

Recolouring planar graphs of girth at least five

Valentin Bartier, Nicolas Bousquet, Carl Feghali, Marc Heinrich, Benjamin Moore, Théo Pierron

▶ To cite this version:

Valentin Bartier, Nicolas Bousquet, Carl Feghali, Marc Heinrich, Benjamin Moore, et al.. Recolouring planar graphs of girth at least five. 2022. hal-03773048

HAL Id: hal-03773048 https://hal.science/hal-03773048

Preprint submitted on 8 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recolouring planar graphs of girth at least five

Valentin Bartier¹, Nicolas Bousquet^{*2}, Carl Feghali^{†1}, Marc Heinrich³, Benjamin Moore^{‡4}, and Théo Pierron²

¹Univ. Lyon, ENS Lyon, CNRS, LIP, F-69342, Lyon Cedex 07, France
²Univ. Lyon, Université Lyon 1, LIRIS UMR CNRS 5205, F-69621, Lyon, France
³School of Computing, University of Leeds, UK
⁴Charles University, Institute of Computer Science, Prague, Czech Republic

September 8, 2022

Abstract

For a positive integer k, the k-recolouring graph of a graph G has as vertex set all proper k-colourings of G with two k-colourings being adjacent if they differ by the colour of exactly one vertex. A result of Dyer et al. regarding graphs of bounded degeneracy implies that the 7-recolouring graphs of planar graphs, the 5-recolouring graphs of triangle-free planar graphs and the 4-recolouring graphs planar graphs of girth at least six are connected. On the other hand, there are planar graphs whose 6-recolouring graph is disconnected, triangle-free planar graphs whose 4-recolouring graph is disconnected and planar graphs of any given girth whose 3-recolouring graph is disconnected.

The main result of this paper consists in showing, via a novel application of the discharging method, that the 4-recolouring graph of every planar graph of girth five is connected. This completes the classification of the connectedness of the recolouring graph for planar graphs of given girth. We also prove some theorems regarding the diameter of the recolouring graph of planar graphs.

1 Introduction and results

Let $k \ge 1$ be an integer, and let G = (V, E) be a graph. A *(proper)* k-colouring of G is a function $\sigma : V \to \{1, \ldots, k\}$ such that, for every edge $xy \in E$, we have $\sigma(x) \ne \sigma(y)$. The k-recolouring graph of G, denoted by $\mathcal{G}(G, k)$, is the graph whose vertices are the kcolourings of G, with two vertices being adjacent whenever the corresponding colourings differ on exactly one vertex.

^{*}Supported by ANR project GrR (ANR-18-CE40-0032)

[†]Supported by the Czech Science Foundation under research grant 19-21082S and by the French National Research Agency under research grant ANR DIGRAPHS ANR-19-CE48-0013-01.

[‡]Supported by a NSERC doctoral grant.

The problem of determining whether $\mathcal{G}(G, k)$ is connected is equivalent to that of determining whether the Glauber dynamics on the k-colourings of G is ergodic (see e.g. [10, 16]). One of the earliest results in this direction is the following theorem by Dyer et al. [13] and rediscovered by Cereceda et al. [9]. For a positive integer d, a graph is d-degenerate if each of its subgraphs contains a vertex of degree at most d.

Theorem 1.1. Let k and d be positive integers such that $k \ge d+2$. If G is a d-degenerate graph, then $\mathcal{G}(G, k)$ is connected.

The bound on k in Theorem 1.1 is best possible (for example, if G is the complete graph on d+1 vertices, then $\mathcal{G}(G, d+1)$ a collection of (d+1)! isolated vertices). Perhaps surprisingly, this bound also turns out to be best possible for planar graphs, triangle-free planar graphs and planar graphs of girth at least six. Indeed, by Euler's formula, planar graphs are 5-degenerate, triangle-free planar graphs are 3-degenerate and planar graphs of girth at least six are 2-degenerate. Thus, Theorem 1.1 implies that if G is a planar graph, then $\mathcal{G}(G,7)$ is connected. Similarly, $\mathcal{G}(G,5)$ is connected if G is triangle-free and planar, and $\mathcal{G}(G,4)$ is connected if G has girth at least six and planar.

On the other hand, there is a planar graph G and a 6-colouring of G where the closed neighbourhood of each vertex contains all 6 colours [2] and hence this 6-colouring forms an isolated vertex in $\mathcal{G}(G, 6)$, implying $\mathcal{G}(G, 6)$ is disconnected. For triangle-free planar graphs, consider the 4-colouring of the cube Q_3 where all colours are used in every face. Again, this colouring forms an isolated vertex in $\mathcal{G}(Q_3, 4)$ and hence $\mathcal{G}(Q_3, 4)$ is disconnected. For planar graphs of girth at least six, the 3-colouring of a cycle C whose length is a multiple of three where the colours 1, 2, 3 alternate in a cyclic fashion also forms an isolated vertex in $\mathcal{G}(C, 3)$. This motivates the following question.

Problem 1.2. What is the smallest integer κ such that for every planar graph G of girth five, the graph $\mathcal{G}(G, \kappa)$ is connected?

Note by our earlier discussion that $\kappa \in \{4, 5\}$ for Problem 1.2. In the main contribution of this paper, we settle the problem by showing that $\kappa = 4$.

Theorem 1.3. For every planar graph G of girth five, the graph $\mathcal{G}(G, 4)$ is connected.

By our earlier discussion, Theorem 1.3 completes the classification of the connectedness of the recolouring graph for planar graphs of given girth. The proof idea of Theorem 1.3 is based on a new application of the discharging method, in which we allow an infinite number of reducible configurations rather than the usual finite number. More precisely, we first prove that a special type of C_5 is forbidden. We then show that any possible gluing of these C_5 in a tree-like structure is also forbidden. We then conclude with a very short and simple discharging argument; see Section 3.1 for a more detailed outline.

On a slightly different track, the question of determining what is the diameter of $\mathcal{G}(G,k)$ (when it is connected) also received considerable attention. In other words, how fast can we reach one k-colouring of G from another by changing the colour of one vertex at a time? This question and, particularly, the following strengthening of Theorem 1.1 conjectured by Cereceda [8] have been the subject of much attention in recent years.

Conjecture 1.4. Let k, d be positive integers, $k \ge d+2$ and let G be a d-degenerate graph on n vertices. Then $\mathcal{G}(G, k)$ has diameter $O(n^2)$.

This bound would be best possible [3]. Although the conjecture has resisted several efforts, there have been some partial results surrounding it. The most important break-through comes from Bousquet and Heinrich [6], who show amongst other results, that $\mathcal{G}(G,k)$ has diameter $O(n^{d+1})$. Nonetheless, Conjecture 1.4 remains open even for d = 2.

When k is substantially larger than d, Bousquet and Perarnau [7] gave the following bound.

Theorem 1.5. Let k and d be positive integers, such that $k \ge 2d+2$. If G is a d-degenerate graph on n vertices, then $\mathcal{G}(G, k)$ has diameter O(n).

We conjecture that the bound on k in Theorem 1.5 can be lowered to d + 3 and this would be best possible [3].

Conjecture 1.6. Let k, d be positive integers, $k \ge d+3$ and let G be a d-degenerate graph on n vertices. Then $\mathcal{G}(G, k)$ has diameter O(n).

To the best of our knowledge, Conjecture 1.6 is only known to hold for outerplanar graphs [1] and 1-degenerate graphs. For partial results, Bartier and Bousquet [5] proved that $\mathcal{G}(G, d+4)$ has diameter O(n) for every *d*-degenerate chordal graph *G* of bounded maximum degree, and Dvořák and Feghali proved that $\mathcal{G}(G, 10)$ has diameter O(n) for every planar graph *G* [11, 12], and if *G* is triangle-free, then so does $\mathcal{G}(G, 7)$ [12].

As a second contribution in this paper, we confirm the conjecture for planar graphs of girth at least 6, improving the bound on k in Theorem 1.5 in this special case.

Theorem 1.7. For every planar graph G on n vertices of girth at least 6, $\mathcal{G}(G,5)$ has diameter O(n).

To prove Theorem 1.7, we show that planar graphs of girth ≥ 6 contain a 2-degenerate 1-island of constant size (see Section 2 for a definition), by using virtually the same arguments as in [14] for a similar result; we then simply show that such a configuration is reducible for the problem. We remark that it may be possible to improve Theorem 1.7 by showing that $\mathcal{G}(G, 4)$ has linear diameter (this would be best possible, since, as discussed, there exists planar graphs G of girth ≥ 6 such that $\mathcal{G}(G, 3)$ is disconnected).

As our final contribution, we address a recent conjecture by Dvořák and Feghali that generalizes their aforementioned result that $\mathcal{G}(G, 10)$ has diameter O(n) for every planar graph G. In order to state the conjecture and our theorem, we require some further definitions. A list assignment L for G is a function that, to each vertex $v \in V$, assigns a set L(v) of colours. An L-colouring of G is (proper) colouring φ of G such that $\varphi(v) \in L(v)$ for each $v \in V$. For a list assignment L of G, the L-recolouring graph of G, denoted by $\mathcal{G}(G, L)$, is the graph whose vertices are the L-colourings of G, with two vertices being adjacent whenever the corresponding colourings differ on exactly one vertex. Dvořák and Feghali [12] made the following conjecture.

Conjecture 1.8. Let G = (V, E) be a planar graph on n vertices, and let L be a list assignment for G with $|L(v)| \ge 10$ for each $v \in V$. Then $\mathcal{G}(G, L)$ has diameter O(n).

In the case where the number 10 of colours is replaced by 12, Conjecture 1.8 follows from [[6], Theorem 1]. We improve this result as follows:

Theorem 1.9. Let G = (V, E) be a planar graph on n vertices, and let L a list assignment for G where $|L(v)| \ge 11$ for every $v \in V$. Then $\mathcal{G}(G, L)$ has diameter O(n).

Our proof of Theorem 1.9 follows the same spirit as our proof of Theorem 1.7, by combining a structural result of Borodin for planar graphs with a standard reducible configuration argument.

Girth/colours	4	5	6	7
3	$+\infty$	$+\infty$	$+\infty$	$O(n^6)$ [6]
4	$+\infty$	$O(n^4)$ [6]	$O(n \log^3(n))$ [15]	O(n) [12]
5	$< +\infty \text{ (Thm 1.3)}$	$O(n\log^2 n) \ [15]$	-	-
6	$O(n^3)$ [6]	O(n) (Thm 1.7)	-	-
7+	$O(n\log n) \ [15]$	-	-	-

Table 1: Existing and open cases for the diameter of the k-recolouring graph of planar graphs with girth g for some combinations of values of k and g. Note that any bound at a given position in the table implies the same bound at its right and below it.

We end this section with a summary in Table 1 of some of the results on the diameter of $\mathcal{G}(G, k)$ for planar graphs G of given girth g. We should remark that there is no known lower bound beating the (trivial) $\Omega(n)$ bound on the diameter of such $\mathcal{G}(G, k)$ whenever it is connected, which suggests an interesting direction of research.

For more results on this topic or on reconfiguration versions of decision problems other than graph colouring, we refer the reader to the surveys by van den Heuvel [19] and by Nishimura [18].

2 The proofs of Theorems 1.7 and 1.9

We start with some definitions. For a subgraph H of a graph G and a colouring σ of G, we denote by $\sigma_{\restriction H}$ the restriction of σ to H. A colouring σ' of G is obtained from σ by a single step recolouring, denoted $\sigma \sim \sigma'$, if σ and σ' differ in the colour of exactly one vertex. Let α and β be two colourings of G. Given a sequence σ' of recolourings $\alpha_{\restriction H} \sim \cdots \sim \beta_{\restriction H}$ in H, we say that σ' lifts to a sequence σ of recolourings $\alpha \sim \cdots \sim \beta$ in G if $\sigma_{\restriction H=\sigma'}$. Note that the recolouring sequences σ and σ' may not be of the same length. For a recolouring sequence $\alpha = \alpha_1 \sim \alpha_2 \sim \cdots \sim \alpha_m = \beta$, the set of new colours of H is the set $\bigcup_{i=2}^m \{\alpha_i(u) : \alpha_i(u) \neq \alpha_1(u), u \in H\}$.

In order to prove Theorems 1.7 and 1.9, we will use the following basic, yet powerful, lemma (whose proof is implicit in a number of papers).

Lemma 2.1. Let G be a graph and v a vertex of G. Let k and C be positive integers, $k \ge d(v) + 2$. Let α and β be k-colourings of G. Suppose that there is a recolouring sequence from $\alpha_{\restriction(G-v)}$ to $\beta_{\restriction(G-v)}$ that recolours the neighbourhood N(v) of v at most C times. Then the sequence lifts to a sequence in G from α to β by at most $\lceil \frac{C}{k-d(v)-1} \rceil + 1$ recolourings of v.

Proof. Since the total number of times the neighbours of v change their colour is at most C, we can let a_1, \ldots, a_C be the new colours in N(v) in order. Let $t = k - d_G(v) - 1$. We start

by recolouring v to a colour not in $\{a_1, \ldots, a_t\}$, followed, in turn, by the first t recolourings a_1, \ldots, a_t in N(v). We then recolour v to a colour not in $\{a_{t+1}, \ldots, a_{2t}\}$, followed by the next t recolourings in N(v); we continue this process until all C recolourings in N(v) have occurred. Notice that it is always possible to recolour v since there are at least t colours not appearing on v or any of its neighbours. Finally, we recolour v to $\beta(v)$. Clearly, v is recoloured at most $\lceil \frac{C}{t} \rceil + 1$ times.

To prove Theorem 1.9, we combine Lemma 2.1 with the following result of Borodin [4].

Theorem 2.2. If G is a planar graph with minimum degree 5, then G has a 3-face T with vertices v_1, v_2, v_3 such that $d(v_1) + d(v_2) + d(v_3) \leq 17$.

We say that (G, α, β) is a counterexample to Theorem 1.9 if

- G is a planar graph,
- α and β are L-colourings of G, where L is as in Theorem 1.9, and
- α cannot be transformed to β by at most 100 recolourings per vertex.

A counterexample (G, α, β) to Theorem 1.7 is said to be *minimal* if |V(G)| is minimum amongst all counterexamples to Theorem 1.7.

Proof of Theorem 1.9. Let (G, α, β) be a minimal counterexample to Theorem 1.9.

We first claim that G has minimum degree at least 5. Otherwise, G has a vertex v with degree at most 4. By minimality, there is a recolouring sequence from $\alpha_{\uparrow(G-v)}$ to $\beta_{\uparrow(G-v)}$ where each vertex gets recoloured at most 100 times. By Lemma 2.1, this sequence lifts to a sequence in G by at most 68 recolourings of v, a contradiction.

We claim that G does not have two vertices v_1 and v_2 such that $v_1v_2 \in E(G)$ and the degrees of v_1 and v_2 are precisely 5. Otherwise, by minimality, there is a recolouring sequence from $\alpha_{\uparrow(G-\{v_1,v_2\})}$ to $\beta_{\uparrow(G-\{v_1,v_2\})}$ where each vertex gets recoloured at most 100 times. By Lemma 2.1, this sequence lifts to a sequence in $G - v_1$ by recolouring v_2 at most 68 times. By the same lemma, the latter sequence lifts to a sequence in G by at most 95 recolouring of v_1 , which is a contradiction. This proves the claim.

A similar argument can be applied to show that G has no triangle v_1, v_2, v_3 such that v_1 has degree 5 and both v_2 and v_3 have degree 6 (consider the graphs $G - \{v_1, v_2, v_3\}$, $G - \{v_1, v_2\}$, $G - v_1$ and G in order). Indeed, v_3 is recoloured 68 times, v_2 is recoloured 95 times and finally v_1 is recoloured at most 96 times.

To complete the proof, by Theorem 2.2, G contains a triangle T with vertices v_1, v_2, v_3 such that $d(v_1) + d(v_2) + d(v_3) \leq 17$. As G does not contain an edge where both vertices have degree 5, it implies that up to relabelling, $d(v_1) = 5$, $d(v_2) = 6$ and $d(v_3) = 6$, which contradicts the preceding argument.

The proof of Theorem 1.7 also follows from Lemma 2.1, but will require some more work.

Given a graph G and an induced subgraph H of G, we say that H is a 2-degenerate 1-island if the following hold:

• each vertex $v \in V(H)$ has at most one neighbour in G - H, and

• there exists an ordering $v_1, \ldots, v_{|V(H)|}$ of the vertices of H such that for each $i \in \{2, \ldots, |V(H)|\}$, the vertex v_i has at most two neighbours in the graph $G - H + \{v_1, \ldots, v_{i-1}\}$.

We say that (G, α, β) is a counterexample to Theorem 1.7 if G is a planar graph of girth at least 6 and α and β are 5-colourings of G that cannot be transformed to one another by at most 6143 recolourings per vertex. A counterexample (G, α, β) to Theorem 1.7 is minimal if |V(G)| is minimum amongst all counterexamples to Theorem 1.7.

It turns out that a minimal counterexample cannot contain a 2-degenerate 1-island of order at most 12.

Lemma 2.3. Let (G, α, β) be a minimal counterexample to Theorem 1.9. Then G does not contain a 2-degenerate 1-island H where $|V(H)| \leq 12$.

Proof. Suppose towards a contradiction that G contains a 2-degenerate 1-island H with $|V(H)| \leq 12$. Let t = |V(H)|, and let v_1, \ldots, v_t be an ordering of V(H) such that for all $i \in \{2, \ldots, t\}$ the vertex v_i has at most two neighbours in the graph $G - H + \{v_1, \ldots, v_{i-1}\}$. By minimality, there is a recolouring sequence σ' in G - H from $\alpha_{\restriction (G-H)}$ to $\beta_{\restriction (G-H)}$ that recolours every vertex at most 6143 times. We lift σ' to a recolouring sequence σ from α to β in G as follows. Let $c_1 = 2049$, and for $i \geq 2$, let $c_i := \lceil \frac{6143+c_{i-1}}{2} \rceil + 1$. Note that $c_{12} = 6143$ and $c_{i+1} \geq c_i$.

To prove the lemma, it suffices to show that σ' lifts to σ by recolouring each v_i at most c_i times. We proceed by induction on i. For i = 1, since v_1 has at most one neighbour in G - H, Lemma 2.1 implies that v_1 is recoloured at most $2049 = c_1$ times. Suppose $i \in \{2, \ldots, t\}$. Since H is a 2-degenerate 1-island, v_i has at most one neighbour in G - H and at most two neighbours in $G - H + \{v_1, \ldots, v_{i-1}\}$, the neighbours of v_i are recoloured at most $c_{i-1} + 6143$ times. By Lemma 2.1, σ' lifts to a recolouring sequence of $G - H + \{1, \ldots, v_i\}$ by recolouring v_i at most c_i times, as desired.

To finish the proof of Theorem 1.7, it suffices to show, by Lemma 2.3, that every planar graph of girth at least 6 contains a 2-degenerate 1-island H where $|V(H)| \leq 12$. This follows nearly immediately from the proof of Theorem 8 in [14]. We include the proof for completeness since the result stated in [14] is slightly weaker (but holds for more general surfaces).

Lemma 2.4. Every planar graph of girth at least 6 contains a 2-degenerate 1-island H where $|V(H)| \leq 12$.

Proof. Let G be a vertex minimal counterexample. Then G has minimum degree 2; if not, the vertex of degree 1 is a 2-degenerate 1-island. A similar argument can also be applied to show that G does not contain

- (A1) a path with at most 12 vertices each of degree at most 3 where the endpoints have degree 2;
- (A2) a cycle C with at most 12 vertices where all vertices in C have degree at most 3 and at least one vertex of C has degree 2.

We now proceed via the classical discharging method. For every $v \in V(G)$, let $\operatorname{ch}(v) = 2d(v) - 6$, and for every face f in G, let $\operatorname{ch}(f) = d(f) - 6$. By Euler's formula, $\sum_{x \in V \cup F} \operatorname{ch}(x) = -12$.

For every face f in G, choose an orientation of f and call it *positive*, and let the other orientation be called *negative*. We now describe a procedure by which charges are redistributed. For any vertex v where d(v) = 2, for any face f incident to v and for any orientation of f, take a maximal facial walk of f starting at v and going around f in the prescribed orientation of f such that the inner vertices of the walk have degree precisely 3. Let u be the other end vertex of the walk (note that u = v is a possibility). Our discharging rules are as follows:

- If the walk has at least 5 inner vertices, f sends a charge of $\frac{1}{2}$ to v.
- Otherwise, u sends a charge of $\frac{1}{2}$ to v.

By the maximality of the walk, if the second case occurs, then by (A1) and (A2) u has degree at least 4 and there are no paths where both endpoints have degree 2 and all internal vertices have degree at least 3.

To reach a contradiction, we show that the final charges $ch^*(f)$ and $ch^*(v)$ of each face f and vertex v is non-negative.

Case 1: d(v) = 2.

In this case, v appears four times in the union of all boundary walks of faces of G (for each face, we consider a boundary walk in the positive orientation and a boundary walk in the negative orientation of the face). Therefore v receives a charge of $4 \times \frac{1}{2} = 2$, and thus $ch^*(v) = -2 + 2 \ge 0$, as required.

Case 2: d(v) = 3.

Vertices of degree 3 start with 0 charge and do not give any charge, hence they end up with a final charge of 0.

Case 3: $d(v) \ge 4$.

Consider the facial walks through which v gives a charge of $\frac{1}{2}$ to some vertices of degree 2. Note that there are at most two such walks for each of the d(v) faces incident to v. Therefore we have $\operatorname{ch}^*(v) \ge 2d(v) - 6 - 2 \times d(v) \times \frac{1}{2} = d(v) - 6$, which is non-negative when $d(v) \ge 6$.

To handle the remaining cases, we first claim that if a neighbour u of v is immediately before v in more than one such walk, then u has degree 2. Assume by contradiction that u has degree ≥ 3 and is just before v in two facial walks ending at v, then by the above procedure u must have degree exactly 3 and there are two paths starting at u each containing at most three inner vertices (each of degree 3) and finishing at a vertex of degree 2 (if there were more than three inner vertices, then together with vertices u and v, it would give at least 5 inner nodes, a contradiction since we assumed that we cannot apply the first discharging rule). It follows that G contains a path on at most 9 vertices each having degree at most 3 and with its endpoints having degree 2, contradicting (A1).

Let us now complete the proof of Case 3. If d(v) = 5, then v cannot be adjacent to at least 4 vertices of degree 2 since otherwise v together with these vertices form a 2-degenerate 1-island of size at most 5. Therefore, by the claim, v gives a charge of at most $3 + 2 \times \frac{1}{2} = 4$, implying $ch^*(v) \ge 0$. Suppose d(v) = 4. Then a similar argument as above (see below) can be applied to show that

- v is not adjacent to more than 2 vertices of degree 2,
- if v is adjacent to two vertices of degree 2, then v does not give any charge through its neighbours of degree at least 3,
- if v is adjacent to one vertex of degree 2, then v does not give any charge through more than one neighbour of degree at least 3

In each case, we obtain $ch^*(v) \ge 0$. Lastly, if v has no neighbour of degree 2, by the claim, $ch^*(v) \ge 2 - 4 \times 1/2 \ge 0$. We now give the details.

- If, for a contradiction, v is adjacent to at least 3 vertices of degree 2, then these vertices together with v form a 2-degenerate 1-island of size at most 5. Thus v is adjacent to at most two vertices of degree 2.
- If v is adjacent to exactly two vertices of degree 2, then v does not send charge through any of its degree 3 neighbours. This is because if v did, then the path from the neighbour to the vertex of degree 2, combined with v and its degree 2 neighbours is a 2-degenerate 1-island of size at most 7, a contradiction. Thus in this case, the final charge of v is at least 0, as desired.
- So we can assume that v is adjacent to at most 1 vertex of degree 2. If v has exactly one neighbour of degree 1, then as before we see that v can send charge through at most one of its degree 3 neighbours. If not, then these two paths from the degree 3 neighbours to the vertices of degree 2 plus v and the degree 2 neighbour form a 2-degenerate 1-island of size at most 10, which is a contradiction. Therefore in this case, again v ends up with final charge at least 0. This completes Case 3.

Case 4: $d(f) = 6.^{1}$

In this case, no vertex receives any charge from f, since otherwise f has a vertex of degree 2 and five vertices of degree 3, which implies that the face is a 2-degenerate 1-island of order 6.

Case 5: $d(f) \ge 7$.

For each vertex of degree 2 that receives a charge of $\frac{1}{2}$ from f in the positive orientation, let $A^+(v)$ be the set consisting of the five vertices of degree exactly 3 following v in the positive orientation of f (notice that these vertices exist). We define $A^-(v)$ similarly for the negative orientation of f. Then $A^+(v)$ and $A^-(u)$ are pairwise disjoint, otherwise there is a 2-degenerate 1-island of size at most 11. Therefore each face f sends a charge of at most $2 \times \frac{1}{2} \times \lfloor \frac{d(f)}{6} \rfloor$, and thus as $d(f) \ge 7$, we get that $\operatorname{ch}^*(f) \ge 0$. This completes Case 5 and hence the proof of the theorem.

¹The degree of a face being the number of vertices incident to it.

3 The proof of Theorem 1.3

3.1 Outline of the proof

We say that (G, α, β) is a special counterexample to Theorem 1.3 if

- G is a planar graph of girth 5,
- α is a 4-colouring of G and β is a 3-colouring of G, and
- there is no sequence of recolourings from α to β .

Note that every special counterexample is indeed a counterexample to Theorem 1.3. Moreover, because of the celebrated Grötszch's Theorem [17] (stating that every triangle-free planar is 3-colourable), if there is no special counterexample, then Theorem 1.3 holds. Therefore, in order to prove Theorem 1.3, it is sufficient to show that there is no special counterexample. For clarity, from now on, we will omit the adjective "special".

We say that a counterexample (G, α, β) is minimal if $|V(G)| \leq |V(H)|$ for any other counterexample (H, α', β') . We prove Theorem 1.3 by contradiction, showing that no minimal counterexample exists.

For a face f of G, the set of vertices incident with f is denoted V(f). Let v be a vertex of G of degree exactly 4 and let f be a 5-face incident with v. A 5-face $f' \neq f$ of G is opposite to f with respect to v if $V(f) \cap V(f') = \{v\}$ (note that since v has degree 4 and G has girth 5, f' is well defined). We say that a vertex u incident with a 5-face f is bad for f if either

- u has degree 3, or
- u has degree 4 and f has an opposite 5-face f' with respect to u such that each vertex in $V(f') \{u\}$ is bad for f'.

If v is bad for f, we say that (v, f) is a *bad pair*. Observe that bad pairs are well-defined inductively, and that the definition gives a natural quasi-order on bad pairs.

To prove Theorem 1.3, we show that every minimal counterexample cannot contain two infinite families of forbidden structures, namely:

- a 5-face f where all vertices are bad for f (Lemma 3.17)
- a 5-vertex v adjacent to four 5-faces f_1, f_2, f_3, f_4 such that each vertex of f_i (except v) is bad for f_i (Lemma 3.19).

We then apply a simple discharging argument to reach a contradiction with the existence of a minimal counterexample.

3.2 Structure of a minimal counterexample

In this section, we show a series of lemmas leading to the sought forbidden structures from Lemmas 3.17 and 3.19 (see the end of this section).

Lemma 3.1. Let (G, α, β) be a minimal counterexample. Then G is connected and has minimum degree at least 3.

Proof. Clearly, G is connected. Suppose that G has a vertex v of degree at most 2. By minimality of (G, α, β) , there is a recolouring sequence from $\alpha_{\uparrow(G-v)}$ to $\beta_{\uparrow(G-v)}$. This sequence can be lifted to a sequence in G by first recolouring v whenever a neighbour of v is recoloured to the colour of v (this is possible since the number 4 of colours implies there is always a colour not appearing on v or any of its neighbours). At the end of the sequence, we recolour v to $\beta(v)$, which is a contradiction.

We say that a vertex v of G is α -frozen if $\alpha(v) = \gamma(v)$ for every colouring γ obtainable from α by a sequence of recolourings. We say that v is α -locked if all colours $\{1, 2, 3, 4\}$ appear in the closed neighbourhood of v. Notice that if a vertex is frozen, then it is also locked, but the converse is not necessarily true. A vertex is said to be α -loose if it is not α -frozen and α -free if it is not α -locked. A locked vertex is said to be unlocked if it becomes free after recolouring at least one of its neighbours.

We have the following simple observations about frozen vertices. The first lemma is obvious.

Lemma 3.2. Let G be a graph, and let φ be a 4-colouring of G. If v is a vertex of G of degree 3 that is φ -locked, then recolouring one of the neighbours of v unlocks v.

The following simple consequence of the lemma will be used repeatedly.

Lemma 3.3. Let G be a graph, and let φ be a 4-colouring of G. Let $P = v_1 v_2 \dots v_k$ be a path in G such that v_1 is φ -free and, for $i \in \{2, \dots, k\}$, v_i is φ -locked and has degree 3. Then, for $j \in \{2, \dots, k\}$ there is a sequence of recolourings from φ to some colouring of G that unlocks v_j that only recolour vertices of P.

Proof. We recolour v_i , $1 \le i \le j-1$ in order, i.e. starting with v_1 and moving towards v_{j-1} . This is possible by Lemma 3.2.

Lemma 3.4. Let G be a graph, φ be a 4-colouring of G and v be a vertex of G with degree 3 that is φ -frozen. Then every neighbour of v is φ -frozen.

Proof. Suppose by contradiction that there is a neighbour w of v that is φ -loose. By definition of w, there is a recolouring sequence starting from φ which recolours w. Up to truncating the recolouring sequence and renaming w, one can assume that w is the only recoloured neighbour of v. Since v is φ -frozen and has degree precisely 3, recolouring w in turn unlocks v by Lemma 3.2, which is a contradiction.

Lemma 3.5. Let G be a graph, φ be a 4-colouring of G and v be a vertex of G with degree 4 that is φ -frozen. If v has two frozen neighbours with the same colour in φ , then all neighbours of v are φ -frozen.

Proof. Let u_1, u_2 be the two φ -frozen neighbours of v such that $\varphi(u_1) = \varphi(u_2)$, and let w_1 and w_2 be the other two neighbours of v. Suppose for a contradiction that one of w_1, w_2 , say w_1 , is φ -loose. So we can let σ be the first colouring obtainable from φ via a sequence of recolourings such that $\sigma(w_1) \neq \varphi(w_1)$.

Note that v being φ -frozen implies $\varphi(v), \varphi(u_1), \varphi(w_1), \varphi(w_2)$ are pairwise distinct. Thus, $\sigma(w_1) \in \{\varphi(u_1), \varphi(w_1), \varphi(w_2)\}$ and since u_1 and u_2 are frozen, $\sigma(u_i) = \varphi(u_i)$ for i = 1, 2. Hence, v has only two colours appearing in its neighbourhood, and hence is σ -loose, which is a contradiction. Using these properties, we can prove the following result. Let G be a plane graph with girth 5, and let f be a face of G. The set of neighbours adjacent to some vertex in V(f) (but not on f) is denoted by N(f). A 5-face f in G with $V(f) = \{v_1, v_2, v_3, v_4, v_5\}$ is bad if $d(v_1) \leq 4$ and $d(v_i) = 3$ for $i \in \{2, \ldots, 5\}$; it is very bad if moreover $d(v_1) = 4$.

Lemma 3.6. Let (G, α, β) be a minimal counterexample. Then G does not contain a bad 5-face with at least one vertex that is α -loose.

Proof. Suppose by contradiction that G contains a bad 5-face f with an α -loose vertex. By definition, α can be transformed to some 4-colouring φ such that f has a vertex that is φ -free.

For a colouring σ of G, we shall slightly abuse notation by saying the face f is σ -free if at least one vertex incident to f is σ -free. Our aim is to show that any recolouring sequence of G - V(f) can be lifted into a sequence in G, which is a contradiction to the minimality of (G, α, β) .

Let u_1 denote the vertex of degree at most 4 in f, and let u_2, \ldots, u_5 denote the vertices of degree precisely 3 on f in a clockwise ordering starting at u_1 around f. By Lemma 3.1, $d(u_1) \ge 3$. We require the following two claims.

Claim 3.7. If f is σ -free and $w \in \{u_1, \ldots, u_5\}$, then there is a recolouring sequence starting from σ that recolours only vertices in V(f) and unlocks w.

Proof of Claim 3.7. If w is σ -free, there is nothing to prove.

So we can assume that w is σ -locked. If $w \neq u_1$, then one of the two facial paths from w to u_1 goes through a free vertex (otherwise all vertices of f are locked). Let u be the first such vertex. We can apply Lemma 3.3 with the path from w to u, and the claim follows.

Therefore we can assume that $w = u_1$. Since f is σ -free and each $u_i \neq w$ has degree precisely 3, the same argument can be applied to show that for $i \in \{2, \ldots, 5\}$ each u_i can be unlocked by recolouring vertices in V(f) - w and if $d(u_1) = 3$, then $w = u_1$ may also be unlocked. To complete the proof, it remains to show that w can also be unlocked if d(w) = 4 by recolouring vertices in V(f).

Let w_1 and w_2 denote the two neighbours of w outside f, and suppose $\sigma(w) = 1$, $\sigma(w_1) = 3$ and $\sigma(w_2) = 4$. Our aim is to try to recolour w to colour 2 (by possibly first recolouring only vertices in V(f) - w). If $\sigma(u_2), \sigma(u_5) \neq 2$, then w is σ -free, a contradiction. Therefore, either $\sigma(u_2) = 2 \neq \sigma(u_5)$ or $\sigma(u_2) = \sigma(u_5) = 2$. We address the two cases separately.

Case 1: $\sigma(u_2) = 2 \neq \sigma(u_5)$.

Since f is σ -free, we can let $i \in \{2, \ldots, 5\}$ be the smallest index such that u_i is σ -free. If i = 2, then we recolour u_2 which, in turn, unlocks v. And if $i \neq 4$, then recolouring u_j unlocks u_{i-1} for $i \ge j \ge 1$ in order as needed.

It only remains to address the case i = 5. We first try to recolour u_5 with a colour distinct from 2; this in turn unlocks u_4 and by recolouring u_4 vertex u_3 is unlocked etc. until u_1 is unlocked. So we can assume that we can only recolour u_5 with colour 2. Without loss of generality, set $\sigma(u_5) = 3$. Since u_4 is locked, $\sigma(u_3) \neq 3 = \sigma(u_5)$ and since u_2 is locked, $\sigma(u_3) \neq 1 = \sigma(u_1)$. Therefore, $\sigma(u_3) = 4 \neq \sigma(u_2) = 2$. Since u_3 is locked, $\sigma(u_4) \neq 2 = \sigma(u_2)$ and thus $\sigma(u_4) = 1$. By applying the recolouring sequence shown in Figure 1, w may be recoloured to 2. This completes Case 1.



Figure 1: Recolouring sequence from Case 1 of Claim 3.7.

Case 2: $\sigma(u_2) = \sigma(u_5) = 2$,

If either u_2 or u_5 , say u_2 , is σ -free, then we recolour u_2 with a colour $\neq 2$ and apply Case 1.

Otherwise, since f is σ -free, either u_3 or u_4 is σ -free. Assume without loss of generality that u_3 is σ -free; then recolouring u_3 unlocks u_2 , in which case the argument from the preceding paragraph can be applied. This completes Case 2 and hence the proof of the claim.

Claim 3.7 will allow us to prove the following claim, from which we will derive Lemma 3.6.

Claim 3.8. If f is σ -free, $z \in V(f)$ and $a \in \{1, 2, 3, 4\}$, then there is a recolouring sequence from σ to a colouring σ' that recolours only vertices in V(f) such that $\sigma'(z) \neq a$ and some vertex in $\{u_2, \ldots, u_5\}$ is σ' -free.

Proof of Claim 3.8. We distinguish three cases (the cases $z = u_4$ and $z = u_5$ are symmetric to respectively $z = u_3$ and $z = u_2$).

Case 1: $z = u_1$.

By Claim 3.7, we can transform σ to some colouring σ' by only recolouring vertices in V(f) so that z is σ -free and so by recolouring z if necessary we can further assume $\sigma'(z) \neq a$. If some vertex in $\{u_1, \ldots, u_5\} \setminus \{z\}$ is σ' -free, then the claim follows. So we can assume that each u_i is σ' -locked. By symmetry, we assume that the neighbours of z outside f are coloured with 2 and 3. Thus, $\sigma'(u_2) = \sigma'(u_5) = 2$ and so $\sigma'(u_3) = 3$, $\sigma'(u_4) = 4 \neq \sigma'(z) = 1$.

Now we proceed with the recolouring sequence shown in Figure 2 (note that, at the end of the sequence, the colour of z is 1 while u_2 is free, as needed).

Case 2: $z = u_2$.

By Claim 3.7, we can transform σ to a colouring σ' by only recolouring vertices in V(f)so that u_1 is σ' -free and by recolouring u_1 if necessary we can further assume that zis σ' -free. If $\sigma'(z) \neq a$ then we are done. Suppose $\sigma'(z) = a$. If u_3 is σ' -locked, then recolouring $z = u_2$ in turn unlocks u_3 by Lemma 3.2, and we are done. If, on the other hand, u_3 is σ' -free, then we recolour u_3 if necessary to unlock u_4 . If, at this point, u_2 is locked, then we can recolour u_1 which is still unlocked, and then u_2 becomes free. We finish the sequence by recolouring $z = u_2$ (to a colour distinct from a) and as u_4 is still free the claim follows. This completes Case 2.



Figure 2: Recolouring sequence that unlocks vertices without changing the colour of v.

Case 3: $z = u_3$,

As before, by Claim 3.7, we can transform σ to a colouring σ' by only recolouring vertices in V(f) so that u_3 is σ' -free and $\sigma'(u_3) \neq a$. Assume that each vertex in $\{u_2, \ldots, u_5\}$ is σ' -locked (else we are done).

Now recolouring u_3 unlocks both u_2 and u_4 , and in turn recolouring u_4 unlocks u_5 . If, at this point, u_3 is locked, we recolour u_2 to unlock u_3 and then recolour u_3 if necessary so that its colour is distinct from a. Since u_5 is free, Case 3 is complete.

We can now finish the proof of the Lemma 3.6. By the minimality of (G, α, β) , there is a sequence s' of recolourings from $\alpha_{\restriction (G-V(f))}$ to $\beta_{\restriction (G-V(f))}$.

We lift s' to a sequence from α to β in G as follows. Each time a vertex $v \in N(f)$ is recoloured to the current colour a of some vertex $z \in V(f)$, we precede the recolouring of v by changing the colour of z to a colour distinct from a and such that another vertex of $\{u_2, \ldots, u_5\}$ is unlocked; this is always possible by Claim 3.8 since (by hypothesis), f contains a free vertex. Observe that this operation only changes the colours of vertices in V(f) and leaves f with a free vertex in $\{u_2, \ldots, u_5\}$. (Note that v indeed has at most one neighbour in f by girth assumption). This shows that s' lifts to a sequence from α to some colouring β' of G such that $\beta'_{|(G-V(f))|} = \beta_{|(G-V(f))|}$ and some vertex in $\{u_2, \ldots, u_5\}$ is β' -free. To finish the proof, we describe a sequence of recolourings from β' to β that recolours only vertices in V(f).

From β' we recolour as many vertices as possible in V(f) to colour 4 (recall that β uses only colours 1, 2 and 3) and let β'' denote the resulting colouring; notice that $\beta''(u_j) = \beta''(u_{j+2}) = 4$ for some $j \in [5]$ (here $u_6 = u_1$ and $u_7 = u_2$). From β , we recolour u_j and u_{j+2} to colour 4 and denote the resulting colouring β^* . Note that u_{j+1} is free in both β'' and β^* unless possibly if $u_{j+1} = u_1$, in which case $\beta''(u_1) = \beta^*(u_1)$. Therefore, by recolouring u_{j+1} if necessary from β'' , we can assume $\beta''(u_{j+1}) = \beta^*(u_{j+1})$. The only cause of difficulty is when $\beta^*(u_{j+3}) = \beta''(u_{j+4})$, $\beta^*(u_{j+4}) = \beta''(u_{j+3})$ and u_{j+4} and u_{j+3} are locked in β^* (all the other cases are easy and left to the reader). In this case, assuming without loss of generality that u_j is the vertex of degree 3, we recolour, from β^* , u_j to colour $\beta(u_j)$, then recolour $u_{j-1} = u_{j+4}$ to colour 4, followed by recolouring $u_{j-2} = u_{j+3}$ to its colour $\beta^*(u_{j+4})$ in β' and finally u_{j+4} to $\beta^*(u_{j+3})$. We finish the sequence by recolouring u_j to colour 4. The proof of the lemma is complete.

In order to prove the next lemma, we make use of the following observation (see also



Figure 3: The unique locked 4-colouring of a 5-face from Lemma 3.9.

Figure 3).

Lemma 3.9. Let f be a 5-face in a plane graph G with girth 5 and $v \in V(f)$. Assume that all vertices of f but v have degree 3. If d(v) = 3, there is no 4-colouring φ of $G[V(f) \cup N(f)]$ such that every vertex in V(f) is φ -locked. Otherwise $d(v) \ge 4$, and there is a unique such colouring (see Figure 3).

The (easy) proof of Lemma 3.9 follows by case analysis and is left to the reader. We are now ready to prove Lemma 3.10.

Lemma 3.10. Let (G, α, β) be a minimal counterexample. Then G does not contain four 5-faces f_1, \ldots, f_4 such that $\bigcap_{i=1}^4 V(f_i)$ contains a vertex v of degree 5 and every vertex in $\bigcup_{i=1}^4 V(f_i) - \{v\}$ has degree 3 (see Figure 4).



Figure 4: Configuration from Lemma 3.10. Numbers indicate degree.

Proof. Suppose otherwise, and let $H = G[\bigcup_{i=1}^{4} V(f_i)]$. For $i \in \{1, \ldots, 4\}$ write $V(f_i) = \{v, u_i, u_{i+1}, z_i, z'_i\}$, where u_i denotes the neighbour of v incident with f_{i-1} and f_i , and $u_i z_i, z_i z'_i, z'_i u_{i+1} \in E(f_i)$.

By the minimality of (G, α, β) , there is a sequence s' of recolourings from $\alpha_{\restriction (G-H)}$ to $\beta_{\restriction (G-H)}$. To reach a contradiction, we show how to lift s' to a sequence from α to β in G. Fix a 4-colouring σ of G.

Claim 3.11. For i = 1, 2, there is a σ -free vertex incident to f_i , f_{i+1} or f_{i+2} .

Proof of Claim 3.11. Suppose by contradiction that all vertices incident to f_i, f_{i+1} and f_{i+2} are σ -locked. By Lemma 3.9, $\sigma(u_i) = \sigma(u_{i+1}), \sigma(u_{i+1}) = \sigma(u_{i+2})$ and $\sigma(u_{i+2}) = \sigma(u_{i+2})$

 $\sigma(u_{i+3})$. Thus only two colours appear in the neighbourhood of v, i.e., v is σ -free, which is a contradiction.

Claim 3.12. There is a recolouring sequence from σ to some colouring σ' such that $\sigma_{\uparrow(G-H)} = \sigma'_{\uparrow(G-H)}$ and $\sigma(v) \neq \sigma'(v)$.

Proof of Claim 3.12. By Claim 3.11, there is a σ -free vertex h incident with f_1, f_2 , or f_3 . Let P denote a shortest path whose internal vertices entirely lie in $V(f_1 \cup f_2 \cup f_3) - \{v\}$ and having ends h and u_1 . Applying Lemma 3.3 to P, we can assume u_1 is σ -free. Similarly, we can assume that u_5 is σ -free. (Note that this keeps u_1 free since u_1 has no neighbour in $V(f_2 \cup f_3 \cup f_4) - \{v\}$). This, in turn, yields $\sigma(u_1) = \sigma(u_5)$. Indeed, if they have different colours and none can be recoloured into the other, they can both be recoloured into $\{1, 2, 3, 4\} \setminus \{\sigma(u_1), \sigma(u_5), \sigma(v)\}$. To prove the claim, we must change the colour of v. We consider two cases.

Case 1: f_2 and f_3 are both σ -locked.

By Lemma 3.9, $\sigma(u_2) = \sigma(u_3) = \sigma(u_4)$. Consequently, only two colours appear in the neighbourhood of v since $\sigma(u_1) = \sigma(u_5)$ by assumption, i.e., v is σ -free so we simply recolour v.

Case 2: there is a σ -free vertex incident to either f_2 or f_3 .

By applying Lemma 3.3 as before, we can assume that u_j is σ -free for some $j \in \{2, 3, 4\}$. Let C be the set of colours appearing on the σ -locked neighbours of v. Then $|C| \leq 2$ by our observations this far and so there is a colour $a \in [4] \setminus (C \cup \{\sigma(v)\})$. We simply recolour the three σ -free neighbours u_1, u_5 and u_j of v if their colour is a and finish the sequence by recolouring v to a. This completes Case 2 and the proof of the claim. \Box

Claim 3.13. There is a recolouring sequence from σ to some colouring σ' such that $\sigma_{\uparrow(G-H)} = \sigma'_{\uparrow(G-H)}$ and $\sigma'(v) \in \{\sigma'(z'_1), \sigma'(z_4)\}.$

Proof of Claim 3.13. Since G is planar, there cannot be an edge between z'_1 and $\{z_3, z'_3\}$ and one between z_4 and $\{z_2, z'_2\}$. Up to exchanging z'_1 and z_4 , assume that the former holds.

Let x be the neighbour of z'_1 which is not z_1 or u_2 . By hypothesis $x \notin \{z_3, z'_3\}$, and since G has girth 5, $x \notin \{z_2, z'_2\}$. If $x \notin H$, then by Claim 3.12, we can assume that $\sigma(x) \neq \sigma(v)$. If $x = z_4$, then $\sigma(x) \neq \sigma(v)$ otherwise the result follows. If $x = z'_4$, we can also assume that $\sigma(x) \neq \sigma(v)$ otherwise u_4 is σ -free, and up to recolouring u_4 , x becomes free and we can recolour it. So we can assume that $\sigma(x) \neq \sigma(v)$.

We try to immediately recolour z'_1 to colour $\sigma(v)$. If this is not possible, then either $\sigma(z_1)$ or $\sigma(u_2)$ is $\sigma(v)$, and as $u_2v \in E(G)$, $\sigma(z_1) = \sigma(v)$. Therefore u_1 is σ -free, and up to applying Lemma 3.3, we can assume that z_1 is σ -free; so we recolour z_1 to a colour different from $\sigma(v)$ and then finally recolour z'_1 with $\sigma(v)$. The claim is proved. \Box

Claim 3.14. For every $w \in V(G - H)$, there is a recolouring sequence from σ to some colouring σ' such that $\sigma_{\uparrow(G-H)} = \sigma'_{\uparrow(G-H)}$ and every vertex in $N_G(w) \cap V(H)$ is σ' -free.

Proof of Claim 3.14. Set $Z_1 = \{u_1, z_1, z_1'\}, Z_2 = \{z_2, z_2'\}, Z_3 = \{z_3, z_3'\}$ and $Z_4 = \{z_4, z_4', u_5\}$. Note that $N_H(w) \subseteq \bigcup_{i=1}^4 Z_i$; moreover, since G has girth 5, w can have at most one neighbour in each Z_i that we denote w_i . By Claim 3.13, we can assume that $\sigma(z'_1) = \sigma(v)$, in turn implying u_2 is σ -free. For $i \in [4]$, let P_i denote the path in H with ends u_2 and w_i and that does not contain v. We apply Lemma 3.3 to subpaths of each P_i (with P_i playing the role of P) starting with P_4 and working our way backwards to P_1 to obtain a sequence of recolourings that unlocks each w_i . At each step, u_2 stays σ -free since z'_1 and v keep their colour. Therefore, if w_i have not become free, one can always find a subpath of P_i ending at w_i and satisfying the hypotheses of Lemma 3.3. This implies the claim.

We can now finish the proof of Lemma 3.10. By the minimality of (G, α, β) , there is a sequence s' of recolourings from $\alpha_{\uparrow (G-H)}$ to $\beta_{\uparrow (G-H)}$.

We lift s' to a sequence from α to some colouring β' of G by, for each vertex $w \in V(G-H)$, recolouring each neighbour w' of w in H whenever w gets recoloured to the colour of w'; this is possible by Claim 3.14 and gives that $\beta'_{\restriction (G-H)} = \beta_{\restriction (G-H)}$. We now conclude the proof by recolouring β' to β in a similar fashion to Lemma 3.6. Note that the only possible vertices coloured 4 are in H since β is a 3-colouring.

We consider two cases depending on whether v has colour 4 or not. Denote by z'_0 (resp. z_5) the neighbour of u_1 (resp. u_5) not in $\{v, z_1\}$ (resp. $\{v, z'_4\}$). Note that z'_0 and z_5 are not in H since otherwise G would contain a C_4 .

Case 1: v has colour 1, 2 or 3.

We first show how to recolour H so that all the u_i 's get colour 4. Assume that this is not the case, and that there exists an index i such that u_i has colour 1, 2 or 3. We prove that we can recolour H in such a way that either the number of u_i 's coloured with 4 increases, or it does not decrease but the number of z_i and z'_i coloured with 4 decreases.

If z_i and z'_{i-1} are not coloured 4, then one can directly recolour u_i with 4, which concludes.

Otherwise, assume by symmetry that z'_{i-1} has colour 4. Then let j < i be the largest index such that z'_{j-1} is not coloured 4 (which exists since z'_0 , which is not in H, has colour 1, 2 or 3). Now z'_j has colour 4, hence z_j does not. Now u_j can be recoloured to 4, which increases the number of u_ℓ 's coloured with 4. So we can assume that u_j is coloured 4. Hence z_j is free (since it has two neighbours coloured with 4) and then can be recoloured with a colour distinct from 4. So either we can immediately recolour z'_j or recolour it after the recolouring of z_j . So we can recolour H in such a way z'_j is not coloured 4, which does not change the number of vertices u_ℓ coloured with 4, but the number of z_ℓ, z'_ℓ coloured with 4 decreased.

By iterating this argument, we obtain a colouring of H where colour 4 lies exactly on the u_i 's. We now claim that the following (\star) property holds: if z_i, z'_i are both locked, we can flip the colours of z_i and z'_i by only recolouring vertices among $\{v, u_i, u_{i+1}, z_i, z'_i\}$. Indeed, if that is the case, the neighbours of z_i and z'_i outside of $\{v, u_i, u_{i+1}, z_i, z'_i\}$ have the same colour. So we can recolour v with the colour of z_i , then recolour u_i with a free colour, then z_i with 4, then z'_i with the previous colour of z_i . Now, up to recolouring v(which only has two colours in its neighbourhood) and freeing u_i , we can colour z_i with the initial colour of z'_i and put back colour 4 on u_i .

Now let us prove that for increasing $i \leq 4$, we can colour z_i, z'_i with their target colour. Note that these target colours do not appear on u_i, u_{i+1} since they are coloured with 4 and β is a 3-colouring. If they appear on another neighbour in H, it should be in $\{z_j, z'_j\}$ for some $j \neq i$. Moreover there is at most one edge between $\{z_i, z'_i\}$ and $\{z_j, z'_j\}$ since *G* has girth 5. Now, using property (\star) if needed, we can recolour z_j, z'_j so that they are not coloured anymore with the target colour of their neighbour in $\{z_i, z'_i\}$. Note that this happens only if j > i since β is a proper colouring and the vertices z_j, z'_j with j < i already have their final colour. In particular, these recolouring steps do not recolour the vertices z_j, z'_j for j < i.

So we can assume that the final colour of z_i (resp. z'_i) does not appear on its neighbourhood except maybe on z'_i (resp. z_i). So we can directly colour z_i, z'_i with their target colour unless both are locked and their target colours are then flipped. But then we apply again (\star) to swap their colours.

We finally set v to $\beta(v)$, then each u_i to $\beta(u_i)$.

Case 2: v has colour 4.

We first recolour vertices to 4 as long as it is possible. We then claim that if u_i and u_{i+1} are locked, we can recolour them (together with z_i, z'_i) so that they have the same colour. Indeed, assume that u_i is coloured with 1 and u_{i+1} with 2. Note that both z_i and z'_i must have a neighbour with colour 4 by maximality of β' . Now either z_i is coloured with 2 or z'_i is coloured with 1 (since both cannot have colour 3). By symmetry we consider only the latter. Either z_i is coloured 2 and we recolour it with 3 so that u_i can be recoloured with 2, or z_i is coloured 3, and we recolour it to 2, so that z'_i can be recoloured to 3 and u_{i+1} to 1. Afterwards, we recolour again vertices with 4 whenever possible.

Therefore, if colours 1, 2 and 3 appear on locked u_i 's, it must be on u_1, u_3 and u_5 . By maximality of β' , the vertices coloured with 4 among $\{z'_1, z_2, z'_3, z_4\}$ must dominate $\{z_1, z'_2, z_3, z'_4\}$. Due to planarity and girth constraints, this implies that at least three vertices among $\{z'_1, z_2, z'_3, z_4\}$ are coloured with 4. Therefore, either u_2 or u_4 (say u_2 by symmetry) sees only colour 4 in its neighbourhood. In that case, either z_1 sees two vertices coloured with 4, and we can recolour it, which unlocks u_1 , or we recolour u_2 so that z'_1 becomes free, then recolour z'_1 and finally recolour z_1 with 4. In both cases, only u_3 and u_5 can be locked in the obtained colouring.

Therefore, we can assume that at most two colours appear on the locked u_i 's. Let c be one of the remaining colours. Now, observe that all the u_i 's coloured with c are unlocked, hence they can be recoloured. We may now recolour v to c, so that we can apply Case 1.

We now will investigate the properties of bad vertices and bad pairs, which were defined in Section 3.1.

Bad vertices have the following property.

Lemma 3.15. Let (G, α, β) be a minimal counterexample. If v is a bad vertex of G for some 5-face f and has degree 4, then v is α -frozen and both neighbours of v which are not incident to f are also α -frozen and are coloured alike under α .

Proof. We prove the result by induction on the pairs (v, f) such that v has degree 4 and is bad for f. Let v be a vertex of degree 4 which is bad for the face f.

Let us assume that all the bad pairs smaller than (v, f) satisfy the statement of the lemma. Let f' be the face opposite of f at v.

Claim 3.16. All the vertices of f' are frozen, and if $z \neq v$ is incident to f', then its two neighbours in f' have different colours.

Proof. In order to prove the first part of this claim, we show that:

- f' contains at least one frozen vertex;
- if z is a frozen vertex of f' different from v, then its neighbours are also frozen.

If all the vertices of f' different from v have degree 3, then f' is very bad and the result of the claim follows immediately from Lemma 3.6. On the other hand, vertices of degree 4 different from v and incident to f' are frozen using the induction hypothesis, hence the first point holds.

In order to prove the second point, consider a frozen vertex z incident to f' and different from v. If z has degree 3, then the neighbours of z are frozen by Lemma 3.4 and in particular its neighbours on f' have different colours. If z has degree 4, then since we know that z is bad for f', by the induction hypothesis, the two neighbours of z which are not incident to f' are frozen and have the same colour. By Lemma 3.5, this implies that all the neighbours of z are frozen. Moreover, since the neighbours of z outside of f' have the same colour, its neighbours incident to f' have different colours.

Hence, the claim above shows that v is frozen. Moreover, the two neighbours of v incident to f' must have the same colour since otherwise it is not possible to colour the vertices of f' with 4 colours such that the property of the claim above holds. This proves the induction step and finishes the proof of the lemma.

We now have all the tools needed to forbid the structures presented in Subsection 3.1.

Lemma 3.17. The graph G does not contain a 5-face f such that every vertex on f is bad for f.

Proof. Let us assume by contradiction that G has a face f satisfying this property. By Lemmas 3.6 and 3.9, f contains a vertex of degree 4, and all of them are frozen by Lemma 3.15. Moreover, as before, if z is a frozen vertex of f, then its neighbours are also frozen. This follows from Lemma 3.4 if z has degree 3, and from Lemma 3.5 and Lemma 3.15 if z has degree 4. Let u and v be two vertices incident to f. We will show that u and v have different colours. This is trivially true if u and v are adjacent. If they are not adjacent, then they have a common neighbour w incident to f. If w has degree 3, then since w is locked all its neighbours have different colours, and in particular $\alpha(u) \neq \alpha(v)$. If w has degree 4, then by Lemma 3.15, its two neighbours not incident to f have the same colour, and consequently u and v must have different colours. Hence, all the vertices incident to f have different colours in α , which contradicts the assumption that α is a 4-colouring of G.

An important consequence of Lemma 3.17 is stated in the following corollary.

Corollary 3.18. Let v be a vertex of G of degree 4, incident to the faces f_1, f_2, f_3, f_4 (in clockwise order). Then v cannot be bad for both f_i and $f_{i+2 \mod 4}$. In particular, every bad 4-vertex is bad for at most two faces.

Proof. Let us assume by contradiction that v is bad for both f_1 and f_3 . Since v is bad for f_1 , then by definition all the vertices of f_3 different from v are bad for f_3 . Since v is also bad for f_3 , we obtain a contradiction with Lemma 3.17.

Lemma 3.19. The graph G does not contain a 5-vertex v adjacent to four 5-faces f_1, f_2, f_3, f_4 such that, for each i, all vertices incident to f_i except v are bad for f_i .

Proof. Let us assume by contradiction that G contains a vertex v with this property. Let H be the graph induced by the vertices incident to the faces f_i . The case where all the vertices of H different from v have degree 3 has already been handled in Lemma 3.10. Hence we can assume that at least one vertex incident to f_i has degree 4. By Lemma 3.15 we know that this vertex is frozen, and so are its neighbours. Moreover, by Remark 3.4, all the neighbours of a frozen vertex of degree 3 are also frozen. Hence, it follows that all the vertices of H are frozen. However, this implies that the colouring α is locked on H, which is not possible by Claim 3.11.

3.3 Discharging

We may now reach a contradiction using a discharging argument. We first give an initial weight of 2d(v) - 6 to each vertex v and $\ell(f) - 6$ to each face f. According to Euler's formula, the total weight is

$$\sum_{v \in V(G)} (2d(v) - 6) + \sum_{f \in F(G)} (\ell(f) - 6) = -12.$$

We then redistribute the weights according to the following rules :

- Every 4-vertex which is bad for at least two faces gives 1 to their opposite faces.
- Every 4-vertex which is bad for only one face f gives 1 to its opposite face, and then splits its remaining weight equally among its two remaining incident 5-faces.
- Every 4-vertex which is not bad and every 5⁺-vertex v gives 1 to each incident 5-face f such that all the vertices of $f \setminus \{v\}$ are bad, and $\frac{1}{2}$ to every other 5-face.

We finally reach a contradiction by showing that every vertex and face ends up with non-negative weight after this process. This contradicts that the sum of the weights is negative and proves Theorem 1.3. First recall that every vertex of G has degree at least 3, and that 3-vertices and 6⁺-faces have non-negative initial weight and do not lose any weight. Therefore, we only have to consider 4⁺-vertices and 5-faces.

We may easily take care of vertices of degree $d \ge 6$: they give at most 1 to each incident face, hence they end up with at least

$$2d - 6 - d = d - 6 \ge 0.$$

Lemma 3.20. Every 5-vertex ends up with non-negative weight.

Proof. By Lemma 3.19, we know that every 5-vertex gives 1 to at most 3 faces and at most $\frac{1}{2}$ to every other face, hence their final weight is at least $4 - 3 \times 1 - 2 \times \frac{1}{2} = 0$. \Box

Lemma 3.21. Every 4-vertex ends up with non-negative weight.

Proof. Let v be a 4-vertex and f_1, f_2, f_3, f_4 be the faces incident to v. If v is bad, then by Corollary 3.18, it is bad for at most two faces. Then either v is bad for two faces and loses 2 by the first rule, or it is bad for one face, and distributes all its weight by the second rule. In both cases, bad 4-vertices end up with weight 0.

We may thus assume that v is not bad. If v gives 1 to an incident face, say f_1 , then f_3 cannot have length 5 otherwise v would be bad for f_3 , and v would not give any weight to f_3 . Otherwise, v gives at most $2 \times \frac{1}{2}$ to f_1 and f_3 . Therefore, v gives at most 1 to $\{f_1, f_3\}$, and at most 1 to $\{f_2, f_4\}$ by symmetry. Then the final weight of v is at least $2 - 2 \times 1 = 0$. This completes the proof of the lemma.

Lemma 3.22. Every 5-face receives a weight of at least 1.

Proof. Let f be a 5-face. By Lemma 3.17, f contains at most four bad vertices for f. Let v be a vertex on f which is not bad for f. If v has degree at least 5 or is a 4-vertex which is not bad, then it gives at least $\frac{1}{2}$ to f by the third rule. If v has degree 4 and is bad, let f, f_1, f', f_2 be the faces incident to v in clockwise order. Observe that if v is bad for f', then it gives 1 to f since f is opposite to f' with respect to v. Otherwise, since v is bad but not bad for f by construction, then v is bad for f_1 or f_2 (but not both since f_1 and f_2 are opposite with respect to v, see Corollary 3.18). Then the second rule applies, and f receives at least $\frac{1}{2}$ from v.

Therefore, every vertex on f which is not bad for f gives at least $\frac{1}{2}$ to f. We may thus assume that f has only one such vertex v. In that case, f receives 1 from v by the third rule unless v is a bad 4-vertex. In that case, again let f, f_1, f', f_2 be the faces incident to v in clockwise order. If f' has length 5, then since v is the only vertex in f which is not bad for f, it implies that v is bad for f' and thus gives 1 to f by the first or second rule since f is a face opposite to f' with respect to v. Otherwise, f' has length at least 6, and v is bad for f_1 or f_2 (but not both, again by Corollary 3.18). In that case, f receives also 1 by the second rule. This completes the proof of the lemma.

Acknowledgements

The authors are grateful to the referee for helpful feedback that improved the presentation of the paper.

References

- V. Bartier, N. Bousquet, and M. Heinrich. Recoloring graphs of treewidth 2. Discrete Mathematics, 344(12):112553, 2021.
- [2] M. Bonamy and N. Bousquet. Recoloring graphs via tree decompositions. *Eur. J. Comb.*, 69:200–213, 2018.
- [3] M. Bonamy, M. Johnson, I. Lignos, V. Patel, and D. Paulusma. Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs. *Journal of Combinatorial Optimization*, pages 1–12, 2012.

- [4] O. Borodin. Solving the Kotzig and Grünbaum problems on the separability of a cycle in planar graphs. *Mat. Zametki*, 46:9–12, 1989.
- [5] N. Bousquet and V. Bartier. Linear transformations between colorings in chordal graphs. In 27th Annual European Symposium on Algorithms, ESA 2019, pages 24:1– 24:15, 2019.
- [6] N. Bousquet and M. Heinrich. A polynomial version of Cereceda's conjecture. arXiv preprint arXiv:1903.05619, 2019.
- [7] N. Bousquet and G. Perarnau. Fast recoloring of sparse graphs. *European Journal* of Combinatorics, 52:1–11, 2016.
- [8] L. Cereceda. *Mixing Graph Colourings*. PhD thesis, London School of Economics and Political Science, 2007.
- [9] L. Cereceda, J. van den Heuvel, and M. Johnson. Mixing 3-colourings in bipartite graphs. *Eur. J. Comb.*, 30(7):1593–1606, 2009.
- [10] S. Chen, M. Delcourt, A. Moitra, G. Perarnau, and L. Postle. Improved bounds for randomly sampling colorings via linear programming. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 2216– 2234, 2019.
- [11] Z. Dvořák and C. Feghali. An update on reconfiguring 10-colorings of planar graphs. The Electronic Journal of Combinatorics, pages P4–51, 2020.
- [12] Z. Dvořák and C. Feghali. A Thomassen-type method for planar graph recoloring. European Journal of Combinatorics, 95:103319, 2021.
- [13] M. Dyer, A. D. Flaxman, A. M. Frieze, and E. Vigoda. Randomly coloring sparse random graphs with fewer colors than the maximum degree. *Random Structures & Algorithms*, 29(4):450–465, 2006.
- [14] L. Esperet and P. Ochem. Islands in graphs on surfaces. SIAM Journal on Discrete Mathematics, 30(1):206–219, 2016.
- [15] C. Feghali. Reconfiguring colorings of graphs with bounded maximum average degree. Journal of Combinatorial Theory, Series B, 147:133 – 138, 2021.
- [16] A. Frieze and E. Vigoda. A survey on the use of Markov chains to randomly sample colourings. Oxford Lecture Series in Mathematics and its Applications, 34:53, 2007.
- [17] H. Grötzsch. A three-color set for three-circle-free nets on the sphere. science Z. Martin Luther Univ. Halle-Wittenberg, Math. Nat. Line, 8:109–120, 1959.
- [18] N. Nishimura. Introduction to reconfiguration. Algorithms, 11(4):52, 2018.
- [19] J. van den Heuvel. The Complexity of change, page 409. Part of London Mathematical Society Lecture Note Series. S. R. Blackburn, S. Gerke, and M. Wildon edition, 2013.