The Gate-2 Inspection over the Transport Construction

Reachability of Admissible Degeneracy Geometry, and the Execution of the Registered K = 7 Search

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

General Reader Summary

The previous paper ran the first of the locality inspection's checks and found that the admissibility rules of the transport construction are local *in form*: read off one at a time, no rule secretly depends on transport history or on global bookkeeping. That cleared the rule most feared to hide a long-range dependence — the one governing how records are conserved — and put the programme on the local-form branch.

But "the rules are local" and "the geometries the rules permit can actually be built" are two different statements, and the gap between them is exactly where this paper works. A blueprint can satisfy every building code line by line and still describe a structure that cannot be assembled — the code checks each rule pointwise, while buildability is a question about whether the pieces fit together at all. A degeneracy geometry can satisfy every local admissibility rule and still fail to arise from any admissible transport. The question this paper makes precise is therefore:

If a geometry satisfies all the local admissibility rules, must there exist an admissible local transport that produces it?

This is Gate 2. The paper has two parts. **Part I** sets up the question and the procedure: it shows that Gate 1 and Gate 2 are two halves of one question (Gate 1 proved the local rule is *sound* — everything buildable obeys it; Gate 2 asks whether it is *complete* — everything that obeys it is buildable), reduces the whole matter to whether a single set of "unbuildable-but-admissible" geometries is empty, locates where such a geometry could hide, fixes the form of the missing condition that would settle it, and — in the discipline the programme adopted — writes down in advance the exact search that would decide it. **Part II** runs that search on the concrete K = 7 geometry. At the smallest scale it passes cleanly: the rules force a single alternating pattern that is automatically conservation-consistent and is realized by a known transport. At the next scale — hubs glued along shared boundaries — a genuine new question appears, the consistency of a sign passed from hub to hub around a loop; we prove the criterion that decides it, prove that the obvious bookkeeping freedom can never cause a failure, and reduce the remaining content to three named, clearly-flagged hypotheses about the shared-boundary rule. Under those hypotheses the registered search contains no failure witness. We state the hypotheses as hypotheses and do not claim them as proved.

Abstract

This paper inspects Gate 2 of the locality decision for the transport construction — whether every germ satisfying the local admissibility predicate L is realized by an admissible local transport — and executes the resulting search on the concrete $K = 7$ geometry. It is in two parts.

Part I (the reduction and the pre-registered search). *The Locality Inspection* returned Gate 1 PASS (conditional on four-sector inventory completeness): every admissibility constraint classified as pointwise or germ-local, and closure-current conservation was certified local in form. We recast Gate 2 as a statement about the **total realizer map** $\rho : \mathcal{T}_L \to \mathcal{D}$: writing $\{L\}$ for the L-satisfying germs and $\mathrm{Im}(\rho)$ for the realizable ones, Gate 1 established **soundness** $\mathrm{Im}(\rho) \subseteq \{L\}$, and Gate 2 is its converse, the **completeness** $\{L\} \subseteq \mathrm{Im}(\rho)$ (Theorem 4.2). The N_2 witnesses are the **unrealizable residue** $\mathfrak{U} \coloneqq \{L\} \setminus \mathrm{Im}(\rho)$, and Gate 2 holds iff $\mathfrak{U} = \emptyset$ (Theorem 5.2). We localize where $\mathfrak{U}$ could live and harden the localization into a conditional reduction (§6): we fix the form of the intrinsic integrability predicate as a four-sector conjunction $I = I_{inc} \wedge I_{hub} \wedge I_{circ} \wedge I_{comp}$, isolate the structural requirements R1–R3 and the matching property, and prove that for any predicate of that form the residue is exactly the hub-completion-failing locus, decomposes sector-wise, and Gate 2 reduces to per-sector hub-completion on $\{L\}$ (Theorem 6.9, [Conditional-on-supply]). The single remaining input is the four cell-completion relations themselves (Open Problem 6.7). We record the non-unique realizer regime (favourable to reachability, Proposition 7.1; raising the falsification burden, Proposition 7.2), prove a minimal unrealizable germ exists if any does (Theorem 8.2), and deposit the pre-registered Gate-2 search $\mathcal{P}_2$ in three strata (Pre-Registration 9.1).

Part II (the execution). Using the concrete $K = 7$ forms — $A_{circ}(\lambda) = (\sum_i \lambda_i)^2$ and $A_{comp}(\lambda) = \sum_i(\lambda_i + \lambda_{i+1})^2$ on the six spoke amplitudes — we run the registered search. **Single-hub stratum:** A_{comp} is a positive sum-of-squares whose kernel on the even cycle C_6 is exactly the alternating ray $\lambda_i = (-1)^i a$, on which A_{circ} vanishes automatically ($N = 6$ even) and the σ-family supplies a realizer; so $\mathfrak{U} \cap \mathcal{P}_2^{\{single\text{-}hub\}} = \emptyset$ (Part II, Theorem 3.6, [Conditional] on C1–C3, one realizing family sufficing for the pass). **Hub-adjacent stratum:** each hub's alternating mode carries a free per-hub sign; the shared-boundary matching defines a $\mathbb{Z}_2$ cochain s on the hub-adjacency graph, and the stratum passes iff s is balanced ($[s] = 0 \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$; Part II, Theorem 4.3, [Proven]). Fixing the construction's physical gauge, $s = \varepsilon^{\{phys\}}\varepsilon^{\{phys\}}\cdot t$; we prove the ε-part is a coboundary, so the verdict depends only on the residue class $[t]$ (Theorem II-5.3, [Proven]) — the hexagon-relabelling freedom is proven harmless. Under **(PBM)** physical-boundary matching the class is trivial (Shared-Boundary Sign Lemma 6.2, [Conditional on PBM]); under **(PBM ∧ AC ∧ AA)** — adding amplitude compatibility and anti-aligning match — the competition stratum also passes (Part II, Lemma 7.2). Assembling: $\mathfrak{U} \cap \mathcal{P}_2 = \emptyset$ [Conditional on C1–C3, PBM, AC, AA] — **no N_2 witness in the registered finite $K = 7$ search**, conditional

on those named hypotheses, which we do not derive. A registered-search pass is a global Gate-2 pass only under Conjecture 8.3; Gate 3 (global assembly) remains the next open gate.

Labelling convention: [Proven], [Conditional], [Conditional-on-gluing], [Conditional-on-supply], [Conditional on PBM], [Conditional on PBM ∧ AC ∧ AA], [Methodological], [Conjectural], [Open — requires inventory], [Open — requires search]. Objects are numbered within each part: Part I objects carry bare section-local numbers (Theorem 6.9, Remark 6.1, …), and Part II objects carry a "II-" prefix (Theorem II-5.3, Remark II-6.1, …) so that the two parts' identically-numbered objects never collide. Cross-references that span the parts name the part explicitly (Part I, Theorem 6.9). References to *The Locality Inspection* and other external papers in the programme are named as such.

Contents

Part I — The Reduction and the Pre-Registered Search

1. From Local Form to Geometric Realization

The Locality Inspection executed Stage 1 of the Inspection Protocol against the transport construction and returned **Gate 1 PASS** on the local-form branch, conditional on four-sector inventory completeness. Every admissibility constraint in the inventory {$\mathcal{C}$_inc, $\mathcal{C}$_hub, $\mathcal{C}$_circ, $\mathcal{C}$_comp} classified as pointwise or germ-local; the decisive constraint, closure-current conservation $\mathcal{C}$_circ, was certified local in form at substrate and continuum levels, with the additivity axiom *forbidding* the non-local kernel a holonomy reading would require. The conservative output N_1 was excluded.

That result settled the **form** of the rules. It did not settle the **reach** of the geometries those rules permit. The two are distinct, and the distinction is the entire subject of this paper.

A local admissibility predicate L determines which germs are *permitted*: those satisfying every local constraint. Reachability determines which permitted germs are *produced* — which actually arise as the degeneracy datum of some admissible transport. Permission is a property of the germ read against the rules; production is the existence of a transport. Nothing in Gate 1 guarantees the two coincide, and the programme's locality verdict turns on whether they do.

The question advances accordingly:

Gate 1. Are the admissibility rules local in form? — Answered: yes (conditional on completeness). **Gate 2.** Are the locally admissible germs reachable by admissible transport? — The present question.

We make Gate 2 precise, reduce it to the emptiness of a single well-defined object, locate where that object could be non-empty, and pre-register the search that decides it. We do not assume reachability, and we do not prove it; we isolate exactly what would refute it and commit, in advance, to how we would look.

---

2. The Output of Gate 1, Recalled

The carrier is the germ. Because $\mathcal{C}$_hub, $\mathcal{C}$_circ, and $\mathcal{C}$_comp classified as germ-level (Gate-1 execution), the germ amendment of *The Locality Decision Theorem* Definition 3.5 is in force: the admissibility carrier is the germ $\hat{G} \in \mathcal{D}$, not the pointwise datum G(R). All predicates below are read at germ level.

The **local admissibility predicate** delivered by Gate 1 is

$$L(\hat{G}) \Leftrightarrow \mathcal{C}\_inc(\hat{G}) \wedge \mathcal{C}\_hub(\hat{G}) \wedge \mathcal{C}\_circ(\hat{G}) \wedge \mathcal{C}\_comp(\hat{G}),$$

and we write its extension

$$\{L\} \coloneqq \{ \hat{G} \in \mathcal{D} : L(\hat{G}) \}$$

for the set of **locally admissible germs**. Gate 1's content is that {L} collects the full local content of admissibility: under Gate 1 PASS, every admissibility constraint is one of the conjuncts of L, so a germ that is the degeneracy datum of an admissible transport necessarily lies in {L}. We isolate that fact in Section 3 as the soundness inclusion; it is the hinge on which the Gate-2 reframing turns.

---

3. The Total Realizer Map

The prior papers defined, for each germ G, the set $\mathcal{R}$(G) of its local realizers and a per-fiber map $\rho : \mathcal{R}(G) \to \mathcal{D}$ that is constant equal to G on that fiber. For Gate 2 we need the realizer relation as a single object over all germs at once. We extend ρ faithfully.

Definition 3.1 (locally admissible transport). A **locally admissible transport** is a transport T defined on a neighbourhood of a base point, satisfying the local-class constraints there. Write $\mathcal{T}\_L$ for the space of locally admissible transports.

Definition 3.2 (total realizer map). The **total realizer map**

$$\rho : \mathcal{T}\_L \to \mathcal{D}, \rho(T) = \text{(the degeneracy germ T produces)},$$

sends each locally admissible transport to the germ it realizes. Its fibers recover the prior paper's realizer sets,

$$\mathcal{R}(\hat{G}) = \rho^{-1}(\hat{G}) = \{ T \in \mathcal{T}\_L : \rho(T) = \hat{G} \},$$

and on each fiber ρ restricts to the constant $\hat{G}$ of *The Locality Inspection* Definition 2.3. The **image**

$$\mathrm{Im}(\rho) = \{ \hat{G} \in \mathcal{D} : \mathcal{R}(\hat{G}) \neq \emptyset \} = \{ \text{realizable germs} \}$$

is the set of germs produced by some locally admissible transport.

Remark 3.3 (consistency with the prior map) [Proven]. The extension is faithful: $\mathcal{R}(\hat{G}) = \rho^{-1}(\hat{G})$ and $\rho|_{\mathcal{R}(\hat{G})} \equiv \hat{G}$ reproduce the prior definition exactly, so every statement of *The Locality Inspection* about fibers is preserved. The only new structure is the assembly of the fibers into one map, which is what lets Gate 2 be stated as a property of ρ rather than of each fiber separately.

---

4. Reachability as Completeness of L

The reframing that organizes this paper is that Gate 1 and Gate 2 are the two inclusions of a single set equality.

Theorem 4.1 (soundness of L) [Proven, inherited from Gate 1]. Under Gate 1 PASS,

$\mathrm{Im}(\rho) \subseteq \{L\}$.

Proof. Let $\hat{G} \in \mathrm{Im}(\rho)$, so $\hat{G} = \rho(T)$ for some $T \in \mathcal{T}_L$. By Definition 3.1, T satisfies every local-class constraint on its neighbourhood; under Gate 1 PASS the local-class constraints are all of $\{\mathcal{C}_{inc}, \mathcal{C}_{hub}, \mathcal{C}_{circ}, \mathcal{C}_{comp}\}$, whose germ-level conjunction is L. Reading these on the germ T produces gives $L(\hat{G})$. Hence $\hat{G} \in \{L\}$.

In words: every realizable germ is locally admissible. This is the necessity direction Gate 1 supplies — L is a *sound* predicate for realizability, an over-approximation that no realizable germ escapes.

Theorem 4.2 (Gate 2 is the completeness of L) [Proven]. Gate 2 holds iff L is **complete** for realizability:

Gate 2 $\Leftrightarrow \{L\} \subseteq \mathrm{Im}(\rho) \Leftrightarrow \{L\} = \mathrm{Im}(\rho) \Leftrightarrow$ ρ surjects onto $\{L\}$.

Proof. Gate 2 is the statement that every L-satisfying germ has a realizer, i.e. $\hat{G} \in \{L\} \Rightarrow \mathcal{R}(\hat{G}) \neq \emptyset \Rightarrow \hat{G} \in \mathrm{Im}(\rho)$; that is precisely $\{L\} \subseteq \mathrm{Im}(\rho)$. Combined with the soundness inclusion (Theorem 4.1) this is the equality $\{L\} = \mathrm{Im}(\rho)$. Surjectivity of ρ onto $\{L\}$ is the same statement, since $\mathrm{Im}(\rho) \subseteq \{L\}$ already holds, so $\mathrm{Im}(\rho)$ meets $\{L\}$ in all of $\{L\}$ iff it equals $\{L\}$.

This is the deliverable's spine. Gate 1 proved one inclusion; Gate 2 *is* the other. The locality question is not two unrelated checks but the soundness and completeness of one local predicate against one realizability relation — and the gap between them, when non-empty, is the precise content of reachability non-locality.

Remark 4.3 (what the reframing does and does not buy) [Methodological]. The soundness/completeness framing does real work, but we are careful about how much. First, it pins the *direction* of the Gate-2 argument: a clean Gate-2 pass is a surjectivity proof, exactly the [Proven] upgrade condition of *The Locality Inspection* Procedure 4.2, and anything short of it is a [Conditional] failed search. Second, it makes the asymmetry between the gates legible: soundness was decidable by reading the constraints (Gate 1 is classification), whereas completeness is an existence claim over a geometry space (Gate 2 is search), which is why the gates are heterogeneous in method and why Principle 3.7 bites only here and at Gate 3. What the reframing does **not** by itself supply is a characterization of the object to search: naming it $\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$ is Definition 5.1, not a description of *where in {L}* a witness sits. That characterization is the work of Section 6 — and Section 6 delivers it only heuristically (a sector-weighted localization), not as an intrinsic condition (Open Problem 6.7). The "more than relabelling" claim for this paper rests on the *direction-pinning and the pre-registered search*, not on an intrinsic residue characterization the paper does not provide.

---

5. The Unrealizable Residue and the Dichotomy

Definition 5.1 (unrealizable residue). The **unrealizable residue** is

$$\mathfrak{U} \coloneqq \{L\} \setminus \mathrm{Im}(\rho) = \{ \hat{G} \in \mathcal{D} : L(\hat{G}) \wedge \boldsymbol{\mathcal{R}}(\hat{G}) = \emptyset \}$$

— the locally admissible germs realized by no admissible transport. Its elements are the **unrealizable admissible germs**.

Theorem 5.2 (Reachability Dichotomy) [Proven]. Exactly one of the following holds.

Case A (Full Reachability). $\mathfrak{U} = \emptyset$. Every locally admissible germ is realizable; Gate 2 passes; ρ surjects onto $\{L\}$; the programme advances to Gate 3 with domain $\{L\}$. **Case B (Reachability Failure).** $\mathfrak{U} \neq \emptyset$. Some $\hat{G}^* \in \{L\}$ has $\boldsymbol{\mathcal{R}}(\hat{G}^*) = \emptyset$; Gate 2 fails; the locality verdict is N_2, witnessed by $\hat{G}^*$.

Proof. $\mathfrak{U} = \emptyset$ or $\mathfrak{U} \neq \emptyset$ are exhaustive and exclusive. By Theorem 4.2, $\mathfrak{U} = \emptyset$ is $\{L\} = \mathrm{Im}(\rho)$, which is Gate 2; by Definition 5.1, $\mathfrak{U} \neq \emptyset$ exhibits an L-satisfying germ with empty realizer set, which is the N_2 cell of the Locality Decision Procedure (*The Locality Decision Theorem* Theorem 3.4). A single such germ suffices, since N_2 is an existential failure.

Remark 5.3 (the dichotomy is trivial; the content is the residue) [Methodological]. Theorem 5.2 is the law of excluded middle applied to $\mathfrak{U}$, and on its own carries no information about which case obtains. We state it only to fix vocabulary and to make explicit that a *single* witness decides Gate 2 negatively. The substantive questions — where $\mathfrak{U}$ could live, what its minimal element looks like, and how to search for one — are Sections 6, 8, and 9. The dichotomy is the frame; it is not the result.

---

6. Where the Residue Can Live: from Localization to Conditional Reduction

The reduction of Sections 4–5 — Gate 2 ⇔ $\mathfrak{U} = \emptyset$ — is exact but offers no grip on *where* a witness could be. This section supplies that grip in two stages. Stages 6.1–6.2 give a **heuristic localization of the search** — a sector-by-sector argument for which conjuncts of L can and cannot populate $\mathfrak{U}$, so the search of Section 9 can be weighted rather than uniform. Stages 6.3–6.4 then harden that localization into a **conditional reduction**: we fix the form of the intrinsic integrability predicate the localization was missing, isolate the structural requirements it must satisfy, and prove (Theorem 6.9) that for any predicate of that form, the residue is the union of per-sector failure loci and Gate 2 reduces to per-sector hub-completion on $\{L\}$. The reduction is conditional on one named, shaped input — the four-sector cell-completion catalogue — which is not supplied here and is stated as Open Problem 6.7. The honest scope is therefore: §6.1–6.2 are heuristic and unconditional; §6.3–6.4 are a proven reduction conditional on the supply of that single object, and on nothing else.

Remark 6.1 (why no intrinsic integrability predicate is offered here, and what is offered instead) [Methodological]. One would like a predicate $I(\hat{G})$, defined without reference to ρ — a Frobenius- or cocycle-type germ condition, or a discrete hub-completion condition — such that realizability factors as "$\hat{G} \in \mathrm{Im}(\rho) \Leftrightarrow L(\hat{G}) \wedge I(\hat{G})$". Such a predicate would turn "does L entail realizability?" into the genuinely interrogable "does L entail I?", conjunct by conjunct, and would upgrade this section from localization to reduction. We do **not** have the predicate itself. Defining integrability as "$\hat{G} \in \mathrm{Im}(\rho)$" is circular — it is realizability under another name — and supplying a non-circular condition requires the explicit $K = 7$ hub adjacency structure, which is transport-programme material not reproduced here and which we do not invent. What we *can* do, and do in §6.3–6.4, is fix the predicate's **form**, isolate the structural requirements it must satisfy, and prove the reduction that follows for any predicate of that form (Theorem 6.9) — leaving only the four cell-completion relations themselves as a named, shaped input (Open Problem 6.7). So the honest scope is: the spine of the paper is the soundness/completeness reframing (§4) and the pre-registered search (§9), neither of which requires an intrinsic I; §6.1–6.2 localize heuristically; and §6.3–6.4 harden that localization into a reduction conditional on the supply of the four-sector catalogue, and on nothing else.

The localization rests on a single structural observation. L is a *germ predicate* — it constrains the degeneracy datum on a neighbourhood — while realizability requires a *transport* over the neighbourhood producing those values. Where the two could diverge is sector-dependent, and the two kinds of conjunct in L behave oppositely.

6.1 The conservation sector is reachability-favourable, in isolation

Proposition 6.2 ($\mathcal{C}$_circ does not by itself obstruct realization) [Conditional]. Taken in isolation, the conservation conjunct $\mathcal{C}$_circ does not obstruct local realization: a germ satisfying the pointwise continuity identity $\partial_\mu C^\mu = s_c$ (substrate), or its continuum analogue $\nabla_\nu C^{\mu\nu} = \hat{C}^\mu$, admits a current realizing the prescribed germ data on a contractible neighbourhood. Hence $\mathcal{C}_{circ}$ alone contributes no element to $\mathfrak{U}$.

Argument. The continuity equation is a first-order underdetermined relation between a current and its source; on a contractible neighbourhood a current with prescribed germ and divergence always exists (ordinary local solvability — the obstruction to such a current is global, not germ-level). The germ amendment keeps the datum read at neighbourhood level, where this solvability is the relevant statement. The label is [Conditional] because the argument assumes the continuum/substrate regularity under which local solvability holds; it is not a bare germ-combinatorial claim.

Remark 6.3 (the in-isolation qualifier is load-bearing). Proposition 6.2 claims only that $\mathcal{C}_{circ}$ acting *alone* does not obstruct. It does **not** claim that the current component contributes nothing to a *joint* obstruction. Section 6's own framing allows a witness to arise from an interaction of conjuncts, and a genuinely joint obstruction — the current germ data interacting with the hub-completion problem, so that no current realizing the germ is *compatible with* an admissible cell configuration — is not excluded here, and indeed could not be cleanly attributed to one sector at all. This is all the search needs: Proposition 6.2 de-weights the conservation axis as an *isolated* source of failure; it does not remove the current component from the joint analysis, and Section 9's search retains it in the germ data.

Remark 6.4 (the gates separate cleanly on $\mathcal{C}_{circ}$). Proposition 6.2 is the Gate-2 face of the gate separation the prior paper insisted on at Gate 3. Conservation is local *in form* (Gate 1), is *locally liftable in isolation* (Gate 2, this proposition), and yet can still obstruct *global assembly* through the transition cocycle on overlaps (Gate 3) — the Maxwell pattern, in which a closed field strength admits a local potential everywhere ($dF = 0 \Longrightarrow F = dA$ locally) yet no global potential need exist on a topologically non-trivial base, the obstruction being a cohomology class of the field strength. The same constraint is favourable at Gates 1 and 2 and is the live candidate only at Gate 3. Conflating these is the error the prior round corrected; Proposition 6.2 keeps them apart on the Gate-2 side.

6.2 The live risk is combinatorial or joint

Conjecture 6.5 (the N_2 risk sits in the combinatorial or joint sectors) [Conjectural]. If $\mathfrak{U} \neq \emptyset$, a witness arises from the combinatorial sectors $\mathcal{C}_{inc}$, $\mathcal{C}_{hub}$, $\mathcal{C}_{comp}$ — or from a joint obstruction coupling those sectors to the current germ data — rather than from $\mathcal{C}_{circ}$ in isolation. The mechanism would be a germ that satisfies the discrete closure/hub/competition predicates germ-wise but whose hub-neighbourhood data cannot be completed to an admissible cell configuration around the $K = 7$ hub — a discrete consistency the germ predicate, evaluated locally, does not enforce.

Basis. Proposition 6.2 removes the conservation sector as an *isolated* source. Among the remaining conjuncts, $\mathcal{C}_{hub}$ references the $K = 7$ hub structure, which imposes adjacency relations between cells meeting at a hub; germ-level satisfaction of a per-cell predicate need not extend to a globally consistent assignment around the hub if the hub combinatorics overconstrain the joint completion. This is the discrete analogue of an integrability obstruction and is where a witness, if one exists, is structurally most natural. The label is [Conjectural]: it is a *prioritization* of the search, grounded in the sector structure, not a demonstration that $\mathfrak{U} \neq \emptyset$, nor an exoneration of the conservation sector against all *joint* completions (Remark 6.3). It directs the search; it does not pre-judge it.

Remark 6.6 (the symmetry with Gate 1). Gate 1 localized its risk to $\mathcal{C}$_circ and then cleared it. Gate 2 localizes its risk *away* from isolated $\mathcal{C}$_circ and onto the combinatorial sectors and their coupling to the current data. The constraint the programme most feared at the level of form is the one most favourable to realization in isolation; the discrete closure sectors that were routine at Gate 1 carry whatever Gate-2 risk exists. The search of Section 9 is weighted accordingly, while retaining the current data so that joint obstructions are not missed.

6.3 The intrinsic integrability condition: a four-sector form

We now do as much of the work of an intrinsic integrability predicate as can be done without the K = 7 hub adjacency structure itself. We cannot exhibit the predicate I(Ĝ) — that requires the adjacency relations, which we do not invent (Remark 6.1) — but we can fix its *form*, isolate the structural requirements any admissible such predicate must carry, and prove the conditional reduction that follows for **any** predicate of that form. The localization of §6.1–6.2 thereby hardens into a reduction conditional on one named, shaped input, rather than remaining a bare open problem. The single missing object is named sharply in Open Problem 6.7 at the end of this subsection; everything above it here is proven modulo that object's supply.

Definition 6.7a (the four-sector cell-completion relations) [Methodological — named input]. For each sector s ∈ {$\mathcal{C}$_inc, $\mathcal{C}$_hub, $\mathcal{C}$_circ, $\mathcal{C}$_comp}, a **cell-completion relation** is a predicate

$I\_s : \mathcal{D} \to \{\top, \bot\}$,

intended to hold at Ĝ when the per-cell s-data that Ĝ assigns to the substrate cells meeting at a K = 7 hub completes to an admissible s-configuration around that hub — i.e. when the *inter-cell* assignment of the sector's closure data is jointly admissible on the closed star of the hub. The four relations {I_inc, I_hub, I_circ, I_comp} are the **named input** of this subsection. They are not constructed here (Remark 6.1); supplying them as explicit conditions is the single prerequisite isolated in Open Problem 6.7.

Remark 6.7b (why the relations are not in hand, and what they are not) [Methodological]. The I_s are conditions on the *inter-cell* assignment of closure data across the substrate lattice. They are categorially distinct from two adjacent objects that share their vocabulary and supply neither: (i) the intra-wheel transition rules governing transitions among the seven internal closure states of a single hub-plus-boundary wheel, where "hub" and "boundary" are *states of one cell*, not adjacent cells; and (ii) the operational ledger-closure inequalities governing the resource cost of fact creation by a channel, where "admissibility" is operational realizability, not inter-cell completion. The cell-completion catalogue is a third object — the inter-cell completion conditions on the substrate lattice — and is the one Open Problem 6.7 requires. The reduction below is structured so that this gap is isolated entirely to the parameter $\{I_s\}$ and propagates nowhere else.

Definition 6.7c (four-sector form). [Methodological] A predicate $I : \mathcal{D} \to \{\top, \bot\}$ is of **four-sector form** if it is the conjunction of one cell-completion relation per sector,

$$I(\hat{G}) \Leftrightarrow I_{inc}(\hat{G}) \wedge I_{hub}(\hat{G}) \wedge I_{circ}(\hat{G}) \wedge I_{comp}(\hat{G}),$$

with $\{I\} \coloneqq \{ \hat{G} : I(\hat{G}) \}$ and $\{I_s\} \coloneqq \{ \hat{G} : I_s(\hat{G}) \}$, so $\{I\} = \bigcap_s \{I_s\}$. The four-sector form mirrors the four-conjunct structure of L itself (Section 2): just as L is the germ-level conjunction of the four sector predicates, I is the *completion-level* conjunction of the four sector relations.

The bare form is too unconstrained to support a reduction. The following three requirements are the minimal structure any *admissible* hub-completion predicate must satisfy — the conditions under which the form behaves as a genuine integrability predicate rather than an arbitrary set membership.

Requirement R1 (locality on the hub neighbourhood). [Methodological] Each I_s is determined by the restriction of $\hat{G}$ to a bounded hub neighbourhood N_h: if $\hat{G}$, $\hat{G}'$ agree on N_h then $I_s(\hat{G}) = I_s(\hat{G}')$ for every s. Hub-completion is decidable from local data within the closure scale of the hub, with no dependence on configuration data outside N_h. R1 is *necessary* for any predicate intrinsically characterizing $\mathrm{Im}(\rho)$, because ρ is assembled from local transport data within the closure scale (Definition 3.1), so its image cannot depend on data outside N_h.

Requirement R2 (sector-additivity). [Methodological] The sectors are the carriers of the failure: the failure locus of I inside $\{L\}$ is the union of per-sector failure loci,

$$\{L\} \setminus \{I\} = \bigcup_{s} (\{L\} \setminus \{I_s\}).$$

This is pure De Morgan — $\neg\bigwedge_s I_s = \bigvee_s \neg I_s$ — *and holds for any four-sector-form predicate with no further assumption; it is what licenses attributing each obstruction to the sector(s) whose relation it fails. We deliberately do* **not** *require the four loci to be disjoint. A germ may fail two or more sectors at once — admissible in each singly but unrealizable only in combination — and such a germ lies in the intersection* $\mathcal{U}_s \cap \mathcal{U}_{s'}$ of two failure loci. This is exactly the *joint obstruction* preserved by Remark 6.3 and Conjecture 6.5: dropping disjointness is what keeps a genuinely cross-sector witness on the table rather than legislating it away. The cost of allowing overlap is only that the residue is a union, not a partition — which affects the *counting* (Proposition 6.12 counts which sectors fail, not a sum of disjoint cardinalities), not the set identity of Theorem 6.9(b).

Requirement R3 (L-compatibility). [Methodological] I refines L without contradicting it: $\{I\} \cap \{L\} \neq \emptyset$, and I is not identically $\bot$ on $\{L\}$. R3 excludes the degenerate predicate that vacuously empties $\{L\}$; it asserts hub-completion is a genuine additional condition, not a covert restatement of $\neg L$. R3 is consistent with Gate 1, which returned $\mathrm{Im}(\rho) \subseteq \{L\}$ non-vacuously (Theorem 4.1).

Remark 6.7d (status of R1–R3) [Methodological]. The three requirements differ in kind, and we are explicit about it. R2 (the union) is free: it is De Morgan and holds for any four-sector-form predicate, so it costs no generality — it is a notational decomposition, not a constraint on I. R1 and R3 are genuine necessary conditions on any predicate that intrinsically characterizes Im(ρ) inside {L}: R1 because ρ is assembled from local transport data within the closure scale (Definition 3.1), so its image cannot depend on data outside N_h; R3 because Gate 1 returned Im(ρ) ⊆ {L} non-vacuously (Theorem 4.1), so the characterizing predicate cannot empty {L}. Both R1 and R3 are therefore satisfied by any genuine characterizing I, costing no generality. What we do **not** assume is disjointness of the per-sector failure loci: that would be a substantive individuation choice on the {I_s} (forcing the sectors to partition the failure), and it is precisely the assumption that would exclude joint obstructions, so we omit it. The sectors are carriers, not a partition.

6.4 The conditional reduction theorem

We can now state the result that hardens §6.1–6.2 from localization to reduction. The matching property is the precise sense in which a four-sector predicate "characterizes" the realizer image.

Definition 6.8 (matching property). [Methodological] A four-sector-form predicate I has the **matching property** for ρ if Im(ρ) = {L} ∩ {I} — among germs passing L, those reachable by ρ are exactly those passing I. This is the non-circular content the integrability predicate of Remark 6.1 was asked to supply: it links I to ρ's image without defining I as membership in that image.

Theorem 6.9 (conditional reduction) [Conditional-on-supply]. Let I be a predicate of four-sector form (Definition 6.7c) satisfying R1–R3 and the matching property (Definition 6.8) for ρ. Then:

(a) The unrealizable residue is exactly the I-failing locus inside {L}: 𝔘 = { Ĝ ∈ {L} : ¬I(Ĝ) } = {L} \ {I}. **(b)** The residue is the union of the per-sector failure loci: 𝔘 = 𝔘_inc ∪ 𝔘_hub ∪ 𝔘_circ ∪ 𝔘_comp, where 𝔘_s ≔ {L} \ {I_s}. The union need not be disjoint: a germ failing two sectors at once lies in 𝔘_s ∩ 𝔘_{s′}, and such joint obstructions are not excluded (Remark 6.3). **(c)** Gate 2 holds iff every sector is hub-complete on {L}: {L} ⊆ Im(ρ) ⇔ ⋀_s ({L} ⊆ {I_s}). **(d)** The §6.1–6.2 localization is a reduction: the reachability obstruction, if any, is carried entirely by the sectors s with {L} ⊄ {I_s}, localized to hub-completion failures within N_h (R1), and by no other structure.

Proof. **(a)** By the matching property Im(ρ) = {L} ∩ {I}; by Theorem 4.1 (Gate-1 soundness) Im(ρ) ⊆ {L}, so Im(ρ) = Im(ρ) ∩ {L} = {L} ∩ {I}. Then by Definition 5.1, 𝔘 = {L} \ Im(ρ) = {L} \ ({L} ∩ {I}) = {L} \ {I} = { Ĝ ∈ {L} : ¬I(Ĝ) }. **(b)** ¬I(Ĝ) = ¬⋀_s I_s(Ĝ) holds iff some sector fails, so {L} \ {I} = ⋃_s ({L} \ {I_s}) = ⋃_s 𝔘_s by De Morgan (R2); no disjointness is claimed or needed. **(c)** By Theorem 5.2 and (a), Gate 2 ⇔ 𝔘 = ∅ ⇔ {L} \ {I} = ∅ ⇔ {L} ⊆ {I} = ⋂_s {I_s} ⇔ ⋀_s ({L} ⊆ {I_s}). **(d)** By (a)–(c) the obstruction set is supported exactly on ⋃_s 𝔘_s; each 𝔘_s is empty unless {L} ⊄ {I_s}; and by R1 each I_s is a local hub-neighbourhood condition, so the obstruction is localized to hub-completion failures within N_h and depends on no configuration data outside it.

Remark 6.10 (what Theorem 6.9 does and does not establish) [Methodological]. Theorem 6.9 is fully proven *modulo the supply of I* — every step is unconditional once a matching, R1–R3-satisfying four-sector predicate exists. It does not prove such an I exists; exhibiting I is Open Problem 6.7. What it establishes is that the *only* missing ingredient between the present localization and a Gate-2 reduction is the supply of the four relations $\{I_s\}$: no further realizer analysis, no additional locality lemma, and no appeal to the wheel or operational-admissibility lineages is required. It converts Conjecture 6.5 — that the live N_2 risk is the combinatorial hub-completion question — from a conjecture about *where* the obstruction lives into a theorem about where it *can* live, conditional only on the supply of I. The label [Conditional-on-supply] marks exactly this: proven but for one named, shaped parameter.

Corollary 6.11 (the conservation sector, reconciled) [Conditional-on-supply]. Under the hypotheses of Theorem 6.9, Proposition 6.2 is the statement $\mathfrak{U}_{circ} = \emptyset$ in isolation: the conservation sector's per-cell data completes locally (continuity identity $\partial_\mu C^\mu = s_c$, local solvability), so $\{L\} \subseteq \{I_{circ}\}$ and $\mathfrak{U}_{circ} = \{L\} \setminus \{I_{circ}\} = \emptyset$. By the union of Theorem 6.9(b) this leaves $\mathfrak{U} = \mathfrak{U}_{inc} \cup \mathfrak{U}_{hub} \cup \mathfrak{U}_{comp}$ — the combinatorial-sector residue of Conjecture 6.5. Crucially, because the union is not required disjoint, a joint obstruction coupling the current data to hub completion is still representable: such a witness fails $\mathcal{C}_{circ}$ *only in combination* with another sector, i.e. it lies in a pairwise intersection $\mathfrak{U}_{circ} \cap \mathfrak{U}_s$. Proposition 6.2 says only that $\mathcal{C}_{circ}$ contributes no obstruction *in isolation* — $\mathfrak{U}_{circ} = \emptyset$ as a standalone locus — which is consistent with $\mathcal{C}_{circ}$ data participating in a joint failure that is attributed, in the union, to the partner sector(s) it co-fails with. Isolation-favourability of $\mathcal{C}_{circ}$ controls the standalone locus $\mathfrak{U}_{circ}$; it does not legislate the intersections, and so does not exclude the joint obstruction Remark 6.3 preserves.

Proof. By Theorem 6.9(b), $\mathfrak{U}_{circ} = \{L\} \setminus \{I_{circ}\}$. Proposition 6.2's in-isolation continuity argument asserts the conservation per-cell data completes locally, i.e. $\{L\} \subseteq \{I_{circ}\}$, whence the standalone locus $\mathfrak{U}_{circ} = \emptyset$. The union then reads $\mathfrak{U} = \mathfrak{U}_{inc} \cup \mathfrak{U}_{hub} \cup \mathfrak{U}_{comp}$. No disjointness is invoked, so any germ whose obstruction genuinely couples conservation to a combinatorial sector remains in $\mathfrak{U}$ via the partner sector's locus.

Proposition 6.12 (the failing-sector set determines the verdict) [Conditional-on-supply]. Under the hypotheses of Theorem 6.9, define the **failing-sector set** and its count

$$F_{\mathfrak{U}} := \{ s : \{L\} \not\subset \{I_s\} \}, \quad n_{\mathfrak{U}} := \#F_{\mathfrak{U}} \in \{0, 1, 2, 3, 4\}.$$

Then $n_{\mathfrak{U}}$ counts *which sectors fail*, not a sum of disjoint residue cardinalities (the loci $\mathfrak{U}_s$ may overlap), and it determines the Gate-2 verdict: $n_{\mathfrak{U}} = 0 \Longrightarrow$ Gate 2 PASS (verdict tends to L); $n_{\mathfrak{U}} \geq 1 \Longrightarrow$ Gate 2 FAIL (verdict N_2), with the failing sectors named explicitly by $F_{\mathfrak{U}}$. A joint obstruction registers as $|F_{\mathfrak{U}}| \geq 2$ with the corresponding $\mathfrak{U}_s$ overlapping; the verdict depends only on whether $F_{\mathfrak{U}}$ is empty, which is insensitive to overlap. The "one residue or two" branch of the downstream chain (Section 10) is the small-$|F_{\mathfrak{U}}|$ reading: how many sectors carry a surviving obstruction.

Proof. By Theorem 6.9(b)–(c), $\mathfrak{U} = \emptyset$ iff every $\mathfrak{U}_s = \emptyset$ iff $F_{\mathfrak{U}} = \emptyset$ iff $n_{\mathfrak{U}} = 0$. When $n_{\mathfrak{U}} \geq 1$, some $\mathfrak{U}_s \neq \emptyset$ so $\mathfrak{U} \neq \emptyset$ and Theorem 5.2 gives Gate 2 FAIL, verdict N_2; the failing sectors are

F_𝔘 by definition. Overlap among the 𝔘_s changes none of these equivalences, since each turns on emptiness of individual loci, not on their cardinalities summing.

Open Problem 6.7 [Open — requires the four-sector cell-completion admissibility catalogue as an explicit stated object]. Supply the four cell-completion relations {I_inc, I_hub, I_circ, I_comp} (Definition 6.7a) as explicit conditions on the per-cell closure data assigned to the substrate cells meeting at a K = 7 hub — for each sector, an explicit predicate deciding when the per-cell s-data completes to an admissible s-configuration on the closed star of the hub — such that the resulting four-sector-form I (Definition 6.7c) satisfies R1–R3 and the matching property (Definition 6.8). This is the single prerequisite for Gate-2 closure. Once supplied, Theorem 6.9 discharges immediately: the residue is the union of per-sector failure loci (6.9b), the failing-sector set F_𝔘 is computed and named (Proposition 6.12), and the verdict routes through the hand-forward of Section 10. The verification of the matching property for the supplied I is itself a finite check at the closure scale of the hub, by R1. The relations have not been located, as an explicit stated object, in the spectral-geometry branch (where "K = 7" denotes the seven *intra-cell* closure states of the canonical wheel) nor in the operational-admissibility line (where "admissibility" denotes operational realizability under finite-distinguishability and irreversible-commitment constraints); the catalogue is a third object, governing *inter-cell* completion across the substrate lattice, and its explicit statement is what would convert Theorem 6.9 from [Conditional-on-supply] into a Gate-2 verdict.

---

7. The Non-Unique Regime and the Falsification Burden

The Locality Inspection Appendix A.6 found the realizer regime **non-unique on L**: a single germ admits at least three ~_G-distinct constitutive realizers — linear (M1), bilinear/energy-current (M2), entropy-current (M3) — each a bona fide Noether current of an exact symmetry, all admissible under the local axioms, with effective uniqueness restored only by the external M2 selection (BCB §9.5.4). We flag at the outset a status question A.6 leaves open: the trichotomy {M1, M2, M3} is *found*, not *proven exhaustive* — no result rules out a further admissible family M4. This matters for the falsification condition below, and we label accordingly. We record the consequence for Gate 2, which is the opposite of the consequence for Gate 3 and must not be confused with it.

Proposition 7.1 (non-uniqueness favours reachability) [Proven]. For a fixed germ Ĝ, the realizer set decomposes over the admissible constitutive families {M_k},

$\mathcal{R}(\hat{G}) = \bigcup k\ \mathcal{R}\{M\_k\}(\hat{G})$,

so $\mathcal{R}(\hat{G}) \neq \emptyset$ iff *at least one* family realizes Ĝ. A multiplicity of admissible families can only enlarge Im(ρ); reachability is therefore weakly easier in the non-unique regime than under any single selected family. This holds for any number of families, enumerated or not, since it is a union argument.

Proof. Im(ρ) is the union of the family images $\mathrm{Im}(\rho_{M_k})$; a union of images contains each. So passing from the M2-selected sub-map to the full non-unique map can only add germs to the image, hence can only shrink $\mathfrak{U}$. The argument does not depend on the family set being finite or exhaustively enumerated.

Proposition 7.2 (non-uniqueness raises the N_2 falsification burden) [Conditional on the M-family classification being complete]. Dually, an N_2 witness $\hat{G}^*$ must have $\mathcal{R}_{M_k}(\hat{G}^*) = \emptyset$ for *every* admissible family M_k — every family must fail to realize it. A germ realized by any single family, even an unselected one, is **not** an N_2 witness. The falsification condition of the Gate-2 search is therefore emptiness across all admissible families, not emptiness of the selected family alone. If {M1, M2, M3} is the complete classification, this is the three-way conjunction below; if a further family M4 is admissible, the condition must extend to it, and a search that checked only M1–M3 would be *incomplete* — a germ defeated on M1–M3 but realized by M4 is not a witness.

Proof. By Proposition 7.1, $\mathcal{R}(\hat{G}^*) = \emptyset$ iff *all* family fibers are empty. The Gate-2 falsification condition is $\mathcal{R}(\hat{G}^*) = \emptyset$ (Theorem 5.2, Case B), which unpacks to emptiness over the full family set. The reduction to a finite three-way conjunction is exactly the content conditional on {M1, M2, M3} being exhaustive; absent that, the conjunction ranges over whatever the complete family set is. The label is therefore [Conditional] on the M-family classification, and the safe alternative — importing the M2 selection and recording it — collapses the burden to a single family and sidesteps the exhaustiveness question entirely (Remark 7.3).

Remark 7.3 (contrast with Gate 3, to forestall the conflation). The non-unique regime makes *both* N_2 and N_3 harder to witness — N_2 because all families must fail to *realize* (Proposition 7.2), N_3 because all families must fail to *extend* (the heavier non-unique burden of *The Locality Inspection* Procedure 5.3). But the two are distinct conditions on distinct objects: realization is non-emptiness of $\mathcal{R}(\hat{G})$ (Gate 2); extension is vanishing of an obstruction class for some member of $\mathcal{R}(\hat{G})$, i.e. $0 \in \mathfrak{o}(\hat{G})$ in the cohomological sense, where $\mathfrak{o}(\hat{G})$ is the prior paper's realizable obstruction image — renamed here from its Im(G) to avoid collision with Im(ρ), the central object of this paper (Gate 3). The shared phrase "all families must fail" refers to two different failures. The Gate-2 search tests the first; it does not touch the second. Either the heavier all-families falsification is run directly — at which point the M-family exhaustiveness question of Proposition 7.2 is live — or the M2 selection is imported and recorded as a precondition, collapsing the conjunction to the single selected family and sidestepping exhaustiveness entirely. The inspector must do one explicitly, not assume uniqueness silently.

---

8. Minimal Unrealizable Germs

If $\mathfrak{U} \neq \emptyset$, the search benefits from being organized around the simplest witness. We make "simplest" precise and prove a minimum exists, so that the search of Section 9 can be bounded by complexity rather than ranging over arbitrary geometries.

Definition 8.1 (complexity grading) [Proven $\mathbb{N}$-valued]. Fix a complexity measure $c : \mathcal{D} \to \mathbb{N}$ on germs built **only** from germ-local, Gate-1-bounded data: the radius (in refinement steps of $\mathscr{R}$) of the smallest neighbourhood on which the germ's data must be specified to **evaluate L**, plus the cell-count of that neighbourhood under the K = 7 architecture. Both quantities are bounded by the germ-locality established at Gate 1 — L is a germ predicate, so its evaluation radius is finite and the enclosed cell-count is finite under the K = 7 finite-presentation character — so c is provably $\mathbb{N}$-valued. Crucially, c does **not** include the radius required to *test realizability* (i.e. to decide $\hat{G} \in \mathrm{Im}(\rho)$), which can in principle be unbounded; that radius is handled separately under Conjecture 8.3, not folded into the grading.

Theorem 8.2 (existence of a minimal unrealizable germ) [Proven]. If $\mathfrak{U} \neq \emptyset$, there exists $\hat{G}\_min \in \mathfrak{U}$ with $c(\hat{G}\_min) \leq c(\hat{G})$ for all $\hat{G} \in \mathfrak{U}$.

Proof. By Definition 8.1, c is $\mathbb{N}$-valued on all of $\mathcal{D}$ — in particular on $\mathfrak{U}$ — using only the L-evaluation radius and cell-count, which are finite for every germ by Gate-1 germ-locality and the finite-presentation character, independently of any realizability-test radius. Hence $c(\mathfrak{U}) \subseteq \mathbb{N}$ is non-empty when $\mathfrak{U} \neq \emptyset$, and by well-ordering of $\mathbb{N}$ it has a least element c_0. Any germ in the (non-empty) preimage $\mathfrak{U} \cap c^{-1}(c_0)$ is a minimal unrealizable germ. The well-ordering does not appeal to Conjecture 8.3: the grading is total and finite-valued by construction, so the [Proven] label is earned.

Conjecture 8.3 (minimality suffices for the search) [Conjectural]. The realizability obstruction is **local and monotone under germ extension**: every unrealizable admissible germ contains, or reproduces the obstruction of, some minimal unrealizable germ, so the obstruction in any witness is already present in a witness of complexity c_0. This is precisely the assumption that the realizability-test radius excluded from Definition 8.1 is in fact bounded by the obstruction's own scale — that a germ unrealizable at large radius is unrealizable because of a sub-germ unrealizable at radius $\leq c_0$. Were it established, the Gate-2 search could be **bounded**: it would suffice to examine germs up to complexity c_0, and exhausting all germs to a complexity bound with no witness found would license a [Proven] (not merely [Conditional]) pass below that bound. We mark it [Conjectural]; it is plausible for the combinatorial sectors (a hub incompatibility persists under enlargement) but not here proven, and it is the *only* place the boundedness assumption enters — Theorem 8.2 does not depend on it. Its status is the difference between a bounded search and an open-ended one, and establishing it is the natural companion result to this paper, alongside Open Problem 6.7.

Remark 8.4 (why minimality is worth the trouble). Theorem 8.2 alone gives a target; Conjecture 8.3, if proven, gives a *stopping criterion*. Without the latter, a failed search over germs up to complexity c remains [Conditional] (a witness might lurk at higher complexity); with it, a failed search to c_0 becomes [Proven] below the bound. The pre-registration of Section 9 is written to be valid either way, but its representativeness argument is materially strengthened by Conjecture 8.3, and we flag the dependency explicitly.

9. The Pre-Registered Gate-2 Search

The Locality Inspection Principle 7.1 makes a Gate-2 pass admissible as evidence only if its search space, representativeness argument, and falsification condition are committed **before** the search. We deposit that record now, so that whatever the execution returns, its standing is auditable rather than asserted. The record is fixed here and is not to be narrowed after the search has seen data.

Pre-Registration 9.1 (Gate-2 search record) [Methodological].

(a) Search space $\mathcal{P}_2$. The L-satisfying germs, parametrized and ordered by the complexity grading c (Definition 8.1), enumerated as:

— single-hub germs (minimal c): all L-satisfying degeneracy data on one K = 7 hub neighbourhood; — hub-adjacent germs: L-satisfying data on two hubs sharing a closure boundary, ranging over admissible transition data; — competition germs: L-satisfying data on the smallest neighbourhood on which $\mathcal{C}$_comp's competing-sector balancing is non-trivial; — and the closure of these under germ extension up to a stated complexity bound c_max, with c_max recorded before searching and raised only by re-registration, never silently.

The conservation sector enters $\mathcal{P}_2$ only through the germ data it constrains, not as an independent search axis, on the strength of Proposition 6.2 (conservation does not obstruct realization in isolation); the search weight is placed on $\mathcal{C}$_hub and $\mathcal{C}$_comp completions per Conjecture 6.5, while the current germ data is retained so that joint obstructions (Remark 6.3) are not missed.

(b) Representativeness argument. $\mathcal{P}_2$ faithfully samples {L} because: (i) by Proposition 6.2 the conservation sector is no *isolated* integrability obstruction, so an isolated-conservation witness is excluded and $\mathfrak{U} \subseteq$ (germs whose obstruction is combinatorial or joint), which $\mathcal{P}_2$ foregrounds while retaining the current data; (ii) by Theorem 8.2 a witness of minimal complexity exists, and $\mathcal{P}_2$ is ordered to reach low-c germs first. We are explicit about what this ordering does and does not buy. The grading c measures L-evaluation cost only (Definition 8.1); membership in $\mathfrak{U}$ is decided by the realizability-test radius, which c does not see. So absent Conjecture 8.3 the c-ordering is a pure *reachability heuristic* — it guarantees that low-c germs are examined early, nothing more — and it is **orthogonal to where the obstruction lives**: a germ can be c-small yet unrealizable only by virtue of structure at large test-radius, so low-c-first does not track the obstruction and confers no coverage of it. (iii) The ordering acquires exhaustiveness relevance *only* through Conjecture 8.3: under 8.3 the obstruction is bounded by its own scale, the test-radius collapses to $\leq c_0$, the two quantities cease to be orthogonal, and exhausting $\mathcal{P}_2$ to c_0 exhausts the obstruction structure — making the sample exhaustive below the bound. Clause (iii), and hence Conjecture 8.3, is the *only* route to a [Proven] pass; absent it, a clean search yields a [Conditional] pass over $\mathcal{P}_2$ up to c_max, and the c-ordering is a convenience of enumeration, not evidence of coverage.

(c) Falsification condition. A germ $\hat{G} \in \mathcal{P}_2$ is an N_2 witness iff $L(\hat{G})$ holds and *no* admissible constitutive family realizes $\hat{G}$. Under the trichotomy {M1, M2, M3} — *conditional on that classification being complete (Proposition 7.2)* — this is

$$L(\hat{G}) \wedge \mathcal{R}_{M1}(\hat{G}) = \emptyset \wedge \mathcal{R}_{M2}(\hat{G}) = \emptyset \wedge \mathcal{R}_{M3}(\hat{G}) = \emptyset.$$

If a further admissible family is identified, the conjunction must extend to it. If instead the M2 selection is imported, the condition collapses to $L(\hat{G}) \wedge \mathcal{R}_{M2}(\hat{G}) = \emptyset$, the import is recorded as a Gate-2 precondition in the same register, and the exhaustiveness question is sidestepped. The search runs the full (non-unique) condition by default and records which option was taken.

(d) Standing claimed on a clean pass. [Proven] iff the search is an exhaustive surjectivity proof of ρ onto {L} (the completeness of Theorem 4.2), available only under Conjecture 8.3 to bound the exhaustion; otherwise [Conditional], a structured failed search over $\mathcal{P}_2$ to c_max. A pass recorded without this register is, per Principle 3.7, not evidence.

Remark 9.2 (what the register buys). The register makes the asymmetry of *The Locality Inspection* Remark 4.3 structural: a failed search over the *characterized* $\mathcal{P}_2$ is admissible evidence; a failed search over an uncharacterized space is not. Because (a)–(d) are fixed before execution, the search cannot be shrunk around a clean result after the fact, and the [Proven]/[Conditional] tag the execution earns is determined by the register, not by how clean the outcome looks.

Part II — The Execution of the Registered K = 7 Search

Part I reduced Gate 2 to the emptiness of the residue $\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$ (Theorem 5.2), fixed the four-sector form of the intrinsic integrability predicate and proved the conditional reduction (Theorem 6.9, [Conditional-on-supply]), and deposited the pre-registered search $\mathcal{P}_2$ in three strata (Pre-Registration 9.1). The single object left unsupplied was the four cell-completion relations (Open Problem 6.7).

The K = 7 constraint forms are now concrete enough to run the registered search. Part II does so, in the sense of the pre-registration: the search space, representativeness, and falsification condition were fixed in Part I before the execution, and what follows is the result of running the registered procedure, with the standing each stratum earns recorded explicitly. The single-hub stratum (§11) closes on the concrete forms; the hub-adjacent stratum (§12) reduces to a $\mathbb{Z}_2$

balance criterion; §13 proves the gauge part of the matching sign is always harmless, so the verdict depends only on a residue class; §14–15 reduce the hub-adjacent and competition strata to named hypotheses about the shared-boundary rule and assemble the conditional verdict; §16 records what the execution narrows in Open Problem 6.7. We do not claim an unconditional or global Gate-2 pass.

10. The Concrete K = 7 Constraint Forms

The K = 7 hub is the wheel W_6: one hub vertex, six outer vertices arranged cyclically, six spokes (hub-to-outer), six outer edges (outer cyclic ring), and six triangular 2-cells (hub, outer i, outer i+1). The degeneracy datum on a single hub is carried by the six spoke amplitudes

$\lambda = (\lambda_0, \lambda_1, \ldots, \lambda_5) \in \mathbb{R}^6$, indices mod 6.

The local admissibility predicate L is the conjunction of four sector constraints (Part I, Section 2). On the single-hub spoke sector the two determining functionals are concrete:

Conservation ($\mathcal{C}$_circ): $A\_circ(\lambda) = (\sum_i \lambda_i)^2$. **Competition ($\mathcal{C}$_comp):** $A\_comp(\lambda) = \sum_i (\lambda_i + \lambda_{i+1})^2$, indices mod 6.

The incidence and hub functionals A_inc, A_hub vanish on the admissible spoke sector (constitutive paper); they restrict the sector to the spoke amplitudes but contribute no further on-sector constraint beyond what A_circ and A_comp impose. We carry them through the execution only to confirm they do not vacuously collapse the admissible set (condition C3, §11.5).

Remark II-2.1 (reading of A_comp) [Methodological]. A_comp is a sum of real squares on real amplitudes $\lambda_i \in \mathbb{R}$ — a positive-semidefinite (PSD) cost functional, not a holomorphic complex quadric. This is the reading confirmed by the constitutive paper, and it is load-bearing: A_comp = 0 forces termwise vanishing only for a PSD functional. Were A_comp the holomorphic $\sum(\lambda_i + \lambda_{i+1})^2$ on complex amplitudes, its zero locus in $\mathbb{C}^6$ would be a single quadric containing far more than the alternating ray, and the single-hub execution below would not close. We record the PSD reading as confirmed and flag it as execution condition C1.

---

11. Execution at the Single-Hub Stratum

3.1 The competition functional and its kernel

Lemma II-3.1 (alternating kernel) [Proven]. As a quadratic form on $\mathbb{R}^6$, $A\_comp(\lambda) = \sum_i(\lambda_i + \lambda_{i+1})^2$ is the nearest-neighbour circulant Q with $Q_{ii} = 2$ and $Q_{\{i,i\pm1\}} = 1$ (indices mod 6). Its eigenvalues are

$\mu_k = 2(1 + \cos(2\pi k/6)), k = 0, 1, \ldots, 5,$

and A_comp(λ) = 0 holds iff λ lies in the kernel of Q.

Proof. Expanding, $\sum_i(\lambda_i + \lambda_{i+1})^2 = \sum_i(\lambda_i^2 + 2\lambda_i\lambda_{i+1} + \lambda_{i+1}^2) = 2\sum_i\lambda_i^2 + 2\sum_i\lambda_i\lambda_{i+1} = \lambda^T Q\lambda$ with Q the stated circulant. A circulant with first row (2, 1, 0, 0, 0, 1) is diagonalized by the Fourier modes of $\mathbb{Z}/6$ with eigenvalues $2 + e^{2\pi ik/6} + e^{-2\pi ik/6} = 2 + 2\cos(2\pi k/6) = \mu_k$. Since A_comp is a sum of real squares it is PSD (Remark II-2.1), so its zero locus is exactly ker Q.

Lemma II-3.2 (the kernel is the alternating ray) [Proven]. The eigenvalues μ_k vanish only at k = 3, where $\mu_3 = 2(1 + \cos \pi) = 0$. Hence ker Q is one-dimensional, spanned by the k = 3 Fourier mode, which is the alternating vector

$\lambda_i = (-1)^i a, a \in \mathbb{R}$.

Proof. $\mu_k = 0 \Leftrightarrow \cos(2\pi k/6) = -1 \Leftrightarrow 2\pi k/6 = \pi \Leftrightarrow k = 3$. The k = 3 mode of $\mathbb{Z}/6$ is $e^{2\pi i\cdot 3\cdot i/6} = e^{\pi i i} = (-1)^i$ (the real Fourier mode at the Nyquist frequency of the even cycle). So ker Q = span$\{ ((-1)^0, \dots, (-1)^5) \} = \{ (-1)^i a : a \in \mathbb{R} \}$.

Remark II-3.3 (even N is what makes the kernel exist) [Methodological]. The Nyquist mode $(-1)^i$ is a genuine eigenvector of the cyclic nearest-neighbour operator only when N is even; for odd N there is no k with $\cos(2\pi k/N) = -1$, so Q is strictly positive-definite and A_comp = 0 forces λ = 0. For the K = 7 hub, N = 6 is even and the alternating ray exists. This even-N fact is **intra-hub** — it concerns the outer cycle C_6 of a single hub — and is logically independent of the inter-hub cycle parity of Section 4. We isolate it here to forestall the conflation: even N licenses the single-hub mode; it says nothing about the hub-adjacent cocycle.

3.2 Conservation is satisfied automatically on the ray

Lemma II-3.4 (conservation on the alternating ray) [Proven]. On the alternating ray $\lambda_i = (-1)^i$ a, the conservation functional vanishes:

$$A\_circ(\lambda) = \left(\sum_i (-1)^i a \right)^2 = \left(a \cdot \sum_{i=0}^{5} (-1)^i \right)^2 = (a \cdot 0)^2 = 0,$$

since $\sum_{i=0}^{5} (-1)^i = (1 - 1 + 1 - 1 + 1 - 1) = 0$ — the alternating signs cancel in pairs because N = 6 is even.

Proof. Direct computation as displayed; the cancellation is exactly the statement that the alternating (k = 3) mode is orthogonal to the constant (k = 0) mode, which is what A_circ = (constant-mode projection)² measures.

Thus on the single-hub spoke sector, A_comp = 0 ⟹ λ ∈ alternating ray ⟹ A_circ = 0. The competition constraint is the binding one; conservation is a consequence of it on this sector, not an independent restriction. This is the Gate-2 face of the Part I, Proposition 6.2 ($\mathcal{C}$_circ is reachability-favourable in isolation): here it is favourable because it is *implied* by A_comp on the ray.

3.3 The realizer on the ray

Lemma II-3.5 (σ-family realizer) [Conditional]. The σ-family identifies the alternating spoke pattern $\lambda_i = (-1)^i a$ as the canonical admissibility-restoring transport response. Hence for each $a \in \mathbb{R}$ there is an admissible local transport realizing the germ on the alternating ray: the ray is contained in Im(ρ) at the single-hub stratum.

Status. [Conditional] on the σ-family alternating response being an admissible member of $\mathcal{T}$_L. By Proposition 7.1 (the union $\mathcal{R}(\hat{G}) = \bigcup k\ \mathcal{R}\{M\_k\}(\hat{G})$) a *single* realizing family suffices for ray ⊆ Im(ρ); the σ-family supplies one, which closes the pass. The all-families condition of Proposition 7.2 is the *falsification* burden — what an N_2 witness must defeat (Case B) — and does **not** bind a pass (Case A); we do not import it here. The single remaining condition is therefore admissibility of the σ-response in $\mathcal{T}$_L (execution condition C2, §11.5), not realization across all families.

3.4 The single-hub residue is empty

Theorem II-3.6 (single-hub Gate-2 pass) [Conditional]. At the single-hub stratum, every germ satisfying L is realizable; the single-hub unrealizable residue is empty:

$\mathfrak{U} \cap \mathcal{P}_2$^{single-hub} = ∅.

Proof. A single-hub germ satisfies L iff it lies on the admissible spoke sector (A_inc = A_hub = 0) and satisfies A_circ = 0 and A_comp = 0. By Lemmas 3.1–3.2, A_comp = 0 restricts to the alternating ray; by Lemma II-3.4, A_circ = 0 holds automatically there; so {L} ∩ {single-hub} = alternating ray (modulo C3, that A_inc ∧ A_hub do not further collapse it — §11.5). By Lemma II-3.5 the ray ⊆ Im(ρ). Hence {L} ∩ {single-hub} ⊆ Im(ρ), i.e. $\mathfrak{U} \cap \mathcal{P}_2$^{single-hub} = ∅. In the language of the Part I, Theorem 6.9, this is the single-hub instantiation $\mathfrak{U}$_comp = $\mathfrak{U}$_circ = ∅ on the single-hub neighbourhood, with I_comp and I_circ read concretely as A_comp = 0 and A_circ = 0.

3.5 Execution conditions

The single-hub pass is [Conditional] on three conditions, recorded so the standing is auditable:

- **C1 (PSD reading of A_comp).** A_comp is a sum of real squares, not a holomorphic complex quadric (Remark II-2.1). **Confirmed** by the constitutive paper.
- **C2 (σ-family admissibility).** The σ-family alternating response is an admissible transport in $\mathcal{T}$_L. By Proposition 7.1 a single realizing family suffices for ray ⊆ Im(ρ), so this one condition discharges the pass; the all-families burden of Proposition 7.2 is the falsification condition for a *witness* (Case B), which does not bind a pass (Case A) and is not required here. Pending only the confirmation that the σ-response lies in $\mathcal{T}$_L.
- **C3 (non-vacuity under A_inc ∧ A_hub).** The incidence and hub functionals restrict the sector to the spoke amplitudes without collapsing the alternating ray to {0}; the six spoke

amplitudes are the complete single-hub germ coordinates. Pending explicit confirmation that A_inc ∧ A_hub leave the ray intact rather than empty (a vacuous pass).

Under C1–C3, Theorem II-3.6 holds. C1 is confirmed; C2 and C3 are finite checks deposited as the single-hub stratum's remaining discharge.

---

12. The Hub-Adjacent Stratum as a $\mathbb{Z}_2$ Balance Problem

The single-hub stratum closes (modulo C1–C3). The pre-registered search prescribes stratum 2 next — hub-adjacent germs, two hubs sharing a closure boundary, ranging over admissible transition data (Pre-Registration 9.1(a)). This stratum is Gate 2, not Gate 3: it tests *realization* across a shared boundary, not *global assembly* over the whole complex.

4.1 The per-hub gauge freedom

Each hub H carries the single-hub alternating mode

λ_i^{(H)} = $(-1)^i a^H$, $a^H \in \mathbb{R}$.

Remark II-4.1 (the amplitude is a gauge-dependent $\mathbb{Z}_2$ quantity) [Methodological]. The label $(-1)^i$ depends on which outer vertex of H is called i = 0. Relabelling the hexagon origin by one step sends $(-1)^i \to (-1)^{i+1} = -(-1)^i$, i.e. $a^H \to -a^H$, with no physical change to the germ. So the *sign* of a^H is defined only up to a per-hub gauge flip; only gauge-invariant combinations carry physical content. This is exactly the structure of a $\mathbb{Z}_2$ gauge field on the hub network, and it is what makes the criterion below a *cohomological* triviality rather than a naive sign condition.

4.2 The shared-boundary matching sign

For two hubs H, H′ sharing a boundary, the alternating modes must agree on the shared spokes/vertices. That agreement fixes a relative sign:

Definition II-4.2 (matching cochain). [Methodological] For each adjacent pair (H, H′) in the hub-adjacency graph Γ_hub, the shared-boundary matching condition imposes

aH′ = s_{HH′} a^H, s_{HH′} ∈ {+1, −1},

where s_{HH′} is read from the cross-boundary completion term — the inter-hub instance of A_comp evaluated on the shared edge. The collection s = (s_{HH′}) is a $\mathbb{Z}_2$-valued cochain on the edges of Γ_hub. Under a per-hub gauge flip $a^H \to -a^H$, every s on an edge incident to H flips;

so individual $s_{HH'}$ are gauge-dependent, but products over closed cycles are gauge-invariant (each hub in a cycle is incident to exactly two cycle edges, and its flip cancels).

4.3 The reachability criterion

Theorem II-4.3 (Hub-Adjacent Reachability Criterion) [Proven]. A consistent hub-adjacent realization of the alternating mode across Γ_{hub} exists if and only if the matching cochain s is **balanced**: every closed hub cycle γ has sign product +1,

$$\prod_{(HH') \in \gamma} s_{HH'} = +1 \text{ for every cycle } \gamma \subseteq \Gamma_{hub},$$

equivalently the matching cohomology class $[s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$ is trivial.

Proof. Fix a basepoint hub H_0 with amplitude $a^{H0} = a$. For any path $H_0 \to H_1 \to \cdots \to H_n$ in Γ_{hub}, transporting the matching condition along the path gives

$$a_{H_n} = \left(\prod_{k=0}^{n-1} s_{H_k H_{k+1}} \right) a.$$

A consistent assignment requires a_{H_n} to be independent of the path chosen. If two paths from H_0 to H_n give different signs, their concatenation (one path followed by the reverse of the other) is a closed cycle with sign product −1. Hence consistency requires every closed cycle to have sign product +1. Conversely, if every closed cycle has product +1, any two paths between the same endpoints differ by closed cycles of product +1, so they assign the same amplitude; the assignment a_H is then globally well-defined by transport from H_0, and on each connected component the amplitude is determined up to the single basepoint gauge choice. The cycle-product condition is precisely the statement that the $\mathbb{Z}_2$ cochain s is a coboundary — $[s] = 0$ in $H^1(\Gamma_{hub}; \mathbb{Z}_2)$ — i.e. s is balanced as a signed graph (Harary balance).

Corollary II-4.4 (frustration is an N_2 witness) [Proven]. If some cycle γ has $\prod_\gamma s_{HH'} = -1$, the alternating mode does not extend consistently across γ: there is a germ satisfying L on the hub-adjacent neighbourhood with no admissible realizer. By the Part I, Theorem 5.2 (Case B), such a germ is an N_2 witness and the hub-adjacent stratum fails Gate 2.

Proof. By Theorem II-4.3, a frustrated cycle obstructs the consistent assignment, so no transport realizes the boundary-matched alternating germ on γ; the germ is in $\mathfrak{U} \cap \mathcal{P}_2^{\text{hub-adjacent}}$, which is then non-empty.

4.4 Balance is not bipartiteness; inter-hub parity is not intra-hub parity

Remark II-4.5 (balance ≠ bipartiteness) [Methodological]. The criterion is the triviality of $[s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$ — *balance* of the signed graph (Γ_{hub}, s) — not bipartiteness of the unsigned graph Γ_{hub}. Bipartiteness is precisely the property that the all-minus signing (every edge −1) is balanced — equivalently that Γ_{hub} has no odd cycles — and is neither necessary nor sufficient for balance of the *actual* matching s: a non-bipartite Γ_{hub} can carry a balanced s (if every odd cycle happens to receive an even number of −1 edges from the matching rule), and a bipartite Γ_hub can carry an unbalanced s (if the matching rule assigns signs whose cycle products are −1). The "hub-adjacent pass ⟺ bipartite" reading is too crude and is **not** what is proven; the correct object is the signed graph (Γ_hub, s) and its balance.

Remark II-4.6 (two parity questions, two graphs) [Methodological]. The even-N fact of Remark II-3.3 — that C_6 admits the alternating mode because 6 is even — is a parity statement about the **intra-hub** outer cycle of a single hub. The balance condition of Theorem II-4.3 is a parity statement about the **inter-hub** cycles of Γ_hub. These are logically independent: the first lives on the six-vertex outer ring of one hub and governs whether the single-hub mode exists; the second lives on the hub-adjacency graph and governs whether existing single-hub modes glue. N = 6 being even does the first job and contributes nothing to the second. Conflating them — arguing the hub-adjacent stratum passes "because 6 is even" — is an error the framing must avoid.

4.5 What the criterion needs to be evaluated

Remark II-4.7 (the two missing inputs, and how the rest of this paper treats them) [Methodological]. Theorem II-4.3 is a proven structural criterion, but evaluating it on the K = 7 transport construction requires two data:

1. **The hub-adjacency graph Γ_hub** of the K = 7 refinement complex — which hubs are adjacent to which.
2. **The cross-boundary sign rule** fixing each s_{HH′} from the shared-edge A_comp — the inter-hub instance of the competition functional, evaluated across the boundary rather than within a single hub.

Both are inter-cell data — precisely the inter-cell completion structure Open Problem 6.7 named. We do not manufacture either. What we do in §13–15 is show that the *graph* Γ_hub turns out not to be needed for the verdict at all, and the *sign rule* reduces to named hypotheses: §13 proves that the gauge (hexagon-relabelling) part of s is always a coboundary, so the verdict depends only on the residue class [t] and not on Γ_hub's connectivity; §14 isolates the hypothesis (PBM) under which [t] is trivial; §15 extends this to the competition stratum and assembles the conditional verdict. The criterion of Theorem II-4.3 is what those sections evaluate; the class [t] is the object they reduce it to.

---

13. The Gauge/Holonomy Split: the Relabelling Freedom Is Harmless

The hub-adjacent criterion (Theorem II-4.3) decides the stratum by the balance of the matching cochain s. Before asking for the cross-boundary rule that fixes s, we separate the part of s that is mere bookkeeping from the part that carries physical content, and prove the bookkeeping part can never frustrate a cycle.

5.1 The two parts of the matching sign

We name the gauge degree of freedom explicitly.

Definition II-5.1 (hub gauge, physical gauge, and the orientation residue). [Methodological] A **hub gauge** is a choice of sign $\varepsilon_H \in \{\pm1\}$ for each hub, encoding the hexagon-origin label; a change $\varepsilon_H \to -\varepsilon_H$ is the one-step origin relabelling of H (Remark II-4.1). Under a hub gauge the matching sign transforms as

$$s_{HH'} \mapsto \varepsilon_H \, s_{HH'} \, \varepsilon_{H'},$$

the standard action of a $\mathbb{Z}_2$ gauge group on a $\mathbb{Z}_2$ edge cochain, since flipping the origin of H flips a^H and hence every s on an edge incident to H. The **physical gauge** is not a free choice but a definite datum of the construction: ε_H^{phys} is the actual relation between hub H's hexagon-origin label and the physical boundary modes incident to H. Fixing the physical gauge, define the **orientation residue** in that gauge by

$$t_{HH'} := \varepsilon_H^{phys} \, s_{HH'} \, \varepsilon_{H'}^{phys} \in \{\pm1\},$$

equivalently $s_{HH'} = \varepsilon_H^{phys} \varepsilon_{H'}^{phys} \cdot t_{HH'}$. We are explicit about what is and is not gauge-invariant. **Per-edge, $t_{HH'}$ is not gauge-invariant:** under a further hub gauge it transforms exactly as s does, $t_{HH'} \mapsto \varepsilon_H \, t_{HH'} \, \varepsilon_{H'}$; the decomposition s = (coboundary)·t is not unique, and any gauge yields a different per-edge t. **What is invariant is the class:** $[t] = [s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$, i.e. the cycle products $\prod_\gamma t_{HH'}$, which are gauge-independent (each hub in a cycle is incident to two cycle edges, so its flip cancels — Lemma II-5.2). The per-edge values $t_{HH'}$ below are therefore always read in the physical gauge, where they are definite numbers fixed by the construction's actual labels; the invariant content of every conclusion is the class [t] and its cycle products. The residue encodes whether the shared boundary contributes a relative orientation beyond the two hub frames — a class-level, not an edge-level, fact.

5.2 The gauge part is always a coboundary

Lemma II-5.2 (coboundaries are balanced; the ε-part is a coboundary) [Proven]. Let $c_{HH'} = \delta_H \delta_{H'}$ for any hub signs $\delta_H \in \{\pm1\}$. Then for every closed cycle $\gamma \subseteq \Gamma_{hub}$, $\prod_{(HH') \in \gamma} c_{HH'} = +1$. In particular the factor $\varepsilon_H^{phys} \varepsilon_{H'}^{phys}$ in Definition II-5.1 contributes +1 to every cycle product, and the balance class of s equals that of the residue alone:

$$\prod_\gamma s_{HH'} = \prod_\gamma t_{HH'} \text{ for every cycle } \gamma, \text{ and } [s] = [t] \in H^1(\Gamma_{hub}; \mathbb{Z}_2).$$

Proof. In a closed cycle each hub is incident to exactly two cycle edges, so each δ_H appears exactly twice in the product, contributing $\delta_H^2 = +1$; hence $\prod_\gamma c_{HH'} = +1$. Applying this with $\delta = \varepsilon^{phys}$ gives $\prod_\gamma \varepsilon_H^{phys} \varepsilon_{H'}^{phys} = +1$, so $\prod_\gamma s_{HH'} = \prod_\gamma(\varepsilon_H^{phys} \varepsilon_{H'}^{phys}) \cdot \prod_\gamma t_{HH'} = \prod_\gamma t_{HH'}$. Since cycle products determine the cohomology class, [s] = [t]. (The conclusion is gauge-independent even though the per-edge $t_{HH'}$ are not: the cycle product is invariant under any further hub gauge by the same two-incidences cancellation.)

Theorem II-5.3 (hub-adjacent pass reduces to the residue class) [Proven]. The hub-adjacent stratum passes — $\mathfrak{U}$_hub-adjacent = ∅ — if and only if the residue class is trivial,

$[t] = 0 \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$, equivalently $\prod_{(HH') \in \gamma} t_{HH'} = +1$ for every cycle $\gamma \subseteq \Gamma_{hub}$.

Proof. By Theorem II-4.3 the stratum passes iff [s] = 0; by Lemma II-5.2 [s] = [t]; hence the stratum passes iff [t] = 0, which is triviality of the residue class. The criterion is stated at class/cycle level, where it is gauge-invariant; it does not depend on the per-edge values $t_{HH'}$ in any particular gauge.

This is the unconditional content of the merge beyond the criterion itself: the hexagon-origin relabelling — the obvious candidate for an apparent sign flip — is proven pure gauge, cancelling around every cycle (Lemma II-5.2). The hub-adjacent obstruction, if any, is carried entirely by the residue class [t], and by nothing in the labelling; in particular the verdict no longer depends on the connectivity of Γ_{hub}, only on whether [t] is trivial. The per-edge residue values $t_{HH'}$ are gauge-dependent and meaningful only once a gauge is fixed (the physical gauge of Definition II-5.1); the verdict itself is a statement about their cycle products, which are gauge-invariant.

---

14. The Shared-Boundary Sign Lemma

The residue $t_{HH'}$ is fixed by what the admissible shared-boundary completion term actually compares. We name the deciding hypothesis and derive the pass under it.

Hypothesis (PBM) — Physical-Boundary Matching. [Hypothesis, not proven here] The admissible K = 7 shared-boundary completion term $A_{comp}^{HH'}$ between adjacent hubs H, H′ compares the **same physical boundary mode** seen from each hub — the boundary spoke/vertex data identified across the shared cell — rather than the two arbitrary local hexagon-origin labels i, j. Equivalently: the shared boundary carries no gauge-invariant relative orientation of its own, independent of the two hub frames; any apparent relative sign arises solely from the two origin labels.

Remark II-6.1 (why PBM is a hypothesis and not a theorem) [Methodological]. PBM is a statement about the transport construction's boundary-completion rule: it asserts the rule is formulated on the physical boundary mode, so the only freedom in $s_{HH'}$ is the two hub gauges. It is plausible — a completion rule comparing arbitrary labels rather than the shared physical data would be ill-posed under relabelling — but plausibility is not proof. Proving PBM requires the explicit shared-boundary $A\_comp^{HH'}$ of the transport construction, the inter-cell datum Open Problem 6.7 named and which we do not have in hand. We therefore state PBM as a hypothesis, label it at every use, and decline to manufacture the boundary rule that would discharge it. The honest status is: **PBM ⟹ pass; PBM is not proven.** We did not assert [s] = 0 as a gauge-free fact derived from nothing — that would beg the question, since "any sign flip is gauge" is the conclusion, not a premise. Lemma II-5.2 proved the ε-part is a coboundary unconditionally (so the verdict depends only on the residue class [t]); PBM is the one hypothesis that fixes that class to the identity, equivalently fixes the physical-gauge residue $t_{HH'}$ to +1.

Lemma II-6.2 (Shared-Boundary Sign Lemma) [Conditional on PBM]. Under PBM, the orientation residue is trivial in the physical gauge:

$t_{HH'} = +1$ for every adjacent pair (H, H′), in the physical gauge of Definition II-5.1.

Consequently the residue class is trivial, [t] = [s] = 0, and the hub-adjacent stratum passes: $\mathfrak{U}\_hub\text{-}adjacent = \emptyset$. (Equivalently and gauge-invariantly: under PBM every cycle product $\prod\gamma$ $s\{HH'\} = +1$.)

Proof. Under PBM, $A\_comp^{HH'}$ compares the same physical boundary mode from both hubs, so the relative sign between a^H and aH' is determined entirely by how each hub's *physical-gauge* origin label relates to that shared mode — i.e. by ε_H^{phys} and $\varepsilon\_{H'}^{phys}$ — with no contribution from the boundary itself. Hence in the physical gauge $s\_{HH'} = \varepsilon\_H^{phys}$ $\varepsilon\_{H'}^{phys}$, so the physical-gauge residue $t\_{HH'} = \varepsilon\_H^{phys}\ s\_{HH'}\ \varepsilon\_{H'}^{phys}$ $= (\varepsilon\_H^{phys}\ \varepsilon\_{H'}^{phys})^2 = +1$ for every edge. By Lemma II-5.2 the cycle products are then all +1, so [t] = [s] = 0; by Theorem II-5.3 the stratum passes. The gauge-invariant content is [s] = 0; the edge-by-edge $t\_{HH'} = +1$ is its reading in the physical gauge, not a separate claim. We have not asserted s = (coboundary) as a gauge-free identity — PBM is precisely the hypothesis that fixes the physical-gauge residue to the identity, and Remark II-6.1 records that this is a hypothesis, not a theorem.

Remark II-6.3 (the lemma is exactly as strong as PBM) [Methodological]. Lemma II-6.2 derives the pass from PBM and nothing else; it adds no independent assumption. If PBM holds, the physical-gauge residue is $t\_{HH'} = +1$ on every edge, hence [t] = 0, and the hub-adjacent stratum is unfrustrated regardless of Γ_hub's connectivity — no balance computation on the graph is even needed, because the residue class is already trivial. If PBM fails — if the shared boundary does carry a relative orientation beyond the two hub frames — then [t] = [s] may be non-trivial, and a frustrated cycle (Corollary II-4.4) is an N_2 witness. The lemma does not adjudicate which; it shows PBM is the single switch, and the thing it switches is the class [t], not a gauge-dependent edge value.

---

15. The Competition Stratum and the Registered-Search Verdict

The third registered stratum (competition germs) tests whether the cross-boundary competition term creates a new unrealizable germ once single-hub alternation and hub-adjacent sign-consistency hold.

Hypothesis (AC) — Amplitude Compatibility. [Hypothesis, not proven here] The cross-boundary completion term $A_{comp}^{HH'}$ forces $|a^H| = |a^{H'}|$ across each shared boundary: the matched alternating mode extends at equal amplitude, not merely equal sign. Equivalently, the cross-boundary term relates the two ray parameters by a unit ratio on the admissible sector.

Remark II-7.1 (why AC is needed beyond PBM) [Methodological]. PBM fixes the *sign* relation $a^{H'} = s_{HH'} a^H$. It does not fix the *magnitude*: the single-hub analysis leaves each $|a^H|$ a free ray parameter, and the cross-boundary $A_{comp}^{HH'}$ term — coupling the boundary spokes of two different hubs — is the first structure that could relate $|a^H|$ to $|a^{H'}|$. If it forces the ratio to 1, competition vanishes on the matched ray; if it forces a ratio $\neq 1$, the alternating mode does not extend at equal amplitude and a residual competition cost survives. AC is the hypothesis that the ratio is 1. Like PBM, it is a statement about the explicit cross-boundary term, stated here and not derived; and it is independent of PBM (signs can match while amplitudes differ).

Hypothesis (AA) — Anti-Aligning Match. [Hypothesis, not proven here] The shared-boundary identification phases the two hubs' boundary modes in anti-phase: the matched boundary spokes satisfy $\lambda_{bd}^{(H)} = -\lambda_{bd}^{(H')}$ on the admissible ray, equivalently the cross-boundary nearest-neighbour pair behaves as an intra-hub adjacent pair (which alternates automatically). AA is logically distinct from PBM and AC: PBM fixes that there is a single shared mode and no intrinsic orientation, AC fixes equal magnitude, and AA fixes that the matching is anti-aligning rather than co-aligning. Intra-hub, anti-alignment is forced by the alternation of adjacent spokes; cross-boundary it is a property of how the identification phases the two modes, and is not delivered by PBM or AC alone.

Lemma II-7.2 (competition completion) [Conditional on PBM ∧ AC ∧ AA]. Suppose every $K = 7$ hub carries the alternating mode $\lambda_i^{(H)} = (-1)^i a^H$, that PBM holds (so signs match on shared boundaries, Lemma II-6.2), that AC holds (so $|a^H| = |a^{H'}|$ across each shared boundary), and that AA holds (so the matching is anti-aligning across the boundary). Then:

(a) **Intra-hub:** for every adjacent spoke pair within a hub, $\lambda_i^{(H)} + \lambda_{i+1}^{(H)} = (-1)^i a^H + (-1)^{i+1} a^H = (-1)^i a^H (1 - 1) = 0$, so $A_{comp}^{H} = 0$.

(b) **Cross-boundary:** the cross-boundary competition term is $A_{comp}^{HH'} = (\lambda_{bd}^{(H)} + \lambda_{bd}^{(H')})^2$, which vanishes iff the two boundary spokes are anti-aligned, $\lambda_{bd}^{(H)} = -\lambda_{bd}^{(H')}$. Under PBM the two spokes carry the same physical boundary mode with relative sign $s_{HH'}$, under AC with equal magnitude, and under AA that relative sign is the anti-aligning one; so $\lambda_{bd}^{(H)} + \lambda_{bd}^{(H')} = 0$ and $A_{comp}^{HH'} = 0$.

Hence no competition-germ obstruction survives, and the competition stratum passes: $\mathfrak{U}_{comp} = \emptyset$.

Proof. (a) is the single-hub computation (Lemma II-3.4), holding on each hub's ray. (b): the cross-boundary $A\_comp^{HH'} = (\lambda\_bd^{(H)} + \lambda\_bd^{(H')})^2$ is a single squared nearest-neighbour sum across the boundary. By PBM the two terms are the same physical mode (one shared spoke amplitude seen from each side), by AC of equal magnitude, and by AA in anti-phase, so $\lambda\_bd^{(H)} + \lambda\_bd^{(H')} = 0$ and the term vanishes — this is the step that the intra-hub case (a) gets for free from alternation and that AA supplies across the boundary. With both intra-hub and cross-boundary competition terms zero on the matched ray, the competition functional is globally minimized at zero and contributes no germ to $\mathfrak{U}$.

Remark II-7.3 (the competition pass inherits three hypotheses) [Methodological]. Lemma II-7.2 is [Conditional on PBM ∧ AC ∧ AA] — strictly weaker standing than the hub-adjacent pass, which needs PBM alone. The three are logically independent: PBM fixes that a single shared mode is compared with no intrinsic orientation (sign structure), AC fixes equal magnitude across the boundary, and AA fixes that the matching is anti-aligning rather than co-aligning. PBM can hold with AC failing (matched signs, mismatched amplitudes — residual competition), with AA failing (a single mode compared but co-aligned, so the nearest-neighbour sum doubles rather than cancels), or with both failing. We record all three explicitly so the competition pass is not read as following from PBM, or from PBM ∧ AC, alone.

Theorem II-7.4 (conditional registered-search pass) [Conditional on C1–C3, PBM, AC, AA]. On the pre-registered finite K = 7 search space $\mathcal{P}_2$ (single-hub, hub-adjacent, competition strata),

$$\mathfrak{U} \cap \mathcal{P}_2 = \emptyset,$$

conditional on: the single-hub execution conditions C1–C3 (§11.5; C1 confirmed); physical-boundary matching PBM (§14); and amplitude compatibility AC together with anti-aligning match AA (§15). No N_2 witness appears in the registered K = 7 local search space under these hypotheses.

Proof. By Theorem II-3.6, $\mathfrak{U} \cap \mathcal{P}_2^{\{single\text{-}hub\}} = \emptyset$ under C1–C3. By Lemma II-6.2, $\mathfrak{U}\_hub\text{-}adjacent = \emptyset$ under PBM. By Lemma II-7.2, $\mathfrak{U}\_comp = \emptyset$ under PBM ∧ AC ∧ AA. The three strata exhaust $\mathcal{P}_2$ (Pre-Registration 9.1(a)), so their union residue is empty under the conjunction of conditions.

Remark II-7.5 (what this verdict is and is not) [Methodological]. The verdict is a clean pass on the *registered finite search space*, conditional on three named, labelled hypotheses. It is **not**: an *unconditional* pass — C1–C3, PBM, AC, AA are not all discharged here (C1 is; the rest are hypotheses or finite checks); a *global* Gate-2 pass — a registered-search pass equals a global pass only if the registered strata exhaust {L}, which requires the prior programme's Conjecture 8.3 (boundedness of the realizability-test radius), absent which the pass is [Conditional] over $\mathcal{P}_2$ at the registered complexity, not over all of {L}; or a statement about Gate 3 — global assembly over the full complex is untouched. What is proven unconditionally (Theorem II-5.3): the gauge part of the matching sign is always a coboundary, so the registered-search pass reduces *exactly* to PBM ∧ AC ∧ AA and the single-hub conditions, with no further hidden content.

16. What the Execution Narrows in Open Problem 6.7

Open Problem 6.7 asked for the full four-sector cell-completion catalogue {I_inc, I_hub, I_circ, I_comp} as explicit conditions around the $K = 7$ hub. This paper narrows that open input at the two executed strata.

At the single-hub stratum, two of the four sector relations are now instantiated concretely:

$I\_circ(\hat{G}) \Leftrightarrow A\_circ(\lambda) = 0$, $I\_comp(\hat{G}) \Leftrightarrow A\_comp(\lambda) = 0$,

with {I_comp} ∩ {single-hub} = alternating ray and {I_circ} ⊇ that ray automatically (Lemma II-3.4). The matching property's ⊇ direction (ray ⊆ Im(ρ)) is supplied by the σ-family (Lemma II-3.5). So at the single-hub neighbourhood, the Part I, Definition 6.7a relations I_circ and I_comp are no longer named-but-unsupplied — they are read off the concrete functionals, and the single-hub residue is empty under C1–C3.

At the hub-adjacent and competition strata, the open input is compressed further still. The hub-adjacent stratum first reduces to balance of the matching cochain s (Theorem II-4.3); §13 then proves the gauge part of s is always a coboundary, so the residue alone matters (Theorem II-5.3); and §14 reduces that residue to a single named hypothesis PBM. So where Open Problem 6.7 asked for a catalogue of completion relations, the hub-adjacent stratum requires only:

the hypothesis that the shared-boundary completion compares physical boundary modes (PBM),

with no balance computation on Γ_hub needed under it (the physical-gauge residue is +1 on every edge, so $[t] = 0$). The competition stratum adds two further hypotheses, amplitude compatibility AC and anti-aligning match AA. The four-sector completion question, across all three registered strata, thus collapses to: the concrete single-hub relations I_circ, I_comp; the hypothesis PBM (fixing the residue class $[t] = 0$); and the hypotheses AC, AA (fixing the cross-boundary amplitude ratio and phase). This is a genuine narrowing — the remaining open content is the stated hypotheses about the explicit shared-boundary term, not a four-sector catalogue. It does not close 6.7: the explicit A_comp^{HH′} that would discharge PBM, AC, AA remains to be supplied, and we do not manufacture it.

Status against Part I, Theorem 6.9. The conditional reduction is now discharged across the registered strata up to PBM ∧ AC ∧ AA: at the single-hub neighbourhood, I_circ and I_comp are concrete and $\mathfrak{U}\_circ = \mathfrak{U}\_comp = \emptyset$ (modulo C1–C3); at the hub-adjacent neighbourhood, the gauge part of the completion is proven harmless (Theorem II-5.3) and the residue is trivial under PBM; at the competition neighbourhood, the cross-boundary term vanishes under PBM ∧ AC ∧ AA. The sectors I_inc and I_hub are confirmed to vanish on the admissible spoke sector (no independent on-sector residue) but their explicit inter-cell forms are not separately exhibited. The

[Conditional-on-supply] label of Theorem 6.9 is, at these strata, reduced to [Conditional on PBM ∧ AC ∧ AA] — the supply needed is exactly the explicit shared-boundary term that decides those three hypotheses.

---

17. Limitations

This paper does not return an unconditional or global Gate-2 verdict. It proves the reduction and runs the registered search to a conditional pass.

The reduction (Part I) is conditional on its single named input. Theorem 6.9 is proven modulo the supply of the four cell-completion relations {I_inc, I_hub, I_circ, I_comp} (Open Problem 6.7); §6.1–6.2 remain a heuristic localization with no unconditional ρ-free integrability predicate. The relations are not manufactured from the wheel transition rules or the operational-admissibility ledger, which govern different objects and would yield a predicate failing the matching property except by coincidence. Proposition 6.2 (conservation reachability-favourable) is [Conditional] on local solvability and is an in-isolation claim only (Part I, Remark 6.3); Proposition 7.2's falsification condition is [Conditional] on the M1/M2/M3 trichotomy being complete; Conjecture 8.3 (minimality suffices) is unproven and is the only route to a [Proven] rather than [Conditional] pass.

The single-hub pass (Part II) is conditional on C1–C3. C1 (PSD reading of A_comp) is confirmed. C2 (admissibility of the σ-family realizer in $\mathcal{T}$_L) requires only one realizing family — by the union of Part I, Proposition 7.1 a single family suffices for a pass; the all-families condition of Proposition 7.2 is the falsification burden for a witness, not a precondition of the pass. C3 (non-vacuity under A_inc ∧ A_hub) is a finite check not discharged here; absent it a vacuous pass is not formally excluded.

The hub-adjacent and competition passes are conditional on PBM, AC, AA — hypotheses, not theorems. Theorem II-5.3 (the ε-part is a coboundary; the hub-adjacent pass reduces to triviality of the residue class [t]) is unconditional. That [t] = 0 holds under PBM; equal-amplitude and anti-aligning cross-boundary matching (AC, AA) are needed for the competition stratum. None of PBM, AC, AA is derived here; discharging them requires the explicit shared-boundary completion term A_comp^{HH′} — the same inter-cell datum Open Problem 6.7 names. If PBM fails the residue class is live and a frustrated cycle is an N_2 witness; if AC or AA fails, residual cross-boundary competition survives. We are explicit that balance is not bipartiteness (Remark II-4.5), and that the inter-hub cycle parity is logically independent of the intra-hub even-N fact (Remark II-4.6).

Registered-search pass ≠ global Gate-2 pass. The three strata exhaust the *registered* $\mathcal{P}_2$, not all of {L}. Promotion to a global pass requires Conjecture 8.3 to bound the realizability-test radius at the registered complexity; absent it, the verdict is [Conditional] over $\mathcal{P}_2$ at the registered complexity, not [Proven] surjectivity of ρ onto {L}.

Inherited and downstream. The whole result rests on Gate 1, itself conditional on four-sector inventory completeness (*The Locality Inspection* A.4). The paper says nothing about Gate 3: global assembly, the realizable obstruction image $\mathfrak{o}(\hat{G})$, and gluing-triviality remain entirely open; a Gate-2 pass fixes Gate 3's domain but does not bear on its verdict. Nor does it address FBI-comp, RC, or the ℓ^2 terminal arena beyond fixing what the residue count would feed them.

---

18. Conclusion

The Locality Inspection established that the admissibility constraints of the transport construction are local in form. The question is therefore no longer about the rules but about the geometries they permit: are the locally admissible germs reachable by admissible transport? This is Gate 2, and the paper both reduces it and runs the resulting search.

The reduction. Gate 2 is not a new check but the second half of the first. Gate 1 proved the local predicate L is *sound* for realizability, $\mathrm{Im}(\rho) \subseteq \{L\}$; Gate 2 is its converse, the *completeness* $\{L\} \subseteq \mathrm{Im}(\rho)$, and the two are the single equality $\{L\} = \mathrm{Im}(\rho)$. Reachability non-locality is the non-emptiness of the gap: the unrealizable residue $\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$, a single element of which fails Gate 2. We located the residue and hardened the localization into a conditional reduction: fixing the four-sector form of the intrinsic predicate, the residue is exactly the hub-completion-failing locus, decomposes sector-wise, and Gate 2 reduces to per-sector hub-completion on {L} (Theorem 6.9) — modulo the one named input, the four cell-completion relations (Open Problem 6.7). We then deposited, in advance, the search that decides it: the complexity-graded space $\mathcal{P}_2$, its representativeness, and the all-families falsification condition, so that whatever the execution returns carries the standing its register earns.

The execution. Running the registered search on the concrete $K = 7$ forms: at the single-hub stratum the competition functional's kernel on the even cycle C_6 is exactly the alternating ray, conservation vanishes there automatically, and the σ-family realizes it — so the single-hub residue is empty (conditional on C1–C3). At the hub-adjacent stratum the alternating mode's free per-hub sign becomes a $\mathbb{Z}_2$ gauge field on the hub network, and the stratum passes iff the matching cochain is balanced; we proved that the gauge (hexagon-relabelling) part of the matching sign is always a coboundary, so the verdict depends only on the residue class [t], and under physical-boundary matching that class is trivial. Under physical-boundary matching at equal amplitude in anti-phase (PBM ∧ AC ∧ AA), the competition stratum passes too. Assembling the three strata, the registered finite $K = 7$ search contains no N_2 witness — conditional on C1–C3, PBM, AC, AA. We kept two distinctions straight (balance is not bipartiteness; inter-hub parity is not the intra-hub even-N fact), stated PBM, AC, AA as hypotheses rather than theorems, and declined to manufacture the shared-boundary rule that would discharge them.

Stated in one line: **Gate 2 is the completeness of the local admissibility predicate against the realizer map; it holds iff the residue $\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$ is empty; the residue reduces sector-wise to a four-sector hub-completion catalogue (Theorem 6.9, the one open input being**

Open Problem 6.7); and on the concrete K = 7 geometry the pre-registered search returns $\mathfrak{U} \cap \mathcal{P}_2 = \emptyset$ conditional on C1–C3 and on the shared-boundary matching being physical-mode, equal-amplitude, and anti-aligning (PBM, AC, AA) — with the unconditional content being that the gauge freedom is harmless and these named hypotheses are the entire remaining gap.

Gate 1: PASS ⟶ Gate 2 single-hub: PASS [Conditional C1–C3] ⟶ hub-adjacent: PASS [Conditional PBM] ⟶ competition: PASS [Conditional PBM ∧ AC ∧ AA] ⟶ registered search: $\mathfrak{U} \cap \mathcal{P}_2 = \emptyset$ [Conditional C1–C3, PBM, AC, AA] ⟶ global Gate 2 [requires Conjecture 8.3] ⟶ Gate 3: open.

The next genuine gate is Gate 3 — global assembly over the full complex, the realizable obstruction image, and gluing-triviality — whose domain is now fixed by the registered-search pass.