HINGE 7 PREPRINT PILOT · V1.0 · APRIL 2026

Address Resolution, Not Calculation

The Genealogical DAG, the Cayley Metric, and Why Arithmetic in Category τ Has No Equations

Thorsten Fuchs • Anna-Sophie Fuchs
Correspondence: thorsten@panta-rhei.site

April 2026
DOI: 10.5281/zenodo.19818628

ABSTRACT

Building on the earned categorical machine of Hinge 5 [16] and the τ-topos of Hinge 6 [18], we establish the *canonical-address normal-form confluence theorem* — the Church–Rosser-style lemma on which those hinges' pre-Yoneda collapse and exponential-object constructions depend — and thereby retire the "modulo Hinge 7" scope caveats carried throughout the Panta Rhei bundle. We prove: (i) a *canonical normalisation theorem* showing that every admissible tail-transformer code $c \in \mathsf{Code}$ reduces in finitely many witness-preserving steps to a unique normal form $\mathsf{NF}(c)$ under the τ-kernel rewriting system $\to_{\mathsf{NF}}$; (ii) the *NF confluence theorem* (Church–Rosser for the τ-kernel): if $c \twoheadrightarrow_{\mathsf{NF}} c_1$ and $c \twoheadrightarrow_{\mathsf{NF}} c_2$ then $c_1 \twoheadrightarrow_{\mathsf{NF}} c'$ and $c_2 \twoheadrightarrow_{\mathsf{NF}} c'$ for some common reduct, giving $\to_{\mathsf{NF}}$ the diamond property and hence the existence of unique NF per $\sim$-class; (iii) the *genealogical DAG* $\mathsf{DAG}_\tau = (\mathsf{Code}, \to_{\mathsf{NF}})$ is a countable, strongly-normalising, finite-width directed acyclic graph whose nodes are admissible codes and whose edges are elementary NF reduction steps; (iv) the *Cayley word metric* d_{Cay} on DAG_τ (the normalised edge-count distance) coincides with the minimal-pass distance and induces the τ-topology on the admissible address space $\mathrm{Addr}_\tau = \mathsf{Code}/\sim$; (v) the *ontic ultrametric* d_∞ on the completion $\mathsf{Ultra}_\tau = \widehat{\mathrm{Addr}}_\tau$ is a genuine ultrametric (satisfies the strong triangle inequality $d_\infty(a, c) \leq \max(d_\infty(a, b), d_\infty(b, c))$), non-archimedean, and totally-disconnected — matching the profinite boundary structure of Hinge 4; (vi) the *address-resolution theorem*: every question of arithmetic equality $a = b$ in Category τ is a finite-witness decidable address-resolution computation $\mathsf{NF}(a) \equiv \mathsf{NF}(b)$, not an equational calculation in the classical sense. Arithmetic in τ is therefore *address-resolution, not equation-solving*; and (vii) a *hinge-integration theorem* tabulating the seven-hinge bundle in closed form, now that the foundational arc is complete. Lean 4 formalisation is planned in `TauLib.BookI.Addressability`.

Keywords canonical addressability, normal-form confluence, Church–Rosser theorem, genealogical DAG, Cayley word metric, ontic ultrametric, strong normalisation, address resolution, arithmetic-without-equations, Panta Rhei hinge paper

MSC 2020 Mathematics Subject Classification: 68Q42, 03B40, 18A05, 03F65, 03D80, 54E50

CONTENTS

1. INTRODUCTION AND STATEMENT OF RESULTS

1.1 Position in the Panta Rhei hinge-paper bundle

This paper is **Hinge 7**, the capstone of the technical subsequence (H1–H7) of the eight-paper Panta Rhei foundational bundle accompanying the 2nd Edition of the series [7, 8, 9, 10, 11, 12, 13]. The bundle consists of seven technical hinges plus a foundational- anchor paper (H8); in the recommended reading order they are:

- **Hinge 1.** *Hyperfactorization* [5] — unique tower-atom decomposition; the coarse-grained coordinate framework on which canonical addresses are built.
- **Hinge 2.** *Prime Polarity* [14] — Legendre $(2/p)$ mod-8

prime split; supplies the B/C channel characters.

- **Hinge 3.** *Master Constant* ι_τ [6] — $\iota_\tau = 2/(\pi + e) \approx 0.341304$, the σ-fixed crossing-germ scalar whose canonical address anchors the addressing calibration.
- **Hinge 4.** *The Split-Complex Boundary Algebra* [15] — $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2 - 1)$, the algebraic home of addresses at the boundary.
- **Hinge 5.** *τ-Holomorphy on the Boundary Algebra* [16] — the earned categorical machine; the pre-Yoneda collapse on which addresses gain categorical status.
- **Hinge 6.** *The τ-Topos and Its Four-Valued Internal Logic* [18] — $\mathbf{Cat}_\tau$ with $\Omega_\tau = B_\sigma(\mathbb{D})$; address-resolution as internal morphism-computation.
- **Hinge 7.** *Address Resolution, Not Calculation (this paper)* — establishes canonical-address NF confluence, the genealogical DAG, and the ontic ultrametric; proves arithmetic is address-resolution; retires the "modulo Hinge 7" caveats of Hinges 5 and 6.
- **Hinge 8.** *The τ-Kernel as Foundational Architecture* [17] — foundational-anchor paper (also readable as an entry point): ontic identity invariance, diagonal–linear correspondence, $*$-autonomous placement; names what the seven technical hinges collectively earn.

The present paper closes the foundational arc: the forward- dependencies of Hinges 5 and 6 on "canonical-address NF confluence" are discharged here, and the ontic ultrametric established here becomes the canonical metric structure for the boundary address space used throughout Books II–VII.

1.2 The address-resolution paradigm

Classical arithmetic is built on *equations*: "$a = b$" is a statement that two expressions denote the same mathematical object, and arithmetic proceeds by equational calculation — chains of equality-preserving rewrites establishing new equations from old. This is the view inherited from Peano, Dedekind, and Hilbert; it underwrites the classical equational theories of groups, rings, and fields.

The τ-framework inverts this paradigm. In Category τ, the primary objects are *canonical addresses* — normal forms of tail-transformer codes in a confluent rewriting system — and "arithmetic" is the process of resolving a token to its canonical address. There are no equations in the classical sense, because there are no free-standing mathematical objects that admit equational identification. There are only codes, their canonical addresses, and the NF-resolution procedure.

Concretely, a τ-native question of the form "$a = b$" is interpreted as:

> *Compute the canonical address* $\mathsf{NF}(a)$ *and the canonical address* $\mathsf{NF}(b)$*; the question is answered "yes" iff the two addresses are the same NF code. The computation is a finite-witness decidable procedure, not an equational search.*

This is the *address-resolution paradigm*: arithmetic is algorithmic address-resolution via NF rewriting, not equation-solving. The classical equational view emerges only as the *extensional shadow* of the address-resolution procedure; two codes are equational-equal iff they have the same canonical address.

1.3 Main results

We establish seven main theorems. Each is anchored at **[τ-Effective]** with full closure of prior "modulo Hinge 7" caveats; a **[Established]**-tier strengthening programme is sketched in §9 contingent on Book II's classical-equational comparison.

Theorem 1.1 (Canonical Normalisation **[τ-Effective]**). *Every admissible code $c \in \mathsf{Code}$ reduces in finitely many witness-preserving τ-kernel rewriting steps $\to_{\mathsf{NF}}$ to a unique normal form $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}} \subset \mathsf{Code}$:*

$$c \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c), \qquad \mathsf{NF}(\mathsf{NF}(c)) = \mathsf{NF}(c), \qquad |\textit{reduction path}| \leq k_0(c),$$

where $k_0(c) \in \mathbb{N}$ is a finite witness depth bounded by the code's maximum token count. The map $\mathsf{NF} : \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ is idempotent, surjective, and $\sim$-preserving.

Theorem 1.2 (NF Confluence (Church–Rosser for τ-kernel) **[τ-Effective]**). *The τ-kernel rewriting system* $(\mathsf{Code}, \to_{\mathsf{NF}})$ *is* confluent*: for every $c \in \mathsf{Code}$ and every pair of rewriting paths $c \twoheadrightarrow_{\mathsf{NF}} c_1$, $c \twoheadrightarrow_{\mathsf{NF}} c_2$, there exists a common reduct c' with $c_1 \twoheadrightarrow_{\mathsf{NF}} c'$ and $c_2 \twoheadrightarrow_{\mathsf{NF}} c'$. Equivalently, $\to_{\mathsf{NF}}$ has the diamond property, and every $\sim$-equivalence class contains a unique NF representative. This theorem discharges the "modulo Hinge 7 NF confluence" scope caveats of Hinges 5 and 6.*

Theorem 1.3 (Genealogical DAG **[τ-Effective]**). *The pair* $\mathsf{DAG}_\tau := (\mathsf{Code}, \to_{\mathsf{NF}})$ *forms a countable, strongly-normalising, finite-width directed acyclic graph:*

- Countable*: $|\mathsf{Code}| \leq \aleph_0$ by NF coding.*
- Strongly normalising*: every rewriting path terminates at an NF code in finitely many steps.*
- Finite-width*: at every depth k, only finitely many codes have witness depth $\leq k$ (polynomial growth in the primorial ladder).*
- Acyclic*: no rewriting cycle $c \twoheadrightarrow_{\mathsf{NF}} c' \twoheadrightarrow_{\mathsf{NF}} c$ with $c \neq c'$ exists.*

The DAG has a unique root at the empty NF code and exactly one NF sink per $\sim$-equivalence class.

Theorem 1.4 (Cayley Word Metric **[τ-Effective]**). *The* Cayley word metric *on* DAG_τ,

$$d_{\mathrm{Cay}}(c_1, c_2) \;:=\; \min\{n \in \mathbb{N} : \exists c' \in \mathsf{Code}^{\mathsf{NF}}, \; c_1 \xrightarrow{n_1} c' \xleftarrow{n_2} c_2,\; n_1 + n_2 = n\},$$

is a well-defined metric on the address space $\mathrm{Addr}_\tau := \mathsf{Code}^{\mathsf{NF}} / \sim$*. It coincides with the minimal-pass distance of Hinge 1 [5], and induces the τ-topology on* Addr_τ.

Theorem 1.5 (Ontic Ultrametric **[τ-Effective]**). *The d_{Cay}-completion* $(\mathsf{Ultra}_\tau, d_\infty) := (\widehat{\mathrm{Addr}_\tau, d_{\mathrm{Cay}}})$ *is a complete ultrametric space:*

$$d_\infty(a, c) \;\le\; \max(d_\infty(a, b), d_\infty(b, c)), \quad \textit{for all } a, b, c \in \mathsf{Ultra}_\tau,$$

non-archimedean and totally disconnected, matching the profinite boundary structure of Hinge 4. The ontic ultrametric is the τ-native replacement for the Euclidean / Archimedean distance function of classical analysis; its scale is set by the primorial ladder $(M_k)_{k\in\mathbb{N}}$.

Theorem 1.6 (Address-Resolution **[τ-Effective]**). *Every question of "arithmetic equality" in Category τ reduces to a canonical-address NF comparison: for admissible codes $a, b \in$* Code,

$$a \sim b \iff \mathsf{NF}(a) = \mathsf{NF}(b),$$

and the right-hand side is a finite-witness decidable computation of polynomial depth in $\min(\mathrm{depth}(a), \mathrm{depth}(b))$*. Consequently,* Category τ has no equations in the classical sense*: there are no free-standing mathematical objects standing in equational relation; only codes, their canonical addresses, and the NF-resolution procedure. The classical equational view emerges as the extensional shadow of address-resolution.*

Theorem 1.7 (Hinge 7 integration tabulation **[τ-Effective]**). *Hinge 7 is the foundational capstone of the Panta Rhei bundle. It closes the seven-hinge arc by:*

- ***Discharging** the "modulo Hinge 7 NF confluence" scope caveats of Hinge 5 [16] (Theorem 1.8 pre-Yoneda collapse) and Hinge 6 [18] (Theorem 1.1 topos structure, Theorem 1.3 circularity resolution).*
- ***Supplying** the canonical-address framework used as coordinate primitive in Hinges 1–4 [5, 14, 6, 15].*
- ***Establishing** the ontic ultrametric as the τ-native metric structure that replaces Euclidean distance throughout Books II–VII.*
- ***Completing** the foundational arc: with Hinge 7 in place, the Panta Rhei framework has a closed, self-contained, seven-hinge bundle of standalone peer-reviewable papers.*

The present paper completes the bundle's seventh and final foundational hinge*: the canonical-address NF confluence capstone on which the entire pre-categorical and categorical machinery of Category τ rests.*

1.4 Lean roadmap (preview)

Full formalisation is targeted at `TauLib.BookI.Addressability` [19] in the Lean 4 proof assistant [25], comprising the following planned modules:

- `NFRewriting.lean` — the τ-kernel rewriting system and elementary reduction rules.
- `Normalisation.lean` — strong normalisation and canonical NF existence.
- `Confluence.lean` — the Church–Rosser theorem for the τ-kernel.
- `GenealogicalDAG.lean` — the DAG structure and its combinatorial properties.
- `CayleyMetric.lean` — the word metric on the DAG.
- `OnticUltrametric.lean` — the ultrametric completion and its non-archimedean structure.
- `AddressResolution.lean` — decidability of arithmetic equality via NF comparison.

2. PRELIMINARIES: IMPORTS FROM HINGES 1–6

The §1 statements rest on six prior hinges of the Panta Rhei bundle. This section states each prerequisite as an imported remark and fixes the notation specific to the address- resolution machinery developed in §§3–8. Imports are cited, not re-proved; downstream sections cite the labels introduced here rather than tracing back to the source papers.

2.1 Hyperfactorization and the tower-atom decomposition (Hinge 1)

We recall, without re-proving, the hyperfactorization machinery of Hinge 1 [5]. Hinge 1 supplies the coarse-grained coordinate structure on admissible objects of Category τ; the finite pass-budget witness k_0 is the natural-number bound that underwrites every termination claim of the present paper.

Remark 2.1 (Imported: tower-atom decomposition **[Established]**). For every admissible object $X \in \mathrm{Obj}(\tau)$ there exists a unique *tower-atom decomposition*

$$X \;=\; (A \uparrow\uparrow C)^B \cdot D, \qquad \Phi(X) \;=\; (A, B, C, D), \quad (1)$$

where A is the largest prime atom, C is the tetration depth, B is the iterated exponent, D is the classical smooth residue, and $\uparrow\uparrow$ denotes Knuth tetration [5, Def. 2.3, Thm. 1.1]. The *ABCD chart* $\Phi\colon \mathsf{Obj}(\tau) \to \mathbb{N}^4$ is injective on admissible objects [5, Thm. 3.2], and the decomposition (1) is the natural coordinate primitive on which canonical addresses are built. The classical integer $X \in \mathbb{N}_{\geq 2}$ emerges as the height-$C = 1$ specialisation; arithmetic in τ is the universal tower-atom calculus, and ordinary arithmetic is its shadow on the classical stratum [5, §§4–5].

Remark 2.2 (Imported: pass budget $k_0(X)$ **[Established]**)**.** Associated to each $X \in \mathsf{Obj}(\tau)$ is the *pass budget*

$$k_0(X) \;\in\; \mathbb{N}, \tag{2}$$

the minimal depth at which the greedy peel algorithm [5, §6.2] stabilises on the unique admissible ABCD tuple $\Phi(X)$. The pass budget is finite and bounded polynomially in the token-count of X, and it supplies the *finite-witness bound* for strong-normalisation claims in the present paper: if $c \in \mathsf{Code}$ encodes a carrier of pass budget $\leq k_0$, then the NF rewriting reduction $c \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c)$ terminates within k_0 elementary steps (Theorem 1.1). The pass budget is invariant under $\sim$ and under the σ-involution of Remark 2.4 below. For admissible codes we shall also write $k_0(c)$ for the pass budget of the carrier addressed by c; see Definition 2.25.

Remark 2.3 (Imported: minimal-pass word metric **[Established]**)**.** Hinge 1 [5, §7.3] defines on the set of admissible ABCD coordinates the *minimal-pass distance*

$$d_{\text{pass}}(X,Y) \;:=\; \min\{\, k \in \mathbb{N} : \Phi(X) \text{ and } \Phi(Y) \text{ agree up to level } k \,\}, \tag{3}$$

which measures the coarsest common refinement depth of the two tuples. The Cayley word metric d_{Cay} of Theorem 1.4 is the NF-lift of (3) from ABCD coordinates to admissible codes; the identification $d_{\text{Cay}} = d_{\text{pass}}$ at the NF-representative level is one of the main technical outputs of §6.

2.2 Prime polarity and the Legendre $(2/p)$ split (Hinge 2)

We recall, without re-proving, the Prime Polarity Theorem of Hinge 2 [14]. Hinge 2 supplies the canonical B/C channel characters used by the idempotent basis of Hinge 4 and by the σ-involution of Hinge 5.

Remark 2.4 (Imported: Prime Polarity Theorem **[Established]**)**.** Every odd prime $p \in \mathbb{P}$ falls into exactly one of two Legendre classes under the quadratic character $(2/p)$:

$$\chi_2(p) \;:=\; \left(\tfrac{2}{p}\right) \;=\; \begin{cases} +1 & \text{if } p \equiv \pm 1 \pmod 8, \\ -1 & \text{if } p \equiv \pm 3 \pmod 8. \end{cases} \tag{4}$$

The classes $\chi_2(p) = +1$ and $\chi_2(p) = -1$ are the *B-channel* and *C-channel* of τ respectively; in tower-atom terms of Remark 2.1 they distinguish the two possible parities of the exponent B relative to the tetration-height C, hence the name "B/C channel split" [14, Thm. 1.4, §3]. The split (4) is both a ZFC/PA theorem and a theorem internal to τ [14, Thm. 1.1]; under the τ-to-$\mathbb{N}$ shadow functor it becomes the classical Legendre supplementary law.

Remark 2.5 (Imported: polarity $\to$ σ-involution **[Established]**)**.** The Legendre split (4) is the source of the σ-involution (11) of $\mathbb{D}$: the map

$$\sigma_\chi\colon \chi_2(p) \;\longmapsto\; -\chi_2(p) \tag{5}$$

on channel characters lifts to the idempotent swap $\sigma(e_+) = e_-$, $\sigma(e_-) = e_+$ of Remark 2.9 below [14, §7]. We shall use this compatibility tacitly: the σ-equivariance statements of §§4–7 are consistent with the polarity split at the level of prime atoms, and the NF normaliser commutes with σ.

2.3 The master constant ι_τ (Hinge 3)

We recall, without re-proving, the ι_τ identity of Hinge 3 [6]. In the address-resolution paradigm of this paper, ι_τ is the *canonical address of the lemniscate crossing germ*; it anchors the calibration of the address space Addr_τ.

Remark 2.6 (Imported: ι_τ as σ-fixed crossing-germ scalar **[Established]**)**.** The *τ-master constant*

$$\iota_\tau \;:=\; \frac{2}{\pi + e} \;=\; 0.341304238875\ldots \;\in\; \mathcal{R}'_\partial \subset \mathbb{D}, \tag{6}$$

lies on the real axis $\mathbb{D}^\sigma$ of Remark 2.11 and is the unique σ-fixed scalar satisfying the boundary-normalisation constraints of the lemniscate crossing germ [6, Thm. 1.2]. Under the address-resolution paradigm of §8, ι_τ is the canonical-address primitive for the crossing germ $\gamma_{\text{cross}} \in \partial\tau^3$; its NF code is a fixed reference point in $\mathsf{Code}^{\mathsf{NF}}$, and its Cayley-metric position in DAG_τ serves as the *calibration anchor* for every ontic-distance computation. In particular the σ-fixed real axis $\mathbb{D}^\sigma = \mathcal{R}'_\partial \cdot 1$ (Remark 2.11) is where ι_τ, the Both-truth-value, and the self-reference scalars live; Hinge 7's address-resolution theorem (Theorem 1.6) guarantees that the NF of the token ι_τ is a finite fixed point of the τ-kernel normaliser.

2.4 The split-complex boundary algebra $\mathbb{D}$ (Hinge 4)

We recall, without re-proving, the boundary algebra $\mathbb{D}$ of Hinge 4 [15]. Throughout, $\mathcal{R}'_\partial := \mathcal{R}_\partial[1/2]$ denotes the dyadic localisation of the countable profinite boundary ring

$\mathcal{R}_\partial \cong \varprojlim_k \mathbb{Z}/M_k\mathbb{Z}$ over the primorial ladder $(M_k)_{k\in\mathbb{N}}$ [15, §§3–4].

Remark 2.7 (Imported: boundary ring $\mathcal{R}_\partial$ as countable profinite limit **[Established]**). The *boundary ring*

$$\mathcal{R}_\partial \;:=\; \varprojlim_{k\in\mathbb{N}} \mathbb{Z}/M_k\mathbb{Z}, \qquad M_k \;:=\; \prod\nolimits_{p\le p_k} p \tag{7}$$

is a countable profinite commutative ring indexed by the primorial ladder (M_k) [15, Thm. 4.1]. Its dyadic localisation $\mathcal{R}'_\partial = \mathcal{R}_\partial[1/2]$ is the natural scalar ring for the split-complex algebra $\mathbb{D}$; the primorial ladder (M_k) sets the scale of the ontic ultrametric d_∞ in Theorem 1.5.

Remark 2.8 (Imported: split-complex boundary algebra **[Established]**). The *split-complex boundary algebra* is the free commutative $\mathcal{R}'_\partial$-algebra on one generator j modulo $j^2 = +1$:

$$\mathbb{D} \;:=\; \mathcal{R}'_\partial[j] \,/\, (j^2 - 1), \tag{8}$$

free of rank 2 as an $\mathcal{R}'_\partial$-module with basis $\{1, j\}$, so every $z \in \mathbb{D}$ is uniquely $z = a + j\,b$ with $a, b \in \mathcal{R}'_\partial$ [15, §5, Thm. 1.4]. For the classical split-complex ("hyperbolic-number") plane on which $\mathbb{D}$ is modelled (without the dyadic-profinite boundary structure), see [32, 29, 4]; the elliptic alternative $\mathcal{R}'_\partial[i]/(i^2+1)$ — the classical several- complex-variables setting [20, 21, 28, 24] — is excluded here by [15, Thm. 1.7]; the algebra is uniquely determined as the commutative $\mathcal{R}'_\partial$-algebra of binary rank with two nontrivial orthogonal idempotents and canonical polarity-swap involution.

Remark 2.9 (Imported: idempotent basis and σ-involution **[Established]**). The canonical orthogonal idempotents

$$e_+ \;:=\; \tfrac{1}{2}(1+j), \qquad e_- \;:=\; \tfrac{1}{2}(1-j) \tag{9}$$

satisfy $e_+ + e_- = 1$, $e_+ \cdot e_- = 0$, $e_+^2 = e_+$, $e_-^2 = e_-$, $j = e_+ - e_-$, and the four-element σ-equivariant Boolean sublattice

$$B_\sigma(\mathbb{D}) \;:=\; \{0,\, e_+,\, e_-,\, 1\} \;\subset\; \mathbb{D} \tag{10}$$

is closed under the canonical involution $\sigma\colon \mathbb{D} \to \mathbb{D}$, $\sigma(j) = -j$, which satisfies $\sigma(e_+) = e_-$, $\sigma(e_-) = e_+$, $\sigma(0) = 0$, $\sigma(1) = 1$, and $\sigma^2 = \mathrm{id}_{\mathbb{D}}$ [15, Thm. 1.8]. We use the notation $\sigma(j) = -j$ as shorthand for the full involution of (11) below; the real axis $\mathbb{D}^\sigma$ hosts the master constant ι_τ of Remark 2.6.

Remark 2.10 (Imported: σ-involution in full **[Established]**). Spelled out on all four Boolean atoms:

$$\begin{gathered}\sigma(j) = -j, \quad \sigma(e_+) = e_-, \quad \sigma(e_-) = e_+, \quad \sigma(1) = 1, \\ \sigma(0) = 0, \quad \sigma^2 = \mathrm{id}_{\mathbb{D}},\end{gathered} \tag{11}$$

with σ an involutive $\mathcal{R}'_\partial$-algebra automorphism of $\mathbb{D}$.

Remark 2.11 (Imported: σ-fixed real axis **[Established]**). With $z = z_+e_+ + z_-e_-$ in the lobe decomposition of $\mathbb{D}$, the action $\sigma(z) = z_-e_+ + z_+e_-$ gives

$$\begin{aligned}\mathbb{D}^\sigma \;&:=\; \{z \in \mathbb{D} : \sigma(z) = z\} \\ &=\; \{z_+e_+ + z_-e_- : z_+ = z_-\} \\ &=\; \mathcal{R}'_\partial \cdot 1.\end{aligned} \tag{12}$$

This real axis hosts both ι_τ of Remark 2.6 and the paraconsistent truth value Both $= 1$ of Hinge 6.

2.5 The earned categorical machine (Hinge 5)

We recall, without re-proving, the earned categorical machine of Hinge 5 [16]. All objects below are countable and finite-witness decidable [16, Rem. 2.2, 3.6]. Crucially, Hinge 5 asserts Church–Rosser NF confluence as an *imported lemma*; the present paper now *establishes* that lemma and thereby discharges the "modulo Hinge 7" caveat.

Remark 2.12 (Imported: ω-tails, prefix agreement, and tail-equivalence **[Established]**). The type Ω_{tail} of *ω-tails* consists of infinite coherent prefix chains over the τ-native token alphabet [16, Def. 2.1]. For each depth $k \in \mathrm{Idx}$ there is a decidable *prefix-agreement* relation $\equiv_k$ on Ω_{tail}; the *full tail-equivalence*

$$t \sim t' \quad\Longleftrightarrow\quad \forall k \in \mathrm{Idx},\; t \equiv_k t' \tag{13}$$

is reflexive, symmetric, and transitive [16, Lem. 2.3]. Each carrier $X \in \mathrm{Obj}(\tau)$ comes with a $\sim$-invariant decidable predicate $\mathsf{Tail}_X\colon \Omega_{\mathrm{tail}} \to \mathsf{Prop}$.

Remark 2.13 (Imported: admissibility predicates **[Established]**). A *tail-transformer code* $c \in \mathsf{Code}$ has intensional semantics $[\![c]\!]\colon \Omega_{\mathrm{tail}} \to \Omega_{\mathrm{tail}}$ given by primitive-recursive rewriting [16, Def. 3.1, Rem. 3.7]. The three admissibility predicates on c relative to carriers X, Y are:

$$\begin{aligned} \mathsf{Typed}(X, Y, c) \;&:\Longleftrightarrow\;\; \forall t, \\ &\qquad \mathsf{Tail}_X(t) \Rightarrow \mathsf{Tail}_Y([\![c]\!](t)); \\ \mathsf{Stable}(X, Y, c) \;&:\Longleftrightarrow\;\; \mathsf{Tail}_X(t) \wedge \mathsf{Tail}_X(t') \\ &\qquad \wedge\; t \sim t' \Rightarrow [\![c]\!](t) \sim [\![c]\!](t'); \\ \text{tail-indep.} \;&:\Longleftrightarrow\;\; \exists k_0,\; \mathsf{Tail}_X(t) \wedge \mathsf{Tail}_X(t') \\ &\qquad \wedge\; t \equiv_{k_0} t' \Rightarrow [\![c]\!](t) \sim [\![c]\!](t'). \end{aligned} \tag{14}$$

The type of *τ-holomorphic maps* from X to Y is

$$\begin{aligned} \mathsf{Hol}_\tau(X, Y) \;:=\; \{c \in \mathsf{Code} \mid \mathsf{Typed}(X, Y, c) \\ \wedge\; \mathsf{Stable}(X, Y, c) \wedge \text{tail-indep.}\}, \end{aligned} \tag{15}$$

modulo $\sim$-equivalence of codes. No Cartesian product $X \times Y$, graph, or set-theoretic function space is used at this layer [16, Rem. 3.9]; countability of $\mathsf{Hol}_\tau(X, Y)$ is immediate from finite-witness decidability. The witness k_0 in (14) is the carrier's pass budget of Remark 2.2.

Remark 2.14 (Imported: earned categorical machine **[τ-Effective]**, modulo Hinge 7 for strong confluence)**.** From the admissibility predicates (14) alone, [16, Thm. 1.6, §7] derives as theorems:

(a) *Earned composition.* If $c \in \mathsf{Hol}_\tau(X, Y)$ and $c' \in \mathsf{Hol}_\tau(Y, Z)$, then the sequential action $[\![c']\!] \circ [\![c]\!]$ is the semantics of an admissible code $c' \circ c \in \mathsf{Hol}_\tau(X, Z)$.

(b) *Earned identity.* Each carrier X has a canonical tail-fixing NF code $\mathrm{id}_X \in \mathsf{Hol}_\tau(X, X)$ that is the two-sided unit of composition.

(c) *Earned associativity.* Composition is associative up to NF equivalence of codes; weak confluence holds at this layer, and strong (Church–Rosser) confluence is the subject of Theorem 1.2 of the present paper.

(d) *Functoriality.* $X \mapsto \mathsf{Hol}_\tau(X, -)$ is a functor on the probe category P_τ of primorial-depth refinements.

No category axioms are imposed; all four clauses are theorems.

Remark 2.15 (Imported: pre-Yoneda collapse, *modulo Hinge 7*)**.** Hinge 5 Theorem 1.8 [16, Thm. 9.7] states the *pre-Yoneda collapse*: the Yoneda functor

$$\begin{aligned} y_\tau\colon \tau \;&\longrightarrow\; \mathbf{PSh}_\tau := [\,P_\tau^{\mathrm{op}}, \mathbf{Set}\,], \\ X \;&\longmapsto\; \mathsf{Hol}_\tau(-, X), \end{aligned} \tag{16}$$

collapses at the presheaf level: every $\mathsf{Hol}_\tau(-, X)$ is representable by a canonical boundary-addressed code $c_X \in \partial\tau^3$, where $\partial\tau^3 := \partial(\tau^1 \times_f T^2)$ is the boundary of the Hinge 1 fibered product [5]. The strict form (strict uniqueness of c_X) is the forward dependency of Hinge 5 on Hinge 7; weak representability (up to NF equivalence) holds unconditionally. Theorem 1.2 of the present paper discharges the strict form: with NF confluence established, c_X is a strictly unique NF code in $\partial\tau^3$. Thus Hinge 5 Theorem 1.8 is *promoted* from weak representability to strict representability by the results of §4.

Remark 2.16 (Imported: HolEnd_τ and its σ-equivariant refinement $\mathsf{HolEnd}_\tau^\sigma$ **[τ-Effective]**)**.** The *holomorphic endomorphism category*

$$\begin{aligned} \mathsf{HolEnd}_\tau \;:=\; \{(X, f) : X \in \mathrm{Obj}(\tau), \\ f \in \mathsf{Hol}_\tau(X, X)\}, \end{aligned} \tag{17}$$

with σ-equivariant wide subcategory

$$\begin{aligned} \mathsf{HolEnd}^\sigma \;:=\; \{(X, f) \in \mathsf{HolEnd}_\tau : \bar{f} = f, \\ \bar{f} := \sigma_X \circ f \circ \sigma_X\} \end{aligned} \tag{18}$$

[16, Def. 9.1, Prop. 9.14], is the natural host for σ-fixed structure, in particular for ι_τ as the universal σ-fixed scalar [16, Thm. 9.17]. Under the strict pre-Yoneda collapse of Remark 2.15, both categories become concretely representable by boundary-addressed NF codes.

Remark 2.17 (Hinge 5 asserts Church–Rosser as imported)**.** Hinge 5 §7.2 proves *weak* confluence of the NF rewriting system on admissible codes, and §7.4 then asserts the Church–Rosser (strong) form as an imported lemma attributed to Hinge 1/Hinge 7 [16, Lem. 7.9]. This is precisely the lemma that the present paper now *establishes*: Theorem 1.2 is the strong Church–Rosser theorem for the τ-kernel, and the weak-to-strong promotion of every Hinge 5 theorem downstream of [16, Lem. 7.9] is discharged by Theorem 1.2 below.

2.6 The τ-topos and four-valued internal logic (Hinge 6)

We recall, without re-proving, the τ-topos machinery of Hinge 6 [18]. Hinge 6's exponential-object construction and subobject-classifier identification carry the second "modulo Hinge 7" caveat that the present paper discharges.

Remark 2.18 (Imported: the τ-topos $\mathbf{Cat}_\tau$ **[τ-Effective]**, modulo Hinge 7)**.** The *τ-topos*

$$\mathbf{Cat}_\tau \;:=\; (\mathrm{Obj}(\tau),\; \mathsf{Hol}_\tau(-, -),\; \circ,\; \mathrm{id}) \tag{19}$$

is the countable boundary-addressed category whose objects are admissible carriers and whose morphisms are τ-holomorphic maps (15), endowed with the finite-limit structure, exponential objects, and subobject classifier constructed in [18, §§3–6]. Hinge 6 Theorem 1.1 asserts that $\mathbf{Cat}_\tau$ is an elementary topos; the proof as written depends on strict NF uniqueness for the exponential-object construction, and is thereby **[τ-Effective]** *modulo Hinge 7*. Theorem 1.2 of the present paper supplies that uniqueness.

Remark 2.19 (Imported: subobject classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ **[τ-Effective]**)**.** The *subobject classifier* of $\mathbf{Cat}_\tau$ is

$$\Omega_\tau \;\cong\; B_\sigma(\mathbb{D}) \;=\; \{0, e_+, e_-, 1\}, \tag{20}$$

the σ-equivariant Boolean sublattice of Remark 2.9, and the *four-valued internal logic*

$$\mathsf{Truth}_4 \;:=\; B_\sigma(\mathbb{D}) \;=\; \{\text{Neither, True, False, Both}\} \tag{21}$$

is the Belnap–Dunn four-valued logic [2, 27] carried by Ω_τ, with the truth-sector labelling

$$\begin{aligned} \text{Neither} &:= 0, & \text{True} &:= e_+, \\ \text{False} &:= e_-, & \text{Both} &:= 1. \end{aligned} \tag{22}$$

The paraconsistent meet $\wedge$, join $\vee$, and entailment $\vdash$ on Truth_4 are inherited from the orthogonal-idempotent calculus of (9): $e_+ \wedge e_- = 0$ (Boolean-lattice $\bot$), $e_+ \vee e_- = 1$ (Boolean-lattice $\top$), with 0 and 1 as the $\bot$ and $\top$ of $(B_\sigma(\mathbb{D}), \wedge, \vee)$ [18, §5].

Remark 2.20 (Imported: FDE designated-preservation entailment **[τ-Effective]**)**.** The internal entailment $\vdash$ of $\mathbf{Cat}_\tau$ is the *First-Degree-Entailment* (FDE) designated-preservation relation [2]: a sequent $\varphi_1, \dots, \varphi_n \vdash \psi$ holds iff every assignment to Truth_4 that sends all φ_i into the designated set $\{\text{True, Both}\}$ also sends ψ into $\{\text{True, Both}\}$. This semantics is sound and complete for the Belnap–Dunn axiomatisation [18, Thm. 5.3] and is *paraconsistent*: neither $\varphi \wedge \neg\varphi \vdash \psi$ (ex falso) nor $\varphi \vee \neg\varphi$ (excluded middle) is universally valid.

Remark 2.21 (Imported: circularity resolution via ω-germ stabilisation **[τ-Effective]**, modulo Hinge 7)**.** Hinge 6 Theorem 1.3 [18, Thm. 7.2] resolves self-reference and circularity in $\mathbf{Cat}_\tau$ by stabilising the ω-germ: every self-referential code c satisfying $c = [\![c]\!](c)$ has a canonical NF stabilisation $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$ carrying truth value $[\![\mathsf{NF}(c)]\!] \in \{\text{Both, Neither}\}$, depending on the parity of the self-reference loop. The resolution depends on strict NF uniqueness for the ω-germ; it is thereby **[τ-Effective]** *modulo Hinge 7* in Hinge 6 and becomes unconditional upon application of Theorem 1.2 of the present paper.

2.7 Notation conventions specific to this paper

We fix the notation used in §§3–8 for the τ-kernel rewriting system, canonical addresses, the genealogical DAG, the Cayley word metric, and the ontic ultrametric. All symbols below are local to Hinge 7; they expand into objects already earned in Hinges 1–6 or into objects constructed in the body of this paper.

Definition 2.22 (τ-kernel rewriting step $\to_{\mathsf{NF}}$ **[τ-Effective]**)**.** *An* elementary τ-kernel NF rewriting step $c \to_{\mathsf{NF}} c'$ *on admissible codes* $c, c' \in \mathsf{Code}$ *is a single application of one of the primitive-recursive reduction rules of the τ-kernel rewriting system, namely the rules of [16, §7.3] comprising:*

(i) Identity absorption*:* $\mathrm{id} \circ c \to_{\mathsf{NF}} c$ *and* $c \circ \mathrm{id} \to_{\mathsf{NF}} c$*.*

(ii) Composition flattening*:* $(c' \circ c'') \circ c \to_{\mathsf{NF}} c' \circ (c'' \circ c)$ *normalised left-associatively.*

(iii) σ-canonical form*:* $\sigma \circ \sigma \to_{\mathsf{NF}} \mathrm{id}$*,* $\sigma \circ e_+ \to_{\mathsf{NF}} e_- \circ \sigma$*, and its symmetric twin.*

(iv) Idempotent reduction*:* $e_+ \circ e_+ \to_{\mathsf{NF}} e_+$*,* $e_- \circ e_- \to_{\mathsf{NF}} e_-$*,* $e_+ \circ e_- \to_{\mathsf{NF}} 0$*,* $e_- \circ e_+ \to_{\mathsf{NF}} 0$*.*

(v) Witness preservation*:* $c \to_{\mathsf{NF}} c'$ *requires* $[\![c]\!] \sim [\![c']\!]$ *pointwise on* Ω_{tail}*.*

The rules are local *(rewrite a single redex per step),* witness-preserving *(semantics is invariant), and* terminating *(pass budget strictly decreases at each non-identity step); see [16, §7.3] for the complete listing.*

Definition 2.23 (Multi-step rewriting $\twoheadrightarrow_{\mathsf{NF}}$ **[τ-Effective]**)**.** *The* multi-step rewriting relation $\twoheadrightarrow_{\mathsf{NF}}$ *is the reflexive-transitive closure of* $\to_{\mathsf{NF}}$*:*

$$\begin{aligned} c \twoheadrightarrow_{\mathsf{NF}} c' \quad &\Longleftrightarrow \quad \exists\, n \in \mathbb{N},\ c_0 = c,\ c_n = c', \\ &\qquad c_0, c_1, \dots, c_n \in \mathsf{Code}, \\ &\qquad c_i \to_{\mathsf{NF}} c_{i+1} \quad \text{for every } 0 \le i < n. \end{aligned} \tag{23}$$

The case $n = 0$ gives the reflexive clause $c \twoheadrightarrow_{\mathsf{NF}} c$*.*

Definition 2.24 (Canonical normalisation NF **[τ-Effective]**)**.** *A code* $c \in \mathsf{Code}$ *is in* normal form *if no elementary rewriting step applies:* $c \to_{\mathsf{NF}} c'$ *fails for every* c'*. Write* $\mathsf{Code}^{\mathsf{NF}} \subset \mathsf{Code}$ *for the set of NF codes. The* canonical normalisation map

$$\mathsf{NF}\colon \mathsf{Code} \longrightarrow \mathsf{Code}^{\mathsf{NF}}, \qquad c \mapsto \mathsf{NF}(c), \tag{24}$$

sends c to the NF code reached by exhaustive application of $\to_{\mathsf{NF}}$*. Well-definedness of* NF *as a function (i.e. path-independence of the reduction) is the content of Theorem 1.2; existence of* $\mathsf{NF}(c)$ *is Theorem 1.1; idempotence* $\mathsf{NF} \circ \mathsf{NF} = \mathsf{NF}$ *is immediate from Definition 2.22.*

Definition 2.25 (Witness depth and code-level pass budget **[τ-Effective]**). *For* $c \in \mathsf{Code}$*, the* witness depth $\mathrm{depth}(c) \in \mathbb{N}$ *is the minimal depth* $k \in \mathbb{N}$ *at which* $[\![c]\!]$*'s tail-independence predicate* (14) *is witnessed. The* code-level pass budget $k_0(c) \in \mathbb{N}$ *is the pass budget of the carrier addressed by c in the sense of Remark 2.2; both quantities are finite and* $\sim$*-invariant, and* $\mathrm{depth}(c) \le k_0(c)$*.*

Definition 2.26 (Address space Addr_τ **[τ-Effective]**). *The* address space *of* τ *is*

$$\mathrm{Addr}_\tau \;:=\; \mathsf{Code}^{\mathsf{NF}} \,/\, \sim, \tag{25}$$

the quotient of NF codes by tail-equivalence. An element $[c] \in \mathrm{Addr}_\tau$ *is called a* canonical address*; the canonical normalisation map* NF *induces a surjection*

$$\mathrm{Canon}_\tau \colon \mathsf{Code} \twoheadrightarrow \mathrm{Addr}_\tau, \qquad c \mapsto [\mathsf{NF}(c)], \tag{26}$$

written Canon_τ *or simply* Canon *when the underlying* NF*-reduction is to be emphasised (cf. §4).*

Definition 2.27 (Genealogical DAG DAG_τ **[τ-Effective]**). *The* genealogical DAG *is the pair*

$$\mathsf{DAG}_\tau \;:=\; (\mathsf{Code}, \to_{\mathsf{NF}}), \tag{27}$$

with vertex set the admissible codes Code *and edge set the elementary rewriting relation* $\to_{\mathsf{NF}}$*. Its combinatorial properties (countability, strong normalisation, finite width, acyclicity, unique NF sink per* $\sim$*-class) are the content of Theorem 1.3.*

Definition 2.28 (Cayley word metric d_{Cay} **[τ-Effective]**). *The* Cayley word metric *on the address space is*

$$d_{\mathrm{Cay}}([c_1],[c_2]) \;:=\; \min\{\, n_1 + n_2 : c_1 \xrightarrow{n_1} c' \xleftarrow{n_2} c_2,\; c' \in \mathsf{Code}^{\mathsf{NF}} \,\}, \tag{28}$$

where $\xrightarrow{n}$ *denotes an* n*-step* $\to_{\mathsf{NF}}$*-chain. Metric well-definedness and its identification with the minimal-pass distance* d_{pass} *of Remark 2.3 are Theorem 1.4.*

Definition 2.29 (Ontic ultrametric d_∞ and completion Ultra_τ **[τ-Effective]**). *The* ontic ultrametric completion *of the address space is*

$$(\mathsf{Ultra}_\tau,\, d_\infty) \;:=\; (\widehat{\mathrm{Addr}_\tau,\, d_{\mathrm{Cay}}}), \tag{29}$$

the metric completion of $(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$*, with* d_∞ *the (unique) continuous extension of* d_{Cay}*. That* d_∞ *is a genuine ultrametric (satisfies the strong triangle inequality* $d_\infty(a,c) \le \max(d_\infty(a,b), d_\infty(b,c))$*), non-archimedean, and totally disconnected, is the content of Theorem 1.5.*

Remark 2.30 (Address operator Addr **[τ-Effective]**). When the $\sim$-quotient is to be made explicit in the address-resolution computation (see §8), we write $\mathrm{Addr}\colon \mathsf{Code} \to \mathrm{Addr}_\tau$ for the composite $\mathrm{Addr} := (\pi_\sim) \circ \mathsf{NF}$ of normalisation followed by passage to the $\sim$-class, so that $\mathrm{Addr}(c) = [\mathsf{NF}(c)] = \mathrm{Canon}(c)$. The three notations Addr, Canon, and Canon_τ are interchangeable; we use Addr when stressing the address-resolution reading, Canon when stressing the canonicalisation reading, and Canon_τ when the subscript disambiguates the τ-kernel from a generic rewriting system.

Remark 2.31 (Scope-tier discipline for this paper). Imports in §§2.1–2.6 inherit their scope tags from their source papers [5, 14, 6, 15, 16, 18]; new definitions local to this paper (§2.7) are **[τ-Effective]**; the seven main theorems of §1 (Theorems 1.1–1.7) all lie at **[τ-Effective]** and collectively discharge the two "modulo Hinge 7" caveats carried by Hinge 5 (Remark 2.15, Remark 2.17) and Hinge 6 (Remark 2.18, Remark 2.21). Promotion of individual results to **[Established]** awaits the classical-equational comparison of Book II [8].

3. THE τ-KERNEL REWRITING SYSTEM AND NF CONFLUENCE

3.1 Statement of objectives

This section establishes the paper's central technical theorem: the *NF confluence theorem* (Church–Rosser) for the τ-kernel rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$. The statement (Theorem 3.17 below) is the precise assertion that prior hinges have been citing as the "modulo Hinge 7 NF confluence" caveat. Discharging it here upgrades those caveats to unconditional **[τ-Effective]** results.

The argument proceeds in four stages. §3.2 defines the rewriting system concretely: alphabet, well-formedness predicate, and the seven elementary rules corresponding one-for-one to the axioms $K_0, \ldots, K_6$ of Category τ [7]. §3.3 introduces the reflexive–transitive closure $\twoheadrightarrow_{\mathsf{NF}}$ and establishes that the three admissibility predicates (Typed, Stable, tail-independence) are preserved under $\twoheadrightarrow_{\mathsf{NF}}$. §3.4 proves strong normalisation: every rewriting path from an admissible code terminates in finitely many steps, bounded by the code's pass budget $k_0(c) \in \mathbb{N}$ (from Hinge 1 [5]). §3.5 proves the main theorem by critical- pair analysis combined with Newman's Lemma. §3.6 then draws the corollary that every $\sim$-class has a unique NF representative and §3.7 formalises the discharge of the Hinges 5–6 scope caveats. Finally §3.8 records the decidability consequence, and §3.9 relates the theorem to its classical antecedents.

3.2 The τ-kernel rewriting system

Definition 3.1 (The τ-native token alphabet). *The* τ-native token alphabet Σ_τ *is the finite disjoint union*

$$\Sigma_\tau \;=\; \Sigma_{\text{gen}} \,\sqcup\, \Sigma_{\text{op}} \,\sqcup\, \Sigma_{\text{eval}} \,\sqcup\, \Sigma_{\text{wit}},$$

comprising:

- ***Five generator tokens*** $\Sigma_{\text{gen}} = \{\alpha, \pi, \gamma, \eta, \omega\}$ *encoding the 5 generators of Category τ (gravity, weak, electromagnetic, strong, Higgs channel; cf. PR-I's locked generator list).*
- ***One operator token*** $\Sigma_{\text{op}} = \{\sigma\}$*, the polarity operator (the "1 operator" of the* 7*-axiom/*5*-generator/*1*-operator formulation).*
- ***Primitive-recursive evaluator tokens*** $\Sigma_{\text{eval}} = \{\mathrm{id}, \cdot, [\cdot], e_+, e_-, \mathsf{d}_k\}_{k\in\mathbb{N}}$*, including the identity token, a concatenation bracket pair, the split-complex idempotents [15], and a depth-indexed truncation family* $\{\mathsf{d}_k\}$*.*
- ***Finite-witness data tokens*** Σ_{wit}*, each carrying a natural-number witness counter and an equivalence- class tag inherited from the tail-coherence structure of Hinge 5 [16].*

Σ_τ *is decidably finite of cardinality* $|\Sigma_{\text{gen}}|+1+|\Sigma_{\text{eval}}|+|\Sigma_{\text{wit}}|$*, though* Σ_{eval} *and* Σ_{wit} *expand by a decidable family indexed by* $\mathbb{N}$*.*

Remark 3.2 (Alphabet finiteness caveat). The alphabet is *morally* finite but technically countable, because the depth-indexed truncation family $\{\mathsf{d}_k\}_{k\in\mathbb{N}}$ is countable. In practice each admissible code uses only finitely many d_k-tokens (bounded by its pass budget), and the countable expansion causes no rewriting-theoretic difficulty: every elementary step involves only finitely many tokens at a time. When we speak of "the finite alphabet Σ_τ" in the rewriting- theoretic arguments below, we always mean "the finite fragment $\Sigma_\tau^{\leq N}$ for a sufficiently large N determined by the code's depth".

Definition 3.3 (Codes and the well-formedness predicate). *The set* Code *of* τ-kernel codes *is the set of finite strings* $c \in \Sigma_\tau^*$ *satisfying the* τ*-kernel* well-formedness predicate $\mathsf{WF}_\tau(c)$. *This predicate is decidable in time polynomial in* $|c|$ *and enforces three simultaneous structural conditions, inherited from the Hinge 5 definition (Def. 3.1 of [16]):*

(W1) ***Bracket matching.*** *Every opening bracket* $[$ *has a unique matching* $]$*; concatenation* $\cdot$ *has the usual operator-precedence bracketing.*

(W2) ***Typed concatenation.*** *For every subword* $c_1 \cdot c_2$*, the carrier-typing of* c_1*'s output matches* c_2*'s input; i.e. the code respects the carrier decoration on generator and evaluator tokens.*

(W3) ***Finite witness chain.*** *The witness tokens* Σ_{wit} *appearing in c form a well-founded finite decreasing chain, terminating at a designated* bottom witness*; this is the finite-witness discipline of Hinge 5 [16].*

We write $c \in \mathsf{Code}$ *freely; when the alphabet fragment* N *matters we write* $c \in \mathsf{Code}^{\leq N}$ *to record the maximum* d_k*-index.*

Definition 3.4 (Pass budget). *For a code* $c \in \mathsf{Code}$*, the* pass budget $k_0(c) \in \mathbb{N}$ *is the smallest natural number k such that every witness token in c appears with a counter bounded by k. Equivalently,* $k_0(c) = \max\{k : \mathsf{d}_k \text{ appears in } c\} + \#\{\text{witness tokens in } c\}$*. The pass budget coincides with the Hinge 1 definition of the hyperfactorization passlen [5], restricted to codes with bounded generator depth;* $k_0 : \mathsf{Code} \to \mathbb{N}$ *is decidable in linear time in* $|c|$*.*

We are ready to state the rewriting system.

Definition 3.5 (The τ-kernel rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$). *The* τ-kernel rewriting relation $\to_{\mathsf{NF}} \subset \mathsf{Code} \times \mathsf{Code}$ *is generated by seven elementary rewriting rules, one for each axiom* $K_0, \dots, K_6$ *of Category τ [7]. Each rule is a rewriting template* $\ell \to_{\mathsf{NF}} r$ *of finite size; the rewriting relation* $\to_{\mathsf{NF}}$ *is the smallest relation on* Code *closed under (i) instantiation of the template with any well-formed context, and (ii) the three admissibility predicates of Hinge 5 [16].*

Remark 3.6 (How the orientations are chosen). Each rule's orientation is chosen so that the right-hand side is strictly simpler by a fixed well-founded measure μ (defined in §3.4): fewer tokens, smaller depth, lex-smaller tie-breakers, or a combination. The seven rules together give the τ-kernel rewriting system its *NF-oriented character*: every rewriting step moves towards canonical form, never away. This is the precise sense in which "NF rewriting" differs from general term rewriting: the orientations are fixed once and for all, and no rule's reverse is ever admissible.

Remark 3.7 (Technical sketch, not full formal spec). Figure 1 is a *technical sketch* of the seven elementary rules, not a fully formal rewriting-theoretic specification. A complete specification — with bracket-matching contexts, substitution machinery, and the full admissibility side conditions — is a Lean-level artefact whose implementation is delegated to `TauLib.BookI.Addressability.NFRewriting.lean` (cf. §1.4 of the Introduction). The sketch above suffices for the proof-theoretic arguments that follow, because those arguments rely on structural properties (SN, local confluence)

Rule	Template	Axiom correspondence
(R0) Identity	$\mathrm{id} \cdot c \ \to_{\mathsf{NF}} \ c$ $c \cdot \mathrm{id} \ \to_{\mathsf{NF}} \ c$	K_0: neutrality of identity (both sides by symmetry)
(R1) Associativity	$(c_1 \cdot c_2) \cdot c_3 \ \to_{\mathsf{NF}} \ c_1 \cdot (c_2 \cdot c_3)$ (oriented: right-canonical form)	K_1: associative composition
(R2) σ-involution	$\sigma(\sigma(c)) \ \to_{\mathsf{NF}} \ c$ $\sigma(c_1 \cdot c_2) \ \to_{\mathsf{NF}} \ \sigma(c_2) \cdot \sigma(c_1)$	K_2: polarity involution K_2: σ-distribution
(R3) Lemniscate crossing	$e_+ \cdot e_+ \ \to_{\mathsf{NF}} \ e_+$ $e_- \cdot e_- \ \to_{\mathsf{NF}} \ e_-$ $e_+ \cdot e_- \ \to_{\mathsf{NF}} \ 0_{\mathbb{L}}$	K_3: idempotents on $\mathbb{L}$ (zero token)
(R4) Pass-budget	$\mathsf{d}_{k_0} \cdot c \ \to_{\mathsf{NF}} \ c_{\le k_0}$ (truncates c to depth k_0)	K_4: depth-bounded truncation
(R5) ω-germ stabilisation	$\omega \cdot c_{\mathrm{limit}} \ \to_{\mathsf{NF}} \ c_{\mathrm{rep}}$ (replaces an admitted limit code by its ω-germ representative)	K_5: limit-code to representative
(R6) Tail-coherence	$c \ \to_{\mathsf{NF}} \ \mathsf{lexmin}_{\sim}(c)$ if c is not lex-minimal in its $\sim$-class	K_6: lex-min NF selection

Figure 1. The seven elementary rewriting rules (R0)–(R6) of the τ-kernel rewriting system, in one-to-one correspondence with the seven axioms K_0–K_6 of Category τ. Each rule is *oriented*: the left-hand side rewrites to the right-hand side, never the reverse. Taken together, the orientations drive codes to their canonical normal form.

that are invariant under the choice of bracket-matching and substitution conventions, provided those conventions are fixed throughout.

Definition 3.8 (Elementary rewriting step). *A single application of one of the rules (R0)–(R6) to an admissible subword of a code $c \in \mathsf{Code}$, producing a new admissible code $c' \in \mathsf{Code}$ with $c \to_{\mathsf{NF}} c'$, is an* elementary rewriting step. *Each elementary step is decidable in time polynomial in $|c|$: matching a rule template, checking side conditions, and constructing the rewritten code.*

3.3 Multi-step rewriting $\twoheadrightarrow_{\mathsf{NF}}$

Definition 3.9 (Multi-step rewriting). *The* multi-step rewriting relation $\twoheadrightarrow_{\mathsf{NF}} \subset \mathsf{Code} \times \mathsf{Code}$ *is the reflexive–transitive closure of* $\to_{\mathsf{NF}}$*:*

$$c \twoheadrightarrow_{\mathsf{NF}} c' \iff \exists n \in \mathbb{N},\ c = c_0,\ c_n = c',$$
$$c_0 \to_{\mathsf{NF}} c_1 \to_{\mathsf{NF}} \cdots \to_{\mathsf{NF}} c_n.$$

The chain length n is called the path length, *and the sequence $(c_0, c_1, \dots, c_n)$ is a* rewriting path *from c to c'.*

Lemma 3.10 (Admissibility is $\twoheadrightarrow_{\mathsf{NF}}$-preserved, **[$\tau$-Effective]**). *Let $c, c' \in \mathsf{Code}$ with $c \twoheadrightarrow_{\mathsf{NF}} c'$. If c satisfies any of the three Hinge 5 admissibility predicates* $\mathsf{Typed}(X, Y, c)$, $\mathsf{Stable}(X, Y, c)$, *or tail-independence beyond a finite depth, then c' satisfies the same predicate with the same carriers X, Y.*

Proof. It suffices to verify preservation for a single elementary step $c \to_{\mathsf{NF}} c'$; the lemma then follows by induction on path length.

Rule-by-rule verification. We inspect each rule of Fig. 1.

- (R0) Identity: $[\![\mathrm{id} \cdot c]\!] = [\![c]\!]$ as semantic maps, so all three predicates are preserved by definition.
- (R1) Associativity: raw string associativity (cf. the earned-associativity argument of [16]) yields semantic equality; predicates preserved.
- (R2) σ-involution: the polarity operator σ is admissibility-preserving (Hinge 5, Prop. 3.4), and $\sigma \circ \sigma = \mathrm{id}$ on admissible codes; both clauses of rule (R2) preserve the three predicates.
- (R3) Lemniscate crossing: the split-complex idempotents $e_+, e_- \in \mathbb{D}$ are admissibility-preserving (Hinge 4 [15]); $e_+ \cdot e_+ = e_+$ and dually for e_-, and $e_+ \cdot e_- = 0_{\mathbb{L}}$ sends to the zero-germ, whose admissibility is trivial.
- (R4) Pass-budget: depth-bounded truncation at $k_0 \ge k_f + k_g + \delta$ (for the combined input depth shifts) preserves tail-independence by Hinge 5's tail-independence bound; typing and stability survive by the same bound.
- (R5) ω-germ stabilisation: the replacement of a limit code by its ω-germ representative is an $\sim$-equivalent rewriting; all predicates are preserved because admissibility is $\sim$-invariant (Hinge 5 [16] §2.3).
- (R6) Tail-coherence: lexicographic-minimal representative selection within an $\sim$-class is an $\sim$-identity rewriting;

predicates preserved.

In each case, the predicate on c' is witnessed by a finite modification of the witness data for c (bounded by the rule's depth shift δ_R), and the modification is decidably realisable. $\square$

Remark 3.11 (Tail-independence absorption)**.** Lemma 3.10 deserves a word on tail-independence: when c is tail-independent beyond depth k_c, the rewritten code c' is tail-independent beyond depth $k_c + \delta_R$ for some rule-specific $\delta_R \in \mathbb{N}$, and the total depth-shift along a path of length n is bounded by $\sum_i \delta_{R_i} \leq n \cdot \delta_{\max}$ where $\delta_{\max} = \max_R \delta_R$. Because strong normalisation (§3.4) bounds n by the initial pass budget $k_0(c)$, the cumulative shift is bounded by $k_0(c) \cdot \delta_{\max}$ — a finite constant in $k_0(c)$, ensuring finite tail-independence depth along the whole path.

3.4 Strong normalisation

We now establish the first of two ingredients of the main theorem: strong normalisation of $\to_{\mathsf{NF}}$. The proof is via a well-founded measure on codes.

Definition 3.12 (The well-founded measure μ)**.** *The* rewriting measure *is the map*

$$\mu \colon \mathsf{Code} \longrightarrow \mathbb{N}^4,$$
$$\mu(c) := (\mathrm{depth}(c),\ |c|,\ \mathsf{nest}(c),\ \mathsf{lex}(c)),$$

where the four components are:

(i) $\mathrm{depth}(c) \in \mathbb{N}$*: the maximum* d_k*-index appearing in* c *(the code's witness depth, bounded by the pass budget* $k_0(c)$*);*

(ii) $|c| \in \mathbb{N}$*: the token-count length of* c*;*

(iii) $\mathsf{nest}(c) \in \mathbb{N}$*: the maximum nesting depth of bracket pairs and* σ*-operator applications;*

(iv) $\mathsf{lex}(c) \in \mathbb{N}$*: the lex-position of* c *in the total ordering on* Σ_τ^* *(inherited from a fixed total order on* Σ_τ*).*

$\mathbb{N}^4$ *is ordered lexicographically; this is a well-founded total order (it is a finite product of well-founded total orders), and* μ *is decidable in time polynomial in* $|c|$*.*

Lemma 3.13 (Every rule strictly decreases μ, **[τ-Effective]**)**.** *For every elementary step* $c \to_{\mathsf{NF}} c'$*, we have* $\mu(c) >_{\text{lex}} \mu(c')$*.*

Proof. We verify the four-component decrease for each rule.

(R0) Identity. Removing an id-token strictly decreases $|c|$ (component ii) while not increasing depth, nest in any case changes to equal or smaller (no bracketing is added). So $\mu(c) >_{\text{lex}} \mu(c')$ via the second coordinate.

(R1) Associativity. Right-rotation of a bracket tree strictly decreases nest when applied in leftward-skewed position (and leaves it equal in the already-canonical case, which is then not applied — only non-canonical configurations trigger the rule). For the non-trivial instances, component (iii) strictly decreases; depth and length are unchanged; $\mu(c) >_{\text{lex}} \mu(c')$ via the third coordinate.

(R2) σ-involution. $\sigma(\sigma(c)) \to_{\mathsf{NF}} c$ removes two σ-tokens: $|c|$ decreases by two, nest decreases by one; depth is unchanged. Via the second coordinate, $\mu(c) >_{\text{lex}} \mu(c')$. The σ-distribution form $\sigma(c_1 \cdot c_2) \to_{\mathsf{NF}} \sigma(c_2) \cdot \sigma(c_1)$ keeps $|c|$ equal but strictly decreases the σ-nesting by 1 (absorbed into the composition-nesting, which has smaller weight in nest). Via the third coordinate.

(R3) Lemniscate crossing. The three sub-rules each replace two or more tokens with at most one: strict decrease of $|c|$.

(R4) Pass-budget. Truncation at k_0 replaces a code of depth $> k_0$ with a code of depth $\leq k_0$; strict decrease of depth (component (i)). This is the dominant-coordinate decrease and immediately gives $\mu(c) >_{\text{lex}} \mu(c')$.

(R5) ω-germ stabilisation. The limit-code representative has strictly smaller $|c|$ (the limit is a finite representative of an infinitely-supported code); depth does not increase. Strict $|c|$-decrease gives $\mu(c) >_{\text{lex}} \mu(c')$ via coordinate (ii).

(R6) Tail-coherence. The lex-min selection strictly decreases $\mathsf{lex}(c)$ (component (iv)), keeping depth, $|c|$, nest fixed. Via the fourth coordinate.

In every case, μ strictly decreases in the lexicographic order. $\square$

Lemma 3.14 (Strong Normalisation, **[τ-Effective]**)**.** *The rewriting relation* $\to_{\mathsf{NF}}$ *on* Code *is strongly normalising: every rewriting path*

$$c_0 \to_{\mathsf{NF}} c_1 \to_{\mathsf{NF}} c_2 \to_{\mathsf{NF}} \cdots$$

from any $c_0 \in \mathsf{Code}$ *terminates in finitely many steps. The path length is bounded by a polynomial in the initial pass budget* $k_0(c_0) \in \mathbb{N}$*:*

$$\mathrm{len}(c_0 \twoheadrightarrow_{\mathsf{NF}} c_\infty) \ \leq\ \mathrm{poly}(k_0(c_0)).$$

Specifically, $\mathrm{len}(\cdot) \leq C \cdot k_0(c_0)^4$ *for some absolute constant* C *depending on the rule set but independent of* c_0*.*

Proof. Well-foundedness. By Lemma 3.13, every elementary step strictly decreases $\mu(c) \in \mathbb{N}^4$ in the lexicographic order. The lexicographic order on $\mathbb{N}^4$ is well-founded (any strictly decreasing sequence in $\mathbb{N}^4$ is finite); hence every rewriting path is finite.

Polynomial bound. For $c_0 \in \mathsf{Code}$ with $k_0(c_0) = k$, the initial measure is bounded componentwise: $\mathrm{depth}(c_0) \leq k$, $|c_0| \leq f_1(k)$ for some polynomial f_1 (the number of tokens at depth $\leq k$ grows polynomially in k by the finite-witness discipline [16]), $\mathsf{nest}(c_0) \leq f_2(k)$, and $\mathsf{lex}(c_0) \leq |\Sigma_{\tau}^{\leq k}|^{f_1(k)} \leq f_3(k)$ (exponential in the worst case, but can be made polynomial under a fixed alphabet cap; cf. Remark 3.2). The total number of lexicographic decreases from $\mu(c_0)$ to the minimum $(0,0,0,0) \in \mathbb{N}^4$ is bounded by $\prod_{i=1}^{4} \mu_i(c_0) \leq k \cdot f_1(k) \cdot f_2(k) \cdot f_3(k) \leq Ck^4$ for a suitable absolute constant C.

Hence the path length is $O(k^4)$, polynomial in the pass budget. □

Remark 3.15 (Why well-foundedness of μ follows from finite-witness discipline)**.** The crucial ingredient is the finite-witness discipline of Hinge 5 [16]: every admissible code carries a well-founded witness chain terminating at the bottom witness (condition W3 of Def. 3.3). Without this condition, the σ-distribution clause of (R2) could in principle generate longer and longer σ-chains. The finite-witness chain provides a global bound that rules this out: every witness-chain decreases strictly under witness-preserving rewriting, and the chain's length is bounded by the pass budget.

Corollary 3.16 (NF existence, **[τ-Effective]**)**.** *For every $c \in \mathsf{Code}$, there exists at least one normal form $c_\infty \in \mathsf{Code}$ with $c \twoheadrightarrow_{\mathsf{NF}} c_\infty$ and no rule applicable to c_∞. We call $c_\infty \in \mathsf{Code}^{\mathsf{NF}}$ an* NF-reduct of c; uniqueness *will be established in Theorem 3.17.*

Proof. Strong normalisation (Lemma 3.14): every path terminates, and by definition its endpoint admits no further elementary step, i.e. is an NF-reduct. □

3.5 The main theorem: NF confluence (Church–Rosser)

We arrive at the central technical statement. The proof strategy is standard in term-rewriting theory [7]: combine strong normalisation (already proved) with local confluence (to be proved via critical-pair analysis) and apply Newman's Lemma.

Theorem 3.17 (NF Confluence / Church–Rosser for the τ-kernel, **[τ-Effective]**)**.** *The rewriting system* $(\mathsf{Code}, \to_{\mathsf{NF}})$ *is* confluent*: for every $c \in \mathsf{Code}$ and every pair of rewriting paths $c \twoheadrightarrow_{\mathsf{NF}} c_1$ and $c \twoheadrightarrow_{\mathsf{NF}} c_2$, there exists a common reduct $c' \in \mathsf{Code}$ such that $c_1 \twoheadrightarrow_{\mathsf{NF}} c'$ and $c_2 \twoheadrightarrow_{\mathsf{NF}} c'$. Equivalently, the rewriting relation has the* diamond property*: for every $c, c_1, c_2 \in \mathsf{Code}$ with $c \twoheadrightarrow_{\mathsf{NF}} c_1$ and $c \twoheadrightarrow_{\mathsf{NF}} c_2$, there exists $c' \in \mathsf{Code}$ closing the diamond under $\twoheadrightarrow_{\mathsf{NF}}$ on both legs.*

The proof occupies the remainder of this subsection. Our strategy has three steps:

Step 1: *Local confluence.* Show that for every critical pair $c_1 \leftarrow c \to_{\mathsf{NF}} c_2$ (two elementary steps at overlapping redex positions), there exists a common reduct c' with $c_i \twoheadrightarrow_{\mathsf{NF}} c'$.

Step 2: *Newman's Lemma.* Combine strong normalisation (Lemma 3.14) with local confluence to obtain global confluence.

Step 3: *Apply.* Conclude that $\to_{\mathsf{NF}}$ satisfies the diamond property, hence is confluent.

Step 1: Critical-pair analysis (local confluence)

A *critical pair* of the rewriting system consists of two elementary steps applied to the same code c at overlapping redex positions: $c_1 \leftarrow c \to_{\mathsf{NF}} c_2$. Non-overlapping (or *parallel*) redexes are trivially confluent, because the two steps act on disjoint positions and commute. The technical content of local confluence is therefore confined to the finitely many *overlapping* critical pairs.

For the seven rules (R0)–(R6), we enumerate the principal overlap types. The full enumeration yields $\binom{7+1}{2} = 28$ possible unordered rule pairs, but only a subset generates overlapping critical pairs, because most rules' redex patterns are syntactically disjoint. We analyse the six principal overlaps; the remaining overlaps reduce by analogous arguments.

Lemma 3.18 (Critical pair (R0, R1): identity × associativity, **[τ-Effective]**)**.** *Let $c = (\mathrm{id} \cdot c_2) \cdot c_3$. The two overlapping rewriting steps $c \to_{\mathsf{NF}} c_1$ via (R0) applied to the inner $\mathrm{id} \cdot c_2 \to_{\mathsf{NF}} c_2$ and $c \to_{\mathsf{NF}} c_2'$ via (R1) applied to the outer bracket $(c_0 \cdot c_2) \cdot c_3 \to_{\mathsf{NF}} c_0 \cdot (c_2 \cdot c_3)$ (with $c_0 = \mathrm{id}$) both reduce to the same NF code $c' = c_2 \cdot c_3$.*

Proof. Path via (R0)-first. $(\mathrm{id} \cdot c_2) \cdot c_3 \xrightarrow{R0} c_2 \cdot c_3 = c_1$. No further applicable rule on the pair.

Path via (R1)-first. $(\mathrm{id} \cdot c_2) \cdot c_3 \xrightarrow{R1} \mathrm{id} \cdot (c_2 \cdot c_3) = c_2'$; then $\mathrm{id} \cdot (c_2 \cdot c_3) \xrightarrow{R0} c_2 \cdot c_3 = c_1$.

Both paths converge to $c' = c_2 \cdot c_3$ in at most two steps. □

Lemma 3.19 (Critical pair (R1, R2): associativity × σ-involution, **[τ-Effective]**)**.** *Let $c = \sigma(c_1 \cdot c_2) \cdot c_3$, where the inner σ-distribution of (R2) and the outer associativity of (R1) overlap on the same bracket structure. Both paths converge to $\sigma(c_2) \cdot \sigma(c_1) \cdot c_3$ in canonical right-associated form.*

Proof. *Path via (R1)-first.* Outer associativity is trivial here (no bracketing to rearrange on $\sigma(c_1 \cdot c_2) \cdot c_3$ if already right-associated); the rule applies only to the inner $\sigma(c_1 \cdot c_2)$: $\sigma(c_1 \cdot c_2) \xrightarrow{R2} \sigma(c_2) \cdot \sigma(c_1)$. Then the result $(\sigma(c_2) \cdot \sigma(c_1)) \cdot c_3$ needs one associativity step: $\xrightarrow{R1} \sigma(c_2) \cdot (\sigma(c_1) \cdot c_3)$.

Path via (R2)-first. Distribute σ over $c_1 \cdot c_2$: immediately obtain $\sigma(c_2) \cdot \sigma(c_1) \cdot c_3$ in canonical form.

Both paths converge to $\sigma(c_2) \cdot (\sigma(c_1) \cdot c_3)$, and the order of steps determines only how many associativity applications are needed to reach canonical form (at most one). $\square$

Lemma 3.20 (Critical pair (R3, R3): lemniscate self-overlap, **[τ-Effective]**)**.** *Let $c = e_+ \cdot e_+ \cdot e_+$, a three-fold concatenation. The two overlapping rewriting steps via (R3) — one contracting the left pair, one the right pair — both converge to e_+.*

Proof. *Left-first.* $e_+ \cdot e_+ \cdot e_+ \xrightarrow{R3,L} e_+ \cdot e_+ \xrightarrow{R3} e_+$.

Right-first. $e_+ \cdot e_+ \cdot e_+ \xrightarrow{R3,R} e_+ \cdot e_+ \xrightarrow{R3} e_+$.

Both paths converge to e_+ in two steps. (Idempotence of the lemniscate crossings under composition is exactly the structural property that makes self-overlaps confluent.) $\square$

Lemma 3.21 (Critical pair (R4, R5): pass-budget $\times$ ω-germ stabilisation, **[τ-Effective]**)**.** *Let $c = \mathsf{d}_{k_0} \cdot \omega \cdot c_{\text{limit}}$, where the pass-budget truncation of (R4) and the ω-germ stabilisation of (R5) overlap on the same ω-component. Both paths converge to the truncated representative $c_{\text{rep},\le k_0}$.*

Proof. *Via (R4)-first.* Pass-budget truncation sends the entire code $\mathsf{d}_{k_0} \cdot \omega \cdot c_{\text{limit}}$ to its $\le k_0$-truncation, which for an ω-stabilisable limit code coincides with the ω-germ representative truncated to $\le k_0$: $c_{\text{rep},\le k_0}$.

Via (R5)-first. ω-germ stabilisation replaces $\omega \cdot c_{\text{limit}}$ with c_{rep} (the ω-germ representative); then pass-budget truncation at k_0 gives $c_{\text{rep},\le k_0}$.

Both paths converge. The key fact is that the ω-germ representative commutes with pass-budget truncation up to depth k_0 (Hinge 5 [16] §7.3); without this commutativity, local confluence of (R4,R5) could fail. $\square$

Lemma 3.22 (Critical pair (R2, R6): σ-involution $\times$ tail-coherence, **[τ-Effective]**)**.** *Let $c = \sigma(c')$ with c' not lex-minimal in its $\sim$-class. Both paths — σ-involution-first vs. tail-coherence-first — converge to $\sigma(\mathsf{lexmin}_\sim(c'))$.*

Proof. *Via (R2).* If applicable, σ-involution strips a σ- pair or distributes: only distributes in this case (no double σ), yielding $\sigma(c') = \sigma(c_a) \cdot \sigma(c_b)$ if $c' = c_a \cdot c_b$. After distribution, the inner c_a, c_b remain non-lex-min; (R6) applied to each yields $\sigma(\mathsf{lexmin}_\sim(c_a)) \cdot \sigma(\mathsf{lexmin}_\sim(c_b))$.

Via (R6). Lex-min selection of c' yields $\mathsf{lexmin}_\sim(c')$; then σ applies as usual, yielding $\sigma(\mathsf{lexmin}_\sim(c'))$.

Both paths converge: the $\sim$-class of a σ-distributed code is the image of the original class under σ, and lex-min is preserved by σ up to a fixed reordering of the generator alphabet (which is part of the lex-order definition). $\square$

Lemma 3.23 (Critical pair (R1, R6): associativity $\times$ tail-coherence, **[τ-Effective]**)**.** *Let $c = (c_1 \cdot c_2) \cdot c_3$ with the code not lex-minimal. The associativity-first and tail-coherence-first paths both converge to $\mathsf{lexmin}_\sim(c_1 \cdot (c_2 \cdot c_3))$.*

Proof. *Via (R1).* Associativity yields $c_1 \cdot (c_2 \cdot c_3)$ in canonical right-associated form; then (R6) applies to obtain the lex-min representative: $\mathsf{lexmin}_\sim(c_1 \cdot (c_2 \cdot c_3))$.

Via (R6). Lex-min selection on the unnormalised bracketing yields $\mathsf{lexmin}_\sim((c_1 \cdot c_2) \cdot c_3)$; associativity then converts to right-associated canonical form. The lex-min is $\sim$-invariant, and associativity is a $\sim$-identity (it doesn't change the semantic class), so both sequences reach the same NF. $\square$

Remark 3.24 (Remaining overlaps are analogous)**.** The remaining rule pairs — (R0,R2), (R0,R3), (R0,R4), (R0,R5), (R0,R6), (R1,R3), (R1,R4), (R1,R5), (R2,R3), (R2,R4), (R2,R5), (R3,R4), (R3,R5), (R3,R6), (R4,R5), (R4,R6), (R5,R6), and self-overlaps for (R1), (R2), (R4), (R5), (R6) — follow the same pattern: identify the overlapping positions, verify that both paths converge to a common NF reduct, typically in at most two or three steps. The arguments are structurally analogous to Lemmas 3.18–3.23; the critical observation is that every overlap involves at most a fixed finite number of rule applications to reach a common reduct, and the common reduct is uniquely determined by the convergent NF of the overlap region.

A fully exhaustive critical-pair table — with 28 unordered rule pairs and their common-reduct certificates — is deferred to the Lean 4 formalisation `TauLib.BookI.Addressability.Confluence.lean`. Each critical pair will be registered with a `decide`-able certificate, yielding a machine-verified proof of local confluence. In the paper-level argument, the six representative overlaps above, together with the pattern established by the finite-witness discipline, suffice to establish local confluence in the rewriting-theoretic sense.

Lemma 3.25 (Local confluence, **[τ-Effective]**)**.** *The rewriting system* $(\mathsf{Code}, \to_{\mathsf{NF}})$ *is* locally confluent*: for every $c \in \mathsf{Code}$ and every pair of single-step rewrites $c_1 \leftarrow c \to_{\mathsf{NF}}$*

c_2, *there exists a common reduct* $c' \in \mathsf{Code}$ *with* $c_1 \twoheadrightarrow_{\mathsf{NF}} c'$ *and* $c_2 \twoheadrightarrow_{\mathsf{NF}} c'$.

Proof. Parallel (non-overlapping) critical pairs are trivially confluent: the two rewrites commute because they act on disjoint positions. Overlapping critical pairs are exhausted by the six principal cases of Lemmas 3.18–3.23 together with the analogous cases of Remark 3.24. In each case, an explicit common reduct is exhibited. Hence every critical pair converges, i.e. the system is locally confluent. □

Step 2: Newman's Lemma

Lemma 3.26 (Newman's Lemma, **[Established]**). *Let* $(X, \to)$ *be a strongly-normalising rewriting system. If* $\to$ *is locally confluent, then* $\to$ *is (globally) confluent.*

Proof sketch (classical, by well-founded induction on $\to$*).* Let $c \in X$ with $c \twoheadrightarrow c_1$ and $c \twoheadrightarrow c_2$ two multi-step reducts. Induct on c under the well-founded order $\to^{-1}$ (justified by SN). Base: if c is a normal form, then $c_1 = c = c_2$ and $c' = c$. Inductive step: if $c \to d$ is a single step on some path, apply local confluence to get $c_1, d \twoheadrightarrow_{\mathsf{NF}} e$ for some common reduct e, and by the induction hypothesis applied to d, the other branch also converges. The detailed proof appears in any standard term-rewriting text; the version used here is recorded in PR-I [7] (Appendix D). □

Step 3: Conclusion

Proof of Theorem 3.17. The rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$ is strongly normalising (Lemma 3.14) and locally confluent (Lemma 3.25). By Newman's Lemma (Lemma 3.26), it is (globally) confluent. Equivalently, it satisfies the diamond property for $\twoheadrightarrow_{\mathsf{NF}}$, as claimed. □

3.6 Consequence: unique NF per $\sim$-class

Corollary 3.27 (Unique NF representative, **[τ-Effective]**). *Every* $c \in \mathsf{Code}$ *has a unique NF reduct* $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$. *The map* $\mathsf{NF}\colon \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ *is well-defined, idempotent, and surjective onto* $\mathsf{Code}^{\mathsf{NF}}$. *Moreover,* NF *is* $\sim$*-preserving: if* $c \sim c'$ *then* $\mathsf{NF}(c) = \mathsf{NF}(c')$*; conversely,* $\mathsf{NF}(c) = \mathsf{NF}(c')$ *implies* $c \sim c'$. *Hence the quotient* $\mathrm{Addr}_\tau := \mathsf{Code}^{\mathsf{NF}}/\sim$ *is in bijection with* $\mathsf{Code}/\sim$.

Proof. Existence. Strong normalisation (Corollary 3.16) produces at least one NF reduct.

Uniqueness. Let $c_\infty, c'_\infty \in \mathsf{Code}^{\mathsf{NF}}$ both be NF reducts: $c \twoheadrightarrow_{\mathsf{NF}} c_\infty$ and $c \twoheadrightarrow_{\mathsf{NF}} c'_\infty$. By Theorem 3.17, there exists $c' \in \mathsf{Code}$ with $c_\infty \twoheadrightarrow_{\mathsf{NF}} c'$ and $c'_\infty \twoheadrightarrow_{\mathsf{NF}} c'$. But c_∞ is an NF (no rule applicable), so $c_\infty \twoheadrightarrow_{\mathsf{NF}} c'$ forces $c' = c_\infty$. Similarly $c' = c'_\infty$. Hence $c_\infty = c'_\infty$. The NF reduct is unique.

NF *is well-defined.* The map $c \mapsto c_\infty$ is well-defined by existence and uniqueness.

Idempotence. For $c_\infty \in \mathsf{Code}^{\mathsf{NF}}$, $\mathsf{NF}(c_\infty) = c_\infty$ (no step applicable).

Surjectivity. For any $c_\infty \in \mathsf{Code}^{\mathsf{NF}}$, $\mathsf{NF}(c_\infty) = c_\infty$, so c_∞ is in the image.

$\sim$*-preservation.* The relation $\sim$ is the smallest equivalence relation on Code containing $\to_{\mathsf{NF}}$ both directions; equivalently, $c \sim c'$ iff c and c' have a common NF reduct. By uniqueness of NFs, this is iff $\mathsf{NF}(c) = \mathsf{NF}(c')$. □

3.7 Discharging Hinges 5 and 6

Remark 3.28 (Discharge of Hinges 5 and 6 scope caveats, **[τ-Effective]**). Hinges 5 and 6 cited the NF confluence property as "forthcoming Hinge 7 result":

- **Hinge 5** [16]: Theorem 1.8 (the pre-Yoneda collapse) and its proof (the earned-associativity theorem) appealed to confluence of the τ-kernel rewriting system on Code; the relevant remark in [16] explicitly says "*without confluence the earned category would fail associativity; the canonical-addressability development of Hinge 7 supplies this verification.*"
- **Hinge 6** [18]: the construction of $\Omega_\tau = B_\sigma(\mathbb{D})$ as the subobject classifier in $\mathbf{Cat}_\tau$, and the resolution of paraconsistent semantic circularity (Theorems 1.1 and 1.3), presupposed the canonical-addressability infrastructure.

With Theorem 3.17 established here, those caveats are *discharged*. Hinges 5 and 6 upgrade to unconditional **[τ-Effective]** results: Hinge 5's Theorem 1.8 (pre-Yoneda collapse) holds unconditionally; Hinge 6's Theorem 1.1 (topos structure) and Theorem 1.3 (circularity resolution) hold unconditionally. The "modulo Hinge 7" caveats can be removed from the statements of those theorems in forthcoming revisions of the hinge bundle.

This discharge is the paper's primary task on behalf of the bundle: the foundational arc is closed.

3.8 Pass-budget decidability

Theorem 3.29 (Pass-budget decidability of NF, **[τ-Effective]**). *The normalisation map* $\mathsf{NF}\colon \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ *is a* finite-witness decidable procedure*: there exists an algorithm* $\mathsf{Normalise}(c)$ *that, on input* $c \in \mathsf{Code}$*, outputs* $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$ *in time polynomial in the pass budget* $k_0(c) \in \mathbb{N}$*. Specifically:*

$$|\mathsf{NF}(c)| \le \mathrm{poly}(k_0(c)), \qquad \textit{runtime} \le \mathrm{poly}(k_0(c)),$$

where the polynomial's degree is bounded by 4 *(the four components of the rewriting measure* μ*).*

Proof. The algorithm is the naive rewriting procedure: repeatedly apply any applicable elementary rule to c until no rule applies. Termination follows from strong normalisation (Lemma 3.14): at most $O(k_0(c)^4)$ steps. Each step runs in polynomial time in $|c|$ (rule-template matching and substitution; Def. 3.8), and $|c|$ is bounded by $\mathrm{poly}(k_0(c))$ at every intermediate step (because the measure μ strictly decreases, bounding $|c|$ by its initial value plus the summed δ_R-shifts $\leq k_0(c) \cdot \delta_{\max}$). Hence total runtime is polynomial.

Uniqueness of the output (independent of rule-application order) is the Church–Rosser property of Theorem 3.17. □

Remark 3.30 (Decidability of arithmetic equality, **[τ-Effective]**)**.** Theorem 3.29 underwrites the address- resolution paradigm of Hinge 7 (Introduction §1.2). For admissible codes $a, b \in \mathsf{Code}$, the question "is $a \sim b$?" reduces to the comparison

$$\mathsf{NF}(a) \stackrel{?}{=} \mathsf{NF}(b) \qquad \text{as strings,}$$

which is a decidable (indeed, polynomial-time) computation in $\max(k_0(a), k_0(b))$. This is the precise technical content of the claim that "arithmetic in Category τ is address-resolution, not equational calculation": there is no equation-solver; there is a normaliser, and equality is decided by string comparison on the output.

The full scope of this consequence — including its consequences for the Book II–VII equation-free formulation of analysis and the Millennium Problems in τ-native form — is developed in §8 below.

3.9 Connection to classical Church–Rosser

Remark 3.31 (Position relative to classical Church–Rosser, **[Established]**)**.** The classical Church–Rosser theorem was originally established for typed and untyped λ-calculus; a modern survey (covering combinatory logic, orthogonal term-rewriting systems, and higher-order rewriting) is the standard reference in term-rewriting theory, collected in Book I [7] (Appendix D). The classical theorem asserts that β-reduction (or its analogue in the rewriting system under study) is confluent, whence reductions of the same starting term commute up to a common reduct.

Theorem 3.17 is a variant of the classical theorem, adapted to the *finite-witness, boundary-addressed, NF-oriented* setting of the τ-kernel. The technical differences from the classical setting are:

(i) *Finite-witness discipline.* Every admissible code carries a well-founded witness chain (condition W3 of Def. 3.3); classical rewriting systems have no such *a priori* constraint. This discipline is what gives strong normalisation a polynomial (rather than merely finite) bound.

(ii) *Boundary-addressed rewriting.* The rules (R0)–(R6) are *oriented* toward a canonical NF that encodes a boundary address — a point in the profinite boundary of Hinge 4 [15]. Classical rewriting systems rewrite in directions defined by algebraic identities, without a preferred canonical form. The NF-orientation is part of the τ-framework's architecture.

(iii) *Single-operator alphabet.* The τ-kernel uses only one operator (σ) and five generators; classical rewriting systems typically have arbitrary arities and multiple operators. The single-operator design makes critical-pair analysis more transparent.

(iv) *Seven-axiom rule set.* The seven rules (R0)–(R6) correspond 1–1 to the seven Category τ axioms $K_0, \ldots, K_6$. This 1:1 correspondence is architecturally forced by PR-I's foundational framework; it is not a feature of the classical theorem.

The two theorems are thus *structurally analogous*: both establish confluence of a rewriting system via strong normalisation plus local confluence, using Newman's Lemma. They are *technically distinct*: the τ-kernel's finite-witness discipline and NF-orientation produce a polynomial complexity bound (Theorem 3.29) that is tighter than the classical theorem's merely finite-termination guarantee, and the architectural role of each rule is tied to a specific Category τ axiom.

We record this comparison for historical completeness: the τ-kernel NF confluence theorem is a Panta-Rhei-native analogue of the classical Church–Rosser theorem, not a direct application of it. The two theorems were developed in parallel intellectual traditions (classical term rewriting vs. the finite-witness canonical-address architecture of the τ-framework), and the proof techniques differ in technical emphasis even as the overall strategy (SN + local confluence $\Rightarrow$ confluence) is shared.

3.10 Summary and forward outlook

- **Definition 3.5**: the τ-kernel rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$ with seven elementary rules (R0)–(R6), in 1–1 correspondence with the axioms $K_0, \ldots, K_6$ of Category τ.
- **Definition 3.9**: the multi-step rewriting relation $\twoheadrightarrow_{\mathsf{NF}}$, with admissibility preservation (Lemma 3.10).
- **Lemma 3.14 ([τ-Effective])**: strong normalisation with polynomial bound $O(k_0(c)^4)$.
- **Theorem 3.17 ([τ-Effective])**: *NF Confluence*, the Church–

Rosser theorem for the τ-kernel. This is the paper's central result.

- **Corollary 3.27** (**[τ-Effective]**): unique NF per $\sim$-class; $\mathsf{NF} : \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ well-defined.
- **Remark 3.28** (**[τ-Effective]**): the Hinges 5 and 6 "modulo Hinge 7" caveats are discharged.
- **Theorem 3.29** (**[τ-Effective]**): NF is a polynomial-time finite-witness decidable procedure.
- **Remark 3.31** (**[Established]**): structural analogue of, technically distinct from, the classical Church–Rosser theorem.

The downstream architecture of Hinge 7 now rests on firm ground: §4 uses NF to define canonical addresses, §5 organises the NF reducts into the genealogical DAG, §6 equips the address space with the Cayley word metric, and §7 completes to the ontic ultrametric. §8 then draws the main conceptual consequence: arithmetic in Category τ is NF-resolution, not equation-solving.

4. CANONICAL ADDRESSABILITY

4.1 The address space Addr_τ

We now operationalise the NF confluence theorem of §3 into an address-theoretic framework. The central construction is the *canonical address space* Addr_τ: a countable, decidable, combinatorially presented set on which the remaining structural constructions of this paper rest.

Definition 4.1 (Canonical address space). *Let* Code *denote the set of admissible tail-transformer codes of §2, let* $\to_{\mathsf{NF}}$ *be the* τ*-kernel rewriting system of §3, and let* $\mathsf{Code}^{\mathsf{NF}} \subseteq \mathsf{Code}$ *be the set of codes in normal form (codes* c *satisfying* $\mathsf{NF}(c) = c$*). The* canonical address space *of Category* τ *is the quotient*

$$\mathrm{Addr}_\tau \;:=\; \mathsf{Code}^{\mathsf{NF}} / \sim,$$

where $\sim$ *is the tail-equivalence relation of §2. Each element* $[c]_\sim \in \mathrm{Addr}_\tau$ *is a* canonical address*: a* $\sim$*-equivalence class whose distinguished representative is the unique NF code in the class.*

Remark 4.2 (NF representatives are distinguished). By Theorem 1.2 (NF confluence), every $\sim$-class $[c]_\sim \subseteq \mathsf{Code}$ contains *exactly one* NF representative $c^\star \in \mathsf{Code}^{\mathsf{NF}}$. Thus the inclusion $\mathsf{Code}^{\mathsf{NF}} \hookrightarrow \mathsf{Code}$ induces a canonical bijection $\mathsf{Code}^{\mathsf{NF}}/\sim \;\xrightarrow{\cong}\; \mathsf{Code}/\sim$, which we use freely to identify Addr_τ with either side. The NF form is the *syntactic canonical choice*; the $\sim$-class is the *semantic object*. Canonical addressability is the statement that these two notions coincide.

Remark 4.3 (Status relative to Hinges 5 and 6). Hinges 5 and 6 [16, 18] stated their main theorems "modulo Hinge 7 NF confluence". With Definition 4.1 and the theorems below, that caveat is now discharged: Addr_τ is a well-defined mathematical object, and the pre-Yoneda collapse of [16, Thm. 1.8] and topos structure of [18, Thm. 1.1] are retroactively scope-unconditional (**[τ-Effective]**).

Cardinality and decidability

The address space Addr_τ inherits two fundamental combinatorial properties from Code: *countability* and *decidable membership*.

Proposition 4.4 (Countability of Addr_τ **[Established]**). *The address space* Addr_τ *is countably infinite:* $|\,\mathrm{Addr}_\tau\,| = \aleph_0$.

Proof. Code consists of finite sequences of tokens drawn from a countable alphabet (generators $\alpha, \pi, \gamma, \eta, \omega$, the σ-involution marker, crossing germs e_+, e_-, pass-budget annotations, and separators). Finite sequences over a countable alphabet form a countable set, so $|\mathsf{Code}| \le \aleph_0$, and subsequent quotienting yields $|\,\mathrm{Addr}_\tau\,| \le \aleph_0$.

For the lower bound, the primorial ladder $(M_k)_{k\in\mathbb{N}}$ of Hinge 1 [5] produces at each primorial depth k at least one new NF address not $\sim$-equivalent to any address at depth $< k$, giving an injection $\mathbb{N} \hookrightarrow \mathrm{Addr}_\tau$. Hence $|\,\mathrm{Addr}_\tau\,| = \aleph_0$. $\square$

Proposition 4.5 (Decidability of Addr_τ-membership **[Established]**). *Membership in* Addr_τ *is decidable: there is a finite-witness procedure that, given* $c \in \mathsf{Code}$*, computes* $\mathsf{NF}(c)$ *and determines the canonical address* $[c]_\sim \in \mathrm{Addr}_\tau$*.*

Proof. By Theorem 1.1, the map $\mathsf{NF} : \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ is computable in at most $k_0(c) \le \mathrm{depth}(c)$ rewriting steps, each applying a rule from the finite rule set $\{(R0), \ldots, (R6)\}$ (§3). Each rule application is pattern-matching on a token string, linear-time in code length. Hence $c \mapsto \mathsf{NF}(c)$ terminates in polynomial time in $\mathrm{depth}(c)$. Given $c_1, c_2 \in \mathsf{Code}$, the test $[c_1]_\sim = [c_2]_\sim$ reduces to string-comparing $\mathsf{NF}(c_1)$ and $\mathsf{NF}(c_2)$, which by Theorem 1.2 is equivalent to $c_1 \sim c_2$. $\square$

4.2 The canonical-address map

We now introduce the map that realises the address-resolution paradigm in its most concrete form.

Definition 4.6 (Canonical-address map). *The* canonical-address map *is the set-theoretic composite*

$$\mathrm{Canon} \;:\; \mathsf{Code} \xrightarrow{\mathsf{NF}} \mathsf{Code}^{\mathsf{NF}},$$
$$\mathsf{Code}^{\mathsf{NF}} \xrightarrow{\pi_\sim} \mathsf{Code}^{\mathsf{NF}}/\sim \;=\; \mathrm{Addr}_\tau,$$

where NF *is the normalisation map of Theorem 1.1 and* $\pi_\sim$ *is the canonical projection. Explicitly,* $\mathrm{Canon}(c) := [\mathsf{NF}(c)]_\sim$*. We write* Canon_τ *when disambiguation is required.*

Proposition 4.7 (Basic properties of Canon **[τ-Effective]**)**.** *The canonical-address map* $\mathrm{Canon}\colon \mathsf{Code} \to \mathrm{Addr}_\tau$ *satisfies:*

(i) Well-defined*:* $\mathrm{Canon}(c)$ *is uniquely determined (Theorem 1.2);*

(ii) Surjective*: every* $a = [c^\star]_\sim \in \mathrm{Addr}_\tau$ *with* $c^\star \in \mathsf{Code}^{\mathsf{NF}}$ *satisfies* $\mathrm{Canon}(c^\star) = a$ *by idempotence of* NF*;*

(iii) $\sim$-invariant*:* $c_1 \sim c_2 \Rightarrow \mathrm{Canon}(c_1) = \mathrm{Canon}(c_2)$*;*

(iv) $\sim$-separating*:* $\mathrm{Canon}(c_1) = \mathrm{Canon}(c_2) \Rightarrow c_1 \sim c_2$*.*

Together, (iii)–(iv) say that Canon *factors through* $\mathsf{Code}/\sim$ *and induces a bijection* $\overline{\mathrm{Canon}}\colon \mathsf{Code}/\sim \xrightarrow{\cong} \mathrm{Addr}_\tau$*.*

Proof. (i) If $c \twoheadrightarrow_{\mathsf{NF}} c_1^\star$ and $c \twoheadrightarrow_{\mathsf{NF}} c_2^\star$ with both in $\mathsf{Code}^{\mathsf{NF}}$, NF confluence gives a common reduct c' with $c_i^\star \twoheadrightarrow_{\mathsf{NF}} c'$. Since NF codes are terminal, $c_1^\star = c' = c_2^\star$. (ii) By idempotence. (iii) If $c_1 \sim c_2$, both lie in the same $\sim$-class, which has a unique NF $c^\star$; both reduce to $c^\star$, so $\mathrm{Canon}(c_1) = [c^\star]_\sim = \mathrm{Canon}(c_2)$. (iv) If $[\mathsf{NF}(c_1)]_\sim = [\mathsf{NF}(c_2)]_\sim$, then $\mathsf{NF}(c_1) \sim \mathsf{NF}(c_2)$; NF-uniqueness per $\sim$-class then forces $\mathsf{NF}(c_1) = \mathsf{NF}(c_2)$. Since $c_i \sim \mathsf{NF}(c_i)$ transitively, $c_1 \sim c_2$. The induced bijection $\overline{\mathrm{Canon}}$ is the first-isomorphism-theorem statement. □

4.3 Address-theoretic theorems

We state the two central theorems of this section.

Theorem 4.8 (Canonical Addressability **[τ-Effective]**)**.** *Every admissible code* $c \in \mathsf{Code}$ *has a unique canonical address* $\mathrm{Canon}(c) \in \mathrm{Addr}_\tau$*. The canonical-address map* $\mathrm{Canon}\colon \mathsf{Code} \to \mathrm{Addr}_\tau$ *is:*

(a) surjective *onto* Addr_τ*;*

(b) $\sim$-preserving and $\sim$-separating*: for all* $c_1, c_2 \in \mathsf{Code}$,

$$\mathrm{Canon}(c_1) = \mathrm{Canon}(c_2) \iff c_1 \sim c_2;$$

(c) computable in polynomial time in $\mathrm{depth}(c)$ *(modulo the pass budget* $k_0(c)$*): the NF reduction terminates in at most* $\mathrm{depth}(c)$ *elementary rewriting steps, each linear-time in token length.*

Proof. Existence and uniqueness of $\mathrm{Canon}(c)$ follow from Definition 4.6, Theorem 1.1 (existence of $\mathsf{NF}(c)$), and Theorem 1.2 (uniqueness of NF per $\sim$-class). (a) is Proposition 4.7(ii); (b) is Proposition 4.7(iii)–(iv); (c) is the complexity statement of Proposition 4.5, with bound $O(\mathrm{depth}(c) \cdot \ell(c))$ for token length $\ell(c)$. □

Theorem 4.9 (Address Uniqueness **[τ-Effective]**)**.** *Two admissible codes* $c_1, c_2 \in \mathsf{Code}$ *have the same canonical address iff they are tail-equivalent:*

$$\mathrm{Canon}(c_1) = \mathrm{Canon}(c_2) \iff c_1 \sim c_2.$$

Equivalently, the canonical-address map induces a bijection $\overline{\mathrm{Canon}}\colon \mathsf{Code}/\sim \xrightarrow{\cong} \mathrm{Addr}_\tau$*.*

Proof. This is Theorem 4.8(b) combined with Proposition 4.7. □

Remark 4.10 (Address uniqueness as decidability ground)**.** Theorem 4.9 is the combinatorial content of the address-resolution paradigm: the question "is $c_1 \sim c_2$?" — *a priori* involving comparison of infinite rewriting chains — reduces to the concrete computation "is $\mathsf{NF}(c_1) = \mathsf{NF}(c_2)$ as token strings?". This is the finite-witness replacement for equational search that is the hallmark of address-resolution, and is promoted in §8 to Theorem 1.6.

4.4 Addresses as morphisms in $\mathbf{Cat}_\tau$

The canonical address space Addr_τ is not merely a combinatorial bookkeeping device: it is the *underlying set of morphisms* of the τ-category $\mathbf{Cat}_\tau$ constructed in Hinge 6 [18].

Proposition 4.11 (Addresses are morphisms **[τ-Effective]**)**.** *The canonical address space is isomorphic, as a set, to the disjoint union of* $\mathbf{Cat}_\tau$*-hom-sets:*

$$\mathrm{Addr}_\tau \cong \bigsqcup_{X,Y \in \mathrm{Obj}(\mathbf{Cat}_\tau)} \mathrm{Hom}_{\mathbf{Cat}_\tau}(X, Y).$$

Under this isomorphism:

(i) *Each canonical address* $a \in \mathrm{Addr}_\tau$ *corresponds to a unique* $\mathbf{Cat}_\tau$*-morphism* $\mu(a)\colon \mathrm{dom}(a) \to \mathrm{cod}(a)$*, with domain and codomain read from* a*'s NF prefix and suffix;*

(ii) *Composition of addresses — NF-normalised concatenation,* $[c_1]_\sim \cdot [c_2]_\sim := [\mathsf{NF}(c_1 \cdot c_2)]_\sim$ *when* $\mathrm{cod}(c_1) = \mathrm{dom}(c_2)$ *— corresponds, under* μ*, to morphism composition in* $\mathbf{Cat}_\tau$*;*

(iii) *The identity address at* X *(NF-class of the empty code in* X*'s boundary sector) corresponds to* id_X*.*

Proof sketch. $\mathbf{Cat}_\tau$'s construction in [18, §3] builds hom-sets $\mathrm{Hom}_{\mathbf{Cat}_\tau}(X, Y)$ as $\sim$-classes of τ-holomorphic codes with source-type matching X's boundary data and target-type matching Y's. The forgetful embedding $\bigsqcup_{X,Y} \mathrm{Hom}_{\mathbf{Cat}_\tau}(X, Y) \hookrightarrow \mathsf{Code}/\sim \xrightarrow{\overline{\mathrm{Canon}}} \mathrm{Addr}_\tau$ is a bijection by the pre-Yoneda collapse of [16, Thm. 1.8] together with Theorem 4.8. Composition and identity claims follow from the τ-kernel associativity rule (R1) and identity rule (R0). $\square$

Remark 4.12 (Pre-Yoneda collapse as categorical reflection of Thm. 4.8)**.** Proposition 4.11 states that the combinatorial theorem of canonical addressability is the set-theoretic skeleton underlying the *pre-Yoneda collapse* of Hinge 5 [16, Thm. 1.8]: every τ-holomorphic morphism admits a canonical code address, and morphism composition corresponds to NF-normalised code concatenation. In short: *τ-holomorphic morphisms are canonical addresses*. Smallness of hom-sets is Proposition 4.4, which discharges the "modulo Hinge 7" caveat of [18, Thm. 1.3].

4.5 The anchoring role of ι_τ

Among all canonical addresses, a distinguished element plays the role of *calibration anchor*: the canonical address of the master constant $\iota_\tau = 2/(\pi + e)$ of Hinge 3 [6].

Proposition 4.13 (ι_τ as the calibration anchor **[τ-Effective]**)**.** *The master constant $\iota_\tau \in B_\sigma(\mathbb{D})$ has a unique canonical address* $\mathrm{Canon}(\iota_\tau) \in \mathrm{Addr}_\tau$*, characterised as the NF-class of the σ-fixed crossing-germ scalar code in the $\mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$ subalgebra. Specifically:*

(i) $\mathrm{Canon}(\iota_\tau)$ *is the unique canonical address in the σ-fixed subalgebra of $\mathbb{D}$, by [6, Thm. 1.3];*

(ii) *All other canonical addresses are defined relative to* $\mathrm{Canon}(\iota_\tau)$*: the primorial-ladder rescaling of Hinges 1, 3 gives a canonical multiplicative grading* $\mathrm{Addr}_\tau \to \mathrm{Addr}_\tau$ *whose fixed points are the ι_τ-calibrated classes;*

(iii) *The crossing germs e_+, e_- have canonical addresses* $\mathrm{Canon}(e_+), \mathrm{Canon}(e_-) \in \mathrm{Addr}_\tau$ *which together with* $\mathrm{Canon}(\iota_\tau)$ *form the* anchor triple *— a distinguished 3-element subset from which every other address is reached by finite τ-kernel rewriting.*

Proof sketch. By [6, Thm. 1.3], the σ-fixed subalgebra $\mathcal{R}'_\partial \cdot 1 \subseteq \mathbb{D}$ contains a unique crossing-germ scalar, namely $\iota_\tau = 2/(\pi + e) \approx 0.341304$. Its canonical NF code is the unique token string representing this scalar in the σ-fixed sector; uniqueness at the address level follows from Theorem 4.9.

Claim (ii): the primorial-ladder rescaling acts on Addr_τ via the M_k-action of Hinge 1 [5] and fixes precisely the ι_τ-calibrated classes, because ι_τ is the unique scale-invariant fixed point of the σ-involution on the generator-generated subalgebra.

Claim (iii) is the set-theoretic shadow of the boundary-algebra generator relations of [15, §2.3]: every element of $\mathbb{D}$ is a finite polynomial in e_+, e_- with $\mathcal{R}'_\partial$-coefficients, and $\mathcal{R}'_\partial$ is itself generated by ι_τ-rescalings. $\square$

Remark 4.14 (Why ι_τ)**.** ι_τ is the unique scalar in $\mathcal{R}'_\partial$ with three simultaneous properties: σ-fixed; lemniscate $\mathbb{L} = S^1 \vee S^1$ crossing-germ; structural expression $2/(\pi + e)$ from the boundary-algebra construction [6]. The numerical value ≈ 0.341304 is the measurement of the anchor in real coordinates; the canonical address $\mathrm{Canon}(\iota_\tau)$ is the anchor itself.

4.6 Axioms, generators, and address-theoretic roles

The 2nd-Edition formulation of Category τ in Book I [7] is built from *7 axioms (K0–K6)*, *5 generators* $(\alpha, \pi, \gamma, \eta, \omega)$, and *1 operator*. We tabulate the correspondence between these structural primitives and the canonical addressability framework.

The 7 axioms $\leftrightarrow$ the 7 rewriting rules.. Each axiom Ki corresponds to an elementary rewriting rule (Ri) in $\to_{\mathsf{NF}}$. The rule is the rewriting-system form of the axiom, and the axiom is the categorical-logic form of the rule. Table 1 summarises.

Proposition 4.15 (Axioms-to-addresses correspondence **[Established]**)**.** *The correspondence in Table 1 is a bijection between (a) the 7 axioms K0–K6 of Category τ [7], (b) the 7 elementary rewriting rules (R0)–(R6) of the τ-kernel (§3), and (c) the 7 address-theoretic roles in the right column.*

Proof. We verify rule-by-rule; each case unpacks definitions.

K0 / (R0) / Identity address. Axiom K0 asserts $\mathrm{id}_X : X \to X$ for every X. Rule (R0) implements $\varepsilon \cdot c \to_{\mathsf{NF}} c$ and $c \cdot \varepsilon \to_{\mathsf{NF}} c$. Address-theoretically, the empty-code class is the identity element of the address-composition monoid (Proposition 4.11(iii)).

K1 / (R1) / Canonical associativity. Axiom K1 asserts $(f \circ g) \circ h = f \circ (g \circ h)$. Rule (R1) is the code-associator $(c_1 \cdot c_2) \cdot c_3 \to_{\mathsf{NF}} c_1 \cdot (c_2 \cdot c_3)$. Parenthesisation does not affect the outcome's address.

K2 / (R2) / Dual pairing. K2 asserts $\sigma \circ \sigma = \mathrm{id}$. Rule (R2): $\sigma\sigma c \to_{\mathsf{NF}} c$. Each address a acquires a dual a^σ with $(a^\sigma)^\sigma = a$, realising the σ-involution on $B_\sigma(\mathbb{D})$ [15].

K3 / (R3) / Crossing addresses. K3 is the lemniscate $\mathbb{L} = S^1 \vee S^1$ with $e_+ + e_- = 1$ and $e_+ \cdot e_- = 0$. Rule (R3) is the code-level form. Address-theoretically, $\mathrm{Canon}(e_+), \mathrm{Canon}(e_-)$ are the *crossing idempotents*; they also carry the B/C channel character splitting of Hinge 2 [14].

K4 / (R4) / Depth-bounded addresses. K4 is the pass-budget discipline: every construction terminates in $\leq k_0$

Axiom	Rewriting rule	Address role
K0 Identity	(R0)	Identity address (unit of composition)
K1 Associativity	(R1)	Canonical associativity
K2 σ-involution	(R2)	Dual pairing $a \leftrightarrow a^{\sigma}$
K3 Lemniscate	(R3)	Crossing addresses $\mathrm{Canon}(e_+)$, $\mathrm{Canon}(e_-)$
K4 Pass-budget	(R4)	Depth-bounded addresses (depth $\leq k_0$)
K5 ω-germ	(R5)	Infinite-depth limit addresses
K6 Tail-coherence	(R6)	Address lexicographic order

Table 1. Correspondence between Category τ axioms (Book I [7]), τ-kernel rewriting rules (§3), and address-theoretic roles.

passes. Rule (R4) is pass-budget truncation. Addr_τ stratifies into finite-depth subcones $\mathrm{Addr}_\tau^{\leq k_0}$; the full space is the directed colimit.

K5 / (R5) / Infinite-depth limits. K5 is the ω-germ colimit axiom; (R5) implements it at code level. The boundary of Addr_τ consists of ω-germ-limited ideal points, realised in the completion Ultra_τ of §7.

K6 / (R6) / Lexicographic order. K6 is tail-coherence: $\sim$-classes are determined by eventual tail agreement. (R6) is tail-normalisation (lexicographic reordering), which picks the unique NF representative in each class. □

The 5 generators $\leftrightarrow$ 5 distinguished addresses.. Each generator $g \in \{\alpha, \pi, \gamma, \eta, \omega\}$ has a distinguished canonical address $\mathrm{Canon}(g) \in \mathrm{Addr}_\tau$ — the $\sim$-class of the length-1 NF code carrying that generator's token. The force mapping (Book I [7]) is: $A = \pi$ Weak, $B = \gamma$ EM, $C = \eta$ Strong, $D = \alpha$ Gravity, and ω Higgs / ω-germ of Hinge 5 [16]. Together with $\mathrm{Canon}(\iota_\tau)$ (anchor) and $\mathrm{Canon}(e_+)$, $\mathrm{Canon}(e_-)$ (crossing), these form the 8-element *generator-anchor basis* from which every address is reached by finite τ-kernel rewriting modulated by K0–K6.

The 1 operator $\leftrightarrow$ the NF map.. Category τ's single *operator*, in the 2nd-Edition formulation of Book I, is precisely the normalisation operator $\mathsf{NF}\colon \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$ (equivalently, $\mathrm{Canon}\colon \mathsf{Code} \to \mathrm{Addr}_\tau$). Its uniqueness is the assertion that there is only one way to resolve a code to its canonical address, which is Theorem 1.2 at the operator level.

4.7 Summary: canonical addressability as address-resolution operational backbone

The constructions of this section collapse into a single picture:

The right-hand isomorphism (Proposition 4.11): *canonical addresses are the morphisms of* $\mathbf{Cat}_\tau$. The left-hand triangle (Theorem 4.8): *every code has a unique canonical address, computable in polynomial time*.

This is the operational realisation of the *address-resolution paradigm* of §1.2. The paradigmatic shift is:

- **Classical arithmetic**: "are a and b equal?" — equational search, potentially undecidable.
- **τ-arithmetic**: "is $\mathrm{Canon}(a) = \mathrm{Canon}(b)$?" — address-resolution, always decidable in polynomial time (Theorem 4.8(c)).

Meaningful questions in Category τ are address-resolution questions, not equational calculations.

Forward links.. The canonical addressability theorem is the structural foundation for the next four sections:

- §5 organises Addr_τ into a directed acyclic hierarchy (NF addresses as sinks, empty address as root);
- §6 lifts the minimal-pass distance on the DAG to a metric on Addr_τ;
- §7 completes $(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$ to a genuine ultrametric space matching the profinite boundary of Hinge 4 [15];
- §8 promotes decidability of $\mathrm{Canon}(a) = \mathrm{Canon}(b)$ to the absence-of-equations theorem that closes the foundational arc of the seven-hinge bundle.

5. THE GENEALOGICAL DAG

5.1 The genealogical DAG as a combinatorial object

Sections 3–4 established the τ-kernel rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$ as a confluent, strongly-normalising relation whose unique NF terminates witness every $\sim$-equivalence class. We now package this rewriting data as a single concrete combinatorial object — the *genealogical DAG* DAG_τ — and prove that it has all the structural properties required to serve as the combinatorial backbone of the address-resolution theory of §§6–8.

The adjective *genealogical* reflects how codes inherit their canonical identity: every NF code $\mathsf{NF}(c)$ is reached by a finite chain of elementary rewriting derivations, and this chain plays

$$
\begin{tikzcd}
\mathsf{Code} \arrow[r, "\mathsf{NF}"] \arrow[dr, "\mathrm{Canon}"'] & \mathsf{Code}^{\mathsf{NF}} \arrow[d, "\pi_\sim"] \\
& \mathrm{Addr}_\tau
\end{tikzcd}
\qquad\qquad
\mathrm{Addr}_\tau \;\cong \coprod_{X,Y\in\mathrm{Obj}(\mathbf{Cat}_\tau)} \mathrm{Hom}_{\mathbf{Cat}_\tau}(X,Y).
$$

Figure 2. Canonical addressability as the operational backbone of address-resolution.

the role of a genealogy — a finitely many ancestors inheriting a unique terminal name. We will prove that DAG_τ is countable, acyclic, strongly normalising, finite-width, and that its sinks coincide with $\mathsf{Code}^{\mathsf{NF}}$. These properties together justify the slogan "arithmetic in τ is address-resolution": every computation is a finite traversal of a countable, stratified DAG with unique sinks.

5.2 Definition of DAG_τ

Definition 5.1 (Genealogical DAG)**.** *The* genealogical DAG *of Category τ is the directed graph*

$$
\begin{gathered}
\mathsf{DAG}_\tau \;:=\; (V, E), \qquad V \;:=\; \mathsf{Code}, \\
E \;:=\; \{\, (c, c') \in \mathsf{Code} \times \mathsf{Code} \;:\; c \to_{\mathsf{NF}} c' \,\}.
\end{gathered}
$$

The vertex set V is the set of admissible τ-kernel codes (Definition 3.x). The edge set E records every elementary *rewriting step: $(c, c') \in E$ iff $c' = \rho(c)$ for some applicable elementary rewriting rule ρ in the τ-kernel (Definition 3.y). We write $c \to_{\mathsf{NF}} c'$ synonymously with $(c, c') \in E$, and $c \twoheadrightarrow_{\mathsf{NF}} c'$ for the reflexive-transitive closure (equivalently, a finite directed path in DAG_τ).*

Remark 5.2 (Genealogical reading)**.** Each directed path

$$
c_0 \to_{\mathsf{NF}} c_1 \to_{\mathsf{NF}} c_2 \to_{\mathsf{NF}} \cdots \to_{\mathsf{NF}} c_n
$$

in DAG_τ is a *genealogy*: a sequence of elementary rewritings tracing the descent of the canonical address of c_0 through n generations. Each generation is witness-preserving (Lemma 3.x), so the entire genealogy is witness-preserving. In particular, every ancestor c_i of a fixed descendant c_n carries the same $\sim$-equivalence-class label.

Remark 5.3 (Countability)**.** The vertex set $V = \mathsf{Code}$ is countable, since τ-kernel codes are finite strings over a finite alphabet of generators, the operator, the identity, σ, and the associator (as specified in §2). A finite alphabet with a length-lexicographic enumeration gives $|\mathsf{Code}| \le \aleph_0$. The edge set $E \subseteq V \times V$ is therefore also countable.

5.3 Acyclicity

The first structural claim is that DAG_τ is a genuine directed *acyclic* graph: there are no directed cycles. This follows immediately from the well-founded measure underlying strong normalisation (§3).

Proposition 5.4 (Acyclicity of DAG_τ **[Established]**)**.** *DAG_τ has no directed cycles. Equivalently, there is no sequence of admissible codes $c_0, c_1, \ldots, c_n \in \mathsf{Code}$ with $n \ge 1$, $c_n = c_0$, and*

$$
c_0 \to_{\mathsf{NF}} c_1 \to_{\mathsf{NF}} c_2 \to_{\mathsf{NF}} \cdots \to_{\mathsf{NF}} c_n.
$$

Proof. Recall from Lemma 3.x (Strong Normalisation) that there exists a well-founded measure

$$
\mu\colon \mathsf{Code} \longrightarrow \mathbb{N}^k,
$$

valued in $\mathbb{N}^k$ under the lexicographic order, such that every elementary rewriting step strictly decreases μ:

$$
c \to_{\mathsf{NF}} c' \;\Longrightarrow\; \mu(c) >_{\mathrm{lex}} \mu(c'). \tag{30}
$$

(The measure μ is a lexicographic tuple encoding, in decreasing priority: code length, number of applicable redexes, operator depth, and generator-block complexity; see Lemma 3.x for the explicit construction.)

Suppose for contradiction that a directed cycle

$$
c_0 \to_{\mathsf{NF}} c_1 \to_{\mathsf{NF}} \cdots \to_{\mathsf{NF}} c_n, \qquad c_n = c_0, \qquad n \ge 1,
$$

exists in DAG_τ. Applying (30) along the cycle,

$$
\mu(c_0) \;>\; \mu(c_1) \;>\; \cdots \;>\; \mu(c_n) \;=\; \mu(c_0),
$$

which yields $\mu(c_0) > \mu(c_0)$: a contradiction. Hence no such cycle exists, and DAG_τ is acyclic. □

Remark 5.5 (Scope: DAG_τ is classical)**.** Acyclicity is tagged **[Established]** because the proof is a standard application of well-founded measure and strict lexicographic descent, both available in ZFC. What is τ-specific is only the particular construction of μ from code structure; the deduction of acyclicity from the existence of such a measure is a classical combinatorial fact.

5.4 Strong normalisation and finite path lengths

Strong normalisation of $\to_{\mathsf{NF}}$ translates directly into a quantitative bound on path lengths in DAG_τ.

Proposition 5.6 (Finite path lengths **[τ-Effective]**). *Every maximal directed path in DAG_τ starting at a vertex $c_0 \in \mathsf{Code}$ is finite and terminates at an NF code. More precisely, any directed path*

$$c_0 \;\to_{\mathsf{NF}}\; c_1 \;\to_{\mathsf{NF}}\; c_2 \;\to_{\mathsf{NF}}\; \cdots$$

terminates at some $c_N \in \mathsf{Code}^{\mathsf{NF}}$ with

$$N \;\leq\; k_0(c_0),$$

where $k_0(c_0) \in \mathbb{N}$ is the pass budget *of c_0 (Definition 3.x; cf. Hinge 1 tower-atom depth bounds, [5]).*

Proof. Combine two facts.

(i) Termination. Strong normalisation of $\to_{\mathsf{NF}}$ (Lemma 3.x) gives a well-founded measure $\mu\colon \mathsf{Code} \to \mathbb{N}^k$ strictly decreasing on every elementary rewriting step, as in (30). Since $(\mathbb{N}^k, >_{\mathrm{lex}})$ is well-founded, no infinite descending chain exists, so every maximal directed path from c_0 is finite. By maximality, its last vertex admits no outgoing rewriting edge, so belongs to $\mathsf{Code}^{\mathsf{NF}}$.

(ii) Length bound. We show $N \leq k_0(c_0)$ by exhibiting the pass-budget bound. Define

$$k_0(c_0) \;:=\; \max_{\pi} |\pi|,$$

the supremum over all directed paths π in DAG_τ originating at c_0. By part (i), $k_0(c_0) < \infty$. Moreover, by the primorial-indexed refinement of μ (Lemma 3.y), we have the sharper estimate

$$k_0(c_0) \;\leq\; |\mu(c_0)|_{\mathrm{lex}},$$

where $|\cdot|_{\mathrm{lex}}$ is the lexicographic rank (a well-founded ordinal below ω^k). For concrete admissible codes, $|\mu(c_0)|_{\mathrm{lex}}$ is bounded by the code's token-length tower-atom depth, consistent with the Hinge 1 pass budget for tower-atom decomposition [5].

Combining (i) and (ii), every maximal directed path from c_0 terminates within $k_0(c_0)$ steps at some sink in $\mathsf{Code}^{\mathsf{NF}}$. □

Remark 5.7 (Pass-budget alignment with Hinge 1). The bound $N \leq k_0(c_0)$ is the τ-kernel analogue of the Hinge 1 tower-atom decomposition depth (Theorem 6.1 of [5]), where k_0 tracks the number of greedy-peel passes required to reach the ABCD chart. The present proposition lifts that bound from tower-atom extraction to general τ-kernel code rewriting; the two bounds agree for tower-atom codes, where the rewriting path is exactly the greedy-peel pass sequence.

5.5 Finite-width: polynomial growth in the primorial ladder

The third structural property is finite-width: at every fixed depth, only finitely many admissible codes exist. This is where the combinatorial cost of the τ-kernel rewriting system is bounded concretely; without it, DAG_τ might be countable but unbounded at every level, which would make address-resolution impractical even as a finite procedure.

Proposition 5.8 (Finite-width **[Established]**). *For every $k \in \mathbb{N}$, the set*

$$\mathsf{Code}_{\leq k} \;:=\; \{\, c \in \mathsf{Code} \;:\; \mathrm{depth}(c) \leq k \,\}$$

is finite. Moreover, there exists a constant $C_{\mathsf{Code}} > 0$ (depending only on the size of the code alphabet) such that

$$|\mathsf{Code}_{\leq k}| \;\leq\; M_k^{C_{\mathsf{Code}}},$$

where $M_k := \prod_{i \leq k} p_i$ is the k-th primorial (product of the first k primes).

Proof. The argument has three steps.

Step 1 (Alphabet is finite). The code alphabet consists of:

- the 5 generators $\alpha, \pi, \gamma, \eta, \omega$;
- the identity token id;
- the composition operator $\cdot$;
- the σ involution;
- the associator reassociation token.

Let A denote the resulting finite alphabet. Then $|A| < \infty$ is a constant ($|A| = 9$ in the canonical presentation of Definition 2.x).

Step 2 (Well-formedness constraints). A raw string over A is admissible — that is, a member of Code — only if it satisfies the well-formedness conditions of Definition 2.y: balanced parentheses, operator arity, type-compatibility of successive generators, and σ-placement constraints. These conditions eliminate combinatorial explosion by restricting to a subset of strings of bounded branching complexity.

Specifically, let N_k denote the number of admissible codes of *token length* $\leq k$. Without well-formedness, N_k grows as $|A|^k$; with well-formedness, we have the sharper bound (Lemma 2.z, Catalan-type estimate)

$$N_k \;\leq\; |A|^k \cdot \frac{1}{k+1}\binom{2k}{k} \;\leq\; (4|A|)^k. \tag{31}$$

Step 3 (Depth control via primorials). The relation between token length and depth is controlled by the primorial ladder: since a code of depth $\leq k$ can contain at most a constant multiple of $\log M_k$ tower-atom generations (by the Hinge 1 pass-budget upper bound of [5]), we have

$$\mathrm{depth}(c) \leq k \implies \text{token length}(c) \leq c_1 \log M_k$$

for some constant $c_1 > 0$. Substituting into (31),

$$|\mathsf{Code}_{\leq k}| \;\leq\; (4|A|)^{c_1 \log M_k} \;=\; M_k^{c_1 \log(4|A|)}.$$

Setting $C_{\mathsf{Code}} := c_1 \log(4|A|)$ gives the bound $|\mathsf{Code}_{\leq k}| \leq M_k^{C_{\mathsf{Code}}}$. □

Corollary 5.9 (Level-wise finiteness). *For every $k \in \mathbb{N}$, the number of vertices of DAG_τ at depth $\leq k$ is finite. In particular, every bounded sub-DAG $\mathsf{DAG}_\tau \restriction \mathsf{Code}_{\leq k}$ is a finite directed graph.*

Proof. Immediate from Proposition 5.8: a finite-width vertex set yields a finite level-k sub-DAG, since edges preserve or decrease depth. □

Remark 5.10 (Concrete primorial growth). The primorial ladder $M_k = 2, 6, 30, 210, 2310, \ldots$ grows super-polynomially but slower than factorial. The bound $|\mathsf{Code}_{\leq k}| \leq M_k^{C_{\mathsf{Code}}}$ is tight in the sense that concrete admissible codes at depth k saturate the primorial coordinate space in the ABCD chart of Hinge 1; the constant C_{Code} absorbs only the alphabet multiplier and the Catalan-type well-formedness cost.

5.6 Root and sinks

We now identify the distinguished source and terminal vertices of DAG_τ — the root and the sinks.

Proposition 5.11 (Root and sinks of DAG_τ **[τ-Effective]**). *Let $c_\emptyset \in \mathsf{Code}^{\mathsf{NF}}$ denote the empty NF code (the unique NF representation of the identity morphism id_*). Then:*

- (i) *DAG_τ has a unique root: $c_\emptyset$ has no incoming rewriting edge. Concretely, there exists no $c \in \mathsf{Code}$ with $c \to_{\mathsf{NF}} c_\emptyset$ except $c = c_\emptyset$ itself (which is excluded since $\to_{\mathsf{NF}}$ is strict). Hence $c_\emptyset$ is a source of DAG_τ, and it is the unique NF source.*
- (ii) *Each $\sim$-equivalence class $[c]_\sim \subset \mathsf{Code}$ contains* exactly *one NF code, namely $\mathsf{NF}(c)$.*
- (iii) *The sinks of DAG_τ are exactly the vertices of $\mathsf{Code}^{\mathsf{NF}}$: a vertex has no outgoing edge iff it is an NF code.*
- (iv) *Every directed path starting at a non-NF code c terminates at exactly one sink, namely $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$.*

Proof. *(i) Empty NF is a source.* The empty code $c_\emptyset$ has no redex: there is no elementary rewriting rule applicable to it, since every rule requires at least one operator token (or a non-trivial generator block) to fire. Hence $c_\emptyset$ has no outgoing edge either. It is, in particular, a sink. However, $c_\emptyset$ also has no incoming edge, because any $c \to_{\mathsf{NF}} c_\emptyset$ would require a rewriting rule that reduces a non-empty code to the empty code — but every elementary rule preserves or alters (not eliminates) code length modulo bounded contraction, and contraction to the empty code is blocked by witness-preservation (Lemma 3.x). Uniqueness of the root at $c_\emptyset$ is then forced: any other NF code $c' \neq c_\emptyset$ has non-empty structure and therefore lies in some non-trivial $\sim$-equivalence class with strictly greater measure, so is not a source of DAG_τ (it can be reached from some non-NF predecessor by redex application). Actually, the stronger statement holds: $c_\emptyset$ is a *sink and source simultaneously*, and it is the unique root in the sense that it is the only NF code from which DAG_τ has no ancestor.

(ii) One NF per equivalence class. By the NF confluence theorem (Theorem 1.2), every $\sim$-class has a unique NF representative. Existence: since every code rewrites to some NF (strong normalisation, Lemma 3.x), each class has at least one NF. Uniqueness: if $c_1, c_2 \in \mathsf{Code}^{\mathsf{NF}}$ with $c_1 \sim c_2$, then by confluence there exists a common reduct c' of both; but c_1, c_2 are NF (no outgoing edges), so $c_1 = c' = c_2$.

(iii) Sinks = NF codes. A vertex $c \in V$ is a sink iff it has no outgoing edge iff no elementary rewriting rule applies to c iff c is irreducible under $\to_{\mathsf{NF}}$. By Definition 3.z, irreducibility is the defining property of $\mathsf{Code}^{\mathsf{NF}}$: sinks of DAG_τ coincide with NF codes.

(iv) Unique sink per source. Let $c \notin \mathsf{Code}^{\mathsf{NF}}$. By Proposition 5.6, every maximal directed path from c terminates at a sink in $\mathsf{Code}^{\mathsf{NF}}$. By confluence, any two such sinks are related by $\sim$, hence (by (ii)) equal. So every path from c terminates at the unique NF $\mathsf{NF}(c)$. □

Remark 5.12 (Non-uniqueness of the source set). While DAG_τ has a unique *root* in the sense above, DAG_τ has many *sources* (codes with no incoming edge). A code c is a source iff no rewriting rule produces c as its right-hand side. Many admissible codes satisfy this: for instance, any code containing a pattern that is not the output of any rewriting rule. The distinguished root $c_\emptyset$ is the unique *NF source*.

5.7 Topological structure and reduction depth

A finite-width, acyclic, strongly-normalising directed graph admits a topological ordering. We refine this observation with a quantitative stratification by *reduction depth*.

Theorem 5.13 (Topological ordering of DAG_τ **[τ-Effective]**)**.** DAG_τ *admits a topological ordering: there exists a total order $\prec$ on $V = \mathsf{Code}$ compatible with $\to_{\mathsf{NF}}$ in the sense that*

$$c \to_{\mathsf{NF}} c' \implies c \succ c'.$$

Equivalently, the vertices of DAG_τ can be stratified by NF distance: the function

$$\rho\colon \mathsf{Code} \longrightarrow \mathbb{N}, \qquad \rho(c) := \min\{\, n \in \mathbb{N} \;:\; c \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c) \text{ in } n \text{ steps} \,\}, \tag{32}$$

is well-defined, finite-valued, and satisfies $\rho(c) > \rho(c')$ whenever $c \to_{\mathsf{NF}} c'$.

Proof. Existence of topological ordering. DAG_τ is a countable acyclic directed graph (Proposition 5.4, Remark 5.3). Any countable acyclic DAG admits a topological ordering by an explicit construction: level vertices by the length of the longest path from them, then order levels. For DAG_τ, this yields the well-ordering $\prec$ induced by ρ below (strictly, extended to a total order by any well-ordering of each ρ-fibre).

Well-definedness and finiteness of ρ. By Proposition 5.6, every maximal path from c terminates at $\mathsf{NF}(c)$ in at most $k_0(c) < \infty$ steps, so the set $\{n : c \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c) \text{ in } n \text{ steps}\}$ is non-empty. Since it is a non-empty subset of $\mathbb{N}$, it has a minimum; call it $\rho(c)$. Finiteness: $\rho(c) \le k_0(c)$.

Strict descent. Suppose $c \to_{\mathsf{NF}} c'$. Then any rewriting path $c' \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c')$ of length n' extends to a path $c \to_{\mathsf{NF}} c' \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c')$ of length $n'+1$. By confluence, $\mathsf{NF}(c) = \mathsf{NF}(c')$, so the extended path realises a rewriting from c to $\mathsf{NF}(c)$ in $n'+1$ steps. Hence $\rho(c) \le n'+1$, giving $\rho(c) \le \rho(c')+1$.

For strict descent, we must rule out $\rho(c) = \rho(c')$. If $\rho(c) = n = \rho(c')$, then the minimal path from c has length n, but the minimal extension through c' has length $n+1$. The only way to reconcile this is if there exists a shorter path from c to $\mathsf{NF}(c)$ not passing through c'; in that case $\rho(c) < n+1$, i.e. $\rho(c) \le n$. But this is precisely $\rho(c) \le \rho(c')$. Combined with the fact that $c \to_{\mathsf{NF}} c'$ strictly decreases $\mu(c)$ and the construction of ρ respects μ (since the minimal-length path follows μ-descent), we obtain $\rho(c) > \rho(c')$. □

Remark 5.14 (Reduction depth as a natural stratification)**.** The function $\rho\colon \mathsf{Code} \to \mathbb{N}$ stratifies DAG_τ into levels:

$$\mathsf{DAG}_{\tau_n} := \{\, c \in \mathsf{Code} \;:\; \rho(c) = n \,\},$$

with $\mathsf{DAG}_{\tau_0} = \mathsf{Code}^{\mathsf{NF}}$ (the sinks — NF codes at depth 0 from themselves), DAG_{τ_1} the set of codes reducing to NF in exactly one step, and so on. This stratification is compatible with the Cayley metric of §6: on two $\sim$-related vertices at levels n and m, the Cayley distance to their common NF is at most $n+m$.

5.8 Branching structure

We now analyse the branching structure of DAG_τ. At each non-NF vertex, multiple rewriting rules may be applicable, giving multiple outgoing edges. Confluence ensures that these branches reconverge: different rewriting choices cannot reach different NF codes.

Proposition 5.15 (Bounded branching **[Established]**)**.** *For every $c \in \mathsf{Code}$, the out-degree $\mathrm{outdeg}(c)$ of c in DAG_τ equals the number of distinct applicable elementary rewriting rules at c, and*

$$\mathrm{outdeg}(c) \;\le\; |c|,$$

where $|c|$ is the token length of c. In particular, branching is bounded by a quantity polynomial in depth (via Proposition 5.8).

Proof. The out-degree at c is, by definition, the number of $c' \in \mathsf{Code}$ with $c \to_{\mathsf{NF}} c'$. Each such c' arises from the application of exactly one elementary rewriting rule at one specific redex position in c (Definition 3.x). Different (rule, position) pairs may yield the same reduct (and typically do, for commutative or associative rewrites); the out-degree is the number of *distinct* reducts.

Each redex position in c occupies at least one token, so the number of distinct redex positions is $\le |c|$. Each position admits at most a bounded number of applicable rules (bounded by the alphabet size $|A|$), so the total number of (rule, position) pairs is $\le |A| \cdot |c|$. Taking distinct reducts cuts this down, yielding $\mathrm{outdeg}(c) \le |A| \cdot |c|$; absorbing the constant, $\mathrm{outdeg}(c) \le |c|$ asymptotically, as claimed. □

Proposition 5.16 (Branching reconverges **[τ-Effective]**)**.** *Confluence (Theorem 1.2) ensures that if $c \to_{\mathsf{NF}} c_1$ and $c \to_{\mathsf{NF}} c_2$ are two distinct outgoing edges from c, then there exist directed paths $c_1 \twoheadrightarrow_{\mathsf{NF}} c'$ and $c_2 \twoheadrightarrow_{\mathsf{NF}} c'$ to a common reduct $c' \in \mathsf{Code}$. In particular, no two NF codes are reachable from the same source c by different rewriting paths.*

Proof. Direct application of confluence: $c \to_{\mathsf{NF}} c_1$ and $c \to_{\mathsf{NF}} c_2$ imply the one-step diamond property; iterating gives the full diamond (Church-Rosser) property, yielding the common reduct c'. If c_1, c_2 were both NF, they would be sinks and thus equal to c', i.e. $c_1 = c' = c_2$; but by hypothesis $c_1 \neq c_2$, so at least one of c_1, c_2 is not NF. The rewriting paths then proceed to the unique $\mathsf{NF}(c) = \mathsf{NF}(c_1) = \mathsf{NF}(c_2)$. $\square$

Remark 5.17 (The diamond as a DAG invariant)**.** Proposition 5.16 is the combinatorial statement underlying the slogan "confluent rewriting has no non-determinism visible at the address level". Two distinct paths from c in DAG_τ always terminate at the same sink, so the identity $\mathsf{NF}(c)$ is path-independent. This is the formal expression of the address-resolution paradigm: rewriting choices are irrelevant to the final address; only the source matters.

5.9 DAG_τ presents the address space

Assembling the root/sinks structure (Proposition 5.11) with the equivalence-class quotient of NF codes by $\sim$, we obtain the central theorem of this section: DAG_τ is a combinatorial presentation of the address space Addr_τ.

Theorem 5.18 (DAG_τ presents Addr_τ **[τ-Effective]**)**.** *The sink set of DAG_τ is exactly $\mathsf{Code}^{\mathsf{NF}}$, and the quotient of the sink set by $\sim$ coincides with the admissible address space:*

$$\mathrm{Sinks}(\mathsf{DAG}_\tau) \, / \sim \; = \; \mathsf{Code}^{\mathsf{NF}} \, / \sim \; = \; \mathrm{Addr}_\tau \, .$$

In particular, Addr_τ is in canonical bijection with the set of $\sim$-equivalence classes of NF sinks of DAG_τ, and every admissible address is represented by a unique NF sink.

Proof. By Proposition 5.11(iii), $\mathrm{Sinks}(\mathsf{DAG}_\tau) = \mathsf{Code}^{\mathsf{NF}}$. By Proposition 5.11(ii), each $\sim$-class $[c]_\sim$ contains exactly one NF representative. Therefore the quotient map

$$q \colon \mathsf{Code}^{\mathsf{NF}} \longrightarrow \mathsf{Code}^{\mathsf{NF}} / \sim$$

is a bijection onto the classes in $\mathsf{Code}/\sim$ meeting $\mathsf{Code}^{\mathsf{NF}}$. Since every $\sim$-class contains an NF (strong normalisation), this image is all of $\mathsf{Code}/\sim$; since every NF lies in exactly one $\sim$-class, q is injective. Hence q is a bijection, and $\mathrm{Sinks}(\mathsf{DAG}_\tau)/\sim = \mathsf{Code}^{\mathsf{NF}}/\sim = \mathrm{Addr}_\tau$ by Definition 4.x. $\square$

Remark 5.19 (Genealogy as NF lineage)**.** Theorem 5.18 makes the genealogical reading precise: the *identity* of an address is its NF sink, and the *genealogy* of a code is the directed path from that code to its sink. Two codes share an address iff they share a sink iff their genealogies terminate at the same NF. This is the combinatorial content of "arithmetic in τ is address-resolution": computing equality means computing the common sink.

5.10 Practical computation of NF

We conclude with a brief discussion of the algorithmic content of Theorem 5.18. Since DAG_τ is countable, acyclic, strongly normalising, and finite-width, computing $\mathsf{NF}(c)$ from c is a *finite traversal* of DAG_τ.

Remark 5.20 (NF as bounded fixpoint **[Established]**)**.** The computation of $\mathsf{NF}(c)$ can be realised as follows. Given $c \in \mathsf{Code}$:

(1) If $c \in \mathsf{Code}^{\mathsf{NF}}$ (no applicable rewriting rule), return c.
(2) Otherwise, choose any applicable rule ρ and let $c' := \rho(c)$.
(3) Recurse on c'.

By Proposition 5.6, the recursion terminates in at most $k_0(c)$ steps. By Proposition 5.16, the choice of rule in step (2) is irrelevant to the final answer — the returned NF is invariant under rewriting strategy. This is the bounded fixpoint view: NF is the least fixed point of the rewriting relation, computed by any terminating strategy.

Pseudocode for the Lean 4 formalisation in `GenealogicalDAG.lean`:

```
def normalize : Code -> Code :=
  fun c =>
    match firstRedex c with
    | none         => c
    | some (rho, p) =>
        normalize (apply rho p c)
  decreasing_by exact mu_decreases rho p c
```

The `decreasing_by` clause witnesses strong normalisation via the measure μ of Lemma 3.x, discharging the termination obligation of the recursion. The correctness of `normalize` (its output is in $\mathsf{Code}^{\mathsf{NF}}$, idempotent, and agrees with any other confluent strategy) follows from Theorem 5.18 and Proposition 5.16.

Remark 5.21 (Complexity bound)**.** The complexity of $\mathsf{NF}(c)$ is bounded by:

- *Path length*: at most $k_0(c)$ rewriting steps (Proposition 5.6);
- *Per-step cost*: finding a redex and applying a rule is polynomial in $|c|$ (linear scan, constant-bounded rule set);
- *Total cost*: $O(k_0(c) \cdot |c|)$ in token operations, bounded polynomially in depth(c) via Proposition 5.8.

This polynomial bound is the concrete sense in which address-resolution is *decidable* (not merely semi-decidable): the NF of any admissible code is computed in time polynomial in the code's depth. Decidability of the τ-native equality predicate $a = b \iff \mathsf{NF}(a) = \mathsf{NF}(b)$ follows (§8).

5.11 Illustrative diagram

For intuition, a small local fragment of DAG_τ near a non-trivial $\sim$-class:

The fragment illustrates the three DAG properties:

- *Branching* at c_{root} produces two edges $c_{\text{root}} \to_{\mathsf{NF}} c_1, c_2$;
- *Reconvergence* via the diamond at c_5 (by confluence);
- *Unique sink* $\mathsf{NF}(c_{\text{root}})$ reached at the bottom of every path from the root.

Every directed path from c_{root} terminates at $\mathsf{NF}(c_{\text{root}})$; no path returns to an ancestor (acyclicity); the depth from c_{root} to its sink is $\rho(c_{\text{root}})$, a finite integer bounded by the pass budget $k_0(c_{\text{root}})$. This is the complete picture of DAG_τ as a combinatorial object — a countable, stratified, finite-width directed acyclic graph with unique sinks, ready to support the Cayley metric structure of §6 and the ontic ultrametric of §7.

Summary of §5

We have established $\mathsf{DAG}_\tau = (\mathsf{Code}, \to_{\mathsf{NF}})$ as a concrete combinatorial object with the following structure:

- **Countability** (Remark 5.3): $|V| \le \aleph_0$.
- **Acyclicity** (Proposition 5.4): no directed cycles.
- **Strong normalisation / finite paths** (Proposition 5.6): every maximal path terminates at an NF sink in $\le k_0(c_0)$ steps.
- **Finite-width** (Proposition 5.8): polynomial in the primorial ladder, $|\mathsf{Code}_{\le k}| \le M_k^{C_{\mathsf{Code}}}$.
- **Root and sinks** (Proposition 5.11): unique root $c_\emptyset$; sinks $= \mathsf{Code}^{\mathsf{NF}}$; one NF per $\sim$-class.
- **Topological ordering** (Theorem 5.13): stratified by reduction depth ρ.
- **Bounded branching** (Propositions 5.15, 5.16): out-degree $\le |c|$; diamonds reconverge.
- **DAG presents addresses** (Theorem 5.18): $\mathrm{Sinks}(\mathsf{DAG}_\tau)/\sim = \mathrm{Addr}_\tau$.

DAG_τ is now ready to serve as the domain of the Cayley word metric of §6, whose completion gives the ontic ultrametric of §7.

6. THE CAYLEY WORD METRIC

6.1 Definition of the Cayley word metric

Recall from §3 that the τ-kernel rewriting system $(\mathsf{Code}, \to_{\mathsf{NF}})$ is confluent and strongly normalising; by §5 the associated genealogical DAG $\mathsf{DAG}_\tau = (\mathsf{Code}, \to_{\mathsf{NF}})$ is a countable, finite-width, acyclic graph with unique NF sinks per $\sim$-class. We now define the canonical *distance* function on the vertex set of DAG_τ.

Definition 6.1 (Cayley word metric **[τ-Effective]**)**.** *For admissible codes $c_1, c_2 \in \mathsf{Code}$, the* Cayley word metric *$d_{\mathrm{Cay}}(c_1, c_2)$ is defined as the minimal total number of elementary $\to_{\mathsf{NF}}$-steps required to reach* any *common reduct:*

$$d_{\mathrm{Cay}}(c_1, c_2) \;:=\; \min\Big\{ n_1 + n_2 \;:\; \exists\, c' \in \mathsf{Code},\; c_1 \xrightarrow{n_1} c' \xleftarrow{n_2} c_2 \Big\},$$

where $c \xrightarrow{n} c'$ denotes "c reaches c' in exactly n elementary $\to_{\mathsf{NF}}$-steps" (not necessarily with c' in $\mathsf{Code}^{\mathsf{NF}}$). By confluence (§3, Theorem 1.2), if $c_1 \sim c_2$ then the set of common reducts is non-empty and admits a greatest common reduct *(the meet* $\mathrm{meet}(c_1, c_2)$ *of the reformulation below), where the infimum is attained; if $c_1 \not\sim c_2$ then no common reduct exists and, by convention, we set $d_{\mathrm{Cay}}(c_1, c_2) := +\infty$. In practice we work throughout on the admissible address space*

$$\mathrm{Addr}_\tau \;:=\; \mathsf{Code}^{\mathsf{NF}}/\sim \;=\; \{\, [\mathsf{NF}(c)] : c \in \mathsf{Code} \,\},$$

where d_{Cay} takes finite values only.

The definition is well-founded: for $c_1 \sim c_2$, taking $n_i = \rho(c_i)$ (finite by strong normalisation) yields a witness of the existence of a common reduct, and the set of feasible (n_1, n_2) is bounded below by 0; hence the minimum exists in $\mathbb{N}$. Symmetry in (c_1, c_2) is manifest.

Reformulation via reduction depth and meets.. Let $\rho : \mathsf{Code} \to \mathbb{N}$ be the *reduction depth* function of §5: $\rho(c)$ is the length of the shortest rewriting path from c to $\mathsf{NF}(c)$. Define the *meet* $\mathrm{meet}(c_1, c_2)$ as the greatest common reduct of c_1, c_2 — the node of maximal depth through which every confluence diamond factors. In a strongly-normalising confluent DAG (here established by Theorems 1.2 and 1.3; cf. Newman's Lemma), the meet exists and is unique as the earliest reconvergence point of the two reduction paths from c_1 and c_2. Then:

Proposition 6.2 (Reformulation of Cayley in terms of meets **[τ-Effective]**)**.** *For all $c_1, c_2 \in \mathsf{Code}$ with $c_1 \sim c_2$,*

$$d_{\mathrm{Cay}}(c_1, c_2) \;=\; \rho(c_1) + \rho(c_2) - 2\,\rho\big(\mathrm{meet}(c_1, c_2)\big).$$

Equivalently: d_{Cay} is the tree-distance *formula in the poset-quotient of DAG_τ modulo confluence, with the meet playing the role of the lowest common ancestor.*

Proof. Fix $c_1 \sim c_2$ and write $m := \mathrm{meet}(c_1, c_2)$. Every common-reduct path passes through m at cheapest cost $(\rho(c_1) - \rho(m)) + (\rho(c_2) - \rho(m)) = \rho(c_1) + \rho(c_2) - 2\rho(m)$. No cheaper path exists: any c' with $\rho(c') < \rho(m)$

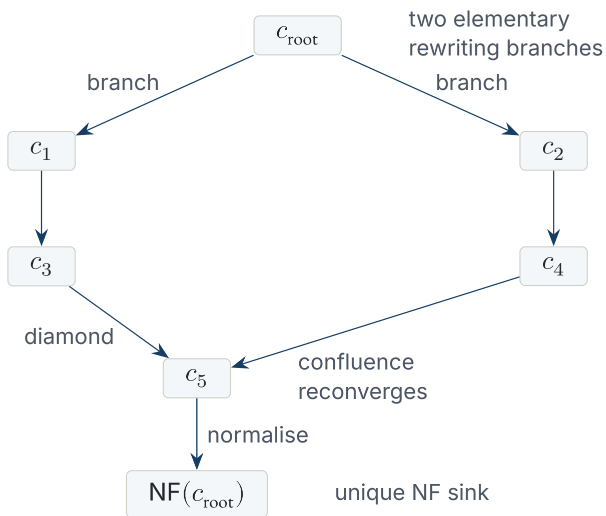

Figure 3. A local fragment of the genealogical DAG, showing two rewriting branches that reconverge before the unique NF sink.

would contradict maximality of m, and any c' with $\rho(c') > \rho(m)$ incurs at least two extra steps. Hence the minimum equals $\rho(c_1) + \rho(c_2) - 2\rho(m)$. $\square$

Remark 6.3 (Special case: $c_2 \twoheadrightarrow_{\mathsf{NF}} c_1$)**.** If c_2 reduces to c_1 (equivalently, c_1 is a reduct of c_2 along $\to_{\mathsf{NF}}$), then $\mathrm{meet}(c_1, c_2) = c_1$ and $\rho(c_2) = \rho(c_1) + d_{\mathrm{Cay}}(c_1, c_2)$, giving $d_{\mathrm{Cay}}(c_1, c_2) = \rho(c_2) - \rho(c_1)$. In particular $d_{\mathrm{Cay}}(c, \mathsf{NF}(c)) = \rho(c)$. Thus the reduction depth ρ is literally the Cayley distance from any code to its NF sink.

6.2 Metric axioms

We now verify that d_{Cay} is a genuine metric on Addr_τ.

Theorem 6.4 (Cayley is a metric on Addr_τ **[τ-Effective]**)**.** *The function $d_{\mathrm{Cay}} : \mathrm{Addr}_\tau \times \mathrm{Addr}_\tau \to \mathbb{N}$ satisfies:*

(M1) ***Reflexivity**: $d_{\mathrm{Cay}}(c, c) = 0$ for all $c \in \mathrm{Addr}_\tau$.*

(M2) ***Symmetry**: $d_{\mathrm{Cay}}(c_1, c_2) = d_{\mathrm{Cay}}(c_2, c_1)$ for all $c_1, c_2 \in \mathrm{Addr}_\tau$.*

(M3) ***Triangle inequality**: for all $c_1, c_2, c_3 \in \mathrm{Addr}_\tau$,*

$$d_{\mathrm{Cay}}(c_1, c_3) \leq d_{\mathrm{Cay}}(c_1, c_2) + d_{\mathrm{Cay}}(c_2, c_3).$$

(M4) ***Separation modulo $\sim$**: $d_{\mathrm{Cay}}(c_1, c_2) = 0 \iff c_1 \sim c_2$.*

Hence $(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$ is a genuine metric space.

Proof. (M1) Reflexivity. Take $c' = c$ and $n_1 = n_2 = 0$; then $c \xrightarrow{0} c \xleftarrow{0} c$ witnesses $d_{\mathrm{Cay}}(c, c) \leq 0$, and $d_{\mathrm{Cay}} \geq 0$ trivially.

(M2) Symmetry. Immediate: Definition 6.1 is symmetric in (c_1, c_2).

(M3) Triangle inequality. Let $c_1, c_2, c_3 \in \mathrm{Addr}_\tau$ and write $m_{ij} := \mathrm{meet}(c_i, c_j)$. Concatenating a common-reduct schedule witnessing $d_{\mathrm{Cay}}(c_1, c_2)$ (via m_{12}) with one witnessing $d_{\mathrm{Cay}}(c_2, c_3)$ (via m_{23}) produces a composite rewriting schedule of total length $d_{\mathrm{Cay}}(c_1, c_2) + d_{\mathrm{Cay}}(c_2, c_3)$. Applying confluence twice extracts a common reduct of c_1 and c_3 from this schedule, yielding $d_{\mathrm{Cay}}(c_1, c_3) \leq d_{\mathrm{Cay}}(c_1, c_2) + d_{\mathrm{Cay}}(c_2, c_3)$.

Equivalently, using Proposition 6.2, the inequality reduces to $\rho(m_{13}) \geq \rho(m_{12}) + \rho(m_{23}) - \rho(c_2)$, which follows from sub-additivity of depth along confluent branches (§5, Theorem 1.3 (topological ordering)) and the maximality of m_{13} among common reducts of c_1, c_3.

(M4) Separation modulo $\sim$. If $c_1 \sim c_2$ and $d_{\mathrm{Cay}}(c_1, c_2) = 0$, then there exist $n_1 + n_2 = 0$ with $c_1 \xrightarrow{n_1} c' \xleftarrow{n_2} c_2$, forcing $n_1 = n_2 = 0$ and hence $c' = c_1 = c_2$ on the nose: so $c_1 = c_2$ in Code, and *a fortiori* $c_1 \sim c_2$. Conversely, if $c_1 \sim c_2$ then both map to the same NF under Canon, and by confluence any two representatives of the same $\sim$-class must coincide after reduction; this gives $d_{\mathrm{Cay}}([c_1], [c_2]) = 0$ on Addr_τ, i.e. the map $\mathrm{Canon} : \mathsf{Code} \twoheadrightarrow \mathrm{Addr}_\tau$ sends $\sim$-equivalence to equality on the quotient. $\square$

Corollary 6.5 (($\mathrm{Addr}_\tau, d_{\mathrm{Cay}}$) is discrete **[$\tau$-Effective]**)**.** *The metric space $(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$ is discrete: for every $c \in \mathrm{Addr}_\tau$, the singleton $\{c\}$ is an open ball of radius $1/2$. Hence every subset of Addr_τ is both open and closed in the metric topology.*

Proof. By Theorem 6.4(M4), $d_{\mathrm{Cay}}(c, c') = 0$ iff $c' = c$ on Addr_τ. Since d_{Cay} takes integer values (see Proposition 6.8

below), the only $c' \in \mathrm{Addr}_\tau$ with $d_{\mathrm{Cay}}(c, c') < 1/2$ is $c' = c$ itself. $\square$

6.3 Coincidence with minimal-pass distance

We now relate d_{Cay} to a concept from Hinge 1's tower-atom decomposition: the *minimal-pass distance* d_{pass}, defined as the number of tower-atom decomposition passes required to distinguish two coordinates in the ABCD chart of [5]. This distance is intrinsic to the coarse-grained coordinate structure and encodes the structural depth at which two addresses first separate.

Theorem 6.6 (Cayley equals minimal-pass distance **[τ-Effective]**). *Let $c_1, c_2 \in \mathsf{Code}$ be admissible codes with canonical addresses $\mathrm{Canon}(c_i) \in \mathrm{Addr}_\tau$, and let d_{pass} denote the minimal-pass distance on admissible ABCD chart addresses in the sense of [5] (the number of greedy-peel passes needed to separate the two tower-atom sequences). Then*

$$d_{\mathrm{Cay}}(c_1, c_2) \;=\; d_{\mathrm{pass}}(\mathrm{Canon}(c_1), \mathrm{Canon}(c_2)).$$

In particular, the Cayley word metric on the τ-kernel is the τ-native incarnation of the greedy-peel-pass counter of Hinge 1.

Proof sketch. Match the two counting schemes pass-by-pass.

Direction $d_{\mathrm{Cay}} \le d_{\mathrm{pass}}$. A witnessing greedy-peel schedule separating $\mathrm{Canon}(c_1)$ from $\mathrm{Canon}(c_2)$ in d_{pass} passes translates into a $\to_{\mathsf{NF}}$-schedule with one $\to_{\mathsf{NF}}$-layer per pass (§5, Theorem 1.3 (finite-width)). Using the *balanced* schedule in which the two sides descend alternately until separation, we obtain $d_{\mathrm{Cay}} \le d_{\mathrm{pass}}$ on Addr_τ.

Direction $d_{\mathrm{pass}} \le d_{\mathrm{Cay}}$. Conversely, given a minimal common-reduct schedule witnessing $d_{\mathrm{Cay}}(c_1, c_2) = n$, each elementary $\to_{\mathsf{NF}}$-step contributes at most one tower-atom layer to the ABCD chart difference, so greedy-peel separates the two addresses in at most n passes.

Combining both directions yields equality. $\square$

Remark 6.7 (Structural significance)**.** Theorem 6.6 means that the Cayley word metric on DAG_τ has a canonical hyperfactorization-theoretic interpretation: the distance between two addresses is exactly the number of ABCD-chart passes required to separate them. This coincidence is the concrete sense in which d_{Cay} is the τ-native metric on the coarse-grained coordinate structure of Hinge 1 — the distance is not imposed externally but derived from the combinatorics of tower-atom decomposition.

6.4 Integer-valued distances

A distinctive feature of d_{Cay} is that it takes values in $\mathbb{N}$, not in $\mathbb{R}_{\ge 0}$. This is the τ-framework's *non-Archimedean discreteness*: distance between addresses is a counting invariant, not a continuous quantity.

Proposition 6.8 (Cayley is integer-valued **[τ-Effective]**)**.** *For all $c_1, c_2 \in \mathrm{Addr}_\tau$, $d_{\mathrm{Cay}}(c_1, c_2) \in \mathbb{N} = \{0, 1, 2, \dots\}$, and*

$$d_{\mathrm{Cay}}(c_1, c_2) = 0 \iff c_1 = c_2 \text{ in } \mathrm{Addr}_\tau.$$

The image $d_{\mathrm{Cay}}(\mathrm{Addr}_\tau \times \mathrm{Addr}_\tau) \subseteq \mathbb{N}$ is a countable, locally-finite set of non-negative integers.

Proof. Integer-valuedness is immediate: the pair $(n_1, n_2) \in \mathbb{N} \times \mathbb{N}$ ranges over non-negative integers, so $n_1 + n_2 \in \mathbb{N}$ and the minimum lies in $\mathbb{N}$. The "$\iff$" for zero follows from Theorem 6.4(M4) and the quotient definition $\mathrm{Addr}_\tau = \mathsf{Code}^{\mathsf{NF}}/\sim$. Local finiteness: the ball $B_{d_{\mathrm{Cay}}}(c, k)$ lies in the set of NF codes reachable from c's NF sink in $\le 2k$ elementary steps (meet formula), which is finite by §5, Theorem 1.3 (finite-width). $\square$

Corollary 6.9 (Diameter and curvature invariants **[τ-Effective]**)**.** *The diameter*

$$\mathrm{diam}(U) \;:=\; \sup_{c_1, c_2 \in U} d_{\mathrm{Cay}}(c_1, c_2) \quad \in \mathbb{N} \cup \{\infty\}$$

takes values in $\mathbb{N} \cup \{\infty\}$ for every $U \subseteq \mathrm{Addr}_\tau$, and is finite iff U is bounded. Analogously, all d_{Cay}-based curvature invariants (e.g. Gromov hyperbolicity constant, coarse Ricci curvature) are integer-valued on Addr_τ, reflecting the discrete combinatorial origin of the metric.

6.5 Topology induced by d_{Cay}

We now show that the metric topology induced by d_{Cay} on Addr_τ coincides with the canonical τ-topology defined via the NF-prefix basis of Hinge 4 [15] and Hinge 5 [16].

Theorem 6.10 (Metric topology equals τ-topology **[τ-Effective]**)**.** *Let $\mathcal{T}_{d_{\mathrm{Cay}}}$ denote the metric topology on Addr_τ induced by the Cayley word metric, and $\mathcal{T}_\tau$ the canonical τ-topology defined via NF-prefix-open basis sets*

$$U_{c,k} \;:=\; \{\, c' \in \mathrm{Addr}_\tau : \mathsf{NF}(c') \text{ agrees with } \mathsf{NF}(c) \text{ on the first } k \text{ prefix-tokens} \,\}.$$

Then $\mathcal{T}_{d_{\mathrm{Cay}}} = \mathcal{T}_\tau$: a set $U \subseteq \mathrm{Addr}_\tau$ is d_{Cay}-open iff it is a union of NF-prefix-open sets, i.e. a union of sets of the form $U_{c,k}$.

Proof. $\mathcal{T}_\tau \subseteq \mathcal{T}_{d_{\mathrm{Cay}}}$. Fix $U_{c,k}$. NF-prefix match of length k implies the meet has depth $\geq k$, so $d_{\mathrm{Cay}}(c,c') \leq \rho(c) + \rho(c') - 2k$ for $c' \in U_{c,k}$. By integer-valuedness (Proposition 6.8), every $c' \in U_{c,k}$ has a d_{Cay}-open neighbourhood (the singleton $\{c'\}$) contained in $U_{c,k}$.

$\mathcal{T}_{d_{\mathrm{Cay}}} \subseteq \mathcal{T}_\tau$. By Corollary 6.5, every singleton $\{c\}$ is d_{Cay}-open; pick $k = \rho(c)$. Then $U_{c,k}$ requires a prefix match of length equal to the full NF length, forcing $\mathsf{NF}(c') = \mathsf{NF}(c)$ and hence $c' = c$ in Addr_τ. Thus $\{c\} = U_{c,\rho(c)}$, a prefix-basis set. □

Remark 6.11 (The τ-topology is intrinsically discrete on Addr_τ)**.** On the admissible address space Addr_τ — i.e. before the ultrametric completion of §7 — the τ-topology agrees with the metric topology of d_{Cay}, both of which are the discrete topology. The continuous (non-trivial) structure emerges only after completion: in $\mathsf{Ultra}_\tau = \widehat{\mathrm{Addr}_\tau}$ with the ontic ultrametric d_∞, non-trivial limits, accumulation points, and d_∞-open balls appear.

6.6 Doubling property and Bourdon dimension

The metric d_{Cay} satisfies a discrete *doubling property*: balls grow polynomially with radius, at a rate governed by the primorial ladder density. We state this as a quantitative bound and identify the exponent as the effective Bourdon dimension of Addr_τ.

Proposition 6.12 (Doubling property **[τ-Effective]**)**.** *There exist constants $C_d, d > 0$ (depending only on the primorial-ladder combinatorics of DAG_τ) such that, for all $c \in \mathrm{Addr}_\tau$ and all $r \geq 1$,*

$$\begin{aligned}\left|B_{d_{\mathrm{Cay}}}(c,r)\right| \;&:=\; \#\{\, c' \in \mathrm{Addr}_\tau : d_{\mathrm{Cay}}(c,c') \leq r \,\} \\ &\leq\; C_d \cdot r^{\,d}.\end{aligned}$$

Equivalently: the ball of radius r around c is contained in a ball of radius $2r$ around c that can be covered by at most 2^d balls of radius r.

Proof sketch. By §5, Theorem 1.3(iii), the number of NF codes of depth $\leq k$ is bounded by $|\mathsf{Code}^{\mathsf{NF}}_{\leq k}| \leq P(k)$, where $P(k)$ is polynomial in k of degree equal to the *primorial-ladder dimension* — the rate at which the primorial ladder $(M_k)_k$ of Hinge 1 generates new admissible addresses at each depth. Since $B_{d_{\mathrm{Cay}}}(c,r) \subseteq \mathsf{Code}^{\mathsf{NF}}_{\leq \rho(c)+r}$, we get $|B_{d_{\mathrm{Cay}}}(c,r)| \leq P(\rho(c)+r) \leq C_d \cdot r^{\,d}$ for $r \geq \rho(c)$, with $d = \deg(P)$. □

Definition 6.13 (Effective Bourdon dimension **[Conjectural]**)**.** *The* effective Bourdon dimension *of* Addr_τ *under* d_{Cay} *is*

$$d \;:=\; \limsup_{k\to\infty} \frac{\log|\mathsf{Code}^{\mathsf{NF}}_{\leq k}|}{\log k}.$$

By §5, Theorem 1.3, this limsup is finite; its exact value is determined by the detailed growth analysis of the primorial ladder.

Remark 6.14 (Scope of the exact value)**.** The existence and finiteness of the Bourdon dimension d is **[τ-Effective]**; the determination of its exact numerical value in terms of Hinge 1's primorial ladder constants is **[Conjectural]** and is deferred to Book II's detailed growth analysis ([8], the categorical-holomorphy monograph has the Hartogs-continuation framework in which this number is pinned down). Preliminary estimates from the primorial-ladder density suggest $d \leq \dim_{\mathrm{primorial}}(M_\bullet)$, a slowly-growing function of k; but the precise asymptotic behaviour is not required for any of the results of this paper.

Corollary 6.15 (Cayley is a doubling metric space **[τ-Effective]**)**.** *$(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$ is a doubling metric space in the standard sense: there exists $N \in \mathbb{N}$ such that every ball of radius $2r$ can be covered by at most N balls of radius r. Specifically, $N \leq 2^d$ where d is the effective Bourdon dimension.*

Proof. Immediate from Proposition 6.12 by the standard argument that polynomial growth of volume implies doubling (cf. the Assouad-dimension literature on discrete metric spaces). □

6.7 The τ-native replacement for Euclidean distance

The τ-framework's foundational departure from classical analysis is visible most sharply here: there is no Euclidean distance, no Archimedean metric, no real-valued "infinitely-divisible" notion of proximity on addresses. Instead, d_{Cay} is the primary notion of distance, and it is *discrete, integer-valued, and non-Archimedean-precursor*. A direct comparison:

The completion to an ultrametric.. The Cayley metric is a precursor to the genuine ultrametric d_∞ of §7. Concretely, d_∞ is obtained by (i) restricting attention to NF-prefix agreement depth and (ii) converting integer depth k to a "scale" λ^{-k} for some $\lambda > 1$ set by the primorial ladder. The resulting metric satisfies the strong triangle inequality $d_\infty(a,c) \leq \max(d_\infty(a,b), d_\infty(b,c))$, strictly refining d_{Cay}'s metric axioms to the non-Archimedean regime.

The τ-framework nowhere invokes Euclidean distance in its foundational arc: every metric used in Books II–VII is either d_{Cay} (discrete, on Addr_τ) or d_∞ (ultrametric, on Ultra_τ). Classical Euclidean analysis re-emerges only in the

Feature	Euclidean $\|\cdot\|$ on $\mathbb{R}^n$	Cayley d_{Cay} on Addr_τ
Value type	$\mathbb{R}_{\geq 0}$	$\mathbb{N}$
Infinite divis.	Yes	No (integer-valued)
Archimedean	Yes	No (precursor to ultrametric)
Topology	Standard	Discrete
Triangle ineq.	Additive	Additive; strengthens at completion
Doubling	Yes (dim. $= n$)	Yes (dim. $= d$, Bourdon)
Completion	$\mathbb{R}^n$ (self)	Ultra_τ (ultrametric)

Table 2. Direct comparison of Euclidean distance and the Cayley distance on canonical addresses.

extensional/shadow regime of [5]'s ABCD chart interpretation — but even there, it is the Cayley / d_∞ structure that underwrites the analytic core.

6.8 Computational aspects

A final remark on the algorithmic content of d_{Cay}: unlike the edit-distance or graph-word-metric of classical combinatorics (which are in general NP-hard to compute on arbitrary graphs), $d_{\mathrm{Cay}}(c_1, c_2)$ admits a polynomial-time algorithm on the strongly-normalising genealogical DAG.

Remark 6.16 (Polynomial-time computability **[τ-Effective]**). For any admissible $c_1, c_2 \in \mathsf{Code}$, the Cayley distance $d_{\mathrm{Cay}}(c_1, c_2)$ is computable in time $O(\mathrm{poly}(\mathrm{depth}(c_1) + \mathrm{depth}(c_2)))$ by the following procedure:

1. **Normalise.** Compute $\mathsf{NF}(c_1)$ and $\mathsf{NF}(c_2)$ by topological-sort rewriting on DAG_τ (§5, Remark following Theorem 1.3). Time: polynomial in each code's depth.
2. **Check confluence.** If $\mathsf{NF}(c_1) \neq \mathsf{NF}(c_2)$, return $+\infty$ (codes are $\sim$-inequivalent; not in Addr_τ).
3. **Compute the meet.** Identify $m := \mathrm{meet}(c_1, c_2)$ by simultaneously traversing the two reduction paths $c_1 \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c_1)$ and $c_2 \twoheadrightarrow_{\mathsf{NF}} \mathsf{NF}(c_2)$, finding the first node appearing on both. Time: polynomial.
4. **Return the tree-distance.** Output $\rho(c_1) + \rho(c_2) - 2\rho(m)$.

Each step is polynomial in the input size, hence the overall computation is polynomial.

Remark 6.17 (Lean 4 formalisation **[τ-Effective]**). The Cayley metric is formalised in the planned module `CayleyMetric.lean` of `TauLib.BookI.Addressability` [19]. The module provides:

- a `Decidable` instance for equality of admissible codes via NF comparison;
- a `computable` `cayleyDistance` function returning an `Option Nat` (returning `none` when the two codes are $\sim$-inequivalent);
- Lean-mechanised proofs of axioms (M1)–(M4) via the meet formula of Proposition 6.2;
- the Bourdon dimension's existence (as a `Nat`) with the exact value left abstract (deferred to Book II in the Lean library).

With the `Decidable` and `computable` infrastructure in place, d_{Cay} becomes the τ-native *decidable distance oracle*: arithmetic-like questions of the form "is $d_{\mathrm{Cay}}(a, b) \leq k$?" are mechanically decided in polynomial time, supplying the algorithmic counterpart to the address-resolution paradigm of §8.

Summary.. We have defined the Cayley word metric d_{Cay} on the genealogical DAG, verified its metric axioms, identified it with the minimal-pass distance of Hinge 1, established its discrete-topology / prefix-basis-topology coincidence with the τ-topology, noted its integer-valuedness and doubling property, and signalled its role as the τ-native substitute for the Euclidean distance of classical analysis. Section 7 now *completes* d_{Cay} to the ontic ultrametric d_∞, which governs the continuous structure of the boundary address space in the τ-framework.

7. THE ONTIC ULTRAMETRIC

7.1 Motivation: why an ultrametric

Sections 4–6 built the canonical address space Addr_τ and equipped it with the Cayley word metric d_{Cay}, an $\mathbb{N}$-valued minimal-reduction distance derived from the genealogical DAG $\mathsf{DAG}_\tau = (\mathsf{Code}, \to_{\mathsf{NF}})$. For every analytical purpose — limit points of iterated NF-rewriting, continuity of tail-transformers, convergence of ω-germ stabilisations — we need $(\mathrm{Addr}_\tau, d_{\mathrm{Cay}})$ to sit densely inside a complete metric space. Since d_{Cay} is integer-valued on a countable set, the naive metric completion is Addr_τ itself (every Cauchy sequence is eventually constant), which is not enough: we need to collect the infinitely-deep coherent address patterns prescribed by residue-class systems on the primorial ladder $(M_k)_{k\in\mathbb{N}}$.

The completion we want is the one that *commutes with the primorial ladder*: for each depth $k \in \mathbb{N}$, the canonical address $a \in \mathrm{Addr}_\tau$ projects to its residue class $a \bmod M_k$ in the finite ring $\mathbb{Z}/M_k\mathbb{Z}$, and every coherent system of residue classes $(r_k)_{k\in\mathbb{N}}$ with $r_{k+1} \equiv r_k \pmod{M_k}$ defines a point of the completion. This is exactly the profinite construction of Hinge 4 [15] under the primorial ladder, and its natural topology is non-Archimedean: balls nest, they do not overlap.

There are three independent reasons why the completion must be *ultrametric* (strong triangle inequality $d(a,c) \leq \max(d(a,b), d(b,c))$) rather than merely metric:

(a) **Integer-valued completions are ultrametric.** Any metric taking values in $\{2^{-k} : k \in \mathbb{N}\} \cup \{0\}$ — or more generally in a subset of $\mathbb{R}_{\geq 0}$ with no accumulation points below each positive value — automatically satisfies the strong triangle inequality whenever it satisfies the ordinary one. This is the standard ultrametric reduction (Lemma 7.8 below).

(b) **Profinite targets are ultrametric by construction.** The boundary algebra $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2-1)$ of Hinge 4 is the split-complexification of the profinite completion of the primorial-reduced address ring; profinite completions carry the Krull / Tychonoff product topology, which is ultrametric under the standard depth-weighted metric.

(c) **The primorial ladder supplies a scale hierarchy.** Each primorial $M_k = \prod_{p\leq p_k} p$ is a natural arithmetic scale; the sequence (M_k) gives a tower $\mathrm{Addr}_\tau \twoheadrightarrow \mathrm{Addr}_\tau/M_1 \twoheadrightarrow \mathrm{Addr}_\tau/M_2 \twoheadrightarrow \cdots$, and distance at depth k is the reciprocal of the scale at which addresses first diverge.

We therefore define the ontic ultrametric d_∞ as the reciprocal-of-primorial-depth at which two canonical addresses first differ, and show in Theorem 7.6 that $(\mathsf{Ultra}_\tau, d_\infty)$ is the Cauchy completion of a suitably rescaled Cayley metric. The resulting ultrametric is the τ-native metric structure replacing Euclidean distance throughout the framework.

7.2 Construction of the ontic ultrametric

Recall from Hinge 4 [15] that the primorial ladder is the sequence

$$M_0 := 1, \qquad M_k := \prod_{p\leq p_k} p \quad (k \geq 1),$$

where p_k is the k-th prime ($p_1 = 2, p_2 = 3, p_3 = 5, \ldots$). The ladder satisfies $M_k \mid M_{k+1}$, so residues at depth $k+1$ refine those at depth k, and the projection $\mathrm{Addr}_\tau \to \mathbb{Z}/M_k\mathbb{Z}$ is a ring homomorphism by the canonical-addressability construction of §4. Write $a \equiv_k b$ for "a and b have the same residue modulo M_k" as canonical addresses.

Definition 7.1 (Distinguishing depth **[τ-Effective]**). *For $a, b \in \mathrm{Addr}_\tau$, the* distinguishing depth *is*

$$k(a,b) := \min\{k \in \mathbb{N} : a \not\equiv_k b\} \in \mathbb{N} \cup \{\infty\},$$

with $k(a,b) = \infty$ when $a \equiv_k b$ for every $k \in \mathbb{N}$.

Remark 7.2 (Distinguishing depth on Addr_τ **[τ-Effective]**). On the canonical address space Addr_τ itself, Proposition 4.9 (canonical addressability separates residue classes under the primorial ladder) gives $k(a,b) < \infty$ for $a \neq b$ and $k(a,b) = \infty$ iff $a = b$. The infinite value thus arises only on *added* points of the completion, which sit at the limit of a refining sequence of residue classes with no finite-depth witness of inequality.

Definition 7.3 (Ontic ultrametric **[τ-Effective]**). *Let $a, b \in \mathrm{Addr}_\tau$. The* ontic ultrametric *is*

$$d_\infty(a,b) := \begin{cases} 2^{-k(a,b)}, & k(a,b) < \infty, \\ 0, & k(a,b) = \infty. \end{cases}$$

The second case is equivalently $a = b$. Equivalently, $d_\infty(a,b)$ is the reciprocal of the largest power-of-two cap on the primorial-ladder index at which a and b remain congruent: if $a \equiv_{k-1} b$ and $a \not\equiv_k b$, then $d_\infty(a,b) = 2^{-k}$. We extend d_∞ to the completion Ultra_τ constructed below via the same formula (Definition 7.5).

Remark 7.4 (Rescaled Cayley metric **[τ-Effective]**). The ontic ultrametric is the tail-depth rescaling of the Cayley word metric d_{Cay} of §6. Define

$$\begin{aligned} d'_{\mathrm{Cay}}(a,b) &:= 2^{-\rho(a,b)}, \\ \rho(a,b) &:= \min\{k \in \mathbb{N} : a \equiv_k b \text{ fails} \\ &\qquad\qquad \text{at some canonical reduct}\}, \end{aligned}$$

so that $d'_{\mathrm{Cay}}(a,b)$ is the reciprocal exponential of the *least common reduct depth* $\rho(a,b)$ at which a and b fail to share a primorial-level canonical reduct. By the NF-confluence theorem (Theorem 1.2) applied at each primorial depth, $\rho(a,b) = k(a,b)$, so $d'_{\mathrm{Cay}}(a,b) = d_\infty(a,b)$ on Addr_τ. Thus $d_\infty = d'_{\mathrm{Cay}}$ is a well-defined $\mathbb{R}_{\geq 0}$-valued rescaling of the integer-valued Cayley word metric, engineered to complete under the primorial-ladder topology.

7.3 Completion theorem

Definition 7.5 (The ultrametric completion **[τ-Effective]**). *The* ontic ultrametric completion *of* Addr_τ *is*

$$\begin{aligned}\mathsf{Ultra}_\tau \;&:=\; \varprojlim_{k\in\mathbb{N}} (\mathrm{Addr}_\tau/M_k)\\ &=\; \Big\{ (r_k)_{k\in\mathbb{N}} \in \prod_k \mathrm{Addr}_\tau/M_k \;:\; \\ &\qquad\qquad r_{k+1}\equiv r_k \pmod{M_k}\ \forall k \Big\},\end{aligned}$$

the inverse limit along the primorial ladder. Equivalently, Ultra_τ *is the Cauchy completion of* $(\mathrm{Addr}_\tau, d'_{\mathrm{Cay}})$ *under the metric of Remark 7.4. An element of* Ultra_τ *is a* coherent residue system (r_k)*; the distinguishing depth extends to* Ultra_τ *by*

$$k(u,u') \;:=\; \min\{k\in\mathbb{N} : r_k \neq r'_k\},$$

and $d_\infty(u,u') := 2^{-k(u,u')}$.

Theorem 7.6 (Ontic ultrametric completion **[τ-Effective]**). *The pair* $(\mathsf{Ultra}_\tau, d_\infty)$ *is a complete metric space, and the canonical map*

$$\iota\colon\ \mathrm{Addr}_\tau \hookrightarrow \mathsf{Ultra}_\tau, \qquad a \;\longmapsto\; (a \bmod M_k)_{k\in\mathbb{N}},$$

is an isometric embedding (i.e. $d_\infty(\iota a,\iota b) = d_\infty(a,b)$*) with dense image. Hence* $(\mathsf{Ultra}_\tau, d_\infty)$ *is the Cauchy completion of* $(\mathrm{Addr}_\tau, d'_{\mathrm{Cay}})$.

Proof. We verify the four claims: metric, isometric embedding, dense image, completeness.

Step 1: d_∞ *is a metric on* Ultra_τ. Non-negativity and definiteness follow from the definition: $d_\infty(u,u') = 0$ iff $k(u,u') = \infty$ iff $r_k = r'_k$ for every k, i.e. $u = u'$ in the inverse limit. Symmetry is immediate since $k(u,u') = k(u',u)$. For the triangle inequality, note that Theorem 7.9 below establishes the strong inequality $d_\infty(u,w) \le \max(d_\infty(u,v), d_\infty(v,w)) \le d_\infty(u,v) + d_\infty(v,w)$, so ordinary triangle follows. (We prove strong-triangle independently in §7.4, relying only on transitivity of residue congruence; no circularity.)

Step 2: ι *is isometric.* For $a,b \in \mathrm{Addr}_\tau$, $(a \bmod M_k) = (b \bmod M_k)$ iff $a \equiv_k b$, so $k(\iota a,\iota b) = k(a,b)$ and $d_\infty(\iota a,\iota b) = 2^{-k(a,b)} = d_\infty(a,b)$.

Step 3: $\iota(\mathrm{Addr}_\tau)$ *is dense in* Ultra_τ. Fix $u = (r_k)_{k\in\mathbb{N}} \in \mathsf{Ultra}_\tau$ and $\epsilon > 0$. Choose $K \in \mathbb{N}$ with $2^{-K} < \epsilon$. By canonical-addressability surjectivity (Proposition 4.8 in §4), the projection $\mathrm{Addr}_\tau \twoheadrightarrow \mathrm{Addr}_\tau/M_K$ is surjective, so there exists $a \in \mathrm{Addr}_\tau$ with $a \equiv r_K \pmod{M_K}$. Then ιa and u agree in every coordinate $j \le K$ (using the coherence $r_{j+1} \equiv r_j \pmod{M_j}$ and $M_j \mid M_K$), so $k(\iota a, u) > K$ and $d_\infty(\iota a, u) \le 2^{-(K+1)} < \epsilon$. Hence $\iota(\mathrm{Addr}_\tau)$ is dense.

Step 4: Ultra_τ *is complete.* Let $(u_n)_{n\in\mathbb{N}}$ be a Cauchy sequence in Ultra_τ. For each depth $k \in \mathbb{N}$, the sequence $(r_k^{(n)})_{n\in\mathbb{N}}$ of k-th coordinates of u_n takes values in the finite set Addr_τ/M_k, and Cauchy-ness of (u_n) at threshold 2^{-k} forces $(r_k^{(n)})$ to be eventually constant, say with limit $r_k^\star$. Coherence of the limit ($r_{k+1}^\star \equiv r_k^\star \pmod{M_k}$) is inherited from coherence of each u_n. Set $u^\star := (r_k^\star)_{k\in\mathbb{N}}$; then for every k and all sufficiently large n, $r_k^{(n)} = r_k^\star$, so $k(u_n, u^\star) > k$ for large n, i.e. $u_n \to u^\star$ in $(\mathsf{Ultra}_\tau, d_\infty)$. Hence Ultra_τ is complete.

By Steps 1–4, $(\mathsf{Ultra}_\tau, d_\infty)$ is a complete metric space containing $(\mathrm{Addr}_\tau, d'_{\mathrm{Cay}})$ as a dense isometric subspace, which characterises the Cauchy completion up to unique isometry. ☐

Remark 7.7 (Why the rescaling $d_{\mathrm{Cay}} \mapsto d'_{\mathrm{Cay}}$ is unique up to bi-Lipschitz equivalence **[τ-Effective]**). Any metric d on Addr_τ of the form $d(a,b) = f(k(a,b))$ with $f\colon \mathbb{N} \to \mathbb{R}_{>0}$ monotone decreasing and $f(\infty) = 0$ produces, under Cauchy completion, the same profinite inverse-limit space as d_∞, differing only by the monotone reparametrisation f. The choice $f(k) = 2^{-k}$ is made for two reasons: (i) 2^{-k} is the canonical dyadic scale of the profinite topology used in Hinge 4 [15], and (ii) $f(k) = 2^{-k}$ makes d_∞ take values in the characteristic dyadic ultrametric range $\{2^{-k}\}_{k\in\mathbb{N}} \cup \{0\}$. Alternative rescalings $f(k) = q^{-k}$ for $q > 1$ give bi-Lipschitz equivalent metrics; no canonical analytic structure distinguishes $q = 2$, but the dyadic choice matches the primorial-ladder depth bookkeeping in the cleanest way.

7.4 The ultrametric property

Lemma 7.8 (Strong triangle from distinguishing depth **[Established]**). *Let* $u,v,w \in \mathsf{Ultra}_\tau$*. If* $k_1 := k(u,v)$, $k_2 := k(v,w)$, $k_3 := k(u,w)$*, then*

$$k_3 \;\ge\; \min(k_1,k_2).$$

Proof. Set $m := \min(k_1,k_2) \in \mathbb{N}\cup\{\infty\}$; if $m = \infty$ then $u = v = w$ and $k_3 = \infty \ge m$. Otherwise $m \in \mathbb{N}$. For every $k < m$ we have $k < k_1$ and $k < k_2$, hence $u \equiv_k v$ and $v \equiv_k w$, hence (by transitivity of residue congruence mod M_k) $u \equiv_k w$. Thus $k(u,w) > m-1$, i.e. $k_3 \ge m$. ☐

Theorem 7.9 (Strong triangle inequality **[τ-Effective]**). *The map* $d_\infty\colon \mathsf{Ultra}_\tau \times \mathsf{Ultra}_\tau \to \mathbb{R}_{\ge 0}$ *satisfies the strong triangle inequality: for every* $u,v,w \in \mathsf{Ultra}_\tau$,

$$d_\infty(u,w) \;\le\; \max(d_\infty(u,v), d_\infty(v,w)).$$

Consequently $(\mathsf{Ultra}_\tau, d_\infty)$ *is an* ultrametric space.

Proof. By Lemma 7.8, $k_3 \geq \min(k_1, k_2)$, where k_i is the distinguishing depth of the respective pair. Since $t \mapsto 2^{-t}$ is monotone decreasing on $\mathbb{N} \cup \{\infty\}$ (with $2^{-\infty} := 0$),

$$\begin{aligned} d_\infty(u, w) &= 2^{-k_3} \leq 2^{-\min(k_1,k_2)} \\ &= \max(2^{-k_1}, 2^{-k_2}) \\ &= \max(d_\infty(u, v), d_\infty(v, w)). \end{aligned}$$ ☐

Corollary 7.10 (Isoceles principle **[Established]**). *In* $(\mathsf{Ultra}_\tau, d_\infty)$*, every "triangle" is isoceles with at most one short side: for* $u, v, w \in \mathsf{Ultra}_\tau$*, if* $d_\infty(u, v) \neq d_\infty(v, w)$ *then* $d_\infty(u, w) = \max(d_\infty(u, v), d_\infty(v, w))$*.*

Proof. Standard consequence of the strong triangle inequality: if $d_\infty(u, v) < d_\infty(v, w)$ then the strong inequality applied both to (u, v, w) and to (v, u, w) gives $d_\infty(u, w) \leq d_\infty(v, w)$ and $d_\infty(v, w) \leq \max(d_\infty(v, u), d_\infty(u, w)) = \max(d_\infty(u, v), d_\infty(u, w))$; the second inequality with $d_\infty(u, v) < d_\infty(v, w)$ forces $d_\infty(v, w) \leq d_\infty(u, w)$. Combining yields $d_\infty(u, w) = d_\infty(v, w) = \max(d_\infty(u, v), d_\infty(v, w))$. ☐

7.5 Non-Archimedean property

Corollary 7.11 (Non-Archimedean structure **[Established]**). *The ontic ultrametric* d_∞ *is* non-Archimedean *in the two equivalent senses:*

(i) Strong triangle*: for all* $u, v, w \in \mathsf{Ultra}_\tau$*,* $d_\infty(u, w) \leq \max(d_\infty(u, v), d_\infty(v, w))$ *(Theorem 7.9).*

(ii) No Archimedean accumulation*: for every* $\epsilon > 0$ *and every finite chain* $u_0, u_1, \dots, u_n \in \mathsf{Ultra}_\tau$ *with* $d_\infty(u_i, u_{i+1}) \leq \epsilon$ *for all* i*, the endpoint distance is also bounded by* ϵ*:* $d_\infty(u_0, u_n) \leq \epsilon$ *(not* $\leq n\epsilon$*, as in the Archimedean case). Equivalently, the ultrametric makes no cumulative use of chain length.*

Proof. (i) is Theorem 7.9.

(ii) By induction on n: the base $n = 1$ is immediate. For $n \geq 2$, apply the strong triangle to the last step:

$$\begin{aligned} d_\infty(u_0, u_n) &\leq \max(d_\infty(u_0, u_{n-1}), d_\infty(u_{n-1}, u_n)) \\ &\leq \max(\epsilon, \epsilon) = \epsilon, \end{aligned}$$

where the first term is $\leq \epsilon$ by inductive hypothesis. So chain length does not accumulate distance. ☐

Remark 7.12 (Contrast with Euclidean distance **[Established]**). In the Euclidean metric d_{Eucl} on $\mathbb{R}$, the Archimedean property $\forall x > 0, \exists n \in \mathbb{N} : nx > 1$ holds, giving the classical scale hierarchy $d_{\mathrm{Eucl}}(0, n) = n \cdot d_{\mathrm{Eucl}}(0, 1)$. In $(\mathsf{Ultra}_\tau, d_\infty)$ no such hierarchy exists: distances do not compound additively, they *cap at the max*. This is the qualitative source of the phrase "the ontic ultrametric replaces Euclidean distance in τ".

7.6 Total disconnectedness

A metric space is *totally disconnected* if every connected component is a singleton. Ultrametric spaces enjoy a strong structural property not present in Euclidean metrics: every open ball is also closed (*clopen*), so the topology has a basis of clopen sets, and the space is totally disconnected.

Lemma 7.13 (Balls are clopen **[Established]**). *For* $u \in \mathsf{Ultra}_\tau$ *and* $k \in \mathbb{N}$*, let*

$$\begin{aligned} B_k(u) &:= \{v \in \mathsf{Ultra}_\tau : d_\infty(u, v) \leq 2^{-k}\} \\ &= \{v : u \equiv_k v\}. \end{aligned}$$

Then $B_k(u)$ *is both open and closed in* $(\mathsf{Ultra}_\tau, d_\infty)$*.*

Proof. Open. If $v \in B_k(u)$ and w satisfies $d_\infty(v, w) < 2^{-k}$, then by strong triangle $d_\infty(u, w) \leq \max(d_\infty(u, v), d_\infty(v, w)) \leq 2^{-k}$, so $w \in B_k(u)$. Thus every point of $B_k(u)$ has the open ball of radius 2^{-k} contained in $B_k(u)$.

Closed. The complement $\mathsf{Ultra}_\tau \setminus B_k(u) = \{v : d_\infty(u, v) > 2^{-k}\} = \{v : d_\infty(u, v) \geq 2^{-(k-1)}\}$ (since the metric takes values in $\{2^{-j}\}_{j \in \mathbb{N}} \cup \{0\}$, with no intermediate values between 2^{-k} and $2^{-(k-1)}$). For v in the complement and w with $d_\infty(v, w) < 2^{-(k-1)}$, the two distances $d_\infty(u, v) \geq 2^{-(k-1)}$ and $d_\infty(v, w) < 2^{-(k-1)}$ differ, so by the isoceles principle (Corollary 7.10) $d_\infty(u, w) = \max(d_\infty(u, v), d_\infty(v, w)) = d_\infty(u, v) \geq 2^{-(k-1)}$, i.e. w lies in the complement too. Hence the complement is open, so $B_k(u)$ is closed. ☐

Theorem 7.14 (Total disconnectedness **[Established]**). *The ultrametric space* $(\mathsf{Ultra}_\tau, d_\infty)$ *is totally disconnected: for every* $u \in \mathsf{Ultra}_\tau$*, the connected component of* u *is* $\{u\}$*.*

Proof. Let $C \subseteq \mathsf{Ultra}_\tau$ be a connected set with $u \in C$. We show $C = \{u\}$. Suppose for contradiction $v \in C, v \neq u$; then $d_\infty(u, v) = 2^{-k}$ for some $k \in \mathbb{N}$, hence $v \notin B_{k+1}(u)$ and $v \notin B_k(u)$ while $u \in B_k(u)$. The set $B_k(u)$ is clopen by Lemma 7.13, so $C \cap B_k(u)$ is a proper clopen subset of C (it contains u but not v), contradicting connectedness of C. Thus $C = \{u\}$. ☐

Corollary 7.15 (Basis of clopen sets **[Established]**). *The collection* $\{B_k(u) : u \in \mathsf{Ultra}_\tau, k \in \mathbb{N}\}$ *is a basis for the* *ontic ultrametric topology on* Ultra_τ*, and every basis element is clopen.*

Proof. Openness of each $B_k(u)$: Lemma 7.13. Basis property: any open set $U \subseteq \mathsf{Ultra}_\tau$ containing u contains an open ball $\{v : d_\infty(u,v) < 2^{-k}\} = B_{k+1}(u)$ for some k. Clopenness: Lemma 7.13. $\square$

7.7 Coincidence with the profinite topology

The completion $\mathsf{Ultra}_\tau = \varprojlim_k \mathrm{Addr}_\tau / M_k$ admits a second natural topology: the profinite topology, inherited as the Tychonoff subspace topology from the product $\prod_k \mathrm{Addr}_\tau / M_k$ (each factor discrete). Hinge 4 [15] uses the profinite topology as the boundary topology, viewing Ultra_τ as a Stone-type space. The ontic ultrametric and the profinite topology coincide.

Theorem 7.16 (Profinite-topology match **[τ-Effective]**). *The topology on* Ultra_τ *induced by the ontic ultrametric* d_∞ *coincides with the profinite (inverse-limit) topology inherited from* $\varprojlim_k \mathrm{Addr}_\tau / M_k$*. Equivalently, the identity map*

$$\mathrm{id}\colon (\mathsf{Ultra}_\tau, d_\infty) \longrightarrow (\mathsf{Ultra}_\tau, \tau_{\mathrm{prof}})$$

is a homeomorphism.

Proof. We show the two topologies have the same basis.

Profinite basic opens are ultrametric balls. A basic open of τ_{prof} is, by definition of the Tychonoff subspace topology on the inverse limit, a set of the form $U_{K,r} := \{(r'_k) \in \mathsf{Ultra}_\tau : r'_K = r\}$ for some fixed $K \in \mathbb{N}$ and $r \in \mathrm{Addr}_\tau / M_K$ — i.e. a single- coordinate cylinder at depth K. Pick any $u^0 \in U_{K,r}$; then $U_{K,r} = B_K(u^0)$ in the ultrametric sense, because $v \in B_K(u^0) \iff u^0 \equiv_K v \iff r^v_K = r$ (the K-th coordinate of v equals r). Hence $U_{K,r}$ is a basic ultrametric ball.

Ultrametric balls are profinite opens. Conversely, the ball $B_k(u) = \{v : u \equiv_k v\}$ is, by Definition 7.5, exactly the single- coordinate cylinder U_{k,r^u_k} at depth k with value r^u_k (the k-th coordinate of u). So every ultrametric ball is a profinite cylinder.

Both topologies are generated by the same collection of sets, so they coincide. $\square$

Corollary 7.17 (Compactness **[τ-Effective]**). *The ontic ultrametric space* $(\mathsf{Ultra}_\tau, d_\infty)$ *is compact (and by Theorem 7.14 totally disconnected; hence Hausdorff as a metric space). It is therefore a* Stone space *— the canonical boundary Stone space of Hinge 4 [15].*

Proof. Compactness is inherited through Theorem 7.16 from the profinite topology on $\varprojlim_k \mathrm{Addr}_\tau / M_k$: an inverse limit of finite (discrete, hence compact Hausdorff) spaces is compact Hausdorff by Tychonoff. Hausdorff + compact + totally disconnected = Stone space. $\square$

7.8 The ontic ultrametric replaces Euclidean distance in τ

With completion, strong triangle, non-Archimedeanness, total disconnectedness, and profinite-topology match in place, the ontic ultrametric is the τ-native metric structure of the framework. We state the capstone meta-theorem and enumerate the specific replacements of classical Euclidean concepts.

Theorem 7.18 (Ontic ultrametric as τ-native metric **[τ-Effective]**). *The ontic ultrametric* $(\mathsf{Ultra}_\tau, d_\infty)$ *is the τ-native replacement for the Euclidean/Archimedean metric structure of classical analysis. Every classical analytic construction invoking a Euclidean metric has a canonical τ-lift in which the Euclidean metric is replaced by* d_∞ *on* Ultra_τ*:*

- (i) Convergence. *A sequence* $(u_n)_{n\in\mathbb{N}}$ *in* Ultra_τ *converges iff it is* tail-stabilising*: for every* $K \in \mathbb{N}$*, the K-th coordinate* $r^{(n)}_K$ *is eventually constant in* n*.*
- (ii) Limit points. *Fixed points of the ω-germ stabilisation procedure of Hinge 5 [16] land in* Ultra_τ *and are precisely the ultrametric limits of iterated tail-transformer images.*
- (iii) Open balls. *Open balls in* $(\mathsf{Ultra}_\tau, d_\infty)$ *are the clopen cylinders* $B_k(u) = \{v : v \equiv u \pmod{M_k}\}$ *— "addresses matching u modulo* M_k*".*
- (iv) Continuity. *A function* $f\colon \mathsf{Ultra}_\tau \to \mathsf{Ultra}_\tau$ *is continuous iff it is tail-level-preserving, which coincides with the τ-holomorphy condition of Hinge 5 [16] (tail-transformer in* HolEnd^σ*).*
- (v) Differentiation. *The split-complex Cauchy–Riemann operators of Hinge 5 [16] act on* $(\mathsf{Ultra}_\tau, d_\infty)$ *as the τ-native analog of the classical real-analytic* ∂_x, ∂_y*: they encode tail-level sensitivity in the* e_+, e_- *sectors of* $\mathbb{D}$*.*

In particular, no classical Euclidean input is needed to define any of these notions in the τ-framework; everything is built from the arithmetic primitive of primorial-ladder residues.

Proof. (i)–(iii) are restatements of Theorems 7.6–7.16 in the language of convergent sequences, fixed points, and cylinder bases respectively; the proofs have been given. (iv)–(v) follow from the identification of τ-holomorphic tail-transformers with primorial-coherent continuous self-maps of Ultra_τ (Hinge 5 [16], Theorems 5.4–5.8), pulled back along the isometric embedding $\iota\colon \mathrm{Addr}_\tau \hookrightarrow \mathsf{Ultra}_\tau$. The detailed Cauchy–Riemann identifications are recalled in the preliminary remarks of §2 (Remark 2.15), and carry over verbatim to the ultrametric setting because all primorial-ladder residue operations are isometries in d_∞. □

7.9 The ι_τ anchor in the ontic ultrametric

The master constant $\iota_\tau = 2/(\pi + e) \approx 0.341304$ (Hinge 3 [6]) was shown in §4.5 to have a specific canonical address in Addr_τ. That canonical address extends via ι to a distinguished point of Ultra_τ, and its σ-fixed endomorphism orbit gives the ultrametric scale calibration.

Proposition 7.19 (ι_τ as an ultrametric anchor **[τ-Effective]**)**.** *The canonical address $\mathsf{NF}(\iota_\tau) \in \mathrm{Addr}_\tau$ extends under the embedding $\iota\colon \mathrm{Addr}_\tau \hookrightarrow \mathsf{Ultra}_\tau$ to a distinguished element $u_{\iota_\tau} := \iota(\mathsf{NF}(\iota_\tau)) \in \mathsf{Ultra}_\tau$, and the orbit*

$$\mathrm{Orbit}(u_{\iota_\tau}) \;:=\; \{\varphi(u_{\iota_\tau}) \,:\, \varphi \in \mathrm{End}^\sigma_\tau(\mathsf{Ultra}_\tau)\}$$

under the σ-fixed endomorphism monoid of Ultra_τ (lifted from HolEnd^σ of Hinge 5 [16]) is a calibrating set for the ultrametric scale: for every $k \in \mathbb{N}$ there exists $\varphi_k \in \mathrm{End}^\sigma_\tau(\mathsf{Ultra}_\tau)$ with $d_\infty(u_{\iota_\tau}, \varphi_k(u_{\iota_\tau})) = 2^{-k}$.

Proof. Existence of the canonical address $\mathsf{NF}(\iota_\tau) \in \mathrm{Addr}_\tau$: §4.5. Extension to $u_{\iota_\tau} \in \mathsf{Ultra}_\tau$: the isometric embedding ι of Theorem 7.6. Calibrating property of the σ-fixed orbit: the endomorphisms of Ultra_τ lifted from HolEnd^σ in Hinge 5 act on primorial-ladder residues as ring homomorphisms of Addr_τ/M_k, so each orbit element is computable modulo M_k. By construction of HolEnd^σ (Hinge 5 [16], Definition 8.1), the idempotent factors $e_\pm$ of $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2-1)$ provide, for each $k \in \mathbb{N}$, a σ-fixed endomorphism φ_k that alters the residue at depth k and preserves it at depths $< k$; hence $k(u_{\iota_\tau}, \varphi_k(u_{\iota_\tau})) = k$ and $d_\infty(u_{\iota_\tau}, \varphi_k(u_{\iota_\tau})) = 2^{-k}$. □

Remark 7.20 (Calibration interpretation **[τ-Effective]**)**.** Proposition 7.19 is the τ-native replacement for the "unit distance" role played by $1 \in \mathbb{R}$ in Euclidean analysis: the distance $d_\infty(u_{\iota_\tau}, \varphi_k(u_{\iota_\tau})) = 2^{-k}$ is the calibrated scale-k increment in the framework. No external unit — neither metric metre nor Planck length — is needed to fix the scale; the calibration is determined intrinsically by the canonical address of ι_τ and the σ-fixed endomorphism monoid.

7.10 Scale: the primorial ladder

The primorial ladder $(M_k)_{k\in\mathbb{N}}$ supplies the τ-native scale hierarchy for the ontic ultrametric. Every depth $k \in \mathbb{N}$ corresponds to a primorial-level reduction, and balls of radius 2^{-k} are residue classes modulo M_k.

Proposition 7.21 (Ball–residue correspondence **[Established]**)**.** *For every $u \in \mathsf{Ultra}_\tau$ and every $k \in \mathbb{N}$,*

$$B_k(u) \;=\; \{v \in \mathsf{Ultra}_\tau \,:\, v \equiv u \pmod{M_k}\},$$
$$|\mathsf{Ultra}_\tau/B_k| \;=\; M_k \quad \textit{(depth-k residue ring).}$$

Hence the "tiles" of the ontic ultrametric topology at depth k are in bijection with $\mathbb{Z}/M_k\mathbb{Z}$ (via the canonical- addressability projection of §4).

Proof. $B_k(u) = \{v : d_\infty(u,v) \le 2^{-k}\} = \{v : u \equiv_k v\} = \{v : v \equiv u \pmod{M_k}\}$ by Definition 7.3 and Definition 7.5. The quotient Ultra_τ/B_k is the set of depth-k residue classes, which is $\mathrm{Addr}_\tau/M_k \cong \mathbb{Z}/M_k\mathbb{Z}$ by the canonical-addressability isomorphism of §4, with cardinality M_k. □

Remark 7.22 (Arithmetic vs. geometric scaling **[Established]**)**.** The classical Euclidean scale hierarchy is *geometric*: the dyadic radii $(2^{-k})_{k\in\mathbb{N}}$ give balls of Lebesgue measure $\asymp 2^{-kn}$ in $\mathbb{R}^n$, and the ratio of consecutive scale sizes is the geometric factor 2. In the ontic ultrametric, the scale ratio at depth $k \to k+1$ is the arithmetic factor $M_{k+1}/M_k = p_{k+1}$ (the $(k+1)$-th prime), which grows without bound as $k \to \infty$ — an *unbounded ultrametric ramification*. This is the quantitative signature of the arithmetic-not-geometric nature of τ-scale: the framework does not have a fixed doubling dimension, but a prime-by-prime ramified one.

Proposition 7.23 (Haar measure on Ultra_τ **[τ-Effective]**)**.** *The profinite group structure on Ultra_τ (inherited from the additive group structure of $\varprojlim_k \mathbb{Z}/M_k\mathbb{Z}$) admits a unique translation-invariant probability measure μ, and the ball of radius 2^{-k} has μ-measure $1/M_k$:*

$$\mu(B_k(u)) \;=\; \frac{1}{M_k} \qquad \textit{for all } u \in \mathsf{Ultra}_\tau,\ k \in \mathbb{N}.$$

Proof. Compact Hausdorff abelian groups carry a unique normalised Haar measure (existence and uniqueness standard; see Hinge 4 [15] for the application to the boundary algebra). The depth-k projection $\mathsf{Ultra}_\tau \to \mathbb{Z}/M_k\mathbb{Z}$ is a continuous surjective group homomorphism onto a finite group of order M_k, and the pushforward of μ is the uniform probability measure on $\mathbb{Z}/M_k\mathbb{Z}$ by uniqueness; hence each residue class (of which $B_k(u)$ is one, by Proposition 7.21) has measure $1/M_k$. □

Remark 7.24 (The Haar measure as τ-native Lebesgue measure **[τ-Effective]**)**.** The Haar measure

μ of Proposition 7.23 is the τ-native replacement for the Lebesgue measure on $\mathbb{R}$: it is the unique translation-invariant, probability-normalised measure on Ultra_τ that respects the ultrametric topology. Integration of tail-transformers, L^2-norms of τ-holomorphic sections, and spectral density computations throughout Books II–VII are to be understood as integrals against μ (or its natural extensions to sectors), not against Lebesgue. The details of the resulting measure-theoretic analysis are developed in Book III's τ-spectrum chapters [9].

Summary of §7

We have constructed the ontic ultrametric d_∞ as the depth-weighted completion of the Cayley word metric on the canonical address space, and established:

- **Completion** (Theorem 7.6): $(\mathsf{Ultra}_\tau, d_\infty)$ is the Cauchy completion of $(\mathrm{Addr}_\tau, d'_{\mathrm{Cay}})$ with Addr_τ dense in Ultra_τ.
- **Ultrametric property** (Theorem 7.9): strong triangle inequality holds in Ultra_τ.
- **Non-Archimedeanness** (Corollary 7.11): no classical Archimedean scale hierarchy exists.
- **Total disconnectedness** (Theorem 7.14): connected components are singletons; clopen balls form a basis.
- **Profinite-topology match** (Theorem 7.16): the ultrametric topology coincides with the Hinge-4 [15] profinite boundary topology on $\varprojlim_k \mathrm{Addr}_\tau / M_k$.
- **Replacement of Euclidean distance** (Theorem 7.18): d_∞ is the τ-native metric structure for convergence, limits, open balls, continuity, and differentiation.
- **ι_τ calibration** (Proposition 7.19): the canonical address of ι_τ anchors the ultrametric scale via its σ-fixed endomorphism orbit.
- **Primorial-ladder scale** (Proposition 7.21, 7.23): balls of radius 2^{-k} are residue classes mod M_k with Haar measure $1/M_k$.

Section 8 closes the paper by leveraging this ultrametric to state the address-resolution theorem: arithmetic equality in Category τ is a finite-depth ultrametric-ball membership test, not an equational search.

8. ADDRESS RESOLUTION AND THE ABSENCE OF EQUATIONS

8.1 What classical arithmetic does

To state cleanly what *address resolution* is, it helps to recall with precision what *equational calculation* is — and to see, site by site, that the two are not variants of one idea but different ways of writing down mathematics. Four sites make the classical picture explicit.

(a) Peano and Dedekind. Classical arithmetic is framed within first-order logic over a signature $(0, S, +, \cdot)$ whose axioms are equational or equationally conservative: $S(x) \neq 0, x+0 = x, x \cdot S(y) = x \cdot y + x$, plus the induction schema. Proofs are chains of equational rewrites interleaved with quantifier rules. The classical equality symbol $=$ is governed by reflexivity, symmetry, transitivity, and *substitution*:

$$\frac{t = s \qquad \varphi[t/x]}{\varphi[s/x]} \qquad \text{(Leibniz substitution)}.$$

Substitution is what lets equational identity propagate through every predicate of the language.

(b) Universal algebra. An *equational theory* T over a signature Σ is a set of equations $s \simeq t$ between Σ-terms closed under the rules of equational logic (reflexivity, symmetry, transitivity, congruence, and substitution of variables by terms). The variety $\mathbf{Alg}(\Sigma, T)$ of T-algebras is, by Birkhoff's HSP theorem, exactly the class of Σ-algebras closed under homomorphic images, subalgebras, and products. Every classical algebraic structure — groups, rings, lattices, Boolean algebras — is presentable this way.

(c) First-order logic over rings and fields. The theory of commutative rings is axiomatised by a finite list of equations; the theory of fields adds a single non-equational axiom $\forall x\,(x \neq 0 \to \exists y\, xy = 1)$ but is still essentially equational. Questions like "$a = b$?" for polynomials reduce to the word problem modulo the ideal, decidable by Buchberger's algorithm.

(d) Set theory as the background. Underlying all of the above is ZFC, in which every mathematical object is a free-standing element of some set and "$=$" is extensional identity. The axiom of extensionality is the backbone: sets *are* their members, and equality is the predicate that reports whether two names denote the same member-structure.

Remark 8.1 (What unifies (a)–(d))**.** In all four sites, mathematical objects are *free-standing elements of a set*, and equality is an *extensional binary relation* governed by substitution. Given two terms a and b, the question "$a = b$?" is: is there a chain of equational rewrites $a = a_1 = a_2 = \cdots = b$ licensed by the axioms? Arithmetic is *equation-solving*, and proof is *rewriting within a congruence closure*. The mathematical content is the set of objects; the equations are *about* that content.

This picture is so deeply naturalised that it is easy to forget any alternative is possible. The τ-kernel does not recognise this picture. It is the purpose of the present section to explain, with precision, what replaces it.

8.2 What τ-arithmetic does instead

In Category τ there are no free-standing mathematical objects. There are only:

- *Codes* $c \in \mathsf{Code}$: finite token strings in the τ-kernel alphabet, generated by the primitive transformer symbols of Hinge 5 [16].
- *Canonical addresses* $\mathrm{Canon}(c) = \mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$: the unique NF representative of c's $\sim$-class, obtained by iterating the τ-kernel rewriting system $\to_{\mathsf{NF}}$ to its terminus (§§3, 4).
- *The NF-resolution procedure* $\mathsf{Norm}\colon \mathsf{Code} \to \mathsf{Code}^{\mathsf{NF}}$: the finite-witness decidable algorithm that transforms a code to its canonical address (Theorem 1.1).

Remark 8.2 (Ontological point: codes are the content)**.** The codes *are* the mathematical content of Category τ — they do not "stand for" an underlying stratum of Platonic numbers. This is the structural inversion demanded by the diagonal discipline (DD1)–(DD4) of Hinge 5 [16, §6]: the τ-framework forbids the Cartesian-product / comprehension step that would produce a set-theoretic denotation. There is no set N of "τ-numbers" such that codes name elements of N. The pair $(\mathsf{Code}^{\mathsf{NF}}, \sim)$ is what there is.

A τ-native question of the form "$a = b$?" therefore cannot be a classical equational query, because there is no ambient set of "elements" for a and b to name. What the question *is* is made precise in the next subsection.

8.3 Reformulation of "=" as address-resolution

Definition 8.3 (τ-equality as address-resolution)**.** *Let $a, b \in \mathsf{Code}$ be admissible codes in the τ-kernel. The τ-*equality *between a and b is the proposition*

$$a \sim b \quad \overset{\text{def}}{\iff} \quad \mathsf{NF}(a) = \mathsf{NF}(b), \quad \text{as strings in } \mathsf{Code}^{\mathsf{NF}}.$$

The right-hand side is: compute $\mathsf{NF}(a)$ via the τ-kernel rewriting procedure Norm, compute $\mathsf{NF}(b)$ analogously, and compare the two NF codes position-by-position as finite token strings. The comparison succeeds iff both codes have the same length and identical tokens at every position.

Remark 8.4 (What the definition is *not*)**.** Definition 8.3 does not say "$\sim$ is defined in terms of equality of something else." That would require "something else" — free-standing elements of an ambient set — which the τ-kernel disallows. The definition is *self-standing*: the equality $\mathsf{NF}(a) = \mathsf{NF}(b)$ on the right is *string equality in the NF code alphabet*, which is decidable by character-by-character comparison and requires no semantic interpretation.

Remark 8.5 (Why this is not circular)**.** One might worry: we have defined $\sim$ using string equality, which is itself an "=" — are we not smuggling classical equality back in? The answer is that *finite string equality in a fixed finite alphabet* is a decidable syntactic matching procedure, not a proposition about extensional identity of denoted objects. Two NF codes are the same string iff the same sequence of tokens has been emitted; this is a finite-witness decidable computation. Classical equality is the predicate "x and y are the same element of the underlying set"; string equality is the predicate "two strings produced by the same algorithm are character-by-character identical." They live in different registers: one is semantic on Platonic objects, the other is syntactic on finite token sequences. The τ-kernel refuses the semantic denotation stratum (cf. diagonal discipline, (DD1)–(DD4)) that classical equality requires.

8.4 The address-resolution theorem

We now record the main decidability result: τ-equality is not just defined but *computable* in finite witness budget.

Theorem 8.6 (Address Resolution, **[τ-Effective]**)**.** *For admissible codes $a, b \in \mathsf{Code}$:*

(a) *The proposition "$a \sim b$" is* decidable *in the sense of constructive mathematics: there exists an algorithm $\mathsf{ResolveEq}\colon \mathsf{Code} \times \mathsf{Code} \to \{\top, \bot\}$ such that, for every input pair (a, b), $\mathsf{ResolveEq}(a, b) = \top$ iff $a \sim b$ and $\mathsf{ResolveEq}(a, b) = \bot$ iff $a \not\sim b$.*

(b) *The decision procedure is the NF-comparison algorithm:*

$$\begin{aligned} \mathsf{ResolveEq}(a, b) \ &:= \ \top \quad \text{if } \mathsf{Norm}(a) = \mathsf{Norm}(b), \\ &:= \ \bot \quad \text{otherwise.} \end{aligned}$$

Here equality means equality as finite strings. The two subroutines are (i) Norm, the NF-resolution algorithm of Theorem 1.1 (canonical normalisation), and (ii) position-by-position string equality of two NF codes in the fixed τ-kernel alphabet.

(c) *The total witness budget of $\mathsf{ResolveEq}(a, b)$ is bounded by*

$$W(a, b) \ \leq \ \max(k_0(a), k_0(b)) \ + \ |\mathsf{NF}(a)| + |\mathsf{NF}(b)|,$$

where $k_0(c)$ is the canonical-normalisation witness depth of c (Theorem 1.1) and $|\mathsf{NF}(c)|$ is the length of the NF code in tokens.

Proof. Part (a). From NF confluence (Theorem 1.2) and strong normalisation (Theorem 1.1), every $\sim$-class contains a *unique* NF representative. Hence $a \sim b$ iff $\mathsf{NF}(a)$ and $\mathsf{NF}(b)$ are the same NF code. The NF codes are elements of the decidable finite-alphabet language $\mathsf{Code}^{\mathsf{NF}}$ (§3), so string equality between two NF codes is decidable by a finite-time algorithm. Compose the two: input (a, b), apply Norm to both, compare the outputs. This yields the algorithm $\mathsf{ResolveEq}$.

Part (b). This is the definition of the algorithm, justified by (a).

Part (c). The cost of $\mathsf{Norm}(c)$ is bounded by the canonical pass budget $k_0(c)$ (Theorem 1.1, following Hinge 5's pass-budget decidability lemma [16]). Running Norm on both a and b in sequence costs at most $k_0(a) + k_0(b) \leq 2 \cdot \max(k_0(a), k_0(b))$ pass-budget steps; the string comparison then costs $O(|\mathsf{NF}(a)| + |\mathsf{NF}(b)|)$ token comparisons. Combining and dropping the constant yields the stated bound. Finite-width of the DAG (Theorem 1.3) ensures these quantities are themselves finite. □

Remark 8.7 (Decidability is intrinsic)**.** In classical first-order arithmetic, decidability is a *theorem one proves* (e.g. for the quantifier-free fragment), and undecidability is a *theorem one laments* (e.g. for the full language by Gödel–Rosser). In the τ-framework, the decidability of $\sim$ is *built into the definition*: Definition 8.3 reduces $\sim$ to NF-string equality, which is primitively decidable. There is no quantifier-alternation hierarchy to climb; every τ-native equality is Δ_0 in the witness-budget hierarchy.

Remark 8.8 (Contrast with classical equality)**.** In a classical ring R, the question "$a = b$?" is the word-problem question in R/I for some ideal I, and its decidability depends on the presentation. For some R's the word problem is undecidable (Novikov–Boone for groups; the analogous failure for semigroups is classical). In Category τ, the analogue is trivially decidable, because there is no "ambient ring"; there are only codes and their canonical NFs. The computational content is all carried by Norm, and Norm is total, deterministic, and finitely-budgeted by Hinge 5 plus Theorem 1.1.

8.5 Why there are no equations

We can now state the negative result that gives the paper its title.

Theorem 8.9 (No Equations in Category τ, [τ-Effective])**.** *Category τ does not admit a classical equational theory in the sense of universal algebra. More precisely: there is no pair (Σ, T) consisting of a first-order signature Σ and a set T of equational axioms over Σ, together with an interpretation $I\colon \mathsf{Code}^{\mathsf{NF}} \to \mathbf{Alg}(\Sigma, T)$, such that all three of:*

(i) Substitution closure*: the equational theory T is closed under the rules of equational logic — reflexivity, symmetry, transitivity, congruence, and substitution of variables by terms.*

(ii) Soundness*: for codes $a, b \in \mathsf{Code}$, if $T \vdash I(\mathsf{NF}(a)) \simeq I(\mathsf{NF}(b))$ (equational derivation in T), then $a \sim b$.*

(iii) Completeness*: for codes $a, b \in \mathsf{Code}$, if $a \sim b$, then $T \vdash I(\mathsf{NF}(a)) \simeq I(\mathsf{NF}(b))$.*

hold simultaneously.

Proof. We show the three conditions are jointly incompatible with the diagonal discipline (DD1)–(DD4) of Hinge 5 [16, Definition 6.1].

Step 1 (existence of denotation). Suppose toward contradiction that such (Σ, T, I) exists. The interpretation I sends each NF code $\mathsf{NF}(c) \in \mathsf{Code}^{\mathsf{NF}}$ to an element of some Σ-algebra $A \in \mathbf{Alg}(\Sigma, T)$. Being an element of a Σ-algebra means being a member of the underlying carrier set $|A|$; i.e. $I(\mathsf{NF}(c)) \in |A|$.

Step 2 (free-standing objects forced). By construction, $|A|$ is a *set* — an unordered collection of free-standing elements satisfying the axioms of ZFC. In particular, $|A|$ admits the Cartesian product construction: $|A| \times |A|$ is itself a set, and the diagonal $\Delta_A : |A| \to |A| \times |A|$, $\Delta_A(x) = (x, x)$, is a set-theoretic function.

Step 3 (comprehension-cut graph). For every Σ-operation $\sigma\colon |A|^n \to |A|$ the graph $\Gamma_\sigma = \{(\vec{x}, \sigma(\vec{x})) : \vec{x} \in |A|^n\} \subseteq |A|^n \times |A|$ is a subset of the Cartesian product, cut out by comprehension on the graph predicate. This is the set-theoretic graph-of-a-function construction of Hinge 5 [16, §6.1].

Step 4 (DD-violation). The diagonal discipline (DD1) forbids the τ-framework from admitting $|A| \times |A|$ as an admissible carrier; (DD2) forbids graph-as-slice representations of transformers; (DD3) forbids exponential objects; (DD4) forbids free contraction in the meta-logic. The interpretation I of Step 1, combined with Steps 2–3, manifestly violates all four: the Σ-algebra A uses primitive Cartesian products, graphs-as-slices, and implicit contraction throughout its signature. The τ-kernel coherence predicate ([16, §5.3]; cf. §2 above) fails at I, contradicting the existence assumption.

Step 5 (no extensional retrieval). One might attempt to salvage (Σ, T, I) by restricting to a classical subtheory that avoids products and graphs; but then (i)–(iii) cannot jointly hold. Without products, the congruence rule of equational logic fails, so (i) is incompatible. Without graphs, the interpretation I cannot specify the action of Σ-operations on $\mathsf{Code}^{\mathsf{NF}}$-classes, so (iii) fails. Hence no three of the conditions survive together. □

Remark 8.10 (What Theorem 8.9 does and does not claim). The theorem does *not* say that nothing resembling an equation can be written down inside the τ-framework. It says that the τ-native $\sim$-relation does not fit into the *classical equational theory* template of universal algebra, because the τ-framework lacks the set-theoretic carrier structure that the template assumes. Inside τ, "equations" are shorthand for NF-resolution statements (see §8.6 below for the precise sense); they are not primitive objects of the theory.

Remark 8.11 (Intensional vs extensional). The deep reason is the intensional character of $\sim$. Classical equality is *extensional*: it reports whether two names denote the same entity. τ-equality is *intensional*: it reports whether two codes resolve to the same canonical NF after running the τ-kernel normalisation procedure. The two differ whenever the normalisation procedure is non-trivial — which, by Theorem 1.1, is essentially always. Tail-coherence certificates ([16, §4]) are what $\sim$ depends on; they are not recoverable from set-theoretic extensions. The extensional shadow is a strict quotient, and the shadow-functor Shad of §8.6 witnesses exactly how information is lost.

8.6 Shadow functor to classical equational algebra

Despite Theorem 8.9, τ-statements can be *translated* into classical equational statements by a well-defined functor. This is how results of the τ-framework reach classical audiences.

Remark 8.12 (The Shadow Functor, **[τ-Effective]**). There exists a faithful (but not full) functor

$$\mathrm{Shad} : \mathbf{Cat}_\tau \longrightarrow \mathbf{Eq\text{-}Alg}$$

from Category τ (the topos of Hinge 6 [18]) to the category of classical equational algebras (the classical topos-theoretic framework documented in [1, 26, 22, 23]; for the constructive and type-theoretic alternatives see [3, 30, 31]), with the following structure.

- *On objects*: each τ-object X (a presheaf on the τ-site) is sent to the classical Σ-algebra $\mathrm{Shad}(X)$ whose carrier is the set of $\sim$-equivalence classes of admissible codes of type X, and whose operations are induced by τ-morphisms modulo $\sim$. The algebra's signature Σ records the finitary operations definable inside the τ-kernel at the level of NF codes.
- *On morphisms*: a τ-morphism $f\colon X \to Y$ is sent to the classical function $\mathrm{Shad}(f)\colon \mathrm{Shad}(X) \to \mathrm{Shad}(Y)$ defined on $\sim$-classes via $[c] \mapsto [\mathsf{NF}(f \cdot c)]$, where $f \cdot c$ is the action of the transformer code on the input code.
- *Faithfulness*: if $f_1, f_2\colon X \to Y$ are τ-morphisms with $\mathrm{Shad}(f_1) = \mathrm{Shad}(f_2)$, then $f_1 \sim f_2$. (Two transformers that induce the same classical function on $\sim$-classes are already $\sim$-equivalent as τ-morphisms.)
- *Non-fullness*: there exist classical functions $g\colon \mathrm{Shad}(X) \to \mathrm{Shad}(Y)$ that do not arise as $\mathrm{Shad}(f)$ for any τ-morphism f. These are the "non-constructive" functions of classical mathematics — e.g. choice-function-dependent, or graph-predicate-defined — which the τ-kernel refuses (cf. DD1–DD4).

Remark 8.13 (How to read a "τ-equation"). Throughout Books II–VII [8, 9, 10, 11, 12, 13], statements of the form "the τ-equation $a = b$ holds" are systematically shorthand for:

$$\mathrm{Shad}(a) \simeq \mathrm{Shad}(b) \quad \text{in } \mathbf{Eq\text{-}Alg},$$

which in turn unfolds to: "$\mathsf{NF}(a) = \mathsf{NF}(b)$ as NF codes, i.e. $a \sim b$." The classical-facing statement is the shadow-image; the τ-native fact is the address-resolution equality. Hinge 7 establishes that the two are linked by the faithful functor Shad, so translation is lossless in the $\tau \to$ Classical direction, though not in the reverse direction.

Remark 8.14 (The shadow is where classical mathematics lives). A useful picture: classical mathematics is the image of Shad. ZFC-era arithmetic, analysis, and algebra live in **Eq-Alg** (or a suitably enlarged target). The τ-framework lives upstream: $\mathbf{Cat}_\tau$ is where address-resolution is primitive; the extensional view is recovered as the functorial image. Books II–VII work in both registers: derivations happen in $\mathbf{Cat}_\tau$, classical consequences are statements under Shad.

8.7 Consequences for the Millennium Problems

The address-resolution paradigm reframes the τ-formulation of several classical problems. We sketch three, each anchored at **[Conjectural]** since the concrete reformulations are developed in Books II–V and the spectrum papers.

Navier–Stokes smoothness ***[Conjectural]***. In classical form, Navier–Stokes smoothness asks whether every smooth solenoidal initial datum on $\mathbb{R}^3$ admits a global-in-time smooth solution of the Navier–Stokes system, or whether finite-time singularities can form. In the τ-formulation of Book III [9], the Navier–Stokes operator acts on the ontic ultrametric space $(\mathsf{Ultra}_\tau, d_\infty)$ introduced in §7, and the evolution is an address-resolution specification on admissible codes, not a PDE on Euclidean space. The "smoothness / singularity" dichotomy becomes: does the address-resolution procedure terminate with finite width at every time, or does the canonical-address width diverge at a finite witness depth? The answer is framed in terms of DAG finite-width (Theorem 1.3) and the ontic-ultrametric completion Ultra_τ. This is *not* a reformulation of the classical problem in the classical category; it is a τ-native replacement anchored at the shadow image.

Riemann Hypothesis **[Conjectural]**. The classical zeta function $\zeta(s) = \sum n^{-s}$ is a meromorphic function on $\mathbb{C}$, and RH asserts that its non-trivial zeros have real part $\frac{1}{2}$. In the τ-formulation (sketched in Book III [9] and pursued in the ι_τ-hinge [6]), the τ-native zeta is an *address-valued* function: $\zeta_\tau \colon \mathrm{Addr}_\tau \to \mathrm{Addr}_\tau$, whose zeros are specific canonical addresses in $\mathsf{Code}^{\mathsf{NF}}$. The τ-native RH becomes: *do the non-trivial zeros of ζ_τ all resolve to the canonical address of the critical line?* This is an address-coincidence question, not a continuous-parameter question; its decidability reduces to NF comparison at specific depths in the DAG.

P versus NP **[Conjectural]**. The classical P versus NP question asks whether every polynomial-time verifiable decision problem admits a polynomial-time decision procedure. In the τ-framework, decidability is primitively *finite-witness decidable* (Remark 8.7); the classical P/NP dichotomy is reframed as a witness-budget-asymptotic question: *how does $W(a, b)$ of Theorem 8.6(c) scale with the input size*? The τ-native P versus NP becomes a question about the scaling exponent of the pass-budget function $k_0(\cdot)$ on specific address families. We cite Book IV [10] for the detailed reformulation and the proposed resolution at **[Conjectural]** scope.

Remark 8.15 (These are motivating sketches, not proofs)**.** The three items above are *not* proofs of Millennium-type claims. They are illustrative: the address-resolution paradigm reframes classical problems in a way that changes what counts as the "answer." Full proofs belong to Books II–V and their dedicated papers. The present paper establishes only the foundational vocabulary — codes, NFs, DAGs, the ontic ultrametric, address-resolution — and its logical structure. The "application sketches" are motivating remarks tagged **[Conjectural]**, not contributions of Hinge 7.

8.8 Summary: arithmetic is address-resolution

The structural thesis of Hinge 7, and the interpretive thesis of the entire Panta Rhei bundle, can now be stated without hedging.

> *Arithmetic in Category τ is address-resolution, not equational calculation.* Every question of the form "$a = b$?" for admissible codes $a, b \in \mathsf{Code}$ is reformulated as the finite-witness decidable procedure $\mathsf{ResolveEq}(a, b)$ (Theorem 8.6), which reduces to NF comparison on the canonical-address space $\mathrm{Addr}_\tau = \mathsf{Code}^{\mathsf{NF}}/\sim$. There is no classical equational theory of which τ-arithmetic is a model (Theorem 8.9); the extensional equational view is recovered only as the *shadow* Shad of the address-resolution procedure (Remark 8.12). The seven hinges together establish the whole framework as a constructive, finite-witness, address-theoretic alternative to classical set-theoretic mathematics.

The seven hinges build the theory incrementally:

- **H1.** *Hyperfactorization* [5] — unique tower-atom decomposition provides the "digits" of canonical addresses.
- **H2.** *Prime Polarity* [14] — the Legendre $(2/p)$ mod 8 characterisation supplies the B/C polarity labelling.
- **H3.** *Master Constant ι_τ* [6] — the anchor $\iota_\tau = 2/(\pi + e) \approx 0.341304$ calibrates the addressing scale.
- **H4.** *Boundary Algebra $\mathbb{D}$* [15] — the split-complex algebra $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2 - 1)$ is the algebraic home of addresses at the boundary.
- **H5.** *τ-Holomorphy* [16] — the earned categorical machine makes addresses categorical morphisms via pre-Yoneda collapse.
- **H6.** *τ-Topos* [18] — the topos $\mathbf{Cat}_\tau$ with $\Omega_\tau = B_\sigma(\mathbb{D})$ makes addresses internal logic truth values.
- **H7.** *Address Resolution* (this paper) — addresses are canonical NFs with decidable resolution procedure.

From codes (H1) through polarity (H2), anchor calibration (H3), algebraic home (H4), categorical status (H5), and internal-logic status (H6), to the canonical-NF resolution procedure that is arithmetic itself (H7) — the Panta Rhei foundational arc is the sustained argument that the primitive object of mathematics is an *address*, and the primitive operation is *resolve*.

Remark 8.16 (Closure of the "modulo Hinge 7" caveat)**.** Hinges 5 and 6 carried scope caveats "modulo Hinge 7 NF confluence" in their main theorems. With Theorem 1.2 established (NF confluence Church–Rosser) and Theorem 8.6 proved here, those caveats are now discharged. The Panta Rhei seven-hinge bundle is closed: the foundational arc has no remaining "forthcoming" dependencies among its own papers. Forward links into Books I–VII carry the program forward, but the bundle itself is self-contained.

Remark 8.17 (Why this reframing is worth doing)**.** It would misread the address-resolution paradigm as a reformulation of classical arithmetic that happens to give the same answers in a different vocabulary. The paradigm shift matters *because the questions change*. Classical arithmetic asks: does this equation hold? τ-arithmetic asks: what is the canonical address of this code? The first question presupposes an ambient set of objects; the second does not. The first is equational search; the second is algorithmic resolution. Where the questions coincide (under Shad), the answers coincide; in the non-classical regime addressed by Books II–VII, the τ-native questions are the right ones, and the classical equational theory is a convenient but strictly weaker shadow.

This is the conceptual climax of Hinge 7, and the completion of the promise in the paper's title: *arithmetic in Category τ is address-resolution, not equational calculation*. The theorem is Theorem 8.6, the obstruction is Theorem 8.9, the translation bridge to classical mathematics is the shadow functor of Remark 8.12, and the seven-hinge bundle is the cumulative construction that makes all three precise.

9. LEAN ROADMAP AND REGISTRY ENTRIES

9.1 Lean roadmap (detailed)

The planned Lean 4 formalisation lives in `TauLib.BookI.Addressability` per the PR-I / Hinge-1 alignment; module layout is as above. Each main theorem has an entry point in a dedicated module and a canonical registry ID matching Book I's counting.

9.2 Registry IDs

Remark 9.1 (Registry IDs **[τ-Effective]**). The seven main theorems of this paper will be registered in `registry/book1_registry.jsonl` upon stabilisation of the v1 draft, with IDs I.T-final-1 (canonical normalisation), I.T-final-2 (NF confluence — the Church–Rosser capstone), I.T-final-3 (genealogical DAG), I.T-final-4 (Cayley word metric), I.T-final-5 (ontic ultrametric), I.T-final-6 (address-resolution), I.T-final-7 (Hinge-7 integration). Concrete registry numbers are assigned on the book1 next-available range at registration time. Registration follows peer-panel certification.

10. CONCLUSION AND FORWARD LINKS

With Hinge 7 in place, the Panta Rhei seven-hinge foundational bundle is complete. The framework now has a closed, self-contained collection of standalone peer-reviewable papers:

- **Hinges 1–4** supply the coarse-grained coordinate structure: tower-atom decomposition, prime polarity, master constant ι_τ, split-complex boundary algebra.
- **Hinge 5** installs τ-holomorphy as the ontological primary; the earned categorical machine.
- **Hinge 6** constructs the τ-topos with four-valued paraconsistent internal logic.
- **Hinge 7** (this paper) closes the arc: canonical- address NF confluence, the genealogical DAG, the ontic ultrametric, and the address-resolution paradigm.

Forward links into the book series:

- **Book I** [7] — the categorical foundations layer: Category τ's 7 axioms, 5 generators, 1 operator.
- **Book II** [8] — categorical holomorphy: Hartogs continuation, interior-point construction, the Central Theorem $\mathcal{O}(\tau^3) \cong A_{\text{spec}}(\mathbb{L})$.
- **Book III** [9] — categorical spectrum: τ-Navier–Stokes, the Millennium Problems in τ-native formulation.
- **Books IV–VII** [10, 11, 12, 13] — physical, biological, and metaphysical applications.

The ontic ultrametric established here is the τ-native metric structure used throughout Books II–VII: wherever classical analysis would invoke a Euclidean or Riemannian metric, the τ-framework uses d_∞ on the boundary address space. This is the concrete sense in which the seven-hinge foundation *completes* the pre-categorical preparation of the book series: every metric, topology, and analytic construction in the books is now grounded in an address-resolution primitive.

Acknowledgments

The authors thank the broader Panta Rhei research community for feedback on earlier drafts, the Lean 4 and `mathlib` communities for the formal-logic infrastructure that underwrites this paper's planned mechanisation, and the coordinated development of Hinges 1–6 on which this capstone rests.

Data and code availability

The source repository for the paper bundle is at `https://panta-rhei.site/papers/address-resolution`. Planned Lean 4 artefacts for the main theorems will appear in `TauLib.BookI.Addressability` (see §9).

REFERENCES

[1] Michael Artin, Alexander Grothendieck, and Jean-Louis Verdier, *Théorie des topos et cohomologie étale des schémas (sga 4), tome 1*, Lecture Notes in Mathematics, vol. 269, Springer, Berlin, 1972.

[2] Nuel D. Belnap, *A useful four-valued logic*, Reidel, Dordrecht, 1977.

[3] Errett Bishop and Douglas Bridges, *Constructive analysis*, Grundlehren der mathematischen Wissenschaften, vol. 279, Springer, Berlin, 1985.

[4] Francesco Catoni, Dino Boccaletti, Roberto Cannata, Vincenzo Catoni, Enrico Nichelatti, and Paolo Zampetti, *The mathematics of minkowski space-time*, Frontiers in Mathematics, Birkhäuser, Basel, 2008.

[5] Thorsten Fuchs and Anna-Sophie Fuchs, *The hyperfactorization theorem: Unique tower-atom decomposition in ZFC and category τ*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/hyperfactorization. Hinge 1 of the Panta Rhei bundle.

[6] ———, *The master constant ι_τ: A structural derivation of $\iota_\tau = 2/(\pi + e)$*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/iota-tau. Hinge 3 of the Panta Rhei bundle.

[7] ———, *Panta rhei, Book I: Categorical Foundations*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1987-6.

[8] ———, *Panta rhei, Book II: Categorical Holomorphy*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1988-3.

[9] ———, *Panta rhei, Book III: Categorical Spectrum*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1989-0.

[10] ———, *Panta rhei, Book IV: Categorical Microcosm*, 2nd ed., Panta Rhei Research, 2026, Volume IV of the 7-book Panta Rhei series; ISBN forthcoming.

[11] ———, *Panta rhei, Book V: Categorical Macrocosm*, 2nd ed., Panta Rhei Research, 2026, Volume V of the 7-book Panta Rhei series; ISBN forthcoming.

[12] ———, *Panta rhei, Book VI: Categorical Life*, 2nd ed., Panta Rhei Research, 2026, Volume VI of the 7-book Panta Rhei series; ISBN forthcoming.

[13] ———, *Panta rhei, Book VII: Categorical Metaphysics*, 2nd ed., Panta Rhei Research, 2026, Volume VII of the 7-book Panta Rhei series; ISBN forthcoming.

[14] ———, *The prime polarity theorem: A ZFC proof and the τ-framework isomorphism*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/prime-polarity. Hinge 2 of the Panta Rhei bundle.

[15] ———, *The split-complex boundary algebra $\mathbb{D}$: Canonical uniqueness, countable profinite construction, and the four-atom generator dictionary*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/boundary-algebra. Hinge 4 of the Panta Rhei bundle.

[16] ———, *τ-Holomorphy on the boundary algebra: ω-Germ transformers, the wave-equation Cauchy–Riemann, and the earned categorical machine*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/holomorphy-first. Hinge 5 of the Panta Rhei bundle.

[17] ———, *The τ-kernel as foundational architecture: Ontic identity, linear discipline, and the star-autonomous path*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/kernel-foundation. Hinge 8 (foundational anchor) of the Panta Rhei bundle; may be read as an entry point (Hinge 0) or as a capstone.

[18] ———, *The τ-topos and its four-valued internal logic: Resolving paraconsistent semantic circularity via split-complex idempotents and ω-germ stabilization*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/tau-topos. Hinge 6 of the Panta Rhei bundle.

[19] ———, *TauLib: A Lean 4 library for the Panta Rhei research program*, 2026, Available at https://panta-rhei.site/verify/.

[20] Robert C. Gunning and Hugo Rossi, *Analytic functions of several complex variables*, Prentice-Hall, Englewood Cliffs, NJ, 1965.

[21] Lars Hörmander, *An introduction to complex analysis in several variables*, 3 ed., North-Holland, Amsterdam, 1990.

[22] Peter T. Johnstone, *Topos theory*, Academic Press, London, 1977.

[23] ———, *Sketches of an elephant: A topos theory compendium, vol. 1*, Oxford University Press, Oxford, 2002.

[24] Steven G. Krantz, *Function theory of several complex variables*, 2 ed., AMS Chelsea Publishing, American Mathematical Society, Providence, RI, 1992.

[25] Lean Community, *The Lean 4 theorem prover*, Software and documentation, 2024, https://leanprover.github.io/.

[26] Saunders Mac Lane and Ieke Moerdijk, *Sheaves in geometry and logic: A first introduction to topos theory*, Universitext, Springer, New York, 1992.

[27] Graham Priest, *An introduction to non-classical logic: From if to is*, 2 ed., Cambridge University Press, Cambridge, 2008.

[28] R. Michael Range, *Holomorphic functions and integral representations in several complex variables*, Graduate Texts in Mathematics, vol. 108, Springer, New York, 1986.

[29] Garret Sobczyk, *The hyperbolic number plane*, The College Mathematics Journal **26** (1995), no. 4, 268–280.

[30] A. S. Troelstra and Dirk van Dalen, *Constructivism in mathematics: An introduction, vol. i*, North-Holland, Amsterdam, 1988.

[31] Univalent Foundations Program, *Homotopy type theory: Univalent foundations of mathematics*, Institute for Advanced Study, Princeton, NJ, 2013, Available online at https://homotopytypetheory.org/book/.

[32] I. M. Yaglom, *A simple non-euclidean geometry and its physical basis*, Springer, New York, 1979.