2024

Graph Coloring Reconfiguration

Reem Mahmoud
Virginia Commonwealth University

GRAPH COLORING RECONFIGURATION

A dissertation submitted in partial fulfillment
of the requirements for the degree
of Doctor of Philosophy in Computer Science
at Virginia Commonwealth University

by
Reem Mahmoud

B.S. in Mathematics, American University of Sharjah, 2019
M.S. in Mathematics, Virginia Commonwealth University, 2021

Director: Daniel Cranston
Associate Professor, Department of Computer Science

Virginia Commonwealth University
Richmond, Virginia
May, 2024

Acknowledgments

In the name of Allah, the Most Gracious, the Most Merciful.

First and foremost, all thanks and praise goes to Allah. It is by Allah's will and grace that I was able to embark on this journey and achieve my goals.

Thanks to my parents for being present every step of the way. You were, and always will be, a constant source of support and encouragement. I can never repay you for everything you have done for me. Mom, dad, without you, I would not be here. Mom, you are my reason. Thank you for giving me the opportunity to make you both proud.

Thanks to my twin sister, Roaa. You are my best friend. I would not have been able to endure these past five years without you by my side. Thanks for being my lifelong partner in everything.

Thanks to my advisor, Daniel Cranston. If I had not taken your class that first semester, I would not have found my passion for graph theory. Thanks for always challenging me and for pushing me to grow and be better. I could not have asked for a better advisor. Thanks for making my journey enjoyable and worthwhile.

Thanks to my committee members Nicolas Bousquet, Thang Dinh, Amer Mouawad, and Hong-Sheng Zhou for useful feedback and comments.

Thanks to my friends for always being there to cheer me up, endure my rants, and celebrate my wins with me. You made me feel right at home.

Finally, I dedicate this to mom and dad.

Table of Contents

List of Figures

List of Tables

Abstract

Reconfiguration is the concept of moving between different solutions to a problem by transforming one solution into another using some prescribed transformation rule (move). Given two solutions s_1 and s_2 of a problem, reconfiguration asks whether there exists a sequence of moves which transforms s_1 into s_2. Reconfiguration is an area of research with many contributions towards various fields such as mathematics and computer science.

The k-COLORING RECONFIGURATION problem asks whether there exists a sequence of moves which transforms one k-coloring of a graph G into another. A move in this case is a type of *Kempe swap*. An α, β-*Kempe swap* in a properly colored graph interchanges the colors on some component of the subgraph induced by vertices colored α and β. Two k-colorings of a graph are k-*equivalent* if we can form one from the other by a sequence of Kempe swaps (never using more than k colors). The k-COLORING RECONFIGURATION problem has applications in statistical physics and mechanics; in particular, it has positive implications on the Markov chains defined for the Ising model in statistical physics and the antiferromagnetic Pott's model in statistical mechanics.

The reconfiguration graph $\mathcal{C}_k(G)$ associated with the k-COLORING RECONFIGURATION problem is defined as follows: The vertices of $\mathcal{C}_k(G)$ are the k-colorings of G and two vertices in $\mathcal{C}_k(G)$ are adjacent if their corresponding k-colorings differ by a single move. We study $\mathcal{C}_k(G)$ for certain classes of graphs G. In particular, we study the connectedness and the diameter of $\mathcal{C}_k(G)$. Indeed, $\mathcal{C}_k(G)$ being connected is equivalent to the k-colorings of G being pairwise k-equivalent. On the other hand, the diameter of $\mathcal{C}_k(G)$ is d if and only if for every two k-colorings of G, there is a sequence of length at most d from one to the other. Our results on the connectedness of $\mathcal{C}_k(G)$ imply the Markov chain for the models mentioned above is ergodic, while lower bound results on the diameter of $\mathcal{C}_k(G)$ imply lower bounds on the mixing time of the Markov chain.

It is conjectured that the diameter of $\mathcal{C}_k(G)$ is $O(n^2)$ for every d-degenerate graph G whenever $k \geq d+2$.

As a step towards proving this conjecture for triangle-free planar graphs, we show that $\mathcal{C}_5(G)$ is $O(n^2)$ for every planar graph G with no 3-cycles and no 5-cycles.

Moreover, we study the connectedness of $\mathcal{C}_5(G)$ for the *triangulated toroidal grid*, $G := T[m \times n]$, which is formed from (a toroidal embedding of) the Cartesian product of C_m and C_n by adding parallel diagonals inside all 4-faces. We prove that all 5-colorings of $T[m \times n]$ are 5-equivalent when $m, n \geq 6$, i.e., $\mathcal{C}_5(G)$ is connected for such toroidal graphs G.

We also explore list-coloring reconfiguration. For a list-assignment L and an L-coloring φ, a Kempe swap is called *L-valid* for φ if performing the Kempe swap yields another L-coloring. Two L-colorings are called *L-equivalent* if we can form one from the other by a sequence of a type of L-valid Kempe swaps. The associated reconfiguration graph in this case is $\mathcal{C}_L(G)$. Let G be a connected k-regular graph with $k \geq 3$. We prove that if L is a k-assignment, then all L-colorings are L-equivalent unless $G \cong K_2 \square K_3$, i.e., $\mathcal{C}_L(G)$ is connected for every k-assignment L. This generalizes an analogous result about the k-colorings of k-regular graphs.

Chapter 1

Introduction

Reconfiguration is the concept of moving between different states of a system by repeatedly applying some prescribed transformation rule. A *state* may refer to the set of configurations of a combinatorial structure or the set of solutions of a problem. The *state space* of a system encompasses all the possible states of the system. A classic example of reconfiguration is the Rubik's cube, invented in 1974 by Hungarian architecture and design professor Ernő Rubik. A player is given an arbitrary configuration of the cube (Figure 1.1a) and tries to reach the monochromatic configuration (Figure 1.1c) by rotating the faces of the cube (Figure 1.1b). Each configuration of the cube is a state, and each rotation corresponds to applying the transformation rule to move from one configuration to another. The state space of the Rubik's cube consists of all the possible configurations of the cube. A single application of the transformation rule, in this case, rotating a face of the cube, is a *reconfiguration step*.

state

state space

reconfiguration step

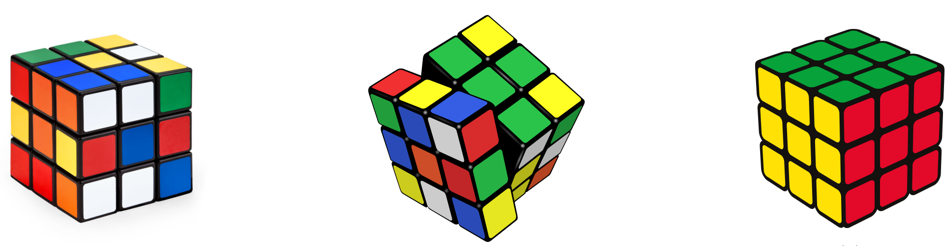

(a) (b) (c)

Figure 1.1: A Rubik's cube that is (a) scrambled[1], (b) rotating[2], and (c) solved[3].

Reconfiguration is an area of research with many contributions to problems in mathematics and computer

[1]wordpress.com, classteaching, 2019.

[2]Wikipedia.com, Booyabazooka, 2006.

[3]youcandothecube.com, Ron Koziol, Pinterest.

science. It draws together challenges and methods from a wide range of different fields including graph theory, combinatorial game theory, probability theory, random sampling, enumeration, complexity theory, discrete geometry, statistical physics, and various others. For example, in computer science, one can study 3-SAT reconfiguration [18] which asks whether it is possible to transform one satisfying assignment of a Boolean formula (with each clause having at most 3 literals) into another using single variable "flips." On the other hand, in geometry, one can study graph morphing [1] which asks whether it is possible to transform one drawing of a graph into another while preserving certain geometric properties.

Reconfiguration problems also appear in many games and puzzles. One example is the Rubik's cube mentioned above. Other examples include sliding puzzles [44] like the 15-puzzle, sliding blocks, Sokoban, and Rush Hour. A player must slide objects on a surface/board starting from an arbitrary configuration with the goal of achieving a desired configuration. Moreover, reconfiguration problems arise in numerous real-world scenarios. One instance is robot motion planning [29] where the goal is to find a sequence of feasible moves that gets the robot from point A to point B. Such problems are used in the development of smart cars and in robot-assisted surgery. Another instance is genome rearrangement [5] where the goal is to determine whether one genome can be transformed into another via a sequence of gene mutations. Such problems contribute to the study of evolution and mutation in plants.

A real-world application that perhaps captures the main aspects of the reconfiguration framework is the POWER SUPPLY RECONFIGURATION problem [52]; see Figure 1.2. Consider a set P of power stations each with a fixed capacity of power supply and a set H of houses each with a fixed amount of power required. Each house has a number of power stations that can supply it with power. Such power stations are *viable* for that house (represented by dashed edges in Figure 1.2). A *valid* assignment $\varphi : H \to P$ assigns a power station for each house (represented by bold edges in Figure 1.2) from its set of viable power stations in such a way that no power station exceeds its capacity. Given an assignment of power φ, a *power transfer* for house h is a reassignment from power station $\varphi(h)$ to some other power station $p \neq \varphi(h)$. Given two valid assignments φ_1 and φ_2, the POWER SUPPLY RECONFIGURATION problem asks whether there is a sequence of power transfers that transforms φ_1 into φ_2 such that each intermediate assignment is also valid.

From a real-world perspective, this situation might represent a power company with multiple generators P which supply power to multiple houses H. From time to time, customers might want to transfer generators (for example, to gain access to extra services) and the company is only able to transfer power one customer at a time so as not to overload the system. Figure 1.2 shows a two-step reconfiguration sequence that transforms the first assignment φ_1 (left) into the last assignment φ_2 (right). For simplicity, assume the power stations

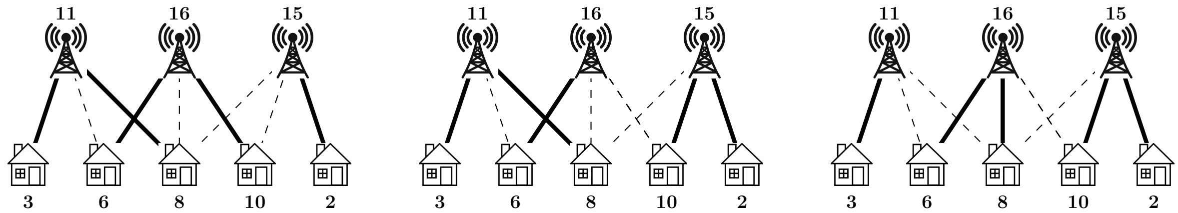

Figure 1.2: The power stations[4] and houses[5] are labeled from left to right p_1, p_2, p_3 and h_1, h_2, h_3, h_4, h_5. A dashed edge indicates a viable power supply station. Bold edges represent a power assignment. The assignment on the left can be transformed into the one on the right in 2 steps by first transferring power for h_4 from p_2 to p_3, then transferring power for h_3 from p_1 to p_2.

and houses are labeled from left to right p_1, p_2, p_3 and h_1, h_2, h_3, h_4, h_5. This sequence might be interpreted as customer h_3 wanting to switch to generator p_2. But in order to do that, the company must first switch h_4 to p_3 (or h_2 to p_1). Otherwise, p_2 would exceed its capacity, resulting in an invalid assignment.

The POWER SUPPLY RECONFIGURATION problem introduces two important aspects of the reconfiguration framework. The first is an underlying problem whose solutions constitute the states of the reconfiguration problem. Given sets P and H along with their capacities and requirements, consider the problem POWER SUPPLY which asks whether there exists a valid assignment for P and H. From that problem we can derive the POWER SUPPLY RECONFIGURATION problem which takes two solutions of POWER SUPPLY along with an appropriate transformation rule and asks whether we can transform one solution into the other. The second aspect is the condition that each term in the sequence must be a valid assignment. For example, transferring power for h_3 from p_1 to p_2, then transferring power for h_4 from p_2 to p_3 is another sequence (different from the one shown in Figure 1.2) which transforms φ_1 into φ_2, by definition of power transfer. However, the first power transfer (reconfiguration step) does not result in a valid solution for POWER SUPPLY since p_2 is overloaded (exceeds its capacity). Thus, such a sequence is not a solution for POWER SUPPLY RECONFIGURATION. In general, our aim is to study the solution space of the underlying problem. So, we are interested in sequences that only use solutions of POWER SUPPLY.

reconfiguration framework

The *reconfiguration framework*, therefore, comprises the following aspects. An underlying problem X, a definition for a solution of X, and a definition for a type of move which transforms a current solution of X into a new solution. Subsequently, the reconfiguration problem for X examines transforming one solution of X into another through a sequence of moves such that each intermediate term in the sequence is also a solution of X. Furthermore, moves are always *reversible* (this is true for all reconfiguration problems mentioned throughout this dissertation). That is, if there exists a move m_1 which can transform state A into

[4] shutterstock.com, Larry-Rains, 2014.
[5] redfin.com, StellarMLS, 2024.

state B, then there exists a move m_2 which can transform B into A; in particular, m_2 "undoes" or *reverses* m_1. Thus, a sequence of moves from A to B is also a sequence from B to A. The reconfiguration framework was properly formalized by Ito et. al. in [45] which sparked an interest in the area and led to a series of follow-up papers.

A popular line of research considers search problems on graphs as underlying problems for the reconfiguration framework. The goal in a search problem is to find a certain structure inside the graph. For example, given a graph G and a positive integer k, the search problem INDEPENDENT SET finds an independent set[6] of size at least k (a k^+*-independent set*) in G. As a result, INDEPENDENT SET RECONFIGURATION asks whether a k^+-independent set of G can be transformed into another via a sequence reconfiguration steps such that each intermediate set is also k^+-independent. Similarly, DOMINATING SET finds a dominating set[7] of size at most k (a k^-*-dominating set*) in G and DOMINATING SET RECONFIGURATION asks if it is possible to move between different k^--dominating sets of G. In fact, both of those problems fall under the *token framework* which states that these sets can be viewed as configurations of tokens placed on the vertices of the graph. Interestingly, this formulation using tokens allows for multiple definitions of the transformation rule. The *Token Addition/Removal* (TAR) model defines a move to be the addition or removal of a token from a vertex in the set. The *Token Jumping* (TJ) model defines a move, called a *jump*, to be removing a token from a vertex and placing it on another vertex. Finally, the *Token Sliding* (TS) model defines a move, called a *slide*, to be removing a token from a vertex and placing it on one of its neighbors. Note that the latter two models are specifically defined for reconfiguration problems where the size of the sets is exactly k and is preserved throughout.

token framework

TAR

TJ

TS

A key example of a well-studied reconfiguration problem with an underlying graph search problem is k-COLORING RECONFIGURATION. Given a graph G and a positive integer k, the k-COLORING[8] search problem finds a k*-coloring* of G, i.e., a coloring of the vertices of G using at most k colors such that no two adjacent vertices get the same color. Naturally, the k-COLORING RECONFIGURATION problem asks if k-colorings of G can be transformed into one another via a sequence of "color swaps" (see Section 1.4 for a formal definition). The purpose of this dissertation is to provide an in-depth study of the k-COLORING RECONFIGURATION problem and present our contributions to the problem.

We now turn to some of the main questions of interest within the reconfiguration community. These

[6] An independent set is a set of vertices that are pairwise nonadjacent.

[7] A dominating set is a set of vertices such that every vertex not in the set has a neighbor in the set.

[8] Typically, graph search problems such as INDEPENDENT SET, DOMINATING SET and k-COLORING are phrased as decision problems (Yes or No questions). However, we phrase them as above to highlight the fact that the states of corresponding reconfiguration problems are, in fact, the Yes-solutions to the search problems. Often, if we can answer a decision problem, it is easy to leverage that to generate a solution.

questions generally pertain to reachability of solutions or connectivity of the solution space of the underlying problem and its associated reconfiguration problem. An *equivalence class* of the solution space is a maximal subset of solutions that are all pairwise reconfigurable to each other. Given an underlying problem with a fixed discrete solution space, we can ask the following:

equivalence class

1. Can we transform every solution into every other through a sequence of reconfiguration steps?
2. What is the number of equivalence classes of the solution space?
3. What is the smallest number of reconfiguration steps k for which every solution can be reconfigured to every other by at most k steps?

Observe that if the answer to Question 1 is No, then the answer to Question 3 is ∞. Moreover, Questions 1 and 2 are closely related. In particular, the answer to Question 1 is Yes if and only if the answer to Question 2 is 1. Furthermore, Question 1 is phrased as a *decision problem*, i.e., a Yes-or-No question. This is because, ideally, we would like an algorithmic proof for Question 1. However, an existence proof is often sufficient for making valuable conclusions about the solution space. A Yes answer for Question 1 means that the solution space is connected, which has positive implications with respect to sampling and enumeration (see Chapter 2.4). In other words, we generally aim to show that a sequence from one solution to another exists as opposed to providing one. However, often our existence proofs could be revised to provide algorithms for constructing this sequence. Note that once a sequence is found, a natural next step is to look for an "optimal" sequence, which motivates Question 3.

decision problem

In the case of the Rubik's cube, the answer to Question 1 is No. This is because some configurations, such as ones formed from the monochromatic cube by rotating a single corner cubelet or by flipping a single edge cubelet, are not reachable through face rotations. In fact, such configurations can only be obtained by manually disassembling the cube and then reassembling it as described[9]. As it turns out, the answer to Question 2 for the Rubik's cube is 12, and each of the aforementioned configurations (monochromatic cube, monochromatic cube with a rotated corner cubelet, monochromatic cube with a flipped edge cubelet) belongs to a different equivalence class. Question 3 is, therefore, most interesting for the equivalence class R containing the monochromatic cube. In 2010, almost three decades after the invention of the puzzle, it was proven that any configuration of the cube belonging to R can be solved, i.e., transformed into the monochromatic cube, in at most 20 face rotations[10] [59].

[9] or peeling off the stickers.

[10] There are configurations for which exactly 20 moves are needed.

We also commonly explore the computational complexity of answering reconfiguration questions. Simply put, how hard is it for a computer, in terms of time and space (computer memory), to answer such questions? The following are widely studied questions on complexity in reconfiguration:

(a) What is the computational complexity of determining whether the solution space is connected?

(b) What is the computational complexity of determining whether two given solutions are reconfigurable to one another?

(c) What is the computational complexity of finding a shortest sequence between two given solutions?

Reconfiguration problems are often modeled using graphs. For a given reconfiguration problem, we define a *reconfiguration graph*, which reformulates the problem in the language of graph theory. Each vertex of the reconfiguration graph represents a solution of the underlying problem, and two vertices are adjacent when their corresponding solutions are reconfigurable to one another by a single reconfiguration step; that is, a single application of the transformation rule. Thus, transforming one solution into another is equivalent to finding a path in the associated reconfiguration graph between their corresponding vertices. This reconfiguration graph formulation allows us to examine Questions 1–3 in the context of graph theory concepts and invariants. In particular, Questions 1–3 can now be stated as follows.

reconfiguration graph

(i) Is the reconfiguration graph connected?

(ii) What is the number of components of the reconfiguration graph?

(iii) What is the diameter of the reconfiguration graph?

Consequently, the reconfiguration graph for the Rubik's cube is disconnected (the answer to Question (i) is No), has 12 components[11] (the answer to Question (ii) is 12), and has diameter 20 (the answer to Question (iii) is 20).

Recall that our main topic of interest for this dissertation is the k-COLORING RECONFIGURATION problem, which considers moving between k-colorings of a graph G using "color swaps." The reconfiguration graph for this problem has a vertex for every k-coloring of G, and two vertices are adjacent whenever their corresponding k-colorings are reconfigurable to one another by a single "color swap." Most of our work on the k-COLORING RECONFIGURATION problem goes into solving (i) above for special classes of graphs G, though

[11] In fact, the 12 classes are all cosets, so the components of the reconfiguration graph they induce are all isomorphic.

we also explore (iii). Further, since all our results for (i) are Yes-answers, they also imply an answer of 1 for (ii). Finally, it is worth noting that the transformation rule for this problem can be defined in two ways, resulting in different variations of the problem (see Chapter 2). We explore both variations and provide new results for each in this dissertation.

1.1 Organization

We now give a general outline of the dissertation. The rest of this chapter recalls basic terminology and definitions of graph theory, graph coloring, and computational complexity. In Section 1.2, we give a brief overview of graph theory notation and concepts as well as relevant classes of graphs. Our notation is standard and can be found in most graph theory textbooks. In Section 1.3, we define proper coloring and list coloring. In Section 1.4, we introduce Kempe swaps, which constitute the key operation we use to transform one coloring of a graph into another. In Section 1.5, we briefly introduce computational complexity and recall relevant complexity classes. A reader who is familiar with the contents of these sections is free to skip them.

In Chapter 2, we take a closer look at the k-COLORING RECONFIGURATION problem and its list-coloring variant, and we examine the associated reconfiguration graph. We also introduce the two main models of transformation for coloring reconfiguration problems. In Section 2.1, we discuss the recoloring model, where a reconfiguration step is a vertex-recoloring, and survey important results in the literature on it. In Section 2.2, we consider the Kempe swap model, where a reconfiguration step is a Kempe swap, and go over previous results on it. In Section 2.3, we discuss computational complexity results for both models. Lastly, in Section 2.4, we examine applications of coloring reconfiguration in statistical mechanics and physics.

Chapters 3–5 contain our own results and contributions to the k-COLORING RECONFIGURATION problem. In Chapter 3, we explore the recoloring model for planar graphs with no short odd cycles. We show that the 5-coloring reconfiguration graph for planar graphs with no 3-cycles and no 5-cycles has quadratic diameter. In Chapter 4 we take a look at the Kempe swap model for toroidal graphs. We prove that the 5-coloring reconfiguration graph for most 6-regular toroidal graphs is connected. Finally, in Chapter 5, we consider the list-coloring Kempe swap model for regular graphs. For $k \geq 3$, we show that k-coloring reconfiguration graph for (almost all) k-regular graphs is connected.

1.2 Graph Theory

In this section, we lay out relevant definitions and concepts from graph theory. Throughout this paper we follow standard graph theory notation used in [63]. For a more detailed review, we refer the reader to any introductory graph theory textbook such as [63] or [31].

The *neighborhood* of v in G, denoted $N_G(v)$, is the set of vertices adjacent to v, or the set of *neighbors* of v. The *closed neighborhood* of v in G, denoted $N_G[v]$, is $N_G(v) \cup \{v\}$. For a set $S \subseteq V(G)$, we define $N(S) := \cup_{s \in S} N(s)$. Similarly, $N[S] := \cup_{s \in S} N[s]$. The *degree* of v in G, denoted $d_G(v)$, is $|N_G(v)|$. When it is clear that we are working in G, we usually drop the subscript. If $d(v) = 0$, then v is *isolated.* If $d_G(v) = k$ for every $v \in V(G)$, then G is *k-regular*. The *minimum degree* of G, denoted $\delta(G)$, is $\min_{v \in V(G)} d(v)$. The *maximum degree*, denoted $\Delta(G)$, is $\max_{v \in V(G)} d(v)$. The *maximum average degree*, denoted $mad(G)$, is $\max_{H \subseteq G} \frac{\sum_{v \in V(G)} d(v)}{|V(G)|}$.

(closed) neighborhood

k-regular

$\delta(G), \Delta(G)$

A graph H is a *subgraph* of G, denoted $H \subseteq G$, if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. Moreover, it is a *spanning* subgraph if $V(H) = V(G)$. A graph H is an *induced subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) = \{vw \in E(G) | v, w \in V(H)\}$. We use $G[V(H)]$ (or $G[H]$) to denote the subgraph induced by $V(H)$. For a set $S \subseteq (V(G) \cup E(G))$, we write $G - S$ to mean $G \setminus S$ and $G + S$ to mean $G \cup S$. If $S = \{s\}$, we write $G - s$ and $G + s$. A graph is *H-free* if it does not contain H as a subgraph. For a set $S \subseteq V(G)$, by *identifying* the vertices in S we mean forming a new graph G' by replacing S in G with a single vertex w' such that $N_{G'}(w') = N_G(S)$ (we replace each maximal set of parallel edges with a single edge). A graph G is *isomorphic* to H, denoted $G \cong H$, if there exists $\varphi : V(G) \to V(H)$ such that $uv \in E(G)$ if and only if $\varphi(u)\varphi(v) \in E(H)$.

(induced/ spanning) subgraph

identifying

isomorphic

A *clique* is a set of pairwise adjacent vertices. An *independent set* is a set of pairwise nonadjacent vertices. A *matching* M of G is a subset of $E(G)$ with no two edges in M sharing an endpoint. A *perfect matching* M is a matching such that every $v \in V(G)$ is an endpoint of some $e \in M$. A *path* P is a sequence of edges $e_1, e_2, \ldots, e_m$ that join a sequence of distinct vertices $v_1, v_2, \ldots, v_{m+1}$ such that $v_i v_{i+1} = e_i$ for every $i \in \{1, \ldots, m\}$. Its length $\ell(P)$ is $|E(P)|$. A *v, w-path* is a path starting at v and ending at w. The non-endpoints of a path are its *internal vertices.* Two paths are *internally disjoint* if, with the possible exception of their endpoints, have no common vertices.

clique

independent set

(perfect) matching

$\ell(P)$

v, w-path

internally disjoint

A graph G is *connected* if there is a v, w-path for every pair $v, w \in V(G)$; otherwise, it is *disconnected.* The *components* of G are its maximal connected subgraphs. Hence, if G is connected, then it has a single

connected

components

component: itself. A component is *trivial* if it is an isolated vertex. For all $v, w \in V(G)$, the length of the shortest v, w-path in G is denoted $dist(v, w)$. If there is no v, w-path, then $\text{dist}(v, w) = \infty$. The *diameter* of G, denoted $diam(G)$, is $\max_{v,w \in V(G)} \text{dist}(v, w)$. For brevity, we denote the set $\{1, 2, \ldots, k\}$ with $[k]$.

trivial

$\text{dist}(v, w)$

$\text{diam}(G)$

1.2.1 Graph Classes

We now introduce relevant classes of graphs.

- Path on n vertices, P_n: A graph with vertices $v_1, v_2, \ldots, v_n$ and edges $v_1v_2, v_2v_3, \ldots, v_{n-1}v_n$; see Figure 1.3e. (path)

- Cycle on n vertices, C_n: A graph with vertices $v_1, v_2, \ldots, v_n$ and edges $v_1v_2, v_2v_3, \ldots, v_{n-1}v_n, v_nv_1$; see Figure 1.3a. A *chord* of a cycle C is an edge $e \notin E(C)$ with endpoints in $V(C)$. A graph is *chordal* if every cycle of length at least 4 contains a chord. The *girth* of a graph is the length of its shortest cycle. A *tree* is a connected graph with no cycles. A collection of trees is a *forest*. (cycle, chordal, girth, tree)

- Complete graph on n vertices, K_n: A graph with vertices $v_1, \ldots, v_n$ and edges $\{v_iv_j : 1 \leq i < j \leq n\}$; see Figure 1.3b. (complete)

- Bipartite graph: A graph whose vertex set can be partitioned into two independent sets, or *parts*, X and Y; see Figure 1.3c. (bipartite)

- Complete bipartite graph, $K_{s,t}$: A bipartite graph with parts X and Y such that $|X| = s$ and $|Y| = t$, and with edge set $\{xy : x \in X, y \in Y\}$; see Figure 1.3d. (complete bipartite)

- Planar graph: A graph that can be embedded (drawn without crossing edges) in the plane. Figures 1.3a, 1.3c, 1.3e, and 1.3f are planar. A planar graph together with an embedding is a *plane graph*. A planar graph divides the plane into a set of regions called *faces*. The unbounded region is the *outer face*. A planar graph with an embedding such that all vertices lie on the outer face is *outerplanar*. Every planar graph has an embedding with all edges straight. (planar, (outer) face)

- Toroidal graph: A graph that can be embedded in the torus. All planar graphs are toroidal; see Figure 1.3g. However, graphs like K_5, K_6, K_7 are toroidal, but nonplanar. (toroidal)

- Cartesian product $G \square H$ of graphs G and H: A graph with vertex set $V(G) \times V(H)$ and edge set $\{(g, h)(g', h') : g = g' \text{ and } hh' \in E(H), \text{ or } h = h' \text{ and } gg' \in E(G)\}$; see Figure 1.3g. (Cartesian product)

- Tensor product $G \times H$ of graphs G and H: A graph with vertex set $V(G) \times V(H)$ and edge set $\{(g,h)(g',h') : gg' \in E(G) \text{ and } hh' \in E(H)\}$; see Figure 1.3c. Tensor product

- Generalized n-dimensional hypercube, $Q_n(S)$: A graph where vertices represent sequences of length n with entries in S, and two vertices are adjacent if their sequences differ in a single entry; see Figure 1.3f. hypercube

- k-degenerate graph: A graph for which there exists a *degeneracy ordering* σ of its vertices; that is, every vertex has at most k neighbors earlier in σ. Equivalently, every subgraph of a k-degenerate graph contains a vertex v with $d(v) \leq k$. The *degeneracy* of a graph G, denoted *deg*(G), is the smallest integer k such that G is k-degenerate. Alternatively, $\deg(G) := \max_{H \subseteq G} \delta(H)$. The following is well-known: Toroidal graphs are 6-degenerate, planar graphs are 5-degenerate, triangle-free planar graphs are 3-degenerate, planar graphs of girth at least 6 are 2-degenerate, and forests are 1-degenerate. k-degenerate deg(G)

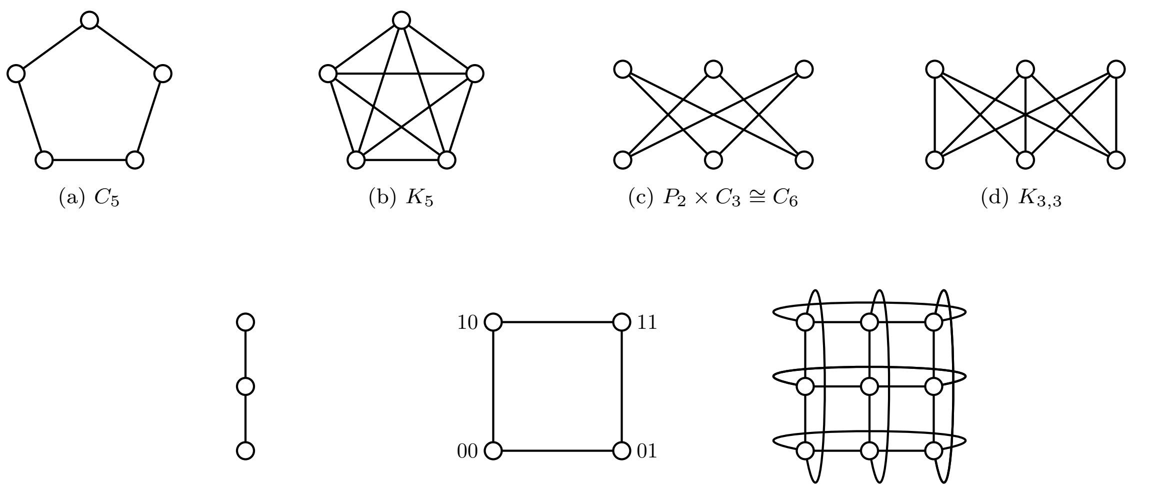

(a) C_5 (b) K_5 (c) $P_2 \times C_3 \cong C_6$ (d) $K_{3,3}$

(e) P_3 (f) $Q_2(\{0,1\})$ (g) $C_3 \square C_3$

Figure 1.3: A gallery of graphs

1.3 Graph Coloring

We now briefly introduce graph coloring and some of its many applications.

Graph coloring first arose when Francis Guthrie was trying to color a map of the counties of England so that every two regions sharing a common border received distinct colors. It is said that he conjectured four colors are enough, and that Guthrie's brother relayed the problem to Augustus De Morgan who, in turn, relayed it to William Hamilton in a letter in 1852. This led to the famous Four-Color problem which considers coloring the faces of a planar graph. A proof for four colors was first claimed by Alfred Kempe in 1879, and was believed to be correct until 1890 when Percy John Heawood pointed out a flaw in Kempe's argument. However, using Kempe's ideas, Heawood proved that five colors are sufficient. In 1976, Kenneth Appel and Wolfgang Haken finally proved the Four-Color problem [4]. Their result, the Four-Color Theorem, is renowned for being the first major result to be proved using a computer.

A *k-coloring* of G is a mapping $\varphi : V(G) \to [k]$. It is *proper* if $\varphi(v) \neq \varphi(w)$ whenever $vw \in E(G)$; see Figure 1.4(left). A *color class* of φ is the preimage of some color α in φ; that is, the set of vertices colored with α in φ. A graph G is *k-colorable* if there exists a proper k-coloring φ of G. The *chromatic number* of G, denoted $\chi(G)$, is the smallest integer k such that G is k-colorable. If $\chi(G) = k$, then G is *k-chromatic*. Further, G is *k-critical* if $\chi(G) = k$ but $\chi(G - v) < k$ for every $v \in V(G)$.

k-coloring

color class

$\chi(G)$

k-critical

Remark 1.1. Throughout this dissertation, we assume all colorings are proper unless stated otherwise.

One widely studied generalization of coloring is list-coloring. List-coloring arises in channel allocation for wireless networks or frequency allocation for stations [62, 42]. A *list-assignment* L of G assigns a list of colors $L(v)$ to every $v \in V(G)$; see Figure 1.4(right). An *L-coloring* of G is a proper coloring φ where $\varphi(v) \in L(v)$ for every $v \in V(G)$. For a given function $f : V(G) \to \mathbb{N}$, an *f-assignment* of G is a list-assignment L such that $|L(v)| = f(v)$ for every $v \in V(G)$. A graph G is *f-choosable* if it has an L-coloring for every f-assignment L. For a positive integer k, a *k-assignment* of G is a list-assignment L such that $|L(v)| = k$ for every $v \in V(G)$. A graph G is *k-choosable* if G has an L-coloring for every k-assignment L. If $|L(v)| = d(v)$ for every $v \in V(G)$, then L is a *degree-assignment*. If G has an L-coloring for every degree-assignment L, then G is *degree-choosable*. The *choice number* (also called *list-chromatic number*), denoted $\chi_l(G)$, is the smallest integer k such that G is k-choosable.

list-assignment

L-coloring

k-/degree-assignment

k-/degree-choosable

$\chi_l(G)$

Figure 1.4: Left: A 3-coloring of K_3. Right: A 3-assignment L of K_3 along with an L-coloring (underlined in lists and shown on vertices).

Graph coloring can be applied to several real-world problems. For example, consider the problem of assigning time slots for a given number of college classes with the caveat that certain pairs of classes cannot be assigned the same time slot (perhaps both classes are core requirements for the same group of students).

The classes can be viewed as the vertices of a graph, and two classes are joined by an edge if they cannot be assigned the same time slot. A coloring of this graph represents a valid time slot assignment. In this case, the colors correspond to time slots. Similarly, various other scheduling problems, such as aircraft-to-flight assignments, can be modeled as graph coloring problems [51]. Graph coloring can also be used to model radio frequency assignments [41] and register allocation [23], among many others.

1.4 Kempe Swaps

In this section we introduce the "color swap" operation used to transform one coloring of a graph into another.

Recall that Alfred Kempe provided a flawed "proof" for the Four-Color problem. Though his attempt was unsuccessful, his ideas were later used by Heawood to prove the Five-Color Theorem and played a crucial role in the eventual correct proofs of the Four-Color problem. One of those ideas introduced by their namesake is *Kempe swaps*. Kempe swaps

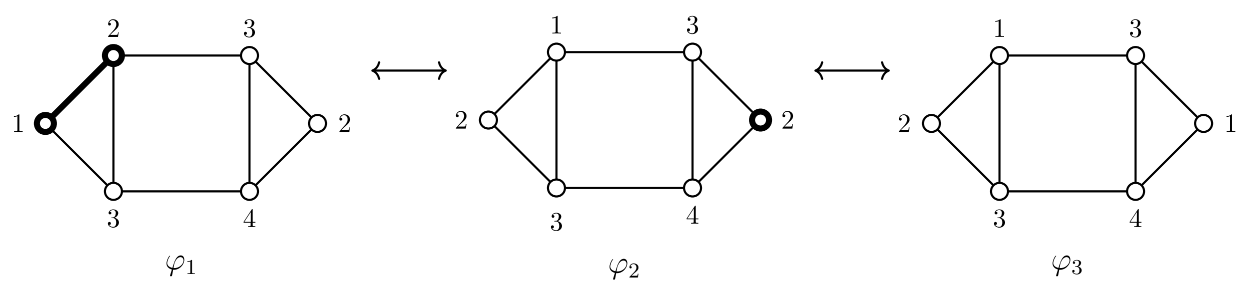

Figure 1.5: A 1,2-Kempe swap is performed on the 1,2-Kempe component (bold) in φ_1 to obtain φ_2. A trivial 1,2-Kempe swap is performed on the trivial 1,2-Kempe component (bold) in φ_2 to obtain φ_3.

Given a coloring φ of G, let G' be the subgraph induced by vertices colored α and β. An *α, β-Kempe swap* exchanges the colors on the vertices of a component of G'. Each component of G' is called an *α, β-Kempe component* (or *Kempe component*); see Figure 1.5. An *α, β-Kempe swap at x* exchanges the colors on the Kempe component of G' containing x (the terminology implicitly implies that the color of x changes from α to β). Performing a Kempe swap on a proper k-coloring φ of G results in a k-coloring φ' that is also proper. Equivalently, the property of being proper for a coloring is preserved under Kempe swaps. For a given coloring φ of G and a vertex $x \in V(G)$ with $\varphi(x) = \alpha$, if $\beta \notin \cup_{y \in N(x)} \varphi(y)$, then we can recolor x with β. Such an α, β-Kempe swap is *trivial* since the α, β-Kempe component containing x is trivial; see Figure 1.5. Analogously, we can also define Kempe swaps for list-coloring. Given a list-assignment L of G,

Kempe component

trivial

an *L-valid Kempe swap* is a Kempe swap that produces an L-coloring; see Figure 1.6. That is, the vertices of the resulting coloring must be colored from their lists. L-valid Kempe swap

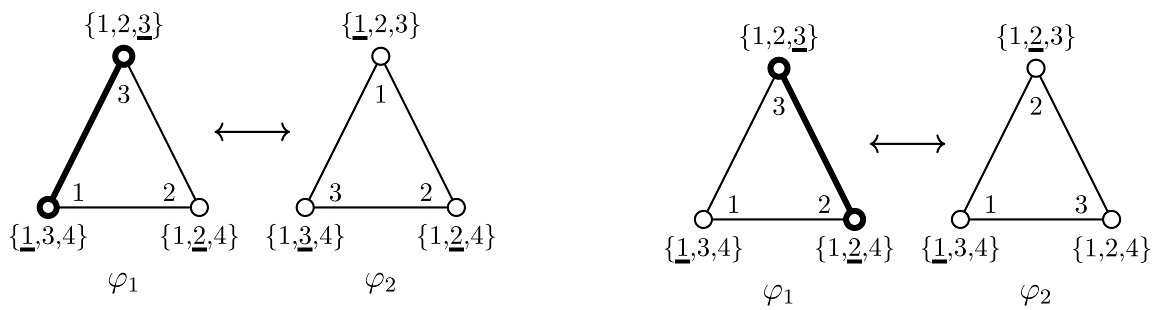

Figure 1.6: Left: An L-valid Kempe swap (bold) from φ_1 to φ_2. Right: Not an L-valid Kempe swap from φ_1 to φ_2 since the vertex with list $\{1, 2, 4\}$ is colored 3 in φ_2, but $3 \notin \{1, 2, 4\}$.

In addition to being used in the proof of the Five-Color Theorem, Kempe swaps are used for extending a coloring of a subgraph of G to a coloring of all of G [2]. They have also been used in problems related to the strong immersion of complete graphs [43], which is the idea of mapping the vertices of the complete graph to those of a graph G such that edges in the complete graph correspond to internally disjoint paths in G. Moreover, Kempe swaps can be defined for edge-colorings[12] and are significant in their study. Kempe components in that context are well-understood because they are always either cycles or paths. In 1964, Vizing famously proved[13] that $\chi'(G) \leq \Delta(G) + 1$ for every simple graph G using Kempe swaps [61]. In fact, he proved more strongly that every k-edge-coloring with $k \geq \Delta(G) + 1$ can be transformed into a $(\Delta(G) + 1)$-edge-coloring via Kempe swaps. A year later he conjectured that every k-edge-coloring with $k \geq \chi'(G)$ can be transformed into a $\chi'(G)$-edge coloring. In 2023, Narboni proved Vizing's conjecture with the help of Kempe swaps [57].

1.5 Computational Complexity

In this section, we briefly discuss computational complexity and introduce some relevant complexity classes.

Given a decision problem D, a fixed input of D defines an *instance* I of D. For example, let D be the problem which asks the following: Given a set of integers S and an integer k, is there a subset S' of S such that $\sum_{s' \in S'} s' = k$? This is called the SUBSET SUM problem. The input of D is a pair (S, k) where S is a set instance

[12]An edge-coloring is an assignment of colors to the edges of a graph such that edges which share an endpoint get different colors. A Kempe swap, therefore, exchanges the colors on the edges of a Kempe component of a subgraph induced by colored edges.

[13]Since the edges of a $\Delta(G)$-vertex must get distinct colors, this implies that every simple graph G satisfies $\chi'(G) = \Delta(G)$ (called *class 1* graphs) or $\chi'(G) = \Delta(G) + 1$ (called *class 2* graphs).

of integers and k is an integer, and $I = (\{-3, 2, 7, 10\}, 9)$ is an instance of D. An instance I is a *YES-instance* (resp. *NO-instance*) if the answer for I is Yes (resp. No). A *certificate* C for an instance I of D is a "proof" for I. For example, the instance I defined above is a YES-instance and $C = \{2, 7\}$ is a certificate for I since $2 + 7 = 9$. A *verifier* is an algorithm that can verify the correctness of certificates. In particular, a verifier for D would be an algorithm V which takes a set of numbers and a value k as input, adds up the numbers, then checks that their sum is equal to k. So, V adds up 2 and 7 and checks that their sum is 9.

YES-/NO-instance

certificate

verifier

The following are the most widely studied and most pertinent complexity classes for this dissertation. The complexity class *P* is the set of all decision problems that can be solved by a deterministic Turing machine in polynomial time. The complexity class *NP* is the set of all decision problems which can be solved by a nondeterministic Turing machine (A "lucky" machine which can always "guess" a correct solution for YES-instances) in polynomial time. Alternatively, NP is the set of all decision problems D such that, for every YES-instance I of D, there exists a certificate C_I and verifier V_D where V_D can verify C_I in polynomial time. In other words, NP is the set of decision problems whose YES-instances can be "checked" in polynomial time. The complexity class *co-NP* is the set of all decision problems D such that, for every NO-instance I of D, there exists a certificate C_I and polynomial-time verifier V_D. Equivalently, a problem belongs to co-NP if and only if its complement belongs to NP. The complexity class *PSPACE* is the set of all decision problems that can be solved by a Turing machine in polynomial space (computer memory). Analogously, we can define complexity classes *NPSPACE* and *co-NPSPACE*.

P

NP

co-NP

PSPACE

Clearly[14], $\mathsf{P} \subseteq (\mathsf{NP} \cap \mathsf{co\text{-}NP})$ and $(\mathsf{NP} \cup \mathsf{co\text{-}NP}) \subseteq \mathsf{PSPACE}$. Furthermore, note that it is not necessary to specify whether the Turing machine for a problem in PSPACE is deterministic or nondeterministic. This is because $\mathsf{PSPACE} = \mathsf{NPSPACE}$, by Savitch's Theorem [60]. In contrast, it is not known whether $\mathsf{P} = \mathsf{NP}$ or whether $\mathsf{NP} = \mathsf{PSPACE}$, though it is widely believed that both statements are false. In fact, the former statement is one of the biggest problems in theoretical computer science and is featured as one of seven problems selected by the Clay Mathematics Institute, called the Millennium Prize Problems, each of which has an attached monetary prize of one million dollars!

Problem A is *reducible* to problem B, or there is a *reduction* from A to B, if every instance of A can be transformed into an instance of B while preserving YES-instances. More precisely, A is reducible to B if there exists an algorithm that can convert every instance I_A of A into an instance I_B of B such that I_A is a YES-instance if and only if I_B is a YES-instance. The notion of reduction brings about the notion of *hardness* in computational complexity, which measures how hard a problem is with respect to other problems. For a

reduction

hardness

[14]Checking all possible computations of a nondeterministic Turing machine can be done in polynomial space.

complexity class X, problem A is X-*hard* if every problem B in X has a reduction to A. In particular, if X is any of the aforementioned complexity classes, then the reduction must be polynomial-time. If A belongs to X and is X-hard, then A is X-*complete*. Observe that if there exists a polynomial-time algorithm which can solve A, and B is polynomial-time reducible to A, then there exists a polynomial-time algorithm which can solve B as well. So, an X-hard problem is viewed as being "at least as hard as" every other problem in X. However, an X-hard problem need not belong to X. For example, the HALTING problem[15] is NP-hard but does not belong to NP. In contrast, SAT[16] is the first proven NP-complete problem.

X-hard

X-complete

[15] Given an encoding of a Turing machine M and an input w, does M halt on w?
[16] Given a Boolean formula, is there a satisfying assignment?

Chapter 2

Coloring Reconfiguration

In this chapter, we discuss two key models for the k-COLORING RECONFIGURATION problem based on different definitions of the transformation rule, and we explore the associated reconfiguration graph for each. We also introduce list-coloring reconfiguration and briefly review the literature for both coloring and list-coloring reconfiguration under both definitions of the transformation rule.

Recall that the k-COLORING RECONFIGURATION problem is derived from the k-COLORING search problem which, given a graph G and a positive integer k, finds a k-coloring of G. The solutions for k-COLORING are the states of k-COLORING RECONFIGURATION, and in Section 1.4 we describe an operation which can transform one k-coloring (state) into another: Kempe swaps. We now introduce two different models for the k-COLORING RECONFIGURATION problem with slight variations in the transformation rule defined using Kempe swaps.

The first model is the *recoloring model*, where the transformation rule is defined to be a trivial Kempe swap. So, any two consecutive colorings in a sequence must differ in color on a single vertex. A move in this case is a *recoloring* since we recolor a single vertex, and a sequence of recolorings is a *recoloring sequence*. Two k-colorings φ and φ_0 are *k-equivalent* if there exists a recoloring sequence S using at most k colors (i.e. $\alpha, \beta \in [k]$ for every trivial α, β-Kempe swap)[1] that transforms one into the other. In other words, S is a sequence $\varphi = \varphi_1, \ldots, \varphi_2, \ldots, \varphi_t = \varphi_0$, where each φ_{i+1} differs from φ_i by a single recoloring.

recoloring model

k-equivalent

The second model is the *Kempe swap model*, where the transformation rule is defined to be a Kempe

Kempe swap model

[1]Note that we bound the number of colors by k since otherwise, we can always construct a recoloring sequence from φ_1 to φ_2 by introducing $|V(G)|$ more colors, recoloring each vertex with a distinct new color in φ_1, then recoloring vertices with their desired color in φ_2.

swap (not necessarily trivial). A move in this case is simply a Kempe swap. Two k-colorings φ_1 and φ_2 are *k-equivalent* if there exists a sequence S of Kempe swaps using at most k colors that transforms one into the other.

For k-colorings φ_1 and φ_2, we write $\varphi_1 \sim_k \varphi_2$ if φ_1 is k-equivalent to φ_2. Further, since our moves are reversible, the relation $\sim_k$ is reflexive ($\varphi_1 \sim_k \varphi_1$), symmetric (if $\varphi_1 \sim_k \varphi_2$, then $\varphi_2 \sim_k \varphi_1$), and transitive (if $\varphi_1 \sim_k \varphi_2$ and $\varphi_2 \sim_k \varphi_3$, then $\varphi_1 \sim_k \varphi_3$). Thus, $\sim_k$ is an equivalence relation, which motivates its name. A graph G is *k-mixing* if its k-colorings are all pairwise k-equivalent. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two sets of k-colorings k-mixing of G. The set $\mathcal{L}_1$ *mixes* if k-colorings in $\mathcal{L}_1$ are pairwise k-equivalent. The set $\mathcal{L}_1$ *mixes with* $\mathcal{L}_2$ if $\varphi_1 \sim_k \varphi_2$ mixes for every $\varphi_1 \in \mathcal{L}_1$ and $\varphi_2 \in \mathcal{L}_2$. The use of the term "mix" in these definitions is motivated by the rapid mixing of Markov chains, a topic which is closely related to coloring reconfiguration and which we discuss in Section 2.4.

Since performing a Kempe swap on a proper coloring of a graph results in another proper coloring, we are guaranteed that each intermediate state in a sequence for k-COLORING RECONFIGURATION is a solution to k-COLORING. However, not every k-coloring of G can be transformed to every other solely using Kempe swaps with at most k colors. In particular, for the recoloring model, it is not always possible to perform a recoloring. A coloring in this model is *frozen* if no vertex can be recolored. Figure 2.1 shows two frozen frozen 3-colorings of $K_2 \square K_3$, or the 3-prism; thus, those colorings are not 3-equivalent. On the other hand, it is always possible to perform a Kempe swap on a coloring. Yet, that still does not guarantee the equivalence of colorings. That is, frozen colorings, though defined differently, also exist for the Kempe swap model. A coloring in this model is *frozen* if for every pair of colors α and β, the subgraph induced by vertices colored α and β consists of a single Kempe component. Both 3-colorings in Figure 2.1 are also frozen for the Kempe swap model; thus, no sequence of Kempe swaps can transform one into the other. This is because the vertex sets of the three color classes (in each coloring) are fixed under Kempe swaps, but differ in these two colorings. Constructing frozen colorings is the primary method we know for showing a graph has nonequivalent colorings.

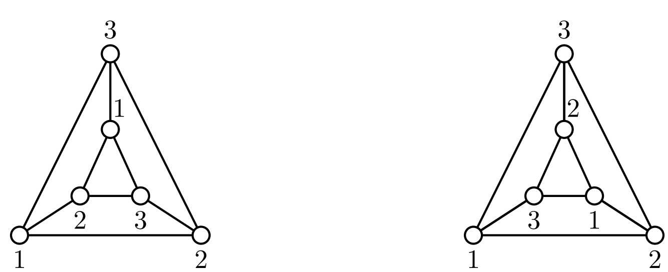

Figure 2.1: Two 3-colorings of the 3-prism that are frozen for both the recoloring and Kempe swap models.

Remark 2.1. By slightly abusing terminology, we use "k-equivalent," "k-mixing," and "frozen" in both models.

The reconfiguration graph associated with the k-COLORING RECONFIGURATION problem is denoted $\mathcal{C}_k(G)$. It has as its vertices the k-colorings of G, and two vertices in $\mathcal{C}_k(G)$ are adjacent if they differ by a single move, where the type of move depends on the model. Let u and v be vertices in $\mathcal{C}_k(G)$ and let φ_u and φ_v be their corresponding k-colorings in G. Under the recoloring model, u is adjacent to v if φ_u can be obtained from φ_v by a single recoloring, i.e., if φ_u differs from φ_v on a single vertex. On the other hand, under the Kempe swap model, u is adjacent to v if φ_u can be obtained from φ_v by a single Kempe swap.

Figure 2.2 shows the reconfiguration graph under both models for all 3-colorings of K_3. Observe that $\mathcal{C}_3(K_3)$ is edgeless in the recoloring model since each coloring is frozen; see Figure 2.2(center). In contrast, $\mathcal{C}_3(K_3)$ is connected in the Kempe swap model, which implies K_3 is 3-mixing; see Figure 2.2(right). By the definition of recoloring, a path in $\mathcal{C}_k(G)$ from u to v in the recoloring model implies a path in the Kempe swap model as well. In fact, let $\mathcal{C}_k^r(G)$ and $\mathcal{C}_k^{Ks}(G)$ denote, respectively, $\mathcal{C}_k(G)$ in the recoloring and Kempe swap models. We have that $\mathcal{C}_k^r(G)$ is a spanning subgraph of $\mathcal{C}_k^{Ks}(G)$. So, if $\mathcal{C}_k^r(G)$ is connected, then so is $\mathcal{C}_k^{Ks}(G)$. The converse, however, is often false, as shown in Figure 2.2. When a statement applies to both models, or the model is understood, we omit the superscripts for the reconfiguration graph.

Analyzing problems of reachability and connectedness in the reconfiguration graph, is equivalent to finding reconfiguration sequences for the k-COLORING RECONFIGURATION problem. That is, $\varphi_u \sim_k \varphi_v$ if there exists a path between u and v in $\mathcal{C}_k(G)$. Furthermore, if $\mathcal{C}_k(G)$ is connected, then every pair of k-colorings of G are k-equivalent. This motivates the following problems with respect to the model.

k-COLOR PATH

Instance: A graph G, a positive integer k, and two k-colorings φ_1 and φ_2 of G.

Question: Is $\varphi_1 \sim_k \varphi_2$, i.e., is there a path between the corresponding vertices of φ_1 and φ_2 in $\mathcal{C}_k(G)$?

k-MIXING

Instance: A graph G and a positive integer k.

Question: Is G k-mixing, i.e., is $\mathcal{C}_k(G)$ connected?

Recall that list-coloring is a generalization of coloring. As a result, we can define the LIST-COLORING RECONFIGURATION problem which, given a graph G and a list-assignment L, asks if it is possible to move between different L-colorings of G. Recall also the operation defined for list-coloring reconfiguration, L-valid

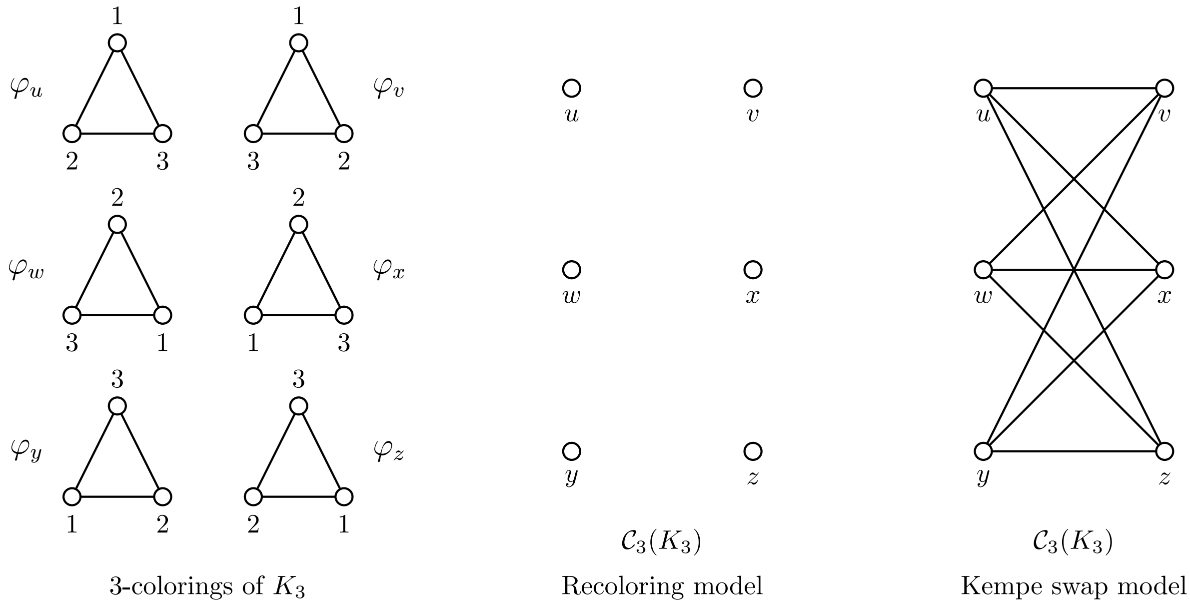

Figure 2.2: Left: All 3-colorings of K_3. Center: The reconfiguration graph $\mathcal{C}_3(K_3)$ under the recoloring model. Right: The reconfiguration graph $\mathcal{C}_3(K_3)$ under the Kempe swap model.

Kempe swaps, which work almost exactly like Kempe swaps except that we must ensure that swaps we perform always recolor vertices with allowable colors. Similar to coloring reconfiguration, we can define an *L-valid recoloring model* and an *L-valid Kempe swap model* for list-coloring reconfiguration. Two L-colorings φ_1 and φ_2 are *L-equivalent*, denoted $\sim_L$, if there exists a sequence of L-valid recolorings (or Kempe swaps, depending on the model) that transforms one into the other. A graph G is *L-swappable* if its L-colorings are pairwise L-equivalent. It is *k-swappable* if it is L-swappable for every k-assignment L. Similarly, it is *degree-swappable* if it is L-swappable for every degree-assignment L. The list-coloring reconfiguration graph for a given list-assignment L (and for each model) is $\mathcal{C}_L(G)$. Finally, the following problems are analogous to those for coloring reconfiguration.

L-equivalent

L-swappable

k-swappable

degree-swappable

LIST-COLOR PATH

Instance: A graph G, a list-assignment L, and two L-colorings φ_1 and φ_2 of G.

Question: Is $\varphi_1 \sim_L \varphi_2$?

L-SWAPPABLE (resp. k-SWAPPABLE)

Instance: A graph G and a list-assignment L (resp. a positive integer k).

Question: Is G L-swappable (resp. k-swappable)?

Next we review relevant results in the literature on coloring and list-coloring reconfiguration for both models. We begin by making a few simple observations which are true for both models.

Observation 2.1. *$V(\mathcal{C}_k(G)) \neq \emptyset$ if and only if $k \geq \chi(G)$. Therefore, we only consider $\mathcal{C}_k(G)$ for $k \geq \chi(G)$. Similarly, we only consider $\mathcal{C}_L(G)$ for list-assignments L such that G admits an L-coloring; in particular, for k-swappability results, we only consider the case $k \geq \chi_l(G)$.*

Observation 2.2. *For every $l < k$, an l-coloring of G is also a k-coloring of G. Therefore, every edge in $\mathcal{C}_k(G)$ is also in $\mathcal{C}_{k+1}(G)$ and $\mathcal{C}_k(G)$ is an induced subgraph of $\mathcal{C}_{k+1}(G)$.*

Observation 2.3. *If $|V(G)| = n$, then $\mathcal{C}_k^r(G)$ is an induced subgraph of $Q_n([k])$.*

Indeed, by labeling the vertices $v_1, \ldots, v_n$, each k-coloring φ corresponding to a vertex in $\mathcal{C}_k^r(G)$ can be viewed as a string of length n with entries in $[k]$, where the entries represent the color of each vertex in φ. By the definition of $Q_n([k])$, each such string corresponds[2] to a vertex in $Q_n([k])$. Since the adjacency relation is defined similarly for both $\mathcal{C}_k(G)$ and $Q_n([k])$, the observation follows. For the remainder of this chapter, let $n := |V(G)|$ for any graph G.

2.1 Recoloring Model

In light of Observation 2.1, a natural first question to ask is whether $\mathcal{C}_k(G)$ is connected for values of k which are sufficiently large compared to $\chi(G)$. Cereceda et al. [20] proved this is false.

Theorem 2.1 ([20])**.** *There exists no function f such that $\mathcal{C}_k(G)$ is connected for every graph G whenever $k \geq f(\chi(G))$.*

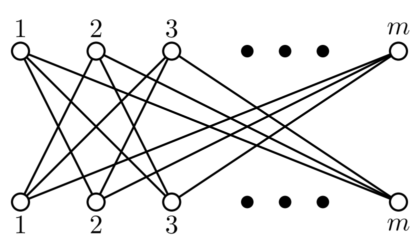

Figure 2.3: A frozen m-coloring of L_m.

In particular, they constructed a class of 2-chromatic graphs L_m with $m \geq 3$ for which $\mathcal{C}_m(L_m)$ is disconnected. For $m \geq 3$, let L_m be the graph formed from $K_{m,m}$ (with parts X and Y) by deleting a perfect

[2]The converse is not true. For example, the string $1, 1, \ldots, 1$ does not correspond to a proper k-coloring of G, i.e., it does not correspond to a vertex in $\mathcal{C}_k(G)$.

matching. More precisely, let $X := \{x_1, x_2, \ldots, x_m\}$ and $Y := \{y_1, y_2, \ldots, y_m\}$, and let $V(L_m) := X \cup Y$ and $E(L_m) := \{x_i y_j | 1 \leq i, j \leq m, i \neq j\}$. Consider the m-coloring φ of L_m where $\varphi(x_i) = \varphi(y_i) = i$; see Figure 2.3. Since every color in $[m]$ appears on $N[v]$ for every $v \in V(L_m)$, the coloring φ is frozen in $\mathcal{C}_m(L_m)$. Thus, L_m is not m-mixing. Interestingly, L_m is k-mixing[3] for all other values of k.

Proposition 2.1 ([20]). *$\mathcal{C}_k(L_m)$ is connected whenever $k \geq 3$ and $k \neq m$.*

We will show that the set of 2-colorings of L_m mixes, and that each k-coloring of L_m is k-equivalent to some 2-coloring.

Proof. Fix k such that $k \geq 3$ and $k \neq m$. Let $\mathcal{L}$ be the set of 2-colorings of L_m in $\mathcal{C}_k(L_m)$. Note that $\mathcal{L}$ mixes. To see this, pick arbitrary $\varphi_1, \varphi_2 \in \mathcal{L}$. For $i \in [2]$, let α_i and β_i be the colors used in φ_i on X and Y, respectively. Since $k \geq 3$, there exists a color $\gamma \notin \{\alpha_1, \beta_1\}$. If $\gamma \in \{\alpha_2, \beta_2\}$, say $\gamma = \alpha_2$, then recolor X with γ and recolor Y with β_2 to obtain φ_2. Otherwise, recolor X with γ, recolor Y with β_2, and recolor X with α_2 to obtain φ_2. So, $\mathcal{L}$ mixes.

Now consider $\varphi \notin \mathcal{L}$. To show that $\mathcal{C}_k(L_m)$ is connected, it suffices to show that $\varphi \sim_k \varphi_0$ for some $\varphi_0 \in \mathcal{L}$. Observe that if $k < m$, then there exists two vertices in X with a common color in φ, by Pigeonhole. Such a color does not appear on Y, by the definition of L_m. Similarly, if $k > m$, then some color does not appear on Y, by Pigeonhole. Call this color α. Recolor X with α in φ then recolor Y with some $\beta \in [k] \setminus \{\alpha\}$. The resulting 2-coloring φ_0 is in $\mathcal{L}$. □

In contrast to Theorem 2.1, if k is sufficiently large compared to $\Delta(G)$, then $\mathcal{C}_k(G)$ is connected. Indeed, Jerrum [46] proved that every graph G is k-mixing whenever $k \geq \Delta(G) + 2$. In fact, a closer look at the proof of this result reveals that we can do better. The proof relies on always being able to recolor a vertex to a color different from the colors on its closed neighborhood. Since $k \geq \Delta(G) + 2$, there is always an extra color to use. Moreover, this is true for every $v \in V(G)$ since $d(v) \leq \Delta(G)$. But, given that the proof is by induction on n, it is enough that G contains one vertex for which this property is true. Fortunately, a class of graphs with such a vertex exists: the class of degenerate graphs. Recall that if G is d-degenerate, then it contains a vertex v with $d(v) \leq d$. Hence, as shown by Dyer et al. [36], replacing $\Delta(G)$ with $\deg(G)$ gives a stronger result. In fact, their result applies, more generally, for list-coloring.

Recall that G is L-swappable if its L-colorings are pairwise L-equivalent.

[3] Note that L_m is clearly not 2-mixing since $\mathcal{C}_2(L_m)$ consists of two isolated vertices; namely, the frozen 2-colorings φ_1 and φ_2 defined by $\varphi_i(X) = i$ and $\varphi_i(Y) = 3 - i$ for each $i \in [2]$.

Theorem 2.2. *For every graph G and list-assignment L, if $|L(v)| \geq \deg(G) + 2$, then G is L-swappable.*

Proof. Recall that $n := |V(G)|$. We use induction on n. If $n = 1$, the result is trivially true. So, assume $n \geq 2$. Since G is $\deg(G)$-degenerate, there exists $v \in V(G)$ with $d(v) \leq \deg(G)$. Let $G' := G - v$. Note that $\deg(G') \leq \deg(G)$, so $|L(w)| \geq \deg(G') + 2$ for every $w \in V(G')$. By induction, G' is L-swappable. Let φ and φ_0 be two L-colorings of G, and let φ' and φ'_0 be their restrictions to G', respectively. It suffices to show that there exists an L-valid recoloring sequence from φ to φ_0.

Since G' is L-swappable, there exists an L-valid recoloring sequence $\varphi' = \varphi'_1, \ldots, \varphi'_t = \varphi'_0$ for some $t \in \mathbb{N}$. We will extend each φ'_i of G' to an L-coloring φ_i of G to obtain an L-valid recoloring sequence from φ to φ_0. Suppose φ'_{i+1} differs from φ'_i by an L-valid recoloring of a vertex u from α to β. If $u \notin N(v)$ or $\varphi_i(v) \neq \beta$, then we perform the same recoloring of u on φ_i to obtain φ_{i+1}. Otherwise, we recolor v with a color $c \in L(v) \setminus \cup_{w \in N[v]} \varphi'_i(w)$. Such a color exists since $|L(v)| \geq \deg(G) + 2$ and $| \cup_{w \in N[v]} \varphi'_i(w)| \leq d(v) + 1 \leq \deg(G) + 1$. Next we perform the same recoloring of u to obtain φ_{i+1}. We repeat this process for each $i \in [t-1]$ to obtain a recoloring sequence $\varphi = \varphi_1, \ldots, \varphi_t$. Finally, if $\varphi_t \neq \varphi_0$, we recolor v in φ_t with its color in φ_0. □

Clearly, Theorem 2.2 subsumes Jerrum's result since $\deg(G) \leq \Delta(G)$. Moreover, this bound is best possible since K_m and L_m are both $(m-1)$-degenerate but not m-mixing. In fact, this bound is sharp even for restricted classes of graphs. Indeed, a class of degenerate graphs often studied for its well-understood structural properties is the class of planar graphs. Thus, it is common to consider the implications of degeneracy results on this class (along with its subclasses). Recall that planar graphs are 5-degenerate, triangle-free planar graphs are 3-degenerate, and planar graphs with girth at least 6 are 2-degenerate. As a result, Theorem 2.2 implies G is k-mixing whenever G is planar and $k \geq 7$, whenever G is triangle-free planar and $k \geq 5$, and whenever G is planar with girth at least 6 and $k \geq 4$. However, Figure 2.4(left) shows a frozen 6-coloring of a planar graph, Figure 2.4(center) shows a frozen 4-coloring of a triangle-free planar graph[4] and Figure 2.4(right) shows a frozen 3-coloring[5] of a planar graph with girth at least 6.

It is interesting to note the techniques used to show that a graph is k-mixing (or L-swappable). Theorem 2.2 uses induction and extension. We delete some vertices of G, get a reconfiguration sequence by induction for the smaller graph G', then extend that sequence to a sequence in G. Typically, the deleted subset of vertices is chosen to possess nice properties that guarantee the extension (for example, a vertex

[4]In contrast, Bartier et al. [7] showed that planar graphs with girth 5 are 4-mixing.
[5]This can be generalized to any cycle with length divisible by 3 and with a (1,2,3)-alternating coloring.

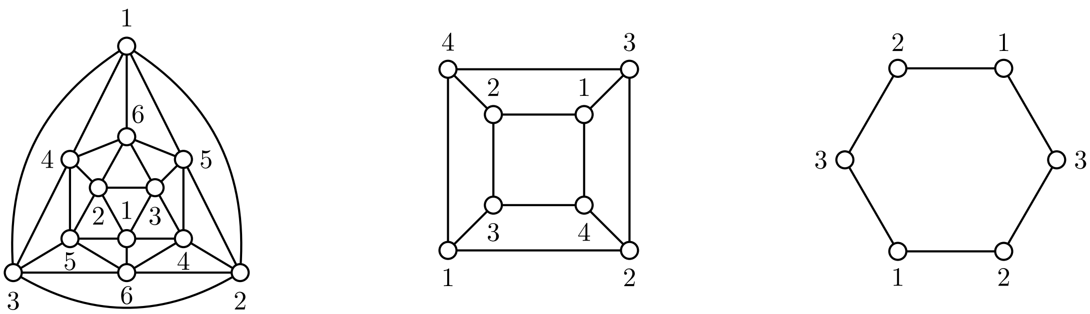

Figure 2.4: Left: A frozen 6-coloring of an icosahedron. Center: A frozen 4-coloring of $Q_3(\{0,1\})$. Right: A frozen 3-coloring of C_6. (In each case, each color used appears in the closed neighborhood of each vertex.)

with "few" neighbors, as in Theorem 2.2). Proposition 2.1 showcases another technique. The goal in this case is to show that a subset of colorings $\mathcal{L}$ mixes, then show that every coloring not belonging to $\mathcal{L}$ mixes with some coloring in $\mathcal{L}$. Another key method uses *reducible configurations*, which are structures that cannot occur in a minimial counterexample. More precisely, we fix a minimal[6] counterexample G to the desired result, then show that some structure H cannot be a subgraph of G; otherwise, an induction (using the minimality of G) and extension (deleting H) approach is possible. Ultimately, to prove the result, we must (1) compile a collection of reducible configurations – as well as, show that they are reducible; that is, their presence in G implies the result – and (2) show that every minimal counterexample contains one. Achieving (1) is usually done through a careful analysis of the structure of G and its colorings. Whereas, achieving (2) is often done via *discharging*, a weight distribution method mainly used on planar graphs. The idea with discharging is to assign weights to elements of G (for example, its vertices), then distribute the weights amongst the elements according to some prescribed rules which designed assuming that G contains no copy of H. Eventually, the aim is to show that the sum of the weights before and after the distribution is different, a contradiction which implies that G does indeed contain some reducible configuration. We use a type of reducible-configurations argument for all our results in Chapters 3–5, and we use discharging in Chapter 3.

As previously mentioned, apart from the connectedness of the reconfiguration graph, its diameter is another area of interest. For $k \geq 4$, Bonsma and Cereceda [13] showed there exists a class[7] of graphs G_k with k-equivalent k-colorings that require a recoloring sequence of length superpolynomial in n. In contrast, Cereceda et al. [21] showed that if G is 3-colorable, then every component[8] of $\mathcal{C}_3(G)$ has diameter $O(n^2)$. Interestingly, there are no known graphs for which the reconfiguration graph is connected and has

[6] "Minimal" here could refer to smallest number of vertices/edges, or smallest value of some function/parameter with respect to the graph.

[7] Such a class of graphs exists even when restricted to bipartite, planar, and planar bipartite graphs [13].

[8] They also exhibit 3-colorings for which a sequence of length $\Omega(n^2)$ is needed; thus, the bound is tight.

superpolynomial diameter. In fact, Cereceda [19] conjectured that, not only is the reconfiguration graph connected for degenerate graphs, but its diameter is actually quadratic.

Cereceda's Conjecture ([19])**.** *For every graph G, if $k \geq \deg(G)+2$, then diam$(\mathcal{C}_k(G))$ is $O(n^2)$.*

Observe that the algorithmic proof of Theorem 2.2 does imply a bound on the length of reconfiguration sequences; unfortunately, that bound is exponential in n. So, the conjecture suggests a major improvement. Furthermore, if the conjecture is true, then the bound is sharp. Indeed, recall that paths, and more generally forests, are 1-degenerate. Bonamy et al. [12] showed that $\text{diam}(\mathcal{C}_3(P_n))$ is $\Omega(n^2)$; in fact, $\Theta(n^2)$. More generally, they showed the bound is tight for chordal graphs. Despite multiple attempts at proving the conjecture, it remains open[9] even for $\deg(G)=2$. Nonetheless, some partial results are known. Cereceda [19] showed the conjecture[10] is true for $k \geq \Delta(G)+2$ and for $k \geq 2\deg(G)+1$. The former case confirms the conjecture for regular graphs G since $\deg(G)=\Delta(G)$. Perhaps the most significant partial result of Cereceda's Conjecture was by Bousquet and Heinrich [15] who showed the following.

Theorem 2.3. *For every graph G, if $k \geq \deg(G)+2$, then diam$(\mathcal{C}_k(G))$ is $O(n^{\deg(G)+1})$. Furthermore, if $k \geq \frac{3}{2}(\deg(G)+1)$, then diam$(\mathcal{C}_k(G))$ is $O(n^2)$.*

For planar graphs, Theorem 2.3 implies diameter $O(n^6)$ whenever $k \geq 7$ and diameter $O(n^2)$ whenever $k \geq 9$. Similarly, for triangle-free planar graphs[11] it implies diameters $O(n^4)$ and $O(n^2)$, respectively, whenever $k \geq 5$ and $k \geq 6$. Lastly, for planar graphs with girth at least 6 it implies diameters $O(n^3)$ and $O(n^2)$, respectively, whenever $k \geq 4$ and $k \geq 5$. Furthermore, their second result improves Cereceda's [19] bound of $2\deg(G)+1$ on k for a quadratic diameter. Surprisingly, Bousquet and Perarnau [16] proved a linear diameter is possible with one more color. What's more, their result can be generalized to list-coloring.

Theorem 2.4. *For every graph G and list-assignment L, if $|L(v)| \geq 2\deg(G)+2$ for every $v \in V(G)$, then diam$(\mathcal{C}_L(G))$ is $O(n)$.*

We include the proof in the spirit of introducing more of the techniques used for solving coloring reconfiguration problems. In particular, one of the tools used for solving problems relating to the diameter of the reconfiguration graph is the following lemma from which the proof of Theorem 2.4 follows.

Lemma 2.1. *Let G be a graph, L be a list-assignment of G, and φ and φ_0 be two L-colorings of G. Fix $v \in V(G)$ and let $G' := G - v$. Further, let φ' and φ'_0 be the restrictions of φ and φ_0 to G', respectively. If*

[9] For $\deg(G)=1$, i.e. for forests, the conjecture follows from [12].

[10] For $k \geq \Delta(G)+2$, the bound on the diameter of the reconfiguration graph was imrpoved to $O(n)$, by Cambie et al. [17].

[11] The authors also proved $\text{diam}(C_5(G))$ is $O(n^2)$ for bipartite planar graphs G; thus, improving the bound $k \geq 6$ for this subclass of triangle-free graphs.

there exists an L-valid recoloring sequence from φ' to φ'_0 which recolors every $w \in N(v)$ at most c times, then it can be extended to an L-valid recoloring sequence from φ to φ_0 which recolors v at most $\lceil \frac{c \cdot d(v)}{|L(v)|-d(v)-1} \rceil + 1$ times.

Proof. As in the proof of Theorem 2.2, the goal is to perform the same recoloring which transforms φ'_i to φ'_{i+1} on φ_i to obtain φ_{i+1} if possible. This is only a problem if φ'_{i+1} differ from φ'_i by a recoloring of vertex u such that $u \in N(v)$ and $\varphi'_{i+1}(u) = \varphi_i(v)$. In this case, we must first recolor v before recoloring u to obtain φ_{i+1}. Note that v must avoid the colors on $N[v]$; that is, it must avoid at most $|N[v]| = d(v) + 1$ colors. So, v can be recolored with at least $|L(v)| - d(v) - 1$ colors. Let T be the set of colors available for v. Since $|T| \geq |L(v)| - d(v) - 1$ and $\varphi_i(v) \notin T$, there exists a color $\alpha \in T$ which is unused by the next $d(v) + 1$ recolorings of $N(v)$. We recolor v with α, and therefore, avoid recoloring v for the next $d(v) + 1$ recolorings of $N(v)$. This gives an L-valid recoloring sequence $\varphi = \varphi_1, \dots, \varphi_t$ of G. Finally, we may need to recolor v in φ_t to its color in φ_0. This implies, in total, v is recolored $\lceil \frac{c \cdot d(v)}{|L(v)|-d(v)-1} \rceil + 1$ many times. □

The idea in Lemma 2.1 is to use extension. But, instead of recoloring v arbitrarily from its list of available colors, we strategically choose a color from the list which allows us to skip a proportion of the next recolorings of $N(v)$. Applying Lemma 2.1, it is now easy to prove Theorem 2.4.

Proof of Theorem 2.4. Note that the connectedness of $\mathcal{C}_L(G)$ follows from Theorem 2.2. Let φ and φ_0 be two L-colorings of G, and let $c := \deg(G) + 1$. We will show that every $v \in V(G)$ is recolored at most c times; thus, the number of total L-valid recolorings is at most cn and $\operatorname{diam}(\mathcal{C}_L(G))$ is $O(n)$.

We use induction on n. If $n = 1$, the result is trivially true. So, assume $n \geq 2$. Since G is $\deg(G)$-degenerate, there exists $v \in V(G)$ with $d(v) \leq \deg(G)$. Let $G' := G - v$, and let φ' and φ'_0 be, respectively, the restrictions of φ and φ_0 to G'. Observe that $\deg(G') \leq \deg(G)$, so $|L(w)| \geq 2\deg(G') + 2$ for every $w \in V(G')$. By induction, there is an L-valid recoloring sequence from φ' to φ'_0 which recolors every $w \in V(G')$ at most c times. By Lemma 2.1, this can be extended to an L-valid recoloring sequence from φ to φ_0 which recolors v at most $\lceil \frac{c \cdot d(v)}{|L(v)|-d(v)-1} \rceil + 1 \leq \lceil \frac{(\deg(G)+1)\deg(G)}{2\deg(G)+2-\deg(G)-1} \rceil + 1 = \deg(G) + 1 = c$ times. □

Applying Theorem 2.4 to the aforementioned classes of planar graphs yields the following: $\operatorname{diam}(\mathcal{C}_k(G))$ is $O(n)$ for: planar graphs G whenever $k \geq 12$, triangle-free planar graphs G whenever $k \geq 8$, and planar graphs G with girth at least 6 whenever $k \geq 6$. Furthermore, the following stronger results have been proved. For planar graphs and triangle-free planar graphs, Dvořák and Feghali [32, 33] improved $k \geq 12$ and $k \geq 8$, respectively, to $k \geq 10$ and $k \geq 7$. For planar graphs with girth at least 6, Bartier et al. [7] improved $k \geq 6$ to $k \geq 5$.

It is natural to wonder whether any of these results on planar graphs can be generalized to list-coloring (of course, the connectedness of $\mathcal{C}_L(G)$ for any list-coloring generalizations follows directly from Theorem 2.2). Bartier et al. [7] generalized the result of Dvořák and Feghali on planar graphs to list-coloring, but with one extra color. More precisely, they showed that for every planar graph G and list-assignment L, if $|L(v)| \geq 11$ for every $v \in V(G)$, then $\mathrm{diam}(\mathcal{C}_L(G))$ is $O(n)$. Following this, Cranston [27] confirmed the list-coloring version of their result for triangle-free planar graphs. That is, he showed that for every triangle-free planar graph and list-assignment L, if $|L(v)| \geq 7$ for every $v \in V(G)$, then $\mathrm{diam}(\mathcal{C}_L(G))$ is $O(n)$.

We end this section with the following questions and open problems. Cereceda's Conjecture is certainly the most intriguing problem to work on given that it is still open even for small values of $\deg(G)$. However, another interesting question to ponder is the following.

Question 2.1. If a $c \cdot n^2$ bound is true for Cereceda's Conjecture, does c depend on $\deg(G)$?

For the most part[12], Theorem 2.3 seems to be the current best bound we have on the diameter of the reconfiguration graph for Cereceda's Conjecture. In particular, recall that if $k \geq \frac{3}{2}(\deg(G)+1)$, then $\mathrm{diam}(\mathcal{C}_k(G))$ is $O(n^2)$. But more than that, Theorem 2.4 gives a linear diameter for $k \geq 2\deg(G)+2$. So, a natural follow up question is:

Question 2.2. What can we conclude about $\mathrm{diam}(\mathcal{C}_k(G))$ for $\frac{3}{2}(\deg(G)+1) \leq k \leq 2\deg(G)+2$? Can we shrink the gap from either end, i.e., can we show that $\mathrm{diam}(\mathcal{C}_k(G))$ is $O(n^2)$ or $O(n)$ for values of k, respectively, larger than $\frac{3}{2}(\deg(G)+1)$ or smaller than $2\deg(G)+2$?

Bartier et al. [7] conjectured the bound on k in Theorem 2.4 can be brought down to $\deg(G)+3$. This is true for $\deg(G)=1$ since $2\deg(G)+2 = \deg(G)+3$, and true for outerplanar graphs [6]. Moreover, the improved bounds of [32, 33, 7] on k for Theorem 2.4 seem to suggest the answer is Yes for all planar graphs.

Finally, we briefly remark on lower bounds for the diameter of the reconfiguration graph. It is easy to see that a linear lower bound holds trivially for every graph G. Indeed, consider two colorings of G which differ in color on every vertex (a coloring and some permutation of it). Clearly, any recoloring sequence must recolor every vertex at least once. In terms of lower bounds on $\mathrm{diam}(\mathcal{C}_k(G))$ for $k \geq \deg(G)+2$, apart from $\Omega(n^2)$ for $k = \deg(G)+2$ given by [12], no nontrivial lower bound is known for $k \geq \deg(G)+3$, which poses

[12] Interestingly, the coefficient for chordal graphs in [12] does not depend on $\deg(G)$.

another intriguing open problem.

2.2 Kempe Swap Model

We now shift our focus to the Kempe swap model. In this section, recall that terms such as "k-mixing" are to be interpreted based on their definition for this model.

Once again, we see that $\chi(G)$ is not the best parameter to consider when studying the connectedness of the reconfiguration graph. Similar to L_m in the recoloring model, for every $s \geq 3$ and $t > s$, the graph $G := K_s \times K_t$ has $\chi(G) = s$ and is s-mixing but not t-mixing. To see this, let $V(K_s) := \{v_1, \ldots, v_s\}$ and $V(K_t) := \{w_1, \ldots, w_t\}$, and let $\varphi_s(v_i, w_j) := i$ and $\varphi_t((v_i, w_j)) := j$; see Figure 2.5. Clearly, a proper coloring of G requires at least s colors[13], so $\chi(G) = s$. Moreover, observe that every subgraph induced by two color classes in φ_s or φ_t is a complete bipartite graph minus a perfect matching, which is connected since $s, t \geq 3$. So, every α, β-Kempe swap in φ_s or φ_t simply permutes color classes α and β; that is, it permutes the colors of two rows in φ_s or two columns in φ_t. Since φ_s is the only s-coloring of G up to permuting color classes, G is s-mixing. However, φ_t is not[14] the only t-coloring of G and φ_t is frozen, so G is not t-mixing.

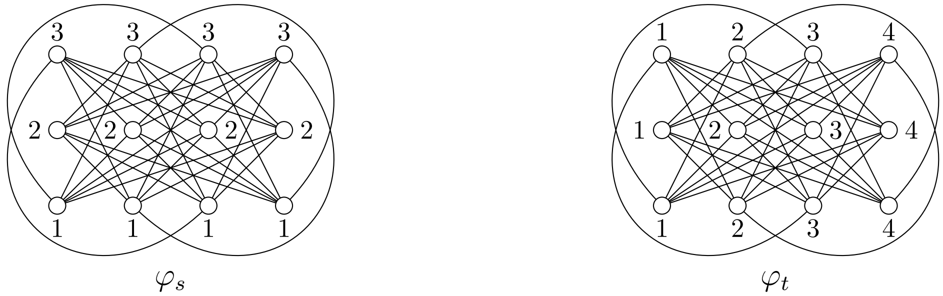

Figure 2.5: Left: The only 3-coloring of $K_3 \times K_4$ up to permuting color classes. Right: A frozen 4-coloring of $K_3 \times K_4$.

Since degenerate graphs seem to show great promise for the recoloring model, it is only reasonable to explore degenerate graphs for the Kempe swap model as well. The most noteworthy result states that, compared with the recoloring model, the Kempe swap model achieves connectedness of the reconfiguration graph for degenerate graphs with one less color. Indeed, Las Vergnas and Meyniel [50] proved that $\deg(G)+1$ colors suffice. Similar to Theorem 2.2, their result can be generalized to list-coloring.

[13] The largest set of vertices in a diagonal forms a clique of size s. More precisely, $\omega(G) = \min\{s, t\} = s$.
[14] Recolor some vertices of φ_s with colors $s + 1$ through t.

Theorem 2.5. *For every graph G and list-assignment L, if $|L(v)| \geq \deg(G)+1$, then G is L-swappable.*

Since the proof closely resembles that of Theorem 2.2, we defer it to Chapter 5 (see Lemma 5.1 and Corollary 5.1) where it is more useful. We remark, however, that the subtlety lies in the local effect each operation has on the graph. Recoloring a vertex only affects that vertex. Thus, in the recoloring model, a vertex must avoid the colors of its neighbors as well as its own color. In contrast, performing a Kempe swap at a vertex typically affects that vertex and some of its neighbors (and possibly some other vertices). So, in the Kempe swap model, a vertex ends up swapping colors with one of its neighbors, thereby only having to avoid one color less than in the recoloring model.

An immediate consequence of Theorem 2.5 is that every planar graph is k-mixing whenever $k \geq 6$. Meyniel [53] extended this by proving that planar graphs are 5-mixing. This is best possible since Mohar [54] proved there exist planar graphs with arbitrarily many 4-colorings which are not 4-equivalent[15]; see Figure 2.6. In contrast, he showed that all 3-colorable planar graphs are 4-mixing (also best possible due to Figure 2.1). Finally, Feghali [38] showed that also every 4-critical planar graph is 4-mixing.

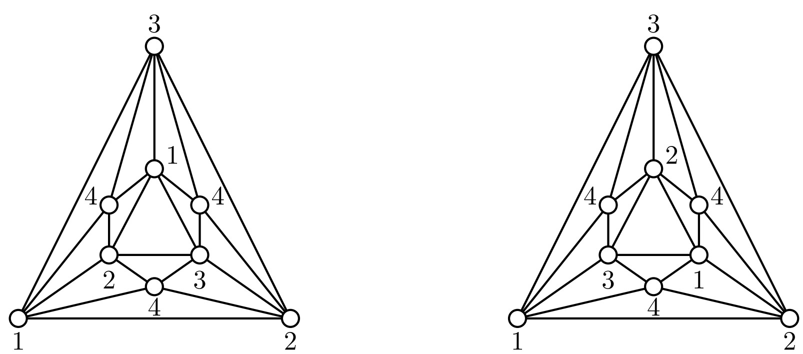

Figure 2.6: Two nonequivalent 4-colorings of a planar graph. Indeed, since the vertex sets of the four color classes (in each 4-coloring) are fixed under Kempe swaps but differ in these two 4-colorings, both 4-colorings are frozen.

A natural continuation of Theorem 2.5 is to analyze the diameter of $\mathcal{C}_k(G)$ for $k \geq \deg(G)+1$. Motivated by Theorem 2.5, Bonamy et al. [11] conjectured something similar to Cereceda's Conjecture.

Conjecture 2.1. *For every graph G, if $k \geq \deg(G)+1$, then diam$(\mathcal{C}_k(G))$ is $O(n^2)$.*

Recall that results for the recoloring model imply the same for the Kempe swap model since $\mathcal{C}_k^r(G)$ is a spanning subgraph of $\mathcal{C}_k^{Ks}(G)$. So, Theorem 2.3 for the recoloring model implies polynomial diameter for

[15] Figure 2.6 is constructed from Figure 2.1 by adding a vertex to every 4-face. This can be used to construct arbitrarily many 4-colorings which are not 4-equivalent by repeatedly adding a copy of Figure 2.6 inside its inner most triangle.

$\mathcal{C}_k(G)$ whenever $k \geq \deg(G)+2$. On the other hand, the case $k = \deg(G)+1$ seems to be much harder to tackle. Nevertheless, for planar graphs and regular graphs, polynomial and quadratic diameters are proved, respectively, even when $k = \deg(G)$. For planar graphs G, Deschamps et al. [30] showed $\text{diam}(\mathcal{C}_5(G))$ is $O(n^c)$ for some positive constant c and remarked that the proof can be adapted for a larger number of colors. On the other hand, recall that $\deg(G) = \Delta(G)$ for regular graphs G. Bonamy et al. [8] showed[16] that if $k \geq \Delta(G)$ and $G \ncong K_2 \square K_3$, then $\text{diam}(C_k(G))$ is $O(n^2)$.

We end this section with the following remark. Note that not much is known for the L-valid Kempe swap model (as far as we know, Theorem 2.5 had not been previously proven, though it may have been intuitively assumed to easily follow from the coloring version proven by Las Vergnas and Meyniel [50]), which suggests an interesting direction of research. Our results in Chapter 5 serve as a first step toward further exploring list-coloring for this model.

2.3 Complexity

Recall from the introduction that a popular line of research is to determine how hard it is to solve coloring reconfiguration problems. In this section, we lay out the most important complexity results for coloring reconfiguration.

As previously mentioned, the use of Kempe swaps for the k-COLORING RECONFIGURATION problem guarantees that colorings are always proper. Moreover, we can check in polynomial time if a k-coloring differs from another by a single Kempe swap (trivial or otherwise). This implies that there exists a verifier which can check whether a sequence (certificate) for k-COLOR PATH is a solution using polynomial space. So, k-COLOR PATH, and therefore k-MIXING[17], belong to NPSPACE. Thus, they also belong to PSPACE, by Savitch's Theorem. As a result, we are often interested in showing that these problems belong to more restrictive classes like P, NP, and co-NP, or showing that they are PSPACE-complete.

For both the recoloring and Kempe swap models, observe that 2-COLOR PATH and 2-MIXING are trivially in P.

For the recoloring model, Cereceda et al. [21] showed that 3-COLOR PATH is in P. They also showed [22] 3-MIXING is co-NP-complete, but is in P when restricted to bipartite planar graphs. For $k \geq 4$, Bonsma and Cereceda [13] proved that k-COLOR PATH is PSPACE-complete. Moreover, the problem remains PSPACE-

[16] The connectedness of $\mathcal{C}_k(G)$ for $k = \Delta(G)$ is a consequence of [40] and [11].
[17] In the case of k-MIXING, we simply run the algorithm for k-COLOR PATH on all pairs of k-colorings.

complete even for bipartite graphs, planar graphs whenever $4 \leq k \leq 6$, and bipartite planar graphs when $k = 4$. For $k \geq 4$, Bousquet [14] proved k-MIXING is co-NP-hard, though he conjectured it is PSPACE-complete.

For the Kempe swap model, Bonamy et al. [10] proved that k-COLOR PATH for $k \geq 3$ is PSPACE-complete. Furthermore, it remains PSPACE-complete even when restricted to $k = 3$ and planar graphs G with $\Delta(G) = 6$. They also showed that both k-COLOR PATH and k-MIXING are in P for chordal graphs, bipartite graphs, and graphs with no induced P_4, or *cographs*. In contrast, no complexity results are known for k-MIXING.

Our Results. We end this section by outlining our main results and contributions in the following chapters towards the relevant decision problems defined for coloring reconfiguration and its list-coloring variant. In Chapter 3, we prove that $\text{diam}(\mathcal{C}_5(G))$ in the recoloring model is $O(n^2)$ for planar graphs G with no 3-cycles and no 5-cycles. In Chapter 4, we prove that 5-MIXING in the Kempe swap model is always a Yes-instance for most 6-regular toroidal graphs. Finally, in Chapter 5, we prove that k-SWAPPABLE in the L-valid Kempe swap model is always a Yes-instance for k-regular graphs with $k \geq 3$.

2.4 Applications and Motivation

In this section we provide motivation for our work and discuss related applications of the k-COLORING RECONFIGURATION problem.

2.4.1 Enumeration

A common goal in reconfiguration is to study the connectedness of the reconfiguration graph. One reason for this emphasis is to find paths between solutions of an underlying problem, and possibly devise algorithms for constructing these paths.

Another key motivation is to approximately enumerate combinatorial structures such as the k-colorings of a given graph G. Exact counting of those k-colorings is unlikely to be possible in polynomial time[18]. However, we can approximately count them by employing an algorithm which generates uniformly random proper k-colorings of G. This is achieved by defining a Markov chain[19] simulation on the set of k-colorings of

[18]This problem is known to be #P-complete, which is at least as hard as NP-complete, but widely believed to be harder.
[19]For information on Markov chains, we refer the reader to [47].

G. The Markov chain represents a random walk on $\mathcal{C}_k(G)$ given by a sequence of random variables $\{X_t\}_{t=0}^{\infty}$. Every variable X_t represents the state of the Markov chain after t steps, and every edge in $\mathcal{C}_k(G)$ is assigned a 'weight' which represents the probability of traversing that edge. Starting with an arbitrary k-coloring of G (a vertex of $\mathcal{C}_k(G)$), we uniformly transition to another k-coloring, one step at a time, by randomly choosing an edge of $\mathcal{C}_k(G)$ using the weights as the probability distribution.

In order to potentially achieve a uniform distribution on all k-colorings, the Markov chain must be *ergodic*, i.e., it must have a unique stationary distribution, and it must converge to this distribution independently of the chosen starting position. This is usually achieved if the reconfiguration graph is connected, which motivates our study of the connectedness of $\mathcal{C}_k(G)$. Moreover, for the method to be efficient, the Markov chain must be *rapidly mixing*, i.e., the time it takes for the distribution to become close enough to the stationary distribution is polynomial in the size of the problem instance. Equivalently, the *mixing time* for the Markov chain must be polynomial. Even though upper bounds on $\text{diam}(\mathcal{C}_k(G))$ do not necessarily tell us much about the mixing time of the Markov chains, lower bounds on $\text{diam}(\mathcal{C}_k(G))$ imply lower bounds on the mixing time, which motivates our study of the diameter of the reconfiguration graph.

2.4.2 Ising and Potts Models

The *Ising* and *k-state Potts* models study the interaction of "spins" of the atoms of a crystalline lattice. They are among the most widely studied in statistical mechanics, additionally offering insight into many areas of solid-state physics. The q-state Potts model (see [64] for further details) has applications in signal and image processing [48], as well as condensed matter systems [49]. The models assign to each vertex v of an underlying graph G a spin $\sigma(v) \in [k]$. So, each spin-assignment of G corresponds to a k-coloring of G, not necessarily proper. In the *ferromagnetic model*, the energy of a spin is low (meaning the spin is more likely) when most edges have endpoints with the same color. In the *antiferromagnetic model*, which is of greater interest for us, the energy of a spin is low when most edges have endpoints with distinct colors. In particular, when the temperature is 0, every edge has endpoints with distinct colors, so the possible spin-assignments of G are precisely its proper k-colorings.

A popular way to simulate the evolution of a graph's spin-assignments over time is using the Glauber dynamics or the Wang–Swendsen–Kotecký (WSK) non-local cluster dynamics. These are Markov chain algorithms for the Ising and Potts models. In order for each algorithm to function properly, we need that each spin of the graph has positive probability of eventually transforming to each other spin; that is, the

Markov chain must be ergodic. In the Glauber dynamics algorithm (resp. WSK algorithm), two states in the Markov chain will be adjacent exactly when they differ by a recoloring (resp. Kempe swap). Thus, in the language of graph coloring reconfiguration, ergodicity in the Glauber dynamics and WSK algorithms, respectively, requires that every two proper k-colorings of G are k-equivalent in the recoloring and Kempe swap models.

We remark that the results of Chapter 4 are of particular significance. Given that $\mathcal{C}_k(G)$ is not guaranteed to be connected for general graphs G, these models have been studied for specific graphs. In particular, highly structural and symmetrical graphs like triangular lattices and Kagomé lattices with boundary conditions, both of which can be formed from toroidal grids (see Chapter 4). The stability of the WSK chain (equivalently, the connectedness of $C_k(G)$) for the Potts model for a Kagomé lattice G has been determined for every value of k [55, 56, 11]. The same authors have determined the stability for the triangular lattice for every value of k except $k = 5$. We solve the case $k = 5$ in Chapter 4.

Chapter 3

The Recoloring Model: 5-Coloring Reconfiguration of Planar Graphs with No Short Odd Cycles

Let $n := |V(G)|$ for any graph G. Recall that Cereceda, van den Heuvel, and Johnson [20, 22] were the first to study the connectedness of $\mathcal{C}_k(G)$. Subsequently, many results [7, 15, 32, 39], as shown in Table 3.1, have focused on the diameter of $\mathcal{C}_k(G)$; in particular, these papers attempt to tackle Cereceda's Conjecture [19] that $\text{diam}(\mathcal{C}_k(G)) = O(n^2)$ whenever G is d-degenerate and $k \geq d + 2$. (The bound $O(n^2)$ is best possible, as shown by Bonamy et al. [12].) Most results on the connectedness of $\mathcal{C}_k(G)$ give an exponential upper bound on its diameter. Bartier et al. [7] proved that $\text{diam}(\mathcal{C}_5(G)) = O(n)$ for every planar graph G of girth at least 6, while Dvořák and Feghali [32] proved $\text{diam}(\mathcal{C}_7(G)) = O(n)$ for every triangle-free planar graph G. Feghali [39] showed that if $d \geq 1$ and $k \geq d + 1$, then for every $\epsilon > 0$ and every graph G with maximum average degree at most $d - \epsilon$, we have $\text{diam}(\mathcal{C}_k(G)) = O(n(\log n)^{d-1})$.

In a recent breakthrough, Bousquet and Heinrich [15] proved that $\text{diam}(\mathcal{C}_k(G)) = O(n^{d+1})$ for every d-degenerate graph G with $k \geq d+2$. Since planar graphs are 5-degenerate, and triangle-free planar graphs are 3-degenerate, their result implies that $\text{diam}(\mathcal{C}_7(G)) = O(n^6)$ for every planar graph G and $\text{diam}(\mathcal{C}_5(G)) = O(n^4)$ for every triangle-free planar graph G. In the same paper, they proved that $\text{diam}(\mathcal{C}_5(G)) = O(n^2)$ for every bipartite planar graph G. They also remarked that Cereceda's Conjecture remains open for triangle-

Graph type / colors	4	5	6	7
Girth 3	∞	∞	∞	$O(n^6)$ [15]
Girth 4	∞	$O(n^4)$ [15]	$O(n \log^3(n))$ [39]	$O(n)$ [32]
Girth 4 and no 5-cycles	∞	$O(n^2)$ (Main Theorem)	-	-
Bipartite	∞	$O(n^2)$ [15]	-	-
Girth 5	$< \infty$ [7]	$O(n \log^2(n))$ [39]	-	-
Girth 6	$O(n^3)$ [15]	$O(n)$ [7]	-	-
Girth 7^+	$O(n \log n)$ [39]	-	-	-

Table 3.1: A summary of known results on the diameter of the reconfiguration graph of planar graphs.

free planar graphs. Our Main Theorem in this chapter is a step towards proving Cereceda's Conjecture for all triangle-free planar graphs.

Main Theorem. *If G is a planar graph with no 3-cycles and no 5-cycles, then diam$(\mathcal{C}_5(G)) \leq 4n^2$.*

3.1 Proof of Main Theorem

Before starting the proof, we review some relevant (mostly standard) definitions.

Let G be a plane graph. Denote by $F(G)$ the set of faces of G and by $d(f)$ the length of each such face f. A *k-vertex* is one with degree k. A *k^+-vertex* (resp. *k^--vertex*) is one with degree at least (resp. at most) k. A *k-neighbor* of a vertex v is a k-vertex adjacent to v. A *k-face* is one of length k. We define *k^+-neighbor, k^--neighbor, k^+-face*, and *k^--face* analogously. A cycle C is *separating* if $G \setminus V(C)$ is disconnected. A vertex v is *inner* if v does not lie on the outer face.

k-/k^+-/k^--vertex/ neighbor/ face

separating

inner

A plane graph G is *Type 1* if $\delta(G) \geq 3$ and *Type 2* if each of the following conditions is satisfied: (i) $\delta(G) \geq 2$, (ii) the outer face f_0 is a 7-face, (iii) $V(f_0) \subsetneq V(G)$, and (iv) every 2-vertex of G lies on f_0; see Figure 3.1. Note that Type 1 and Type 2 graphs are not necessarily mutually exclusive. That is, a graph G could satisfy the criteria to be both Type 1 and Type 2. For example, G could satisfy (ii)-(iv) and have $\delta(G) \geq 3$, so also satisfy (i). In that case, we will always define G to be Type 2, unless we explicitly say otherwise, since our conclusions for Type 2 graphs will be stronger.

Type 1, 2

Let $T := \emptyset$ if G is Type 1 and $T := V(f_0)$ if G is Type 2. Let $V_1 := \{v \in V(G) \setminus T : d_G(v) \leq 3\}$, $V_2 := \{v \in V(G) \setminus T : d_{G \setminus V_1}(v) \leq 3\}$, $V_3 := \{v \in V(G) \setminus T : v \notin V_1 \cup V_2\}$, and $V_4 := T$. For all $i \in \{1, 2, 3, 4\}$, each vertex $v \in V_i$ has *level i*.

level

Let v be a 3-vertex with neighbors v_1, v_2, v_3, and assume that v is incident with 3 distinct faces. For each

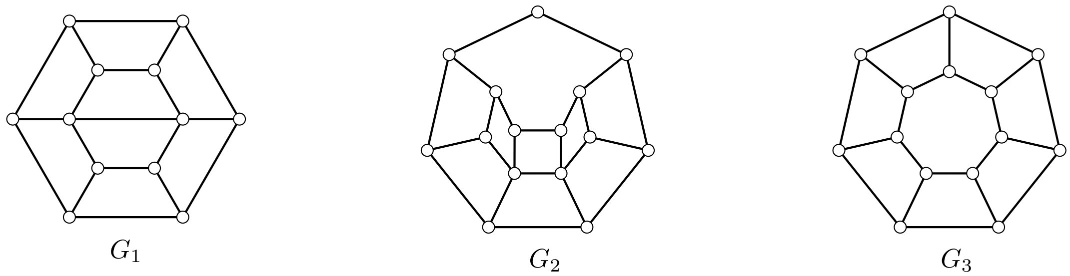

Figure 3.1: The graph G_1 is Type 1 since $\delta(G_1) \geq 3$. The graph G_2 is Type 2 since (i) $\delta(G_2) \geq 2$, (ii) the outer face f_0 is a 7-face, (iii) G_2 contains inner vertices (so, $V(f_0) \subsetneq V(G_2)$), and (iv) the only 2-vertex of G_2 lies on f_0. The graph G_3 is Type 1 and Type 2 since (i) $\delta(G_3) \geq 2$, (ii) f_0 is a 7-face, (iii) $V(f_0) \subsetneq V(G_3)$, and (iv) G_3 contains no 2-vertices.

$i \in [3]$, the *face opposite* v_i (with respect to v) is the face incident with vertex v that is not incident with edge vv_i. A 3-vertex v is *good* if it has a neighbor w of level at most 2 whose opposite face is a 4-face. We call w a *good neighbor* of v.

face opposite

good (neighbor)

Bousquet and Heinrich [15], in their proof for bipartite planar graphs, introduced the notion of the *level of a vertex.* They used it to identify a pair of vertices in G at distance 2 along a common face that could be identified to proceed by induction. We will apply a similar technique for planar graphs with no 3-cycles and no 5-cycles. However, identifying such vertices in these graphs might create a 5-cycle, if the vertices lie on a 7-cycle (it will not create any 3-cycle since G has no 5-cycle).

To avoid this problem, we show how to find good vertices inside a subgraph of our graph which does not contain any separating 7-cycles; note that identifying a pair of vertices in such a subgraph cannot create a 5-cycle. In particular, if our graph has a separating 7-cycle, then we pick an innermost such cycle and, using discharging, we show that we can use the extra charge from Euler's formula to ensure we find a good vertex away from the outer face. This technique has been used previously; for example, see [34, 35]. Our next lemma and subsequent corollary establish the existence of good vertices inside separating 7-cycles.

Key Lemma. *Let G be a connected plane graph with no 3-cycles, no 5-cycles, and no separating 7-cycles. If either G is Type 1 or G is Type 2, then G contains a good vertex v with a good neighbor w such that $v, w \notin T$.*

Proof. Let $\widehat{V} := V(f_0)$. The proofs for Type 1 graphs and Type 2 graphs are similar, though for Type 2 graphs the proof is harder. Recall that if G is Type 1, then $T := \emptyset$ and if G is Type 2, then $T := \widehat{V}$. We use discharging to show that $V(G) \setminus T$ contains some good vertex and its good neighbor. Assume instead that G is a counterexample to the lemma. Denote by $ch(v)$ and $ch(f)$ (resp. $ch^*(v)$ and $ch^*(f)$) the initial (resp. final) charges of each vertex v and each face f. Let $\mathrm{ch}(v) := d(v) - 4$ for every $v \in V(G)$ and $\mathrm{ch}(f) := d(f) - 4$ for every $f \in F(G)$. Using Euler's formula, the total initial charge is -8:

$$\begin{aligned}\sum_{v\in V(G)}(d(v)-4)+\sum_{f\in F(G)}(d(f)-4) &= \sum_{v\in V(G)} d(v) - 4|V(G)| + \sum_{f\in F(G)} d(f) - 4|F(G)| \\ &= 2|E(G)| - 4|V(G)| + 2|E(G)| - 4|F(G)| \\ &= 4(|E(G)| - |V(G)| - |F(G)|) \\ &= -8. \end{aligned} \tag{3.1}$$

Now we redistribute charge according to rules (R1) and (R2) below and show that the total final charge is greater than -8, a contradiction.

(R1): Every 6^+-face f gives $\frac{d(f)-4}{d(f)}$ to every incident vertex.[1]

(R2): Every 3-vertex not in T takes $\frac{1}{3}$ from every neighbor that has level at least 3.

Assume first that G is Type 1. We show that $\mathrm{ch}^*(v) \geq 0$ and $\mathrm{ch}^*(f) \geq 0$ for all v and all f. Each 4-face f loses no charge, so $\mathrm{ch}^*(f) = 4 - 4 = 0$. By (R1), each 6^+-face f has $\mathrm{ch}^*(f) = (d(f) - 4) - d(f)\left(\frac{d(f)-4}{d(f)}\right) = 0$. Note that the only vertices that receive charge from their neighbors are 3-vertices that are not in T, and the only vertices that give charge to their neighbors are vertices of level at least 3, by (R2). Further, observe that each vertex v of level 3 has $d(v) \geq 4$ and at least four 4^+-neighbors; otherwise, v has level at most 2. So, each of those four 4^+-neighbors receives no charge from v, and $\mathrm{ch}^*(v) \geq d(v) - 4 - \frac{1}{3}(d(v)-4) = \frac{2}{3}(d(v)-4) \geq 0$, by (R2). Moreover, each vertex v of level 2 has $d(v) \geq 4$ and loses no charge, so $\mathrm{ch}^*(v) = \mathrm{ch}(v) = d(v) - 4 \geq 0$. Finally, every 3-vertex v has $\mathrm{ch}(v) = 3 - 4 = -1$, so needs to receive at least 1.

Let v be a 3-vertex. First assume that v appears on three distinct faces. Let M_v be the set of neighbors of v of level at least 3. Note that, for every neighbor x of v, either (i) $x \in M_v$ or (ii) the face opposite x is a 6^+-face; otherwise, v is a good vertex with a good neighbor x, a contradiction. Now $\mathrm{ch}^*(v) \geq -1 + \frac{1}{3}|M_v| + \frac{1}{3}(3 - |M_v|) = 0$ by (R1) and (R2).

Now assume v appears on at most two distinct faces. So some edge e incident with v is a cut-edge. Let f' be the face containing e. Since $\delta(G) \geq 3$, each component of $G - e$ contains a cycle that lies in the boundary walk of f'. Further, these cycles are vertex disjoint; see Figure 3.2. Because G is triangle-free,

[1] If a vertex v appears multiple times on a boundary walk of f, then f gives v charge $\frac{d(f)-4}{d(f)}$ for each such incidence.

$d(f') \geq 2(4) = 8$. Since v appears at least twice on the boundary walk of f', the charge v receives from f' is at least $2(8-4)/8 = 1$, and $\text{ch}^*(v) \geq 0$. Thus, the lemma holds when G is Type 1.

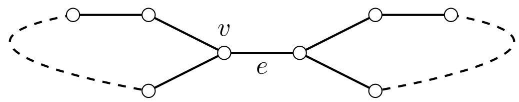

Figure 3.2: An example of the case when G is Type 1.

Now we assume that G is Type 2. By the arguments above, $\text{ch}^*(f) \geq 0$ for all f and $\text{ch}^*(v) \geq 0$ for every vertex v of level 1, 2, or 3, i.e., for every $v \notin T$. It remains to show that $\sum_{v \in \widehat{V}} \text{ch}^*(v) > -8$. Let n_2 be the n_2
number of 2-vertices on f_0. Note that $n_2 \leq 6$; otherwise, since G is connected, $\widehat{V} = V(G)$, a contradiction.

Claim 3.1. *If $w_1, \ldots, w_m$ are 2-vertices on f_0 with $m \geq 2$ and $w_i w_{i+1} \in E(G)$ for all $i \in [m-1]$, then their incident face f (other than f_0) is a 6^+-face. Thus, each w_i, as well as the 3^+-neighbors of w_1 and w_m on f_0, receives $\frac{d(f)-4}{d(f)}$ from f.*

Proof of claim. The second statement follows directly from the first via (R1). To see the first statement, let v_1 and v_2 be the neighbors of $w_{\lfloor \frac{m}{2} \rfloor}$ and $w_{\lfloor \frac{m}{2} \rfloor + 1}$ on f_0 (other than each other), respectively, and assume f is a 4-face. This means $m = 2$ since $d(v_1) \geq 3$ and $d(v_2) \geq 3$, and also $f = G[\{v_1, w_1, w_2, v_2\}]$. Denote by v_3, v_4, v_5 the remaining vertices on f_0. Now $G[\{v_1, v_2, v_3, v_4, v_5\}]$ is a 5-cycle, a contradiction. □

Recall that $\text{ch}^*(f) \geq 0$ for all f and $\text{ch}^*(v) \geq 0$ for all $v \notin \widehat{V}$. So, to reach a contradiction to (3.1), it suffices to show that $\sum_{v \in \widehat{V}} \text{ch}^*(v) > -8$. Note that each vertex on f_0 takes $\frac{d(f_0)-4}{d(f_0)} = \frac{3}{7}$ from f_0 by (R1). Moreover, by (R2), every 3^+-vertex v on f_0 satisfies $\text{ch}^*(v) \geq (d(v) - 4) - \frac{1}{3}(d(v) - 2) \geq \frac{-4}{3}$. Finally, every 2-vertex v satisfies $\text{ch}^*(v) \geq 2 - 4 = -2$. Let $g(\widehat{V})$ be the charge gained via (R1), from faces other than f_0, $g(\widehat{V})$
in total by all vertices on f_0. Now $\sum_{v \in \widehat{V}} \text{ch}^*(v) \geq -2n_2 - \frac{4}{3}(7 - n_2) + 7(\frac{3}{7}) + g(\widehat{V}) = \frac{-2}{3}n_2 - \frac{19}{3} + g(\widehat{V})$. Let $h(\widehat{V}) := \frac{2}{3}n_2 - \frac{5}{3}$. Now, it suffices to show that $g(\widehat{V}) > h(\widehat{V})$ since this implies $\sum_{v \in \widehat{V}} \text{ch}^*(v) > -8$. $h(\widehat{V})$

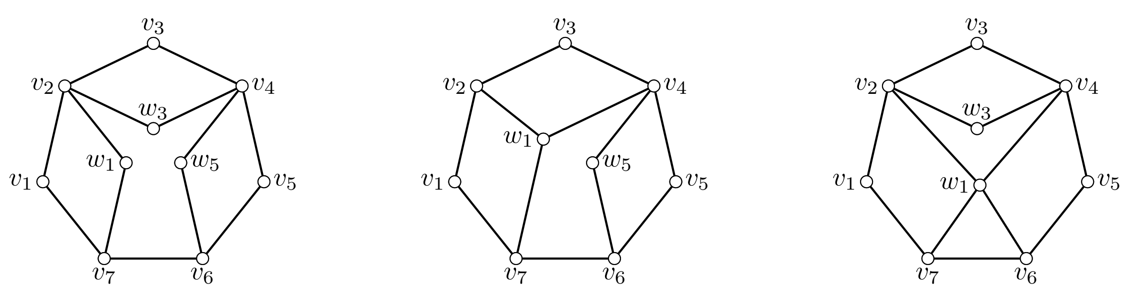

Figure 3.3: The three instances of Case 2, in the proof of the Key Lemma when no 2-vertices on f_0 are adjacent. Left: The nonneighbors of v_1, v_3, and v_5 on their incident 4-faces are distinct. Center: The nonneighbors of v_1 and v_3 on their incident 4-faces are the same, i.e., $w_1 = w_3$. Right: The nonneighbors of v_1 and v_5 on their incident 4-faces are the same, i.e., $w_1 = w_5$.

Case 1: $n_2 \leq 2$. Now $h(\widehat{V}) \leq \frac{-1}{3}$, so $g(\widehat{V}) \geq 0 > h(\widehat{V})$.

Case 2: $n_2 = 3$. If f_0 has at least two adjacent 2-vertices, then $g(\widehat{V}) \geq \frac{1}{3}(4) > \frac{1}{3} = h(\widehat{V})$, by Claim 3.1. So, we assume no pair of 2-vertices are adjacent on f_0. Let $v_1, \ldots, v_7$ be the vertices on f_0 in cyclic order. By symmetry, assume v_1, v_3, and v_5 are 2-vertices. Note that at least one 2-vertex is incident with a 6^+-face. To see this, assume every 2-vertex is incident with a 4-face. Let w_1, w_3, and w_5 be the inner vertices incident with the 4-faces of v_1, v_3, and v_5, respectively; see Figure 3.3. (Note that the boundary cycle of f_0 cannot have a chord, since any chord would create either a 3-cycle or a 5-cycle, both of which are forbidden by hypothesis.) Observe that w_1, w_3, and w_5 are distinct. Otherwise, by symmetry, either $w_1 = w_3$ or $w_1 = w_5$. In the former case $G[\{w_1, v_4, v_5, v_6, v_7\}]$ is a 5-cycle, and in the latter $G[\{w_1, v_6, v_7\}]$ is a 3-cycle, both of which are contradictions. But now $G[\{w_1, v_2, w_3, v_4, v_5, v_6, v_7\}]$ is a 7-cycle that separates v_3 from w_5, a contradiction. Hence, by symmetry, v_1 is incident with a 6^+-face f, so v_1, v_2, and v_7 each get at least $\frac{1}{3}$ from f by (R1). Thus, $g(\widehat{V}) \geq \frac{1}{3}(3) > \frac{1}{3} = h(\widehat{V})$.

Case 3: $n_2 = 4$. By Pigeonhole, at least two 2-vertices are adjacent on f_0. So, $g(\widehat{V}) \geq \frac{1}{3}(4) > 1 = h(\widehat{V})$ by Claim 3.1.

Case 4: $n_2 = 5$. By Pigeonhole, either (a) at least four 2-vertices on f_0 induce a path, or (b) three 2-vertices on f_0 induce a path P_1 and the other two 2-vertices on f_0 induce another path P_2, and no vertex on P_1 is adjacent to a vertex on P_2. In both cases, at least 6 vertices on f_0 get at least $\frac{1}{3}$ via (R1) from faces other than f_0, by Claim 3.1. So, $g(\widehat{V}) \geq \frac{1}{3}(6) > \frac{5}{3} = h(\widehat{V})$.

Case 5: $n_2 = 6$. Now the face f incident to all 2-vertices (other than f_0) is a 7^+-face. So, each of the six 2-vertices (as well as the single 3^+-vertex) on f_0 gets at least $\frac{3}{7}$ by Claim 3.1. Thus, $g(\widehat{V}) \geq \frac{3}{7}(7) > \frac{7}{3} = h(\widehat{V})$. □

Corollary 3.1. *If G is a plane graph with no 3-cycles, with no 5-cycles, and with $\delta(G) \geq 3$, then G contains a good vertex v. Furthermore, identifying v and its nonneighbor on an incident 4-face results in a smaller plane graph with no 3-cycles and no 5-cycles.*

Proof. Assume the corollary is false and G is a smallest counterexample. Now G is connected; otherwise, the result holds, by the minimality of G, for each component of G. Suppose first that G has no separating 7-cycles. By the Key Lemma for Type 1 graphs[2], G contains a good vertex v. Let v_1, v_2, v_3 be the neighbors of v with v_3 being a good neighbor of v and f being the 4-face opposite of v_3. Let w be the nonneighbor of

[2]Although G may also satisfy the criteria to be a Type 2 graph, by hypothesis it satisfies the criteria to be a Type 1 graph, so we invoke the Key Lemma for Type 1 graphs.

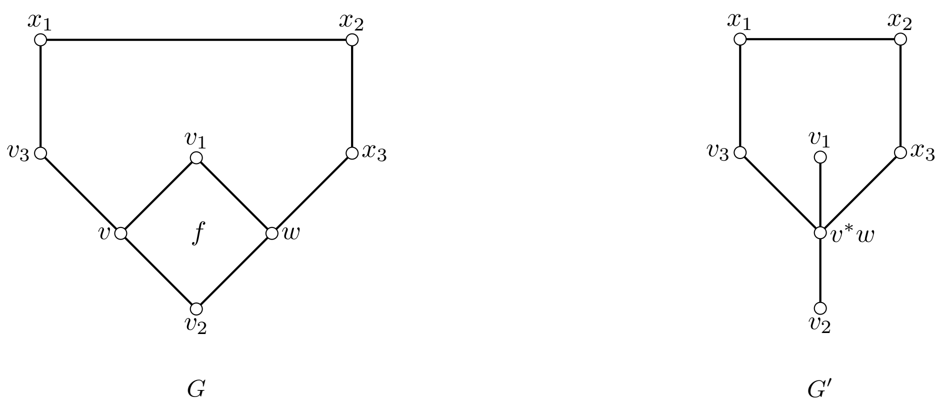

Figure 3.4: Now $vv_3x_1x_2x_3wv_1$ is a separating 7-cycle in G if $(v^*w)v_3x_1x_2x_3$ is a 5-cycle in G'.

v on f; see Figure 3.4. Form G' from G by identifying v with w to form a new vertex v^*w. Since v and w lie on a 4-face, G' is a plane graph. If G' has a 3-cycle, then G has a 5-cycle, a contradiction. Suppose G' has a 5-cycle, $(v^*w)v_3x_1x_2x_3$; call it C'. Since $d_G(v_1) \geq 3$, we note that v_1 has a neighbor y in G (other than v and w) and in G' we know y lies in the interior of C'. Thus, in G the 7-cycle $vv_3x_1x_2x_3wv_1$ separates y from v_2, again a contradiction. Hence, G' has no 3-cycles and no 5-cycles. So, the corollary holds when G has no separating 7-cycles.

Suppose instead that G has a separating 7-cycle and let C be an innermost such cycle. Denote by C_{in} C_{in} the subgraph induced by vertices that lie on C and inside C. Observe that C_{in} is a plane graph with no 3-cycles, no 5-cycles, and no separating 7-cycles. Further, $\delta(C_{\text{in}}) \geq 2$ with every 2-vertex of C_{in} lying on the outer face of C_{in}, which is induced by $V(C)$. By the Key Lemma, C_{in} contains a good vertex v. Define v_1, v_2, v_3, and w as in the previous paragraph. Note that $v, v_3 \notin V(C)$, by the Key Lemma. So, $d_G(v) = 3$ and v_3 has level at most 2 in G, i.e., v is a good vertex in G. We note that v_1 and v_2 cannot both lie on C. Suppose the contrary. Since C has length 7, some v_1, v_2-path P in C has odd length at most 5. If P has length 1 or 3, then Pv_1vv_2 is a 3-cycle or 5-cycle, a contradiction. So assume P has length 5. Now Pv_1vv_2 is a 7-cycle in C_{in} that separates v_3 from w, a contradiction. So assume by symmetry that v_1 does not lie on C.

Form G' from G by identifying v with w to form a new vertex v^*w. Clearly G' is a plane graph with no 3-cycles. Assume G' has a 5-cycle. So G has a (v, w)-path P of length five. It is easy to check that v_1 is not an interior vertex of P. Let $Q := Pvv_1w$, and denote its vertices by v, v_3, x_1, x_2, x_3, x_4, v_1 (in order), where $x_4 = w$. See Figure 3.5. Since $d(v_1) \geq 3$, let y be a neighbor of v_1 other than v and x_4. Note that y is not on Q, since this would give a 7-cycle with a chord, and thus a 3-cycle or a 5-cycle in G, a contradiction. So Q separates v_2 from y. Since C_{in} has no separating cycles, Q must not lie entirely in C_{in}. Since v, v_1, and v_3 are inner vertices of C_{in}, the vertices of Q outside C are contained in $\{x_2, x_3\}$.

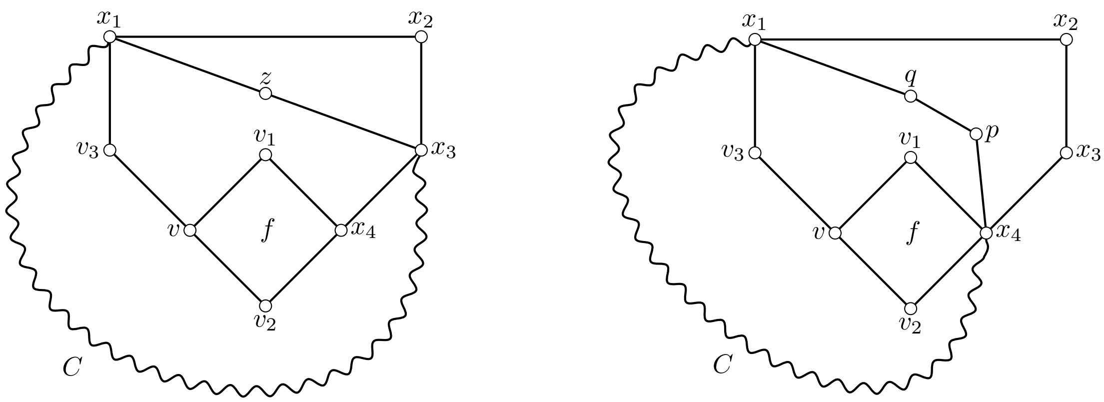

Figure 3.5: Two instances where the 7-cycle Q does not lie entirely in C. Left: Here x_2 lies outside C and z is not part of Q. Right: Here x_2 and x_3 both lie outside C, and both p and q are not part of Q.

Suppose that Q has exactly one vertex x_i with $i \in \{2, 3\}$ outside C. The path $x_{i-1}x_ix_{i+1}$ together with the shorter path in C between x_{i-1} and x_{i+1} forms a cycle of length at most 5. Thus, this must be a 4-cycle $x_{i-1}x_ix_{i+1}z$. If z is not on Q, then $(Q \setminus \{x_i\}) \cup x_{i-1}zx_{i+1}$ is a 7-cycle of C_{in}, which must be facial, since C_{in} has no separating 7-cycles; see the left of Figure 3.5. But now the path vv_1w lies on both a 7-face and 4-face in C_{in}, so v_1 is an inner 2-vertex of C_{in}, a contradiction. So instead z must lie on Q. But now either $x_{i-1}z$ or $x_{i+1}z$ must be a chord of Q, a contradiction.

So assume instead that Q has exactly two vertices outside C, and these are x_2 and x_3. Similar to above, the path $x_1x_2x_3x_4$ together with the shorter path on C between x_1 and x_4 forms a 4-cycle or 6-cycle. If this cycle has length 4, then x_1x_4 is a chord of Q, a contradiction. So the cycle must have length 6; we denote its vertices (in order) by $x_1x_2x_3x_4pq$. Note that neither p nor q is on Q, since p and q lie on C, but all vertices of Q other than x_1, x_2, x_3, x_4 are inner vertices of C_{in}. Thus, $(Q \setminus \{x_2, x_3\}) \cup x_4pqx_1$ is a 7-cycle in C_{in} as on the right in Figure 3.5; again it must be a facial 7-cycle. But now the path vv_1w lies on both a 7-face and a 4-face, so v_1 is an inner 2-vertex of C_{in}, a contradiction. □

The proof of our main-thm-proj1]Main Theorem (restated below) is now similar to a proof of Bousquet and Heinrich [15] who showed the same conclusion for the smaller class of all bipartite planar graphs.

Main Theorem. *If G is a plane graph with no 3-cycles and no 5-cycles, then diam$(\mathcal{C}_5(G)) \leq 4n^2$.*

Proof. Let G be a plane graph with no 3-cycles and no 5-cycles. Let φ_A and φ_B be two 5-colorings of G. (Recall that $n := |V(G)|$.) We will show that we can transform φ_A into φ_B by recoloring each vertex at most $4n$ times. We use induction on n.

Case 1: G contains a vertex v such that $d(v) \leq 2$. Assume $n > 1$; otherwise, the result follows trivially. Let φ'_A and φ'_B be the restrictions of φ_A and φ_B to $G - v$. By induction, there exists a sequence $\mathcal{S}'$ of recolorings that transforms φ'_A into φ'_B such that each vertex is recolored at most $4(n-1)$ times. We extend $\mathcal{S}'$ to a sequence $\mathcal{S}$ of recolorings in G. To form $\mathcal{S}$ in G, we can perform each recoloring step from $\mathcal{S}'$, except when a neighbor w of v is to be recolored with the current color α of v. In that case, we need to recolor v before recoloring its neighbor w. The number of colors unused on $N(v) \cup \{v\}$ is at least $5 - (1 + d(v)) \geq 2$. We recolor v with one of these colors that is not the target color in the next recoloring of a neighbor of v. So, we only need to recolor v at most once for every two successive recolorings of neighbors of v. Finally, we may need to recolor v to its target color in φ_B. Since $d(v) \leq 2$ and each neighbor of v is recolored at most $4(n-1)$ times in $\mathcal{S}'$, the total number of times we recolor v is at most $\frac{2(4(n-1))}{2} + 1 < 4n$.

Case 2: $\delta(G) \geq 3$. Assume $n > 1$; otherwise, the result follows trivially. By Corollary 3.1, G contains a good vertex v. Moreover, if v_1, v_2, v_3 are the neighbors of v, where v_3 is a good neighbor, and w is the nonneighbor of v on the 4-face opposite v_3, then G', which is formed from G by identifying v with w into a vertex v^*w, is a planar graph with no 3-cycles and no 5-cycles. v, w v_1, v_2, v_3

Note that if $\varphi_A(v) \neq \varphi_A(w)$, then we can transform φ_A into a coloring φ'_A such that $\varphi'_A(v) = \varphi'_A(w)$ and every vertex is recolored by this transformation at most twice. To see this, let $\alpha := \varphi_A(w)$. Note that $\varphi_A(v_1) \neq \alpha$ and $\varphi_A(v_2) \neq \alpha$ since $v_1, v_2 \in N(w)$. If $\varphi_A(v_3) \neq \alpha$, then we recolor v with α and we are done. So, assume $\varphi_A(v_3) = \alpha$. Recall that v_3 has level at most 2, so v_3 has at most 3 neighbors of degree greater than 3. Fix $x_1, x_2, x_3 \in N(v_3)$ such that each 4^+-neighbor of v_3 is among $\{x_1, x_2, x_3\}$. Now there exists a color $\beta \notin \{\alpha, \varphi_A(x_1), \varphi_A(x_2), \varphi_A(x_3)\}$; see Figure 3.6. If need be, we first recolor every β-colored 3-neighbor of v_3 (of which there may be arbitrarily many); this is possible since at least one color does not appear on the closed neighborhood of each 3-vertex. Observe that v might possibly be recolored, but w is not recolored since $\varphi_A(w) = \alpha \neq \beta$. We now recolor v_3 with β, recolor v with α, and are done (since w and v now both have color α, we can identify them, to form a new vertex).

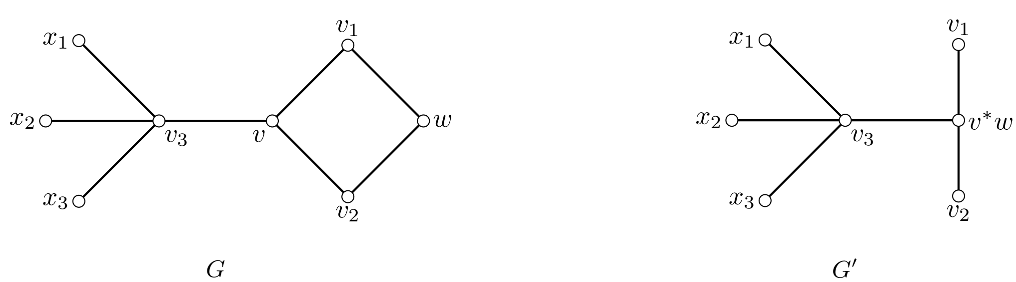

Figure 3.6: An instance of Case 2 in the proof of the Main Theorem.

Similarly, if $\varphi_B(v) \neq \varphi_B(w)$, then we can transform φ_B into a coloring φ'_B such that $\varphi'_B(v) = \varphi'_B(w)$. Note that φ'_A and φ'_B are proper 5-colorings of G'. By induction, there exists a sequence $\mathcal{S}'$ that transforms φ'_A into φ'_B in G' such that each vertex is recolored at most $4(n-1)$ times. It is easy to see that $\mathcal{S}'$ extends to a sequence $\mathcal{S}$ in G such that each vertex is recolored at most $4(n-1)$ times (for each recoloring of v^*w in G' we recolor both v and w in G).

Recall that we recolor each vertex at most twice when forming φ'_A from φ_A and at most twice when forming φ'_B from φ_B. Thus, the total number of times we recolor each vertex is at most $4(n-1)+4 \leq 4n$. □

Chapter 4

The Kempe Swap Model: In Most 6-regular Toroidal Graphs All 5-colorings are Kempe Equivalent

The $a \times b$ *toroidal grid* is the cartesian product of cycles of lengths a and b. The $a \times b$ *triangulated toroidal grid*, $T[a \times b]$, is formed from the toroidal grid by adding a diagonal inside each 4-face, so that all diagonals are parallel; see Figure 4.1. Clearly each $T[a \times b]$ is a 6-regular toroidal graph. And when $a \geq 3$ and $b \geq 3$, the graph $T[a \times b]$ is 4-colorable (see Lemma 4.3). Mohar and Salas [55] showed that not all 4-colorings are 4-equivalent (when a and b are divisible by 3). In contrast, Bonamy, Bousquet, Feghali, and Johnson [11] showed that all 6-colorings of $T[a \times b]$ are 6-equivalent. Further, they asked whether all 5-colorings of $T[a \times b]$ are 5-equivalent. This question motivates this chapter's main result. We answer the question affirmatively when $a \geq 6$ and $b \geq 6$. (triangulated) toroidal grid

Theorem 4.1. *If G is a triangulated toroidal grid $T[a \times b]$ with $a \geq 6$ and $b \geq 6$, then all 5-colorings of G are 5-equivalent.*

Our proof holds in more generality, but to state our main result we need one more definition. For an embedding of a graph G in the torus, the *edge-width* is the length of the shortest non-contractible cycle. Negami showed that if a toroidal graph is 6-regular, then it has a unique embedding[1] in the torus. So,

[1]We omit a formal definition of unique embedding; informally, it means that we can transform any embedding to any other by "sliding" the graph around the torus, keeping it embedded throughout this sliding process.

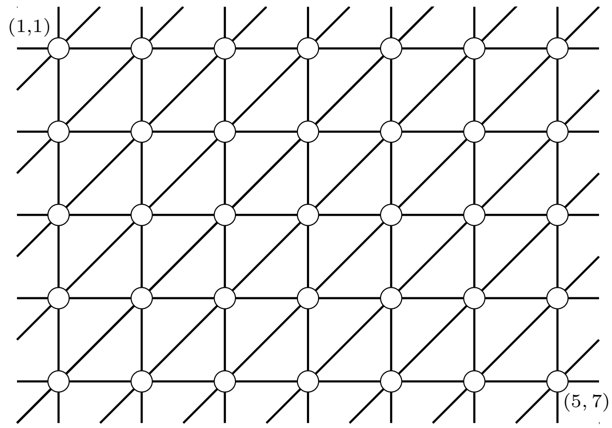

Figure 4.1: A triangulated toroidal grid $T[5 \times 7]$.

for each 6-regular toroidal graph G, by the *edge-width of* G we mean the edge-width of the unique toroidal embedding of G. The purpose of this paper is to prove the following. edge-width

Main Theorem. *If G is a 6-regular toroidal graph with edge-width at least 7, then all 5-colorings of G are 5-equivalent.*

Our Main Theorem implies the following corollary, the proof of which can also be found in Section 4.1.

Main Corollary. *If G is a 6-regular toroidal graph on n vertices chosen uniformly at random, then asymptotically almost surely all 5-colorings of G are 5-equivalent.*

We now mention another direction of research where the present work is relevant. Recall the following. A graph is d-degenerate if each of its subgraphs has minimum degree at most d. Las Vergnas and Meyniel proved [50] (see Lemma 4.1) that if G is d-degenerate, then all k-colorings of G are k-equivalent[2], whenever $k > d$. Thus, every planar graph G has all k-colorings equivalent whenever $k > 5$. Meyniel [53] extended this result to the case $k = 5$. In contrast, Mohar [54] constructed planar graphs with arbitrarily many 4-colorings no two of which are 4-equivalent (Figure 2.6).However, he showed that if G is planar and $\chi(G) = 3$, then all 4-colorings of G are 4-equivalent. It is easy to check that if G is bipartite, then all 3-colorings of G are 3-equivalent. By combining these results, Mohar showed that if G is planar, then all $(\chi(G) + 1)$-colorings of G are $(\chi(G) + 1)$-equivalent. It is natural to ask whether the same result holds for all toroidal graphs. Our results can be viewed as a step toward answering this question affirmatively.

We prove the Main Theorem in Sections 4.1 and 4.2. Since the proof of Theorem 4.1 is very similar to

[2]Generalized as Theorem 2.5 in Chapter 2.

that of the Main Theorem, we only discuss it in Section 4.3, where we sketch how to adapt the proof of the Main Theorem. It is worth noting that the Main Theorem immediately implies the case where $a \geq 7$ and $b \geq 7$. By symmetry, we can assume that $a \leq b$. So the discussion in Section 4.3 just handles the case when $a = 6$.

4.1 Proof Outline and Preliminaries

4.1.1 An Introduction to Good Templates

Given a coloring φ of G, our idea is to identify the vertices in one or more independent sets, each of which receives a common color under φ. If the resulting graph is 4-degenerate, then all of its 5-colorings are 5-equivalent, as shown in Lemma 4.1; and these 5-colorings correspond to some of the 5-colorings of G (precisely those 5-colorings of G where the vertices in each identified independent set receive a common color). A *good 4-template* in G is an independent set T of size 4 such that identifying all vertices of T yields a 4-degenerate graph; see Figure 4.2. We show that if φ_1 and φ_2 are 5-colorings that each use a common color on some good 4-template, say T_1 and T_2, then φ_1 and φ_2 are 5-equivalent. We also show that every 5-coloring is 5-equivalent to a 5-coloring that uses a common color on the vertices of some good 4-template. Together, these two steps prove our Main Theorem. To formalize this approach, we introduce more terminology. good 4-template

A *template* T in G is a collection of disjoint independent sets; each set in T is a *color* of T. A template with a single color is *monochromatic*. A template T *appears* in a coloring φ of G if the vertices in each color of T receive a common color under φ. If T appears in φ, then also φ *contains* T. By *contracting a template* T, we mean identifying the vertices in each color of T. When we contract template T in G, the resulting graph is denoted G_T. A template T is *good* for G if G_T is 4-degenerate. We will see that every 6-regular toroidal graph G is vertex transitive. So if G has a single good template, then it has many of them. Thus, given a 5-coloring φ, our focus will be on finding a good 4-template contained in φ (or some 5-coloring that is Kempe 5-equivalent to φ). Good templates play a central role in our proof of the Main Theorem. This is due to the following three easy lemmas. The first was originally proved in [50]; but for completeness, we include the proof. The third, Lemma 4.3, holds in more generality, which we discuss in Section 4.3. template color appears contains contracting a template good

Lemma 4.1. *Let G be a graph with a vertex w such that $d(w) < k$. If all k-colorings of $G - w$ are k-equivalent, then also all k-colorings of G are k-equivalent. Thus, if G is d-degenerate and $d < k$, then all*

k-colorings of G are k-equivalent.

Proof. The second statement follows from the first by induction on $|V(G)|$, where the base case $|V(G)| = 1$ holds trivially. Now we prove the first. Let w be a vertex with $d(w) < k$. Let $G' := G - w$. Let φ_1 and φ_2 denote k-colorings of G and let φ_1' and φ_2' denote their restrictions to G'. By hypothesis, all k-colorings of G' are k-equivalent. So there exists a sequence $\psi_1', \ldots, \psi_s'$ of k-colorings of G' such that $\psi_1' = \varphi_1'$, $\psi_s' = \varphi_2'$ and each ψ_{i+1}' differs from ψ_i' by only a single Kempe swap. We extend each ψ_i' to a k-coloring ψ_i of G as follows. Let $\psi_1 = \varphi_1$. Suppose that ψ_{i+1}' differs from ψ_i' by an α/β swap at a vertex v_i. To construct ψ_{i+1} from ψ_i we use the same α/β swap at v_i, unless w lies in the same α/β component as v_i and has at least 2 neighbors in that component. In that case, some color $\gamma \in \{1, \ldots, k\}$ is not used on the closed neighborhood of w. Now we first recolor w with γ, and then use the α/β swap at v_i. We call the resulting coloring ψ_{i+1}. By induction on i, each ψ_i restricts to ψ_i' on G'. Thus, ψ_s agrees with φ_2 on all vertices except for possibly w. If needed, recolor w with its color in φ_2. □

Lemma 4.2. *If T is a good template in a graph G, then all 5-colorings of G containing T are 5-equivalent.*

Proof. Form G_T from G by contracting T. Note that each 5-coloring φ of G containing T corresponds to a 5-coloring φ_T of G_T (formed by contracting T in φ). Since T is good, G_T is 4-degenerate. By Lemma 4.1, all 5-colorings of G_T are 5-equivalent. If η and ζ are 5-colorings of G containing T, then η_T and ζ_T are 5-equivalent colorings of G_T. Further, this is witnessed by a sequence of Kempe swaps. This same sequence of Kempe swaps witnesses the 5-equivalence of η and ζ. To simulate in G an α/β swap at a vertex v in G_T, we simply perform an α/β swap at each vertex in G that was identified to form v. (If $v \in V(G)$, then this is a single swap; but if v represents some non-singleton color of T, then this may be multiple swaps.) □

Lemma 4.3. *Let G be a 4-colorable graph. If φ_1 and φ_2 are 5-colorings of G that contain monochromatic good templates T_1 and T_2, then φ_1 and φ_2 are 5-equivalent.*

Proof. Let φ_0 be a 4-coloring of G; for concreteness, assume that φ_0 does not use green. Form φ_1' and φ_2' from φ_0 by recoloring the vertices of T_1 and T_2, respectively, with green. Now φ_1 and φ_1' are 5-equivalent, by Lemma 4.2, since they both contain the good template T_1. Similarly, φ_2 and φ_2' are 5-equivalent. Finally, φ_1' and φ_2' are both 5-equivalent to φ_0, since each is formed from φ_0 by recoloring an independent set with green (and each such recoloring step is a valid Kempe swap). □

We will see that all 6-regular toroidal graphs with edge-width at least 7 are 4-colorable (with one exception, which we handle separately). This was proved by Yeh and Zhu [65], building on work of Collins and Hutchinson [25]. The latter also proved that $T[a \times b]$ is 4-colorable whenever $a \geq 6$ and $b \geq 6$. So most of the work in proving our Main Theorem (as well as Theorem 4.1) goes to showing that if G is a 6-regular toroidal graph with edge-width at least 7, then every 5-coloring of G is 5-equivalent to a coloring that contains a monochromatic good template. That is the content of Section 4.2.

In view of Lemmas 4.2 and 4.3, and ideas in the previous paragraph, we need a tool to prove that certain templates are good. A *4-degeneracy order* of a graph G is an order σ of $V(G)$ such that each vertex in σ has at least $d(v) - 4$ neighbors that appear earlier in σ. A *4-degeneracy prefix* of a graph G is an order σ of some subset of $V(G)$ such that each vertex in σ has at least $d(v) - 4$ neighbors that appear earlier in σ. A subgraph H of G is *locally connected* if each pair of vertices in H that is at distance two in G is also at distance two in H. A subgraph H of G is *well-behaved* if (i) it is locally connected and (ii) $G - H$ is connected.

4-degeneracy order

4-degeneracy prefix

locally connected

well-behaved

Lemma 4.4. *Let H be a well-behaved subgraph of G and let V_H denote its vertex set. Let T be a template with all vertices in V_H; denote the new vertices in G_T by V_T. Suppose there exist $v_1, v_2 \in V_H \setminus V_T$ and there exist $v_3, v_4 \in V(G_T) \setminus V_H$ such that v_3 is a common neighbor of v_1, v_2 and that v_4 is a common neighbor of v_3, v_i for some $i \in \{1, 2\}$. If there exists a 4-degeneracy prefix σ for $V_H \setminus V_T$, then G_T has a 4-degeneracy order, with all vertices of V_T coming last in the degeneracy order.*

Proof. We will extend σ to all of $V(G_T) \setminus V_H$. Let $R := \emptyset$. We iteratively add vertices of $V(G_T) \setminus V_H$ to R, so that at each step (i) $G[R]$ is connected and (ii) σ can be extended to $(V_H \setminus V_T) \cup R$. Since $v_1, v_2 \in V_H \setminus V_T$ and v_3 is their common neighbor, we can add v_3 to R and append it to σ; similarly, we can add v_4. Let $S := V(G_T) \setminus (V_H \cup R)$. Now suppose that $S \neq \emptyset$. Choose $v \in R$ such that v has a neighbor in S. (This is possible since $G - H$ is connected.) Denote the neighbors of v in cyclic order by $w_1, \ldots, w_6$. Since $G[R]$ is connected, and $|R| \geq 2$, we can assume that $w_1 \in R$ and $w_6 \notin R$. If $w_6 \in S$, then we add w_6 to R and append it to the prefix, since its neighbors v and w_1 are already in R. So assume that $w_6 \in V_H$. Now $w_2 \notin V_H$, since $w_1, v \notin V_H$ and H is locally connected. If $w_2 \in S$, then we add w_2 to R, since its neighbors w_1 and v are in R. So assume that $w_2 \in R$. Note that $w_3 \notin V_H$, since H is locally connected and $w_6 \in V_H$ but $v \notin V_H$. Again, if $w_3 \in S$, then we add w_3 to R; so assume that $w_3 \in R$. If $w_4 \in S$, then we add w_4 to R, so assume that $w_4 \notin S$. If $w_4 \in V_H$, then also $w_5 \in V_H$, since H is locally connected and $w_4, w_6 \in V_H$. This contradicts our choice of v as having a neighbor in S. So $w_4 \notin V_H$, which implies that $w_4 \in R$. Finally, $w_5 \in S$ since v has some neighbor in S, by our choice of v. Thus, we can always grow R and σ, as desired. □

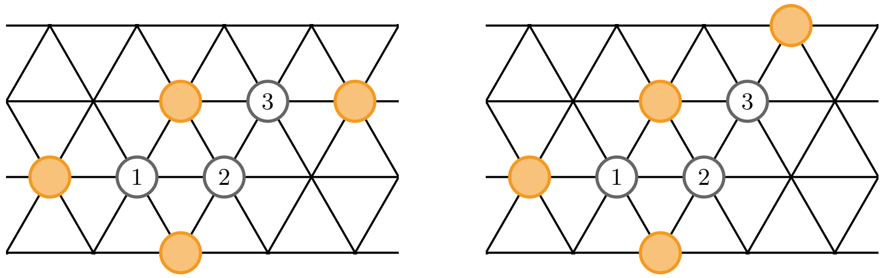

Figure 4.2: Two examples of good 4-templates in 6-regular toroidal graphs; each includes a triple centered at 1. In both examples, 2 and 3 serve as v_1 and v_2 in Lemma 4.4. It is easy to check that the subgraph induced by the four orange vertices and the three numbered vertices is locally connected; thus, it is well-behaved.

To apply Lemma 4.4, we will want to verify that some subgraphs H are locally connected. This is fairly easy, but is complicated slightly by the presence of non-contractible cycles. So the following lemma is useful.

Lemma 4.5. *Let G be a 6-regular toroidal graph with edge-width at least 7. If H is a subgraph of G with diameter at most 4, then H is locally connected unless there exists a vertex w such that H contains at least four neighbors of w but excludes w.*

First, suppose instead that there exists a vertex w such that H contains at least four neighbors of w, but excludes w. By Pigeonhole, H contains some pair of non-adjacent vertices such that w is their only common neighbor. So, H is *not* locally connected. Thus, the hypothesis on w in Lemma 4.5 is necessary.

Proof. Suppose there exist $w, x \in V(H)$ with $d_G(w, x) = 2$ but $d_H(w, x) > 2$. Denote the common neighbor(s) in G of w and x by y_1 and y_2 (if it exists). Since $d_H(w, x) > 2$, we have $y_1, y_2 \notin V(H)$. Since H has diameter at most four, a short case analysis shows that H contains at least four neighbors of either y_1 or y_2. (Since H has diameter at most 4, but G has edge-width at least 7, we can essentially ignore the possibility of non-contractible cycles creating problems.) □

Remark 4.1. Often when we claim, in later proofs, that a template is good, we will be implicitly using Lemmas 4.4 and 4.5. A *triple* is an independent set of size 3 with a common neighbor. For example, in each picture in Figure 4.2 a triple comprises the three orange neighbors of vertex 1. A triple itself is not a good 4-template. However, each triple is a subset of 12 good 4-templates. These include the two good 4-templates shown in Figure 4.2, along with 10 others that arise by rotation and reflection. Further, we show in Lemma 4.10 that if φ contains a triple, then φ is 5-equivalent to some 5-coloring that contains a good 4-template. Because triples have fewer vertices than any good template, they are easier to work with. Thus, we aim to show that every 5-coloring is 5-equivalent to a 5-coloring containing a triple. This motivates the following lemma.

triple

Lemma 4.6. *Let T be a good template for G, as witnessed by a subgraph H in Lemma 4.4. Let $v \notin V_H$ be a vertex such that $T \cup \{v\}$ is also a good template, with v as its own color. For each 5-coloring φ_0 of T and each $\alpha \in \{1, \ldots, 5\}$ such that using α on v extends φ_0 to a proper 5-coloring φ'_0 of $T+v$, there exists a proper 5-coloring of G that extends φ'_0. This v is called a* bonus vertex *for T. In particular, if there exists some other color C in T such that $C \cup \{v\}$ includes a triple, then every 5-coloring φ containing T is 5-equivalent to a 5-coloring containing a triple.* bonus vertex

Proof. For an example, see Figure 4.5, where T has the two colors $\{1, 2, 3, 4\}$ and $\{6, 9\}$, H is induced by $1, \ldots, 11$ and $v = 12$. The hypothesis that v is its own color implies that v is not identified with any other vertex of T. Thus, v can appear last in the degeneracy order, and is not required to have the same color as any other vertex in T. Since $T \cup \{v\}$ is a good template, every proper 5-coloring of $T \cup \{v\}$, that respects its colors, extends to a 5-coloring of G (we simply color the vertices outside $T \cup \{v\}$ in the reverse of the 4-degeneracy order). Fix a proper 5-coloring φ'_0 of $T \cup \{v\}$, and let φ be an extension of this coloring of $T \cup \{v\}$ to a 5-coloring of G. Let φ_1 be any other 5-coloring of G that contains T. Since φ_1 and φ both contain T, they are 5-equivalent. If $C \cup \{v\}$ includes a triple, then φ contains this triple, so φ_1 is indeed 5-equivalent to a 5-coloring that contains a triple. □

4.1.2 Shifted Triangulated Toroidal Grids

We denote each vertex in a triangulated toroidal grid $T[a \times b]$ by an ordered pair (i, j) with $i \in \{1, \ldots, a\}$ and $j \in \{1, \ldots, b\}$; vertices are numbered as entries of a matrix, with $(1, 1)$ in the top left and (a, b) in the bottom right; see Figure 4.1. So (i, j) is adjacent to $(i, j-1), (i-1, j), (i-1, j+1), (i, j+1), (i+1, j), (i+1, j-1)$ with arithmetic modulo a and b, as appropriate. Let $T[a \times b, c]$ denote a *triangulated $a \times b$ toroidal grid with shift c*. The vertex set is the same as that for $T[a \times b]$, and the edge set is the same except that edges from column b to column 1 are "shifted" by c. More precisely, each vertex (i, b) is adjacent to $(i+c-2, 1), (i+c-1, 1), (i+1, b), (i+1, b-1), (i, b-1), (i-1, b)$. So $T[a \times b] = T[a \times b, 1]$. The following useful result[3] characterizes all 6-regular toroidal graphs. triangulated $a \times b$ toroidal grid with shift c

Theorem 4.2 ([3, 58]). *Every 6-regular toroidal graph has the form $T[a \times b, c]$ for some positive integers a, b, c with $c \leq a$.*

This is Theorem 3.4 in [58]; the condition $c \leq a$ draws on the comment in [58] at the bottom of page 169.

[3]This characterization was stated by Altshuler [3], but his paper did not give a complete proof. About 10 years later, Negami published a proof [58], apparently unaware of the work of Altshuler.

A *circulant* $C_n[1, r, r+1]$ is a 6-regular graph[4] with vertex set $\{1, \ldots, n\}$ and with i and j adjacent if $i - j \in \{\pm 1, \pm r, \pm(r+1)\}$. Each circulant $C_n[1, r, r+1]$ has a natural embedding as a triangulation of the torus: begin with the hamiltonian cycle consisting of edges of "length" 1, wrapping around the torus in one direction, and now embed the remaining edges, each wrapping around the torus in the other direction. Building on work of Collins and Hutchinson [25], Yeh and Zhu [65] proved the following. Circulant $C_n[1, r, r+1]$

Theorem 4.3 ([65, 25])**.** *All 6-regular toroidal graphs are 4-colorable, with the following exceptions:*

1. $G \in \{T[3 \times 3, 2], T[3 \times 3, 3], T[5 \times 3, 2], T[5 \times 3, 3], T[5 \times 5, 3], T[5 \times 5, 4]\}$.

2. $G = T[m \times 2, 1]$ *with* m *odd.*

3. $G = C_n[1, r, r+1]$ *and* $n \in \{2r+2, 2r+3, 3r+1, 3r+2\}$ *and* n *is not divisible by 4.*

4. $G = C_n[1, 2, 3]$ *and* n *is not divisible by 4.*

5. $G = C_n[1, r, r+1]$ *and* $(r, n) \in \{(3, 13), (3, 17), (3, 18), (3, 25), (4, 17), (6, 17), (6, 25), (6, 33), (7, 19), (7, 25), (7, 26), (9, 25), (10, 25), (10, 26), (10, 37), (14, 33)\}$.

It is easy to prove that the graphs in Theorem 4.3 are not 4-colorable. More specifically, if $\alpha(G)$ denotes the indpendence number of G, then simple arguments show that each of the graphs in parts (1–4) have $\alpha(G) < |G|/4$, and thus, $\chi(G) \geq |G|/\alpha(G) > 4$; details are given in Section 3 of [65]. (The proof for graphs in (5) is ad hoc.) Thus, the result in Theorem 4.3 is best possible.

Lemma 4.7. *If* G *is a 6-regular toroidal graph with edge-width at least 6, and* $G \notin \{C_{26}[1, 10, 11], C_{37}[1, 10, 11]\}$, *then* G *is 4-colorable.*

Proof. This is an easy consequence of Theorem 4.3. Each graph in (1), (2), (3), and (4) has edge-width at most 5, 3, 3, and 3 (respectively). So we only consider graphs in (5). Note that $C_n[1, r, r+1]$ always has edge-width at most $r+1$. So we assume that $r \geq 6$. When $(r, n) = (6, 17)$, we have the non-contractible cycle with successive edge lengths $6, 6, 6, -1$. When $(r, n) = (6, 25)$, we have $6, 6, 6, 7$. When $(r, n) = (6, 33)$, we have $6, 6, 7, 7, 7$. When $(r, n) = (7, 19)$, we have $7, 7, 7, -1, -1$. When $(r, n) = (7, 25)$, we have $8, 8, 8, 1$. When $(r, n) = (7, 26)$, we have $8, 8, 8, 1, 1$. When $(r, n) = (9, 25)$, we have $9, 9, 9, -1, -1$. When $(r, n) = (10, 25)$, we have $11, 11, 1, 1, 1$. When $(r, n) = (14, 33)$, we have $15, 15, 1, 1, 1$. $\square$

[4]When $n \in \{2r, 2r+1, 2r+2\}$ it is actually a 6-regular multigraph.

Lemma 4.8. *Let G be a 6-regular toroidal graph with edge-width at least 7. If φ_1 and φ_2 are 5-colorings of G and each contains some good monochromatic template (possibly different templates in φ_1 and φ_2), then φ_1 and φ_2 are 5-equivalent.*

Proof. When G is 4-colorable, the result follows from Lemma 4.3. So, by Lemma 4.7, we only need to consider $G \in \{C_{26}[1, 10, 11], C_{37}[1, 10, 11]\}$. Further, $C_{26}[1, 10, 11]$ has a non-contractible cycle of length 6; it has edge lengths $11, 11, 1, 1, 1, 1$. So assume that $G = C_{37}[1, 10, 11]$. Let φ_0 denote the 5-coloring of G that uses color $i \bmod 4$ on vertex i, for each $i < 37$, and uses a fifth color on vertex 37, call it green. An *s-rotation* of φ_0 uses green on vertex s and uses color $i \bmod 4$ on vertex $i + s$, for each $i < 37$; a *rotation* is rotation
an s-rotation for some value of s. If T_1 is a 4-template, then some rotation φ_0^s of φ_0 uses green only on a vertex that is not in $\cup_{v \in T_1} N[v]$, since $|T_1| * 7 = 28 < 37$; in fact, there exist at least $37 - 28 = 9$ of these. So there exists a 5-coloring $\varphi_{1,s}$ of G that agrees with φ_0^s outside of T_1 and uses green on T_1; again, there are at least 9 of these. By Lemma 4.2, we know that all of these 5-colorings containing T_1 are 5-equivalent; further, from each of them, we can recolor the vertices of T_1 to reach φ_0^s. By the transitivity of equivalence, and the fact that G is vertex transitive, every 5-coloring containing T_1 is 5-equivalent to every rotation of φ_0. The same is true of 5-colorings containing T_2. Thus, φ_1 and φ_2 are 5-equivalent. □

Now we prove the Main Theorem, assuming results in Section 4.2. For reference, we restate it.

Main Theorem. *If G is a 6-regular toroidal graph with edge-width at least 7, then all 5-colorings of G are 5-equivalent.*

Proof. Let G be a 6-regular toroidal graph with edge-width at least 7. In Lemma 4.16 of Section 4.2 we prove that every 5-coloring of G is 5-equivalent to a 5-coloring that contains a good 4-template. Now Lemma 4.8 proves that all 5-colorings of G are 5-equivalent. □

An event $E(n)$ happens *asymptotically almost surely* if $\Pr(E(n)) \to 1$ as $n \to \infty$.

Lemma 4.9. *Fix a positive integer g. If G is a 6-regular toroidal graph on n vertices chosen uniformly at random, then asymptotically almost surely the edge-width of G is at least g.*

Proof. Fix a positive integer n. We prove an upper bound on the fraction of 6-regular toroidal n-vertex graphs that have edge-width less than g, and we show that this fraction tend to 0 as n tends to infinity.

Claim 4.1. *Each 6-regular toroidal graph can be written as $T[a \times b, c]$ for at most 6 choices of positive integers (a, b, c), such that $c \leq a$.*

Claim 4.1 follows immediately from Corollary 3.7 of [58]. The idea of that proof is to start from a vertex and follow edges along a "straight line" until returning to that vertex. More formally, to follow a straight line, the successor of each vertex v is 3 later than its predecessor in the cyclic order of neighbors of v. The factor 6 in Claim 4.1 arises from the 6 choices of neighbors to leave our initial vertex. (Since each graph $T[a \times b, c]$ is vertex-transitive, our choice of initial vertex has no effect.)

Claim 4.2. *There exists a constant C_g such that, for all positive integers a and b, at most C_g graphs of the form $T[a \times b, c]$ (with $c \leq a$) have edge-width less than g.*

Fix positive integers a and b. We bound (independently of a and b) the number of integers c with $1 \leq c \leq a$ such that $T[a \times b, c]$ has edge-width less than g. We make no effort to optimize this bound, only to show that it is independent of a and b. Suppose that $T[a \times b, c]$ has a non-contractible cycle C of length less than g. Let r denote the number of number of times that C wraps around the torus in the direction of the dimension of length a. Let s denote the number of edges on C of length 1 in this direction, and let t_1 and t_2 denote the numbers of edges on C from column b to column 1 of lengths c and $c + 1$, respectively. More precisely, let s, t_1, t_2 denote the net numbers of these edges traversed in the positive direction as we traverse C (for some arbitrary orientation). So $s + ct_1 + (c + 1)t_2 = ra$. Thus, $c = (ra - s - t_2)/(t_1 + t_2)$. Since s, t_1, t_2, and r are all bounded in terms of g (e.g. $g > s + t_1 + t_2 > -g$ and $g + 1 \geq r \geq -(g + 1)$), so is the number of choices for c. This proves the claim.

Let $d(n)$ denote the number of divisors of n. To build a 6-regular toroidal graph $T[a \times b, c]$ on n vertices, we have $d(n)$ choices for the pair (a, b). So the number of 6-regular toroidal graphs with edge-width less than g is at most $d(n)C_g$, by Claim 4.2. The number of 6-regular toroidal graphs for the pair (a, b) is precisely a, since $1 \leq c \leq a$. So the number for the pairs (a, b) and (b, a) is $a + b \geq \max(a, b) \geq \sqrt{n}$. By Claim 4.1, this overcounts the total number of 6-regular toroidal graphs on n vertices by a factor of at most 6. Thus, the fraction of 6-regular toroidal graphs on n vertices with edge-width less than g is at most $d(n)C_g/(((d(n)/2)\sqrt{n})/6) = 12C_g/\sqrt{n}$, which tends to 0 as n tends to infinity. □

This theorem and lemma immediately imply the following corollary.

Main Corollary. *If G is a 6-regular toroidal graph on n vertices chosen uniformly at random, then asymptotically almost surely all 5-colorings of G are 5-equivalent.*

4.2 Good Templates

Throughout this section we assume that G is a 6-regular toroidal graph with edge-width at least 7. A coloring φ of G is *nice* if it is 5-equivalent to some coloring that contains a good 4-template. In this section nice we prove that every 5-coloring φ of G is nice. Our proof uses a sequence of lemmas, with increasingly weaker hypotheses, culminating in the desired result. We adopt conventions common to all figures in this section. We generally number vertices in the order that we determine their colors. Sometimes, we do not determine the color of a vertex, but rather determine a color not used there. For such vertices, we restart the numbering, using higher numbers. To denote that a vertex is not colored green, we lightly shade it green.

We often argue by symmetry. For example, we assume that any triple centered at 30 in Figure 4.3 is $\{1,2,3\}$, rather than $\{28,31,32\}$. It is possible that the only automorphism of G that fixes vertex 30 is the identity map. So we do not quite mean that we can apply the same argument to G under some automorphism. But we do mean something close. If we consider the 6-regular triangulation of the plane, the set of automorphisms that fix a given vertex v has size 12; it has 6-fold rotational symmetry, together with reflection through a line containing v. The essential point is that our arguments do not make use of the global structure of G, but only of subgraphs of G of diameter at most 6. Each time that we invoke symmetry, we appeal to a map from one such subgraph to another, typically with many vertices in common. And we always mean one of these 12 symmetries mentioned above. To highlight these symmetries, we draw our pictures as subgraphs of the 6-regular triangulation of the plane. However, a moment's reflection shows that these are isomorphic to subgraphs of a triangulated toroidal grid (as it is drawn in Figure 4.1).

Along similar lines, it is possible that two vertices drawn as distinct in some figure are in fact the same vertex, when G has a short non-contractible cycle. Such a pair must be drawn at distance at least 7. Similarly, vertices may be joined by an edge that is not shown; but such vertices must be drawn at distance at least 6. In some figures, the subgraph shown may have diameter 7 or more. However, the attentive reader will note that whenever we conclude that φ contains a good template T, the vertices of T are contained in some subgraph H with diameter at most 4. Thus, we avoid these potential complications.

Lemma 4.10. *If a 5-coloring φ of G contains a triple, then φ is nice.*

Proof. Assume the lemma is false, and let φ be a counterexample with a triple centered at 30, as in Figure 4.3.

(1-3) By symmetry, assume 1, 2, 3 are orange. Neither 6 nor 7 is orange; otherwise $\{1,2,3,6\}$ or $\{1,2,3,7\}$ is a good template. By symmetry, none of the following are orange: 9, 10, 12, 13, 15, 16, 18, 19, 21, 22.

(4) We will show that 8 is not orange; suppose instead that it is. If 23 or 24 is orange, then φ has a good orange template at $\{1, 2, 8, w\}$, where $w \in \{23, 24\}$. So neither 23 nor 24 is orange; similarly, neither 25 nor 26 is orange. Let α denote the color of 7. If neither 28 nor 29 is α, then an α/orange swap at 7,8 makes 7 orange (and does not change the color anywhere else except 8, and possibly 9). This gives a good template at $\{1, 2, 3, 7\}$, which shows that 8 is not orange. So it suffices to ensure that neither 28 nor 29 is α.

First suppose that 28 is α. If 27, 29, and 30 do not have distinct colors, then some color β is absent from the closed neighborhood of 28, so we use an α/β swap at 28, ensuring that neither 28 nor 29 is α. Thus, we assume that 27, 29, and 30 have distinct colors. If 28, 31, 32 have distinct colors, then let γ denote the color on 30. Now an α/γ swap at 28,30 reduces to the case where 28 (and also 29) is not color α, handled above. So assume that 28, 31, 32 do not have distinct colors. Thus, some color is absent from the closed neighborhood of 30. Now recoloring 30 ensures that 30 has the same color as 27 or 29, which we have handled above. Thus, we assume that 28 is not α.

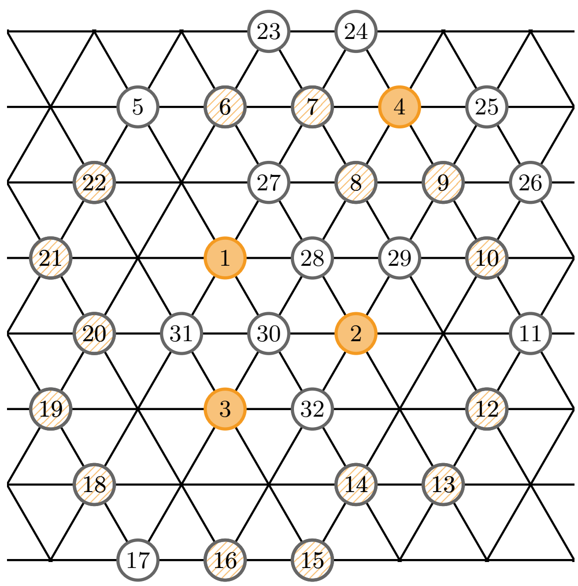

Figure 4.3: The proof of Lemma 4.10.

Suppose instead that 29 is α. By symmetry, we can assume that 9 and 27 are both some color β; otherwise, we reflect across the line through 8 and 28. Now we will recolor 8 so that it is not orange. If 28 and 4 have a common color, then some color is unused on the closed neighborhood of 8, so we can recolor 8. So assume that 28 and 4 use distinct colors; let γ be the color on 28. If 30 is α or β, then some color is unused on the closed neighborhood of 28, so we can recolor 28 using the color on 4. This allows us to recolor 8, which ensures that neither 8 nor 4 is orange. Now we can recolor 7 orange, reducing to a case above. So assume that 30 is neither α nor β. If γ does not appear on $\{31, 32\}$, then we swap colors on 28 and 30 and recolor 8 with γ. So assume that γ appears on $\{31, 32\}$. Now some color is unused on the closed neighborhood of 30, so we can recolor 30, ensuring that it uses either α or β; this reduces to a case above. Thus, we assume that 29 is not α. All of this shows that we can assume that 8 is not orange. By symmetry, neither 14 nor 20 is orange. So 4 is orange; otherwise we recolor 8 with orange and reduce to a previous case.

By symmetry, assume that green does not appear on the unique common neighbor of any of the sets of vertices $\{1,5\}$, $\{2,11\}$, $\{3,17\}$. We use green/orange swaps at each green/orange component containing vertices 1, 2, and 3. In the resulting coloring, 1, 2, and 3 are green but 4 is orange. So either one of 7, 8, or 9 is green or we can recolor 8 with green. For each possibility, we reduce to an earlier case. □

Let *parallel pairs* denote the template, with two colors and two vertices in each color, induced by vertices $\{1,2,3,4\}$ in Figure 4.4. parallel pairs

Lemma 4.11. *If a 5-coloring φ of G contains parallel pairs, then φ is nice.*

Proof. Assume the lemma is false, and let φ be a counterexample. By Lemma 4.10, it suffices to show that φ is equivalent to a 5-coloring that contains a triple. We assume 1,2 are green and 3,4 are blue, as in Figure 4.4.

(5-7) Note that 5 must be green, as follows. If some other neighbor v of 19 is green, then φ contains a good template at $\{v,1,2,3,4\}$. The 4-degeneracy prefix begins 20, 14 and the bonus vertex is 15 (which we make green). If 19 has no green neighbor, then we can recolor 19 green and get a green triple at $\{1,2,19\}$. So 5 is green, as claimed. By the same argument, 6 is blue. If 10 is green, then φ contains a good template at $\{1,2,3,4,10\}$, with prefix 20, 14, 21 and bonus vertex 15. So 10 is not green; by symmetry, neither is 11. Similarly, neither 12 nor 13 is blue. Suppose that none of 7, 8, 16, and 25 is green or blue. Now we use green/blue swaps at 1,3 and at 2,4. Each of these green/blue components that we recolored contains at most 4 vertices; in particular, neither contains 5, so 5 is still green. If 19 is blue, then we have a blue triple centered at 14. If a neighbor v of 19 other than 5 is blue, then we have a good template $\{v,1,2,3,4\}$, as above (with 15 as bonus vertex). Otherwise, we can recolor 19 blue to get a blue triple centered at 14. Thus, we assume that at least one of 7, 8, 16, and 25 is blue or green. By symmetry, assume that 7 is blue.

(8) Note that 14 is not blue, since then we have a blue triple at $\{3,4,14\}$. Similarly, 15 is not green. Suppose, to reach a contradiction, that 16 is blue. Let α be a color that is not green and not blue and that does not appear on 23 or 24. We use an α/green swap at 1 and an α/green swap at 2 (possibly a single swap, if 1 and 2 are in the same α/green component). It is easy to check that 1 and 2 are the only green vertices that get recolored α. In particular, 5 is still green.

If a neighbor v of 19 is α, then we have a good template at $\{v,1,2,3,4\}$. Otherwise, either 19 is α or we can recolor it α. In each case, we get an α triple at $\{1,2,19\}$. Thus, 16 is not blue. So 8 must be green. Otherwise, the blue/green component containing 2,4 has at most two more vertices: 21 and 22. After a blue/green swap at 2,4, we can recolor 21 green. This gives a good template at $\{1,2,3,4,5,21\}$, with bonus vertex 15. Thus, 8 is green.

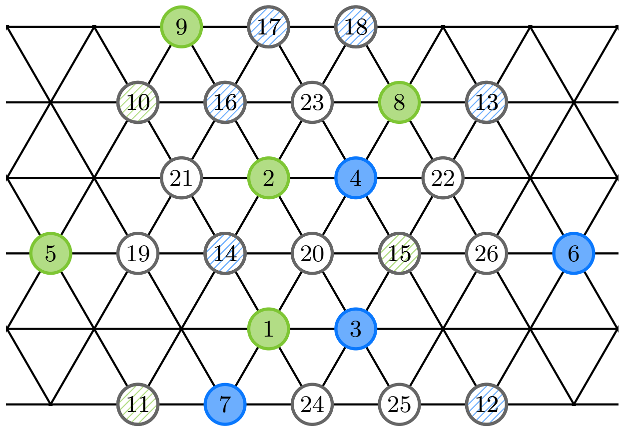

Figure 4.4: The proof of Lemma 4.11.

(9) Note that 17 is not green (or blue, as we will see); otherwise, it would form a green triple with 2 and 8. Suppose that 9 is not green. Let α be a color that is not green or blue and is also not used on 23 or 24. Now we use a green/α swap at 1 and also use a green/α swap at 2. It is easy to check that the only green vertices that get uncolored are 1 and 2. In particular, 5 is still green. If a neighbor v of 19 is colored α, then we have a good template at $\{v,1,2,3,4\}$, with bonus vertex 15. So either 19 is α or we can recolor it α. In either case, we get a triple colored α. Thus, 9 is green.

If 17 is blue, then we have a good template at $\{1,2,3,4,8,9,17\}$, with bonus vertex 15; the degeneracy prefix is 23, 16, 20, 14, 21. So 17 is not blue. Similarly, 18 is not blue; this good template replaces 17 in the previous one with 18, and the degeneracy prefix is 23, 16, 17, 20, 14, 21. Let β be a color that is not green or blue and is unused on 24 and 25. We use β/blue swaps at 3 and at 4. It is easy to check that the only blue vertices that get uncolored are 3 and 4. This new coloring again satisfies the hypotheses of the lemma, but 3 and 4 have a common color β different from the color on 6. We repeat our argument above that showed that 6 was blue. Now it shows that 6 is β, but in fact 6 is blue. This contradiction finishes the proof. $\square$

Lemma 4.12. *If a 5-coloring φ of G uses the same color on vertices 1, 2, 3, 4 in Figure 4.5, then φ is nice.*

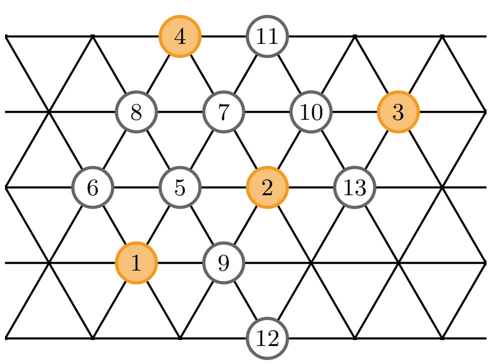

Figure 4.5: The proof of Lemma 4.12.

Proof. By symmetry, assume that 1, 2, 3, 4 are orange. By Pigeonhole, some color other than orange is used on at least two neighbors of 5; call this color blue. If these neighbors are 6 and 7, then φ contains parallel pairs centered at 5, so we are done by Lemma 4.11. Thus, we assume that blue is used on vertices 9 and v, where $v \in \{6,7,8\}$. If $v = 8$, then we have the good template $\{1,2,3,4,8,9\}$; the prefix is 5, 7, 10, 11, and the bonus vertex is 12. If

$v = 6$, then we have the good template $\{1, 2, 3, 4, 6, 9\}$; the prefix is 5, 7, 8, 10, 11 and the bonus vertex is 12. It is easy to check that the subgraphs induced by $\{1, \ldots, 11\}$ and $\{1, \ldots, 12\}$ are well-behaved, since each of them has diameter 4, but every non-contractible cycle has length at least 7. So we assume that blue is used on 7 and 9. Repeating the same argument, with 10 in place of 5, shows that blue is used on 7 and 13. But now $\{7, 9, 13\}$ is a blue triple, so φ is nice. □

Let *crossing pairs* denote the template, with two colors and two vertices in each color, induced by vertices $\{1, 2, 3, 4\}$ in Figure 4.6. crossing pairs

Lemma 4.13. *If a 5-coloring φ of G contains crossing pairs, then φ is nice.*

Proof. We assume the lemma is false, and let φ be a counterexample. Further, we assume that 1 is blue, 2 is red, 3 is blue, and 4 is red. We may also assume that 5 is green, 6 is orange, and 7 is purple. By symmetry between red and blue, either 8 is green or 8 is blue.

Case 1: 8 is green. Note that 9 is not red, or we have red/green parallel pairs centered at 7. Similarly, 11 is not blue. Neither 10 nor 12 is green, since φ has no green triple. So 9 and 10 are orange and blue (in some order), and 11 and 12 are red and purple (in some order).

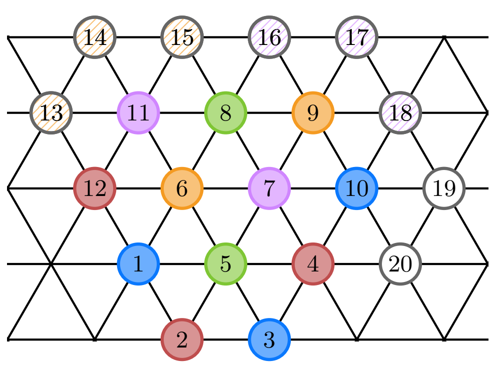

Figure 4.6: Case 1.1: 9 is orange and 10 is blue.

Case 1.1: 9 is orange, 10 is blue. (By horizontal symmetry, and permuting colors, this also includes the case that 9 is blue and 10 is orange, but 11 is purple and 12 is red.) None of 13, 14, 15 is orange; otherwise we have an orange triple centered at 8 (when 15 is orange) or we have a good template. When 13 is orange, the template is $\{1, 2, 3, 4, 6, 9, 13\}$, the prefix is 5, 7, 8, 11, 12, and the bonus vertex is the neighbor of 2 that forms a triple with 1 and 3. When 14 is orange, the template is $\{1, 2, 3, 4, 6, 9, 12, 14\}$ (with 12 in its own color class) the prefix is 5, 7, 8, 11; and the bonus vertex[5] is the neighbor w of 12 and 1 other than 6. After we color w orange, we are done by Lemma 4.12. If 11 is red and 12 is purple, then a red/orange swap at 11,6 gives a red triple centered at 5. So 11 is purple and 12 is red. By horizontal symmetry, none of 16, 17, 18 is purple. One of 19 and 20 is purple; otherwise a purple/blue

[5]We must include 12 in the first template; otherwise, adding the bonus vertex w creates a subgraph that is not locally connected, since $d_G(11, w) = 2$, but $d_{H+w}(11, w) > 2$.

swap at 7,10 gives a blue triple centered at 5. Similarly, one of 19 and 20 is orange; otherwise we use a purple/orange swap at 11,6,7,9 followed by a blue/orange swap at 10,7, which gives a blue triple centered at 5. So 19 and 20 are purple and orange (in some order). Now either we have purple/orange parallel pairs centered at 10, or we get them after a purple/orange swap at 11,6,7,9. This contradicts Lemma 4.11.

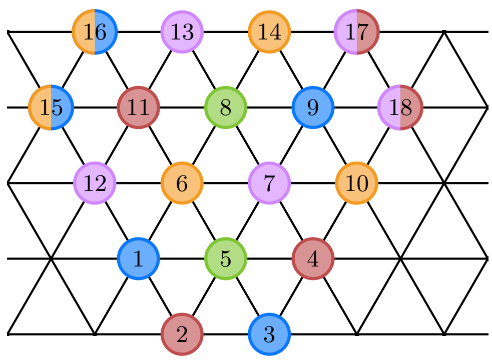

Figure 4.7: Case 1.2: 9 is blue, 10 is orange, 11 is red, 12 is purple.

Case 1.2: 9 is blue, 10 is orange, 11 is red, 12 is purple. One of 13 and 14 is orange; otherwise a green/orange swap at 5,6,8 gives an orange triple centered at 7. Similarly, one of 13 and 14 is purple; otherwise, a purple/green swap at 5,7,8 gives a purple triple centered at 6. By Lemma 4.11, φ has no parallel pairs, so 13 is purple and 14 is orange. If neither 15 nor 16 is orange, then an orange/red swap at 6,11 gives a red triple centered at 5. If neither 15 nor 16 is blue, then we recolor 11 blue, and recolor 6 red, which gives a red triple centered at 6. So 15 and 16 are orange and blue (in some order). By symmetry, 17 and 18 are purple and red (in some order). Now we use a blue/green swap at 8,9, recolor 11 green, and recolor 6 red; this gives a red triple centered at 5.

Case 2: 8 is blue. Now 9 is not red; otherwise, we have a good template at $\{1, 2, 3, 4, 9\}$, with prefix 5,7, and the bonus vertex is the neighbor of 2 that forms a triple with 1 and 3. So 9 is either orange or green.

Case 2.1: 9 is orange. (10-12) If 10 is green, then we have green/orange parallel pairs centered at 7. So 10 is blue. None of 15, 16, 17 is orange; otherwise we have an orange triple centered at 8 or a good template as in Case 1. (If $v = 15$, then our bonus vertex gives a blue triple. But if $v = 16$, then we use Lemma 4.12, as in Case 1.)

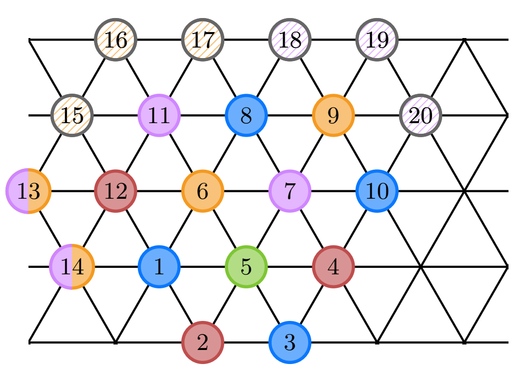

Figure 4.8: Case 2.1: 9 is orange.

If 11 is red, then a red/orange swap at 11,6 gives a red triple centered at 5. So 12 is red; otherwise, we recolor 6 red, getting a red triple centered at 5. And 11 is green or purple. Assume 11 is purple; otherwise, a green/purple swap at 5,7 reduces to this case (with green and purple interchanged). Similar to 15, 16, 17, none of 18, 19, 20 is purple. (13-14) Now 13 or 14 is orange; otherwise an orange/red swap at 12,6 gives a red triple centered at 5. Also, 13 or 14 is purple; otherwise a purple/orange swap at 11,6,7,9, followed by a red/purple swap at 12,6, again gives a red triple centered at 5. So 13 and 14 are orange and purple (in some order). But now we either have orange/purple parallel pairs centered at 12, or else we get them after a purple/orange swap at 11,6,7,9.

Case 2.2: 9 is green. (11-12) Now 11 or 12 is red; otherwise, we recolor 6 red and get a red triple centered at 5. And 11 or 12 is green; otherwise a green/orange swap at 5,6 reduces to case 2.1, with green and orange interchanged. Note that 10 is either blue or orange.

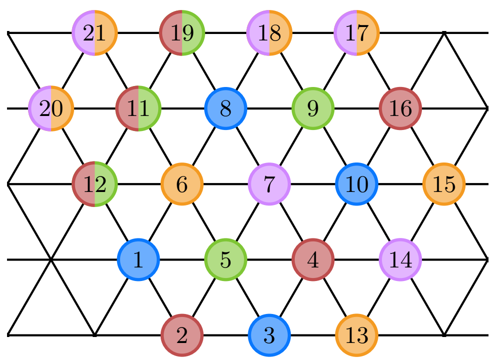

Figure 4.9: Case 2.2a: 10 is blue.

Case 2.2a: 10 is blue. (13-16) One of 13 and 14 is purple; otherwise we use a purple/red swap at 7,4, and recolor 6 purple, which gives purple/blue parallel pairs centered at 5. Similarly, one of 13 and 14 is orange; otherwise, a purple/orange swap at 6,7 reduces to the previous sentence. So 13 and 14 are purple and orange in some order. We assume 13 is orange and 14 is purple; otherwise a purple/orange swap at 6,7 reduces to this case. If 15 is green, then we have green/purple parallel pairs centered at 10. And if 15 is red, then we have the good template $\{1, 2, 3, 4, 8, 15\}$; the prefix is 5, 6, 7, 10 and the bonus vertex is the neighbor of 3 that forms a triple with 2 and 4. So 15 is orange. Now 16 is red; otherwise we have a purple triple centered at 10.

(17-19) Since we have no blue triple, 17 is purple or orange. If 18 is red, then we have red/blue parallel pairs centered at 9. Thus, 18 is purple or orange. If 19 is purple or orange, then we either have orange/purple parallel pairs centered at 8, or we get them after an orange/purple swap at 6,7. Thus, 19 is red or green.

(20-21) We show that 20 and 21 are both purple or orange. Suppose 11 is red. If neither 20 nor 21 is orange, then a red/orange swap at 11,6 gives a red triple centered at 5. If neither 20 nor 21 is purple, then a purple/orange swap at 6,7 followed by a red/purple swap at 11,6 again gives a red triple centered at 5. Instead, assume 11 is green. If neither 20 nor 21 is orange, then an orange/green swap at 11,6,5 gives orange/purple parallel pairs centered at 4. And if neither 20 nor 21 is purple, then we use a purple/orange swap at 6,7, followed by a green/purple swap at 11,6,5, followed by a purple/orange swap at 5,7; this again gives orange/purple parallel pairs centered at 4. So 20 and 21 are orange and purple. If 11 is green, then a green/blue swap at 11,8,9,10 gives a green triple centered at 7. So assume 11 is red, and use a red/blue swap at 11,8. If 18 is purple, then a red/orange swap at 8,6 gives a red triple centered at 5; so assume 18 is orange. Now a purple/orange swap at 6,7 and a purple/red swap at 8,6 again gives a red triple centered at 5.

Case 2.2b: 10 is orange. (13-14) Note that 13 or 18 is purple; otherwise, we use a purple/red swap at 7,4 and recolor 6 purple, which gives purple/blue parallel pairs centered at 5. Suppose that 18 is purple. Now 13 is not green; otherwise we have green/purple parallel pairs centered at 4. So 13 is orange. Now we reduce to Case 1 with 13, 18, 10, 7, in place of 1, 2, 3, 4 (and orange, purple, red in place of blue, red, green). Thus, 13 is purple. If 19 is red, then we have the good template $\{1, 2, 3, 4, 8, 19\}$; the prefix is 5, 6, 7, 10 and the bonus vertex is the neighbor of 3 that forms a triple with 2 and 4. So 19 is not red. Thus, 14 is red; otherwise a red/orange swap at 4,10 gives orange/blue parallel pairs centered at 5.

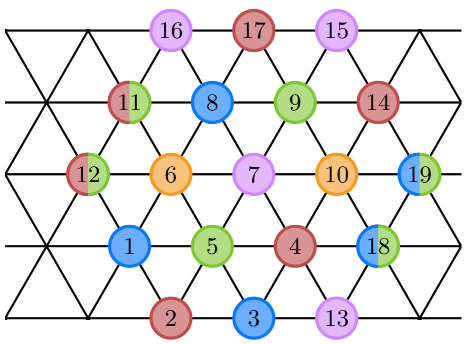

Figure 4.10: Case 2.2b: 10 is orange.

(15-17) If 15 is orange, then we have the good template $\{1, 2, 3, 4, 6, 10, 15\}$, with prefix 5, 7, 9. Our bonus vertex is 13, which allows us to finish by Lemma 4.12. Suppose 15 is blue. If 19 is not purple, then a purple/orange swap at 6,7,10 and a red/orange swap at 4,7 gives red/blue parallel pairs centered at 9. So assume 19 is purple. Now 18 is blue; otherwise we recolor 10 blue and get a blue triple centered at 9. Now 17 is orange; otherwise an orange/green swap at 10,9 gives parallel orange/green pairs centered at 7. Finally, a green/purple swap at 5,7,9 and a red/green swap at 4,7 gives red/blue parallel pairs centered at 9. Thus, 15 is not blue. So 15 is purple, since it sees red and green and is not orange or blue. Thus, 16 is purple; otherwise a purple/blue swap at 7,8 gives a blue triple centered at 5. If 11 is red and 12 is green, then a green/blue swap at 8,9 gives a green triple centered at 6. So 11 is green and 12 is red. If 17 is orange, then we recolor 8 red and recolor 7 blue, which gives a blue triple centered at 5. So 17 is red.

(18-19) Now 18 or 19 is green; otherwise a green/orange swap at 9,10 gives green/orange parallel pairs centered at 7. And 18 or 19 is blue; otherwise we recolor 10 blue, which gives blue/red parallel pairs centered at 9. So 18 and 19 are blue and green (in some order). Now we use a purple/orange swap at 6, 7, 10 followed by a red orange swap at 4, 7. This gives a red triple centered at 9. □

Let a *pair* denote the template, with one color and two vertices, induced by vertices $\{1, 3\}$ in Figure 4.10. pair

Lemma 4.14. *Every 5-coloring of G is equivalent to a 5-coloring with a pair.*

Proof. Six vertices in a coloring form *alternating sets* if they are colored as vertices 1,2,3,4,5,6 in Figure 4.12 (not Figure 4.11), ignoring all other colors in that picture. We first show that each 5-coloring is equivalent to a 5-coloring that contains either a pair or alternating sets. Assume instead that φ is a counterexample to this statement. alternating sets

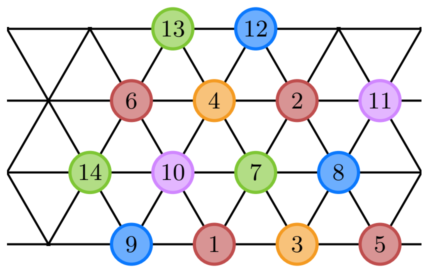

Figure 4.11: A claim in the proof of Lemma 4.14.

(1-6) Since we have no pair, we also have no triple. By Pigeonhole, two colors appear exactly twice on the neighborhood of 7. Since we have no pair, these repeated colors appear "across" from each other; that is, each color appears on two vertices with 7 as their unique common neighbor. By rotational symmetry, we assume 1 and 2 are red, and 3 and 4 are orange, as in Figure 4.11. If no neighbor of 3 is red, and no neighbor of 1 is orange, then an orange/red swap at 1,3 gives orange on both 1 and 4, a pair. Thus, either 5 is red or 9 is orange; by symmetry, we assume that 5 is red. By repeating the same argument for 2 and 4, we assume that 6 is red; if instead 11 is orange, then we have alternating sets, as claimed.

(7-14) Now 7 and 8 are new colors, green and blue, respectively. If 9 is orange, then we have alternating sets; and if 9 is green, then we have a green pair, 9 and 7. If 9 is purple, then we recolor 7 with purple and get a purple pair. So 9 is blue. This implies that 10 is purple. Clearly, 11 is not red or blue. It is also not green, since φ has no pair, and it is not orange, since φ does not have alternating sets. So 11 is purple. If 12 is green or purple, then φ contains a pair; so 12 is blue. Now 13 and 14 are both green; for each vertex, three colors appear on its neighbors and a fourth is forbidden, since φ contains no pair. Now we use a purple/orange swap at 10,4. This gives an orange pair at 10 and 3. This proves the claim that φ is equivalent to a 5-coloring that contains either a pair or alternating sets.

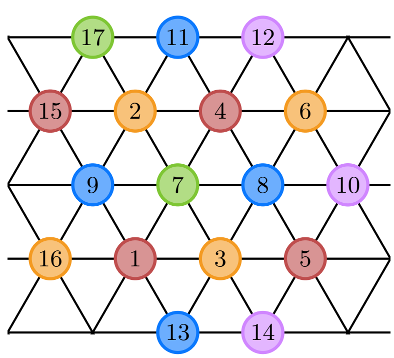

Figure 4.12: Finishing the proof of Lemma 4.14.

(1-10) Now we assume that φ contains alternating sets, with 1, 4, 5 red and with 2, 3, 6 orange, as in Figure 4.12. We also assume 7 is green and 8 is blue. (a) First suppose that 9 and 10 are both purple. At least one of 11 and 12 is blue or green. So either φ contains as a pair one of $\{7, 11\}$ or $\{8, 12\}$ or the resulting coloring contains one of these pairs after a green/blue swap at 7,8. So assume that 9 and 10 are not both purple. (b) Next suppose that 9 is blue and 10 is green. If none of 11, 12, 13, 14 is purple, then we recolor both 3 and 4 purple, which gives a pair. If instead one of 11, 12, 13, 14 is purple, then recoloring 7 or 8 purple gives a purple pair. Thus, we assume that exactly one of 9 and 10 is purple. (c) By reflectional symmetry, we assume that 9 is blue and 10 is purple.

(11-18) Since 7 is green, 11 is not green. Similarly, 11 is not purple, since then we can recolor 7 purple. So 11 is blue. If 12 is green, then we can recolor 7 purple and recolor 8 green. So 12 is purple. By reflectional symmetry, 13 is blue and also 14 is purple. Note that 15 is not green, since 7 is green. And 15 is not purple, since then we could recolor 7 purple to get a purple pair. So 15 is red. Similarly, 16 is orange. Next, 17 is green; otherwise, an orange/green swap at 2,7,3 gives the orange pair $\{6, 7\}$. Now recoloring 2 purple gives a purple pair at $\{2, 12\}$. □

Lemma 4.15. *If φ is a 5-coloring of G with a pair, then φ is equivalent to a coloring that contains either a triple, or parallel pairs, or crossing pairs. In particular, φ is nice.*

Proof. The second statement follows immediately from the first, combined with Lemmas 4.10, 4.11, and 4.13. So we now prove the first; instead assume that φ is a counterexample.

(1-8) By assumption, 1 and 2 have a common color, red; see Figure 4.13. Now 3 and 4 have new colors: purple and green, respectively. Further, 5 is also purple. Suppose instead that 5 is orange; now each of 7 and 8 is either purple or blue. If either is purple, then we have red/purple crossing pairs centered at 4. Otherwise, both 7 and 8 are blue, so we have red/blue parallel pairs centered at 4. Thus, we assume that 5 is purple; similarly, we assume that 6 is green. Below we frequently use this argument implicitly. Finally, 7 and 8 have new colors; if they are not distinct, then we have parallel pairs centered at 4. So assume 7 is orange and 8 is blue. Now we consider 4 cases, depending on the color of 9.

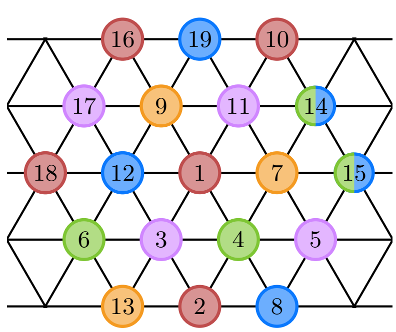

Figure 4.13: Case 1: 9 is orange.

Case 1: 9 is orange. (10-13) Since 7 and 9 are both orange, 10 is red, like 1, and 11 is purple, like 3. Now 12 is blue. And 13 is orange; otherwise, a purple/orange swap at 3 gives an orange triple centered at 1.

(14-16) One of 14 and 15 is blue; otherwise a blue/orange swap at 7 gives blue/red parallel pairs centered at 4. And one of of 14 and 15 is green; otherwise, an orange/green swap at 7,4 gives orange/red crossing pairs centered at 3. Note that 17 is not red, since 1 is red and

6 and 9 have distinct colors. So 16 is red; otherwise a red/orange swap at 9,1,7 gives a red triple centered at 11.

(17-19) Note that 18 cannot be purple, since 3 is purple, but 6 and 9 have distinct colors. So 17 is purple; otherwise, a blue/purple swap at 12,3 gives blue/red crossing pairs centered at 4. Further, 18 cannot be orange, since then a red/blue swap at 12,1 gives red/orange crossing pairs centered at 17. So 18 is red. If 19 is green, then an orange/blue swap at 9,12 gives orange/red parallel pairs centered at 3. So 19 is blue. But now we recolor 9 with green, and recolor 12 with orange. This also gives orange/red parallel pairs centered at 3.

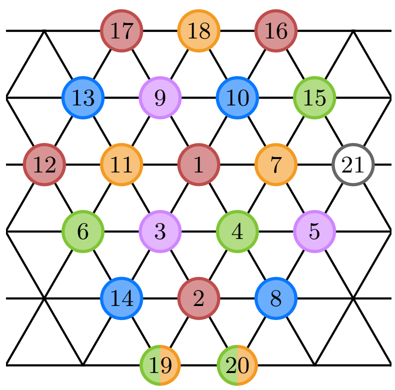

Figure 4.14: Case 2: 9 is purple.

Case 2: 9 is purple. (10-14) If 10 is green, then we have green/purple parallel pairs centered at 1, so 10 is blue. Since 3 and 9 are purple, 11 is orange and 12 is red. If 13 is green, then we have green/purple parallel pairs centered at 11; so 13 is blue. If 14 is orange, then we have orange/red parallel pairs centered at 3; so 14 is blue.

(15-16) Note that 21 is not green, since 4 is green but 5 and 10 are distinct colors. Now 15 is green; otherwise, a green/orange swap at 4,7 gives orange/purple crossing pairs centered at 1. Note that 18 is not red, since 1 is red but 9 and 15 are distinct colors. So 16 is red; otherwise a red/blue swap at 1,10 gives blue/purple crossing pairs centered at 11.

(17) Suppose 17 is not red. If neither 19 nor 20 is purple, then a red/purple swap at 2,3,1,9 gives a red triple centered at 11. If neither 19 nor 20 is green, then a green/red swap at 2,4,1 gives a green triple centered at 3. So 19 and 20 are purple and green. Now we recolor 2 with orange and do a green/red swap at 4,1, which gives green/purple crossing pairs centered at 11. So 17 is red.

(18-20) Now 18 is orange; otherwise an orange/purple swap at 9,11,3 gives an orange triple centered at 1. If neither 19 nor 20 is green, then a green/red swap at 2,4,1 gives a green triple centered at 3. If neither 19 nor 20 is orange, then we recolor 2 with orange and use a green/red swap at 4,1. This gives green/purple crossing pairs centered at 11. So 19 and 20 are green and orange. Now we recolor 9 with green. Afterwards, a red/purple swap at 1,3,2 gives a purple triple centered at 4.

Case 3: 9 is blue. (10-11) Now 10 is orange, since all other colors are used on its neighborhood. If 11 is orange, then we have orange/red parallel pairs centered at 3. So 11 is blue.

(12-13) If 12 is green, then we are actually in Case 2, by reflecting horizontally and interchanging colors green and purple. So we assume that 12 is not green; thus, 12 is purple. Similarly, 17 is not green, by reflecting both vertically and horizontally, and interchanging both orange/blue and green/purple. Now 13 is green; otherwise a green/red swap at 1,4,2 gives a green triple centered at 3.

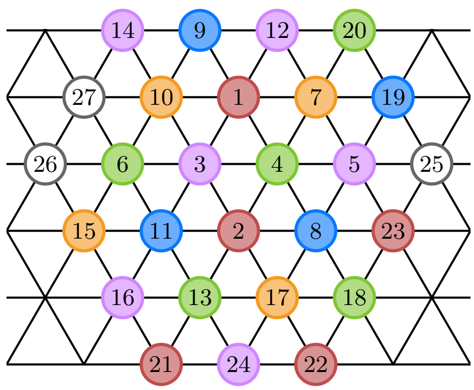

Figure 4.15: Case 3 in the proof of Lemma 4.15.

(14) Note that 27 is not purple, since 3 is purple but 6 and 9 have distinct colors. So, if 14 is not purple, then a purple/orange swap at 10,3 gives purple/orange parallel pairs, centered at 1. Thus, 14 is purple.

(15-16) Now 15 and 16 must be orange and purple. If neither is orange, then recoloring 11 with orange gives orange/red parallel pairs centered at 3. If neither is purple, then a purple/blue swap at 11,3 gives blue/red crossing pairs centered at 4. So 15 and 16 are orange and purple. Suppose that 15 is purple. If 26 is blue, then we have blue/purple crossing pairs centered at 6. And if 26 is orange, then we have orange/purple parallel pairs centered at 6. So 26 is red. This means that 27 is blue. Now a green/orange swap at 6,10 gives green/red crossing pairs centered at 3. So we assume 15 is not purple. Thus, 15 is orange and 16 is purple.

(17-18) If 17 is purple, then we recolor 2 with orange, recolor 11 with red, and use a red/green swap at 1,4. This gives red/green parallel pairs centered at 3. Thus, 17 is orange. Now 23 is not green, since 4 is green but 5 and 17 have distinct colors. So 18 is green; otherwise, a green/blue swap at 8,4 gives green/blue parallel pairs centered at 2.

(19-20) If neither 19 nor 20 is blue, then a blue/orange swap at 7 gives blue/red parallel pairs centered at 4. If neither 19 nor 20 is green, then a green/orange swap at 7,4 gives orange/red crossing pairs centered at 3. So 19 and 20 are blue and green. If 19 is green, then 23 is orange since 7 is orange. Thus, 25 is red; otherwise, we can recolor 5 red and get a red triple centered at 4. Now a blue/red swap at 8,2,11 gives red/green parallel pairs centered at 5. Thus, 19 is blue and 20 is green.

(21-24) Note that 24 cannot be red, since 2 is red and 16 and 17 have distinct colors. So 21 is red; otherwise, a red/green swap at 13,2,4,1 gives a green triple centered at 3. Similarly, 22 is red; otherwise, a red/orange swap at 2,17 and a red/green swap at 1,4 gives a red triple centered at 8. Now 23 must be red;

otherwise a red/blue swap at 8,2,11 gives 11 and 21 red but 16 and 17 with distinct colors, a contradiction. Finally, 24 is purple. If not, then we recolor 17 purple and repeat the argument above showing that 17 is orange, getting a contradiction. Now an orange/blue swap at 17,8 gives orange/red parallel pairs centered at 4.

Case 4: 9 is green. (9-11) We assume that 10 is green; otherwise reflecting vertically (interchanging orange/blue) reduces to an earlier case. Further, either 11 or 24 is purple. Otherwise, a purple/red swap at 3 gives a purple triple centered at 4. By possibly reflecting vertically, we assume 11 is purple.

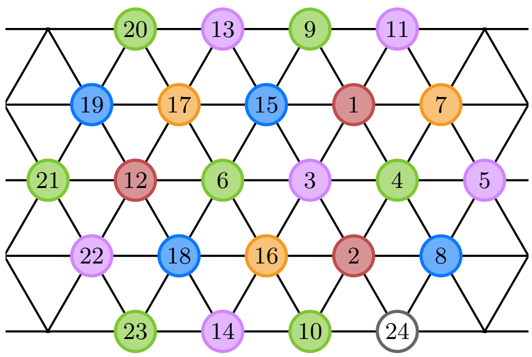

Figure 4.16: Case 4 in the proof of Lemma 4.15.

(12) Note that 15 and 16 must each be either orange or blue, and they must have distinct colors; otherwise, we have parallel pairs centered at 3. Denote the colors of 15 and 16 by α and β, respectively. If neither 13 nor 17 is purple, then an α/purple swap at 15,3 gives α/red crossing pairs centered at 4. And if neither 13 nor 17 is β, then an α/β swap at 15 gives β/red parallel pairs centered at 3. So 13 and 17 are β and purple (in some order). By a similar argument, 18 and 14 are α and purple (in some order). Thus, 12 is red; otherwise, we recolor 6 red, which gives a red triple centered at 3.

(13-14) If 17 is purple, then either 13 is orange and 15 is blue or vice versa. In the first case, a blue/red swap at 15,1 gives red/purple crossing pairs centered at 6. In the second case, we recolor 1 blue, which gives blue/purple parallel pairs centered at 15. So 13 is purple. Suppose 18 is purple. If 15 is blue, then a red/blue swap at 15,1 gives red/purple parallel pairs centered at 6. If 15 is orange, then we recolor 1 blue and recolor 15 red. Again, we get red/purple parallel pairs centered at 6. So 14 is purple.

(15-18) Suppose 15 is orange and 16 is blue. If 24 is purple, then recolor 1 blue, 2 orange, 15 red, and 16 red; this gives a red triple centered at 6. So 24 is orange. Recolor 1 blue. Use a red/purple swap at 2,3, and a red/orange swap at 3,15. This gives orange/purple parallel pairs centered at 16. So 15 is blue and 16 is orange. This implies that 17 is orange and 18 is blue.

(19-23) Note that 19 cannot be green, since 6 is green, but 12 and 13 have distinct colors. If 19 is purple, then a red/blue swap at 15,1 gives red/purple parallel pairs centered at 17. So 19 is blue. If 20 is red, then a red/blue swap at 1,15 gives a red triple centered at 17. So 20 is green. Note that 21 or 22 must be green; otherwise, a red/green swap at 12,6 gives a red triple centered at 3. And 22 cannot be green, since 6 is green but 12 and 14 have distinct colors. So 21 is green. If 22 is orange, then we recolor 12 purple and recolor 6 red, which gives a red triple centered at 3. So 22 is purple. Finally, 23 must be green; otherwise we have purple/red or purple/orange crossing pairs centered at 18. Now a blue/orange swap at 18,16 gives blue/orange parallel pairs centered at 6. □

Now we prove the main result of this section.

Lemma 4.16. *If G is a 6-regular toroidal graph with edge-width at least 7, then every 5-coloring of G is 5-equivalent to a 5-coloring that contains a good 4-template.*

Proof. Let G be a graph satisfying the hypothesis, and let φ_0 be a 5-coloring of G. By Lemma 4.14, φ_0 is 5-equivalent to some 5-coloring φ_1 that contains a pair. By Lemma 4.15, φ_1 is 5-equivalent to some 5-coloring φ_2 that contains either a triple, or parallel pairs, or crossing pairs. By Lemmas 4.10, 4.11, and 4.13. φ_2 is 5-equivalent to some 5-coloring φ_3 the contains a triple, and thus contains a good 4-template. □

4.3 Extensions and Open Questions

We begin the section by sketching the proof of Theorem 4.1. For reference, we repeat the statement below.

Theorem 4.1. *If G is a triangulated toroidal grid $T[a \times b]$ with $a \geq 6$ and $b \geq 6$, then all 5-colorings of G are 5-equivalent.*

Proof Sketch. We sketch how to modify the proof of the Main Theorem to prove Theorem 4.1. Recall that our Main Theorem handles the case that $a \geq 7$ and $b \geq 7$, since then G has edge-width at least 7. By symmetry, we can assume that $a \leq b$. So assume that $a = 6$. By Lemma 4.7, we know G is 4-colorable; so Lemma 4.3 still applies. In fact, the only place that our proof uses edge-width 7 is when we apply Lemma 4.5 to show that a subgraph H is locally connected. However, in the proof of Lemma 4.5 we only need the fact that H does not contain five vertices on a non-contractible cycle of length 6. The graph $T[6 \times 6]$ has exactly 18 non-contractible cycles of length 6. These run along the 6 rows, 6 columns, and 6 diagonals. And for $T[6 \times b]$ with $b \geq 7$, we have b non-contractible cycles of length 6; each runs along a column. So, to adapt the proof to our present situation, it suffices to check that H does not contain exactly 5 vertices on any of these non-contractible 6-cycles. In the proofs of Lemmas 4.10 and 4.12, this is true. However, in the proofs of Lemmas 4.11 and 4.13, there is sometimes a single non-contractible cycle C of length 6 with 5 of its vertices in H. In this case, we simply add the final vertex v of C to T and to H, making v the sole vertex in its color in T. It is straightforward to check that the subgraph $G - H - \{v\}$ is still connected. Thus, the proof of Lemma 4.5 still applies. □

In the introduction we mentioned that Mohar proved that if G is a planar graph with $\chi(G) = k$, then all $(k+1)$-colorings of G are $(k+1)$-equivalent. Mohar also constructed, for any positive integers k, ℓ with $k < \ell$ a k-chromatic graph with a single k-coloring and with two ℓ-colorings that are not ℓ-equivalent (Figure 2.5). We would still like to find larger classes of graphs G for which all $(\chi(G)+1)$-colorings are $(\chi(G)+1)$-equivalent. Our conjectures below focus on graphs embedded in the torus, and other surfaces. We first state a lemma that will likely be useful in studying this problem. We call a template T in a graph G *k-good* if G_T is k-degenerate.

Lemma 4.17. *Let G be a graph with $\chi(G) \leq k$. If φ_1 and φ_2 are $(k+1)$-colorings of G that contain monochromatic k-good templates T_1 and T_2, respectively, then φ_1 and φ_2 are $(k+1)$-equivalent.*

We omit the proof, since it is nearly identical to that of Lemma 4.3.

Conjecture 4.1. *If G is a triangulated toroidal grid $T(a \times b)$ (say with $a \geq 3$ and $b \geq 3$), then all 5-colorings of G are 5-equivalent.*

Conjecture 4.2. *If G is a 4-chromatic toroidal graph, then all 5-colorings of G are 5-equivalent.*

Conjecture 4.3. *If G is a toroidal graph, then all 5-colorings of G are 5-equivalent.*

Conjecture 4.4. *For every surface S there exists c_S such that if G embeds in S with edge-width at least c_S, then all 5-colorings of G are 5-equivalent.*

Chapter 5

The L-valid Kempe Swap Model: Kempe Equivalent List Colorings

Recall that Meyniel [53] proved that if G is a planar graph, then all of its 5-colorings are 5-equivalent; see also [30]. This was extended by Las Vergnas and Meyniel [50], who proved the same conclusion for all K_5-minor-free graphs[1]. They also showed that if G is $(k-1)$-degenerate, then all of its k-colorings are k-equivalent; Lemma 5.1 below (equivalently, Theorem 2.5 from Chapter 2) generalizes this result to list coloring. Mohar [54] conjectured that if G is connected and k-regular, then all of its k-colorings are k-equivalent. This is a natural next step, since the sparsest graphs that are not $(k-1)$-degenerate are k-regular. Mohar's Conjecture was proved for $k = 3$ by Feghali, Johnson, and Paulusma [40] (with a single exception $K_2 \square K_3$) and for $k \geq 4$ by Bonamy, Bousquet, Feghali, and Johnson [11]; this was reproved in a stronger form in [9]. Our Main Theorem in this chapter (stated on the next page) is an analogous result for list-coloring.

Remark 5.1. Recall from Observation 2.1 that we only consider list-assignments L for which G admits an L-coloring. Thus, the existence of an L-coloring for G should be implicitly assumed in the definition of L-swappability.

To illustrate some of the key ideas in this paper, we now reprove a helpful lemma of Las Vergnas and Meyniel [50], generalized to the context of list-coloring.

[1] An *H-minor* of a graph is a subgraph isomorphic to H obtained through a combination of vertex deletions, edge deletions, and edge contractions. A graph is *H-minor-free* if it contains no H-minor.

Lemma 5.1. *Let G be a connected graph, let L be a list-assignment for G, and fix $v \in V(G)$ with $|L(v)| > d(v)$. Let $G' := G - v$ and let L' denote L restricted to G'. If G' is L'-swappable, then G is L-swappable.*

Proof. Assume G' is L'-swappable. Let φ_1 and φ_2 denote L-colorings of G, and let φ_1' and φ_2' denote their restrictions to G'. Since G' is L'-swappable, there exist L'-colorings $\psi_0', \ldots, \psi_t'$ of G' such that $\psi_0' = \varphi_1'$ and $\psi_t' = \varphi_2'$ and ψ_i' differs from ψ_{i-1}' by a single L'-valid Kempe swap, for each $i \in [t]$. Now we extend each ψ_i' to an L-coloring ψ_i of G such that ψ_i and ψ_{i-1} are L-equivalent. Suppose that ψ_i' differs from ψ_{i-1}' by an α, β-swap at v_i, for some $v_i \in V(G)$ and some colors α and β. If $\psi_{i-1}(v) \notin \{\alpha, \beta\}$ or if v is not in the same α, β-component of ψ_{i-1} as v_i, then we form ψ_i from ψ_{i-1} by performing the same α, β-swap at v_i. This approach also works if $\{\alpha, \beta\} \subseteq L(v)$ and v is in the same α, β-component as v_i, but v has degree 1 in that component. So suppose that v is in the same α, β-component as v_i, but either $|L(v) \cap \{\alpha, \beta\}| = 1$ or v has degree at least 2 in that α, β-component. Since $|L(v)| > d(v)$, there exists $\gamma \in L(v)$ that is unused by ψ_{i-1} on the closed neighborhood of v. We first recolor v with γ, and then perform the α, β-swap at v_i. By induction on i, this gives an L-coloring ψ_t that restricts to ψ_t'. Lastly, if $\psi_t(v) \neq \varphi_2(v)$, then we recolor v with $\varphi_2(v)$. □

Corollary 5.1. *A graph G is L-swappable whenever there exists a vertex order in which each vertex x is preceded by fewer than $|L(x)|$ neighbors. In particular, G is L-swappable when G is $(k-1)$-degenerate and L is a k-assignment. This includes the special case that $k := \Delta$ and G is connected, but not regular.*

Proof. We prove the first statement by induction on $|V(G)|$. The base case $|V(G)| = 1$ holds trivially. The induction step holds by Lemma 5.1, taking v to be the final vertex in the order. The second statement follows from the first, using any order that witnesses that G is $(k-1)$-degenerate. Finally, the third statement obviously follows from the second, when we order the vertices by non-increasing distance from some vertex of degree less than k. □

Now we can state our Main Theorem and outline its proof.

Main Theorem. *If G is a connected graph with $\Delta \geq 3$ and $G \notin \{K_2 \square K_3, K_{\Delta+1}\}$, then G is Δ-swappable. If $G \in \{K_2 \square K_3, K_{\Delta+1}\}$ and L is a Δ-assignment that is not identical everywhere, then G is L-swappable.*

Note that the interesting case in our Main Theorem is when G is regular, since otherwise the result is included in Corollary 5.1. A crucial step in proving the Main Theorem is verifying the following Key Lemma. Its proof mirrors that of Lemma 5.1.

Key Lemma. *If H is degree-swappable and G is a connected graph containing H as an induced subgraph, then G is degree-swappable.*

In view of the Key Lemma, we have a natural plan to prove the Main Theorem. (i) Compile a collection $\mathcal{H}$ of known degree-swappable graphs. (ii) Show that if G is connected with $\Delta \geq 3$ and $G \notin \{K_2 \square K_3, K_{\Delta+1}\}$, then G contains as an induced subgraph some $H \in \mathcal{H}$. (For brevity, we omit from this sketch the details of handling $K_2 \square K_3$ and $K_{\Delta+1}$, but they are not hard.) In fact, Erdös, Rubin, and Taylor [37] used a similar approach to characterize all degree-choosable graphs. They showed that if G is connected and not a Gallai tree, then G contains as an induced subgraph an even length cycle with at most one chord[2], which we call a *good cycle*. It is easy to check that every good cycle is degree-choosable. A *cut set* $S \subset V(G)$ is such that $G - S$ is disconnected. A graph G is *k-connected* if $|V(G)| > k$ and every cut set S satisfies $|S| \geq k$. The *connectivity* of G is the largest integer k such that G is k-connected. A *block* in a graph is a maximal 2-connected subgraph. A *Gallai tree* is a connected graph in which each block is an odd cycle or a clique. These results on good cycles imply that every connected graph is degree-choosable unless it is a Gallai tree; this implication uses an analogue of our Key Lemma that holds for degree-choosability.[3] So perhaps we might hope to even characterize degree-swappable graphs. But when we take this approach, we quickly find many degree-choosable graphs that are not degree-swappable.

good cycle

k-connected

block

Gallai tree

Example 5.1. Denote the vertices of an n-cycle C_n by $v_1, \ldots, v_n$. Let $L(v_i) = \{i, i+1\} \bmod n$. See the left side of Figure 5.1. Note that C_n has two L-colorings: (1) $\varphi(v_i) = i$, for all v_i, and (2) $\varphi(v_i) = i+1$, for all v_i. However, $|L(v_i) \cap L(v_j)| \leq 1$ for all distinct i, j. Thus, neither φ_1 nor φ_2 admits any L-valid Kempe swap. So C_n is not degree-swappable.

More generally, let G be any Gallai tree. For each block B_i of G, let $d_i := d_{B_i}(v)$ for all $v \in V(B_i)$. Assign to each block B_i a list L_i of size d_i such that $L_{i_1} \cap L_{i_2} = \emptyset$ whenever $i_1 \neq i_2$. Let $L(v) := \cup_{B_i \ni v} L_i$. Now G is not L-choosable, as is easy to verify by induction on its number of blocks. (This is the standard construction showing that Gallai trees are not degree-choosable.) Form G' from G by adding some edge xy such that $x, y \in V(G)$ but $xy \notin E(G)$, x and y are in distinct blocks, and G' is not a Gallai tree. Let $L'(x) := L(x) \cup \{\alpha\}$, $L'(y) := L(y) \cup \{\alpha\}$, and $L'(z) := L(z)$ for all $z \in V(G) \setminus \{x, y\}$, and some $\alpha \notin \bigcup_{v \in V(G)} L(v)$. See the right side of Figure 5.1. Note that G' is L'-choosable (by the result of Erdös, Rubin, and Taylor [37], since G' is not a Gallai tree). Furthermore, G' has some L'-coloring φ_1 with $\varphi_1(x) = \alpha$ and has some other L'-coloring φ_2 with $\varphi_2(y) = \alpha$. Moreover, since G is not L-choosable, every L'-coloring of G' uses α on either x or y.

[2] The general case of this result easily reduces to the case when G is 2-connected, which is known as Rubin's Block Lemma. For a shorter proof, see Section 9 of [28].

[3] To prove this analogue, we greedily color the vertices of $G \setminus H$ in order of non-increasing distance from H. Afterward, we can extend this coloring to H precisely because H is degree-choosable.

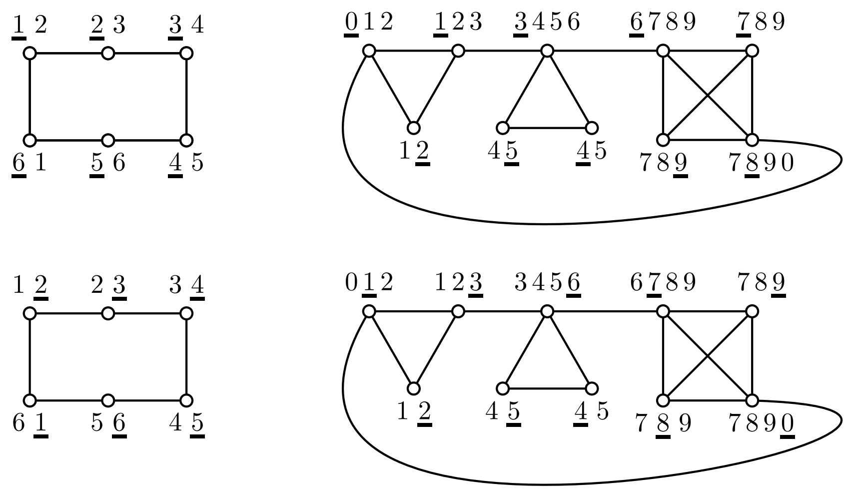

Figure 5.1: Left: A 6-cycle and a 2-assignment showing that it is not degree-swappable. Right: Another "Gallai tree plus edge" and a degree-assignment showing that it is not degree-swappable.

Note that $L'(x) \cap L'(y) = \{\alpha\}$, so no L-valid Kempe swap can move α from x to y, or vice versa. Thus, φ_1 and φ_2 are not L-equivalent. This implies that G' is not degree-swappable, even though, as noted above, G' is degree-choosable.

Note that Example 5.1 includes an even cycle with a single chord e whenever the two cycles containing e each have odd length (deleting any edge besides the chord gives a Gallai tree). Thus, many good cycles are not degree-swappable. In fact, we have discovered further graphs that are not degree-swappable[4], and we do not yet have a conjectured description of all such graphs. So we suggest the following problem.

Problem 5.1. Characterize all degree-swappable graphs.

To prove the Main Theorem, we split into two cases: (i) connectivity at most 3 and (ii) connectivity at least 4. In the first case, which takes most of the work, we use a small vertex cut to show that G contains an induced subgraph H from a family of known degree-swappable graphs. In the second case, the higher connectivity allows us to more explicitly construct a sequence of L-valid Kempe swaps to transform any L-coloring φ_1 into any other φ_2. The rest of the paper is organized as follows. In Section 5.1 we prove the Key Lemma, as well as a few other helpful results on swappability. In Section 5.2 we compile a family of known degree-swappable graphs. Finally, in Section 5.3 we prove the Main Theorem.

[4]Interestingly, all such graphs have connectivity 2. Moreover, in a follow-up paper [24], we show that every 4-connected graph is degree-swappable (in fact, more strongly, we show it is L-swappable for every list-assignment L satisfying $|L(v)| \geq d(v)$ for every $v \in V(G)$). Thus, we conjecture in [24] that every 3-connected graph is L-swappable for every such list-assignment L.

5.1 Swappability Lemmas

In this section, we prove a number of lemmas about swappability. More precisely, each lemma considers a graph G and a list assignment L and identifies a set of L-colorings of G that are pairwise L-equivalent. Each of these results can be viewed as extending Lemma 5.1.

Lemma 5.2. *(a) If $|L(v)| \geq d(v)$ for all v and $|L(w)| > d(w)$ for some w, then G is L-swappable (assuming that G is connected). (b) If $x \in V(G)$ and $G - x$ is connected, then the same result holds even if we only require $|L(x)| \geq 1$.*

Proof. The first statement follows directly from Corollary 5.1. We order the vertices by non-increasing distance from w; thus, every vertex other than w has a neighbor later in the order. For the second statement, we simply put x first in the order, and apply the previous result to $G - x$, which is still connected. □

Fix a graph G, a vertex $v \in V(G)$, a list assignment L for G, and a color $\alpha \in L(v)$. Let $\mathcal{L}$ denote the set of all L-colorings of G. Let $\mathcal{L}_{v,\alpha}$ denote the set of L-colorings φ such that $\varphi(v) = \alpha$. Recall that if $\mathcal{L}'$ is a set of L-colorings of G that are pairwise L-equivalent, then $\mathcal{L}'$ *mixes*. If also φ is an L-coloring that is L-equivalent to some $\varphi' \in \mathcal{L}'$, then we say that φ *mixes with* $\mathcal{L}'$; often it is the case that $\varphi \notin \mathcal{L}'$.

$\mathcal{L}$, $\mathcal{L}_{v,\alpha}$
mixes
mixes with $\mathcal{L}'$

Lemma 5.3. *Let G be a connected graph such that $w, x \in V(G)$, $wx \in E(G)$, and $G - x$ is connected. If L is a degree-assignment for G such that there exists $\alpha \in L(x) \setminus L(w)$, then $\mathcal{L}_{x,\alpha}$ is nonempty and mixes. More generally, $\cup_{\alpha \in L(x)\setminus L(w)} \mathcal{L}_{x,\alpha}$ mixes.*

Proof. We let $G' := G - wx$, let $L'(x) := L(x) \setminus L(w)$, and let $L'(v) := L(v)$ for all $v \in V(G) \setminus \{x\}$. Now we apply Lemma 5.2(b) to G' and L'. □

Lemma 5.4. *Let G be a graph with $v, w_1, w_2 \in V(G)$ such that $G - \{w_1, w_2\}$ is connected, $w_1, w_2 \in N(v)$, and $w_1w_2 \notin E(G)$. Fix a degree-assignment L for G.*

(1) If there exists $\alpha \in L(w_1) \cap L(w_2)$, then $\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}$ is nonempty and mixes.

(2) If there exist $\alpha \in L(w_1) \cap L(w_2)$ and $\beta \in L(w_1) \setminus L(v)$, then $(\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}) \cup \mathcal{L}_{w_1,\beta}$ is nonempty and mixes.

(3) If there exist $\alpha \in L(w_1) \cap L(w_2)$ and $\beta \in L(w_1) \setminus L(v)$ and also $N(w_1) = N(w_2)$, then $\bigcup_{\alpha \in L(w_1)\cap L(w_2)} (\mathcal{L}_{w_1,\alpha} \cup \mathcal{L}_{w_2,\alpha}) \cup \bigcup_{\beta \in L(w_1)\setminus L(v)} \mathcal{L}_{w_1,\beta}$ is nonempty and mixes.

Proof. To prove (1), let $G' := G - \{w_1, w_2\}$, let $L'(y) := L(y) \setminus \{\alpha\}$ for all $y \in N(w_1) \cup N(w_2)$, and let $L'(z) := L(z)$ for all other $z \in V(G)$. Note that $|L'(v)| \geq |L(v)| - 1 > d_G(v) - 2 = d_{G'}(v)$. Thus, we can apply Lemma 5.2(a) to G' and L'.

Now we prove (2). If $\alpha = \beta$, this holds by Lemma 5.3 (with $x := w_1$ and $w := v$). So assume $\alpha \neq \beta$. Let $G' := G - vw_1 - w_2$ and let $L'(w_1) := \{\alpha, \beta\}$ and $L'(z) := L(z) \setminus \{\alpha\}$ for all $z \in N(w_2)$ and $L'(z) := L(z)$ otherwise. Now L' mixes for G', by Lemma 5.2(a), with $w := v$. These L'-colorings of G' are in bijection with colorings of G in $\mathcal{L}_{w_2,\alpha} \cap (\mathcal{L}_{w_1,\alpha} \cup \mathcal{L}_{w_1,\beta})$, and each L'-valid Kempe swap in G' maps to an L-valid Kempe swap in G that respects this bijection. Since $\mathcal{L}_{w_1,\beta}$ mixes by Lemma 5.3 and $\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}$ mixes by (1), the result follows. (Note that $L_{w_1,\beta} \cap L_{w_2,\alpha} \neq \emptyset$, since we can color w_1 with β and color w_2 with α, and then color $G - \{w_1, w_2\}$ greedily towards v.)

Finally, we prove (3). Consider $\varphi \in \mathcal{L}_{w_1,\alpha} \cup \mathcal{L}_{w_2,\alpha}$. If $\varphi(w_1) \neq \varphi(w_2)$, then we simply recolor w_1 or w_2 so that they both use color α; this is possible because $N(w_1) = N(w_2)$, so α is unused on $N(w_1)$ (and on $N(w_2)$). Thus, φ mixes with $(\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}) \cup \mathcal{L}_{w_1,\beta}$ by (2). Finally, if there exist distinct $\alpha, \alpha' \in L(w_1) \cap L(w_2)$, then we repeat the argument above with α' in place of α. Similarly, if there exist distinct $\beta, \beta' \in L(w_1) \setminus L(v)$, then we repeat the argument above with β' in place of β. This proves (3). □

We will often want to prove that a graph G is L-swappable, for some list assignment L. When we want to prove that a graph is degree-swappable, the following lemma significantly restricts the possibilities for L that we must consider.

Lemma 5.5. *Fix a graph G, a degree-assignment L, and an edge vw such that $G - vw$ is connected and degree-choosable. If $|L(v) \cap L(w)| \leq 1$, then G is L-swappable.*

Proof. If $|L(v) \cap L(w)| = 0$, then G is L-swappable if and only if $G - vw$ is L-swappable, and the latter statement holds by Lemma 5.2(a). So assume instead that $|L(v) \cap L(w)| = 1$ and that $L(v) \cap L(w) = \{\alpha\}$. Form L_1 from L by removing α from $L(v)$. Form L_2 from L by removing α from $L(w)$. Form L_3 from L by removing α from both $L(v)$ and $L(w)$. Note that G is L_1-swappable if and only if $G - vw$ is L_1-swappable, and the latter is true by Lemma 5.2(a). The same is true for L_2-swappable. Since L_3 is a degree-assignment for $G - vw$, by assumption $G - vw$ has an L_3-coloring φ. Note that φ is both an L_1-coloring and an L_2-coloring. Thus, L_1-colorings mix with L_2-colorings. Since every L-coloring of G is either an L_1-coloring or an L_2-coloring (or both), we conclude that G is L-swappable. □

The rest of this section is dedicated to proving the Key Lemma (from the introduction). In fact, we prove a more general version, Lemma 5.6.

Lemma 5.6. *Fix a graph G and a function $f : V(G) \to \mathbb{Z}^+$. Let H be an induced subgraph of G such that $G - H$ is f-swappable. Let $f'(x) := f(x) - (d_G(x) - d_H(x))$ for all $x \in V(H)$. If $f'(x) \geq d_H(x)$ for all $x \in V(H)$ and H is f'-swappable, then G is f-swappable.*

To see that Lemma 5.6 generalizes the Key Lemma, let $f(v) := d_G(v)$ for all $v \in V(G)$, and note that $f'(v) = d_H(v)$ for all $v \in V(H)$ and $G - H$ is f-swappable by Lemma 5.2(a). (Each component of $G - H$ has a vertex w with a neighbor in H.)

In this paper we only need the version of the lemma from the introduction (when $f(v) = d(v)$), but the more general version is no harder to prove, and it is useful elsewhere [26]. The proof of Lemma 5.6 mirrors that of Lemma 5.1. We start with L-colorings φ'_1 and φ'_2 of $G \setminus H$ and a sequence of L-colorings $\psi'_0, \ldots, \psi'_t$, showing that φ'_1 and φ'_2 are L-equivalent. And we seek to extend these L-colorings of $G \setminus H$ to L-colorings of G. Our main obstacle is the possibility that at some step colors used on H might interfere with our desired Kempe swap. So first we prove, for every L-coloring φ' of $G \setminus H$ and every L-valid Kempe swap for φ', that some extension of φ' to an L-coloring of G does not interfere. This notion of "non-interference" motivates Definition 5.1 and Lemma 5.7.

Definition 5.1. For a graph G and a list assignment L for G, an L-coloring φ of G is *(α, β)-versatile at w* if $\varphi(w) \in \{\alpha, \beta\}$ and an (α, β)-swap at w is L-valid for φ. (α, β)-versatile

Lemma 5.7. *Fix a graph G, a connected subgraph H, and a list assignment L for $V(G)$. Let φ' be an L-coloring for $G - H$ that is (α, β)-versatile at a vertex w. If $|L(v)| \geq d_G(v)$ for all $v \in V(H)$, and H is not a Gallai tree, then there exists an L-coloring φ of G that extends φ' such that φ is (α, β)-versatile at w. Further, there exists such an L-coloring φ with the property that each (α, β)-component of φ contains the vertex set of at most one (α, β)-component of φ'.*

Proof. Since H is not a Gallai tree, it contains an induced even cycle, C, with at most one chord. We first show how to extend φ' to $G - C$, and then how to extend it to all of G.

A key step is to show that if $|V(H)| \geq 2$ and $x \in V(H)$, then there exists an L-coloring φ of $G - (H - x)$ that extends φ' to x and that is (α, β)-versatile at w. Suppose that $|V(H)| \geq 2$ and fix $x \in V(H)$. When choosing a color for x (for brevity, we denote it by $\varphi'(x)$), to ensure that the resulting extension of φ' is an L-coloring that is (α, β)-versatile at w, we need to check the following three properties: (a) $\varphi'(x) \neq \varphi'(y)$ for all $y \in N(x) \setminus H$, (b) if $\varphi'(x) \in \{\alpha, \beta\}$ and x has a neighbor $y \notin H$ with $\varphi'(y) \in \{\alpha, \beta\}$, then $\{\alpha, \beta\} \subseteq L(x)$, and (c) if $\varphi'(x) \in \{\alpha, \beta\}$, then x has at most one neighbor $y \notin H$ with $\varphi'(y) \in \{\alpha, \beta\}$. Now (a) ensures the extension is a proper L-coloring; (b) ensures that an (α, β)-swap at w will not create a problem at x; and (c) ensures that each (α, β)-component of φ contains at most one (α, β)-component of φ' and that an (α, β)-swap at w will not create a problem at y.

We form a list assignment $L'(x)$ from $L(x)$ by first removing each color that is used by φ' on a neighbor of x. Further, if α is used on a neighbor of x and $\alpha \notin L(x)$, then we remove β from $L(x)$. Similarly, if α is used on two neighbors of x, then we remove β from $L(x)$, regardless of whether or not $\alpha \in L(x)$. We also remove α from $L(x)$ if either of these situations occurs, but with β and α interchanged. Since $|L(x)| \geq d_G(x)$, we must have $|L'(x)| \geq d_H(x) \geq 1$, because H is connected and $|V(H)| \geq 2$. To extend φ' to x, we simply choose any color in $L'(x)$. This completes the key step, started in the previous paragraph. To extend φ' to $G - C$, we repeatedly apply the key step, coloring vertices in order of non-increasing distance from C. Now we show how to extend φ' to C.

Case 1: C has no chord. For each $v \in V(C)$, form $L'(v)$ as in the previous paragraph. Again, $|L'(v)| \geq 2$ for all $v \in V(C)$. First suppose there exists γ and $x, y \in V(C)$ such that $\gamma \notin \{\alpha, \beta\}$ and $xy \in E(C)$ and $\gamma \in L'(x) \setminus L'(y)$. Now we color x with γ and proceed around C finishing with y. Each time that we color a vertex on C, we ensure that its degree in the subgraph induced by vertices colored α and β (among all colored vertices, both inside and outside C) is at most 1. This ensures that each (α, β)-component of φ contains the vertex set of at most one (α, β)-component of φ'. This process succeeds because each time we color another vertex z_1 we reduce the number of allowable colors on its uncolored neighbor z_2 by at most one (even if we completely repeat the process of removing colors as in the previous paragraph, treating z_1 as though it is outside H); this is virtually the same as extending φ' to $G - C$. We can finish at y because the color on x does not restrict our choice of color for y.

Suppose instead there exists $\gamma \notin \{\alpha, \beta\}$ such that $\gamma \in L'(v)$ for all $v \in C$. Now we use γ on one maximum independent set in C and color each remaining vertex v from $L'(v) \setminus \{\gamma\}$.

Finally, suppose that $L'(v) = \{\alpha, \beta\}$ for each $v \in V(C)$. Now we alternate α and β around C. In each case, it is easy to check that the resulting coloring φ is (α, β)-versatile at w. The key observation for this last case is that no vertex of C has a neighbor in $G - C$ colored α or β, by construction of $L'(v)$.

Case 2: C has a chord. Let x denote one endpoint of the chord and let y and z denote the neighbors of x on C, besides the other endpoint of the chord. Form $L'(v)$ for each $v \in V(C)$, as above; again $|L'(v)| \geq d_H(v)$ for all $v \in V(C)$. Since $|L'(x)| \geq d_H(x) = 3$, there exists $\gamma \in L'(x) \setminus \{\alpha, \beta\}$. If $\gamma \in L'(y) \cap L'(z)$, then use γ on y and z and color greedily toward x in the remaining uncolored subgraph. So assume instead that $\gamma \notin L'(y) \cap L'(z)$; by symmetry, assume that $\gamma \notin L'(z)$. Now use γ on x, then color the remaining uncolored subgraph greedily in order of non-increasing distance from z. Again, we can finish at z because using γ on x does not restrict the choice of color for z. As above, each time that we color a vertex on C, we ensure that its degree in the subgraph induced by vertices colored α and β (both inside and outside C) is at most 1. Again, this ensures that each (α, β)-component of φ contains the vertex set of at most one (α, β)-component of φ'. $\square$

Now we prove Lemma 5.6. As mentioned above, the proof mirrors that of Lemma 5.1, but now Lemma 5.7 ensures some extension to H is versatile for the next Kempe swap in $G - H$.

Proof of Lemma 5.6. Fix a graph G and a function $f : V(G) \to \mathbb{Z}^+$. Let H be an induced subgraph of G such that $G - H$ is f-swappable. Let $f'(x) := f(x) - (d_G(x) - d_H(x))$ for all $x \in V(H)$. Assume that $f'(x) \geq d_H(x)$ for all $x \in V(H)$ and that H is f'-swappable.

Let L be an f-assignment for G. Let $G' := G - H$. By assumption, each component of G' is L-swappable. So G' is L-swappable. Let φ_0 and φ be two L-colorings of G, and let φ'_0 and φ' denote their restrictions to G'. Since G' is L-swappable, there exists a sequence $\varphi'_0, \varphi'_1, \ldots, \varphi'_k = \varphi'$ of L-colorings of G' such that every two successive L-colorings differ by a single L-valid Kempe swap. By induction on k, we extend each φ'_i to an L-coloring φ_i of G such that every two successive L-colorings in the sequence $\varphi_0, \varphi_1, \ldots, \varphi_k = \varphi$ are L-equivalent. The case $k = 0$ is easy because $\varphi_0 = \varphi$, so we are done.

So assume that $k \geq 1$. Suppose that φ'_{i+1} differs from φ'_i by an α, β-swap at a vertex v_i. By Lemma 5.7, there exists an L-coloring $\widetilde{\varphi}_i$ of G, such that the restriction of $\widetilde{\varphi}_i$ to G' is φ'_i and an α, β-swap at v_i is L-valid in $\widetilde{\varphi}_i$. Furthermore, the restriction of this new coloring (after performing the α, β-swap at v_i) is φ'_{i+1}. It now suffices to show that φ_i and $\widetilde{\varphi}_i$ are L-equivalent. We do this by a sequence of Kempe swaps that recolors H but never changes the colors on $V(G - H)$.

For each $v \in V(G - H)$, remove $\varphi'_i(v)$ from $L(w)$ for each $w \in N(v) \cap V(H)$; denote the resulting list assignment on H by L_H. Note that $|L_H(x)| \geq f'(x) \geq d_H(x)$ for all $x \in V(H)$. If $|L_H(x)| > f'(x)$ for some x, then H is L_H-swappable by Corollary 5.1, since $|L_H(y)| \geq f'(y) \geq d_H(y)$ for all $y \in H$. So assume $|L_H(y)| = f'(y)$ for all $y \in H$. Since H is f'-swappable, the restrictions of φ_i and $\widetilde{\varphi}_i$ to H (which are both L_H-colorings) are L_H-equivalent. Consider a sequence of Kempe swaps that witnesses this. Note that performing the same Kempe swaps in G transforms φ_i to $\widetilde{\varphi}_i$ (this is because for each edge vw with $v \in V(H)$

and $w \notin V(H)$, we have $\varphi_i'(w) \notin L_H(v)$). Now performing an α, β-swap at v_i in $\widetilde{\varphi}_i$ yields an L-coloring of G that restricts to φ_{i+1}' on H; denote this L-coloring of G by φ_{i+1}.

The previous two paragraphs show that we can use L-valid Kempe swaps to transform φ_0 into an L-coloring $\widetilde{\varphi}$ that agrees with φ on $G - H$. Finally, we transform $\widetilde{\varphi}$ to φ. This is possible precisely because H is f'-swappable. □

5.2 Degree-swappable Graphs

In this section we prove that various graphs are degree-swappable. In view of Lemma 5.6, if a connected k-regular graph G contains an induced copy of any degree-swappable graph, then G is k-swappable. Our first example of degree-swappable graphs requires a new definition. A *theta graph*, $\Theta_{a,b,c}$, consists of two 3-vertices, x and y, that are linked by internally disjoint paths of lengths a, b, and c. theta graph

Lemma 5.8. *If G is a bipartite theta graph, $\Theta_{a,b,c}$, then G is degree-swappable.*

Proof. Fix a degree-assignment L for G. Note that $G - vw$ is degree-choosable for all $vw \in E(G)$. So, by Lemma 5.5, $|L(v) \cap L(w)| \geq 2$ for all $vw \in E(G)$. Let x and y denote the 3-vertices of G, and let P_a, P_b, and P_c denote the internally disjoint x, y-paths of lengths a, b, and c, respectively. Since G has an x, y-path with all internal vertices of degree 2, by Lemma 5.5 we can assume, by transitivity, that $|L(x) \cap L(y)| \geq 2$.

Suppose $L(x) \neq L(y)$. Specifically, suppose $L(x) = \{1, 2, 3\}$ and $L(y) = \{1, 2, 4\}$. This implies that every 2-vertex has the list $\{1, 2\}$. Fix an arbitrary L-coloring φ of G with $\varphi(x) = 3$ and $\varphi(y) = 4$. Starting from an arbitrary L-coloring φ', we can recolor x with 3 and recolor y with 4, then reach φ by using at most one 1, 2-swap on each path of $G - \{x, y\}$. Thus, we assume that $L(x) = L(y) = \{1, 2, 3\}$.

Suppose two disjoint x, y-paths, say P_a and P_b, have the same list for their 2-vertices. Form G' from G by deleting all 2-vertices of P_a. Note that $d_{G'}(x) = 2 < 3 = |L(x)|$, so G' is L-swappable by Lemma 5.2(a) with $w := x$. This implies that G is also L-swappable, as follows. Since P_a and P_b have lengths of the same parity, given any L-coloring of G, we can recolor the internal vertices of P_a so the neighbors of x and y on P_a have the same colors as their neighbors on P_b. Thus, any sequence of Kempe swaps in G' extends to a sequence in G. So we assume that no two x, y-paths have the same list for their 2-vertices.

Denote by x_a, x_b, and x_c the 2-neighbors of x, if they exist (when one x, y-path has no 2-vertices, assume it is P_c, and let x_c denote y). Assume $L(x_a) = \{1, 2\}$, $L(x_b) = \{1, 3\}$, and $L(x_c) = \{2, 3\}$. Note that the sets

$\mathcal{L}_{x,1}, \mathcal{L}_{x,2}$, and $\mathcal{L}_{x,3}$ each mix by Lemma 5.3.

Consider $\varphi \in \mathcal{L}_{x_a,1} \cap \mathcal{L}_{x_b,1}$. Note that the 2,3-component containing x is a path (also containing all 2-vertices of P_c and possibly y, depending on the parities of the x, y-paths). So a 2,3-swap at x shows that $\mathcal{L}_{x,2}$ mixes with $\mathcal{L}_{x,3}$. By symmetry, $\mathcal{L}_{x,1}$ mixes with $\mathcal{L}_{x,2}$. Thus, $\mathcal{L}_{x,1} \cup \mathcal{L}_{x,2} \cup \mathcal{L}_{x,3}$ mixes; that is, $\mathcal{L}$ mixes. □

Recall that an even length cycle with at most one chord is a *good cycle*.

Lemma 5.9. *Let G be a graph that contains as induced subgraphs two good cycles, H_1 and H_2. Assume that H_1 and H_2 intersect in at most one vertex and that all but at most one edge induced by $V(H_1) \cup V(H_2)$ either lies in H_1 or lies in H_2; further, if there exists such an edge, then $V(H_1) \cap V(H_2) = \emptyset$. If P is a shortest path from H_1 to H_2, then $G[V(H_1) \cup V(H_2) \cup V(P)]$ is degree-swappable. Thus, if G is a connected graph with at least two degree-choosable blocks, then G is degree-swappable.*

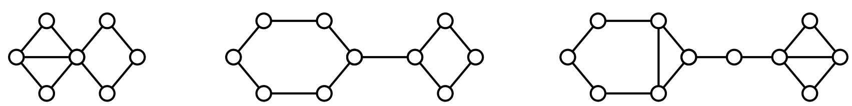

Figure 5.2: Three examples of $G[V(H_1) \cup V(H_2) \cup V(P)]$ in Lemma 5.9.

Proof. To begin, we prove the second statement from the first. Let B_1 and B_2 be two degree-choosable blocks of G. So neither B_1 nor B_2 is a complete graph or odd cycle. By Rubin's Block Lemma [37], there exist induced even cycles H_1 and H_2, each with at most one chord, in B_1 and B_2, respectively. Let P be a shortest path from H_1 to H_2. By the first statement and the Key Lemma, G is degree-swappable.

Now we prove the first statement. If possible, choose H_1 so that it is an even cycle with at most one chord that is closest to H_2. By Lemma 5.6, it suffices to consider the case that $V(G) = V(H_1) \cup V(H_2) \cup V(P)$. Let L be a degree-assignment for G. Let u be the endpoint of P in H_1. For each path Q, let $\ell(Q)$ denote the length (number of edges) of Q. We often write the *parity of* Q to mean the parity of $\ell(Q)$. If $\ell(P) \neq 0$, let v be the neighbor of u on P. Suppose $N_{H_1}(v) = \{w_1, \ldots, w_t\}$ with $t \geq 3$. Let P_i be the w_i, w_{i+1}-path in H_1 for every $i \in [t]$. If P_i is even for some $i \in [t]$, then $G[v \cup V(P_i)]$ is an even cycle closer to H_2, contradicting our choice of H_1. So assume P_i is odd for every $i \in [t]$. Since H_1 is even, t must also be even; thus, $t \geq 4$. But now $G[v \cup V(P_1) \cup V(P_2)]$ is an even cycle closer to H_2. Thus, when $\ell(P) \neq 0$, we assume $|N_{H_1}(v)| \leq 2$.

$\ell(Q)$

parity of Q

Also, since $G - xy$ is degree-choosable for every $xy \in E(H_1)$, we assume $|L(x) \cap L(y)| \geq 2$ by Lemma 5.5. Thus, if $w, z \in V(H_1)$ and some w, z-walk consists only of 2-vertices (including w and z), then transitivity implies $L(w) = L(z)$.

Case 1: $\ell(P) = 0$ or $|N_{H_1}(v)| = 1$.

Case 1.1: H_1 has no chord. Pick $u_1, u_2 \in N_{H_1}(u)$ and let $L(u_1) = \{a, b\}$; see Figure 5.3a. As observed above, $L(u_2) = L(u_1) = L(u_i)$ for all $u_i \in V(H_1) \setminus \{u\}$, by transitivity. Since H_1 is even, $\varphi(u_1) = \varphi(u_2)$ for every $\varphi \in \mathcal{L}$. By Corollary 5.1, $\mathcal{L}_{u_1,a}$ and $\mathcal{L}_{u_1,b}$ each mix. To see this, order the vertices of $G - H_1$ by non-increasing distance from u (with u last), and recall that always $\varphi(u_1) = \varphi(u_2)$. Now we show that $\mathcal{L}_{u_1,a} \cup \mathcal{L}_{u_1,b}$ mixes. Since H_2 is degree-choosable and $|L(u)| = d(u) = 3$, there exists an L-coloring φ' with $\varphi'(u) \notin \{a, b\}$. Performing an a, b-swap at u_1 in φ' shows that $\mathcal{L}_{u_1,a}$ mixes with $\mathcal{L}_{u_1,b}$. Thus, $\mathcal{L}$ mixes.

Case 1.2: H_1 has a chord xy with $u \notin \{x, y\}$. Let P_1, P_2, and P_3 be the x, y-path, x, u-path, and y, u-path on H_1 avoiding u, y, and x, respectively; see Figure 5.3b. If $\ell(P_1)$ is odd, then $\ell(P_2 \cup P_3)$ is also odd since H_1 is even. Thus, H_1 induces a bipartite theta graph, and we are done by Lemma 5.8. So we instead assume $\ell(P_1)$ is even.

Assume $L(w) = \{a, b\}$ for every 2-vertex w on P_1. Recall that $\{a, b\} \subseteq L(x) \cap L(y)$. Let $L(x) = \{a, b, c\}$. Suppose $L(y) \neq L(x)$; so assume $L(y) = \{a, b, d\}$. By Lemma 5.3, both $\mathcal{L}_{x,c}$ and $\mathcal{L}_{y,d}$ mix (and are nonempty). Further, $\mathcal{L}_{x,c} \cap \mathcal{L}_{y,d} \neq \emptyset$, since H_2 is degree-choosable; thus, $\mathcal{L}_{x,c} \cup \mathcal{L}_{y,d}$ mixes. Since $G[V(P_1)]$ is an odd cycle, $\mathcal{L} = \mathcal{L}_{x,c} \cup \mathcal{L}_{y,d}$, and we are done. So we instead assume that $L(x) = L(y) = \{a, b, c\}$.

Since H_1 and $\ell(P_1)$ are both even, $\ell(P_2)$ and $\ell(P_3)$ have the same parity. By Lemma 5.3, $\mathcal{L}_{x,c}$ and $\mathcal{L}_{y,c}$ each mix (and are nonempty). Further, $\mathcal{L} = \mathcal{L}_{x,c} \cup \mathcal{L}_{y,c}$. So it suffices to show that $\mathcal{L}_{x,c}$ mixes with $\mathcal{L}_{y,c}$. Pick $x_1 \in N_{P_2}(x)$ and $y_1 \in N_{P_3}(y)$, if they exist; again, see Figure 5.3b.

Suppose first that neither P_2 nor P_3 has any internal vertex. Since $|L(u) \cap L(x)| \geq 2$, assume that $a \in L(u)$. Since $\ell(P_1)$ is even, we can color $V(H_1)$ so that all neighbors of x and y use color a. Since H_2 is degree-choosable, we can extend this coloring to an L-coloring φ of G. Now performing a b, c-swap at x shows that $\mathcal{L}_{x,c}$ mixes with $\mathcal{L}_{y,c}$, and we are done.

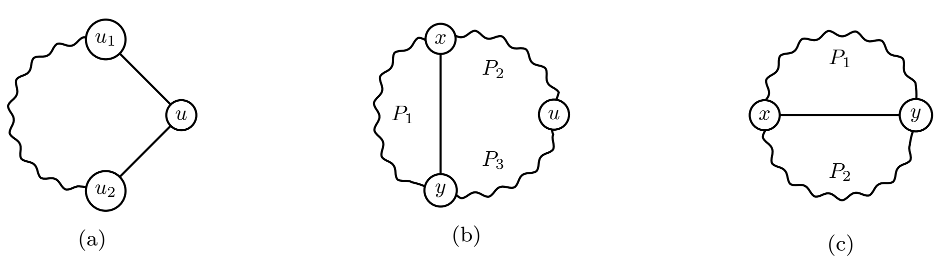

Figure 5.3: The 3 subcases in Case 1. (a) Case 1.1: H_1 has no chord. (b) Case 1.2: H_1 has a chord xy. (c) Case 1.3: H_1 has a chord uy (i.e. $x = u$).

Assume instead that P_2 or P_3 contains an internal 2-vertex. By symmetry, assume x_1 exists. Let w_1 be a 2-vertex adjacent to x on P_1. Recall that $L(w_1) = \{a, b\}$. As above, we assume $a \in L(x_1)$. Now $\mathcal{L}_{w_1,a} \cap \mathcal{L}_{x_1,a}$ mixes by Lemma 5.4. To construct an L-coloring in this set, we can color x arbitrarily from $L(x) \setminus \{a\}$ and color greedily toward H_2. Thus, $\mathcal{L}_{w_1,a} \cap \mathcal{L}_{x_1,a}$ contains φ_1 and φ_2 such that $\varphi_1(x) = c$ and such that $\varphi_2(x) = b$, so $\varphi_2(y) = 2$. This again proves that $\mathcal{L}_{x,c}$ mixes with $\mathcal{L}_{y,c}$, so we are done.

Case 1.3: H_1 has a chord xy with $x = u$ (by symmetry). Let P_1 and P_2 be the u, y-paths in H_1; see Figure 5.3c. Since H_1 is even, $\ell(P_1)$ and $\ell(P_2)$ have the same parity. If $\ell(P_1)$ is odd, then H_1 induces a bipartite theta graph, and we are done by Lemma 5.8. So $\ell(P_1)$ is even. Pick $y_1 \in N_{P_1}(y)$ and $y_2 \in N_{P_2}(y)$. By symmetry, assume $L(y) = \{a, b, c\}$ and $L(y_1) = \{a, b\}$. By Lemma 5.3, $\mathcal{L}_{y,c}$ is nonempty and mixes. By Lemma 5.3, $\mathcal{L}_{y,c}$ mixes.

Suppose $L(y_2) = L(y_1) = \{a, b\}$. Pick $\varphi \in \mathcal{L}_{y,a}$. We show that φ is L-equivalent to some $\varphi' \in \mathcal{L}_{y,c}$. If $\varphi(x) \neq c$, then recoloring y with c gives a coloring in $\mathcal{L}_{y,c}$. So assume $\varphi(x) = c$. If $\varphi(w) = b$ for some $w \in N_{G-E(H_1)}(x)$, then there exists $\alpha \in L(x)$ which does not appear on $N[x]$. We recolor x with α then recolor y with c to get a coloring in $\mathcal{L}_{y,c}$. So assume $\varphi(w) \neq b$ for every $w \in N_{G-E(H_1)}(x)$. Now we perform an a, b-swap at y followed by a b, c-swap at y to get a coloring in $\mathcal{L}_{y,c}$. The same argument shows that every $\varphi \in \mathcal{L}_{y,b}$ is L-equivalent to some coloring in $\mathcal{L}_{y,c}$. Since $\mathcal{L}_{y,c}$ mixes, also $\mathcal{L}_{y,a} \cup \mathcal{L}_{y,b} \cup \mathcal{L}_{y,c}$ mixes; that is, $\mathcal{L}$ mixes.

So we assume $L(x_1) \neq L(y_1)$. By symmetry, assume $L(y_1) = \{a, c\}$. By Lemma 5.3, both $\mathcal{L}_{y,b}$ and $\mathcal{L}_{y,c}$ mix (and are nonempty). Further, $\mathcal{L}_{y,a}$ is nonempty. And for every $\varphi \in \mathcal{L}_{y,a}$, we can perform an a, b-swap (resp. a, c-swap) at y to get a coloring in $\mathcal{L}_{y,b}$ (resp. $\mathcal{L}_{y,c}$). Thus, $\mathcal{L}_{y,a} \cup \mathcal{L}_{y,b} \cup \mathcal{L}_{y,c}$ mixes; that is, $\mathcal{L}$ mixes.

Case 2: $|N_{H_1}(v)| = 2$. Let $N_{H_1}(v) = \{v_1, v_2\}$; denote the v_1, v_2-paths in H_1 by P_1 and P_2. Recall, from the start of the proof, that P_1 and P_2 are both odd.

Case 2.1: H_1 has no chord. By symmetry, assume $\ell(P_2) > 1$. Pick $w_1 \in N_{P_2}(v_1)$ and $w_2 \in N_{P_2}(v_2)$; see Figure 5.4a. Assume $L(v_1) = \{a, b, c\}$ and $L(w_1) = \{a, b\}$, by symmetry. Suppose $L(v_1) \neq L(v_2)$; specifically, suppose $L(v_2) = \{a, b, d\}$. By Lemma 5.3, both $\mathcal{L}_{v_1,c}$ and $\mathcal{L}_{v_2,d}$ mix. By Lemma 5.5, if $\ell(P_1) > 1$, then $L(w) = \{a, b\}$ for every $w \in V(P_1)$. Thus, $\mathcal{L}_{v_1,c} \cap \mathcal{L}_{v_2,d} \neq \emptyset$. Further, for every $\varphi \notin \mathcal{L}_{v_1,c} \cup \mathcal{L}_{v_2,d}$, there exists $\gamma \in (\{c, d\} - \varphi(v))$. So we can perform a Kempe swap either at v_1 or at v_2 to get a coloring in $\mathcal{L}_{v_1,c} \cup \mathcal{L}_{v_2,d}$. Thus, $\mathcal{L}$ mixes, and we are done. So we assume instead that $L(v_1) = L(v_2) = \{a, b, c\}$.

By Lemma 5.3, $\mathcal{L}_{v_1,c}$ and $\mathcal{L}_{v_2,c}$ each mix. If $\ell(P_1) = 1$ or $L(w) = \{a, b\}$ for every $w \in V(P_1)$, then $\mathcal{L}_{v_1,a}$ and $\mathcal{L}_{v_1,b}$ each mix by Lemma 5.2(a), with $w := v_2$. This is because $\varphi(w_2) = \varphi(v_1)$ for every

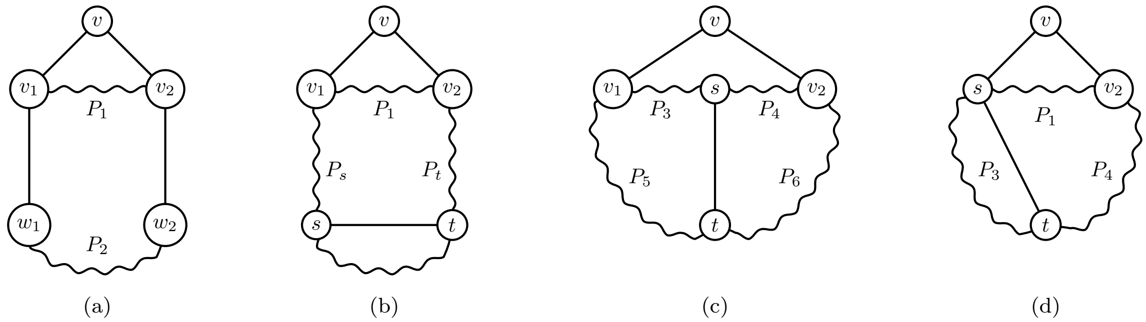

Figure 5.4: The 4 instances of Case 2, when H_1 contains 2 neighbors of v: the first comprises Case 2.1 and the remaining three comprise Case 2.2. (a) H_1 has no chord. (b) H_1 has a chord st on P_2. (c) H_1 has a chord st on $P_1 \cup P_2$. (d) H_1 has a chord v_1t (i.e. $s = v_1$).

$\varphi \in \mathcal{L}_{v_1,a} \cup \mathcal{L}_{v_1,b}$. Similarly, $\mathcal{L}_{v_2,a}$ and $\mathcal{L}_{v_2,b}$ each mix. Further, $\mathcal{L}_{v_1,a} \cap \mathcal{L}_{v_2,b} \neq \emptyset$. Also, for every $\gamma \in \{a, b\}$, the set $\mathcal{L}_{v_1,c} \cap \mathcal{L}_{v_2,\gamma} \neq \emptyset$ and $\mathcal{L}_{v_2,c} \cap \mathcal{L}_{v_1,\gamma} \neq \emptyset$. So $\mathcal{L}_{v_1,a} \cup \mathcal{L}_{v_1,b} \cup \mathcal{L}_{v_1,c}$ mixes; that is, $\mathcal{L}$ mixes. Thus, we assume $\ell(P_1) > 1$. Pick $y_1 \in N_{P_1}(v_1)$ and $y_2 \in N_{P_1}(v_2)$. By the above, we may assume $L(y_1) = L(y_2) = \{a, c\}$.

By Lemma 5.3, $\mathcal{L}_{v_1,b}$ and $\mathcal{L}_{v_2,b}$ each mix. Also, $\mathcal{L}_{v_1,c} \cap \mathcal{L}_{v_2,b} \neq \emptyset$ and $\mathcal{L}_{v_1,b} \cap \mathcal{L}_{v_2,c} \neq \emptyset$. We note that $\mathcal{L} = \mathcal{L}_{v_1,c} \cup \mathcal{L}_{v_1,b} \cup \mathcal{L}_{v_2,c} \cup \mathcal{L}_{v_2,b}$. So it suffices to show that $\mathcal{L}_{v_2,c} \cup \mathcal{L}_{v_2,b}$ mixes. As before, $\mathcal{L}_{v_1,a}$ mixes by Lemma 5.2(a), with $w := v_2$. Moreover, $\mathcal{L}_{v_1,a} \cap \mathcal{L}_{v_2,b} \neq \emptyset$ and $\mathcal{L}_{v_1,a} \cap \mathcal{L}_{v_2,c} \neq \emptyset$. Thus, $\mathcal{L}_{v_2,c} \cup \mathcal{L}_{v_2,b}$ mixes, and we are done.

Case 2.2: H_1 has a chord st. Recall, from the start of the proof, that H_1 is an even cycle (with at most one chord) closest to H_2; further, $\ell(P_1)$ and $\ell(P_2)$ are odd. Assume $s, t \in V(P_2) - \{v_1, v_2\}$ (with s closer to v_1). Let P_s (resp. P_t) be the s, v_1-path (resp. t, v_2-path) avoiding t (resp. avoiding s); see Figure 5.4b. If P_s and P_t have the same parity, then $G[H_1]$ is a bipartite theta graph (since H_1 is even), and we are done by Lemma 5.8. So assume P_s and P_t have opposite parities. Now $H_1[V(P_s) \cup V(P_t) \cup v]$ is an even cycle closer to H_2, contradicting our assumption. The same argument works (interchanging P_1 and P_2) if $s, t \in V(P_1) - \{v_1, v_2\}$.

Assume instead that $s \in V(P_1) - \{v_1, v_2\}$ and $t \in V(P_2) - \{v_1, v_2\}$. Let P_3, P_4, P_5, and P_6 be the v_1, s-path, v_2, s-path, v_1, t-path, and v_2, t-paths forming P_1 and P_2; see Figure 5.4c. By symmetry, assume P_3 is even and P_4 is odd. Now P_5 is even and P_6 is odd; otherwise, $G[H_1]$ is a bipartite theta graph, and we are done by Lemma 5.8. But $H_1[V(P_4) \cup V(P_5) \cup v]$ is an even cycle closer to H_2, contradicting our choice of H_1.

So the chord st must have an endpoint in $\{v_1, v_2\}$; say $s = v_1$. Note that $t \neq v_2$; otherwise, $G[H_1]$ is a bipartite theta graph, and we are done by Lemma 5.8. By symmetry, assume t is on P_2. Let P_3 and P_4 be the s, t and v_2, t-paths forming P_2; see Figure 5.4d. Now P_3 is even and P_4 is odd; otherwise, $G[H_1]$ is a bipartite theta graph, and we are done by Lemma 5.8. But again $H_1[V(P_4) \cup v \cup s]$ is an even cycle closer to H_2, contradicting our choice of H_1. $\square$

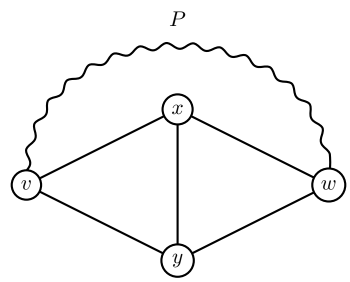

Figure 5.5: A K_4^+ formed by subdividing an edge one or more times, and the resulting path P.

Lemma 5.10. *The graph K_4^+ formed from K_4 by subdividing a single edge one or more times is degree-swappable.*

Proof. Let P be the path formed by subdividing the edge in K_4 one or more times. Denote the 3-vertices of G by v, w, x, y, where v and w each have a 2-neighbor; see Figure 5.5. Let L be a degree-assignment for G. By symmetry and by Lemma 5.5, we assume that $L(z) = \{1, 2\}$ for every $z \in V(P)$ and $\{1, 2\} \subseteq L(v) \cap L(w)$. By symmetry, assume that $L(v) = \{1, 2, 3\}$.

Assume first that $L(w) = L(v) = \{1, 2, 3\}$. By Lemma 5.3, each of $\mathcal{L}_{v,3}$ and $\mathcal{L}_{w,3}$ mix. Moreover, $\mathcal{L}_{v,3} \cap \mathcal{L}_{w,3} \neq \emptyset$. Thus, $\mathcal{L}_{v,3} \cup \mathcal{L}_{w,3}$ mixes. Let $\mathcal{L}_1 := \mathcal{L}_{v,3} \cup \mathcal{L}_{w,3}$ and $\mathcal{L}_2 := \mathcal{L} \setminus \mathcal{L}_1$. We now show that every $\varphi \in \mathcal{L}_2$ mixes with $\mathcal{L}_1$; thus $\mathcal{L}$ mixes. Pick $\varphi \in \mathcal{L}_2$ and assume by symmetry that $\varphi(v) = 1$. If $3 \notin \{\varphi(x), \varphi(y)\}$, then we recolor v with 3 to get a coloring in $\mathcal{L}_1$, and we are done. So assume by symmetry that $\varphi(x) = 3$. Now suppose $\varphi(w) = 1$. If $1 \notin L(x)$, there exists $\gamma \in L(x)$ with $\gamma \notin \cup_{z \in N[x]} \varphi(z)$. So we recolor x with γ then recolor v with 3, and we are done. If, instead, $1 \in L(x)$, then a 1,3-swap at v gives a coloring in $\mathcal{L}_1$, and we are done. So assume instead that $\varphi(w) = 2$. If $1 \notin L(x)$, then $2 \in L(x)$ by Lemma 5.5. Now a 2,3-swap at w gives a coloring in $\mathcal{L}_1$, and we are done. If, instead, $1 \in L(x)$, then a 1,3-swap at v gives a coloring in $\mathcal{L}_1$, and we are done.

Instead assume $L(w) \neq L(v)$; specifically, assume $L(w) = \{1, 2, 4\}$. By Lemma 5.3, the sets $\mathcal{L}_{v,3}$ and $\mathcal{L}_{w,4}$ each mix. Let $\mathcal{L}_1 := \mathcal{L}_{v,3} \cup \mathcal{L}_{w,4}$ and $\mathcal{L}_2 =: \mathcal{L} \setminus \mathcal{L}_1$. We show that $\mathcal{L}_1$ mixes. If $\{3, 4\} \not\subseteq L(x) \cap L(y)$ or $L(x) \neq L(y)$, then $\mathcal{L}_{v,3} \cap \mathcal{L}_{w,4} \neq \emptyset$; thus, $\mathcal{L}_1$ mixes. Otherwise, $L(x) = L(y) = \{3, 4, \alpha\}$ with $\alpha \in \{1, 2\}$ by Lemma 5.5. By symmetry, assume $L(x) = L(y) = \{3, 4, 1\}$. We show again that L_1 mixes.

Suppose $\ell(P)$ is even. Pick $\varphi \in \mathcal{L}_{v,2} \cap \mathcal{L}_{w,2} \cap \mathcal{L}_{x,3} \cap \mathcal{L}_{y,4}$. Now we can recolor either x with 1 then v with 3, or y with 1 then w with 4 to get colorings in $\mathcal{L}_{v,3}$ and $\mathcal{L}_{w,4}$, respectively. Thus, $\mathcal{L}_1$ mixes. Instead assume that $\ell(P)$ is odd. Pick $\varphi_1 \in \mathcal{L}_{v,1} \cap \mathcal{L}_{w,2} \cap \mathcal{L}_{x,3} \cap \mathcal{L}_{y,4}$ and $\varphi_2 \in \mathcal{L}_{v,2} \cap \mathcal{L}_{w,1} \cap \mathcal{L}_{x,3} \cap \mathcal{L}_{y,4}$. Now a 1,3-swap at v in φ_1 or a 1,4-swap at w in φ_2 give colorings in $\mathcal{L}_{v,3}$ and $\mathcal{L}_{w,4}$, respectively. Also, φ_1 and φ_2 are L-equivalent, as witnessed by a 1,2-swap at v. Thus, $\mathcal{L}_1$ mixes.

Finally, we show that every $\varphi^* \in \mathcal{L}_2$ mixes with $\mathcal{L}_1$; thus, $\mathcal{L}$ mixes. Pick $\varphi^* \in \mathcal{L}_2$ with $\varphi^*(v) = 1$, by symmetry. Note that $\{\varphi^*(x), \varphi^*(y)\} = \{3, 4\}$; otherwise, we can either recolor v with 3 or w with 4 to get a coloring in $\mathcal{L}_1$, and we are done. So assume $\varphi^*(x) = 3$ and $\varphi^*(y) = 4$. If $\ell(P)$ is odd, then we assume $\varphi^*(v) = 1$ and $\varphi^*(w) = 2$; if not, then we achieve this with a 1,2-swap at v. Now a 1,3-swap at v shows that φ^* mixes with $\mathcal{L}_1$. Instead assume that $\ell(P)$ is even. Now assume $\varphi^*(v) = \varphi^*(w) = 2$; if not, then we achieve this by a 1,2-swap at v. Recolor x with 1, then recolor v with 3. □

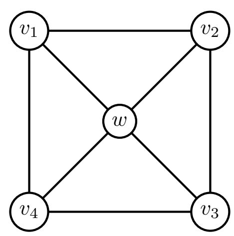

Figure 5.6: The 4-wheel.

Let $W_4 := C_4 \vee K_1$; this is the "4-wheel", or wheel with 4 spokes (see Figure 5.6).

Lemma 5.11. *The graph W_4 is degree-swappable.*

Proof. Let $G := W_4$. Denote the 3-vertices by v_1, v_2, v_3, v_4, in order along a 4-cycle, and denote the dominating vertex by w; see Figure 5.6. Fix a degree assignment L for G. Note that $|L(x) \cap L(y)| \geq 2$ for all $xy \in E(G)$, by Lemma 5.5. Since $|L(v_1) \cap L(v_2)| \geq 2$ and $|L(v_2) \cap L(v_3)| \geq 2$ and $|L(v_2)| = 3$, we conclude from Pigeonhole that $|L(v_1) \cap L(v_3)| \geq 1$. We consider the three cases $|L(v_1) \cap L(v_3)| \in \{1, 2, 3\}$.

Case 1: $|L(v_1) \cap L(v_3)| = 3$. Assume $L(v_1) = L(v_3)$. If $L(v_2) \neq L(v_1)$, then there exists $\beta \in L(v_1) \setminus L(v_2)$ since $|L(v_2)| = |L(v_1)|$. Further, since $L(v_1) = L(v_3)$, we have $L(v_1) \cap L(v_3) = L(v_1)$ so $\cup_{\alpha \in L(v_1) \cap L(v_3)} \mathcal{L}_{v_1,\alpha} = \mathcal{L}$. By Lemma 5.4(3), $\mathcal{L}$ mixes. So assume $L(v_2) = L(v_1)$ and, by symmetry, $L(v_4) = L(v_1)$. So there exists $\alpha \in L(w) \setminus \bigcup_{i=1}^{4} L(v_i)$, and clearly $\mathcal{L}_{w,\alpha}$ mixes. Given an L-coloring φ with $\varphi(w) \neq \alpha$, we can simply recolor w with α. Thus, $\mathcal{L}$ mixes.

Case 2: $|L(v_1) \cap L(v_3)| = 2$. Assume that $L(v_1) = \{a, b, c\}$ and $L(v_3) = \{a, b, d\}$. If $\{c, d\} \not\subseteq$

$L(v_2)\cap L(v_4)\cap L(w)$, then $\mathcal{L}$ mixes by Lemma 5.4(3), as above. So assume $\{c,d\}\subseteq L(v_2)\cap L(v_4)\cap L(w)$. By Case 1 and symmetry, we assume $L(v_2)\neq L(v_4)$. Interchanging v_2 and v_4 with v_1 and v_3 (and interchanging c and d with a and b) shows that also $a,b\in L(w)$. Further, $L(v_2)\cup L(v_4)\subseteq L(w)$. Thus, $L(v_2)=\{a,c,d\}$ and $L(v_4)=\{b,c,d\}$, up to possibly swapping v_2 and v_4. Now by Lemma 5.4(3), both $\bigcup_{\alpha\in\{a,b\}}\mathcal{L}_{v_1,\alpha}\cup\mathcal{L}_{v_3,\alpha}$ and $\bigcup_{\alpha\in\{c,d\}}\mathcal{L}_{v_2,\alpha}\cup\mathcal{L}_{v_4,\alpha}$ mix. The union of these two sets is all of $\mathcal{L}$. Further, the two sets mix, since some coloring φ lies in both; namely, $\varphi(v_1)=\varphi(v_3)=a$, $\varphi(v_2)=\varphi(v_4)=c$, and $\varphi(w)=d$. Thus, $\mathcal{L}$ mixes.

Case 3: $|L(v_1)\cap L(v_3)|=1$. Assume $L(v_1)=\{a,b,c\}$ and $L(v_3)=\{a,d,e\}$. By symmetry between v_1,v_3 and v_2,v_4, we assume that $|L(v_2)\cap L(v_4)|=1$. Since $|L(v_i)\cap L(v_j)|\geq 2$ for all $i\in\{1,3\}$ and $j\in\{2,4\}$, we assume that $L(v_2)=\{a,b,d\}$ and $L(v_4)=\{a,c,e\}$. Since $c\in L(v_1)\setminus L(v_2)$, Lemma 5.4(3) shows that $\mathcal{L}_{v_1,a}\cup\mathcal{L}_{v_1,c}$ mixes. Since $b\in L(v_1)\setminus L(v_4)$, Lemma 5.4(3) shows that $\mathcal{L}_{v_1,a}\cup\mathcal{L}_{v_1,b}$ mixes. Thus, $\mathcal{L}$ mixes. $\square$

5.3 Proof of Main Theorem

In this section, we prove our main result: Every connected k-regular graph (with $k\geq 3$) is k-swappable unless $G=K_{k+1}$ or $G=K_2\,\square\,K_3$. Further, these exceptional graphs are L-swappable whenever L is a Δ-assignment that is not identical everywhere. We split the main result into 4 cases depending on the connectivity of G. In Theorem 5.1, we prove the result for connectivity 1. In Theorem 5.2, we prove the result for connectivity 2 by using a cut set of size 2 (a *2-cut*), and showing that G contains an induced degree-swappable subgraph from the family of graphs compiled in Section 5.2.

For connectivity 3, we split into three cases: (i) $k\geq 5$, which we prove in Lemma 5.12, (ii) $k=4$, which we prove in Lemma 5.13, and (iii) $k=3$, which we prove in Lemma 5.14. For cases (i) and (ii) we use a 3-cut and show that G contains an induced degree-swappable subgraph from Section 5.2. For case (iii), we first show that for every 3-assignment L of G, the lists must be identical for all vertices. Then, in Theorem 5.3, we invoke the result of Feghali, Johnson, and Paulusma for 3-colorings of 3-regular graphs [40, Theorem 1] (for completeness, we state it below). Finally, we prove the result for 4-connected graphs in Theorem 5.4 as follows. In Lemmas 5.15 and 5.16, we handle the case of k-assignments L that are not identical for all vertices (Lemma 5.16 handles the case $G=K_{k+1}$). In Lemma 5.17, we show that either G contains a 4-wheel (so we are done by Lemma 5.11), or the absence of a 4-wheel restricts the possible L-colorings enough that G must be L-swappable.

Theorem A. *[40, Theorem 1] If G is a connected 3-regular graph that is neither K_4 nor the graph $K_2\,\square\,K_3$, then all 3-colorings of G are 3-equivalent.*

Theorem 5.1. *For $k \geq 3$, if G is k-regular with connectivity 1, then G is k-swappable.*

Proof. Since G has connectivity 1, it contains a cut-vertex. Thus, G contains at least two endblocks. If G contains at most one degree-choosable block, then some endblock B is not degree-choosable. Let v be the cut-vertex in B. Since $k \geq 3$ and G is k-regular, $B = K_{k+1}$ and $d_B(v) = k$. But now $G = B$ which contradicts that G has connectivity 1. Thus, G contains at least two degree-choosable endblocks. By Lemma 5.9, G is k-swappable. □

Now we consider the case that G has a vertex cut of size at most 3. In a graph G, a block is *Gallai* if it an odd cycle or a clique; otherwise it is *non-Gallai*. Gallai non-Gallai

Theorem 5.2. *For $k \geq 3$, if G is k-regular with connectivity 2, then G is k-swappable.*

Proof. Let G satisfy the hypotheses of the theorem, and let L be a k-assignment for G. Let S be a 2-cut in G. Let $G_S := G - S$. Now every endblock B of G_S has order at least $k-1$, since each vertex v of B has $d_G(v) = k$ and v has at most two edges to S. Further, if an endblock B is non-Gallai, then B contains an even cycle with at most one chord.[5] We call such a cycle a *good cycle*. If G_S has at least two endblocks each of which contains a good cycle, then Lemma 5.9 implies that G is L-swappable. So we assume that at most one endblock of G_S contains a good cycle. Thus, at most one endblock of G_S is non-Gallai. G_S good cycle

Claim 5.1. *G_S has at most two Gallai endblocks.*

Proof of Claim 5.1. Observe that every Gallai endblock B is regular of degree either $k-1$ or $k-2$, since $|S| = 2$. Further, every such B sends at least $k-1$ edges to S. This holds because either (i) B is regular of degree $k-1$, so B has at least k vertices, and at least $k-1$ of these each send an edge to S, or (ii) B is regular of degree $k-2$, so B has at least $k-1$ vertices, and at least $k-2$ of these each sends two edges to S.

Suppose, contrary to the claim, that S contains at least 3 Gallai endblocks, B_1, B_2, B_3. As noted above, each B_i sends S at least $k-1$ edges. So $3(k-1) \leq 2k$; thus, $k = 3$. Further, G_S has no other endblocks. But now some B_i is its own component, so it contains no cut-vertex, and thus sends more edges to S than counted above, which gives a contradiction. □

Claim 5.2. *G_S has exactly one non-Gallai endblock; call it B_0.*

[5]Recall that Rubin's Block Lemma [37] says a block contains such a cycle if and only if it is non-Gallai.

Proof of Claim 5.2. Suppose G_S has no non-Gallai endblocks. Since G_S has at most two (Gallai) endblocks, each endblock is its own component. So each endblock sends S at least $\min\{2 \cdot (k-1), 1 \cdot k\} = k$ edges, for a total of at least $2k$ edges to S. If either component of G_S is $(k-2)$-regular, then S has too many incident edges, so we get a contradiction. Thus, each endblock is K_k. If either vertex in S has at least 2 neighbors in each component of G_S, then we are done by Lemma 5.9. So assume that each vertex of S sends one edge to one component and sends $k-1$ edges to the other. Now G contains a copy of K_4^+, and we are done by Lemma 5.10. Thus, instead G_S has exactly one non-Gallai endblock; call it B_0. □

Case 1: G_S has exactly one Gallai endblock; call it B_1. If B_1 is $(k-1)$-regular, then some $v \in S$ has at least two neighbors in B_1, but does not dominate B_1. Thus, $B_1 + v$ contains a good cycle, and so does B_0. Now we are done by Lemma 5.9. So assume instead that B_1 is $(k-2)$-regular and that B_1 has order at least $k-1$. In fact, each vertex of S is adjacent to all of B_1. Now counting edges shows that $B_1 = K_{k-1}$, each vertex of S is adjacent to all of B_1, and S is an independent set. Hence, G contains a copy of K_4^+, and we are done by Lemma 5.10.

Case 2: G_S has exactly 2 Gallai endblocks, B_1 and B_2. Suppose that B_1 or B_2 is its own component; by symmetry, say that it is B_1. If B_1 is $(k-1)$-regular, then some vertex $x \in S$ has at least two neighbors in B_1, but does not dominate it. So $B_1 + x$ and B_0 each contain good cycles, and we are done by Lemma 5.9. If B_1 is $(k-2)$-regular, then it sends at least $2(k-1)$ edges to S. So the total number of edges incident to S is at least $2(k-1) + (k-1) + 1 = 3k - 2 > 2k$, a contradiction.

Thus, B_1 and B_2 are in the same component G_2 of G_S and B_0 is its own component G_1. Now S sends $2(k-1)$ edges to non-cut-vertices in endblocks of G_2. This implies, by counting edges to S, that every other vertex of G_2 is non-adjacent to every vertex in S.

If some vertex $x \in S$ sends at least two edges to B_1, then (since x sends no edges to the cut-vertex in B_1) G_S has a good cycle in B_0 and another good cycle in $B_1 + x$, so we are done by Lemma 5.9. The same is true if some $x \in S$ sends at least two edges to B_2. So assume that each $x \in S$ sends at most one edge to B_1 and one edge to B_2. Since B_1 and B_2 each receive at least $k-1$ edges from S, this implies that $k = 3$ and each $x \in S$ sends exactly one edge to each of B_1 and B_2. Now $G_2 + x$ must contain a good cycle, unless $B_1 = B_2 = K_2$ and $G_2 + x$ is an odd cycle. But in this case, each vertex of G_2 that is not a non-cut-vertex of B_1 or B_2 has too few incident edges, a contradiction. □

Lemma 5.12. *For $k \geq 5$, if G is k-regular with connectivity 3, then G is k-swappable.*

Proof. Let G satisfy the hypotheses of the lemma, and let L be a k-assignment for G. Let S be a minimal 3-cut of G, and let $G_S := G - S$. G_S

Claim 5.3. *Every Gallai endblock B of G_S is regular of degree $k-1$, $k-2$, or $k-3$ and sends at least $k-1$ edges to S. If B is a component of G_S, then B sends at least k edges to S.*

Proof of Claim 5.3. Since G is k-regular and $|S| = 3$, every non-cut-vertex in an endblock B sends at most 3 edges to S, so sends at least $k-3$ edges to vertices of B. Since B is a Gallai endblock, B is regular. And since G is 3-connected, some non-cut-vertex of B sends an edge to S. Thus, B is regular of degree $k-1$, $k-2$, or $k-3$.

Note that every Gallai endblock B of G_S sends at least $k-1$ edges to S. This is because either (i) B is $(k-1)$-regular, so at least $k-1$ vertices of B each send 1 edge to S, or (ii) B is $(k-2)$-regular, so at least $k-2$ vertices of B each send 2 edges to S, or (iii) B is $(k-3)$-regular, so at least $(k-3)$ vertices of B each send 3 edges to S. Thus, the number of edges that B sends to S is at least $\min\{1\cdot(k-1), 2\cdot(k-2), 3\cdot(k-3)\} = k-1$. Further, if a component of G_S consists of a single block, say B, then a similar computation shows that the number of edges B sends to S is at least k. □

Claim 5.4. *G_S has no endblock B that is an odd cycle with length at least 5.*

Proof of Claim 5.4. If such a B exists, then it sends at least $4(k-2)$ edges to S. By Claim 5.3, each component not containing B sends at least k edges to S. But $4\cdot(k-2)+k = 5k-8 > 3k$, a contradiction. □

Claim 5.5. *G_S has at most 1 non-Gallai endblock and at most 3 Gallai endblocks.*

Proof of Claim 5.5. If G_S contains at least 2 non-Gallai endblocks, then it contains 2 good cycles; so G is degree-swappable, by Lemma 5.9. So instead G_S contains at most one non-Gallai endblock. If G_S has at least 4 Gallai endblocks, then by Claim 5.3 the number of edges from S to endblocks of G_S is at least $4(k-1) > 3k$, a contradiction. So G_S has at most 3 Gallai endblocks. □

We will show G contains an induced degree-swappable subgraph, so we are done by Lemma 5.6. Specifically, we will often show G contains two good cycles, and invoke Lemma 5.9.

Case 1: G_S has 0 non-Gallai endblocks. Let B_1 be an endblock that is a whole component G_1. If $B_1 = K_{k-2}$, then each vertex in S is adjacent to all of B_1 (and some pair of vertices in S is non-adjacent), so G contains an induced K_4^+; thus, we are done by Lemma 5.10. So instead $B_1 \in \{K_{k-1}, K_k\}$. Suppose

$B_1 = K_k$. Now some vertex $x \in S$ is adjacent to at least two vertices of B_1 and also non-adjacent to some vertex $y \in B_1$; see Figure 5.7a. (This follows from Pigeonhole and the fact that each vertex of S has a neighbor in B_1.) Choose $z \in N(y) \cap S$. Since x and z each have a neighbor in G_2, the subgraph $G[S \cup V(G_2)]$ contains an x, z-path. Note that $N(x) \cap N(z) \cap V(B_1) = \emptyset$, since each vertex of B_1 has a unique neighbor in S. Thus, G contains a copy of K_4^+ (possibly using the edge xz, rather than an x, z-path through G_2), and we are done by Lemma 5.10.

Assume instead that $B_1 = K_{k-1}$. Now each vertex of B_1 has exactly two neighbors in S. Denote S by $\{x_1, x_2, x_3\}$. By symmetry, there exist vertices $y_1, y_2, y_3 \in B_1$ such that y_1 and y_2 are both adjacent to each of x_1 and x_2; and y_3 is adjacent to x_2 and x_3. Now $\{x_1, y_1, y_2, y_3\}$ induces $K_4 - e$. Since x_3 is adjacent to y_3 (but not y_1 or y_2), again G contains K_4^+ as above, so we are done by Lemma 5.10.

Case 2: G_S has exactly 1 non-Gallai endblock, B_1. Let G_1 be the component of G_S containing B_1, and let B_2 be a regular endblock of another component G_2. Since S sends at least $k-1$ edges to B_2, by Pigeonhole some vertex $x \in S$ sends at least 2 edges to B_2. If B_2 has at least k vertices, then x is not adjacent to all of B_2, since x sends an edge to G_1. Thus, $B_2 + x$ is non-Gallai, so it contains a good cycle. Since $B_2 + x$ and B_1 each contain a good cycle, we are done by Lemma 5.9.

Assume instead that B_2 has order $k-1$ or $k-2$. We will show that $G_2 = B_2$. Assume the contrary; so B_2 has a cut-vertex y; see Figure 5.7b. If B_2 has order $k-2$, then $d_{G_2}(y) \geq d_{B_2}(y) + 1 = (|B_2| - 1) + 1 = k - 2$; so $d_{G[S \cup \{y\}]}(y) = k - d_{G_2}(y) \leq 2$ However, each other vertex of B_2 is adjacent to all of S. So some $x \in S$ is not adjacent to y but sends at least two edges to B_2. Thus, $B_2 + x$ and B_1 each contain a good cycle, so we are done by Lemma 5.9.

Assume instead that B_2 has order $k-1$; again, see Figure 5.7b. The number of edges from S to B_2 is at least $2(k-2)$, and $d_{G_2}(y) \geq 1 + d_{B_2}(y) = 1 + ((k-1) - 1) = k - 1$, so $d_{G[S \cup \{y\}]}(y) \leq 1$. Thus, S sends y at most 1 edge and in total sends all other vertices of B_2 at least $2(k-2)$ edges. If two vertices of S each send only one edge to B_2, then some non-cut-vertex of B_2 has too few incident edges. This is because each non-cut-vertex of B_2 needs an edge from at least 2 vertices of S, and B_2 has at least $(k-2) \geq 3$ non-cut-vertices. So at least 2 vertices of S send at least 2 edges to B_2, and at least one of them is non-adjacent to y. We let x denote such a vertex. Now B_1 and $B_2 + x$ each contain a good cycle, and we are done by Lemma 5.9. Thus, we assume that $G_2 = B_2$.

Suppose B_2 has $k-1$ vertices. Now S sends B_2 exactly $2(k-1)$ edges. Thus, some vertex $x \in S$ sends at least two edges to B_2 and is not adjacent to all of B_2; this is because $(k-1) + 2(1) < 2(k-1) < 2(k-1) + 1$.

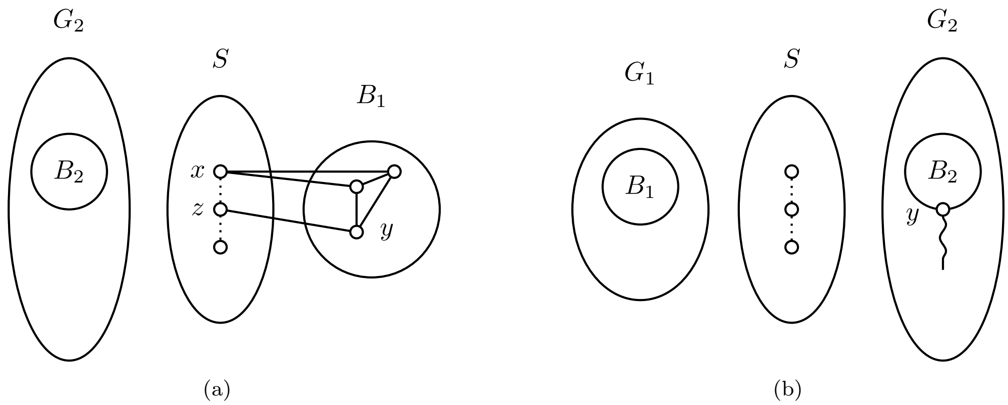

Figure 5.7: S and the components of G_S when (a) G_1 is a single Gallai endblock B_1, and (b) G_1 contains a non-Gallai endblock B_1 and G_2 has more than one endblock.

(The first term represents the possibility of at most one vertex of S sending at least two edges to B_1 and the third term represents the possibility of at least two vertices of S dominating B_1.) Now $B_2 + x$ contains a good cycle, as does B_1, so we are done by Lemma 5.9. Hence, B_2 has $k-2$ vertices, and every vertex in S is adjacent to all of B_2. At least one pair in S is not adjacent; call it x_1, x_2. Each of x_1 and x_2 has a neighbor in G_1, so G contains a copy of K_4^+, and we are done by Lemma 5.10. □

Lemma 5.13. *If G is a 4-regular graph with connectivity 3, then G is k-swappable.*

Proof. Let S be a vertex cut of size 3, and let $G_S := G - S$. As in the proof of Lemma 5.12, if G_S has at least two non-Gallai endblocks, then we are done by Lemma 5.9. So we assume that G_S has at most one non-Gallai endblock. G_S

Claim 5.6. *Each Gallai endblock of G_S is K_2, K_3, or K_4.*

Proof of Claim 5.6. The proofs of Claims 5.3 and 5.4 in Lemma 5.12 also work here essentially unchanged. So each Gallai endblock is regular of degree $k-1$, $k-2$, or $k-3$. (Here $k=4$.) As before, if G_S has an endblock B that is an odd cycle with length at least 5, then the number of edges going to S is at least $4(k-2)+k$. However, we can strengthen this bound as follows. If B is its own component, then the number of edges it sends to S is at least $5(k-2)$. And if B is not its own component, then some other endblock in its component sends S at least $k-1$ edges. Thus, we get $\min\{4\cdot(k-2)+(k-1), 5(k-2)\}+k > 3k$, again a contradiction. □

Case 1: G_S has exactly 1 non-Gallai endblock B_1. Let G_1 be the component of G_S containing B_1. Let B_2 be a Gallai endblock in some other component G_2. B_1, G_1 B_2, G_2

Case 1.1: B_2 is K_4. Now $G_2 + x$ contains a good cycle, for some $x \in S$, as follows (so we are done by Lemma 5.9). If some vertex x of S sends at least 2 edges to B_2, then $B_2 + x$ is non-Gallai, since x is also non-adjacent to some vertex in B_2 (this includes the case that $G_2 = B_2$). Thus, $x + B_2$ contains a good cycle. Otherwise, each vertex of S sends one edge to B_2 and some vertex x sends an edge to another endblock of G_2. Now $x + G_2$ is 2-connected and irregular (so it is non-Gallai and contains a good cycle).

Case 1.2: B_2 is K_3. Suppose that no vertex of S sends edges to exactly 2 vertices of B_2. Now $G_2 \neq B_2$, two vertices of S each send a single edge to B_2, and at least one of them, call it x, sends an edge to another endblock of G_2. Now $G_2 + x$ is 2-connected and irregular; thus, it contains a good cycle. So we assume some vertex $x \in S$ sends exactly two edges to B_2. Now $B_2 + x$ is non-Gallai and contains a good cycle. Since B_1 is non-Gallai, we are done by Lemma 5.9.

Case 1.3: B_2 is K_2. By symmetry, we assume that every Gallai endblock in G_2 is K_2. If $G_2 = K_2$, then G contains a copy of K_4^+ (with its non-adjacent 3-vertices in S), so we are done by Lemma 5.10. If G_2 contains at least 3 endblocks, then each non-cut-vertex in each of these endblocks is adjacent to all of S, so G contains an induced $K_{2,3}$ (with two vertices in S), which is a bipartite theta graph, so we are done by Lemma 5.8. Thus, G_2 has exactly two endblocks. If $G_2 = P_3$, then some x in S is adjacent to all vertices in G_2, so $G_2 + x$ contains a good cycle, and we are done by Lemma 5.9. If some interior block B_3 of G_2 is non-Gallai (so contains a good cycle), then each x in S has a path to B_3 in G_2 (and sends at most one edge to B_3), so we are done by considering B_1, B_3, and a shortest path between them. Instead, assume that every interior block of G_2 is Gallai. Since G_2 only has two endblocks, some $x \in S$ sends an edge to the interior block B_3. Again, $G_2 + x$ contains a good cycle.

Case 2: G_S has 0 non-Gallai endblocks.

Case 2.1: G_S has at least 4 endblocks. As in the proof of Lemma 5.12, each endblock sends at least $k - 1 = 3$ edges to S. So the number of endblocks in G_S is at most $|S|k/(k-1) = 3 \cdot 4/3 = 4$. Suppose that G_S has 4 Gallai endblocks. If some endblock is a K_3, then it sends S at least 4 edges, so S has too many incident edges (as above), a contradiction. Now G_S has only 2 components, and no component is a single endblock; otherwise, S sends that component at least 4 edges, and $4 + 3 \cdot 3 > 4|S|$, a contradiction. So let B_1, B_2 be endblocks of G_1 and let B_3, B_4 be endblocks of G_2, with $B_1, B_2, B_3, B_4 \in \{K_2, K_4\}$. Similarly, S sends edges only to non-cut-vertices of B_1, B_2, B_3, B_4. This implies that if some vertex $x \in S$ sends two edges each to G_1 and G_2, then $x + G_1$ and $x + G_2$ each contain a good cycle, and we are done by Lemma 5.9. That is because, for each $i \in [2]$, either some endblock of G_i is K_4 or both endblocks of G_i are K_2 and there is an interior endblock B_5 that is irregular. Now by Pigeonhole, some x in S must send two edges each to

G_1 and G_2, so we are done.

Case 2.2: G_S has a component that is neither K_3 nor K_4. By Case 2.1, G_S has at most 3 endblocks. So let B_1 be a block of G_S that is a whole component of G_1. If B_1 is an odd cycle of length 5 or more, then B_1 sends S at least 10 edges, and $G_S - B_1$ sends S at least $k-1 = 3$ edges, but $10+3 > 4|S| = 12$, a contradiction. If $B_1 = K_2$, then every vertex of S is adjacent to all of B_1, and some pair of vertices in S is non-adjacent. Thus, G contains a K_4^+, and we are done by Lemma 5.10.

Instead assume that $B_1 \in \{K_3, K_4\}$. This implies that some vertex $x \in S$, sends at least 2 edges to B_1 and is not adjacent to all of B_1. If x sends an edge to two endblocks in G_2, then $G_2 + x$ contains a good cycle (either some endblock of G_2 is not K_2, or else $G_2 = P_3$, or G_2 contains an interior block that is not K_2). Thus, x sends an edge to at most one endblock of G_2. So if G_2 contains two endblocks, call them B_2 and B_3, then we may assume that x sends no edges to B_3 and that $B_3 \neq K_2$. This implies that some vertex $y \in S$ sends at least 2 edges to B_3 and is not adjacent to all of B_3, so $B_3 + y$ contains a good cycle. Now we are done, by considering the good cycles in $B_1 + x$ and in $B_3 + y$. Hence G_2 has a single endblock; that is, $G_2 = B_2$. By the previous paragraph, $G_2 \neq K_2$, a contradiction.

Case 2.3: G_S has a component that is K_3 or has 3 components. Suppose that $B_1 = B_2 = K_3$. If some vertex $x \in S$ sends exactly 2 edges to each G_i, then each $G_i + x$ contains a good cycle, so we are done by Lemma 5.9. But such x must exist because (a) each B_i gets 3 edges from S, (b) S sends out at most 12 edges, and (c) we cannot write 6 as a sum of 3 terms, each 1 or 3. Assume instead that $B_1 = K_4$. If G_S has a third component G_3, then G_3 is a single block B_3, and counting edges from S shows that $B_1 = B_2 = B_3 = K_4$ and S is an independent set. Now G contains a bipartite theta graph (take 2 vertices in S and a neighbor of each in each component). Assume instead that G_S has only 2 components.

By symmetry, we assume that $B_1 = K_4$ and $B_2 = K_3$. If some vertex $x_i \in S$ sends two edges to each of B_1 and B_2, then we are done; so assume this does not happen. By symmetry, we assume that $d_{B_2}(x_1) = 3$, $d_{B_2}(x_2) = 2$, and $d_{B_2}(x_3) = 1$, so $d_{B_1}(x_1) = d_{B_1}(x_2) = 1$ and $d_{B_1}(x_3) = 2$; see Figure 5.8a. This implies that $x_2x_3 \in E(G)$. But now x_2 and x_3 lie in a 4-cycle with two vertices of B_1, and they also lie in a 4-cycle with two vertices of B_2. The union of these 4-cycles is a bipartite theta graph; so we are done, by Lemma 5.8.

Case 2.4: G_S has 2 components and both are K_4. Suppose that $B_1 = B_2 = K_4$ (and G_S has only two components). Denote S by x_1, x_2, x_3. By symmetry (and counting edges), we assume that $d_{B_1}(x_1) = d_{B_2}(x_2) = 2$ and $d_{B_2}(x_1) = d_{B_1}(x_2) = d_{B_1}(x_3) = d_{B_2}(x_3) = 1$. This implies that $x_1x_3, x_2x_3 \in E(G)$ and $x_1x_2 \notin E(G)$; see Figure 5.8b. If y_1 is a neighbor of x_1 in B_1 and y_2 is a neighbor of x_3 in B_1, then

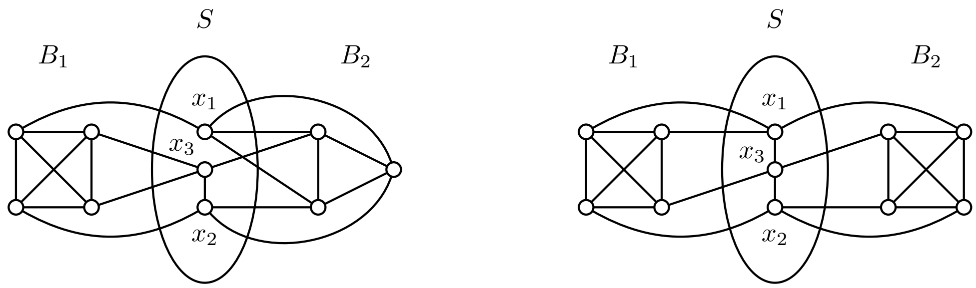

(a) (b)

Figure 5.8: S and the components of G_S when (a) $B_1 = K_4$ and $B_2 = K_3$, and (b) $B_1 = B_2 = K_4$.

$\{x_1, x_3, y_1, y_2\}$ induces a 4-cycle in $V(B_1) \cup S$. Similarly, there exists an induced 4-cycle in $V(B_2) \cup S$ that uses edge x_2x_3. Each of these 4-cycles is a good cycle, and they intersect in a single vertex x_3. So we are done. □

Now we handle the case when G is 3-regular and has connectivity 3, and two vertices of G have distinct lists.

Lemma 5.14. *Let G be 3-connected and 3-regular, but not K_4. Let L be a 3-assignment for G. If there exist $v, w_1 \in V(G)$ with $L(w_1) \neq L(v)$, then G is L-swappable.*

Proof. Since G is connected, we assume $vw_1 \in E(G)$. Since G is not a clique, G has an L-coloring by the list version of Brooks' Theorem. Thus, $\mathcal{L} \neq \emptyset$. Denote $N(v)$ by $\{w_1, w_2, w_3\}$. If possible, also choose v so that $L(w_2) \neq L(v)$ or $L(w_3) \neq L(v)$ (or both); this choice is used only once, near the end of the proof. If $w_1w_2, w_2w_3 \in E(G)$, then $\{w_1, w_3\}$ is a vertex cut, which contradicts that G is 3-connected. Thus, by symmetry, $\{w_1, w_2, w_3\}$ induces at most one edge.

We note that $G - e$ is degree-choosable for each $e \in E(G)$, as follows. Assume instead that $G - e$ is a Gallai tree for some $e \in E(G)$. Since $G - e$ is 2-connected, it is is a Gallai block, i.e., a complete graph or an odd cycle. But this is impossible since $G - e$ is irregular. Thus, $G - e$ is degree-choosable. By Lemma 5.5, this implies that $|L(x) \cap L(y)| \geq 2$ for all $xy \in E(G)$. In particular, $|L(w_i) \cap L(v)| \geq 2$ for each $i \in [3]$.

Let $\mathcal{A}_1 := \cup_{\alpha \in L(w_1) \cap L(w_2)} \mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}$ and $\mathcal{B}_1 := \cup_{\alpha \in L(w_1) \setminus L(v)} \mathcal{L}_{w_1,\alpha}$. Define $\mathcal{A}_2$, $\mathcal{B}_2$, $\mathcal{A}_3$, and $\mathcal{B}_3$ analogously, as in Table 5.9. Let $\mathcal{A} := \cup_{i=1}^3 \mathcal{A}_i$ and $\mathcal{B} := \cup_{i=1}^3 \mathcal{B}_i$. Observe that $\mathcal{L} = \mathcal{A} \cup \mathcal{B}$. Moreover, by Lemma 5.3, $\mathcal{B}_1$ is nonempty and mixes. Recall from above that $G[w_1, w_2, w_3]$ has at most one edge.

$\mathcal{A}_1$, $\mathcal{B}_1$
$\mathcal{A}_2$, $\mathcal{B}_2$
$\mathcal{A}_3$, $\mathcal{B}_3$
$\mathcal{A}$, $\mathcal{B}$

To help clarify the arguments in this proof, we will often draw an auxiliary graph that has a vertex for

$\mathcal{A}$	$\mathcal{B}$
$\mathcal{A}_1 : \cup_{\alpha \in L(w_1)\cap L(w_2)} \mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha}$	$\mathcal{B}_1 : \cup_{\alpha \in L(w_1)\setminus L(v)} \mathcal{L}_{w_1,\alpha}$
$\mathcal{A}_2 : \cup_{\alpha \in L(w_1)\cap L(w_3)} \mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_3,\alpha}$	$\mathcal{B}_2 : \cup_{\alpha \in L(w_2)\setminus L(v)} \mathcal{L}_{w_2,\alpha}$
$\mathcal{A}_3 : \cup_{\alpha \in L(w_2)\cap L(w_3)} \mathcal{L}_{w_2,\alpha} \cap \mathcal{L}_{w_3,\alpha}$	$\mathcal{B}_3 : \cup_{\alpha \in L(w_3)\setminus L(v)} \mathcal{L}_{w_3,\alpha}$

Figure 5.9: Every L-coloring is in $\mathcal{A}$ or $\mathcal{B}$.

each $\mathcal{A}_i$ and $\mathcal{B}_j$ that is non-empty. (Since $|L(w_i) \cap L(v)| \geq 2$ for all i, we have $|L(w_i) \setminus L(v)| \leq 1$. Thus, each nonempty $\mathcal{B}_j$ mixes.) If $\mathcal{A}_i \cup \mathcal{B}_j$ mixes, then we draw an edge between the vertices $\mathcal{A}_i$ and $\mathcal{B}_j$. So, to show that $\mathcal{L}$ mixes, it suffices to show that this auxiliary graph is connected.

Case 1: $G[w_1, w_2, w_3]$ has one edge. By symmetry between w_2 and w_3, assume that $w_1w_3 \notin E(G)$. So either $w_2w_3 \in E(G)$ or $w_1w_2 \in E(G)$.

First suppose $w_2w_3 \in E(G)$. This implies that $\mathcal{A}_3 = \emptyset$. Since $w_1w_2 \notin E(G)$, the sets $\mathcal{A}_1$ and $\mathcal{A}_2$ are both nonempty. By Lemma 5.4(2), $\mathcal{A}_1 \cup \mathcal{B}_1$ and $\mathcal{A}_2 \cup \mathcal{B}_1$ both mix; see Figure 5.10a. Moreover, for each $i \in \{2, 3\}$, if $\mathcal{B}_i$ is nonempty, then $\mathcal{B}_i \cup \mathcal{A}_{i-1}$ mixes by Lemma 5.4(2). So $\mathcal{A} \cup \mathcal{B}$ mixes; that is, $\mathcal{L}$ mixes.

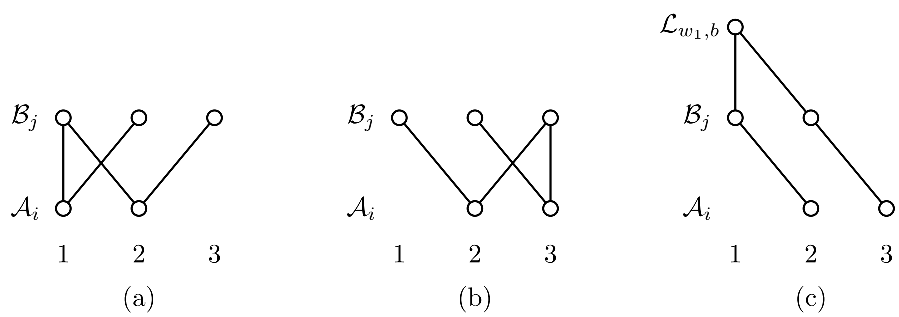

Figure 5.10: (a) The case that $w_2w_3 \in E(G)$. (b) The case that $w_1w_2 \in E(G)$ and $\mathcal{B}_3 \neq \emptyset$. (c) The case that $w_1w_2 \in E(G)$, $\mathcal{B}_3 = \emptyset$, and $\mathcal{B}_2 \neq \emptyset$.

Instead assume $w_1w_2 \in E(G)$. So $\mathcal{A}_1 = \emptyset$. Since $w_2w_3 \notin E(G)$, the sets $\mathcal{A}_2$ and $\mathcal{A}_3$ are both nonempty. By Lemma 5.4(2), $\mathcal{A}_2 \cup \mathcal{B}_1$ mixes. Now we show that $\mathcal{B}_2 = \mathcal{B}_3 = \emptyset$. Suppose first that $\mathcal{B}_3 \neq \emptyset$. By Lemma 5.4(2), $\mathcal{A}_2 \cup \mathcal{B}_3$ and $\mathcal{A}_3 \cup \mathcal{B}_3$ both mix; see Figure 5.10b. If $\mathcal{B}_2 = \emptyset$, then we are done, since $\mathcal{A} \cup \mathcal{B}$ mixes. So assume $\mathcal{B}_2 \neq \emptyset$. Now by Lemma 5.4(2) $\mathcal{A}_3 \cup \mathcal{B}_2$ mixes. Thus, $\mathcal{A} \cup \mathcal{B}$ mixes; that is, $\mathcal{L}$ mixes. So we assume $\mathcal{B}_3 = \emptyset$.

Now suppose that $\mathcal{B}_2 \neq \emptyset$. By Lemma 5.4, $\mathcal{B}_2 \cup \mathcal{A}_3$ mixes. Since $\mathcal{B}_1$ and $\mathcal{B}_2$ are nonempty, there exists $\alpha \in L(w_1) \setminus L(v)$ and $\beta \in L(w_2) \setminus L(v)$. If $\alpha \neq \beta$, then $\mathcal{B}_1 \cap \mathcal{B}_2 \neq \emptyset$; thus, $\mathcal{B}_1 \cup \mathcal{B}_2$ mixes. So $\mathcal{A} \cup \mathcal{B}$ mixes, and we are done. So assume $\alpha = \beta$. If $L(w_1) = L(w_2)$, then there exists $\gamma \in L(v) \setminus (L(w_1) \cap L(w_2))$. By Lemma 5.3, the set $\mathcal{L}_{v,\gamma}$ mixes. Moreover, $\mathcal{B}_1 \cap \mathcal{L}_{v,\gamma}$ and $\mathcal{B}_2 \cap \mathcal{L}_{v,\gamma}$ are both nonempty. Thus, $\mathcal{B}_1 \cup \mathcal{L}_{v,\gamma} \cup \mathcal{B}_2$

mixes. So $\mathcal{A} \cup \mathcal{B}$ mixes, and we are done. So assume $L(w_1) \neq L(w_2)$. Pick $a \in L(v) \cap L(w_1) \cap L(w_2)$ and $b \in L(w_1) \setminus L(w_2)$ and $c \in L(w_2) \setminus L(w_1)$; see Figure 5.11a. Note that $\mathcal{L}_{w_1,b}$ mixes by Lemma 5.3, and $\mathcal{B}_2 \cap \mathcal{L}_{w_1,b} \neq \emptyset$. So $\mathcal{B}_2 \cup \mathcal{L}_{w_1,b}$ mixes; see Figure 5.10c. Moreover, there exists $\varphi \in \mathcal{L}_{w_1,b} \cap \mathcal{L}_{w_3,b}$, and φ mixes with $\mathcal{B}_1$ by Lemma 5.4(2). Since $\varphi \in \mathcal{L}_{w_1,b}$, the set $\mathcal{L}_{w_1,b} \cup \mathcal{B}_1$ mixes; hence, $\mathcal{B}_1 \cup \mathcal{B}_2$ mixes. So $\mathcal{A} \cup \mathcal{B}$ mixes, and we are done. So we instead assume $\mathcal{B}_2 = \emptyset$.

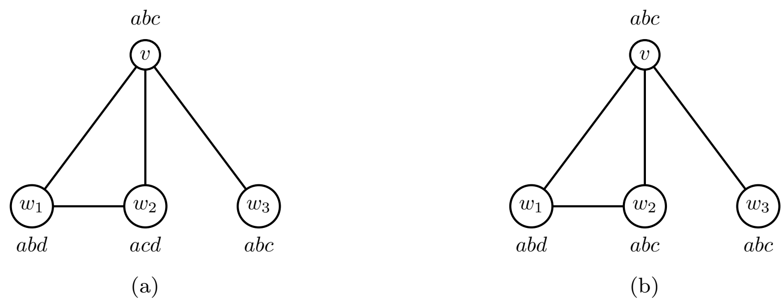

Figure 5.11: Two cases when $w_1 w_2 \in E(G)$. (a) A 3-assignment for $N[v]$ when $d \in L(w_1) \cap L(w_2)$ and $L(w_1) \neq L(w_2)$ and $L(v) = L(w_3)$. (b) A 3-assignment for $N[v]$ when $L(v) = L(w_2) = L(w_3)$.

Now $\mathcal{B}_2 = \emptyset = \mathcal{B}_3$, so $L(w_2) = L(w_3) = L(v)$; see Figure 5.11b. From above $\mathcal{A}_1 = \emptyset$ and $\mathcal{A}_2 \cup \mathcal{B}_1$ mixes. So it suffices to show that $\mathcal{A}_3$ mixes with $\mathcal{B}_1$. By Lemma 5.3, we know that $\mathcal{L}_{w_2,c}$ and $\mathcal{L}_{w_1,d}$ and $\mathcal{L}_{v,c}$ are nonempty sets that mix. Note that $\mathcal{B}_1 = \mathcal{L}_{w_1,d}$. Since $\mathcal{L}_{w_1,d} \cap \mathcal{L}_{w_2,c} \neq \emptyset$, we see that $\mathcal{L}_{w_2,c} \cup \mathcal{L}_{w_1,d}$ mixes. Similarly, $\mathcal{L}_{w_1,d} \cap \mathcal{L}_{v,c} \neq \emptyset$, so $\mathcal{L}_{w_1,d} \cup \mathcal{L}_{v,c}$ mixes. Finally $\mathcal{L}_{v,c} \cap \mathcal{L}_{w_2,\alpha} \cap \mathcal{L}_{w_3,\alpha} \neq \emptyset$ for each $\alpha \in \{a, b\}$. Thus, $\mathcal{L}_{v,c} \cup \bigcup_{\alpha \in \{a,b\}} (\mathcal{L}_{w_2,\alpha} \cap \mathcal{L}_{w_3,\alpha})$ mixes. Combining all these observations gives that $\mathcal{L}_{w_2,c} \cup \mathcal{L}_{w_1,d} \cup \mathcal{L}_{v,c} \cup \bigcup_{\alpha \in \{a,b\}} (\mathcal{L}_{w_2,\alpha} \cap \mathcal{L}_{w_3,\alpha})$ mixes. But this set contains $\mathcal{A}_3 \cup \mathcal{B}_1$, so we are done.

Case 2: $G[w_1, w_2, w_3]$ has no edges. So $\mathcal{A}_1, \mathcal{A}_2$, and $\mathcal{A}_3$ are all nonempty. Moreover, $\mathcal{B}_1 \cup \mathcal{A}_1$ and $\mathcal{B}_1 \cup \mathcal{A}_2$ both mix by Lemma 5.4(2). If $\mathcal{B}_2 \neq \emptyset$, then $\mathcal{B}_2 \cup \mathcal{A}_1$ and $\mathcal{B}_2 \cup \mathcal{A}_3$ both mix by Lemma 5.4(2); see Figure 5.12a. If, in addition, $\mathcal{B}_3 = \emptyset$, then $\mathcal{A} \cup \mathcal{B}_1 \cup \mathcal{B}_2$ mixes, and we are done. Otherwise, $\mathcal{B}_3 \neq \emptyset$. So $\mathcal{B}_3 \cup \mathcal{A}_2$ and $\mathcal{B}_3 \cup \mathcal{A}_3$ both mix by Lemma 5.4(2). Thus, $\mathcal{A} \cup \mathcal{B}$ mixes, and we are done. So assume $\mathcal{B}_2 = \emptyset$; by symmetry, also $\mathcal{B}_3 = \emptyset$.

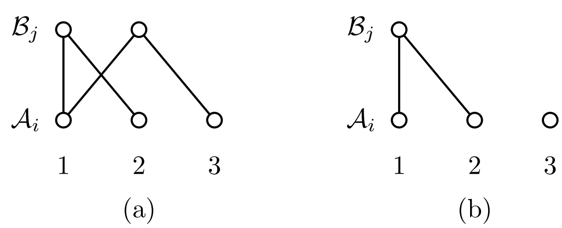

Figure 5.12: Here $G[w_1, w_2, w_3]$ has no edges. (a) The case that $\mathcal{B}_2 \neq \emptyset$. (b) The case that $\mathcal{B}_2 = \emptyset = \mathcal{B}_3$.

Now it suffices to show that $\mathcal{A}_3 \cup \mathcal{B}_1$ mixes. Fix z_1 and z_2 in $N(w_1) \setminus \{v\}$. By symmetry between w_1 and

v, we assume $N(w_1)$ induces no edges and $L(w_1) = L(z_1) = L(z_2)$; see Figure 5.13a. By Lemma 5.3, $\mathcal{L}_{v,c}$ is nonempty and mixes. Now we show that $\mathcal{A}_3$ mixes. Form G' from G by deleting v then identifying w_2 and w_3. Call the new vertex w_{23}. Note that G' is 2-connected; equivalently, w_{23} is not a cut-vertex in G'. To see this, note that $\{w_2, w_3\}$ is not a vertex cut in G, since G is 3-connected. So $\{w_2, w_3, v\}$ is not a vertex cut in G, since v is a leaf in $G - \{w_2, w_3\}$. But the components of $G - \{v, w_2, w_3\}$ are the same as those of $G' - w_{23}$. So w_{23} is not a cut-vertex in G', as desired.

Let L' be a 3-assignment for G' with $L'(w_{23}) = L(w_2)$ and $L'(x) = L(x)$ for all other $x \in V(G')$. Now G' is L'-swappable, by Lemma 5.2(a) with $w := w_1$, since $d_{G'}(w_1) = 2 < 3 = |L'(w_1)|$. We note that every coloring φ' of G' can be extended to a coloring φ in G since $|\cup_{x \in N(v)} \varphi'(x)| \leq 2$. Thus, L'-colorings are in bijection with colorings in $\mathcal{A}_3$. Moreover, every α, β-swap performed in φ' can also be performed in φ as follows. If the swap does not involve $N(v)$, or $\varphi(v) \notin \{\alpha, \beta\}$, then we perform the same swap as in φ' (possibly a swap at w_2 and w_3 each). If $\varphi(w_1) = \varphi(w_2) = \varphi(w_3)$ or $\beta = d$ (so the swap is at w_1), then we can recolor v with $\gamma \notin \{\alpha, \beta\}$ then perform the swap. Otherwise, suppose $\varphi(v) = \alpha$ and (i) $\varphi(w_1) = \beta$, or (ii) $\varphi(w_2) = \varphi(w_3) = \beta$. In case (i), the swap is valid since $\{a, b\} \in L(v)$, and in case (ii), the swap is valid since $L(v) = L(w_2) = L(w_3)$. Thus, $\mathcal{A}_3$ mixes, as claimed.

For every $\varphi \in \cup_{\gamma \in \{a,b\}} \mathcal{L}_{w_2,\gamma} \cap \mathcal{L}_{w_3,\gamma}$, either (i) $\varphi(w_1) = \varphi(w_2) = \varphi(w_3)$, or (ii) $\varphi(w_1) = d$, or (iii) $\varphi(v) = c$. In case (i), we have $\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3 \neq \emptyset$, so we are done. In case (ii), $\varphi \in \mathcal{B}_1$; and in case (iii), $\varphi \in \mathcal{L}_{v,c}$. Thus, $\mathcal{L}_{w_2,\gamma} \cap \mathcal{L}_{w_3,\gamma}$ mixes either with $\mathcal{B}_1$ or $\mathcal{L}_{v,c}$, for each $\gamma \in \{a, b\}$. So it suffices to show that $\mathcal{B}_1$ mixes with $\mathcal{L}_{v,c}$. By our choice of v at the beginning of the proof, there exists $u \in N(z_1) \setminus \{w_1\}$ with $L(u) = L(z_1)$; otherwise, we would have chosen z_1 instead of v. Since G is 3-connected, $G - z_1 - z_2$ contains a u, v-path P. Note that $L(x) \neq L(y)$ for some x and y that are successive on P. By construction, either x or y is not in $N(v) \cup N(w_1)$; see Figure 5.13b. By symmetry, assume $x \notin N(v) \cup N(w_1)$. Since there exists $\gamma \in L(x) \setminus L(y)$, by Lemma 5.3, the set $\mathcal{L}_{x,\gamma}$ mixes. Further, $\mathcal{L}_{x,\gamma} \cap \mathcal{B}_1 \neq \emptyset$ and $\mathcal{L}_{x,\gamma} \cap \mathcal{L}_{v,c} \neq \emptyset$; thus, $\mathcal{B}_1 \cup \mathcal{L}_{v,c}$ mixes, and we are done. ☐

Theorem 5.3. *Let G be 3-regular with connectivity 3. If L is a 3-assignment for G, then G is L-swappable unless either (a) $G = K_4$ or (b) $G = K_2 \square K_3$ and $L(v) = L(w)$ for all $v, w \in V(G)$.*

Proof. Let L be a 3-assignment for G. If there exists $v, w \in V(G)$ such that $L(v) \neq L(w)$, then G is L-swappable by Lemma 5.14 (or Lemma 5.16 when $G = K_4$)[6]. Otherwise, $L(v) = L(w)$ for every $v, w \in V(G)$. If $G \neq K_2 \square K_3$, then by Theorem A [40], G is L-swappable. ☐

[6] Although we have not yet proved this lemma, its proof is independent of the present theorem, so invoking it now is logically consistent. We make this choice to preserve the narrative flow of Section 5.3.

Figure 5.13: (a) A 3-assignment for $N(v) \cup N(w_1)$ when $L(w_1) = L(z_1) = L(z_2)$ and $L(v) = L(w_2) = L(w_3)$. (b) The first instance, along a u, v-path, of a consecutive pair x, y with distinct lists.

Lemma 5.15. *Let G be 4-connected, k-regular, but not a clique. Let L be a k-assignment for G. If there exist $v, w_1 \in V(G)$ with $vw_1 \in E(G)$ and $L(w_1) \neq L(v)$, then G is L-swappable.*

Proof. By the list version of Brooks' Theorem [37], G has an L-coloring; that is, $\mathcal{L} \neq \emptyset$. Denote the neighbors of v by $w_1, \ldots, w_k$. By Lemma 5.3, for each $i \in [k]$ and each $\alpha \in L(w_i) \setminus L(v)$, the set $\mathcal{L}_{w_i,\alpha}$ mixes. By Lemma 5.4(2), for all distinct $j, \ell \in [k]$ and each $\beta \in L(w_j) \cap L(w_\ell)$, the set $\mathcal{L}_{w_j,\beta} \cap \mathcal{L}_{w_\ell,\beta}$ mixes. Let $\mathcal{L}_1 := \cup_{i\in[k]} \cup_{\alpha\in L(w_i)\setminus L(v)} \mathcal{L}_{w_i,\alpha}$. Let $\mathcal{L}_2 := \cup_{j,\ell\in[k],j\neq\ell} \cup_{\beta\in L(w_j)\cap L(w_\ell)} \mathcal{L}_{w_j,\beta} \cap \mathcal{L}_{w_\ell,\beta}$. By Pigeonhole, for $\mathcal{L}_1$, $\mathcal{L}_2$
every $\varphi \in \mathcal{L}$, either $\varphi \in \mathcal{L}_1$ or $\varphi \in \mathcal{L}_2$, or both. So $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$.

We will often want to L-color a small set of vertices, S, and show that our coloring φ of S extends to an L-coloring of G. This is possible whenever $S \subset N(v)$ and $|S| \leq 3$ and $|(\cup_{x\in S}\varphi(x)) \cap L(v)| < |S|$. After coloring S, we greedily color $G - S$ in order of nonincreasing distance from v. This uses that $G - S$ is connected, since G is 4-connected and $|S| < 4$. And the same argument shows that all such L-colorings (with a fixed coloring of S) mix.

Case 1: $|L(w_1)\setminus L(v)| \geq 2$. Note, for each $\alpha \in L(w_1)\setminus L(v)$, that $\mathcal{L}_{w_1,\alpha} \neq \emptyset$. Further, $\cup_{\alpha\in L(w_1)\setminus L(v)}\mathcal{L}_{w_1,\alpha}$ mixes, by Lemma 5.3. Suppose there exists $i \in [k] \setminus \{1\}$ with $L(w_i) \setminus L(v) \neq \emptyset$. Similar to above, $\cup_{\beta\in L(w_i)\setminus L(v)}\mathcal{L}_{w_i,\beta}$ mixes. Further, given $\beta \in L(w_i) \setminus L(v)$, there exists $\alpha \in L(w_1) \setminus L(v)$ such that $\alpha \neq \beta$. Since $\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_i,\beta} \neq \emptyset$, we see that $\mathcal{L}_1$ mixes.

Fix $i, j \in [k]$ with $w_iw_j \notin E(G)$ and $L(w_i) \cap L(w_j) \neq \emptyset$; fix $\alpha \in L(w_i) \cap L(w_j)$. Recall that $\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha}$ mixes, by Lemma 5.4(1). Suppose that $i, j \in [k] \setminus \{1\}$. Fix $\alpha \in L(w_i) \cap L(w_j)$. Now there exists $\beta \in L(w_1) \setminus (L(v) \cup \{\alpha\})$. Thus, there exists an L-coloring φ with $\varphi(w_i) = \varphi(w_j) = \alpha$ and $\varphi(w_1) = \beta$, by the remark before Case 1. Note that $\varphi \in \mathcal{L}_{w_1,\beta} \cap (\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha})$. Thus, $\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha}$ mixes with $\mathcal{L}_1$ for each such choice of i, j, and $\alpha \in L(w_i) \cap L(w_j)$.

Finally, suppose there exists $i \in [k] \setminus \{1\}$ with $\alpha \in L(w_1) \cap L(w_i)$ and $w_1 w_i \notin E(G)$. Let $G' := G - vw_1$. Let $L'(w_i) := \{\alpha\}$, $L'(w_1) = \{\alpha\} \cup (L(w_1) \setminus L(v))$, and let $L'(x) := L(x)$ for all other $x \in V(G')$. Note that G' has an L'-coloring and all L'-colorings of G' mix, by Lemma 5.1. Further, the L'-colorings φ' of G' are in bijection with L-colorings φ of G with $\varphi(w_i) = \alpha$ and $\varphi(w_1) \in \{\alpha\} \cup (L(w_1) \setminus L(v))$. This latter set includes a coloring in $\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_i,\alpha}$ and also a coloring in $\mathcal{L}_1$. Thus, $\mathcal{L}_1 \cup \mathcal{L}_2$ mixes; that is, $\mathcal{L}$ mixes.

By symmetry among the w_i's, we henceforth assume that $|L(w_i) \setminus L(v)| \le 1$ for all $i \in [k]$.

Case 2: $|L(w_i) \setminus L(v)| \le 1$ for all $i \in [k]$ and $|(L(w_1) \cup L(w_2)) \setminus L(v)| \ge 2$. As above, $\cup_{\alpha \in L(w_i) \setminus L(v)} \mathcal{L}_{w_i,\alpha}$ mixes, for each $i \in [k]$, by Lemma 5.3. By hypothesis, there exist $\alpha \in L(w_1) \setminus L(v)$ and $\beta \in L(w_2) \setminus L(v)$ with $\alpha \neq \beta$. So there exists $\varphi \in \mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\beta}$. Further, for every $i \in [k]$ with $L(w_i) \neq L(v)$, there exists $\gamma \in L(w_i) \setminus L(v)$ such that either (a) $\mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_i,\gamma} \neq \emptyset$ (and $\gamma \neq \alpha$) or (b) $\mathcal{L}_{w_2,\beta} \cap \mathcal{L}_{w_i,\gamma} \neq \emptyset$ (and $\gamma \neq \beta$). Thus, $\mathcal{L}_1$ mixes.

Now instead fix any distinct $i, j \in [k]$ with $w_i w_j \notin E(G)$. For each $\alpha \in L(w_i) \cap L(w_j)$, recall that $\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha}$ mixes by Lemma 5.4(1). For each such i, j there exists $\alpha \in L(w_i) \cap L(w_j)$ and such a coloring φ either with $\varphi(w_1) \notin L(v)$ or with $\varphi(w_2) \notin L(v)$, unless $\{i, j\} = \{1, 2\}$. However, in this case, $\bigcup_{\alpha \in L(w_1) \cap L(w_2)} \mathcal{L}_{w_1,\alpha} \cap \mathcal{L}_{w_2,\alpha} \cup \bigcup_{\beta \in L(w_1) \setminus L(v)} \mathcal{L}_{w_1,\beta}$ mixes, by Lemma 5.4(2). Thus, $\mathcal{L}_1 \cup \mathcal{L}_2$ mixes; that is, $\mathcal{L}$ mixes.

By symmetry among the w_i's, we now assume $|(L(w_i) \cup L(w_j)) \setminus L(v)| \le 1$ for all $i, j \in [k]$.

Case 3: $|\cup_{i \in [k]} L(w_i) \setminus L(v)| = 1$. Let $\{\alpha\}$ denote $\cup_{i \in [k]} L(w_i) \setminus L(v)$. Note that $\mathcal{L}_{w_h,\alpha}$ mixes whenever $h \in [k]$ and $\alpha \in L(w_h)$, by Lemma 5.3. Since G is not a clique, there exist $i, j \in [k]$ such that $w_i w_j \notin E(G)$. For each such i, j and $\beta \in L(w_i) \cap L(w_j)$, the set $\mathcal{L}_{w_i,\beta} \cap \mathcal{L}_{w_j,\beta}$ mixes, by Lemma 5.4(1). Now we show that $\mathcal{L}_{w_h,\alpha}$ mixes with $\mathcal{L}_{w_i,\beta} \cap \mathcal{L}_{w_j,\beta}$ for all such h, i, j, α, and β. First suppose that $\alpha \neq \beta$. If $h \notin \{i, j\}$, then there exists an L-coloring φ with $\varphi(w_h) = \alpha$ and $\varphi(w_i) = \varphi(w_j) = \beta$. If $h \in \{i, j\}$, then this claimed mixing follows from Lemma 5.4(2). So assume instead that $\alpha = \beta$. If $h \in \{i, j\}$, then this is trivial, since $\mathcal{L}_{w_i,\beta} \cap \mathcal{L}_{w_j,\beta} \subseteq \mathcal{L}_{w_h,\alpha}$. So assume $h \notin \{i, j\}$. Pick $\beta' \in L(w_i) \cap L(w_j) \setminus \{\alpha\}$. There exists an L-coloring φ with $\varphi(w_h) = \alpha$ and $\varphi(w_i) = \varphi(w_j) = \beta'$. But now we are again done by Lemma 5.4(2). More specifically, $\mathcal{L}_{w_h,\alpha}$ mixes and $\mathcal{L}_{w_i,\alpha}$ mixes. By Lemma 5.4(2), also $\mathcal{L}_{w_i,\beta} \cap (\mathcal{L}_{w_j,\beta} \cup \mathcal{L}_{w_j,\alpha})$ mixes. Since $\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha} \subseteq \mathcal{L}_{w_i,\alpha}$, we are done. $\square$

The next lemma handles the case that $G = K_{k+1}$. The main ideas in the proof are similar to those in the previous proof, but the details are a bit different.

Lemma 5.16. *Let $G = K_{k+1}$, where $k \geq 3$, and let L be a k-assignment. If there exist $v, w \in V(G)$ such that $L(v) \neq L(w)$, then G is L-swappable.*

Proof. We denote the vertices of G by $v_1, \ldots, v_{k+1}$, and we consider three cases.

Case 1: There exist $i, j \in [k+1]$ such that $|L(v_i) \setminus L(v_j)| \geq 2$. Fix $\varphi \in \mathcal{L}$. By symmetry, assume that $i = 1$ and $j = 2$. Since $|L(v_2)| = k$ and $|V(G)| = k+1$, there exists $\ell \in [k+1]$ such that $\varphi(v_\ell) \notin L(v_2)$. Let $\alpha := \varphi(v_\ell)$. Let $\mathcal{L}_1 := \cup_{\beta \in L(v_1) \setminus L(v_2)} \mathcal{L}_{v_1,\beta}$. By Lemma 5.3, we know that $\mathcal{L}_{v_\ell,\alpha}$ mixes and also $\mathcal{L}_1$ mixes. We will construct an L-coloring φ' such that $\varphi' \in \mathcal{L}_1 \cap \mathcal{L}_{v_\ell,\alpha}$. This will prove that $\mathcal{L}_1$ mixes with $\mathcal{L}_{v_\ell,\alpha}$. Since $\varphi \in \mathcal{L}_{v_\ell,\alpha}$ and φ is an arbitrary L-coloring, we conclude that G is L-swappable. Let $\varphi'(v_\ell) = \alpha$, color v_1 from $L(v_1) \setminus (L(v_2) \cup \{\beta\})$, and thereafter color greedily, finishing with v_2.

Case 2: There exist distinct $i, j, \ell \in [k+1]$ and distinct colors α_i, α_j such that $\alpha_i \in L(v_i) \setminus L(v_\ell)$ and $\alpha_j \in L(v_j) \setminus L(v_\ell)$. We can assume that Case 1 does not hold, so $|L(w) \setminus L(x)| \leq 1$ for all $w, x \in V(G)$. By symmetry, we assume that $i = 1$, $j = 2$, and $\ell = 3$. Let $\mathcal{L}_1 := \mathcal{L}_{v_1,\alpha_1}$ and $\mathcal{L}_2 := \mathcal{L}_{v_2,\alpha_2}$. By Lemma 5.3, note that $\mathcal{L}_1$ mixes and $\mathcal{L}_2$ mixes. We construct $\varphi' \in \mathcal{L}_1 \cap \mathcal{L}_2$. Let $\varphi'(v_1) := \alpha_1$, let $\varphi'(v_2) := \alpha_2$, and thereafter color greedily, finishing with v_3. Fix an arbitrary L-coloring φ. Now it suffices to show that either φ is L-equivalent to some L-coloring in $\mathcal{L}_1$ or φ is L-equivalent to some L-coloring in $\mathcal{L}_2$. Since $|L(v_3)| = k$ and $|V(G)| = k+1$, there exists h such that $\varphi(v_h) \notin L(v_3)$. Further, either $\varphi(v_h) \neq \alpha_1$ or $\varphi(v_h) \neq \alpha_2$; by symmetry, assume the former. Now, as in the previous case, there exists an L-coloring φ' such that $\varphi' \in \mathcal{L}_1 \cap \mathcal{L}_{v_h,\varphi(v_h)}$. Since $\mathcal{L}_{v_h,\varphi(v_h)}$ mixes, by Lemma 5.3, we are done.

Case 3: There exist $i, j, \ell \in [k+1]$ with $i \notin \{j, \ell\}$ and a color α such that $\alpha \in (L(v_j) \cap L(v_\ell)) \setminus L(v_i)$. We can assume neither Case 1 nor Case 2 holds. So, $|L(v_i) \cap L(v_j)| = |L(v_i) \cap L(v_\ell)| = k-1$. If $j = \ell$, i.e., there is only one vertex v_j such that $\alpha \in L(v_j)$ (implying that $L(v_l) = L(v_i)$ for all $l \in [k+1] \setminus \{j\}$), then every coloring φ lies in $\mathcal{L}_{v_j,\alpha}$. This is because $|L(v_i)| = k$ and $|V(G)| = k+1$, which implies that some vertex (namely v_j) has $\varphi(v_j) \notin L(v_i)$. By Lemma 5.3, $L_{v_j,\alpha}$ mixes. So, we conclude that G is L-swappable. Thus, we assume that $j \neq \ell$.

By symmetry, assume that $i = 1, j = 2$, and $\ell = 3$. If $L(v_2) = L(v_3)$, then pick $\beta \in L(v_1) \setminus L(v_2)$. By Lemma 5.3, each of $\mathcal{L}_{v_1,\beta}, \mathcal{L}_{v_2,\alpha}$, and $\mathcal{L}_{v_3,\alpha}$ mixes. Moreover, $\mathcal{L}_{v_1,\beta} \cap \mathcal{L}_{v_2,\alpha} \neq \emptyset$ (color v_1 with β, color v_2 with α, then color greedily, finishing with v_3). So, $\mathcal{L}_{v_1,\beta}$ mixes with $\mathcal{L}_{v_2,\alpha}$. By symmetry between v_2 and v_3, the set $\mathcal{L}_{v_1,\beta}$ also mixes with $\mathcal{L}_{v_3,\alpha}$. Further, if $\alpha \in L(v_h)$ for some $h \in [k+1] \setminus \{2, 3\}$, then $\mathcal{L}_{v_h,\alpha}$ mixes by Lemma 5.3. And as above, $\mathcal{L}_{v_h,\alpha} \cap \mathcal{L}_{v_1,\beta} \neq \emptyset$, which implies that $\mathcal{L}_{v_h,\alpha}$ mixes with $\mathcal{L}_{v_1,\beta}$. Since $|L(v_1)| = k$ and $|V(G)| = k+1$, for every L-coloring φ there exists $h \in [k+1]$ such that $\varphi(v_h) \notin L(v_1)$. In particular,

$\varphi(v_h) = \alpha$. So, $\varphi \in \mathcal{L}_{v_h,\alpha}$, and we conclude that G is L-swappable. Thus, we assume $L(v_2) \neq L(v_3)$.

By symmetry, assume $L(v_1) = [k], L(v_2) = (\{\alpha\} \cup [k]) \setminus \{1\}$, and $L(v_3) = \{\alpha\} \cup [k-1]$. By Lemma 5.3, each of $\mathcal{L}_{v_1,k}, \mathcal{L}_{v_1,1}, \mathcal{L}_{v_2,\alpha}, \mathcal{L}_{v_2,k}, \mathcal{L}_{v_3,\alpha}$, and $\mathcal{L}_{v_3,1}$ mixes. Moreover, $\mathcal{L}_{v_1,k} \cap \mathcal{L}_{v_2,\alpha} \neq \emptyset$ (color greedily, finishing with v_3). Similarly, $\mathcal{L}_{v_1,k} \cap \mathcal{L}_{v_3,1} \neq \emptyset, \mathcal{L}_{v_1,1} \cap \mathcal{L}_{v_2,k} \neq \emptyset$, and $\mathcal{L}_{v_1,1} \cap \mathcal{L}_{v_3,\alpha} \neq \emptyset$. Thus, $\mathcal{L}_{v_1,k}$ mixes with each of $\mathcal{L}_{v_2,\alpha}$ and $\mathcal{L}_{v_3,1}$; similarly, $\mathcal{L}_{v_1,1}$ mixes with each of $\mathcal{L}_{v_2,k}$ and $\mathcal{L}_{v_3,\alpha}$. By coloring greedily, finishing with v_1, there exists $\varphi \in \mathcal{L}_{v_2,\alpha} \cap \mathcal{L}_{v_3,2}$. And a $2,\alpha$-swap in φ at v_2 gives a coloring in $\mathcal{L}_{v_3,\alpha}$. Thus, $\mathcal{L}_{v_2,\alpha}$ mixes with $\mathcal{L}_{v_3,\alpha}$. Finally, if $\alpha \in L(v_l)$ for some $l \in [k+1] \setminus \{2,3\}$, then $\mathcal{L}_{v_l,\alpha}$ mixes by Lemma 5.3. And, $\mathcal{L}_{v_l,\alpha}$ mixes with $\mathcal{L}_{v_1,1}$ since $\mathcal{L}_{v_l,\alpha} \cap \mathcal{L}_{v_1,1} \neq \emptyset$. As above, for every L-coloring φ, there exists $h \in [k+1]$ such that $\varphi(v_h) = \alpha$. Thus, we conclude that G is L-swappable. □

Lemma 5.17. *Let G be a 4-connected graph that is k-regular, but does not contain an induced 4-wheel, W_4, and is not K_{k+1}. If $L(v) = [k]$ for all $v \in V(G)$, then G is L-swappable.*

Proof. Let G satisfy the hypothesis. Fix an arbitrary vertex $v \in V(G)$ and denote $N(v)$ by $\{w_1, \ldots, w_k\}$. By Pigeonhole, for every $\varphi \in \mathcal{L}$ there exist distinct $i, j, \alpha \in [k]$ with $w_iw_j \notin E(G)$ and $\varphi(w_i) = \varphi(w_j) = \alpha$. Since $L(x) = [k]$ for all $x \in V(G)$, for all distinct i, j such that $w_iw_j \notin E(G)$ and all distinct $\alpha, \beta \in [k]$, the sets $\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha}$ and $\mathcal{L}_{w_i,\beta} \cap \mathcal{L}_{w_j,\beta}$ mix with each other; we simply use α, β-swaps at w_i and w_j. So, in this proof, let $\mathcal{L}_{i,j} := \cup_{\alpha \in [k]}(\mathcal{L}_{w_i,\alpha} \cap \mathcal{L}_{w_j,\alpha})$. For convenience, when $w_iw_j \in E(G)$, let $\mathcal{L}_{i,j} := \emptyset$. So $\mathcal{L}_{i,j}$
$\mathcal{L} = \cup_{i,j \in [k], i \neq j} \mathcal{L}_{i,j}$.

By Lemma 5.4, if $w_iw_j \notin E(G)$, then $\mathcal{L}_{i,j}$ mixes. Fix distinct $i, j, \ell \in [k]$ such that $w_iw_j, w_jw_\ell \notin E(G)$. We show that $\mathcal{L}_{i,j}$ mixes with $\mathcal{L}_{j,\ell}$. Let $G' := G - vw_j$, let $L'(w_i) := \{1\}$, $L'(w_\ell) = \{2\}$, $L'(w_j) := \{1, 2\}$, and $L'(x) := [k]$ for all $x \in V(G) - \{w_i, w_j, w_\ell\}$. Let $G'' := G - \{w_i, w_j, w_\ell\}$. Note that G'' is connected, since G is 4-connected. Let σ'' be a vertex order of $V(G'')$ by non-increasing distance from v, and let $\sigma' := w_i, w_\ell, w_j, \sigma''$. Now σ' shows that $\mathcal{L}'$ mixes, since each vertex x is preceded by fewer than $|L'(x)|$ neighbors in σ'.

Now consider distinct $h, i, j, \ell \in [k]$ such that $w_hw_i, w_jw_\ell \notin E(G)$. We must show that $\mathcal{L}_{h,i}$ mixes with $\mathcal{L}_{j,\ell}$. If $w_hw_j \notin E(G)$, then $\mathcal{L}_{h,i}$ mixes with $\mathcal{L}_{h,j}$ and $\mathcal{L}_{h,j}$ mixes with $\mathcal{L}_{j,\ell}$, as above, so we are done. Assume instead that $w_hw_j \in E(G)$. By symmetry, also $w_iw_j, w_iw_\ell, w_\ell w_h \in E(G)$. However, now G contains an induced 4-wheel, a contradiction. □

Now we combine the previous four lemmas to completely handle the 4-connected case.

Theorem 5.4. *Let G be 4-connected. If G is k-regular and not K_{k+1}, then G is k-swappable. If $G = K_{k+1}$, L is a k-assignment, and L is not identical everywhere, then G is L-swappable.*

Proof. If G contains W_4 as an induced subgraph, then we are done by Lemma 5.11 and Lemma 5.6. So assume it does not. Fix a k-assignment L for G. If there exist $x, y \in V(G)$ such that $L(x) \neq L(y)$, then G is L-swappable, by Lemma 5.15 or Lemma 5.16. So assume instead that $L(x) = L(y)$ for all $x, y \in V(G)$. By symmetry, we assume that $L(v) = [k]$ for all $v \in V(G)$. Now we are done by Lemma 5.17. □

Bibliography

[1] S. Alamdari et al., *How to morph planar graph drawings*, in: SIAM J. Comput. 46.2 (2017), pp. 824–852.

[2] M. O. Albertson, *You can't paint yourself into a corner*, in: J. Combin. Theory Ser. B 73.2 (1998), pp. 189–194.

[3] A. Altshuler, *Construction and enumeration of regular maps on the torus*, in: Discrete Math. 4 (1973), pp. 201–217.

[4] K. Appel and W. Haken, *The solution of the four-color-map problem*, in: Sci. Amer. 237.4 (1977), pp. 108–121, 152.

[5] L. Bailey et al., "Complexity and Enumeration in Models of Genome Rearrangement", in: *Computing and Combinatorics*, vol. 14422, Lecture Notes in Comput. Sci. Springer, Cham, 2024, pp. 3–14.

[6] V. Bartier, N. Bousquet, and M. Heinrich, *Recoloring graphs of treewidth 2*, in: Discrete Math. 344.12 (2021), Paper No. 112553, 7.

[7] V. Bartier et al., *Recolouring planar graphs of girth at least five*, in: (2021).

[8] M. Bonamy, V. Delecroix, and C. Legrand–Duchesne, *Kempe changes in degenerate graphs*, in: European Journal of Combinatorics (2023), p. 103802.

[9] M. Bonamy, V. Delecroix, and C. Legrand–Duchesne, *Kempe changes in degenerate graphs*, in: European Journal of Combinatorics (2023), p. 103802.

[10] M. Bonamy et al., *Diameter of colorings under Kempe changes*, in: Theoret. Comput. Sci. 838 (2020), pp. 45–57.

[11] M. Bonamy et al., *On a conjecture of Mohar concerning Kempe equivalence of regular graphs*, in: J. Combin. Theory Ser. B 135 (2019), pp. 179–199.

[12] M. Bonamy et al., *Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs*, in: J. Comb. Optim. 27.1 (2014), pp. 132–143.

[13] P. Bonsma and L. Cereceda, *Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances*, in: Theoret. Comput. Sci. 410.50 (2009), pp. 5215–5226.

[14] N. Bousquet, *A Note on the Complexity of Graph Recoloring*, in: arXiv:2401.03011 (2024).

[15] N. Bousquet and M. Heinrich, *A polynomial version of Cereceda's conjecture*, in: J. Combin. Theory Ser. B 155 (2022), pp. 1–16.

[16] N. Bousquet and G. Perarnau, *Fast recoloring of sparse graphs*, in: European J. Combin. 52 (2016), pp. 1–11.

[17] S. Cambie, W. Cames van Batenburg, and D. W. Cranston, *Optimally reconfiguring list and correspondence colourings*, in: European J. Combin. 115 (2024), Paper No. 103798, 20.

[18] J. Cardinal et al., *Reconfiguration of satisfying assignments and subset sums: easy to find, hard to connect*, in: Theoret. Comput. Sci. 806 (2020), pp. 332–343.

[19] L. Cereceda, "Mixing graph colourings", in: PhD thesis, London School of Economics and Political Science, 2007.

[20] L. Cereceda, J. van den Heuvel, and M. Johnson, *Connectedness of the graph of vertex-colourings*, in: Discrete Math. 308.5-6 (2008), pp. 913–919.

[21] L. Cereceda, J. van den Heuvel, and M. Johnson, *Finding paths between 3-colorings*, in: J. Graph Theory 67.1 (2011), pp. 69–82.

[22] L. Cereceda, J. van den Heuvel, and M. Johnson, *Mixing 3-colourings in bipartite graphs*, in: European J. Combin. 30.7 (2009), pp. 1593–1606.

[23] G. J. Chaitin, *Register Allocation & Spilling via Graph Coloring*, in: SIGPLAN Not. 17.6 (1982), pp. 98–101.

[24] D. Chakraborty, C. Feghali, and R. Mahmoud, *Kempe Equivalent List Colorings Revisited*, in: arXiv:2211.16728 (2022).

[25] K. L. Collins and J. P. Hutchinson, "Four-coloring six-regular graphs on the torus", in: *Graph colouring and applications (Montréal, QC, 1997)*, vol. 23, CRM Proc. Lecture Notes, Amer. Math. Soc., Providence, RI, 1999, pp. 21–34.

[26] D. W. Cranston, *Kempe equivalent list edge-colorings of planar graphs*, in: Discrete Math. 346.11 (2023), Paper No. 113180, 13.

[27] D. W. Cranston, *List-recoloring of sparse graphs*, in: European Journal of Combinatorics 105 (2022), p. 103562.

[28] D. W. Cranston and L. Rabern, *Brooks' theorem and beyond*, in: J. Graph Theory 80.3 (2015), pp. 199–225.

[29] E. D. Demaine et al., "Coordinated motion planning: reconfiguring a swarm of labeled robots with bounded stretch", in: *34th International Symposium on Computational Geometry*, vol. 99, LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2018, Art. No. 29, 15.

[30] Q. Deschamps et al., *Strengthening a theorem of Meyniel*, in: SIAM J. Discrete Math. 37.2 (2023), pp. 604–611.

[31] R. Diestel, *Graph theory*, Fifth, vol. 173, Graduate Texts in Mathematics, Springer, Berlin, 2018, pp. xviii+428.

[32] Z. Dvořák and C. Feghali, *A Thomassen-type method for planar graph recoloring*, in: European J. Combin. 95 (2021), Paper No. 103319, 12.

[33] Z. Dvořák and C. Feghali, *An update on reconfiguring 10-colorings of planar graphs*, in: Electron. J. Combin. 27.4 (2020), Paper No. 4.51, 21.

[34] Z. Dvořák, D. Král, and R. Thomas, "Coloring triangle-free graphs on surfaces. Extended abstract", in: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, Philadelphia, PA, 2009, pp. 120–129.

[35] Z. Dvořák and L. Postle, *Correspondence coloring and its application to list-coloring planar graphs without cycles of lengths 4 to 8*, in: J. Combin. Theory Ser. B 129 (2018), pp. 38–54.

[36] M. Dyer et al., *Randomly coloring sparse random graphs with fewer colors than the maximum degree*, in: Random Structures & Algorithms 29.4 (), pp. 450–465.

[37] P. Erdős, A. L. Rubin, and H. Taylor, "Choosability in graphs", in: *Proceedings of the West Coast Conference on Combinatorics, Graph Theory and Computing (Humboldt State Univ., Arcata, Calif., 1979)*, Congress. Numer., XXVI, Utilitas Math., Winnipeg, Man., 1980, pp. 125–157.

[38] C. Feghali, *Kempe equivalence of 4-critical planar graphs*, in: J. Graph Theory 103 (2023), pp. 139–147.

[39] C. Feghali, *Reconfiguring colorings of graphs with bounded maximum average degree*, in: J. Combin. Theory Ser. B 147 (2021), pp. 133–138.

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

[41] D. Fotakis et al., “Frequency assignment in mobile and radio networks”, in: *Networks in Distributed Computing*, 1997.

[42] N. Garg, M. Papatriantafilou, and P. Tsigas, “Distributed list coloring: how to dynamically allocate frequencies to mobile base stations”, in: *Proceedings of SPDP ’96: 8th IEEE Symposium on Parallel and Distributed Processing*, 1996, pp. 18–25.

[43] T. A. Gibson, *Strengthening strong immersions with Kempe chains*, in: arXiv:1805.03729 (2018).

[44] R. A. Hearn and E. D. Demaine, *PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation*, in: Theoret. Comput. Sci. 343.1-2 (2005), pp. 72–96.

[45] T. Ito et al., “On the complexity of reconfiguration problems”, in: *Algorithms and computation*, vol. 5369, Lecture Notes in Comput. Sci. Springer, Berlin, 2008, pp. 28–39.

[46] M. Jerrum, *A very simple algorithm for estimating the number of k-colorings of a low-degree graph*, in: Random Structures Algorithms 7.2 (1995), pp. 157–165.

[47] M. Jerrum, *Counting, sampling and integrating: algorithms and complexity*, Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 2003, pp. xii+112.

[48] L. Kiefer, M. Storath, and A. Weinmann, *Iterative Potts Minimization for the Recovery of Signals with Discontinuities from Indirect Measurements: The Multivariate Case*, in: Found Comput Math 21 (3 2021), pp. 649–694.

[49] M. Krasnytska, B. Berche, and Y. Holovatch, *Phase transitions in the Potts model on complex networks*, in: Condensed Matter Physics 16 (2013), p. 23602.

[50] M. Las Vergnas and H. Meyniel, *Kempe classes and the Hadwiger conjecture*, in: J. Combin. Theory Ser. B 31.1 (1981), pp. 95–104.

[51] D. Marx, *GRAPH COLOURING PROBLEMS AND THEIR APPLICATIONS IN SCHEDULING*, in: Periodica Polytechnica Electrical Engineering 48 (2004), pp. 11–16.

[52] A. Merlin and H. Back, “Search for a minimal-loss operating spanning tree configuration in an urban power distribution system”, in: *Proc. 5th Power System Computation Conf., Cambridge, UK*, vol. 5, 1975, pp. 1–18.

[53] H. Meyniel, *Les* 5*-colorations d’un graphe planaire forment une classe de commutation unique*, in: J. Combin. Theory Ser. B 24.3 (1978), pp. 251–257.

[54] B. Mohar, “Kempe equivalence of colorings”, in: *Graph theory in Paris*, Trends Math. Birkhäuser, Basel, 2007, pp. 287–297.

[55] B. Mohar and J. Salas, *A new Kempe invariant and the (non)-ergodicity of the Wang-Swendsen-Kotecký algorithm*, in: J. Phys. A 42.22 (2009), pp. 225204, 32.

[56] B. Mohar and J. Salas, *On the non-ergodicity of the Swendsen–Wang–Kotecký algorithm on the kagomé lattice*, in: Journal of Statistical Mechanics: Theory and Experiment 2010 (2010), p. 05016.

[57] J. Narboni, *Vizing's edge-recoloring conjecture holds*, in: arXiv:2302.12914 (2023).

[58] S. Negami, *Uniqueness and faithfulness of embedding of toroidal graphs*, in: Discrete Math. 44.2 (1983), pp. 161–180.

[59] T. Rokicki et al., *God's Number is 20*, URL: https://cube20.org/.

[60] W. J. Savitch, *Relationships between nondeterministic and deterministic tape complexities*, in: J. Comput. System Sci. 4 (1970), pp. 177–192.

[61] V. G. Vizing, *On an estimate of the chromatic class of a p-graph*, in: Diskret. Analiz 3 (1964), pp. 25–30.

[62] W. Wang and X. Liu, "List-coloring based channel allocation for open-spectrum wireless networks", in: *VTC-2005-Fall. 2005 IEEE 62nd Vehicular Technology Conference, 2005.* Vol. 1, 2005, pp. 690–694.

[63] D. B. West, *Introduction to graph theory*, Prentice Hall, Inc., Upper Saddle River, NJ, 1996, pp. xvi+512.

[64] F. Y. Wu, *The Potts model*, in: Rev. Modern Phys. 54.1 (1982), pp. 235–268.

[65] H.-G. Yeh and X. Zhu, "4-colorable 6-regular toroidal graphs", in: vol. 273, 1-3, EuroComb'01 (Barcelona), 2003, pp. 261–274.