Two approaches to the chain-pigeonhole step: cyclic 2-SAT vs. König lift

Background

The chain-pigeonhole step asks whether, for any adjacent SP tire pair (T_1, T_2) sharing cycle γ of length k, the projection-supports overlap non-trivially:

$$\pi_D(\mathcal{C}(T_1)) \cap \pi_U(\mathcal{C}(T_2)) \neq \emptyset.$$

Empirically (`tire_fiber_step2.tex` + `tire_fiber_step2_large.tex`) the overlap is always non-empty, and in the worst case has exactly 6 elements forming a single S_3-orbit.

Two independent proof attempts now exist in the notes, attacking slightly different statements with very different techniques. This note compares them and assesses which is more promising.

Approach 1: cyclic 2-SAT (`rainbow_proof.tex`)

What it tries to prove

A single-tire characterisation: for an antipodal-chord SP tire T with $m_1 \geq m - 1$,

$$\pi_D(\mathcal{C}(T)) \; = \; \mathcal{P}_m,$$

where $\mathcal{P}_m \subseteq \{1,2,3\}^m$ is the "perms-per-face" set $\left((m/2)!^2 \cdot \binom{3}{m/2}^2 = 36\right.$ elements at both $m \in \{4,6\}$).

Technique

1. **$\subseteq$ direction:** clean from the O-face dual proper-colouring constraint.
2. **Reduce $\supseteq$ to 2-SAT:** For each $\sigma \in \mathcal{P}_m$, define orientation bits $o_0, \ldots, o_{m-1} \in \{0,1\}$ at the D-positions. The proper-colouring constraint on T'_{ann} becomes a cyclic 2-SAT on these orientations.
3. **Open step (Conjecture 1.5):** the cyclic 2-SAT is satisfiable for every $\sigma \in \mathcal{P}_m$. Empirically true (6–18 satisfying orientations per σ at $m = 6$), but a clean structural proof is open.

Status

The $\subseteq$ direction is fully proven. The $\supseteq$ direction reduces to a 2-SAT solvability claim that remains open as Conjecture 1.5; a naive "all-zero orientation" construction fails, and a correct general argument needs implication-graph or S_3-equivariant case analysis.

Approach 2: König lift (`worst_case_proof_sketch.tex`)

What it tries to prove

A two-tire overlap statement: for adjacent SP tires (T_1, T_2) sharing γ with both sides above their saturation thresholds,

$$|\pi_D(\mathcal{C}(T_1)) \cap \pi_U(\mathcal{C}(T_2))| \ \geq\ 6.$$

This is the actual chain-pigeonhole question (no need to characterise π_D as a whole).

Technique

1. **Build a bipartite face-incidence graph G:** left vertices $= T_1$'s γ-face partition $\mathcal{F}_1$; right vertices $= T_2$'s γ-face partition $\mathcal{F}_2$; edges $= \gamma$-edges (each lies in one $\mathcal{F}_1$-face and one $\mathcal{F}_2$-face). Each face has exactly 3 γ-edges, so G is 3-regular bipartite.

2. **Apply König's theorem:** every bipartite graph admits a proper Δ-edge-colouring. G gets a proper 3-edge-colouring $\chi : E(G) \to \{1, 2, 3\}$.

3. **Lift back to γ:** define $\sigma(e) := \chi(e)$ for each γ-edge. Then at every $\mathcal{F}_1$-face (and $\mathcal{F}_2$-face), the three incident edges have three distinct colours, so $\sigma|_F$ is a permutation of $\{1, 2, 3\}$.

4. Hence σ lies in the Latin overlap, and its S_3 orbit (size 6) is in $\pi_D \cap \pi_U$.

What's still open

The construction works directly when *both* tires give a chord on γ (so both $\mathcal{F}_1, \mathcal{F}_2$ are direct γ-face partitions). In the actual chain-pigeonhole setup, T_2's chord is on $B^{(2)}_{\text{in}}$, not on γ.

Open conjecture (Conj. *t2-induces-partition* of the worst-case note): T_2 nonetheless induces a γ-face partition $\widetilde{\mathcal{F}_2}$ such that $\pi_U(\mathcal{C}(T_2)) \supseteq \mathcal{L}(\gamma, \widetilde{\mathcal{F}_2})$ (the Latin subset of γ-colourings compatible with $\widetilde{\mathcal{F}_2}$). A candidate construction exists: group γ-edges by which $B^{(2)}_{\text{in}}$-face their T_2-side annular triangles share an edge with.

Side-by-side comparison

	Approach 1: cyclic 2-SAT	Approach 2: König lift
What it proves	Single-tire: $\pi_D = \mathcal{P}_m$ (a full structural characterisation, $\lvert\pi_D\rvert = 36$).	Two-tire: $\lvert S_1 \cap S_2\rvert \geq 6$ (a lower bound on the overlap, witnessed by the rainbow S_3-orbit).
Hard step	Cyclic 2-SAT solvability for every $\sigma \in \mathcal{P}_m$.	Showing T_2's side induces a γ-face partition when T_2's chord is on $B^{(2)}_{\text{in}}$ rather than γ.
Tooling	Custom orientation-bit machinery; implication-graph analysis on a cyclic 2-SAT.	Classical: König's edge-colouring theorem for bipartite graphs.
Strength of conclusion	Stronger than needed: characterises π_D exactly.	Exactly what's needed for chain pigeonhole.
What's proven	$\subseteq$ direction (Lemma 1.2); 2-SAT reduction (Prop 1.4); sharpness counterexample (Prop 1.7).	König-overlap prop (when both chords on γ); S_3-invariance lower-bound argument.
What's open	2-SAT solvability (Conj 1.5); empirically verified.	"T_2 induces γ-partition" (Conj *t2-induces-partition*); plausibility check via candidate construction.

Which approach is more promising

Approach 2 (König lift) is more promising, for three reasons:

1. The hard step is already proven

König's theorem is a 100-year-old textbook result. The König- overlap proposition is a clean lift via that theorem and is fully written down. The remaining open piece (T$_2 inducesa\gamma$-partition) is a geometric/structural claim about how T_2's annular triangulation distributes $B^{(2)}_{\text{in}}$ faces across γ-edges — likely amenable to a direct construction.

Approach 1's open piece (2-SAT solvability) is a combinatorial satisfiability claim with no obvious leveraged tool. Empirically true, but the structural reason isn't yet clear.

2. Approach 2 proves exactly what we want

Chain pigeonhole asks: is the overlap non-empty? Approach 2 directly addresses this with a lower bound of 6, which is also empirically tight. Approach 1 proves a strictly stronger statement (the full $\pi_D = \mathcal{P}_m$ characterisation) that is more than chain pigeonhole needs. Proving more than necessary is a strictly harder task and unnecessary for the goal.

3. The König lower bound has the right structure

Worst-case overlap is empirically a single S_3-orbit of size 6, which is exactly what the König lift produces (lift one 3-edge-colouring, then act by S_3 on colours). The proof mirror the empirical phenomenon. Approach 1's machinery is more general and doesn't naturally explain why the worst case is an S_3-orbit.

When Approach 1 still pays off

Approach 1 *does* give a stronger result: a complete characterisation of π_D for the antipodal-chord SP tire, including the upper bound $\lvert\pi_D\rvert \leq 36$. This is useful if we ever need finer control over π_D's shape (e.g. if we want to prove that π_D does not contain certain S_3-orbits, or if we want to control the chain-pigeonhole overlap above the floor of 6). For just establishing non-empty overlap, Approach 2 suffices.

Reconciliation: the 6 is the same 6

Both approaches witness the same canonical 6-element worst-case intersection:

- Approach 1: the rainbow S_3-orbit $(a, b, c, b, c, a) \cdot S_3$ sits inside $\pi_D \cap \pi_U$ when $\pi_D = \mathcal{P}_m$ contains it and π_U does too.
- Approach 2: the König-lifted Latin colouring's S_3-orbit sits inside the intersection directly, by construction.

For the case T_1 antipodal-chord SP, T_2 chordless SR, $k = 6$ (the worst tested case), these are literally the same 6 elements.

Recommended next move

Attack the open conjecture in Approach 2 (Conj *t2-induces-partition*):

1. Write down the candidate induced γ-partition $\widetilde{\mathcal{F}_2}$ explicitly (the "group γ-edges by which $B_{\text{in}}^{(2)}$-face's neighbours they share annular edges with" construction).
2. Verify $\pi_U(\mathcal{C}(T_2)) \supseteq \mathcal{L}(\gamma, \widetilde{\mathcal{F}_2})$ computationally for $k \in \{3, 4, 5, 6\}$ and several T_2 structures.
3. If empirical fit is exact (as in the worst-case data), prove inclusion via a transfer-matrix / fibre-lifting argument.

This is more leveraged than continuing to refine the 2-SAT implication-graph analysis, which would prove a stronger statement that we do not need.