Admissibility and the $\mathbb{Z}_7$ Closure Connection

Flatness, Holonomy, and Occupancy in the K = 7 Transport Channel

General Reader Summary

The previous paper identified the transport group of the K = 7 closure construction and showed that any closure content beyond orientation cannot appear on a single overlap. Whatever transport content survives must instead live on loops, as a possible $\mathbb{Z}_7$-valued holonomy.

That result sharply reduced the problem but did not solve it. Knowing where a surplus could live is not the same as knowing whether the substrate actually permits it.

This paper studies that question directly.

The central observation is that identifying the transport group only determines the possible closure channel. It does not determine how the substrate's admissibility rules act on that channel. The hub-completion, conservation, and competition constraints may either permit nontrivial $\mathbb{Z}_7$ holonomy or force every loop to be trivial.

We therefore construct the closure connection itself and ask three questions.

1. Does admissibility force the $\mathbb{Z}_7$ transport to be flat on elementary plaquettes?
2. If flatness holds, what cohomology class remains on non-bounding loops?
3. Does admissibility populate that class or force it to vanish?

The answer determines whether the closure channel survives as genuine transport structure or collapses to redundancy.

There is one subtlety the construction must handle before any of this is well-posed. The transport group is non-commutative — orientation reversals invert rotations — so a naive sum of rotation labels around a loop is not even well-defined until the orientation behaviour of that loop is pinned down. The earlier orientation-coherence result supplies exactly the missing ingredient: along every admissible loop the orientation holonomy is trivial, which removes the inversions and lets the rotation labels add cleanly. Only after this step does the closure curvature become a genuine cochain.

The result is a reformulation of Gate 3. The problem is no longer one of identifying the transport group. The transport group is already known. The problem is whether admissibility populates the $\mathbb{Z}_7$ channel.

Contents

Abstract

The transport-group analysis identified the strong-model transport group as

$G \cong D_7$ (dihedral, order 14),

with orientation projection

$\pi\_or : D_7 \to \mathbb{Z}_2$

and rotation kernel

$N = \ker(\pi\_or) \cong \mathbb{Z}_7$.

Here the *strong model* denotes the K = 7 closure construction equipped with the full admissibility set (hub completion, conservation, competition); the transport group is the group generated by admissible single-face transport elements modulo gauge. That analysis further proved that single-face transport is gauge-equivalent to orientation alone, so any surviving closure content must occur as loop holonomy rather than as a local overlap degree of freedom. The coefficient group of the resulting connection is therefore not a modelling choice — it is forced to be $\mathbb{Z}_7$ by the transport group together with single-face rigidity (Section 2).

The present paper studies whether admissibility populates that holonomy sector. We work on the vacuum transport 2-complex Γ_vac (vertices = hubs, edges = admissible transport faces, 2-cells = minimal admissible closure cycles), constructed in Section 3 as the combinatorial shadow of admissible closure transport, and assign a rotation label

$\rho\_f \in \mathbb{Z}_7$

to each admissible transport face f as the rotation coordinate of its transport element $g\_f \in D_7$ — with $g\_f = s^\{\rho\_f\} \in N$ on orientation-even faces and $g\_f = r\cdot s^\{\rho\_f\} \in r\cdot N$ on orientation-odd faces.

Because D_7 is non-abelian, the rotation labels do not in general add around a loop: an orientation reversal conjugates $s \mapsto s^{-1}$, so the $\mathbb{Z}_7$ fibre carries a $\mathbb{Z}_2$ action twisted by the orientation character. We show — using the Gate-2 orientation-coherence condition (H-OC), which forces trivial orientation holonomy on every admissible loop — that this twist is trivial along all admissible loops (the cycle-level instance H-OC(γ)). The twisted coefficient system therefore reduces to ordinary $\mathbb{Z}_7$, and the closure curvature is the well-defined valued cochain

$\Omega(P) = \Sigma\_\{f \subset \partial P\}\ \varepsilon\_\{P,f\}\ \rho\_f \pmod 7$,

with $\varepsilon\_\{P,f\} \in \{+1, -1\}$ the orientation of face f in the boundary ∂P.

Vanishing closure curvature,

$\Omega(P) = 0$ for every elementary plaquette P,

is the exact $\mathbb{Z}_7$ analogue of the orientation-coherence condition. Under this condition the closure connection defines a cohomology class

$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$.

The remaining Gate-3 question is therefore reduced to occupancy:

$\kappa = 0$ or $\kappa \neq 0$?

We derive the reduction but do not assume the verdict. The admissibility action on the labels ρ_f is set up here; its explicit computation — and hence the verdict on flatness and occupancy — is deferred to the sequel. The paper converts Gate 3 from a transport-group problem into a flatness-and-occupancy problem.

---

1. Introduction

The transport-group paper settled three structural questions.

1. The transport group, $G \cong D_7$ (order 14).
2. The orientation projection, $\pi\_or : D_7 \to \mathbb{Z}_2$.
3. The location of any surviving surplus — loop holonomy, not single faces.

What remains is whether admissibility populates the closure channel. The objective of this paper is therefore not to identify new transport structure but to determine the mathematical object that carries the residue, and to state precisely what must be computed to decide whether that residue survives.

Let

$N = \ker(\pi\_or) \cong \mathbb{Z}_7$.

The transport-group analysis established that any observable element of N must occur through loop transport. Accordingly the natural question becomes:

Does admissibility permit nontrivial $\mathbb{Z}_7$ loop holonomy?

This is the occupancy problem.

Before it can be posed cleanly, one structural hazard must be cleared. N sits inside the non-abelian group D_7, where reflections invert rotations. Section 4 isolates this hazard; Section 5 removes it using the Gate-2 coherence result and only then defines the curvature.

1.1 Scope and novelty

The cohomological machinery used below is standard. In particular, the classification of flat abelian connections by first cohomology, $H^1(\Gamma\_vac; \mathbb{Z}_7)$, is a classical result of cellular cohomology and discrete gauge theory. The novelty of the present paper is not the cohomological classification itself.

The novelty lies in identifying the correct transport object to which that classification applies. Three programme-specific steps are required before the standard machinery becomes available:

1. determining the transport group from the closure-frame primitives, $G \cong D_7$;
2. identifying the transport kernel $N \cong \mathbb{Z}_7$ and showing the surviving residue must live there (Section 2);
3. showing that orientation coherence trivialises the non-abelian twist inherited from D_7 (Section 5).

Only after these steps does the closure residue become an ordinary $\mathbb{Z}_7$-valued flat connection. The contribution of the paper is therefore not the cohomology theorem (Theorem 5.3) but the reduction of Gate 3 to the hypotheses under which that theorem applies.

---

2. Why a $\mathbb{Z}_7$ Connection Is Forced

The introduction of a $\mathbb{Z}_7$-valued connection is not an additional modelling assumption. It follows directly from the transport-group construction.

The transport-group analysis established that, under the strong model, $G \cong D_7$ with orientation projection $\pi\_or : D_7 \to \mathbb{Z}_2$ and kernel $N = \ker(\pi\_or) \cong \mathbb{Z}_7$. The same analysis further proved that single-face transport is gauge-equivalent to its orientation parity alone (single-face rigidity). Consequently any transport information surviving beyond Gate 2 cannot reside in local overlaps and must instead appear through transport around loops.

A loop-supported transport residue must therefore take values in the part of the transport group invisible to orientation — namely the kernel N. Since $N \cong \mathbb{Z}_7$, the surviving transport datum is necessarily a $\mathbb{Z}_7$-valued object.

The appearance of a $\mathbb{Z}_7$ connection is thus not postulated independently. It is forced by three earlier results, in combination:

- the transport group, $G \cong D_7$;
- the kernel identification, $N = \ker(\pi\_or) \cong \mathbb{Z}_7$;
- single-face rigidity.

The only remaining freedom concerns occupancy of the resulting connection, not its coefficient group. This is made precise as Proposition 5.1a, once the connection has been constructed.

---

3. The Vacuum Transport Complex

The transport complex Γ_vac is not introduced as an independent geometric object. It is the combinatorial shadow of admissible closure transport, and we take it throughout to be the *admissible* transport 2-complex. Its cells arise from the closure architecture:

- 0-cells: admissible vacuum hubs.
- 1-cells: admissible transport faces between neighbouring hubs, each oriented.
- 2-cells: minimal admissible closure cycles generated by hub completion, each carrying a boundary orientation.

The appearance of 2-cells is therefore not optional. Whenever admissible transport permits a minimal closed cycle of faces whose completion is enforced by the closure rules, that cycle defines a plaquette. Accordingly Γ_vac is naturally a cellular 2-complex rather than merely a graph.

This distinction is essential. If the transport structure were only a 1-complex, every loop observable would automatically define a homomorphism

$\pi_1(\Gamma\_vac) \to \mathbb{Z}_7$,

flatness would be vacuous, and there would be no notion of local curvature. The existence of admissibility-generated plaquettes introduces the possibility of local closure curvature, and therefore of a genuine flatness condition. The Gate-3 problem is inherently a problem on a transport 2-complex, not on a transport graph.

By the admissible-complex convention every loop of Γ_vac is an admissible loop. This is what reconciles the scope of Lemma 5.1 (detwisting, stated on admissible loops) with that of Theorem 5.3 (classification, stated over the full complex Γ_vac): the admissible subcomplex *is* Γ_vac. Were one to enlarge Γ_vac to include inadmissible loops, the detwisting and the classification would have to be carried over the admissible subcomplex Γ_vac^adm ⊆ Γ_vac instead, and $H^1(\Gamma\_vac; \mathbb{Z}_7)$ replaced by $H^1(\Gamma\_vac\text{^}adm; \mathbb{Z}_7)$ throughout.

---

4. Closure Transport as a Twisted $\mathbb{Z}_7$ Connection

On the vacuum transport 2-complex Γ_vac of Section 3, each admissible face f carries a transport element g_f ∈ D_7, and a rotation coordinate

$\rho\_f \in \mathbb{Z}_7$,

defined as the rotation part of g_f relative to the generators $\langle r, s \mid s^7 = r^2 = 1, r\, s\, r^{-1} = s^{-1} \rangle$, with s the order-7 rotation. Concretely, the extraction is unambiguous: every element of D_7 has a unique normal form, either s^k or r·s^k with k ∈ $\mathbb{Z}_7$, and ρ_f is defined to be that exponent k. (The convention that reflections are written as r·s^k, rather than s^k·r = r·s^{−k}, fixes the sign once and for all.) The faces split by orientation parity:

orientation-even faces: g_f = s^{ρ_f} ∈ N, orientation-odd faces: g_f = r · s^{ρ_f} ∈ r·N.

So ρ_f is read directly off the normal form of the closure-transport element on face f; nothing is assumed about its existence beyond the group law of D_7. It is essential that ρ_f is the *coordinate* of g_f, not the element itself. Writing g_f = s^{ρ_f} for all faces would place every transport element in the rotation subgroup N, collapse D_7 to its abelian part, and erase the entire non-abelian hazard this paper exists to handle. Odd faces live in the non-trivial coset r·N; only their rotation coordinate is recorded by ρ_f. The collection

ρ = { ρ_f }

is then a candidate $\mathbb{Z}_7$ connection on Γ_vac, twisted by the orientation parities of the faces.

4.1 The non-abelian hazard

Loop holonomy is the *group* product of the g_f around a loop, taken inside D_7, not an arithmetic sum of labels. Write the orientation character

$$w : \pi_1(\Gamma\_vac) \to \mathbb{Z}_2 \cong \{\pm 1\},$$

induced by π_or, recording the net orientation of a loop. In D_7 a reflection conjugates rotations by inversion,

$$r\, s\, r^{-1} = s^{-1},$$

so transporting a rotation label across a reflection sends $\rho \mapsto -\rho$. Consequently the rotation labels of a loop combine with signs dictated by the partial orientation accumulated along the loop. Equivalently, the $\mathbb{Z}_7$ fibre carries a $\mathbb{Z}_2$ action through w, and the correct coefficient system is the **twisted** group $\mathbb{Z}_7\_w$, not plain $\mathbb{Z}_7$.

This is the precise reason the additive curvature formula cannot be written down yet. On a loop with nontrivial orientation holonomy the expression $\Sigma\ \rho\_f$ is not gauge-invariant and not even well-defined as a $\mathbb{Z}_7$ quantity, because the signs ε depend on where the reflections fall.

4.2 Gauge action

Gauge transformations act by conjugation at hubs, $g\_f \mapsto h\_\beta \cdot g\_f \cdot h\_\alpha^{-1}$ with hub elements h = s^{α}. On the rotation coordinate this gives

$\rho\_f \mapsto \rho\_f + (\alpha\_\beta - \alpha\_\alpha)$ on orientation-even edges, $\rho\_f \mapsto \rho\_f - (\alpha\_\beta + \alpha\_\alpha)$ across orientation-odd edges,

for adjacent hubs α, β with hub parameters $\alpha\_\alpha, \alpha\_\beta \in \mathbb{Z}_7$. The odd-edge sign is forced by the coset structure: for an odd face g_f = r·s^{ρ},

$$s^{\alpha\_\beta} \cdot (r \cdot s^{\rho}) \cdot s^{-\alpha\_\alpha} = r \cdot s^{-\alpha\_\beta} \cdot s^{\rho} \cdot s^{-\alpha\_\alpha} = r \cdot s^{\rho - \alpha\_\alpha - \alpha\_\beta},$$

using $s^a\, r = r\, s^{-a}$. The rotation coordinate of the conjugated odd element is therefore $\rho - \alpha\_\alpha - \alpha\_\beta$, exactly the odd-edge rule above. This is the operational confirmation that odd faces are coset elements r·N, not rotations.

The gauge group here is rotation-valued ($h = s^{\alpha} \in N$), not the full D_7. This is the correct residual freedom, not an arbitrary restriction: the reflection part was already fixed at Gate 2, where orientation was gauged to coherence. A gauge group equal to all of D_7 would act with reflections, mix orientation-even and orientation-odd faces, and reopen exactly the twist this section is isolating; the rotation-valued residual is what remains once orientation is pinned. Only loop holonomy survives this action — the content of the transport-group reduction — but the holonomy is, a priori, valued in the twisted coefficients.

---

5. Detwisting via Orientation Coherence, and the Closure Curvature

The twist of Section 4 is removed by the Gate-2 result, not by assumption.

In plain terms, the section turns on a single distinction. H-OC(∂) — coherence on plaquette boundaries — is enough to *define flatness*. H-OC(γ) — coherence on all loops — is what is needed to *define a cohomology class*. The reader who keeps only that sentence in mind will not lose the thread of what follows.

5.1 Orientation coherence trivialises the twist

Lemma 5.1 (Detwisting). [proven, conditional on H-OC] The Gate-2 orientation-coherence condition H-OC has two logically separable instances, distinguished by the class of loops on which orientation holonomy is required to vanish:

H-OC(∂): $w(\partial P) = +1$ for every elementary plaquette P, H-OC(γ): $w(\gamma) = +1$ for every loop γ of Γ_vac.

Since each plaquette boundary ∂P is a loop, H-OC(γ) ⟹ H-OC(∂); the converse does not hold, because plaquette boundaries span only B_1, whereas loops span all of Z_1. The two instances do distinct work:

(i) Under H-OC(∂), the boundary holonomy of every plaquette lies in the rotation kernel N: an even number of reflections multiplies to a pure rotation, so the parity obstruction vanishes. This makes the *flatness predicate* of Definition 5.2 well-posed. It does **not**, however, make Ω a well-defined valued $\mathbb{Z}_7$ cochain. The rotation coordinate of the boundary holonomy is fixed only up to an overall sign, depending on basepoint and traversal, because a basepoint shift across a reflection edge negates it; and the raw traversal-sign sum $\Sigma\ \varepsilon\_\{P,f\}\ \rho\_f$ is not equal to that coordinate when odd faces are present, since the correct signs must also carry the reflection parity accumulated to each face — which is path data, not boundary-orientation data.

(ii) Under H-OC(γ), the action is trivial around every loop, so the local system has trivial monodromy on the connected complex Γ_vac and admits a global trivialization:

$\mathbb{Z}_7\_w = \mathbb{Z}_7$ on Γ_vac.

In that frame the connection is rotation-valued throughout — the reflection parts absorbed — the per-plaquette sign ambiguity of (i) is reconciled coherently across shared edges, the signs ε_{P,f} reduce to the honest cellular incidence signs, and the closure curvature Ω(P) is a well-defined valued $\mathbb{Z}_7$ 2-cochain. The cohomology class κ of a flat connection is then well-defined in ordinary $H^1(\Gamma\_vac; \mathbb{Z}_7)$.

Proof. The $\mathbb{Z}_2$ action on the fibre is pulled back from {±1} along w, so a loop acts as the identity exactly when its orientation holonomy is +1. For (i), H-OC(∂) gives $w(\partial P) = +1$, hence an even number of odd faces on each boundary and a boundary holonomy in N. But the coordinate is basepoint-sensitive: for an oriented boundary with faces a (odd), b (even), c (odd) — net parity even — reading from a gives g_a g_b g_c = (r s^{ρ_a})(s^{ρ_b})(r s^{ρ_c}) = s^{ρ_c − ρ_a − ρ_b}, whereas reading the same boundary from b gives g_b g_c g_a = s^{ρ_b}(r s^{ρ_c})(r

$s^{\rho_a}) = s^{\rho_a + \rho_b - \rho_c} = s^{-(\rho_c - \rho_a - \rho_b)}$, using $s^k r = r s^{-k}$. The coordinate negates across a reflection edge; only the predicate (coordinate ≡ 0), which is sign-robust since $0 = -0$, is unambiguous. For (ii), H-OC(γ) gives $w(\gamma) = +1$ on all loops, so the monodromy of the local system is trivial; a local system with trivial monodromy on a connected complex is isomorphic to the constant system with the same fibre, which furnishes the global trivialization. Reconciling the per-plaquette sign choices of (i) across edges shared by adjacent plaquettes is precisely the monodromy condition, i.e. H-OC(γ) — which is why the valued cochain belongs to (ii) and not to (i).

The marker [conditional on H-OC] is essential: detwisting is *proven given* the Gate-2 coherence result and is not an independent assumption of this paper. The dependence is graded. If H-OC(∂) fails on some plaquette, the boundary holonomy there carries a residual reflection, no rotation coordinate exists, and the flatness predicate of Definition 5.2 is not even well-posed. If H-OC(∂) holds but H-OC(γ) fails on some non-bounding loop, flatness is well-posed plaquette-by-plaquette, but the per-plaquette sign choices do not reconcile across shared edges, Ω is not a globally well-defined valued cochain, and the class lives in twisted cohomology $H^1(\Gamma\_vac; \mathbb{Z}_7\_w)$ rather than ordinary $H^1(\Gamma\_vac; \mathbb{Z}_7)$. Both contingencies are recorded but not pursued here.

Proposition 5.1a (Uniqueness of the closure coefficient group). [proven, given the transport-group theorem] Under the strong model, any loop-supported transport residue invisible to orientation takes values in $\mathbb{Z}_7$, and no alternative coefficient group is available.

Proof. A residue invisible to orientation lies in ker(π_or). By the transport-group theorem ker(π_or) = N ≅ $\mathbb{Z}_7$. Hence the fibre of any such residue is $\mathbb{Z}_7$; an alternative coefficient group would require a different transport group, contradicting G ≅ D_7.

This fixes the coefficient *group* — the fibre — but not its twisting. Whether the fibre appears as the twisted system $\mathbb{Z}_7\_w$ or the constant system $\mathbb{Z}_7$ is the separate question settled by Lemma 5.1; Proposition 5.1a asserts only that, twisted or not, the group is $\mathbb{Z}_7$ and nothing else. Together with Section 2, it removes any impression that $\mathbb{Z}_7$ was chosen rather than derived.

5.2 The closure-curvature cochain

Work in the global trivialization furnished by Lemma 5.1(ii), where the connection is rotation-valued throughout and the reflection parts have been absorbed. In that frame, for each plaquette P and each face f ⊂ ∂P the boundary-orientation sign

$\varepsilon_{P,f} \in \{+1, -1\}$

is the honest cellular incidence sign — equal to +1 when f is traversed with the orientation of ∂P and −1 otherwise — and the per-plaquette sign ambiguity of Lemma 5.1(i) has been reconciled across shared edges. Define the closure-curvature 2-cochain

$\Omega(P) = \Sigma_{f \subset \partial P} \varepsilon_{P,f} \rho_f \pmod 7$.

By Lemma 5.1(ii), H-OC(γ) is what makes this a well-defined valued $\mathbb{Z}_7$ 2-cochain; H-OC(∂) alone secures only the flatness predicate, not the value. In the trivialized frame Ω is moreover gauge-invariant, the additive gauge shifts of Section 4.2 reducing to coboundaries whose contribution telescopes around any closed boundary and cancels.

Definition 5.2 (Closure flatness)

Assume H-OC(∂), so that every plaquette boundary holonomy is a pure rotation and the predicate below is well-posed. The connection is *flat* iff the boundary holonomy of every elementary plaquette is trivial, equivalently

$\Omega(P) = 0$ for every elementary plaquette P.

The equivalence uses only that the rotation coordinate, though fixed only up to sign under H-OC(∂), vanishes regardless of sign — the predicate is sign-robust ($0 = -0$), so it does not require the valued cochain of Section 5.2.

Three conditions are in play and should not be conflated. H-OC(∂) — orientation-triviality on plaquette boundaries — makes the flatness *predicate* well-posed, and is logically prior to flatness. Flatness, $\Omega(P) = 0$, is a further and independent condition: a predicate can be well-posed yet false. H-OC(γ) — orientation-triviality on all loops — is what makes Ω a well-defined *valued* cochain and promotes a flat connection to a well-defined class κ. Flatness is the direct $\mathbb{Z}_7$ analogue of the orientation-sector vanishing condition; the coherence conditions H-OC(∂), H-OC(γ) are the analogues of the orientation-sector well-definedness conditions, at boundary level (predicate) and cycle level (value) respectively.

Theorem 5.3 (Classification of flat closure connections). [proven, standard cellular cohomology]

Assume H-OC(γ), and let ρ denote the connection in the global trivialization of Lemma 5.1(ii). If the closure connection is flat, then ρ defines a cohomology class

$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$,

and gauge-equivalent flat connections define the same class. Conversely every class in $H^1(\Gamma\_vac; \mathbb{Z}_7)$ arises from some flat connection.

Proof. This is the standard discrete-gauge statement, recorded for completeness rather than as a programme-specific result. In the global trivialization the connection is rotation-valued and the reflection-conjugation has been absorbed, so the incidence signs $\varepsilon\_\{P,f\}$ are the honest cellular coboundary signs; this is why the theorem may take H-OC(γ) for granted. Identify ρ with a 1-cochain in $C^1(\Gamma\_vac; \mathbb{Z}_7)$. The curvature of Section 5.2 is then the cellular coboundary, $\Omega = d^1\rho$. Flatness, $\Omega(P) = 0$ on every 2-cell, is exactly $d^1\rho = 0$, so ρ is a 1-cocycle. The gauge shifts of Section 4.2, post-detwisting, are exactly the image of d^0 acting on the hub parameters $\{ \alpha\_\alpha \} \in C^0(\Gamma\_vac; \mathbb{Z}_7)$, i.e. coboundaries. Hence gauge classes of flat connections are

$\ker d^1 / \mathrm{im}\, d^0 = H^1(\Gamma\_vac; \mathbb{Z}_7)$,

with the cohomology taken on the full 2-complex Γ_vac, so that plaquette flatness — not merely the 1-skeleton loop condition — enters.

The content of Theorem 5.3 is definitional; its role is to name the object precisely so that the genuinely open questions of Sections 6–7 can be stated against it. As noted in Section 1.1, the theorem itself is classical; the programme-specific work is everything that licenses its application — Sections 2 through 5.

---

6. The Occupancy Question

Two possibilities remain. Both are **open** pending the admissibility computation of Section 7.

Branch A. [conjectural]

Admissibility forces

$\kappa = 0$.

Then the closure channel survives algebraically but not physically: the $\mathbb{Z}_7$ structure is present in the construction yet carries no realised holonomy. The closure residue collapses to redundancy.

Branch B. [conjectural]

Admissibility permits

$\kappa \neq 0$.

Then a genuine closure-holonomy sector survives, supported on non-bounding loops of Γ_vac.

The transport-group paper identified the algebraic sector $N\_alg \cong \mathbb{Z}_7$. The present analysis reduces the remaining problem to determining the holonomy sector N_hol, defined as the image of admissible flat connections in $H^1(\Gamma\_vac; \mathbb{Z}_7)$.

A structural remark sharpens the stakes, and corrects a tempting error. Plaquette boundaries ∂P span the boundary group $B_1(\Gamma\_vac)$, not the homology $H_1(\Gamma\_vac) = Z_1/B_1$; flatness is a statement about B_1 (the curvature vanishes on every plaquette boundary), while the surviving class lives on H_1. Because $\mathbb{Z}_7$ is a field, the universal-coefficient theorem gives $H^1(\Gamma\_vac; \mathbb{Z}_7) \cong \mathrm{Hom}(H_1(\Gamma\_vac; \mathbb{Z}_7), \mathbb{Z}_7)$. Hence:

- If $H_1(\Gamma\_vac; \mathbb{Z}_7) = 0$, then $H^1(\Gamma\_vac; \mathbb{Z}_7) = 0$, so $\kappa = 0$ necessarily — Branch A holds by topology alone. Note carefully that $\kappa = 0$ here is forced by the *vanishing of H^1*, not by flatness: flatness only places ρ in the cocycle group, and an empty H^1 leaves no nontrivial class for it to represent.

- Branch B is available only if $H_1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$, i.e. only if Γ_vac carries non-bounding loops.

Whether the vacuum transport complex has nontrivial H_1 is itself an admissibility question, not a free choice, and is part of what Section 7 must settle.

The two branches and the topology remark are best compressed into a single chain.

Proposition 6.1 (Occupancy trichotomy). [proven, given Theorem 5.3 and universal coefficients] Assume H-OC(γ). The survival of a nontrivial closure-holonomy sector — a nonzero class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ — is governed by three conditions, in a descending chain in which each presupposes the previous:

1. Flatness: $\Omega(P) = 0$ for every elementary plaquette P;
2. Topology: $H^1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$;
3. Population: $\kappa \neq 0$.

These are not logically independent — (3) presupposes (2), which presupposes (1) — but hierarchical, and the Gate-3 verdict is their conjunction. Failure at each level collapses the closure channel by a distinct mechanism:

- if (1) fails, ρ is not a cocycle and κ is undefined as a class;
- if (1) holds but (2) fails, κ is defined but forced to vanish, since $H^1(\Gamma\_vac; \mathbb{Z}_7) = 0$ admits no nonzero class;
- if (1) and (2) hold but (3) fails, κ is defined and a nonzero class is available, yet admissibility populates $\kappa = 0$.

Only when all three hold does a nontrivial closure-holonomy sector survive:

$$\Omega = 0 \longrightarrow H^1(\Gamma\_vac; \mathbb{Z}_7) \neq 0 \longrightarrow \kappa \neq 0.$$

Proof. Condition (1) is Theorem 5.3: a flat connection, and only a flat connection, defines a class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$. Condition (2) is the universal-coefficient identification $H^1(\Gamma\_vac; \mathbb{Z}_7) \cong \mathrm{Hom}(H_1(\Gamma\_vac; \mathbb{Z}_7), \mathbb{Z}_7)$ over the field $\mathbb{Z}_7$, by which a trivial H^1 forces every class to zero. Condition (3) is the definition of occupancy. The descending presupposition is immediate: a nonzero element requires a nonzero group, which requires the class to be defined.

The three conditions of Proposition 6.1 are precisely the three computational targets of Section 7, in order. The proposition is proven as a logical skeleton; which of its cases obtains is the open verdict.

---

7. What Must Be Computed

The remaining missing ingredient is the action of the three admissibility rules,

- hub completion,
- conservation,
- competition,

on the rotation labels ρ_f and on the 2-cell structure of Γ_vac. These rules determine, in order:

1. whether closure flatness (Definition 5.2) holds — i.e. whether $\Omega(P) = 0$ is forced on every admissible plaquette;
2. whether admissibility generates a transport complex whose topology contains non-bounding loops — i.e. whether $H_1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$ — so that nontrivial classes are available at all;
3. whether admissibility populates such a class or annihilates it.

The hypothesis levels carry forward as in Section 5. Item (1) is a predicate question and needs only the boundary-level instance H-OC(∂), under which flatness is well-posed plaquette-by-plaquette. Items (2)–(3) concern κ as a class and therefore presuppose the cycle-level instance H-OC(γ), under which Ω is a valued cochain and the class is defined.

The Gate-3 problem is thereby reduced to the map

$N\_alg \longrightarrow N\_hol$.

The subsequent operational-quotient programme then determines the further reduction

$N\_hol \longrightarrow N\_obs$,

from holonomy sector to observable sector.

The computation in (1)–(3) is the substance of the sequel and is not attempted here. What this paper fixes is the *target* of that computation: a single $\mathbb{Z}_7$ cohomology class on the coefficient system, untwisted under H-OC(γ), with the orientation hazard discharged.

Conclusion

The transport-group analysis identified the closure channel. The present paper identifies the mathematical object that carries it, and clears the one obstruction — the orientation twist inherited from the non-abelian transport group D_7 — that stood between that object and a well-posed curvature.

The coefficient group is forced, not chosen: the surviving residue lives in $\ker(\pi\_or) \cong \mathbb{Z}_7$, and no other group is available (Section 2, Proposition 5.1a). The carrying complex is forced, not assumed: Γ_vac is the combinatorial shadow of admissible closure transport, and its plaquettes are the minimal admissible closure cycles enforced by hub completion (Section 3).

The detwisting is likewise not assumed. It follows from the Gate-2 orientation-coherence condition: with orientation holonomy trivial on every admissible loop — the cycle-level instance H-OC(γ) — the $\mathbb{Z}_2$ action on the $\mathbb{Z}_7$ fibre trivialises, and the twisted coefficient system reduces to ordinary $\mathbb{Z}_7$. Only then does the closure curvature

$$\Omega(P) = \Sigma\_\{f \subset \partial P\}\ \varepsilon\_\{P,f\}\ \rho\_f \pmod 7$$

become a cochain, and only then does flatness define a class

$$\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7).$$

The cohomological classification of that class is standard; the programme-specific contribution is the chain of reductions that makes the standard machinery applicable. The remaining question is occupancy. Gate 3 is reduced to a single problem:

Does admissibility populate the $\mathbb{Z}_7$ holonomy sector or force it to vanish?

The reduction is proven (conditional on H-OC). The verdict is open.