WARWICK
THE UNIVERSITY OF WARWICK

Manuscript version: Author's Accepted Manuscript
The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:
http://wrap.warwick.ac.uk/132230

How to cite:
Please refer to published version for the most recent bibliographic citation information.
If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:
The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Publisher's statement:
Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

RECONFIGURING GRAPH HOMOMORPHISMS ON THE SPHERE

JAE-BAEK LEE, JONATHAN A. NOEL, AND MARK SIGGERS

ABSTRACT. Given a loop-free graph H, the reconfiguration problem for homomorphisms to H (also called H-colourings) asks: given two H-colourings f of g of a graph G, is it possible to transform f into g by a sequence of single-vertex colour changes such that every intermediate mapping is an H-colouring? This problem is known to be polynomial-time solvable for a wide variety of graphs H (e.g. all C_4-free graphs) but only a handful of hard cases are known. We prove that this problem is PSPACE-complete whenever H is a $K_{2,3}$-free quadrangulation of the 2-sphere (equivalently, the plane) which is not a 4-cycle. From this result, we deduce an analogous statement for non-bipartite $K_{2,3}$-free quadrangulations of the projective plane. This include several interesting classes of graphs, such as odd wheels, for which the complexity was known, and 4-chromatic generalized Mycielski graphs, for which it was not.

If we instead consider graphs G and H with loops on every vertex (i.e. reflexive graphs), then the reconfiguration problem is defined in a similar way except that a vertex can only change its colour to a neighbour of its current colour. In this setting, we use similar ideas to show that the reconfiguration problem for H-colourings is PSPACE-complete whenever H is a reflexive K_4-free triangulation of the 2-sphere which is not a reflexive triangle. This proof applies more generally to reflexive graphs which, roughly speaking, resemble a triangulation locally around a particular vertex. This provides the first graphs for which H-RECOLOURING is known to be PSPACE-complete for reflexive instances.

1. INTRODUCTION

All graphs in this paper are assumed to be finite, undirected and without multiple edges, unless otherwise specified. A vertex is said to be *reflexive* if it has a loop and *irreflexive* otherwise. A graph is said to be *reflexive* if all of its vertices are reflexive and *irreflexive* if all of its vertices are irreflexive. A *homomorphism* from a graph G to a graph H is a mapping $f : V(G) \to V(H)$ such that $f(u)f(v) \in E(H)$ for every $uv \in E(G)$. For the sake of brevity, a homomorphism f from G to H will sometimes be referred to as an H-*colouring* of G and, for $v \in V(G)$, we call $f(v)$ the *colour* of v. We denote the set of all H-colourings of a graph G by $\mathrm{Hom}(G, H)$.

Given an irreflexive graph H and two H-colourings f and g of a graph G, a *reconfiguration sequence* taking f to g is a sequence $f_0, \ldots, f_m \in \mathrm{Hom}(G, H)$ such that $f_0 = f$, $f_m = g$ and f_i differs from f_{i+1} on a unique vertex for $0 \le i \le m-1$. If there exists a reconfiguration sequence taking f to g, then we say that f *reconfigures* to g. We are interested in the complexity of the following decision problem, called H-RECOLOURING:

> **Instance:** A graph G and $f, g \in \mathrm{Hom}(G, H)$.
> **Question:** Does f reconfigure to g?

The H-RECOLOURING problem is part of a growing area known as "combinatorial reconfiguration," a central focus of which is to determine the complexity of deciding whether a given solution to a combinatorial problem can be transformed into another by applying a sequence of allowed modifications.

For further background on combinatorial reconfiguration in general, see [1, 9–11, 14, 16, 20, 21] and the surveys of van den Heuvel [13], Ito and Suzuki [15] and Nishimura [23].

An interesting special case of the H-Recolouring problem is when H is a complete graph on k vertices, in which case H-colourings are nothing more than proper k-colourings. Cereceda, van den Heuvel and Johnson [8] showed that K_3-Recolouring can be solved in time $O(|V(G)|^2)$. This came as some surprise, given that it is NP-complete to decide whether a graph admits a K_3-colouring. On the other hand, Bonsma and Cereceda [2] showed that the complexity jumps drastically for larger cliques: for every fixed $k \geq 4$, the K_k-Recolouring problem is PSPACE-complete.[1] Later, Brewster, McGuinness, Moore and Noel [5] extended this dichotomy to the case when H is a "circular clique."

Wrochna [26] developed ideas inspired by algebraic topology to prove the remarkably general result that H-Recolouring is solvable in polynomial time whenever H does not contain a cycle of length 4. By further refining his topological approach, Wrochna [27] (see also [28]) proved a "multiplicativity" result for graphs without cycles of length 4 which is closely connected to Hedetniemi's Conjecture [12]; very recently, Tardif and Wrochna [24] have extended these methods beyond the setting of C_4-free graphs.

On the hardness side, however, only a few examples are known. As we have mentioned, the results of [2, 5] show that the problem is PSPACE-complete for certain cliques and circular cliques. In addition, Wrochna [25] proved that there exists a graph H such that H-Recolouring is PSPACE-complete even when the instance graph G is just a cycle and Brewster, Lee, Moore, Noel and Siggers [3] proved that H-Recolouring is PSPACE-complete if H is an odd wheel; for $k \geq 3$, the *wheel* W_k is the graph consisting of an irreflexive cycle of length k and a vertex adjacent to every vertex of the cycle and it is *odd* if k is odd.

Our goal in this paper is to obtain a rich class of graphs H for which H-Recolouring is PSPACE-complete. Throughout the paper, a *quadrangulation* is a connected irreflexive graph admitting an embedding in the 2-sphere (or, equivalently, the plane) in which every face is bounded by four edges. Our main result for irreflexive graphs is the following.

Theorem 1.1. *If H is a finite irreflexive quadrangulation not containing $K_{2,3}$ as a subgraph and not isomorphic to the 4-cycle, then H-Recolouring is PSPACE-complete.*

To prove Theorem 1.1, we will reduce K_4-Recolouring to H-Recolouring and apply the result of Bonsma and Cereceda [2] mentioned above. However it will be apparent in the proof that essentially the same approach could have been used to reduce F-Recolouring to H-Recolouring for any graph F and so there is nothing particularly special about the choice of K_4 (except that K_4-Recolouring is known to be PSPACE-complete).

As an application of Theorem 1.1, we will derive an analogous statement for non-bipartite quadrangulations of the projective plane. We are indebted to an anonymous referee who pointed out that one of our original proofs (which was only stated for odd wheels) holds in this generality.

Theorem 1.2. *If H is a non-bipartite quadrangulation of the projective plane not containing $K_{2,3}$ as a subgraph, then H-Recolouring is PSPACE-complete. Moreover, it remains PSPACE-complete when restricted to instances (G, f, g) such that G is bipartite.*

This applies to interesting families of graphs such as odd wheels and 4-chromatic generalized Mycielski graphs. Therefore, it recovers, and vastly extends, the result of Brewster et al. [3] that H-Recolouring is PSPACE-complete for any odd wheel H. The key to the proof of Corollary 1.2 is that the "bipartite double cover" of a non-bipartite quadrangulation of the projective plane is a quadrangulation of the sphere. If H is an even wheel, then the bipartite double cover is no longer

---

[1]Note that it is not hard to see that H-Recolouring is in PSPACE for every finite graph H.

a quadrangulation of the sphere itself, but it can be "retracted" to one. Using this, we obtain the following explicit extension of the result for odd cycles from [3].

Theorem 1.3. *For $k \geq 3$ and $k \neq 4$, W_k-Recolouring is PSPACE-complete. Moreover, it remains PSPACE-complete when restricted to instances (G, f, g) such that G is bipartite.*

Note that the condition that $k \neq 4$ in Theorem 1.3 is necessary unless P = PSPACE. To see this, observe that the graph W_4 contains two pairs of vertices with identical neighbourhoods and identifying both of these pairs "folds" W_4 to K_3. This observation can be used to show that W_4-Recolouring and K_3-Recolouring are polynomially equivalent and so the former is solvable in polynomial time by the result of [8]. For details on how these "folding reductions" work in general, see Wrochna [25, Proposition 4.3].

Theorem 1.1 fits a general theme, which first emerged in the topological approach of Wrochna [26], that the complexity of H-Recolouring may be closely related to the structure of a topological complex in which vertices, edges and 4-cycles of H (and, more generally, complete bipartite subgraphs) are faces. That is, Wrochna's result [26] says that, if this complex is "thin" in the sense that all of its faces are 0- or 1-dimensional, then its simple topological structure can be exploited to obtain a polynomial-time algorithm, whereas our result says that if this complex has basically the same topology as a 2-sphere, then the problem is PSPACE-complete. We further discuss the likely connections between H-Recolouring and the "topology" of H in the setting of reflexive graphs H, which we discuss next, as the connections are more natural in this setting.

Consider now the reconfiguration problem for H-colourings of G where both of the graphs G and H are reflexive, studied previously in [4]. In this setting, the definition of a reconfiguration sequence is somewhat different; it is defined to be a sequence $f_0, \ldots, f_m \in \text{Hom}(G, H)$ such that f_i and f_{i+1} differ on a unique vertex u_i for $0 \leq i \leq m - 1$ and $f_{i+1}(u_i)$ is a neighbour of $f_i(u_i)$. That is, it is the same as the definition for irreflexive graphs, but with an additional restriction that a vertex can only change its colour to a neighbour of its current colour.

Let us justify this extra condition by redefining the notion of homomorphism reconfiguration in terms of paths in the well known Hom-graph. For general graphs G and H, the *Hom-graph*, denoted $\mathbf{Hom}(G, H)$, is the graph with vertex set $\text{Hom}(G, H)$ in which two homomorphisms ϕ and ψ in $\text{Hom}(G, H)$ are adjacent if $\phi(x)\psi(y) \in E(H)$ for every $xy \in E(G)$. The Hom-graph can be viewed as the 2-skeleton of the Hom-complex which was first introduced by Lovász [19] in his celebrated proof of the Kneser Conjecture.

When G and H are irreflexive, it is easily shown that one can reconfigure between two homomorphisms ϕ and ψ from G to H if and only if they are in the same component of $\mathbf{Hom}(G, H)$; see, e.g., [6, Proposition 3.2] for a proof. However, for reflexive graphs, the presence of a loop on every vertex of G means that, for any two homomorphisms ϕ and ψ which are adjacent in $\mathbf{Hom}(G, H)$ and vertex $v \in V(G)$, we must have that $\phi(v)$ and $\psi(v)$ are neighbours in H. Using this observation, it is not hard to see that, when G and H are reflexive, there is a path between two H-colourings of G if and only if one can be reconfigured to the other with the extra condition mentioned above. This definition in terms of the Hom-graph allows us to define the reconfiguration not only for reflexive graphs, but also for graphs in which loops are allowed, but not required. Further, defining it this way helps preserve connections between homomorphism reconfiguraton and important concepts in graph homomorphism theory, such as the connection to the Hom-complex.

Basic among these concepts is the fact from Brightwell and Winkler [7] that a graph H is "dismantlable" if and only if $\mathbf{Hom}(G, H)$ is connected for all H. From this we get that H-Recolouring is trivial if H is dismantlable. The notion of "folding" mentioned above can be seen as a irreflexive version of dismantling. Using the same proof as for folding, it is easily shown that if H dismantles to H', then the problems H-Recolouring and H'-Recolouring are polynomially equivalent.

We consider such properties to be topological properties as there is a close connection between $\mathbf{Hom}(G, H)$ being disconnected for various G and non-trivial homotopy in the *clique complex* of H—the simplicial complex on the vertices of H whose k-simplices are the reflexive $(k + 1)$-cliques of H. While H being dismantlable implies that its clique complex deformation retracts to a single vertex, and so has trivial homotopy, non-trivial homomotopy in the clique complex of H is the main source of disconnectedness in $\mathbf{Hom}(C, H)$ for various cycles C or higher dimensional analogues.

In [17], Larose, shows that any non-trivial homotopy in the clique complex of a reflexive graph H yields an NP-complete retraction problem $\mathrm{Ret}(H)$ (which is the usual analogue of the homomorpism problem for reflexive graphs). Wrochna's result suggests that non-trivial homotopy of dimension 1 will not suffice to make a hard reconfiguration problem, and indeed, in a forthcoming paper [18], we obtain an analogue of Wrochna's result in the reflexive setting by showing that H-RECOLOURING is polynomial time solvable if H contains no triangles—that is, if the clique complex has only 0- and 1-simplices. To find H for which H-RECOLOURING is not polynomial time solvable, it seems that the natural candidates are graphs H with 2-dimensional holes.

We believe that any two dimensional hole in the clique complex of a reflexive graph H yields a hard reconfiguration problem. A reflexive graph is called a *triangulation* if the underlying irreflexive graph is connected and can be embedded in the plane so that all faces are bounded by three edges. Triangulations of a sphere that are not K_4 are the simplest examples of simplicial complexes with two dimensional holes. To make an explicit conjecture, we expect for a reflexive graph H that if $\mathbf{Hom}(S, H)$ is disconnected for any reflexive triangulation S of a sphere, then H-RECOLOURING is PSPACE-complete.

The following result, our main result for reflexive graphs, is a step towards resolving this, and to our knowledge, provides the first examples of graphs for which H-RECOLOURING is PSPACE-complete when restricted to reflexive instances.

Theorem 1.4. *If H is a finite reflexive triangulation not containing K_4 as a subgraph and not isomorphic to a reflexive triangle, then H-RECOLOURING is PSPACE-complete when restricted to instances (G, f, g) such that G is reflexive.*

As is the case in [17] we do not expect the converse of our conjecture to hold. Indeed, it does not. We will prove a more general result than Theorem 1.4 which only requires H to have the "local" structure of a triangulation near a particular vertex, and the "global" property of "stiffness" which effectivly stops this local structure from dismantling to the particular vertex. This more general result is stated and proved in Section 5 (Theorem 5.5).

The rest of the paper is organized as follows. In the next section, we give an overview of the proofs of Theorems 1.1 and 1.4. In particular, we describe the types of gadgets used in the proofs and show that the existence of such gadgets is sufficient to prove the main theorems. In Section 3, we build up several basic structural properties about graphs satisfying the hypotheses of Theorem 1.1 and use them to construct the required gadgets. We then deduce Theorem 1.2 in Section 4 and the using a reduction for even wheels, apply it to prove Theorem 1.3. In Section 5, we state and prove a generalization of Theorem 1.4.

2. OVERVIEW OF THE MAIN REDUCTION

The goal of this section is to introduce the main gadgets used in the proofs of Theorems 1.1 and 1.4. We will mainly focus on Theorem 1.1, only commenting briefly at the end of the section about how the ideas can be adapted to the setting of Theorem 1.4. For the time being, we let H be any $K_{2,3}$-free finite quadrangulation other than the 4-cycle. We claim that H must contain a vertex of degree three. Indeed, by Euler's Polyhedral Formula, it must contain a vertex of degree at most three and since it is a quadrangulation which is $K_{2,3}$-free and not isomorphic to C_4, it cannot have a vertex of degree less than three (see Lemma 3.2 for a formal proof). So, we may

choose an arbitrary vertex of degree three and label it 0. Also, label the three faces incident with 0 by f_1, f_2 and f_3 and, for $i \in \{1, 2, 3\}$, label the unique vertex incident to f_i and not adjacent to 0 by i. Since the three neighbours of 0 are distinct and H is $K_{2,3}$-free, the vertices $1, 2$ and 3 must be distinct. We let $\alpha_{1,2}, \alpha_{2,3}$ and $\alpha_{3,1}$ be defined so that $\alpha_{i,j}$ is the common neighbour of vertices i, j and 0, which is unique because H is $K_{2,3}$-free; see Figure 1.
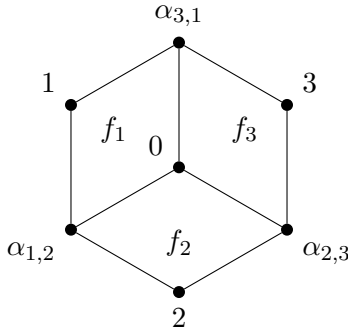

FIGURE 1. The local structure near vertex 0.

Our aim is to reduce the K_4-Recolouring problem, which was shown to be PSPACE-complete in [2], to the H-Recolouring problem. To this end, we let (G, f, g) be an instance of K_4-Recolouring, where $V(K_4) = \{1, 2, 3, 4\}$ and will construct an instance (G', f', g') of H-Recolouring such that $|V(G')| = O\left(|V(G)|^2\right)$ and f reconfigures to g if and only if f' reconfigures to g'. The construction is broken down into four steps. We describe the first two steps now and postpone the description of the third and fourth until after some additional discussion.

Step 1. Each vertex $u \in V(G)$ is represented by four vertices u_1, u_2, u_3 and u_4 in G'.

Step 2. For each $u \in V(G)$ and $i \in \{1, 2, 3, 4\}$, define

$$f'(u_i) := \begin{cases} 1 & \text{if } f(u) = i, \\ 0 & \text{otherwise} \end{cases}$$

and define $g'(u_i)$ analogously.

As one may be able to glean from Step 2, the images of the vertices u_1, u_2, u_3, u_4 under an H-colouring of G' will be used to encode the colour of u under an associated K_4-colouring in a simple way. That is, we think of u_i mapping to $1/0$ as meaning that the ith colour is "turned on/turned off" at the vertex u. Note that we will not mind if more than one colour is turned on at u (in fact, it is necessary to allow this in order for u to transition between colours in the associated K_4-colourings). What we need in order to make the reduction work is to design gadgets which force the following properties to be maintained throughout any reconfiguration sequence starting with f':

  (i) For each $u \in V(G)$ and $i \in \{1, 2, 3, 4\}$, the colour of u_i is either 0 or 1.
  (ii) For each $uv \in E(G)$ and $i \in \{1, 2, 3, 4\}$, the vertices u_i and v_i cannot map to 1 at the same time.
 (iii) For each $u \in V(G)$ at least one of the vertices u_1, u_2, u_3, u_4 maps to 1.

Given that these properties are maintained, it will easily follow that if f' reconfigures to g', then f reconfigures to g. Indeed, for each H-colouring in the reconfiguration sequence taking f' to g', we define a K_4-colouring by assigning each $u \in V(G)$ to the minimum $i \in \{1, 2, 3, 4\}$ such that u_i is mapped to 1. By (iii), such an i always exists and, by (ii), this choice will always produce a

K_4-colouring. Clearly, any two consecutive such K_4-colourings will differ on at most one vertex. Also, applying this transformation to f' itself yields f, and applying it to g' yields g. Therefore, we obtain a reconfiguration sequence taking f to g, as desired.

However, when trying to prove the other direction, one soon realizes that it is important to make the gadgets sufficiently "flexible" so that we can mimic any reconfiguration sequence taking f to g by a reconfiguration sequence of H-colourings taking f' to g'. That is, we need to not only block the "undesirable configurations" (e.g. u_2 and v_2 both mapping to 1 for $uv \in E(G)$), but also to allow any sorts of "allowed transitions" between configurations (e.g. changing the colour of u_3 from 1 to 0 while u_4 is coloured with 1). This discussion is an attempt to motivate condition (c) of the following technical-looking definition. For a function f on a set X and a vector $x = (x_1, \ldots, x_k) \in X^k$, we write $f(x)$ for $(f(x_1), \ldots, f(x_k))$.

Definition 2.1. Let H be a graph, k be a positive integer and $\mathcal{P} \subseteq V(H)^k$ be a k-ary relation on $V(H)$ (the elements of which we call k-patterns). A $\mathcal{P}$-gadget is a graph $Y = Y(x)$ where $x = (x_1, \ldots, x_k)$ is any ordered set of signal vertices $x_1, \ldots, x_k \in V(Y)$ such that the following hold.
(a) For each $p \in \mathcal{P}$ there is a canonical $\zeta_p \in \text{Hom}(Y, H)$ such that $\zeta_p(x) = p$.
(b) If $\psi \in \text{Hom}(Y, H)$ reconfigures to ζ_p for some $p \in \mathcal{P}$, then $\psi(x) \in \mathcal{P}$.
(c) If $p, q \in \mathcal{P}$ differ on at most one coordinate, then there exists a reconfiguration sequence taking ζ_p to ζ_q such that $\xi(x) \in \{p, q\}$ for every element ξ of this sequence.

As we will see shortly, the proof of Theorem 1.1 boils down to establishing the following two lemmas. Given distinct vertices a, b of H, say that a is across from b if a is not adjacent to b and there exists a face of H incident to both a and b.

Lemma 2.2. Let H be a finite $K_{2,3}$-free quadrangulation which is not a 4-cycle. If a_0 is across from a_1 and b_0 is across from b_1, then there there exists an $\{(a_0, b_0), (a_1, b_0), (a_0, b_1)\}$-gadget.

Lemma 2.3. Let H be a finite $K_{2,3}$-free quadrangulation which is not a 4-cycle and let $0, 1 \in V(H)$ such that 0 has degree three and 1 is across from 0. Then there exists a $\left(\{0,1\}^4 \setminus \{(0,0,0,0)\}\right)$-gadget.

In general, for a graph H and distinct $0, 1 \in V(H)$, we refer to a $\{(0,0), (1,0), (0,1)\}$-gadget as a not-both-one gadget and a $\left(\{0,1\}^4 \setminus \{(0,0,0,0)\}\right)$-gadget as a not-all-zero gadget. Given these lemmas, we complete the construction of (G', f', g') by applying the following steps. See Figure 2 for an illustration of the full construction of (G', f', g').

Step 3. For each ordered pair (u, v) with $uv \in E(G)$ and $i \in \{1, 2, 3, 4\}$, we add a not-both-one gadget $Y(u_i, v_i)$ to G' which is disjoint from all vertices added so far except for the signal vertices u_i and v_i. We define f' on Y to agree with $\zeta_{(f'(u_i), f'(v_i))}$. The definition of g' on Y is analogous.[2]

Step 4. For each $u \in V(G)$ we add a not-all-zero gadget $Z(u_1, u_2, u_3, u_4)$ to G', disjoint from all vertices added so far except for the signal vertices u_1, u_2, u_3, u_4. We define f' on Z to agree with $\zeta_{(f'(u_1), f'(u_2), f'(u_3), f'(u_4))}$. The definition of g' on Z is analogous.

This completes the construction of the instance (G', f', g') of H-Recolouring. Next, we prove a general lemma which says that the existence of a not-both-one gadget and a not-all-zero gadget is enough to prove that H-Recolouring is PSPACE-complete.

Lemma 2.4. If H is a finite graph and $0, 1$ are distinct vertices of H such that there exists a not-both-one gadget and a not-all-zero gadget, then H-Recolouring is PSPACE-complete.

---

[2]Technically, we only require one of the gadgets $Y(u_i, v_i)$ or $Y(v_i, u_i)$ in order for the reduction to work. However, adding both gadgets provides symmetry which is convenient in the exposition of the proof of Lemma 2.4.
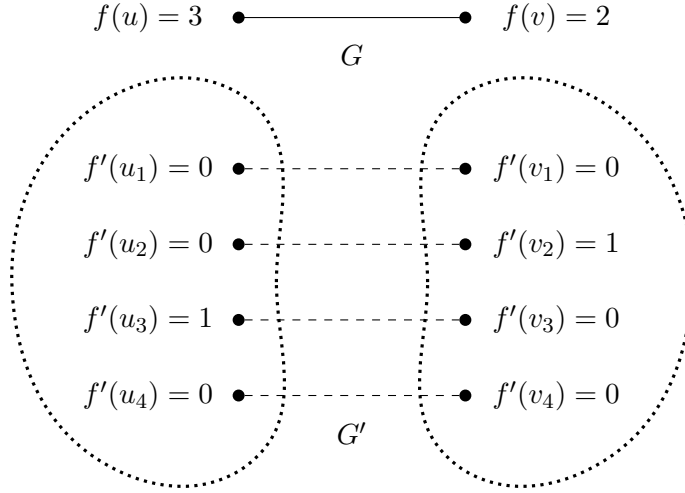
FIGURE 2. An illustration of the way in which two adjacent vertices u, v of G are represented in the graph G'. Each dashed line connects signal vertices of a pair of not-both-one gadgets and each thick dotted curve encloses the four signal vertices of a not-all-zero gadget.

Proof. Given an instance (G, f, g) of the K_4-Recolouring problem, we let (G', f', g') be an instance for H-Recolouring constructed using Steps 1-4 outlined above. Note that the size of these gadgets depends only on H and so $|V(G')| = O\left(|V(G)| + |E(G)|\right) = O\left(|V(G)|^2\right)$. We show that f reconfigures to g if and only if f' reconfigures to g'.

First suppose that there is a reconfiguration sequence $f'_0, \ldots, f'_m \in \mathrm{Hom}(G', H)$ taking f' to g'. By Step 4 and the definition of the not-all-zero gadget, we know that, for each $u \in V(G)$ and $0 \le j \le m$, there exists $i \in \{1, 2, 3, 4\}$ such that $f'_j(u_i) = 1$. For $1 \le j \le m$, define $f_j : V(G) \to \{1, 2, 3, 4\}$ by

$$f_j(u) := \min\left\{i : f'_j(u_i) = 1\right\}$$

for $u \in V(G)$. By Step 3 and the definition of the not-both-one gadget, we have that each f_j is a K_4-colouring of G; that is, adjacent vertices of G receive distinct colours. Clearly, by construction, $f_0 = f$ and $f_m = g$. Also, since, for $0 \le j \le m-1$, the mappings f'_j and f'_{j+1} differ on exactly one vertex, the mappings f_j and f_{j+1} differ on at most one vertex. Thus, we can take a subsequence of $f_0, \ldots, f_m$ obtained by deleting repetitions (i.e. removing one of f_i or f_{i+1} when $f_i = f_{i+1}$) to get a reconfiguration sequence taking f to g.

For the other direction, suppose that there is a reconfiguration sequence $f_0, \ldots, f_m \in \mathrm{Hom}(G, K_4)$ taking f to g. We construct a reconfiguration sequence taking f' to g'. We may assume that $m = 1$, i.e. that f and g differ on a unique vertex u, since the general case follows by induction on m. So, without loss of generality, we assume that $f(u) = 1$ and that $g(u) = 2$ and that $f(v) = g(v)$ for all $v \in V(G) \setminus \{u\}$.

Since f and g are K_4-colourings which differ only on u, we have that $f(v), g(v) \in \{3, 4\}$ for every vertex v adjacent to u. By construction, this means that f' assigns

- the colouring $\zeta_{(1,0)}$ to $Y(u_1, v_1)$,
- the colouring $\zeta_{(0,1)}$ to $Y(v_1, u_1)$,
- the colouring $\zeta_{(0,0)}$ to both $Y(u_2, v_2)$ and $Y(v_2, u_2)$, and
- the colouring $\zeta_{(1,0,0,0)}$ to $Y(u_1, u_2, u_3, u_4)$.

Similarly, g' assigns

- the colouring $\zeta_{(0,0)}$ to both of $Y(u_1, v_1)$ and $Y(v_1, u_1)$,
- the colouring $\zeta_{(1,0)}$ to $Y(u_2, v_2)$,
- the colouring $\zeta_{(0,1)}$ to $Y(v_2, u_2)$, and
- the colouring $\zeta_{(0,1,0,0)}$ to $Y(u_1, u_2, u_3, u_4)$.

On all other vertices of G', the colourings f' and g' agree with one another. By condition (c) of Definition 2.1, we know that we can reconfigure $\zeta_{(0,0)}$ to $\zeta_{(1,0)}$ in such a way that the colour of the first signal vertex stays in $\{0, 1\}$ the second is mapped to 0 throughout. For each neighbour v of u, one at a time, we apply the first part of this reconfiguration sequence on $Y(u_2, v_2)$, stopping just before the first step in which the colour of u_2 changes from 0 to 1. Similarly, on $Y(v_2, u_2)$, apply the first steps of a reconfiguration sequence from $\zeta_{(0,0)}$ to $\zeta_{(0,1)}$ and, on $Y(u_1, u_2, u_3, u_4)$, apply the first steps of a reconfiguration sequence from $\zeta_{(1,0,0,0)}$ to $\zeta_{(1,1,0,0)}$, in all cases stopping just before the first time the colour of u_2 changes from 0 to 1. Next, we go through each of these gadgets again, one by one, and continue the reconfiguration sequence, this time stopping at the last step in which the colour of u_2 is 0. Then, go through each gadget one last time to complete the reconfiguration sequence. Note that this procedure maintains an H-colouring throughout since any two of these gadgets only intersect on u_2 and possibly a vertex v_2 which does not change colour.

Thus, we have arrived at a colouring h' which assigns

- the colouring $\zeta_{(1,0)}$ to $Y(u_1, v_1)$,
- the colouring $\zeta_{(0,1)}$ to $Y(v_1, u_1)$,
- the colouring $\zeta_{(1,0)}$ to $Y(u_2, v_2)$,
- the colouring $\zeta_{(0,1)}$ to $Y(v_2, u_2)$, and
- the colouring $\zeta_{(1,1,0,0)}$ to $Y(u_1, u_2, u_3, u_4)$.

and, on all other vertices of G', agrees with both f' and g'. By applying the same steps as above with g' in the place of f' and swapping the roles of colours 1 and 2, we see that g' reconfigures to h' as well and so, by symmetry and transitivity of the "reconfigures to" relation, f' reconfigures to g'. This completes the proof. $\square$

Before closing this section, let us make a few remarks about Theorem 1.4. Let H be a finite K_4-free reflexive triangulation which is not a reflexive triangle. This time, we will let 0 be an arbitrary vertex and label the neighbours of 0 apart from 0 itself by $1, \ldots, k$ in clockwise order with respect to the embedding of H, where k is the degree of 0 (we follow the convention that the loop on 0 does not count towards its degree). Note that, as H is a triangulation and not a triangle, the neighbours of 0 apart from 0 itself must form a reflexive cycle in H; this is proved formally in Lemma 5.3. The following lemmas are the crux of the proof of Theorem 1.4.
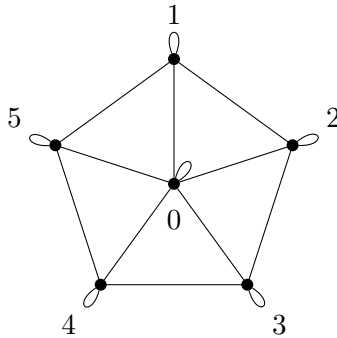

FIGURE 3. The subgraph of H induced by the vertex 0 and its neighbours in the case $k = 5$.

Lemma 2.5. *If H be a finite reflexive K_4-free triangulation which is not a reflexive triangle and $01 \in E(H)$ with $0 \neq 1$, then there there exists a not-both-one gadget.*

Lemma 2.6. *If H is a finite reflexive K_4-free triangulation which is not a reflexive triangle and $01 \in E(H)$ with $0 \neq 1$, then there there exists a not-all-zero gadget.*

As was mentioned in the introduction, we will actually prove a more general result than Theorem 1.4 and, thus, will require more general lemmas than Lemmas 2.5 and 2.6. These lemmas are stated and proved in Section 5.

3. Gadgets for Quadrangulations

Throughout this section, let H be a finite $K_{2,3}$-free quadrangulation which is not a 4-cycle. We begin by obtaining some basic structural properties of H which will be useful in the proofs of Lemmas 2.2 and 2.3.

3.1. **Basic Structural Properties of H.** The following two lemmas highlight some of the main ways in which we exploit the fact that H is $K_{2,3}$-free.

Lemma 3.1. *If $abcd$ and $a'b'cd$ are distinct cycles of H, then a is not adjacent to b'.*

Proof. Suppose that a is adjacent to b'. Then b, d and b' are all common neighbours of a and c. As H is $K_{2,3}$-free, a and c cannot have more than two common neighbours, and so it must be the case that some of these vertices coincide. Since b and d are distinct vertices on a cycle (by hypothesis), and so are b' and d, it must be the case that $b = b'$. However, now we get that a, c and a' are common neighbours of b and d, and these three vertices are distinct from one another because a and c are on a cycle, as are a' and c, and the cycles $abcd$ and $a'b'cd$ are distinct by hypothesis. This contradicts the assumption that H is $K_{2,3}$-free and completes the proof. $\square$

Lemma 3.2. *Every vertex of H has degree at least three.*

Proof. Suppose not. As H is a quadrangulation, it is clear that it has no vertex of degree zero or one. So, let y be a vertex of degree two and let x and z be its two neighbours. Since H is a quadrangulation and y has degree two, the faces f_1 and f_2 incident to the edge xy must be incident to x, y, z. For $i \in \{1, 2\}$, let w_i be the fourth vertex on the boundary of f_i. If $w_1 \neq w_2$, then x and z have three distinct common neighbours, contradicting the fact that H is $K_{2,3}$-free. On the other hand, if w_1 and w_2 coincide, then f_1 and f_2 have the same boundary which implies that H is a 4-cycle, and is again a contradiction. $\square$

The following definitions are useful for stating the next lemma.

Definition 3.3. A vertex v of a graph G is *frozen* by an F-colouring f if $g(v) = f(v)$ for every F-colouring $g \in \mathrm{Hom}(G, F)$ which reconfigures to f.

Definition 3.4. An F-colouring f of a graph G is said to be *frozen* if every vertex of G is frozen by f.

Definition 3.5. Say that a graph F is *stiff* if the identity map on $V(F)$ is a frozen F-colouring of F.

Observation 3.6. A graph F is not stiff if and only if there exists distinct vertices $u, v \in V(F)$ such that $N(u) \subseteq N(v)$.

Lemma 3.7. *H is stiff.*

Proof. Suppose not. Then, by Observation 3.6, there exists $u, v \in V(H)$ such that $N(u) \subseteq N(v)$. By Lemma 3.2, u has degree at least three. Thus, u and v have at least three distinct common neighbours, which contradicts the fact that H is $K_{2,3}$-free. $\square$

3.2. **Not-Both-One Gadget for Quadrangulations.** The basic idea underlying the proof of Lemma 2.2 is that if we take a homomorphism of a path $P = y_1 \cdots y_m$ to the "ladder" graph L as in Figure 4 such that the vertex y_i is constrained to map to c_i or d_i, then the pair (y_1, y_m) can map only to one of (d_1, d_m), (d_1, c_m) or (c_1, c_m), and one can easily reconfigure between such patterns. As we will see, constraining y_i to map to c_i or d_i is easy using frozen colourings and $K_{2,3}$-freeness. The harder part is finding an image of the graph L in H with vertices $(c_1, d_1, c_m, d_m) = (a_1, a_0, b_0, b_1)$. This takes most of the subsection, and is done with the directed graph Φ defined in Definition 3.9. Note that a path in the directed graph Φ corresponds to an image of the graph L in H, stronger than a homographic image: we also insist pairs in L connected by dashed edges map to vertices that are across from one another on a face of H.
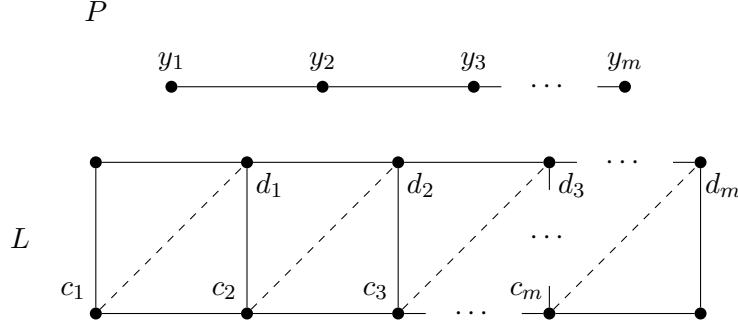

FIGURE 4. Graphs P and L motivating Φ of Defintion 3.9.

We start with the following definitions.

Definition 3.8. Let $A \subseteq V(H)^2$ be the set of all ordered pairs (a, b) such that a is across from b.

Note that because H is $K_{2,3}$-free, there is a unique face incident to both a and b for any pair $(a, b) \in A$.

Definition 3.9. Let Φ be the directed graph with vertex set A where there is an arc from (a, b) to (c, d) if $ac, bc, bd \in E(H)$ and $ad \notin E(H)$.

Figure 4 shows a 'path' in Φ from (c_1, d_1) to (c_m, d_m). Observe that the following is immediate from the definition of Φ.

Observation 3.10. There is an arc from (a, b) to (c, d) in Φ if and only if there is an arc from (d, c) to (b, a) in Φ.

The following lemma highlights the utility of Φ in proving Lemma 2.2.

Lemma 3.11. *A directed path from (a_1, a_0) to (b_0, b_1) in Φ yields a $\{(a_0, b_0), (a_1, b_0), (a_0, b_1)\}$-gadget.*

Proof. Let $(c_1, d_1), \ldots, (c_m, d_m)$ be a directed path in Φ with $(c_1, d_1) = (a_1, a_0)$ and $(c_m, d_m) = (b_0, b_1)$. We initiate the construction of the gadget $Y(x_1, x_2)$ with a copy H^* of H where the vertex of H^* corresponding to a vertex $v \in V(H)$ is denoted by v^*. We then add a path $y_1 \cdots y_m$ and edges from y_i to the copies of the two common neighbours of c_i^* and d_i^* in H^* for $1 \le i \le m$. Define the signal vertex x_1 to be y_1 and the signal vertex x_2 to be y_m. Define the mappings $\zeta_{(a_0, b_0)}, \zeta_{(a_1, b_0)}$ and $\zeta_{(a_0, b_1)}$ so that the copy of H^* is coloured by the identity map $v^* \mapsto v$ and the path $y_1 \cdots y_m$ is coloured by

$$\zeta_{(a_0, b_0)}(y_i) := \begin{cases} d_i & \text{if } i = 1, \\ c_i & \text{otherwise.} \end{cases}$$

10

$$\zeta_{(a_1, b_0)}(y_i) := c_i \text{ for } 1 \le i \le m,$$

$$\zeta_{(a_0, b_1)}(y_i) := d_i \text{ for } 1 \le i \le m.$$

Note that these mappings are indeed H-colourings of Y by definition of Φ and that they satisfy condition (a) of Definition 2.1.

We observe that, by Lemma 3.7, every vertex of H^* is frozen by each of the homomorphisms ζ_p. Thus, if ψ reconfigures to ζ_p for some p, then $\psi(v^*) = v$ for each $v \in V(H)$. Since H is $K_{2,3}$-free, the two vertices of H adjacent to both c_i and d_i do not have a third common neighbour, and so we get that $\psi(y_i) \in \{c_i, d_i\}$. In particular, $\psi(y_1) \in \{a_0, a_1\}$ and $\psi(y_m) \in \{b_0, b_1\}$. So, to verify condition (b) of Definition 2.1, we need only to show that it cannot be the case that $\psi(y_1) = a_1$ and $\psi(y_m) = b_1$. If we have $\psi(y_1) = a_1 = c_1$, then we must have $\psi(y_2) = c_2$ because c_1 is not adjacent to d_2 (by construction of Φ). Repeating the same argument, we get $\psi(y_3) = c_3$, $\psi(y_4) = c_4$, and so on. In particular, we must have $\psi(y_m) = c_m = b_0$, as desired.

Finally, we check condition (c) of Definition 2.1. By symmetry of reconfiguration sequences, it suffices to consider $p = (a_0, b_0)$ and $q = (a_1, b_0)$ or (a_0, b_1). The former case is trivial as $\zeta_{(a_0, b_0)}$ and $\zeta_{(a_1, b_0)}$ differ only on y_1. In the latter case, we start with $\zeta_{(a_0, b_0)}$ and change the colours of each y_i for $2 \le i \le m$, one by one, from c_i to d_i. Each of the intermediate mappings is an H-colouring of Y which differs from the previous one on a unique vertex. After all of these changes have been made, we arrive at $\zeta_{(a_0, b_1)}$. Also, every H-colouring in this sequence maps $x_1 = y_1$ to a_0 and $x_2 = y_m$ to either b_0 or b_1, as required. This completes the proof. $\qquad\square$

Therefore, our goal in proving Lemma 2.2 will be to show that there is a directed path in Φ between any two elements of A. To this end, we build up further useful properties of Φ.

Lemma 3.12. *Every vertex of Φ has in-degree and out-degree equal to 2.*

Proof. Let (a, b) be a vertex of Φ and let f be the face of H incident to a and b. If there is an arc from (a, b) to (c, d), then c must be adjacent to both a and b. Thus, since H is $K_{2,3}$-free, there are precisely two choices for c; namely, the two other vertices incident to f. Now, given a choice of c, we may let f' be the unique face with $f' \ne f$ such that the f' is incident to bc. Letting d be the non-neighbour of c on the boundary of f', we have that a is not adjacent to d by Lemma 3.1 and so there is an arc from (a, b) to (c, d) in Φ.

The argument above shows that the out-degree of (a, b) is at least two. To prove that it is exactly two, we need to show that the choice of d is unique once c has been chosen. If not, let $d' \ne d$ be a vertex which is across from c, adjacent to b and not adjacent to a. Let f'' be the face whose boundary contains c and d' and note that $f'' \notin \{f, f'\}$ as d' is not adjacent to a and is not equal to d. Let x, y be the other two vertices on the boundary of f''. Then c and d' are adjacent to all three of b, x and y and so, since H is $K_{2,3}$-free, these three vertices cannot be distinct. Since $x \ne y$, we get that, without loss of generality, $b = x$. However, we now have that the edge bc is on the boundary of three distinct faces, namely f, f' and f'', contradicting the fact that H is a planar quadrangulation and completing the proof.

Finally, to see that the in-degree of each vertex is also equal to two, we simply use the fact that the out-degree of every vertex is two and apply Observation 3.10. $\qquad\square$

The following definition will be helpful in further analysing Φ.

Definition 3.13. Let Γ be the graph with vertex set A where (a, b) is adjacent to (c, d) if

(a) the face f incident to a and b is distinct from the face f' incident to c and d,
(b) there is exactly one edge incident to both f and f',
(c) a is adjacent to c, and
(d) b is adjacent to d.

From this definition, if is clear that for (a, b) in A, if the face f incident to a and b shares an edge with a face f', then there is a pair (c, d) A with f' incident to c and d, such that (a, b) is adjacent to (c, d) in Γ.

As it turns out, the undirected graph underlying Φ is precisely Γ.

Lemma 3.14. *Let $(a, b), (c, d) \in A$. Then (a, b) is adjacent to (c, d) in Γ if and only if there is an arc from (a, b) to (c, d) or an arc from (c, d) to (a, b) in Φ.*

Proof. First, suppose that (a, b) is adjacent to (c, d) in Γ. Let f and f' be the faces of H incident to a, b and c, d, respectively, and let e be the common edge of the boundaries of f and f'. Note that the fact that a is adjacent to c and b is adjacent to d implies that $\{a, b\} \cap \{c, d\} = \emptyset$. Thus, we must have that exactly one of c or d is incident with e and, likewise, exactly one of a or b is incident with e. By Lemma 3.1, the vertex of $\{a, b\}$ that is not incident with e cannot be adjacent to the vertex of $\{c, d\}$ that is not incident with e. Thus, by definition of Γ, this pair cannot be a and c, nor can it be b and d. On the other hand, it is easily checked (using Lemma 3.1) that if $e = bc$, then there is an arc from (a, b) to (c, d) in Φ and, if $e = ad$, then there is an arc from (c, d) to (a, b) in Φ.

Now, for the other direction, suppose, without loss of generality, that there is an arc from (a, b) to (c, d) in Φ. Let f be the face incident to a, b and f' be the face incident to c, d. Since a is adjacent to c but neither c nor a is adjacent to d, we must have $f \neq f'$. Also, since H has no $K_{2,3}$, it must be the case that c is incident to f and b is incident to f'. Thus, the boundaries of f and f' share the edge bc and, by Lemma 3.1, this is the unique such edge. So, we can conclude that (a, b) is adjacent to (c, d) in Γ. $\square$

Next, we prove that Γ is connected. Note that there are some subtleties here to be aware of. In particular, the fact that every finite quadrangulation has a vertex of degree at most three will be crucial. For infinite quadrangulations of the plane, the graph Γ is not connected in general; e.g. if H were the infinite square grid, then Γ would contain exactly four connected components, each of which is itself isomorphic to an infinite square grid. Also, the graph Γ can sometimes be disconnected if H is a finite quadrangulation of the torus; consider, for example, the Cartesian product of two cycles. Thus, in some sense, the structure of the sphere (in particular, the fact that it has positive curvature) is important for our proof to go through.

Lemma 3.15. *Γ is connected.*

Proof. Let vertices $0, 1, 2, 3, \alpha_{1,2}, \alpha_{2,3}$ and $\alpha_{3,1}$ and faces f_1, f_2 and f_3 of H be defined as in the previous section. We show that every $(a, b) \in A$ admits a path to $(0, 1)$ in Γ, which will complete the proof.

Let $(a, b) \in A$ be arbitrary and let f' be the face whose boundary contains a and b. We denote the four faces of H whose boundaries share a unique edge with the boundary of f' by $f'_1, \ldots, f'_4$. In light of Lemma 3.14, Lemma 3.12 gives us that (a, b) has four neighbours $(c_1, d_1), \ldots, (c_4, d_4)$ in Γ, where c_i and d_i are on the boundary of f'_i. Thus, by the connectedness of the planar dual of H, we get that there is a path in Γ starting at (a, b) and terminating at a pair (s, t) on the boundary of the face f_1. We show that there is a path from every such (s, t) to $(0, 1)$.

Of course, if $(s, t) = (0, 1)$, then we are simply done. If $(s, t) = (\alpha_{3,1}, \alpha_{1,2})$, then it is adjacent to $(0, 2)$, which is adjacent to $(\alpha_{2,3}, \alpha_{3,1})$, which is adjacent to $(0, 1)$ and we are done. The proof in the case $(s, t) = (\alpha_{1,2}, \alpha_{3,1})$ is similar. If $(s, t) = (1, 0)$, then we see that $(1, 0)$ is adjacent to $(\alpha_{3,1}, \alpha_{2,3})$, which is adjacent to $(0, 2)$, which is adjacent to $(\alpha_{3,1}, \alpha_{1,2})$ which was already shown to admit a path to $(0, 1)$. This completes the proof. $\square$

We are now ready to prove Lemma 2.2.

Proof of Lemma 2.2. By Lemmas 3.14 and 3.15, the undirected graph underlying Φ is connected and, by Lemma 3.12, every vertex of Φ has in-degree equal to its out-degree. Thus, Φ has an Eulerian circuit. In particular, this implies that there is a directed path between any two vertices of Φ, and so we are done by Lemma 3.11. $\qquad\square$

3.3. Not-All-Zero Gadget for Quadrangulations.

We will now use Lemma 2.2 to prove Lemma 2.3. We remark that, in a subtle way, this construction relies strongly on the fact that 0 was chosen to be a vertex of degree *exactly* three in H. In topological language, the proof relies strongly on the fact that the "link" of the vertex 0 is a triangle.

Proof of Lemma 2.3. We construct a not-all-zero gadget $Z(z_1, z_2, z_3, z_4)$. First, applying Lemma 2.2, we construct two $\{(1,2),(0,2),(1,0)\}$-gadgets $W_1(z_1, w_1)$ and $W_3(z_3, w_3)$, and two $\{(1,3),(0,3),(1,0)\}$-gadgets $W_2(z_2, w_2)$ and $W_4(z_4, w_4)$ disjointly. Next, add a disjoint copy H^* of H in which the vertex of H^* corresponding to a vertex $v \in V(H)$ is denoted v^* and add four new vertices $x_{1,2}, y_{1,2}, x_{3,4}$ and $y_{3,4}$ such that

- $x_{1,2}$ is adjacent to $w_1, w_2, 0^*$ and $y_{1,2}$ and
- $y_{1,2}$ is adjacent to $\alpha^*_{1,2}, \alpha^*_{3,1}$ and $x_{1,2}$.

Finally, we add a $\{(1,0),(0,1),(1,1)\}$-gadget $Y(y_{1,2}, y_{3,4})$. See Figure 5.
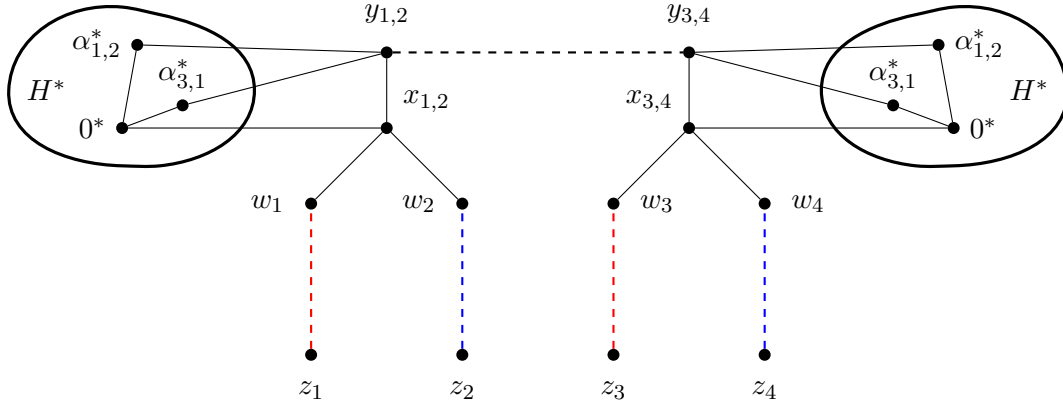

FIGURE 5. An illustration of the not-all-zero gadget. Thin solid black lines represent edges, thick solid black closed curves represent the copy H^* of H (drawn twice for clarity), red dashed lines connect signal vertices of $\{(1,2),(0,2),(1,0)\}$-gadgets, blue dashed lines connect signal vertices of $\{(1,3),(0,3),(1,0)\}$-gadgets and the black dashed line connects signal vertices of a $\{(1,0),(0,1),(1,1)\}$-gadget.

Now, for $p = (p_1, \ldots, p_4) \in \{0,1\}^4 \setminus \{(0,0,0,0)\}$, we define ζ_p as follows. The copy H^* of H is coloured according to the identity colouring $v^* \mapsto v$. If $p_1 = 0$, then we colour $W_1(z_1, w_1)$ according to $\zeta_{(0,2)}$ and similar for $W_3(z_3, w_3)$ if $p_3 = 0$. Similarly, if $p_2 = 0$, then we colour $W_2(z_2, w_2)$ according to $\zeta_{(0,3)}$ and similar for $W_4(z_4, w_4)$ if $p_4 = 0$. For each i such that $p_i = 1$, we colour the gadget $W_i(z_i, w_i)$ with $\zeta_{(1,0)}$. We colour $x_{1,2}$ with $\alpha_{1,2}$ if $p_2 = 1$, with $\alpha_{3,1}$ if $p_2 = 0$ and $p_1 = 1$ and with $\alpha_{2,3}$ otherwise. Also, colour $y_{1,2}$ with 1 if at least one of p_1 or p_2 is equal to 1 and 0 otherwise. The colouring of $x_{3,4}$ and $y_{3,4}$ is similar. Since at least one of $p_1, \ldots, p_4$ is equal to 1, we know that one of $y_{1,2}$ or $y_{3,4}$ is mapped to 1. We colour the gadget $Y(y_{1,2}, y_{3,4})$ with one of the colourings $\zeta_{(1,0)}, \zeta_{(0,1)}$ or $\zeta_{(1,1)}$ depending on which of $y_{1,2}$ or $y_{3,4}$ has already been coloured with 1. One can easily check that this definition of ζ_p is a valid H-colouring and that $\zeta_p(z_i) = p_i$ for $i = 1, 2, 3, 4$;

thus condition (a) of Definition 2.1 is satisfied. The rest of the proof consists of verifying that the other two conditions of Definition 2.1 hold.

Since each of the vertices z_1, z_2, z_3, z_4 is contained in either a $\{(1,2),(0,2),(1,0)\}$-gadget or a $\{(1,3),(0,3),(1,0)\}$-gadget which is mapped to a canonical colouring, we know that $\psi(z_i) \in \{0,1\}$ for all ψ which reconfigures to any ζ_p for $p \in \{0,1\}^4 \setminus \{(0,0,0,0)\}$. Thus, to verify condition (b) of Definition 2.1 we need only show that no H-colouring ψ which reconfigures to some ζ_p can map all of z_1, z_2, z_3, z_4 to zero. By definition of the $\{(1,2),(0,2),(1,0)\}$-gadget and $\{(1,3),(0,3),(1,0)\}$-gadget, if $\psi(z_1) = \psi(z_2) = 0$, then $\psi(w_1) = 2$ and $\psi(w_2) = 3$. Also, by Lemma 3.7, ψ must colour H^* according to the identity map. Thus, $x_{1,2}$ must map to a common neighbour of $0, 2, 3$, which implies that $\psi(x_{1,2}) = \alpha_{2,3}$ as H is $K_{2,3}$-free. From this, we get that $y_{1,2}$ must map to 0, as this is the only common neighbour of $\alpha_{1,2}, \alpha_{3,1}$ and $\alpha_{2,3}$, again by $K_{2,3}$-freeness. Applying the same argument starting with z_3 and z_4 gives us that $y_{3,4}$ maps to 0 as well, but this contradicts the definition of the $\{(1,0),(0,1),(1,1)\}$-gadget. Therefore, condition (b) of Definition 2.1 holds.

Finally, we move on to condition (c) of Definition 2.1. Let $p, q \in \{0,1\}^4 \setminus \{(0,0,0,0)\}$ differ on a unique coordinate. By symmetry and without loss of generality we can assume that $p_1 = 0$ and $q_1 = 1$ (note that this is indeed without loss of generality since we can always swap the names of vertices 2 and 3 of H). We begin by reconfiguring the colouring of the gadget between z_1 and w_1 from $\zeta_{(0,2)}$ to $\zeta_{(1,2)}$ without changing the colour of w_1. Then, without changing the colour of z_1, we reconfigure it to $\zeta_{(1,0)}$. At this point, the current colouring differs from ζ_q only on $x_{1,2}$ and vertices of the gadget $Y(y_{1,2}, y_{3,4})$. However, we may simply change the colour of $x_{1,2}$ to $\zeta_q(x_{1,2})$ at this point since, by definition of ζ_q, it is either $\alpha_{1,2}$ or $\alpha_{3,1}$, both of which are compatible with the colour of $y_{1,2}$. As a final step, we reconfigure the colouring of $Y(y_{1,2}, y_{3,4})$ to match ζ_q. This completes the proof. □

4. Quadrangulations of Projective Planes

Our first goal in this section is to deduce Theorem 1.2 from Theorem 1.1. We then adapt the argument to even wheels to finish off Theorem 1.3. We start with some background about embeddings into projective plane, and recall a couple of useful ideas from [25].

Recall that the projective plane $\mathcal{P}^2$ can be expressed as the quotient $q : S^2 \to \mathcal{P}^2$ of the sphere S^2 modulo the equivalence relation identifying antipodal pairs of points. Though we use this point of view in our proof, for intuition and figures, we use the equivalent representation of $\mathcal{P}^2$ as the disk B^2 modulo the equivalence identifying antipodal pairs of points on the boundary.

One of the main features of interest of the non-orientable surface $\mathcal{P}^2$ is that one can embed the cycle C_r in it, for r odd, so that there is exactly one face of length $2r$. Indeed, recall that to find the length of a face in an embedding of a graph, we walk around it, keeping our right hand on the boundary, and count the number of vertices we encounter, with repetition, until we get back to where we started. Referring to the first picture of Figure 6 in which we have embedded C_5 in the projective plane, we see that there is one face of length 10 as follows. Start just below vertex 1 on the northern (top) hemisphere of the figure, and keeping your right hand on the edge, traverse the first edge to arrive at vertex 2. From there, the edge on your right is the edge from 2 to 3e drawn in the southern hemisphere. Traversing that edge, you then continue from 3 to 4 on the west. You do not get back to you starting position until you have traversed ten edges. Playing the same game with the second picture, you see that you arrive back at vertex 1 after five steps, but you are on the other side of it. You must traverse ten edges to get back to where you started.

Recall that an embedding of a graph H in a surface is a quadrangulation if every face is of length four. An odd wheel can be drawn as a quadrangulation of $\mathcal{P}^2$ by embedding its outer cycle in the manner described above; see Figure 7.
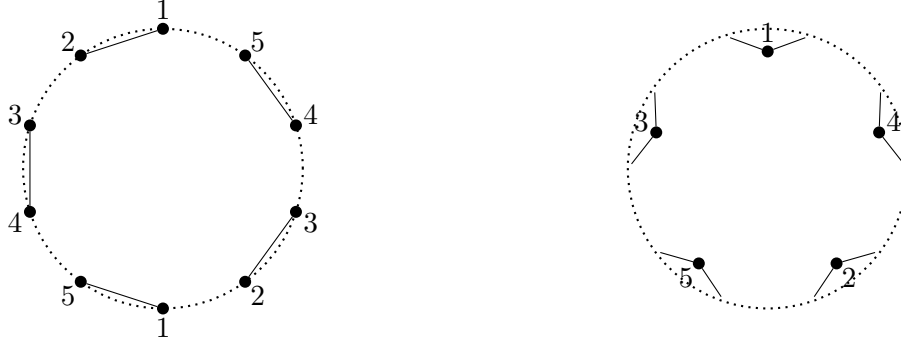
14

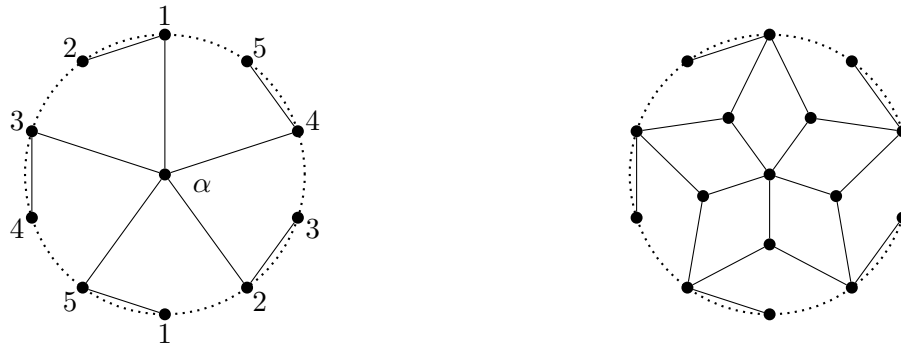FIGURE 6. Two drawings of the same embedding of C_5 in the projective plane

FIGURE 7. Embedding of odd wheel and 4-chromatic Myceilski graph in projective plane

Definition 4.1. For graphs F_1 and F_2, the *categorical* product (sometimes called the *direct* or *tensor* product) of F_1 and F_2, denoted $F_1 \times F_2$, is the graph on vertex set $V(F_1) \times V(F_2)$ where (u_1, u_2) is adjacent to (v_1, v_2) if and only if $u_1 v_1 \in E(F_1)$ and $u_2 v_2 \in E(F_2)$.

The product $H \times K_2$ of H with an edge is often referred to as the "bipartite double cover" of H, among other names. It is well known and easy to see that if H is bipartite, then $H \times K_2$ is two disjoint copies of H. It is well known, and easy to show, that for a graph G and homomorphisms $\phi_1 : G \to F_1$ and $\phi_2 : G \to F_2$, the map $\phi_1 \times \phi_2 : G \to F_1 \times F_2$ defined by $v \mapsto (\phi_1(v), \phi_2(v))$ is a homomorphism. (In fact it is known to be the unique homomorphism that commutes with the projections of the product onto its factors.)

Lemma 4.2 (Wrochna [25]). *For every graph H, the $(H \times K_2)$-RECOLOURING problem is polynomially equivalent to the restriction of the H-RECOLOURING problem to bipartite instances.*

Proof. Let $V(K_2) = \{1, 2\}$, let $\pi_1 : V(H) \times V(K_2) \to V(H)$ be the projection onto the first coordinate and let $\pi_2 : V(H) \times V(K_2) \to V(K_2)$ be the projection onto the second coordinate. Note that π_1 and π_2 are homomorphisms from $H \times K_2$ to H and K_2, respectively. Thus, if (G, f, g) is an instance of $(H \times K_2)$-RECOLOURING, then G admits a K_2-colouring (equivalently, G is bipartite) by composing f with π_2.

Now, let G be any bipartite graph and let A, B be the two sets of the bipartition. Given a homomorphism f from G to $H \times K_2$, we can compose f with π_1 to get a homomorphism f' to H. On the other hand, if f' is a homomorphism from G to H, we can let $f : V(G) \to V(H) \times V(K_2)$ be defined so that $f(v) = (f'(v), 1)$ if $v \in A$ and $f(v) = (f'(v), 2)$ if $v \in B$. It is not hard to show

that both of these transformations preserve the "reconfigures to" relation, and so the $(H \times K_2)$-RECOLOURING problem is polynomial-time equivalent to the H-RECOLOURING problem restricted to bipartite instances. $\square$

Theorem 1.2 is now immediate from Theorem 1.1 and Lemma 4.2 by the following fact.

Fact 4.3. *Let H be a non-bipartite quadrangulation of the projective plane. The graph $H \times K_2$ can be embedded as a quadrangulation of the sphere.*

Proof. Given a embedding of H in the projective plane, by making small perterbations we may assume that no vertex of H is embedded on the boudary of $\mathcal{P}^2$. Consider the preimage under $q : S^2 \to \mathcal{P}^2$ of this embedding of H. It is clearly an embedding on S^2 of some graph H'; and q clearly induces a graph homomorphism $q : H' \to H$. For any vertex v of H, let $q^{-1}(v) = \{v_N, v_S\}$ where the subscript designates the hemisphere of S^2 in which the vertex is embedded. What remains to be proven is that H' is a quadrangulation of S^2 and that it is isomorphic to $H \times K_2$. To see that it is a quadrangulation, observe that a closed (topological) path P in S^2 defines a closed path $q(P)$ in $\mathcal{P}^2$. The boundary of a face of H' therefore maps to the boundary of a face in H and, by assumption, this is always a quadrangle.

The graphs H' and $H \times K_2$ clearly have the same numbers of vertices and edges, so to show they are isomorphic it is enough to show there is a homomorphism between them that is a bijection on the vertices. As H' is a quadrangulation of S^2, it is bipartite, so has a K_2-colouring ϕ. We claim that the homomorphism

$$q \times \phi : H' \to H \times K_2 \text{ defined by } x \mapsto (q(x), \phi(x))$$

is a bijection. Indeed we have that $[q \times \phi](v_N)$ and $[q \times \phi](v_S)$ are both in $\{(v, 0), (v, 1)\}$ for each v, so it is enough to show that $\phi(v_N) \neq \phi(v_S)$ for all v in $V(H)$. Assume, towards contradiction, that $\phi(v_N) = \phi(v_S)$. Then for any neighbour u of v we have that each of v_N and v_S are adjacent to one of u_N and u_S, and so $\phi(u_N) = \phi(u_S)$. As H is connected this therefore holds for all vertices of H and so ϕ induces a K_2-colouring of H, which is impossible. Thus $\phi(v_N) \neq \phi(v_S)$ for all v in $V(H)$ as needed. $\square$

As we noted above, this tells us that H-RECOLOURING is PSPACE-complete for H an odd wheel. One can construct many non-bipartite quadrangulations of $\mathcal{P}^2$. Starting with an odd cycle embedded as we show in Figure 6 one simply has to quadrangulate it, as in the graph on the right of Figure 7. This is the Grötzsch graph, which is obtained from a 5-cycle by applying the Mycielski construction. More generally, every 4-chromatic "generalized Mycielski graph" (see [22] for a definition) can be embedded as a quadrangulation of the projective plane, yielding another interesting special case of Theorem 1.2.

We now provide an extra argument which allows us to extend the argument to even wheels of lenth at least six, thereby allowing us to prove Theorem 1.3. We need to recall one more basic idea about retractions.

Definition 4.4. Given a graph F and an induced subgraph H of F, a map $\varphi : V(F) \to V(H)$ is called a *retraction* from F to H if φ is a homomorphism from F to H and $\varphi(v) = v$ for every $v \in V(H)$.

Definition 4.5. An induced subgraph H of a graph F is called a *retract* of F if there is a retraction from F to H.

Lemma 4.6. *If H is a retract of F, then H-RECOLOURING reduces to F-RECOLOURING.*

Proof. Let φ be a retraction from F to H and define $\iota : V(H) \to V(F)$ by $\iota(v) = v$ for all $v \in V(H)$. Clearly, ι is a homomorphism from H to F. Note that $\varphi \circ \iota$ is nothing more than the identity map on H.

Now, given an instance (G, f, g) of H-Recolouring, we consider the instance $(G, \iota \circ f, \iota \circ g)$ of F-Recolouring. Clearly, if f reconfigures to g, then composing each map on the reconfiguration sequence with ι yields a reconfiguration sequence taking $\iota \circ f$ to $\iota \circ g$. On the other hand, if $\iota \circ f$ reconfigures to $\iota \circ g$, then composing each map on the reconfiguration sequence with φ (and possibly taking a subsequence) yields a reconfiguration sequence taking f to g. This completes the proof. $\square$

With Theorem 1.2, this and the following fact show that H-Recolouring is PSPACE-complete for even wheels $H = W_{2n}$ as long as $n \geq 3$, proving Theorem 1.3.

Fact 4.7. *For even $k \geq 6$, the graph $W_k \times K_2$ retracts to a non-trivial quadrangulation of S^2.*

Proof. The graph $W_k \times K_2$ consists of two cycles

$$C_1 := (x_1, 1)(x_2, 2), \ldots, (x_k, 2),$$

$$C_2 := (x_1, 2)(x_2, 1), \ldots, (x_k, 1)$$

and vertices $(\alpha, 1)$ and $(\alpha, 2)$ where the neighbourhood of $(\alpha, 1)$ is $\{(x_1, 2), \ldots, (x_k, 2)\}$ and the neighbourhood of $(\alpha, 2)$ is $\{(x_1, 1), \ldots, (x_k, 1)\}$. Let H_k be the subgraph of $W_k \times K_2$ induced by $V(C_1) \cup \{(\alpha, 1), (\alpha, 2)\}$. The graph H_k is easily seen to be a non-trivial quadrangulation of the sphere, so we will be done if we can find a retraction from $W_k \times K_2$ to H_k. Let $\varphi : V(W_k) \times V(K_2) \to V(H_k)$ be defined so that, for $1 \leq i \leq k$ and $j \in \{1, 2\}$, we have

$$\varphi(x_i, j) := (x_{i+1}, j) \text{ if } i \not\equiv j \bmod 2,$$

$$\varphi(x_i, j) := (x_i, j) \text{ if } i \equiv j \bmod 2,$$

$$\varphi(\alpha, j) := (\alpha, j)$$

(where indices are viewed modulo k). It is easily observed that this is, indeed, a retraction from $W_k \times K_2$ to H_k and so we are done. $\square$

5. Gadgets for Reflexive Triangulations

In this section, all graphs are assumed to be reflexive, unless otherwise stated. Our goal is to prove that H-Recolouring is PSPACE-complete when H is a reflexive graph which contains a substructure which resembles a triangulation near a vertex. The following definition is useful for defining the class of graphs that we consider.

Definition 5.1. Given a stiff reflexive graph H, we say that a set $S \subseteq V(H)$ is *listable* if there exists a reflexive S-gadget $X(x_1)$.

The following, somewhat cumbersome but quite broad, definition describes the class of graphs that we will consider.

Definition 5.2. A stiff reflexive graph H is *locally triangulated* around a vertex $0 \in V(H)$ if H contains a reflexive subgraph F such that

(a) $0 \in V(F)$,
(b) the subgraph of H induced by $V(F)$ is K_4-free,
(c) the neighbourhood of 0 in F contains a spanning cycle, say $12 \cdots k$,
(d) for $1 \leq i \leq k$, the neighbours i and $i + 1$ of 0 (where vertex labels are viewed modulo k) have a common neighbour β_i in $V(F)$ which is distinct from 0,
(e) for $1 \leq i \leq k$, the neighbourhood of β_i in F contains a spanning cycle,
(f) every pair $\{u, v\}$ where $uv \in E(F)$ and $u \neq v$ is listable and
(g) the sets $\{0, 2, 3\}$ and $\{0, 3, 4\}$ are listable.

The main focus of this section is on proving variants of Lemmas 2.5 and 2.6 for graphs which are locally triangulated around a vertex (Lemmas 5.6 and 5.7 below). Next, we show that locally triangulated graphs generalize K_4-free triangulations distinct from the reflexive triangle and, thus, Lemmas 2.5 and 2.6 follow from the results of this section.

Lemma 5.3. *Let H be a finite reflexive K_4-free triangulation which is not isomorphic to a reflexive triangle. Then H is locally triangulated around every vertex $0 \in V(H)$.*

Proof. We let $F := H$. Given this, it is trivial that conditions (a) and (b) of Definition 5.2 hold. In any triangulation, apart from the triangle, the neighbourhood of any given vertex induces a cycle. So, condition (c) of Definition 5.2 holds. As in Definition 5.2, we denote the neighbours of 0 by $1, \ldots, k$ where $12 \cdots k$ is a cycle.

If H is not stiff, then, by Observation 3.6 and since 0 is an arbitrary vertex and H is reflexive, we have, without loss of generality, $N(0) \subseteq N(1)$. However, if this were the case, then the vertices $0, 1, 2, 3$ would form a reflexive K_4, which is a contradiction. So, H is stiff.

The edge from i to $i + 1$ separates the face f incident to $0, i$ and $i + 1$ from another face f'. Let β_i be the vertex incident to f', distinct from i and $i + 1$. Since H is not a reflexive triangle, we must have that β_i is distinct from vertex 0 and so condition (d) of Definition 5.2 holds. The same argument that was used to prove (c) shows that (e) holds as well.

Now, given an edge $u, v \in E(H)$, we let x, y be the two vertices, distinct from u and v, which are incident to the two faces whose boundary contains the edge uv. Then u, v and x are the only common neighbours of u, v and x, and so

$$N(u) \cap N(v) \cap N(x) \cap N(y) = \{u, v\}.$$

Given this, and the fact that H is stiff, a $\{u, v\}$-gadget can be constructed by simply taking a copy of H coloured by the identity map in all canonical colourings and adding a reflexive signal vertex x_1 adjacent to u, v, x and y. So condition (f) of Definition 5.2 holds. Also, for any triangle x, y, z in H, we have

$$N(x) \cap N(y) \cap N(z) = \{x, y, z\}$$

which, via a similar argument to that which was used for (f), gives us condition (g). This completes the proof. $\qquad \square$

Remark 5.4. Of course, the condition that H is locally triangulated around a vertex covers a wide variety of graphs beyond K_4-free reflexive triangulations. For example, many reflexive graphs embedded on other surfaces (which need not even be triangulations of that surface, i.e., they can have larger faces) satisfy the criteria of Definition 5.2.

Thus, by Lemma 5.3, the following theorem, proven in the rest of this section, generalizes Theorem 1.4.

Theorem 5.5. *If H is a finite reflexive graph which is locally triangulated around a vertex 0, then H-RECOLOURING is PSPACE-complete when restricted to instances (G, f, g) such that G is reflexive.*

The aim in the rest of the section is to establish the following two lemmas which imply Theorem 5.5 via Lemma 2.4.

Lemma 5.6. *Let H be a finite reflexive graph which is locally triangulated around a vertex 0 and let $1, \ldots, k$ and $\beta_1, \ldots, \beta_k$ be as in Definition 5.2. Then, for $1 \leq i \leq k$, there exists*

- *a $\{(0, 0), (0, 1), (1, 0)\}$-gadget,*
- *a $\{(1, i), (1, \beta_i), (0, \beta_i)\}$-gadget and*
- *a $\{(1, i + 1), (1, \beta_i), (0, \beta_i)\}$-gadget.*

Lemma 5.7. *If H is a finite reflexive graph which is locally triangulated around a vertex 0, then for every neighbour 1 of 0 with $0 \neq 1$, there exists a not-all-zero gadget.*

5.1. **Not-Both-One Gadget for Locally Triangulated Graphs.** In the rest of this section H, always denotes a finite reflexive graph which is locally triangulated around a vertex 0. We define a directed graph analogously to Definition 3.9, guided by the graphs in Figure 8 rather than in Figure 4.

Definition 5.8. Let A be the set of all pairs $(x,y) \in V(H)^2$ such that $x \neq y$, $xy \in E(H)$ and $\{x,y\}$ is listable.

Definition 5.9. Let Φ be a directed graph on vertex set A where there is an arc from (a,b) to (c,d) if $ac, bc, bd \in E(H)$ and $ad \notin E(H)$.
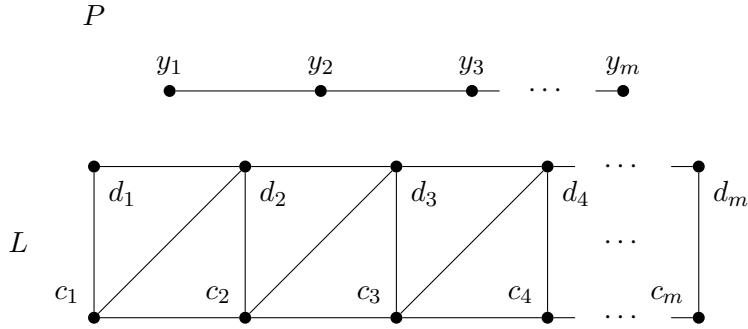

FIGURE 8. Graphs P and L motivating Φ of Defintion 5.9.

We remark that an important difference between the digraph Φ defined in this section and the one for irreflexive quadrangulations in Section 3 is that, for given an arc $(a,b)(c,d)$, the vertices b and c may actually coincide. This is the key property which allows us to get away with local arguments in this section.

Observation 5.10. There is an arc from (a,b) to (c,d) in Φ if and only if there is an arc from (d,c) to (b,a).

The following lemma is proved in a way which is analogous to the proof of Lemma 3.11. We omit the details.

Lemma 5.11. *If there is a directed path from (a_1, a_0) to (b_0, b_1) in Φ, then there exists a reflexive $\{(a_0, b_0), (a_1, b_0), (a_0, b_1)\}$-gadget.*

Thus, Lemma 5.6 is implied by the following lemma, via Lemma 5.11.

Lemma 5.12. *For $1 \leq i \leq k$, there is a path from (a_1, a_0) to (b_0, b_1) in Φ for the following choices of a_0, a_1, b_0, b_1*
  (i) $a_0 = b_0 = 0$ and $a_1 = b_1 = 1$,
 (ii) $a_0 = 1$, $a_1 = 0$, $b_0 = \beta_i$ and $b_1 = i$ and
(iii) $a_0 = 1$, $a_1 = 0$, $b_0 = \beta_i$ and $b_1 = i + 1$.

Proof. Since the subgraph of H induced by $V(F)$ is K_4-free, we know that vertex i is not adjacent to $i+2$ for $1 \leq i \leq k$. Therefore, for (i), we simply observe that the following is a directed path in Φ:

$$(1,0)(2,3)(3,4)\cdots(k-1,k)(0,1).$$

For $1 \leq j \leq k$, label the neighbours of β_j by $x_1^j, \ldots, x_{t_j}^j$ so that $x_1^j = j$ and $x_1^j \cdots x_{t_j}^j$ is a cycle. Again, since the subgraph of H induced by $V(F)$ is K_4-free, we know that x_ℓ^j and $x_{\ell+2}^j$ are non-adjacent. Also, 0 and β_j are non-adjacent. Thus, the following is a directed path in Φ:

$$(0,1)(2,\beta_1)(x_2^1, x_3^1) \cdots (x_{t_1-1}^1, x_{t_1}^1)(\beta_1, 2)(1,0).$$

We follow this by the directed path

$$(1,0)(2,3)(3,4) \cdots (i-2, i-1)(0,i)(i, \beta_i).$$

We complete the proof of (ii) by continuing along the following path

$$(i, \beta_i)(x_2^i, x_3^i) \cdots (x_{t_i-1}^i, x_{t_i}^i)(\beta_i, i).$$

The argument used to prove (ii) applies *mutatis mutandis* to prove (iii). □

5.2. **Not-All-Zero Gadget for Locally Triangulated Graphs.** Our next goal is to prove Lemma 5.7, which, when combined with Lemma 5.6 (i), will complete the proof of Theorem 5.5 via Lemma 2.4.

Proof of Lemma 5.7. Our goal is to construct a not-all-zero gadget $Z(z_1, z_2, z_3, z_4)$. Recall that the neighbours of the vertex 0 in H are labelled $1, \ldots, k$ for some $k \geq 4$ where consecutive neighbours modulo k are adjacent. In what follows, we will often refer to vertex 5 which, if $k = 4$, is regarded as the same as vertex 1.

As a first step, we apply Lemmas 5.6 and 5.11 to disjointly add

- a reflexive $\{(1, \beta_i), (0, \beta_i), (1, i)\}$-gadget $W_i(z_i, w_i)$ for $1 \leq i \leq 3$ and
- a reflexive $\{(1, \beta_4), (0, \beta_4), (1, 5)\}$-gadget $W_4(z_4, w_4)$.

Then, disjointly from the construction so far and from one another, use conditions (f) and (g) of Definition 5.2 to add

- a reflexive $\{0, 1\}$-gadget with signal vertex y_1,
- a reflexive $\{0, 2, 3\}$-gadget with signal vertex y_2,
- a reflexive $\{0, 3, 4\}$-gadget with signal vertex y_3 and
- a reflexive $\{0, 5\}$-gadget with signal vertex y_4.

Finally, add an edge from y_i to w_i for $1 \leq i \leq 4$ and add edges $y_1 y_2$, $y_2 y_3$ and $y_3 y_4$. See Figure 9.
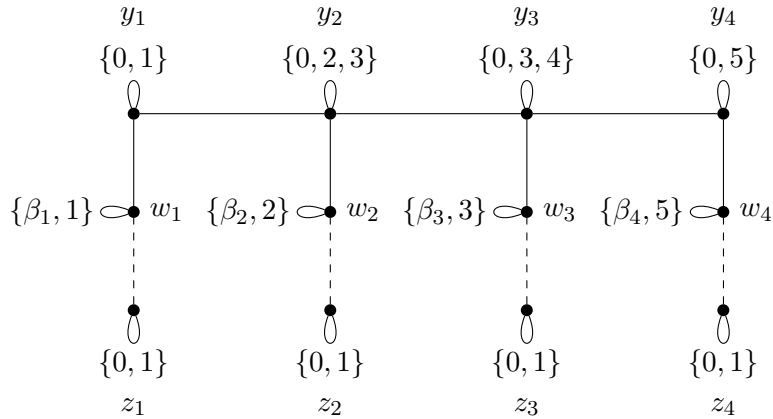

FIGURE 9. The vertices z_i, w_i and y_i for $1 \leq i \leq 4$ with their lists. Solid lines represent edges of Z and dashed lines represent gadgets which force w_i to map to β_i when z_i maps to 0.

Now, for $p = (p_1, p_2, p_3, p_4) \in \{0, 1\}^4 \setminus \{(0, 0, 0, 0)\}$, we define ζ_p as follows. If $p_i = 1$ for $1 \leq i \leq 3$, then we colour $W_i(z_i, w_i)$ with $\zeta_{(1,i)}$ and y_i with 0. Similarly, if $p_4 = 1$, colour $W_4(z_4, w_4)$ with $\zeta_{(1,5)}$ and y_4 with 0. On the other hand, if $p_i = 0$, then we colour $W_i(z_i, w_i)$ with $\zeta_{(0,\beta_i)}$. If there exists $j > i$ with $p_j = 1$, then we colour y_i with the smallest non-zero colour in its list and, otherwise, colour it with the largest such colour.

Let us now show that if ψ reconfigures to ζ_p for some $p \in \{0, 1\}^4 \setminus \{(0, 0, 0, 0)\}$, then at least one of $\psi(z_1), \ldots, \psi(z_4)$ is one. The gadgets $W_1, \ldots, W_4$ imply that, if $\psi(z_1) = 0$, then $\psi(w_1) = \beta_1$ which implies $\psi(y_1) = 1$ and if $\psi(z_2) = 0$, then $\psi(y_2) \in \{2, 3\}$ which means that it must be equal to 2 because $13 \notin E(H)$ since $V(F)$ induces a K_4-free subgraph of H. By the same argument, $\psi(y_4)$ must be 5 and $\psi(y_3)$ must be 4, which is a contradiction because $24 \notin E(H)$ since $V(F)$ induces a K_4-free subgraph of H.

Now, suppose that $p, q \in \{0, 1\}^4 \setminus \{(0, 0, 0, 0)\}$ differ on exactly one coordinate. Without loss of generality, there is some i such that $p_i = 0$ and $q_i = 1$. Let j be any coordinate such that $p_j = q_j = 1$. Starting with ζ_p, we can reconfigure the colouring on the gadget $W_i(z_i, w_i)$ so that z_i maps to 1 and w_i does not map to β_i. We can then change the colour of y_i to 0. After doing this, we change the colours on the vertices $y_{i'}$ for i' between j and i so that they match their colours under ζ_q (and we can do that without modifying the colourings on the non-signal vertices of the gadgets $W_{i'}(z_{i'}, w_{i'})$). This completes the proof. $\qquad\square$

References

[1] P. Bonsma. Independent set reconfiguration in cographs and their generalizations. *J. Graph Theory*, 83(2):164–195, 2016.

[2] P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoret. Comput. Sci.*, 410(50):5215–5226, 2009.

[3] R. C. Brewster, J.-B. Lee, B. Moore, J. A. Noel, and M. Siggers. Graph homomorphism reconfiguration and frozen H-colourings. arXiv:1712.00200v1, preprint, 2017.

[4] R. C. Brewster, J.-B. Lee, and Mark Siggers. Recolouring reflexive digraphs. *Discrete Math.*, 341(6):1708–1721, 2018.

[5] R. C. Brewster, S. McGuinness, B. Moore, and J. A. Noel. A dichotomy theorem for circular colouring reconfiguration. *Theoret. Comput. Sci.*, 639:1–13, 2016.

[6] R. C. Brewster and J. A. Noel. Mixing homomorphisms, recolorings, and extending circular precolorings. *J. Graph Theory*, 80(3):173–198, 2015.

[7] G. R. Brightwell and P. Winkler. Gibbs measures and dismantlable graphs. *J. Combin. Theory Ser. B*, 78(1):141–166, 2000.

[8] L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *J. Graph Theory*, 67(1):69–82, 2011.

[9] C. Feghali, M. Johnson, and D. Paulusma. A reconfigurations analogue of Brooks' theorem and its consequences. *J. Graph Theory*, 83(4):340–358, 2016.

[10] C. Feghali, M. Johnson, and D. Paulusma. Kempe equivalence of colourings of cubic graphs. *European J. Combin.*, 59:1–10, 2017.

[11] P. Gopalan, P. G. Kolaitis, E. Maneva, and C. H. Papadimitriou. The connectivity of Boolean satisfiability: computational and structural dichotomies. *SIAM J. Comput.*, 38(6):2330–2355, 2009.

[12] S. T. Hedetniemi. Homomorphisms of graphs and automata. Technical report, University of Michigan, 1966.

[13] J. van den Heuvel. The complexity of change. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 127–160. Cambridge Univ. Press, Cambridge, 2013.

[14] T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoret. Comput. Sci.*, 412(12-14):1054–1065, 2011.

[15] T. Ito and A. Suzuki. Web survey on combinatorial reconfiguration. Retrieved June 12, 2018 from `http://www.ecei.tohoku.ac.jp/alg/coresurvey/`.

[16] M. Johnson, D. Kratsch, S. Kratsch, V. Patel, and D. Paulusma. Finding shortest paths between graph colourings. *Algorithmica*, 75(2):295–321, 2016.

[17] B. Larose. Taylor operations on finite reflexive structures. *Int. J. Math. Comput. Sci.*, 1(1):1–21, 2006.

[18] J.-B. Lee, J. A. Noel, and M. Siggers. Reconfiguring homomorphisms to triangle-free reflexive graphs. in preparation.

[19] L. Lovász. Kneser's conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25(3):319–324, 1978.

[20] B. Mohar. Kempe equivalence of colorings. In *Graph theory in Paris*, Trends Math., pages 287–297. Birkhäuser, Basel, 2007.

[21] M. Mühlenthaler. Degree-constrained subgraph reconfiguration is in P. In *Mathematical foundations of computer science 2015. Part II*, volume 9235 of *Lecture Notes in Comput. Sci.*, pages 505–516. Springer, Heidelberg, 2015.

[22] N. V. Ngoc and Z. Tuza. 4-chromatic graphs with large odd girth. volume 138, pages 387–392. 1995. 14th British Combinatorial Conference (Keele, 1993).

[23] N. Nishimura. Introduction to reconfiguration. *Algorithms (Basel)*, 11(4):Paper No. 52, 25, 2018.

[24] C. Tardif and M. Wrochna. Hedetniemi's conjecture and strongly multiplicative graphs. arXiv:1808.04778v1, preprint, 2018.

[25] M. Wrochna. Reconfiguration and structural graph theory. Master's thesis, University of Warsaw, Poland, 2014.

[26] M. Wrochna. Homomorphism reconfiguration via homotopy. In *32nd International Symposium on Theoretical Aspects of Computer Science*, volume 30 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 730–742. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2015.

[27] M. Wrochna. Square-free graphs are multiplicative. *J. Combin. Theory Ser. B*, 122:479–507, 2017.

[28] M. Wrochna. *The topology of solution spaces of combinatorial problems.* PhD thesis, University of Warsaw, Poland, 2018.

(Jae-Baek Lee and Mark Siggers) College of Natural Sciences, Kyungpook National University, Daegu 702-701, South Korea

E-mail address: `dlwoqor0923@gmail.com, mhsiggers@knu.ac.kr`

(Jonathan A. Noel) Department of Computer Science and DIMAP, University of Warwick, Coventry CV4 7AL.

Current address: Mathematics Institute and DIMAP, University of Warwick, Coventry CV4 7AL.

E-mail address: `j.noel@warwick.ac.uk`