Structural Route Closure in the UNNS Substrate

From Finite Obstruction to Transfinite Refinement

Project: UNNS_COMMON_REFINEMENT · extension UNNS_ADMISSIBILITY_PERCOLATION

Corpus: 500 verified positive-integer cases · 793 rank-one generator families
7 higher-rank affine systems · 30-system admissibility–percolation pilot corpus (v2.5.1 audited)

Formal verification: Lean 4.31.0 · kernel build 837/837, cached replay 829/829

Status: Foundational theory manuscript · Date: 2026

Abstract

We study the transition from an equal-endpoint identity $ab = cd$ to the strictly stronger requirement of a common four-factor refinement $a = ef$, $b = gh$, $c = eg$, $d = fh$, across the successive cancellative commutative regimes considered in the UNNS Substrate research program: the positive integers, rank-one numerical monoids, higher-rank positive affine monoids, and the Conway/omnific (surreal) integers. We establish the **Structural Route-Closure Law**: primal factor traceability is equivalent to global four-factor route closure. This equivalence is classical algebra; the contribution here is its identification as a single cross-regime law, complete proofs in each finite regime, and three theorem-level results classifying how it succeeds or fails. The positive-integer case is recorded as a canonical, unconditionally successful multiplicative control, built from the classical gcd/coprime witness. **Principal Theorem I** gives the exact additive rank-one criterion $\mathrm{RCI}(H) = 1$, sharpened by an explicit constructive failure witness when $\mathrm{RCI}(H) > 1$ and completed by a proved *Cofinal Primal-Tail Theorem* showing every rank-one failure is confined to a finite obstruction core. **Principal Theorem II** gives the corresponding higher-rank affine criterion $\mathrm{ARD}(H) = 0$, and shows — with exact counterexamples — that lattice saturation, sufficient at rank one, is no longer sufficient once rank exceeds one. **Principal Theorem III**, the Persistent Defect-Ray Theorem, converts this insufficiency into an explicit geometry of obstruction: whenever $\mathrm{rank}(H) \geq 2$ and $\mathrm{ARD}(H) > 0$, there exist $a \in H$ and $0 \neq \rho \in H$ such that $a + n\rho$ is non-primal for every $n \geq 0$, in sharp contrast to the finite obstruction core guaranteed at rank one. The finite results are computationally validated under their stated protocols, with zero observed exceptions: exhaustive search over the stated rank-one generator range, and fixed certificate/regression testing in the affine regimes. We then examine whether a comparable persistent channel can survive in the transfinite (omnific) setting, present the audited structural chain of the candidate Conway/omnific route-closure mechanism, and give its extracted induction proposition an independent, Lean-kernel-verified formalization, explicitly scoped short of a full reconstruction of the Hahn-series/surreal machinery. Finally, using a 30-system admissibility–percolation pilot corpus — re-audited from instrument version 2.5.0 (25 `FULL_PERCOLATION`, 4 `HARD_FRAGMENTATION`, 1 `TAIL_FRAGMENTATION`) to an exact-probe version 2.5.1 that corrects exactly these 5 verdicts to `FULL_PERCOLATION` — we show that route closure, perturbative admissibility (A_κ), and percolative connectivity scale ($\kappa_{\text{connect\_exact}}$) are distinct, non-equivalent structural coordinates: admissibility is saturated at $A_\kappa = 1$ across all 30 systems regardless of traceability label, while the rank-one free/non-free contrast produces an exact connectivity-scale signature that does not extend to a monotone ordering at higher affine rank. The manuscript's backbone is LAW $\to$ THEOREMS $\to$ FAILURE GEOMETRY $\to$ TRANSFINITE MECHANISM $\to$ FORMAL VERIFICATION $\to$ STRUCTURAL EXTENSION.

Contents

Main Results

Structural Route-Closure Law. In the cancellative commutative setting considered here,

$$\text{primal factor traceability} \iff \text{global four-factor route closure.}$$

A note on numbering. This manuscript labels its three classifying results **Principal Theorem I, II, III** throughout the prose and section headings below. These are distinct from the sequential `theorem` environment numbers (Theorem 1–4) used at their point of statement and proof, since a fourth, independently numbered result — the Cofinal Primal-Tail Theorem (Theorem 2) — is interposed between Principal Theorems I and II in that sequence. Concretely: Principal Theorem I is Theorem 1, Principal Theorem II is Theorem 3, and Principal Theorem III is Theorem 4.

Positive-integer control. In $(\mathbb{Z}_{>0}, \times)$, the canonical gcd/coprime witness $e = \gcd(a, c)$, $f = a/e$, $g = c/e$, $h = b/g$ refines every equal-endpoint identity $ab = cd$ unconditionally.

Principal Theorem I (Rank-One Common Refinement Criterion). For nonzero additive $H \subseteq \mathbb{N}_0$ with $m = \min(H \setminus \{0\})$ and $\gamma = \gcd(H)$: H has global refinement $\iff m = \gamma \iff H = \gamma\mathbb{N}_0$; equivalently $\mathrm{RCI}(H) := m/\gamma = 1$.

Principal Theorem II (Affine Route-Closure Criterion). For a positive affine monoid H: global refinement $\iff$ every atom of H is prime $\iff H \cong \mathbb{N}_0^r \iff \mathrm{ARD}(H) := |\mathrm{Atoms}(H)| - \mathrm{rank}(\mathrm{gp}(H)) = 0$.

Principal Theorem III (Persistent Defect-Ray Theorem). If $\mathrm{rank}(H) \geq 2$ and $\mathrm{ARD}(H) > 0$, there exist $a \in H$ and $0 \neq \rho \in H$ such that $a + n\rho$ is non-primal for every $n \geq 0$.

Transfinite mechanism. A primal base, together with an exact locally routed factor block whose complementary residual strictly decreases a well-founded structural rank, forces every element to be primal and hence forces global four-factor refinement (Route-Closure Induction Proposition, Section 9); the extracted mechanism is independently Lean-kernel verified (Section 10).

This box states the manuscript's load-bearing results before their derivation, so that what follows can be read as detailed development of an already-stated architecture rather than as an accumulation of separate claims.

1 Introduction — From Equal Endpoints to Common Structural Routes

The starting point is an elementary observation about equal products. If

$$ab = cd \tag{1}$$

in a commutative monoid, the two factorizations $a \cdot b$ and $c \cdot d$ reach the same endpoint, but nothing in (1) says the routes that produced that endpoint are related. A strictly stronger requirement is that they share a *common four-factor refinement*: elements e, f, g, h with

$$a = ef, \qquad b = gh, \qquad c = eg, \qquad d = fh. \tag{2}$$

When such e, f, g, h exist, the equality (1) is not merely an endpoint coincidence; it is the visible trace of a single underlying routing of factors, and (1) is a two-dimensional slice of the refinement square built from e, f, g, h. The UNNS Substrate Research Program's `UNNS_COMMON_REFINEMENT` project asks, across successively larger algebraic settings, exactly when (1) upgrades to (2) for *every* equal-endpoint pair, and what happens structurally when it does not.

This manuscript reports four levels of result, matching the levels at which the project's own internal records distinguish its findings:

1. a single **cross-regime law** (Section 2), identifying the exact algebraic state that is equivalent to global refinement in every regime considered;
2. a **canonical positive-integer control** (Section 3), in which the classical multiplicative gcd/coprime witness refines every equal-endpoint identity unconditionally;
3. three **finite-regime theorems** (Sections 4–6) that classify exactly when the law holds or fails for rank-one numerical monoids and higher-rank positive affine monoids, together with the resulting **failure geometry** (Section 7);
4. a **transfinite realization and verification program** (Sections 8–10) asking whether a persistent finite-rank obstruction channel can survive into the Conway/omnific (surreal) integers, and to what extent the candidate mechanism that prevents this is independently checkable.

A further, later section (Section 11) brings in a structural extension, `UNNS_ADMISSIBILITY_PERCOLATION`, which asks a different question of the same corpus: does losing route closure carry any consequence for a system's perturbative admissibility or its percolative connectivity? The answer — that these are distinct, non-equivalent structural coordinates, not that they are statistically independent — is presented as a corollary to the structural interpretation established by the law and the three theorems, not as a fourth theorem of factorization.

We are precise, throughout, about a distinction the project's own internal synthesis insists on: the underlying primality/refinement equivalence stated as the Structural Route-Closure Law is *classical algebra*. Refinement ("interpolation", or the 2×2 Riesz decomposition property) and its equivalence with every-element primality have been studied, in essentially the generality used here, in the theory of refinement monoids and partially ordered abelian groups with interpolation [5], in the general theory of commutative semigroups [6], in the closely related theory of Schreier and pre-Schreier domains and rings originating with Cohn [7, 8], and specialized to numerical semigroups [9], to affine (finitely generated, normal or non-normal) semigroups [10], and to general factorization theory in commutative monoids [11]. The contribution of the present work is not a new algebraic theorem at this level; it is the identification of this equivalence as a single structural law that organizes four otherwise disconnected regimes, the complete proofs (not sketches) of the three exact finite-regime theorems that classify its failure, the resulting geometry of obstruction, and an independently verified transfinite repair mechanism. Well-founded structural descent, in particular, is a *mechanism* used to prove the law holds in the omnific candidate realization; it is not itself the common law, a distinction we return to explicitly in Section 8 and restate in the conclusion.

The manuscript is organized as follows. Section 2 gives formal definitions, states the Structural Route-Closure Law, and gives a first cross-regime overview figure. Section 3 records the positive-integer control. Sections 4–6 state and prove Principal Theorems I–III in full, with their computational validation. Section 7 synthesizes Principal Theorems I–III geometrically. Section 8 introduces the Conway/omnific case and its audited structural chain. Section 9 states the extracted Route-Closure Induction Proposition and its Lean formalization. Section 10 separates the mathematical theorem sections from the independent verification of the candidate transfinite mechanism. Section 11 presents the admissibility–percolation structural extension, including the instrument audit from v2.5.0 to v2.5.1. Section 12 states limitations and open problems. Section 13 concludes by restating the full hierarchy. Section 14 records data, code, analytics, and reproducibility information.

2 Definitions and the Structural Route-Closure Law

Throughout, monoids are commutative and cancellative unless stated otherwise.

Definition 1 (Global four-factor refinement)**.** A commutative cancellative monoid M has *global four-factor refinement* if for every $a, b, c, d \in M$ with $ab = cd$ there exist $e, f, g, h \in M$ satisfying (2).

Definition 2 (Primal element; primal factor traceability)**.** An element $p \in M$ is *primal* if whenever $p \mid ab$ there exist $p_1, p_2 \in M$ with $p = p_1 p_2$, $p_1 \mid a$, $p_2 \mid b$. M has *primal factor traceability* (equivalently, *every-element primality*, or the *pre-Schreier* property [7]) if every element of M is primal.

Definition 3 (Route closure; route defect)**.** M has *route closure* at an equal-endpoint pair ($ab = cd$) if that pair admits a refinement (2); M has *global route closure* if it has route closure at every such pair. A *route defect* is an equal-endpoint pair with no refinement.

Figure 1 makes Definitions 1–3 visually concrete before the theorem machinery below states exactly when they hold: an equal-endpoint identity $ab = cd$ is a two-dimensional shadow of a hidden four-factor square, and the manuscript's recurring question is exactly when that square's corners e, f, g, h can be recovered.

With these definitions, the central structural claim of the project is the following law, which the underlying algebra makes an equivalence rather than a one-directional implication.

Law 1 (Structural Route-Closure Law)**.** In the cancellative commutative setting considered here,

$$\text{primal factor traceability} \iff \text{global four-factor route closure.}$$

$$\begin{tikzcd}
 & ab = cd \arrow[d, dashed, "\text{hidden four-factor square}"] & \\
 & e \arrow[dl, dash, "a = ef"'] \arrow[dr, dash, "c = eg"] & \\
f \arrow[dr, dash, "b = gh"'] & & g \arrow[dl, dash, "d = fh"] \\
 & h &
\end{tikzcd}$$

$$\begin{tikzcd}
a \mid cd \arrow[d, "\text{primality}"] \\
a = ef,\ e \mid c,\ f \mid d \arrow[d] \\
\text{common refinement}
\end{tikzcd}$$

Figure 1: The refinement square. An equal-endpoint identity $ab = cd$ (Definition 1) is the visible trace of a hidden square with corners e, f, g, h satisfying (2); primality of the distinguished factor a (Definition 2, Observation 1) is exactly what lets that square be recovered from $ab = cd$ alone.

Thus factor traceability is not merely correlated with route closure: it is the exact algebraic state corresponding to it. This equivalence is classical algebra [5, 6, 7, 8]; the identification of primal factor traceability as *the* invariant that survives across the positive integers, rank-one numerical monoids, higher-rank positive affine monoids, and the omnific integers — and the analysis of how each regime obtains or loses that state, with complete proofs in each case — is the structural contribution of the UNNS Substrate program's investigation reported here.

Remark 1 (Multiplicative notation versus additive notation)**.** Definitions 1–3 and Law 1 are written multiplicatively, matching the classical primality/refinement language of (1)–(2) and the positive-integer control of Section 3. From Section 4 onward the ambient monoids ($H \subseteq \mathbb{N}_0$, then positive affine $H \subseteq \mathbb{Z}^d$) are additive, and the same law is restated with $+$ in place of $\cdot$, $\leq_H$ ("$x \leq_H y$" meaning $y - x \in H$) in place of $|$, and $e + f$, $g + h$, $e + g$, $f + h$ in place of (2): an atom is prime exactly as in Definition 2, with sums instead of products. These are the same law and the same proof pattern in two notations for the same algebraic structure (a commutative cancellative monoid), never two different laws. The two notations must nonetheless never be silently mixed inside a single proof: the additive rank-one construction of Lemma 4 ($e = \min(a, c)$, $f = a - e$, $g = c - e$, with h recovered from the endpoint equality) has no multiplicative analogue and must not be presented as a restatement of the positive-integer gcd/coprime witness of Proposition 1, nor conversely; the two witnesses answer the same question in two different regimes, using two different mechanisms, and Section 3 is stated separately from Theorem 1 for exactly this reason.

A useful local refinement of Law 1, used repeatedly in the proofs below, isolates a single distinguished factor:

Observation 1 (Local route-closure criterion)**.** For a nonzero distinguished factor a in an equal-endpoint identity $ab = cd$ running through a: a is primal if and only if every such identity through a admits a four-factor refinement.

2.1 Cross-regime architecture

Figure 2 previews how the four regimes examined in this manuscript relate to one another as a cross-regime progression of successive structural comparisons, each of increasing structural scope and each carrying its own exact criterion for when Law 1's right-hand side holds.

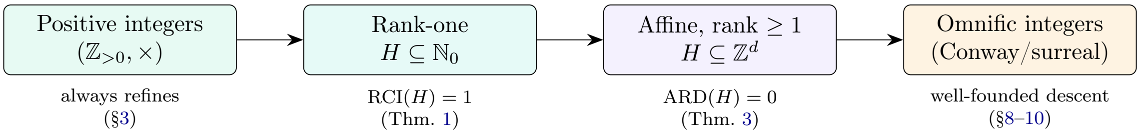

Figure 2: Cross-regime route-closure architecture. Each arrow marks a successive structural regime of increasing scope, not a literal algebraic inclusion of the previous one; each regime's label below it names its exact route-closure criterion, proved or audited in the section indicated.

Sections 3–6 instantiate Law 1 and Observation 1 across these regimes, and identify exactly where the equivalence's right-hand side — route closure — fails, and what that failure looks like geometrically.

3 Positive Integers — Canonical Route-Closure Control

Before turning to the additive numerical-monoid setting of Principal Theorem I, it is worth recording explicitly why the classical multiplicative case, $(\mathbb{Z}_{>0}, \times)$, is not merely a special case folded into the rank-one theorem: it uses a different (multiplicative, gcd-based) witness construction, and it holds *unconditionally* on its own domain — this section studies the full monoid $(\mathbb{Z}_{>0}, \times)$ itself, not an arbitrary multiplicative sub-monoid, and on that domain the gcd/coprime construction below succeeds universally, with no analogue of a defective case. (Proper multiplicative sub-monoids of $\mathbb{Z}_{>0}$ certainly exist — e.g. the odd positive integers, or the positive perfect squares — but they are outside the scope of this control.) It is the canonical control against which the additive rank-one theorem's own sufficiency direction (Lemma 4 below) is compared.

Proposition 1 (Positive-Integer Route Closure)**.** *Let $a, b, c, d \in \mathbb{Z}_{>0}$ with $ab = cd$. Set*

$$e = \gcd(a, c), \qquad f = a/e, \qquad g = c/e, \qquad h = b/g.$$

Then $e, f, g, h \in \mathbb{Z}_{>0}$, $\gcd(f, g) = 1$, and

$$a = ef, \qquad c = eg, \qquad d = fh, \qquad h = d/f,$$

so (e, f, g, h) is a four-factor refinement of $ab = cd$. Consequently $(\mathbb{Z}_{>0}, \times)$ has global four-factor refinement.

Proof. By construction $e = \gcd(a, c)$ divides both a and c, so $f = a/e$ and $g = c/e$ are positive integers with $\gcd(f, g) = 1$ (dividing a and c by their gcd leaves coprime quotients), and $a = ef$, $c = eg$ hold by definition. From $ab = cd$ we get $efb = egd$, and cancelling the common factor e gives

$$fb = gd.$$

Since $g \mid fb$ and $\gcd(f, g) = 1$, Euclid's lemma gives $g \mid b$, so $h := b/g$ is a positive integer, i.e. $b = gh$. Substituting back into $fb = gd$ gives $fgh = gd$, and cancelling g gives $fh = d$. Thus (e, f, g, h) satisfies all four identities of (2), and $h = b/g = d/f$ is forced, matching the construction. Since a, b, c, d were arbitrary subject to $ab = cd$, every equal-endpoint identity in $(\mathbb{Z}_{>0}, \times)$ refines. □

Unlike every other route-closure result in this manuscript, Proposition 1 has no failure mode to classify: it holds for *every* equal-endpoint identity of positive integers, with no analogue of Principal Theorem I's index $\mathrm{RCI}(H)$ or Principal Theorem II's defect $\mathrm{ARD}(H)$. This is precisely why it is presented as a control rather than as a fourth theorem alongside Principal Theorems I–III: it fixes the canonical witness construction (a gcd/coprime routing) against which the strictly different additive construction of Lemma 4 is deliberately contrasted, and it is the base case instantiating the "exact routed local block" hypothesis of the Route-Closure Induction Proposition (Section 9) in its simplest possible form.

3.1 Computational validation

The project's frozen Phase-1 protocol fixes exactly the witness of Proposition 1 — $e = \gcd(a, c)$, $f = a/e$, $g = c/e$, $h = b/g = d/f$ — as the canonical construction, together with a route-closure defect indicator D_R (0 for a verified refinement, 1 for none found). Applying this frozen witness to 500 randomly generated equal-endpoint quadruples (a, b, c, d) with a, b, c, d drawn from a wide range (products up to six digits), every single case produced a valid refinement: $D_R = 0$ in all 500/500 cases, with zero exceptions. This is exactly the outcome Proposition 1 predicts as a general theorem, not a statistical regularity; the 500-case run is a machine-checked regression confirming the general construction rather than a source of the claim itself.

4 Principal Theorem I — Rank-One Common Refinement Criterion

Theorem 1 (Rank-One Common Refinement Criterion)**.** *Let $H \subseteq \mathbb{N}_0$ be a nonzero additive submonoid, let $m = \min(H \setminus \{0\})$, and let $\gamma = \gcd(H)$ (equivalently, the positive generator of $\mathrm{gp}(H) = \gamma\mathbb{Z}$). The following are equivalent:*

(i) *H has the 2×2 (Riesz) refinement property: for all $a, b, c, d \in H$ with $a + b = c + d$ there exist $e, f, g, h \in H$ with $a = e + f$, $b = g + h$, $c = e + g$, $d = f + h$;*

(ii) *m is prime with respect to the algebraic preorder of H (i.e. $m \leq_H x$ or $m \leq_H y$ whenever $m \leq_H x + y$);*

(iii) *$H = m\mathbb{N}_0$;*

(iv) $H = (\mathbb{R}_{\geq 0}H) \cap \mathrm{gp}(H) = \gamma\mathbb{N}_0$, *i.e.* H *equals its own rank-one lattice saturation;*

(v) $m = \gamma$.

Define the route-closure index $\mathrm{RCI}(H) := m/\gamma \geq 1$. *Then* $\mathrm{RCI}(H) = 1$ *if and only if* H *has global common refinement, and* $\mathrm{RCI}(H) > 1$ *forces an explicit, constructible route defect.*

Proof. The equivalence is established by a cycle of four lemmas, followed by the lattice form and an explicit converse.

Lemma 1 (The least positive element is an atom)**.** *m cannot be written as $u + v$ with $u, v \in H$ both nonzero.*

Proof. If $u, v > 0$ then $u, v < m$ by minimality of m among nonzero elements of H, but $u \in H \setminus \{0\}$ would then contradict $m = \min(H \setminus \{0\})$. ☐

Lemma 2 (Refinement forces the least element to be prime)**.** *If H satisfies (i), then m is prime in the sense of (ii).*

Proof. Suppose $m \leq_H x+y$, i.e. $m+z = x+y$ for some $z \in H$. Refinement gives $m = r+s$, $z = t+u$, $x = r+t$, $y = s+u$ with $r, s, t, u \in H$. By Lemma 1, $r = 0$ or $s = 0$. If $r = 0$ then $m = s$, so $y = s+u = m+u$ and $m \leq_H y$. If $s = 0$ then $m = r$, so $x = r+t = m+t$ and $m \leq_H x$. Either way m is prime. ☐

Lemma 3 (A prime least element forces $H = m\mathbb{N}_0$)**.** *If m is prime in the sense of (ii), then $H = m\mathbb{N}_0$.*

Proof. Let $x \in H$, $x > 0$. As ordinary integers, m and x are commensurable: there exist positive integers r, s with $rm = sx$. Then $sx = m + (r-1)m$, so $m \leq_H sx$. Writing sx as the s-fold sum $x + \cdots + x$ and applying primality of m repeatedly (at each step m divides the running sum, hence divides one of the two summands in a binary split of that sum) gives $m \leq_H x$, i.e. $x = m + y$ for some $y \in H$. If $y > 0$, repeat with y in place of x; since the ordinary-integer value strictly decreases at each step, the process terminates after finitely many steps at $x = qm$ for some $q \in \mathbb{N}_0$. Hence $H \subseteq m\mathbb{N}_0$. The reverse inclusion holds because $m \in H$ and H is closed under addition, so $H = m\mathbb{N}_0$. ☐

Lemma 4 ($m\mathbb{N}_0$ has refinement)**.** *For every $m \geq 1$, the monoid $m\mathbb{N}_0$ satisfies (i).*

Proof. Scaling by m identifies $m\mathbb{N}_0$ with $\mathbb{N}_0$, so it suffices to prove refinement in $\mathbb{N}_0$. For $a + b = c + d$ in $\mathbb{N}_0$, set

$$e = \min(a, c), \qquad f = a - e, \qquad g = c - e.$$

Both e and (since $a \geq e$, $c \geq e$) f, g are nonnegative integers, and $a = e+f$, $c = e+g$ hold by construction. The endpoint equality $a+b = c+d$ then gives $b - g = d - f$ (both sides equal $a + b - c = b - (c - a) = b - g + f - f$, i.e. a direct substitution shows this common value, call it h, is nonnegative): concretely, if $a \leq c$ then $e = a$, $f = 0$, $g = c - a \geq 0$, and $b = d + (c-a) = d + g$ forces $h := b - g = d \geq 0$; if $c \leq a$ the symmetric computation gives $h := d - f \geq 0$ with $b = g + h$. In either case $h \in \mathbb{N}_0$ and $b = g + h$, $d = f + h$, completing the refinement (e, f, g, h). ☐

Closing the cycle. Lemma 2 gives (i)⇒(ii); Lemma 3 gives (ii)⇒(iii); Lemma 4 gives (iii)⇒(i) (via the scaled copy of $\mathbb{N}_0$), closing (i)–(iii) into an equivalence. For the lattice form (iv): for a nonzero submonoid of $\mathbb{N}_0$, $\mathrm{gp}(H) = \gamma\mathbb{Z}$ and $(\mathbb{R}_{\geq 0}H) \cap \mathrm{gp}(H) = \gamma\mathbb{N}_0$, so $H = m\mathbb{N}_0$ is equivalent to $m\mathbb{N}_0 = \gamma\mathbb{N}_0$, i.e. $m = \gamma$, which is (v); and $m = \gamma$ says exactly that H already equals its rank-one lattice saturation $\gamma\mathbb{N}_0$, which is (iv). This closes (iii) $\iff$ (iv) $\iff$ (v), completing the full equivalence, with $\mathrm{RCI}(H) = m/\gamma = 1$ exactly encoding (v). □

Corollary 1 (Constructive failure)**.** *If* $\mathrm{RCI}(H) > 1$ *(equivalently* $m > \gamma$*), an explicit route defect can be constructed: let* x *be the least element of* H *not a multiple of* m*, and let* $k \geq 2$ *be the least integer with* $kx - m \in H$ *(such* k *exists because a common multiple of* m *and* x *eventually supplies one). Then*

$$m + (kx - m) \;=\; x + (k-1)x$$

is an equal-endpoint identity in H *with no refinement: any refinement would split the atom* m *(Lemma 1) as* $m = e + f$ *with* $e = m, f = 0$ *or* $e = 0, f = m$*; the first forces* $x - m \in H$*, contradicting minimality of* x*; the second forces* $(k-1)x - m \in H$*, contradicting minimality of* k*.*

Example 1 (Corollary 1 made concrete)**.** Let $H = \langle 4, 6\rangle = \{0, 4, 6, 8, 10, 12, \ldots\} \subset \mathbb{N}_0$ (all nonnegative integer combinations of 4 and 6). Then $\gamma = \gcd(4, 6) = 2$ and $m = \min(H \setminus \{0\}) = 4$, so $\mathrm{RCI}(H) = 4/2 = 2 > 1$ and H is globally non-refinable by Theorem 1. The least element of H not a multiple of $m = 4$ is $x = 6$; the least $k \geq 2$ with $kx - m \in H$ is $k = 2$, since $2 \cdot 6 - 4 = 8 \in H$. Corollary 1 then produces the explicit equal-endpoint identity

$$\underbrace{4}_{m} + \underbrace{8}_{kx-m} = \underbrace{6}_{x} + \underbrace{6}_{(k-1)x}, \qquad \text{i.e.} \qquad 4 + 8 = 6 + 6 = 12,$$

with no valid four-factor refinement: the atom $m = 4$ forces $(e, f) = (4, 0)$ or $(0, 4)$ (Lemma 1); the first would require $g = 6 - 4 = 2 \in H$, and the second would require $h = 6 - 4 = 2 \in H$, and $2 \notin H$ in either case. By contrast, $H' := 2\mathbb{N}_0 = \{0, 2, 4, 6, \ldots\}$ has $m' = \gamma' = 2$, so $\mathrm{RCI}(H') = 1$ and every equal-endpoint identity in H' refines by Lemma 4: e.g. $4 + 8 = 6 + 6$ itself, read inside H' rather than $H = \langle 4, 6\rangle$, refines as $e = 4, f = 0, g = 2, h = 6$ ($4 = 4 + 0$, $8 = 2 + 6$, $6 = 4 + 2$, $6 = 0 + 6$), since $2 \in H'$ even though $2 \notin H$. The single missing generator 2 is exactly the gap that separates a globally refinable rank-one monoid from a non-refinable one with the same ambient group $2\mathbb{Z}$.

4.1 Computational validation

Theorem 1 was validated by exhaustive search over all 793 rank-one generator families with generators in the range 1–12 and generator-set sizes 1–4. The predicted classification — 280 globally refinable families ($\mathrm{RCI}(H) = 1$) and 513 globally non-refinable families ($\mathrm{RCI}(H) > 1$) — was confirmed exactly: for every one of the 513 predicted-non-refinable systems, the canonical counterexample construction of Corollary 1 was verified, and an exact finite search found *zero* systems among the 513 for which any valid refinement witness nonetheless existed. The classification is therefore exact, not merely a majority statistical pattern.

4.2 The Cofinal Primal-Tail Theorem

Theorem 1 classifies *whether* a rank-one monoid has global refinement, but says nothing about the internal distribution of failure inside a non-refinable H. A sharper, elementwise analysis answers this: global failure never means failure everywhere.

Normalize by γ: write $S = H/\gamma$, a primitive numerical monoid. If $S = \mathbb{N}_0$, Theorem 1 already gives global refinement and every element is primal. Otherwise S has a finite *conductor* $c = \min\{n : n + \mathbb{N}_0 \subseteq S\}$, so every gap of S lies in $\{1, \dots, c-1\}$.

Lemma 5 (Gap-localization)**.** *If an equal-endpoint identity $x + t = y + z$ in S (with distinguished element x) admits no 2×2 refinement, then $|x - y| \notin S$ and $|x - z| \notin S$; hence $|x - y| < c$ and $|x - z| < c$.*

Proof. If $y \geq x$ and $y - x \in S$, the identity refines immediately via $x = x + 0$ with complementary pieces $y - x$ and z. If $x \geq y$ and $x - y \in S$, it refines via $x = y + (x - y)$ with the remaining corner equal to t. The symmetric argument applies to z. So a failed square requires both distances to be gaps of S, and every gap is strictly below the conductor c. □

Theorem 2 (Cofinal Primal-Tail Theorem)**.** *Let S be a primitive numerical monoid with conductor $c > 0$. Then every $x \in S$ with $x \geq 4c$ is primal. Equivalently, in the unnormalized monoid H, every $x \geq 4\gamma c$ is primal.*

Proof. Suppose for contradiction that $x + t = y + z$ is non-refinable with $x \geq 4c$. By Lemma 5, a failed square requires y, z within gap-distance of x. If $y \geq x$: write $x = (x - c) + c$; since $x - c \geq 3c$, both pieces lie in S, and $y - (x - c) = y - x + c \geq c$ while $z - c \geq x - c \geq 2c$ (using $z > x - c$), so all complementary pieces lie in S, giving a refinement — contradiction. The symmetric argument rules out $z \geq x$. So a failed square requires $y = x - r$, $z = x - s$ with $0 < r, s < c$ (Lemma 5), whence $t = x - r - s \geq 4c - 2(c-1) > 2c$. Now split $x = (y - c) + (r + c)$: both pieces lie in S since $y - c \geq 2c$ and $r + c \geq c$, with complementary pieces c and $t - c \geq c$, both in S. This gives a refinement, again a contradiction. Hence every $x \geq 4c$ is primal. □

4.3 The finite obstruction core

Theorem 2 guarantees, for *every* non-refinable rank-one system, a cofinal all-primal tail beyond $4c$ (equivalently $4\gamma c$ in H); a finer elementwise analysis of the 513 non-refinable rank-one systems in the corpus confirms this bound is comfortable in practice. Among these systems, 414 (80.7%) have a non-primal element already near the bottom of H, but every system's non-primal locus is confined to a *finite obstruction core*: the median count of non-primal elements is 18 (maximum 130 across the corpus), and the ratio of the empirically observed tail onset to the conductor ranges from 2.1 to 3.5 (median 2.67), comfortably inside the proved universal bound of $4c$ from Theorem 2. In short: a rank-one system can fail global route closure, but the failure is always eventually repaired — it cannot recur indefinitely far out along H.

Example 2 (The corpus extremes)**.** The corpus's two extremes make the bound concrete. At the small end, $S = \langle 2, 3 \rangle = \{0, 2, 3, 4, 5, 6, \dots\}$ (conductor $c = 2$) has the smallest obstruction core in the corpus: its non-primal locus is exactly the four elements $\{2, 3, 4, 5\}$, with the exact primal tail already beginning at 6, comfortably inside the guaranteed bound $4c = 8$. At the large end, $S = \langle 11, 12\rangle$ (conductor $c = 110$) has the corpus's largest recorded obstruction core, 130 non-primal elements, with the exact tail onset at $231 = 2.1c$ — still well inside the guaranteed universal bound $4c = 440$ from Theorem 2, even though the core itself is over thirty times larger than in the $\langle 2, 3\rangle$ case. Both are governed by the same proved bound; only the conductor changes.

Figure 3 summarizes Principal Theorem I and its two companion results in a single picture: the algebraic criterion on the left, and the elementwise geometry it produces — a finite non-primal core with a guaranteed cofinal all-primal tail — on the right.

$\mathrm{RCI}(H) := m/\gamma$
$\mathrm{RCI}(H) = 1$
$\iff m = \gamma \iff H = \gamma\mathbb{N}_0$
$\iff$ global refinement
(Thm. 1)
$\mathrm{RCI}(H) > 1$
$\Rightarrow$ explicit constructive route defect
(Cor. 1)
normalize $S = H/\gamma$; conductor $c = \min\{n : n + \mathbb{N}_0 \subseteq S\}$

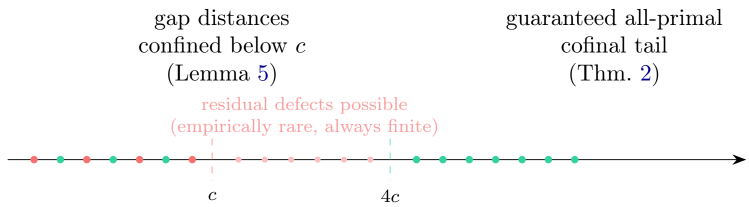

Figure 3: Rank-one criterion and finite obstruction core. Left: the route-closure index $\mathrm{RCI}(H)$, its two exhaustive outcomes, and the corollary each produces. Right: the elementwise geometry a non-refinable H always exhibits. What Lemma 5 confines below c is the pairwise *distance* inside a failed square, not the location of every non-primal element itself: individual non-primal elements can and do occur above c (the possible residual band shown here), but Theorem 2 guarantees the entire non-primal locus is finite and vanishes by $4c$, beyond which the tail is entirely primal (realized already at $3c$ for $\langle 2, 3\rangle$ and at $2.1c$ for $\langle 11, 12\rangle$, Example 2).

5 Principal Theorem II — Affine Route-Closure Criterion

Theorem 1 is a rank-one statement. The natural question is whether an analogous criterion survives once the ambient monoid has rank two or more. It does, but the criterion strengthens in a way that reveals higher rank as fundamentally different from rank one.

Theorem 3 (Affine Route-Closure Criterion)**.** *Let H be a positive affine monoid: a finitely generated additive submonoid of $\mathbb{Z}^d$ with $H \cap (-H) = \{0\}$ (commutative, cancellative, torsion-free, conical, finitely generated). Let $\mathcal{A}(H)$ be its finite atom set, $r = \operatorname{rank}\operatorname{gp}(H)$, $s = |\mathcal{A}(H)| \geq r$, and $\mathrm{ARD}(H) = s - r$. The following are equivalent:*

(i) *H has the Riesz 2×2 refinement property;*

(ii) *every atom of H is prime in the algebraic preorder;*

(iii) *factorization into atoms is unique;*

(iv) *H is a free commutative monoid on its atoms, i.e. $H \cong \mathbb{N}_0^r$;*

(v) $s = r$, *i.e.* $\mathrm{ARD}(H) = 0$.

Proof. **(i)⇒(ii).** Let p be an atom and $p+z = x+y$. Refinement gives $p = e+f$, $z = g+h$, $x = e+g$, $y = f+h$ with $e, f, g, h \in H$. Since p is an atom, $e = 0$ or $f = 0$. If $e = 0$ then $p = f$ and $p \leq_H y$; if $f = 0$ then $p = e$ and $p \leq_H x$. So p is prime.

(ii)⇒(iii). H is atomic (every nonzero element is a finite sum of atoms). Suppose $p_1 + \cdots + p_m = q_1 + \cdots + q_n$ are two atomic factorizations. Since p_1 is prime, it lies below some q_j; since q_j is itself an atom, $p_1 = q_j$. Cancelling this atom from both sides and inducting gives equality of the two multisets of atoms.

(iii)⇒(iv). Let the distinct atoms be $u_1, \ldots, u_s$. Unique factorization makes $\mathbb{N}_0^s \to H$, $(n_1, \ldots, n_s) \mapsto \sum_i n_i u_i$, bijective, so $H \cong \mathbb{N}_0^s$. The Grothendieck group is then $\mathrm{gp}(H) \cong \mathbb{Z}^s$, so $s = r$ and (iv) holds with $H \cong \mathbb{N}_0^r$.

(iv)⇒(v). If $H \cong \mathbb{N}_0^r$ then the standard basis vectors are exactly the atoms, so $s = r$, i.e. $\mathrm{ARD}(H) = 0$. (Conversely, $s = r$ atoms generating $\mathrm{gp}(H) \cong \mathbb{Z}^r$ must be $\mathbb{Z}$-linearly independent, ruling out any nontrivial relation between nonnegative atomic combinations, so unique factorization and freeness follow — closing (v)⇒(iv)⇒(iii) as well.)

(iv)⇒(i). In $\mathbb{N}_0^r$, refinement is coordinatewise: for $a+b = c+d$, set $e_i = \min(a_i, c_i)$, $f_i = a_i - e_i$, $g_i = c_i - e_i$; the endpoint equality yields a nonnegative h_i with $b_i = g_i + h_i$, $d_i = f_i + h_i$ (exactly the rank-one construction of Lemma 4, applied coordinatewise). Thus $\mathbb{N}_0^r$, and every monoid isomorphic to it, has refinement.

This closes the full cycle of equivalences. □

Corollary 2 (Geometric form). *Let $L = \mathrm{gp}(H)$ and $C = \mathbb{R}_{\geq 0} H$. Refinement is equivalent to the existence of a lattice basis $u_1, \ldots, u_r$ of L with $C = \sum_i \mathbb{R}_{\geq 0} u_i$ and $H = \sum_i \mathbb{N}_0 u_i$; equivalently, $H = C \cap L$ and C is unimodular simplicial with respect to L.*

5.1 Saturation versus atomic independence

The rank-one theorem was $H \subseteq \mathbb{N}_0$: refinement $\iff$ H equals its own lattice saturation. In rank one every pointed rational cone has exactly one ray, so the cone is automatically simplicial, and the primitive ray generator is automatically a lattice basis relative to the group it generates; unimodular simpliciality is therefore *invisible* in rank one. Once rank exceeds one, lattice saturation (normality of the affine semigroup) remains necessary but is no longer sufficient: the higher-rank criterion additionally demands atomic independence, i.e. unimodular simpliciality of the generating cone (Corollary 2), which can fail even when H is normal. Four exact examples make this precise, in increasing order of subtlety, drawn from the seven-system test suite underlying the computational validation below.

- *Not even lattice-saturated.* $H = \langle (2,0), (1,1), (0,1) \rangle \subsetneq \mathbb{N}_0^2$ already fails to equal its own lattice saturation: the saturation contains $(1,0)$ (since $2(1,0) = (2,0) \in H$ and $(1,0) \in \mathrm{gp}(H)$), but $(1,0) \notin H$, a genuine lattice hole. The same identity $(2,0)+(0,2) = (1,1)+(1,1)$ (with $(0,2) = 2(0,1) \in H$) again has all three atoms $(0,1), (1,1), (2,0)$, so $\mathrm{ARD}(H) = 3 - 2 = 1$; here failure is visible already from missing lattice saturation, before unimodularity is ever in question.

- *Saturated but non-unimodular.* $H = \langle (2,0), (1,1), (0,2) \rangle$ equals $\{(x,y) \in \mathbb{N}_0^2 : x \equiv y \pmod 2\}$, so it is lattice-saturated relative to its own generated group. But $(2,0) +$

$(0,2) = (1,1) + (1,1)$ with all three generators atoms, so $\mathrm{ARD}(H) = 3 - 2 = 1$ and refinement fails: saturation alone is not sufficient beyond rank one.

- *Saturated and simplicial, but non-unimodular.* $H = \langle (1,0), (1,1), (1,2) \rangle$ is the full lattice-point monoid of the cone $0 \le y \le 2x$, yet $(1,0) + (1,2) = 2(1,1)$, so again $\mathrm{ARD}(H) = 1$: even normality plus simpliciality is insufficient without unimodularity.
- *Saturated but non-simplicial.* $H = \langle (0,0,1), (1,0,1), (0,1,1), (1,1,1) \rangle$, the lattice-point monoid of $0 \le x \le z$, $0 \le y \le z$, has cone with four extreme rays in rank three, and $(0,0,1) + (1,1,1) = (1,0,1) + (0,1,1)$ gives $\mathrm{ARD}(H) = 4 - 3 = 1$.

In each case the listed generators are genuinely atoms (irreducible: not a sum of two nonzero elements of H), which is what makes the reported $\mathrm{ARD}(H)$ correct rather than an artifact of a redundant generating set. For the saturated example $H = \{(x, y) \in \mathbb{N}_0^2 : x \equiv y \pmod 2\} = \langle (2,0), (1,1), (0,2) \rangle$, take $(2,0) = (x_1, y_1) + (x_2, y_2)$ with both summands nonzero and in H: from $y_1 + y_2 = 0$ and $y_1, y_2 \ge 0$ we get $y_1 = y_2 = 0$, so $x_1 \equiv y_1 = 0$ and $x_2 \equiv y_2 = 0 \pmod 2$, i.e. x_1, x_2 are even; being nonzero with $y_i = 0$ forces $x_i \ge 2$ for both i, but then $x_1 + x_2 \ge 4 > 2$, a contradiction, so no such decomposition exists and $(2,0)$ is an atom. The same parity argument, applied to $(1,1) = (x_1, y_1) + (x_2, y_2)$, forces (since $x_1 + x_2 = 1 = y_1 + y_2$ with all four entries nonnegative integers) one of the two summands to equal $(0,0)$, again ruling out a decomposition into two nonzero elements of H; by the symmetric argument $(0,2)$ is an atom as well. All three generators are therefore atoms, and $|\mathcal{A}(H)| = 3 > 2 = r$ genuinely witnesses $\mathrm{ARD}(H) = 1$ rather than double-counting a reducible generator. The same direct verification (checked computationally for all four retained examples) is what the frozen `UNNS_ADMISSIBILITY_PERCOLATION` records certify for the higher-rank counterexamples used throughout Sections 5–6.

So the affine higher-rank picture is

$$\boxed{\text{route closure} = \text{lattice closure} + \text{atomic independence}}$$

or, algebraically, route closure $\iff$ every irreducible is prime $\iff$ free atomic architecture.

Figure 4 makes this collapse of sufficient conditions visible: what alone suffices at rank one (Theorem 1) splits into two independently insufficient conditions at higher rank, and only their conjunction — unimodular simpliciality, Corollary 2 — restores the criterion.

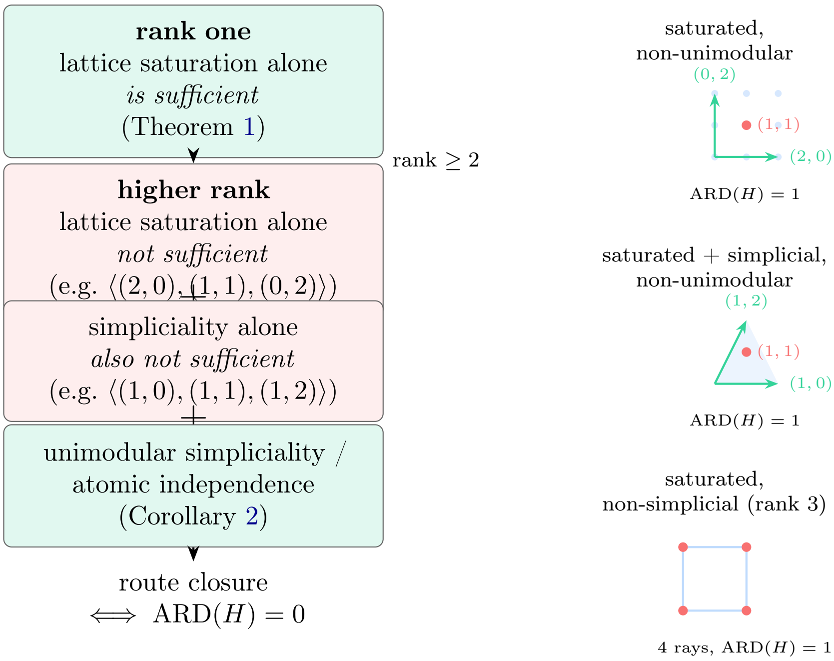

Figure 4: From rank-one saturation to higher-rank atomic independence. Left: what suffices at rank one collapses into two separately insufficient conditions once rank exceeds one, restored only by their conjunction. Right: the three qualitatively distinct affine counterexamples of Section 5.1 that make each insufficiency concrete — a non-unimodular saturated cone, a simplicial but non-unimodular cone, and a saturated non-simplicial cone in rank three, each with one redundant atom beyond the rank (ARD$(H) = 1$).

5.2 Computational validation

The frozen test suite fixes exactly seven higher-rank affine systems, named and classified in Table 1.

Table 1: The seven-system affine test suite. $s = |\mathcal{A}(H)|$ is the atom count, $r = \operatorname{rank} \operatorname{gp}(H)$.

System	Generators	s	r	ARD(H)	Verdict
`FREE_N2`	$(1,0),(0,1)$	2	2	0	refines
`FREE_SKEW`	$(2,0),(1,1)$	2	2	0	refines
`FREE_N3`	$(1,0,0),(0,1,0),(0,0,1)$	3	3	0	refines
`NONNORMAL_2D`	$(2,0),(1,1),(0,1)$	3	2	1	fails
`NORMAL_PARITY`	$(2,0),(1,1),(0,2)$	3	2	1	fails
`NORMAL_CONE2`	$(1,0),(1,1),(1,2)$	3	2	1	fails
`NORMAL_SQUARE3`	$(0,0,1),(1,0,1),(0,1,1),(1,1,1)$	4	3	1	fails

Elementwise verification against 12 certificate families and 748 regression certificate instances confirmed the predicted classification exactly: the three free systems (`FREE_N2`, `FREE_SKEW`, `FREE_N3`) refine globally, matching ARD$(H) = 0$, while the four systems of Section 5.1 (`NONNORMAL_2D`, `NORMAL_PARITY`, `NORMAL_CONE2`, `NORMAL_SQUARE3`) all fail global refinement, matching ARD$(H) = 1$ in every case. In every one of the four failing systems the non-primal locus was found to be *unbounded* — propagating along rays, congruence classes, or strips, or occupying all but a lower-dimensional primal spine of the monoid (sample windows ranged from 313 elements with 145 primal, to 2,470 graded elements with only

10 primal). This computational finding is sharpened into an unconditional theorem next.

6 Principal Theorem III — Persistent Defect-Ray Theorem

The finding above — that affine failure, unlike rank-one failure, is never confined to a finite core — is strong enough to state and prove as a theorem in its own right, converting an empirical pattern into a guaranteed geometric consequence of $\mathrm{ARD}(H) > 0$.

Theorem 4 (Persistent Defect-Ray Theorem)**.** *Let H be a positive affine monoid with* $\mathrm{rank}(H) \geq 2$ *and* $\mathrm{ARD}(H) > 0$*. Then there exist an atom $a \in \mathcal{A}(H)$ and $0 \neq \rho \in H$ such that*

$$a + n\rho \notin P(H) \qquad \text{for every } n \in \mathbb{N}_0.$$

In particular $H \setminus P(H)$ contains an infinite affine ray, hence is unbounded.

The proof proceeds through four lemmas, each isolating one stage of the obstruction's geometry.

6.0.1 Non-prime atom

Lemma 6. $\mathrm{ARD}(H) > 0$ *forces at least one atom $p \in \mathcal{A}(H)$ to be non-prime (hence non-primal, since for an atom primality and primal divisibility coincide).*

Proof. H is atomic with atoms generating H. If every atom were prime, cancellation through atomic factorizations would force unique factorization (Theorem 3, (ii)⇒(iii)), and the atom map $\mathbb{N}_0^{|\mathcal{A}(H)|} \to H$ would be an isomorphism, forcing $|\mathcal{A}(H)| = \mathrm{rank}(H)$, i.e. $\mathrm{ARD}(H) = 0$. Contrapositively, $\mathrm{ARD}(H) > 0$ forces some atom p to be non-prime. □

6.0.2 Irredundant witness

Lemma 7. *A non-prime atom p yields distinct atoms p, q and $t, w \in H$ with*

$$p + t = q + w, \qquad p \not\leq_H w, \qquad q \not\leq_H t.$$

Proof. Since p is non-prime, there are $y, z \in H$ with $p \leq_H y + z$ but $p \not\leq_H y$ and $p \not\leq_H z$. Factor $y + z$ into atoms and, among all sub-multisets of those atomic factors whose sum is divisible by p, choose one minimal under inclusion, with sum $S = p + t$. No atom in the minimal sub-multiset equals p (else p would divide y or z), so it has at least two atoms; let q be one of them and w the sum of the rest, giving $p + t = q + w$. Minimality gives $p \not\leq_H w$. If $t = q + u$ for some u, then $p + q + u = q + w$ cancels to $p + u = w$, contradicting $p \not\leq_H w$; hence $q \not\leq_H t$. □

6.0.3 Violated facet

Lemma 8. *Let $C = \mathbb{R}_{\geq 0}H$ (a pointed cone, since H is positive). For the witness of Lemma 7, after possibly exchanging p, q, the difference $p - q \notin C$.*

Proof. If both $p - q \in C$ and $q - p \in C$, the pointed cone C would contain the nonzero line through $p - q$, contradicting pointedness. So at most one of $p - q, q - p$ lies in C; exchange

p, q if necessary so that, writing $a = p$, $b = q$, $u = t$, $v = w$, the witness reads $a + u = b + v$ with $a \not\leq_H v$, $b \not\leq_H u$, and $a - b \notin C$. □

6.0.4 Persistent direction

Lemma 9. *Since* $\operatorname{rank}(H) = r \geq 2$ *and* $a - b \notin C$*, there is a facet functional* λ *with* $\lambda(C) \geq 0$ *and* $\lambda(a - b) < 0$*; the facet* $F = C \cap \ker\lambda$ *has* $\dim F = r - 1 \geq 1$*, so* $H \cap F$ *contains some* $0 \neq \rho \in H$ *with* $\lambda(\rho) = 0$*, and hence* $\lambda(a - b + n\rho) < 0$ *for every* $n \geq 0$*, giving* $a - b + n\rho \notin C \supseteq H$*, i.e.* $b \not\leq_H a + n\rho$ *for every* $n \geq 0$*.*

Proof. Immediate from the displayed construction: since $\rho \in F \subseteq \ker\lambda$, $\lambda(\rho) = 0$, so $\lambda(a - b + n\rho) = \lambda(a - b) < 0$ for every n, and every element of $H \subseteq C$ has $\lambda \geq 0$, so $a - b + n\rho \notin H$. This is precisely the propagation mechanism unavailable in rank one, where a pointed cone has no nonzero facet direction at all. □

Proof of Theorem 4. For $n \geq 0$ set $X_n = a + n\rho$ (using the notation of Lemmas 8–9). Translating the witness of Lemma 7 along ρ: $X_n + u = b + (v + n\rho)$. Suppose for contradiction X_n is primal. Since $X_n \leq_H b + (v + n\rho)$, primality gives $X_n = x_1 + x_2$ with $x_1 \leq_H b$, $x_2 \leq_H v + n\rho$; since b is an atom, $x_1 = 0$ or $x_1 = b$. If $x_1 = 0$ then $X_n \leq_H v + n\rho$, and cancelling the common translation gives $a \leq_H v$, contradicting irredundancy (Lemma 7). If $x_1 = b$ then $b \leq_H X_n$, so $a - b + n\rho \in H$, contradicting Lemma 9. Both cases are impossible, so $X_n \notin P(H)$ for every $n \geq 0$. Since $\rho \neq 0$ and H is torsion-free, $a, a + \rho, a + 2\rho, \ldots$ are distinct and unbounded, so $H \setminus P(H)$ contains an infinite affine ray. □

Example 3 (The mechanism made concrete)**.** Instantiate Lemmas 6–9 on the saturated affine counterexample of Section 5.1, $H = \langle (2, 0), (1, 1), (0, 2) \rangle = \{(x, y) \in \mathbb{N}_0^2 : x \equiv y \pmod 2\}$, with $\operatorname{ARD}(H) = 1$. The atom $(1, 1)$ is non-prime: taking $a = (1, 1)$, $u = (1, 1)$, $b = (2, 0)$, $v = (0, 2)$ gives the equal-endpoint identity $a + u = b + v$, i.e. $(1, 1) + (1, 1) = (2, 0) + (0, 2)$, with $a \not\leq_H v$ and $b \not\leq_H u$ (both differences, $(1, -1)$, lie outside $\mathbb{N}_0^2$), so this is exactly an irredundant witness in the sense of Lemma 7. Here $a - b = (-1, 1) \notin C = \mathbb{R}^2_{\geq 0}$, witnessed by the facet functional $\lambda(x, y) = x$ (so $\lambda(C) \geq 0$ but $\lambda(a - b) = -1 < 0$), and the violated facet is $F = C \cap \ker\lambda = \{(0, y) : y \geq 0\}$, on which $H \cap F = \{(0, 2k) : k \geq 0\}$ supplies the nonzero direction $\rho = (0, 2)$. Theorem 4 then guarantees the explicit infinite family

$$X_n = a + n\rho = (1,\, 2n + 1), \qquad n = 0, 1, 2, \ldots,$$

is non-primal in H for every n: at each n, $X_n + u = b + (v + n\rho)$ reproduces the same obstruction, since $\lambda(a - b + n\rho) = \lambda(a - b) = -1$ regardless of n, so $b \not\leq_H X_n$ can never be repaired by moving further out along the ray. This is the two-dimensional cone-geometry picture of Figure 5 realized on an explicit, checkable family of lattice points.

Figure 5 illustrates the geometry assembled by Lemmas 6–9: the witness pair a, b sits astride a violated facet of the cone C, and the facet direction ρ translates the witness forever without ever re-entering C.

This construction was verified computationally against the four retained affine failure controls, instantiating 100 regression instances of the ray construction with no exception.

Corollary 3 (Dimension-sensitive contrast)**.** *For a nonfree positive affine monoid* H*: at rank one,* $H \setminus P(H)$ *is finite (Theorem 2); at rank* ≥ 2*,* $H \setminus P(H)$ *contains an infinite*

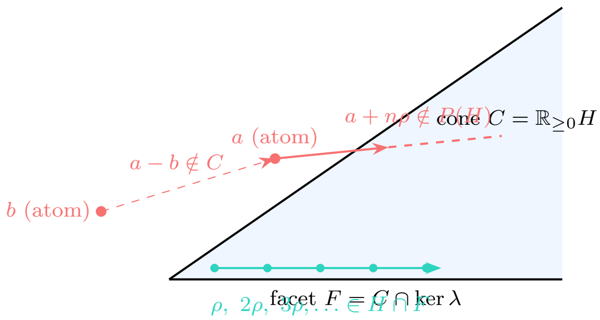

Figure 5: Persistent Defect-Ray proof geometry. The irredundant witness atoms a, b (Lemma 7) are separated by a violated facet of the cone C (Lemma 8); translating a along a nonzero facet direction $\rho \in H \cap F$ (Lemma 9) keeps every $a + n\rho$ outside the primal locus $P(H)$ forever (main proof).

affine ray (Theorem 4). This is an exact structural transition, not a feature of the retained examples alone: the rank hypothesis is essential, since a rank-one pointed cone has no nonzero facet direction along which the obstruction could be preserved.

7 Finite Obstruction Geometry: Core versus Ray

Only after Theorems 1–4 are in hand can their geometric content be synthesized; the geometry below is an interpretation of those theorems, not a substitute for them.

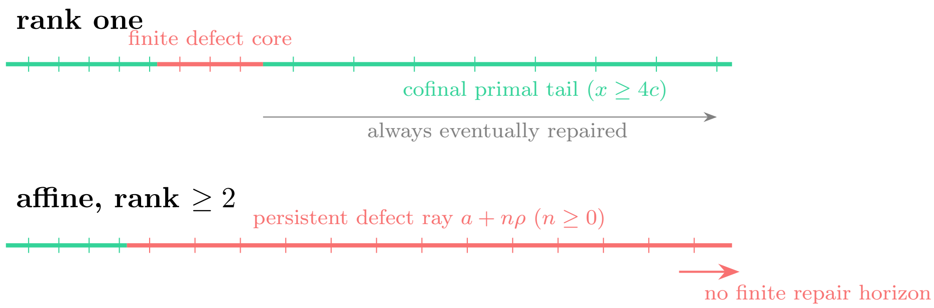

Rank one: finite defect core $\to$ cofinal primal tail. A failing rank-one monoid's non-primal locus is always finite and eventually repaired (Theorem 2, Section 4.3).

Affine, rank ≥ 2: persistent defect ray $\to$ no finite repair horizon. Whenever $\mathrm{ARD}(H) > 0$, Theorem 4 guarantees an infinite ray of non-primal elements that is never repaired.

The qualitative content of Theorems 1–4 is therefore that *route defect at rank one is a bounded, self-healing phenomenon, while route defect at higher affine rank is an unbounded, permanent one.* This sharp dichotomy is what motivates asking, in the next section, whether a comparable persistent channel can survive once the ambient system is enlarged past the finite affine setting entirely.

8 From Finite Obstruction to Transfinite Refinement

Theorem 4 shows that a persistent, permanently unrepaired route defect is possible once rank and non-primality combine at the finite affine level. This raises the natural next question, which the project frames explicitly:

> *If persistent finite-rank obstruction is possible, what prevents such a channel from surviving in the omnific (Conway/surreal) setting?*

The omnific integers are a transfinite extension of the ordinary integers within Conway's surreal number system [3, 4]. The candidate mechanism audited by the project — against the source repository `gaearon/conway-refinement` at commit `264445c9...` (full commit id in [2] and Section 10) — answers the question by exhibiting an explicit repair chain that forbids a persistent non-primal channel from propagating unchecked. Figure 6 gives the audited source-level chain, organized into its three functional phases:

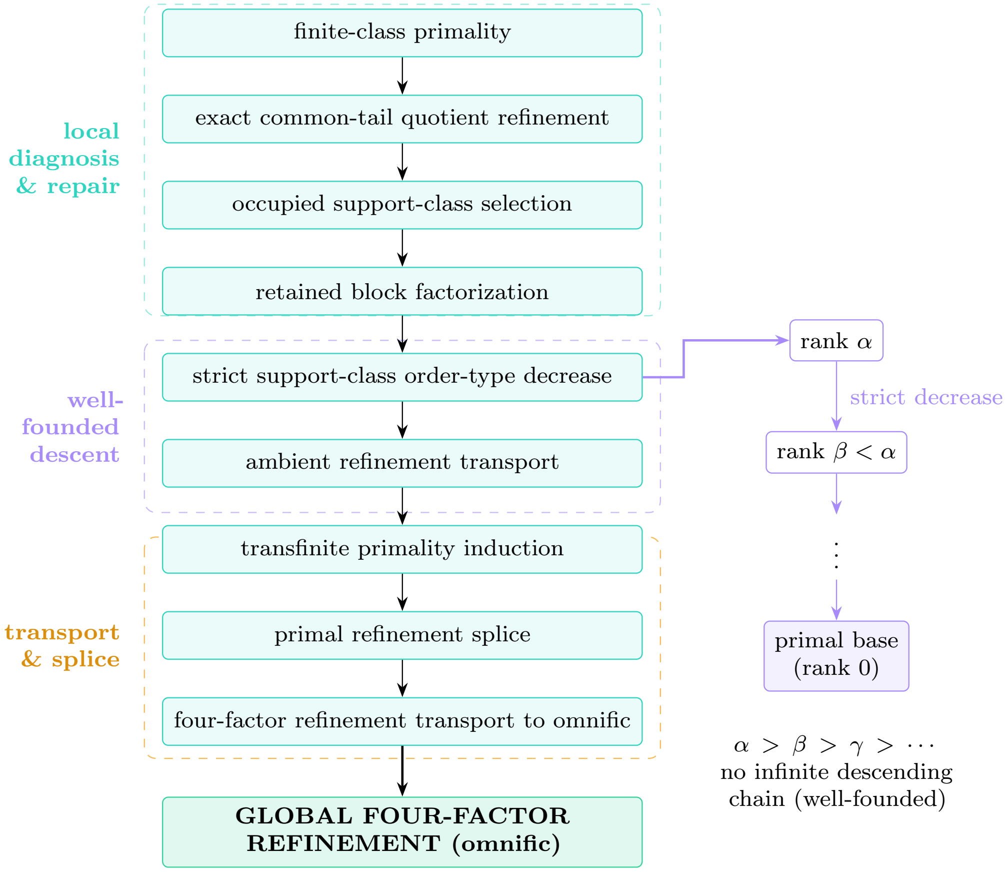

Figure 6: Transfinite repair architecture. The nine-step audited chain (Section 8) organized into its three functional phases: local diagnosis and repair of an unresolved residual, well-founded descent to a strictly lower structural complexity, and transport of the repaired block back into a global splice. The side branch makes explicit what the middle phase actually supplies: not a proof of primality by itself, but a strictly decreasing, well-founded rank sequence that cannot descend forever, forcing termination at a primal base after finitely many steps.

The mechanism this chain implements can be summarized as: an unresolved residual is never permitted to remain at unchanged structural complexity. Every divisibility not already covered by a primal base is repaired by an exact local block whose complement is forced to *strictly lower* well-founded complexity, and this process is shown to terminate at a primal base, after which the local repairs are transported back into the ambient (omnific) system and spliced into a global four-factor refinement. This is the sense in which the manuscript's subtitle, *From Finite Obstruction to Transfinite Refinement*, does real work: it is not a rhetorical arc but a description of an audited, step-by-step mechanism that explicitly forecloses the persistent-channel failure mode exhibited by Theorem 4 at the finite affine level.

We record explicitly the project's own status label for this result: *structural extraction completed; the abstract, from-scratch theorem for the general Hahn-series/surreal setting is not separately (re-)proved here.* What is extracted and independently verified is the induction *mechanism* underlying this chain, stated and formalized next. We also record explicitly, as promised in the introduction, the distinction the project's own synthesis insists on: *well-founded structural descent is a mechanism, not the common law.* A decreasing complexity measure alone is inert if local routing fails, if descent terminates in a non-primal base, or if a quotient repair cannot be reconstructed in the ambient system; what does the work, in every regime of this manuscript, is Law 1. The nine-step chain above is one audited realization of that mechanism for one specific source, not a restatement of the law itself.

9 Route-Closure Induction Proposition and Lean Formalization

The audited chain of Section 8 is source-specific to the Conway/omnific realization. What can be extracted from it, and stated and proved independently of that specific source, is a general induction principle.

Proposition 2 (Route-Closure Induction Proposition)**.** *Let M be a commutative cancellative monoid equipped with a well-founded structural rank ρ. Suppose a designated base regime of M is primal. Suppose further that for every non-base element $a \in M$ and every divisibility $a \mid bc$ witnessing a potential failure of primality at a, there exist $t, w, e, f \in M$ with*

$$a = tw, \qquad t = ef, \qquad e \mid b, \qquad f \mid c, \qquad \rho(w) \prec \rho(a)$$

(that is: a factors into an exact ambient routed local block t — itself splitting as $e \mid b$, $f \mid c$ — and a complementary residual w of strictly lower structural rank than a). Then every element of M is primal, and hence every product equality in M admits a four-factor refinement.

Proof idea. This is a structural induction on the well-founded rank ρ. The base case is immediate: elements of the base regime are primal by hypothesis. For the inductive step, consider a non-base element a and a divisibility $a \mid bc$ witnessing a potential failure of primality at a. By hypothesis this produces t, w, e, f with $a = tw$, $t = ef$, $e \mid b$, $f \mid c$, and $\rho(w) \prec \rho(a)$. By the induction hypothesis (well-foundedness of ρ guarantees no infinite strictly-decreasing chain), the residual w is primal, and the exact routing $t = ef$ with $e \mid b$, $f \mid c$ transports w's primality back through $a = tw$ to a full routing of $a \mid bc$, establishing that a is primal. Since a was an arbitrary non-base element, every element of M is primal, and Law 1 then gives global four-factor refinement. □

Schematically:

$$\text{primal base} + \text{exact routed local block} + \text{strict well-founded residual descent}$$
$$\implies \text{every element primal} \implies \text{four-factor refinement.}$$

Note that the positive-integer control of Section 3 instantiates this schema in its simplest possible form, with a single base step and no descent needed: the gcd/coprime block of Proposition 1 *is* the exact routed local block, applied once.

9.1 Lean formalization

The Route-Closure Induction Proposition, extracted from the audited chain of Section 8, was independently formalized in Lean 4 (v4.31.0) against the locked Conway algebraic interface, in `04_PROOF_MAP/lean/RouteClosureInduction.lean`, as two theorems:

- `UNNS.CommonRefinement.forall_isPrimal_of_wellFoundedRouteClosure`
- `UNNS.CommonRefinement.hasFourFactorRefinement_of_wellFoundedRouteClosure`

Independent kernel verification of this formalization is reported in Section 10. As stated explicitly there, this is a scoped result: an independent formalization of the *extracted* route-closure induction mechanism, not a from-scratch reimplementation of the complete concrete Hahn-series/surreal machinery underlying the original candidate proof.

10 Verification of the Conway Candidate Mechanism

The mathematical content of Sections 8–9 is deliberately kept separate here from its *verification*, so that the manuscript never conflates "Principal Theorems I–III are proved" (Sections 4 –6, self-contained proofs above) with "the candidate transfinite realization has been independently checked" (this section). The latter is a verification protocol executed against an external source repository, not a theorem of this manuscript.

Figure 7 gives a compact overview of the seven-item protocol before Table 2 reports its full detail.

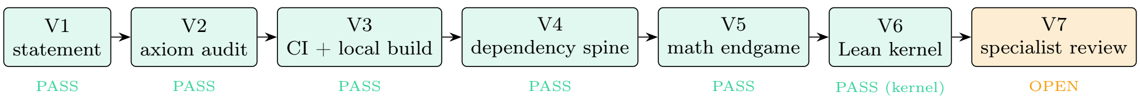

Figure 7: Verification evidence chain. Items V1–V6 close with PASS, including the independently Lean-kernel-verified extracted mechanism (V6); only V7 (independent specialist mathematical review) remains open.

10.1 Verification protocol and outcomes

Table 2 summarizes the seven-item verification protocol (V1–V7) executed against `gaearon/conway-refinement` at the pinned commit `264445c9...`, using Lean 4.31.0, Mathlib (`fabf563a...`), and the project's Combinatorial Games dependency (`3c6dcdbc...`).

Table 2: Conway candidate mechanism verification protocol (V1–V7).

Check	Outcome
V1 Statement fidelity	PASS
V2 Axiom / trust audit	PASS (13,114 axiom declarations audited; only `propext`, `Classical.choice`, `Quot.sound` allowed)
V3 Exact-commit CI build	PASS (3,145 `lake` build jobs, CI conclusion *success*)
V3 Independent local targeted final-theorem build	PASS (2,649 jobs, target `ConwayRefinement.Surreal.OmnificInteger.Refinement.ConwayRefinement`, clean checkout)
V4 Dependency-spine audit	PASS (51 statement declarations checked, 32 proof links verified)
V5 Independent mathematical endgame	PASS
V5 Independent Lean endgame compilation (standalone)	WRITTEN; standalone minimal-environment kernel build not separately executed
V6 Critical transfinite step source reconstruction	PASS
V6 Independent Lean formalization of the extracted critical mechanism	**PASS — kernel verified** (see below)
V7 Independent specialist mathematical review	**OPEN — external**

10.2 Kernel-verified formalization of the extracted mechanism

The route-closure induction formalization of Section 9 (`RouteClosureInduction.lean`) was independently kernel-compiled:

```
[837/837] Built ConwayRefinement.UNNS.RouteClosureInduction; Build
completed successfully (837 jobs).
[829/829] Replayed ConwayRefinement.UNNS.RouteClosureInduction; Build
completed successfully (829 jobs).
```

Both the full target build (837 jobs) and the cached recheck/replay (829 jobs) completed successfully, with only informational linter notices from diagnostic `#check` commands and no compilation failure. As stated in Section 9, the explicit scope of this verification is the *extracted* route-closure induction mechanism against the locked Conway algebraic interface, not a from-scratch reimplementation of the full concrete Hahn-series/surreal machinery.

10.3 What remains open

Of the seven verification items, six (V1–V6, including the independently Lean-kernel-verified extracted mechanism) are closed. The single remaining required confirmation is **V7, independent specialist mathematical review** of the Hahn-series/surreal realization by a domain expert external to the project. Two further items are noted as optional strengthenings rather than required closures: a standalone minimal-environment kernel build of the independent endgame formalization, and a from-scratch independent reimplementation of the concrete Hahn/surreal critical mechanism as a mini-project. The project's internally assessed status is therefore: *mechanically verified candidate proof, with independent local* *final-theorem build, kernel-verified independent formalization of the extracted route-closure induction principle, and a strong source/statement/axiom audit*, pending only external specialist review.

11 Structural Extension — Admissibility and Percolation

The results above are entirely internal to factorization algebra: they establish when route closure holds and, when it fails, what its failure looks like. A natural further question, addressed by the `UNNS_ADMISSIBILITY_PERCOLATION` extension project, is whether losing route closure has any consequence for a system's behavior under two further UNNS structural diagnostics: perturbative admissibility and percolative connectivity. The result is a corollary to the structural interpretation established above, not a further theorem of factorization:

$$\text{route closure} \neq \text{perturbative admissibility} \neq \text{percolation connectivity scale.}$$

11.1 The frozen 30-system corpus

A frozen 30-system pilot corpus was constructed from five classes of six systems each — `INT_FREE`, `R1_FREE`, `R1_FAIL`, `AFF_FREE`, `AFF_FAIL` — carrying 18 `TRACEABLE` and 12 `NON_TRACEABLE` route-closure labels drawn from Theorems 1–3. Three coordinates were then measured on every system:

(i) *Route closure* (native property: primal factor traceability; failure: exact route obstruction), inherited from Sections 4–5;

(ii) *Perturbative admissibility*, via the STRUC-I observable A_κ (does the ordered ladder remain inside its structural vulnerability budget under perturbation?);

(iii) *Percolative connectivity scale*, via the STRUC-PERC-I observable $\kappa_{\text{connect\_exact}}$ (at what normalized scale does the gap-vulnerability graph become connected?).

11.2 STRUC-I result: admissibility is saturated

Across all 30 systems, STRUC-I admissibility was measured at $\text{mean}(A_\kappa) = \min(A_\kappa) = A_\kappa$ at $\kappa_{\max} = 1$, identically, *regardless of `TRACEABLE`/`NON_TRACEABLE` label.* Because A_κ takes the constant value 1 across the entire corpus, it carries no variance for route-closure status to correlate with, so the correct statement is not that the two are "uncorrelated" (a claim a constant variable cannot support) but the stronger and more precise one: *route-closure status does not separate STRUC-I admissibility in this corpus*, because STRUC-I admissibility is already saturated at its maximum value for every system tested, including all 12 systems that are provably route-defective by Theorems 1 and 3. Loss of route closure carries no admissibility consequence in this corpus.

11.3 Original STRUC-PERC-I v2.5.0

The percolative-connectivity coordinate was first measured with instrument version 2.5.0, which reports an original connectivity-scale observable κ_{connect} together with a scale-threshold eventual-connectivity verdict (FULL_PERCOLATION, TAIL_FRAGMENTATION, or HARD_FRAGMENTATION) computed from an interquartile-range gap scale. The exact-threshold observable $\kappa_{\text{connect\_exact}}$ used throughout the rest of this section is *not* a v2.5.0 quantity: it is introduced by the later v2.5.1 audit/repair instrument described next, and reported here retroactively for all 30 systems once that instrument exists. Applied to the 30-system corpus, v2.5.0 returned FULL_PERCOLATION for 25 of the 30 systems — all 12 rank-one systems (R1_FREE, R1_FAIL) and 13 of the 18 higher-dimensional systems — and reported the remaining 5 higher-dimensional systems as HARD_FRAGMENTATION (INT03, AFF01, AFF04, AFN06) or TAIL_FRAGMENTATION (AFN01), suggesting — at face value — a localized rather than pervasive eventual-connectivity failure outside rank one.

11.4 Audit and repair: v2.5.1

That face-value reading did not survive an internal audit. Figure 8 summarizes the resulting audit flow and its outcome.

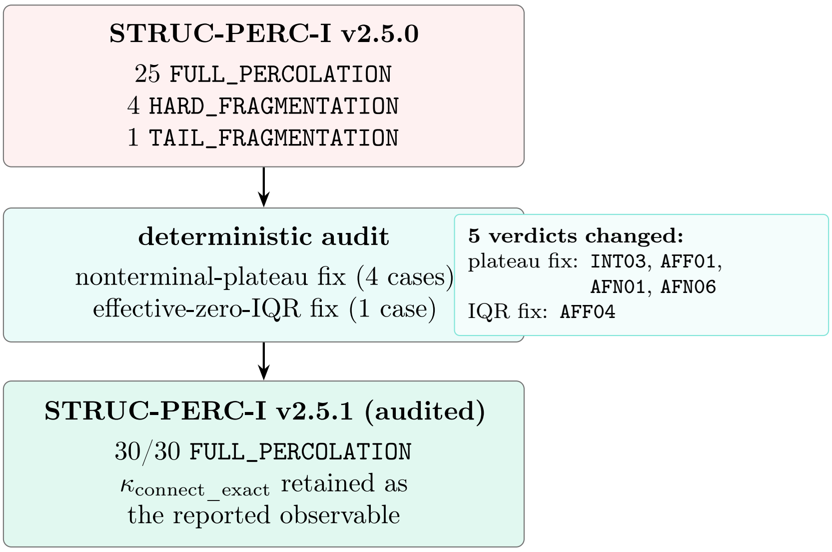

Figure 8: The v2.5.0-to-v2.5.1 audit flow. v2.5.0's coarser interquartile-range scale and plateau-stopping rule mistook a finite-scale threshold effect in five higher-dimensional systems for permanent fragmentation; the v2.5.1 audit/repair instrument applies two deterministic fixes and every corrected system reaches FULL_PERCOLATION, with no other verdict disturbed.

Instrument version 2.5.1 adds an exact connectivity probe that continues the search past the scale-threshold cutoff used by v2.5.0 whenever a hard cutoff, rather than a true absence of connectivity, is suspected, and this is what actually accounts for the five changed verdicts. Four of the five (INT03, AFF01, AFN01, AFN06) trace to a nonterminal-plateau fix in the adaptive connectivity-extension search: v2.5.0 treated an early plateau as a terminal fragmentation signal, whereas v2.5.1 treats a plateau as a marker only and continues to an exact threshold or to exhaustion of the scheduled probes. The fifth, AFF04, traces to a separate numerical fix: its raw interquartile gap scale was effectively zero (within numerical tolerance), which v2.5.1's median-fallback rule resolves. The remaining 25 verdicts, already FULL_PERCOLATION under v2.5.0, are unchanged. This audit-and-repair episode is itself part of the record: every $\kappa_{\text{connect\_exact}}$ value reported below is the exact, v2.5.1-audited value, and the corrected instrument is what the analytics dashboard and this manuscript both cite (Section 14).

11.5 Exact connectivity-scale result

Table 3 and Figure 9 report the audited $\kappa_{\text{connect\_exact}}$ class statistics.

Table 3: $\kappa_{\text{connect\_exact}}$ by class across the 30-system corpus (6 systems per class), under the audited v2.5.1 instrument.

Class	label	mean	median	min	max
`R1_FREE`	TRACEABLE	0.000	0.000	0.000	0.000
`R1_FAIL`	NON_TRACEABLE	1.000	1.000	1.000	1.000
`INT_FREE`	TRACEABLE	5.423	5.139	1.270	10.412
`AFF_FREE`	TRACEABLE	5.794	3.133	1.707	15.074
`AFF_FAIL`	NON_TRACEABLE	3.225	2.475	1.112	8.400

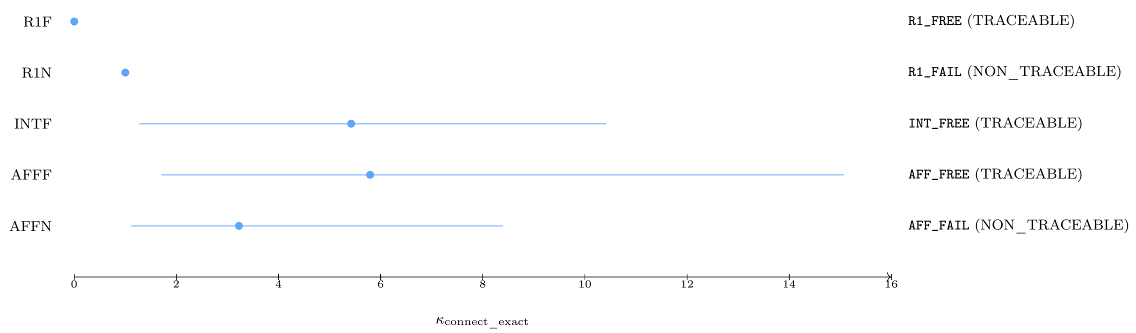

Figure 9: $\kappa_{\text{connect\_exact}}$ by class: min–max range (bar) with class mean (dot), audited v2.5.1 values. The rank-one classes give a clean, zero-spread, non-overlapping signature; the higher-rank affine classes overlap and do not preserve the same ordering.

The rank-one classes give an exact, reproducible, zero-spread signature: $\kappa_{\text{connect\_exact}} = 0$ for 6/6 route-closed (`R1_FREE`) systems, versus $\kappa_{\text{connect\_exact}} = 1$ for 6/6 route-defective (`R1_FAIL`) systems, with no overlap and no within-class variance.

11.6 Rank-one signature and affine non-generalization

This clean rank-one contrast does *not* extend to a monotone route-closed/route-defective ordering at higher affine rank: `AFF_FREE` (route-closed) has *higher* mean and maximum $\kappa_{\text{connect\_exact}}$ than `AFF_FAIL` (route-defective), and both affine classes show substantial within-class spread and mutual overlap (Table 3, Figure 9). The secondary geometric signature that route structure leaves in connectivity scale is therefore real at rank one but domain-specific rather than universal.

11.7 Three distinct, non-equivalent coordinates

Figure 10 summarizes the section's overall structural finding: route closure, perturbative admissibility, and percolative connectivity scale behave as three separately measurable coordinates of a system, not as a single axis with route closure at one end. The three coordinates can interact — the rank-one connectivity-scale split tracks route closure exactly, for instance — but they are not equivalent: no one of them is a stand-in for the others, and (as the affine case shows) the tracking itself is not universal.

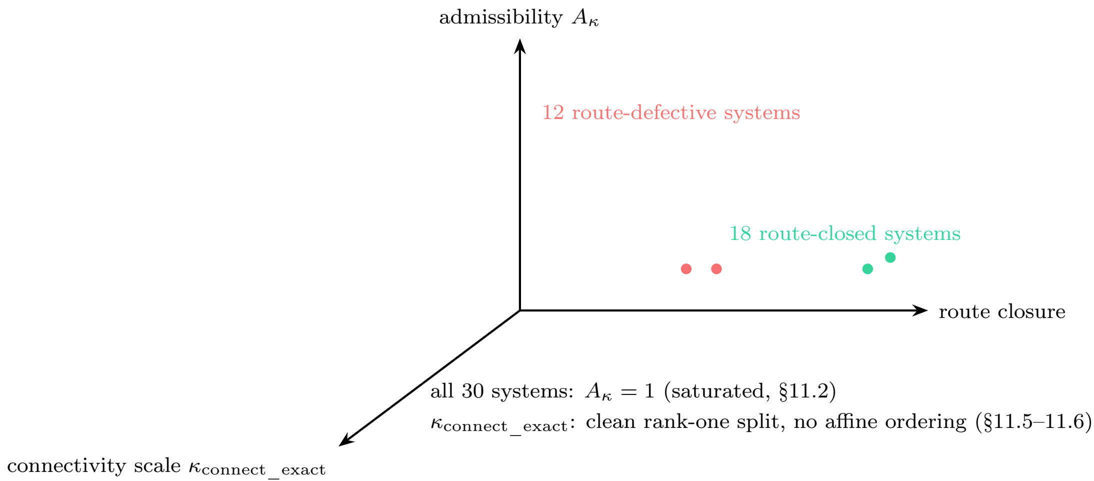

Figure 10: Three-coordinate UNNS structural space (schematic). Route closure (horizontal), perturbative admissibility (vertical), and percolative connectivity scale (diagonal) are measured as distinct, non-equivalent coordinates on the same 30-system corpus; admissibility is saturated on this axis for every system, while connectivity scale separates cleanly at rank one but not at higher affine rank.

11.8 Combined interpretation

The extension project's own combined conclusion, which this manuscript adopts, is:

> Primal factor traceability governs exact algebraic route closure, but route closure is not a universal proxy for structural admissibility or percolative connectivity. A system may be algebraically route-defective while remaining perturbatively admissible and eventually connected in gap space. The algebraic defect can nevertheless leave secondary geometric signatures, as the exact rank-one connectivity-scale contrast demonstrates.

In compact form: `UNNS_COMMON_REFINEMENT` identifies the exact law of algebraic route closure (Sections 2–10); `UNNS_ADMISSIBILITY_PERCOLATION` locates that law inside a broader, multi-coordinate structural landscape. Non-refinement does not, by itself, mean non-admissibility or permanent fragmentation; route closure is one coordinate of a larger structural state space, not a stand-in for the others.

12 Limitations and Open Problems

- **Independent specialist mathematical review** of the Hahn-series / surreal realization underlying the candidate Conway/omnific mechanism (Section 8) remains the single required open item in the verification protocol (V7, Table 2). The present manuscript's transfinite claims should be read at the verification level stated in Section 10: mechanically verified candidate proof with an independently Lean-kernel-verified extracted mechanism, not yet externally peer-reviewed at the level of the full Hahn-series/surreal construction.

- **The Canonical Refinement Problem.** Where a common refinement exists (Sections 3–5), this manuscript does not address whether there is a *structurally privileged* refinement among possibly several valid ones; this is identified as an open branch beyond the present existence/traceability investigation, and is closely related to uniqueness questions studied for refinement monoids in general [5].

- **Bounded-scale percolation and connectivity-threshold geometry.** Section 11 reports the exact connectivity scale $\kappa_{\text{connect\_exact}}$ under the audited v2.5.1 instrument but does not develop bounded-scale percolation behavior, component-growth profiles, or a theory of when the rank-one connectivity signature should versus should not be expected to generalize; these are identified as separate, potentially interacting structural coordinates for future work.

- **Broader non-affine systems.** The present theorems are stated for additive submonoids of $\mathbb{N}_0$ (Theorem 1) and positive affine monoids (Theorems 3–4). Extending the classification to non-affine finitely generated cancellative monoids, and relating it to the general non-unique factorization theory [11] and to Schreier/pre-Schreier structure beyond the affine setting [7], is noted as an optional extension rather than a gap in the results reported here.

- **Scope of the computational validation.** The three computational checks reported in Sections 3, 4, and 5 are of three different kinds, and it matters which is which. The 500 integer-quadruple cases (Section 3) are a *randomized regression corpus*: quadruples drawn at random subject to $ab = cd$, not an exhaustive search over any stated range, with zero exceptions to Proposition 1 observed across the sample. The 793 rank-one generator families (Section 4) are *exhaustive within their stated finite generator range* (generators up to 12): within that range every family was checked, with zero exceptions to Theorem 1. The affine work (Section 5) rests on a *fixed seven-system control/certificate suite* (Table 1), not a search at all, chosen to instantiate each qualitatively distinct affine case (free, non-normal, normal with different atom counts). None of these three corpora, by construction, bounds the behavior of families with much larger generators, much higher affine rank, or outside the randomized sample's distribution. No mechanism in the proofs of Theorems 1–4 depends on the size of the generators or on rank beyond $r \geq 2$, so no range- or sample-dependent failure is expected; this is a statement about the honesty of the evidence's scope, not about a suspected gap in the theorems themselves, which are proved unconditionally above and do not depend on any of these corpora for their validity.

13 Conclusion — From Route Defect to Route Closure

This manuscript has traced a single structural question — what survives, and what fails, when an equal-endpoint identity $ab = cd$ is asked to upgrade to a common four-factor refinement — across four regimes of increasing structural complexity, and across one further structural coordinate. The complete hierarchy established is:

LAW: factor traceability $\iff$ route closure
POSITIVE-INTEGER CONTROL: the gcd/coprime witness always refines
PRINCIPAL THEOREM I: rank-one closure $\iff$ $\mathrm{RCI}(H) = 1$ (cofinal primal tail otherwise)
PRINCIPAL THEOREM II: affine closure $\iff$ $\mathrm{ARD}(H) = 0$
PRINCIPAL THEOREM III: affine defect can persist along an infinite ray
TRANSFINITE MECHANISM: strict well-founded descent prevents a persistent residual channel

Two consequences follow directly from Principal Theorems I–III. First, **route defect is regime-dependent, not a matter of degree**: it is always finite and self-healing at rank one (Theorem 1, Theorem 2), while at higher affine rank it is permanent and unbounded (Theorem 4) — a sharp qualitative discontinuity. Second, **the transfinite setting is not a loophole for the higher-rank failure mode**: the audited Conway/omnific chain, and its independently Lean-kernel-verified extracted induction mechanism, show explicitly how a persistent channel of the kind Theorem 4 guarantees at finite affine rank is foreclosed *within the audited candidate realization* once well-founded descent to a primal base is available — the extracted induction mechanism is independently kernel-verified, while independent specialist review of the complete Hahn-series/surreal realization remains the open item recorded in Section 12. A third consequence follows from the structural extension: **route closure is a coordinate, not a proxy**: Section 11 shows that losing it need not cost a system its perturbative admissibility or its eventual percolative connectivity, even as it can, in the cleanest case tested, leave an exact and reproducible trace in that connectivity's scale.

All theorem-level results (Sections 3–6) are proved in full, not sketched, and computationally validated within this manuscript. The transfinite mechanism (Sections 8–10) is reported at its actual, scoped verification status: an audited candidate realization with an independently kernel-verified extracted mechanism, pending external specialist review. The structural extension (Section 11) is reported as a corollary to the structural interpretation established by the law and the three theorems, together with the v2.5.0-to-v2.5.1 instrument audit that corrected five of its thirty connectivity verdicts, all to `FULL_PERCOLATION`. Exploration of the open problems in Section 12, and in particular independent specialist review of the transfinite realization, is identified as the priority direction for future work.

14 Data, Code, Analytics, and Reproducibility

All computational records, formal verification material, frozen chamber instruments, source audits, derived datasets, and synthesis documents supporting this study are available in the `UNNS_COMMON_REFINEMENT` project repository [1], including the `UNNS_ADMISSIBILITY_PERCOLATION` extension used in Section 11. This includes: the positive-integer control records underlying Section 3 (`01_INTEGER/output/`); the rank-one and affine validation records underlying Theorems 1–4 (`outputs/records/`, `04_PROOF_MAP/output/`); the Lean 4 formalization and kernel build evidence underlying Section 10 (`04_PROOF_MAP/lean/`, `outputs/records/ROUTE_CLOSURE_KERNEL_VERIFICATION.md`); and the 30-system admissibility–percolation corpus, both v2.5.0 and audited v2.5.1 STRUC-PERC-I outputs, and its STRUC-I chamber outputs underlying Section 11 (UNNS_ADMISSIBILITY_PERCOLATION/corpus/, UNNS_ADMISSIBILITY_PERCOLATION/outputs/). An interactive, self-contained reproducibility instrument, UNNS_COMMON_REFINEMENT_ANALYTICS.html, generated directly from these same frozen records, is included at the repository root as of the commit recorded in [1]. The analytics includes SHA-256 provenance for the frozen project records used in its build and provides source drawers and claim–evidence mappings for the principal results reported here. At the time of writing, the builder script that produces this analytics page from the frozen records (analytics/BUILD_ANALYTICS.py) has not yet been committed to the repository; only the generated, self-contained HTML page is currently public, and this section cites that page rather than its builder accordingly. The candidate transfinite mechanism audited in Sections 8–10 is external source material, gaearon/conway-refinement [2], pinned at the commit stated therein for reproducibility of the verification protocol.

References

[1] UNNS Substrate Research Program. *UNNS Common Refinement: Common Refinement and Structural Route Closure in the UNNS Substrate.* GitHub repository, 2026. https://github.com/ukbbi/UNNS/tree/main/projects/UNNS_COMMON_REFINEMENT Repository state used in this work: commit 760a81d6df84a89fcf293e27bba6c49c1af04ecf (includes the generated UNNS_COMMON_REFINEMENT_ANALYTICS.html reproducibility instrument at the repository root). Accessed 20 September 2026.

[2] gaearon. *conway-refinement*: Lean 4 formalization of Conway/omnific refinement mechanisms. GitHub repository, 2026. Commit audited in this work: 264445c93b78554c408e99e4e7f663693b4e91ab.

[3] J. H. Conway. *On Numbers and Games.* Academic Press, London, 1976.

[4] H. Gonshor. *An Introduction to the Theory of Surreal Numbers.* London Mathematical Society Lecture Note Series 110. Cambridge University Press, Cambridge, 1986.

[5] K. R. Goodearl. *Partially Ordered Abelian Groups with Interpolation.* Mathematical Surveys and Monographs 20. American Mathematical Society, Providence, RI, 1986.

[6] P. A. Grillet. *Commutative Semigroups.* Advances in Mathematics 2. Kluwer Academic Publishers, Dordrecht, 2001.

[7] P. M. Cohn. Bezout rings and their subrings. *Mathematical Proceedings of the Cambridge Philosophical Society*, 64(2):251–264, 1968.

[8] P. M. Cohn. *Free Rings and Their Relations*, 2nd edition. Academic Press, London, 1985.

[9] J. C. Rosales and P. A. García-Sánchez. *Numerical Semigroups.* Developments in Mathematics 20. Springer, New York, 2009.

[10] W. Bruns and J. Gubeladze. *Polytopes, Rings, and K-Theory.* Springer Monographs in Mathematics. Springer, New York, 2009.

[11] A. Geroldinger and F. Halter-Koch. *Non-Unique Factorizations: Algebraic, Combinatorial and Analytic Theory*. Pure and Applied Mathematics 278. Chapman & Hall/CRC, Boca Raton, 2006.

[12] L. de Moura and S. Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, Lecture Notes in Computer Science, vol. 12699, pages 625–635. Springer, 2021.

[13] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020.