

Explicit Refinement Dynamics and Stability Flow in VERSF

Closure Evolution Operators, Refinement Coherence, and the Emergence of Lipschitz Continuum Geometry

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The previous paper showed that the regularity assumptions needed for a smooth continuum are not arbitrary. They can be traced to substrate-engineering constraints: bounded local coupling, controlled refinement multiplicity, transport sparsity, and bounded refinement distortion.

That left one deeper question:

What actual substrate dynamics produce those engineering constraints?

This paper answers that question by introducing explicit refinement dynamics. Instead of treating refinement as an abstract map from one level to the next, we model it as a concrete evolution process acting on local closure states, transport directions, and admissible cone structures.

The central idea is that the substrate carries a *finite closure catalogue* — in VERSF, the $K = 7$ closure structure — and refinement acts through admissibility-preserving transitions inside this catalogue. Because the catalogue is finite and the transitions preserve admissibility, refinement cannot drift chaotically. Nearby substrate regions remain closure-compatible as they refine.

The key technical object is the **refinement evolution operator**:

$$\mathcal{E}_\ell : (\mathcal{P}_\ell, \mathcal{T}_\ell, \mathcal{B}_\ell, \mathcal{C}_\ell, \mathcal{K}) \rightarrow (\mathcal{P}_{\ell+1}, \mathcal{T}_{\ell+1}, \mathcal{B}_{\ell+1}, \mathcal{C}_{\ell+1}, \mathcal{K}).$$

It describes how the substrate updates from one refinement level to the next.

The main goal is to prove that admissible refinement evolution generates **Lipschitz refinement stability**:

$$\Xi_\ell(x, y) \leq K \cdot d_\ell(x, y) + \varepsilon_\ell.$$

This is the special case $\alpha = 1$ of the previous paper's refinement-distortion bound. It means that nearby substrate regions refine in proportionally nearby ways — microscopic structure does not inject uncontrolled roughness into the emerging continuum.

The result matters because $\alpha = 1$ removes the weak-norm bridge problem left in the continuum-limit paper. It allows the geometry-emergence chain to proceed with Lipschitz-controlled cone fields and pointwise-stable transport currents, rather than relying on weaker Sobolev/ L^2 corrections.

The result moves the programme from

regular TPB substrate \Rightarrow Lorentzian geometry

to

admissible refinement dynamics \Rightarrow regular TPB substrate \Rightarrow Lorentzian geometry.

That is the next major step in turning VERSF into a genuine substrate-to-geometry reconstruction framework.

Contents

- [General-Reader Summary](#)
 - [Abstract](#)
 - [1. Introduction](#)
 - [2. Inputs from the Previous Paper](#)
 - [3. Closure States and Substrate Refinement Dynamics](#)
 - [3.1 The finite closure catalogue](#)
 - [3.2 The closure-generated cone map](#)
 - [3.3 The refinement evolution operator](#)
 - [4. Admissibility Conditions on Refinement Evolution](#)
 - [5. The Refinement-Coherence Functional](#)
 - [6. Theorem 1 — Coherence Contraction Implies Lipschitz Closure Stability](#)
 - [7. Theorem 2 — Lipschitz Closure Stability Implies Lipschitz Cone Stability](#)
 - [8. Theorem 3 — Explicit Refinement Dynamics Imply R4' with \$\alpha = 1\$](#)
 - [9. Corollary — Lipschitz Continuum Cone Field](#)
 - [10. Main Theorem](#)
 - [11. Stability Flow Interpretation](#)
 - [12. Why This Matters](#)
 - [13. Limitations and Open Problems](#)
 - [14. Conclusion](#)
-

Abstract

The preceding paper, *Deriving Regularity Hypotheses from Admissible TPB Dynamics in VERSF*, introduced the substrate-engineering constraints R1–R4' and showed that they imply the continuum regularity hypotheses H6, H6', H7, and H8 required by *Continuum-Limit Regularity and Cone Convergence in VERSF*. The most delicate of these engineering constraints was R4', bounded refinement distortion, expressed through the refinement-distortion functional

$$\Xi_{\ell}(x, y) = d_H(R_{\{\ell \rightarrow \ell+1\}}(C_{\ell}(x)), R_{\{\ell \rightarrow \ell+1\}}(C_{\ell}(y))).$$

That paper identified the high-value target $\alpha = 1$, corresponding to Lipschitz refinement distortion,

$$\Xi_{\ell}(x, y) \leq K \cdot d_{\ell}(x, y) + \varepsilon_{\ell},$$

because Lipschitz cone stability removes the weak-norm robustness gap in the smooth-regularity bridge of the continuum-limit paper.

The present paper supplies the next substrate-dynamical layer. We replace the abstract refinement functor $R_{\{\ell \rightarrow \ell+1\}}$ with an explicit **refinement evolution operator**

$$\mathcal{E}_{\ell} : (P_{\ell}, T_{\ell}, B_{\ell}, C_{\ell}, \mathcal{K}) \rightarrow (P_{\{\ell+1\}}, T_{\{\ell+1\}}, B_{\{\ell+1\}}, C_{\{\ell+1\}}, \mathcal{K}),$$

acting on a *fixed ℓ -independent* finite closure catalogue \mathcal{K} with $K = 7$. The operator carries a **position-indexed closure-update family** $\{\mathcal{E}_{\{\ell, x\}^K}\}_{x \in P_{\ell}}$ of maps $\mathcal{K} \rightarrow \mathcal{K}$, with the *equivariance* $\chi_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x) = \mathcal{E}_{\{\ell, x\}^K}(\chi_{\ell}(x))$. Position-indexing is essential: a single ℓ -indexed closure-update map would force the L-term below to vanish, since equal parent closures would refine to equal children regardless of substrate distance.

We introduce a **refinement-coherence functional** $\mathcal{C}_{\ell}(x, y)$, measuring closure mismatch between neighbouring substrate regions, and decompose the closure-coherence condition into three independently meaningful sub-conditions:

- **E4a** — pointwise contractivity: $\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, x\}^K}(\kappa')) \leq \beta \cdot \rho(\kappa, \kappa')$;
- **E4b** — position-coherence: $\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, y\}^K}(\kappa)) \leq L \cdot d_{\ell}(x, y)$;
- **E4c** — uniform admissibility gap: $\beta < 1$ uniformly in ℓ .

A lemma combines E4a + E4b via the triangle inequality to give the closure-coherence inequality

$$\mathcal{C}_{\{\ell+1\}}(\mathcal{E}_{\ell} x, \mathcal{E}_{\ell} y) \leq \beta \cdot \mathcal{C}_{\ell}(x, y) + L \cdot d_{\ell}(x, y).$$

Together with a two-sided refinement-distance compatibility (D') and a cone-refinement compatibility axiom E5 connecting $\{\mathcal{E}_{\{\ell, x\}^K}\}$ to the abstract refinement functor R via Φ , this yields Lipschitz control of the refinement-distortion functional:

$$\Xi_{\ell}(x, y) \leq K_{\Xi} \cdot d_{\ell}(x, y) + \varepsilon_{\ell}, \quad K_{\Xi} = L_{\Phi} \cdot L \cdot A^2 / (A_{-} \cdot (1 - \beta)).$$

Consequently, admissible refinement dynamics generate $R4'$ with $\alpha = 1$. Combined with the previous paper, this proves that explicit closure-stable TPB refinement dynamics satisfy H8 with Lipschitz regularity, and hence feed into the continuum-limit chain with stronger smoothness control.

The principal result is therefore

finite ℓ -independent closure catalogue + position-indexed admissible refinement evolution + closure-coherence sub-conditions + cone-refinement compatibility $\Rightarrow \alpha = 1 \Rightarrow$ Lipschitz continuum cone stability.

This moves the VERSF geometry programme from a substrate-engineering conditional result to a dynamical refinement-flow result.

1. Introduction

The geometry branch of the VERSF programme has developed in three stages.

Stage I. *Structural Necessity of Lorentzian Geometry in VERSF* showed that, given a smooth strongly causal continuum equipped with invariant causal cones and conserved transport density, Lorentzian geometry is forced.

Stage II. *Continuum-Limit Regularity and Cone Convergence in VERSF* supplied the convergence bridge: under regularity hypotheses H6, H6', H7, H8, and H9, admissible TPB refinement sequences converge to a strongly causal Lorentzian length space admitting C^k Lorentzian metric realisations for every prescribed $k \geq 2$.

Stage III. *Deriving Regularity Hypotheses from Admissible TPB Dynamics in VERSF* showed that H6, H6', H7, and H8 follow from substrate-engineering constraints R1–R4'. The most important of these was R4', bounded refinement distortion.

R4' remained, however, an *engineering constraint*. It identified the quantity to control,

$$\Xi_\ell(x, y),$$

but it did not yet derive that control from explicit substrate update laws.

Stage IV (this paper). We supply the missing dynamical layer.

The central question is:

What refinement dynamics force $\Xi_\ell(x, y) \leq K \cdot d_\ell(x, y) + \varepsilon_\ell$ — that is, $\alpha = 1$?

We argue that Lipschitz refinement distortion follows from three structural facts:

1. the local closure catalogue is **finite**;
2. refinement acts through **admissibility-preserving** closure transitions;
3. closure mismatch between neighbouring regions is **contractive** under refinement.

Together, these imply that nearby substrate regions cannot diverge chaotically under refinement: their admissible transport cones remain Lipschitz-close.

The strategy is to introduce a refinement-coherence functional measuring closure mismatch, prove it contracts under refinement, lift this contraction to cone-distance through Lipschitz response of the closure-to-cone map, and finally identify the lifted bound with $R4'$ in its $\alpha = 1$ form. The result is a chain of three theorems whose composition yields the substrate-to-geometry result.

Epistemic status. The main theorem chain (Theorems 1, 2, 3) is *proven conditional* on the substrate primitives E1–E3 (kinematic; derived from admissibility, H5, BCB), the dynamical conditions E4a, E4b, E4c, E5 (the load-bearing hypotheses of this paper), the two-sided refinement-distance compatibility (D'), and the Lipschitz response of Φ (automatic from the finite catalogue, Proposition 7.0). E4a–E4c constrain the position-indexed family $\{ \mathcal{E}_{\ell, x}^K \}$ as it acts within the closure catalogue (\mathcal{K}, ρ) ; E5 constrains how this family interfaces with the abstract refinement functor R via the closure-to-cone map Φ . The non-collapse half of (D') is conditional. All conditional hypotheses are physically motivated by the $K = 7$ closure algebra and the substrate's local-coupling architecture; their derivation from those primitives is an open structural problem identified in §13.

2. Inputs from the Previous Paper

The previous paper introduced the **refinement-distortion functional**

$$\Xi_{\ell}(x, y) = d_H(R_{\ell \rightarrow \ell+1}(C_{\ell}(x)), R_{\ell \rightarrow \ell+1}(C_{\ell}(y))),$$

where d_H denotes the Hausdorff distance between admissible cone-direction sets at level $\ell + 1$.

It introduced $R4'$ in the general Hölder form

$$\Xi_{\ell}(x, y) \leq K \cdot d_{\ell}(x, y)^{\alpha} + \varepsilon_{\ell}, \quad \alpha \in (0, 1],$$

and established the implication

$$R4' + H5 + R2 + \text{bounded refinement overlap} \Rightarrow H8.$$

The high-value target was identified as $\alpha = 1$ — Lipschitz stability,

$$\Xi_{\ell}(x, y) \leq K \cdot d_{\ell}(x, y) + \varepsilon_{\ell},$$

because Lipschitz cone stability upgrades the continuum cone field to Lipschitz regularity and eliminates the need for a weak-norm bridge in the continuum-limit paper.

The task of the present paper is therefore to derive R4' with $\alpha = 1$ from explicit refinement dynamics.

3. Closure States and Substrate Refinement Dynamics

3.1 The finite closure catalogue

The VERSF substrate carries a fixed finite catalogue of admissible local closure states

$$\mathcal{K} = \{ \kappa_1, \dots, \kappa_K \}, K = 7,$$

independent of refinement level. The catalogue is equipped with a fixed finite-metric structure

$$\rho : \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}_{\geq 0},$$

normalised so that

$$\rho(\kappa, \kappa) = 0, \rho_{\min} := \min_{\{\kappa \neq \kappa'\}} \rho(\kappa, \kappa') > 0, \rho_{\max} := \max_{\{\kappa \neq \kappa'\}} \rho(\kappa, \kappa') < \infty.$$

The ℓ -independence of (\mathcal{K}, ρ) is a structural commitment of the present paper: the closure catalogue is identified across refinement levels, and refinement acts *within* this fixed catalogue rather than between distinct level-indexed catalogues. The closure-level data carried by a substrate region still varies with ℓ — but only through the level-indexed *assignment* χ_ℓ defined below, not through any ℓ -dependence in (\mathcal{K}, ρ) themselves. This identification is essential for the uniformity statements in §4 and Proposition 7.0 below: without ℓ -independent ρ_{\min} , the constant L_Φ in Theorem 2 could blow up as the mesh shrinks.

Physical meaning of ρ . The metric ρ is not a free auxiliary device chosen merely to make Φ Lipschitz. It must encode *physically meaningful* closure separation: two closure states $\kappa, \kappa' \in \mathcal{K}$ are close in ρ only when their admissible transport directions, local compatibility constraints, and refinement-child sets are structurally similar. Thus ρ should ultimately be derived from the $K = 7$ closure graph itself — for example by closure-graph distance, weighted admissibility-transition distance, or mismatch count over admissibility constraints. Proposition 7.0 uses only finiteness of \mathcal{K} , but the *physical interpretation* of the Lipschitz constant L_Φ depends on ρ being tied to substrate structure rather than chosen arbitrarily. A free choice of ρ that minimised L_Φ artificially would yield a Lipschitz bound without physical content; ρ derived from the closure graph yields a Lipschitz bound whose constant reflects genuine substrate structure.

Each closure state $\kappa \in \mathcal{K}$ determines

- admissible local transport directions,
- allowed closure transitions,
- local compatibility constraints with adjacent closure states,
- admissible refinement children.

Define the **closure assignment**

$$\chi_\ell : P_\ell \rightarrow \mathcal{K},$$

where $\chi_\ell(x)$ is the closure state associated with substrate node x at level ℓ .

3.2 The closure-generated cone map

Each closure state determines a cone-direction set through a *level-independent* map

$$\Phi : \mathcal{K} \rightarrow \mathcal{C},$$

where \mathcal{C} is the space of admissible cone-direction sets, equipped with Hausdorff distance d_H . Since \mathcal{K} is finite and Φ is ℓ -independent, the image $\Phi(\mathcal{K})$ is a subset of \mathcal{C} of cardinality at most $K = 7$, with bounded diameter

$$d_{H,\max} := \max_{\{\kappa, \kappa' \in \mathcal{K}\}} d_H(\Phi(\kappa), \Phi(\kappa')) < \infty.$$

The cone field is then a *derived* quantity:

$$C_\ell(x) = \Phi(\chi_\ell(x)).$$

This relegates cone structure to a secondary role: cones are produced by closure states, not posited independently. The closure assignment carries the primary substrate information.

Architectural cost (cones-as-derived). This closure-generated formulation is a structural commitment that should be flagged. Prior geometry papers in the programme treated the cone field C_ℓ as a *primary* substrate object, with the family $\{C_\ell(x) : x \in P_\ell\}$ ranging over the full space of admissible cone fields. The present paper restricts C_ℓ to the image of $\Phi \circ \chi_\ell$, which is in general a strict subset of admissible cone fields. Earlier results that quantified over arbitrary admissible cone fields therefore continue to hold under the present formulation — the image is a subset of the original domain — but the substrate degrees of freedom are narrowed. Readers of the broader programme should treat this as a load-bearing assumption rather than a notational convenience; it is the *price* that the cones-as-derived architecture pays for the automatic Lipschitz response of Φ established in Proposition 7.0.

3.3 The refinement evolution operator

Define the **refinement evolution operator**

$$\mathcal{E}_\ell : (P_\ell, T_\ell, B_\ell, C_\ell, \mathcal{K}) \rightarrow (P_{\ell+1}, T_{\ell+1}, B_{\ell+1}, C_{\ell+1}, \mathcal{K}).$$

It decomposes into four components

$$\mathcal{E}_\ell = (\mathcal{E}_\ell^{\wedge P}, \mathcal{E}_\ell^{\wedge T}, \mathcal{E}_\ell^{\wedge B}, \{ \mathcal{E}_{\ell, x}^{\wedge K} \}_{x \in P_\ell}),$$

with

- $\mathcal{E}_\ell^{\wedge P}$ refining substrate nodes ($P_\ell \rightarrow P_{\ell+1}$);
- $\mathcal{E}_\ell^{\wedge T}$ refining admissible transport relations ($\ll \ell \rightarrow \ll \ell+1$);
- $\mathcal{E}_\ell^{\wedge B}$ updating the bit measure subject to BCB;
- $\{ \mathcal{E}_{\ell, x}^{\wedge K} \}_{x \in P_\ell}$ updating closure states, *position-indexed* by the substrate site x .

The closure update is the **position-indexed equivariance**

$$\chi_{\ell+1}(\mathcal{E}_\ell^{\wedge P} x) = \mathcal{E}_{\ell, x}^{\wedge K}(\chi_\ell(x)).$$

Each $\mathcal{E}_{\ell, x}^{\wedge K}$ is a map $\mathcal{K} \rightarrow \mathcal{K}$. The family is allowed to vary with substrate position x , subject to a coherence condition (E4b below) that controls the variation by substrate distance.

Why position-indexing is essential. A single ℓ -indexed map $\mathcal{E}_\ell^{\wedge K} : \mathcal{K} \rightarrow \mathcal{K}$ would force

$$\chi_\ell(x) = \chi_\ell(y) \Rightarrow \chi_{\ell+1}(\mathcal{E} x) = \chi_{\ell+1}(\mathcal{E} y)$$

for *all* x, y , regardless of substrate distance. That would make the refined closure-mismatch $\rho(\chi_{\ell+1}(\mathcal{E} x), \chi_{\ell+1}(\mathcal{E} y))$ depend only on the closure pair $(\chi_\ell(x), \chi_\ell(y))$, independent of $d_\ell(x, y)$. The $L \cdot d_\ell(x, y)$ term in E4 below would then be vacuous. Position-indexing carries the genuine physical content that distant substrate regions can refine differently even when their parent closures match — and the coherence condition E4b supplies the L -term by bounding how much. The single-map formulation is recovered as the degenerate case in which $\mathcal{E}_{\ell, x}^{\wedge K}$ is independent of x ; the L -term is then forced to zero, and Theorem 1 reduces to a pure geometric-decay bound.

4. Admissibility Conditions on Refinement Evolution

We impose five admissibility conditions on \mathcal{E}_ℓ . Conditions E1–E3 are *derived* from established VERSF substrate primitives. Conditions E4 (split into E4a, E4b, E4c) and E5 are the new dynamical hypotheses that carry the central work of this paper.

E1 — Closure preservation (proven from admissibility)

For every $x \in P_\ell$,

$$\mathcal{E}_{\ell, x}^{\wedge K}(\mathcal{K}) \subseteq \mathcal{K}.$$

Each position-indexed closure-update map sends admissible closure states to admissible closure states. This is a consequence of the closure catalogue being defined as the admissibility-stable set under refinement transitions.

E2 — Transport preservation (proven from H5)

If $x \preceq_{\ell} y$, then

$$\mathcal{E}_{\ell}^{\wedge P}(x) \preceq_{\{\ell+1\}} \mathcal{E}_{\ell}^{\wedge P}(y).$$

This is the dynamical realisation of H5 (faithful refinement of the transport relation).

E3 — BCB preservation (proven from BCB, uniform in bounded sets)

There exists a function $\omega : \mathbb{R}_{\{>0\}} \times \mathbb{N} \rightarrow \mathbb{R}_{\{\geq 0\}}$ with $\omega(V, \ell) \rightarrow 0$ as $\ell \rightarrow \infty$ for each fixed V , such that for every bounded region $U \subset P_{\ell}$ with bit-volume $B_{\ell}(U) \leq V$,

$$|B_{\{\ell+1\}}(\mathcal{E}_{\ell}^{\wedge P} U) - B_{\ell}(U)| \leq \omega(V, \ell).$$

The $o(1)$ error is *uniform over U at fixed bit-volume bound* — the rate of convergence may degrade with V but is otherwise independent of which U of given volume is chosen. This is the dynamical realisation of bit conservation under refinement.

E4 — Closure coherence (conditional; physically motivated)

This is the central new dynamical condition, decomposed into three independently meaningful sub-conditions.

E4a — Pointwise contractivity of $\mathcal{E}_{\{\ell, x\}}^{\wedge K}$. There exists $0 \leq \beta < 1$, independent of (ℓ, x) , such that for every $x \in P_{\ell}$ and every $\kappa, \kappa' \in \mathcal{K}$,

$$\rho(\mathcal{E}_{\{\ell, x\}}^{\wedge K}(\kappa), \mathcal{E}_{\{\ell, x\}}^{\wedge K}(\kappa')) \leq \beta \cdot \rho(\kappa, \kappa').$$

Each position-fixed closure-update map is a strict contraction on (\mathcal{K}, ρ) .

Scope of E4a. Strict contractivity on a finite catalogue should not be read as an arbitrary global compression imposed for convenience. The *intended domain* of E4a is the **admissible local-transition sector** of the closure graph: pairs (κ, κ') of closure states that can arise as locally compatible alternatives within the same refinement neighbourhood — i.e., the physically admissible comparison relation on \mathcal{K} . If one extends ρ to all of $\mathcal{K} \times \mathcal{K}$, the global statement above is a convenient shorthand; the load-bearing physical requirement is contraction on *admissibly comparable* closure pairs. E4a then says that refinement filters locally inconsistent alternatives; it does *not* say that every abstract pair of catalogue labels is dynamically forced together. The $K = 7$ catalogue paper that derives E4a should specify this admissible comparison relation explicitly (see §13).

E4b — Position-coherence of the family $\{\mathcal{E}_{\ell,x}^{\wedge K}\}$. There exists $L < \infty$, independent of (ℓ, κ) , such that for every $x, y \in P_{\ell}$ and every $\kappa \in \mathcal{K}$,

$$\rho(\mathcal{E}_{\ell,x}^{\wedge K}(\kappa), \mathcal{E}_{\ell,y}^{\wedge K}(\kappa)) \leq L \cdot d_{\ell}(x, y).$$

The position-indexed family varies Lipschitz-continuously in the substrate position. This is the condition that supplies the L -term in the closure-coherence inequality below; without it, distant substrate regions could refine identically-labelled closures to wildly different children.

E4c — Uniform admissibility gap (uniformity of β). There exists $\beta \in [0, 1)$ such that the position- and level-uniform contraction rate β_{ℓ} in E4a satisfies

$$\beta_{\ell} \leq \beta \text{ for every } \ell.$$

Equivalently, the admissibility-rejection rate of the closure-transition graph is bounded below by $1 - \beta$ at every level — the closure graph does not develop fixed-point-like neighbourhoods as the mesh shrinks. This is a *spectral-gap-style* condition in the refinement limit: it forbids the contraction rate from approaching unity as $\ell \rightarrow \infty$. The weaker " $\sup_{\ell} \beta_{\ell} < 1$ " reading is equivalent when the supremum is attained or bounded away from 1, but the explicit existence form removes any ambiguity in the borderline case $\beta_{\ell} \rightarrow 1$ with $\sup = 1$. Without E4c, the geometric series in Theorem 1 would not converge uniformly across ℓ , and $K_{\mathcal{C}}$ could blow up at the continuum limit.

Lemma 4.1 (E4-inequality as a consequence of E4a + E4b). Under the position-indexed equivariance $\chi_{\ell+1}(\mathcal{E}_{\ell}^{\wedge P} x) = \mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(x))$, conditions E4a and E4b together imply

$$\rho(\chi_{\ell+1}(\mathcal{E}_{\ell} x), \chi_{\ell+1}(\mathcal{E}_{\ell} y)) \leq \beta \cdot \rho(\chi_{\ell}(x), \chi_{\ell}(y)) + L \cdot d_{\ell}(x, y).$$

Proof. Apply the triangle inequality on (\mathcal{K}, ρ) through the intermediate point $\mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(y))$:

$$\rho(\mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(x)), \mathcal{E}_{\ell,y}^{\wedge K}(\chi_{\ell}(y))) \leq \rho(\mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(x)), \mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(y))) + \rho(\mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(y)), \mathcal{E}_{\ell,y}^{\wedge K}(\chi_{\ell}(y))).$$

The first term is bounded by $\beta \cdot \rho(\chi_{\ell}(x), \chi_{\ell}(y))$ by E4a; the second term is bounded by $L \cdot d_{\ell}(x, y)$ by E4b. The LHS equals $\rho(\chi_{\ell+1}(\mathcal{E}_{\ell} x), \chi_{\ell+1}(\mathcal{E}_{\ell} y))$ by the position-indexed equivariance. ■

Remark on path-symmetry. The intermediate point $\mathcal{E}_{\ell,x}^{\wedge K}(\chi_{\ell}(y))$ was chosen; the symmetric path through $\mathcal{E}_{\ell,y}^{\wedge K}(\chi_{\ell}(x))$ yields the same numerical bound (E4b dominates the first triangle leg, E4a the second). The result is independent of path choice.

Status of E4a–E4c. All three sub-conditions are *conditional*. Their physical motivation comes from the finite $K = 7$ closure algebra: admissibility-preserving transitions between distant closure states require multi-step paths in the catalogue graph, and at each refinement step admissibility filtering rejects a constant fraction of locally inconsistent transitions (E4a + E4c); the position-coherence E4b reflects the local-coupling architecture of the substrate, in which spatially close

regions cannot independently choose closure children. Derivation of E4a–E4c from the explicit $K = 7$ catalogue topology is an open structural problem identified in §13.

E5 — Cone-refinement compatibility (axiom; status discussed below)

The abstract refinement functor $R_{\{\ell \rightarrow \ell+1\}}$ acting on cone-direction sets and the substrate-level refinement evolution \mathcal{E}_ℓ acting on substrate sites are compatible up to a vanishing error: there exists a sequence $\delta_\ell \rightarrow 0$ such that for every $x \in P_\ell$,

$$d_H(R_{\{\ell \rightarrow \ell+1\}}(C_\ell(x)), C_{\{\ell+1\}}(\mathcal{E}_\ell \wedge P x)) \leq \delta_\ell.$$

That is, $R_{\{\ell \rightarrow \ell+1\}} \circ C_\ell = C_{\{\ell+1\}} \circ \mathcal{E}_\ell \wedge P$ up to $\delta_\ell \rightarrow 0$ in d_H . The sequence δ_ℓ is inherited from the substrate-mesh contraction.

Status of E5. Under the closure-generated cone formulation $C_\ell = \Phi \circ \chi_\ell$, E5 is essentially a *structural compatibility* between the position-indexed family $\{\mathcal{E}_{\{\ell, x\}^K}\}$ and the action of the abstract refinement functor R on cone sets through Φ . It can be read in either of two ways:

- (i) *as an axiom* — \mathcal{E} is defined to be R -compatible, and $\delta_\ell \rightarrow 0$ is part of the definition;
- (ii) *as a derivation* — given an abstract R and a Φ -compatible choice of position-indexed $\{\mathcal{E}_{\{\ell, x\}^K}\}$, the equality holds with δ_ℓ inherited from substrate-mesh contraction.

The present paper treats E5 as an axiom; deriving it from a deeper specification of \mathcal{E} in relation to R is an open structural problem (§13).

Load-bearing status of E5. E5 is not a harmless technical condition. It is the *axiom that identifies* the abstract cone-refinement functor R with the concrete closure evolution \mathcal{E} through Φ . Without E5, Theorems 1 and 2 still prove Lipschitz stability of the *closure-generated* cone field $C_\ell = \Phi \circ \chi_\ell$ along explicit refinement lineages, but they do *not* yet prove that the abstract refinement-distortion functional Ξ_ℓ — which is defined in terms of R , not \mathcal{E} — is controlled. Theorem 3 therefore depends essentially on E5 to bridge from "closure evolution is stable" to "refinement distortion is bounded". The present paper should be read as a *conditional compatibility theorem*: once closure evolution and cone refinement are shown to agree up to $\delta_\ell \rightarrow 0$, explicit admissible refinement dynamics imply $R4'$ with $\alpha = 1$. Deriving E5 from a deeper construction of R or \mathcal{E} — or constructing one from the other so that E5 is identically satisfied — remains one of the central next tasks (§13).

5. The Refinement-Coherence Functional

Define the **refinement-coherence functional**

$$\mathcal{C}_\ell(x, y) := \rho(\chi_\ell(x), \chi_\ell(y)).$$

This is the closure-mismatch measure between two substrate regions at level ℓ . It is bounded above by ρ_{\max} uniformly and is zero precisely when x and y carry the same closure state.

Under Lemma 4.1 (E4a + E4b + position-indexed equivariance), the functional satisfies the one-step inequality

$$\mathcal{C}_{\{\ell+1\}}(\mathcal{E}_{\ell} x, \mathcal{E}_{\ell} y) \leq \beta \cdot \mathcal{C}_{\ell}(x, y) + L \cdot d_{\ell}(x, y).$$

The first term contracts geometrically with each refinement step (with β bounded uniformly below 1 by E4c). The second term is small when x and y are close in substrate distance. Together they prevent closure mismatch from amplifying under refinement.

We will also need a **two-sided refinement-compatibility** estimate on substrate distance: distances are approximately preserved under refinement, neither blowing up nor collapsing.

(D') — Two-sided refinement-distance compatibility. There exist a constant $A \in [1, A_{\max}]$, a *non-collapse constant* $A_{-} \in (0, 1]$ (so that A_{-}^{-1} controls how much distances can shrink under refinement), and a sequence $\eta_j(x, y) \rightarrow 0$ as the mesh shrinks (uniformly over bounded substrate pairs), such that for all $j \geq 0$ and refinement-descendant pairs (x_j, y_j) ,

$$A_{-} \cdot d_{\ell}(x, y) - \eta_j(x, y) \leq d_{\{\ell+j\}}(x_j, y_j) \leq A \cdot d_{\ell}(x, y) + \eta_j(x, y).$$

The upper bound is *proven* in the previous paper as part of bounded refinement multiplicity (R3) combined with bounded refinement overlap. The lower bound — a *non-collapse condition* preventing refinement descendants from being identified faster than the rate A_{-}^{-1} — is *conditional*. It is physically natural: refinement should not merge spatially distinct parent regions into coincident descendants. A derivation of the non-collapse bound from R3 plus a separation axiom on the refinement functor is an open problem (§13).

Uniformity remark on (D'). The η_j sequence is required to be uniform over bounded substrate pairs — pairs (x, y) with $d_{\ell}(x, y)$ bounded by some fixed scale d_{\max} . This uniformity is what makes the ε_{ℓ} -terms in Theorem 1 uniform in (x, y) .

6. Theorem 1 — Coherence Contraction Implies Lipschitz Closure Stability

Statement. Assume E1, E2, E3, E4a–E4c, and (D'). Then for every ℓ sufficiently large there exist constants $K_{\mathcal{C}} < \infty$ and a sequence $\varepsilon_{\ell} \rightarrow 0$, *uniform in (x, y) over bounded substrate pairs*, such that

$$\mathcal{C}_{\ell}(x, y) \leq K_{\mathcal{C}} \cdot d_{\ell}(x, y) + \varepsilon_{\ell}.$$

Explicitly, one may take

$$K_{\mathcal{C}} = L \cdot A / (A_{-} \cdot (1 - \beta)).$$

Lineage qualification. The theorem is most naturally stated for **refinement-generated pairs**: pairs (x, y) at level ℓ that arise as descendants of a bounded pair (x_0, y_0) at some seed level ℓ_0 . This is the native setting for a refinement-flow theorem and is what the proof below establishes directly. To promote the bound from refinement-generated pairs to *arbitrary* bounded pairs at level ℓ , a standard **lineage-covering condition** is required: every level- ℓ substrate site lies in the image of some refinement chain from a bounded seed level, and every bounded pair at level ℓ admits bounded ancestors up to the uniform η -error of (D') . Under this lineage-covering condition — which is plausible for genuine refinement systems but is *external* to the present argument — the descendant-pair result extends to all bounded substrate pairs at level ℓ . The statement above should therefore be read as conditional on lineage coverage when applied to arbitrary level- ℓ pairs.

Proof. Fix substrate points x_0, y_0 at some seed level ℓ_0 , and let

$$x_j = (\mathcal{E}_{\{\ell_0 + j - 1\}} \circ \dots \circ \mathcal{E}_{\{\ell_0\}})^{\mathcal{P}}(x_0), y_j = (\mathcal{E}_{\{\ell_0 + j - 1\}} \circ \dots \circ \mathcal{E}_{\{\ell_0\}})^{\mathcal{P}}(y_0)$$

denote their refinement descendants at level $\ell_0 + j$. Apply Lemma 4.1 once:

$$\mathcal{C}_{\{\ell_0 + 1\}}(x_1, y_1) \leq \beta \cdot \mathcal{C}_{\{\ell_0\}}(x_0, y_0) + L \cdot d_{\{\ell_0\}}(x_0, y_0).$$

Apply Lemma 4.1 again at level $\ell_0 + 1$:

$$\mathcal{C}_{\{\ell_0 + 2\}}(x_2, y_2) \leq \beta \cdot \mathcal{C}_{\{\ell_0 + 1\}}(x_1, y_1) + L \cdot d_{\{\ell_0 + 1\}}(x_1, y_1) \leq \beta^2 \cdot \mathcal{C}_{\{\ell_0\}}(x_0, y_0) + \beta \cdot L \cdot d_{\{\ell_0\}}(x_0, y_0) + L \cdot d_{\{\ell_0 + 1\}}(x_1, y_1).$$

Iterating n times yields

$$\mathcal{C}_{\{\ell_0 + n\}}(x_n, y_n) \leq \beta^n \cdot \mathcal{C}_{\{\ell_0\}}(x_0, y_0) + L \cdot \sum_{j=0}^{n-1} \beta^{n-1-j} \cdot d_{\{\ell_0 + j\}}(x_j, y_j).$$

By the uniform bound $\mathcal{C}_{\{\ell_0\}}(x_0, y_0) \leq \rho_{\max} < \infty$ and by the *upper* half of (D') ,

$$\mathcal{C}_{\{\ell_0 + n\}}(x_n, y_n) \leq \beta^n \cdot \rho_{\max} + L \cdot A \cdot d_{\{\ell_0\}}(x_0, y_0) \cdot \sum_{j=0}^{n-1} \beta^{n-1-j} + L \cdot \sum_{j=0}^{n-1} \beta^{n-1-j} \cdot \eta_j(x_0, y_0).$$

By E4c, $\beta < 1$ uniformly in ℓ , so the geometric series satisfies

$$\sum_{j=0}^{n-1} \beta^{n-1-j} = (1 - \beta^n) / (1 - \beta) \leq 1 / (1 - \beta).$$

Since $\eta_j \rightarrow 0$ uniformly over bounded pairs (uniformity clause of (D')), the tail term $\sum \beta^{n-1-j} \cdot \eta_j$ is itself $o(1)$ as $n \rightarrow \infty$ *uniformly in* (x_0, y_0) over bounded pairs by the following standard estimate. Fix $\varepsilon > 0$; choose N such that $\eta_j(x_0, y_0) < \varepsilon$ for all $j \geq N$ uniformly over bounded pairs (possible by (D')). Split the sum:

$$\sum_{j=0}^{n-1} \beta^{n-1-j} \cdot \eta_j = \sum_{j=0}^{N-1} \beta^{n-1-j} \cdot \eta_j + \sum_{j=N}^{n-1} \beta^{n-1-j} \cdot \eta_j.$$

The first sum is bounded by $\beta^{n-N} \cdot N \cdot \eta_{\max} \rightarrow 0$ as $n \rightarrow \infty$ (the geometric prefactor β^{n-N} eats the bounded $N \cdot \eta_{\max}$). The second sum is bounded by $\varepsilon / (1 - \beta)$. Since ε was arbitrary, the total is $o(1)$ uniformly. Define

$$K_{\mathcal{C}}^{\{\text{upper}\}} := L \cdot A / (1 - \beta), \varepsilon^{\{\ell_0 + n\}} := \beta^n \cdot \rho_{\max} + L \cdot \sum_{j=0}^{n-1} \beta^{n-1-j} \cdot \eta_j(x_0, y_0).$$

Both contributions to $\varepsilon^{\{\ell_0 + n\}}$ are uniform in (x_0, y_0) over bounded pairs (the first by ρ_{\max} being independent of (x_0, y_0) ; the second by the uniformity clause of (D')). Thus

$$\mathcal{C}_{\{\ell_0 + n\}}(x_n, y_n) \leq K_{\mathcal{C}}^{\{\text{upper}\}} \cdot d_{\{\ell_0\}}(x_0, y_0) + \varepsilon^{\{\ell_0 + n\}}.$$

Restating at the refined level. To convert the bound from one in $d_{\{\ell_0\}}$ to one in $d_{\{\ell_0 + n\}}$, we invoke the *lower* half of (D') — *once only*, at this final step. (Inside the iteration above, only the upper half of (D') was used, to dominate $d_{\{\ell_0 + j\}}(x_j, y_j)$ by $A \cdot d_{\{\ell_0\}}(x_0, y_0) + \eta_j$; the lower half plays no role until the very end.) Apply the lower half with $j = n$:

$$d_{\{\ell_0\}}(x_0, y_0) \leq A_-^{-1} \cdot (d_{\{\ell_0 + n\}}(x_n, y_n) + \eta_n(x_0, y_0)).$$

Substituting,

$$\mathcal{C}_{\{\ell_0 + n\}}(x_n, y_n) \leq (K_{\mathcal{C}}^{\{\text{upper}\}} / A_-) \cdot d_{\{\ell_0 + n\}}(x_n, y_n) + \varepsilon^{\{\ell_0 + n\}} + (K_{\mathcal{C}}^{\{\text{upper}\}} / A_-) \cdot \eta_n(x_0, y_0).$$

Define

$$K_{\mathcal{C}} := K_{\mathcal{C}}^{\{\text{upper}\}} / A_- = L \cdot A / (A_- \cdot (1 - \beta)), \varepsilon_{\ell} := \varepsilon^{\{\ell_0 + n\}} + K_{\mathcal{C}} \cdot \eta_n(x_0, y_0)$$

for $\ell = \ell_0 + n$.

Then $\varepsilon_{\ell} \rightarrow 0$ uniformly in (x_0, y_0) over bounded pairs, and

$$\mathcal{C}_{\ell}(x, y) \leq K_{\mathcal{C}} \cdot d_{\ell}(x, y) + \varepsilon_{\ell}.$$

This is the asserted Lipschitz closure stability. ■

Remark 6.1 (constants). The Lipschitz constant $K_{\mathcal{C}} = LA / (A_- (1 - \beta))$ is uniform across substrate regions and across refinement levels $\ell \geq \ell_0$. It depends on five substrate primitives: the position-coherence rate L (E4b), the upper refinement-distance scaling A and the non-collapse constant A_- ((D')), and the uniform contraction rate β (E4c). The asymmetry between A and A_- in the constant reflects the *non-collapse* requirement: in the limit $A_- \rightarrow 0$ the Lipschitz constant diverges, recovering the failure mode in which refinement collapses distinct parents to coincident descendants.

Remark 6.2 (uniformity). Three uniformity statements are load-bearing. (a) $\rho_{\max} < \infty$ uniformly across ℓ — guaranteed by ℓ -independence of \mathcal{K} (§3.1). (b) $\beta < 1$ uniformly across ℓ — guaranteed by E4c (uniform admissibility gap). (c) $\eta_j \rightarrow 0$ uniformly over bounded substrate pairs — clause of (D'). Failure of any of the three would make the ε_{ℓ} -tail non-uniform in (x_0, y_0) and the Lipschitz bound merely pointwise.

Remark 6.3 (alternative formulation). An alternative reading of Theorem 1 states the bound at the *coarse* level — $\mathcal{C}_{\ell_0}(x_0, y_0) \leq K_{\mathcal{C}^{\{\text{coarse}\}}} \cdot d_{\ell_0}(x_0, y_0) + \varepsilon_{\ell_0}$ — and lets the iteration deliver this bound by descending from $\ell_0 + n$ with $n \rightarrow \infty$. That formulation avoids invoking the lower half of (D') entirely and produces

$$K_{\mathcal{C}^{\{\text{coarse}\}}} = L \cdot A / (1 - \beta),$$

without the A_{-}^{-1} prefactor. The trade-off is that the continuum-limit constant K_{∞} would then need a separate A_{-} analysis to convert the substrate-level coarse bound into its continuum counterpart. The formulation adopted here places A_{-} explicitly in the constant where the continuum limit will need it; the coarse-level formulation defers that work to the limit step. Both readings give the same continuum $K_{\infty} = L_{\Phi} \cdot L \cdot A^2 / (A_{-} \cdot (1 - \beta))$.

7. Theorem 2 — Lipschitz Closure Stability Implies Lipschitz Cone Stability

The closure-to-cone map $\Phi : \mathcal{K} \rightarrow \mathcal{C}$ is defined on a finite metric space. We need its Lipschitz response in d_H to lift Theorem 1 to the cone field.

Proposition 7.0 (Lipschitz response of Φ from finite catalogue)

Statement. Let \mathcal{K} be a finite metric space with cardinality K and $\rho_{\min} := \min_{\{\kappa \neq \kappa'\}} \rho(\kappa, \kappa') > 0$. Then any map $\Phi : \mathcal{K} \rightarrow \mathcal{C}$ is Lipschitz with constant

$$L_{\Phi} := \max_{\{\kappa \neq \kappa'\}} d_H(\Phi(\kappa), \Phi(\kappa')) / \rho(\kappa, \kappa') \leq d_{H,\max} / \rho_{\min},$$

where $d_{H,\max} := \max_{\{\kappa, \kappa' \in \mathcal{K}\}} d_H(\Phi(\kappa), \Phi(\kappa'))$ is the cone-side diameter of the image.

Proof. For any $\kappa \neq \kappa'$ in \mathcal{K} , $d_H(\Phi(\kappa), \Phi(\kappa')) \leq d_{H,\max}$ and $\rho(\kappa, \kappa') \geq \rho_{\min}$, so the ratio is uniformly bounded by $d_{H,\max} / \rho_{\min}$. Equality with the maximum defines L_{Φ} . For $\kappa = \kappa'$ the bound is trivial. ■

This proposition is *automatic* for finite closure catalogues, and is the technical reason that the finiteness of \mathcal{K} is essential to the entire programme. With a continuous closure space the Lipschitz constant of Φ would need to be derived from a regularity condition on Φ ; with \mathcal{K} finite it follows from the discrete topology alone.

Uniformity in ℓ (Remark 7.1). Since (\mathcal{K}, ρ) and Φ are ℓ -independent by §3.1–§3.2, the constants ρ_{\min} , $d_{H,\max}$, and L_{Φ} are *all* ℓ -independent. This is what makes the Lipschitz response uniform across refinement levels. Were ρ_{\min} permitted to shrink with ℓ — for example via an ℓ -dependent rescaling of the catalogue metric — L_{Φ} could blow up as the mesh shrinks, and Theorem 2 would fail uniformly. The ℓ -independence of $(\mathcal{K}, \rho, \Phi)$ is therefore not merely a notational simplification but a structural prerequisite for the continuum-limit application.

Theorem 2

Statement. Assume E1, E2, E3, E4a–E4c, (D'), and ℓ -independence of $(\mathcal{K}, \rho, \Phi)$ (whence $L_{\Phi} < \infty$ uniformly by Proposition 7.0). Then the cone field is Lipschitz in substrate distance:

$$d_H(C_{\ell}(x), C_{\ell}(y)) \leq K_{\Phi} \cdot d_{\ell}(x, y) + \varepsilon_{\ell'},$$

with $K_{\Phi} := L_{\Phi} \cdot K_{\mathcal{C}}$ and $\varepsilon_{\ell'} \rightarrow 0$ uniformly over bounded substrate pairs.

Proof. By the closure-generated cone map (§3.2),

$$d_H(C_{\ell}(x), C_{\ell}(y)) = d_H(\Phi(\chi_{\ell}(x)), \Phi(\chi_{\ell}(y))).$$

By Lipschitz response of Φ (Proposition 7.0),

$$d_H(\Phi(\chi_{\ell}(x)), \Phi(\chi_{\ell}(y))) \leq L_{\Phi} \cdot \rho(\chi_{\ell}(x), \chi_{\ell}(y)) = L_{\Phi} \cdot \mathcal{C}_{\ell}(x, y).$$

By Theorem 1,

$$\mathcal{C}_{\ell}(x, y) \leq K_{\mathcal{C}} \cdot d_{\ell}(x, y) + \varepsilon_{\ell}.$$

Combining,

$$d_H(C_{\ell}(x), C_{\ell}(y)) \leq L_{\Phi} \cdot K_{\mathcal{C}} \cdot d_{\ell}(x, y) + L_{\Phi} \cdot \varepsilon_{\ell}.$$

Set $K_{\Phi} := L_{\Phi} \cdot K_{\mathcal{C}}$ and $\varepsilon_{\ell'} := L_{\Phi} \cdot \varepsilon_{\ell}$. Uniformity of $\varepsilon_{\ell'}$ over bounded pairs inherits from uniformity of ε_{ℓ} (Remark 6.2). ■

Remark 7.2 (lineage propagation). The lineage qualification of Theorem 1 propagates to Theorem 2 directly. Theorem 2 is established for refinement-generated pairs (descendants of bounded seed pairs), and extends to arbitrary bounded level- ℓ pairs under the same lineage-covering condition stated in §6.

8. Theorem 3 — Explicit Refinement Dynamics Imply R4' with $\alpha = 1$

Statement. Assume E1, E2, E3, E4a–E4c, E5, (D'), and ℓ -independence of $(\mathcal{K}, \rho, \Phi)$. Then admissible refinement dynamics satisfy bounded refinement distortion R4' with $\alpha = 1$:

$$\Xi_{\ell}(x, y) \leq K_{\Xi} \cdot d_{\ell}(x, y) + \varepsilon_{\ell}'' ,$$

with explicit constant

$$K_{\Xi} = L_{\Phi} \cdot K_{\mathcal{C}} \cdot A = L_{\Phi} \cdot L \cdot A^2 / (A_{-} \cdot (1 - \beta)) .$$

Proof. By definition,

$$\Xi_{\ell}(x, y) = d_{\text{H}}(R_{\{\ell \rightarrow \ell+1\}}(C_{\ell}(x)), R_{\{\ell \rightarrow \ell+1\}}(C_{\ell}(y))) .$$

By E5 (cone-refinement compatibility), there exists $\delta_{\ell} \rightarrow 0$ such that for all $x \in P_{\ell}$,

$$d_{\text{H}}(R_{\{\ell \rightarrow \ell+1\}}(C_{\ell}(x)), C_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x)) \leq \delta_{\ell} .$$

Applying the triangle inequality for d_{H} twice (once for x , once for y),

$$\Xi_{\ell}(x, y) \leq d_{\text{H}}(C_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x), C_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P y)) + 2 \delta_{\ell} .$$

Apply Theorem 2 at level $\ell + 1$ to the pair $(\mathcal{E}_{\ell} \wedge P x, \mathcal{E}_{\ell} \wedge P y)$. The Theorem 2 uniformity is over bounded substrate pairs at level $\ell + 1$; boundedness of $(\mathcal{E}_{\ell} \wedge P x, \mathcal{E}_{\ell} \wedge P y)$ follows from boundedness of (x, y) via the upper half of (D') at $j = 1$, so the uniformity premise is inherited correctly:

$$d_{\text{H}}(C_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x), C_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P y)) \leq K_{\Phi} \cdot d_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x, \mathcal{E}_{\ell} \wedge P y) + \varepsilon_{\{\ell+1\}}' .$$

Apply the upper half of (D') at $j = 1$ — this is the *only* invocation of (D') in the proof of Theorem 3, in contrast to Theorem 1 where the full strength of (D') (all j , both halves) is required:

$$d_{\{\ell+1\}}(\mathcal{E}_{\ell} \wedge P x, \mathcal{E}_{\ell} \wedge P y) \leq A \cdot d_{\ell}(x, y) + \eta_{\ell+1}(x, y) .$$

Combining,

$$\Xi_{\ell}(x, y) \leq K_{\Phi} \cdot A \cdot d_{\ell}(x, y) + K_{\Phi} \cdot \eta_{\ell+1}(x, y) + \varepsilon_{\{\ell+1\}}' + 2 \delta_{\ell} .$$

Set

$$K_{\Xi} := K_{\Phi} \cdot A = L_{\Phi} \cdot K_{\mathcal{C}} \cdot A = L_{\Phi} \cdot L \cdot A^2 / (A_{-} \cdot (1 - \beta)) , \varepsilon_{\ell}'' := K_{\Phi} \cdot \eta_{\ell+1}(x, y) + \varepsilon_{\{\ell+1\}}' + 2 \delta_{\ell} .$$

Each error contribution vanishes uniformly over bounded substrate pairs by a distinct uniformity premise: $K_{\Phi} \cdot \eta_{\ell+1}$ by the uniformity clause of (D'); $\varepsilon_{\{\ell+1\}}'$ by Theorem 2 (which propagates

back through Theorem 1's three uniformity premises of Remark 6.2); and 2 δ_ℓ by E5's uniformity clause. Therefore $\varepsilon_\ell'' \rightarrow 0$ uniformly, and

$$\Xi_\ell(x, y) \leq K_\Xi \cdot d_\ell(x, y) + \varepsilon_\ell''.$$

This is precisely R4' in the form $\alpha = 1$, with the explicit constant K_Ξ expressed in terms of five primitive constants $(\beta, L, A, A_-, L_\Phi)$. ■

Remark 8.1 (architectural roles of (D')). The full strength of (D') — bound for all $j \geq 0$, both upper and lower halves — is consumed by Theorem 1, where the iteration ranges over j and the lower half is needed at the level-restatement step. Theorem 3 only invokes the $j = 1$ trace of the upper half. So (D') is fundamentally a Theorem-1 hypothesis; Theorem 3 inherits its conclusion through K_C (which carries the A and A_- dependence) and additionally requires only the cheapest one-step refinement-distance bound.

Remark 8.2 (lineage propagation). As with Theorem 2, the lineage qualification of Theorem 1 propagates to Theorem 3: the R4' bound is established for refinement-generated pairs at level ℓ , and extends to arbitrary bounded level- ℓ pairs under the lineage-covering condition of §6.

9. Corollary — Lipschitz Continuum Cone Field

The previous paper established

$$R4' + H5 + R2 + \text{bounded refinement overlap} \Rightarrow H8.$$

Theorem 3 supplies R4' with $\alpha = 1$. Hypotheses H5, R2, and bounded refinement overlap are already established for admissible TPB substrates. Therefore H8 holds with Lipschitz regularity, and the continuum cone field satisfies

$$d_H(C(x), C(y)) \leq K_\infty \cdot d_\infty(x, y),$$

with K_∞ inheriting the explicit constant structure $K_\infty = L_\Phi \cdot L \cdot A^2 / (A_- \cdot (1 - \beta))$ (modulo the limit operation that converts the L -Lipschitz substrate bound to its continuum counterpart).

The continuum cone field is therefore Lipschitz. This implies

- stronger continuum regularity,
- improved metric control,
- pointwise stability of transport-current smoothing,
- removal of the weak-norm robustness gap,
- simplification of the Lorentzian-emergence chain.

10. Main Theorem

Theorem (Substrate-to-Geometry with Lipschitz cone stability). Let a TPB substrate satisfy

1. fixed finite closure catalogue (\mathcal{K}, ρ) with $K = 7$, ℓ -independent;
2. closure-generated cone map $\Phi : \mathcal{K} \rightarrow \mathcal{C}$ (ℓ -independent; need not be injective);
3. admissible refinement evolution conditions E1, E2, E3, E4a, E4b, E4c, and E5;
4. two-sided refinement-distance compatibility (D');
5. finite propagation;
6. BCB;
7. bounded refinement overlap;
8. controlled refinement multiplicity.

(Lipschitz response of Φ with constant $L_\Phi < \infty$ is automatic from items 1–2 by Proposition 7.0, and is not a separate hypothesis.)

Then R4' holds with $\alpha = 1$ for refinement-generated bounded pairs at level ℓ , and for arbitrary bounded level- ℓ pairs under the lineage-covering condition of §6. Consequently, the refinement sequence satisfies H8 with Lipschitz regularity, and the full regularity set is

H6, H6', H7, H8_Lip.

Conditional on H9, the continuum-limit paper yields a strongly causal Lorentzian continuum with Lipschitz cone field and C^k Lorentzian metric realisations for every prescribed $k \geq 2$, *without* requiring the weak-norm bridge correction as a load-bearing step.

The explicit Lipschitz constant for the continuum cone field is

$$K_\infty = L_\Phi \cdot L \cdot A^2 / (A_- \cdot (1 - \beta)),$$

determined by five substrate primitives: the closure-coherence contraction rate β (E4a, uniform by E4c), the position-coherence rate L (E4b), the upper and lower refinement-distance scalings A and A_- ((D')), and the closure-to-cone response constant L_Φ (Proposition 7.0).

Positive corollary (Sobolev-free emergence). Beyond the negative statement above ("without weak-norm bridge correction"), the Lipschitz regime supports three positive structural simplifications of the continuum-limit construction:

(i) *Pointwise smoothing of transport currents without Sobolev detours.* Transport currents may be smoothed by pointwise mollification of the Lipschitz cone field, with the mollification error controlled in d_H rather than in L^2 - or H^s -norms. The Sobolev intermediaries used in the original continuum-limit paper to bridge L^2 substrate convergence to pointwise cone convergence become unnecessary.

(ii) *Direct C^k -realisability of the Lorentzian metric without H^s intermediaries.* The C^k realisation theorem in the structural-necessity paper may be applied directly to the Lipschitz cone

field via standard partition-of-unity smoothing, without first lifting the substrate-level bound to a Sobolev space and then taking traces.

(iii) *Pointwise stability under refinement.* Two refinement sequences seeded from substrate pairs (x_n, y_n) with $d_\ell(x_n, y_n) \rightarrow d_\infty(x, y)$ yield cone fields whose pairwise Hausdorff distance is itself Lipschitz-controlled — no weak-norm averaging is required to convert refinement stability into limit stability.

In each case, the simplification is structural rather than merely technical: the Lipschitz bound eliminates an entire layer of analytic machinery from the continuum-limit pipeline.

11. Stability Flow Interpretation

The theorem chain composes into a *stability flow* on substrate refinement:

(E4a: pointwise contractivity) + (E4b: position-coherence) + (E4c: uniform gap) + (E5: cone-refinement compatibility) \rightarrow (Lemma 4.1: closure-coherence inequality, $\mathcal{C}_{\ell+1}(\mathcal{E}_x, \mathcal{E}_y) \leq \beta \cdot \mathcal{C}_\ell(x, y) + L \cdot d_\ell(x, y)$) \rightarrow (Thm 1: Lipschitz closure stability, $\mathcal{C}_\ell \leq K_C \cdot d_\ell + \varepsilon_\ell$) \rightarrow (Thm 2: Lipschitz cone stability, $d_H(C_\ell(x), C_\ell(y)) \leq K_\Phi \cdot d_\ell + \varepsilon_\ell'$) \rightarrow (Thm 3: Lipschitz refinement distortion, $\Xi_\ell \leq K_\Xi \cdot d_\ell + \varepsilon_\ell''$) \rightarrow (R4' with $\alpha = 1$) \rightarrow (H8_Lip) \rightarrow (Lipschitz continuum cone field) \rightarrow (C^k Lorentzian geometry, $k \geq 2$).

Each arrow is a proven implication conditional on the named substrate inputs. The flow can be read in two directions.

Forward — start from the four dynamical hypotheses (E4a, E4b, E4c, E5) and propagate stability up to continuum geometry: contractivity + position-coherence + uniform gap give Lipschitz closure stability; cone-refinement compatibility lifts this to Lipschitz refinement distortion; the previous paper's substrate-engineering chain then delivers Lipschitz continuum cone stability and Lorentzian geometry.

Backward — start from the regularity needed by the continuum-limit paper and trace it down to substrate primitives: smooth Lorentzian geometry requires a Lipschitz cone field, which requires $\alpha = 1$, which requires Lipschitz refinement distortion, which requires Lipschitz cone stability, which requires Lipschitz closure stability, which requires E4a (pointwise contractivity) for the geometric decay, E4b (position-coherence) for the spatial coupling, E4c (uniform gap) for the level-uniformity, and E5 (cone-refinement compatibility) for the lift from closure to cone.

In either direction, the flow identifies the *quadruple* (E4a, E4b, E4c, E5) as the dynamical hypotheses carrying the geometry-regularity content of the substrate. E4a–E4c constrain the position-indexed family $\{ \mathcal{E}_{\ell, x}^K \}$ as it acts within (\mathcal{K}, ρ) ; E5 constrains how this family interfaces with the abstract refinement functor R via Φ . All other ingredients are either kinematic (E1, E2, E3), topological (Lipschitz response of Φ from finite catalogue), or geometric ((D')).

12. Why This Matters

This paper upgrades the geometry programme in a structural way.

The previous paper showed

$R1-R4' \Rightarrow H6, H6', H7, H8.$

The present paper shows

explicit closure refinement dynamics $\Rightarrow R4'$ with $\alpha = 1.$

The most delicate regularity condition — bounded refinement distortion — is therefore no longer merely an engineering assumption. It becomes a *consequence of explicit refinement flow* under a small, well-localised set of dynamical hypotheses (E4a, E4b, E4c, E5) that have clear physical motivation in the $K = 7$ closure algebra and the substrate's local-coupling architecture.

The composite result is

admissible refinement dynamics \Rightarrow Lipschitz continuum cone stability \Rightarrow clean Lorentzian emergence.

This is a substantial step in turning VERSF from a *kinematic* substrate framework (cones posited, constraints imposed) into a *dynamical* substrate-to-geometry reconstruction framework (cones derived from closure states, constraints derived from refinement flow).

The remaining external dependencies are now narrower and more concrete: derive E4a–E4c and E5 from the $K = 7$ closure algebra; give an explicit construction of the position-indexed family $\{ \mathcal{E}_{\ell, x}^K \}$; derive the non-collapse half of (D'); verify the topology condition H9 for explicit substrate models; promote Lipschitz to C^k regularity.

13. Limitations and Open Problems

Several issues remain.

Derivation of E4a–E4c from $K = 7$ closure algebra. The closure-coherence sub-conditions are the load-bearing dynamical hypotheses of this paper. A deeper paper should derive them directly from the structure of the $K = 7$ closure catalogue: E4a (pointwise contractivity) from the admissibility-rejection rate of the closure-transition graph; E4b (position-coherence) from the local-coupling architecture of substrate refinement; E4c (uniform gap) from a structural spectral-

gap condition on the closure-transition operator across refinement levels. The contraction constants β , L should ideally be computed from catalogue topology.

Explicit form of the position-indexed family $\{ \mathcal{E}_{\{\ell, x\}^K} \}$. The position-indexed formulation introduced in §3.3 resolves the equivariance/ L -term inconsistency at the structural level but does not yet specify the construction of $\mathcal{E}_{\{\ell, x\}^K}$ from substrate primitives. An explicit derivation — for example as a measurable selection from a relational admissibility predicate, or as a deterministic construction tied to a local neighbourhood of x in P_ℓ — would close this gap and replace the structural commitment of E4b with a derived bound.

Derivation of the non-collapse half of (D') from R3 + a separation axiom. The upper bound of (D') follows from R3 plus bounded overlap as in the previous paper. The lower (non-collapse) bound is physically natural — refinement should not merge distinct parents to coincident descendants — but is currently *conditional*. A derivation from R3 plus an explicit separation axiom on the refinement functor would promote (D') to fully proven status.

E5 as axiom or derivation. The cone-refinement compatibility E5 is treated as an axiom in this paper. Under the closure-generated cone formulation $C = \Phi \circ \chi$, E5 reduces to a compatibility between $\{ \mathcal{E}_{\{\ell, x\}^K} \}$ and the action of R on $\Phi(\mathcal{K})$. Two derivation paths are open: (i) construct \mathcal{E} from $R + \Phi$ such that E5 is identically satisfied; (ii) construct R from $\{ \mathcal{E}_{\{\ell, x\}^K} \} + \Phi$ and show $\delta_\ell \rightarrow 0$ from substrate-mesh contraction. Either path would replace the axiom with a structural derivation.

The two paths have asymmetric philosophical consequences. Path (i) makes E5 trivial but pushes the dynamical content of \mathcal{E} onto R , which must then be specified independently — and which carries less direct substrate meaning. Path (ii) makes \mathcal{E} primary and constructs R from the substrate dynamics, which is the more VERSF-natural direction (substrate-first, geometry-derived); it keeps the substrate as the primitive and the abstract refinement functor as a derived bookkeeping device. Path (ii) is therefore the architecturally consistent direction for the programme as a whole, and is the preferred target for future work.

Uniformity conditions on β , L_Φ , BCB error. Three uniformity statements are made in §4 (E4c for β , Remark 7.1 for L_Φ , E3 for the BCB error). Each requires a structural input: a *spectral-gap-style* closure-graph condition for β (E4c forbids $\beta_\ell \rightarrow 1$ as the mesh shrinks); ℓ -independence of $(\mathcal{K}, \rho, \Phi)$ for L_Φ ; uniform BCB error rate in bounded sets for E3. The first two are *conditional* on the $K = 7$ substrate carrying the appropriate uniformity; the third should be checked explicitly for the substrate-engineering construction of the previous paper.

Quantitative computation of L_Φ . The closure-to-cone response constant L_Φ is bounded above by $d_{H, \max} / \rho_{\min}$ (Proposition 7.0), but its precise value depends on the explicit cone-direction sets $\Phi(\kappa)$ for each $\kappa \in \mathcal{K}$. These should be computed from the $K = 7$ catalogue topology.

Numerical determination of A and A_- . Both refinement-distance scaling constants — the upper A and the non-collapse A_- — control the propagation of distance bounds through

refinement. They should be measured in explicit TPB substrate simulations. The ratio A / A_-^- enters $K_{\mathcal{C}}$ directly.

The topology condition H9. H9 remains external to the present derivation chain. Its derivation from substrate primitives is a separate open programme.

C^∞ regularity. The present paper achieves Lipschitz regularity for the continuum cone field. Higher regularity (C^k for $k \geq 2$ in the metric, C^∞ as a desideratum) requires either a stronger contraction structure or an external smoothing procedure. The interaction with H8'-type higher-order bounds should be analysed.

Numerical refinement simulations. Direct numerical computation of $\Xi_\ell(x, y)$, $\mathcal{C}_\ell(x, y)$, the constants $(\beta, L, A, A_-^-, L_\Phi)$, and the effective α exponent should be developed for explicit TPB substrate models. Such simulations would also test the sharpness of the Lipschitz bound versus possible super-Lipschitz behaviour.

Generation-dependence of E4. It is conceivable that the contraction constant β has a generation-resolved structure tied to the $K = 7$ wheel architecture used elsewhere in VERSF. A $K = 7$ -resolved version of E4a–E4c would tighten the link between the geometry programme and the substrate-dynamics programme.

Sharpness of Lemma 4.1. The triangle-inequality combination of E4a + E4b gives the closure-coherence inequality with coefficients (β, L) . Whether this is sharp — that is, whether some closure-graph topologies admit a tighter joint analysis yielding effective coefficients strictly smaller than (β, L) — is open. A non-trivial joint analysis could lower K_{Ξ} and hence the continuum Lipschitz constant K_∞ .

Structural condition on A / A_-^- . The continuum Lipschitz constant $K_\infty \propto A^2 / A_-^-$ depends on the *ratio* A / A_-^- more than on either constant individually. A structural condition forcing $A / A_-^- \leq \text{const}$ — for example, "refinement preserves substrate distance ratios up to a bounded factor" — would be a cleaner geometric hypothesis than separate bounds on A and A_-^- , and would tie the continuum geometry more directly to the metric structure of the refinement sequence.

Architectural narrowing from cones-as-derived. The closure-generated cone formulation (§3.2) restricts the substrate cone field to the image of $\Phi \circ \chi$, narrowing the substrate degrees of freedom compared to earlier programme papers in which cone fields were primary. A consistency audit of earlier results — confirming that quantification over arbitrary admissible cone fields specialises correctly to the image of Φ — is straightforward but should be carried out explicitly.

Lineage coverage. The proof of Theorem 1 is naturally formulated along refinement descendants of bounded seed pairs. To promote the result from refinement-generated pairs to *arbitrary* bounded pairs at a fixed level ℓ , a **lineage-covering axiom** is required: every bounded level- ℓ pair must admit bounded ancestors at some earlier seed level, up to vanishing refinement error. This condition is plausible for genuine refinement systems with surjective refinement maps

but is external to the present argument. It should be stated explicitly when Theorem 1 is invoked outside the descendant-pair setting, and eventually derived from a structural condition on the refinement functor (for example, surjectivity-up-to- η of \mathcal{E}_ℓ^P at every level).

Admissible-pair scope of E4a. E4a is *strongest* if read as a global contraction over all of $\mathcal{K} \times \mathcal{K}$. The physically intended version is contraction over **admissibly comparable closure pairs** within a common local refinement context, as discussed under "Scope of E4a" in §4. A future K = 7 catalogue paper should (i) define this admissible comparison relation explicitly — for example as the edge relation of a closure-transition graph or as the support of a local admissibility predicate; (ii) prove the contraction bound on the admissible sector; and (iii) verify that the proof of Lemma 4.1 carries through when ρ is restricted to admissibly comparable pairs (the triangle path used there stays within admissible comparisons, since the intermediate point $\mathcal{E}_{\{\ell,x\}^K}(\chi_\ell(y))$ is itself an admissible local update at x ; the symmetric path through $\mathcal{E}_{\{\ell,y\}^K}(\chi_\ell(x))$ likewise stays within admissibility, so Lemma 4.1 survives the restriction independently of path choice).

A subsidiary structural point: the metric ρ , extended from the physically meaningful admissibly comparable pairs to all of $\mathcal{K} \times \mathcal{K}$, should satisfy $\rho \geq \rho_{\min}$ on every distinct pair so that Proposition 7.0's $L_\Phi \leq d_{H,\max} / \rho_{\min}$ bound remains uniform. The admissible-pair scope of E4a is then a contraction condition on a sub-metric of (\mathcal{K}, ρ) , with the global ρ providing the uniform lower bound that Proposition 7.0 requires. The present paper handles this correctly by treating ρ as a global metric and letting E4a be a dynamics-side property; a future K = 7 catalogue paper should make the admissible-sector / global-extension distinction explicit at the level of ρ itself.

14. Conclusion

This paper introduces explicit refinement dynamics into the VERSF geometry programme. The main new objects are

- the refinement evolution operator \mathcal{E}_ℓ with position-indexed closure-update family $\{\mathcal{E}_{\{\ell,x\}^K}\}$;
- the closure assignment χ_ℓ ;
- the (ℓ -independent) closure-to-cone map $\Phi : \mathcal{K} \rightarrow \mathcal{C}$;
- the refinement-coherence functional \mathcal{C}_ℓ ;
- the closure-coherence sub-conditions E4a (pointwise contractivity), E4b (position-coherence), E4c (uniform admissibility gap);
- the cone-refinement compatibility axiom E5;
- the two-sided refinement-distance compatibility (D').

The principal result is that if closure mismatch contracts under admissible refinement evolution, then refinement distortion is Lipschitz:

$$\Xi_\ell(x, y) \leq K_\Xi \cdot d_\ell(x, y) + \varepsilon_\ell, \quad K_\Xi = L_\Phi \cdot L \cdot A^2 / (A_- \cdot (1 - \beta)).$$

Therefore R4' holds with $\alpha = 1$. This achieves the highest-value target identified in the previous paper. Lipschitz refinement distortion yields Lipschitz continuum cone stability and removes the weak-norm robustness gap in the continuum-to-Lorentz bridge — and positively, supports pointwise smoothing of transport currents without Sobolev detours, direct C^k -realisability of the metric without H^s intermediaries, and pointwise stability of refinement sequences.

The geometry-emergence chain therefore becomes

closure refinement dynamics + cone-refinement compatibility $\Rightarrow \alpha = 1 \Rightarrow H8_Lip \Rightarrow$ controlled (i.e., Lipschitz-stable) Lorentzian continuum limit $\Rightarrow C^k$ Lorentzian metric realisations, conditional on H9 and the continuum-limit hypotheses.

The remaining work is now specific and concrete: derive E4a–E4c and E5 from the finite $K = 7$ closure algebra, give an explicit construction of the position-indexed family $\{ \mathcal{E}_{\ell,x}^K \}$, derive the non-collapse half of (D') from R3 plus a separation axiom, compute the constants $(\beta, L, A, A_-, L_\Phi)$ for explicit substrate models, verify H9 for those models, and build numerical refinement simulations to test the predicted Lipschitz regime.