

# COMMENTARY ON THE 4-COLORING MOTIVATION FOR PLANE DEPTH SEQUENCING

## PURPOSE

This note records analytic and empirical responses to the motivation stated in the companion paper: the hope that the quadrilateral sequencing of a deep embedding can be combined with an inductive local choice rule to 4-color the underlying maximal planar graph, with ring completion as the only structural obstruction. The reasoning below is informal and is meant to scope where the 4-coloring story is delicate and where empirical probing is most likely to be informative.

### 1. WHY A PURE PIGEONHOLE ARGUMENT IS UNLIKELY TO CLOSE THIS

A purely counting / pigeonhole argument showing “every ring completion succeeds for at least one coloring of the preceding moves” would, if it worked, give a constructive 4-coloring of every maximal planar graph by following the canonical sequence. The 4-color theorem of Appel–Haken and the cleaner Robertson–Sanders–Seymour–Thomas proof exist because no such argument has been found, and substantial structural machinery (unavoidable configurations, reducibility, Kempe-chain swaps) is necessary. A short pigeonhole would therefore be too good to be true.

There is also an online-coloring obstacle. Even with optimal play, deterministic online algorithms on planar graphs can be forced past 4 colors when the algorithm sees the input one vertex at a time and cannot revise earlier choices. The quadrilateral sequencing fixes a vertex-revelation order in advance, and any local rule that colors a new vertex against its already-colored  $G'$ -neighbors is effectively such an online algorithm.

### 2. WHAT RING COMPLETION ACTUALLY CONSTRAINS

Under a strictly local online rule that colors each new vertex against its already-colored  $G'$ -neighbors, every  $G'$ -edge constraint is enforced at the moment its second endpoint is colored. By the time a ring completion fires, all five edges of its quadrilateral (four perimeter edges and the level edge) have already had their constraints enforced, since their endpoints were colored at earlier steps. The ring completion adds no new edge constraints and therefore cannot fail at the moment it fires — provided that the earlier non-ring-completion moves never got stuck.

This shifts the question. The interesting failure mode is not at ring completion. It is at the earlier moves where a new vertex is introduced and the rule cannot supply a valid color because all four colors are forbidden by the new vertex’s already-colored  $G'$ -neighbors. In other words: the relevant obstruction lives in the anchor drop, level add, and join steps, not in the ring completions.

## 3. EMPIRICAL CHECK ON SMALL TRIANGULATIONS

An empirical check was set up under the simplest concrete instance of an online rule: for each non-ring-completion move, color each new vertex with the smallest unused color among its already-colored  $G'$ -neighbors, in canonical (depth, vertex-id) order. The check enumerates every 3-connected planar triangulation of order  $n \in [4, 9]$  (via Sage’s `graphs.planar_graphs`) and classifies each run as one of `success`, `no_color_available` (greedy exhausted the palette), or `ring_completion_monochromatic` (a ring completion found a monochromatic edge).

The observation on small orders is consistent with the theoretical analysis above:

- Every failure observed so far is of type `no_color_available`.
- No failure of type `ring_completion_monochromatic` occurred.
- Failures appear already at  $n = 5$  and dominate at  $n = 6, 7$ .

The empirical picture, in short: the framing “ring completion is where the trouble lies” is not supported by this discipline. Pure greedy fails at the introduction of vertices well before any ring completion has a chance to be inconsistent. If the local rule is to succeed at all, it must either be smarter than greedy in choosing colors at introductions, or must permit backtracking / Kempe-chain repair when the palette runs out.

## 4. DIRECTIONS THAT HAVE A CHANCE

- **Kempe-chain repair at the failing step.** When a new vertex cannot be colored because its colored neighborhood uses all four colors, swap colors along a Kempe chain to free one. This is the engine of Heawood’s 5-color proof and a load-bearing component of the 4-color proof. It is not pigeon-hole; it is a local recoloring argument whose success depends on the chain structure, not on cardinalities alone.
- **Reducibility / unavoidable-configuration argument.** Identify a finite list of local patches that can appear at a failing introduction, and show each is reducible in the Appel–Haken sense. This is essentially a miniature 4-color proof restricted to the configurations that arise under the canonical sequence.
- **Restricting to subclasses.** For stacked / Apollonian triangulations or triangulations with no separating triangles, an inductive coloring may line up cleanly with the canonical sequence. If empirically pure greedy succeeds on a subclass, that suggests a tractable target subtheorem.
- **Choosing a smarter introduction rule.** Greedy with smallest-color is a particularly weak choice. A more structural rule — e.g., one that consults depth parity, the move type, and the local triangle type — might avoid the `no_color_available` dead-ends. The check infrastructure can be re-parameterized with such a rule to test.

## 5. WHAT THE EMPIRICAL FAILURES MIGHT BE USEFUL FOR

Even though the failures of pure greedy are not the ring-completion failures originally hoped for, they have potential value as a structural probe. Each failing graph identifies a small triangulation on which the canonical sequence forces an introduction step into a 4-color dead-end. The structure of the dead-end — which

vertex, which move type, which neighborhood — is recorded by the saved canonical graphs in `data/graphs/`. These small examples could:

- suggest the form of a smarter introduction rule;
- identify a candidate “unavoidable configuration” for a reducibility argument;
- or, in conjunction with the related colored-edge flip-class work, identify a structural witness that the sequencing interacts non-trivially with proper 4-colorings.

The empirical pipeline is short enough that broader sweeps over  $n$  or over alternative introduction rules are cheap to run, and the failure list can be regenerated whenever the rule changes.