The Topology of the Vacuum Transport Complex

Boundary Ranks, the Cyclomatic Criterion, and the Availability of the $\mathbb{Z}_7$ Closure Sector

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The earlier transport papers found a hidden sevenfold transport channel and showed that, if anything of it survives, it can survive only as transport around loops. A later paper reduced the question further: any surviving residue would be carried by a single cohomology class

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

and Gate 3 became the question of whether the substrate's admissibility rules populate that class.

But there is a prior question, cheaper than occupancy and capable of ending the matter on its own. A class can be populated only if it exists. And whether it exists is not a question about the admissibility *dynamics* at all — it is a question about the *shape* of the transport complex. If the complex has no loops that fail to close, there is no class to populate, and the occupancy programme stops before it starts.

This paper isolates that prior question and reduces it to a single computation.

The key idea is elementary. A network of hubs and transport faces has a definite number of independent loops — its cyclomatic number, fixed by simple counting. The completion rule fills some of those loops with plaquettes, and a filled loop no longer counts. The sevenfold sector is available precisely when the plaquettes fail to fill every loop — when some independent loop is left open. So the topological verdict comes down to comparing two integers: the number of independent loops in the transport graph, and the number of them that completion actually bounds.

We do not compute those integers here, because that requires the explicit cell structure of the vacuum, which is the task of the sequel. What we establish is that *those two integers are the only data the topology depends on* — everything else is settled. The paper separates two questions that had been tangled together: whether a sevenfold class is topologically available, and whether admissibility populates it. Only the second is the occupancy problem. The first is arithmetic.

Contents

Abstract

The closure-connection analysis reduced Gate 3 to the possible existence of a cohomology class

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

where Γ_vac is the admissible vacuum transport complex: hubs as 0-cells, admissible transport faces as 1-cells, and the minimal admissible closure cycles enforced by hub completion as 2-cells. The reduction left two logically distinct questions entangled — whether such a class is *topologically available*, and whether admissibility *populates* it. This paper resolves the first independently of the second.

Working over the field $\mathbb{Z}_7$, we set up the cellular chain complex

$C_2(\Gamma\_vac) \xrightarrow{\partial_2} C_1(\Gamma\_vac) \xrightarrow{\partial_1} C_0(\Gamma\_vac)$,

and reduce the availability question to rank arithmetic. For a connected complex the first Betti number is

$\dim_{\mathbb{Z}_7} H_1(\Gamma_{vac}; \mathbb{Z}_7) = \mu - \mathrm{rank}_{\mathbb{Z}_7} \partial_2, \mu = |E| - |V| + 1,$

where μ is the cyclomatic number of the transport graph (the loop content of the 1-skeleton) and rank ∂_2 is the bounding power of admissible completion. The central result is the **Cyclomatic Availability Criterion**: a nontrivial $\mathbb{Z}_7$ closure sector is topologically available if and only if

$\mathrm{rank}_{\mathbb{Z}_7} \partial_2 < \mu,$

equivalently, the plaquette boundaries fail to span the cycle space of the transport graph. The universal-coefficient theorem over the field $\mathbb{Z}_7$ transfers this from homology to cohomology without an Ext correction.

The paper does not compute μ or rank ∂_2, since both require the explicit cell structure of the vacuum, deferred to the sequel. What it fixes is that these two integers are the *only* data on which the topological verdict depends, and that the verdict reduces to a single $\mathbb{Z}_7$-rank computation on the plaquette boundary operator. The criterion is proven; the cell counts, and hence the verdict, are open.

---

1. Introduction

The transport-group programme has proceeded through three stages.

The first identified the transport group

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

and the orientation-blind transport kernel

$N = \ker(\pi_{or}) \cong \mathbb{Z}_7.$

The second established single-face rigidity: no kernel content survives at the level of local overlaps, so any residue must be loop-supported.

The third constructed the closure connection and proved — conditional on the Gate-2 orientation-coherence condition — that any surviving residue is carried by a class

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

That paper closed with a three-condition chain (its Proposition 6.1): a nontrivial closure-holonomy sector survives only if **(1)** the connection is flat, **(2)** the topology is nontrivial, $H^1(\Gamma_{vac}; \mathbb{Z}_7) \neq 0$, and **(3)** admissibility populates a nonzero class. The conditions are hierarchical — each presupposes the previous — and the Gate-3 verdict is their conjunction.

This paper addresses condition (2) alone, and addresses it completely up to one explicit computation. The distinction it enforces is between two questions that the reduction left adjacent:

Topological availability. Does Γ_vac support nontrivial classes at all?

Occupancy. If it does, does admissibility populate one?

Availability is the cheaper question, and it is a potential terminator. If the topology is trivial, condition (2) fails, and by Proposition 6.1 the closure channel collapses to Branch A — redundancy — regardless of anything the admissibility dynamics might do. The occupancy computation, which is the expensive semantic bridge from each admissibility rule to a constraint on rotation labels, never has to be attempted. It is therefore worth settling availability first, on its own terms, before any dynamical work begins.

The remark that makes this tractable is that availability depends on Γ_vac only as a *cellular complex* — its cells and their incidences — and not on the rotation labels ρ_f, the curvature Ω, or the action of the rules. It is pure topology, and topology over a field is rank arithmetic. The contribution of this paper is to carry out that reduction to the point where a single integer comparison remains.

1.1 Scope and boundary of this paper

Two limits are deliberate.

First, this paper takes the *topological* content of hub completion only — that completion generates 2-cells. Whether transport around a completed hub is dynamically required to trivialise (which would force flatness, condition (1)) is a separate matter belonging to the occupancy paper. Here a completed hub contributes a plaquette to C_2 and nothing more; its boundary is an element of C_1, and its bounding action is measured by rank ∂_2. No claim is made about the *values* ρ_f carried on those cells.

Second, the cohomological machinery — cellular chain complexes, rank-nullity, the universal-coefficient theorem — is entirely standard. The novelty is not the machinery but its application: identifying that the Gate-3 topology question reduces to the single comparison rank $\partial_2 < \mu$, and identifying both sides of that comparison with programme-meaningful quantities.

Third — and this is a hypothesis, not a limit, because the legitimacy of the whole availability-before-dynamics programme depends on it — we assume that **cell-admissibility is purely combinatorial**: whether a given hub, face, or plaquette belongs to Γ_vac is determined by the closure rules without reference to the rotation labels ρ_f. Call this hypothesis (CA). Under (CA) the complex Γ_vac is fixed before any label is assigned, availability is genuinely label-free, and the topology of this paper can be settled before the dynamics of the occupancy paper. If (CA) failed — if admissibility of a cell were itself a condition on the labels it carries — then which cells exist would depend on ρ_f, the factorisation of availability from occupancy would leak, and the two could not be separated cleanly. We flag (CA) as the load-bearing assumption of the separation and as something the sequel's explicit derivation of Γ_vac must **confirm by**

construction — exhibiting cell-admissibility as a combinatorial predicate on the closure architecture — rather than inherit. The proven results below hold for whatever complex the sequel derives; what (CA) secures is that "the topology of Γ_vac" is a well-posed object independent of the labels at all.

---

2. The Cellular Chain Complex over $\mathbb{Z}_7$

Definition 2.1 (Admissible vacuum transport complex)

Γ_vac is the cellular 2-complex generated by admissible closure transport, with cells:

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

We take Γ_vac connected throughout — the admissible transport graph is connected, which is a property of the substrate established upstream and assumed here. (If the substrate were disconnected into c components, the analysis applies componentwise: the cyclomatic number becomes $\mu = |E| - |V| + c$ and the boundary rank rank $\partial_1 = |V| - c$, so the rank-deficit formula of Theorem 3.3 is unchanged in form. This contingency is recorded but not pursued.)

Write

$|V|$ = number of hubs, $|E|$ = number of transport faces, $|F|$ = number of plaquettes.

These three integers are the cell counts, finite for a finite vacuum patch.

Definition 2.1a (The periodic case)

If admissible transport is not finite but doubly or triply periodic under a translation group Λ, the object to which the proven results apply is **the finite quotient complex Γ_vac / Λ**, with cell counts taken per fundamental domain. This substitution is not innocent, and the sequel must make it deliberately rather than by default, for two reasons.

First, the homology of the quotient is not the homology of the infinite periodic complex; it is the homology of a different, compact object. Every statement below — Lemma 3.1, Theorem 3.3, the use of $b_0 = 1$ — is a statement about that quotient, and applies only once the quotient is confirmed to be connected with $H_0 = \mathbb{Z}_7$.

Second, and more dangerously, the quotient of a doubly-periodic planar lattice by its translation group is a closed 2-torus, which has $b_1 = 2$ over any field **by construction of the quotient**, independently of completion. A naive "per fundamental domain" computation would then report T1 — a nonzero sector — for a reason that has nothing to do with admissible completion failing to bound transport cycles, and everything to do with the topology imposed by periodic identification. This would be a false positive, and the two contributions are not separated by any bookkeeping: on the torus the identification homology and the completion-deficit homology live in the *same* H_1, with no canonical splitting available without extra structure.

The remedy is therefore prescriptive, not a matter of careful accounting.

Recommended default — the finite patch. The sequel should work with a genuinely finite vacuum patch wherever possible. On a finite patch no periodic identification is imposed, the false positive cannot arise, and Theorem 3.3 applies directly with $b_0 = 1$. This is the clean route and avoids the ill-posedness entirely.

Contingency — the baseline-subtraction recipe. If the sequel must work on a quotient Γ_vac / Λ, it should not attempt to "separate" the two homologies inside H_1, which is not well-posed. Instead it should subtract a fixed topological baseline: compute $b_1(\Gamma\_vac / \Lambda; \mathbb{Z}_7)$ for the actual complex, compute the baseline b_1 of the underlying closed manifold the quotient is modelled on ($b_1 = 2$ for the 2-torus from a doubly-periodic lattice, $b_1 = 3$ for the 3-torus from a triply-periodic one), and attribute to the $\mathbb{Z}_7$ closure sector only the **excess**

b_1^{closure} = $b_1(\Gamma\_vac / \Lambda; \mathbb{Z}_7) - b_1$^{baseline}.

The closure sector is available iff this excess is positive; the baseline classes are artefacts of identification and carry no closure holonomy. This recipe is well-posed because it subtracts a fixed number fixed by the identification topology alone, not an ill-defined subspace of H_1.

The criterion of this paper is sound on either object. What it does not do, and must not be read as doing, is license treating raw quotient homology as completion-deficit homology; on a quotient the baseline must be subtracted first, and on a finite patch the question does not arise.

Definition 2.2 (Chain groups and boundary operators)

Because $\mathbb{Z}_7$ is prime, $\mathbb{Z}_7$ is a field, and the chain groups are $\mathbb{Z}_7$-vector spaces:

$C_0(\Gamma\_vac) = \mathbb{Z}_7$^{|V|}, $C_1(\Gamma\_vac) = \mathbb{Z}_7$^{|E|}, $C_2(\Gamma\_vac) = \mathbb{Z}_7$^{|F|},

freely generated by hubs, faces, and plaquettes respectively. The cellular boundary operators are the $\mathbb{Z}_7$-linear maps

$\partial_2 : C_2 \to C_1$, $\partial_1 : C_1 \to C_0$,

with ∂_1 sending an oriented face to (head − tail) and ∂_2 sending a plaquette to the signed sum of its boundary faces, the signs being the incidence orientations ε_{P,f} of the closure-connection paper. The defining relation of any chain complex,

$\partial_1 \partial_2 = 0$,

holds by construction: the boundary of a plaquette is a closed loop, hence in ker ∂_1.

The complex is truncated at the 2-skeleton, and this is without loss for the availability question. The first homology $H_1 = \ker \partial_1 / \operatorname{im} \partial_2$ depends only on the segment $C_2 \to C_1 \to C_0$; any 3-cells the substrate might carry act through $\partial_3 : C_3 \to C_2$, which alters ker ∂_2 — hence H_2 — but leaves ker ∂_1 and im ∂_2 untouched, so H_1 is unaffected. Whatever the intrinsic dimension of the substrate, availability of the $\mathbb{Z}_7$ sector is a property of the 2-skeleton alone.

Working over the field has one immediate payoff. All the homology data are dimensions of vector spaces — ranks and nullities — with no torsion to track inside the computation and no Ext term in the universal-coefficient theorem. The only place torsion re-enters is the comparison with integral homology, addressed in §4.3.

---

3. Homology and the Rank Formula

The homology groups are, as usual,

$H_0 = C_0 / \operatorname{im} \partial_1$, $H_1 = \ker \partial_1 / \operatorname{im} \partial_2$, $H_2 = \ker \partial_2$,

and we write $b_i = \dim_{\mathbb{Z}_7} H_i$ for the $\mathbb{Z}_7$-Betti numbers. The whole topological question is the size of H_1, and over a field this is forced by counting.

Lemma 3.1 (Boundary ranks of a connected complex). [proven, standard cellular homology]

For Γ_{vac} connected,

$\operatorname{rank} \partial_1 = |V| - 1$.

Proof. H_0 measures connected components: $\dim H_0 = \dim C_0 - \operatorname{rank} \partial_1 = |V| - \operatorname{rank} \partial_1$. A connected complex has $H_0(\Gamma_{vac}; \mathbb{Z}_7) = \mathbb{Z}_7$, so $\dim H_0 = 1$, giving $\operatorname{rank} \partial_1 = |V| - 1$. Connectedness of the complex is determined by its 0- and 1-cells alone, so this uses only the transport graph, not the plaquettes.

Definition 3.2 (Cyclomatic number)

The **cyclomatic number** of the transport graph is

$\mu = |E| - |V| + 1$.

It is the first Betti number of the 1-skeleton — the number of independent loops in the transport graph before any plaquette is attached. Equivalently $\mu = \dim \ker \partial_1$, the dimension of the cycle space Z_1.

The identification is immediate from Lemma 3.1:

$\dim Z_1 = \dim \ker \partial_1 = |E| - \operatorname{rank} \partial_1 = |E| - (|V| - 1) = \mu$.

Theorem 3.3 (First Betti number as a rank deficit). [proven, rank-nullity]

For Γ_vac connected,

$\dim_{\mathbb{Z}_7} H_1(\Gamma\_vac; \mathbb{Z}_7) = \mu - \operatorname{rank}_{\mathbb{Z}_7} \partial_2$.

Proof. By definition $\dim H_1 = \dim \ker \partial_1 - \dim \operatorname{im} \partial_2$. The first term is μ by Definition 3.2; the second is $\operatorname{rank} \partial_2$. Hence $\dim H_1 = \mu - \operatorname{rank} \partial_2$. Since $\dim H_1 \geq 0$, the identity also records the constraint $\operatorname{rank} \partial_2 \leq \mu$ — completion cannot bound more independent loops than the graph contains, which is the statement that plaquette boundaries are themselves cycles, $\operatorname{im} \partial_2 \subseteq Z_1$.

The content of Theorem 3.3 is that H_1 is a *deficit*: the loops of the transport graph (μ of them, independent) minus the loops that completion bounds ($\operatorname{rank} \partial_2$ of them, independent). What survives in homology is exactly the loops left open. This is the precise sense in which the plaquettes act as relators on the cycle space:

$H_1(\Gamma\_vac; \mathbb{Z}_7) = Z_1 / \operatorname{im} \partial_2 = \mathbb{Z}_7^\mu / \langle \text{plaquette boundaries} \rangle$.

A nonzero H_1 means the relators fail to span the cycles.

Remark 3.4 (The second homology is a separate sector)

The remaining homology is $H_2 = \ker \partial_2$, of dimension

$\dim H_2 = |F| - \operatorname{rank} \partial_2$.

This is the space of closed 2-chains — combinations of plaquettes with no boundary, i.e. closed surfaces in the transport complex. It is logically independent of the $\mathbb{Z}_7$ holonomy sector, which lives in H^1 and is dual to H_1. We record H_2 for completeness and for the Euler-characteristic cross-check of §4.2, but it plays no role in the availability of κ.

---

4. The Availability Criterion

4.1 The criterion

We say a nontrivial $\mathbb{Z}_7$ closure sector is **topologically available** if there exists a nonzero class $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, i.e. if $H^1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$.

Theorem 4.1 (Cyclomatic Availability Criterion). [proven, given Theorem 3.3 and universal coefficients]

The following are equivalent:

1. a nontrivial $\mathbb{Z}_7$ closure sector is topologically available, $H^1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$;
2. $H_1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$;
3. $\text{rank}\_\{\mathbb{Z}_7\}\ \partial_2 < \mu$;
4. the plaquette boundaries do not span the cycle space Z_1 of the transport graph.

Proof. (1) ⟺ (2): the universal-coefficient theorem over the field $\mathbb{Z}_7$ gives

$$H^1(\Gamma\_vac; \mathbb{Z}_7) \cong \text{Hom}(H_1(\Gamma\_vac; \mathbb{Z}_7), \mathbb{Z}_7),$$

with no Ext term. The operative reason is that $\mathbb{Z}_7$ is a field: every $\mathbb{Z}_7$-module is projective, so $\text{Ext}^1\_\{\mathbb{Z}_7\}(-, \mathbb{Z}_7)$ vanishes identically, and cohomology is the literal dual of homology. (Freeness of the chain groups is true but is not what kills the Ext term here.) A vector space is nonzero iff its dual is nonzero, so $H^1 \neq 0 \Leftrightarrow H_1 \neq 0$.

(2) ⟺ (3): by Theorem 3.3, $\dim H_1 = \mu - \text{rank}\ \partial_2$, which is nonzero exactly when $\text{rank}\ \partial_2 < \mu$.

(3) ⟺ (4): $\text{rank}\ \partial_2 = \dim \text{im}\ \partial_2$ is the dimension of the span of the plaquette boundaries inside Z_1; since $\dim Z_1 = \mu$, that span is proper exactly when $\text{rank}\ \partial_2 < \mu$.

The equivalence (1) ⟺ (2) is the only step the prior paper had reached, and on its own it is empty — true for any complex. The strengthening is the chain through (3) and (4), which converts the question into the comparison of two computable integers and identifies each with a structural quantity: μ counts the loop content generated by admissible transport, $\text{rank}\ \partial_2$ measures how much of that content admissible completion bounds.

4.2 Euler-characteristic cross-check

The cell counts give the Euler characteristic two independent ways, which must agree — a consistency check on any explicit cell structure the sequel proposes.

Combinatorially,

$$\chi(\Gamma\_vac) = |V| - |E| + |F|.$$

Homologically,

$$\chi(\Gamma\_vac) = b_0 - b_1 + b_2 = 1 - (\mu - \text{rank}\ \partial_2) + (|F| - \text{rank}\ \partial_2) = 1 - \mu + |F|.$$

Substituting $\mu = |E| - |V| + 1$ returns $|V| - |E| + |F|$, so the two expressions coincide identically. The practical use is the rearrangement

$b_1 = 1 + b_2 - \chi$,

which lets the sequel obtain b_1 from χ (pure counting) together with $b_2 = \dim\ker\partial_2$. Combined with Theorem 3.3, the two routes to b_1 — via rank ∂_2 directly, and via χ and b_2 — must give the same integer, catching arithmetic errors in the boundary matrices. Note that the b_2 here is the second Betti number *of the truncated 2-skeleton*, $b_2 = \dim\ker\partial_2$, and not the substrate's intrinsic H_2 (which, with 3-cells present, would be $\ker\partial_2 / \operatorname{im}\partial_3$). The cross-check is internally consistent because every term in it — χ, b_0, b_1, b_2 — is computed on the same 2-skeleton; b_2 should not be read as a substrate invariant.

4.3 The torsion mechanism: a $\mathbb{Z}_7$ sector invisible to rational topology

Working over $\mathbb{Z}_7$ rather than $\mathbb{Q}$ is not a convenience; it is the physically correct choice, since the transport fibre is $\mathbb{Z}_7$. It also changes what "nontrivial topology" can mean, and the difference is structural enough to state as a proposition rather than a remark, because it supplies a mechanism by which the entire closure sector could exist even where ordinary topology appears trivial.

Proposition 4.2 (Torsion availability). [proven, universal coefficients]

For Γ_vac a connected CW complex,

$H_1(\Gamma\_vac; \mathbb{Z}_7) \cong H_1(\Gamma\_vac; \mathbb{Z}) \otimes \mathbb{Z}_7$,

with no Tor correction in degree 1. Consequently

$b_1(\mathbb{Z}_7) = \operatorname{rank}\_\mathbb{Z}\ H_1(\Gamma\_vac; \mathbb{Z}) + t_7$,

where t_7 is the **7-rank of the integral torsion** — the number of cyclic summands of order *divisible by* 7, equivalently the count of 7-primary summands $\mathbb{Z}/7^k$ with $k \geq 1$, equivalently $\dim\_\{\mathbb{Z}_7\}(H_1(\Gamma\_vac; \mathbb{Z})\_tors \otimes \mathbb{Z}_7)$. Each such summand contributes exactly one $\mathbb{Z}_7$ dimension. In particular the $\mathbb{Z}_7$ closure sector can be nonzero, $b_1(\mathbb{Z}_7) > 0$, even when the rational first homology vanishes, $\operatorname{rank}\_\mathbb{Z}\ H_1 = 0$ — namely whenever $t_7 > 0$.

Proof. The universal-coefficient theorem for homology gives the natural short exact sequence

$0 \to H_1(\Gamma\_vac; \mathbb{Z}) \otimes \mathbb{Z}_7 \to H_1(\Gamma\_vac; \mathbb{Z}_7) \to \operatorname{Tor}(H_0(\Gamma\_vac; \mathbb{Z}), \mathbb{Z}_7) \to 0$.

For a connected complex $H_0(\Gamma\_vac; \mathbb{Z}) = \mathbb{Z}$ is free, so the Tor term vanishes and the first map is an isomorphism. Decomposing the integral homology as $H_1(\Gamma\_vac; \mathbb{Z}) = \mathbb{Z}^\{r\} \oplus$ (torsion) with $r = \operatorname{rank}\_\mathbb{Z}\ H_1$, tensoring with $\mathbb{Z}_7$ preserves the free part and sends each cyclic torsion summand $\mathbb{Z}/m$ to $\mathbb{Z}/m \otimes \mathbb{Z}_7 = \mathbb{Z}/\gcd(m,7)$, which is $\mathbb{Z}_7$ when $7 \mid m$ and 0 otherwise. Crucially, the test is divisibility by 7, not the presence of a $\mathbb{Z}/7$ factor: a $\mathbb{Z}/49$ summand has $\gcd(49,7) = 7$, so $\mathbb{Z}/49 \otimes \mathbb{Z}_7 = \mathbb{Z}/7 \neq 0$ and contributes one $\mathbb{Z}_7$ dimension, even though it contains no $\mathbb{Z}/7$ direct summand. Counting one dimension per torsion summand of order divisible by 7 gives exactly t_7, so $b_1(\mathbb{Z}_7) = r + t_7$, positive whenever $t_7 > 0$ regardless of r.

The free rank r is the ordinary (rational) loop count — the number of geometrically open loops a rational eye would see. The term t_7 is invisible to that eye: it is generated not by an open loop but by a loop that closes only after seven traversals, an order-7 cycle in the integral homology. **This is the precise mechanism by which the $\mathbb{Z}_7$ closure sector can be carried entirely by 7-torsion in a complex that is rationally acyclic in degree 1.** Given that the entire substrate architecture is built on $\mathbb{Z}_7$ closure — order-7 hubs, the $\mathbb{Z}_7$ kernel, sevenfold completion — the appearance of integral 7-torsion in Γ_vac is not a curiosity to be hoped against but a natural structural possibility, and it is exactly the contribution the rational picture would miss.

Two consequences for the sequel. First, Theorem 3.3 already incorporates the torsion correctly: computing rank_{$\mathbb{Z}_7$} ∂_2 over the field from the outset folds both r and t_7 into the single $\mathbb{Z}_7$-rank deficit $\mu - \text{rank } \partial_2$, so no separate torsion bookkeeping is needed provided the boundary matrices are taken over $\mathbb{Z}_7$ and not over $\mathbb{Q}$ or $\mathbb{Z}$. Second, and as a warning: availability must not be checked rationally. A computation of rank_$\mathbb{Q}$ ∂_2 that returned "full rank, $H_1(\mathbb{Q}) = 0$" would license the false conclusion that the sector is empty, when a $\mathbb{Z}_7$ computation on the same complex could return $b_1(\mathbb{Z}_7) > 0$ through torsion. The field of coefficients is part of the criterion, not a detail of its evaluation.

---

5. Worked Examples: The Criterion on Realizable Cell Data

The criterion of Theorem 4.1 is a formula relating three integers; it carries conviction only once one has seen it return both verdicts on cell data one can actually draw. This section does that. Two small complexes are exhibited, one giving $H_1 = 0$ and one giving $H_1 \cong \mathbb{Z}_7$, with every rank computed explicitly.

A guard first, because the examples must not reintroduce at the level of illustration the circularity they are meant to dispel. **The complexes below are not claims about the vacuum.** Their plaquettes are not asserted to be admissible closure cycles; their hubs are not order-7 closure hubs. They are abstract cellular 2-complexes, chosen minimal, on which the criterion is evaluated to demonstrate that it is well-defined and discriminating. Which complex — if either — resembles Γ_vac is precisely the question the sequel's derivation must answer. The examples show that the machinery has content; they do not pre-judge its application.

5.1 Example T0 — the filled triangle

Take three hubs v_1, v_2, v_3; three oriented transport faces

$e_1 = (v_1 \to v_2), e_2 = (v_2 \to v_3), e_3 = (v_3 \to v_1)$;

and one plaquette P whose boundary is the 3-cycle, $\partial_2 P = e_1 + e_2 + e_3$.

Counting: $|V| = 3, |E| = 3, |F| = 1$, so

$\mu = |E| - |V| + 1 = 3 - 3 + 1 = 1$.

The 1-skeleton is a single triangle, with exactly one independent loop — the cycle $e_1 + e_2 + e_3$ — so $\dim Z_1 = 1$, consistent with $\mu = 1$.

The boundary operator $\partial_2 : \mathbb{Z}_7^1 \to \mathbb{Z}_7^3$ has the single column $(1, 1, 1)^T$, which is nonzero, so

$\mathrm{rank}\_\{\mathbb{Z}_7\}\ \partial_2 = 1$.

By Theorem 3.3,

$\dim\_\{\mathbb{Z}_7\}\ H_1 = \mu - \mathrm{rank}\ \partial_2 = 1 - 1 = 0$,

so $H_1(\Gamma\_vac; \mathbb{Z}_7) = 0$ and $H^1 = 0$. The one loop of the graph is exactly the one the plaquette bounds; nothing survives. This is **Branch T0**: completion is loop-complete, and on a complex of this shape the closure sector would be unavailable. The Euler cross-check agrees: $\chi = 3 - 3 + 1 = 1 = b_0 - b_1 + b_2 = 1 - 0 + 0$. ✓

5.2 Example T1 — the diagonal-split square, one triangle filled

Take four hubs v_1, v_2, v_3, v_4 on a square, with four boundary faces and one diagonal:

$e_1 = (v_1 \to v_2)$, $e_2 = (v_2 \to v_3)$, $e_3 = (v_3 \to v_4)$, $e_4 = (v_4 \to v_1)$, $e_5 = (v_2 \to v_4)$ (the diagonal).

The diagonal splits the square into two triangles: the lower triangle (v_1, v_2, v_4) with loop $e_1 + e_5 + e_4$, and the upper triangle (v_2, v_3, v_4) with loop $e_2 + e_3 - e_5$. Attach a single plaquette P filling only the upper triangle:

$\partial_2 P = e_2 + e_3 - e_5$.

Counting: $|V| = 4$, $|E| = 5$, $|F| = 1$, so

$\mu = |E| - |V| + 1 = 5 - 4 + 1 = 2$.

The 1-skeleton has two independent loops — one per triangle — so $\dim Z_1 = 2$, consistent with $\mu = 2$. The boundary operator $\partial_2 : \mathbb{Z}_7^1 \to \mathbb{Z}_7^5$ has the single column $(0, 1, 1, 0, -1)^T$, nonzero, so

$\mathrm{rank}\_\{\mathbb{Z}_7\}\ \partial_2 = 1$.

By Theorem 3.3,

$\dim\_\{\mathbb{Z}_7\}\ H_1 = \mu - \mathrm{rank}\ \partial_2 = 2 - 1 = 1$,

so

$H_1(\Gamma\_vac; \mathbb{Z}_7) \cong \mathbb{Z}_7$.

The surviving generator is the lower-triangle loop $e_1 + e_5 + e_4$, which no plaquette bounds: completion has filled one of the two independent loops and left the other open. This is **Branch T1**: completion is loop-deficient, a single $\mathbb{Z}_7$ class is topologically available, and on a complex of this shape the occupancy question would become live. The Euler cross-check agrees: $\chi = 4 - 5 + 1 = 0 = 1 - 1 + 0$. ✓

5.3 What the examples establish

The two complexes differ by one hub, two faces, and the placement of a single plaquette, yet they sit on opposite sides of the fork. This is the point. The verdict is not a property of the criterion — which is fixed — but of the cell data fed to it, and small changes in how completion bounds the available loops flip the outcome between T0 and T1. The criterion is therefore both well-defined (it returns a definite integer on any explicit complex) and discriminating (it returns different integers on different complexes). Neither property was visible while ∂_2 remained an abstract symbol. The examples also exhibit the Euler cross-check doing its job: in both cases the two independent routes to b_1 agree, which is the validation any explicit Γ_vac computation must pass before its rank is trusted.

What the examples do *not* establish is which side Γ_vac falls on. That requires the actual completion cycles, and obtaining those is the content of the sequel — to which §7 now turns.

---

6. The Topological Fork

Theorem 4.1 splits the programme cleanly. Both branches are **open** pending the cell computation of §7, but each is now stated as a definite arithmetic condition rather than an abstract dichotomy, and §5 has exhibited each on an explicit complex.

Branch T0 — completion is loop-complete. [conditional on the cell computation]

$\mathrm{rank}_{\mathbb{Z}_7}\, \partial_2 = \mu$, equivalently $H_1(\Gamma\_vac; \mathbb{Z}_7) = 0$.

The plaquette boundaries span the entire cycle space: every independent loop of the transport graph is bounded by admissible completion. Then $H^1(\Gamma\_vac; \mathbb{Z}_7) = 0$, no nonzero class exists, and by Proposition 6.1 of the closure-connection paper the closure channel collapses to Branch A — the $\mathbb{Z}_7$ structure is present algebraically but carries no realised holonomy. **In this branch the occupancy problem is settled negatively without any dynamical computation**, and the Gate-3 verdict is decided here. Example 5.1 is the minimal instance.

Branch T1 — completion is loop-deficient. [conditional on the cell computation]

$\mathrm{rank}_{\mathbb{Z}_7}\, \partial_2 < \mu$, equivalently $H_1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$,

with

$\dim_{\mathbb{Z}_7} H_1 = \mu - \operatorname{rank} \partial_2 > 0$

independent non-bounding loops surviving. Then nontrivial classes are topologically available, condition (2) of Proposition 6.1 holds, and the occupancy question becomes meaningful: admissibility must be examined to determine whether it populates one of the available classes (Branch B) or annihilates all of them (Branch A by dynamics rather than by topology). Only in this branch does the expensive semantic bridge of the occupancy paper have to be built. Example 5.2 is the minimal instance. Note that by Proposition 4.2 this branch can be entered through integral 7-torsion alone, with no rationally visible open loop — a route the toy examples, being torsion-free, do not display but which Γ_{vac} may.

The fork is therefore not symmetric in cost. T0 decides Gate 3 here — closing the $\mathbb{Z}_7$ closure-holonomy sector via Branch A, conditional on Proposition 6.1 of the closure-connection paper; T1 hands the question on to the occupancy paper. This asymmetry is the reason to compute the topology first.

6.1 Why loop-deficiency is the structurally less restrictive outcome

Theorem 4.1 is neutral between T0 and T1. Nevertheless the architecture of the transport programme gives a structural reason to regard T1 as the natural outcome and T0 as the exceptional one — though the reason is heuristic, and the heuristic must first be stated against the right invariant.

The starting observation is exact but light. Since $\operatorname{im} \partial_2 \subseteq Z_1$ with $\dim Z_1 = \mu$, one has $\operatorname{rank} \partial_2 \leq \mu$ always; T0 is the single maximal value $\operatorname{rank} \partial_2 = \mu$ (corank zero, every transport loop bounded), while T1 is any corank ≥ 1 (one loop surviving suffices). Maximality is demanded on one side, a single failure on the other. That is scene-setting, not yet an argument: completion is a structured map, not a random one, and a structured map can hit full rank deliberately, so the bare count of values carries no weight on its own.

The weight comes from identifying what actually controls the fork, and here a tempting reading must be corrected. One might say "local completion versus global connectivity" — completion acting through minimal closure cycles while the transport graph carries larger cycles local cells cannot reach. That is imprecise, because locality does not by itself force T1: dense overlapping plaquettes routinely span large cycles, and a finely subdivided disk has $H_1 = 0$ however large it grows, every long loop being a $\mathbb{Z}_7$-combination of small plaquette boundaries. The true determinant is not the scale of completion but the global shape of Γ_{vac} — whether it is **acyclic**, every cycle bounding (the tiled-disk case, T0), or carries **noncontractible cycles**, handles or holes no collection of plaquettes can fill (the torus-or-annulus case, T1). Noncontractibility, not locality, is what makes a cycle survive.

With the determinant correctly identified, the heuristic has two independent legs, and is stronger for it.

Geometric route (weak, possibly reversed). If transport connectivity introduces any noncontractible loop — any cycle the admissible closure cells do not collectively bound — then

T1 holds. One might expect such a loop to survive in a complex whose connectivity is generated globally while its 2-cells are generated by purely local closure. But this expectation is *not* robust against the structural picture of §6.2, and the honest reading is that the geometric leg may point the wrong way. The §6.2 constraints — sparse, locally generated, closure-origin hubs, local plaquettes — are entirely consistent with a finite vacuum patch that is disk-like and **simply connected**, and on such a patch there are no noncontractible loops at all: the geometric leg then yields T0, not T1. So "survival of at least one noncontractible loop is the less demanding expectation" is true only if the patch already has handles or holes, and false for a contractible patch — which is a live possibility for exactly the kind of complex §6.2 anticipates. The geometric leg is therefore weak, and weak in a direction that could favour T0.

Arithmetic route (robust). The torsion leg does not share this weakness, and it is where the weight of the heuristic properly rests. By Proposition 4.2, integral 7-torsion in H_1 — a loop that closes only after seven traversals — contributes to $b_1(\mathbb{Z}_7)$ even when the complex is rationally acyclic and carries no geometrically open loop, and even when it is simply connected in the rational sense. It survives precisely the disk-like patch that defeats the geometric leg. Given a substrate built throughout on order-7 closure — order-7 hubs, the $\mathbb{Z}_7$ kernel, sevenfold completion — integral 7-torsion is the structurally motivated route, not an incidental one, and it lands in the $\mathbb{Z}_7$ sector by arithmetic rather than by any failure of completion to span. The two legs are therefore not co-equal: the geometric leg is contingent on global shape the programme does not guarantee and may well lack, while the arithmetic leg is robust against the §6.2 expectations and is tied directly to the $\mathbb{Z}_7$ architecture. To the extent the heuristic favours T1 at all, it does so through torsion.

We emphasise that neither leg is a proof, and that T0 is correspondingly not one condition but two: it requires Γ_vac to be *both* rationally acyclic in degree 1 *and* free of integral 7-torsion in degree 1. The phrase "loop-complete" hides this conjunction, and the heuristic should not be over-read on the strength of the geometric leg — which, as just noted, may not even point toward T1 on the patch §6.2 expects. The purpose of the observation is only to locate the burden of computation: T0 demands that completion both bound every transport cycle and leave no order-7 cycle behind, whereas T1 needs a single survivor — and on a simply-connected patch that survivor can only be a torsion class. Accordingly the sequel should be read not merely as a calculation of rank ∂_2, but as a determination of whether admissible completion is **globally loop-complete** in this two-fold sense, with the integral 7-torsion of Γ_vac the contribution most likely to decide it. The topological availability of the $\mathbb{Z}_7$ sector hinges precisely on that question.

6.2 Structural expectations for Γ_vac

The analysis above treats Γ_vac abstractly, as it must: the topological reduction depends only on incidence structure, not on the physical reading of the cells. But the transport programme already constrains what the eventual complex can be, and recording those constraints sharpens what the sequel is expected to derive. Nothing in this subsection is used in any proof above; it is heuristic scene-setting, marked as such.

Four constraints come from the programme architecture.

Closure origin of the 0-cells. The fundamental K = 7 closure unit consists of six boundary relations returning through a seventh completion relation. Transport begins not from arbitrary vertices but from closure hubs with a fixed local combinatorial structure, and the 0-cells of Γ_vac inherit that origin.

Locality of the 1-skeleton. The transport-group analysis showed that transport occurs only through admissible neighbouring closure relations, and that single-face transport carries no gauge-invariant closure content. The 1-skeleton is therefore expected to be generated by local admissible transport between neighbouring hubs, not to be a dense or arbitrary graph.

Locality of the 2-cells. Completion generates plaquettes by enforcing minimal admissible closure cycles. These are expected to be local objects, arising from local closure consistency rather than large-scale constraint. The paper has assumed no specific plaquette shape, and does not do so now.

Conservation and competition as cell selection. These rules are expected to remove transport configurations that violate admissibility, reducing the allowed cells and incidences. It is worth being precise about how this touches availability, since §1.1 and §7 state that availability does not depend on the rules' action on the rotation labels. Under hypothesis (CA) of §1.1 there is no tension: conservation and competition enter the *topology* only through **cell selection** — determining which hubs, faces, and plaquettes are admissible, hence what Γ_vac *is* — and this selection is, by (CA), a combinatorial predicate that does not consult the labels ρ_f, whose dynamics are reserved for the flatness and population conditions. The resulting complex is then a constrained admissible subcomplex fixed by the closure rules at the level of which cells exist. The caveat from §1.1 applies in full force here: if cell-admissibility for these two rules secretly referenced a label — if, say, a plaquette were admissible only for certain ρ_f on its boundary — then (CA) would fail precisely at conservation and competition, and the separation would leak exactly at the rules this paragraph is about. That cell-admissibility under conservation and competition is genuinely label-free is therefore the specific instance of (CA) the sequel most needs to verify by construction.

Taken together these suggest Γ_vac is a sparse, locally generated transport complex rather than a densely triangulated one. This is an expectation about the derivation, not a theorem, and it must be read against the determinant established in §6.1: the fork is decided by noncontractibility (and torsion), not by density as such — recall that a *dense* subdivided disk is still T0. Sparsity is therefore not a direct route to T1. What sparsity does control is a cell count, and there the intuition has an exact and provable core.

Proposition 6.2 (Cell-count sufficient condition for T1). [proven, rank bound]

If the number of plaquettes is strictly less than the cyclomatic number,

$|F| < \mu = |E| - |V| + 1$ (connected case),

then $H_1(\Gamma\_vac; \mathbb{Z}_7) \neq 0$ and Branch T1 holds. Equivalently, $|F| \geq \mu$ is a *necessary* condition for T0.

Proof. The plaquette boundaries number $|F|$, so their span has dimension $\operatorname{rank} \partial_2 \le |F|$. If $|F| < \mu$ then $\operatorname{rank} \partial_2 \le |F| < \mu$, whence by Theorem 3.3 $\dim_{\mathbb{Z}_7} H_1 = \mu - \operatorname{rank} \partial_2 \ge \mu - |F| > 0$. The contrapositive gives the necessary condition for T0. The bound is over $\mathbb{Z}_7$ and so already accounts for any 7-torsion contribution; no rational computation is involved.

This is the rigorous residue of the sparsity heuristic: too few plaquettes cannot span the cycle space, and the deficit is at least $\mu - |F|$. The converse fails — $|F| \ge \mu$ does not give T0, since the plaquettes may be redundant (several bounding the same cycles, so $\operatorname{rank} \partial_2 < |F|$) or may leave a 7-torsion class unbounded even at full rational rank. So Proposition 6.2 settles the sparse extreme outright and leaves the dense regime to the explicit rank computation, exactly where §6.1 located it.

The significance for the sequel is then sharp rather than vague. If the derived complex turns out to have fewer completion plaquettes than independent transport loops, the $\mathbb{Z}_7$ sector is available with no further computation — Proposition 6.2 closes it. Only if $|F| \ge \mu$ does the verdict require the full $\operatorname{rank} \partial_2$, and only then does the dense-versus-deficient question of §6.1 become live. Either way the controlling object is the derived complex itself: once Γ_vac is known, $|F|$, μ, the matrix ∂_2, its rank, and the verdict follow in order.

---

7. What the Sequel Must Compute

A point of order, to dispel any impression that the sequel merely evaluates a rank on a matrix already in hand. **The matrix is not given; it is the second output of the sequel, not the first.** The primary task of the sequel is the derivation of the admissible vacuum transport complex Γ_vac itself — establishing, from the closure rules, which hubs are admissible, which transport faces connect them, and which minimal cycles hub completion enforces as plaquettes. The boundary operator ∂_2 is then read off that derived cell structure; it does not exist independently of it. The dependency order is

admissibility rules ⟶ derived cell structure of Γ_vac ⟶ boundary operator ∂_2 ⟶ $\operatorname{rank}_{\mathbb{Z}_7} \partial_2$ ⟶ verdict,

and the present paper has supplied only the last two arrows. The first two — the construction of the complex — are the substance of the next paper and are where its difficulty lies.

With that order fixed, Theorem 4.1 reduces topological availability to data of three kinds, all read off the derived cell structure of Γ_vac:

1. **The cell counts** $|V|$, $|E|$, $|F|$ — of a finite vacuum patch, or of the finite quotient $\Gamma_\text{vac} / \Lambda$ in the periodic case, with the quotient caveats of Definition 2.1a observed. These fix $\mu = |E| - |V| + 1$ (connected case) by counting alone.
2. **The boundary operator** $\partial_2 : \mathbb{Z}_7^{|F|} \to \mathbb{Z}_7^{|E|}$ as an explicit matrix over $\mathbb{Z}_7$, its entries the incidence signs $\varepsilon_{P,f}$ of each plaquette against each face — obtained from the derived complex, not posited.

3. **The single integer** $\mathrm{rank}_{\mathbb{Z}_7}\,\partial_2$, computed by Gaussian elimination over $\mathbb{Z}_7$.

The verdict is then immediate:

$\mathrm{rank}_{\mathbb{Z}_7}\,\partial_2 = \mu \Longrightarrow$ T0 (sector unavailable, Gate 3 closed), $\mathrm{rank}_{\mathbb{Z}_7}\,\partial_2 < \mu \Longrightarrow$ T1 (sector available, occupancy live),

and in the T1 case $\dim H_1 = \mu - \mathrm{rank}\,\partial_2$ counts the available classes. The Euler cross-check of §4.2 validates the matrices before the rank is trusted, exactly as it did on the two worked complexes of §5.

This is the entire remaining topological content once the complex is in hand: **one matrix and one rank.** No part of it requires the rotation labels ρ_f, the curvature Ω, or the action of conservation and competition on transport — those belong to the flatness and population conditions, not to availability. The derivation of the complex, and the computation of $\mathrm{rank}\,\partial_2$, is the substance of the next paper and is not attempted here.

---

8. Relationship to the Occupancy Programme

This paper supplies condition (2) of the closure-connection paper's Proposition 6.1 and nothing else. It does not establish flatness (condition 1), does not compute κ, and does not determine whether admissibility populates any class (condition 3). Its place in the chain is precise:

$\Omega = 0 \longrightarrow H^1(\Gamma_{\mathrm{vac}}; \mathbb{Z}_7) \neq 0 \longrightarrow \kappa \neq 0$, (1) (2 — this paper) (3)

with the present paper resolving the middle link up to the single rank computation of §7. The hypothesis levels carry forward as established in the closure-connection paper: availability is a statement about the complex and inherits the boundary-level coherence H-OC(∂) needed only to make the cells well-defined as oriented objects; the class κ that availability makes room for requires the cycle-level H-OC(γ) once one passes to occupancy. Availability itself does not depend on the rotation dynamics at all.

The next paper must then determine, in order: whether admissibility forces flatness; in the T1 case, which of the available classes admissibility populates; and whether the surviving classes survive the further operational reduction $N_{\mathrm{hol}} \longrightarrow N_{\mathrm{obs}}$. Only the conjunction of all three with topological availability decides the fate of the $\mathbb{Z}_7$ closure sector.

Conclusion

The closure-connection paper reduced Gate 3 to a cohomology class $\kappa \in H^1(\Gamma_{\mathrm{vac}}; \mathbb{Z}_7)$ and identified three hierarchical conditions for that class to survive. The present paper resolves the second — topological availability — up to a single explicit computation, and in doing so separates it cleanly from occupancy.

Availability is not a question about the admissibility dynamics. It is a question about the shape of the transport complex, and over the field $\mathbb{Z}_7$ that question is rank arithmetic. The first Betti number is a deficit,

$$\dim_{\mathbb{Z}_7} H_1(\Gamma_{vac}; \mathbb{Z}_7) = \mu - \mathrm{rank}_{\mathbb{Z}_7} \partial_2,$$

the loop content of admissible transport minus the bounding power of admissible completion. The Cyclomatic Availability Criterion follows: the $\mathbb{Z}_7$ closure sector is topologically available if and only if rank $\partial_2 < \mu$ — if and only if completion fails to bound every loop of the transport graph. The universal-coefficient theorem over $\mathbb{Z}_7$ carries the statement from homology to the cohomology where κ lives, with no torsion correction inside the field computation and a favourable torsion contribution when integral 7-torsion is present.

The verdict reduces to one matrix and one rank, and it forks asymmetrically. If completion is loop-complete, rank $\partial_2 = \mu$, the topology is trivial, and Gate 3 closes here — the occupancy problem never arises. If completion is loop-deficient, rank $\partial_2 < \mu$, the sector is available and the occupancy question becomes meaningful for the first time. The criterion is proven; the cell counts that decide which branch obtains are the task of the sequel.

While Theorem 4.1 is neutral between the two branches, the global nature of transport connectivity against the local nature of admissible completion, together with the $\mathbb{Z}_7$-torsion route of Proposition 4.2, makes loop-deficiency the structurally less restrictive outcome: T0 requires completion to bound every transport cycle and leave no order-7 cycle behind, whereas T1 requires only one survivor by either route. Whether this heuristic survives the explicit cell computation — whether admissible completion is in fact globally loop-complete — is the central question of the sequel.

The result establishes the final topological prerequisite for the Gate-3 programme and isolates the one computation on which the availability of the entire $\mathbb{Z}_7$ closure sector depends.