Finite Distinguishability Packing in the VERSF Framework

Structural Factorisation of Admissible Sector Cardinality through Operational Density, Refinement Rank, and the Entropy-Partition Gap

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

Abstract

Prior VERSF work established:

(i) discrete ledger realizability under bounded closure support, (ii) uniqueness *and* quantization of the admissible entropy partition, (iii) a unique closure scale ξ with $k_{max} = 2\pi/\xi$, (iv) finite admissible spectral channels N_{spec} from the $\mathbb{Z}_7$ closure architecture, (v) refinement-stable admissibility under the finite-rank closure projector Ω_{max}.

From these alone, $|\Sigma(M)| < \infty$ follows immediately: bounded closure support admits only finitely many ledger-realizable states (§2.1), and admissible sectors are equivalence classes of ledger states under entropy quotient and refinement projection. The cardinality of admissible sectors is therefore already finite before any packing machinery is invoked. The substantive question is not *whether* $|\Sigma(M)|$ is finite — it is — but *what structurally controls it*.

The present paper supplies that structural content. We derive an explicit **factorised cardinality ceiling** whose every factor is an admissibility invariant:

$|\Sigma(M)| \leq \mathrm{Vol}_{op}(M) / \Delta_{op}^d$, $d = N_{spec} \cdot C_{max}$,

where $\Delta_{op} > 0$ is the operational distinguishability quantum (the entropy-partition gap, §2.3), $\mathrm{Vol}_{op}(M)$ is the d-dimensional operational Hausdorff measure of M, N_{spec} is the spectral-channel count (§2.4), and C_{max} is the maximum per-channel rank of Ω_{max} (§2.5). The dimensional exponent d reflects the geometry of the admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_{max})$, with the spectral-channel decomposition $\mathrm{Im}(\Omega_{max}) = \bigoplus_\alpha V_\alpha$ made explicit.

The result upgrades the prior finite-spectrum theorem from a structural finiteness claim to a quantitative **dimensional-packing factorisation**. It exposes the *structural origin* of finiteness — not its mere fact — by identifying which admissibility invariants control the ceiling and at what dimensional scaling. Each factor is independently testable through its own programme; the ceiling itself is an admissibility-invariant structural quantity of M.

Applied to particle physics, the factorisation supplies the missing mathematical bridge between closure-scale finiteness and finite particle-family structure. The charged-lepton generation bound

becomes interpretable as a realisation of the universal factorisation rather than a sector-specific coincidence, with each closed sector family inheriting its own dimensional exponent $d_{\mathscr{F}}$ from its operational sub-manifold $M_{\mathscr{F}}$, sharing the universal entropy gap Δ_{op} across all admissible sectors of M.

Contents

---

1. Introduction

The previous VERSF paper *Finite Admissible Spectra from Closure-Scale Distinguishability* established that stable admissible physical sectors terminate structurally under finite distinguishability, ledger completeness, refinement persistence, and closure-scale spectral termination.

Finiteness is not the claim of this paper. It follows directly from the discrete commitment ledger: bounded closure support admits only finitely many ledger-realizable states (§2.1), and admissible sectors are equivalence classes of ledger states under entropy quotient (§2.3) and refinement projection (§2.5), hence

$$|\Sigma(M)| \le N_{ledger}(M) < \infty.$$

The "infinite distinguishable sectors within bounded support" loophole familiar from ordinary analysis is closed by the discrete ledger alone, before any packing geometry enters.

What §2.1 does not supply is the **structural origin** of this finiteness. The crude bound $|\Sigma(M)| \leq N\_ledger(M)$ tells us how many ledger states fit, not:

- how the cardinality factorises through admissibility invariants;
- what dimensional exponent governs its scaling;
- whether the ceiling is itself a structural invariant of M;
- which closure-architecture quantities control the count.

The present paper supplies this factorisation. The substantive contribution is the **explicit factorised cardinality ceiling**

$$|\Sigma(M)| \leq Vol\_op(M) / \Delta\_op^d, \quad d = N\_spec \cdot C\_max,$$

whose every factor is an admissibility invariant. The packing geometry exposes which structural quantities of M control the cardinality and at what dimensional scaling, transforming a finiteness claim into a quantitative dimensional-packing claim.

Scope clarification. The hypotheses required for the factorisation are stronger than discrete ledger alone — they include the finite spectral rank of Ω_max (§2.5), the spectral-channel decomposition (§2.4), and the quantized entropy partition (§2.3). Discrete ledger gives finiteness; the additional structural inputs give the *factorisation*.

The argument is structurally reminiscent of coding-theoretic sphere-packing bounds, finite phase-space distinguishability, finite-dimensional Hilbert-space distinguishability, and entropy-capacity limits — translated into the VERSF closure framework. Crucially, the bound here is **ontological rather than statistical**: it follows from admissibility, not from sampling.

The paper proceeds in five steps:

1. finite operational d-volume $Vol\_op(M) < \infty$ from §2.1 (Lemma 4.0);
2. a strictly positive operational distinguishability quantum $\Delta\_op > 0$ from §2.1 (Lemma 4.1);
3. bounded admissible spectral support $\sigma(S) \subseteq [0, k\_max]$ (Lemma 5.1);
4. the spectral-channel decomposition $Im(\Omega\_max) = \bigoplus\_\alpha V\_\alpha$ with $d \leq N\_spec \cdot C\_max$ (Lemma 6.1);
5. the sphere-packing bound $|\Sigma(M)| \leq Vol\_op(M) / \Delta\_op^d$ (Lemma 7.1).

---

2. Prior Structural Results

The argument depends on six previously established VERSF results, summarised here for self-containment.

2.1 Discrete Ledger Realizability

The closure ontology admits only ledger-realizable states. The commitment ledger is discrete, and within bounded closure support only finitely many ledger-realizable states exist. Continuous physically realized distinguishability is forbidden by the discrete commitment ledger.

2.2 Unique Closure Scale

The closure-scale theorem establishes a unique scale

$k\_max = 2\pi / \xi$,

with ξ the unique admissible closure length. Modes above k_max are phase-redundant rather than physically distinguishable.

2.3 Entropy Admissibility (Uniqueness *and* Quantization)

The admissible entropy programme established that

operational equivalence ≡ closure equivalence ≡ entropy equivalence

must coincide uniquely, and that the resulting partition $\Pi\_S(M)$ is **operationally quantized** — cells are discretely separated rather than densely accumulating. Uniqueness and quantization are independent results within §2.3: the former asserts that the partition is determined by admissibility; the latter asserts that its cells admit a positive minimum separation.

2.4 Finite K = 7 Closure Architecture

The admissible closure geometry possesses:

- finite $\mathbb{Z}_7$ Fourier structure,
- finite projected closure algebra,
- a finite number $N\_spec < \infty$ of admissible spectral channels.

2.5 Ω_max-Projected Dynamics

Projected closure operators Ω_max possess:

- finite admissible support,
- **finite spectral rank** $R\_max = \mathrm{rank}(\Omega\_max) < \infty$,
- stable refinement projection.

The finite rank of Ω_max is established in the projected closure dynamics paper of the VERSF programme via the $\mathbb{Z}_7$ closure architecture, **without reference to admissible sector cardinality** — so its invocation here does not produce circularity.

2.6 Prior Finite Spectrum Theorem

Stable admissible sectors terminate structurally. Formal factorisation of the cardinality bound was not yet derived. The present paper supplies the factorisation with correct dimensional scaling.

3. Definitions

Definition 3.1 (Closure manifold). The closure manifold M is the admissibility-complete distinguishability substrate equipped with closure scale ξ, admissible spectral channels N_spec, and projected closure operator Ω_max.

Definition 3.2 (Admissible sector). An admissible sector S ∈ Σ(M) is a ledger-realizable structure that is simultaneously

(a) distinguishable, (b) entropy-consistent under the unique quantized partition Π_S, (c) refinement-stable under Ω_max, (d) spectrally admissible: σ(S) ⊆ [0, k_max].

Definition 3.3 (Operational distinguishability metric). Let D_op: Σ(M) × Σ(M) → ℝ_≥0 denote the operational distinguishability distance, induced by the entropy-equivalence quotient of §2.3.

Definition 3.4 (Distinguishability quantum). The distinguishability quantum

Δ_op := min { D_op(S_i, S_j) : S_i ≠ S_j, S_i, S_j ∈ Σ(M) }.

(The infimum is shown to be attained as a minimum in Lemma 4.1.)

Definition 3.5 (Refinement complexity). For S ∈ Σ(M),

C_ref(S) := min { dim V : V ⊆ Im(Ω_max), S ∈ V, V Ω_max-invariant }.

Definition 3.6 (Operational d-volume). Let 𝒜 := Im(Ω_max) and d := dim 𝒜. The operational d-volume is

Vol_op(M) := (2^d / ω_d) · ℋ^d_{D_op}(M ∩ 𝒜),

where ℋ^d_{D_op} denotes the d-dimensional Hausdorff measure under the operational metric D_op, and ω_d is the unit-ball volume in d-dimensional Euclidean space. The prefactor 2^d / ω_d normalises Vol_op so that Vol_op(M) / Δ_op^d is directly the cell count of a maximally tight Δ_op-pairwise-separated packing in 𝒜.

Definition 3.7 (Admissible packing number).

P(M) := max { |F| : F ⊆ Σ(M), D_op(S_i, S_j) ≥ Δ_op ∀ S_i ≠ S_j ∈ F }.

---

4. Finite Admissible Content and the Operational Distinguishability Quantum

Finiteness of |Σ(M)| and Vol_op(M) both flow from §2.1 (discrete ledger). The role of these lemmas in the paper is *not* to establish finiteness — that is immediate — but to fix the structural quantities (Δ_op, Vol_op) that enter the factorised ceiling of Theorem 8.1.

Lemma 4.0 (Finite admissible content under bounded closure support). *Let M be a closure manifold with bounded closure support. Then:*

(i) *the number of ledger-realizable states within M is finite;* (ii) *|Σ(M)| ≤ N_ledger(M) < ∞ (crude bound);* (iii) *the operational d-volume Vol_op(M) < ∞.*

Proof. By §2.1, the commitment ledger is discrete and bounded closure support admits only finitely many ledger-realizable states; denote this count N_ledger(M). Admissible sectors are equivalence classes of ledger states under entropy quotient (§2.3) and refinement projection (§2.5), so |Σ(M)| ≤ N_ledger(M) < ∞, establishing (i) and (ii).

For (iii), Vol_op(M) is the d-dimensional Hausdorff measure of M ∩ 𝓐 — the operational d-volume of the closure manifold restricted to the admissible subspace — *not* of the finite set Σ(M). (The latter would in fact be zero for d ≥ 1, since finite point sets carry zero d-dimensional Hausdorff measure.) The relevant geometric object is the continuous substrate M ∩ 𝓐 ⊆ 𝓐, into which the finitely many admissible sectors are embedded.

By the operational compatibility of D_op with the closure metric — Ω_max is continuous in the closure topology and inherits D_op-boundedness on bounded closure support — bounded closure support of M implies that M ∩ 𝓐 = Ω_max(M) is a D_op-bounded subset of the finite-dimensional Euclidean space 𝓐. Bounded subsets of finite-dimensional Euclidean spaces have finite d-dimensional Hausdorff measure, so 𝓗^d_{D_op}(M ∩ 𝓐) < ∞, and Vol_op(M) < ∞ by Definition 3.6.

Remark 4.0.1. Lemma 4.0(ii) already supplies finiteness of |Σ(M)|. The substantive content of Theorem 8.1 is therefore not finiteness, which is immediate from §2.1, but the **factorised cardinality ceiling** through admissibility invariants (Vol_op, Δ_op, N_spec, C_max). The crude bound N_ledger(M) is opaque about which structural quantities of M control the count; the factorised bound is structurally transparent.

Lemma 4.1 (Existence of a distinguishability quantum). *On any admissibility-complete closure manifold M with bounded closure support, Δ_op > 0.*

Proof. By Lemma 4.0(ii), Σ(M) is finite. The operational distinguishability metric D_op therefore takes only finitely many distinct values on Σ(M) × Σ(M). The set

$$\{ D_{op}(S_i, S_j) : S_i \neq S_j \in \Sigma(M) \} \subset \mathbb{R}_{>0}$$

is a finite subset of $\mathbb{R}_{>0}$, hence its infimum is attained at a strictly positive minimum.

Remark 4.1.1. This derivation uses only the discrete ledger (§2.1). §2.3 is invoked only to *identify* Δ_{op} with the entropy-partition gap — quantization of Π_S confirms that Δ_{op} is the universal operational gap between admissible entropy cells, not a quantity dependent on the choice of bounded closure region. This identification is what makes Δ_{op} an admissibility invariant of M rather than a region-local constant.

---

5. Closure-Scale Spectral Boundedness

The closure-scale theorem of §2.2 establishes the upper spectral edge $k_{max} = 2\pi / \xi$.

Lemma 5.1 (Finite admissible spectral support). *For every admissible sector $S \in \Sigma(M)$,*

$$\sigma(S) \subseteq [0, k_{max}].$$

Proof. If $\sigma(S)$ contained admissible content at $k > k_{max}$, the closure-scale theorem would be violated, contradicting the uniqueness of ξ.

Combined with the finite $\mathbb{Z}_7$ closure architecture of §2.4, the admissible spectral content of any sector decomposes into at most $N_{spec} < \infty$ admissible channels indexed within $[0, k_{max}]$.

---

6. Refinement Complexity and the Spectral-Channel Decomposition

We make explicit the spectral-channel decomposition of $\mathrm{Im}(\Omega_{max})$. The previous version conflated two possible readings of the $N_{spec} \cdot C_{max}$ factor; we now state which is intended.

Spectral-channel decomposition. The finite-rank operator Ω_{max} respects the $\mathbb{Z}_7$ spectral channel decomposition of §2.4: its image decomposes as a direct sum over admissible spectral channels,

$$\mathrm{Im}(\Omega_{max}) = \bigoplus_{\alpha = 1}^{N_{spec}} V_\alpha,$$

with each V_α the Ω_{max}-image within channel α. Define the per-channel rank $C_{max}^{(\alpha)} := \dim V_\alpha$, and

$$C_{max} := \max_\alpha C_{max}^{(\alpha)}.$$

Then the total admissible-subspace dimension satisfies

$$d := \dim \mathcal{A} = \sum_\alpha \dim V_\alpha \leq N_{spec} \cdot C_{max},$$

where $\mathcal{A} := \mathrm{Im}(\Omega_{max})$ is the admissible subspace.

Lemma 6.1 (Finite refinement complexity). *There exists $C_{max} < \infty$ such that for every admissible sector $S \in \Sigma(M)$,*

$$C_{ref}(S) \leq C_{max}, \text{ and } d \leq N_{spec} \cdot C_{max} < \infty.$$

Proof. By §2.5, $\mathrm{rank}(\Omega_{max}) < \infty$, established via the $\mathbb{Z}_7$ closure architecture independently of admissible sector counting. Refinement stability requires $\Omega_{max}(S) \equiv S$ in admissible distinguishability classes, so $S \in \mathrm{Im}(\Omega_{max})$. The spectral-channel decomposition splits $\mathrm{Im}(\Omega_{max})$ into the direct sum of finite-dimensional V_α; any $S \in V_\alpha$ has $C_{ref}(S) \leq \dim V_\alpha \leq C_{max}$. The total dimension bound follows.

Corollary 6.3 (Admissible sectors form a finite-dimensional family). *The admissible sector space $\Sigma(M)$ is contained in $\mathcal{A} = \mathrm{Im}(\Omega_{max})$ of dimension $d \leq N_{spec} \cdot C_{max}$.*

---

7. Finite Admissible Sector Density: a Sphere-Packing Argument

Metric structure of $\mathcal{A}$. The admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_{max})$ is a finite-dimensional vector space — the image of a finite-rank operator on the operational state space — and inherits a natural inner product from the entropy-equivalence quotient (§2.3). The operational metric D_{op} restricts to a Euclidean metric on $\mathcal{A}$, giving $\mathcal{A}$ the structure of a d-dimensional Euclidean space in which the standard ball-volume identity $\mathrm{vol}(B(r)) = \omega_d \cdot r^d$ holds. The sphere-packing argument below invokes this Euclidean structure. Were $\mathcal{A}$ endowed with a genuinely non-Euclidean metric (e.g. an ultrametric on a ledger-discrete lattice), the dimensional scaling Δ_{op}^d would persist but the geometric constant $2^d / \omega_d$ in Definition 3.6 would require modification.

Lemma 7.1 (Packing bound). *Let $d := \dim \mathcal{A} \leq N_{spec} \cdot C_{max}$. The admissible packing number satisfies*

$$P(M) \leq \mathrm{Vol}_{op}(M) / \Delta_{op}^d.$$

Proof. Endow $\mathcal{A}$ with the operational metric D_{op}. Let $F \subseteq \Sigma(M)$ be any pairwise-Δ_{op}-separated family, i.e. $D_{op}(S_i, S_j) \geq \Delta_{op}$ for $S_i \neq S_j \in F$. Around each $S \in F$ place an open D_{op}-ball $B(S, \Delta_{op} / 2)$ inside $\mathcal{A}$. By the triangle inequality and the pairwise separation, these balls are pairwise disjoint.

Each ball has d-volume $\omega_d \cdot (\Delta_{op} / 2)^d$. Disjointness gives

$$|F| \cdot \omega_d \cdot (\Delta_{op} / 2)^d \le \mathscr{H}^d_{D_{op}}(M \cap \mathcal{A}).$$

By Definition 3.6, $\mathscr{H}^d_{D_{op}}(M \cap \mathcal{A}) = (\omega_d / 2^d) \cdot \mathrm{Vol}_{op}(M)$. Substituting,

$$|F| \cdot \omega_d \cdot (\Delta_{op} / 2)^d \le (\omega_d / 2^d) \cdot \mathrm{Vol}_{op}(M),$$

so

$$|F| \cdot (\Delta_{op})^d \le \mathrm{Vol}_{op}(M),$$

hence

$$|F| \le \mathrm{Vol}_{op}(M) / \Delta_{op}^d.$$

Taking the supremum over admissible families F yields the bound on P(M).

Remark 7.1.1. The dimensional exponent $d = N_{spec} \cdot C_{max}$ is the central structural quantity. It is *not* a multiplicative prefactor — it appears in the *exponent* of Δ_{op}, where it controls how tightly admissible sectors can be packed in $\mathcal{A}$. Packing in d dimensions costs Δ_{op}^d per cell, not Δ_{op}^1.

Remark 7.1.2. Every factor in the bound is an admissibility invariant: N_{spec} from the $\mathbb{Z}_7$ closure architecture, C_{max} from the per-channel rank of Ω_{max}, Δ_{op} from the discrete ledger / entropy-partition gap, and $\mathrm{Vol}_{op}(M)$ from the operational geometry of M.

---

8. The Finite Distinguishability Packing Theorem

Theorem 8.1 (Factorised Cardinality Ceiling). *Let M be a closure manifold satisfying:*

1. *discrete ledger realizability with bounded closure support* (§2.1),
2. *unique closure scale ξ with $k_{max} = 2\pi/\xi$* (§2.2),
3. *unique quantized admissible entropy partition Π_S* (§2.3),
4. *finite admissible spectral channels N_{spec} from the $\mathbb{Z}_7$ closure architecture* (§2.4),
5. *refinement-stable admissibility under the finite-rank projector Ω_{max}* (§2.5).

Then $|\Sigma(M)| < \infty$ (already from (1) alone), and admits the explicit factorised ceiling

$$|\Sigma(M)| \le P(M) \le \mathrm{Vol}_{op}(M) / \Delta_{op}^d, \quad d = N_{spec} \cdot C_{max},$$

where each of Δ_{op}, $\mathrm{Vol}_{op}(M)$, N_{spec}, C_{max} is an admissibility invariant of M.

Proof. By Lemma 4.0, |Σ(M)| < ∞ and Vol_op(M) < ∞. By Lemma 4.1, Δ_op > 0. By Lemma 5.1, σ(S) ⊆ [0, k_max] with admissible content in N_spec < ∞ channels. By Lemma 6.1 and Corollary 6.3, admissible sectors lie in 𝒜 of dimension d ≤ N_spec · C_max. By Lemma 7.1, P(M) ≤ Vol_op(M) / Δ_op^d.

Since |Σ(M)| ≤ P(M) by Definition 3.7, the factorised bound follows.

Role of hypothesis (3). Hypothesis (3) enters twofold. *Uniqueness* of Π_S underlies the definition of D_op (Definition 3.3) and hence the well-definedness of Δ_op as the minimum operational separation. *Quantization* of Π_S, via Remark 4.1.1, ensures Δ_op is a universal admissibility invariant across all of M rather than a region-dependent constant — without quantization, Δ_op could vary across operational sub-manifolds, and the ceiling would fail to be a global structural invariant. Hypotheses (1), (2), (4), (5) carry the structural weight of the cardinality bound itself; hypothesis (3) makes the bound a *universal* invariant.

Remark 8.1.1. The theorem is stronger than the prior finite-spectrum result in four respects:

(i) it gives a *factorised* cardinality ceiling rather than mere finiteness, exposing the structural origin of the count; (ii) the dimensional scaling Δ_op^d is correct, with d = N_spec · C_max in the exponent; (iii) each factor is an admissibility invariant, so the ceiling is itself a structural invariant of M; (iv) the role of each hypothesis (1)–(5) is now sharply delineated: (1) supplies finiteness and Δ_op > 0, (2), (4), (5) supply the dimensional architecture of the ceiling, and (3) supplies *both* the entropy-equivalence quotient defining D_op *and* the universality of Δ_op across M.

---

9. Cardinality Corollaries

Corollary 9.1 (Universal finiteness). Every admissible sector family on M is finite. (Already from §2.1.)

Corollary 9.2 (Termination of stable towers). No infinite tower of mutually distinguishable refinement-stable admissible sectors can exist on M.

Corollary 9.3 (Ceiling as admissibility invariant). The factorised ceiling Vol_op(M) / Δ_op^d is invariant under admissibility-preserving deformations of M and is determined by (Vol_op, Δ_op, N_spec, C_max). The packing number P(M) inherits this invariance.

Corollary 9.4 (Sector-family bound). For any closed admissible sector family ℱ ⊆ Σ(M) (e.g. a generation family, a flavour family, an isospin family),

|ℱ| ≤ Vol_op(M_ℱ) / Δ_op^{d_ℱ}, d_ℱ := dim 𝒜_ℱ ≤ N_spec^{(ℱ)} · C_max^{(ℱ)},

where M_ℱ is the operational sub-manifold supporting ℱ, and (N_spec^{(ℱ)}, C_max^{(ℱ)}) are the spectral-channel parameters of M_ℱ.

Remark 9.4.1. The exponent Δ_op^{d_𝓕} uses the *global* universal Δ_op established in Remark 4.1.1, not a family-local Δ_op^{(𝓕)}. Universality of Δ_op across M — supplied by quantization of Π_S (hypothesis (3) of Theorem 8.1) — is what makes the family-level ceiling well-defined and comparable across sub-manifolds. If hypothesis (3) failed, each family could carry its own Δ_op^{(𝓕)} and inter-family comparisons would lose structural meaning.

Corollary 9.4 is the principal physical specialisation: each closed sector family inherits its own dimensional exponent from its operational sub-manifold, while sharing the universal entropy gap with all other admissible structure on M.

10. Consequences for Physical Sector Structure

The factorisation implies, as structural consequences:

- stable particle sectors are finite in number (already from §2.1);
- the cardinality of each sector family factorises through (Vol_op, Δ_op, d_𝓕);
- stable generation towers terminate at a packing-determined ceiling whose dimensional exponent comes from the spectral-channel architecture;
- admissible closure families have admissibility-invariant cardinality ceilings;
- infinite distinguishable spectra are already excluded by the discrete ledger; the factorisation identifies *what controls the resulting finite count*.

Together these strengthen the earlier finite-spectrum programme from structural finiteness to a quantitative dimensional-packing factorisation.

11. Application to Particle Families

The charged-lepton closure programme established n ∈ {1, 2, 3}. The present theorem does not, on its own, derive n = 3 — that requires the sector-specific closure mechanism of the lepton programme. What the theorem *does* supply is the global structural factorisation that governs why finite generation structure should exist at all and what controls its dimensional scaling.

A hypothetical fourth stable charged-lepton generation would require either

(a) distinguishability finer than Δ_op (forbidden by §2.1 ledger discreteness, which underlies Lemma 4.1), (b) refinement complexity beyond C_max (forbidden by Lemma 6.1), (c) spectral support beyond k_max (forbidden by Lemma 5.1),

or a violation of the admissible packing ceiling P(M_lepton) of Corollary 9.4. All four are admissibility violations.

Thus:

The finite charged-lepton hierarchy is one realisation of a universal factorisation, not a sector-specific accident.

By Corollary 9.4, the same factorisation applies to quark generations, neutrino families, and any other closed admissible sector family — each with its own operational sub-manifold M_ℱ and dimensional exponent d_ℱ, sharing the universal Δ_op across all admissible sectors of M.

12. Relation to Entropy Admissibility

The role of entropy admissibility is now sharply circumscribed.

The discrete ledger (§2.1) alone is sufficient to establish |Σ(M)| < ∞ and Δ_op > 0. The role of §2.3 (entropy admissibility) is to:

- supply the entropy-equivalence quotient defining D_op (uniqueness of Π_S);
- identify Δ_op as the operational gap of the quantized partition (quantization of Π_S);
- ensure Δ_op is a universal admissibility invariant across M rather than a region-dependent constant.

Without entropy partition quantization, Δ_op could vary across regions of M, and the global cardinality ceiling would fail to be a structural invariant. §2.3 therefore supplies *universality* of the ceiling rather than its existence.

13. Relation to Information-Theoretic Packing

The factorisation is structurally reminiscent of:

- coding-theoretic sphere-packing bounds (Hamming-type),
- Shannon distinguishability bounds on channel capacity,
- finite-dimensional Hilbert-space distinguishability limits,
- finite phase-space information density à la Bekenstein.

One important structural difference deserves emphasis. In information theory, packing bounds are typically *statistical*: they reflect the impossibility of reliably distinguishing more than a certain number of signals subject to noise.

The VERSF packing bound is *not* statistical. It follows from:

- closure-scale geometry,

- admissibility,
- refinement persistence,

each of which is an ontological constraint on what can exist, not a statistical constraint on what can be measured. The bound is therefore best understood as **ontological factorisation** rather than informational packing.

The analogy to the Hamming bound $q^n / V(n, \lfloor (d-1)/2 \rfloor)$ is loose but evocative: the dimensional exponent here is supplied by the spectral-channel architecture rather than by codeword length, and the role of "minimum distance" is played by the universal entropy gap Δ_{op} rather than by a code design parameter. The structural family resemblance is geometric, not isomorphic.

---

14. Predictions and Falsifiability

The factorisation predicts:

- finite stable particle sectors universally (already from §2.1);
- the cardinality of every sector family factorises through (Vol_{op}, Δ_{op}, $d_{\mathscr{F}}$);
- the dimensional exponent $d_{\mathscr{F}}$ traces to the spectral-channel architecture of the relevant operational sub-manifold;
- Δ_{op} is universal across all admissible sectors of M.

Potential falsifiers include:

- discovery of an arbitrarily extendable stable particle hierarchy;
- experimental demonstration of refinement-stable admissible sectors at infinite tower depth;
- observation of operational distinguishability below Δ_{op};
- detection of stable admissible structure beyond k_{max};
- a closed admissible sector family violating Corollary 9.4;
- variation of Δ_{op} across operational sub-manifolds (would invalidate hypothesis (3) and the universality of the ceiling).

Any such observation would invalidate either the discrete ledger, the entropy partition, the closure scale, or the finite rank of Ω_{max} — each of which is independently testable through its own programme.

---

15. Open Problems

The factorisation is now established with correct dimensional scaling. The following remain open:

1. explicit closed-form evaluation of Vol_op(M) for the physical closure manifold;
2. operator-level enumeration of admissible sectors within Im(Ω_max);
3. quark-sector packing structure under the appropriate sub-manifold M_quark and exponent d_quark;
4. neutrino-sector admissibility and its operational sub-manifold;
5. gauge-sector closure and admissible vector-bundle packing;
6. microscopic derivation of Δ_op from the discrete commitment ledger;
7. derivation of the exact Standard Model sector cardinalities as specific instances of Corollary 9.4.

Each open problem is now a *quantitative* problem about evaluating the explicit factorisation on a specific operational sub-manifold, rather than a qualitative problem about whether finiteness holds.

16. Conclusion

The previous finite-spectrum theorem established structural termination of stable admissible sectors. Finiteness of |Σ(M)| is in fact immediate from the discrete ledger (§2.1); bounded closure support admits only finitely many ledger-realizable states. The "infinite distinguishable sectors within bounded support" loophole is closed by §2.1 alone.

The present paper supplies what §2.1 does not: the **structural origin** of this finiteness, expressed as an explicit factorised cardinality ceiling

|Σ(M)| ≤ Vol_op(M) / Δ_op^d, d = N_spec · C_max,

whose every factor is an admissibility invariant — Δ_op from the operational entropy gap (universal under hypothesis (3)), Vol_op from the operational d-volume of M, N_spec from the $\mathbb{Z}_7$ closure architecture, and C_max from the per-channel rank of Ω_max. The dimensional exponent d enters via the geometry of the admissible subspace 𝒜 = Im(Ω_max), with the spectral-channel decomposition Im(Ω_max) = ⊕_α V_α exhibited explicitly.

The result upgrades the VERSF finite-sector programme from **structural finiteness** to **structural factorisation**: not the claim that stable physical sectors are finite (which is immediate from §2.1), but the identification of which admissibility invariants control the cardinality, at what dimensional scaling, and through what geometric mechanism.

Applied to particle physics, the factorisation supplies the missing mathematical bridge between closure-scale finiteness and finite particle-family structure — each closed sector family inheriting its own dimensional exponent d_ℱ and operational sub-manifold M_ℱ, with the universal Δ_op shared across all admissible sectors of M.

Stable physical sectors are finite by the discrete ledger; their cardinality is *structured* by the factorisation Vol_op / Δ_op^d. The packing theorem identifies not the fact of finiteness, but the architecture that controls it.