

Chain pigeonhole on cut tires (pendant-redefined version)

Setup recap

Take G' the cubic planar dual of a maximal planar graph G . Suppose G' is a minimum counterexample to 4CT. Pick a 6-edge cut $C \subseteq E(G')$ (a matching cut), form G'_0, G'_1 with pendant edges as in `cut_depth_label.tex`, BFS-label edges by depth from the pendants.

The crucial change under the *redefined* cut tire definition: each cut tire $T_d^{(i,f)}$ is structurally isomorphic to a partial tire dual $D(T)$ in `paper.tex`. Its face boundary is the T'_{ann} -analogue, and its labelled pendants (out spokes at degree-2 vertices incident to a depth- $(d-1)$ edge; in spokes for depth- $(d+1)$) are the analogue of $D(T)$'s leaves.

The chain pigeonhole argument

Reduction by minimality

By minimality of G' , each G'_i (after pendant addition) is strictly smaller than G' and admits a proper 3-edge-colouring $\chi_i : E(G'_i) \rightarrow \{1, 2, 3\}$. Restricting χ_i to the 6 depth-0 pendants gives a boundary configuration $\sigma_i \in \{1, 2, 3\}^6$.

A proper 3-edge-colouring of G' exists iff some σ is achievable as both σ_0 and σ_1 , i.e. $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ where $\mathcal{R}_i := \{\sigma_i \mid \chi_i \text{ proper}\}$. G' a counterexample $\mathcal{R}_0 \cap \mathcal{R}_1 = \emptyset$.

Layered decomposition via cut tires

The depth-0 pendants of G'_i are precisely the *out spokes* of the cut tires at depth 1 (one out spoke per depth-0 edge). So $\sigma_i = \pi_{\text{out}}(\{T_1^{(i,f)}\})$, where π_{out} projects the global colouring onto all out spokes of depth-1 cut tires combined.

More generally, the cut-tire chain at depths $d = 1, 2, \dots, d_{\max}$ partitions the edges of G'_i by depth. Each cut tire $T_d^{(i,f)}$ has χ_i -restricted colouring that:

- Properly colours the face boundary (a closed walk in H_d),
- Colours out spokes (representing depth- $(d-1)$ edges in G'_i),
- Colours in spokes (representing depth- $(d+1)$ edges).

Compatibility between adjacent tire layers

An out spoke of T_d at boundary vertex v represents a specific depth- $(d-1)$ edge $e^* = vw$ in G'_i . This edge e^* is a *face boundary edge* of some cut tire at depth $d-1$ (since e^* has depth $d-1$ and so lies in H_{d-1}).

Therefore: the colour of an out spoke at depth d equals the colour of the corresponding face boundary edge at depth $d-1$.

This is the chain compatibility constraint: T_d 's out-spoke projection $\pi_{\text{out}}(T_d) \subseteq \{1, 2, 3\}^{|\text{out spokes}|}$ corresponds, via the bijection $\{\text{out spokes of } T_d\} \rightarrow \{\text{specific edges on face boundaries of } T_{d-1}\text{'s}\}$, to a projection of T_{d-1} 's face-boundary colouring.

Symmetrically: in spokes of T_d specific face-boundary edges of some T_{d+1} .

Result inheritance via the partial-tire-dual identification

Because each cut tire *is* a partial tire dual (up to graph isomorphism), every counting/structural result for $D(T)$ in `paper.tex` applies directly:

- **Proposition 1.13 (edge–vertex coloring bijection):** For spoke-only cut tires (face boundary a simple cycle of length n , all n spokes pendant), the number of proper 3-edge-colourings is $2^n + 2(-1)^n$.
- **Definitions 1.15–1.17:** The face boundary plays the role of T'_{ann} ; the spokes play the role of the tire annular face connector $T'_{f'}$ and its inner/outer spokes.
- **Rainbow conjecture (rainbow_proof.tex):** For cut tires whose face boundary has the antipodal-chord SP structure, the spoke projection support equals the perms-per-face set \mathcal{P}_m , conditional on Conjecture 1.5.
- **Face-pair-connection structural description (k9_surviving_partitions.tex):** Cut tires with r all-3 face structure admit exactly 2^r valid “induced partitions” for König-lift purposes.

These results provide quantitative bounds on π_{in} and π_{out} at each layer.

The chain pigeonhole step

For the chain $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_{d_{\text{max}}}$ on side i :

- At the deep interior (depth d_{max}), $T_{d_{\text{max}}}$'s in spokes either don't exist or terminate at the deepest vertex. In the example G'_1 of Holton-McKay #0, T_6 has 2 in spokes still going somewhere, but the chain does eventually terminate.
- Each T_d constrains T_{d-1} and T_{d+1} via the in/out-spoke correspondence with face boundary edges.
- $\mathcal{R}_i = \pi_{\text{out}}(\{T_1^{(i,f)}\})$ as restricted by the chain of compatibility conditions running from depth 1 inward to d_{max} .

Chain pigeonhole says: if at each layer the realisable projection contains enough structure (e.g. a full S_3 -orbit), the chain condition propagates through and yields a non-empty \mathcal{R}_i of substantial size. If both \mathcal{R}_0 and \mathcal{R}_1 contain a common S_3 -orbit, they intersect, contradicting that G' is a counterexample.

What's actually new under the redefinition

1. Direct result transfer

Previously (with the original cut tire definition involving incident edges in G'_i): we had to re-derive each result for cut tires, because cut tires weren't formally $D(T)$.

Now: each cut tire *is* $D(T)$ for some virtual tire T , so Prop 1.13, the rainbow proof, the König-lift framework, all apply without re-derivation.

2. Cubicity restored at face-boundary vertices

Each face-boundary vertex of a cut tire has degree 2 (in H_d) + 1 pendant (in spoke or out spoke) = 3 in the cut tire. So *cut tires are cubic* (at face-boundary vertices), restoring the cubic structure that the previous definition broke (where H_d was not cubic).

This means classical cubic-graph results (König, Tait, chromatic polynomial of $L(C_n)$, etc.) apply directly.

3. The chain-defined structure is more rigid

Under the previous definition, “the cut tire” included variable numbers of incident edges depending on G'_i ’s local structure. Under the new definition, the cut tire is determined by:

- The face f (in H_d).
- Which of f ’s boundary vertices are degree-2 in H_d .
- For each such vertex, whether its non- H_d neighbour is at depth $d - 1$ or $d + 1$.

This is finite, combinatorial, and exactly the data needed for chain compatibility.

What stays open

The fundamental obstacles *don’t* change:

(a) Chain pigeonhole at each layer

The same open conjectures (rainbow, König-lift, 2-SAT solvability) gate the chain argument. They now apply to cut tires directly rather than via translation, but their statements are unchanged.

(b) Chain well-definedness

- Each H_d must have ≥ 1 face. Empirically this holds in the Holton-McKay example at all depths 1–7, but no general guarantee.
- Some H_d faces have length 2 (just a single edge of H_d with two adjacent faces sharing it). These “trivial” cut tires have empty face boundary cycle interior and contribute very little structure.
- Some boundary vertices have degree > 2 in H_d (cut-vertices or branch points of H_d); these don’t get pendants under the strict reading. The cut tire then has fewer than $|f|$ spokes.

(c) Chain length and depth-by-depth variability

In the example, the chain has variable face counts (1–3) and face lengths (2–14) across depths. A uniform bound on $|\pi_{\text{out}}|$ across depths still seems unlikely; the chain would have to handle depth-by-depth structure individually.

Net assessment

The redefinition (cut tire = face boundary + labelled pendants at degree-2 vertices) is a **strict improvement** over the previous formulation, because:

1. It identifies each cut tire with a partial tire dual $D(T)$, so all results from `paper.tex` and downstream notes transfer directly.
2. It restores cubicity at face-boundary vertices, which the previous definition broke.
3. The data needed for chain compatibility (face f , degree-2 boundary vertices, in/out classification) is finite and combinatorial.

But the fundamental hard step is unchanged: chain pigeonhole at each layer reduces to the same open conjectures (rainbow / König- lift) that gate the partial-tire-dual chain pigeonhole. The redefinition doesn't make those easier; it makes the transfer free of charge.

Concrete next step. Empirically verify on the 6 Holton-McKay graphs: for each, pick a matching 6-cut, build cut tire chains on both sides, compute \mathcal{R}_0 and \mathcal{R}_1 , check whether $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$. By 4CT (which we're *not* assuming for the proof) we know it does; the question is whether the cut-tire chain machinery successfully identifies the overlap structurally, or whether (after composition) the chain pigeonhole inequality fails empirically. This is the analogue of `tire_fiber_step2.tex` for the cut-tire setting.