Chain pigeonhole on cut tires: a sketch and honest assessment

Setup

Take G' the cubic planar dual of a maximal planar graph G. Suppose G' is a minimum counterexample to the 4-colour theorem — no proper 3-edge-colouring of G' exists, but every smaller cubic planar graph admits one.

By cyclic edge-connectivity (G internally 6-connected implies G' cyclically 6-edge-connected), pick a 6-edge cut $C \subseteq E(G')$ partitioning $V(G')$ into S and $V \setminus S$, both non-trivial. Form G'_0 and G'_1 as in `cut_depth_label.tex` by removing C and attaching pendant edges at degree-2 vertices.

For the cleanest setting, assume C is a *matching cut* (each boundary vertex on each side has exactly 1 cut edge, so each side attaches 6 pendants). Apply the BFS depth labelling: pendants get depth 0; each edge adjacent (sharing a vertex) to a depth-d edge gets depth $d+1$.

For each $d > 0$, the cut tires at (d, f) (Definition in `cut_depth_label.tex`) layer G'_i concentrically around the cut.

The argument, step by step

(Step 1) Reduction by minimality. Each G'_i has $|S_i|+6 < |V(G')|$ vertices (assuming $|V \setminus S_i| > 6$, which holds in the matching-cut case for any non-degenerate cut). By minimality of G', each G'_i is properly 3-edge-colourable. Let $\chi_i : E(G'_i) \to \{1, 2, 3\}$ be any such colouring.

(Step 2) The induced cut configuration. The 6 depth-0 pendant edges of G'_i correspond bijectively to the 6 cut edges of C (each pendant replaces a cut edge). Let $\sigma_i := \chi_i|_{\text{depth-0 edges of } G'_i} \in \{1, 2, 3\}^6$, indexed by the cut edges. This is the *boundary configuration* at the cut.

(Step 3) Gluing. A proper 3-edge-colouring of G' exists iff there exists a single colouring of C that extends to both sides, i.e. iff some $\sigma \in \{1, 2, 3\}^6$ is achievable as both σ_0 for some χ_0 and σ_1 for some χ_1. Let

$$\mathcal{R}_i := \{\sigma_i : \chi_i \text{ a proper edge 3-colouring of } G'_i\}.$$

Then G' has a proper 3-edge-colouring iff $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$. Assuming G' is a counterexample, this intersection is empty.

(Step 4) Layered description of $\mathcal{R}_i$ via cut tires. Each cut tire $T_d^{(i,f)}$ has its own "ring projection" constraints. Define:

- $\pi_{\text{in}}(T_d^{(i,f)})$: the projection of χ_i onto the depth-$(d-1)$ inner spokes of the cut tire. For $d = 1$, this is exactly σ_i restricted to those pendants whose boundary vertex sits on f's boundary.

- $\pi_{\text{out}}(T_d^{(i,f)})$: projection onto depth-$(d+1)$ outer spokes.

Adjacent cut tires share layers: outer spokes of T_d are inner spokes of T_{d+1} (when their faces are appropriately adjacent in the embedding). So the chain of cut tires at depths $1, 2, \ldots, d_{\max}$ has consistency constraints $\pi_{\text{out}}(T_d) = \pi_{\text{in}}(T_{d+1})$ along each chain.

(Step 5) Chain pigeonhole at the cut. $\mathcal{R}_i$ is determined by the chain of cut tires on side i: a σ_i is achievable iff there exists a consistent sequence of cut-tire colourings from the deep interior outward to the cut that projects to σ_i on the depth-0 pendants.

Chain pigeonhole says: if at each step d, the cut tire's achievable inner-projection set $\pi_{\text{in}}(T_d)$ is sufficiently large (containing enough S_3-orbits), then $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$, contradicting that G' is a counterexample.

What this needs to be a proof

The argument above sketches the *shape* of a proof. The non-trivial parts are:

(a) Chain consistency: well-definedness of the chain

For the chain $T_1 \to T_2 \to \cdots \to T_{d_{\max}}$ to be well-defined, each T_d must have ≥ 1 face, and adjacent tires must share layers cleanly. Obstacles:

- The depth-d subgraph H_d may have no faces (if H_d is a tree or empty). The empirical example (`cut_depth_label.tex`) shows H_d has 1–3 faces at each depth $1 \leq d \leq 7$, but this is not guaranteed in general.
- Multiple faces at the same depth mean the chain forks; chain pigeonhole becomes a tree-pigeonhole.
- Face boundary walks need not be simple cycles — they can revisit vertices, as in the existing tire definition's treatment of cut-vertices.

(b) Quantitative chain pigeonhole

Even with a clean chain, we need a quantitative argument that $|\pi_{\text{in}}(T_d)|$ is large enough at each step to force non-empty intersection with the adjacent tire. This is the *same* chain pigeonhole question studied in `rainbow_proof.tex` and `worst_case_proof_sketch.tex`, now applied to the cut-tire chain.

The two open conjectures that would close this step are:

- **Rainbow conjecture (`rainbow_proof.tex`, Conj 1.5):** for the antipodal-chord SP case, the inner-spoke projection support equals the perms-per-half set $\mathcal{P}_m$. For cut tires this would mean each cut tire's π_{in} saturates a known S_3-symmetric set.
- **König-lift conjecture (`worst_case_proof_sketch.tex`, Conj *t2-induces-partition*):** adjacent tires induce γ-face partitions whose König lifts give a Latin intersection of size ≥ 6. The face-pair-connection refinement (`k9_surviving_partitions.tex`) corrects the naive candidate partition.

Neither is fully proved. See `two_approaches_comparison.tex` for the comparison.

(c) Cut-tire-specific issues

Cut tires differ from the tires of `paper.tex` in important ways:

- Cut tires are derived from the *depth labelling* on G'_i, not from a primal level structure on G. The correspondence to primal tires (Defs 1.15–1.17) is by analogy, not by direct identification.
- The depth-d subgraph H_d is generally *not* cubic and may not even be connected. Its faces may behave differently from G''s faces.
- The cut tire's T'_{ann}-analogue (= face boundary of f in H_d) is a closed walk in H_d, not a cycle of G'. Its structure depends on the depth labelling.

So even with the rainbow/König conjectures proved for primal tires, their transfer to cut tires requires verification.

Empirical check: chain length and tire structure

For the example tire chain on G'_1 of Holton-McKay #0 ($d = 1, \ldots, 7$):

d	# faces in H_d	largest face length	inner spokes	outer spokes
1	2	12	5	4
2	2	7	4	3
3	3	2	2	2
4	2	8	2	5
5	2	14	4	6
6	1	12	7	1

Observations:

- Chain length ≤ 7 in this example.
- Face counts and sizes vary irregularly — depth 3 has three small (2-edge) faces; depth 5 has a face of length 14.
- Total inner + outer spokes at each depth ranges from 4 to 10.

This irregularity is the structural obstacle to clean chain pigeonhole: a uniform bound on $|\pi_{\text{in}}|$ across depths seems unlikely. The actual chain pigeonhole would need to handle depth-by-depth structure.

Net assessment

The chain pigeonhole argument on cut tires is *structurally sound but technically open.* It:

- Gives a clean reformulation of the 4CT reducibility question in terms of cut-derived layered structure.
- Maps directly onto the existing `paper.tex` tire framework via the depth-as-distance-to-cut analogy.
- Inherits all the open conjectures from `rainbow_proof.tex` and `worst_case_proof_sketch.tex` (chain pigeonhole at each layer, intersection non-emptiness, etc.).
- Adds new technical issues specific to the cut-tire setting (depth subgraphs H_d may degenerate, irregular face structure across depths).

Concrete next steps.

1. Verify that the cut-tire chain is well-defined on the 6 Holton-McKay graphs and a few other test cases (each G'_i has cut tires at every depth d in some range; no empty H_d).

2. Compute $\pi_{\text{in}}(T_d^{(i)})$ for each cut tire and check the rainbow S_3-orbit appears at the cut layer ($d = 1$).

3. Check pairwise compatibility at the cut between G'_0 and G'_1: do $\mathcal{R}_0$ and $\mathcal{R}_1$ overlap? If empirically yes for all 6 Holton-McKay graphs, that's evidence; if no for some, that's a falsification of the chain argument as currently stated.

Step 3 is the cleanest empirical test — it's just an extension of the step-2 pairwise compatibility analysis (`tire_fiber_step2.tex`) to the cut-tire / 4CT setting.