Substrate Emergence of the Lorentz Transport Group in VERSF

K = 7 Closure Symmetry, Bit-Conservation Bilinear Preservation, Hyperbolic Signature Selection, and the Continuum-Limit Emergence of SO(1, D) from the Substrate

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The recent geometric papers of the VERSF programme established gravity as the continuum-limit shadow of irreversible commitment transport on the K = 7 simplicial commitment foam. The Lorentzian Completion paper produced an emergent Lorentzian continuum geometry. The transport-geometry papers produced refinement-stable parallel transport on that continuum. The Einstein–Hilbert Emergence paper showed that the continuum connection, the Riemann curvature tensor, the invariant volume form, and the Einstein–Hilbert action itself all emerge as continuum-limit consequences of refinement-stable transport. The recent paper *Microscopic Origin of Refinement Transport in VERSF* then constructed the refinement transport operator T_γ itself explicitly from K = 7 commitment-foam combinatorics, closing the existence-and-stability problem for refinement transport.

However, that paper inherited one structural input that it could not derive from the substrate: the transport group itself. The substrate-level edge transport operators U_ij ∈ 𝒢 were defined with 𝒢 = SO(1, D), the local Lorentz group of the inherited continuum-limit signature. This identification was acknowledged in §2 (Inherited Structures), the Theorem 8.2 Remark, and recorded as OP9 — the residual problem of constructing an admissible substrate-level transport group from K = 7 / BCB combinatorics independent of the inherited continuum-limit signature, with the Lorentz structure emerging in the continuum limit rather than being imposed.

That is the gap this paper closes.

The construction proceeds in three structural steps. First, we define the **abstract substrate transport group** 𝒢_sub purely combinatorially — as the maximal group of admissible edge transport operators consistent with the seven cell-level closure constraints C_1–C_7 of the K = 7 architecture, with no inner-product structure on tangent fibres assumed. Second, we show that BCB (Bit Conservation and Balance) at the substrate level forces 𝒢_sub to preserve a substrate-level quadratic invariant whose continuum limit, by refinement-Cauchy convergence, is a unique-up-to-scale invariant bilinear form η_μν. Third, we determine its signature: the inherited CCC framework supplies finite-speed coherence propagation forcing indefinite signature; the substrate ontology supplies a single irreversible-commitment direction per edge forcing at most one negative eigenvalue; finite distinguishability of orthogonal directions forces exactly D

positive eigenvalues. The combined constraints give signature (−, +, …, +) — i.e. Lorentzian (1, D) in the mostly-plus convention of the dynamical paper.

The remaining work is to show that SO(1, D)-compatible transport is a *refinement-stable* fixed point. We establish that anisotropic transport deviations — substrate-level transport operators violating SO(1, D) compatibility — are exponentially suppressed under closure-compatible refinement, with suppression rate $\eta > 0$ inherited from the same K = 7 / TPB refinement-flow contraction structure that gives geometric refinement scaling $\varepsilon_n \leq C \lambda^n$ in the predecessor paper.

Assembling the three: the abstract substrate transport group $\mathcal{G}_{sub}$, in the continuum limit, preserves a Lorentzian bilinear form (Theorems 4.1 and 5.1), and SO(1, D) is the unique refinement-stable fixed point (Theorem 6.2). Therefore $\mathcal{G}_\infty \cong SO(1, D)$ (Theorem 7.1), with continuum refinement transport operators $T_\gamma \in SO(1, D)$. The first-order generator of T_γ is then the Levi-Civita connection of the inherited Lorentzian continuum geometry (Theorem 8.1) — now with the transport-group identification substrate-derived rather than inherited.

The paper closes OP9 of the predecessor at the *existence and identification* level: $\mathcal{G}_\infty \cong SO(1, D)$ is now derived from the K = 7 closure architecture, the BCB bit-conservation principle, the CCC framework, the substrate-level irreversible-commitment ontology, and the refinement-stable closure structure. What the paper does *not* derive: the quantitative one-loop suppression rate η as an explicit function of the K = 7 bare-coupling integer $N_{loop} = 14$ (the Wilson-paper-style quantitative matching analysis, recorded as OP7); the overall scale of $\eta_{\mu\nu}$ (a normalisation problem inherited from the bilinear universality class, recorded as OP8); and a fully non-abelian Wilson-style derivation paralleling the abelian U(1) gauge-bundle template of the K = 7 Wilson Limit paper (recorded as OP9). Several further problems inherited from the predecessor paper — C_λ, matter, quantum completion, foam dynamics — are recorded unchanged.

The result completes the transport side of the VERSF geometric emergence programme at the symmetry-group level. The Lorentz group is now no longer inherited as an external geometric symmetry; it emerges as the stable transport symmetry of closure-compatible irreversible commitment transport on the K = 7 simplicial substrate.

Table of Contents

Abstract

The recent paper *Microscopic Origin of Refinement Transport in VERSF* constructed the refinement transport operator T_γ from K = 7 simplicial commitment-foam combinatorics and closed the existence-and-stability problem for refinement transport. However, the substrate-level edge transport group was inherited as $\mathcal{G}$ = SO(1, D), the local Lorentz group of the inherited continuum-limit signature — recorded as OP9 of that paper.

The present paper closes OP9.

We construct the abstract substrate transport group $\mathcal{G}_{sub}$ from the K = 7 cell-level closure constraints C_1–C_7 with no inner-product structure on tangent fibres assumed. Bit Conservation and Balance (BCB) at the substrate level — combined with finite-distinguishability conservation under transport — forces $\mathcal{G}_{sub}$ to preserve a substrate-level quadratic invariant whose continuum limit is a refinement-stable invariant bilinear form $\eta_{\mu\nu}$ unique up to overall scale. The signature of $\eta_{\mu\nu}$ is then forced to Lorentzian (1, D) by three independent structural obstructions: CCC finite-speed propagation combined with irreversibility-as-distinguishability-consumption (indefinite signature); the multi-timelike refinement instability of ultra-hyperbolic signatures (at most one negative eigenvalue); and finite distinguishability of orthogonal directions (exactly D positive eigenvalues). Stability follows from explicit spectral analysis of the anisotropy refinement operator $\mathscr{R}_{aniso}$, with $\rho(\mathscr{R}_{aniso}) < 1$ forced by foundational C_5–C_7 refinement-cell rigidity, and SO(1, D) emerging as the $\rho = 1$ eigenspace $\mathcal{A}^{Lor}$ of the spectral decomposition $\mathcal{A} = \mathcal{A}^{Lor} \oplus \mathcal{A}^{aniso}$. Continuum metric geometry is then *induced from* this substrate transport geometry via an explicit refinement-frame soldering $g_{\mu\nu} = e^a{}_\mu e^b{}_\nu \eta_{ab}$, not inherited from a continuum-level Lorentzian structure.

The paper establishes eight principal results.

(i) Bit-Conservation Bilinear Preservation (§4, Theorem 4.1). Closure-compatible refinement transport on the K = 7 commitment foam preserves, in the continuum limit, an invariant bilinear form $\eta_{\mu\nu}$ unique up to overall scale (Lemma 4.3). BCB bit-conservation lifted to the bilinear sector via pairwise distinguishability conservation supplies the substrate-level mechanism.

(ii) Multi-Timelike Refinement Instability (§5, Lemma 5.1). Ultra-hyperbolic signatures (p, q) with $p \geq 2$ generate non-equivalent refinement orderings on the continuum manifold, violating the closure-compatible refinement framework. Lorentzian signature is therefore a *structural necessity* of the refinement architecture, not merely an ontological inheritance.

(iii) Lorentzian Signature Selection (§5, Theorem 5.2). Combined inputs — CCC + irreversibility-as-distinguishability-consumption (an articulation of the source-paper §2.1 irreversibility ontology, recorded as OP12) (Step 1), multi-timelike refinement instability with point-level coherence (Step 2), finite distinguishability of orthogonal directions (Step 3) — force $\eta_{\mu\nu}$ to have Lorentzian signature (1, D).

(iv) Spectral Suppression of Transport Anisotropy (§6, Theorem 6.2). The anisotropy refinement operator $\mathscr{R}_{aniso}$ has spectral radius $\rho(\mathscr{R}_{aniso}) = \lambda_{aniso} < 1$, with the strict inequality forced by foundational C_5–C_7 refinement-cell rigidity. The continuum-limit anisotropy mode space admits the spectral decomposition $\mathcal{A} = \mathcal{A}^{Lor} \oplus \mathcal{A}^{aniso}$, with Lorentz-compatible modes refinement-invariant and anisotropic modes exponentially suppressed at rate $\eta = -\log \lambda_{aniso} > 0$.

(v) Transport-Group Emergence and Uniqueness (§7, Theorems 7.1 and 7.2). The continuum-limit admissible transport group is

$$\mathcal{G}_\infty \cong SO(1, D),$$

with continuum refinement transport operators $T_\gamma \in SO(1, D)$. Moreover, $SO(1, D)^\uparrow$ is the *unique* admissible refinement-stable transport group: every alternative candidate is excluded by one of eight structurally distinct substrate-level obstructions (O1)–(O8).

(vi) Substrate-Induced Continuum Metric (§8, Definition 8.0). Continuum metric geometry is induced from substrate transport geometry via the refinement-frame soldering $g_{\mu\nu} = e^a_\mu e^b_\nu \eta_{ab}$, with the refinement frame bundle $F \to \mathscr{M}$ a principal SO(1, D)-bundle over the emergent continuum and $\nabla_\mu e^a_\nu = 0$ the vielbein compatibility condition under refinement-compatible transport.

(vii) Substrate-Derived Levi-Civita Emergence (§8, Theorem 8.1). Both metric compatibility $\nabla_\mu g_{\nu\rho} = 0$ *(via vielbein compatibility and η-preservation under SO(1, D)) and torsion-freeness* $\Gamma^\alpha_{[\beta\mu]} = 0$ (via BCB triangle closure C_4) are now substrate-derived. The first-order generator of T_γ is the Levi-Civita connection of the soldered continuum geometry, with no Lorentz-structure inheritance invoked.

(viii) Microscopic Lorentzian Recovery (§9, Theorem 9.1). The Lorentzian geometric structure of the continuum-limit VERSF gravitational sector emerges as the stable refinement fixed point of closure-compatible irreversible commitment transport on the K = 7 simplicial substrate.

The paper therefore closes OP9 of the predecessor at the existence-uniqueness-identification level and unifies the transport-group structure of the continuum-emergence cascade with the K = 7 / BCB substrate architecture.

The paper does **not** derive:

- the closure-normalisation factor C_λ,

- the quantitative one-loop suppression rate η as an explicit function of K = 7 bare-coupling integers,
- the overall scale of $\eta_{\mu\nu}$,
- a fully non-abelian Wilson-style derivation paralleling the K = 7 Wilson Limit paper's abelian U(1) template,
- Standard-Model matter coupling,
- quantum fluctuations of T_γ or of the commitment foam itself,
- subleading higher-curvature corrections.

Each is recorded as an explicit open problem.

The contribution is structural: it supplies the substrate-level derivation of the Lorentz transport group within the VERSF continuum-emergence cascade, closing the residual transport-group inheritance gap below the geometric programme.

---

1. Introduction

The VERSF gravity programme has now established, in successive layers:

- emergent Lorentzian continuum geometry (Lorentzian Completion paper),
- refinement-stable parallel transport on the emergent continuum (transport-geometry papers),
- effective stress–energy structure (source-structure paper),
- Einstein-compatible field equations (dynamical-geometry paper),
- a unified covariant action $\mathcal{S}_{VERSF}$ (variational paper),
- substrate emergence of the Einstein–Hilbert action (Einstein–Hilbert Emergence paper),
- substrate construction of the refinement transport operator T_γ (Microscopic Origin paper).

The Microscopic Origin paper carried the inversion of the geometric programme through the deepest layer reached so far: T_γ itself is now constructed from K = 7 simplicial commitment-foam holonomy under closure-compatible refinement, rather than inherited from the transport-geometry programme as a continuum-level object. Theorems 5.3 and 6.2 of that paper closed OP2 and OP8 of the Einstein–Hilbert Emergence paper (the existence and stability questions for T_γ).

The residual gap that paper explicitly recorded was the substrate origin of the **transport group itself**. Definition 3.1 of the Microscopic Origin paper identified the substrate-level edge transport group $\mathcal{G}$ with SO(1, D), the local Lorentz group of the inherited continuum-limit signature. Although the §2 (Inherited Structures) ledger and the Theorem 8.2 Remark were honest about this inheritance, the identification was not derived from substrate combinatorics. This was recorded as **OP9 — Substrate derivation of the transport group $\mathcal{G}$**:

Definition 3.1 identifies the substrate-level edge transport group with $\mathcal{G} = SO(1, D)$... Closing OP9 would mean constructing an admissible transport group $\mathcal{G}$_sub from substrate combinatorics independent of the inherited signature, with $\mathcal{G}$_sub → SO(1, D) emerging in the continuum limit by the same coherence-propagation mechanism that the CCC framework supplies for the metric.

The present paper closes OP9.

The construction. The argument proceeds in three structural steps.

Step 1 — abstract substrate transport group. We define $\mathcal{G}$_sub combinatorially as the maximal group of admissible edge transport operators consistent with all C_1–C_7 closure constraints at every cell of the K = 7 foam, with no inner-product structure on tangent fibres assumed (§3). This is genuinely substrate-defined: C_1–C_7 are cell-level combinatorial constraints inherited from the foundational papers, with no reference to Lorentz structure or continuum-limit signature.

Step 2 — bilinear preservation. BCB (Bit Conservation and Balance) at the substrate level supplies a substrate-level quadratic invariant — the bit-count form — that $\mathcal{G}$_sub preserves at every refinement order modulo plaquette-bounded corrections. By refinement-Cauchy convergence (Theorem 5.3 of the Microscopic Origin paper), the continuum limit of this substrate quadratic invariant is a well-defined bilinear form η_μν, unique up to overall scale (§4).

Step 3 — signature. The signature of η_μν is determined by combining several substrate-level inputs: (a) CCC finite-speed coherence propagation together with irreversibility-as-distinguishability-consumption (forcing indefinite signature, since commitment-propagation directions consume bit-resource while spatial directions preserve it, with opposite-sign contributions to the BCB bilinear pairing); (b) multi-timelike refinement instability — ultra-hyperbolic signatures (p, q) with $p \geq 2$ generate non-equivalent refinement orderings on the continuum manifold, violating closure-compatible refinement — together with the coherent point-level commitment orientation supplied by closure-compatible coarse-graining of BCB-balanced edge orientations (jointly forcing at most one negative eigenvalue); (c) finite distinguishability of orthogonal directions (forcing exactly D positive eigenvalues). The combined constraints give signature (1, D) — i.e. (−, +, …, +) in mostly-plus convention (§5).

The remaining sections establish that SO(1, D)-compatible transport is the unique refinement-stable fixed point (§6), assemble the main theorem (§7), and re-derive the Levi-Civita identification now with substrate-derived transport-group structure (§§8–9).

Scope and limits. The paper closes OP9 of the Microscopic Origin paper at the *existence-uniqueness-identification* level: $\mathcal{G}_\infty \cong SO(1, D)$ is now derived from substrate combinatorics + BCB + CCC + the substrate-level irreversible-commitment ontology. What it does *not* close: the quantitative one-loop suppression rate η in Theorem 6.2 is inherited from the foundational refinement-flow structure rather than computed from K = 7 bare-coupling integers (OP7); the overall scale of η_μν is inherited from the bilinear universality class (OP8); a fully non-abelian Wilson-style derivation paralleling the K = 7 Wilson Limit paper's abelian U(1) template — with explicit one-loop matching at the substrate scale — is a further successor problem (OP9). Further inherited residual problems — C_λ, matter coupling, quantum completion, foam dynamics, anisotropic Wilson coefficients, higher-curvature corrections — are recorded unchanged.

Methodological context. OP9 of the Microscopic Origin paper noted the methodological precedent of the K = 7 Wilson Limit paper, which supplied a worked U(1) gauge-bundle template for the structurally parallel question — showing that the K = 7 closure structure dynamically constrains the effective gauge transport theory toward Lorentz compatibility via one-loop matching, with the rate fixed by the bare coupling $\beta_{K=7} = 2^7 \cdot 15/14$ inherited from the closure-counting integer $N_{loop} = 14$. The present paper is the **SO(1, D) tangent-bundle sector** analogue of that gauge-bundle template. The two are different transport problems on the same substrate (gauge bundle vs tangent bundle); the structural parallel — closure-counting integer $N_{loop} = 14$, K = 7 cell architecture, refinement-flow contraction — is the same. What differs is the non-abelian, (D + 1)(D + 2)/2-dimensional nature of SO(1, D), which makes a direct transcription of the abelian Wilson one-loop machinery non-trivial (recorded as OP9 below). The present paper achieves the existence-uniqueness-identification result; the quantitative Wilson-style derivation of the suppression rate is left to a successor.

Epistemic discipline. Results are labelled *proven*, *conditional*, or *conjectural*, with conditional results stating the additional assumptions explicitly.

Clarification — what this paper is not. The paper does not re-derive the Lorentzian Completion paper's continuum geometry: the (−, +, +, +) signature of the inherited continuum-limit metric $g_{\mu\nu}$ is a continuum-limit result of the CCC framework, inherited here without re-derivation. What the present paper derives is the substrate-level identification of the abstract transport group $\mathcal{G}_{sub}$ with SO(1, D) in the continuum limit. The continuum-limit signature and the substrate-level transport group are different objects — related by the natural pairing of the bilinear form $\eta_{\mu\nu}$ on tangent fibres with the metric $g_{\mu\nu}$ on the continuum manifold. Theorem 8.2 of the Microscopic Origin paper established compatibility between the two; the present paper supplies the substrate construction of $\eta_{\mu\nu}$ and $\mathcal{G}_{sub}$ independent of $g_{\mu\nu}$, with their compatibility now a derived consequence.

---

2. Inherited Structures

We inherit without re-derivation:

- **K = 7 simplicial closure architecture.** The foundational papers establish that uniform isotropic substrates admitting BCB-compliant bit-closure are uniquely hexagonal, with each cell carrying exactly K = 7 closure vertices and $N_{loop} = 2K = 14$ oriented closure traversals. The constraint decomposition C_1–C_3 (edge consistency) + C_4 (triangle / fold completion) + C_5–C_7 (hexagonal embedding) is the cell-level structural input of the present paper.
- **BCB and TPB.** Bit Conservation and Balance (BCB) supplies the substrate-level bit-count form preserved by closure-compatible transport. Ticks-Per-Bit (TPB) supplies the

local refinement-step structure with geometric refinement scaling $\varepsilon_n \leq C \lambda^n$, $0 < \lambda < 1$, inherited from the foundational papers.

- **Finite distinguishability and irreversible commitment.** Any bounded region of the substrate admits only finitely many admissible distinctions; commitment events are irreversible state changes of the void substrate. Each commitment edge carries a **single irreversibility direction** — the commitment orientation, used in §5 below to constrain the number of negative eigenvalues of $\eta_{\mu\nu}$ to exactly one. Inherited from the source-structure paper §2.1 and the foundational papers.
- **CCC (Causal–Coherence Coordination) framework.** Finite-speed coherence-preserving propagation at the substrate level, with the substrate-level cone of coherence-preserving directions lifting to the Lorentzian causal cone in the continuum limit. The CCC framework forces *indefinite* signature on the continuum-limit bilinear form (§5 below); it does *not*, on its own, fix the number of negative eigenvalues, which requires the further substrate-ontology input of single-direction commitment orientation.
- **Emergent Lorentzian continuum geometry $g_{\mu\nu}$**, with signature (−, +, +, +), from the Lorentzian Completion paper. The present paper derives the substrate-level transport bilinear form $\eta_{\mu\nu}$ independently of $g_{\mu\nu}$; their compatibility ($\eta_{\mu\nu}$ of the tangent fibre $\cong g_{\mu\nu}$ of the continuum manifold) is then a derived consequence rather than an assumed identification.
- **Microscopic refinement transport.** The refinement transport operator T_γ of the Microscopic Origin paper, constructed as the continuum-limit holonomy of K = 7 simplicial commitment-foam edge transport operators U_{ij} under closure-compatible refinement. Theorems 5.3 (existence), 6.2 (Levi-Civita emergence), and 7.2 (Riemann emergence) of that paper are used throughout, with the new content of the present paper being that the substrate-level transport group $\mathcal{G}$ is now derived rather than inherited.
- **Refinement-flow contraction structure.** The K = 7 / TPB cell-subdivision combinatorics produces a refinement-flow operator on the substrate-level closure structure, with spectral-gap contraction rate inherited from the foundational papers. The present paper uses this contraction structure to establish refinement suppression of transport anisotropies (§6).
- **K = 7 Wilson Limit paper.** Methodological precedent for the structurally parallel gauge-bundle problem: the K = 7 closure structure dynamically constrains the effective U(1) gauge transport theory toward Lorentz compatibility, with rate fixed by the bare coupling $\beta_{K=7} = 2^7 \cdot 15/14$ inherited from $N_{loop} = 14$. The present paper inherits the methodological framework but not the specific one-loop matching computation, which remains a quantitative residual problem (OP9 below).

These are treated as the input architecture of the present construction.

Notation. We write S for the simplicial commitment foam, $S^{(n)}$ for the foam at refinement order n, ε_n for the characteristic edge length at refinement order n, U_{ij} for the edge transport operator on edge e: $v_i \to v_j$, T_γ for the continuum-limit refinement transport operator, $\mathcal{G}_{sub}$ for the abstract substrate transport group, $\mathcal{G}_\infty$ for its continuum-limit, $\eta_{\mu\nu}$ for the substrate-derived continuum-limit bilinear form on tangent fibres, $g_{\mu\nu}$ for the inherited continuum-limit metric on the continuum manifold (Lorentzian Completion), $Q_{sub}^{(n)}$ for the substrate-level bit-count quadratic invariant at refinement order n, Δ_n for the cell-averaged transport anisotropy measure (Definition 6.1 below), and $\eta > 0$ for the substrate-level refinement-flow contraction rate. Greek indices run over (0, 1, …, D); the physically central case is $D + 1 = 4$. Signature mostly-plus, inherited from the dynamical paper.

---

3. The Abstract Substrate Transport Group

We define the substrate-level transport group combinatorially, with no inner-product structure on tangent fibres assumed.

Definition 3.1 — Admissible substrate transport structure

An **admissible substrate transport structure** on the $K = 7$ commitment foam S is an assignment, to each oriented edge $e: v_i \to v_j$ of S, of a transport operator U_{ij} in some abstract group, satisfying:

- **(T1) Edge consistency (C_1–C_3).** $U_{ji} = U_{ij}^{-1}$; composition of edge operators along any path is associative and well-defined.
- **(T2) Triangle / fold closure (C_4).** For any triangle T_{ijk} on the foam, the loop holonomy $H(T_{ijk}) = U_{ki} U_{jk} U_{ij}$ satisfies $H(T_{ijk}) = I + \mathcal{O}(\varepsilon^2)$.
- **(T3) Hexagonal embedding closure (C_5–C_7).** For each $K = 7$ cell, the seven closure traversals (six hexagonal + one central) satisfy the BCB cell-level closure constraints inherited from the foundational papers.
- **(T4) Finite distinguishability.** U_{ij} acts on a finite-dimensional space of admissible commitment configurations at v_i, mapping to admissible configurations at v_j.
- **(T5) Closure-compatible refinement.** Under refinement of the foam (Definition 3.2 of the Microscopic Origin paper), refined edge operators are consistent with parent edge operators in the sense of Lemma 4.4 of that paper.
- **(T6) Path-reversal consistency.** $U_{ji} = U_{ij}^{-1}$ extends to inverse-path consistency at every refinement order.
- **(T7) BCB bit-conservation.** The substrate-level bit-count form $Q_{sub}^{(n)}$ (defined below in §4) is preserved by every $U_{ij}^{(n)}$ at every refinement order, modulo plaquette-bounded subleading corrections.

Note on (T7). $Q_{sub}^{(n)}$ is defined in §4 Theorem 4.1 Step 1 below; (T7) is a forward reference. The pair ($\mathcal{G}_{sub}$, Q_{sub}) is a fixed-point characterisation: $\mathcal{G}_{sub}$ is the maximal group of admissible edge transport operators preserving the BCB bit-count form Q_{sub}, and Q_{sub} is the BCB bilinear form preserved by closure-compatible transport. The fixed point is non-trivial — i.e. Q_{sub} exists and $\mathcal{G}_{sub}$ is non-empty — by Theorem 4.1 below. This forward reference is unavoidable: the closure-compatibility constraints of Definition 3.1 and the BCB bilinear form of §4 are jointly characterised, not sequentially defined.

The crucial point of Definition 3.1: **no inner-product structure on tangent fibres is assumed**. The U_{ij} are abstract group elements satisfying combinatorial closure constraints. The bilinear form $\eta_{\mu\nu}$ that $\mathcal{G}_{sub}$ turns out to preserve is *derived* in §4, not postulated here.

Definition 3.2 — Abstract substrate transport group $\mathcal{G}$_sub

The **abstract substrate transport group** $\mathcal{G}$_sub is the maximal group of admissible edge transport operators satisfying (T1)–(T7) of Definition 3.1, considered as a single abstract group acting on the substrate-level tangent fibre at each vertex of the K = 7 commitment foam.

At refinement order n, we write $\mathcal{G}$_sub^(n) for the substrate transport group on S^(n).

Definition 3.3 — Continuum-limit transport group $\mathcal{G}$_∞

The **continuum-limit transport group** is

$\mathcal{G}$_∞ := lim_{n → ∞} $\mathcal{G}$_sub^(n)

in the closure-compatible refinement-Cauchy sense of Theorem 5.3 of the Microscopic Origin paper, acting on the continuum-limit tangent fibre at each point of the emergent continuum.

Remark — substrate-derivation discipline

Definitions 3.1–3.3 are formulated *without* invoking any continuum-limit signature, Lorentz structure, or inner-product structure on tangent fibres. The closure constraints (T1)–(T7) are inherited from the K = 7 foundational papers and the BCB principle; none of them references Lorentz structure. The substrate transport group is therefore genuinely substrate-defined: the identification $\mathcal{G}$_∞ ≅ SO(1, D) is what §§4–7 below derive, not what §3 above postulates.

This is the structural improvement over the Microscopic Origin paper's Definition 3.1, where the substrate-level transport group was explicitly identified with SO(1, D) at definition time. The present paper postpones that identification to Theorem 7.1 below, where it is derived rather than postulated.

---

4. Bit-Conservation Bilinear Preservation

We now establish the existence and uniqueness (up to scale) of an invariant continuum-limit bilinear form on tangent fibres, preserved by the abstract substrate transport group $\mathcal{G}$_sub.

Theorem 4.1 — Bilinear preservation [proven, conditional on BCB bit-conservation and closure-compatible refinement]

Closure-compatible refinement transport on the K = 7 commitment foam preserves, in the continuum limit, an invariant bilinear form

η_μν

on the continuum-limit tangent fibre at each point, unique up to overall scale.

Proof

Three steps.

Step 1 — existence of substrate-level quadratic invariant $Q_{sub}^{(n)}$.

BCB at the substrate is a cell-level *linear* conservation law: total bit-count flowing into a cell equals bit-count flowing out, with the $K = 7$ closure structure (C_5–C_7 hexagonal embedding) supplying the balance constraint. This linear statement is not yet sufficient for bilinear-form preservation; the lift to a quadratic invariant on tangent fibres requires a further substrate-level input.

The required input is *pairwise bit-distinguishability conservation*. Let $\langle u, v\rangle_{sub}^{(n)}$ denote the substrate-level pairing measuring distinguishability between configurations u and v at $v_i^{(n)}$, defined as the bit-count differential between the joint configuration (u, v) and the separated configurations $(u) \sqcup (v)$ under BCB cell-level accounting. By finite distinguishability (inherited from source paper §2.1), closure-compatible transport must preserve pairwise distinguishability: if $U_{ij}^{(n)}$ could collapse two distinguishable configurations into one, the bounded region containing the transported pair would contain strictly fewer distinguishable configurations than the bounded region containing the source pair, contradicting finite-distinguishability inheritance under closure-compatible transport.

Pairwise distinguishability conservation is therefore a *bilinear* conservation law: every $U_{ij} \in \mathcal{G}_{sub}^{(n)}$ preserves $\langle \cdot, \cdot\rangle_{sub}^{(n)}$ up to plaquette-bounded subleading corrections,

$$\langle U_{ij} u, U_{ij} v\rangle_{sub}^{(n)} = \langle u, v\rangle_{sub}^{(n)} + \mathcal{O}(\varepsilon_n^2),$$

with the $\mathcal{O}(\varepsilon_n^2)$ bound from BCB triangle closure C_4 as in Lemma 6.3 of the Microscopic Origin paper. By polarisation, this is equivalent to preservation of the quadratic form $Q_{sub}^{(n)}(v) := \langle v, v\rangle_{sub}^{(n)}$.

The substrate-level bilinear form is therefore the genuine lift of BCB linear bit-conservation to the bilinear sector, obtained by combining cell-level bit-balance (BCB) with finite-distinguishability conservation under transport (inherited). The lift is non-trivial: BCB alone, without finite-distinguishability inheritance, gives only the linear sector; finite distinguishability alone, without BCB, gives no cell-level conserved quantity. The two together produce $Q_{sub}^{(n)}$.

Step 2 — continuum-limit existence.

By the refinement-Cauchy convergence of Theorem 5.3 of the Microscopic Origin paper, applied jointly to the substrate quadratic form rather than to path holonomy alone, the sequence

$$\{Q_{sub}^{(n)}\}$$

is Cauchy in the appropriate operator norm on the substrate-level tangent fibre, with the continuum limit

$$Q := \lim_{n \to \infty} Q_{sub}^{(n)}$$

existing and being uniquely determined by the closure-compatibility data. The bilinear form $\eta_{\mu\nu}$ is then obtained by polarisation:

$$\eta_{\mu\nu} u^\mu v^\nu := \frac{1}{2} [Q(u + v) - Q(u) - Q(v)].$$

Continuum-limit invariance of $\eta_{\mu\nu}$ under $\mathcal{G}_\infty$ follows from substrate-level invariance of $Q_{sub}^{(n)}$ under $\mathcal{G}_{sub}^{(n)}$ at every refinement order, combined with the convergence $\mathcal{G}_{sub}^{(n)} \to \mathcal{G}_\infty$ in the closure-compatible refinement-Cauchy sense.

Step 3 — uniqueness up to scale.

Any second refinement-stable bilinear form $\eta'_{\mu\nu}$ *preserved by* $\mathcal{G}_\infty$ corresponds, in the substrate-level limit, to a second substrate-level quadratic invariant $Q'_{sub}^{(n)}$ preserved by every $U_{ij} \in \mathcal{G}_{sub}^{(n)}$ up to $\mathcal{O}(\varepsilon_n^2)$ corrections. By the bilinear-form-sector rigidity result of Lemma 4.3 below, $Q'_{sub}^{(n)} = c_n Q_{sub}^{(n)} + \mathcal{O}(\varepsilon_n^2)$ for some scalar c_n at each refinement order; taking the continuum limit gives $\eta'_{\mu\nu} = c\, \eta_{\mu\nu}$ for some $c \in \mathbb{R} \setminus \{0\}$.

Lemma 4.2 — Nondegeneracy [proven, conditional on finite distinguishability]

The continuum-limit bilinear form $\eta_{\mu\nu}$ of Theorem 4.1 is nondegenerate:

$$\eta_{\mu\nu} V^\mu = 0 \text{ for all } V \Rightarrow V = 0.$$

Proof

Suppose $\eta_{\mu\nu} V^\mu = 0$ for some non-zero V. Then $Q(V) = \eta_{\mu\nu} V^\mu V^\nu = 0$, but moreover $Q(V + W) = Q(W)$ for all W. The configuration V is therefore *bit-count-equivalent* to the zero configuration in the substrate-level sense: at every refinement order, V and 0 are indistinguishable under the substrate-level bit-count form $Q_{sub}^{(n)}$.

By finite distinguishability (inherited from the source paper §2.1): any bounded region of the substrate admits only finitely many distinguishable admissible configurations, and these are preserved by closure-compatible transport. A non-zero V indistinguishable from 0 under $Q_{sub}^{(n)}$ at every refinement order corresponds to a non-trivial admissible configuration collapsing to 0 under continuum-limit transport, contradicting finite distinguishability inheritance.

Hence $\eta_{\mu\nu}$ is nondegenerate.

Lemma 4.3 — Uniqueness up to scale [proven, conditional on foundational C_5–C_7 bilinear-form-sector rigidity inheritance]

The continuum-limit bilinear form $\eta_{\mu\nu}$ of Theorem 4.1 is unique up to overall scale: any other refinement-stable bilinear form $\eta'_{\mu\nu}$ *preserved by* $\mathcal{G}_\infty$ satisfies $\eta'_{\mu\nu} = c\, \eta_{\mu\nu}$ for some scalar $c \in \mathbb{R} \setminus \{0\}$.

Proof

Any second refinement-stable bilinear form $\eta'_{\mu\nu}$ *preserved by* $\mathcal{G}_\infty$ corresponds, in the substrate-level limit, to a second substrate-level quadratic invariant $Q'_{sub}^{(n)}$ preserved by every $U_{ij} \in \mathcal{G}_{sub}^{(n)}$ up to $\mathcal{O}(\varepsilon_n^2)$ corrections.

The C_5–C_7 hexagonal embedding constraint imposes a six-fold cyclic consistency relation on bilinear forms restricted to the six adjacency-coding directions of the K = 7 cell, with the central-radial direction constrained by (T7) BCB admissibility. The reduction of the resulting bilinear-form sector to a one-parameter family (modulo orientation) is a foundational rigidity result distinct from — though structurally adjacent to — the C_5–C_7 cell-architecture rigidity that establishes uniqueness of the K = 7 substrate itself.

Inheritance distinction. Two C_5–C_7 rigidity results from the foundational papers are invoked in the present paper, and they should not be conflated:

- *Cell-architecture rigidity*: the K = 7 hexagonal substrate is the unique BCB-admissible cell architecture (foundational papers, used throughout the VERSF programme).
- *Bilinear-form-sector rigidity*: the space of refinement-stable bilinear forms preserved by closure-compatible transport on the K = 7 cell is at most one-dimensional modulo orientation (foundational papers, used here in Lemma 4.3 and §4 Step 3).

Both are inputs to the present paper. The second is the load-bearing inheritance for the present Lemma; it is cited as a separate foundational result rather than as a consequence of the first.

Therefore $Q'_{sub}^{(n)} = c_n\, Q_{sub}^{(n)} + \mathcal{O}(\varepsilon_n^2)$; taking the continuum limit, $\eta'_{\mu\nu} = c\, \eta_{\mu\nu}$ for some $c \in \mathbb{R} \setminus \{0\}$.

Structural significance

Theorem 4.1 supplies the substrate-level bilinear-form preservation that the Microscopic Origin paper's Definition 3.1 inherited as $\mathcal{G} = SO(1, D)$'s defining metric. The existence of $\eta_{\mu\nu}$ is now derived from BCB bit-conservation at the substrate level rather than inherited from the continuum-limit signature.

The uniqueness-up-to-scale is the strongest result the substrate construction can deliver without further input: the overall scale of $\eta_{\mu\nu}$ is not fixed by the closure architecture alone. This is recorded as OP8 below — substrate derivation of the bilinear-form scale, which would correspond, after dimensional analysis, to the substrate origin of the coherence-scale ξ of the CCC framework.

---

5. Hyperbolic Signature Selection

We now determine the signature of the bilinear form $\eta_{\mu\nu}$.

Lemma 5.0 — Coherent point-level commitment orientation [proven, conditional on closure-compatible refinement]

Under closure-compatible refinement, the commitment orientations of edges incident to a substrate vertex contract to a coherent single direction at each continuum-limit point: the continuum-limit cone of commitment-propagation directions at each point is one-dimensional modulo positive scaling.

Proof

At refinement order n, the edges incident to a vertex $v_i^{(n)}$ carry commitment orientations forming a discrete angular distribution determined by the $K = 7$ cell architecture: six adjacency-coding boundary edges plus the central edge, with the angular separation between distinct edges *scale-invariant* under closure-compatible refinement. The angular structure is a structural feature of the $K = 7$ closure constraints C_5–C_7, not a property that contracts as $\varepsilon_n \to 0$.

The continuum-limit point-level commitment orientation is therefore *not* obtained by angular contraction of the discrete cell-level orientations, but by **coarse-graining of BCB-balanced edge-orientations** across the increasingly fine cell structure under closure-compatible refinement.

Concretely: each substrate-level edge incident to $v_i^{(n)}$ contributes its commitment orientation with weight given by the BCB-balanced bit-flow through that edge. The cell-level bit-balance constraints C_5–C_7 force the weighted vector sum of incident-edge orientations to align with a coherent direction at the continuum point — the **net commitment-flow direction** at $v_i^{(n)}$ under BCB balance. This coherent direction is well-defined at each substrate vertex by the $K = 7$ cell-balance constraints (the six hexagonal closure traversals plus the central traversal produce a single net flow direction per cell, not seven independent ones).

Under closure-compatible refinement, the substrate-level net commitment-flow directions at refined vertices $v_i^{(n+1)}$ inherit consistently from parent vertices $v_i^{(n)}$ via the refined-edge-orientation inheritance of Definition 3.2 of the Microscopic Origin paper: each refined edge inherits the commitment direction of its parent edge, and the BCB balance at the refined cell level propagates the parent's net flow direction to the refined cells.

As $\varepsilon_n \to 0$, the coarse-graining averages over increasingly many $K = 7$ cells per continuum-limit unit volume, with the substrate-level net commitment-flow directions converging to a continuous timelike vector field on the emergent continuum manifold. At each continuum point, the resulting timelike direction is *one-dimensional modulo positive scaling*: the substrate ontology supplies one commitment orientation per edge (the deep one-dimensionality of time, recorded as Remark (b) below), and the cell-level BCB balance supplies one net flow direction per cell, with the closure-compatible refinement averaging preserving this one-dimensionality through the $n \to \infty$ limit.

The continuum-limit cone of commitment-propagation directions at each point is therefore one-dimensional modulo positive scaling, inherited from the one-dimensionality of commitment orientation per edge via closure-compatible coarse-graining of BCB-balanced cell flows — not from angular contraction of the substrate-level discrete cone.

Lemma 5.1 — Multi-timelike refinement instability [proven, conditional on closure-compatible refinement and irreversible-commitment ontology]

Suppose the continuum-limit bilinear form $\eta_{\mu\nu}$ has signature (p, q) with $p \geq 2$. Then the substrate commitment-order partial order $\prec$ on $S^{(n)}$ cannot be lifted to a globally refinement-compatible causal ordering on the continuum-limit manifold under closure-compatible refinement.

Equivalently: only signatures with exactly one negative eigenvalue ($p = 1$) admit globally refinement-compatible irreversible ordering. Ultra-hyperbolic signatures ($p \geq 2$) are *refinement-incompatible*, not merely ontologically peculiar.

Proof

Five steps.

Step 1 — multiple timelike foliations. For $\eta_{\mu\nu}$ with signature (p, q), $p \geq 2$, the η-negative-norm cone at each continuum point is $(p - 1)$-dimensional modulo positive scaling: there exists a $(p - 1)$-dimensional family of independent timelike directions at each point. For $p = 2$ this gives a one-parameter family of distinct timelike foliations; for $p \geq 3$, a higher-dimensional family.

Step 2 — multiple potential causal orderings. Each independent timelike direction generates its own potential continuum causal ordering: for V_1, V_2 two independent timelike vectors with $\eta(V_1, V_1) < 0$ and $\eta(V_2, V_2) < 0$, the V_1-foliation and V_2-foliation produce distinct partial orders on the continuum manifold, with distinct light cones, distinct past-and-future sets, and distinct causal closures.

Step 3 — substrate provides only one partial order. By Lemma 5.0, the substrate commitment-order partial order $\prec$ on $S^{(n)}$ lifts coherently to a single timelike direction at each continuum point — there is *one* commitment orientation, not p. The substrate therefore provides only one canonical partial order, not p independent ones.

Step 4 — refinement non-equivalence. Closure-compatible refinement sequences along different timelike foliations (V_1 vs V_2) would generate distinct continuum-limit causal orderings, each consistent with the substrate-level $\prec$ on substrate vertices but distinct as continuum-level partial orders on the manifold. These refinement-limit orderings are *non-equivalent*: they disagree on which continuum points lie in the causal future of a given point, on light-cone boundaries, and on past-history sets.

The non-equivalence is genuine, not gauge: each foliation generates a different geometric structure, with no canonical map between them in the absence of a preferred timelike direction.

Step 5 — refinement-compatibility violation. The closure-compatible refinement framework of the Microscopic Origin paper requires that all closure-compatible refinement sequences approximating a given continuum manifold yield the *same* continuum-limit causal structure — equivalently, that the refinement-limit causal ordering is independent of the refinement-order sequence (the structural analogue of Theorem 5.3 Step 2 of the Microscopic Origin paper, applied to causal structure rather than to path holonomy).

The lift from path-holonomy refinement-Cauchy convergence (Microscopic Origin paper Theorem 5.3) to causal-structure refinement-Cauchy convergence is itself part of the closure-compatible refinement framework: closure-compatible refinement of the commitment foam preserves the commitment-order partial order $\prec$ under subdivision (Definition 3.2 of the Microscopic Origin paper preserves commitment orientation; commitment-order $\prec$ is generated by oriented commitment edges via Definition 8.1 of that paper), and the continuum-limit coarse-graining of $\prec$ is consistent with the inherited Lorentzian causal structure by Theorem 8.2 of the Microscopic Origin paper. The refinement-Cauchy convergence of causal structure is therefore inherited from the same closure-compatible refinement architecture that produces refinement-Cauchy convergence of path holonomy, not a separately-postulated property.

Multiple non-equivalent refinement orderings violate this refinement-compatibility requirement: the continuum-limit causal structure would depend on the refinement sequence used to approximate it, contradicting the closure-compatible refinement framework.

Therefore $p \geq 2$ is incompatible with globally refinement-compatible irreversible ordering. Combined with the indefiniteness result of Theorem 5.2 Step 1 ($p \geq 1$), the only admissible signature is $p = 1$.

Remark — what Lemma 5.1 contributes structurally

This is the strongest structural argument the paper makes for single-time-dimensionality. A purely ontology-led argument — invoking only Lemma 5.0 plus the substrate ontology of one commitment orientation per edge — would be ontologically correct but somewhat intuition-led.

Lemma 5.1 replaces this with a *refinement-compatibility* argument: $p \geq 2$ is not merely incompatible with the substrate ontology, it is incompatible with the closure-compatible refinement framework itself. The signature constraint $p = 1$ is therefore not an ontological choice but a *structural necessity* of the refinement architecture.

This is the substrate-level analogue of the standard mathematical-physics fact that ultra-hyperbolic geometries do not admit globally well-posed causal evolution: the substrate construction now exhibits this not merely as a continuum-limit consequence of the inherited Lorentzian Completion structure but as a refinement-architectural requirement at the substrate level. Lorentzian signature (1, D) is the unique signature compatible with globally refinement-stable causal ordering.

Theorem 5.2 — Lorentzian signature [proven, conditional on CCC framework, Lemma 5.1 multi-timelike instability, Lemmas 5.0 and 4.2, and finite distinguishability of orthogonal directions]

The continuum-limit bilinear form $\eta_{\mu\nu}$ of Theorem 4.1 has Lorentzian signature with exactly one negative eigenvalue and exactly D positive eigenvalues:

$$\text{signature}(\eta_{\mu\nu}) = (1, D),$$

corresponding to $(-, +, +, \ldots, +)$ in the mostly-plus convention of the dynamical paper §2.1.

Proof

Three substrate-level inputs jointly determine the signature.

Step 1 — indefinite signature from CCC + irreversibility.

The inherited CCC framework supplies finite-speed coherence-preserving propagation at the substrate level, dividing admissible directions into a propagation interior cone and a non-propagation exterior. CCC alone is logically compatible with a positive-definite $Q_{sub}^{(n)}$ plus an *externally imposed* preferred timelike vector field. The additional substrate-level input required for the cone structure to be built *into* the bilinear form rather than imposed externally is the irreversibility-of-commitment ontology of the source paper §2.1.

The substrate-level mechanism is *irreversibility as distinguishability consumption*. Each commitment event converts potential distinguishability into actualised commitment, irreversibly. At the BCB bit-count level:

- bit-count along *spatial* (non-propagation, orthogonal) directions measures distinguishability *preserved*: configurations distinguishable at fixed commitment-time remain distinguishable;
- bit-count along the *commitment-propagation* direction measures distinguishability *consumed*: each unit of propagation expends one unit of potential-distinguishability into actualised commitment.

The two contributions to the BCB bilinear pairing therefore enter with opposite signs:

$$\langle u, v\rangle_{sub}^{(n)} = (\text{spatial distinguishability of } u, v) - (\text{commitment-propagation overlap of } u, v).$$

The opposite-sign structure is not a postulated metric signature but a substrate-level consequence of the irreversibility ontology. Propagation directions and spatial directions contribute to the BCB bit-balance with opposite signs because they are doing opposite things to the same bit-resource: propagation consumes, spatial directions preserve.

The articulation "distinguishability consumption" is treated as an interpretation of the irreversibility ontology of source paper §2.1 — structurally implicit there but not, to the author's knowledge, isolated as a separate foundational theorem. It is the substrate-level mechanism by which irreversibility produces the indefinite-signature structure of the BCB bilinear pairing; closing it to a separate foundational theorem (rather than an articulation of inherited ontology) is a residual structural problem at the foundational layer rather than within the present paper.

In the continuum limit, $\eta_{\mu\nu}$ inherits this opposite-sign structure by Theorem 4.1, giving an indefinite-signature bilinear form: at least one negative eigenvalue (propagation directions) and at least one positive eigenvalue (spatial directions). A purely elliptic $\eta_{\mu\nu}$ would require either no propagation directions (contradicting CCC) or propagation directions with the same sign as spatial directions (contradicting the irreversibility consumption mechanism). A degenerate $\eta_{\mu\nu}$ is excluded by Lemma 4.2.

Step 2 — at most one negative eigenvalue from refinement-compatibility and substrate ontology.

Two independent substrate-level inputs jointly force at most one negative eigenvalue:

(a) **Refinement-compatibility (Lemma 5.1).** Any signature (p, q) with $p \geq 2$ generates multiple non-equivalent continuum-limit causal orderings under closure-compatible refinement, violating the refinement-stability of substrate commitment ordering. The signature constraint $p \leq 1$ is therefore a structural necessity of the refinement architecture, not merely an ontological inheritance.

(b) **Point-level commitment coherence (Lemma 5.0).** Independent of (a), Lemma 5.0 establishes that the continuum-limit cone of commitment-propagation directions at each point is one-dimensional modulo positive scaling, with the substrate-level one-dimensionality of commitment orientation per edge lifting coherently to a single timelike direction at each continuum point.

A bilinear form $\eta_{\mu\nu}$ with two or more negative eigenvalues would generate either non-equivalent refinement orderings (contradicting Lemma 5.1) or a multi-dimensional family of η-negative directions at each point (contradicting Lemma 5.0). The two inputs are complementary: Lemma 5.0 establishes that the substrate *provides* only one direction; Lemma 5.1 establishes that the refinement architecture *requires* only one direction.

Therefore $\eta_{\mu\nu}$ has at most one negative eigenvalue.

Step 3 — exactly D positive eigenvalues from finite distinguishability.

The remaining directions — orthogonal to the commitment orientation — form the spacelike-like sector. By finite distinguishability (inherited from source paper §2.1): the substrate admits finitely many distinguishable orthogonal directions per bounded region; in the continuum limit, these lift to a continuous D-dimensional manifold of spacelike directions ($D = 3$ in the physically central case $D + 1 = 4$).

By Lemma 4.2 (nondegeneracy), $\eta_{\mu\nu}$ cannot have null orthogonal directions: every spacelike-like direction has $Q(V) \neq 0$. Combined with Step 1 (indefinite signature with at least one positive eigenvalue) and Step 2 (at most one negative eigenvalue): the remaining D-dimensional orthogonal sector consists of exactly D positive eigenvalues.

Assembling: $\eta_{\mu\nu}$ has exactly one negative eigenvalue and exactly D positive eigenvalues. Signature (1, D), conventionally written (−, +, +, …, +) in the mostly-plus convention.

Remark — what Theorem 5.2 does and does not derive

What it does derive: the *signature* of $\eta_{\mu\nu}$ as (1, D), from three substrate-level inputs — CCC framework (Step 1), single-direction commitment orientation (Step 2), and finite distinguishability of orthogonal directions (Step 3).

What it does not derive:

(a) **The CCC framework itself** is inherited from the Lorentzian Completion paper. The substrate-level mechanism by which coherence propagation produces finite-speed causality is a continuum-emergence question of that paper, not re-derived here.

(b) **The single-direction commitment orientation per edge** is the substrate-level statement of *time being one-dimensional*. This is the deepest structural input of the irreversibility ontology of VERSF and is a foundational structural feature of the framework rather than a derivable consequence within the present paper. The substrate-level question — *why* is commitment orientation one-dimensional rather than multi-dimensional? — corresponds at the continuum level to the question of why time is one-dimensional, and is an open foundational problem distinct from the OP9 transport-group derivation closed here.

(c) **The dimension D + 1 = 4 itself** is the physically central case, with D = 3 corresponding to the three spatial dimensions of observed geometry. The substrate construction allows arbitrary $D \geq 1$; the selection D + 1 = 4 is a phenomenological input, not a substrate derivation.

The three inputs (a), (b), (c) are recorded as inherited / external to the present derivation. The honest summary: given the CCC framework, the single-direction commitment ontology, and finite distinguishability, the signature is forced to be (1, D) — but each of these three inputs is itself a structural inheritance.

Remark — K = 7 local architecture vs. continuum dimensionality

The K = 7 closure architecture is the substrate's *local cell* structure: six adjacency-coding boundary vertices plus one gauge-anchoring central vertex per cell, with $N_{loop} = 14$ oriented closure traversals. The cell carries a hexagonal boundary structure plus a central gauge anchor: combinatorially two-dimensional in its boundary adjacency (the six-fold hexagonal pattern of boundary vertices), but with a distinguished centre that supplies the BCB gauge anchoring. The cell is therefore not a planar tiling unit in a strict sense — it is a hexagonal-boundary-plus-centre configuration whose adjacency structure is two-dimensional but whose gauge structure is fixed by the central vertex.

The substrate's *global* dimensionality — the spatial dimension D of the continuum-limit manifold — is a separate property, established as $D = 3$ (giving $D + 1 = 4$ with the commitment-propagation direction) for the physically central case. The lift from $K = 7$ local hexagonality to $D = 3$ global dimensionality is the content of the foundational papers' substrate-embedding arguments: each 2D-hexagonal cell is embedded in a higher-dimensional substrate combinatorial structure, with the global $D = 3$ spatial dimensionality emerging from the tiling pattern of $K = 7$ cells across the substrate.

The present paper inherits the $K = 7 \leftrightarrow D = 3$ correspondence without re-derivation. The two are structurally independent inheritances: $K = 7$ is the unique BCB-admissible *local* cell architecture (foundational rigidity), while $D = 3$ is the *global* spatial dimensionality of the continuum-limit substrate manifold (foundational substrate-embedding). Closing the global-dimensionality inheritance — deriving $D = 3$ from substrate combinatorics independent of phenomenological selection — would be a deeper foundational problem distinct from the OP9 transport-group derivation closed here.

Structural significance

Theorem 5.2 supplies the substrate-level *signature selection* that the Microscopic Origin paper inherited via the identification $\mathcal{G} = SO(1, D)$. The signature is now derived from CCC + substrate ontology rather than imposed by inheriting the Lorentzian metric. The relationship between substrate-level $\eta_{\mu\nu}$ and continuum-limit metric $g_{\mu\nu}$ is therefore the natural one: $\eta_{\mu\nu}$ is the bilinear form on tangent fibres that the substrate produces, and the continuum-limit metric $g_{\mu\nu}$ on the continuum manifold is its natural lift via the inherited frame structure.

---

6. Refinement Stability of Lorentz-Compatible Transport

The previous sections established the existence and signature of $\eta_{\mu\nu}$, and hence identified the bilinear-preserving group as $O(1, D)$ (the full orthogonal group preserving $\eta_{\mu\nu}$). To complete the derivation of $\mathcal{G}_\infty \cong SO(1, D)$, we must show that:

(a) anisotropic transport deviations from $SO(1, D)$ compatibility are exponentially suppressed under refinement (Theorem 6.2); (b) the connected, orientation-preserving and time-orientation-preserving subgroup $SO(1, D)^\uparrow$ is the refinement-stable fixed point (Corollary 6.3 and Theorem 7.1).

Definition 6.1 — Transport anisotropy measure

Let $U_{ij}^{(n)}$ be a substrate-level edge transport operator at refinement order n. The **transport anisotropy** of $U_{ij}^{(n)}$ is

$$\delta(U_{ij}^{(n)}) := \inf_{\Lambda \in O(1, D)} \|U_{ij}^{(n)} - \Lambda\|,$$

i.e. the operator-norm distance from $U_{ij}^{(n)}$ to the nearest element of the bilinear-preserving group $O(1, D)$ of $\eta_{\mu\nu}$. The **cell-averaged transport anisotropy** at refinement order n is

$\Delta_n := \sup_{\text{cells } C \text{ of } S^{(n)}}$ (cell-average of $\delta(U_{ij}^{(n)})$ over edges in C).

Δ_n is uniformly bounded over bounded regions by the closure-compatibility constraints (T1)–(T7) of Definition 3.1.

Definition 6.2 — Anisotropy mode space and refinement operator $\mathscr{R}_{aniso}$

The **anisotropy mode space** at refinement order n is

$\mathcal{A}_n :=$ the vector space of cell-level transport-anisotropy configurations on $S^{(n)}$,

equipped with the cell-averaging norm. An element $\psi \in \mathcal{A}_n$ is a configuration of cell-averaged anisotropy values $\delta(U_{ij}^{(n)})$ over the cells of $S^{(n)}$, considered as a linear-space element under closure-compatible cell averaging.

The closure-compatible refinement map $S^{(n)} \to S^{(n+1)}$ induces a linear operator

$\mathscr{R}_{aniso} : \mathcal{A}_n \to \mathcal{A}_{n+1}$,

the **anisotropy refinement operator**, defined by averaging parent-cell anisotropies across refined cells under closure-compatible subdivision. Concretely, for $\psi \in \mathcal{A}_n$ a parent-cell anisotropy configuration, $\mathscr{R}_{aniso}\, \psi \in \mathcal{A}_{n+1}$ is the refined-cell anisotropy configuration induced by closure-compatible subdivision of each parent cell into refined cells, with the parent anisotropy distributed across the refined cells according to the $K = 7$ cell-subdivision combinatorics.

In the continuum-limit refinement-stable sector, $\mathscr{R}_{aniso}$ acts as a linear operator on the inverse limit

$\mathcal{A} := \lim \leftarrow_n \mathcal{A}_n$,

the **continuum-limit anisotropy mode space**, with spectral content determined by the $K = 7$ / TPB refinement-flow structure.

Theorem 6.2 — Spectral suppression of transport anisotropy [proven, conditional on foundational C_5–C_7 refinement-cell rigidity inheritance]

The anisotropy refinement operator $\mathscr{R}_{aniso}$ of Definition 6.2 has spectral radius

$\rho(\mathscr{R}_{aniso}) = \lambda_{aniso} < 1$,

with the strict inequality $\lambda_{aniso} < 1$ forced by foundational C_5–C_7 refinement-cell rigidity. The continuum-limit anisotropy mode space admits a spectral decomposition

$\mathcal{A} = \mathcal{A}^{Lor} \oplus \mathcal{A}^{aniso}$,

where:

- $\boldsymbol{\mathcal{A}^{Lor}}$ is the Lorentz-compatible (kernel) sector — eigenvalue $\rho = 1$, refinement-invariant. Elements of $\mathcal{A}^{Lor}$ represent transport configurations preserved exactly by closure-compatible refinement; these are precisely the SO(1, D)-compatible configurations.
- $\boldsymbol{\mathcal{A}^{aniso}}$ is the anisotropic sector — eigenvalues $\rho_k < 1$, refinement-suppressed. Elements of $\mathcal{A}^{aniso}$ represent transport configurations that decay exponentially under refinement.

The cell-averaged transport anisotropy at refinement order n satisfies

$$\|\Delta_n\| = \|\mathscr{R}_{aniso}^n \Delta_0\| \leq C_0 \cdot \lambda_{aniso}^n,$$

with $\eta := -\log \lambda_{aniso} > 0$ the **refinement-flow spectral gap**.

Proof

Four steps.

Step 1 — $\mathscr{R}_{aniso}$ is well-defined and linear.

The refinement of $S^{(n)}$ to $S^{(n+1)}$ is closure-compatible by Definition 3.2 of the Microscopic Origin paper, with refined edge transport operators consistent with parent edge transport operators via Lemma 4.4 of that paper. The map $\mathscr{R}_{aniso} : \mathcal{A}_n \to \mathcal{A}_{n+1}$ on cell-averaged anisotropy configurations is therefore well-defined, single-valued, and linear in the cell-averaging structure. The $\mathcal{O}(\varepsilon_n^2)$ corrections of Lemma 4.4 contribute only to the higher-order (non-linear) sector and are subleading to the linear-operator behaviour.

Step 2 — strict contractivity from C_5–C_7 refinement-cell rigidity.

The C_5–C_7 hexagonal embedding constraint at each refined cell imposes a six-fold hexagonal consistency relation on the refined-cell anisotropies induced from a parent cell. The foundational *refinement-cell rigidity* lemma establishes that this consistency relation has trivial kernel — i.e. the only refined-cell anisotropy configuration satisfying C_5–C_7 that does not reduce the parent-cell anisotropy norm is the zero configuration.

The non-triviality of the kernel ({0}) is the foundational input that forces strict contraction:

$$\|\mathscr{R}_{aniso}\, \psi\| < \|\psi\| \text{ for all } \psi \neq 0 \text{ in the anisotropic sector},$$

which is equivalent to $\rho(\mathscr{R}_{aniso}) < 1$.

Inheritance distinction. This is a *third* C_5–C_7 rigidity statement from the foundational papers, structurally adjacent to but distinct from the two invoked in §4:

- *Cell-architecture rigidity*: uniqueness of the K = 7 substrate as the BCB-admissible cell architecture.
- *Bilinear-form-sector rigidity* (Lemma 4.3 inheritance): one-parameter rigidity of the bilinear-form sector preserved by closure-compatible transport on the K = 7 cell.
- *Refinement-cell rigidity* (present inheritance): trivial kernel of the C_5–C_7 consistency relation on refined-cell anisotropies, forcing strict contraction of $\mathscr{R}$_aniso.

The present paper invokes all three as separate foundational inheritances. The numerical value of λ_aniso requires the quantitative Wilson-style one-loop analysis recorded as OP7.

Step 3 — spectral decomposition.

The continuum-limit anisotropy mode space $\mathcal{A}$ admits the spectral decomposition $\mathcal{A}$ = $\mathcal{A}$^{Lor} ⊕ $\mathcal{A}$^{aniso} according to the spectrum of $\mathscr{R}$_aniso:

- **$\mathcal{A}$^{Lor} = ker($\mathscr{R}$_aniso − I)**: the eigenspace of eigenvalue 1, consisting of refinement-invariant configurations. By construction, these are anisotropy configurations preserved exactly by closure-compatible refinement — equivalently, transport configurations consistent with O(1, D)-compatibility (and hence with SO(1, D)-compatibility after orientation/time-orientation restriction via Lemmas 6.3a/b below).
- **$\mathcal{A}$^{aniso} = the complementary spectral sector**: eigenspaces of eigenvalue |ρ_k| < 1, consisting of configurations contracted under refinement. These are anisotropy configurations strictly suppressed under closure-compatible refinement.

The decomposition is direct because $\mathscr{R}$_aniso is a closed linear operator on $\mathcal{A}$ with spectrum bounded by 1 (eigenvalue 1 attained on $\mathcal{A}$^{Lor}, all other eigenvalues strictly inside the unit disk by Step 2).

Step 4 — exponential suppression iterate.

For any initial anisotropy configuration Δ_0 ∈ $\mathcal{A}$, write the spectral decomposition Δ_0 = Δ_0^{Lor} + Δ_0^{aniso}. Iterating $\mathscr{R}$_aniso:

Δ_n = $\mathscr{R}$_aniso^n Δ_0 = Δ_0^{Lor} + $\mathscr{R}$_aniso^n Δ_0^{aniso}.

The Lorentz-compatible component is refinement-invariant; the anisotropic component is bounded by

‖$\mathscr{R}$_aniso^n Δ_0^{aniso}‖ ≤ ρ($\mathscr{R}$_aniso)^n ‖Δ_0^{aniso}‖ = λ_aniso^n ‖Δ_0^{aniso}‖.

Setting C_0 := ‖Δ_0^{aniso}‖ and restricting attention to the anisotropic sector (the part of Δ_n that is *not* refinement-invariant):

$\|\Delta_n - \Delta_0^{Lor}\| \le C_0 \cdot \lambda_{aniso}^n$,

with $\eta := -\log \lambda_{aniso} > 0$ the refinement-flow spectral gap. The Lorentz-compatible component $\Delta_0^{Lor} \in \mathcal{A}^{Lor}$ is preserved exactly; the anisotropic component vanishes exponentially.

Remark — structural significance of the spectral upgrade

The spectral statement on the well-defined linear operator $\mathcal{R}_{aniso}$ replaces qualitative refinement-flow descriptions ("$\Delta_n \le C_0 e^{-\eta n}$") with explicit operator-theoretic content. The pay-off is structural:

1. **SO(1, D) is the dominant fixed-point eigenspace** of $\mathcal{R}_{aniso}$, not merely a stable configuration: $\mathcal{A}^{Lor}$ is the eigenvalue-1 eigenspace of an explicitly-defined linear operator.
2. **Anisotropy suppression is spectral contraction**, not an inherited RG-flow heuristic: the rate η is the spectral gap $\log(1/\rho(\mathcal{R}_{aniso}))$.
3. **Mode decomposition** $\Delta = \Delta^{Lor} + \Delta^{aniso}$ cleanly separates Lorentz-compatible content (preserved) from anisotropic content (suppressed), with refinement acting by orthogonal projection in the limit $n \to \infty$.

This is the structural cleanup the §6 framework requires. What remains inherited is the *trivial-kernel* property of the C_5–C_7 refinement-cell consistency relation — the foundational result that forces $\lambda_{aniso} < 1$ rather than $\lambda_{aniso} \le 1$.

Remark — what Theorem 6.2 does and does not derive

What it derives: the *existence* of a spectral gap $\eta > 0$ in the anisotropy refinement operator, the spectral decomposition $\mathcal{A} = \mathcal{A}^{Lor} \oplus \mathcal{A}^{aniso}$, and the exponential suppression of anisotropic modes under refinement.

What it does not derive: the *quantitative value* of λ_{aniso} (equivalently η) as an explicit function of the $K = 7$ bare-coupling integer $N_{loop} = 14$ or other foundational substrate parameters. This is the quantitative one-loop matching analysis that the $K = 7$ Wilson Limit paper performs for the U(1) gauge-bundle sector; the analogous quantitative computation for the non-abelian SO(1, D) tangent-bundle sector is recorded as OP7 below. The present paper establishes that $\rho(\mathcal{R}_{aniso}) < 1$; the explicit numerical value of ρ is left open.

Lemma 6.3a — Orientation preservation [proven, conditional on closure-compatible cell-orientation inheritance]

Closure-compatible refinement on the $K = 7$ foam preserves cell orientation: refined cells inherit the parent cell's orientation under subdivision. The refinement-stable transport fixed point therefore lies in the proper subgroup of O(1, D), i.e. $\det \Lambda = +1$.

Proof

By Definition 3.2 of the Microscopic Origin paper, closure-compatible refinement preserves the $K = 7$ closure structure on each cell, with refined cells inheriting the parent cell's structural orientation under subdivision. The cell-orientation inheritance is part of the foundational $K = 7$ closure architecture, which fixes a consistent orientation on each maximal cell (via the central-vertex gauge anchoring of the $K = 7$ closure structure).

Under refinement, orientation-reversing transport operators would generate cell-orientation inconsistencies between parent and refined cells, contradicting closure-compatible refinement. The refinement-stable fixed point therefore lies in the orientation-preserving subgroup of O(1, D), i.e. $\det \Lambda = +1$.

Lemma 6.3b — Time-orientation preservation [proven, conditional on irreversibility of commitment]

The irreversible-commitment ontology forbids reversal of commitment-propagation direction under refinement. The refinement-stable transport fixed point therefore lies in the orthochronous subgroup of O(1, D), i.e. $\Lambda^0{}_0 \geq 1$.

Proof

By the irreversibility-of-commitment ontology (source paper §2.1 and Definition 3.3 (F2) of the Microscopic Origin paper), commitment orientation cannot be reversed at the substrate level: once a commitment event occurs at v_j with $v_j \prec v_i$, the commitment-order partial order is irrevocable.

Closure-compatible refinement preserves commitment orientation by Definition 3.2 of the Microscopic Origin paper: refined edges inherit the commitment direction of the parent edge. Therefore time-orientation-reversing transport operators — those mapping future-directed commitment-propagation to past-directed — would generate commitment-orientation inconsistencies under refinement, contradicting closure-compatible refinement.

The refinement-stable fixed point therefore lies in the time-orientation-preserving (orthochronous) subgroup of O(1, D), characterised by $\Lambda^0{}_0 \geq 1$ in the mostly-plus convention.

Corollary 6.3 — SO(1, D) as refinement fixed point [proven, conditional on Theorem 6.2, Lemma 6.3a, and Lemma 6.3b]

SO(1, D)-compatible transport — i.e. substrate transport structures with $U_{ij} \in SO(1, D)^\uparrow$ (the connected, orientation- and time-orientation-preserving subgroup of O(1, D)) at every edge — is the unique refinement-stable fixed point of closure-compatible refinement.

Proof

By Theorem 6.2, anisotropic deviations from O(1, D)-compatibility are exponentially suppressed under refinement; the refinement-stable fixed point is therefore contained in O(1, D). By Lemma 6.3a, the fixed point lies in the proper subgroup ($\det \Lambda = +1$). By Lemma 6.3b, the fixed point lies in the orthochronous subgroup ($\Lambda^0_0 \geq 1$). The connected, proper, orthochronous subgroup of O(1, D) is the proper orthochronous Lorentz group SO(1, D)↑.

Henceforth we write SO(1, D) for SO(1, D)↑ — the connected, proper, orthochronous Lorentz group — with the orientation and time-orientation conditions implicit.

---

7. Transport-Group Emergence Theorem

We now assemble the previous results into the main theorem of the paper.

Theorem 7.1 — Lorentz transport-group emergence [proven, conditional on Theorems 4.1, 5.2, 6.2, Lemmas 5.0, 5.1, and Corollary 6.3]

The continuum-limit admissible transport group on the K = 7 simplicial commitment foam is

$\mathcal{G}_\infty \cong \mathrm{SO}(1, D)$,

with continuum refinement transport operators

$T_\gamma \in \mathrm{SO}(1, D)$

for every continuum path γ on the emergent Lorentzian continuum.

Proof

Direct assembly.

Step 1 — bilinear preservation. By Theorem 4.1, the continuum-limit transport group $\mathcal{G}_\infty$ preserves a refinement-stable invariant bilinear form $\eta_{\mu\nu}$, unique up to overall scale (Lemma 4.3).

Step 2 — signature. By Theorem 5.2 (via Lemmas 5.0 and 5.1), $\eta_{\mu\nu}$ has Lorentzian signature (1, D). The maximal group preserving $\eta_{\mu\nu}$ is therefore the full orthogonal group O(1, D).

Step 3 — refinement stability and connectedness. By Theorem 6.2, anisotropic deviations from O(1, D)-compatibility are spectrally suppressed under refinement, with SO(1, D)-compatible transport corresponding to the ρ = 1 eigenspace $\mathcal{A}^{Lor}$ of $\mathcal{R}_{aniso}$. By Corollary 6.3 (via Lemmas 6.3a and 6.3b), the refinement-stable fixed point is the connected, proper, orthochronous subgroup SO(1, D)↑.

Combining the three: $\mathcal{G}_\infty \cong$ SO(1, D)↑, henceforth written simply as SO(1, D) with orientation and time-orientation conditions implicit.

By construction, $T_\gamma \in \mathcal{G}_\infty$ for every continuum path γ (refinement transport operators are by definition elements of the continuum-limit transport group). Hence $T_\gamma \in SO(1, D)$.

Theorem 7.2 — Fixed-point uniqueness [proven, conditional on Theorems 4.1, 5.2, 6.2, Lemmas 4.3, 5.0, 5.1, 6.3a, 6.3b]

The proper orthochronous Lorentz group SO(1, D)↑ is the **unique** admissible refinement-stable transport group on the K = 7 simplicial commitment foam.

Equivalently: under the closure-compatibility constraints (T1)–(T7) of Definition 3.1, BCB bit-conservation, finite distinguishability, CCC, irreversible-commitment ontology, and refinement-stability, no transport group other than SO(1, D)↑ survives as the continuum-limit transport group.

Proof

Each of the substrate inputs eliminates a distinct class of alternative transport groups; together they admit only SO(1, D)↑ as the refinement-stable fixed point.

Obstruction inventory. The following classes of alternative transport groups are excluded:

(O1) **Alternative bilinear forms.** Any transport group preserving a bilinear form $\eta'_{\mu\nu}$ not proportional to $\eta_{\mu\nu}$ is excluded by Lemma 4.3 (bilinear-form-sector rigidity from C_5–C_7): the bilinear-form sector of refinement-stable transport admits only the one-parameter family $\eta_{\mu\nu}$ up to overall scale.

(O2) **Degenerate bilinear forms.** Any transport group preserving a degenerate bilinear form is excluded by Lemma 4.2 (nondegeneracy from finite distinguishability).

(O3) **Multi-timelike signatures.** Any transport group with bilinear-form signature (p, q), $p \geq 2$, is excluded by Lemma 5.1 (multi-timelike refinement instability): ultra-hyperbolic signatures generate non-equivalent refinement orderings, violating closure-compatible refinement.

(O4) **Multi-direction commitment incompatibility.** Any transport group whose preserved bilinear form admits a multi-dimensional family of timelike directions is excluded by Lemma 5.0 (coherent point-level commitment orientation): the substrate provides only one commitment-orientation direction per continuum point.

(O5) **Elliptic (positive-definite) signatures.** Any compact transport group (e.g. SO(D + 1)) preserving a positive-definite form is excluded by Theorem 5.2 Step 1 (CCC + irreversibility-as-distinguishability-consumption): the BCB bilinear pairing has indefinite signature by the opposite-sign contributions of commitment-propagation and spatial directions.

(O6) **Anisotropic non-O(1, D) groups.** Any transport group failing O(1, D)-compatibility — i.e. lying outside the closure of O(1, D) — is excluded by Theorem 6.2 (spectral suppression): such groups lie in the anisotropic mode sector $\mathcal{A}^{aniso}$ with eigenvalue $\rho_k < 1$, hence are refinement-suppressed and cannot survive as the $n \to \infty$ fixed point.

Note on (O6). Logically, (O6) is a *spectral consistency check* rather than an obstruction independent of (O1)–(O5): any candidate transport group failing to lie within the closure of O(1, D) already fails to preserve $\eta_{\mu\nu}$ (O1), fails to admit signature (1, D) (O3, O5), or fails one of the substrate-ontology constraints. (O6) is retained in the enumeration because it confirms — via the explicit spectral content of Theorem 6.2 — that there are no anisotropic refinement-stable survivors among candidates that *would* preserve $\eta_{\mu\nu}$ in some weakened sense. The eight-obstruction enumeration is therefore robust at the spectral level, even though (O6) overlaps logically with (O1)–(O5).

(O7) **Orientation-reversing extensions.** Any transport group containing improper ($\det \Lambda = -1$) elements is excluded by Lemma 6.3a (orientation preservation from K = 7 cell-orientation closure).

(O8) **Antichronous extensions.** Any transport group containing antichronous ($\Lambda^0_0 < -1$) elements is excluded by Lemma 6.3b (time-orientation preservation from irreversibility of commitment).

Assembly. The intersection of the surviving classes — non-degenerate (O2), of bilinear-form universality class η (O1), of signature (1, D) (O3, O5) with point-level coherence (O4), of O(1, D)-compatible (O6), proper (O7), and orthochronous (O8) — is precisely $SO(1, D)^\uparrow$.

No alternative admissible refinement-stable transport group exists: every group satisfying the substrate-level constraints lies in $SO(1, D)^\uparrow$, and $SO(1, D)^\uparrow$ satisfies all of them by Theorem 7.1.

Structural significance

Theorem 7.1 closes OP9 of the Microscopic Origin paper at the *existence-and-identification* level. Theorem 7.2 strengthens this to the *existence-uniqueness-identification* level: SO(1, D) is not merely *a* stable fixed point but *the* admissible refinement-stable transport group, with all alternative candidates explicitly excluded by structurally distinct substrate-level obstructions.

The structural inversion this completes: in the Microscopic Origin paper, the substrate-level edge transport operators U_{ij} were defined as elements of SO(1, D) — i.e. the Lorentz structure was *input* at the substrate level. In the present paper, U_{ij} are abstract elements of an abstract group $\mathcal{G}_{sub}$ defined purely by closure compatibility (Definition 3.1, with no reference to Lorentz structure), and the identification $\mathcal{G}_\infty \cong SO(1, D)$ is an *output* of the continuum limit, with uniqueness established by Theorem 7.2.

This is the substrate analogue, for the tangent-bundle / transport-group sector, of what Theorem 5.3 of the Microscopic Origin paper did for the refinement transport operator itself: an inheritance-level identification has been replaced by a substrate-level construction with the inheritance recovered as a derived — and uniquely-derived — consequence.

8. Connection Compatibility

We now confirm that the substrate-derived transport group is compatible with the Levi-Civita connection emergence of the Microscopic Origin paper Theorem 6.2, with the $\eta_{\mu\nu} \leftrightarrow g_{\mu\nu}$ compatibility now realised via an explicit soldering construction rather than inherited as a continuum-level identification.

Definition 8.0 — Refinement frame bundle and soldering

Let V_{sub} denote the abstract fibre on which $\mathcal{G}_\infty$ acts (the continuum-limit tangent fibre of the substrate transport bundle), equipped with the substrate-derived bilinear form η_{ab} of Theorem 4.1 / Theorem 5.2.

A **refinement frame** at a continuum point x of the emergent manifold is a linear isomorphism

$e^a{}_\mu(x) : T_x \to V_{sub}$,

mapping the continuum-limit tangent space T_x at x to the substrate-level fibre V_{sub}, with the frame determined by closure-compatible refinement up to a left SO(1, D) gauge action: two frames at the same point related by $\Lambda^a{}_b \in SO(1, D)$ are physically equivalent.

Construction of $e^a{}_\mu(x)$. The refinement frame is constructed as the continuum limit of substrate-level frames under closure-compatible refinement. At each refinement order n, every substrate vertex $v_i^{(n)}$ carries a substrate-level frame $e^a{}_\mu^{(n)} : T_{v_i^{(n)}} \to V_{sub}$ determined by the local edge transport operators $U_{ij}^{(n)}$ at $v_i^{(n)}$ — concretely, the substrate-level frame is the linear map identifying the local tangent structure at $v_i^{(n)}$ (the space of admissible commitment configurations) with the abstract fibre V_{sub}, with the identification fixed up to a $\mathcal{G}_{sub}^{(n)}$ gauge action. As $n \to \infty$, the substrate-level frames at the converging substrate-vertex sequence $v_i^{(n)} \to x$ converge in the closure-compatible refinement-Cauchy sense (parallel to Theorem 5.3 of the Microscopic Origin paper, applied to frames rather than to path holonomy) to a continuum-limit refinement frame $e^a{}_\mu(x)$.

Well-definedness up to SO(1, D) gauge is the substrate analogue of Theorem 7.1: the continuum-limit transport group $\mathcal{G}_\infty \cong SO(1, D)$ acts as the natural gauge group on the continuum-limit refinement frame, with two substrate-vertex sequences converging to the same continuum point yielding refinement frames related by an SO(1, D) gauge transformation. The continuum-limit refinement frame bundle is therefore the principal SO(1, D)-bundle whose fibres are the continuum-limit refinement-frame gauge orbits.

The collection of refinement frames at all continuum points forms the **refinement frame bundle** $F \to \mathscr{M}$, a principal SO(1, D)-bundle over the emergent continuum manifold $\mathscr{M}$. Closure-compatible refinement transport defines a connection on F whose first-order generator is $\Gamma^\alpha{}_{\beta\mu}$ of the Microscopic Origin paper Definition 3.4.

The **soldering of η to g** is the construction

$$g_{\mu\nu}(x) := e^a_\mu(x)\, e^b_\nu(x)\, \eta_{ab},$$

pulling the substrate-derived bilinear form η_{ab} on V_{sub} back to a continuum-limit metric $g_{\mu\nu}(x)$ on T_x. The metric $g_{\mu\nu}$ is independent of the choice of refinement frame within an SO(1, D)-gauge orbit, by η_{ab} being SO(1, D)-invariant.

The **vielbein compatibility condition** under refinement-compatible transport is

$$\nabla_\mu e^a_\nu(x) = 0,$$

i.e. the refinement frame is parallel-transported by the substrate-derived Levi-Civita connection — which is the substrate-level statement that the refinement frame bundle is the natural soldered tangent bundle of the emergent geometry.

Remark — what the soldering achieves

Definition 8.0 upgrades the $\eta_{\mu\nu} \leftrightarrow g_{\mu\nu}$ compatibility from an *inherited identification* (the Microscopic Origin paper's Theorem 6.2 Step 3 framing) to an *explicit construction*: $g_{\mu\nu}$ is now *induced* from η_{ab} via the refinement-frame soldering, with no Lorentz-structure inheritance required at the continuum level.

The structural pay-off:

- **Substrate transport bundle V_{sub} → continuum tangent bundle T $\mathscr{M}$** via the refinement frame e^a_μ. The substrate transport geometry is the *internal* geometry; the continuum geometry is the *external* geometry; the soldering connects them.
- **η is internal, g is external.** η_{ab} is the bilinear form on the substrate fibre V_{sub} (the transport-internal geometry); $g_{\mu\nu}$ is the bilinear form on the continuum tangent space T_x (the emergent external geometry). They are related by pull-back, not by inheritance.
- **Metric compatibility** $\nabla_\mu g_{\nu\rho} = 0$ follows by direct computation: $\nabla_\mu g_{\nu\rho} = \nabla_\mu(e^a_\nu e^b_\rho \eta_{ab}) = 0$ by $\nabla_\mu e^a_\nu = 0$ (vielbein compatibility) and η_{ab} constant (no derivative in V_{sub}).

This is the cleanest geometric realisation of the substrate-emergent structure: continuum metric geometry is *induced from* substrate transport geometry through refinement-frame soldering, not merely *compatible with* it.

Theorem 8.1 — Substrate-derived Levi-Civita emergence [proven, conditional on Theorem 7.1, Definition 8.0 soldering, and Theorem 6.2 of the Microscopic Origin paper]

The first-order generator of refinement transport,

$$\Gamma^\alpha_{\beta\mu} := -\lim_{\varepsilon \to 0} \left[(T_{(\varepsilon e_\mu)})^\alpha_\beta - \delta^\alpha_\beta \right] / \varepsilon,$$

coincides with the Levi-Civita connection of the soldered continuum metric $g_{\mu\nu} = e^a{}_\mu e^b{}_\nu \eta_{ab}$ of Definition 8.0, with both metric-compatibility and torsion-freeness following from substrate-derived structures.

Proof

Three steps, now invoking the explicit soldering construction of Definition 8.0.

Step 1 — soldered metric is well-defined and Lorentzian.

By Theorem 7.1, $T_\gamma \in SO(1, D)$; by Theorem 4.1 / Theorem 5.2, η_{ab} is the substrate-derived bilinear form on V_{sub} with Lorentzian signature (1, D). By Definition 8.0, the soldered metric

$$g_{\mu\nu}(x) := e^a{}_\mu(x)\, e^b{}_\nu(x)\, \eta_{ab}$$

is well-defined on each tangent space T_x (independent of refinement-frame choice within SO(1, D) gauge orbit) and inherits Lorentzian signature (1, D) from η_{ab}. The identification $g_{\mu\nu} \leftrightarrow$ inherited continuum metric of the Lorentzian Completion paper is then a layered-consistency check (Theorem 8.2 of the Microscopic Origin paper), with both metrics now produced from substrate-level inputs.

Step 2 — metric compatibility from η-preservation and vielbein compatibility.

By Theorem 7.1, the continuum-limit refinement transport operators $T_\gamma \in SO(1, D)$ preserve η_{ab}: $\eta_{ab}(T_\gamma u, T_\gamma v) = \eta_{ab}(u, v)$ for all $u, v \in V_{sub}$. Equivalently, T_γ acts on V_{sub} as an η-isometry at every continuum path γ.

The first-order generator $\Gamma^\alpha{}_{\beta\mu}$ of T_γ — equivalently, the connection on the refinement frame bundle F — satisfies the vielbein compatibility condition $\nabla_\mu e^a{}_\nu = 0$ of Definition 8.0 by construction: refinement-compatible transport preserves the refinement frame up to SO(1, D) gauge, which is exactly the statement that the frame is parallel-transported by the substrate-derived connection.

Therefore:

$$\nabla_\mu g_{\nu\rho} = \nabla_\mu(e^a{}_\nu e^b{}_\rho \eta_{ab}) = (\nabla_\mu e^a{}_\nu) e^b{}_\rho \eta_{ab} + e^a{}_\nu (\nabla_\mu e^b{}_\rho) \eta_{ab} + e^a{}_\nu e^b{}_\rho (\partial_\mu \eta_{ab}) = 0,$$

with the first two terms vanishing by $\nabla_\mu e^a{}_\nu = 0$ and the last by η_{ab} being constant on V_{sub} (it carries no μ-dependence as a substrate-internal object). Metric compatibility $\nabla_\mu g_{\nu\rho} = 0$ is therefore a *direct calculational consequence* of the substrate construction, not an inheritance.

Step 3 — torsion-freeness.

By Lemma 6.3 of the Microscopic Origin paper, BCB triangle closure C_4 at the substrate level forces $\Gamma^\alpha_{[\beta\mu]} = 0$ — torsion-freeness of the limit connection. This is inherited from the Microscopic Origin paper unchanged.

Assembling Steps 1–3: $\Gamma^\alpha_{\beta\mu}$ is the unique torsion-free metric-compatible connection on $(g_{\mu\nu}, \nabla)$ — the Levi-Civita connection of the soldered continuum geometry. Both the metric (Step 1) and the metric-compatibility (Step 2) are now substrate-derived via the refinement-frame soldering of Definition 8.0, with no Lorentz-structure inheritance invoked.

Structural significance

Theorem 8.1 closes the substrate-level Levi-Civita emergence at the deepest layer reached so far: the metric $g_{\mu\nu}$, the connection $\Gamma^\alpha_{\beta\mu}$, the metric-compatibility $\nabla_\mu g_{\nu\rho} = 0$, *and the torsion-freeness* $\Gamma^\alpha_{[\beta\mu]} = 0$ are all substrate-derived via the refinement-frame soldering construction. The continuum geometric structure is *induced from* substrate transport geometry, not *inherited from* a continuum-level Lorentzian metric.

The Microscopic Origin paper's Theorem 6.2 Step 3 (metric compatibility) was stated as "inherited from the Lorentzian Completion paper: refinement transport at every refinement order preserves the inner product on the inherited Lorentzian geometry under the BCB-admissibility constraint $U_{ij} \in \mathcal{G} = SO(1, D)$." The substrate inheritance acknowledged there is now closed at two levels:

1. **Bilinear-form substrate origin (Theorem 4.1)**: η_{ab} is constructed from BCB bit-conservation + finite distinguishability + closure-compatible refinement, not inherited from the Lorentzian Completion paper.
2. **Soldering substrate origin (Definition 8.0 / Theorem 8.1)**: $g_{\mu\nu}$ is constructed by pull-back from η_{ab} via the refinement frame bundle, not inherited as a continuum-level Lorentzian metric.

This is the strictly-strengthened version of Microscopic Origin paper Theorem 6.2 Step 3 noted in the §10 strictly-strengthened-results subsection below.

---

9. Microscopic Lorentzian Recovery

Theorem 9.1 — Microscopic Lorentzian recovery [proven, conditional on Theorems 4.1, 5.2, 6.2, 7.1, 7.2, 8.1, Lemmas 5.0, 5.1, Definition 8.0 soldering, and the cascade theorems of the Microscopic Origin and Einstein–Hilbert Emergence papers]

The Lorentzian geometric structure of the continuum-limit VERSF gravitational sector emerges as the stable refinement fixed point of closure-compatible irreversible commitment transport on the $K = 7$ simplicial substrate, with every component of the Lorentzian geometry — the metric

$\eta_{\mu\nu}$ / $g_{\mu\nu}$, the transport group SO(1, D), the causal cone structure, the Levi-Civita connection, the Riemann curvature, and the Einstein–Hilbert action — substrate-derived from K = 7 / BCB / CCC / refinement-flow inputs.

Proof

Direct assembly across the substrate construction chain.

Bilinear form $\eta_{\mu\nu}$. By Theorem 4.1, BCB bit-conservation at the substrate level + refinement-Cauchy convergence produces a unique-up-to-scale bilinear form on tangent fibres in the continuum limit.

Signature. By Theorem 5.2, CCC finite-speed propagation + single-direction commitment orientation + finite distinguishability of orthogonal directions force signature (1, D) on $\eta_{\mu\nu}$.

Transport group. By Theorems 6.2, 7.1, and Corollary 6.3, refinement suppression of anisotropies + closure-compatible preservation of orientation and time-orientation force the continuum-limit transport group to be $SO(1, D)\uparrow$.

Levi-Civita connection. By Theorem 8.1 (combined with Theorem 6.2 of the Microscopic Origin paper), the first-order generator of refinement transport is the Levi-Civita connection of ($g_{\mu\nu}$, ∇), with metric compatibility now substrate-derived from $\eta_{\mu\nu}$ preservation and torsion-freeness inherited from BCB triangle closure C_4.

Riemann curvature and Einstein–Hilbert. By Theorem 7.2 (discrete plaquette curvature → Riemann tensor) and Theorem 9.1 (microscopic Einstein–Hilbert recovery) of the Microscopic Origin paper, the substrate-derived connection produces the Riemann tensor and the Einstein–Hilbert action with substrate origins on the K = 7 commitment foam.

Causal cone. By Theorem 8.2 (Lorentzian compatibility) of the Microscopic Origin paper, the commitment-order causal structure on the K = 7 foam is consistent with the inherited Lorentzian causal cone of $g_{\mu\nu}$. The present paper strengthens this: with $\eta_{\mu\nu}$ now substrate-derived, the causal cone is itself derivable from the substrate ($\eta_{\mu\nu} V^\mu V^\nu = 0$ defines the null cone), not merely compatible with an inherited continuum-limit cone.

The full Lorentzian geometric structure — metric, transport group, causal cone, connection, curvature, action — is therefore the stable refinement fixed point of closure-compatible irreversible commitment transport on the K = 7 simplicial substrate. Each component has an explicit substrate origin; none is inherited as primitive geometric structure at the substrate layer.

Structural significance

Theorem 9.1 completes the substrate construction of the Lorentzian geometric sector of the VERSF programme. Combined with the Microscopic Origin paper's microscopic GR recovery (its Theorem 11.1), the present Theorem 9.1 establishes that the entire geometric structure of continuum gravity is the continuum-limit shadow of K = 7 commitment-foam combinatorics —

with the Lorentz transport group, previously the residual inheritance of the geometric programme, now derived alongside every other geometric object.

---

10. Structural Interpretation

The present paper completes the inversion of the geometric side of the VERSF programme at the transport-symmetry layer. The Lorentz group is now no longer a primitive geometric symmetry at any layer of the construction:

- *Variational layer* (variational paper Theorem 5.1): $\mathcal{S}$_EH is the unique admissible action within the continuum-limit operator basis. Lorentz structure enters through the continuum-limit operator basis.
- *Continuum-emergence layer* (Einstein–Hilbert Emergence paper cascade): the continuum-limit operator basis — Γ, R^α_{βμν}, R, √|g| — emerges from refinement transport geometry. Lorentz structure enters through the transport group SO(1, D) preserving the metric.
- *Microscopic substrate layer* (Microscopic Origin paper): the refinement transport operator T_γ is constructed from K = 7 commitment-foam combinatorics. Lorentz structure enters through the substrate-level edge transport group $\mathcal{G}$ = SO(1, D) (recorded as OP9).
- *Substrate symmetry-group layer (present paper)*: the transport group itself, $\mathcal{G}$_∞ ≅ SO(1, D), is the stable refinement fixed point of the closure-compatible refinement flow on the abstract substrate transport group $\mathcal{G}$_sub defined purely combinatorially by C_1–C_7 closure constraints. **No Lorentz structure is inherited at any layer below this.**

The geometric programme of VERSF is therefore now structurally closed at the symmetry-group layer for the inheritance gap recorded as OP9 of the Microscopic Origin paper.

Derivational chain through the present paper

Substrate Layer	Continuum-Emergent Quantity	Theorem
Abstract group $\mathcal{G}$_sub from C_1–C_7 closure	Continuum-limit transport group $\mathcal{G}$_∞	Definitions 3.1–3.3
BCB bit-conservation + finite distinguishability	Refinement-stable bilinear form η_μν	Theorem 4.1
Finite distinguishability	Nondegeneracy of η_μν	Lemma 4.2
C_5–C_7 bilinear-form-sector rigidity	Uniqueness up to scale of η_μν	Lemma 4.3
Closure-compatible commitment orientation	One-dimensional timelike direction per point	Lemma 5.0
Closure-compatible refinement requires unique continuum-limit causal ordering	Multi-timelike refinement instability for $p \geq 2$	Lemma 5.1

Substrate Layer	Continuum-Emergent Quantity	Theorem
CCC + irreversibility-as-consumption + Lemmas 5.0, 5.1	Lorentzian signature (1, D)	Theorem 5.2
Anisotropy mode space + closure-compatible refinement	Anisotropy refinement operator $\mathscr{R}$_aniso	Definition 6.2
C_5–C_7 refinement-cell rigidity	Spectral gap ρ($\mathscr{R}$_aniso) < 1, mode decomposition $\mathcal{A} = \mathcal{A}$^Lor ⊕ $\mathcal{A}$^aniso	Theorem 6.2
K = 7 cell-orientation closure	det Λ = +1 (orientation preservation)	Lemma 6.3a
Irreversibility of commitment	$\Lambda^0{}_0 \geq 1$ (time-orientation preservation)	Lemma 6.3b
Combined obstructions	SO(1, D)↑ fixed point	Corollary 6.3
Bilinear + signature + stability assembly	$\mathcal{G}$_∞ ≅ SO(1, D)	Theorem 7.1
Eight obstruction classes (O1)–(O8)	Uniqueness of SO(1, D)↑	Theorem 7.2
Refinement frame bundle + soldering g = e e η	Substrate-induced continuum metric	Definition 8.0
η-preservation + vielbein compatibility ∇e = 0	Metric compatibility ∇_μ g_νρ = 0	Theorem 8.1
Full substrate cascade	Microscopic Lorentzian recovery	Theorem 9.1

Cross-strand correspondence with the foundational programme

K = 7 / BCB / TPB / CCC Substrate	Continuum Transport-Geometry Output
C_1–C_7 closure constraints	Abstract transport group $\mathcal{G}$_sub
BCB bit-conservation	Invariant bilinear form η_μν
CCC finite-speed propagation	Indefinite signature
Single-direction commitment orientation	Exactly one negative eigenvalue
Finite distinguishability of orthogonal directions	Exactly D positive eigenvalues
K = 7 / TPB cell-subdivision contraction	Refinement suppression rate η
Closure orientation preservation	SO(1, D)↑ subgroup selection
Combined substrate inputs	$\mathcal{G}$_∞ ≅ SO(1, D)

In each row, the right-hand entry is the continuum-limit consequence of the left-hand substrate structure under closure-compatible refinement. The Lorentz transport group is therefore now the *output* of the substrate construction, not an *input*.

Layered consistency with the Microscopic Origin paper

The relationship between the present paper and the Microscopic Origin paper is layered consistency in the sense of Proposition 8.1 of the Einstein–Hilbert Emergence paper or Theorem 8.2 of the Microscopic Origin paper. The Microscopic Origin paper inherits $\mathcal{G}$ = SO(1, D) at Definition 3.1; the present paper derives $\mathcal{G}$_∞ ≅ SO(1, D) at Theorem 7.1. Both papers refer to the same continuum-limit transport group, with the present paper supplying the substrate construction that the Microscopic Origin paper inherited as input architecture.

This is the standard VERSF-programme pattern: a continuum-level inheritance is closed by a substrate-level construction one layer down, with consistency between the two layers as the structural quality-check.

Strictly-strengthened results vs. relocated inheritances

The present paper produces *three structurally distinct kinds* of result against the Microscopic Origin paper:

Type A — substrate-level constructions (relocated inheritances). Most theorems supply a substrate construction of an object that paper inherited as continuum-level input: Theorem 4.1 (bilinear form $\eta_{\mu\nu}$ from BCB), Theorem 5.2 (Lorentzian signature from CCC + ontology), Theorem 6.2 (spectral suppression of anisotropies), Theorem 7.1 (transport group $\mathcal{G}_\infty \cong SO(1, D)$). Each closes an inheritance acknowledgement of the Microscopic Origin paper by supplying the missing substrate construction.

Type B — strictly-strengthened theorems. Theorem 8.1 is structurally different: it is a *strictly strengthened version* of a Microscopic Origin paper theorem. The Microscopic Origin paper Theorem 6.2 Step 3 stated metric compatibility as *inherited from the Lorentzian Completion paper*, conditional on the inheritance $\mathcal{G} = SO(1, D)$. Theorem 8.1 closes this inheritance at two levels: $\eta_{\mu\nu}$ is substrate-derived (Theorem 4.1) and $g_{\mu\nu}$ is induced via the refinement-frame soldering (Definition 8.0); metric compatibility then follows by direct calculation, with no Lorentz inheritance invoked at any layer.

This is the structurally cleanest result of the present paper. It *removes* an inheritance entirely from the chain rather than relocating it one layer down.

Type C — new structural results. Three results of the present paper close the structural argument beyond mere substrate-relocation of inheritances:

- **Lemma 5.1 (multi-timelike refinement instability)** replaces any ontology-led argument for "at most one negative eigenvalue" with a refinement-architectural theorem: ultra-hyperbolic signatures are *refinement-incompatible*, not merely ontologically peculiar. This elevates the single-time-dimension constraint from an ontological inheritance to a structural necessity of the refinement framework.
- **Theorem 7.2 (fixed-point uniqueness)** strengthens Theorem 7.1's identification statement to a uniqueness statement: $SO(1, D)^\uparrow$ is not merely *a* stable fixed point but *the* admissible refinement-stable transport group, with eight structurally distinct obstruction classes (O1)–(O8) excluding every alternative candidate.
- **Definition 8.0 (refinement frame bundle and soldering)** upgrades the $\eta_{\mu\nu} \leftrightarrow g_{\mu\nu}$ compatibility from an *inherited identification* to an *explicit construction*: $g_{\mu\nu}$ is now *induced from* η_{ab} via the refinement-frame soldering $g_{\mu\nu} = e^a{}_\mu e^b{}_\nu \eta_{ab}$, not

merely *compatible with* it. The substrate transport bundle V_{sub} and the continuum tangent bundle $T\mathscr{M}$ are now connected by an explicit soldering map, not by inheritance.

The three Type C results upgrade the paper's structural maturity at the symmetry-group layer: alternative signatures become refinement-inconsistent (not merely ontologically forbidden), SO(1, D) is uniquely the dominant fixed-point eigenspace (not merely a stable configuration), and continuum metric geometry is induced from substrate transport geometry (not merely compatible with it).

---

11. Open Problems

OP numbering convention. The OP numbers in this section are local to the present paper. Where an OP corresponds to a problem recorded in an inherited paper under a different number, the cross-walk is given inline. The local numbering preserves cross-paper coherence while keeping §11 internally readable.

The paper does **not** close:

- the closure-normalisation factor C_λ,
- the form and amplitude of the memory kernel $\mathcal{K}(x, x')$,
- Standard-Model matter coupling,
- quantum fluctuations of T_γ or of the commitment foam itself,
- subleading higher-curvature corrections,
- the quantitative one-loop suppression rate η in Theorem 6.2,
- the overall scale of $\eta_{\mu\nu}$ in Theorem 4.1,
- a fully non-abelian Wilson-style derivation paralleling the K = 7 Wilson Limit paper's U(1) template,
- foam dynamics and global existence,
- variational microscopic completion of the non-geometric sectors.

Each is recorded below.

OP1 — Closure-normalisation factor C_λ

The substrate origin of the closure-normalisation factor C_λ in $\kappa_{eff} = 8\pi C_\lambda \hbar\xi^2/c^3$ remains open. Inherited unchanged from Microscopic Origin paper OP1 / Einstein–Hilbert Emergence paper OP1 / variational paper OP6 / dynamical paper OP6.

OP2 — Form and amplitude of the memory kernel

The form, decay law, and amplitude of the bilocal memory kernel $\mathcal{K}(x, x')$ remain open. Inherited unchanged from Microscopic Origin paper OP2 / variational paper OP3.

OP3 — Anisotropic Wilson coefficients

The substrate origin of the anisotropic transport-curvature Wilson coefficients ($\alpha_{\hat{Q}}$, $\beta_{\hat{Q}}$, $\gamma_{\hat{Q}}$) of variational paper §6 remains open. Inherited unchanged from Microscopic Origin paper OP3 / variational paper OP4.

OP4 — Higher-order refinement corrections

Subleading refinement orders generate higher-curvature corrections (Gauss–Bonnet at second-order Lovelock, R^2, $R_{\mu\nu} R^{\mu\nu}$, $R_{\mu\nu\rho\sigma} R^{\mu\nu\rho\sigma}$ at next derivative order). Substrate structural form open. Inherited from Microscopic Origin paper OP4 / Einstein–Hilbert Emergence paper OP3, lifted to the transport-symmetry layer here.

Structural-pattern remark. The present paper exhibits the same second-order closure pattern noted in Microscopic Origin paper OP4: leading-order refinement (C_4 triangle closure at order ε, bit-conservation at order ε^2, refinement suppression at order $e^{-\eta n}$) identifies the admissible SO(1, D) structure, while higher-order corrections from subleading refinement orders may be needed for full physical closure. The two cautions of K = 7 Wilson Limit paper §11.8 apply unchanged: "RG-order" and "closure-cycle" senses of "second order" are mathematically distinct, and the structural resonance is interpretive rather than theorematic.

OP5 — Matter coupling

Standard-Model matter remains not substrate-derived. Inherited unchanged from Microscopic Origin paper OP5 / Einstein–Hilbert Emergence paper OP5 / variational paper OP7 / source paper OP7.

OP6 — Quantum completion of refinement transport

Quantum fluctuations of T_γ, of the substrate transport group $\mathcal{G}_{sub}$ itself, and of the foam combinatorics remain open. Inherited from Microscopic Origin paper OP6 / Einstein–Hilbert Emergence paper OP6 / variational paper OP8, with the additional present-paper question: what is the quantum substrate transport group whose classical limit is the SO(1, D) of Theorem 7.1?

OP7 — Substrate derivation of the suppression rate η

Theorem 6.2 establishes the existence of an exponential refinement-suppression rate $\eta > 0$ for transport anisotropies. The quantitative value of η as an explicit function of the K = 7 bare-coupling integer $N_{loop} = 14$ (and other foundational substrate parameters) is not derived here; only its positivity. Closing OP7 would mean computing η explicitly via the non-abelian extension of the K = 7 Wilson Limit paper's one-loop matching analysis, with the abelian U(1) one-loop coefficient $\beta_{K=7} = 2^7 \cdot 15/14$ as the methodological starting point.

This is the quantitative Wilson-style analogue of Theorem 6.2 for the SO(1, D) tangent-bundle sector. OP7 is structurally parallel to the quantitative content of the K = 7 Wilson Limit paper §§5–10, lifted to the non-abelian case.

OP8 — Substrate derivation of the bilinear-form scale

Theorem 4.1 establishes $\eta_{\mu\nu}$ as unique up to overall scale. The substrate origin of the overall scale of $\eta_{\mu\nu}$ — equivalently, the substrate origin of the CCC coherence scale ξ inherited from the Lorentzian Completion paper — is not derived here. Closing OP8 would yield a substrate-level derivation of the coherence scale, paralleling the substrate derivation of $\kappa_{eff} = 8\pi C_\lambda \hbar\xi^2/c^3$ (OP1) but for the metric scale rather than the gravitational coupling scale.

OP9 — Non-abelian Wilson-style derivation

The K = 7 Wilson Limit paper supplies a worked U(1) gauge-bundle template for the structurally parallel transport-group inheritance problem, with explicit one-loop matching giving the bare coupling $\beta_{K=7} = 2^7 \cdot 15/14$ from $N_{loop} = 14$. The present paper achieves the existence-uniqueness-identification result $\mathcal{G}_\infty \cong SO(1, D)$ without performing the analogous non-abelian one-loop matching. Closing OP9 would mean extending the Wilson methodology to the non-abelian SO(1, D) tangent-bundle case, computing the analogous bare-coupling expression and the corresponding suppression rate.

The technical challenge: SO(1, D) is non-abelian and (D + 1)(D + 2)/2-dimensional (ten generators at D + 1 = 4), so the abelian U(1) one-loop machinery of the Wilson paper does not transfer directly. The natural framing: develop a non-abelian Wilson construction at the substrate scale, with one-loop matching across the K = 7 cell structure on the SO(1, D) tangent bundle. This is genuinely substrate-level dynamical content, distinct from the kinematic existence-uniqueness-identification result of Theorems 7.1 and 7.2 here.

This OP closes the previous-paper OP9 only at the *kinematic existence-uniqueness-identification* level. The full *quantitative dynamical-level* closure — with explicit numerical rates derived from K = 7 bare-coupling integers via non-abelian Wilson methodology — is the natural successor problem.

OP10 — Foam dynamics and global existence

The foam-level dynamical structure under which closure-compatible refinement sequences emerge as a derived rather than postulated property remains open. Inherited unchanged from Microscopic Origin paper OP7.

OP11 — Variational microscopic completion

The full microscopic substrate construction of every sector of $\mathcal{S}_{VERSF}$ on the K = 7 foam remains open. Inherited unchanged from Microscopic Origin paper OP8.

OP12 — Distinguishability-consumption articulation of irreversibility

Theorem 5.2 Step 1 articulates the substrate-level mechanism producing indefinite signature on the BCB bilinear pairing as *irreversibility-as-distinguishability-consumption*: commitment-propagation directions consume bit-resource (negative-sign contribution); spatial directions preserve bit-resource (positive-sign contribution); the opposite-sign structure is built into the BCB bilinear pairing rather than imposed externally.

This articulation is treated in the present paper as an interpretation of the source-paper §2.1 irreversibility ontology — structurally implicit there but not, to the author's knowledge, isolated as a separate foundational theorem. Closing OP12 would mean elevating distinguishability-consumption from an interpretive articulation of the irreversibility ontology to a separately-derived foundational theorem: a foundational-layer result establishing that commitment-propagation and spatial bit-counting contribute with opposite signs to the BCB bilinear pairing as a consequence of the irreversible-commitment ontology, independent of the indefinite-signature continuum-limit consequence.

This is a foundational-layer problem rather than a residual of the present paper: the articulation is load-bearing for Theorem 5.2 Step 1, and future readers (especially referees of successor papers) should know that this is one of the substrate-level mechanisms cited as inherited articulation rather than separately-proven foundational theorem. OP12 is recorded separately from OP11 (variational microscopic completion) because it concerns the *substrate-ontology* layer — the irreversibility-of-commitment ontology of the source paper — rather than the *variational-completion* layer of the higher-level action structure.

OP cross-walk to inherited papers

For readers cross-referencing the OPs of the present paper against those of the predecessor papers, the correspondence is:

Present paper	**Microscopic Origin paper**	**Higher-layer cross-references**
OP1 (C_λ)	OP1 (C_λ)	Einstein–Hilbert Emergence OP1, variational OP6, dynamical OP6
OP2 (memory kernel form/amplitude)	OP2	variational OP3
OP3 (anisotropic Wilson coefficients)	OP3	variational OP4
OP4 (higher-curvature corrections)	OP4	Einstein–Hilbert Emergence OP3
OP5 (matter coupling)	OP5	Einstein–Hilbert Emergence OP5, variational OP7, source OP7
OP6 (quantum completion)	OP6	Einstein–Hilbert Emergence OP6, variational OP8
OP7 (substrate derivation of η)	—	new (quantitative SO(1, D) Wilson analogue)
OP8 (substrate derivation of bilinear-form scale)	—	new (substrate origin of CCC scale ξ)

Present paper	Microscopic Origin paper	Higher-layer cross-references
OP9 (non-abelian Wilson-style derivation)	—	new (closes Microscopic Origin OP9 at identification level only)
OP10 (foam dynamics)	OP7	—
OP11 (variational microscopic completion)	OP8	—
OP12 (distinguishability-consumption articulation)	—	new (foundational source paper §2.1 layer)

The present paper's OP9 has a different content from the Microscopic Origin paper's OP9: the Microscopic Origin paper's OP9 (substrate derivation of $\mathcal{G}$) is closed at the existence-uniqueness-identification level by Theorems 7.1 and 7.2 of the present paper, while the present paper's OP9 (non-abelian Wilson-style derivation) is the quantitative-rate residual problem left open by the present paper's identification-level closure. The numerical coincidence "OP9" reflects local OP numbering and does not indicate a single open problem persisting across papers.

The three OPs new to the present paper — OP7, OP8, OP9 (quantitative Wilson extension), and OP12 — concern, respectively: the quantitative spectral rate, the metric scale, the non-abelian Wilson methodology, and the substrate-ontology articulation of distinguishability-consumption. Each is structurally distinct from the inherited residual problems and each has a clear successor-paper home.

---

12. Conclusion

The previous VERSF papers established:

- the K = 7 simplicial closure architecture (foundational papers),
- BCB / TPB and finite distinguishability,
- irreversible commitment ontology,
- Lorentzian causal completion (Lorentzian Completion paper),
- refinement-stable transport geometry,
- effective stress–energy structure,
- Einstein-compatible dynamics,
- a unified covariant action $\mathcal{S}$_VERSF with action-level GR recovery,
- substrate emergence of the Einstein–Hilbert action,
- substrate construction of refinement transport T_γ (Microscopic Origin paper).

The present paper closes the residual transport-group inheritance problem recorded as OP9 of the Microscopic Origin paper.

The central results are:

- **Theorem 4.1.** Closure-compatible refinement transport on the $K = 7$ commitment foam preserves, in the continuum limit, an invariant bilinear form $\eta_{\mu\nu}$ unique up to overall scale (Lemma 4.3). BCB bit-conservation lifted to the bilinear sector via pairwise distinguishability conservation (Step 1) supplies the substrate-level mechanism.
- **Lemma 5.1.** Ultra-hyperbolic signatures (p, q) with $p \geq 2$ generate non-equivalent refinement orderings under closure-compatible refinement. The Lorentzian signature constraint $p = 1$ is therefore a *structural necessity* of the refinement architecture, not merely an ontological inheritance.
- **Theorem 5.2.** The combined inputs — CCC + irreversibility-as-distinguishability-consumption (indefinite signature), multi-timelike refinement instability with point-level coherence (at most one negative eigenvalue), finite distinguishability of orthogonal directions (exactly D positive eigenvalues) — force $\eta_{\mu\nu}$ to have Lorentzian signature (1, D).
- **Theorem 6.2.** The anisotropy refinement operator $\mathscr{R}_{aniso}$ has spectral radius $\rho(\mathscr{R}_{aniso}) = \lambda_{aniso} < 1$, with strict contraction forced by foundational C_5–C_7 refinement-cell rigidity. The continuum-limit anisotropy mode space admits the spectral decomposition $\mathcal{A} = \mathcal{A}^{Lor} \oplus \mathcal{A}^{aniso}$, with Lorentz-compatible modes refinement-invariant (eigenvalue 1) and anisotropic modes exponentially suppressed (eigenvalues < 1).
- **Theorems 7.1 and 7.2.** The continuum-limit admissible transport group is

  $\mathcal{G}_\infty \cong SO(1, D)$,

  with continuum refinement transport operators $T_\gamma \in SO(1, D)$. Moreover, $SO(1, D)^\uparrow$ is the *unique* admissible refinement-stable transport group: every alternative candidate is excluded by one of eight structurally distinct substrate-level obstructions (O1)–(O8).

- **Definition 8.0 and Theorem 8.1.** Continuum metric geometry is *induced from* substrate transport geometry via the refinement-frame soldering $g_{\mu\nu} = e^a_\mu e^b_\nu \eta_{ab}$ on a principal SO(1, D)-bundle $F \to \mathscr{M}$. Metric compatibility $\nabla_\mu g_{\nu\rho} = 0$ follows from η-preservation under SO(1, D) and vielbein compatibility $\nabla_\mu e^a_\nu = 0$; *torsion-freeness* $\Gamma^\alpha_{[\beta\mu]} = 0$ follows from BCB triangle closure C_4. The first-order generator of refinement transport is the Levi-Civita connection of the soldered continuum geometry, with no Lorentz-structure inheritance invoked.
- **Theorem 9.1.** The Lorentzian geometric structure of the continuum-limit VERSF gravitational sector emerges as the stable refinement fixed point of closure-compatible irreversible commitment transport on the $K = 7$ simplicial substrate.

The result therefore shifts the standing of the Lorentz transport group in the VERSF framework from

"the inherited continuum-limit transport group preserving the inherited Lorentzian metric"

— that was the framing through the Microscopic Origin paper —

to

"the unique refinement-stable fixed point of the closure-compatible refinement flow on the abstract substrate transport group $\mathcal{G}$_sub defined purely combinatorially by C_1–C_7 closure constraints, with bilinear preservation from BCB bit-conservation lifted via pairwise distinguishability conservation, signature from CCC + irreversibility-as-consumption + multi-timelike refinement instability, spectral stability from C_5–C_7 refinement-cell rigidity giving SO(1, D)-compatibility as the dominant fixed-point eigenspace $\mathscr{A}$^{Lor} of $\mathscr{R}$_aniso, uniqueness from the eight obstruction classes (O1)–(O8), and continuum metric induced from substrate transport via refinement-frame soldering g_μν = e^a_μ e^b_ν η_ab on a principal SO(1, D)-bundle."

— which is the present paper. The two framings are layered consistently: the inherited continuum-limit SO(1, D) of the Microscopic Origin paper and the substrate-derived SO(1, D) of the present paper are the same group, with the present paper supplying the substrate construction that the Microscopic Origin paper inherited.

Honest limitation. The present paper closes OP9 of the Microscopic Origin paper at the existence-uniqueness-identification level: $\mathcal{G}$_∞ ≅ SO(1, D) is now derived from K = 7 / BCB / CCC / substrate-ontology / refinement-flow inputs rather than imposed. What it does *not* close: the quantitative one-loop suppression rate η as an explicit function of N_loop = 14 (OP7); the overall scale of η_μν (OP8); a fully non-abelian Wilson-style dynamical derivation paralleling the K = 7 Wilson Limit paper's U(1) template (OP9); the distinguishability-consumption articulation of irreversibility (OP12, foundational-layer); and the inherited residual problems C_λ (OP1), memory kernel form (OP2), anisotropic Wilson coefficients (OP3), higher-curvature corrections (OP4), matter coupling (OP5), quantum completion (OP6), foam dynamics (OP10), and variational microscopic completion (OP11). Each is sharply isolated, each has a clear successor-paper home, and none remains a structural existence question at the symmetry-group layer.

The structural progression of the geometric side of the VERSF programme is now:

commitment ontology → K = 7 simplicial closure architecture → BCB / TPB closure structure → CCC framework → abstract substrate transport group $\mathcal{G}$_sub → bit-conservation bilinear preservation → hyperbolic signature selection → refinement-flow suppression of anisotropies → SO(1, D) refinement fixed point → continuum refinement transport → continuum connection → discrete plaquette curvature → Riemann emergence → Lorentzian compatibility → Einstein–Hilbert recovery → memory kernel existence → microscopic GR recovery.

The geometric side is now structurally closed at the substrate symmetry-group layer for the transport-group inheritance problem specifically, with the quantitative dynamical content (OP7 / OP8 / OP9) and the deeper inherited problems (OP1 / OP2 / OP3 / OP4 / OP5 / OP6 / OP10 / OP11) explicitly recorded as remaining.

The picture this leaves is: the Lorentz group is the stable transport symmetry of closure-compatible irreversible commitment transport on the K = 7 simplicial substrate. SO(1, D) is no longer a primitive geometric symmetry inherited from continuum physics; it is the coarse-grained refinement-limit shadow of the K = 7 closure architecture, with BCB bit-conservation supplying the bilinear-preservation mechanism, the CCC framework and substrate ontology supplying the signature selection, and the K = 7 / TPB refinement-flow contraction supplying the stability that selects $SO(1, D)\uparrow$ as the unique refinement fixed point.

---

13. References to Inherited VERSF Papers

The present paper carries inline citations to several earlier VERSF papers. For a stand-alone read, each is identified below by title and the specific results invoked.

- **Foundational Papers** — *K = 7 Simplicial Closure Architecture, BCB / TPB, and the Hexagonal Substrate.* Inherited: the K = 7 closure architecture (six adjacency-coding boundary vertices plus one gauge-anchoring central vertex per cell, fourteen oriented closure traversals $N_{loop} = 14$), the constraint decomposition C_1–C_3 (edge consistency) + C_4 (triangle / fold completion) + C_5–C_7 (hexagonal embedding) with C_5–C_7 supplying the rigidity argument for one-parameter bilinear form uniqueness (Lemma 4.3 / Theorem 4.1 Step 3), Bit Conservation and Balance (BCB) supplying the substrate-level bit-count form preserved by closure-compatible transport (Theorem 4.1 Step 1), Ticks-Per-Bit (TPB) supplying the geometric refinement scaling $\varepsilon_n \leq C \lambda^n$, and the refinement-flow contraction structure underlying Theorem 6.2.
- **Source-Structure Paper** — *Effective Source Structure and Admissibility Closure in VERSF.* Inherited: finite distinguishability (its §2.1; used in Lemma 4.2 nondegeneracy and Theorem 5.2 Step 3 finite distinguishability of orthogonal directions), the irreversible-commitment ontology (used in Theorem 5.2 Step 2 single-direction commitment orientation per edge), and the four-sector source decomposition.
- **Lorentzian Completion Paper** — *Lorentzian Causal Completion of Emergent Commitment Geometry in VERSF.* Inherited: the CCC framework supplying finite-speed coherence-preserving propagation (used in Theorem 5.2 Step 1 indefinite signature derivation), the coherence scale ξ (the scale of $\eta_{\mu\nu}$, recorded as OP8), the (−, +, +, +) signature convention of the inherited continuum-limit metric $g_{\mu\nu}$ (consistent with the substrate-derived signature of Theorem 5.2), and the natural pairing $\eta_{\mu\nu} \leftrightarrow g_{\mu\nu}$ via the inherited frame structure (Theorem 8.1).
- **Transport-Geometry Papers** (cited as the inherited transport-geometry framework of the VERSF programme) — refinement-stable holonomy framework at the continuum-limit level, the convergence of refinement transport, and the closure-compatibility condition. The Microscopic Origin paper supplies the substrate construction below this framework; the present paper supplies the substrate construction of the transport group used by that framework.
- **Dynamical-Geometry Paper** — *Bianchi-Compatible Geometry and Einstein-Compatible Dynamics in VERSF.* Inherited: the substrate-scale coupling form $\kappa_{eff} = 8\pi C_\lambda \hbar\xi^2/c^3$, the mostly-plus signature convention (its §2.1), and the GR Recovery Theorem with (R1)–(R3) and (B1)–(B2).
- **Variational Paper** — *Unified Covariant Action for Emergent Commitment Geometry in VERSF.* Inherited: the unified action $\mathcal{S}_{VERSF}$, the Variational Closure Theorem, the action-level Einstein–Hilbert uniqueness, and the four-sector decomposition with anisotropic Wilson coefficients ($\alpha_{\hat{Q}}$, $\beta_{\hat{Q}}$, $\gamma_{\hat{Q}}$) recorded as OP3.

- **Einstein–Hilbert Emergence Paper** — *Substrate Derivation of the Einstein–Hilbert Action in VERSF.* Inherited: the continuum-emergence cascade Theorem 3.5 → 4.1 → 5.2 → 6.2 / 6.3 → 7.1 → 8.2, with the present paper supplying substrate-level transport-group derivation below the cascade.
- **Microscopic Origin Paper** — *Microscopic Origin of Refinement Transport in VERSF.* Inherited extensively: the $K = 7$ simplicial commitment foam (Definition 3.1 of that paper), the closure-compatible refinement framework (Definition 3.2), the closure-compatible foam properties (Definition 3.3, conditions F1–F4), the edge transport operators U_{ij} (Definition 4.1), the path holonomy T_γ (Definition 4.2 and Definition 8.1 transport-vs-causal-graph distinction), the simplicial loop holonomy (Definition 4.3), the composition consistency lemma (Lemma 4.4), the refinement-Cauchy convergence framework (Theorem 5.3), the Levi-Civita emergence (Theorem 6.2), the BCB-triangle-closure torsion-free lemma (Lemma 6.3), the discrete plaquette curvature → Riemann tensor convergence (Theorem 7.2), the commitment-order causal structure (Definition 8.1 and Theorem 8.2 Lorentzian compatibility), the microscopic Einstein–Hilbert recovery (Theorem 9.1), the residual closure-history operator and bilocal memory kernel emergence (Definition 10.1 and Theorem 10.2), and the microscopic GR recovery (Theorem 11.1). The present paper closes OP9 of the Microscopic Origin paper.
- **K = 7 Wilson Limit Paper** — methodological precedent for the structurally parallel U(1) gauge-bundle transport-group inheritance problem. Inherited: the methodological framework of $K = 7$ closure-counting integer $N_{loop} = 14$ controlling the bare coupling ($\beta_{K=7} = 2^7 \cdot 15/14$ in the abelian U(1) case), the one-loop matching analysis at fixed scale supplying the suppression rate, the structural-pattern caution that "second-order" closure may manifest at the closure-cycle level rather than RG-order level (§11.8 of that paper), and the honesty lesson that a one-step substrate derivation may not deliver full Lorentz compatibility without near-isotropy of the bare structure (§9.2 R1 of that paper). The non-abelian extension to the SO(1, D) tangent-bundle case is recorded as OP9 of the present paper.

Numerical results and structural conventions inherited from these papers — including the $K = 7$ closure architecture, $N_{loop} = 14$, the geometric refinement scaling $\varepsilon_n \leq C \lambda^n$, the coherence scale ξ, the substrate-scale coupling $\kappa_{eff} = 8\pi C_\lambda \hbar\xi^2/c^3$, and the mostly-plus signature convention — are treated throughout the present paper as established inputs.