Birkhoff–Heesch reducibility and the fiber view

A dictionary between classical 4CT reducibility theory and our spoke-fiber decomposition

Purpose

The conversation around the fiber decomposition note (`fiber_decomposition.tex`) flagged that the technique of "split a coloring into a boundary configuration plus an extension-counting fiber" is exactly the framework of classical reducibility theory for the Four-Color Theorem (4CT) — developed by Birkhoff (1913), then Heesch, Bernhart, Allaire, Swart, Appel–Haken, and finally Robertson–Sanders–Seymour–Thomas (RSST, 1997). This note (a) summarises the relevant pieces of that machinery in modern notation following Thomas's survey [5], then (b) gives an explicit dictionary to our spoke-fiber language and an honest assessment of what carries over.

Part I: classical reducibility, in modern notation

The 4CT and its Tait dual

Theorem (4CT, vertex form)**.** *Every plane graph G has a proper 4-vertex-coloring.*

Theorem (Tait 1880, edge form)**.** *Every cubic plane graph H with no cut-edge has a proper 3-edge-coloring. This is equivalent to the 4CT.*

The equivalence goes through planar duality: a 4-coloring of a triangulation G, with colors $(0,0),(1,0),(0,1),(1,1)$ in $\mathbb{Z}_2 \times \mathbb{Z}_2$, induces a 3-edge-coloring of the dual G^* by assigning each dual edge the sum of the colors of the two regions it separates. This is exactly the side our paper works on: a triangulation G has a cubic dual $G' \subseteq G^*$ (specifically: G^* minus the unbounded face), and edge 3-colorings of G' *are* (up to handling of the outer face) 4-colorings of G.

Birkhoff's minimum-counterexample setup (1913)

The reducibility programme assumes for contradiction that there is a *minimum counterexample* to the 4CT — a plane graph T that is not 4-colorable but every smaller plane graph is. Birkhoff showed:

Theorem (Birkhoff 1913)**.** *Every minimum counterexample to the 4CT is an internally 6-connected triangulation.*

("Internally 6-connected" means: removing fewer than 5 vertices keeps the graph connected, and removing any 5 vertices leaves at most a single isolated vertex.) All subsequent work assumes T has this form.

Configurations, rings, free completions

A *configuration* captures a small local piece of T together with degree information.

Definition (Configuration, Thomas [5])**.** *A configuration is a pair $K = (H, \gamma)$ where H is a near-triangulation (one face is designated as "special", and every other face is a triangle) and $\gamma : V(H) \to \mathbb{Z}_{\geq 5}$ satisfies:*

(i) for interior vertices v (not on the special face), $\gamma(v) = \deg_H(v)$;

(ii) for boundary vertices v (on the special face), $\gamma(v) > \deg_H(v)$;

(iii) ring-size$(K) := \sum_{v \text{ boundary}, H \setminus v \text{ connected}} (\gamma(v) - \deg_H(v) - 1) \geq 2$.

K appears *in a triangulation T if H is an induced subgraph of T, every non-special face of H is a face of T, and $\gamma(v)$ equals the degree of v in T for every $v \in V(H)$.*

Definition (Free completion)**.** *The* free completion *of K is the (essentially unique) plane graph S obtained from H by adding a single cycle R — the* ring *of K — around the special face, plus the unique set of edges making S a triangulation of the disk bounded by R in which every vertex of H has degree exactly γ. R has length equal to the ring-size of K.*

Color sets and reducibility

Let $\mathcal{K}(R)$ be the set of all proper 4-colorings of the ring R (often considered up to the S_4-action on colors, leaving roughly $|\mathcal{K}(R)|/24$ orbits).

Definition (Good and bad colorings)**.** *A coloring $\varphi \in \mathcal{K}(R)$ is* good *if it extends to a proper 4-coloring of the free completion S, i.e. to all of $H \cup R$. Otherwise it is* bad*. Write $\mathcal{C} \subseteq \mathcal{K}(R)$ for the set of good colorings.*

Now suppose K appears in a minimum counterexample T, with free completion $S \subseteq T$. Let $T' := T \setminus V(H)$ (the "outside"); since T is a minimum counterexample, T' has a 4-coloring, which restricts to some $\mathcal{C}' \subseteq \mathcal{K}(R)$. For T to be a counterexample,

$$\mathcal{C}' \subseteq \mathcal{K}(R) \setminus \mathcal{C}$$

(otherwise a coloring extending both into S and into T' gives a 4-coloring of all of T). The goal of reducibility is to show $\mathcal{C}' = \emptyset$, contradicting that T' is colorable.

Definition (Reducibility, classical taxonomy)**.**

- *A* *K is* ***A-reducible*** *if $\mathcal{C} = \mathcal{K}(R)$ (every ring coloring extends). A-reducibility immediately gives $\mathcal{C}' = \emptyset$, but is too strong to hold for any nontrivial configuration.*
- *D* *K is* ***D-reducible*** *if every bad coloring $\varphi \in \mathcal{K}(R) \setminus \mathcal{C}$ can be transformed — via a sequence of Kempe-chain swaps on T' — into a good coloring. D-reducibility is checkable by computer on $\mathcal{K}(R)$ alone (no knowledge of T' beyond that it is 4-colorable), because the Kempe-swap closure operation can be applied at the level of ring colorings.*
- *C* *K is* ***C-reducible*** *if it is D-reducible* after *replacing H by a smaller graph H' (obtained from H by contracting up to four edges in RSST). This is strictly more powerful than D-reducibility.*

Birkhoff's diamond (4 mutually adjacent pentagons surrounded by a 6-cycle ring) was the first nontrivial D-reducible configuration: of the 31 ring colorings up to S_4-symmetry, 16 are good directly, and the other 15 each admit a Kempe-swap to a good one.

Discharging and unavoidability

Reducibility tells you that *if* a good configuration appears in T, then T is not a counterexample. The other half of the proof is to show that some good configuration *must* appear:

Definition (Unavoidable set)**.** *A set $\mathcal{U}$ of configurations is* unavoidable *if every internally* 6*-connected triangulation contains at least one $K \in \mathcal{U}$ as a subconfiguration.*

Heesch's *discharging method* proves unavoidability:

1. Assign a *charge* $\mathrm{ch}(v) := 6 - \deg(v)$ to each vertex. By Euler's formula, $\sum_v \mathrm{ch}(v) = 12$ on any triangulation of the sphere.
2. Define *discharging rules* that redistribute charge between vertices without changing the total.
3. Show that after discharging, no vertex carries positive charge *unless* it lies in (the neighborhood of) some configuration in $\mathcal{U}$.
4. Since total charge is $12 > 0$, some positive-charge vertex exists post-discharging, so some configuration in $\mathcal{U}$ must appear.

The two halves combine: 4CT = "every reducible configuration in $\mathcal{U}$ blocks counterexamples" + "$\mathcal{U}$ is unavoidable."

The proofs

- **Appel–Haken (1976/77).** $|\mathcal{U}| = 1936$ configurations (later reduced to 1482); 487 discharging rules. Computer verification needed.
- **RSST (1997).** $|\mathcal{U}| = 633$ configurations; 32 discharging rules; quadratic-time 4-coloring algorithm. Still computer-assisted, but the unavoidability part was written in a formal language and machine-checked.

For configurations of ring-size r, the number of colorings of R modulo S_4 is roughly $3^{r-1}/3 + O(2^r)$ from the standard chromatic-polynomial formula for C_r. RSST's largest ring-size is 14, with $\sim$200,000 ring colorings per configuration to check.

Part II: dictionary to the fiber-decomposition view

Side-of-the-duality conventions

Our work is on the *edge* side of Tait: we color edges of G' with 3 colors, where G' is essentially the dual of a triangulation G. Birkhoff–Heesch is on the *vertex* side of G. By Tait's theorem these are equivalent in principle, but the shape of a "configuration boundary" looks different on the two sides:

Vertex side (Birkhoff–Heesch)	**Edge side (us)**
Triangulation T	Cubic dual G'
Configuration $K = (H, \gamma)$ in T	Tire annular face connector $T'_{f'} \subseteq G'$
Inner-face triangles of H	Vertices of $V(f')$ (dual to annular faces of the tire)
Ring R (a cycle around H)	Boundary of $T'_{f'}$, consisting of spoke edges E_S *plus* (in the multi-tire chain) the cycle $V(f')$ itself
Ring coloring $\varphi : V(R) \to \{1, 2, 3, 4\}$	Spoke configuration $\sigma : E_S \to \{1, 2, 3\}$
$\mathcal{K}(R)$ = all proper 4-colorings of R	$\Sigma :=$ all spoke configurations σ (whether realisable or not)
$\mathcal{C}$ = good (extendable to $H \cup R$) ring colorings	Realisable spoke configurations: $\{\sigma : N(T'_{f'}; \sigma) > 0\}$
$\mathcal{K}(R) \setminus \mathcal{C}$ = bad ring colorings	Unrealisable σ: the boundary of $T'_{f'}$'s reachability
Fiber over a ring coloring: # extensions of φ to $H \cup R$	Fiber count $N(T'_{f'}; \sigma)$
Kempe-chain swap on $T \setminus V(H)$	Tait-Kempe swap on $G' \setminus V(T'_{f'})$ (two-color alternating-edge swap)

The fiber identity, in Birkhoff language

Our identity $P_e(T'_{f'}, 3) = \sum_\sigma N(T'_{f'}; \sigma)$ is literally Birkhoff's "good colorings" decomposition on the edge side, with *the same scaffold*: count interior extensions fibered over boundary state. The classical literature does not emphasise this identity per se (it's the trivial sum), but every reducibility analysis is built around the fiber distribution $\{N(T'_{f'}; \sigma)\}_\sigma$.

What "reducible" means in our world

Translating definitions of A-, D-, C-reducibility through the dictionary:

- **Edge-A-reducibility of** $T'_{f'}$: every spoke configuration $\sigma \in \Sigma$ is realisable. This would mean $T'_{f'}$ supports an edge 3-coloring extension from *any* boundary input — a very strong locally-flexible property.
- **Edge-D-reducibility of** $T'_{f'}$: every unrealisable σ can be Tait-Kempe-swapped (in the outside graph $G' \setminus T'_{f'}$) to a realisable one. This is the natural form of the conductivity step from our chain-pigeonhole sketch.
- **Edge-C-reducibility of** $T'_{f'}$: D-reducibility after replacing $T'_{f'}$ by a smaller subgraph (e.g. contracting some spoke edges or merging adjacent faces).

The chain-pigeonhole step as a reducibility statement

Our conductivity step — "the middle tire T_B's $\phi_B : \mathcal{P}_{AB} \to 2^{\mathcal{P}_{BC}}$ takes any input to a $>|\mathcal{P}_{BC}|/2$-sized output" — is the *quantitative* analogue of D-reducibility, but applied to *composition* of configurations rather than reduction. In particular:

- Birkhoff–Heesch D-reducibility says "bad inputs to a single configuration can be Kempe-moved to good ones."
- Our chain-pigeonhole says "the bad inputs to the inner tire and the bad inputs to the outer tire together don't cover $\mathcal{P}_\gamma$, so something good remains."

These are not the same statement, but they are about the same data $\mathcal{C}$ vs. $\mathcal{K} \setminus \mathcal{C}$ on the shared cycle.

Part III: does the machinery apply?

What carries over cleanly

1. **Vocabulary and scaffolding.** Configurations, rings, free completions, color sets, fiber-of-extensions are all there; the menagerie/fiber notes are using the same objects.
2. **Kempe-chain machinery.** Tait-Kempe chains (two-color alternating edge paths in G') are a real, well-developed tool. Any conductivity argument we want is fundamentally Tait-Kempe in flavour.
3. **Computer-verifiability for small tires.** Just as D-reducibility of an individual configuration is verified by enumerating $\sim$200,000 ring colorings, the fiber distribution $\{N(T'_{f'}; \sigma)\}$ for a fixed small tire is a finite Sage computation. We can test conductivity empirically before trying to prove it.

What does *not* carry over straightforwardly

1. **The reducibility scale.** Classical reducibility was only ever practical for ring-size ≤ 14 because the number of ring colorings explodes. Our tires can have arbitrarily large boundary cycles (in a multi-layer triangulation, the annular ring is whatever the level structure gives), so the *single-configuration* reducibility approach hits the same wall as A&H/RSST: only small tires are tractable directly.
2. **Compositionality.** Birkhoff–Heesch operates on *one* configuration at a time. Our chain pigeonhole/nesting is fundamentally about *composing* configurations along shared boundaries. This is a structural feature classical reducibility does *not* engage with — their unavoidability argument (discharging) replaces it. If we want to prove a statement of the form "every nested chain of tires admits a global 4-coloring," Birkhoff–Heesch does not directly give us the tool; we need either a transfer-matrix / monotonicity argument across nestings or a structural result about how the realisable supports behave under composition.
3. **Unavoidability is automatic / different.** In 4CT, unavoidability is a separate hard problem solved by discharging. In our setup the tire decomposition of G is given by the level structure (Bauerfeld [7]), so there is no analogous unavoidability question — the tires are already there. This means *half* of the classical apparatus (discharging) is not the point of contact; the contact is purely on the reducibility/color-set side.
4. **Tait correspondence is global, not local.** Edge 3-coloring of G' globally encodes a vertex 4-coloring of G, but the local correspondence between an edge coloring of $T'_{f'}$ and a vertex coloring of the corresponding piece of G is subtle: edge swaps on the dual side do not always correspond to single Kempe chains on the primal side, and the "ring" as a vertex cycle in G may differ in length and structure from the "boundary" of $T'_{f'}$ on the edge side. Any time we want to import a vertex-side Kempe argument we will have to do the bookkeeping carefully.

Assessment

Yes, the Birkhoff–Heesch framework applies, but only as a language and as a verification tool for small instances. It does *not* hand us a proof of the nesting/chain-pigeonhole conjecture for free, because:

- Classical reducibility is single-configuration; our argument is multi-configuration / compositional.
- Classical reducibility's quantitative input is exhaustive enumeration of ring colorings up to size 14; we want statements uniform over arbitrarily large tires.

What it *does* hand us:

- A precise vocabulary — good/bad colorings, free completion, D-/C-reducibility — so that the conductivity step can be stated in established terms.
- A concrete computational template: for any specific tire, compute $\{N(T'_{f'};\sigma)\}_\sigma$ in Sage and check whether the realisable support $\mathcal{C}$ is large enough that two adjacent tires must overlap. This is the analogue of mechanical D-reducibility checking.
- Strong evidence about the difficulty: nothing in 80+ years of reducibility work has reduced 4CT to a structural argument across all configuration sizes. If our chain argument succeeds, it will be *because tires are a more structured class than arbitrary configurations*, not because the reducibility apparatus suddenly gives uniform results.

Concrete next steps

1. **Pick a small tire family** (B_{in} a k-cycle for $k \in \{3,4,5,6\}$, B_{out} a small cycle, no O-chords) and compute the full fiber distribution $\{N(T'_{f'};\sigma)\}_\sigma$ in Sage. This gives the realisable support $\mathcal{C}$ and, in classical terms, tests whether small tires are A-reducible (full support), D-reducible (Kempe-recoverable), or neither.
2. **Check support overlap for adjacent tires.** Given two small tires sharing a cycle γ, do the realisable supports $\mathcal{C}_{\text{out}}, \mathcal{C}_{\text{in}} \subseteq \mathcal{K}(\gamma)$ always intersect? This is the empirical version of the chain-pigeonhole step. If they sometimes miss, the simple form of the argument fails and we need a Kempe-chain (D-reducibility-style) escape route.
3. **Look up monosystems / Tutte.** Tutte [6] formulated 4CT-adjacent counting problems in terms of his *chromial*; this is the closest existing transfer-matrix-style framing of these color sets, and may give a cleaner composition rule than the raw fiber sum.
4. **Look up Sokal / Chang–Shrock** on chromatic polynomials of strip graphs. They use transfer matrices for infinite families of small-ring configurations, which is structurally similar to nested tires with fixed ring size.

References

[1] G. D. Birkhoff, *The reducibility of maps*, Amer. J. Math. **35** (1913), 115–128.

[2] H. Heesch, *Untersuchungen zum Vierfarbenproblem*, B. I. Hochschulskripten 810/810a/810b, Bibliographisches Institut, Mannheim, 1969.

[3] K. Appel and W. Haken, *Every planar map is four colorable. Part I: Discharging*; *Part II: Reducibility* (with J. Koch), Illinois J. Math. **21** (1977), 429–567.

[4] N. Robertson, D. P. Sanders, P. Seymour, R. Thomas, *The Four-Colour Theorem*, J. Combin. Theory Ser. B **70** (1997), 2–44.

[5] R. Thomas, *An update on the Four-Color Theorem*, Notices Amer. Math. Soc. **45** (1998), 848–859.

[6] W. T. Tutte, *Chromials*, in: *Hypergraph Seminar* (C. Berge, D. Ray-Chaudhuri, eds.), Lecture Notes in Math. 411, Springer, 1974, pp. 243–266.

[7] E. Bauerfeld, *Plane Depth Sequencing*, manuscript (math-research repository), 2026.