpet. Its eight branches are $\phi_\ell(x) = (x + \ell)/3$, with $\ell \in \{0,1,2\}^2 \setminus \{(1,1)\}$. Equip P with

$$\mathrm{d}_P(x,y) = \tfrac{1}{4} d_P^{\mathrm{int}}(x,y),$$

where lengths for d_P^{int} use the ℓ^1 norm. Every deleted open square has boundary contained in P. On the boundary of one such square, points on the same or adjacent sides can be joined at ℓ^1 cost equal to their norm distance, by using the common corner when necessary. For points



Dyadic square Carpet, first subdivision

Figure 1: First-stage seam spaces. The outer perimeters carry the output markings. The interior lines record full edge coincidences between adjacent retained tiles or a retained boundary facing the central hole.

on opposite sides, one of the two routes around the boundary has cost at most twice the side length. The chord has ℓ^1 norm at least the side length. The residual-set observation gives

$$\tfrac{1}{4}\|x-y\|_1 \le \mathrm{d}_P(x,y) \le \tfrac{1}{2}\|x-y\|_1. \tag{9.3}$$

Thus d_P is a one-bounded compact metric with the ordinary carpet topology.

Take $B = \partial[0,1]^2 \subseteq P$ with the metric induced by d_P, and $K = \bigcup_\ell \phi_\ell(B)$, with $q(\ell, b) = \phi_\ell(b)$ and output inclusion of B. The weight of every branch is $1/3$. The coordinates clamped to the closed first-level square of a tile define an ℓ^1-short map. A point outside that square is sent into its boundary, and a point inside the square belonging to P belongs to the tile's carpet. Thus clamping restricts to a length-nonincreasing retraction of P onto the tile. The intrinsic metric inside the tile therefore agrees with the restricted intrinsic metric from P, and scales by $1/3$. Distinct first-level carpets meet only along their square boundaries, and all their common boundary points are included in the prescribed seam K. Therefore the pushout map onto P is bijective. Equation (9.1) shows that it is isometric for d_P, just as in the gasket argument. Its restriction to K is the seed metric by Theorem 3.2, proving admissibility.

The initial chain starts with the outer square perimeter and adds the perimeters of all retained cells at each depth. By Lemma 4.3, its completion is the carpet with d_P. The Euclidean or box metric gives an equivalent continuous realization and a bi-Lipschitz comparison by (9.3). Noquez and Moss establish an independent coalgebraic presentation with a category of square metric spaces [16]. The calculation here locates the carpet within the same compact-seam and higher-finality theorem as the interval and the gasket. The intrinsic metric on the geometric carpet provides the boundary metric used in this identification.

The finite-point obstruction in [18] can be seen directly. Use the eight fixed-point marks $\alpha_\ell = \ell/2$ and the associated finite relation. In the left lower first-level tile and its right neighbor, the points $(1, 1/3)$ and $(0, 1/3)$ of the tail carpet have the same image $(1/3, 1/9)$ under their respective branches. Both tail points lie in the finite marked orbit: they are $\phi_{(2,0)}(\alpha_{(2,2)})$ and $\phi_{(0,0)}(\alpha_{(0,2)})$, respectively. Neither is one of the fixed-point marks. Consequently their two representatives survive as distinct elements of the finite point-gluing functor by Lemma 2.1. This violates orbit injectivity and cannot be remedied by changing the finite gluing scale. With the compact marks both tails lie in B and their equal images are recorded by q. The same holds along the full common edge. Hence the completed comparison meets the exact-overlap criterion of Theorem 6.4.

9.8 Nested radicals as a continuous coalgebra

The continuous result covers structures whose geometric inverse branches fail to be Lipschitz. On $X = [-2, 2]$ with marked endpoints, let

$$\psi_\delta(x) = \delta\sqrt{2 + \delta x}, \qquad \delta \in \{-1, 1\}.$$

These maps have images the two closed halves and agree at the prescribed seam. The gluing algebra on X is a continuous bijection between compact Hausdorff spaces. Its inverse defines a continuous coalgebra. Let $L = [-1, 1]$ with the usual rescaled interval metric and endpoints, and let $\varphi_\delta(u) = (u + \delta)/2$. Set $g(x) = \frac{2}{\pi}\arcsin(x/2)$. The elementary half-angle identities give $g\psi_\delta = \varphi_\delta g$ for each δ, including the endpoints. Thus g is the unique coalgebra morphism from X to the balanced interval. Since its inverse is $u \mapsto 2\sin(\pi u/2)$, it gives another topological realization of the final coalgebra.

If the successive signs in a nested radical are $\varepsilon_1, \varepsilon_2, \ldots$, put $\delta_i = \prod_{j=1}^{i} \varepsilon_j$. The radical truncated after n signs is $\psi_{\delta_1} \cdots \psi_{\delta_n}(0)$ by expanding the definitions, since $\delta_i\delta_{i+1} = \varepsilon_{i+1}$. Its balanced interval value is $\sum_{i=1}^{n} \delta_i 2^{-i}$. Passing to the limit gives the formula

$$2\sin\left(\frac{\pi}{2}\sum_{i=1}^{\infty}\frac{\varepsilon_1 \cdots \varepsilon_i}{2^i}\right).$$

Every point in the balanced interval has a signed binary expansion, and the corresponding ε_i are recovered from the successive δ_i. This is the same coinductive identification as in the source, and the explicit formula agrees with the classical radical analysis [6].

10 Changing the overlap relation

A specification can impose additional seam identifications while keeping its marks and branches. This produces a morphism of completed systems. The comparison makes precise which part of an attractor changes when one modifies its declared overlaps.

Proposition 10.1 (Comparison under seam quotients)**.** *Suppose two admissible specifications have the same sorts, branches, weights and metric boundary spaces. Write their seam data as (K, q, b) and (K', q', b'). Assume there are continuous surjections $k_s \colon K_s \to K'_s$ such that $q'_s = k_s q_s$ and $b'_s = k_s b_s$. Under eventual contraction, there is a canonical short surjection $H \colon \widehat{G} \to \widehat{G}'$. If $\tau \colon \mathcal{F} \to \mathcal{F}'$ is the natural comparison of gluing functors, then*

$$\gamma' H = \mathcal{F}'(H)\tau_{\widehat{G}}\gamma. \tag{10.1}$$

On addresses, $H_s\pi_s = \pi'_s$. In particular, equality of canonical addresses for the first specification implies equality for the second, and H is an isomorphism of topological carriers if and only if the two address relations agree.

Proof. On each pushout define τ_X by the same branch representative, using k_s on the seam. The conditions on k_s make this well defined and marked, and it is natural in X. It is surjective because the branch quotient is surjective. Every chain admissible for the first gluing is admissible after applying the comparison, with the same paid costs. Therefore τ_X is short.

Start with $u_0 = \mathrm{id}_B$ and recursively define $u_{n+1} = \tau_{A'_n}\mathcal{F}(u_n) \colon A_{n+1} \to A'_{n+1}$. The maps commute with the initial-chain inclusions by naturality and the defining equations of the marks.

Each is short. Each is also onto because $\mathcal{F}$ preserves surjections of marked sets by the branch formula and τ is onto. Thus the induced map $G \to G'$ is a short surjection. Its unique short extension to the completions has compact image containing G', hence is onto. The stage equations yield $Ha = a'\mathcal{F}'(H)\tau_{\widehat{G}}$ on the dense stage union of $\mathcal{F}\widehat{G}$. All maps are continuous, so the equation holds everywhere, and gives (10.1). It also shows that H commutes with each branch. The nested images of any address shrink to the same point as its evaluation in the second specification, giving $H_s\pi_s = \pi'_s$. The last assertion follows from the surjectivity of both address evaluations and compactness. □

The morphism H can equivalently be obtained by sending the first final coalgebra to an $\mathcal{F}'$-coalgebra via $\tau_{\widehat{G}}$ and using finality of $\widehat{G}'$. This characterization immediately gives composition for successive seam quotients. In the condensed category the same construction gives the map yH of relative final objects, with the first structure changed by τ. The comparison commutes with descent for equivariant seam quotients by the uniqueness in the coalgebra mapping spaces.

10.1 Role of the hypotheses

Compactness of the boundaries supplies the fixed compact seam, the infimum attainment in the entry and exit variables, and the closedness of every finite-depth piece. The Cantor-function example in Section 3 shows that a compact Hausdorff equivalence quotient can still have a degenerate path pseudometric. In the finite point case this issue disappears because the relevant graph has finitely many positive lengths for each fixed pair of endpoints.

Isometry of the seed embedding ensures that a stage carries the same distances after it is inserted into a later stage. When it fails, the algebraic chain can still exist, but the metric colimit used here has no such distance characterization. Likewise, eventual contraction is a substantive condition. With two marks, one branch of weight one, and $q = b = \mathrm{id}_B$, the functor is the identity and the seed is isometric. The completion of its initial algebra is B. The category of sets with two distinct marks has no terminal object, since a third unmarked element in a test object can be mapped to either mark. Therefore the identity functor has no final coalgebra on that category. The scale condition cannot be dropped as a general existence theorem.

In the higher result, the representing object is the condensed set of continuous plots into the compact completion, with an arbitrary higher parameter object allowed in the base. The descent theorem applies to locally given compact-seam data and their coherent transports. Its proof uses effective finite compact covers, separation of the compact metric quotients, and monicity of the markings. The open-cover and collapsed-mark examples of Section 7 identify two distinct failures of the representing argument if those conditions are changed.

10.2 Conclusion

The initial chain, its metric, and the final universal property can all be retained when the overlaps are compact spaces. The seam metric reduces the admissibility calculation to the initial boundary. The closed-cover argument then gives both ordinary continuity and continuity on condensed plots. The geometric comparison identifies exactly which extra overlaps are responsible when an attractor is a proper quotient of the canonical completion. In the point specialization these are statements about the same functors and maps as in [18]; the compact examples and higher descent apply to overlap and parameter data outside that specialization.

A Lean verification of the specialization

The accompanying source file `lean/Specialization.lean` is checked with Lean 4.19.0 and uses its standard library [10]. It formalizes the comparison for arbitrary types M, I and an arbitrary relation R on $M \times I$. Thus the algebraic proof does not enumerate a finite model. Finiteness is used in the paper to specialize the metric construction, not in the quotient isomorphism. The two presentations in the development are defined independently. The original carrier is the quotient of $M \times X$ by the equivalence generated by replacements of related boundary labels. The new carrier is a quotient of $K \sqcup (M \times X)$ by the equations identifying each boundary representative with its class in K.

Lean declaration	Mathematical content
`backward_forward` and `forward_backward`	The quotient maps in the two directions are inverse on every element.
`marking_preserved`	The chosen output labels correspond.
`naturality`	The carrier equivalence commutes with every pointed map.
`pushout_universal`	The new quotient has the universal property of the span pushout.
`algebra_iff`	The algebra morphism equation is preserved and reflected by transport.
`unique_coalgebra_map_iff`	For each source coalgebra, transport preserves and reflects the existence of a unique pointed coalgebra morphism.
`chain_cost_iff`	The same cost is realized by a finite paid chain between corresponding endpoints in either quotient.
`quotient_scores_agree`	Applying a common infimum operation to the two predicates of realizable costs gives identical distances.

The chain predicate is generated by paid steps and concatenation at a common quotient point. Equality of quotient endpoints is the zero-cost jump. Thus it formalizes the comparison of admissible chains with identical scores used in Theorem 8.1. The costs, their binary sum operation, and the operation taking an infimum over realizable costs are parameters of this theorem. In the metric instance they are real distances, addition, and the corresponding real infimum. The file proves equality for any such instantiation. The numerical and analytic facts about those infima are supplied by the paper proofs in Sections 3 and 4. Completion, higher sheaf finality, and the geometric examples are proved in the manuscript.

To check the development, run `lean lean/Specialization.lean` with the bundled toolchain specification. The file prints the axiom dependencies for the main declarations. The checks use the ordinary quotient and propositional extensionality axioms of Lean where displayed in the check output. There are no admitted declarations or additional axioms. The archived check log records the version and the reported dependencies.

Authorship and evidence. The responsible submitter uses the anonymous byline. AI assistance was used to develop, draft, and typeset the arguments and to produce and check the Lean companion. The mathematical evidence consists of the proofs in the manuscript, the explicit comparison with the cited preprint, and the Lean kernel check of the declarations above.

B Quantitative comparison formulas

B.1 Hölder bound for a coalgebra map

Take the setting of (5.7), with $K > 1$ a Lipschitz bound on c and $\epsilon_n \le C\theta^n$ for $0 < \theta < 1$. Put

$$\alpha = \frac{\log(1/\theta)}{\log(K/\theta)} \in (0, 1).$$

For $0 < d = \mathrm{d}_X(x, y) < 1$, choose $n = \lceil \log(1/d)/\log(K/\theta) \rceil$. Then $\theta^n \le d^\alpha$ and $K^n d \le K d^\alpha$. Using that $c^{[n]}$ is K^n-Lipschitz, (5.7) gives

$$\mathrm{d}_{\widehat{G}}(f(x), f(y)) \le \big(2(1 + L)C + LK\big) d^\alpha.$$

The endpoint cases $d = 0, 1$ follow from equality or the one-boundedness of the target. For $K = 1$, passing to the limit in (5.7) gives the Lipschitz estimate with constant L.

B.2 Iterated orbit estimates

Write $u\colon B \to O_0$ for the marking isomorphism of sets and let $\Phi^{(n)}\colon \mathcal{F}^n O_0 \to O_n$ be the iterated branch evaluation. When the orbit algebra is injective, all these maps are bijective. From the recursion in Propositions 2.2 and 6.1,

$$h_0 \iota_n = \Phi^{(n)} \mathcal{F}^n u, \qquad h_0^{-1}|_{O_n} = \iota_n \mathcal{F}^n(u^{-1})(\Phi^{(n)})^{-1}.$$

Suppose the orbit gluing metrics are separated at these stages and u^{-1} is L_--Lipschitz, for $L_- \ge 1$. If the inverse iterated orbit evaluations have common Lipschitz bound J, the displayed identity and preservation of Lipschitz bounds give JL_- on every stage. Passing to a common stage for any two points of O proves that h_0^{-1} is JL_--Lipschitz. Conversely, bounds on h_0^{-1} give uniform iterated bounds if u is Lipschitz: restrict h_0^{-1} to O_n, identify A_n by ι_n, and apply $\mathcal{F}^n u$. Thus, under bi-Lipschitz boundary comparison, uniform iterated inverse bounds and a Lipschitz inverse on the orbit are equivalent.

In the point specialization, $L_- = 1/\Delta(O_0)$. If $\lambda \le \rho$, u is short and $\Phi^{(1)}$ is short. It follows that h_0 is short and the uniform inverse bound yields the exact bi-Lipschitz criterion from the source. In that case the orbit structure and its inverse are bi-Lipschitz by the identities

$$\Phi_O = h_0 a_G \mathcal{F}(h_0^{-1}), \qquad \Phi_O^{-1} = \mathcal{F}(h_0) a_G^{-1} h_0^{-1}.$$

For a general weighted graph the same implication holds with the branchwise hypotheses $\lambda_e \le r_e$ and with L_0 as in Section 6. A fixed Lipschitz bound on a one-step inverse alone does not give a uniform bound on its iterates, since the constants may multiply.

B.3 Forward bounds when the scales differ

Let $L_0 \ge 1$ bound the Lipschitz constant of u, and put $K_\Phi = L_\Phi$ from (6.3). Under the same separation hypotheses, the depth-n comparison $h_0 \iota_n$ is $L_0 K_\Phi^n$-Lipschitz, by iteration of the stage recursion. For $x, y \in \widehat{G}_s$, select depth-n representatives and replace their tails by marks to obtain $x_n, y_n \in A_{n,s}$. Their distances to x, y are bounded by ϵ_n. In the attractor the images under h and h_0 of each such pair lie in the same geometric piece, of diameter at most η_n. It follows that

$$\mathrm{d}_A(h(x), h(y)) \le 2\eta_n + L_0 K_\Phi^n \big(\mathrm{d}_{\widehat{G}}(x, y) + 2\epsilon_n\big).$$

This is useful when $K_\Phi^n \epsilon_n \to 0$. For common scales $r_e = \rho$, $\lambda_e = \lambda$ with both below one, $K_\Phi^n \epsilon_n = \max(\rho, \lambda)^n$. The estimate therefore gives an explicit modulus of continuity even when the stage comparison has no uniform Lipschitz bound. The address proof of Proposition 6.3 establishes continuity for every compatible system satisfying the two eventual contraction hypotheses.

References

[1] J. Adámek. On terminal coalgebras derived from initial algebras. *8th Conference on Algebra and Coalgebra in Computer Science*, LIPIcs 139, 12:1–12:21, 2019. DOI 10.4230/LIPIcs.CALCO.2019.12.

[2] P. America and J. J. M. M. Rutten. Solving reflexive domain equations in a category of complete metric spaces. *Journal of Computer and System Sciences* 39(3), 343–375, 1989. Publisher record.

[3] M. Barr. Terminal coalgebras in well-founded set theory. *Theoretical Computer Science* 114(2), 299–315, 1993. Publisher record.

[4] P. Bhattacharya, L. S. Moss, J. K. Ratnayake, and R. Rose. Fractal sets as final coalgebra obtained by completing an initial algebra. In *Horizons of the Mind*, Lecture Notes in Computer Science 8464, 146–167, 2014. Publisher record.

[5] P. Scholze, based on joint work with D. Clausen. Lectures on condensed mathematics. Updated lecture notes, 2026. arXiv:2605.03658.

[6] J. Fernández-Sánchez and W. Trutschnig. Nested square roots of 2 revisited. *The American Mathematical Monthly* 127(4), 344–351, 2020. DOI 10.1080/00029890.2020.1707059.

[7] P. Freyd. Real coalgebra. Posting to the Categories mailing list, 22 December 1999. Archived posting.

[8] I. Hasuo, B. Jacobs, and M. Niqui. Coalgebraic representation theory of fractals. *Electronic Notes in Theoretical Computer Science* 265, 351–368, 2010. Publisher record.

[9] J. E. Hutchinson. Fractals and self-similarity. *Indiana University Mathematics Journal* 30(5), 713–747, 1981. DOI 10.1512/iumj.1981.30.30055.

[10] Lean Project. *The Lean Language Reference*. Reference manual. The accompanying development is pinned to Lean 4.19.0.

[11] T. Leinster. A general theory of self-similarity. *Advances in Mathematics* 226(4), 2935–3017, 2011. DOI 10.1016/j.aim.2010.10.009. arXiv:1010.4474.

[12] J. Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170. Princeton University Press, 2009. arXiv:math/0608040.

[13] J. Lurie. *Kerodon*. Online reference for homotopy-coherent mathematics. Reference site.

[14] A. Manokaran, J. Ratnayake, and R. Jayewardene. $\mathbb{Z}_p$ as a final coalgebra obtained by Cauchy completing the initial algebra. *Journal of the National Science Foundation of Sri Lanka* 47(1), 2019. DOI 10.4038/jnsfsr.v47i1.8932.

[15] R. D. Mauldin and S. C. Williams. Hausdorff dimension in graph directed constructions. *Transactions of the American Mathematical Society* 309(2), 811–829, 1988. Publisher PDF.

[16] V. Noquez and L. S. Moss. The Sierpiński carpet as a final coalgebra. *Theory and Applications of Categories* 45(2), 33–129, 2026. Journal record.

[17] J. Ratnayake, A. Manokaran, R. Jayewardene, V. Noquez, and L. S. Moss. Presenting the Sierpinski gasket in various categories of metric spaces. *Applied Categorical Structures* 32, article 21, 2024. DOI 10.1007/s10485-024-09773-0.

[18] J. Ratnayake and A. Manokaran. Presenting point-glued self-similar objects as final coalgebras in metric spaces via completion. A categorical Hutchinson theorem. Preprint, 5 November 2026. arXiv:2611.07584v1.

[19] J. J. M. M. Rutten and D. Turi. On the foundations of final coalgebra semantics. Non-well-founded sets, partial orders, metric spaces. *Mathematical Structures in Computer Science* 8(5), 481–540, 1998. DOI 10.1017/S0960129598002588.