# Deriving Regularity Hypotheses from Admissible TPB Dynamics in VERSF

**Local Coupling, Transport Sparsity, Refinement Stability, and the Substrate Origin of Continuum Geometry**

*Keith Taylor — VERSF Theoretical Physics Programme*

---

## General-Reader Summary

The companion paper *Continuum-Limit Regularity and Cone Convergence in VERSF* proved a conditional result: *if* the discrete TPB substrate satisfies five external regularity hypotheses — uniform doubling (H6), uniform local transport sparsity (H6′), bounded combinatorial dimension (H7), refinement Hölder compatibility (H8), and smooth-structure existence (H9) — then its refinement sequence converges to a smooth Lorentzian continuum equipped with exactly the structure required by the Lorentzian-emergence machinery of *Structural Necessity of Lorentzian Geometry in VERSF*. That paper honestly identified its hypotheses as external and named their substrate-level derivation as the principal remaining mathematical risk.

This paper addresses that question.

The intuition is straightforward. Regularity at the continuum level cannot be wished into existence: it has to come from constraints on the substrate. A discrete substrate that grows wildly under refinement will not produce a finite-dimensional continuum. A substrate whose admissible local directions proliferate uncontrollably will not produce a finite-dimensional cone field. A substrate whose admissible causal trajectories can return arbitrarily close to themselves will produce almost-closed causal curves at the continuum scale. A substrate whose refinement maps act discontinuously on cone-direction sets will produce a continuum cone field that is too rough to support a smooth metric.

The paper identifies seven substrate-level engineering constraints — bounded local coupling, controlled refinement multiplicity, uniform distinguishability separation, bounded refinement distortion, bounded refinement overlap, refinement bi-Lipschitz behaviour on cone-direction sets, and uniform Ahlfors d-regularity in pseudometric balls — that prevent each of these failure modes. The fourth constraint, *bounded refinement distortion*, is formulated in terms of an explicit substrate-engineering observable: the refinement-distortion functional $\Xi_\ell$, which measures how differently two nearby substrate regions evolve under refinement. Bounded refinement distortion says that this functional grows no faster than a Hölder modulus — equivalently, that the substrate cannot refine chaotically. The seventh constraint, *Ahlfors d-regularity*, says that the substrate's volume profile in pseudometric balls is bounded both above and below by a fixed power $r^d$ of the radius; the upper bound rules out tree-like architecture (whose volume grows like $\Delta^r$ and which fails the doubling property), and the lower bound rules out arbitrarily sparse substrates

whose small balls become too empty to anchor the packing argument that underlies doubling. Together the two bounds give the standard Ahlfors-regularity-implies-doubling input needed for finite-dimensional continuum limits. We call a TPB substrate satisfying these seven constraints (in addition to the baseline VERSF substrate axioms) a *regular TPB substrate*. The main theorem shows that regular TPB substrates satisfy four of the five continuum regularity hypotheses (H6, H6′, H7, H8) automatically — with the substrate-level Hölder cone stability of the discrete admissible-direction sets *derived from* bounded refinement distortion and the bi-Lipschitz behaviour of refinement on cones, rather than postulated directly. The fifth hypothesis, H9, is structurally different — it concerns the smooth-manifold topology of the limit and intersects the Donaldson–Freedman exotic-4-manifold theory — and is not derived here.

The result converts the geometry-emergence chain from a long conditional statement about external regularity hypotheses to a shorter statement about substrate engineering. The remaining tasks — exhibiting concrete TPB substrate models that satisfy the engineering constraints, sharpening the substrate-level origin of H9, deriving Ahlfors d-regularity (in both directions, upper and lower) from the K = 7 closure architecture rather than assuming it, and tightening the constants delivered by each of the seven engineering constraints — are now concrete substrate-engineering problems rather than open mathematical questions about continuum convergence.

---

# Abstract

We derive four of the five regularity hypotheses required by *Continuum-Limit Regularity and Cone Convergence in VERSF* from seven substrate-level engineering constraints on admissible TPB refinement dynamics. The constraints — bounded local coupling (R1), controlled refinement multiplicity (R2), uniform distinguishability separation (R3), bounded refinement distortion (R4′), bounded refinement overlap (R5), refinement bi-Lipschitz on cone-direction sets (R6), and uniform Ahlfors d-regularity in pseudometric balls (R7) — are strengthened forms of conditions implicit in the VERSF substrate axioms A0–A4 and BCB, together with the K = 7 closure architecture established in the *sequential-interface-transport* paper. Substrates satisfying R1–R7 in addition to A0–A4, BCB, H1, and H5 (with A1 strengthened to A1* and A2 strengthened to A2*) are called *regular TPB substrates*.

We prove four theorems and a main synthesis.

- **Theorem 1 (Uniform volume growth implies uniform doubling — proven, conditional on R7).** Polynomial volume growth (R7) in pseudometric balls — uniform across refinement levels — implies H6. R1 (bounded local degree) and R2 (controlled refinement multiplicity) alone give at most exponential growth (the Δ-regular-tree regime), which is consistent with the failure of doubling; R7 is therefore a genuine additional engineering constraint, not derivable from R1 + R2 alone, and its substrate-level grounding (the conjecture that K = 7 + R4′ entail polynomial growth) is open.
- **Theorem 2 (Finite closure catalogue implies bounded combinatorial dimension — proven, conditional on R1 and the K = 7 closure structure of the sequential-interface-transport paper).** Finite admissible local transport branching implies H7.

- *Theorem 3 (Finite distinguishability plus irreversibility implies transport sparsity — proven, conditional on A1, A2*, R3, R5).** The substrate-level reality of irreversible commitment combined with a uniform lower bound on operational distinguishability implies H6′.
- **Theorem 4 (Bounded refinement distortion implies Hölder cone stability — proven, conditional on R4′, R5, R6, and H5, via the refinement-distortion-propagation Lemma 4.1).** Bounded refinement distortion — quantified by the explicit refinement-distortion functional $\Xi_\ell(x, y) := d_H(\mathcal{R}(\mathcal{C}_\ell(x)), \mathcal{R}(\mathcal{C}_\ell(y)))$ introduced in §8.2 — together with bi-Lipschitz refinement on cone-direction sets (R6) implies H8. Hölder cone stability of the discrete sets $\mathcal{C}_\ell$ is no longer postulated but *derived* (Lemma 4.1) from the substrate-engineering bounds on $\Xi_\ell$ and on the inverse-refinement action.
- **Theorem 5 (Main synthesis — proven by conjunction of Theorems 1–4).** Every regular TPB substrate satisfies H6, H6′, H7, and H8. Conditional on H9 (smooth-structure existence on the limit), the substrate refinement sequence converges, by the companion paper's Theorems 5 and 8, to a strongly causal Lorentzian length space admitting a $C^k$ Lorentzian metric for every prescribed $k \geq 2$.

The principal methodological move is the explicit factoring of the geometry-emergence chain into a *substrate-engineering layer* (the seven engineering constraints R1, R2, R3, R4′, R5, R6, R7, derivable from strengthened forms of A0–A4 and the K = 7 closure architecture) and a *continuum-convergence layer* (the companion paper's H6, H6′, H7, H8, H9). R4′ is stated as a bound on a substrate-engineering observable — the refinement-distortion functional $\Xi_\ell$ — that can in principle be computed directly on any concrete TPB substrate model; the discrete Hölder cone stability of $\mathcal{C}_\ell$ used downstream is derived from this bound (Lemma 4.1) rather than postulated. R7 (polynomial volume growth) is a strengthening required to deliver doubling: R1 alone gives at most exponential volume growth (the Δ-regular-tree regime), and doubling is a polynomial-growth property, so R7 is genuinely a separate engineering constraint rather than a consequence of R1 + R2. With this paper in place, the residual risk in the VERSF geometry programme is concentrated in three named open problems: (i) the substrate-level derivation of R1–R7 from the underlying VERSF closure-architecture papers in fully tightened form — particularly the conjecture that the K = 7 closure architecture entails R4′ with $\alpha = 1$ (Lipschitz refinement distortion), which would eliminate the §10.10 weak-norm-robustness gap of the companion paper entirely; (ii) the topological question of H9, which intersects 4-manifold theory and is not a dynamical condition; and (iii) the weak-norm-robustness gap flagged in §10.10 of the companion paper. The geometry-emergence chain is therefore now *substrate-engineering conditional* rather than *regularity-hypothesis conditional*.

We close by stating, as the principal next-paper task, the construction of explicit TPB substrate models satisfying R1–R7 with quantitative constants ($\Delta$, M, $\sigma_*$, $\alpha$, K, the bi-Lipschitz constants ($c_{\mathcal{R}}$, $C_{\mathcal{R}}$), and the Ahlfors-regularity constants (d, $c_{vol}$, $C_{vol}$)), since better constants — particularly $\alpha = 1$ in R4′ and the conjecture that K = 7 + R4′ entail two-sided Ahlfors d-regularity (R7) — substantially shorten the §10 chain of the companion paper and could potentially eliminate the wave-equation correction step entirely.

---

# Notation glossary

Notation extends the companion *Continuum-Limit Regularity* paper. New symbols introduced in this paper:

| Symbol | Meaning |
|---|---|
| R1 | Uniform bounded local degree (engineering constraint, §3.1) |
| R2 | Controlled refinement multiplicity (engineering constraint, §3.2) |
| R3 | Uniform distinguishability separation (engineering constraint, §3.3) |
| R4′ | Bounded refinement distortion (engineering constraint, §3.4); stated as a Hölder bound on the refinement-distortion functional $\Xi_\ell$ |
| R5 | Bounded refinement overlap on children (engineering constraint, §3.5) — promoted from a technical hypothesis to a named substrate-engineering constraint |
| R6 | Refinement bi-Lipschitz on cone-direction sets (engineering constraint, §3.6); supplies the inverse-refinement Lipschitz constant $c_{\mathcal{R}}^{-1}$ required by Lemma 4.1 |
| R7 | Polynomial volume growth in pseudometric balls — uniform Ahlfors d-regularity (engineering constraint, §3.7); strengthens R1 + R2 to deliver doubling (R1 alone gives at most exponential growth, the Δ-regular-tree regime) |
| A1* | Strengthened finite distinguishability with uniform lower bound (introduced in §4.1) |
| A2* | Strengthened irreversible commitment with operational-distinctness clause (introduced in §4.2) |
| $\Delta$ | Uniform bound on A4-local-coupling degree at every refinement level |
| M | Uniform bound on refinement multiplicity per node per refinement step |
| $\sigma_{op}$ | Uniform operational floor on distinguishability across refinement levels (from A1*) |
| $\sigma_*$ | Continuum lower bound on rescaled distinguishability scale (from R3); satisfies $\sigma_* \geq \sigma_{op}$ |
| K, $\alpha$ | Constants in the refinement-distortion bound (from R4′) |
| $c_{\mathcal{R}}$, $C_{\mathcal{R}}$ | Bi-Lipschitz constants for the refinement functor's action on cone-direction sets in Hausdorff distance (from R6): $c_{\mathcal{R}} \cdot d_H(\mathcal{C}, \mathcal{C}') \leq d_H(\mathcal{R}(\mathcal{C}), \mathcal{R}(\mathcal{C}')) \leq C_{\mathcal{R}} \cdot d_H(\mathcal{C}, \mathcal{C}')$ |
| d, $C_{vol}$, $c_{vol}$ | Polynomial-growth dimension and two-sided Ahlfors constants (from R7): $c_{vol} \cdot r^d \leq$ #(operationally distinguishable nodes in rescaled ball of radius r) $\leq C_{vol} \cdot r^d$ |
| $\Xi_\ell(x, y)$ | Refinement-distortion functional (§8.2): $d_H(\mathcal{R}_{\ell\to\ell+1}(\mathcal{C}_\ell(x)), \mathcal{R}_{\ell\to\ell+1}(\mathcal{C}_\ell(y)))$ |
| $K_{eff}$ | Effective constant in the derived Hölder cone-stability bound (Lemma 4.1): $K_{eff} := c_{\mathcal{R}}^{-1} \cdot K$ |

| Symbol | Meaning |
|---|---|
| $K_{\mathcal{C}}$ | Continuum Hölder constant for the limiting cone field $\mathcal{C}$ (Theorem 4): $K_{\mathcal{C}} = K_{eff} \cdot C_{GH}$, with $C_{GH}$ the Lorentzian-GH-embedding distortion constant inherited from the companion paper's Theorem 6 |
| $\mathcal{C}_\ell(x)$, $\mathcal{C}(x)$ | Discrete and continuum admissible-direction sets (as in companion paper) |
| $d_H$ | Hausdorff distance on the space of admissible direction sets |
| $Ext(\mathcal{C})$ | Set of extremal directions of a convex cone $\mathcal{C}$ |
| $dim_{comb}(\mathcal{C})$ | Combinatorial dimension of a convex cone (number of extremal generators) |
| $\omega(r)$ | Refinement modulus of continuity, $\omega(r) = K \cdot r^\alpha$ |
| K = 7 closure catalogue | The finite set of admissible local closure configurations established in the sequential-interface-transport paper |
| $K_{max}$ | Uniform cardinality bound on the local closure catalogue $\mathcal{K}_\ell$; $K_{max} = 7$ for VERSF, established in the sequential-interface-transport paper |

All other notation as in the companion paper and prior VERSF papers.

---

# Table of Contents


1. Introduction
2. Inputs from the companion paper
3. Definition: Regular TPB substrate
4. Substrate-axiom strengthening: A1 → A1* and A2 → A2*
5. Theorem 1 — Uniform Ahlfors d-regularity implies uniform doubling (H6)
6. Theorem 2 — Finite closure catalogue implies bounded combinatorial dimension (H7)
7. Theorem 3 — Finite distinguishability plus irreversibility implies transport sparsity (H6′)
8. Theorem 4 — Closure-stable refinement implies Hölder cone stability (H8)
9. Why H9 is structurally different
10. Theorem 5 — Main synthesis
11. Architectural consequences
12. Falsification paths and quantitative tests
13. Open problems and dependencies
14. Conclusion


---

# 1. Introduction

The companion paper *Continuum-Limit Regularity and Cone Convergence in VERSF* established the following conditional theorem (combining its Theorems 5 and 8):

*Given* an admissible TPB refinement sequence satisfying H1–H9 and BCB, the continuum limit ($\mathcal{M}$_∞, d_∞, ≺_∞, J^μ) is a strongly causal Lorentzian length space admitting a C^k Lorentzian metric for every prescribed k ≥ 2.

Five of the nine hypotheses (H6, H6′, H7, H8, H9) were introduced as external regularity hypotheses on the refinement family. The companion paper's §14 explicitly identified the substrate-level derivation of these five hypotheses as the principal remaining mathematical risk in the VERSF geometry programme. Its §15 conclusion named the substrate-engineering paper — exhibiting discrete substrate models for which the five hypotheses hold — as the natural next work-item.

The present paper supplies the substrate-engineering layer for four of those five hypotheses. The fifth, H9, is structurally different and is treated separately (§9).

The central question is:

Which substrate-level engineering constraints, applied to admissible TPB refinement dynamics, are *jointly sufficient* to entail H6, H6′, H7, H8 at the refinement-sequence level?

Our answer factors into seven engineering constraints — bounded local coupling, controlled refinement multiplicity, uniform distinguishability separation, bounded refinement distortion, bounded refinement overlap, refinement bi-Lipschitz on cone-direction sets, and uniform Ahlfors d-regularity in pseudometric balls — formalised below as R1, R2, R3, R4′, R5, R6, R7. The main theorem (§10) shows that these seven constraints, in conjunction with the baseline VERSF axioms A0–A4 (with A1 strengthened to A1* and A2 strengthened to A2*), BCB, and the companion paper's H1 and H5, entail H6, H6′, H7, H8.

A reader tracking the open-problem landscape should note that the *kind* of residual risk has changed. The companion paper's residual risk was distributed across five named external regularity hypotheses, each open to substrate-level derivation. After this paper, four of the five are derived from seven named substrate-engineering constraints, and the residual risk concentrates in three places: (i) the tighter substrate-level grounding of R1–R7 in the prior VERSF closure-architecture papers (sequential-interface-transport, σ-duality, admissible-coarse-graining), which we sketch but do not fully discharge — particularly the conjecture (§8.7) that the K = 7 closure architecture entails R4′ with α = 1; (ii) the topological H9, which is properly outside the dynamical scope of this paper; and (iii) the weak-norm-robustness gap in §10.10 of the companion paper, which is properly a revision to the *companion-companion* (Lorentzian-emergence) paper rather than to either continuum-regularity paper.

The mathematical-risk concentration has therefore moved one further layer down — from *which substrate models satisfy H6–H9* (companion paper, open) to *which substrate models satisfy R1–R7 + H9* (this paper, partially answered) to *which substrate models satisfy the underlying VERSF closure-architecture conditions that imply R1–R7* (next paper, open).

---

## 2. Inputs from the companion paper

We restate the five external regularity hypotheses from the companion paper *verbatim*, with the original §2.2 numbering preserved for cross-reference.

- **H6 (Uniform local doubling).** *There exists $D < \infty$ such that for every $\ell \in \mathbb{N}$, every $x \in P_\ell$, and every $r > 0$, the transport-pseudometric ball $B_\ell(x, 2r)$ admits a covering by at most D balls of radius r.*
- **H6′ (Uniform local transport sparsity).** *There exists $\sigma > 0$, independent of refinement level $\ell$, such that for every admissible commitment-transport sequence $\gamma_\ell = (s_0, s_1, \ldots, s_n)$ at level $\ell$, any two distinct committed states $s_i \neq s_j$ along $\gamma_\ell$ satisfy $d_\ell(s_i, s_j) \geq \sigma_\ell$, where $\sigma_\ell$ is the minimum distinguishability scale at level $\ell$, and the rescaled quantity $\tilde{\sigma}_\ell := \lambda_\ell \cdot \sigma_\ell$ has a strictly positive continuum lower bound $\liminf_{\ell\to\infty} \tilde{\sigma}_\ell = \sigma_- > 0$.**
- **H7 (Bounded combinatorial dimension).** *There exists $N < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x \in P_\ell$, the number of admissible transport directions at x is bounded by N.*
- **H8 (Refinement Hölder compatibility).** *There exist $\alpha > 0$ and $K < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x, y \in P_\ell$, the discrete cone-direction sets $\mathcal{C}_\ell(x)$, $\mathcal{C}_\ell(y)$ satisfy $d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \leq K \cdot d_\ell(x, y)^\alpha + \varepsilon_\ell$ with $\varepsilon_\ell \to 0$.*
- **H9 (Smooth-structure existence).** *The topological 4-manifold structure inherited by $(\mathcal{M}_\infty, d_\infty)$ from companion-paper Theorem 1 admits at least one compatible $C^\infty$ atlas.*

This paper derives H6, H6′, H7, H8 from substrate-level engineering constraints R1, R2, R3, R4′, R5, R6, R7 (introduced in §3) plus strengthened forms of the baseline VERSF axioms. H9 is structurally different (topological rather than dynamical) and is treated separately in §9.

---

## 3. Definition: Regular TPB substrate

A *regular TPB substrate* is a refinement family

$\{P_\ell, T_\ell, B_\ell, \mathcal{R}_{\ell\to\ell+1}\}_{\ell\in\mathbb{N}}$

satisfying:

- the baseline VERSF substrate axioms A0, A1*, A2*, A3, A4, and BCB (with A1* and A2* the strengthened forms of A1 and A2 introduced in §4);
- the companion paper's H1 (finite invariant propagation) and H5 (refinement compatibility);
- the seven engineering constraints R1–R7 introduced below.

The constraints split naturally into local-structure constraints (R1, R2 — about the substrate graph and refinement at a single node), distinguishability-preservation constraints (R3, R5 —

about how operational distinctness propagates along transport sequences and under refinement), refinement-geometry constraints on cone-direction sets (R4′, R6 — bi-directional control of the refinement functor's Hausdorff action), and a global volume-growth constraint (R7 — required to deliver doubling, since local-structure constraints alone permit exponential growth).

### 3.1 R1 — Uniform bounded local degree

*There exists $\Delta < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x \in P_\ell$, the number of A4-local-coupling neighbours of x in $P_\ell$ is at most $\Delta$.*

R1 is the engineering strengthening of A4 (local coupling). A4 by itself requires that interactions propagate only between neighbouring admissible states; R1 adds that the number of such neighbours is uniformly bounded *across all refinement levels*. The bound Δ is a substrate-engineering parameter — its value depends on the concrete substrate model and on the $K = 7$ closure architecture of the *sequential-interface-transport* paper.

### 3.2 R2 — Controlled refinement multiplicity

*There exists $M < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x \in P_\ell$, the refinement map $\boldsymbol{\mathcal{R}}_{\ell\to\ell+1}$ sends x to at most M locally distinguishable child nodes at level $\ell+1$, after quotienting by operational equivalence.*

R2 is the engineering strengthening of the local-degree-control clause in companion-paper §3.2 (which only required uniform-in-ℓ bounded degree, not uniform-in-ℓ bounded refinement multiplicity). The phrase "locally distinguishable, after quotienting by operational equivalence" is operationally important: refinement may technically produce many child nodes per parent, but the *operationally relevant* count is the number of nodes that can be distinguished by some admissible measurement. R2 bounds this latter count.

### 3.3 R3 — Uniform distinguishability separation

*There exists a constant* $\sigma\_ > 0$ such that for every $\ell \in \mathbb{N}$ and every pair of distinct committed substrate states $s, s' \in P_\ell$ lying on a common admissible commitment-transport sequence, the rescaled metric satisfies $\lambda_\ell \cdot d_\ell(s, s') \geq \sigma\_$.

R3 is the engineering form of H6′ — but stated as a substrate-level constraint on what counts as an admissible refinement family, rather than as a continuum-limit hypothesis. The relationship between R3 and the strengthened A1* is developed in §7: A1* gives the *operational* lower bound on distinguishability, and R3 promotes it to a *geometric* lower bound on the rescaled metric distance between distinct committed states along admissible transport sequences.

### 3.4 R4′ — Bounded refinement distortion

*Define the local refinement-distortion functional*

$$\Xi_\ell(x, y) := d_H(\boldsymbol{\mathcal{R}}_{\ell\to\ell+1}(\mathcal{C}_\ell(x)), \boldsymbol{\mathcal{R}}_{\ell\to\ell+1}(\mathcal{C}_\ell(y))),$$

*where $\mathcal{R}_{\ell\to\ell+1}$ is the refinement functor, $\mathcal{C}_\ell(\cdot)$ the discrete admissible-transport-direction set, and $d_H$ the Hausdorff distance on cone-direction sets. Intuitively, $\Xi_\ell(x, y)$ measures how differently the substrate refines two nearby regions $x, y \in P_\ell$. The substrate-engineering hypothesis is:*

*There exist constants $\alpha > 0$ and $K < \infty$ such that for all sufficiently large refinement levels $\ell \in \mathbb{N}$ and every pair $x, y \in P_\ell$,*

$$\Xi_\ell(x, y) \le K \cdot d_\ell(x, y)^\alpha + \varepsilon_\ell,$$

*with $\varepsilon_\ell \to 0$ as $\ell \to \infty$.*

R4′ is the engineering form of "refinement is not chaotic," stated as a quantitative Hölder bound on the substrate-engineering observable $\Xi_\ell$. The constants K and α depend on the substrate's closure-stability properties under refinement, particularly on whether nearby substrate regions refine through equivalent $K = 7$ closure catalogues. Closure stability is the engineering content of "refinement is not chaotic"; bounded refinement distortion is its quantitative form.

The Hölder cone-stability of $\mathcal{C}_\ell$ used downstream is not a hypothesis here. It is a *derived intermediate result* (Lemma 4.1 in §8.3), obtained from R4′ together with R5 (bounded refinement overlap, §3.5), R6 (refinement bi-Lipschitz on cones, §3.6), and H5. The substrate-engineering content is the bound on $\Xi_\ell$ itself, a measurable property of the refinement functor's action on cone-direction sets.

## 3.5 R5 — Bounded refinement overlap

*There exists $\sigma_{op} > 0$, uniform in $\ell \in \mathbb{N}$, such that for every parent $x \in P_\ell$ and every pair of distinct children $y, y' \in \mathcal{R}_{\ell\to\ell+1}(x)$, the rescaled metric distance satisfies*

$$\lambda_{\ell+1} \cdot d_{\ell+1}(y, y') \ge \sigma_{op}.$$

R5 is the refinement-step version of A1*: distinct children of a parent under refinement remain separated above the same operational floor $\sigma_{op}$ that bounds distinct-state separations within any single level. Without R5, the $\sigma_{op}$ floor at level $\ell$ would not propagate to level $\ell+1$, and the chain of operational distinguishability would break across refinement.

R5 was introduced in earlier drafts as a "technical hypothesis" of Theorem 3 and Lemma 4.1 (then Theorem 4a). The substrate-engineering content is the same as those proofs require — distinct children stay distinct above $\sigma_{op}$ — but the role is large enough that R5 deserves a named slot among the engineering constraints rather than being relegated to a technical footnote. R5 is used directly in Theorem 3 (via §7.4) and as a component of Lemma 4.1's inverse-refinement injectivity argument (via §8.3).

## 3.6 R6 — Refinement bi-Lipschitz on cone-direction sets

R6 is most naturally read as a *one-sided* engineering constraint: a uniform lower bi-Lipschitz constant $c_{\mathcal{R}} > 0$ for the refinement functor's action on cone-direction sets, with the upper constant arising as a consequence of R4′ + H5 + R2.

*Formal statement.* Provided R5 holds (so that the Hausdorff metric $d_H$ is a genuine metric, not a degenerate pseudo-metric, on the space of post-refinement cone-direction sets), there exist constants $c_{\mathcal{R}}, C_{\mathcal{R}} \in (0, \infty)$, uniform in $\ell$, such that for every $\ell \in \mathbb{N}$ and every pair $x, y \in P_\ell$, the position-indexed cone-direction sets $\mathcal{C} = \mathcal{C}_\ell(x)$, $\mathcal{C}' = \mathcal{C}_\ell(y)$ satisfy

$$c_{\mathcal{R}} \cdot d_H(\mathcal{C}, \mathcal{C}') \leq d_H(\mathcal{R}(\mathcal{C}), \mathcal{R}(\mathcal{C}')) \leq C_{\mathcal{R}} \cdot d_H(\mathcal{C}, \mathcal{C}').$$

This position-indexed quantification is exactly what Lemma 4.1 invokes. We do not assert the broader form (the bi-Lipschitz inequality on *every* pair of admissible cone-direction sets at level $\ell$ regardless of base point), since the proofs in this paper do not require it. A broader form, if needed for downstream substrate-engineering arguments, would be a strictly stronger hypothesis whose derivation from σ-duality would correspondingly require more work.

*The R5 precondition.* R5 (bounded refinement overlap, §3.5) guarantees that distinct level-($\ell$+1) children of a parent remain operationally distinguishable above the $\sigma_{op}$ floor. Without R5, distinct post-refinement cone-direction sets could fail to be operationally distinguishable, and $d_H$ on the space of post-refinement cone-direction sets would degenerate to a pseudo-metric in which R6's lower bound $c_{\mathcal{R}} > 0$ would be vacuous (one could artificially make $d_H(\mathcal{R}(\mathcal{C}), \mathcal{R}(\mathcal{C}'))$ arbitrarily small by identifying operationally indistinguishable post-refinement directions). R5 secures the operational floor on which $d_H$ is a genuine metric on cone-direction sets, and R6's bi-Lipschitz inequality then becomes a substantive statement.

*Upper constant $C_{\mathcal{R}}$ is essentially derived.* The upper inequality $d_H(\mathcal{R}(\mathcal{C}), \mathcal{R}(\mathcal{C}')) \leq C_{\mathcal{R}} \cdot d_H(\mathcal{C}, \mathcal{C}')$ follows from R4′ specialised to the case where the level-$\ell$ cones in question are themselves $d_H$-close: R4′ bounds $\Xi_\ell(x, y)$ by a Hölder modulus of $d_\ell(x, y)$, and H5 + R2 + bounded local degree together control the $d_H$-comparability of level-$\ell$ cones at nearby base points. R6's upper bound is therefore best read as a consequence of R4′ + H5 + R2 rather than as an independent engineering constraint.

*Lower constant $c_{\mathcal{R}} > 0$ is the substantive content.* The substantive engineering content of R6 is the lower bound $c_{\mathcal{R}} > 0$ on the refinement functor's action — the statement that $\mathcal{R}$ cannot *collapse* cone-direction-set differences arbitrarily; nearby post-refinement cones must reflect quantitatively nearby pre-refinement cones, with no arbitrary loss of resolution under refinement. This is precisely the inverse-refinement bound $c_{\mathcal{R}}^{-1} = C_{\mathcal{R}}^{inv}$ that Lemma 4.1's Step (iii) requires.

*Conjectural derivation.* The natural conjecture (§13) is that $c_{\mathcal{R}} > 0$ follows from the σ-duality refinement-functor structure together with R2 + R5, but this is not established here; bi-Lipschitz on cone-direction sets is genuinely an additional substrate-engineering constraint and is not derivable from R2 + R5 alone (which together give injectivity but not a uniform inverse-Lipschitz constant).

### 3.7 R7 — Polynomial volume growth in pseudometric balls (Ahlfors d-regularity)

*There exist constants $d < \infty$, $C_{vol} < \infty$, and $c_{vol} > 0$, uniform in $\ell$, such that for every $\ell \in \mathbb{N}$, every $x \in P_\ell$, and every $r > 0$, the number of operationally distinguishable nodes in the rescaled pseudometric ball $B_\ell(x, r)$ at level $\ell$ satisfies*

$$c_{vol} \cdot r^d \leq \#\{ y \in P_\ell : \lambda_\ell \cdot d_\ell(x, y) \leq r, \text{ y operationally distinguishable } \} \leq C_{vol} \cdot r^d.$$

R7 is the engineering form of "the substrate is uniformly Ahlfors d-regular under refinement." The upper bound bounds volume growth from above (polynomial, not exponential); the lower bound bounds it from below (the substrate cannot be arbitrarily sparse at small scales). Both bounds are required for the standard polynomial-growth-implies-doubling argument: the upper bound caps the volume of large balls, and the lower bound supplies an r-uniform packing density inside small balls.

R7 is genuinely an additional engineering constraint, not derivable from R1 + R2 alone. R1 + R2 control the *local* density of nodes at each scale, but they do not prevent the substrate from having a *tree-like* global structure in which ball volume grows exponentially. The standard counterexample is the Δ-regular tree: each node has Δ neighbours (R1 holds trivially), refinement multiplicity is bounded (R2 holds trivially), but the volume of a ball of radius r grows like $\Delta^r$ (exponentially, not polynomially). Trees fail doubling; doubling requires polynomial volume growth. R7's upper bound rules out tree-like substrates; the lower bound rules out substrates whose small-ball packing density degenerates as $r \to 0$ or as $\ell \to \infty$, since such degeneracy would defeat the r-uniform packing argument that underlies doubling. R7 together selects the lattice-like / manifold-like architectures needed for finite-dimensional continuum limits.

The conjecture (§13) is that the K = 7 closure architecture combined with R4′ entails R7 — that finite-closure refinement dynamics with bounded refinement distortion cannot produce tree-like volume growth, and that bounded local degree combined with closure-stable refinement enforces a uniform lower density at small scales. This is open. For now R7 must be assumed as a separate engineering constraint.

R7 supplies the Ahlfors-d-regularity input that Theorem 1 (§5) requires, via the standard result that uniformly Ahlfors d-regular metric measure spaces are doubling (Coifman–Weiss 1971; cf. Heinonen, *Lectures on Analysis on Metric Spaces*, 2001, §3, where Ahlfors d-regular spaces are shown to be doubling with constant depending only on ($C_{vol}$, $c_{vol}$, d)).

### 3.8 Status of R1–R7

R1 strengthens A4. R2 strengthens the local-degree-control clause of companion-paper §3.2. R3 promotes A1* to a metric statement on admissible transport sequences. R4′ states a Hölder bound on the substrate-engineering observable $\Xi_\ell$. R5 promotes bounded refinement overlap from a technical hypothesis to a named engineering constraint propagating the $\sigma_{op}$ floor across refinement steps. R6 states refinement bi-Lipschitz on cone-direction sets, supplying the inverse-refinement Lipschitz constant required by Lemma 4.1. R7 is a polynomial-volume-growth constraint required to deliver doubling, ruling out the Δ-regular-tree regime that R1 + R2 alone admit.

In each case, the engineering constraint is *implicit* in prior VERSF papers but is made *explicit and quantitative* here. The substrate-engineering question — which concrete TPB substrate models satisfy R1–R7 with what constants — is sketched in §13 and is the principal next-paper task. Several entries are conjecturally derivable from finer substrate-architectural conditions (notably R6 from σ-duality, R7 from K = 7 + R4′); these conjectures are flagged individually in §13.

---

# 4. Substrate-axiom strengthening: A1 → A1* and A2 → A2*

Two of the baseline VERSF substrate axioms — A1 (finite distinguishability) and A2 (irreversible commitment) — as stated in prior VERSF papers carry the operationally correct content at each fixed refinement level but do not directly entail the uniformly-in-ℓ statements required by the substrate-engineering arguments of this paper. We therefore introduce two surgical strengthenings: A1 → A1* (uniform operational floor) and A2 → A2* (committed-state distinctness coincides with operational distinctness). Both strengthenings are conservative — they make explicit content that A1 and A2 already imply at each fixed level, but promote it to a uniform-across-levels statement that is required for the continuum limit to preserve substrate-level operational reality.

## 4.1 A1* — Finite distinguishability with uniform operational lower bound

*There exists $\sigma_{op} > 0$, independent of refinement level ℓ, such that for every $\ell \in \mathbb{N}$ and every pair of operationally distinct substrate states $s, s' \in P_\ell$, the rescaled metric distance satisfies $\lambda_\ell \cdot d_\ell(s, s') \geq \sigma_{op}$.*

A1* says that two states are operationally distinguishable only if they are separated by at least $\sigma_{op}$ in the rescaled continuum metric. The constant $\sigma_{op}$ is the operational floor on distinguishability — below this scale, no admissible measurement can distinguish the states.

A1 alone permits the minimum distinguishability scale $\sigma_\ell$ to shrink under refinement in such a way that the rescaled quantity $\tilde{\sigma}_\ell = \lambda_\ell \cdot \sigma_\ell$ tends to zero. In that case, the distinction between committed substrate states would fail to propagate to the operationally accessible coarse-graining at the continuum scale: pairs of states distinguishable at fine refinement levels would fall below the operational floor of any coarse-graining and become operationally indistinguishable in the limit, leaving "irreversible commitment" (A2) without operational content at the continuum scale. A1* prevents this degeneracy by enforcing a uniform operational floor that survives coarse-graining.

The conceptual point is that A1 is a *local* finiteness condition (finitely many distinguishable states at each level) while A1* is a *uniform* finiteness condition (finite operational floor uniform in ℓ). For the continuum limit to preserve the operational reality of commitment, A1* is required.

## 4.2 A2* — Irreversible commitment with operational distinctness

*A2 holds, and additionally: for every refinement level $\ell \in \mathbb{N}$, any two distinct committed substrate states $s, s' \in P_\ell$ correspond to operationally distinguishable facts — i.e., the equivalence relation of operational indistinguishability at level $\ell$ cannot identify states arising from distinct admissible commitment events.*

A2* makes explicit the operational content that A2 already carries. A2's substantive claim is that commitment events produce *new facts*; A2* fixes the substrate-level criterion for what "new fact" means at the operational level. The justification is the following structural argument internal to A2's claim:

A2 asserts that an admissible commitment event produces a new committed state distinct from the prior state, and that this distinction is irreversible. The substantive content of "new fact" requires that the distinction between the prior and the new states be operationally accessible — otherwise the commitment event has produced nothing operationally new, and A2's existence claim is vacuous. Hence A2's own internal logic forces the operational-distinctness clause: a commitment event that produced an operationally indistinguishable outcome would not, by A2's own criterion, be a commitment event at all.

A2* makes this internal logic explicit at the level of the axiom, parallel to A1*'s explicit-uniform-floor strengthening of A1. The pair (A1*, A2*) together encode the substrate-level operational reality of distinguishability and commitment in the form required by the substrate-engineering arguments of §7.

## 4.3 Relationship between A1*, A2*, R3, and H6′

A1* gives a uniform lower bound on *operational* distinguishability between *any* pair of distinct states. A2* says that distinct committed states correspond to operationally distinguishable facts at each level. R3 promotes the A1* + A2* conjunction to a uniform lower bound on *geometric* separation between *committed* states *along a common admissible transport sequence*. H6′ is the continuum-limit statement that this geometric separation persists in the limit.

The implication chain is:

A1* + A2* + R5 ⇒ R3 ⇒ H6′.

The middle step (A1* + A2* + R5 ⇒ R3) is the substantive content of Theorem 3 (§7), proved in two stages: A1* + A2* gives the within-level $\sigma_{op}$ lower bound (via the direct contradiction between operational indistinguishability at the $\sigma_{op}$ scale and committed-state distinctness), and R5 propagates the bound across refinement levels. The R3 ⇒ H6′ step is essentially a re-reading: R3 is the substrate-level form of H6′.

---

# 5. Theorem 1 — Polynomial volume growth implies uniform doubling (H6)

## 5.1 Statement

**Theorem 1 (Uniform Ahlfors d-regularity implies uniform doubling).** *Under R7 (two-sided Ahlfors d-regularity in pseudometric balls with constants $c_{vol}$, $C_{vol}$, $d$), every regular TPB refinement sequence satisfies H6: there exists $D < \infty$ such that for every $\ell \in \mathbb{N}$, every $x \in P_\ell$, and every $r > 0$, the rescaled pseudometric ball $B_\ell(x, 2r)$ admits a covering by at most D balls of radius r. The doubling constant satisfies the explicit r-uniform bound $D \leq (C_{vol} / c_{vol}) \cdot 12^d$, depending only on the engineering constants of R7.*

## 5.2 Proof

The proof is the standard Ahlfors-regularity-implies-doubling argument from analysis on metric measure spaces (Coifman–Weiss 1971; cf. Heinonen, *Lectures on Analysis on Metric Spaces*, 2001, §3, where Ahlfors d-regular spaces are shown to be doubling with a constant depending only on the regularity constants).

Fix $\ell \in \mathbb{N}$, $x \in P_\ell$, and $r > 0$. Work throughout in rescaled coordinates so that the rescaled pseudometric is $\rho_\ell := \lambda_\ell \cdot d_\ell$.

*Step 1 — Cover $B_\ell(x, 2r)$ by r-balls via a maximal r/2-separated set.* By a standard greedy construction, choose a maximal r/2-separated set $\{ y_1, \ldots, y_N \} \subset B_\ell(x, 2r)$ of operationally distinguishable nodes — i.e., $\rho_\ell(y_i, y_j) \geq r/2$ for $i \neq j$, and every operationally distinguishable node in $B_\ell(x, 2r)$ lies within $\rho_\ell$-distance $r/2$ of some $y_i$. Then the r-balls $B_\ell(y_i, r)$ cover $B_\ell(x, 2r)$:

$$B_\ell(x, 2r) \subseteq \cup_{i=1}^{N} B_\ell(y_i, r).$$

The covering number is at most N. It suffices to bound N r-uniformly.

*Step 2 — Packing argument using two-sided R7.* The points $\{y_i\}$ are r/2-separated, so the balls $B_\ell(y_i, r/4)$ are pairwise disjoint. Each $B_\ell(y_i, r/4)$ is contained in the enlarged ball $B_\ell(x, 2r + r/4) \subseteq B_\ell(x, 3r)$. Counting operationally distinguishable nodes and using disjointness:

$$\sum_{i=1}^{N} \#(\text{operationally distinguishable nodes in } B_\ell(y_i, r/4)) \leq \#(\text{operationally distinguishable nodes in } B_\ell(x, 3r)).$$

Now apply R7 in both directions. The *lower* bound of R7, applied to each $B_\ell(y_i, r/4)$, gives

$$\#(\text{operationally distinguishable nodes in } B_\ell(y_i, r/4)) \geq c_{vol} \cdot (r/4)^d.$$

The *upper* bound of R7, applied to $B_\ell(x, 3r)$, gives

$\#(\text{operationally distinguishable nodes in } B_\ell(x, 3r)) \leq C_{vol} \cdot (3r)^d$.

Combining:

$N \cdot c_{vol} \cdot (r/4)^d \leq C_{vol} \cdot (3r)^d$.

*Step 3 — Cancel r-dependence.* Divide through by $c_{vol} \cdot (r/4)^d$. The r-dependence cancels identically:

$N \leq (C_{vol} / c_{vol}) \cdot (3r)^d / (r/4)^d = (C_{vol} / c_{vol}) \cdot (12)^d$.

This bound is r-uniform: it depends only on the Ahlfors constants ($C_{vol}$, $c_{vol}$, d), not on r, x, or ℓ.

*Step 4 — Conclusion.* The covering number of $B_\ell(x, 2r)$ by r-balls is at most $D := (C_{vol} / c_{vol}) \cdot 12^d$, uniformly in ℓ, x, r. This is precisely H6 with explicit doubling constant $D = D(C_{vol}, c_{vol}, d)$.

This completes the proof of Theorem 1.

### 5.3 Why R1 + R2 alone are insufficient

The previous draft of this theorem attempted to derive doubling from R1 (bounded local degree) and R2 (controlled refinement multiplicity) alone via a node-density argument. That argument is incorrect: R1 + R2 give at most *exponential* node-density growth (bounded out-degree at each coupling step, iterated through the propagation radius), and exponential growth is consistent with the failure of doubling. The canonical counterexample is the Δ-regular tree: every node has Δ neighbours (R1 holds trivially), refinement multiplicity is uniformly bounded (R2 holds trivially), but the ball of radius r contains $\sim\Delta^r$ nodes (exponential growth), and the tree is not doubling.

Doubling is a *polynomial-growth* property; it requires the substrate's volume profile to be polynomial in r (with a fixed polynomial degree d, the doubling dimension), not merely exponentially-bounded. R7 is the engineering constraint that supplies this polynomial-growth content. R1 + R2 alone select substrates with bounded local interaction structure but do not exclude tree-like global structure.

Note that the two-sided form of R7 is essential for the argument. R7's upper bound alone (the one-sided form) is consistent with arbitrarily sparse substrates whose small balls contain few or no operationally distinguishable nodes; in that case the packing argument of Step 2 collapses because the left-hand side $\sum \#(\ldots) \geq N \cdot c_{vol} \cdot (r/4)^d$ would have $c_{vol} \to 0$. The lower bound $c_{vol} > 0$ is the r-uniform packing-density input that makes the cancellation in Step 3 work. Substrate-engineering models discharging R7 must therefore exhibit *both* polynomial-bounded growth *and* uniform lower packing density.

The substrate-engineering question is therefore: *what substrate-architectural conditions deliver two-sided Ahlfors d-regularity?* The natural conjecture (§13) is that the K = 7 closure architecture combined with R4′ (bounded refinement distortion) entails R7 in both directions — that finite-closure refinement dynamics with bounded distortion cannot produce tree-like proliferation at coarse scales (upper bound) and cannot produce arbitrarily sparse refinement at small scales (lower bound, since each closure configuration generates at least one operationally distinguishable child). Establishing this conjecture is open and is a principal next-paper task. Until it is established, R7 must be assumed as a separate engineering constraint, and Theorem 1 is *proven conditional on R7* rather than on R1 + R2 alone.

### 5.4 Interpretation

H6 is the continuum-limit expression of the substrate rule:

The substrate's volume profile is Ahlfors d-regular in pseudometric radius, uniformly across refinement levels.

Equivalently, the substrate has a bounded doubling dimension, in the sense of analysis on metric measure spaces. This rules out tree-like substrate architectures (exponential volume growth, no doubling) and also arbitrarily sparse architectures (where small balls contain too few distinguishable nodes to anchor the packing argument), selecting lattice-like or manifold-like architectures whose continuum limits can carry finite-dimensional smooth structure.

### 5.5 Where the proof concentrates its risk

The risk is concentrated entirely in R7. If R7's upper bound fails — i.e., the substrate exhibits exponential or super-polynomial volume growth at some level — then doubling fails, H6 fails, and the precompactness conclusion of the companion paper's Theorem 1 fails as well. If R7's lower bound fails — i.e., the substrate becomes arbitrarily sparse at small scales — then the packing argument in Step 2 collapses, and even with polynomial-bounded growth from above, the doubling constant cannot be made r-uniform. *Illustrative failure mode.* A substrate whose refinement at small scales produces operationally distinguishable nodes only on a positive-codimension subset of pseudometric balls — e.g., concentrated on a fractal of Hausdorff dimension strictly less than d — would satisfy R7's upper bound with the full-dimensional d but violate the lower bound, since c_vol · r^d would exceed the actual node count in small balls. The substrate-engineering question is whether the K = 7 closure architecture, perhaps combined with R4′, automatically delivers two-sided R7 and rules out such fractal-concentration pathologies. We do not establish this here; it is one of the principal next-paper tasks for the substrate-engineering layer.

---

## 6. Theorem 2 — Finite closure catalogue implies bounded combinatorial dimension (H7)

## 6.1 Statement

**Theorem 2 (Finite closure catalogue implies bounded combinatorial dimension).** *Under R1 (uniform bounded local degree) and the $K = 7$ closure architecture of the sequential-interface-transport paper, every regular TPB refinement sequence satisfies H7: there exists $N < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x \in P_\ell$, the number of admissible transport directions at x is bounded by N.*

## 6.2 The K = 7 closure architecture

The *sequential-interface-transport* paper establishes that admissible local transport on a TPB substrate is generated by a finite catalogue of *closure configurations* — admissible local arrangements of substrate states that support transport.

*Brief recap (for readers without the sequential-interface-transport paper to hand).* Closure configurations are the admissible local "transport-channel templates" — discrete patterns of substrate states arranged around a node x such that the pattern can support an admissible commitment-transport step out of x. The *sequential-interface-transport* paper derives, from the substrate-level commitment dynamics combined with the σ-duality refinement structure, that the admissible closure configurations at any node form a finite catalogue $\mathcal{K}_\ell$ with $|\mathcal{K}_\ell| \leq 7$. The specific count $K_{max} = 7$ arises from the seven admissible positions on the K-wheel intertwiner structure (the *bigrading wheel* of σ-duality), with the seven positions corresponding to the distinct ways the commitment-transport tick can be operationally distinguished from its substrate context. *Notation caveat:* the "K" in "K-wheel" is the σ-duality wheel-structure label and is unrelated to $K_{max}$ except numerologically — both happen to take the value 7 because $K_{max}$ is the cardinality bound on closure positions on the K-wheel, but they are distinct quantities (the K-wheel is a structural object; $K_{max}$ is a cardinality of a closure catalogue defined on that object). Where the distinction matters in what follows we use $K_{max}$ for the cardinality and "K-wheel" or "K = 7 closure architecture" for the structural label. The substrate-engineering content of Theorem 2 is that the catalogue is *finite*; the specific value $K_{max} = 7$ reflects the K-wheel cardinality and is conditional on the sequential-interface-transport paper's substrate-level derivation. For the present theorem, only finiteness matters; the value 7 enters only when we identify the explicit constant in H7.

Formally:

- (Finite local closure catalogue) For every $\ell \in \mathbb{N}$ and every $x \in P_\ell$, the set of admissible local closure configurations at x is contained in a finite catalogue $\mathcal{K}_\ell$ of cardinality $|\mathcal{K}_\ell| \leq K_{max}$, with $K_{max} < \infty$ uniform in $\ell$. For VERSF this is $K_{max} = 7$.
- (Closure-generated transport) Every admissible local transport channel at x is generated by some configuration in $\mathcal{K}_\ell$.

## 6.3 Proof of Theorem 2

By the closure-generated transport clause, every admissible transport direction $v \in \mathcal{C}_\ell(x)$ at node x is generated by some closure configuration $\kappa \in \mathcal{K}_\ell$. Distinct generators $\kappa, \kappa' \in \mathcal{K}_\ell$ may produce distinct primitive transport directions, but the total number of *primitive* directions is at most $|\mathcal{K}_\ell| \leq K_{max}$.

The discrete admissible-direction set $\mathcal{C}_\ell(x)$ may contain non-primitive directions obtained by convex combination or by closure of primitive directions under admissible composition. However, these derived directions do not increase the *combinatorial dimension* of the cone — they are all contained in the convex hull (or appropriate closure) of the primitive generators:

$$\mathcal{C}_\ell(x) \subseteq \mathrm{conv}\{v_\kappa : \kappa \in \mathcal{K}_\ell\}.$$

The number of *extremal* directions of $\mathcal{C}_\ell(x)$ is therefore at most $|\mathcal{K}_\ell| \leq K_{max}$:

$$\#\mathrm{Ext}(\mathcal{C}_\ell(x)) \leq K_{max}.$$

For the purposes of H7, we identify $N := K_{max}$. Since H7 asks for a bound on the number of admissible transport directions (interpreted as extremal generators of the cone), the bound is established.

This completes the proof of Theorem 2.

## 6.4 Interpretation

H7 is the continuum-limit expression of the substrate rule:

The local closure catalogue is finite.

For VERSF the catalogue cardinality is $K_{max} = 7$. Other substrate architectures with finite closure catalogues would also satisfy H7, with $N = K_{max}$ for that architecture.

A substrate with an *infinite* closure catalogue — one that admits arbitrarily many primitive transport channels at each node — would violate H7 and would produce a continuum cone field with infinite combinatorial dimension. Such substrates would not converge to a finite-dimensional Lorentzian manifold in the companion paper's sense.

## 6.5 Where the proof concentrates its risk

The risk is concentrated in the $K = 7$ closure architecture itself, which is established in the *sequential-interface-transport* paper and inherited here. If that architecture is wrong — if substrate dynamics in fact admit more (or fewer, or different) admissible closures than $K = 7$ — then the constant $N$ in H7 changes, but the *finite-N* conclusion does not, so long as the closure catalogue remains finite. The structural conclusion of Theorem 2 is robust; only the explicit constant $N = K_{max}$ depends on the architecture.

The qualitative claim — *finite closure catalogue implies H7* — is therefore a robust substrate-engineering principle. The quantitative claim — *$N = 7$ for VERSF* — is inherited from the sequential-interface-transport paper and is conditional on that paper's result.

# 7. Theorem 3 — Finite distinguishability plus irreversibility implies transport sparsity (H6′)

## 7.1 Statement

**Theorem 3 (Finite distinguishability plus irreversibility implies transport sparsity).** *Under A1* (strengthened finite distinguishability with uniform operational lower bound $\sigma_{op} > 0$), A2* (irreversible commitment with operational distinctness), R3 (uniform distinguishability separation $\sigma_* > 0$), and R5 (bounded refinement overlap, §3.5), every regular TPB refinement sequence satisfies H6′ with $\sigma_* \geq \sigma_{op}$.*

## 7.2 Why A2 alone is not enough

A2 (irreversibility) forbids any admissible commitment-transport sequence to revisit the *same* committed state:

$s_i \neq s_j$ for all $i < j$ in $\gamma_\ell$.

This is the exact-acyclicity content of A2. However, A2 alone does not prevent two *distinct* committed states from being arbitrarily close in the rescaled continuum metric:

A2: $s_i \neq s_j \Longrightarrow$ no exact revisit. Want: $\lambda_\ell \cdot d_\ell(s_i, s_j) \geq \sigma_* > 0 \Longrightarrow$ no almost-revisit.

The gap between "no exact revisit" and "no almost-revisit" is the quantitative gap that H6′ closes. The substrate-engineering question is: what additional condition forces distinct committed states to remain geometrically separated? The §7.3 argument shows that A2* (committed-state distinctness coincides with operational distinctness) together with A1* (uniform operational floor) closes this gap at each fixed level — and R5 propagates the closure across refinement levels.

## 7.3 The reality-of-commitment argument

The within-level lower bound is a direct contradiction between A1* and A2*. By A1* (uniform operational floor), two states with rescaled distance below $\sigma_{op}$ are operationally indistinguishable. By A2* (committed-state distinctness coincides with operational distinctness), two distinct committed states cannot be operationally indistinguishable. The conjunction "operationally indistinguishable but committedly distinct" is therefore forbidden by A1* + A2*.

Formally:

*Proof of Theorem 3, first stage (within-level $\sigma_{op}$ lower bound).* Suppose for contradiction that $\gamma_\ell = (s_0, s_1, \ldots, s_n)$ is an admissible commitment-transport sequence at level $\ell$ containing distinct $s_i \neq s_j$ with $\lambda_\ell \cdot d_\ell(s_i, s_j) < \sigma_{op}$.

- By A1*, the operational floor $\sigma_{op}$ is the minimum scale at which any admissible measurement can distinguish two states. Two states with $\lambda_\ell \cdot d_\ell(s, s') < \sigma_{op}$ are therefore operationally indistinguishable.
- By A2*, distinct committed states are operationally distinguishable.
- The conjunction is a contradiction: $s_i$ and $s_j$ cannot simultaneously be operationally indistinguishable (by A1*) and operationally distinguishable (by A2*).

Therefore the hypothesised $\gamma_\ell$ does not exist, and every distinct pair $s_i \neq s_j$ on an admissible level-$\ell$ transport sequence satisfies

$$\lambda_\ell \cdot d_\ell(s_i, s_j) \geq \sigma_{op} > 0.$$

This is the within-level form of R3 with $\sigma_* := \sigma_{op}$. The proof is completed by the across-level propagation argument of §7.4 (which uses R5) and the conclusion in §7.5.

## 7.4 Across-level propagation via R5

The within-level bound of §7.3 holds at each fixed $\ell$ but does not by itself give a *uniform* ($\ell$-independent) lower bound. The pathology to rule out is that a single committed state at level $\ell$ might refine into multiple operationally distinguishable children at level $\ell+1$, two of which are visited consecutively in some level-($\ell+1$) transport sequence. If the children were arbitrarily close in the rescaled metric, the $\sigma_{op}$ lower bound at level $\ell$ would fail to propagate to level $\ell+1$.

R5 (bounded refinement overlap, §3.5) is exactly the engineering constraint that rules this out: for every parent $x \in P_\ell$ and every pair of distinct children $y, y' \in \mathcal{R}_{\ell\to\ell+1}(x)$,

$$\lambda_{\ell+1} \cdot d_{\ell+1}(y, y') \geq \sigma_{op}.$$

R5 is the refinement-step analogue of A1*. Under R5, the within-level $\sigma_{op}$ lower bound at level $\ell$ propagates to level $\ell+1$, and hence by induction to every level. (In earlier drafts R5 was referred to as a "bounded refinement overlap" technical hypothesis; this paper promotes it to a named engineering constraint, since its role in Theorem 3 and in Lemma 4.1 of §8 is on equal footing with R1, R2, R3, R4′.)

## 7.5 Conclusion of proof

Combining the within-level bound of §7.3 with the across-level propagation via R5 (§7.4) gives a uniform $\sigma_* \geq \sigma_{op} > 0$ such that for every $\ell$ and every distinct pair $s_i \neq s_j$ on an admissible commitment-transport sequence at level $\ell$,

$$\lambda_\ell \cdot d_\ell(s_i, s_j) \geq \sigma_*.$$

This is R3, which is the substrate-level form of H6′. Passing to the limit, $\liminf_{\ell\to\infty} \tilde{\sigma}_\ell = \sigma_* > 0$, which is the continuum statement of H6′.

This completes the proof of Theorem 3.

### 7.6 Interpretation

H6′ is the continuum-limit expression of the substrate rule:

The substrate cannot create irreversible facts that later collapse to operationally identical facts along the same causal trajectory.

This is the operational content of "no fake commitment." It is not derivable from A2 (exact acyclicity) alone — it requires the operational-distinctness clause of A2*, the uniform operational floor of A1*, and the propagation guarantee of R5.

A natural way to state the result: A1* + A2* + R5 ⟹ H6′, with the implication routed through the within-level contradiction (A1* + A2*) and the across-level propagation (R5).

### 7.7 Where the proof concentrates its risk

The risk is concentrated in three places:

1. *A1* itself.* If the operational floor $\sigma_{op}$ cannot be established uniformly in $\ell$ for a given substrate, the within-level argument collapses. The substrate-engineering question is whether physically motivated TPB substrates have such a floor.
2. *A2* itself.* If A2's "irreversible commitment" content does not in fact entail operational distinctness — i.e., if some substrate model admits commitment events whose outcomes are operationally indistinguishable from the prior state — A2* fails, and the within-level $\sigma_{op}$ lower bound cannot be established. The substrate-engineering question is whether the operational-distinctness clause is intrinsic to all physically motivated TPB substrate models or whether some closure architectures admit operationally vacuous commitment events.
3. *R5 (bounded refinement overlap).* If the refinement map can produce arbitrarily close child states for a parent, the floor at level $\ell$ does not propagate to level $\ell+1$. The substrate-engineering question is whether the refinement maps of the *σ-duality* paper satisfy R5.

All three risks are concrete and targetable — A1* and A2* concern operational structure of the substrate (uniform floor and operational distinctness of committed states), R5 concerns refinement-map properties; all reduce to specific properties of concrete substrate models.

---

# 8. Theorem 4 — Bounded refinement distortion implies Hölder cone stability (H8)

### 8.1 Statement

**Theorem 4 (Bounded refinement distortion implies Hölder cone stability).** *Under R4′ (bounded refinement distortion with constants $\alpha > 0$, $K < \infty$), R5 (bounded refinement overlap), R6 (refinement bi-Lipschitz on cone-direction sets with constants $c_{\mathcal{R}}$, $C_{\mathcal{R}} \in (0, \infty)$), and H5 (refinement compatibility), every regular TPB refinement sequence satisfies H8 with the same Hölder exponent $\alpha$ and continuum Hölder constant*

$$K_{\mathcal{C}} = K_{eff} \cdot C_{GH}, \quad K_{eff} := c_{\mathcal{R}}^{-1} \cdot K,$$

*where $C_{GH}$ is the Lorentzian-Gromov–Hausdorff-embedding distortion constant inherited from the companion paper's Theorem 6 proof.*

The proof proceeds in two stages: first, Lemma 4.1 (refinement-distortion propagation, §8.3) derives Hölder cone stability of the discrete sets $\mathcal{C}_\ell$ at each refinement level from R4′ + R6 (with R5 securing the metric setting); second, Theorem 4 lifts this discrete Hölder stability to the continuum H8 via Lorentzian-GH convergence and absorbs the $C_{GH}$ factor into $K_{\mathcal{C}}$.

## 8.2 The refinement-distortion functional

The substrate-engineering content of R4′ is captured by the refinement-distortion functional introduced in §3.4:

$$\Xi_\ell(x, y) := d_H(\mathcal{R}_{\ell\to\ell+1}(\mathcal{C}_\ell(x)), \mathcal{R}_{\ell\to\ell+1}(\mathcal{C}_\ell(y))).$$

The functional measures how differently the substrate refines two nearby regions $x, y \in P_\ell$ — equivalently, how much the post-refinement cone-direction sets at level $\ell + 1$ (when viewed as images of the level-$\ell$ cones) differ from each other in the Hausdorff topology.

R4′ asserts that $\Xi_\ell$ is bounded by a Hölder modulus of continuity:

$$\Xi_\ell(x, y) \le K \cdot d_\ell(x, y)^\alpha + \varepsilon_\ell, \text{ with } \varepsilon_\ell \to 0 \text{ as } \ell \to \infty.$$

The condition says, equivalently:

Refinement cannot amplify local transport differences faster than a Hölder modulus.

This is the substrate-level mechanism generating H8. The companion paper's H5 (refinement compatibility, qualitative form) requires only that admissible directions persist under refinement: $T_\ell \subseteq \mathcal{R}^{-1}(T_{\ell+1})$. H5 by itself is sufficient for the *qualitative* convergence of cone-direction sets but not for the *quantitative* Hölder regularity required by the companion paper's Theorem 6. R4′ supplies the missing quantitative content as a constraint on the refinement-distortion observable $\Xi_\ell$, rather than as a postulate about the cones themselves.

## 8.3 Lemma 4.1 — Refinement-distortion propagation

**Lemma 4.1 (Refinement-distortion propagation).** *Under R4′, H5, R5 (bounded refinement overlap), and R6 (refinement bi-Lipschitz on cone-direction sets, with constants $c_{\mathcal{R}}$, $C_{\mathcal{R}}$), the* *discrete admissible-direction sets $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$ are uniformly Hölder-stable across refinement levels: there exists an effective constant $K_{eff} := c_{\mathcal{R}}^{-1} \cdot K < \infty$ such that for every sufficiently large $\ell \in \mathbb{N}$ and every pair $x, y \in P_\ell$,*

$$d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \leq K_{eff} \cdot d_\ell(x, y)^\alpha + c_{\mathcal{R}}^{-1} \cdot \varepsilon_\ell.$$

*Proof.* The proof factors into three substantive steps: (i) refinement preserves admissible-direction structure (H5); (ii) R4′ gives a Hölder bound on the refinement-distortion functional $\Xi_\ell$; (iii) the lower bi-Lipschitz constant of R6 translates the post-refinement bound back to a bound on the pre-refinement cones. R5 enters in justifying the use of R6 in the setting at hand: bounded refinement overlap guarantees that distinct children of a parent remain operationally distinguishable, which is what makes the Hausdorff metric on post-refinement cone-direction sets a well-defined object on which R6's bi-Lipschitz inequality can act.

*Step (i) — Refinement preserves admissible-direction structure.* By H5, admissible transport directions persist under refinement:

$$T_\ell \subseteq \mathcal{R}^{-1}(T_{\ell+1}).$$

Equivalently, the image $\mathcal{R}(\mathcal{C}_\ell(x))$ of the level-$\ell$ cone is contained in the level-$(\ell+1)$ cone at the corresponding refinement-image node: $\mathcal{R}(\mathcal{C}_\ell(x)) \subseteq \mathcal{C}_{\ell+1}(\mathcal{R}(x))$. The refinement functor therefore acts in a direction-preserving manner; in particular, the post-refinement cones are well-defined admissible cone-direction sets at level $\ell + 1$.

*Step (ii) — Post-refinement Hölder bound from R4′.* Applying R4′ directly:

$$d_H(\mathcal{R}(\mathcal{C}_\ell(x)), \mathcal{R}(\mathcal{C}_\ell(y))) = \Xi_\ell(x, y) \leq K \cdot d_\ell(x, y)^\alpha + \varepsilon_\ell.$$

This bounds the Hausdorff distance between the *refinement images* of the level-$\ell$ cones at x and y.

*Step (iii) — Inverse-refinement bound from R6.* By R6 (refinement bi-Lipschitz on cone-direction sets, with lower constant $c_{\mathcal{R}} > 0$):

$$c_{\mathcal{R}} \cdot d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \leq d_H(\mathcal{R}(\mathcal{C}_\ell(x)), \mathcal{R}(\mathcal{C}_\ell(y))).$$

Rearranging:

$$d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \leq c_{\mathcal{R}}^{-1} \cdot d_H(\mathcal{R}(\mathcal{C}_\ell(x)), \mathcal{R}(\mathcal{C}_\ell(y))) = c_{\mathcal{R}}^{-1} \cdot \Xi_\ell(x, y).$$

The inverse-refinement amplification constant $C_{\mathcal{R}} := c_{\mathcal{R}}^{-1} < \infty$ is supplied directly by R6 and is uniform in $\ell$. *Note on the role of R5.* R5 is what makes R6 substantively applicable to the cone-direction sets relevant here: without bounded refinement overlap, distinct level-$(\ell+1)$ directions could fail to be operationally distinguishable, and the $d_H$ on post-refinement cones would collapse to a degenerate pseudo-distance on which R6's lower bound $c_{\mathcal{R}} > 0$ would be vacuous. R5 secures the operational floor on which the $d_H$ on cone-direction sets is a genuine metric, and R6 then provides the bi-Lipschitz inversion.

*Combining steps (i)–(iii).* From Steps (ii) and (iii):

$$d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \le c_{\mathcal{R}}^{-1} \cdot \Xi_\ell(x, y) \le c_{\mathcal{R}}^{-1} \cdot K \cdot d_\ell(x, y)^\alpha + c_{\mathcal{R}}^{-1} \cdot \varepsilon_\ell.$$

Defining $K_{eff} := c_{\mathcal{R}}^{-1} \cdot K$ gives the claimed bound. The constant $K_{eff}$ is finite, uniform in $\ell$, and depends only on the engineering constants of R4′ and R6 (the role of R5 is to validate the metric setting, not to enter the constants).

This completes the proof of Lemma 4.1.

## 8.4 Proof of Theorem 4

Given the discrete Hölder stability established in Lemma 4.1, the lift to the continuum H8 is mechanical via the companion paper's Theorem 2 and the Arzelà–Ascoli-type argument used in its Theorem 6 proof.

The Hausdorff-Hölder-continuous family $\{\mathcal{C}_\ell\}_{\ell\in\mathbb{N}}$ *delivered by Lemma 4.1, with uniform modulus* $K_{eff} \cdot r^\alpha$ *and additive error* $c_{\mathcal{R}}^{-1} \cdot \varepsilon_\ell \to 0$*, converges along Lorentzian-GH-convergent subsequences to a Hausdorff-Hölder-continuous limit* $\mathcal{C}: \mathcal{M}_\infty \to$ (non-degenerate convex double cones), with the same exponent $\alpha$ and continuum Hölder constant $K_{\mathcal{C}} = K_{eff} \cdot C_{GH}$, where $C_{GH} < \infty$ is the Lorentzian-GH-embedding constant inherited from the companion paper's Theorem 6 (under the bounded Lorentzian-GH distortion guaranteed by the H1 + H5 + H6 + H6′ + H7 hypothesis cluster). The discrete additive error $c_{\mathcal{R}}^{-1} \cdot \varepsilon_\ell$ contributes the H8 additive error term $\varepsilon_\ell \to 0$ (absorbing the $c_{\mathcal{R}}^{-1}$ factor into the redefinition $\varepsilon_\ell \mapsto c_{\mathcal{R}}^{-1} \cdot \varepsilon_\ell$, which preserves $\varepsilon_\ell \to 0$) and vanishes in the limit; the continuum Hölder constant $K_{\mathcal{C}}$ captures only the leading-order modulus and does not absorb the additive error.

For any $x_\infty, y_\infty \in \mathcal{M}_\infty$,

$$d_H(\mathcal{C}(x_\infty), \mathcal{C}(y_\infty)) \le K_{\mathcal{C}} \cdot d_\infty(x_\infty, y_\infty)^\alpha, \quad K_{\mathcal{C}} = K_{eff} \cdot C_{GH}.$$

This is precisely H8.

This completes the proof of Theorem 4.

## 8.5 Interpretation

H8 is not fundamentally a continuum regularity assumption.

It is the continuum manifestation of *bounded refinement distortion*.

The substrate-level quantity controlling regularity is not the cone field itself but the rate at which refinement changes admissible transport structure between neighbouring regions. The refinement-distortion functional $\Xi\_\ell(x, y)$ therefore becomes the central geometric-engineering observable of the substrate programme — a quantity that can in principle be measured directly on concrete TPB substrate models, without first having to compute the continuum cone field.

A substrate satisfying bounded refinement distortion cannot refine chaotically: nearby regions remain transport-compatible under refinement, and the continuum cone field inherits Hölder regularity automatically via Lemma 4.1 and Theorem 4. Without bounded refinement distortion, arbitrarily small substrate perturbations could generate order-one changes in the local cone structure, destroying smooth continuum emergence — the Chruściel–Grant (2012) $C^0$ pathologies of the companion paper's §10 would then appear in the smoothing limit.

The substrate-engineering content of H8 lives in R4′ as a bound on $\Xi\_\ell$, not in a postulated Hölder-cone-stability constraint on $\mathcal{C}\_\ell$: the discrete Hölder stability of the cone-direction sets used downstream is *derived* from R4′ + R5 + R6 + H5 via Lemma 4.1, with no Hölder postulate on the cones themselves.

## 8.6 Threshold-vs-quantitative character of R4′

R4′ is a *threshold hypothesis* in the sense that for any $\alpha \in (0, 1]$ — however small — the companion paper's §10 chain delivers a C^k Lorentzian metric for any prescribed $k \geq 2$. The exponent α controls the *rate* of regularity convergence but not the *existence* of the smooth structure.

However, *better* α values have downstream consequences. Specifically:

- **α = 1 (Lipschitz refinement distortion).** The induced Hölder cone stability becomes Lipschitz. The companion paper's wave-equation correction in §10.7 Step 3c becomes pointwise small in $C^0$, eliminating the weak-norm-robustness gap of §10.10 entirely. The bridge to the companion-companion (Lorentzian-emergence) paper's §6.1 conformal-factor argument operates with pointwise smallness rather than only $L^2$/Sobolev control.
- **α = 1 with Lipschitz cone derivatives, i.e., C^{1,1}-class refinement distortion.** This would give a C^{1,1} continuum cone field and shorten the §10 chain further. The wave-equation correction is no longer needed because the naive mollification already produces a pointwise-divergence-free current. (We caveat that α > 1 is not the right way to phrase this: Hausdorff-distance Hölder moduli with α > 1 are formally meaningful on cone bundles but on connected Euclidean domains they collapse real-valued Hölder maps to constants, so the correct continuum strengthening is C^{1,1} via Lipschitz cone derivatives, not α > 1 in the Hölder modulus itself.)

Identifying substrate-engineering hypotheses that yield α = 1 in R4′ (or C^{1,1}-class refinement distortion as the stronger version) — equivalently, substrate architectures whose refinement-distortion functional $\Xi\_\ell$ has Lipschitz or C^{1,1} stability — is the natural sharpening question for the next paper, with substantial downstream payoff.

## 8.7 Why finite closure architecture naturally supports R4′

The link between the *sequential-interface-transport* paper's K = 7 closure architecture and the present R4′ is direct, and is the key reason for expecting concrete TPB substrate models to satisfy R4′ rather than violate it.

Because admissible local transport is generated by a finite closure catalogue (cardinality $K_{max}$ = 7 for VERSF), neighbouring substrate regions cannot refine through arbitrarily different local closure classes. Finite closure structure therefore *discretizes refinement response*: at each refinement level there are only finitely many distinguishable closure configurations through which a node can be refined, so the refinement-distortion functional $\Xi_\ell$ takes values in a structured set rather than varying continuously over an unbounded range.

More precisely:

- The number of primitive admissible transport generators is finite ($\leq K_{max} = 7$).
- Refinement acts through a finite transition algebra on closure states — i.e., the action of the refinement functor $\mathcal{R}$ on the closure catalogue $\mathcal{K}_\ell$ is a discrete operation between finite sets.
- Local transport sectors remain closure-equivalent under bounded refinement: a node refining through closure configuration $\kappa \in \mathcal{K}_\ell$ produces children whose closure configurations lie in a finite, structurally-determined subset of $\mathcal{K}_{\ell+1}$.

This implies that refinement distortion $\Xi_\ell(x, y)$ cannot grow arbitrarily fast with decreasing scale: nearby regions either share a common closure configuration (in which case $\Xi_\ell \approx 0$ up to refinement noise) or pass through closely related configurations (in which case $\Xi_\ell$ is bounded by the finite-algebra distance between configurations, which is uniformly bounded above).

The natural conjecture is therefore:

K = 7 closure architecture + R5 + closure stability ⟹ R4′ (and conjecturally also R6).

In this interpretation, *Hölder continuum regularity (H8) is the macroscopic shadow of finite closure-algebra stability at the substrate level.* The geometry of the continuum inherits its quantitative regularity from the discreteness and finiteness of the underlying closure algebra. This is a much stronger conceptual claim than "Hölder stability is assumed"; it is "Hölder stability is generated by the K = 7 closure structure," and it is the form of the §10 chain we expect to verify in concrete substrate models.

The derivation of R4′ from the K = 7 closure architecture, with explicit values of $\alpha$ and K, is open and is the principal next-paper task for the substrate-engineering layer. The conjecture $\alpha$ = 1 (Lipschitz refinement distortion, with the §10.7 wave-equation correction collapsing as a consequence) is the highest-value version of this open problem.

## 8.8 Where the proof concentrates its risk

The risk is concentrated entirely in R4′ — equivalently, in the substrate-level boundedness of the refinement-distortion functional $\Xi_\ell$. Three risk-loci are worth distinguishing:

1. *Existence of a Hölder modulus on $\Xi\_\ell$ (any $\alpha > 0$).* This is the *threshold* version of the risk. If no positive exponent works — if $\Xi\_\ell$ exhibits genuinely chaotic behaviour under refinement — R4′ fails and the entire §10 regularity-upgrade chain of the companion paper collapses. The natural conjecture (§8.7) is that the K = 7 closure architecture rules this out.
2. *Value of the exponent $\alpha$.* For $\alpha < 1$, R4′ holds but with the §10.7 wave-equation correction's weak-norm-robustness gap inherited. For $\alpha \geq 1$, the weak-norm-robustness gap disappears entirely. The substrate-engineering question is which $\alpha$ the K = 7 architecture actually delivers.
3. *Value of the lower bi-Lipschitz constant $c\_\mathcal{R}$ in R6.* If $c\_\mathcal{R} \to 0$ under refinement — i.e., the refinement functor's action on cone-direction sets degenerates toward a collapsing rather than a bi-Lipschitz embedding — then $K\_eff = c\_\mathcal{R}^{-1} \cdot K$ diverges and Lemma 4.1 fails to deliver discrete Hölder cone stability even when R4′ itself holds at the post-refinement level. The substrate-engineering question is whether the σ-duality refinement functor of the *σ-duality* paper has a uniformly positive lower bi-Lipschitz constant on cone-direction sets in Hausdorff distance.

All three loci are concrete and targetable in concrete substrate models. The risk-relocation effected by this paper is that the previously-postulated H8 has been replaced by a substrate-engineering observable $\Xi\_\ell$ whose properties can in principle be computed directly from any concrete TPB substrate model.

---

# 9. Why H9 is structurally different

H9 (smooth-structure existence on the limit) is qualitatively different from H6, H6′, H7, H8. The latter four are all *dynamical* conditions — they constrain how the substrate evolves under refinement. H9 is a *topological* condition — it constrains the smooth-manifold structure of the limit ($\mathcal{M}\_\infty$, $d\_\infty$).

### 9.1 The Donaldson–Freedman obstruction

By results of Donaldson (1983) and Freedman (1982), in dimension 4:

- Some topological 4-manifolds admit *no* compatible smooth structure (e.g., the E8-manifold).
- Some topological 4-manifolds admit *infinitely many* inequivalent smooth structures (the exotic $\mathbb{R}^4$ phenomena and exotic structures on closed 4-manifolds).
- Smooth-structure existence is *not* a local condition — it depends on global topological invariants (Donaldson polynomials, Seiberg–Witten invariants) that cannot be computed from local transport dynamics alone.

This means that no local substrate-engineering constraint — no matter how stringent — can entail H9. Smooth-structure existence is a *global topological* property of the limiting 4-manifold and is not derivable from substrate dynamics that operate locally.

### 9.2 What we can and cannot say

What this paper does *not* do is derive H9 from substrate-level constraints. What it *can* say is:

- For continuum limits that are compact perturbations of standard cosmological 4-manifolds (Minkowski, FLRW, asymptotically flat black-hole geometries), smooth structures are unique up to diffeomorphism (cf. Freedman–Quinn 1990 for the precise classification), and H9 is essentially automatic.
- For more exotic continuum topologies — should they arise from concrete substrate models — H9 is a genuine restriction whose status must be analysed case-by-case using 4-manifold topology methods.
- The substrate-engineering question of which TPB substrate models produce continuum limits with the cosmologically natural topology (and hence automatic H9) is open.

### 9.3 Forward pointer

The substrate-engineering paper that exhibits concrete TPB substrate models should, in addition to verifying R1–R7 with explicit constants, characterise the topology of the continuum limit for each model. For models producing cosmologically natural topologies, H9 follows from Freedman–Quinn classification. For models producing exotic topologies, H9 must be analysed using 4-manifold-topological invariants and may require either substrate-architecture restrictions or topological hypotheses inherited from the $K = 7$ closure structure.

The interaction between substrate-dynamical conditions and the 4-manifold topology of the limit is, at present, a programmatic open problem. The present paper does not attempt it.

---

# 10. Theorem 5 — Main synthesis

### 10.1 Statement

**Theorem 5 (Main synthesis).** *Let* $\{P_\ell, T_\ell, B_\ell, \mathcal{R}_{\ell\to\ell+1}\}_{\ell\in\mathbb{N}}$ *be a regular TPB substrate satisfying:*

- *baseline VERSF substrate axioms A0, A1*, A2*, A3, A4 (with A1* and A2* the strengthened forms of A1 and A2 introduced in §4) and BCB;*
- *the companion paper's H1 (finite propagation) and H5 (refinement compatibility);*
- *the seven engineering constraints R1 (uniform bounded local degree, §3.1), R2 (controlled refinement multiplicity, §3.2), R3 (uniform distinguishability separation, §3.3), R4' (bounded refinement distortion, §3.4), R5 (bounded refinement overlap, §3.5), R6 (refinement bi-Lipschitz on cone-direction sets, §3.6), R7 (polynomial volume growth in pseudometric balls, §3.7);*
- *the K = 7 closure architecture of the sequential-interface-transport paper (§6.2).*

*Then the refinement sequence satisfies H6 (Theorem 1, from R7), H7 (Theorem 2, from R1 + K = 7), H6′ (Theorem 3, from A1* + A2* + R3 + R5), and H8 (Theorem 4, with Hölder cone stability of $\mathcal{C}\_\ell$ derived from R4′ + R5 + R6 + H5 via the refinement-distortion-propagation Lemma 4.1).*

*If additionally the continuum limit ($\mathcal{M}\_\infty$, $d\_\infty$) satisfies H9 (smooth-structure existence on the limiting 4-manifold), then the refinement sequence satisfies the full hypothesis set H1–H9 of the companion paper Continuum-Limit Regularity and Cone Convergence in VERSF. Consequently, by Theorems 5 and 8 of that paper, the refinement sequence converges to a strongly causal Lorentzian length space admitting a C^k Lorentzian metric for every prescribed k ≥ 2.*

## 10.2 Proof

The proof is the conjunction:

- Theorem 1 (§5) gives H6 from R7 (polynomial volume growth in pseudometric balls).
- Theorem 2 (§6) gives H7 from R1 and the K = 7 closure architecture.
- Theorem 3 (§7) gives H6′ from A1*, A2*, R3, R5 (via the within-level A1*+A2* contradiction in §7.3 and across-level R5 propagation in §7.4).
- Theorem 4 (§8) gives H8 from R4′, R5, R6, H5 (via the refinement-distortion-propagation Lemma 4.1 in §8.3).

The conjunction H6 + H6′ + H7 + H8 plus the assumed H9 yields the full hypothesis set of the companion paper. Applying the companion paper's Theorem 5 (continuous Lorentzian length space) and Theorem 8 (C^k Lorentzian metric for any k ≥ 2) completes the proof.

This completes the proof of Theorem 5.

## 10.3 Significance

Theorem 5 converts the companion paper's *regularity-hypothesis-conditional* result:

A0–A4 + BCB + H1–H9 ⟹ smooth Lorentzian continuum

into a *substrate-engineering-conditional* result:

regular TPB substrate + H9 ⟹ smooth Lorentzian continuum,

where "regular TPB substrate" is a precisely specified class of substrate models satisfying explicit engineering constraints R1–R7, and H9 is the residual topological condition that the present paper does not address.

The mathematical-risk concentration has moved one further layer down: from *external regularity hypotheses* (companion paper) to *substrate engineering constraints* (this paper) to *concrete substrate models satisfying R1–R7 with quantitative constants* (next paper). Within R1–R7, the highest-value sharpenings are: (i) achieving α = 1 in R4′ (Lipschitz refinement distortion) —

would close the §10.10 weak-norm-robustness gap of the companion paper entirely; (ii) deriving R7 (Ahlfors d-regularity, in both directions) from K = 7 + R4′ — would consolidate H6's foundation and replace an independent assumption with a substrate-architectural consequence; (iii) deriving R6 (refinement bi-Lipschitz on cones, particularly the lower constant $c_{\mathcal{R}} > 0$) from the σ-duality refinement-functor properties — would consolidate Lemma 4.1's inverse-refinement step on the σ-duality side rather than as a parallel engineering hypothesis.

---

# 11. Architectural consequences

## 11.1 Risk relocation

The companion paper's *Continuum-Limit Regularity* §15 conclusion identified the substrate-engineering question as the principal remaining mathematical risk. The present paper sharpens that risk:

| Layer | Status before this paper | Status after this paper |
|---|---|---|
| Continuum-convergence (H6, H6′, H7, H8) | External hypotheses requiring substrate derivation | Derived from engineering constraints R1–R7 |
| Substrate engineering (R1–R7) | Not isolated as a named layer | Explicitly named; R4′ stated as a Hölder bound on the substrate-engineering observable $\Xi_\ell$ (rather than on the cones themselves); R5 (bounded refinement overlap), R6 (refinement bi-Lipschitz on cones), R7 (polynomial volume growth) promoted to named engineering constraints; derivation from underlying VERSF closure-architecture papers sketched but not fully discharged |
| Topological (H9) | External hypothesis | Confirmed as genuinely external (Donaldson–Freedman obstruction); cosmologically natural topologies give automatic H9 |
| Wave-equation bridge (§10.10 of companion paper) | External robustness question | Inherited unchanged, but contingent: α = 1 in R4′ would eliminate this gap entirely |

The residual mathematical risk now lives at three named loci:

1. The substrate-level grounding of R1–R7 in the underlying VERSF closure-architecture papers (*sequential-interface-transport*, *σ-duality*, *admissible-coarse-graining*), with explicit constants — particularly the conjectures that the K = 7 closure architecture entails R4′ (§8.7), R6 (the σ-duality refinement functor is bi-Lipschitz on cones), and R7 (polynomial volume growth follows from K = 7 + R4′).

2. The topological question of H9 for non-cosmologically-natural continuum topologies.
3. The weak-norm-robustness gap in §10.10 of the companion paper, properly a revision to the Lorentzian-emergence paper.

## 11.2 Constant tracking is a substantive sharpening question

The seven engineering constraints R1–R7 are each parametric in explicit constants: Δ (R1), M (R2), $\sigma_*$ (R3), (α, K) (R4′), $\sigma_{op}$ (R5), ($c_{\mathcal{R}}$, $C_{\mathcal{R}}$) (R6), (d, $c_{vol}$, $C_{vol}$) (R7) — with the doubling constant in H6 depending on R7 only through the dimension d and the Ahlfors-regularity ratio $C_{vol} / c_{vol}$. The continuum-convergence chain inherits these constants but does not require them to take specific values — for any positive choice of all of them, the chain delivers a smooth Lorentzian limit (modulo H9 and the §10.10 gap).

However, *better* constants substantially shorten the downstream chain:

- Better Δ, M (smaller values) ⟹ tighter local-structure control on the substrate.
- Better $\sigma_*$, $\sigma_{op}$ (larger values) ⟹ tighter quantitative-acyclicity bound in Lemma 7.2 of the companion paper ⟹ stronger exclusion of almost-closed causal curves.
- Better α in R4′ (closer to 1, or $C^{1,1}$) ⟹ closer-to-Lipschitz refinement distortion, hence closer-to-Lipschitz cone stability of $\mathcal{C}$ ⟹ potential elimination of the wave-equation correction step in §10.7.
- Better $c_{\mathcal{R}}$ in R6 (larger value, i.e., refinement closer to a Lipschitz-isometry on cone-direction sets) ⟹ smaller $K_{eff} = c_{\mathcal{R}}^{-1} \cdot K$ in Lemma 4.1 ⟹ tighter constants throughout the §10 chain.
- Better (d, $C_{vol}$, $c_{vol}$) in R7 (smaller polynomial-growth dimension d, smaller upper constant $C_{vol}$, larger lower constant $c_{vol}$) ⟹ tighter doubling constant $D \le (C_{vol} / c_{vol}) \cdot 12^d$ in H6 ⟹ tighter dimensional control on the limit.

The constants are coupled and the geometry-emergence chain's final constants are products of these. A larger $c_{\mathcal{R}}$ (R6) directly reduces $K_{eff}$ and hence improves the H8 constant $K_{\mathcal{C}}$; a smaller d (R7) directly reduces the H6 doubling constant; an α = 1 in R4′ collapses the §10.7 wave-equation correction step entirely. The substrate-engineering paper should therefore exhibit concrete TPB substrate models satisfying R1–R7 and track the constants (α, $c_{\mathcal{R}}$, d, $C_{vol} / c_{vol}$, $\sigma_{op}$) delivered by each model — jointly rather than in isolation — since the downstream geometry simplifies considerably for models with strong constants across all engineering constants simultaneously.

## 11.3 The architectural picture

Combining the present paper with the companion paper and the Lorentzian-emergence paper, the full geometry-emergence chain in VERSF is:

```
Layer 0 — Substrate engineering (this paper):
   R1  (bounded local degree)             ← strengthens A4
   R2  (refinement multiplicity)          ← strengthens §3.2 local-degree
control
   R3  (distinguishability separation)    ← promotes A1* to metric statement
```

```
   R4′ (bounded refinement distortion)     ← substrate-engineering observable
                                              Ξ_ℓ(x,y); Hölder bound on Ξ_ℓ
   R5  (bounded refinement overlap)        ← refinement-step analogue of A1*;
                                              propagates σ_op across levels
   R6  (refinement bi-Lipschitz on cones)  ← lower constant c_𝓡 > 0 supplies
                                              inverse-refinement bound in
Lemma 4.1
   R7  (Ahlfors d-regularity, two-sided)   ← upper bound rules out tree-like
Δ^r
                                              growth, lower bound prevents
small-ball
                                              packing degeneration; required
for doubling

Layer 0.5 — Substrate-axiom strengthening:
   A1* (uniform operational floor)        ← strengthens A1
   A2* (operational distinctness)         ← strengthens A2

Layer 1 — Continuum-regularity hypotheses (this paper derives from R1–R7):
   H6  (uniform doubling)                  ← R7 (Theorem 1)
   H7  (bounded combinatorial dimension)   ← R1 + K = 7 architecture (Theorem
2)
   H6′ (transport sparsity)                ← A1* + A2* + R3 + R5 (Theorem 3)
   H8  (refinement Hölder compatibility)   ← R4′ + R5 + R6 + H5, via
                                              Lemma 4.1 (refinement-
distortion
                                              propagation) and Theorem 4

Layer 1.5 — Topological hypothesis (this paper does not address):
   H9 (smooth-structure existence)        ← Topological; assumed separately

Layers 2–4 — Continuum convergence (companion paper):
   Theorem 5 (continuous Lorentzian       ← H1, H5, H6, H6′, H7 + BCB
              length space)
   Theorems 6, 7, 8 (regularity upgrade)  ← Theorem 5 + H8, H9 + Sämann
(2016)

Layers 5–7 — Lorentzian emergence (companion-companion paper):
   Lemma 5.2, §6.1 chain                  ← Theorem 8 + smooth structure
                                            (modulo §10.10 weak-norm gap)
   Theorems 1L, 1G, 2, 3                  ← downstream Lorentzian results
```

The architecture is now layered cleanly: *substrate engineering* (this paper) → *continuum regularity* (companion paper) → *Lorentzian emergence* (companion-companion paper). Each layer takes the output of the previous as input and concentrates a different category of residual risk.

---

# 12. Falsification paths and quantitative tests

The framework is falsified if any of the following are observed. Paths E1–E7 concern the substrate-engineering layer (this paper), with E4 falsifying R4′ and E4′ falsifying the lower bi-

Lipschitz constant of R6 (so E4 and E4′ are structurally distinct rather than alternative names for the same failure); E6 falsifies A1*, E6′ falsifies A2*; E7 covers failure of R7 in either the upper-bound or lower-bound direction. Paths G1–G9 of the companion paper and F1–F7 of the Lorentzian-emergence paper are inherited.

**(E1) Failure of bounded local degree (R1).** If a physically motivated TPB substrate model exhibits A4-local-coupling degree that grows unboundedly under refinement, R1 fails. The substrate would produce a continuum limit with unbounded local node density, which together with E7 (failure of polynomial volume growth) would violate H6 and break the precompactness conclusion of the companion paper's Theorem 1.

**(E2) Failure of controlled refinement multiplicity (R2).** If a physically motivated TPB substrate model exhibits refinement multiplicity that grows unboundedly per refinement step (operationally distinguishable child nodes per parent), R2 fails. Combined with E1 + E7, this would produce uncontrolled node-density growth in pseudometric balls and break H6.

**(E3) Failure of uniform distinguishability separation (R3).** If a physically motivated TPB substrate model produces refinement families with $\tilde{\sigma}_\ell \to 0$ — i.e., distinct committed states along admissible transport sequences become arbitrarily close in rescaled metric — R3 fails. By Theorem 3, A1* or R5 must also be violated. The substrate would produce a continuum limit with almost-closed causal curves, breaking H6′ and hence the strong-causality conclusion of the companion paper's Theorem 3.

**(E4) Failure of bounded refinement distortion (R4′).** If a physically motivated TPB substrate model exhibits a refinement-distortion functional $\Xi_\ell(x, y)$ that fails to be bounded by any Hölder modulus — i.e., refinement amplifies local transport differences faster than any positive-$\alpha$ power of $d_\ell(x, y)$ — R4′ fails. By Lemma 4.1, Hölder cone stability of $\mathcal{C}_\ell$ then fails, and the substrate produces a continuum cone field that is continuous (companion paper Theorem 2 still works) but not Hölder, breaking the H8-dependent regularity upgrade of the companion paper's §10. The Chruściel–Grant (2012) $C^0$ pathologies could then occur in the smoothing limit. This is the qualitative *threshold* form of the falsifier; even with R4′ holding at small $\alpha$, the §10.7 weak-norm gap of the companion paper persists, and only $\alpha \geq 1$ fully eliminates it.

**(E4′) Failure of the lower bi-Lipschitz constant in R6.** Even if R4′ holds at the post-refinement level, the lower bi-Lipschitz constant $c_{\mathcal{R}}$ of R6 may fail to be uniform in $\ell$ — for example, if the σ-duality refinement functor's action on cone-direction sets degenerates toward a collapsing rather than a bi-Lipschitz embedding. In that case, $K_{eff} = c_{\mathcal{R}}^{-1} \cdot K$ in Lemma 4.1 diverges, the lemma fails to translate the post-refinement R4′ bound back to a bound on the pre-refinement cones, and H8 is not delivered even though R4′ holds. This is structurally distinct from E4: E4 falsifies the post-refinement bound on $\Xi_\ell$; E4′ falsifies the inverse-refinement bi-Lipschitz constant required by R6. R5 may simultaneously fail in such substrates (since the operational floor required for the $d_H$ on post-refinement cone-direction sets to be a genuine metric also degenerates), in which case E4′ and a failure of R5 are coupled.

**(E5) Failure of K = 7 closure architecture.** If the *sequential-interface-transport* paper's K = 7 closure architecture turns out to be incorrect — i.e., if admissible local closures form an infinite or pathological catalogue rather than a finite seven-element one — then Theorem 2 fails to deliver an explicit bound on H7. The qualitative conclusion (some finite N for H7) may still hold if the closure catalogue is finite but with a different cardinality; the quantitative $N = 7$ claim fails.

*(E6) Failure of A1 (strengthened distinguishability).** If a physically motivated TPB substrate model fails to admit a uniform operational floor $\sigma_{op} > 0$ — i.e., the minimum distinguishability scale collapses under refinement — A1* fails, the $\sigma_* \geq \sigma_{op}$ chain in Theorem 3 collapses, and H6′ fails. R5 fails simultaneously (its content is the refinement-step version of A1*). This would mean that "irreversible commitment" in the substrate becomes operationally meaningless in the continuum limit, which would be a deep structural failure of the VERSF programme.

*(E6') Failure of A2 (operational distinctness of committed states).** If a physically motivated TPB substrate model admits commitment events whose outcomes are operationally indistinguishable from the prior state — i.e., commitment events that produce nothing operationally new — A2* fails, the within-level $\sigma_{op}$ contradiction step in §7.3 collapses, and H6′ fails. The §4.2 structural argument suggests this should not happen — A2's own "new fact production" content would be violated — but the substrate-engineering question is whether all physically motivated closure architectures intrinsically enforce A2* or whether some admit operationally vacuous commitment events. This is structurally distinct from E6: E6 falsifies the operational-floor scale, E6′ falsifies the alignment between committed-state distinctness and operational distinctness.

**(E7) Failure of Ahlfors d-regularity (R7).** R7 fails in two distinct directions. (a) *Failure of the upper bound* — the substrate exhibits super-polynomial (in particular exponential) volume growth in pseudometric balls (the Δ-regular-tree regime, or any tree-like architecture). By Theorem 1, doubling H6 then fails, and the precompactness conclusion of the companion paper's Theorem 1 fails too. (b) *Failure of the lower bound* — the substrate becomes arbitrarily sparse at small scales, so that $B_\ell(y, r)$ contains $o(r^d)$ operationally distinguishable nodes as $r \to 0$. In that case the packing argument in Theorem 1 Step 2 collapses (the lower-bound side $c_{vol} \cdot (r/4)^d \to 0$ invalidates the cancellation), and even with R7's upper bound intact, the doubling constant cannot be made r-uniform. Note that E7 in either direction is *consistent with* R1 + R2 holding: bounded local degree and bounded refinement multiplicity are not enough to rule out tree-like global volume growth (failure mode (a)) or arbitrarily sparse refinement at small scales (failure mode (b)). E7 is therefore a genuinely distinct failure mode that must be ruled out by additional substrate-architectural conditions (the conjecture being that K = 7 + R4′ does so in both directions, but this is open — see §13).

---

# 13. Open problems and dependencies

**Open problems flagged in this paper.**

*Substrate-engineering origin of R1–R7.*

- **Substrate-level derivation of R1 from the K = 7 closure architecture.** R1 (bounded local degree) should follow from the finite-closure-catalogue structure of the *sequential-interface-transport* paper, with Δ determined by the maximum coupling degree compatible with K = 7 closure configurations. Making this implication explicit, with a quantitative value of Δ, is open.
- **Substrate-level derivation of R2 from refinement-functor properties.** R2 (controlled refinement multiplicity) should follow from properties of the refinement functor $\boldsymbol{\mathcal{R}}$ established in the *σ-duality* paper. The σ-duality structure ought to bound the multiplicity of operationally distinguishable children per parent; identifying the explicit constant M from σ-duality is open.
- *Substrate-level derivation of A1 (and hence R3) from finite distinguishability and refinement compatibility.** A1* strengthens A1 to a uniform operational floor. Whether A1 + refinement compatibility implies A1* automatically — perhaps with σ_op determined by the minimum distinguishability scale at the coarsest refinement level and propagated forward — is open.
- *Substrate-level derivation of A2 from A2 and admissible-commitment dynamics.** A2* strengthens A2 to include the operational-distinctness clause: distinct committed states correspond to operationally distinguishable facts. The structural argument in §4.2 shows that A2's "new fact production" content forces this clause logically — a commitment event that produced an operationally indistinguishable outcome would not, by A2's own criterion, be a commitment event at all. Whether this structural derivation can be tightened into a formal one within the BCB/closure axiomatic setting of the *sequential-interface-transport* paper is open.
- **Explicit derivation of the refinement-distortion modulus (the principal next-paper task for the §10 chain).** The present paper introduces the refinement-distortion functional Ξ_ℓ(x, y) := d_H($\boldsymbol{\mathcal{R}}$(𝓒_ℓ(x)), $\boldsymbol{\mathcal{R}}$(𝓒_ℓ(y))) and the bounded-distortion hypothesis R4′ as a substrate-engineering observable. The next substrate-engineering task is to derive explicit bounds on Ξ_ℓ from the K = 7 closure algebra and the σ-duality refinement functor. The key question is whether the closure architecture yields:
    - *Hölder stability with α ∈ (0, 1)* — sufficient for the §10 chain of the companion paper to deliver a C^k Lorentzian metric for any k ≥ 2;
    - *Lipschitz stability with α = 1* — particularly important because it would eliminate the weak-norm-robustness gap of the companion paper's §10.10 by making the naive mollification pointwise divergence-stable, thereby removing the wave-equation correction of §10.7 Step 3c as a load-bearing step;
    - *C^{1,1}-class refinement distortion* — Lipschitz cone derivatives at α = 1; would deliver a C^{1,1} continuum cone field directly, shortening the §10 chain further. The conjecture (§8.7) is that K = 7 closure architecture + R5 + closure stability ⟹ R4′ with α = 1, but this is not established here.
- **Substrate-level derivation of R5 from σ-duality refinement-functor structure.** R5 (bounded refinement overlap, §3.5) — distinct children of a parent stay separated above σ_op under refinement — should follow from the categorical lift provided by the σ-duality refinement functor. Establishing this is open and is required for the §7.4 across-level propagation argument in Theorem 3.
- **Substrate-level derivation of R6 from σ-duality refinement-functor structure.** R6 (refinement bi-Lipschitz on cone-direction sets, §3.6), particularly the lower bi-Lipschitz

constant $c\_\mathcal{R} > 0$, should follow from the σ-duality refinement functor's action on cone-direction sets. R2 + R5 together give injectivity but not bi-Lipschitz, so R6 is a genuine additional substrate-engineering constraint and its derivation is open. The conjecture is that the σ-duality functor's structural properties (categorical adjointness, finite intertwiner cardinality) deliver $c\_\mathcal{R} > 0$ uniformly in ℓ.

- **Substrate-level derivation of R7 (Ahlfors d-regularity in both directions) from K = 7 + R4′.** R7 — uniform Ahlfors d-regularity in pseudometric balls, with both an upper bound C_vol and a lower bound c_vol > 0 — is required for Theorem 1's doubling conclusion. R1 + R2 alone give at most exponential growth from above (the Δ-regular-tree regime, which fails doubling) and admit arbitrarily sparse refinement from below, so R7 in both directions is genuinely a separate engineering constraint. The conjecture is that the K = 7 closure architecture combined with bounded refinement distortion (R4′) rules out tree-like volume growth at coarse scales (upper bound, since finite-closure refinement with bounded distortion cannot sustain the unbounded branching required for Δ^r ball volumes) and arbitrarily sparse refinement at small scales (lower bound, since each closure configuration generates at least one operationally distinguishable child, so closure-stable refinement maintains a uniform packing density). Establishing this in both directions is open and is the principal substrate-engineering route to discharging R7.

*Topological gap.*

- **H9 for cosmologically natural topologies.** For continuum limits that are compact perturbations of Minkowski or FLRW, H9 is essentially automatic by Freedman–Quinn classification. Confirming that physically motivated TPB substrate models actually produce continuum limits in this topological class — rather than exotic 4-manifolds — is open. This is a substrate-engineering question that intersects 4-manifold topology and is qualitatively different from R1–R7.
- **Substrate-level constraints implying cosmologically natural topology.** Whether the K = 7 closure architecture combined with R1–R7 automatically constrains the limit topology to the cosmologically natural class is open. If so, H9 would be a consequence rather than a separate assumption. If not, the substrate-engineering paper must analyse each concrete substrate's limit topology individually.

*Constant tracking.*

- **Quantitative bounds Δ, M, σ_op, σ_*, α, K, $c\_\mathcal{R}$, $C\_\mathcal{R}$, d, c_vol, C_vol for concrete TPB substrate models** — with the downstream doubling constant in H6 depending on R7 only through the dimension d and the Ahlfors-regularity ratio C_vol / c_vol. The downstream geometry simplifies considerably for substrates with strong constants — particularly α = 1 in R4′ (Lipschitz refinement distortion), which would eliminate the wave-equation correction step in the companion paper's §10.7. Identifying substrate models with strong constants is open.

*Inherited gaps.*

- **Weak-norm-robustness gap in companion paper's §10.10.** Inherited unchanged. Best addressed by a revision to the Lorentzian-emergence paper's §6.1 along the lines of the three candidate fixes in §10.10 of the companion paper. *Contingent gap:* if the $K = 7$ closure architecture turns out to entail R4′ with $\alpha \geq 1$, this gap closes automatically.
- **C^∞ desideratum in companion paper's §10.8.** Inherited unchanged. $C^2$ suffices for the geometry-emergence chain.

**Principal remaining structural challenge.** The single most important next-paper task is the explicit substrate-engineering paper exhibiting concrete TPB substrate models satisfying R1–R7 with quantitative constants. This paper should: (i) verify R1–R7 with explicit values of Δ, M, σ_op, σ_*, α, K, c_𝓡, C_𝓡, d, c_vol, and C_vol — particularly by computing the refinement-distortion functional Ξ_ℓ directly on concrete substrate dynamics; (ii) characterise the topology of the continuum limit and verify H9 (or note where it must be assumed); (iii) ideally aim for $\alpha = 1$ in R4′ to eliminate the wave-equation correction in the companion paper's §10.7; (iv) attempt to discharge R7 in both directions (upper bound C_vol from finite-closure global structure, lower bound c_vol from closure-stability of refinement at small scales) from the K = 7 + R4′ conjecture rather than assuming it. With such a paper in place, the geometry-emergence chain in VERSF becomes *substrate-model-specific* rather than substrate-engineering-conditional, completing the chain from concrete substrate dynamics to smooth Lorentzian geometry.

**Status summary.**

| Result | Status |
|---|---|
| Theorem 1 (R7 ⟹ H6) | Proven; explicit doubling constant D ≤ (C_vol / c_vol) · 12^d, depending only on the Ahlfors constants of R7. R7 is genuinely independent of R1 + R2 — see §5.3 — and its derivation in both directions from K = 7 + R4′ is open |
| Theorem 2 (R1 + K = 7 ⟹ H7) | Proven, conditional on K = 7 closure architecture from sequential-interface-transport paper |
| Theorem 3 (A1*, A2*, R3, R5 ⟹ H6′) | Proven via the direct A1*+A2* contradiction for the within-level σ_op bound and R5 for across-level propagation |
| Lemma 4.1 (R4′, H5, R5, R6 ⟹ Hölder cone stability of 𝒞_ℓ) | Proven; *derives* discrete Hölder cone stability from the substrate-engineering observable Ξ_ℓ, with effective constant K_eff = c_𝓡⁻¹ · K |
| Theorem 4 (R4′ + R5 + R6 + H5 ⟹ H8) | Proven via Lemma 4.1 + companion paper's GH machinery; continuum constant K_𝒞 = K_eff · C_GH; H8 is derived rather than postulated |
| Theorem 5 (Main synthesis) | Proven by conjunction; gives regular TPB substrate + H9 ⟹ smooth Lorentzian continuum |
| R1 derivation from K = 7 architecture | Open; Δ should follow from closure-catalogue structure |
| R2 derivation from σ-duality refinement-functor structure | Open; M should follow from refinement-functor properties |

| Result | Status |
|---|---|
| A1* derivation from A1 | Open; conjecture A1 + refinement compatibility ⟹ A1* |
| A2* derivation from A2 | Open; the §4.2 internal-logic argument (A2's "new fact production" content forces operational distinctness, since a commitment event with an operationally indistinguishable outcome would not satisfy A2's own criterion) is essentially deductive but has not yet been formally discharged within the BCB/closure axiomatic setting of the *sequential-interface-transport* paper |
| R4′ derivation from K = 7, with α = 1 desideratum | Open; α = 1 highest-value sharpening (eliminates §10.10 weak-norm gap entirely) |
| R5 derivation from σ-duality refinement-functor structure | Open |
| R6 derivation from σ-duality refinement-functor structure | Open; lower bi-Lipschitz constant c_𝓡 > 0 is the substantive content (the upper constant follows from R4′ + H5 + R2) |
| R7 derivation from K = 7 + R4′ | Open; R1 + R2 alone admit Δ-regular-tree counterexamples, so R7 is independent of R1 + R2 in both the upper-bound and lower-bound directions |
| H9 for cosmological topologies | Automatic by Freedman–Quinn; verification for concrete models open |
| H9 for exotic topologies | Genuinely external; case-by-case 4-manifold analysis required |
| Quantitative constants for concrete substrates | Open; primary content of next paper |
| Companion §10.10 weak-norm robustness | Inherited unchanged from companion paper; *contingent* — closes automatically if K = 7 ⟹ α = 1 in R4′ |
| Companion §10.8 C^∞ desideratum | Inherited unchanged from companion paper |

---

# 14. Conclusion

The companion paper *Continuum-Limit Regularity and Cone Convergence in VERSF* established that admissible TPB refinement converges to a smooth Lorentzian continuum conditional on five external regularity hypotheses H6, H6′, H7, H8, H9. Its §15 conclusion named the substrate-level derivation of these hypotheses as the principal remaining mathematical risk in the VERSF geometry programme.

The present paper supplies that derivation for four of the five hypotheses.

The substrate-engineering layer consists of seven explicit constraints — bounded local coupling (R1), controlled refinement multiplicity (R2), uniform distinguishability separation (R3), bounded refinement distortion (R4′), bounded refinement overlap (R5), refinement bi-Lipschitz on cone-direction sets (R6), and uniform Ahlfors d-regularity in pseudometric balls (R7) — applied in addition to the baseline VERSF substrate axioms A0–A4 (with A1 strengthened to A1* and A2 strengthened to A2*), BCB, H1, and H5. Substrates satisfying R1–R7 in this enlarged setting are called *regular TPB substrates*. The main theorem (§10) shows that every regular TPB substrate satisfies H6 (Theorem 1, from R7 — with R1 + R2 alone insufficient since they admit Δ-regular-tree exponential growth, hence R7 is a separately-named engineering constraint), H7 (Theorem 2, from R1 + K = 7 closure architecture), H6′ (Theorem 3, from A1* + A2* + R3 + R5 via the direct A1*+A2* contradiction and across-level R5 propagation), and H8 (Theorem 4, derived from R4′ + R5 + R6 + H5 via the refinement-distortion-propagation Lemma 4.1, which uses R6's lower bi-Lipschitz constant $c_{\mathcal{R}} > 0$ to translate the Hölder bound on the substrate-engineering observable $\Xi_\ell$ back to a Hölder bound on the pre-refinement cones). Conditional on H9 (smooth-structure existence on the continuum limit), regular TPB substrates therefore generate the full hypothesis set H1–H9 of the companion paper and hence the smooth Lorentzian continuum.

The substrate-engineering content of R4′ is concentrated in a quantitative bound on a measurable substrate-engineering observable — the *refinement-distortion functional* $\Xi_\ell(x, y) := d_H(\mathcal{R}(\mathcal{C}_\ell(x)), \mathcal{R}(\mathcal{C}_\ell(y)))$ — rather than in any postulated regularity of the cones themselves. The discrete Hölder cone stability of $\mathcal{C}_\ell$ used downstream is derived from R4′ + R5 + R6 + H5 via Lemma 4.1 in §8.3. $\Xi_\ell$ can in principle be computed directly on any concrete TPB substrate model, making R4′ a substrate-engineering observable rather than a postulate about the limit. The §8.7 conjecture that K = 7 closure architecture + R5 + closure stability ⟹ R4′ (with α = 1 as the highest-value target) makes the link to the *sequential-interface-transport* paper's closure architecture explicit and gives the next paper a concrete target. R6's lower bi-Lipschitz constant $c_{\mathcal{R}} > 0$ and R7's polynomial growth dimension d are conjecturally further consequences of the K = 7 + σ-duality architecture, but neither has been discharged here.

The fifth hypothesis H9 is structurally different: it is topological rather than dynamical, intersects the Donaldson–Freedman exotic-4-manifold theory, and cannot be derived from local substrate dynamics. For continuum limits that are compact perturbations of cosmologically natural 4-manifolds (Minkowski, FLRW, asymptotically flat black-hole geometries), H9 is automatic by Freedman–Quinn classification. For exotic continuum topologies it is a genuine restriction.

The principal mathematical risk in the VERSF geometry programme therefore moves one further layer down: from *external regularity hypotheses* (companion paper) to *substrate engineering constraints* (this paper) to *concrete substrate models satisfying R1–R7 with quantitative constants — particularly an explicit value of the refinement-distortion exponent α, the lower bi-Lipschitz constant $c_{\mathcal{R}}$, and the Ahlfors-regularity constants ($d$, $c_{vol}$, $C_{vol}$) — and producing cosmologically natural continuum topology* (next paper). The natural next-paper task is to exhibit such concrete substrate models, with particular attention to the refinement-distortion functional $\Xi_\ell$ — substrates achieving α = 1 (Lipschitz refinement distortion) would eliminate the wave-equation correction step in the companion paper's §10.7 and substantially simplify the bridge to the Lorentzian-emergence machinery.

The key methodological move is the explicit factoring of the geometry-emergence chain into four layers — substrate engineering (this paper) → continuum regularity (companion paper) → Lorentzian emergence (companion-companion paper) → field-theoretic structure (downstream papers) — with each layer concentrating a different category of residual risk on top of a mechanically tight derivation chain. With the present paper in place, the substrate-engineering layer is the first explicitly named and the first whose contents are concrete substrate-architectural constraints rather than abstract regularity hypotheses. The geometry-emergence chain is, in this layered architecture, *conditionally closed* end-to-end from substrate dynamics to Lorentzian geometry, modulo the three named residual gaps — substrate-level derivation of R1–R7 with explicit constants (particularly the value of α in R4′, $c_{\mathcal{R}}$ in R6, and the Ahlfors-regularity ratio $C_{vol} / c_{vol}$ and dimension d in R7), the topological condition H9 for non-cosmologically-natural limits, and the §10.10 weak-norm-robustness gap inherited from the companion paper (contingent: closes automatically if R4′ holds with α ≥ 1).

The remaining work is to construct concrete TPB substrate models verifying R1–R7 quantitatively — computing the refinement-distortion functional $\Xi_\ell$, the R6 bi-Lipschitz constants, and the R7 Ahlfors-regularity constants directly on those models, and ideally discharging R7 in both directions from the K = 7 + R4′ conjecture rather than assuming it — to confirm H9 for those models' continuum limits, and to close the §10.10 weak-norm-robustness gap via revision to the Lorentzian-emergence paper's §6.1 along the lines of the three candidate fixes identified there (or by establishing α ≥ 1 in R4′, which closes the gap automatically). None of these residual tasks requires further geometric-emergence machinery to be developed; all are concrete substrate-engineering or revision tasks that the layered architecture now makes individually targetable.