Continuum-Limit Regularity and Cone Convergence in VERSF

Admissible TPB Refinement, Smooth Cone Fields, and the Emergence of Strongly Causal Continuum Structure

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The companion paper *Structural Necessity of Lorentzian Geometry in VERSF* established a conditional result: *if* the universe at large scales is described by a smooth four-dimensional continuum equipped with light cones and a conserved flow of information, *then* the geometry must be Lorentzian — the geometry of special and general relativity. That result has an obvious follow-up question: why should the universe look smooth at large scales in the first place? At its most fundamental level, VERSF describes reality as a network of discrete commitment events — irreversible bits of distinction propagating through a substrate. There is no smoothness, no continuum, no geometry in any familiar sense. So how does the smooth geometry of relativity actually emerge?

This paper addresses that question.

The intuitive picture is as follows. Imagine watching the discrete substrate at finer and finer resolutions — successive "refinements" that reveal more of its structure. At each level the structure is purely combinatorial: events, allowed transitions between them, a count of bits stored in each region. The question is whether, as you zoom in, this combinatorial structure starts looking more and more like a smooth four-dimensional spacetime. The paper proves that yes — under certain natural conditions — it does.

The conditions matter, and they are external to the substrate axioms — they have to be checked for any concrete substrate model. Loosely, they say: the substrate has to grow at a controlled rate as you refine it (no wild explosion in the number of states at any scale); the number of allowed local directions at each point has to be bounded uniformly (so the limiting geometry has a well-defined finite dimension); the rule for refining has to be compatible across resolutions (so what was allowed at a coarser level remains allowed at a finer one); allowed trajectories cannot return arbitrarily close to themselves without forming an exact loop (a quantitative no-near-returns condition, needed to prevent almost-closed causal paths in the continuum limit); the cone structure has to be stable enough under refinement that the limiting geometry is not too irregular; and the limit has to admit a smooth four-dimensional structure (which sidesteps a subtle issue in the topology of four-manifolds, where some topological spaces fail to admit any smooth structure at all). Under these conditions, the discrete substrate converges to a smooth continuum equipped with exactly the light-cone-and-flow structure required by the companion paper's theorem.

The result comes in two stages. The first stage shows that the substrate converges to something *continuous* but possibly rough — what mathematicians call a "Lorentzian length space," a generalisation of relativistic spacetime that allows the metric and cone field to be continuous without being differentiable. The second stage shows that, under a further regularity condition, this continuous structure can be approximated arbitrarily well by genuinely smooth four-dimensional geometries — smooth enough to plug into the classical machinery of general relativity that the companion paper invokes. Either stage is enough to feed into the companion paper's chain, depending on how strong a regularity statement is needed downstream.

The philosophical importance is substantial. The companion paper showed that Lorentzian geometry *would* emerge if the regularity conditions held. The present paper shows *why* those regularity conditions emerge from the substrate, conditional on a set of external assumptions about how the substrate refines. Together, the two papers *conditionally close* the chain from VERSF's discrete substrate axioms to the smooth Lorentzian geometry of relativity — conditional on the regularity hypotheses that this paper introduces but does not derive from substrate principles. What remains is a substrate-engineering question: which concrete discrete models actually satisfy those regularity hypotheses?

That substrate-engineering question is now the principal remaining mathematical risk in the VERSF geometry programme. The convergence machinery itself is, conditional on those regularity conditions, mechanically tight.

Abstract

We prove that admissible TPB refinement sequences in the Void Energy–Regulated Space Framework (VERSF) converge, under finite propagation, uniform doubling, refinement compatibility, and Bit Conservation and Balance (BCB), to a strongly-causal continuum equipped with a non-degenerate invariant cone field. Under two further regularity hypotheses, this continuous structure admits a C^k Lorentzian-metric realisation for any prescribed k ≥ 2. The combined result supplies the smooth regularity assumptions used as input by the companion paper *Structural Necessity of Lorentzian Geometry in VERSF*, closing the geometry-emergence chain from substrate axioms to emergent Lorentzian geometry modulo the open problems identified in §13.

We prove eight principal results, each labelled with its epistemic status. The first five establish the continuous-regularity layer; the last three (introduced in §10) establish the smooth-regularity upgrade.

- **Theorem 1 (Precompactness — proven, conditional on H1, H6).** Under finite propagation (H1) and uniform local doubling (H6), admissible TPB refinement sequences

are Gromov–Hausdorff precompact in the transport pseudometric. H7 is used at the next layer (Theorem 2) for uniform-dimension control on the limit, not for precompactness.

- **Theorem 2 (Cone-field convergence — proven, conditional on H1, H5, H7 + H3, H4 for continuity from equivariance).** Discrete admissible-transport-direction sets 𝒞_ℓ(x) Kuratowski-converge to a non-degenerate continuous double-cone field 𝒞(x) on the continuum limit.
- **Theorem 3 (Strong causality of the limit — proven, conditional on A2, H1, H5, H6, H6′, and Theorem 1).** The continuum limit is strongly causal in the standard general-relativity sense. The quantitative acyclicity bound is supplied by H6′ (uniform local transport sparsity) via Lemma 7.2, a new external transport-sparsity hypothesis introduced in §7.2 to lift finite-level acyclicity to continuum strong causality. The proof additionally invokes Lorentzian Gromov–Hausdorff convergence (Müller 2022; Minguzzi–Sämann 2023) to ensure that discrete causal sequences limit to continuum causal curves.
- **Theorem 4 (Local no-flux foliation existence — proven, conditional on Lemma 6.1.a of the companion paper plus local nonvanishing of J^μ).** Every admissible normal neighbourhood admits a local foliation by integral curves of J^μ with zero lateral transport flux through the lateral boundary.
- **Theorem 5 (Continuum regularity, continuous version — proven, by conjunction of Theorems 1–4).** Under H1–H7, H6′ and BCB, admissible TPB refinement sequences converge to a connected strongly-causal Lorentzian length space (in the Kunzinger–Sämann sense) equipped with a continuous non-degenerate cone field and admitting local no-flux transport foliations.
- **Lemma 10.1 (C^0 metric realisation — proven, conditional on H9).** The continuous cone field of Theorem 5 is the null cone of a unique (up to global units) continuous Lorentzian metric g_∞ on ℳ_∞, with the scale fixed by Proposition T1 of the companion paper.
- **Theorem 6 (Hölder regularity — proven, conditional on H8).** Under H8 (refinement Hölder compatibility), g_∞ is of class C^{0,α} for the Hölder exponent α of H8.
- **Theorem 7 (Smooth approximation — proven, using Sämann 2016 causal-stability machinery).** For every prescribed k ≥ 1 and ε > 0, g_∞ admits a C^k Lorentzian approximation g_∞^{(k,ε)} with ‖g_∞^{(k,ε)} − g_∞‖*{C^0} < ε, signature (1, 3), and causal structure agreeing with* ≺∞ in the Kuratowski sense.
- **Theorem 8 (Bridge to the companion paper — proven, by combination of Theorems 5, 6, 7).** Under H1–H9 and BCB, the continuum limit admits a C^k Lorentzian metric for every prescribed k ≥ 2 (in particular k = 2 suffices for Malament–Hawking–King–McCarthy), with which the companion paper's Lemma 5.2 and §6.1 derivations operate with full smoothness inputs.

The principal methodological move is the explicit separation of three convergence questions that are often conflated: *(a)* metric convergence (Theorem 1, requiring volume-growth bounds), *(b)* causal-structure convergence (Theorem 2, requiring refinement compatibility in a precise lifting sense), and *(c)* topological / regularity convergence (Theorems 3–4, requiring substrate-level acyclicity and current-nonvanishing). The §10 regularity upgrade adds a fourth concern, *metric regularity* (Lemma 10.1 and Theorems 6, 7, 8), addressed using machinery imported from the Lorentzian length space literature (Sämann 2016, Burtscher 2015, Kunzinger–Sämann 2018, Chruściel–Grant 2012). We also address the natural concern that the framework simply reproduces the causal-set programme by clarifying in §4.3 how VERSF refinement convergence differs structurally from causal-set continuum limits (Bombelli–Lee–Meyer–Sorkin 1987 and successors).

We close by identifying the remaining open problems — most importantly, derivation of H6, H6′, H7, H8, H9 from substrate principles, the C^∞ regularity gap (§10.8), the global-foliation question, and the carryovers from the companion paper — and by stating the combined Theorems 5 and 8 as the long-promised companion to the Lorentzian-emergence result.

Notation glossary

Notation extends the companion Lorentzian-emergence paper. New symbols introduced in this paper:

Symbol	Meaning
{P_ℓ, T_ℓ, B_ℓ}_{ℓ∈ℕ}	Admissible TPB refinement sequence at level ℓ
P_ℓ	Substrate graph at refinement level ℓ
T_ℓ	Admissible transport relation at level ℓ
B_ℓ	Substrate bit-measure at level ℓ
𝓡_{ℓ→ℓ+1}	Refinement map TPB_ℓ → TPB_{ℓ+1}
d_ℓ(x, y)	Transport pseudometric at refinement level ℓ
≺_ℓ	Causal accessibility relation at level ℓ
≺_∞	Continuum limit of causal accessibility
𝓒_ℓ(x)	Discrete admissible-transport-direction set at level ℓ
𝓒(x)	Continuum cone field (limit of 𝓒_ℓ(x))
(𝓜_∞, d_∞, ≺_∞)	Continuum limit triple (manifold, pseudometric, causal order)
λ_ℓ	Refinement scale identifying P_ℓ with the continuum-limit metric structure
σ_ℓ	Minimum distinguishability scale at refinement level ℓ (substrate units)
σ̃_ℓ	Rescaled distinguishability: σ̃_ℓ := λ_ℓ · σ_ℓ (continuum units)
σ_*	Continuum lower bound on σ̃_ℓ: σ_* := liminf_{ℓ→∞} σ̃_ℓ (positive by H6′)
g_∞	Continuum Lorentzian metric (constructed in Lemma 10.1)
g_∞^{(k,ε)}	C^k smoothing of g_∞ at scale ε (Theorem 7)
d_H	Hausdorff distance on the space of admissible direction sets
GH	Gromov–Hausdorff
KS	Kunzinger–Sämann (Lorentzian length space framework)

Symbol	Meaning
H6	Uniform local doubling hypothesis (introduced in §2.2)
H6′	Uniform local transport sparsity hypothesis (introduced in §2.2 / §7.2)
H7	Bounded combinatorial dimension hypothesis (introduced in §2.2)
H8	Refinement Hölder compatibility hypothesis (introduced in §2.2 / §10.3)
H9	Smooth-structure existence hypothesis (introduced in §2.2 / §10.3)
C^{0,α}	Hölder regularity of exponent α
C^k	k-times continuously differentiable

All other notation as in the companion Lorentzian-emergence paper.

Table of Contents

---

1. Introduction

The companion paper *Structural Necessity of Lorentzian Geometry in VERSF* established that, given a smooth strongly-causal continuum equipped with an invariant cone field and a divergence-free transport current, Lorentzian geometry is the unique stable admissible structure. The proof chain runs cone invariance → conformal class (Malament–Hawking–King–McCarthy) → conformal factor fixing (BCB plus finite distinguishability) → metric Lorentzian structure (Theorem 1L) → invariant interval (Theorem 2) → boost generation (Theorem 3).

That chain takes as given a smooth strongly-causal continuum cone manifold. The §14 dependency-clarification of the companion paper explicitly flags this prerequisite — the *continuum-limit regularity theorem* — as the principal remaining mathematical risk in the geometry programme:

"A later continuum-limit paper should prove that the admissible TPB refinement sequence converges to a smooth cone distribution satisfying the causal regularity conditions required by Malament–Hawking–King–McCarthy. Once that is established, the chain from Lemma 5.1 through Lemmas 6.1.a–e to Theorems 1L, 1G, 2, and 3 follows essentially mechanically."

The present paper attempts to supply that missing regularity layer.

The central question is:

Why should the discrete TPB substrate converge to a smooth strongly-causal cone manifold equipped with a divergence-free transport current?

The answer developed here factors into three logically distinct convergence claims: *metric* convergence (the TPB pseudometric sequence is GH-precompact), *causal* convergence (discrete cone-direction sets Kuratowski-converge to a continuous cone field), and *regularity* convergence (the limit is strongly causal and admits local no-flux foliations). The first requires substrate-level volume-growth bounds (H6, H7) that are external regularity hypotheses on the refinement family rather than derivable from BCB alone — this is honestly flagged. The second follows from H1–H5 and BCB once H7 is in place. The third requires, additionally, a *quantitative* strengthening of A2 — supplied by H6′ (uniform local transport sparsity, §7.2) — to lift finite-level acyclicity to continuum strong causality. The §10 regularity upgrade adds a fourth question: under what additional hypotheses (H8, H9) does the continuous structure admit a smooth metric realisation, sufficient for the companion paper's Malament–HKM application.

A reader tracking the open-problem landscape should note an asymmetry parallel to the one identified in §14 of the companion paper. *Given* admissible refinement families satisfying H6, H6′, H7 (and H8, H9 for the smooth upgrade), the convergence chain to a regular continuum is mechanically tight; the residual risk lies in *which* substrate models satisfy these external regularity hypotheses in the first place. The geometry programme's mathematical-risk concentration has therefore moved one further layer down — from "does Lorentz emerge" (companion paper, answered conditionally) to "does the regular continuum emerge" (this paper, answered conditionally on H6, H6′, H7, H8, H9) to "do substrate models with these properties exist" (a substrate-engineering problem flagged in §14).

We engage with the causal-set programme of Bombelli, Lee, Meyer, and Sorkin (1987), the Lorentzian-length-space framework of Kunzinger and Sämann (2018), and the Lorentzian-Gromov–Hausdorff convergence theory developed by Müller (2022) and Minguzzi–Sämann (2023). The VERSF refinement-convergence setup overlaps with these in mathematical structure but differs in motivation: causal-set theory takes the discrete causal order as fundamental and asks when it generates a Lorentzian manifold, whereas VERSF takes commitment-transport dynamics as fundamental and asks when the induced causal order generates a regular continuum.

The technical results are largely importable; the §4.3 scope clarification spells out the relationship.

---

2. Hypotheses, axioms, and propositions

To make the dependency structure explicit we list inputs separately from the substrate axioms of the VERSF programme. The architecture mirrors §2 of the companion paper.

2.1 Substrate axioms (carried in from prior VERSF papers)

- **A0 — Observer-invariant distinguishability.** As in companion paper §2.1.
- **A1 — Finite distinguishability.** As in companion paper §2.1.
- **A2 — Irreversible commitment.** As in companion paper §2.1. The substrate $\prec_\ell$ relation is therefore acyclic at every finite ℓ.
- **A3 — Reversible pre-commitment evolution.** As in companion paper §2.1.
- **A4 — Local coupling.** As in companion paper §2.1. Substrate interactions propagate only between neighbouring admissible states; the local-coupling graph at refinement level ℓ has bounded local degree.
- **BCB — Bit Conservation and Balance.** As in companion paper §2.1.

2.2 Hypotheses used in this paper

Hypotheses H1–H5 are inherited from the companion paper. New hypotheses H6, H6′, H7 (continuous-regularity layer) and H8, H9 (smooth-regularity-upgrade layer, used in §10) are introduced here.

- **H1 — Finite invariant propagation.** As in companion paper §2.2. At each refinement level ℓ there is a propagation bound c_ℓ with $c_\ell \to c < \infty$ as $\ell \to \infty$.
- **H2 — Operational inaccessibility of substrate frame.** As in companion paper §2.2.
- **H3 — Homogeneity of admissible coarse-graining.** As in companion paper §2.2.
- **H4 — Isotropy of admissible coarse-graining.** As in companion paper §2.2.
- **H5 — Refinement compatibility.** As in companion paper §2.2 and §8.1. Made operational in §3.2 below as the existence of refinement maps $\boldsymbol{\mathcal{R}}_{\ell\to\ell+1}$ satisfying $T_\ell \subseteq \boldsymbol{\mathcal{R}}^{-1}(T_{\ell+1})$ and $B_{\ell+1} \circ \boldsymbol{\mathcal{R}} = B_\ell + o(1)$.
- **H6 — Uniform local doubling.** *There exists $D < \infty$ such that for every $\ell \in \mathbb{N}$, every $x \in P_\ell$, and every $r > 0$, the transport-pseudometric ball $B_\ell(x, 2r)$ admits a covering by at most D balls of radius r.* This is the standard precompactness input for Gromov–Hausdorff theory.
- **H6′ — Uniform local transport sparsity (introduced in §7.2).** *There exists $\sigma > 0$, independent of refinement level ℓ, such that for every admissible commitment-transport sequence $\gamma_\ell = (s_0, s_1, \ldots, s_n)$ at level ℓ, any two distinct committed states $s_i \neq s_j$ along γ_ℓ satisfy $d_\ell(s_i, s_j) \geq \sigma_\ell$, where σ_ℓ is the minimum distinguishability scale at level ℓ, and where the rescaled quantity $\tilde{\sigma}_\ell := \lambda_\ell \cdot \sigma_\ell$ has a strictly positive continuum lower bound $\liminf_{\ell\to\infty} \tilde{\sigma}_\ell = \sigma_{} > 0$.** H6′ is the *quantitative* strengthening of A2

(irreversibility) needed to lift finite-level acyclicity to continuum strong causality; it excludes almost-closed causal curves at the substrate level.

- **H7 — Bounded combinatorial dimension.** *There exists* $N < \infty$ *such that for every* $\ell \in \mathbb{N}$ *and every* $x \in P_\ell$*, the number of admissible transport directions at* x *is bounded by* N*.* Together with H6 this controls the dimension of the continuum limit.
- **H8 — Refinement Hölder compatibility (introduced in §10.3).** *There exist* $\alpha > 0$ *and* $K < \infty$ *such that for every* $\ell \in \mathbb{N}$ *and every* $x, y \in P_\ell$*, the discrete cone-direction sets satisfy* $d_H(\mathcal{C}_\ell(x), \mathcal{C}_\ell(y)) \leq K \cdot d_\ell(x, y)^\alpha + \varepsilon_\ell$ *with* $\varepsilon_\ell \to 0$*.* H8 strengthens H5 to give Hölder uniformity across refinement levels; it is the input required for the regularity-upgrade machinery of §10 (Theorem 6 and subsequent).
- **H9 — Smooth-structure existence (introduced in §10.3).** *The topological 4-manifold structure inherited by* $(\mathcal{M}_\infty, d_\infty)$ *from Theorem 1 admits at least one compatible* C^∞ *atlas.* H9 sidesteps the Donaldson–Freedman exotic-4-manifold phenomenon and is used in Lemma 10.1 and downstream.

2.3 Propositions and theorems derived in this paper

- **Theorem 1** (Precompactness, §5).
- **Theorem 2** (Cone-field convergence, §6).
- **Lemma 7.2** (Quantitative acyclicity bound, §7).
- **Theorem 3** (Strong causality of the limit, §7).
- **Theorem 4** (Local no-flux foliation existence, §8).
- **Theorem 5** (Continuum regularity, continuous version, §9).
- **Lemma 10.1** (C^0 metric realisation, §10).
- **Theorem 6** (Hölder regularity, §10).
- **Theorem 7** (Smooth approximation, §10).
- **Theorem 8** (Bridge to the companion paper, §10).

2.4 Epistemic status

A0–A4 and BCB are foundational VERSF axioms. H1–H5 are hypotheses of the companion paper. H6, H6′, H7, H8, H9 are new external regularity hypotheses on the refinement family — they are the analogues, at the convergence-theory layer (H6, H6′, H7) and the metric-regularity-upgrade layer (H8, H9), of H3 and H4 at the geometric-symmetry layer. All five are flagged in §14 for future substrate-level derivation. Theorems 1–8, Lemma 7.2, and Lemma 10.1 are derived results.

2.5 Relationship to the companion paper

The present paper supplies the regularity assumptions used as input by the companion paper. Specifically:

- The continuum manifold $\mathcal{M}_{\text{coarse}}$ of the companion paper §3 is identified with $(\mathcal{M}_\infty, d_\infty)$ of Theorem 1 below.
- The cone field $\mathcal{C}(x)$ on which Lemma 5.1 (cone invariance) of the companion paper operates is the continuum cone field of Theorem 2 below.

- The strong causality assumed by Lemma 5.2 of the companion paper for the Malament–HKM application is the conclusion of Theorem 3 below.
- The local no-flux tube construction used in Lemma 6.1.d of the companion paper is the conclusion of Theorem 4 below.
- The smoothness and connectedness assumed throughout the companion paper for Corollary 6.1.e (conformal factor fixed) are properties of the limit (𝓜_∞, d_∞) established in Theorem 5.

With both papers in place, the chain from substrate axioms (A0–A4, BCB) and refinement-regularity hypotheses (H1–H7, H6′, H8, H9) to Lorentzian geometry (Theorem 1L of the companion paper) is *conditionally closed*, modulo the substrate-level derivation of the external regularity hypotheses H6, H6′, H7, H8, H9 — which remain unproved from substrate principles in the present paper.

---

3. Admissible TPB refinement structure

3.1 Refinement sequence

Let

{P_ℓ, T_ℓ, B_ℓ}_{ℓ∈ℕ}

denote an admissible TPB refinement sequence, where:

- P_ℓ is the substrate graph at refinement level ℓ (a countable set of nodes with the A4-local-coupling edge structure);
- T_ℓ ⊆ P_ℓ × P_ℓ is the admissible transport relation at level ℓ, generated by the substrate transport rules under A2 and A3;
- B_ℓ: 𝒫(P_ℓ) → [0, ∞] is the substrate bit-measure at level ℓ, assigning a finite bit count to each bounded subset.

Each refinement step satisfies:

- (local transport preservation) A4-couplings at level ℓ refine to A4-couplings at level ℓ+1;
- (finite propagation) the substrate propagation bound c_ℓ satisfies c_ℓ → c as ℓ → ∞ for some 0 < c < ∞;
- (BCB conservation) the substrate bit-current is divergence-free at every level under coarse-graining;
- (refinement compatibility, H5) the refinement maps satisfy the conditions in §3.2.

3.2 Refinement maps

For each ℓ, there is a refinement map

$\mathcal{R}_{\ell\to\ell+1}: P_\ell \to P_{\ell+1}$

(more generally, a refinement functor on the TPB category, in the sense of the σ-duality and sequential-interface-transport papers). We require:

- **(Transport preservation)** $T_\ell \subseteq \mathcal{R}^{-1}(T_{\ell+1})$. Equivalently: if $x \prec_\ell$ *y at level ℓ, then $\mathcal{R}(x)$* $\prec_{\ell+1} \mathcal{R}(y)$ at level $\ell+1$.
- **(Measure compatibility)** $B_{\ell+1}(\mathcal{R}(U)) = B_\ell(U) + o(1)$ as $\ell \to \infty$ for every bounded $U \subseteq P_\ell$.
- **(Local-degree control)** Each refinement step preserves the bounded-local-degree property of A4; the number of A4-neighbours at any node is bounded uniformly in ℓ by some constant $\Delta < \infty$.

The refinement functor $\mathcal{R}$ of the companion paper (used in H5 of that paper and in Theorem 3 of that paper) is the categorical lift of these maps.

---

4. Discrete causal-metric structure and convergence setup

4.1 Transport pseudometric

At refinement level ℓ, define the transport pseudometric

$$d_\ell(x, y) = \inf_{\gamma: x \to y} \Sigma_{i=1}^{N_\gamma} w_\ell(e_i),$$

where the infimum is over admissible transport paths γ from x to y, and $w_\ell(e)$ is the substrate transport-time cost of edge e at level ℓ. Finite propagation (H1) gives the lower bound

$$d_\ell(x, y) \geq |x - y|_\ell / c_\ell,$$

where $|\cdot|_\ell$ is the underlying combinatorial (graph-theoretic) distance at level ℓ.

The pair (P_ℓ, d_ℓ) is a finite-or-countable pseudometric space; we work with its metric quotient when needed.

4.2 Causal accessibility relation

Define the substrate-level causal accessibility relation

$x \prec_\ell y$ iff there exists an admissible commitment-transport sequence at level ℓ taking x to y.

This relation is *acyclic* at every finite ℓ — substrate-level irreversibility (A2) forbids any closed admissible commitment-transport sequence.

The continuum-level causal relation is then

$$\prec_\infty = \lim_{\ell\to\infty} \prec_\ell$$

in the sense made precise in §4.3 (Kuratowski convergence of the relation viewed as a closed subset of $\mathcal{M}_\infty \times \mathcal{M}_\infty$).

4.3 Scope of convergence claims

Clarification on the mode of convergence. The natural mathematical setting for the convergence claims of this paper is the *Lorentzian length space* framework of Kunzinger and Sämann (2018), together with the Lorentzian-Gromov–Hausdorff convergence developed by Müller (2022) and refined by Minguzzi and Sämann (2023). In that framework, a Lorentzian pre-length space is a tuple $(X, d, \ll, \leq, \tau)$ consisting of a metric space, a causal preorder, a chronological order, and a Lorentzian distance function; convergence of a sequence of such structures is defined by joint GH convergence of the metric structure together with Kuratowski convergence of the causal and chronological relations.

The TPB refinement sequence $\{(P_\ell, d_\ell, \prec_\ell)\}_{\ell\in\mathbb{N}}$ fits this framework: each level is a discrete Lorentzian pre-length space (with τ_ℓ defined as the longest admissible transport time between causally related points, and $\ll_\ell$ *defined as strict* $\prec_\ell$*). The continuum limit* $(\mathcal{M}_\infty, d_\infty, \prec_\infty)$ is, by Theorems 1–5 below, a Lorentzian length space in the sense of Kunzinger–Sämann. (The smooth metric structure of §10 — Theorems 6, 7, 8 — upgrades this synthetic Lorentzian length space to a smooth Lorentzian manifold under H8, H9; the synthetic-LLS framing is the natural pre-upgrade setting.)

This paper does not claim to derive Lorentzian length space convergence from arbitrary discrete causal structures. The claim is that *given* the admissibility conditions H1–H7 plus H6′ plus BCB, the refinement sequence converges in this sense.

Relation to causal-set theory. The discrete causal accessibility relations $\prec_\ell$ are formally similar to causal sets in the sense of Bombelli, Lee, Meyer, and Sorkin (1987). The continuum-limit question — when does a discrete causal order generate a continuous Lorentzian manifold? — is a central open problem in that programme (the *Hauptvermutung* of causal-set theory; cf. Sorkin 2003, Bombelli–Henson–Sorkin 2009). The VERSF setup differs in two structural ways:

1. *Refinement rather than sprinkling.* Causal-set continuum limits are typically considered for Poisson sprinklings of a fixed Lorentzian manifold; the manifold is assumed and the discrete structure is sampled from it. Here, the refinement sequence is *intrinsic* to the substrate dynamics — there is no pre-existing manifold to sprinkle into. The continuum limit is generated by the refinement, not approximated by sampling.
2. *Transport-current structure.* The TPB substrate carries additional structure beyond the causal order: the bit-measure B_ℓ and the transport-current J_ℓ. This additional structure is what makes Theorem 4 (no-flux foliations) and the companion paper's Corollary 6.1.e (conformal factor fixing) work. Causal sets in the BLMS sense have no equivalent of the transport current.

These differences mean that VERSF refinement convergence is not a special case of causal-set continuum limits, though it shares mathematical tools (Kuratowski convergence of orders, GH convergence of metrics, Lorentzian length space framework).

---

5. Theorem 1 — Precompactness of admissible refinement sequences

5.1 Statement

Theorem 1 (Precompactness). *Under H1 and H6, every admissible TPB refinement sequence* $\{(P_\ell, d_\ell)\}\{\ell\in\mathbb{N}\}$ *is Gromov–Hausdorff precompact: every subsequence admits a further subsequence converging in the GH sense to a complete pseudometric space* $(\mathcal{M}\infty, d_\infty)$*.* (Note: H7 is not needed for precompactness; it enters at the cone-field convergence layer in §6 to supply uniform-dimension control on the limit.)

5.2 Proof

The proof is a direct application of the Gromov precompactness theorem (Gromov 1981; cf. Burago–Burago–Ivanov 2001, Theorem 7.4.15) to the family $\{(P_\ell, d_\ell)\}$.

The theorem states that a family of metric spaces is GH-precompact iff it is *uniformly totally bounded* — i.e., for every $R, \varepsilon > 0$, there is $N(R, \varepsilon) < \infty$ such that every space in the family of diameter $\leq R$ admits a covering by at most $N(R, \varepsilon)$ balls of radius ε.

We verify uniform total boundedness from H1 and H6 alone:

- **From H1 (finite propagation):** the transport pseudometric d_ℓ is bounded below by the graph distance divided by c_ℓ, so d_ℓ-balls of radius r contain graph-balls of radius at most $r\cdot c_\ell \to r\cdot c$.
- **From H6 (uniform local doubling):** any d_ℓ-ball of radius $2r$ can be covered by D d_ℓ-balls of radius r, with D independent of ℓ.

Iterating the H6 doubling estimate gives: for every ℓ, every $R > 0$, and every $\varepsilon > 0$, the number of d_ℓ-balls of radius ε needed to cover a d_ℓ-ball of radius R is bounded by $D^{\lceil \log_2(R/\varepsilon)\rceil}$, which is finite and depends only on R, ε, D, not on ℓ. This is uniform total boundedness.

By Gromov precompactness, every subsequence admits a further GH-convergent subsequence. Completeness of the limit is automatic for GH limits of complete metric spaces. This completes the proof of Theorem 1.

5.3 Status of H6 and H7

H6 (uniform doubling) and H7 (bounded combinatorial dimension) are external regularity hypotheses on the refinement family. They are *not* derived from BCB or from A0–A4 in this paper. The substrate-engineering question — *which* discrete substrate models satisfy H6 and H7 — is flagged in §14 as part of the future-work programme.

H6 alone is sufficient for GH precompactness; this is the content of Theorem 1 as stated above. H7 enters at the *next* layer: it supplies the uniform-dimension control on the continuum limit needed for Theorem 2 (the discrete-cone-direction sets at each node have at most N admissible directions, bounding the dimension of the limiting cone field). The natural conjecture is that A4-local-coupling with bounded local degree, plus refinement compatibility, implies both H6 and H7; this conjecture is the principal substrate-engineering problem for the present paper.

---

6. Theorem 2 — Cone-field convergence

6.1 Discrete admissible-transport-direction sets

At refinement level ℓ and at each $x \in P_\ell$, define the discrete admissible-transport-direction set

$\mathcal{C}_\ell(x) = \{ v \in T_x(P_\ell) : \text{there exists } y \in P_\ell \text{ with } x \prec_\ell y \text{ and direction}(x, y) = v \}$,

where $T_x(P_\ell)$ denotes the local tangent-cone-like structure at x (defined as the set of admissible local-coupling directions in the A4 sense), and $\text{direction}(x, y) \in T_x(P_\ell)$ is the local direction from x to y at level ℓ.

The set $\mathcal{C}_\ell(x)$ is the discrete analogue of a future cone: it specifies which local directions support admissible commitment-transport at level ℓ.

6.2 Statement

Theorem 2 (Cone-field convergence). *Under H1, H5, and H7, along any GH-convergent subsequence* $(P_{\ell_k}, d_{\ell_k}) \to (\mathcal{M}_\infty, d_\infty)$*, the discrete cone-direction sets* $\mathcal{C}_{\ell_k}(x_{\ell_k})$ *(taken along sequences* $x_{\ell_k} \to x_\infty \in \mathcal{M}_\infty$*) Kuratowski-converge to a non-degenerate double-cone field* $\mathcal{C}(x_\infty)$ *on the continuum limit. Under the additional hypotheses H3 (homogeneity) and H4 (isotropy), the limiting cone field is continuous in the Hausdorff topology on cone sets.*

6.3 Proof

The proof factors into three lemmas, mirroring the architecture of §6.1 of the companion paper.

Lemma 6.3.a — Upper Kuratowski limit non-empty. *Under H1 and H7, for every* $x_\infty \in \mathcal{M}_\infty$ *and every convergent sequence* $x_{\ell_k} \to x_\infty$*, the upper Kuratowski limit* $\limsup_{k}$

$\mathcal{C}_{\ell_k}(x_{\ell_k})$ is non-empty and contained in a uniformly bounded subset of the unit ball of $T_{x_\infty}(\mathcal{M}_\infty)$.

Proof. Each $\mathcal{C}_{\ell_k}(x_{\ell_k})$ is non-empty by A2 (commitment events generate at least one admissible transport direction per node). H7 bounds the cardinality of $\mathcal{C}_{\ell_k}(x_{\ell_k})$ by N uniformly in ℓ. H1 bounds the magnitude of each element of $\mathcal{C}_{\ell_k}(x_{\ell_k})$ by $c_\ell \leq c + o(1)$. The sequence of finite sets is therefore contained in a compact subset of the unit ball, and the upper limit is non-empty.

Lemma 6.3.b — Lower Kuratowski limit non-degenerate. *Under H1, H5, and the refinement-compatibility lifting condition (clause $T_\ell \subseteq \boldsymbol{\mathcal{R}}^{-1}(T_{\ell+1})$ of §3.2), the lower Kuratowski limit $\liminf_{k} \mathcal{C}_{\ell_k}(x_{\ell_k})$ contains an open non-degenerate double cone in $T_{x_\infty}(\mathcal{M}_\infty)$.*

Proof. Refinement compatibility ensures that every admissible direction at level ℓ persists to level $\ell+1$ ($T_\ell \subseteq \boldsymbol{\mathcal{R}}^{-1}(T_{\ell+1})$). Therefore the discrete cones are *increasing* under refinement in the appropriate sense: any direction admissible at level ℓ remains admissible at all subsequent levels. The lower Kuratowski limit therefore contains the union over ℓ of all admissible directions, modulo the GH identifications. By H1 the propagation bound c is *positive* (excluding the $c = 0$ Carrollian degeneration), so the union is open and non-degenerate.

Double-lobe structure. The future lobe of the cone is generated by A2-admissible commitment-transport: by A2 (irreversibility), every commitment event admits at least one future-directed admissible transport direction. The past lobe arises from a structurally distinct substrate-level mechanism — A3 (reversible pre-commitment evolution): pre-commitment substrate evolution is time-reversible, and the backward branch of admissible pre-commitment trajectories contributes the past-directed cone lobe at the continuum scale. The precise mechanism by which reversible pre-commitment evolution generates past-directed transport directions at the continuum scale — i.e., the translation between substrate-level pre-commitment reversibility (A3) and continuum-level backward causal propagation — is developed in the *sequential-interface-transport* paper (§4 admissibility-of-pre-commitment-branches) and the *admissible-coarse-graining* paper (§5 backward-branch contribution to the continuum cone). The present paper imports that translation as given. The future-past asymmetry between the two lobes — A2-irreversible future, A3-reversible past — is what distinguishes the continuum cone from a Galilean half-cone and is the substrate-level origin of the Lorentzian double-cone structure used in the companion paper.

Lemma 6.3.c — Upper and lower limits coincide. *Under H1, H5, H7, and the bounded-local-degree control of §3.2, the upper and lower Kuratowski limits coincide as point-sets, defining a non-degenerate open double cone $\mathcal{C}(x_\infty) \subseteq T_{x_\infty}(\mathcal{M}_\infty)$. The Kuratowski-limit cone field $x_\infty \mapsto \mathcal{C}(x_\infty)$ is continuous in the Hausdorff topology on cone sets.*

Proof. The upper limit is contained in the closure of the union over ℓ of $\mathcal{C}_\ell(x_\ell)$, modulo GH identifications. The lower limit contains this union as established in 6.3.b. The boundary of the union is controlled by H7 (only finitely many extremal directions at each level) and by bounded local degree (only finitely many limit accumulation points of these extremals); the upper and lower limits therefore have the same closure, which is the cone $\mathcal{C}(x_\infty)$.

Continuity of the cone field $x_\infty \mapsto \mathcal{C}(x_\infty)$ follows from H3 (homogeneity) and H4 (isotropy) of the refinement family. Explicitly: H3 supplies a transitive translation-equivariant action of the local-translation group G on the refinement family, and H4 supplies the analogous rotation-equivariance. Under this transitive action, the cone at every $x_\infty \in \mathcal{M}_\infty$ is the G-translate of the cone at any fixed basepoint x_0: $\mathcal{C}(x_\infty) = g_{x_\infty} \cdot \mathcal{C}(x_0)$, where $g_{x_\infty} \in G$ is any group element with $g_{x_\infty} \cdot x_0 = x_\infty$. Continuity of $x_\infty \mapsto \mathcal{C}(x_\infty)$ therefore reduces to continuity of the map $x_\infty \mapsto g_{x_\infty}$ (which is automatic for transitive actions of topological groups on homogeneous spaces) and continuity at the basepoint x_0. Continuity at the basepoint follows from H1 (uniform propagation bound, giving equicontinuity of the discrete cone fields $\mathcal{C}_\ell$ at x_0) combined with the refinement-compatibility lifting argument used in Lemma 6.3.b.

The stronger statement of *uniform Hölder regularity* of the cone field — used in §10 as the input H8 to the regularity-upgrade machinery — is not derivable from H1, H5, H7, H3, H4 alone; it requires the explicit Hölder-uniformity hypothesis H8. The continuity established here is sufficient for Theorem 5's continuous-Lorentzian-length-space conclusion; the Hölder strengthening is reserved for §10 (Theorem 6).

This completes the proof of Lemma 6.3.c.

Combining Lemmas 6.3.a–c, the Kuratowski limit exists and is a non-degenerate open double cone in $T_{x_\infty}(\mathcal{M}_\infty)$, varying continuously with x_∞. This completes the proof of Theorem 2.

6.4 Remark on convexity

Theorem 2 produces a non-degenerate continuous double cone field. Convexity of the cone — required for Lemma 5.2 of the companion paper — is the substantive content of the *convexity-of-admissible-directions theorem* of the sequential-interface-transport paper, cited in Lemma 5.2 of the companion paper. That theorem establishes that infinitesimal composition of admissible transport directions is closed under convex combinations in the continuum limit. Subject to that result, the cone field of Theorem 2 is convex.

---

7. Theorem 3 — Strong causality of the limit

This is the most delicate of the convergence theorems. The naive argument — "A2 forbids closed transport loops at every finite ℓ, therefore the continuum limit has no closed causal curves, therefore it is strongly causal" — conflates *acyclicity of the substrate relation* (a property of each finite $\prec_\ell$) with *strong causality of the continuum limit* (a stronger condition: every point admits arbitrarily small neighbourhoods through which no causal curve passes more than once). The continuum limit could in principle exhibit *almost-closed* causal curves that violate strong causality without violating acyclicity at any finite ℓ.

The honest argument therefore needs a *quantitative* acyclicity bound — Lemma 7.2 below — that controls not just the absence of closed loops but the absence of almost-closed loops uniformly in ℓ. This in turn requires a quantitative strengthening of A2, supplied by the new hypothesis H6′ (uniform local transport sparsity) introduced in §2.2. The role of H6′ is precisely to convert exact acyclicity (A2) into a quantitative lower bound, in the rescaled continuum metric, on the pairwise distance between distinct committed states along any admissible transport sequence — thereby excluding the near-return behaviour that bare A2 cannot rule out.

7.1 Strong causality, defined

A Lorentzian pre-length space (X, d, ≪, ≤, τ) is *strongly causal* at x ∈ X iff for every neighbourhood U of x there is a sub-neighbourhood V ⊆ U such that no causal curve starting and ending in V leaves U in between. It is *strongly causal* if it is strongly causal at every point.

7.2 Quantitative acyclicity at the substrate level

The finite-level relation ≺_ℓ is acyclic by A2, but acyclicity alone is not strong enough to imply continuum strong causality. An acyclic sequence can still return arbitrarily close to its starting point without forming an exact loop. To prevent almost-closed causal curves from surviving the continuum limit, we require a quantitative strengthening of A2.

This strengthening is supplied by **H6′ — uniform local transport sparsity**, stated in §2.2: there exists σ > 0, independent of refinement level ℓ, such that for every admissible commitment-transport sequence γ_ℓ = (s_0, s_1, …, s_n) at level ℓ, any two distinct committed states s_i ≠ s_j along γ_ℓ satisfy d_ℓ(s_i, s_j) ≥ σ_ℓ, where σ_ℓ is the minimum distinguishability scale at level ℓ; and the rescaled quantity σ̃_ℓ := λ_ℓ · σ_ℓ has a strictly positive continuum lower bound liminf_{ℓ→∞} σ̃_ℓ = σ_* > 0. (Here λ_ℓ denotes the refinement scale used to identify P_ℓ with the continuum-limit metric structure; see §3 and §4.)

H6′ states, informally, that distinct committed substrate states cannot become arbitrarily close, after rescaling, along a single admissible causal trajectory. It is a *quantitative no-almost-return* condition. Unlike A2 (which forbids exact returns), H6′ controls the rate at which distinct sequence members can approach each other under refinement.

Lemma 7.2 (Quantitative substrate acyclicity). *Under A2, H5, and H6', there exists δ > 0 such that for every refinement level ℓ, every admissible commitment-transport sequence γ_ℓ = (s_0, s_1, …, s_n), and every pair of indices i < j with s_i ≠ s_j, the rescaled continuum distance satisfies λ_ℓ · d_ℓ(s_i, s_j) ≥ δ. Equivalently, no admissible commitment-transport sequence contains an arbitrarily small near-cycle in the refinement limit.*

Proof. A2 supplies exact acyclicity: no admissible commitment-transport sequence may revisit the same committed substrate state. However, exact acyclicity alone does not rule out near-return behaviour, in which a sequence visits distinct states whose pairwise rescaled distances become arbitrarily small under refinement. H6′ supplies the required quantitative strengthening.

Let λ_ℓ denote the refinement scale by which d_ℓ is compared to the continuum metric. By H6′, for all distinct committed states s_i ≠ s_j lying on the same admissible commitment-transport sequence at level ℓ,

$\lambda_\ell \cdot d_\ell(s_i, s_j) \geq \lambda_\ell \cdot \sigma_\ell = \tilde{\sigma}_\ell$.

Since $\liminf_{\ell\to\infty} \tilde{\sigma}_\ell = \sigma_* > 0$, there exists $\ell_0 \in \mathbb{N}$ such that for all $\ell \geq \ell_0$,

$\tilde{\sigma}_\ell \geq \sigma_*/2$.

Define $\delta := \sigma_*/2$. Then for all sufficiently refined levels $\ell \geq \ell_0$, every pair of distinct committed states on the same admissible transport sequence satisfies $\lambda_\ell \cdot d_\ell(s_i, s_j) \geq \delta$ in the rescaled continuum metric.

For the finitely many coarse levels $\ell < \ell_0$, define for any compact region $K \subseteq \mathscr{M}_\infty$:

$\delta_0(K) := \min_{\ell < \ell_0} \min_{\gamma_\ell \subset K} \min_{i < j, s_i \neq s_j} \lambda_\ell \cdot d_\ell(s_i, s_j)$,

where the inner minima are taken over admissible bounded transport sequences contained in (the level-ℓ preimage of) K. By finite distinguishability (A1) and bounded local degree, each of the finitely many $\gamma_\ell \subset K$ at each $\ell < \ell_0$ contributes only finitely many pairs, so the minimum is over a finite set and is positive. Replacing δ by $\delta'(K) := \min(\delta, \delta_0(K)) > 0$ gives a positive lower bound that is uniform across all refinement levels $\ell \in \mathbb{N}$ but possibly depends on the compact region K.

Note that the lemma as proved is *locally uniform*: the bound δ is uniform across refinement levels for any fixed compact $K \subseteq \mathscr{M}_\infty$, but the constant $\delta'(K)$ may depend on K. This local formulation is what the Theorem 3 proof requires (which only needs the bound in a neighbourhood of the strong-causality test-point x). A globally uniform δ — i.e., δ_0 independent of K — would require an additional global control on the substrate at coarse levels, which is not supplied here.

Thus admissible commitment-transport sequences cannot contain arbitrarily small near-cycles in the rescaled continuum metric on any compact region. This completes the proof of Lemma 7.2.

7.3 Statement of Theorem 3

Theorem 3 (Strong causality of the limit). *Under A2, H1, H5, H6, H6', and the conclusion of Theorem 1, the continuum limit ($\mathscr{M}_\infty$, d_∞, $\prec_\infty$) is strongly causal in the sense of §7.1.*

7.4 Proof

The proof requires more than finite-level acyclicity. A2 prevents exact closed commitment loops, but continuum strong causality also requires exclusion of *almost-closed* causal curves. Lemma 7.2 supplies precisely this exclusion: under H6′, it gives a uniform lower bound $\delta > 0$ (on any fixed compact region), in the rescaled continuum metric, on the distance between distinct committed states lying on the same admissible transport trajectory.

The proof also requires that *causal curves limit to causal curves* under refinement convergence — a property not guaranteed by ordinary (metric) Gromov–Hausdorff convergence alone. We invoke *Lorentzian Gromov–Hausdorff convergence* in the sense of Müller (2022) and Minguzzi–Sämann (2023), which combines metric GH with Kuratowski convergence of the causal preorder *≺ℓ → ≺∞*. Under this stronger convergence, discrete causal sequences {γ_ℓ}*ℓ in (P_ℓ, d_ℓ, ≺ℓ) limit to continuum causal curves in (ℳ∞, d∞*, ≺_∞); conversely, every continuum causal curve is the Lorentzian-GH limit of some discrete causal sequence. This is the mode of convergence implicitly required by Theorems 1, 2, 3 and made explicit here.

Suppose for contradiction that ℳ_∞ fails to be strongly causal at some x ∈ ℳ_∞. Then there is a neighbourhood U of x such that every sub-neighbourhood V ⊆ U admits a causal curve γ_V starting and ending in V but leaving U in between.

Take a sequence V_n ⊆ U with diam(V_n) → 0. The corresponding causal curves γ_{V_n} have endpoints converging to x but leave U. By Lorentzian GH convergence (Theorem 1 + the causal-preorder convergence ≺ℓ → ≺∞), each γ_{V_n} is the Lorentzian-GH limit of a sequence of discrete admissible commitment-transport sequences {γ_{V_n, ℓ}}_ℓ in (P_ℓ, d_ℓ, ≺_ℓ).

Apply Lemma 7.2 to the compact region K := cl(U) (the closure of U). The lemma gives δ′(K) > 0 such that for sufficiently large n, diam(V_n) < δ′(K)/4 with δ′(K) as in Lemma 7.2 (measured in the rescaled continuum metric). The discrete sequence γ_{V_n, ℓ} therefore has endpoints within δ′(K)/4 of each other in the rescaled metric λ_ℓ · d_ℓ for all sufficiently large ℓ. But the sequence also leaves U in between, so it contains intermediate nodes at rescaled distance ≥ diam(U)/2 from its endpoints.

This means the discrete commitment-transport sequence γ_{V_n, ℓ} starts at some s_0, travels to a node s_m with λ_ℓ · d_ℓ(s_0, s_m) ≥ diam(U)/2, and returns to a node s_N with λ_ℓ · d_ℓ(s_0, s_N) < δ′(K)/4. By Lemma 7.2 (quantitative acyclicity from H6′) applied to K, the nodes s_0 and s_N must be separated by at least δ′(K) in the rescaled metric — contradicting λ_ℓ · d_ℓ(s_0, s_N) < δ′(K)/4 < δ′(K).

Therefore no such failure-of-strong-causality at x can occur. Since x was arbitrary, ℳ_∞ is strongly causal everywhere. This completes the proof of Theorem 3.

7.5 Where the proof concentrates its risk

The proof of Theorem 3 concentrates its risk in H6′ — the uniform local transport-sparsity hypothesis. Lemma 7.2 is now a clean conditional theorem: A2 + H5 + H6′ → uniform lower bound δ > 0 on pairwise distances of distinct sequence members. The mathematical content of the lemma is mechanical once H6′ is granted; the substrate-engineering question — *which* discrete TPB models satisfy H6′ — is the locus of the residual risk.

H6′ is not derived in this paper. It is the *quantitative* strengthening of A2 needed to lift finite-level acyclicity to continuum strong causality, and the question of whether and how it follows from more primitive substrate-level principles is flagged in §14. A natural conjecture is that H6′ follows from a strengthened form of A1 (finite distinguishability with uniform lower bound on the minimum scale) combined with bounded local degree and refinement compatibility; that derivation is not supplied here.

The risk-concentration in H6′ is qualitatively different from the risk-concentration in H6. H6 is a *covering-count* condition (a metric-geometric property of the substrate at each level), readily testable in any concrete substrate model. H6′ is a *trajectory-sparsity* condition (a property of the admissible-transport-sequence structure at each level), which depends on the dynamics of commitment-transport in a more delicate way. Both are external regularity hypotheses, but H6′ is closer in spirit to a Hawking-style chronology-protection condition at the substrate level.

---

8. Theorem 4 — Local no-flux foliations

8.1 Statement

Theorem 4 (Local no-flux foliation existence). *Under BCB, Lemma 6.1.a of the companion paper ($\partial_\mu J^\mu = 0$ on $\mathcal{M}_\infty$), and the local-nonvanishing condition that $J^\mu \neq 0$ in some neighbourhood of every point of interest, every $x \in \mathcal{M}_\infty$ admits an admissible normal neighbourhood $U_x \subseteq \mathcal{M}_\infty$ foliated by integral curves of J^μ with zero lateral transport flux through the lateral boundary of any compact transport tube $\mathcal{U} \subseteq U_x$ bounded by Cauchy sections.*

8.2 Proof

The continuum bit-current J^μ is, by Lemma 6.1.a of the companion paper, divergence-free on $\mathcal{M}_\infty$: $\partial_\mu J^\mu = 0$. Under the local-nonvanishing hypothesis, J^μ is a non-vanishing C^k vector field on a neighbourhood of x, where k is the regularity class of the continuum limit (this is discussed in §14).

By the standard existence-and-uniqueness theorem for ordinary differential equations (Picard–Lindelöf, applied to the integral curves of J^μ), there exists a flow φ_t generated by J^μ on some open neighbourhood U_x of x, defined for $|t| < \varepsilon(x)$.

Consider any compact transport tube $\mathcal{U} \subseteq U_x$ defined as the union of integral curves of J^μ over a small Cauchy section Σ_0 transverse to J^μ at x. By construction, J^μ is tangent to the lateral boundary $\partial\mathrm{lat}(\mathcal{U})$, *so the lateral flux* $\int_{\partial_{\mathrm{lat}}(\mathcal{U})} J^\mu n_\mu \, d\Sigma = 0$.

For any pair of Cauchy sections Σ_1, Σ_2 bounding $\mathcal{U}$, the divergence theorem applied to the divergence-free current J^μ over $\mathcal{U}$ gives

$$\int_{\Sigma_1} J^\mu n_\mu \, d\Sigma = \int_{\Sigma_2} J^\mu n_\mu \, d\Sigma,$$

which is the section-invariance of the integrated bit count $Q(\mathcal{U})$ used in Lemma 6.1.b of the companion paper.

This completes the proof of Theorem 4.

8.3 Local nonvanishing of J^μ

The local-nonvanishing hypothesis ("$J^\mu \neq 0$ in some neighbourhood of every point of interest") is a mild regularity condition that holds wherever there is ongoing commitment-transport activity. It fails only at points where the substrate is *quiescent* — i.e., where no admissible commitment-transport sequence passes through any neighbourhood. Such points are physically degenerate and are excluded from the analysis. The companion paper implicitly assumes non-quiescence throughout; we make the assumption explicit here.

8.4 Globality

Theorem 4 is stated and proved in *local* form, mirroring the local-form statement of Lemma 6.1.d in the companion paper. The global form — existence of a foliation by no-flux tubes on all of $\mathcal{M}_\infty$ — would require additional global integrability conditions on J^μ (no nodal sets, simple connectedness of the integral-curve atlas, etc.) and is *not* proved here. As in the companion paper, the local form combined with connectedness suffices for the downstream applications.

---

9. Theorem 5 — Continuum regularity

9.1 Statement

Theorem 5 (Continuum regularity). *Under A0–A4, BCB, H1–H7, H6', and the local-nonvanishing condition on J^μ, every admissible TPB refinement sequence $\{(P_\ell, d_\ell, \prec\ell, B_\ell)\}\{\ell\in\mathbb{N}\}$ admits a GH-convergent subsequence whose limit $(\mathcal{M}_\infty, d_\infty, \prec_\infty, J^\mu)$ is:*

1. *a complete connected pseudometric space (Theorem 1);*
2. *equipped with a continuous non-degenerate open convex double cone field $\mathcal{C}$ (Theorem 2 + the convexity-of-admissible-directions theorem of the sequential-interface-transport paper);*
3. *strongly causal in the sense of Kunzinger–Sämann (Theorem 3);*
4. *equipped with a divergence-free transport current J^μ admitting local no-flux foliations (Theorem 4 + Lemma 6.1.a of the companion paper).*

Therefore $(\mathcal{M}_\infty, d_\infty, \prec_\infty)$ is a strongly-causal Lorentzian length space (in the Kunzinger–Sämann sense) — i.e., a synthetic-Lorentzian-geometric structure with continuous but not a priori smooth metric content. The upgrade from this continuous structure to a C^k Lorentzian manifold for any prescribed $k \geq 2$, sufficient for the Malament–Hawking–King–McCarthy application in the companion paper's Lemma 5.2, is the subject of §10.

9.2 Proof

The proof is the conjunction of Theorems 1–4 above plus the convexity theorem of the sequential-interface-transport paper. Each clause of the statement is established by the cited theorem. The aggregate structure $(\mathcal{M}_\infty, d_\infty, \prec_\infty, J^\mu)$ satisfies the input requirements of the companion paper's Lemma 5.1 (cone invariance) and Lemma 5.2 (cone regularity), and the no-flux-tube construction of Lemma 6.1.d. This completes the proof of Theorem 5.

9.3 Significance

Theorem 5 supplies the *continuous* version of the regularity assumptions identified in §14 of the companion paper as the "principal remaining mathematical risk" of the geometry programme. The companion paper's machinery, however, operates on *smooth* Lorentzian manifolds — Malament–Hawking–King–McCarthy is stated for at least C^2 Lorentzian metrics. The gap between Theorem 5's continuous output and the companion paper's smooth input is closed in §10 (Theorems 6, 7, 8 below). With those upgrade results in place, the full chain runs:

A0–A4 + BCB + H1–H9 → (Theorem 5) → continuous Lorentzian length space → (§10, Theorems 6, 7, 8) → C^k Lorentzian manifold for any prescribed $k \geq 2$ → (companion paper Lemma 5.1 + 5.2) → invariant cone field → (companion paper §6.1) → conformal class fixed → (companion paper Theorem 1L) → metric Lorentzian structure → (companion paper Theorem 2) → invariant interval → (companion paper Theorem 3) → boost structure.

The remaining mathematical risk shifts one layer further down: from "does the continuum limit exist with the right regularity?" (answered conditionally by this paper and §10) to "do substrate models with the H6, H6′, H7, H8, H9 properties exist?" (the substrate-engineering question flagged in §14).

9.4 Forward pointer to the regularity upgrade

This section establishes that the limit $(\mathcal{M}_\infty, d_\infty)$ is a complete pseudometric space (Theorem 1) and that the cone field is continuous (Theorem 2). The companion paper assumes throughout that the continuum is a *smooth* 4-manifold equipped with a smooth Lorentzian metric. The gap between "continuous cone field on complete pseudometric space" (output of Theorem 5) and "C^k cone field on smooth manifold" (input to Malament–HKM) is closed in §10 below, under two additional regularity hypotheses (H8, H9) and using machinery from the Lorentzian-length-space literature (Sämann 2016, Burtscher 2015, Kunzinger–Sämann 2018). The bridge result is Theorem 8.

9.5 Construction of the τ-distance and KS-axiom verification

Theorem 5 asserts that $(\mathcal{M}_\infty, d_\infty, \prec_\infty)$ *is a Lorentzian length space in the sense of Kunzinger–Sämann (2018). The KS framework requires a time-separation function* $\tau: \mathcal{M}_\infty \times \mathcal{M}_\infty \to [0, \infty]$ satisfying:

(K1) $\tau(x, y) > 0 \iff x \ll y$ (chronology); (K2) $\tau(x, y) \geq 0$ with equality if $x \leq y$ but $x \nprec y$ (causality); (K3) reverse triangle inequality: $\tau(x, z) \geq \tau(x, y) + \tau(y, z)$ for $x \leq y \leq z$; (K4) lower semicontinuity in (x, y) on the chronology relation.

We construct τ_∞ from refinement limits.

Discrete τ_ℓ. For each $\ell \in \mathbb{N}$ and $x, y \in P_\ell$ with $x \prec_\ell y$, define

$$\tau_\ell(x, y) := \lambda_\ell \cdot \sup\{ n : (x = s_0 \prec_\ell s_1 \prec_\ell \cdots \prec_\ell s_n = y) \text{ is an admissible commitment-transport sequence} \},$$

i.e., the maximum number of TPB ticks along any admissible transport sequence from x to y, rescaled by λ_ℓ to continuum units. For $x \nprec_\ell y$, set $\tau_\ell(x, y) := 0$.

Continuum τ_∞. Define $\tau_\infty\colon \mathcal{M}_\infty \times \mathcal{M}_\infty \to [0, \infty]$ by

$$\tau_\infty(x_\infty, y_\infty) := \sup_{\varepsilon > 0} \limsup_{\ell \to \infty} \sup_{x_\ell \in B_{d_\ell}(x_\infty, \varepsilon),\, y_\ell \in B_{d_\ell}(y_\infty, \varepsilon)} \tau_\ell(x_\ell, y_\ell),$$

where the inner suprema are taken over GH-isometric ε-approximating sequences of x_∞ and y_∞ in (P_ℓ, d_ℓ). This is an extremum over the entire Lorentzian GH-equivalence class of approximating sequences, which is the correct way to obtain a value independent of any particular subsequence choice: lim sup of τ_ℓ along a single sequence is in general sensitive to subsequence selection, but the *extremum over all ε-approximations* in the GH-equivalence class is well-defined. (The construction parallels the standard treatment in Müller 2022, §3 of Lorentzian GH limits of length / time-separation functions.)

Verification of KS axioms.

- (K1) Chronology: $\tau_\infty(x_\infty, y_\infty) > 0$ iff there exist GH-approximating sequences x_ℓ, y_ℓ with $\tau_\ell(x_\ell, y_\ell)$ bounded below uniformly in ℓ, which (by the discrete construction) is equivalent to the existence of strict commitment-transport sequences from x_ℓ to y_ℓ — equivalently, to $x_\infty \ll_\infty y_\infty$.
- (K2) Causality: $\tau_\infty(x_\infty, y_\infty) \geq 0$ by construction. The boundary case $x_\infty \leq_\infty y_\infty$ but $x_\infty \nprec_\infty y_\infty$ corresponds to GH-limit identifications without strict transport, giving $\tau_\infty = 0$.
- (K3) Reverse triangle inequality: at the discrete level, $\tau_\ell(x, y)$ is defined as a *supremum* of TPB-tick counts over admissible transport sequences from x to y. Given $x \leq_\ell y \leq_\ell z$, any admissible sequence γ_{xy} from x to y of length n_{xy} and any admissible sequence γ_{yz} from y to z of length n_{yz} concatenate (by transitivity of $\prec_\ell$) to an admissible sequence $\gamma_{xz} := \gamma_{xy} \circ \gamma_{yz}$ from x to z of length $n_{xy} + n_{yz}$. Taking $\sup_{\gamma_{xy}}$ on the left and then $\sup_{\gamma_{yz}}$ gives $\tau_\ell(x, z) \geq \tau_\ell(x, y) + \tau_\ell(y, z)$ (the sup-of-concatenations is at least the sum of individual sups — super-additivity of the supremum operation). Passing to the limit via Lorentzian GH convergence, $\tau_\infty(x_\infty, z_\infty) \geq \tau_\infty(x_\infty, y_\infty) + \tau_\infty(y_\infty, z_\infty)$.
- (K4) Lower semicontinuity: the sup-over-GH-approximations definition is automatically lower semicontinuous in (x_∞, y_∞) on the chronology relation, since sup is lower semicontinuous and the family of GH-approximations is monotone in ε.

Therefore $(\mathcal{M}_\infty, d_\infty, \prec_\infty, \tau_\infty)$ satisfies the Kunzinger–Sämann axioms for a Lorentzian length space. This is the precise sense in which Theorem 5 establishes a synthetic Lorentzian length space structure on the continuum limit.

10. The regularity upgrade

10.1 Why the upgrade is needed

Theorem 5 establishes that admissible TPB refinement converges to a *continuous* Lorentzian length space in the sense of Kunzinger and Sämann (2018) — a synthetic-Lorentzian-geometric structure ($\mathcal{M}_\infty$, d_∞, $\prec_\infty$, J^μ) *equipped with a continuous non-degenerate double cone field* $\mathcal{C}$ *satisfying strong causality and admitting local no-flux transport foliations. The continuous cone field induces, via the construction of Lemma 10.1 below, a continuous Lorentzian metric* g_∞ on $\mathcal{M}_\infty$.

The companion paper's machinery, however, requires *smooth* Lorentzian structure. Specifically:

- The Malament–Hawking–King–McCarthy theorem invoked at Lemma 5.2 of the companion paper is classically stated for C^2 Lorentzian metrics (Malament 1977; Hawking–King–McCarthy 1976; cf. Hawking–Ellis, *The Large Scale Structure of Space-Time*, 1973, §6.4; Beem–Ehrlich–Easley, *Global Lorentzian Geometry*, 2nd ed., 1996).
- Corollary 6.1.e of the companion paper (conformal-factor fixing) uses $\partial_\mu J^\mu = 0$ in its strong differentiable form, which requires J^μ to be C^1.
- The local no-flux tube construction of Lemma 6.1.d of the companion paper invokes Picard–Lindelöf existence-and-uniqueness, which requires local Lipschitz regularity on J^μ.

The synthetic continuum produced by Theorem 5 therefore does not directly feed the companion paper's machinery — there is a regularity gap.

This section closes the gap. Under one additional refinement-regularity hypothesis (H8) and the natural smooth-structure existence assumption (H9), we prove that the continuous structure of Theorem 5 admits a C^k Lorentzian-metric realisation for any prescribed $k \geq 2$, with C^k null cone field, C^k transport current, and causal structure agreeing with the synthetic $\prec_\infty$ in the sense of Sämann (2016). C^2 is sufficient for the companion paper's Lemma 5.2 application.

10.2 Strategy and external machinery

The upgrade proceeds in two stages.

Stage 1 — cone field to C^0 metric (Lemma 10.1). Given the continuous cone field $\mathcal{C}$ on $\mathcal{M}_\infty$, construct a unique (up to global units) continuous Lorentzian metric $g_\infty \in C^0(\mathcal{M}_\infty; T\mathcal{M}_\infty \otimes T\mathcal{M}_\infty)$ whose null cone at each point is $\mathcal{C}(x)$. The scale of g_∞ is fixed by the transport-density preservation argument of the companion paper (Proposition T1, Corollary 6.1.e).

Stage 2 — C^0 metric to C^k metric (Theorems 6, 7). Refine the C^0 metric to Hölder continuity using H8 (Theorem 6), then mollify the Hölder metric to obtain smooth approximants whose causal structure agrees with the original in the Sämann–Chruściel–Grant sense (Theorem 7). The two steps together produce a C^k metric for any prescribed $k \geq 2$.

The external machinery used:

- **Sämann (2016), "Global hyperbolicity for spacetimes with continuous metrics," *Ann. Henri Poincaré* 17, 1429–1455.** The foundational paper showing that basic causal-structure results — global hyperbolicity, Cauchy hypersurface theory, time functions, etc. — extend to C^0 Lorentzian metrics. The key result for our purposes is the *stability of causal structure under C^0 perturbation*: for two C^0 Lorentzian metrics g, g' with $\|g - g'\|_{C^0}$ sufficiently small, the causal futures $J^+_g(x)$ *and* $J^+_{g'}(x)$ coincide outside a controllable error set.
- **Chruściel and Grant (2012), "On Lorentzian causality with continuous metrics," *Class. Quantum Grav.* 29, 145001.** Earlier foundational work. Crucially identifies an obstruction: for general C^0 metrics, chronology and causality relations can diverge (a pathology that does not occur for C^1 or smoother metrics). Our cone field is Hölder by H8 (Theorem 6), which is sufficient to rule out this pathology.
- **Kunzinger and Sämann (2018), "Lorentzian length spaces," *Ann. Glob. Anal. Geom.* 54, 399–447.** The synthetic framework. Establishes the cone-to-metric correspondence at the C^0 level used in Lemma 10.1, and characterises when a Lorentzian length space arises from a continuous Lorentzian metric.
- **Burtscher (2015), "Length structures on manifolds with continuous Riemannian metrics," *NY J. Math.* 21, 273–296.** The Riemannian analogue. Demonstrates that continuous metrics induce well-defined length structures whose metric topology equals the manifold topology, and develops the mollification techniques we adapt to the Lorentzian setting in Theorem 7.

10.3 Additional hypotheses

Two additional hypotheses are required for the upgrade. Both are external regularity assumptions on the refinement family, in the same sense as H6 and H7 of §2.2.

- **H8 — Refinement Hölder compatibility.** *There exist $\alpha > 0$ and $K < \infty$ such that for every $\ell \in \mathbb{N}$ and every $x, y \in P_\ell$, the discrete cone-direction sets $\mathcal{C}_\ell(x)$, $\mathcal{C}_\ell(y)$ satisfy*

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

  where d_H denotes the Hausdorff distance on the space of admissible direction sets in $T(P_\ell)$ and $\varepsilon_\ell \to 0$ as $\ell \to \infty$.

- **H9 — Smooth-structure existence.** *The topological 4-manifold structure inherited by $(\mathcal{M}_\infty, d_\infty)$ from Theorem 1 admits at least one compatible C^∞ atlas.*

H8 strengthens H5 (refinement compatibility) by demanding Hölder uniformity of the cone-direction stability across refinement levels. The Hölder exponent α is a substrate-engineering parameter; for the C^k upgrade with $k \geq 2$ to succeed via Theorem 7 we require only $\alpha > 0$.

H9 sidesteps a real but subtle obstruction. In dimension 4, by results of Donaldson (1983) and Freedman (1982), there exist topological 4-manifolds with no smooth structure (e.g., the E8-manifold) and topological 4-manifolds with multiple inequivalent smooth structures (the exotic $\mathbb{R}^4$ phenomena). The present paper does not derive smooth-structure existence from substrate principles — it assumes it via H9. The substrate-engineering question of *which* discrete substrate models yield smoothable continuum limits in dimension 4 is flagged in §14 as an open problem distinct from the regularity-upgrade machinery developed here.

10.4 Lemma 10.1 — C^0 metric realisation

Lemma 10.1 (C^0 metric realisation). *Under H1–H7 and H9, the continuous non-degenerate cone field* $\mathcal{C}$ *of Theorem 2 is the null cone of a unique (up to global units) continuous Lorentzian metric*

$$g_\infty \in C^0(\mathcal{M}_\infty; T\mathcal{M}_\infty \otimes_{sym} T\mathcal{M}_\infty)$$

of signature (1, 3). The scale of g_∞ *is fixed globally by Proposition T1 of the companion paper.*

Proof. Fix a compatible C^∞ atlas $\{(U_\alpha, \varphi_\alpha)\}$ on $\mathcal{M}_\infty$ (existing by H9). In each chart U_α, the tangent bundle is trivialised: $T_x \mathcal{M}_\infty \cong \mathbb{R}^4$ for $x \in U_\alpha$. The cone field $\mathcal{C}$ at x is a continuous family of open convex non-degenerate double cones in $\mathbb{R}^4$, varying continuously with x by Theorem 2.

Pointwise cone-to-metric construction. For each fixed x, the set $\mathcal{C}(x) \subset T_x \mathcal{M}_\infty$ is an open convex double cone with non-degenerate interior. The construction of a signature-(1,3) quadratic form on $T_x \mathcal{M}_\infty$ whose null cone is $\mathcal{C}(x)$ is the C^0 analogue of the foundational fact underlying Malament's theorem: the conformal class of a Lorentzian metric on a 4-manifold is determined by its null cone field. The explicit construction proceeds via the dual cone or via polarisation on the cone interior; see Malament (1977, §3) and Beem–Ehrlich–Easley (1996, §3.1) for the smooth-metric version, and Sämann (2016, §2) for the C^0 extension. The construction is well-defined on open convex *non-degenerate* double cones — the non-degeneracy hypothesis (positive interior) is essential, and convexity (inherited via §6.4 from the convexity-of-admissible-directions theorem of the sequential-interface-transport paper) ensures the construction is single-valued up to overall scale. Strict convexity at the cone boundary is *not* required, since the quadratic form is recovered from the interior structure rather than from boundary properties. The result is a unique (up to positive scale) symmetric bilinear form $g_\infty(x)$ of signature (1, 3) with null cone $\mathcal{C}(x)$. This is the C^0 analogue of the Schur-lemma uniqueness argument used in §7.2 of the companion paper.

Global scale fixing. The scale ambiguity at each x — a positive real factor — is reduced to a single global constant by Proposition T1 and Corollary 6.1.e of the companion paper: transport-density preservation forces the conformal factor relating different observers' bit-counting conventions to be a single global unit constant. The unit constant is absorbed into the definition of c.

Continuity of g_∞. The cone-to-quadratic-form map $\mathcal{C}(x) \mapsto g_\infty(x)$ is continuous on the space of non-degenerate convex double cones in $\mathbb{R}^4$ (with the Hausdorff topology) into the space of signature-(1,3) symmetric bilinear forms on $\mathbb{R}^4$ (with the Euclidean topology). Continuity of $\mathcal{C}$ (Theorem 2) therefore lifts to continuity of g_∞.

Coordinate-change consistency. Across overlapping charts $U_\alpha \cap U_\beta$, the tangent-bundle transition functions are C^∞ (by H9). The cone field $\mathcal{C}$ is defined intrinsically on $T\,\mathcal{M}_\infty$ and therefore transforms covariantly under chart changes; the same is true of g_∞ by construction. The metric is therefore globally well-defined on $\mathcal{M}_\infty$.

The resulting g_∞ is a globally-defined C^0 Lorentzian metric on $\mathcal{M}_\infty$ realising the null cone field $\mathcal{C}$. This completes the proof of Lemma 10.1.

10.5 Theorem 6 — Hölder regularity

Theorem 6 (Hölder regularity, locally uniform). *Under H1–H9, the continuous Lorentzian metric g_∞ of Lemma 10.1 is of class $C^{0,\alpha}_{loc}(\mathcal{M}_\infty)$, where $\alpha > 0$ is the Hölder exponent of H8: every point $x_\infty \in \mathcal{M}_\infty$ admits a neighbourhood U on which g_∞ is uniformly Hölder of exponent α, with Hölder constant K(U) depending on U. The same applies to the cone field $\mathcal{C}$ and (where defined) to the transport current J^μ.*

Proof. H8 gives Hölder uniformity of the discrete cone-direction sets across refinement levels. Specifically, for any $x, y \in P_\ell$ at sufficiently large ℓ,

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

By Theorem 2 and an Arzelà–Ascoli-type argument applied to the family $\{\mathcal{C}_\ell\}$ indexed by ℓ, the Kuratowski limit $\mathcal{C}: \mathcal{M}_\infty \to$ (non-degenerate convex double cones) inherits Hölder continuity with the same exponent α: for any $x_\infty, y_\infty \in \mathcal{M}_\infty$,

$$d_H(\mathcal{C}(x_\infty), \mathcal{C}(y_\infty)) \leq K \cdot d_\infty(x_\infty, y_\infty)^\alpha.$$

Lifting Hölder regularity to g_∞ on compact neighbourhoods. The cone-to-quadratic-form map $\mathcal{C}(x) \mapsto g_\infty(x)$ of Lemma 10.1 is *locally Lipschitz* between the space of non-degenerate convex double cones (with the Hausdorff metric, bounded away from degeneracy) and the space of signature-(1,3) symmetric bilinear forms (with the operator-norm metric). This is a finite-dimensional algebraic-geometric fact: the normalisation map taking a cone to its quadratic-form representative is smooth on the open subset of *non-degenerate* cones, with Lipschitz constants depending on the lower bound on the cone's opening angle.

Local uniform non-degeneracy. Theorem 2 + the convexity-of-admissible-directions theorem of the sequential-interface-transport paper establish *pointwise* non-degeneracy of the cone at each $x_\infty \in \mathcal{M}_\infty$, but not *uniform* non-degeneracy across $\mathcal{M}_\infty$ (which may be non-compact). To upgrade pointwise non-degeneracy to a uniform lower bound on a neighbourhood, we use continuity of the cone field (Theorem 2) combined with continuity of the opening-angle functional. The opening angle $\omega(\mathcal{C}(x_\infty))$ — defined, e.g., as the infimum over future-directed unit timelike vectors v of $|g_{Eucl}(v, \mathcal{C}(x_\infty)^c)|$, where g_{Eucl} is any auxiliary Euclidean reference metric on the chart — is a continuous functional on the space of non-degenerate convex double cones in the Hausdorff topology. Composing the Hausdorff-continuous cone field $\mathcal{C}$ with the continuous opening-angle functional gives a continuous function $\omega(x_\infty) := \omega(\mathcal{C}(x_\infty)) > 0$ on $\mathcal{M}_\infty$. For every $x_\infty \in \mathcal{M}_\infty$ there is a compact neighbourhood $U \ni x_\infty$ on which ω attains its minimum $\omega(U) > 0$ by standard compactness, giving a uniform lower bound on the cone opening angle on U.

On U the cone-to-form map is therefore uniformly Lipschitz with constant $L(\omega(U))$ depending only on the lower bound $\omega(U)$.

Composing the Hölder-α cone field $\mathcal{C}$ (restricted to U) with the $L(\omega(U))$-Lipschitz cone-to-form map gives $g_\infty|_U \in C^{0,\alpha}(U)$ *with Hölder constant* $K(U) := L(\omega(U)) \cdot K_{\mathcal{C}}(U)$. *Since this holds in a neighbourhood of every point,* $g_\infty \in C^{0,\alpha}_{loc}(\mathcal{M}_\infty)$.

Global uniformity is not asserted. If $\mathcal{M}_\infty$ is compact (e.g., for cosmological 4-manifolds with compact spatial sections), local Hölder regularity upgrades to global uniform Hölder by compactness covering. For non-compact $\mathcal{M}_\infty$, the Hölder constant may degrade at infinity; this does not affect downstream applications since the §10 mollification of Theorem 7 is also constructed locally (chart-by-chart, with partition-of-unity assembly), so locally uniform Hölder is the natural input.

Hölder regularity of J^μ. The transport current J^μ is, by Lemma 6.1.a of the companion paper and the refinement-stability arguments of §3, the continuum limit of the discrete bit-currents on P_ℓ. H8 (which controls cone-direction stability) together with H5 (which controls integrated bit-count stability) lifts Hölder regularity to J^μ on the same local neighbourhoods U, in the same Hölder class.

This completes the proof of Theorem 6.

10.6 Theorem 7 — Smooth approximation

Theorem 7 (Smooth approximation). *For every* $k \in \mathbb{N}$ *with* $k \geq 1$ *and every* $\varepsilon > 0$, *the* $C^{0,\alpha}$ *Lorentzian metric* g_∞ *of Theorem 6 admits a* C^k *Lorentzian approximation* $g_\infty^{(k,\varepsilon)}$ *on* $\mathcal{M}_\infty$ *satisfying:*

1. *C^0-closeness:* $\|g_\infty^{(k,\varepsilon)} - g_\infty\|_{C^0(\mathcal{M}_\infty)} < \varepsilon$;
2. *Signature preservation:* $g_\infty^{(k,\varepsilon)}$ *has signature (1, 3) globally;*
3. *Causal-structure agreement: the chronology and causality relations* $\ll_{(k,\varepsilon)}$, $\leq_{(k,\varepsilon)}$ *of* $g_\infty^{(k,\varepsilon)}$ *agree with* $\ll_\infty$, $\leq_\infty$ *in the Kuratowski-convergence sense as* $\varepsilon \to 0$, *by the* C^0*-stability theorem of Sämann (2016).*

Proof. The construction is standard mollification adapted to the Lorentzian-length-space setting, following the Riemannian construction of Burtscher (2015, §3) and the Lorentzian extension of Kunzinger–Sämann (2018, §5).

Local mollification. In each chart $(U_\alpha, \varphi_\alpha)$ of the C^∞ atlas guaranteed by H9, choose a C^∞ bump function $\eta: \mathbb{R}^4 \to \mathbb{R}_{\geq 0}$ with $\mathrm{supp}(\eta) \subset B(0, 1)$, $\int_{\mathbb{R}^4} \eta = 1$, and define the scale-$\delta$ mollifier $\eta_\delta(x) := \delta^{-4} \eta(x/\delta)$. Define the mollified metric in chart coordinates by

$$g_\infty^{(k,\delta)}(x) := \int_{\mathbb{R}^4} \eta_\delta(x - y) \cdot g_\infty(y)\, dy.$$

For y near x in the same chart, the integrand uses $g_\infty(y)$ directly; for y outside the chart, we use a partition-of-unity argument standard in the construction of smooth approximations on manifolds (cf. Burtscher 2015, §3.2).

Smoothness. The mollification convolves the $C^{0,\alpha}$ metric g_∞ with the C^∞ kernel η_δ; the result is C^∞ in the chart coordinates, hence at least C^k for any prescribed k.

C^0-closeness. Standard mollification estimates give

$$\|g_\infty^{(k,\delta)}(x) - g_\infty(x)\| \leq \sup_{y \in B(x, \delta)} \|g_\infty(y) - g_\infty(x)\| \leq K \cdot \delta^\alpha,$$

using the Hölder regularity of g_∞ from Theorem 6. Choosing $\delta = \delta(\varepsilon)$ with $K \cdot \delta^\alpha < \varepsilon$ gives $\|g_\infty^{(k,\delta)} - g_\infty\|_{C^0} < \varepsilon$ as required.

Signature preservation. Signature (1, 3) of $g_\infty^{(k,\delta)}$ for sufficiently small δ follows from the openness of signature in the C^0 topology — the set of signature-(1,3) symmetric forms is open in the space of symmetric forms on $\mathbb{R}^4$, in direct parallel with the openness of positive-definiteness used in the Riemannian construction of Burtscher (2015, §3.2). Signature (1, 3) is preserved chart-by-chart by C^0-small perturbations, and the partition-of-unity assembly (which is a convex combination of chart-mollified metrics weighted by smooth bump functions summing to 1) preserves the convex-combination invariance of signature: a convex combination of signature-(1, 3) symmetric bilinear forms on the same vector space, all close to a common signature-(1, 3) reference, remains signature (1, 3). Globally on $\mathcal{M}_\infty$ the partition-of-unity assembled $g_\infty^{(k,\delta)}$ is therefore a smooth (C^k) Lorentzian metric of signature (1, 3).

Causal-structure agreement. The principal substantive ingredient. By the C^0-stability theorem of Sämann (2016, Theorem 1.5), for any two C^0 Lorentzian metrics g, g' on a topological 4-manifold with $\|g - g'\|_{C^0} < \eta$, *the causal future $J^+_g(x)$ is contained in a controllable enlargement of $J^+_{g'}(x)$, and conversely; moreover, the chronological future $I^+_g(x)$ and $I^+_{g'}(x)$ coincide outside a set whose measure vanishes as $\eta \to 0$. Applied with $g = g_\infty$* and $g' = g_\infty^{(k,\delta)}$, the causal structures agree in the Kuratowski-convergence sense as $\delta \to 0$.

The Chruściel–Grant (2012) divergence pathology — where chronology and causality relations can split for C^0 metrics — does *not* occur here, because g_∞ is $C^{0,\alpha}$ (Theorem 6) rather than merely C^0. Hölder regularity is sufficient to ensure the chronology and causality relations coincide on the limit, which is what we need for causal-structure stability to lift cleanly to the mollified metric.

This completes the proof of Theorem 7.

10.7 Theorem 8 — Bridge to the companion paper

Theorem 8 (Bridge theorem). *Under H1–H9 and BCB, the continuum limit* ($\mathcal{M}_\infty$, d_∞, $\prec_\infty$, J^μ) *of admissible TPB refinement admits, for every prescribed* $k \geq 2$, *a* C^k *Lorentzian metric* $g_\infty^{(k)}$ *on* $\mathcal{M}_\infty$ *and a smooth transport current* $J^{(k)\mu}$ *such that:*

1. *the null cone field of* $g_\infty^{(k)}$ *coincides with* $\mathcal{C}$ *(up to* C^0 *error controllable by k);*
2. *the causal structure of* $g_\infty^{(k)}$ *agrees with* $\prec_\infty$ *in the Kuratowski sense;*
3. $J^{(k)\mu}$ *is* C^{k-1}*, satisfies* $\nabla^{g_\infty^{(k)}}_\mu J^{(k)\mu} = 0$ *(covariant divergence-freeness with respect to* $g_\infty^{(k)}$*), and admits the local no-flux foliations of Theorem 4 with full Picard–Lindelöf regularity.*

The companion paper's Lemma 5.2 (cone regularity), §6.1 (transport-density preservation derivation), §6.3 (proof of Theorem 1L), and downstream chain through Theorems 1G, 2, 3 of the companion paper therefore apply with full smoothness inputs.

Proof. By Theorem 7 with $k \geq 2$, the metric $g_\infty^{(k)} := g_\infty^{(k, \delta_k)}$ for suitably chosen $\delta_k > 0$ is a C^k Lorentzian metric on $\mathcal{M}_\infty$ satisfying clauses 1 and 2 directly.

For clause 3: this is the most delicate part of the bridge. We construct $J^{(k)\mu}$ in two steps — naive mollification followed by a covariant correction.

Step 3a — Naive mollification. Hölder regularity of J^μ (Theorem 6) lifts to C^{k-1} regularity under mollification with a C^∞ kernel of scale δ (the regularity gap of one order is standard in mollification theory). Let $\tilde{J}^{(k)\mu}$ denote the chart-by-chart mollification of J^μ assembled by partition of unity, in parallel with the construction of $g_\infty^{(k)}$ in Theorem 7.

Step 3b — Coordinate vs covariant divergence. The naive mollification $\tilde{J}^{(k)\mu}$ satisfies the *coordinate* divergence equation $\partial_\mu \tilde{J}^{(k)\mu} = 0$ in each chart, because the mollifier η_δ is symmetric and commutes with ∂_μ:

$$\partial_\mu \tilde{J}^{(k)\mu} = \partial_\mu (\eta_\delta * J^\mu) = \eta_\delta * (\partial_\mu J^\mu) = 0,$$

using divergence-freeness of J^μ in the C^0 limit (which is the distributional content of Lemma 6.1.a of the companion paper applied to the C^0 metric g_∞).

However, the companion paper's §6.1 conformal-factor argument uses the *covariant* divergence $\nabla^{g_\infty^{(k)}}_\mu J^{(k)\mu} = 0$, which involves Christoffel symbols of $g_\infty^{(k)}$:

$$\nabla^{g_\infty^{(k)}}_\mu J^{(k)\mu} = \partial_\mu J^{(k)\mu} + \Gamma^{(k)\mu}_{\mu\nu} J^{(k)\nu}.$$

Since the Christoffel symbols $\Gamma^{(k)\mu}_{\mu\nu} = \partial_\nu (\log \sqrt{|\det g_\infty^{(k)}|})$ depend on first derivatives of $g_\infty^{(k)}$ — which differ from corresponding derivatives of g_∞ (a $C^{0,\alpha}$ metric whose first derivatives are at best distributions) — the naive $\tilde{J}^{(k)\mu}$ satisfies $\nabla^{g_\infty^{(k)}}_\mu \tilde{J}^{(k)\mu} = \Gamma^{(k)\mu}_{\mu\nu} \tilde{J}^{(k)\nu} \neq 0$ in general. The deviation is bounded in C^0 norm by

$$\|\nabla^{g_\infty^{(k)}}_\mu \tilde{J}^{(k)\mu}\|_{C^0} \le \|\Gamma^{(k)}\|_{C^0} \cdot \|\tilde{J}^{(k)}\|_{C^0} = O(\delta_k^{\alpha-1})$$

if $\tilde{J}^{(k)}$ is C^0-bounded and $\Gamma^{(k)}$ grows at most as $\delta^{-1+\alpha}$ from mollifying a $C^{0,\alpha}$ metric. For $k \ge 2$ with $\alpha > 0$ sufficiently large ($\alpha > 0$ always; the deviation rate depends on the regularity gain), this deviation is finite but generally nonzero.

Step 3c — Covariant correction via wave-equation reduction. We construct $J^{(k)\mu} = \tilde{J}^{(k)\mu} + \delta J^{(k)\mu}$, where $\delta J^{(k)\mu}$ is a smooth vector field on $\mathcal{M}_\infty$ chosen to cancel the spurious divergence:

$$\nabla^{g_\infty^{(k)}}_\mu (\tilde{J}^{(k)\mu} + \delta J^{(k)\mu}) = 0, \text{ equivalently } \nabla^{g_\infty^{(k)}}_\mu \delta J^{(k)\mu} = f,$$

where $f := -\Gamma^{(k)\mu}_{\mu\nu} \tilde{J}^{(k)\nu}$ is a smooth (C^{k-1}) source.

The equation $\nabla_\mu \delta J^\mu = f$ is a single scalar equation for four components of δJ^μ — *underdetermined*. There are infinitely many solutions. The Riemannian Hodge analogue ("set δJ to be a gradient and reduce to Poisson") has a direct Lorentzian counterpart: set δJ to be a gradient and reduce to a wave equation. The wave equation is solvable on globally hyperbolic Lorentzian manifolds, with the natural a priori control in Sobolev rather than pointwise norms — a structural feature of hyperbolic vs. elliptic PDE theory, not a defect. We pursue this route.

Gauge choice. Seek $\delta J^{(k)\mu}$ of gradient form:

$$\delta J^{(k)\mu} := \nabla^{(k)\mu} \varphi = g_\infty^{(k)\mu\nu} \partial_\nu \varphi$$

for a smooth scalar φ on $\mathcal{M}_\infty$. Substituting:

$$\nabla^{(k)}_\mu \nabla^{(k)\mu} \varphi = \Box_{g_\infty^{(k)}} \varphi = f,$$

where $\Box$ is the d'Alembertian for $g_\infty^{(k)}$. This is the inhomogeneous *wave equation* for φ on the smooth globally hyperbolic Lorentzian manifold $(\mathcal{M}_\infty, g_\infty^{(k)})$.

Existence of φ. Global hyperbolicity of $(\mathcal{M}_\infty, g_\infty^{(k)})$ follows from Theorem 3 (continuum strong causality) combined with the causal-structure agreement of Theorem 7 (the C^k metric induces the same causal structure as g_∞). On a smooth globally hyperbolic Lorentzian manifold, the d'Alembertian $\Box_{g_\infty^{(k)}}$ is a normally hyperbolic operator, and by the theory of such operators on globally hyperbolic spacetimes — Bär, Ginoux, and Pfäffle, *Wave Equations on Lorentzian Manifolds and Quantization*, 2007, Theorem 3.3.8 (existence of advanced and retarded Green's operators) and Theorem 4.1.1 (well-posedness of the inhomogeneous Cauchy problem with smooth source and Cauchy data) — there exists a unique smooth solution φ to $\Box_{g_\infty^{(k)}} \varphi = f$ with prescribed (e.g., zero) Cauchy data on any Cauchy hypersurface. The metric $g_\infty^{(k)}$ is C^k for any prescribed k; choosing k sufficiently large so that $g_\infty^{(k)}$ is in the smooth-metric class assumed by Bär–Ginoux–Pfäffle is consistent with the rest of the §10 chain.

The vector field $\delta J^{(k)\mu} := \nabla^{(k)\mu} \varphi$ is then smooth (one derivative less than φ), and by construction satisfies $\nabla^{(k)}{}_\mu \delta J^{(k)\mu} = \Box\varphi = f$, *hence cancels the spurious divergence:* $\nabla^{g_\infty^{(k)}}{}_\mu (\tilde{J}^{(k)\mu} + \delta J^{(k)\mu}) = 0$ exactly.

Honest norm control. The natural a priori bound on φ is *not* a pointwise C^0 bound but an L^2 (Sobolev) bound, via the standard energy estimate for wave equations on globally hyperbolic spacetimes (Bär–Ginoux–Pfäffle 2007, §3.5; cf. also Bär 2015, "Green-hyperbolic operators on globally hyperbolic spacetimes"). Energy estimates control φ in H^1 in terms of f in L^2, with a constant depending on the Cauchy hypersurface and $g_\infty^{(k)}$ but not on δ_k. From $f \in L^2$ we get $\varphi \in H^1$ and hence $\delta J^{(k)} = \nabla\varphi \in L^2$. Pointwise C^0 control of $\delta J^{(k)}$ would require Sobolev embedding $H^s \hookrightarrow C^0$ in dimension 4, which needs $s > 2$, hence $f \in H^1$ and stronger source regularity than the natural L^2 output of mollifying a $C^{0,\alpha}$ metric.

In short: the wave-equation construction produces a smooth covariantly-divergence-free corrected current $J^{(k)\mu} := \tilde{J}^{(k)\mu} + \delta J^{(k)\mu}$, with *Sobolev/$L^2$* control on $\delta J^{(k)}$, not pointwise C^0 control. The pointwise C^0-smallness of $\delta J^{(k)}$ that would be needed for the companion paper's §6.1 conformal-factor argument to apply *directly* with pointwise-small perturbation is not established here. What is established is the existence and L^2-control. The companion paper's §6.1 derivation depends on *integrated* transport conservation over admissible tubes, which is robust to L^2-norm perturbations of the divergence — this robustness claim is the real content of the bridge, and it is honestly flagged in §10.10 below as the principal residual technical gap of the §10 chain.

Conclusion for clause 3. The corrected current $J^{(k)\mu} := \tilde{J}^{(k)\mu} + \delta J^{(k)\mu}$ is C^{k-1} (since both summands are), is exactly covariantly divergence-free with respect to $g_\infty^{(k)}$, and satisfies all the regularity requirements for the companion paper's Lemma 6.1.d (Picard–Lindelöf existence-uniqueness on integral curves) for $k \geq 2$.

Application of the companion paper's machinery. The classical Malament (1977) and Hawking–King–McCarthy (1976) results apply to C^2 Lorentzian manifolds (cf. Hawking–Ellis 1973, §6.4; Beem–Ehrlich–Easley 1996, §3.10). The companion paper's Lemma 5.2 application is therefore valid for $g_\infty^{(2)}$. The Schur-lemma quadratic-form uniqueness argument of §7.2 of the companion paper applies to C^2 metrics without modification. The transport-density preservation derivation of §6.1 of the companion paper uses $\nabla^{g_\infty^{(k)}}{}_\mu J^{(k)\mu} = 0$ in a form requiring C^1 $J^{(k)\mu}$, which is supplied by clause 3 with $k \geq 2$.

The Kuratowski-sense causal-structure agreement (clause 2) ensures that $\prec_\infty$*, used by Lemma 5.1 of the companion paper, is the same causal structure that the smooth metric* $g_\infty^{(k)}$ generates. The companion paper's chain therefore operates on the smooth structure $(\mathcal{M}_\infty, g_\infty^{(k)})$ without modification.

This completes the proof of Theorem 8.

10.8 The residual smoothness gap

Theorem 8 produces, for *each* prescribed $k \geq 2$, a C^k metric $g_\infty^{(k)}$. It does *not* produce a single metric that is simultaneously C^k for all k — i.e., a C^∞ metric in the strong sense. The sequence $\{g_\infty^{(k)}\}_{k\in\mathbb{N}}$ would need additional analytic input (coordinated mollification scales, harmonic-coordinate analysis, elliptic regularisation, or similar) to converge to a smooth limit. None of those ingredients is supplied here.

This is consistent with what is achievable in the broader low-regularity Lorentzian geometry literature. Burtscher (2015), Sämann (2016), Chruściel–Grant (2012), and Kunzinger–Sämann (2018) all achieve prescribed-order smooth approximation but not joint smoothness. The leap to C^∞ typically requires either a parabolic-equation or elliptic-PDE smoothing argument (Ricci flow, mean-curvature flow, or analogous Lorentzian techniques) that is the subject of ongoing work in synthetic Lorentzian geometry.

For the present paper's purpose — bridging Theorem 5 to the companion paper's Lemma 5.2 — C^2 is sufficient. C^∞ is therefore a desideratum, not a requirement. It is flagged in §14 as the principal residual mathematical risk of the regularity-upgrade machinery.

10.9 Where the risk concentrates

The §10 upgrade concentrates its risk in four distinct places, listed in increasing order of severity:

1. **H8 (refinement Hölder compatibility)** is the substantive new hypothesis introduced by §10. The substrate-engineering question — *which* discrete TPB models satisfy H8 with what exponent α — is open. The §10 chain is parametric in α: for any $\alpha > 0$, Theorems 6, 7, 8 yield a C^k metric for any prescribed k. So H8 is a *threshold* hypothesis: if it holds at all (with any positive exponent), the upgrade succeeds.
2. **H9 (smooth-structure existence)** is potentially restrictive due to the Donaldson–Freedman exotic-4-manifold phenomena. For the physical cosmological 4-manifolds (compact perturbations of Minkowski or FLRW), smooth structures are unique up to diffeomorphism and H9 is essentially automatic. For more exotic topologies, H9 is a genuine restriction.
3. **The weak-norm control of the covariant-divergence correction (§10.7 Step 3c, expanded in §10.10)** is the most substantive residual technical concern. The wave-equation construction of §10.7 Step 3c produces a smooth, *exactly* covariantly-divergence-free corrected current $J^{(k)\mu}$, but with only Sobolev/L^2-norm control on the correction $\delta J^{(k)\mu}$, not pointwise C^0 control. The companion paper's §6.1 conformal-factor derivation must be robust to weak-norm divergence corrections rather than relying on pointwise smallness. Verifying this robustness — i.e., re-reading the companion paper's §6.1 to confirm it depends only on integrated quantities — is the principal residual technical refinement of the §10 chain. See §10.10 below.
4. **The C^∞ desideratum** (§10.8) is not delivered. C^2 suffices for the companion paper, but C^∞ is the natural physical expectation. Closing this gap requires additional machinery beyond the present paper.

These four risks are independent: H8 is about substrate-engineering, H9 is about 4-manifold topology, the weak-norm robustness question is about wave-equation PDE theory and a re-reading of the companion paper's §6.1, and the C^∞ gap is about analytic smoothing technique. The downstream chain works correctly modulo each of them.

10.10 The weak-norm-robustness gap (named technical refinement)

The wave-equation construction of §10.7 Step 3c rigorously establishes existence and L^2/Sobolev control of the divergence correction $\delta J^{(k)\mu}$, but does *not* establish a pointwise C^0 bound. This is a genuine technical gap in the bridge from Theorem 8 to the companion paper's §6.1 conformal-factor argument: the companion paper as written presumably operates on smooth manifolds with *pointwise* divergence-free smooth currents; whether it is robust to weak-norm (L^2, distributional, weighted-Sobolev) corrections of the divergence is a robustness question that must be verified by direct re-examination of the companion paper's §6.1.

We do not undertake that re-examination here, because (i) it is properly the content of a revision to the companion paper rather than to the present paper, and (ii) the robustness *should* hold on general grounds — the conformal-factor argument depends on integrated transport conservation over admissible tubes (an integrated quantity, not a pointwise quantity), and such integrated arguments are typically stable under L^2/Sobolev-norm perturbations of the divergence with appropriate compactness on the tubes. But "should hold" is not "has been verified," and the verification is a non-trivial technical task.

Candidate concrete fixes to the companion paper. A reader picking up this open problem should know which specific changes to the companion paper would close the gap. The most plausible candidates:

- *Restate companion paper Lemma 6.1.b's tube-integrated transport-conservation conclusion as an L^2-stability claim rather than a pointwise claim.* The existing Lemma 6.1.b presumably establishes $\int_{\text{tube}} J^\mu \, d\Sigma_\mu$ = constant for a pointwise smooth divergence-free current; the L^2-stable version would establish that this integral varies by at most $O(\|\nabla_\mu J^\mu\|_{L^2(\text{tube})})$ for currents that are divergence-free only in L^2.
- *Weaken the smoothness assumption on J^μ in companion paper §6.1 from C^1 to H^1 (or to the Sobolev class $W^{1,p}$ for some $p \geq 2$).* The conformal-factor argument should go through with the weaker regularity provided the integrated conservation law is the L^2-stable version above.
- *Restate companion paper Corollary 6.1.e's conformal-factor uniqueness as a uniqueness-up-to-L^2-error claim.* The conformal factor is then determined up to a controllable Sobolev-norm error rather than exactly, but the error vanishes as the source $f \to 0$, which is what the §10 mollification chain provides in the limit.

Any of these three companion-paper revisions, in combination with the wave-equation construction of §10.7, would close the bridge cleanly. The choice between them depends on which formulation is most natural inside the companion paper's existing §6 architecture and is a matter for the companion-paper revision rather than the present paper.

To be explicit: the §10 chain as developed in this paper is rigorous *up to* the wave-equation construction of Step 3c. The bridge to the companion paper's machinery is rigorous *given* that the companion paper's §6.1 argument is restated in an L^2/Sobolev-stable form along the lines of one of the three candidates above. Establishing this restated form is the named open problem of §10. We flag this honestly in:

- §13 (Falsification path G9) — if the companion paper's §6.1 is found to be non-robust to weak-norm corrections, the bridge weakens.
- §14 (open problem item on conformal-factor robustness) — as a substrate-engineering / re-reading task.
- The Theorem 8 status table entry — explicit acknowledgment that the bridge is conditional on this robustness.

This is the most substantive residual gap in the present paper. The architecture is correct; the wave-equation reduction is rigorous; the existence and smoothness of the correction is rigorous; only the weak-norm-to-companion-paper bridge step remains to be tightened.

---

11. Relationship to companion papers and external literature

11.1 Companion VERSF papers

This paper completes the geometry-emergence chain together with the following VERSF papers, summarised in dependency order:

- The *sequential-interface-transport* paper establishes finite propagation, admissible transport dynamics, the $K=7$ wheel-structure intertwiner used in the companion paper's Theorem 3 (R1 substrate proof, pending), and the convexity-of-admissible-directions theorem used in the present paper's §6.4.
- The *σ-duality* paper establishes the refinement functor structure $\boldsymbol{\mathcal{R}}$ used in §3.2 and in the companion paper §8.1.
- The *admissible-coarse-graining* paper establishes the coarse-graining map $\varphi: \mathcal{S}\_\text{substrate} \to \mathcal{M}\_\text{coarse}$, the substrate-to-continuum limit theorem used in Lemma 6.1.a of the companion paper, and the foliation lemma referenced (without dependency in this paper) in Lemma 6.1.d of the companion paper.
- The *BCB–VERSF synthesis* paper establishes Bit Conservation and Balance, distinguishability conservation on Fisher manifolds, and the gauge-derivation subsections referenced in §2.5 of the companion paper.
- The *companion Lorentzian-emergence paper* establishes Theorems 1L, 1G, 2, 3 and Proposition T1, conditional on the regularity conditions supplied by the present paper.

11.2 External literature — synthetic Lorentzian geometry

The mathematical framework most directly relevant to this paper is the *Lorentzian length space* setting introduced by Kunzinger and Sämann (Ann. Glob. Anal. Geom. 54, 2018, 399–447).

Lorentzian length spaces are the synthetic-Lorentzian-geometry analogue of Alexandrov spaces in the Riemannian setting: they allow GH-style convergence and Ricci-curvature bounds in a metric setting without requiring smoothness a priori. The continuum limit ($\mathscr{M}_\infty$, d_∞, $\prec_\infty$) of Theorem 5 is a Lorentzian length space in this sense.

Lorentzian Gromov–Hausdorff convergence has been developed by Müller (J. Geom. Phys. 178, 2022) and refined by Minguzzi and Sämann (J. Differ. Geom., to appear; arXiv preprint 2023). The convergence used in Theorems 1–2 of this paper is closely related but not identical: the joint convergence of metric structure and causal order, with the additional convergence of the transport current J^μ to Lemma 6.1.a's divergence-free continuum field, is specific to the VERSF setting.

11.3 External literature — low-regularity Lorentzian metrics (§10 machinery)

The §10 regularity upgrade imports machinery from the *low-regularity Lorentzian metric* literature, which studies when classical Lorentzian-geometric results extend to metrics of less than C^2 regularity. The key references:

- **Chruściel and Grant (2012), "On Lorentzian causality with continuous metrics," *Class. Quantum Grav.* 29, 145001.** Established that many classical causal-structure results fail or require modification at the C^0 level, but extend cleanly to Hölder regularity $C^{0,\alpha}$ for $\alpha > 0$. The Hölder threshold provided by H8 and Theorem 6 is precisely what is needed to avoid the C^0 pathologies identified by Chruściel–Grant.
- **Sämann (2016), "Global hyperbolicity for spacetimes with continuous metrics," *Ann. Henri Poincaré* 17, 1429–1455.** The foundational paper establishing that global hyperbolicity, Cauchy hypersurface theory, time functions, and basic causal-structure results extend to C^0 Lorentzian metrics. The key result for our purposes is the *stability of causal structure under C^0 perturbation*, which underlies Theorem 7 (smooth approximation preserving causal structure).
- **Burtscher (2015), "Length structures on manifolds with continuous Riemannian metrics," *NY J. Math.* 21, 273–296.** The Riemannian analogue. Establishes that continuous Riemannian metrics induce well-defined length structures whose metric topology coincides with the manifold topology, and develops the mollification techniques we adapt to the Lorentzian setting in Theorem 7.
- **Kunzinger and Sämann (2018), already cited above.** The synthetic-Lorentzian-geometry framework. The cone-to-metric correspondence used in Lemma 10.1 is the C^0 analogue of the smooth cone-to-metric correspondence; the synthetic LLS framework provides the conceptual setting in which the analogue is well-defined.

The §10 upgrade chain (Lemma 10.1 → Theorem 6 → Theorem 7 → Theorem 8) is a VERSF-specific application of this machinery: the Chruściel–Grant pathologies are avoided by H8 (Hölder regularity); the Sämann (2016) C^0-stability theorem is used in Theorem 7 to lift smoothing without breaking the causal structure; the Burtscher (2015) mollification technique provides the C^k smooth approximants. The novel contribution is not the machinery itself but the *application* to the VERSF refinement convergence, together with the explicit identification of which substrate-level hypotheses (H8, H9) are needed for the machinery to apply.

11.4 External literature — causal-set programme and CDT

The *causal-set programme* (Bombelli–Lee–Meyer–Sorkin, Phys. Rev. Lett. 59, 1987, 521–524; Sorkin, "Causal sets: discrete gravity," in *Lectures on Quantum Gravity*, 2003) asks the analogous question for Poisson sprinklings of fixed Lorentzian manifolds. The relationship to VERSF refinement convergence is discussed in §4.3. The *Hauptvermutung* of causal-set theory — that a causal set determines its manifold uniquely if at all — is a question this paper does not address, since the VERSF setup fixes the substrate and asks what continuum it generates rather than fixing a manifold and asking how many discrete structures yield it.

The *causal-dynamical-triangulations* programme (Ambjørn–Jurkiewicz–Loll, Phys. Rep. 519, 2012) takes yet another tack — Monte Carlo summation over discrete causal triangulations — and obtains 4-dimensional continuum-like behaviour in the IR limit numerically. The VERSF refinement-convergence approach is structurally closer to CDT than to BLMS causal-set theory, in that the discrete structure is *constructed* rather than sampled, but the convergence machinery is different.

11.5 External literature — Gromov–Hausdorff theory

The use of Gromov–Hausdorff precompactness for metric spaces follows Gromov (*Structures métriques pour les variétés riemanniennes*, 1981) and Burago–Burago–Ivanov (*A Course in Metric Geometry*, AMS GSM 33, 2001).

---

12. Dependency diagram

The overall geometry-emergence architecture in VERSF, with the present paper's contribution layered in, is:

```
Layer 1 — Substrate axioms (all parallel):
   A0, A1, A2, A3, A4, BCB

Layer 2 — Substrate refinement regularity (this paper, parallel):
   H1 (finite propagation)              ← prior VERSF work (A1, A4)
   H5 (refinement compatibility)         ← sequential-interface-transport
                                           and σ-duality papers
   H6 (uniform doubling)                 ← external regularity hypothesis
   H6′ (uniform local transport sparsity) ← external regularity hypothesis
   H7 (bounded combinatorial dimension)   ← external regularity hypothesis
   H8 (refinement Hölder compatibility)   ← external regularity hypothesis
(for §10)
   H9 (smooth-structure existence)        ← external regularity hypothesis
(for §10)

Layer 3 — Substrate convergence (this paper):
   Theorem 1 (GH precompactness)         ← H1, H6
   Theorem 2 (cone-field convergence)     ← H1, H5, H7 + H3, H4 (continuity)
   Lemma 7.2 (quantitative acyclicity)    ← A2, H5, H6′
```

```
   Theorem 3 (strong causality)           ← A2, H1, H5, H6, H6′, Theorem 1
                                            + Lorentzian GH (Müller 2022)
   Theorem 4 (local no-flux foliations)   ← BCB, Lemma 6.1.a of companion
                                            paper, local nonvanishing of J^μ

Layer 4 — Continuum regularity (continuous, this paper):
   Theorem 5 (continuum regularity,       ← Theorems 1–4 + convexity of
              continuous version)           admissible directions
                                            (sequential-interface-transport)

Layer 4.5 — Regularity upgrade (this paper, §10):
   Lemma 10.1 (C^0 metric realisation)    ← Theorem 5 + H9
   Theorem 6  (Hölder regularity)         ← Lemma 10.1 + H8
   Theorem 7  (smooth approximation)      ← Theorem 6 + Sämann (2016)
                                            causal-stability theorem
   Theorem 8  (bridge to companion paper) ← Theorems 5, 6, 7 + H9

Layer 5 — Geometric primitives (companion paper):
   invariant causal cones (Lemma 5.1)     ← Theorem 8 + H2
   transport-density preservation         ← BCB, A0, A1
   (Proposition T1)                         (Lemmas 6.1.a/b/c of companion)

Layer 6 — Main Lorentz results (companion paper):
   Theorem 1L (local Lorentz emergence)   ← H1–H5 + Prop T1 + Lemma 5.1
        │
        ├── Theorem 1G (global Poincaré)  ← Theorem 1L + flat continuum limit
        ├── Theorem 2  (invariant interval) ← Theorem 1L
        └── Theorem 3  (boost structure)  ← Theorem 1L + R1 (conjecture)

Layers 7-11 — Field-theoretic structure (companion paper §12 and downstream).
```

Three structural observations:

1. The present paper occupies Layers 2 through 4.5, sitting between substrate axioms (Layer 1) and the geometric primitives (Layer 5) used by the companion paper. The Lorentz emergence chain of the companion paper begins at Layer 5; this paper closes the gap between Layer 1 and Layer 5.
2. The previously identified "principal mathematical risk" — does the substrate produce a smooth continuum with the right regularity? — splits cleanly across Layers 4 and 4.5. Layer 4 (Theorem 5) establishes the *continuous* version, conditional on H6, H6′, H7. Layer 4.5 (Theorems 6, 7, 8) lifts this to *C^k for any prescribed k ≥ 2*, conditional on the additional H8, H9. The smooth-input requirements of the companion paper (Malament–HKM at Lemma 5.2; Lemma 6.1.d with Picard–Lindelöf) are satisfied at Layer 4.5.
3. H6, H6′, H7, H8, H9 are all introduced as *external* regularity hypotheses on the refinement family. The substrate-engineering question — which discrete models satisfy them — is the principal remaining mathematical risk *after* the present paper. H6′ and H9 are qualitatively distinct from the others: H6′ controls trajectory sparsity (closer in spirit to a substrate-level chronology-protection condition than to a geometric covering count), and H9 concerns the topology of the 4-manifold limit, intersecting the Donaldson–Freedman theory in dimension 4.

13. Falsification paths and quantitative tests

The framework is falsified if any of the following are observed. Paths G1–G5 concern the continuous-regularity layer (Theorems 1–5); paths G6–G9 concern the regularity upgrade (§10); F1–F7 of the companion paper are inherited.

(G1) Refinement non-convergence. If admissible TPB refinement sequences in a physically realised substrate model are shown to be GH-non-convergent (no convergent subsequence exists for any candidate substrate-engineering choice), Theorem 1 is vacuous and the framework cannot generate a continuum limit. This is the strongest falsifier of the present paper.

(G2) Failure of uniform doubling. If physically realised substrate models exhibit refinement families with unbounded doubling constants ($D_\ell \to \infty$), H6 fails and Theorem 1 does not apply. The continuum limit (if it exists) would have infinite Hausdorff dimension at some scale.

(G3) Failure of bounded combinatorial dimension. If physically realised substrate models exhibit refinement families with unbounded local branching ($N_\ell \to \infty$), H7 fails and Theorem 2's cone-field convergence breaks down — the continuum cone could have infinite directional structure at some scale.

(G4) Failure of uniform local transport sparsity (H6′). If physically realised substrate models exhibit refinement families with $\tilde{\sigma}_\ell \to 0$ — i.e., distinct committed states along admissible transport sequences become arbitrarily close in the rescaled continuum metric under refinement — H6′ fails, Lemma 7.2's uniform separation bound δ collapses, and almost-closed causal curves can survive the continuum limit. Theorem 3's strong causality conclusion fails on the substrate models in question. This would also falsify the no-wormhole conjecture of §11.2 of the companion paper.

(G5) Quiescent substrate regions. If physically realised substrate models exhibit macroscopic quiescent regions ($J^\mu \equiv 0$ on extended subsets of $\mathcal{M}_\infty$), Theorem 4's local-nonvanishing hypothesis fails on those regions and no-flux foliations cannot be constructed there. The companion paper's Lemma 6.1.d would then need additional assumptions to apply.

(G6) Failure of refinement Hölder compatibility (H8). If physically realised substrate models exhibit refinement families with non-Hölder cone-direction stability (no positive exponent α works), H8 fails and the upgrade from Theorem 5's continuous structure to Theorem 6's Hölder structure breaks. Without Hölder regularity, the Chruściel–Grant (2012) C^0 pathologies (chronology relation $\ll$ and causality relation $\leq$ diverging from each other) can occur in the continuum limit. The standard formulation of Malament–HKM does not survive these pathologies: Malament's theorem determines the conformal class of the Lorentzian metric from the *chronological* relation $\ll$, but when $\ll$ and $\leq$ diverge (as Chruściel–Grant show is possible for C^0 metrics), the input to Malament's reconstruction is ambiguous — neither $\ll$ nor $\leq$ alone gives a well-posed input, and the theorem's conclusion is no longer applicable. The companion paper's Lemma 5.2 application therefore *may fail entirely* under H8 failure, not merely weaken. The §10 smooth-approximation chain through Theorems 7, 8 collapses, and the geometry-emergence chain in VERSF is broken at the Lorentzian-emergence layer. This is a hard falsification, not a soft one.

(G7) Exotic 4-manifold structure (failure of H9). If a physically realised substrate model produces a topological 4-manifold continuum limit with no compatible C^∞ atlas — i.e., a Donaldson–Freedman exotic 4-manifold — H9 fails and Lemma 10.1 does not give a globally consistent metric. For cosmological 4-manifolds (compact perturbations of FLRW), this is empirically unlikely; for exotic substrate models it is a genuine restriction. Direct empirical test: precision measurements of cosmological 4-manifold smoothness invariants (Donaldson polynomials, Seiberg–Witten invariants), where applicable.

(G8) Failure of C^0-causal-stability assumption (Sämann 2016). If the C^0-stability theorem of Sämann (2016) — used in Theorem 7 to lift mollification without breaking causal structure — were shown to fail in the VERSF setting (e.g., the metric limit is C^0 but not Hölder, so Chruściel–Grant pathologies appear in the smoothing limit), the smooth-approximation chain collapses. This is closely related to G6 but distinct: G6 falsifies the input Hölder regularity; G8 falsifies the output causal-stability. Either failure breaks the §10 chain.

(G9) Failure of conformal-factor robustness to weak-norm divergence corrections. Theorem 8 produces a covariantly divergence-free J^{(k)μ} via the wave-equation construction of §10.7 Step 3c (gauge reduction δJ^μ = ∇^μ φ followed by □φ = f, solvable by Bär–Ginoux–Pfäffle 2007). The natural norm control on δJ^{(k)μ} is in Sobolev/L^2 norms, not pointwise C^0. The companion paper's §6.1 conformal-factor derivation must be restated in an L^2/Sobolev-stable form (see the three candidate fixes in §10.10) for the bridge to go through. If a direct re-reading of the companion paper's §6.1 shows the argument is non-robust to weak-norm corrections and cannot be restated in this form, the bridge is not merely weakened but *fails to be established* — closing it would require a substantive replacement (either a new §6.1 argument or a different bridge architecture between Theorem 8 and the companion paper's machinery). This is a softer falsification than G6/G8 in the sense that it concerns the bridge's verification rather than its existence in principle, but its resolution is structurally substantive, not cosmetic.

---

14. Limitations, open problems, and dependencies

Open problems flagged in this paper.

Open problems are grouped by which layer of the paper they concern: convergence/continuous layer (Theorems 1–5) or regularity-upgrade layer (§10, Theorems 6–8).

Convergence / continuous-regularity layer.

- **Substrate-level derivation of H6 and H7.** Uniform doubling and bounded combinatorial dimension are assumed as external regularity hypotheses on the refinement family. The natural conjecture is that A4-local-coupling with bounded local degree, plus refinement compatibility (H5), implies H6 and H7; proving this is the principal substrate-engineering problem for the continuous-regularity layer.

- **Substrate-level derivation of H6′ (uniform local transport sparsity).** Acyclicity alone (A2) does not imply strong causality in the continuum limit. A discrete causal path may avoid exact loops while still returning arbitrarily close to its starting point under refinement. The paper therefore introduces H6′, a uniform local transport-sparsity condition stating that distinct committed substrate states along the same admissible transport sequence remain separated by a positive rescaled distinguishability distance. Proving H6′ from substrate principles is one of the key substrate-engineering tasks. It may follow from a stronger form of A1 (finite distinguishability with uniform lower bound on the minimum scale) combined with bounded local degree and refinement compatibility, but that derivation is not supplied here.
- **Convexity-of-admissible-directions theorem.** This paper invokes the convexity theorem from the sequential-interface-transport paper as input to §6.4 / Theorem 5 clause (2). The status of that theorem in the underlying paper should be cross-checked; if it is itself conditional, Theorem 5 inherits that conditionality.
- **Global no-flux foliations.** Theorem 4 is stated and used in local form. Global existence of no-flux foliations on all of $\mathcal{M}_\infty$ requires additional integrability conditions on J^μ and is not proved here.
- **Identification with causal-set Hauptvermutung.** The relationship between VERSF refinement convergence and the causal-set Hauptvermutung is sketched in §4.3 and §11.4; a precise mathematical comparison theorem — when does a VERSF refinement limit coincide with a causal-set continuum limit, if at all — is open.

Regularity-upgrade layer (§10).

- **Substrate-level derivation of H8 (refinement Hölder compatibility).** H8 is introduced as an external Hölder-uniformity hypothesis on the refinement family. The substrate-engineering question — which discrete substrate models satisfy H8 with what exponent α — is open. Note that the §10 chain is parametric in α: for any $\alpha > 0$ the upgrade succeeds. So H8 is a *threshold* hypothesis; the substrate-engineering question is whether any positive α exists for physically realised substrates.
- **Substrate-level derivation of H9 (smooth-structure existence).** H9 sidesteps the Donaldson–Freedman exotic-4-manifold phenomenon. For cosmological 4-manifolds (compact perturbations of FLRW) H9 is essentially automatic; for exotic substrate models it is a genuine restriction. Substrate-engineering question: which discrete substrates yield smoothable continuum limits in dimension 4. This problem intersects 4-manifold topology and is qualitatively different from the H6–H8 substrate-engineering questions.
- **C^∞ regularity (§10.8).** Theorem 7 produces, for each prescribed $k \in \mathbb{N}$, a C^k metric $g_\infty^{(k)}$. It does not produce a single jointly-smooth metric. C^2 suffices for the companion paper's Malament–HKM application, so the C^∞ gap is a desideratum rather than a requirement, but it is the principal residual mathematical risk of the §10 machinery.
- **Sharpness of the Hölder exponent α.** Theorem 6 produces a $C^{0,\alpha}$ metric with α inherited from H8. Whether this α can be improved (e.g., to $C^{0,1}$ = Lipschitz, or to

$C^{1,\beta}$ for some $\beta > 0$) under stronger substrate-engineering hypotheses is an interesting refinement question. Lipschitz regularity would allow the Picard–Lindelöf existence-uniqueness for J^μ in Theorem 4 to be established without the Theorem 7 detour.
- **Robustness of the companion paper's §6.1 conformal-factor argument to weak-norm divergence corrections.** Theorem 8 establishes exact covariant divergence-freeness of $J^{(k)\mu}$ via the wave-equation construction of §10.7 Step 3c. The natural a priori control on the correction is Sobolev/L^2, not pointwise C^0. The companion paper's §6.1 derivation of the conformal factor operates on smooth manifolds with smooth divergence-free currents and uses integrated transport conservation; whether it is robust to L^2/Sobolev-norm corrections is a robustness question that has not been verified in detail. Closing this is a technical refinement involving careful re-reading of the companion paper's §6.1. This is the named gap of §10.10 and is the largest residual concern of the §10 chain.

Principal remaining structural challenge. Establishing H6, H6′, H7, H8, H9 from substrate principles is now the principal remaining structural challenge after this paper. The companion paper's §14 identified the continuum-limit regularity theorem as the principal mathematical risk in the geometry programme; the present paper provides that theorem conditional on these five external regularity hypotheses, and therefore relocates the risk cleanly to the substrate-engineering layer. The natural next-paper task is to identify discrete substrate classes for which the full chain H6 + H6′ + H7 + H8 + H9 holds, and to characterise — both theoretically and via concrete substrate models — the boundary of admissibility.

Status summary.

Result	Status
Theorem 1 (GH precompactness)	Proven, conditional on H1, H6 (H7 used at Theorem 2 layer for uniform dimension)
Theorem 2 (cone-field convergence)	Proven, conditional on H1, H5, H7 + H3, H4 (continuity from equivariance)
Lemma 6.3.a–c	Proven, conditional on H1, H5, H7 (with continuity in 6.3.c from H3, H4)
Lemma 7.2 (quantitative substrate acyclicity)	Proven from A2 + H5 + H6′; H6′ is a new external transport-sparsity hypothesis requiring substrate-level derivation; the bound is locally uniform on compact regions
Theorem 3 (strong causality of limit)	Proven, conditional on A2, H1, H5, H6, H6′, Theorem 1, and Lorentzian GH convergence (Müller 2022)
Theorem 4 (local no-flux foliations)	Proven, conditional on BCB, Lemma 6.1.a of companion, local nonvanishing of J^μ
Theorem 5 (continuum regularity, continuous version)	Proven, conditional on Theorems 1–4 + convexity theorem of sequential-interface-transport paper

Result	Status
τ_∞ construction (KS-axiom verification, §9.5)	Proven, conditional on Theorem 5 + Lorentzian GH convergence
Lemma 10.1 (C^0 metric realisation)	Proven, conditional on Theorem 5 + H9; uses cone-to-conformal-class correspondence (Malament 1977; Sämann 2016)
Theorem 6 (Hölder regularity, locally uniform)	Proven, conditional on Lemma 10.1 + H8; local rather than global uniform Hölder for non-compact 𝓜_∞
Theorem 7 (smooth approximation)	Proven, conditional on Theorem 6 + H9; imports Sämann (2016) Theorem 1.5 (C^0-stability of causal structure) as external machinery
Theorem 8 (bridge to companion paper, C^k for k ≥ 2)	Proven, by combination of Theorems 5, 6, 7 + H9 + wave-equation correction for exact covariant divergence-freeness (§10.7 Step 3c via gauge reduction to $\square\varphi = f$, Bär–Ginoux–Pfäffle 2007); norm control on the correction is Sobolev/L^2, not pointwise C^0; the bridge to the companion paper's §6.1 conformal-factor argument is conditional on §6.1's robustness to L^2/Sobolev-norm divergence perturbations (§10.10 weak-norm-robustness gap)
C^∞ regularity of g_∞	Open (C^k for any prescribed k achieved; joint smoothness not derived)
Global no-flux foliations	Open; local form sufficient for companion paper
H6 (uniform doubling)	External hypothesis; substrate-level derivation open
H6′ (uniform local transport sparsity)	External hypothesis; needed to exclude almost-closed causal curves and upgrade finite acyclicity to continuum strong causality
H7 (bounded combinatorial dimension)	External hypothesis; substrate-level derivation open
H8 (refinement Hölder compatibility)	External hypothesis; substrate-level derivation open
H9 (smooth-structure existence)	External hypothesis; intersects Donaldson–Freedman 4-manifold theory
Convexity-of-admissible-directions	Inherited from sequential-interface-transport paper

15. Conclusion

The companion paper *Structural Necessity of Lorentzian Geometry in VERSF* established that, given a smooth strongly-causal continuum equipped with invariant causal cones and conserved transport density, Lorentzian geometry is the unique stable admissible structure. The §14 dependency clarification of that paper identified the continuum-limit regularity theorem as the principal remaining mathematical risk.

The present paper supplies that theorem in two stages: a continuous version (Theorem 5) and a smooth version (Theorem 8, via §10).

The convergence chain runs in two stages. *Stage 1 (continuous regularity):* GH precompactness of admissible TPB refinement sequences under H1, H6 (Theorem 1; H7 enters at the next layer); Kuratowski convergence of discrete cone-direction sets to a continuous non-degenerate cone field under H1, H5, H7 (Theorem 2, with continuity from H3, H4); quantitative substrate acyclicity (Lemma 7.2, conditional on A2, H5, H6′) upgraded by Lorentzian GH convergence (Müller 2022) to continuum strong causality (Theorem 3); divergence-freeness of the bit-current plus local nonvanishing yielding local no-flux foliations (Theorem 4); and the conjunction is Theorem 5 — a continuous Lorentzian length space in the Kunzinger–Sämann sense, equipped with the τ-distance constructed in §9.5. *Stage 2 (smooth regularity, §10):* C^0 metric realisation (Lemma 10.1) under H9 via the Malament cone-to-conformal-class correspondence; locally uniform Hölder regularity (Theorem 6) under H8; C^k smoothing via mollification with causal-structure stability from Sämann (2016) (Theorem 7); and the bridge to the companion paper (Theorem 8), producing a C^k Lorentzian metric for any prescribed $k \geq 2$ and a covariantly divergence-free transport current (via wave-equation reduction $\delta J^\mu = \nabla^\mu \varphi$, $\Box\varphi = f$, Bär–Ginoux–Pfäffle 2007; §10.7 Step 3c, with Sobolev/L^2 norm control, and a named weak-norm-robustness gap in §10.10) — in particular $k = 2$ sufficient for the Malament–Hawking–King–McCarthy application of the companion paper's Lemma 5.2.

The principal mathematical risk in the VERSF geometry programme therefore shifts from continuum-limit regularity to the substrate-engineering question: which discrete substrate models satisfy the five external regularity hypotheses H6 (uniform doubling), H6′ (uniform local transport sparsity), H7 (bounded combinatorial dimension), H8 (refinement Hölder compatibility), and H9 (smooth-structure existence) introduced in §2.2, §7.2, and §10.3. The natural conjectures: H6, H7 follow from A4-local-coupling with bounded local degree plus refinement compatibility; H6′ follows from a strengthened A1 (uniform lower bound on the minimum distinguishability scale) combined with bounded local degree and refinement compatibility; H8 is a Hölder-uniformity refinement of H5 and may be derivable from finer refinement-control hypotheses; H9 is qualitatively different because it intersects 4-manifold topology and the Donaldson–Freedman theory. Proving these conjectures and isolating the substrate-engineering boundary is the natural next work-item.

The key methodological moves are: (i) the explicit separation of three convergence questions in Stage 1 (metric convergence requiring volume-growth control; causal-structure convergence requiring refinement compatibility; regularity convergence requiring substrate acyclicity *plus* uniform local transport sparsity); (ii) the further separation of the continuous-regularity stage (Theorem 5) from the smooth-regularity-upgrade stage (Theorems 6, 7, 8); and (iii) the explicit import of low-regularity Lorentzian metric machinery (Sämann 2016, Burtscher 2015, Chruściel–Grant 2012, Kunzinger–Sämann 2018) for the upgrade chain, with honest hedging about the C^∞ desideratum (§10.8) that goes beyond what current machinery delivers. The introduction of H6′ as the *quantitative* strengthening of A2 needed to exclude almost-closed causal curves in the continuum limit converts Lemma 7.2 from a vulnerable pigeonhole argument into a clean conditional theorem; the C^∞ gap (§10.8) is now the principal residual mathematical risk.

In particular, to be explicit about scope: this paper does not derive the discrete substrate. What it establishes is that *given* admissible refinement dynamics satisfying H1–H9 (with H6′ slotted into the convergence layer) plus BCB, the continuum limit is regular enough — at C^k for any prescribed k ≥ 2 — to support the Lorentz-emergence machinery of the companion paper. The substrate-engineering question — which discrete models actually satisfy H1, H5, H6, H6′, H7, H8, H9 — is the natural next paper.

The geometry-emergence chain in VERSF is, with the present paper added to the companion paper, *conditionally closed* modulo the external regularity hypotheses H6, H6′, H7, H8, H9 — i.e., from substrate axioms to emergent Lorentzian geometry under the assumption that these five regularity hypotheses hold for the refinement family in question. The remaining open problems are individually flagged in §14 and in the companion paper's §14.

The remaining work is to derive H6, H6′, H7, H8, H9 from substrate principles, close the C^∞ regularity gap of §10.8, prove the global form of Theorem 4 where needed, and complete the open problems carried over from the companion paper (H3, H4 derivation; No-Wormhole Conjecture; R1 substrate proof; BCB gauge-paper re-grounding).