Constitutive Completion and the Gate-2 Verdict

Deriving the K = 7 Boundary-Completion Rule, Identifying the Orientation Obstruction, and Returning a Verdict on the Registered Reachability Search

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

General Reader Summary

The previous paper narrowed the second locality check — Gate 2 — to a very small set of remaining questions. It showed that the single-hub K = 7 structure works: the local rules force an alternating pattern around a hub, and that pattern is realized by transport. It then showed that the next level — matching neighbouring hubs across a shared boundary — depended on three things being true of the shared-boundary rule: that it compares the same physical mode from both sides (PBM), that it forces equal amplitude (AC), and that the two sides cancel in the correct anti-aligned way (AA). Those three were left as hypotheses.

This paper asks whether those conditions are extra assumptions, or whether they follow from the construction itself. The honest answer turns out to be sharper and more interesting than a flat "they follow."

Two of the three — physical-mode matching and equal amplitude — *do* follow, with no further input, from requiring the boundary rule to be local, label-independent, and a genuine cost (a positive quantity that is zero only at exact completion). The third — anti-alignment — does **not** follow from those requirements alone. The local requirements leave exactly two candidate rules: one whose lowest-cost state is anti-aligned (the two sides cancel) and one whose lowest-cost state is co-aligned (the two sides reinforce). They are mirror images, equally local and equally valid as cost functions. Something outside the local cost has to choose between them.

That something is the **orientation** of the construction — the same structural fact that, inside a single hub, makes neighbouring spokes cancel rather than reinforce. Inside a hub the two spokes bounding a cell meet it with opposite induced orientation, which is why the rule reads "sum, then square" rather than "difference, then square." Across a shared boundary the analogous question is whether gluing two hubs identifies the shared spoke with the same induced orientation from both sides or with opposite induced orientations. If the complex is *orientation-coherent* — built so that adjacent cells induce opposite orientations on their shared face, the ordinary condition for an oriented surface — then anti-alignment is forced across the boundary by exactly the mechanism that forces it inside a hub, and the gate passes. If the complex is *not* orientation-coherent — if

following a loop of hubs around the complex returns you with a flipped orientation, the way a Möbius band does — then some boundary carries the co-aligning rule, the alternating pattern cannot be assigned consistently around that loop, and that frustration is a genuine unbuildable-but-admissible geometry: a Gate-2 failure.

So the verdict is not unconditional. It is a clean function of one named, binary structural property of the construction:

Gate 2 passes for the K = 7 nearest-neighbour transport construction if and only if the substrate gluing is oriented — built as an oriented 2-cell complex, with adjacent cells inducing opposite orientations on their shared faces. Under oriented gluing the registered search contains no failure witness; under a frustrated orientation loop, we exhibit the failure witness explicitly.

A point worth stating plainly, because it fixes exactly what kind of input this is. The K = 7 construction gives the local wheel W_6 — one hub, six boundary states, six spokes, a cyclic ring — and that proves each hub is *locally* orientable and carries the alternating mode. It does **not**, by itself, prove that all hubs are glued orientation-coherently across the whole substrate: local orientability never implies global orientability, the way a Möbius band is built from perfectly orientable strips. So oriented gluing is not derivable from K = 7 alone; it is an additional, standard, structural premise about how the wheels are assembled into the substrate. K = 7 supplies local orientability; oriented gluing supplies the global coherence.

The signed sign-cocycle the previous paper introduced is, we show, precisely the orientation class of the hub complex. Its triviality *is* oriented gluing. The previous papers had located the obstruction; this paper identifies what it is — and identifies the single structural premise (oriented gluing) that decides it.

We can also say something about how that orientation condition is checked, though here we are careful not to claim more than we have shown. The condition itself is clean and unconditional: the gate passes exactly when the vacuum carries no hidden orientation twist — a definite, checkable property of the substrate, owing nothing to the rest of the machinery. What the wider framework *might* add is a labour-saving shortcut. The framework already contains a "loop-consistency" rule: any small loop that can be shrunk to a point must close up trivially. If — and this is the part not yet established — the orientation information lives inside the same object that this rule constrains, then the rule disposes of all the shrinkable loops for free, and one need only check the loops that *cannot* be shrunk (loops that wrap around the whole space, the way a loop around a doughnut's hole cannot be tightened away). That "if" is a specific, stateable identification between two pieces of the framework, and the relevant background paper does not yet prove it; we record it as a clearly-marked assumption to be settled later. The honest summary: the pass condition is a single checkable property of empty space, established here and depending on nothing further; the framework's loop-consistency rule would, *if* one identification is confirmed, reduce the work of checking it to just the non-shrinkable loops — a convenience, not a change to the condition. A pleasant by-product, carrying the same "if," is that orientation defects (should they exist) would have to sit on exactly the non-shrinkable loops — precisely how a localized, persistent object ought to behave, and the thread a separate paper picks up.

We also prove the half of the boundedness question that is genuinely available: for a construction whose admissibility cost couples no more than two hubs at a time, the *constraint* content is exhausted by the single-hub and hub-adjacent neighbourhoods, so the registered search space captures every way the local rules can fail. The remaining half — that local realizers actually glue into global ones — is not free; it is the orientation-coherence condition again, and it is the natural hinge into Gate 3.

The verdict is therefore a closure: a definite answer as a function of orientation-coherence, with the pass and the witness both made explicit — the orientation condition being a single checkable property of the vacuum, with the framework offering (pending one assumption) a shortcut for checking most of it. Gate 3 — whether locally realizable pieces assemble over the full complex — remains open, and orientation-coherence is exactly where it reconnects.

Abstract

The preceding Gate-2 inspection reduced reachability of admissible degeneracy geometry to the emptiness of the unrealizable residue $\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$, executed the registered $K = 7$ search to a conditional pass, and left three shared-boundary hypotheses — Physical-Boundary Matching (PBM), Amplitude Compatibility (AC), Anti-Aligning Match (AA) — together with the boundedness question Conjecture 8.3. This paper supplies the constitutive step and returns a verdict.

The boundary rule is determined up to one binary datum. We require the shared-boundary completion functional to be (B0) exchange-symmetric — neither hub privileged — (B1) local in the two matched boundary amplitudes, (B2) label-independent, (B3) a positive-semidefinite cost, and (B4) built from the same squared-linear competition primitive as the intra-hub functional $A_{comp}^H = \Sigma_i (\lambda_i^H + \lambda_{i+1}^H)^2$. These requirements reduce the general symmetric quadratic to a **two-element family** — the two PSD perfect squares

$A_\partial^{HH'} = (\lambda_\partial^H + \lambda_\partial^{H'})^2$ (anti-aligning, zero set $\lambda_\partial^{H'} = -\lambda_\partial^H$), $A_\partial^{HH'} = (\lambda_\partial^H - \lambda_\partial^{H'})^2$ (co-aligning, zero set $\lambda_\partial^{H'} = +\lambda_\partial^H$).

The division of labour is exact: B4 collapses the form to rank one ($\beta^2 = \alpha\gamma$); B0 equalizes the diagonal ($\alpha = \gamma$); together they fix $|\beta| = \alpha$. B0–B4 do **not** select between the two squares; the selection is a sign $\beta = \pm\alpha$ that the local cost cannot fix. We make this explicit rather than assert anti-alignment.

The selector is oriented gluing (B5), which $K = 7$ does not supply on its own. The intra-hub anti-aligning sign is forced by orientation: adjacent spokes meet their shared cell with opposite induced orientation, giving the sum-square. The cross-boundary sign is the same datum at the level of the gluing map. We isolate it as **(B5) oriented gluing** — the substrate is an oriented 2-cell complex, adjacent cells inducing opposite orientations on their shared face — and prove that **B0–B5 uniquely fix the anti-aligning term** (Theorem 4.4, [Proven]); thereunder the three prior

hypotheses trace to three distinct premises — PBM follows from B2 (Theorem 5.1, [Proven]), AC from the exchange symmetry B0 (Theorem 5.2, [Proven]), and AA from B5 (Theorem 5.3, [Proven]) — none redundant. We are explicit that B5 is **not** derivable from K = 7 alone: K = 7 proves each wheel locally orientable, but global oriented gluing — that local orientations cohere across the whole substrate — does not follow from local orientability (Remark 4.5), and is the one additional structural premise the verdict carries.

The sign-cocycle is the orientation class. We identify the previous paper's signed matching cochain s with the orientation cochain of the hub complex Γ_hub: $[s] \in H^1(\Gamma\_hub; \mathbb{Z}_2)$ is the first obstruction to a coherent orientation, the w_1-type class of the gluing. Hence (Theorem 6.3, [Proven]) [s] = 0 ⟺ B5 holds ⟺ the substrate gluing is oriented. Under B5 the residue is empty across all four sectors (F_𝔘 = ∅, Theorem 8.5); under a frustrated orientation loop γ with ∏_γ s = −1 we exhibit an explicit N_2 witness (Theorem 6.4, [Proven]).

Boundedness, split honestly. For a nearest-neighbour construction — admissibility cost coupling at most two hubs — we prove the *constraint-locality* half of Conjecture 8.3 unconditionally: L-satisfaction is decidable on single-hub and hub-adjacent neighbourhoods, so $\mathcal{P}_2$ exhausts the constraint content (Theorem 9.1, [Proven]). The *realizability-globalization* half — that pairwise-local realizers glue — is **not** free; it is exactly cocycle triviality [s] = 0, i.e. B5 (Theorem 9.3, [Conditional on B5]). We do not conflate the two; the gluing half is the orientation condition and the natural hinge into Gate 3.

Verdict.

𝔘 = ∅ ⟺ B5 (oriented gluing of the substrate), under C1 (PSD reading, confirmed), C2 (σ-family admissibility, finite check), C3 (non-vacuity, finite check). Equivalently (under C1–C3): **Gate 2 PASS for the oriented K = 7 nearest-neighbour substrate** — the entire remaining content being the oriented-gluing premise, which K = 7 supplies locally but not globally; under a frustrated orientation loop, Gate 2 FAIL with the witness of Theorem 6.4.

B5 is the checkable condition [s]_vac = 0; the interface axiom buys a verification shortcut. Unconditionally, under C1–C3, B5 ⟺ [s]*vac = 0 (Theorem 6.2/10.1 restricted to the vacuum) — the pass criterion is the triviality of the vacuum orientation class, and nothing from the interface dynamics enters it. Separately, the interface loop-consistency axiom — contractible loops close, U*{δγ} = 𝟙 — would discharge the *contractible-loop* part of [s]*vac = 0 automatically if the interface update's orientation (σ) component is identified with the Gate-2 matching cochain s (Lemma 10A.0'). That identification is **not yet in hand**: the interface-dynamics paper supplies loop-consistency and a two-component (σ, ω) update with σ linked to orientation perturbation, but does not currently establish that U*{δγ} acts on the Gate-2 boundary mode by the matching-sign product. We therefore deposit it as a named condition (Lemma 10A.0′, [Conditional on the interface-orientation identification]) — but it conditions only the *shortcut* (contractible loops for free, leaving non-contractible vacuum loops to check), **not** the pass criterion, which stands on Theorem 10.1 alone. We are equally careful that loop-consistency does *not* give [s] = 0 unqualified even granting the identification (the Möbius band satisfies it yet is non-orientable, its non-contractible loop carrying holonomy −1): simple connectivity of the vacuum is one *sufficient* route to [s]_vac = 0, not the criterion, and a non-simply-connected vacuum passes equally if its non-contractible loops carry trivial holonomy (Remark 10A.3). As a by-product — conditional on the same identification — the vacuum/defect split is sharpened: a non-trivial $[s]_\gamma$ can live only on a non-contractible loop, so orientation defects are defect-encircling content, never trivial-holonomy vacuum (§10B).

This closes the Gate-2 reachability branch as a verdict-valued function of one checkable condition — **[s]_vac = 0**, the triviality of the vacuum orientation class, unconditional under C1–C3 — with the interface axiom offered separately as the shortcut that makes most of that check automatic once the identification of Lemma 10A.0′ is established. It does not settle Gate 3; the residual vacuum-holonomy condition and the global gluing are where Gate 3 reconnects.

Gate 2 PASS iff [s]_vac = 0 (the vacuum orientation class is trivial), under C1–C3. Simple connectivity of the vacuum guarantees [s]_vac = 0 via loop-consistency but is not required; a non-simply-connected vacuum passes provided all its non-contractible loops carry trivial orientation holonomy.

Labelling convention retained: [Proven], [Conditional], [Conditional on B5], [Conditional on the interface-orientation identification], [Methodological], [Conjectural], [Open]. Where a result's topology is standard but an identification is definitional, the label states both halves (e.g. Theorem 6.2). The interface-orientation identification (Lemma 10A.0′) is deposited, not established — see §10A′ and Limitations.

Contents

---

1. From Hypotheses to a Constitutive Datum

Gate 1 established that the admissibility constraints of the transport construction are local in form. Gate 2 asks whether the geometries satisfying those local constraints are reachable by admissible transport, reduced by the prior inspection to the emptiness of the unrealizable residue

$\mathfrak{U} = \{L\} \setminus \mathrm{Im}(\rho)$,

where $\{L\}$ is the set of locally admissible germs and $\mathrm{Im}(\rho)$ the set of germs realized by admissible transport. The prior paper executed the registered $K = 7$ search to a conditional pass, leaving three shared-boundary hypotheses — PBM, AC, AA — and the boundedness question Conjecture 8.3.

The discipline of the inspection arc was to refuse to manufacture the construction's unstated boundary rule. That discipline produced ever-sharper conditionals but could not close, because closure requires *evaluating* the rule, not hypothesizing its properties. This paper descends to the construction and evaluates. The result is not a flat pass: the local cost requirements determine the boundary rule only up to a single binary sign, and that sign is fixed by an orientation datum of the complex. The verdict is therefore a clean function of that datum. We state it as such, prove the pass on the orientation-coherent branch, and exhibit the witness on the other.

---

2. Prior State of Gate 2

The $K = 7$ hub is the wheel W_6: one hub vertex, six outer vertices, six spokes, six outer edges, six triangular cells. The single-hub degeneracy datum is the six spoke amplitudes

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

The competition functional is

$A\_comp^H = \Sigma_i (\lambda_i + \lambda_{i+1})^2$, i mod 6,

a positive-semidefinite quadratic form whose kernel — established in the prior paper as the nearest-neighbour circulant on the even cycle C_6 with the Nyquist mode at $k = 3$ — is the alternating ray

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

The conservation functional is $A\_circ^H = (\Sigma_i \lambda_i)^2$. On the alternating ray,

$\Sigma_{i=0}^5 (-1)^i a = a(1 - 1 + 1 - 1 + 1 - 1) = 0$, so $A_{circ}^H = 0$ automatically ($N = 6$ even). The single-hub stratum passes [Conditional on C1–C3], with one realizing family (the σ-family) sufficing for the pass. The open question was the shared-boundary term between adjacent hubs, on which the three hypotheses PBM/AC/AA turned. This paper supplies that term and identifies what decides it.

The notational conventions of the prior paper are in force: each hub H carries $\lambda_i^H = (-1)^i a^H$ with a per-hub gauge flip $a^H \to -a^H$ (one-step hexagon-origin relabelling); the shared-boundary matching defines a $\mathbb{Z}_2$ cochain s on the hub-adjacency graph Γ_{hub}; in the construction's physical gauge $s_{HH'} = \varepsilon^{H}{}^{phys}\, \varepsilon^{phys}_{H'} \cdot t_{HH'}$, with the ε-part a proven coboundary, so the verdict is the triviality of the class $[t] = [s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$.

---

3. Constitutive Requirements for Boundary Completion

Let two $K = 7$ hubs H, H′ share a boundary. Write λ_∂^H, $\lambda_\partial^{H'}$ for the boundary spoke amplitude as carried by each hub. We require the boundary-completion functional $A_\partial^{HH'}$ to satisfy:

B0 (exchange symmetry — neither hub privileged). $A_\partial^{HH'}$ is symmetric under interchange of the two sides of the shared boundary, $H \leftrightarrow H'$: the shared face privileges neither hub. In the quadratic form below this is the statement $\alpha = \gamma$. B0 is a distinct requirement from locality and label-independence — it asserts the two hubs enter the boundary cost on equal footing — and it is the premise from which equal-amplitude matching (AC) will follow.

B1 (locality). $A_\partial^{HH'}$ depends only on the two matched boundary amplitudes λ_∂^H, $\lambda_\partial^{H'}$ — not on distant hubs, transport history, or global bookkeeping.

B2 (label-independence). $A_\partial^{HH'}$ is invariant under the per-hub hexagon-origin relabelling that defines the gauge of §2. It therefore depends on the boundary data only through the gauge-invariant physical boundary mode, not through either hub's arbitrary origin label.

B3 (positive-semidefinite cost). $A_\partial^{HH'} \geq 0$, with equality exactly at admissible completion. Zero cost means exact completion, not cancellation inside an indefinite form. This is the same reading certified for A_{comp}^H (the PSD reading C1).

B4 (competition primitive). $A_\partial^{HH'}$ is built from the same squared-linear competition primitive as the intra-hub functional — a perfect square of a linear form in the two boundary amplitudes — so that boundary completion is the cross-cell instance of the intra-hub competition rule, not a new species of constraint.

Remark 3.1 (what B0–B4 deliberately do not include) [Methodological]. B0–B4 are the *local cost* requirements. They say nothing about how the two hubs' boundary frames are identified across the shared cell — whether the shared spoke is seen with the same or opposite induced orientation from each side. That identification is not a property of a cost functional; it is a property of the gluing map of the complex. We isolate it separately as B5 in §4, precisely because conflating it with B4 is the error that would make the verdict assert what it should derive.

---

4. Derivation of the Boundary Term up to One Sign, and Its Fixing by Orientation

4.1 B0–B4 reduce the term to a two-element family

The most general real local quadratic in the two boundary amplitudes is

$$A_\partial^{HH'} = \alpha\,(\lambda_\partial^H)^2 + 2\beta\,\lambda_\partial^H\,\lambda_\partial^{H'} + \gamma\,(\lambda_\partial^{H'})^2.$$

Exchange symmetry (B0 — no side privileged) forces $\alpha = \gamma$:

$$A_\partial^{HH'} = \alpha\,[\,(\lambda_\partial^H)^2 + (\lambda_\partial^{H'})^2\,] + 2\beta\,\lambda_\partial^H\,\lambda_\partial^{H'}.$$

Positive semidefiniteness (B3) requires $\alpha \geq 0$ and $|\beta| \leq \alpha$. The competition primitive (B4) requires $A_\partial^{HH'}$ to be a perfect square of a linear form, $A_\partial^{HH'} = (p\,\lambda_\partial^H + q\,\lambda_\partial^{H'})^2$; this forces the **rank-1 condition** $\beta^2 = \alpha\gamma$ (the discriminant of the form vanishes), and nothing more by itself. It is the combination of B4 ($\beta^2 = \alpha\gamma$) *with* B0 ($\alpha = \gamma$) that gives $|\beta| = \alpha$: from $\beta^2 = \alpha\gamma = \alpha^2$ we get $|\beta| = \alpha$. We are explicit about the division of labour — B4 collapses the form to rank one, B0 equalizes the diagonal, and only together do they fix $|\beta| = \alpha$. With $\alpha > 0$ (a non-degenerate cost; $\alpha = 0$ is the trivial no-cost term, excluded by B3's "zero only at completion") this leaves exactly two terms:

$\beta = +\alpha$: $A_\partial^{HH'} = \alpha\,(\lambda_\partial^H + \lambda_\partial^{H'})^2$, zero set $\lambda_\partial^{H'} = -\lambda_\partial^H$ (anti-aligning). **$\beta = -\alpha$:** $A_\partial^{HH'} = \alpha\,(\lambda_\partial^H - \lambda_\partial^{H'})^2$, zero set $\lambda_\partial^{H'} = +\lambda_\partial^H$ (co-aligning).

Proposition 4.1 (the local-cost requirements determine the term up to one sign) [Proven]. Under B0–B4, $A_\partial^{HH'}$ is, up to positive normalization, one of the two perfect squares above. B0–B4 do not select between them: both are local, exchange-symmetric, PSD, and built from the squared-linear primitive, and they are exchanged by the relabelling $\lambda_\partial^{H'} \to -\lambda_\partial^{H'}$ of one hub's boundary frame.

Proof. B0 forces $\alpha = \gamma$; B3 gives $\alpha \geq 0$, $|\beta| \leq \alpha$; B4 forces $\beta^2 = \alpha\gamma = \alpha^2$, so $|\beta| = \alpha$. The two solutions $\beta = \pm\alpha$ are the two stated squares. Each satisfies B0–B4. They differ only by the sign of the relative orientation of the two boundary frames, which B0–B4 do not constrain. Without B0, B4 alone gives only $\beta^2 = \alpha\gamma$ — a perfect square $(p\lambda_\partial^H + q\lambda_\partial^{H'})^2$ with $\alpha = p^2$, $\beta = pq$, $\gamma =$

q^2 and $|p| \neq |q|$ permitted — whose zero set $\lambda_\partial^{H'} = -(p/q)\lambda_\partial^H$ has unequal amplitudes; equal amplitude is therefore the content of B0, not of B4. ■

This is the crux the prior framing obscured. AA — anti-alignment — is **not** a consequence of the local cost requirements; it is the choice $\beta = +\alpha$, and B0–B4 leave it open. Asserting anti-alignment at this stage would beg exactly the question the paper exists to answer.

4.2 The intra-hub sign is fixed by orientation

The same two-element ambiguity exists intra-hub and is resolved there by a fact we can name. The intra-hub term is $(\lambda_i + \lambda_{i+1})^2$ — the *sum*-square, anti-aligning — rather than $(\lambda_i - \lambda_{i+1})^2$. The reason is orientation: spokes i and i+1 bound the triangular cell (hub, outer i, outer i+1), and the oriented boundary of that cell traverses the two spokes with *opposite* induced orientation. The competition primitive compares induced-oriented boundary data, so two spokes entering a shared cell with opposite induced orientation enter the primitive with opposite sign of orientation and the same sign of amplitude — giving the sum, $\lambda_i + \lambda_{i+1}$, whose square is the term. The sum-square is the orientation-coherent reading of the primitive on an oriented cell.

Proposition 4.2 (the intra-hub anti-aligning sign is the oriented reading) [Proven]. On an oriented W_6, the competition primitive applied to the two spokes of each cell is $(\lambda_i + \lambda_{i+1})^2$, because the cell's oriented boundary induces opposite orientations on its two bounding spokes. The co-aligning alternative $(\lambda_i - \lambda_{i+1})^2$ is the reading on a cell with reversed relative orientation of its two spokes, which an oriented W_6 does not present. ■

4.3 The cross-boundary sign is the gluing orientation: B5

The cross-boundary term faces the identical choice, and it is resolved by the identical kind of datum, now at the level of the gluing map rather than a single cell.

Requirement B5 (oriented gluing of the substrate). The substrate complex is oriented: adjacent cells induce opposite orientations on their shared face. Equivalently, the gluing map identifying the boundary spoke of H with that of H′ is orientation-reversing on the shared cell, exactly as the two spokes within a single cell carry opposite induced orientation. B5 is a property of how wheels are *assembled* into the substrate, not of the wheel itself: K = 7 gives the local wheel W_6 and hence each hub's local orientability, but oriented gluing is the global-coherence statement that local orientations agree across every shared face. The two are genuinely distinct, and §4.4 makes the gap precise.

Theorem 4.4 (B0–B5 uniquely fix the anti-aligning term) [Proven]. Under B0–B4 together with B5, the shared-boundary completion term is, up to positive normalization,

$$A_\partial^{HH'} = (\lambda_\partial^H + \lambda_\partial^{H'})^2, \text{ with zero set } \lambda_\partial^{H'} = -\lambda_\partial^H.$$

Proof. By Proposition 4.1, B0–B4 leave $\beta = \pm\alpha$. B5 fixes the relative induced orientation of the two boundary frames to be opposite — the same configuration that, by Proposition 4.2, yields the sum-square intra-hub. The competition primitive read on the oriented shared cell therefore compares λ_∂^H and $\lambda\_\partial^{H'}$ with the same sign of amplitude and opposite induced orientation, giving the linear form $\lambda\_\partial^H + \lambda\_\partial^{H'}$ and hence $\beta = +\alpha$. The co-aligning term $\beta = -\alpha$ is the reading under an orientation-reversed gluing, excluded by B5. ■

Remark 4.5 (B5 is a global-gluing premise, derivable from K = 7 only at the local level, not at the global) [Methodological]. B5 is neither manufactured nor a cost requirement: it is the orientability of the substrate *gluing*, a definite structural fact about how the K = 7 complex is assembled. We are careful about which layer it sits at. K = 7 alone proves **local** orientability — each wheel W_6 is an oriented disk, and Proposition 4.2's intra-hub sum-square is the oriented reading on a single oriented cell; this much is derived outright. B5 is the **global** statement that these local orientations cohere across every shared face, and it does *not* follow from local orientability: a Möbius band is assembled from perfectly orientable strips yet admits no global orientation, and the same gap is logically open here. So B5 is the one additional, standard, structural premise the verdict carries — not derivable from K = 7, but the ordinary condition any oriented 2-complex satisfies. What B5 is *not* is a property of the cost functional, which is why it cannot be folded into B4, and why the prior framing's attempt to derive anti-alignment from a competition principle alone could not have succeeded. The honest content: the local cost gives a two-element family; oriented gluing picks the element; K = 7 supplies the local orientability but not the global coherence. We carry both branches forward.

4.4 What K = 7 gives, and what oriented gluing adds

We record the local/global decomposition explicitly, since it is what fixes the standing of the verdict.

Proposition 4.6 (local orientability is derived; global coherence is a premise) [Proven for the local half]. From K = 7:

each hub is the oriented wheel W_6, so the substrate is **locally orientable** — every hub carries a consistent local orientation and the intra-hub anti-aligning mode (Proposition 4.2). [Proven]

Global oriented gluing (B5) — that the local orientations agree across every shared face, equivalently that no orientation-reversing loop exists — is not entailed by local orientability and is the additional premise. [Premise, not derivable from K = 7]

Proof of the local half. The wheel W_6 is a disk triangulated by six cells about a hub; it is orientable, and an orientation is fixed by the cyclic ring order, giving Proposition 4.2's sum-square. The global half is the orientability of the assembled complex, which by §6 is the triviality of the orientation class [s] and is not a consequence of each piece being orientable (the Möbius obstruction). ■

---

5. PBM, AC, AA: Three Distinct Premises (B2, B0, B5)

We resolve the three prior hypotheses against the derivation, distinguishing which follow from the local cost alone and which require B5.

Theorem 5.1 (PBM follows from B2) [Proven]. The boundary term compares the same physical boundary mode from each side. Indeed B2 (label-independence) requires $A_\partial^{HH'}$ to depend on the boundary data only through the gauge-invariant physical mode; the derived term, a function of λ_∂^H and $\lambda_\partial^{H'}$ as the single shared mode carried by each hub, does exactly this. PBM is a consequence of B2 and does **not** require B5.

Proof. By B2 the term is invariant under each hub's origin relabelling, so it factors through the gauge-invariant boundary mode; that is the statement of PBM. The reduction of §4.1 already takes the two arguments to be the one physical mode as represented from each side, which is the content B2 forces. ■

Theorem 5.2 (AC follows from exchange symmetry B0) [Proven]. Equal amplitude across the boundary, $|\lambda_\partial^{H'}| = |\lambda_\partial^H|$ on the zero set, holds iff the boundary cost is exchange-symmetric (B0, $\alpha = \gamma$). Under B0 the zero set of either branch is $\lambda_\partial^{H'} = \mp\lambda_\partial^H$, whence equal magnitude; AC holds.

Proof. B4 alone gives a rank-1 form $(p\,\lambda_\partial^H + q\,\lambda_\partial^{H'})^2$ with zero set $\lambda_\partial^{H'} = -(p/q)\,\lambda_\partial^H$, i.e. $|\lambda_\partial^{H'}| = |p/q|\cdot|\lambda_\partial^H|$ — unequal in general. Equal amplitude is exactly $|p| = |q|$, i.e. $p^2 = q^2$, i.e. $\alpha = \gamma$, which is B0. So AC is the content of the exchange symmetry, not of the perfect-square structure: B4 fixes the zero set to a single line through the origin; B0 fixes that line's slope to ∓1. AC is independent of B5 — both the anti-aligning ($\beta = +\alpha$) and co-aligning ($\beta = -\alpha$) branches give equal magnitude once B0 holds, so AC is the one of the three prior hypotheses insensitive to orientation. ■

Theorem 5.3 (AA is the content of B5) [Proven]. Under B0–B5 the zero set is anti-aligning, $\lambda_\partial^{H'} = -\lambda_\partial^H$. AA holds **iff** B5 holds: it is the anti-aligning branch $\beta = +\alpha$, which Theorem 4.4 selects exactly from orientation-coherence. Under the co-aligning branch (B5 failing on the relevant edge) AA fails and the boundary is co-aligning, $\lambda_\partial^{H'} = +\lambda_\partial^H$.

Proof. By Theorem 4.4, $\beta = +\alpha \Leftrightarrow$ B5, and $\beta = +\alpha$ is precisely AA. The co-aligning branch $\beta = -\alpha$ is the negation. ■

Summary. The three prior hypotheses trace to three distinct premises, none redundant: **PBM ← B2** (label-independence), **AC ← B0** (exchange symmetry — the shared face privileges neither hub), **AA ← B5** (oriented gluing). PBM and AC are theorems of the local cost requirements, with no orientation input; AA is exactly the orientation-coherence datum B5 — neither a free hypothesis nor a consequence of the cost, but a named structural property of the complex. B4 (the perfect-square primitive) does *not* on its own deliver AC: it collapses the form to rank one

($\beta^2 = \alpha\gamma$), and equal amplitude is the separate content of B0. The prior paper's three-hypothesis conjunction therefore collapses to a single binary structural question: is the gluing orientation-coherent?

---

6. The Sign-Cocycle Is the Orientation Class

We now identify what the prior paper's signed matching cochain *is*. The identification is definitional — it runs through Theorem 6.1, Theorem 5.3, and the construction of B5 — rather than an independently discovered coincidence; its value is organizational, converting "evaluate [s]" into a statement about a standard topological invariant whose triviality is decidable structurally. Nothing is smuggled: the topology is standard, and the identification of s with the orientation cochain holds by the construction of B5.

6.1 Identification

Each hub carries $\lambda_i^H = (-1)^i a^H$; across a shared boundary the derived rule (Theorem 4.4, under B5) gives $\lambda_\partial^{H'} = -\lambda_\partial^H$, i.e. the alternating modes match anti-aligningly and the relative amplitude sign is

$$a^{H'} = s_{HH'} a^H, \quad s_{HH'} \in \{\pm 1\}.$$

In the physical gauge, $s_{HH'} = \varepsilon^{H}{}^{phys} \varepsilon^{phys}_{H'} \cdot t_{HH'}$ with the ε-part a coboundary (prior paper, proven). The residue $t_{HH'}$ is +1 on an orientation-coherent edge (B5 holding there) and −1 on an orientation-reversed edge (B5 failing there, the co-aligning branch).

Theorem 6.1 (the residue is the local orientation datum) [Proven]. In the physical gauge, $t_{HH'} = +1$ iff the gluing across (H, H′) is orientation-coherent (B5 on that edge), and $t_{HH'} = -1$ iff it is orientation-reversing.

Proof. By Theorem 5.3, the anti-aligning branch (t = +1 in the physical gauge, by the prior paper's Lemma) is exactly B5 on that edge; the co-aligning branch (t = −1) is its failure. ■

Theorem 6.2 (the matching class is the orientation class of the substrate complex) [Proven for the topology; identification by construction of B5]. The cohomology class $[s] = [t] \in H^1(\Gamma\_hub; \mathbb{Z}_2)$ is the first obstruction to a coherent orientation of the substrate 2-complex tiled by the wheels — the w_1-type orientation class of the gluing. Its triviality is orientability:

[s] = 0 ⟺ the substrate complex admits a coherent orientation ⟺ B5 holds globally.

Proof. t is the $\mathbb{Z}_2$ cochain assigning to each shared face whether the gluing there preserves (+1) or reverses (−1) induced orientation (Theorem 6.1). A coherent global orientation of the assembled 2-complex exists iff this cochain is a coboundary — iff there is a per-hub orientation sign making every shared face coherent — which is [t] = 0. This is the standard characterization of the first orientation obstruction (w_1) of a 2-cell complex via its $\mathbb{Z}_2$ orientation cochain, and is the [Proven] content. That the matching cochain s coincides with this orientation cochain t is the identification [s] = [t]: it holds *by the construction of B5* — Theorem 6.1 defines t as the orientation sign and Theorem 5.3 ties the matching sign to it — not as an independent discovery. The theorem is therefore the standard topology applied to a cochain the construction has arranged to be the orientation cochain; the work is in the arrangement (§§4–5), not in the topology. ■

Remark 6.2a (it is the substrate complex, not the bare adjacency graph, whose orientability is meant) [Methodological]. The class lives on Γ_hub, but the datum t_{HH′} it carries is the induced-orientation sign of the *geometric* gluing across each shared face of the substrate 2-complex — not a property of the abstract adjacency graph. The two must not be conflated: a graph-theoretic property of Γ_hub (bipartite or not) does not track t, which is fixed by how the wheels are glued, not by which wheels are adjacent. When B5 is evaluated against the construction (§10A), the object whose orientability is in question is the substrate 2-complex the wheels tile, carrying the induced-orientation data on shared faces; Γ_hub is only the index over which the resulting cochain is recorded.

6.2 The two branches

Theorem 6.3 (orientation-coherent branch: hub-adjacent pass) [Proven]. If B5 holds globally — Γ_hub orientation-coherent — then [s] = [t] = 0, every cycle product $\prod\gamma\, s$\{HH′} = +1, the signed graph is balanced, and the hub-adjacent stratum contains no N_2 witness.

Proof. By Theorem 6.2, B5 global ⟺ [s] = 0. By the prior paper's Hub-Adjacent Reachability Criterion, [s] = 0 is exactly the hub-adjacent pass. ■

Theorem 6.4 (frustrated branch: explicit N_2 witness) [Proven]. If B5 fails along a loop — there is a cycle γ ⊆ Γ_hub with an odd number of orientation-reversing edges, so $\prod\gamma\, s$\{HH′} = −1 — then the alternating mode cannot be assigned consistently around γ, and the boundary-matched alternating germ on a neighbourhood of γ is an unrealizable admissible germ:

Ĝ_γ ∈ 𝔘 ∩ $\mathcal{P}_2$^{hub-adjacent}, 𝔘 ≠ ∅, Gate 2 FAIL, verdict N_2.

Proof. Transporting a^H around γ multiplies by ∏_γ s = −1, returning a → −a: no consistent amplitude assignment exists on γ (the prior paper's frustration corollary). Every hub on γ individually satisfies L (each carries the admissible alternating ray), and each edge individually satisfies the boundary rule, so the germ on the closed star of γ lies in {L}; but no transport realizes it, since realization requires the consistent assignment γ forbids. Hence Ĝ_γ ∈ {L} \ Im(ρ) = 𝔘, witnessed at hub-adjacent complexity. ■

This is the closure on the failing branch: not merely "the gate might fail" but the explicit witness — the alternating germ on an orientation-frustrated loop — named and located at minimal complexity. The verdict is genuinely two-valued, and we have both values.

---

7. The Four Cell-Completion Relations

We record the four cell-completion relations of Open Problem 6.7 at the $K = 7$ nearest-neighbour scale, distinguishing those that carry a failure locus from those satisfied by construction.

7.1 Incidence (I_inc) — definitional, no failure locus. $I\_inc(\hat{G})$ holds iff the germ assigns amplitudes exactly to the spoke coordinates of the wheel, $\lambda \in \mathbb{R}^6$, with no non-spoke coordinate. This is satisfied by construction of a $K = 7$ germ; it is the statement that the germ *is* a W_6 datum, not an extra condition that could fail on $\{L\}$. It carries no on-sector residue.

7.2 Hub (I_hub) — definitional, no failure locus. $I\_hub(\hat{G})$ holds iff the six spoke amplitudes close cyclically mod 6, i.e. the local neighbourhood is W_6 with $i \sim i+1 \bmod 6$. Again satisfied by construction of the germ; no on-sector residue.

We flag 7.1–7.2 as definitional explicitly: they are not evaluated conditions that happen to pass, but properties built into "$\hat{G}$ is a $K = 7$ germ." The substantive content is carried entirely by the two sectors below.

7.3 Conservation (I_circ) — content-bearing. $I\_circ(\hat{G})$ holds iff $A\_circ = (\Sigma_i \lambda_i)^2 = 0$, i.e. $\Sigma_i \lambda_i = 0$. This is a genuine condition; it is satisfied automatically on the alternating ray (§2), so it carries no residue *on $\{L\}$*, but it is not definitional — a non-alternating germ would violate it.

7.4 Competition (I_comp) — content-bearing, intra-hub and cross-boundary. $I\_comp(\hat{G})$ holds iff both

intra-hub: $\lambda_i + \lambda_{i+1} = 0$ for all i mod 6, cross-boundary: $\lambda\_\partial\text{^}H + \lambda\_\partial\text{^}\{H'\} = 0$ for every adjacent pair (under B5; the cross-boundary clause is $\lambda\_\partial\text{^}H - \lambda\_\partial\text{^}\{H'\} = 0$ on any orientation-reversed edge).

The cross-boundary clause is the locus where B5 enters the completion relations: under orientation-coherence it is the anti-aligning condition satisfied on the matched ray; under frustration it is the co-aligning condition, which the alternating ray violates around a frustrated loop, producing the witness of Theorem 6.4.

---

8. Evaluation of the Unrealizable Residue

The failing-sector set is $F\_\mathfrak{U} = \{ s : \{L\} \not\subset \{I\_s\} \}$. We evaluate the two content-bearing sectors; the two definitional sectors contribute no locus (§7.1–7.2).

8.1 Conservation. On $\{L\}$ (the alternating ray, single-hub) $\Sigma_i \lambda_i = 0$, so $\{L\} \subseteq \{I\_circ\}$ and $\mathfrak{U}\_circ = \emptyset$. [Proven]

8.2 Competition, intra-hub. On the alternating ray $\lambda_i + \lambda_{i+1} = (-1)^i a + (-1)^{i+1} a = (-1)^i a (1 - 1) = 0$, so the intra-hub clause holds on {L}. [Proven]

8.3 Competition, cross-boundary — the B5-sensitive locus. Under B5 (Theorem 4.4) the cross-boundary clause is $\lambda\_\partial\text{^}H + \lambda\_\partial\text{^}\{H'\} = 0$, satisfied on the matched anti-aligned ray (Theorem 6.3), so {L} ⊆ {I_comp} and 𝔘_comp = ∅. Under a frustrated loop (Theorem 6.4) the alternating ray cannot satisfy the cross-boundary clause consistently around the loop, so {L} ⊄ {I_comp} and 𝔘_comp ≠ ∅. [Proven, branch-dependent]

Theorem 8.5 (residue verdict, branch-dependent) [Proven].

Under B5: F_𝔘 = ∅ and $𝔘 \cap \mathcal{P}_2 = \emptyset$. Under a frustrated orientation loop: F_𝔘 = {𝒞_comp} and $𝔘 \cap \mathcal{P}_2 \supseteq \{\hat{G}\_\gamma\} \neq \emptyset$ (Theorem 6.4).

Proof. The definitional sectors carry no locus; I_circ and the intra-hub competition clause are satisfied on {L} unconditionally (§8.1–8.2). The cross-boundary competition clause is satisfied on {L} iff B5 (§8.3). So F_𝔘 = ∅ iff B5, with the sole possible failing sector being 𝒞_comp via its cross-boundary clause, witnessed by Ĝ_γ. ■

The residue is empty precisely on the orientation-coherent branch, and its only possible non-emptiness is the competition sector's cross-boundary clause, carrying the orientation witness. This is the sharp form of the prior paper's "the live risk is combinatorial": the risk is real, it lives in exactly one sector's cross-boundary clause, and it is governed by exactly one structural property.

---

9. Boundedness, Split: Constraint-Locality and Realizability-Gluing

Conjecture 8.3 asked whether the registered search exhausts the Gate-2 obstruction space. We prove the half that is genuinely available unconditionally and identify the other half honestly as the orientation condition — declining to conflate constraint-locality with realizability-gluing.

9.1 The constraint-locality half (unconditional)

Theorem 9.1 ($\mathcal{P}_2$ exhausts the constraint content) [Proven]. For the nearest-neighbour $K = 7$ construction — admissibility cost composed of incidence, hub, conservation, and competition terms, none coupling more than two hubs and the inter-hub coupling factoring through pairwise shared boundaries — membership in {L} is decidable on single-hub and hub-adjacent neighbourhoods. No primitive constraint couples three or more hubs except through pairwise shared-boundary completion. Hence every way the local rules can *fail to be satisfied* appears within $\mathcal{P}_2$.

Proof. By inspection of the functional: incidence, hub, conservation are single-hub; competition is single-hub plus pairwise shared-boundary. A germ violates L iff some term is nonzero, and every term is evaluable on a single hub or a single adjacent pair. So L-violation is detectable at single-hub or hub-adjacent complexity, and $\mathcal{P}_2$ — which enumerates exactly these — captures the full constraint content. ■

This is a real and unconditional result: it bounds *where the constraints live*. It does not yet bound where *realizability* fails, and we are explicit that it does not.

9.2 The realizability-gluing half (conditional on B5)

Remark 9.2 (why constraint-locality does not give realizability-globalization) [Methodological]. Gate 2 concerns realizability, not L-satisfaction. Theorem 9.1 shows L is pairwise-checkable; it does not show that pairwise-local realizers assemble into a realizer of a larger germ. That assembly is a compatibility-on-overlaps condition — a gluing statement — and gluing is not delivered by the functional's locality. Asserting it would import a Gate-3 step into the Gate-2 bound. We instead identify it precisely.

Theorem 9.3 (realizability globalizes iff the orientation class is trivial) [Conditional on B5]. For the nearest-neighbour K = 7 construction, the pairwise-local realizers of an L-satisfying germ assemble into a global local realizer iff the matching class is trivial, [s] = 0 — equivalently iff B5 holds (Theorem 6.2). Under B5, $\mathfrak{U} \cap \mathcal{P}_2 = \emptyset \Longrightarrow \mathfrak{U} = \emptyset$: the registered search globalizes and the constraint-locality bound of Theorem 9.1 becomes a realizability bound. Under a frustrated loop, globalization fails at γ and the witness Ĝ_γ (Theorem 6.4) is the obstruction.

Proof. A larger germ's realizer is built by assigning each hub its alternating amplitude consistently with every shared-boundary matching. The obstruction to a consistent global assignment is exactly the holonomy of s around cycles — the class [s] (prior paper's criterion). By Theorem 6.2 this is the orientation class, trivial iff B5. Under [s] = 0 every cycle closes, the pairwise realizers agree on overlaps, and the assignment extends globally: any L-germ is realized by patching its pairwise-realized pieces, so $\mathcal{P}_2$-emptiness of $\mathfrak{U}$ forces global emptiness. Under [s] ≠ 0 the holonomy around the frustrated loop obstructs assembly, exhibited as Ĝ_γ. ■

Remark 9.4 (this is the honest form of Conjecture 8.3, and the hinge into Gate 3) [Methodological]. Theorem 9.1 (constraint-locality) is unconditional and is the part of Conjecture 8.3 that holds outright. Theorem 9.3 (realizability-gluing) is conditional on B5 and is the part that is *not* free — it is a genuine gluing statement, and it is the same orientation condition that governs the residue. We do not claim "Conjecture 8.3 holds" unqualified; we claim its constraint half outright and its realizability half exactly when the orientation class vanishes. That the gluing half is a gluing statement is the natural pointer to Gate 3, which is global gluing over the full complex: Gate 2's realizability-globalization is the orientation (w_1) layer of that question, and Gate 3 will take up the higher layers.

---

10. The Gate-2 Verdict

Assembling the branches:

Theorem 10.1 (Gate-2 verdict for the $K = 7$ nearest-neighbour construction) [Proven, branch-stated]. Under C1 (PSD reading of A_comp, confirmed), C2 (σ-family admissibility in $\mathcal{T}$_L, finite check), and C3 (non-vacuity under $A\_inc \wedge A\_hub$, finite check):

$\mathfrak{U} = \emptyset \Leftrightarrow$ B5 (the substrate gluing is oriented, equivalently $[s] = 0 \in H^1(\Gamma\_hub; \mathbb{Z}_2)$).

Equivalently, with $\mathrm{Im}(\rho) \subseteq \{L\}$ from Gate 1:

Under oriented gluing (B5): $\mathrm{Im}(\rho) = \{L\}$, **Gate 2 PASS** for the $K = 7$ nearest-neighbour transport construction. **Under a frustrated orientation loop:** $\mathfrak{U} \ni \hat{G}\_\gamma \neq \emptyset$, **Gate 2 FAIL**, verdict N_2, witnessed by the alternating germ on the frustrated loop (Theorem 6.4).

Proof. By Theorem 8.5, $F\_\mathfrak{U} = \emptyset \Leftrightarrow$ B5, with the only candidate failing sector the cross-boundary competition clause. By Theorem 9.1 the constraint content is exhausted by $\mathcal{P}_2$; by Theorem 9.3, under B5 the realizability content globalizes, so $\mathfrak{U} \cap \mathcal{P}_2 = \emptyset \Rightarrow \mathfrak{U} = \emptyset$, giving $\{L\} \subseteq \mathrm{Im}(\rho)$, hence with Gate-1 soundness $\mathrm{Im}(\rho) = \{L\}$ and Gate 2 PASS. Under a frustrated loop, $\hat{G}\_\gamma \in \mathfrak{U}$ by Theorem 6.4, so Gate 2 FAIL. ■

The closure. This is a verdict, not a narrowing. The prior arc had reduced Gate 2 to three shared-boundary hypotheses and an open boundedness question. We have shown: PBM and AC are theorems of the local cost; AA, the boundedness-gluing condition, and the residue's only possible non-emptiness are all the *same single structural premise* — oriented gluing of the substrate, $[s] = 0$ — which is the orientation class of the substrate complex. Gate 2 passes iff that class is trivial, and we exhibit the witness when it is not. The verdict is a clean function of one named binary premise, with both values realized explicitly. What remains is to say what $K = 7$ does and does not settle about that premise.

10A. Evaluation of B5 against the K = 7 Construction

We evaluate B5 against the existing construction and return a dependency rather than a manufactured value, since manufacturing it would be the one move the arc has refused throughout.

The local half is derived. $K = 7$ furnishes the wheel W_6 — one hub, six boundary states, six spokes, the cyclic ring — and thereby proves each hub *locally* orientable, with the intra-hub alternating mode and its sum-square competition reading (Propositions 4.2, 4.6). This is settled by $K = 7$ with no further input.

The global half is a premise, not derivable from $K = 7$. B5 asks how one wheel is glued to another across a shared boundary, and whether those gluings cohere over the whole substrate. $K = 7$ describes the local wheel; it does not, by itself, fix the global gluing orientation. Local orientability does not entail global orientability — the Möbius obstruction is precisely a coherent local structure with no global orientation — so a frustrated orientation loop ($[s] \neq 0$) remains logically possible on the strength of K = 7 alone. Oriented gluing is therefore an additional structural premise, not a theorem of K = 7.

The strongest honest statement. Combining the two halves:

If the K = 7 transport substrate is built as an oriented 2-cell complex — adjacent cells inducing opposite orientations on their shared faces — then B5 holds, [s] = 0, PBM/AC/AA hold, and Gate 2 PASSES. K = 7 supplies the local orientability; oriented gluing supplies the global coherence; the latter is not derivable from the former.

So the verdict is not "B5 follows from K = 7" — it does not — but "B5 follows from the *oriented* K = 7 substrate construction." This is a genuine closure of the conditional kind: the entire remaining Gate-2 content is the single, standard, named premise that the substrate is assembled with oriented gluing. Two consequences worth stating. First, the ordinary case — a coherent wheel-tiled 2-complex built with the standard opposite-induced-orientation convention — satisfies B5 by construction, and there the verdict is unqualified PASS; the orientation convention may already be present in the substrate-construction layer, in which case §10A is discharged by citation. Second, if the substrate construction does *not* fix an orientation convention, then it is underspecified at exactly the gluing, and Gate 2 forces the convention: oriented gluing is the orientation datum the construction must carry if the gate is to pass. Either way the open content is one binary structural premise about the substrate, evaluated at the geometric 2-complex (Remark 6.2a), not the abstract adjacency graph.

10A′. Reducing B5 to an Interface-Orientation Identification via Loop-Consistency

§10A leaves B5 as a structural premise on the gluing, equivalently the condition $[s]_{vac} = 0$. That equivalence is unconditional; what the interface loop-consistency axiom adds is not a reduction of the criterion but a *shortcut* for verifying it — and even that shortcut rests on one identification the interface-dynamics paper does not yet establish. We state both at exactly the precision available: the pass criterion $[s]_{vac} = 0$ stands alone; the loop-consistency shortcut is conditional on a deposited identification.

Definition 10A.0 (vacuum sector and restricted orientation class). By the *vacuum sector* $\Gamma_{vac} \subseteq \Gamma_{hub}$ we mean the ground-state sub-complex of the K = 7 substrate: the region carrying the admissible alternating transport pattern with no localized defect insertion or defect-encircling puncture. The orientation cocycle s on Γ_{hub} restricts to a cocycle $s_{vac} := s|_{\Gamma_{vac}}$, and its cohomology class

$$[s]_{vac} := [s|_{\Gamma_{vac}}] \in H^1(\Gamma_{vac}; \mathbb{Z}_2)$$

is the *vacuum orientation class*. The Gate-2 vacuum pass condition is the triviality of this restricted class, $[s]_{vac} = 0$. Defect-bearing regions, if any, are not part of the vacuum-sector verdict; they belong to the separate configuration-space question of §10B. Throughout §10A′, "[s]_vac = 0" is this restricted-class condition, and the restriction map $H^1(\Gamma\_hub; \mathbb{Z}_2) \to H^1(\Gamma\_vac; \mathbb{Z}_2)$ is what carries the whole-complex class of §6 to the vacuum-sector condition used here.

What the interface axiom supplies. The interface dynamics requires infinitesimally contractible loops to close trivially, $U_{\delta\gamma} = \mathbb{1}$, so that local update structure coarse-grains coherently (interface loop-consistency). In the closure-phase language, the holonomy around a loop is the accumulated committed closure phase, an operational observable — the committed phase difference recorded after a closed circuit — and loop-consistency is the statement that this observable is trivial on every contractible loop.

What that does and does not give [Methodological]. Triviality of the holonomy on *contractible* loops is precisely the condition for the orientation holonomy to descend to a well-defined class $[s] \in H^1(\Gamma\_hub; \mathbb{Z}_2)$ — it is what makes [s] depend only on the homotopy class of a loop. It is **not** the statement [s] = 0, which is triviality on *all* loops including the non-contractible. The gap between the two is exactly the non-contractible loops, and that gap is not an oversight to be closed but the precise location of the vacuum/defect split. We do not overstate: loop-consistency forbids non-trivial holonomy on contractible loops; it says nothing about non-contractible ones. The reversibility route failed because a sign flip is invertible (involution); the loop-consistency route does not repeat that error, because it constrains the *committed-record* holonomy — a single-valuedness condition on an operational observable — not mere invertibility. But it constrains it only where the axiom reaches: contractible loops.

Lemma 10A.0′ (the orientation sign appears in the interface holonomy on an elementary plaquette) [Conditional on the interface-orientation identification — a deposited lemma]. Let δγ be an elementary contractible loop — the boundary of a single shared-face plaquette between adjacent hubs. *Provided* the interface update $U_{\delta\gamma}$ acts on the Gate-2 boundary mode through its orientation component, $U_{\delta\gamma}$ carries, as its action on the boundary amplitude, the same induced-orientation sign that the competition primitive reads on that plaquette (Proposition 4.2 intra-hub, Theorem 4.4 cross-boundary): an orientation-*preserving* elementary loop returns $a \mapsto +a$ (the sum-square / anti-aligning reading), an orientation-*reversing* elementary loop returns $a \mapsto -a$ (the difference-square / co-aligning reading). Under that proviso an orientation-reversing elementary loop has $U_{\delta\gamma} \neq \mathbb{1}$ — its action on the boundary mode is multiplication by −1, not the identity.

Standing — what is established and what is deposited. The interface-dynamics construction provides loop-consistency ($U_{\delta\gamma} = \mathbb{1}$ on contractible loops) and a two-component cell update U_{ij} acting on a state (σ, ω), with σ linked to orientation perturbation and ω to closure parity. What it does **not** currently provide is a result identifying the Gate-2 matching cochain s with the σ (orientation) component of the interface update — i.e. a statement that $U_{\delta\gamma}$ acts on the boundary mode by the matching-sign product $\prod_{\delta\gamma} s$ rather than through a closure phase whose orientation part is a mere quotient, or through ω alone. That identification is the exact link Lemma 10A.0′ needs, and it is **not in hand**: we therefore deposit it as a named condition rather than assert it.

Interface-orientation identification (deposited). The Gate-2 orientation cocycle s coincides with the orientation (σ) component of the interface update U: on the boundary mode, $U_{\delta\gamma}$ acts as multiplication by $\prod_{\delta\gamma} s$.

If this identification holds — if the interface σ-component is the Gate-2 orientation sign — then $U_{\delta\gamma} = \mathbb{1}$ forces $\prod_{\delta\gamma} s = +1$ *a fortiori, and Lemma 10A.0′ follows; on a contractible loop the orientation holonomy is then trivial. We do not claim the identification; we state precisely what it is and what would establish it (an extension of the interface-dynamics paper, or a result pinning U's action on the σ-component to s). Absent it, the implication "*$U_{\delta\gamma} = \mathbb{1} \Longrightarrow$ trivial orientation holonomy" is conditional, and §10A′ is a reduction *to* this identification, not a derivation through it. ■

With the deposited identification in force, $U_{\delta\gamma} = \mathbb{1}$ on every contractible loop forces the orientation sign to +1 there, which is trivial orientation holonomy on contractible loops. Without it, the contractible-loop triviality of the *orientation* class is not yet delivered by loop-consistency.

Theorem 10A.1 (vacuum pass criterion, unconditional; and the loop-consistency shortcut, conditional). With Γ_{vac} and $[s]_{vac}$ as in Definition 10A.0:

(a) Pass criterion [Proven, under C1–C3]. By Theorem 6.2 applied to the restricted sub-complex Γ_{vac}, B5 holds on the vacuum sector iff $[s]_{vac} = 0$, and Gate 2 passes on the vacuum iff this holds:

$[s]_{vac} = 0 \Leftrightarrow$ B5 holds on $\Gamma_{vac} \Leftrightarrow$ Gate 2 PASS on the vacuum, with PBM/AC/AA, $\mathfrak{U} = \emptyset$ following by §§5–9.

This is Theorem 6.2/10.1 applied to Γ_{vac} via the restriction map $H^1(\Gamma_{hub}; \mathbb{Z}_2) \to H^1(\Gamma_{vac}; \mathbb{Z}_2)$ of Definition 10A.0; it does not use the interface dynamics. The pass criterion is therefore unconditional on the interface identification.

(b) Loop-consistency shortcut [Conditional on the interface-orientation identification (Lemma 10A.0′)]. Granting the deposited identification, loop-consistency ($U_{\delta\gamma} = \mathbb{1}$ on contractible loops) forces trivial orientation holonomy on every *contractible* vacuum loop, so verifying $[s]_{vac} = 0$ reduces to checking the *non-contractible* vacuum loops only. This shortens the verification; it does not change the criterion of (a).

Simple connectivity is a sufficient route, not the criterion. If every vacuum loop is contractible then — granting (b) — $[s]_{vac} = 0$ follows with no further check; but a non-simply-connected vacuum (e.g. an orientable torus) passes equally, provided its non-contractible loops carry trivial orientation holonomy. The criterion is the triviality of the class (a), not the topology that happens to guarantee it.

Proof. (a) By Theorem 6.2 applied to Γ_{vac}, $[s]_{vac} = 0 \Leftrightarrow$ *B5 on the vacuum sector; by Theorem 10.1 (under C1–C3) this is* $\mathfrak{U} = \emptyset$ *and Gate 2 PASS there. The restriction to* Γ_{vac} *is by Definition 10A.0; no interface input is used. (b) Granting Lemma 10A.0′,* $U_{\delta\gamma} = \mathbb{1}$ on each contractible elementary loop forces its orientation sign to +1; since the elementary plaquettes generate the contractible part of the loop space, composing over a generating set extends this to every contractible loop, which then carries trivial orientation holonomy — leaving only the non-contractible loops of Γ_vac to determine [s]_vac. ■

Remark 10A.2 (the Möbius check, kept honest) [Methodological]. Theorem 10A.1 must not be misread as "loop-consistency ⟹ [s]_vac = 0" unqualified — that statement is false, and the Möbius band is the standing counterexample: it satisfies loop-consistency (every contractible loop closes, local update coarse-grains fine) yet is non-orientable, its non-contractible core loop carrying holonomy −1. Loop-consistency reaches only the contractible loops; what survives without further input is the descent of [s] to a cohomology class, not its vanishing. The non-contractible vacuum loops, if any, must be checked separately — which is exactly why the condition is [s]_vac = 0 and not the topology.

Remark 10A.3 (what is unconditional, and what the identification actually buys) [Methodological]. Two statements must be kept separate, because conflating them understates the verdict. **First, unconditionally (under C1–C3): B5 ⟺ [s]_vac = 0.** This is Theorem 6.2/10.1 restricted to the vacuum sector — the realized substrate the wheels tile — and the interface identification plays no role in it. The pass criterion is therefore not conditional on anything from the interface dynamics: Gate 2 PASS iff the vacuum orientation class is trivial, full stop. **Second, separately: the interface-orientation identification (Lemma 10A.0′) buys a verification shortcut, not a reduction of the criterion.** If the interface update's σ-component is the Gate-2 matching cochain s, then loop-consistency (contractible loops close, U_{δγ} = 𝟙) discharges the *contractible-loop* content of [s]_vac = 0 automatically, shrinking the check from "all vacuum loops" to "non-contractible vacuum loops only." That is a statement about how much remains to verify, conditional on the identification — it does **not** enter the equivalence B5 ⟺ [s]_vac = 0.

So B5 reduces to the single condition [s]_vac = 0, unconditionally. The identification is a separate convenience premise: granting it, the contractible part of [s]_vac = 0 is free and only non-contractible vacuum loops need checking; absent it, [s]_vac = 0 must be verified across all vacuum loops directly — but the criterion itself is unchanged either way. We are explicit that simple connectivity is **not** the criterion: it is one sufficient route to [s]_vac = 0 (every loop contractible, so — granting the identification — loop-consistency forces triviality everywhere), but a non-simply-connected vacuum passes equally if its non-contractible loops carry trivial holonomy. This is a strict improvement over the prior framing: B5 was "the gluing is oriented," a bare convention on shared faces; it is now the checkable topological condition [s]_vac = 0 on the vacuum sector, with the interface identification offered separately as the thing that makes most of that check automatic. The honest chains are therefore:

Pass criterion (unconditional, under C1–C3): B5 ⟺ [s]_vac = 0 ⟺ Gate 2 PASS on the vacuum. **Verification shortcut (conditional on the interface-orientation identification):** loop-consistency ⟹ contractible vacuum loops carry trivial orientation holonomy automatically ⟹ only non-contractible vacuum loops remain to check for [s]_vac = 0.

For §10B the relevant consequence is the part (b) shortcut, not the part (a) criterion: *granting* the interface-orientation identification, no non-trivial orientation holonomy can sit on a contractible loop, so any non-trivial configuration holonomy is confined to non-contractible loops. §10B uses exactly this, and inherits its conditionality — the defect-localization statement is conditional on the same deposited identification, while the pass criterion of part (a) is not.

10B. Hand-Forward: Configurations on the Realized Vacuum

We close with a forward-pointing observation, sharply separated from the verdict above and labelled accordingly. Nothing in this subsection conditions, weakens, or recolours the Gate-2 result; it identifies a distinct object for a distinct paper.

Two different classes, two different objects [Methodological]. The Gate-2 verdict turns on the orientation class of the *complex* — $[s] \in H^1(\Gamma\_hub; \mathbb{Z}_2)$ read as a property of the **gluing**, i.e. whether the substrate is assembled coherently (B5). This is settled, under B5, as [s] = 0: the vacuum gluing is orientation-coherent and Gate 2 passes. A logically separate object is the $\mathbb{Z}_2$ sector label of a **configuration** of the degeneracy datum on that fixed, already-oriented substrate — the holonomy [s]_γ of the alternating mode's per-hub sign-assignment around a loop γ. The first is a property of how the substrate is *built*; the second is a property of a *state* on the built substrate. They are not the same class, and the often-feared tension — "[s] = 0 globally yet [s]_γ ≠ 0 around a loop" — is no contradiction, because the two assertions concern different objects: a coherent gluing ([s] = 0) and a frustrated configuration on it ([s]_γ ≠ 0). This is the ordinary distinction between an orientable crystal and a dislocation field on it: the lattice is orientable; the configuration carries the defect.

The candidate particle mechanism [Conjectural], localization conditional on the interface identification. On a realized, orientation-coherent vacuum (B5; [s] = 0), the configuration space of the alternating mode splits into $\mathbb{Z}_2$ sectors labelled by [s]_γ. The vacuum is the trivial sector. A localized configuration that is vacuum-like (alternating) everywhere except for a sign-frustration that cannot be flattened around one loop sits in a nontrivial sector: a **localized orientation defect** on an orientation-coherent substrate. The *localization* — that such a defect cannot sit on a contractible loop — follows from Theorem 10A.1(b), hence inherits its condition: *granting the interface-orientation identification*, loop-consistency forbids non-trivial holonomy on any contractible loop, so a nontrivial [s]_γ can live only on a non-contractible loop encircling a defect. Under that identification, any orientation defect is necessarily non-contractible-loop content, never a property of a trivial-holonomy ([s]_vac = 0) vacuum — the precise sense in which "space is orientation-coherent, but may contain localized orientation defects" is consistent rather than contradictory. Absent the identification, the localization is conjectural rather than forced. Such a defect would be a candidate for a persistent transport defect — the long-standing VERSF reading of particles as persistent structures in the substrate rather than fundamental point objects — carrying exactly one $\mathbb{Z}_2$ quantum number (self-conjugate, pair-annihilating, one species), the charge inherited from the $\mathbb{Z}_2$ coefficient.

Why this is a separate paper, and what its hinge is [Open]. The Gate-2 verdict is about the vacuum (the ground-state gluing) and does not test whether localized excitations exist on it; the configuration question lives in a layer Gate 2 does not reach, so it neither helps nor harms the verdict here. We flag explicitly that a topological sector label is *not* by itself a persistent particle: the $\mathbb{Z}_2$ class gives the sectors for free, but persistence — that the defect cannot relax to the vacuum sector under the substrate's own refinement dynamics $\mathscr{R}$ — is a *separate* theorem, the discrete analogue of topological protection, and is the load-bearing open problem of any defect programme. It is not established here and is not assumed; "survives refinement" is the conclusion such a paper must earn, not a premise it may take. We therefore defer the defect picture to a dedicated paper — *Localized Orientation Defects of the K = 7 Transport Substrate* — whose first theorem is the existence of nontrivial configuration sectors, whose load-bearing theorem is refinement-stability of those sectors under $\mathscr{R}$, and whose charge analysis and any particle identification are downstream of both, the latter to be labelled as the speculative step it is. The present paper claims none of this; it records only that the obstruction class it identified for the vacuum is also the natural home, *one layer up in configuration space*, for such a defect — and stops there.

---

11. Limitations

The pass criterion is unconditional; the interface identification conditions only a verification shortcut. Gate 2 PASS holds iff B5, and B5 ⟺ [s]_vac = 0 unconditionally under C1–C3 (Theorem 6.2/10.1 on the vacuum, Theorem 10A.1(a)). §10A shows B5 is not derivable from K = 7 alone (K = 7 gives local orientability; oriented gluing is the global-coherence premise), so [s]_vac = 0 is the checkable condition the verdict rests on — a topological property of the vacuum sector, evaluated at the geometric 2-complex per Remark 6.2a. The interface-orientation identification (Lemma 10A.0′) does **not** enter this criterion; it buys only the shortcut of Theorem 10A.1(b) — granting it, loop-consistency discharges the contractible-loop content of [s]_vac = 0 automatically, leaving only non-contractible vacuum loops to check. That identification is not yet established by the interface-dynamics paper and is deposited (next entry). The honest summary: the verdict stands on [s]_vac = 0 alone (unconditional); the identification is a convenience that shrinks the verification, not a premise the pass depends on. B5 remains the *only* remaining content — PBM, AC, AA, the residue, and the realizability-gluing half of boundedness all reduce to it — and the failure witness is exhibited should [s]_vac ≠ 0.

"Nearest-neighbour" is load-bearing. Theorems 9.1 and 9.3 assume the admissibility cost couples at most two hubs, with inter-hub coupling factoring through pairwise shared boundaries. A later extension introducing genuine three-hub or higher-radius primitive coupling would require rerunning the constraint-locality bound with those terms, and could in principle introduce obstructions beyond the orientation class.

The interface-orientation identification (Lemma 10A.0′) is a deposited condition on the verification shortcut, not on the pass criterion, and is not currently established. The pass criterion B5 ⟺ [s]*vac = 0 (Theorem 10A.1(a)) is unconditional under C1–C3 and uses no interface input. What the identification conditions is the shortcut of Theorem 10A.1(b): that loop-consistency discharges the contractible-loop content of [s]vac = 0 automatically. The shortcut rests on the interface update U{δγ} carrying the Gate-2 orientation sign as its action on the* *boundary mode — concretely, that the σ-component of the two-component (σ, ω) interface state (σ linked to orientation perturbation, ω to closure parity) is the Gate-2 matching cochain s, so that* $U\{\delta\gamma\} = \mathbb{1}$ forces $\prod\_\{\delta\gamma\}$ s = +1 a fortiori. The interface-dynamics construction supplies loop-consistency and the (σ, ω) update, but does **not** currently contain a result identifying σ's action on the boundary mode with s rather than with a closure phase whose orientation part is a quotient, or with ω. This identification is therefore **not in hand**: it is a deposited prerequisite, not a citable result. Discharging it requires either extending the interface-dynamics paper to pin U's action on the orientation component, or a separate result establishing the σ = s identification. Until then Lemma 10A.0′ and the *shortcut* part 10A.1(b) are [Conditional on the interface-orientation identification], while the *criterion* part 10A.1(a) is [Proven, under C1–C3]. The honest reading: the verdict [s]_vac = 0 stands unconditional; the identification only governs whether checking it can skip the contractible loops. Everything downstream of *granting* the identification (the contractible-loop discharge, the defect localization) inherits the condition; everything in the criterion itself does not.

The PSD reading (C1) is load-bearing. The single-hub kernel, and the perfect-square reduction of §4, assume the real positive-semidefinite reading of the competition functional. A holomorphic complex reading changes the kernel and the boundary-term derivation; C1 is confirmed for the present construction but is a genuine premise.

C2 and C3 remain finite checks. The σ-family realizer's membership in $\mathcal{T}$_L (C2) and the non-vacuity of the alternating ray under A_inc ∧ A_hub (C3) are finite verifications deposited, not discharged here. One realizing family suffices for the pass (the all-families burden is the falsification condition, not a precondition).

Gate 3 is untouched, and is where the gluing half reconnects. This paper closes the Gate-2 reachability branch: local realizability over the registered K = 7 strata, as a function of orientation-coherence. It does not address whether locally realizable completions glue over the full complex beyond the orientation (w_1) layer — that is Gate 3. Theorem 9.3 makes explicit that Gate 2's realizability-globalization is precisely the orientation layer of the global gluing question; the higher layers are Gate 3's.

Downstream questions are untouched. The compatibility/capstone questions, FBI-comp, RC, and the ℓ^2 terminal arena are not addressed beyond fixing what the residue verdict feeds them.

---

12. Conclusion

The Gate-2 branch asked whether every locally admissible K = 7 degeneracy geometry is realized by admissible transport. The prior arc reduced this to three shared-boundary hypotheses and an open boundedness question, but could not close while it declined to read the construction's boundary rule. This paper read it.

The local cost requirements — locality, label-independence, positive-semidefinite cost, the squared-linear competition primitive — determine the boundary term only **up to one sign**: the anti-aligning square $(\lambda_{\partial}^{H} + \lambda_{\partial}^{H'})^2$ and the co-aligning square $(\lambda_{\partial}^{H} - \lambda_{\partial}^{H'})^2$ are both admissible, and the cost cannot choose between them. The selector is orientation: the same fact that makes neighbouring spokes cancel inside a hub — opposite induced orientation on a shared cell — makes neighbouring hubs cancel across a shared boundary, *provided the gluing is oriented*. We named that premise B5, proved B0–B5 fix the anti-aligning term uniquely, and were explicit that B5 is not a theorem of K = 7: K = 7 proves each wheel locally orientable, while oriented gluing is the global-coherence premise that local orientability does not entail (the Möbius gap).

Against this derivation, the three prior hypotheses trace to three distinct premises: PBM from label-independence (B2), AC from the exchange symmetry that privileges neither hub (B0), and AA — the third — is exactly the orientation datum B5. The prior paper's signed sign-cocycle is, we showed, the orientation class of the hub complex: $[s] \in H^1(\Gamma_{hub}; \mathbb{Z}_2)$ is the first obstruction to a coherent orientation, and its triviality *is* orientability. The residue is empty in exactly one sector-clause under exactly one condition; the boundedness question splits into an unconditional constraint-locality half and a realizability-gluing half that is again orientation-coherence. Everything converges on one binary structural property.

The verdict is therefore a clean function of one named premise:

Gate 2 PASS for the K = 7 nearest-neighbour substrate iff $[s]_{vac} = 0$ — the vacuum orientation class is trivial — unconditionally under C1 (confirmed) and the finite checks C2, C3. K = 7 supplies local orientability; $[s]_{vac} = 0$ is the global-coherence condition the verdict rests on, not derivable from K = 7 but checkable on the vacuum sector. **Under a frustrated orientation loop ($[s]_{vac} \neq 0$), Gate 2 FAIL with the explicit N_2 witness $\hat{G}_{\gamma}$ — the alternating germ on the loop.**

This closes the branch as a two-valued verdict with both values realized, rather than a narrowing. The condition that decides it is **$[s]_{vac} = 0$**, the triviality of the vacuum orientation class — and this pass criterion is *unconditional* under C1–C3 (Theorem 10A.1(a)): §10A shows it is not derivable from K = 7, but it is a checkable topological property of the vacuum, and B5 $\Leftrightarrow$ $[s]_{vac} = 0$ needs nothing from the interface dynamics. The interface loop-consistency axiom enters only as a *verification shortcut* (Theorem 10A.1(b)): granting the deposited interface-orientation identification — that the interface update's σ-component is the Gate-2 matching cochain s (Lemma 10A.0′), which the interface-dynamics paper does not yet establish — loop-consistency discharges the contractible-loop part of $[s]_{vac} = 0$ automatically, leaving only non-contractible vacuum loops to check. The shortcut is conditional; the criterion is not. The Möbius caveat and the sufficient-not-necessary role of simple connectivity are stated in full at Remark 10A.3; the one-liner below governs.

Gate 2 PASS iff $[s]_{vac} = 0$, under C1–C3. Simple connectivity of the vacuum guarantees this via loop-consistency, but is not required; a non-simply-connected vacuum passes provided all its non-contractible loops carry trivial orientation holonomy. (Full statement and caveats: Remark 10A.3; the load-bearing interface identification: Lemma 10A.0′ and Limitations.)

The programme now advances to Gate 3: global assembly and gluing over the full complex. Theorem 9.3 has already located where Gate 2 hands off — the orientation (w_1) layer of the global gluing question is settled here; the higher layers, and the residual vacuum-holonomy premise, are Gate 3's to take up.

A second forward direction, distinct from Gate 3 and from the verdict, is recorded in §10B and stated here only as a pointer. The orientation class this paper identified for the *gluing* of the vacuum is also the natural home, one layer up in *configuration* space, for a localized orientation defect on the already-oriented substrate — a candidate realization of the long-standing VERSF reading of particles as persistent transport defects rather than fundamental objects. Granting the same interface-orientation identification §10A′ deposits, the loop-consistency shortcut (Theorem 10A.1(b)) makes the localization a theorem rather than a hope: a non-trivial configuration holonomy [s]_γ can live only on a non-contractible loop, so an orientation defect is defect-encircling content, never an intrinsic property of a trivial-holonomy vacuum; absent that identification the localization is conjectural. This is deferred to a dedicated paper; its load-bearing open theorem is refinement-stability of the defect sectors under $\mathscr{R}$, without which a topological sector label is not yet a persistent particle. The present verdict neither depends on nor anticipates that result: Gate 2 concerns the vacuum gluing, and configurations on the realized vacuum are a separate question in a layer Gate 2 does not test.