Must a minimum 4CT counterexample have a separating n-cycle with n even and $n \geq 6$?

The question

Statement. Let G be a hypothetical minimum 4-colour counterexample (a minimum planar triangulation requiring ≥ 5 colours). Must G contain a separating n-cycle with n *even* and $n \geq 6$?

Short answer. I do not know of a proof either way. The question is subtle: a parity computation (below) shows that cuts in the cubic dual G^* have a definite parity tied to the side sizes, but I do not see this parity forcing the existence of even-length separating cycles.

Parity computation in the cubic dual

Let G be a planar triangulation and G^* its cubic planar dual. G has V vertices, $3V - 6$ edges, $2V - 4$ faces; G^* has $2V - 4$ vertices, $3V - 6$ edges, V faces. In particular, $|V(G^*)| = 2V - 4$ is *always even.*

Lemma (Cut size parity in cubic graphs)**.** *Let G^* be a cubic graph and let C be an edge cut separating $V(G^*)$ into sides S and $T = V(G^*) \setminus S$. Then*

$$|C| \equiv |S| \equiv |T| \pmod 2.$$

Proof. Counting degree on side S: $3|S| = 2e_S + |C|$, where e_S is the number of edges with both endpoints in S. Hence $|C| = 3|S| - 2e_S \equiv |S| \pmod 2$. Since $|S| + |T| = |V(G^*)|$ is even, $|S| \equiv |T| \pmod 2$. □

Translating to primal cycles. Via the $G \leftrightarrow G^*$ duality, an n-cycle in G corresponds to an n-edge cut in G^*. Lemma says: an *even* n-cycle in G ($n \geq 6$) corresponds to a G^* cut with both sides having an *even* number of vertices. An *odd* n-cycle in G ($n \geq 7$) corresponds to both sides having an *odd* number of vertices.

What Birkhoff gives us

For a minimum 4CT counterexample G:

- G is internally 6-connected (no separating 3-cycle, no separating 4-cycle, no separating 5-cycle with ≥ 2 vertices on each side).
- Equivalently in G^*: cyclic edge connectivity ≥ 6.

So all sufficiently small cuts in G^* are excluded. The smallest non-trivial cuts can be 6-edge ($n = 6$, even) or 7-edge ($n = 7$, odd), and Birkhoff alone permits both.

A potential "no" — can G^* have only odd cuts ≥ 7?

In principle, the minimum non-trivial cut could be 7 (and all non-trivial cuts could be odd-length). In this case the minimum primal separating cycle has length 7 (odd), and the question's answer is "no."

Is this realisable? We need a planar cubic graph G^* satisfying:

- cyclic edge connectivity ≥ 7;
- all non-trivial cyclic edge cuts have odd size.

Lemma gives a constraint: if a G^* has *any* cut at all of size ≥ 6, its parity is fixed by side size. For "all cuts odd" to hold, no even-sized cut would separate. In a cubic graph this corresponds to the side sizes being *odd*. Since $|V(G^*)|$ is even (for any planar triangulation dual), the sides $(|S|, |V(G^*)| - |S|)$ have matching parity — both even or both odd. "All cuts odd" means "all cyclic separations have odd side sizes," which is possible in principle.

However: I don't know of a planar cubic graph that has *only* odd cyclic cuts. The known internally 6-connected examples (icosahedron's dual = dodecahedron, etc.) all have many even cuts of size 6 as well as odd cuts.

A heuristic suggesting "yes" — vertex links

Every vertex v in a planar triangulation G has a *link*: the cycle formed by its neighbours. By Birkhoff, this link is a 5-cycle isolating v. In G^*, this corresponds to a 5-edge cut isolating a single triangle face of G^* (= one vertex of G).

Next layer. Consider the "second link" of v: the set of vertices at G-distance exactly 2 from v. This forms a cycle around the link, of length depending on the degrees of link vertices.

If all link vertices have degree 5 (= minimum for internally 6-connected): the link's 5 vertices contribute $5 \cdot 5 = 25$ incidences, of which 5 are to v (the centre) and $2 \cdot 5 = 10$ are between link vertices (the 5-cycle). The remaining $25 - 5 - 10 = 10$ incidences go to second-link vertices. But the 5 link vertices share their second-link vertices in pairs (each triangle face containing v, a link vertex, and a second-link vertex), so the second link has ≤ 5 distinct vertices.

Working through Euler more carefully: if all degrees are ≥ 5 and the link of v has length 5, the second link has length exactly $5 \cdot (5 - 4) = 5$ (in the icosahedron, second link = link of antipodal vertex). But in larger triangulations (degrees of link vertices ≥ 5, some higher), the second link is generically a cycle of length $\sum_{u \in \text{link}} (\deg(u) - 4) = 5 + \sum_u (\deg u - 5) \geq 5$ with equality only when all link degrees are exactly 5 (icosahedron case).

So in larger internally 6-connected triangulations, second-link length ≥ 6, often equal to 6, often even (especially in "vertex-transitive enough" graphs). This is a heuristic for why 6-cycle separators are abundant, but it's not a proof.

What I can conclude

- **Parity is determined by side size** (Lemma). Even n-cycle separators in G correspond to even-sided cuts in G^*.

- **Birkhoff doesn't rule out odd cuts.** Minimum non-trivial cyclic cut in G^* could in principle be of any size ≥ 6.
- **No known proof that even $n \geq 6$ separators must exist** in min 4CT counterexamples.
- **Empirically**, in all tested internally 6-connected planar triangulations (icosahedron, pentakis dodecahedron, Holton–McKay duals), even 6-cycle separators with both sides ≥ 2 vertices exist in abundance.

Conjecture

Conjecture. *Every internally* 6*-connected planar triangulation* G *with* $|V(G)| \geq 12$ *has a separating* n*-cycle with* n *even and* $n \geq 6$.

Equivalently: every planar cubic graph with cyclic edge connectivity ≥ 6 and $|V| \geq 20$ has a cyclic edge cut of size 6.

This conjecture seems plausible based on the second-link heuristic, but I don't have a proof. A planar cubic graph that violates it would be a structural curiosity worth a name — a "cyclically 7-edge-connected planar cubic graph" — and I do not know an example.

Relevance to the cut-tire framework. If the conjecture holds, our cut-tire framework's domain assumption (= cyclic edge connectivity exactly 6 in G^*) is automatically satisfied by every minimum 4CT counterexample. If it doesn't hold, we'd need to either prove that the counterexample is not of the violating type, or extend the framework to higher-size cuts.