

Grötzsch's Theorem

by

Nabiha Asghar

An essay
presented to the University of Waterloo
in fulfillment of the
research essay requirement for the degree of
Master of Mathematics
in
Combinatorics & Optimization

Waterloo, Ontario, Canada, 2012

© Nabiha Asghar 2012

I hereby declare that I am the sole author of this essay. This is a true copy of the essay, including any required final revisions, as accepted by my examiners.

I understand that my essay may be made electronically available to the public.

Abstract

Grötzsch's Theorem is one of the most famous theorems in graph colouring theory. Its original proof, given in German, in 1958, was fairly complex. In 1989, Steinberg and Younger [17] gave the first correct proof, in English, of the dual version of this theorem. This essay studies the Steinberg-Younger proof in detail, putting special emphasis on improved presentation of their arguments and clarity of exposition. It also gives a new, much simpler proof that is inspired by Carsten Thomassen's [19], but is due to an unpublished work of C. Nunes da Silva, R.B. Richter and D. Younger.

Acknowledgements

First and foremost, I thank Allah Almighty for giving me the ability and strength to complete this task successfully.

My advisor, Dr. R. Bruce Richter, you gave me much-needed encouragement and confidence. Working under you has made me an improved thinker and a writer, and I have only taken on a little from your own excellent teaching and research skills. Dr. Penny Haxell, thank you for reading through this essay and giving your valuable feedback.

My parents and sisters, I am grateful to you for keeping me steady throughout, and for always praying for me. Farrukh Bhai, Najma Apa, Areej and Arslan, I am indebted to you for life.

My friends, you know who you are, keep the support coming. All of you have unknowingly brightened up so many of my days here. No thank you can be enough for you.

Two people, who never lost hope in me, I pray that I become worthy of knowing you.

Dedication

This work is dedicated to Ammi, Abba, Aalia Apa and Sadia Apa.

Table of Contents

List of Figures	vii
1 Introduction	1
1.1 Flows on Graphs	2
1.2 Graph Colourings	5
1.3 The Flow-Colouring Duality	8
1.4 History of Grötzsch's Theorem	9
2 Steinberg-Younger Proof of Grötzsch's Theorem	13
2.1 Preliminaries & Main Theorem	13
2.2 Proof Outline	14
2.3 Complete Proof of Theorem 2.1.1	15
3 A new proof of Grötzsch's Theorem	37
3.1 Preliminaries	37
3.2 Proof of Grötzsch's Theorem	39
References	56

List of Figures

1.1	A graph with a NWZ 5-flow.	3
1.2	A 3-colourable graph and a list-colouring of a graph.	7
1.3	An auxiliary graph used in the construction of G . It contains 8 vertices; every vertex other than P has a 3-list available to it.	11
1.4	The interior of one 4-cycle of G_a containing a copy of PE , and an example of a invalid list-colouring.	12
2.1	Reducing a 2-vertex.	16
2.2	Forming G^{11} from G^1 (Case 2). The squiggly edges lie in the 3-cut of G^2 . This example shows exactly two edges in the 3-cut of G^2	19
2.3	Different cases in the reduction of a digon.	21
2.4	Reducible scenarios in Case 1 when a 4-vertex d is a cut-vertex in G' . H, J are subgraphs in G' such that $H \cap J = \{d\}$ with $v' \in H$ and $v'' \in J$. In both cases, the dashed line shows a nonperipheral 4-cut contained in G	23
2.5	Forming G' from G to reduce a 4-valent distinguished vertex d	24
2.6	Reducing a triangle that contains exactly two 3-vertices.	25
2.7	Reducing a facial triangle containing exactly one 3-vertex.	26
2.8	A zigzag in a planar embedding of G . It consists of edges (t, u) , (u, v) and (v, w) . Vertices u, v are distinct from d and d^*	28
2.9	Reducing a 6-cut that contains a zigzag. The side S^2 corresponds to the segment G^2 of G by the 6-cut c	29
2.10	Graph G with a Grötzsch Configuration, and the reduced and contracted graphs G', G^* respectively.	34

2.11	$m3$ -orientations of the segment of G' by δY that contains vertex v , for all possible sets of directions on δY . The directions inherited from G^* and z' are shown with solid arrows.	36
3.1	Example depicting the reduction of a map M that satisfies the 5-0 hypothesis and does not contain any $(S \cup T)$ - $N(S)$ chords or S - $(S \cup T)$ 2-paths.	45
3.2	Reducing chords with one end in $S \cup T$ in a map M that satisfies the 3-1 hypothesis. Here, $a, b, c \in S$, (d, e) is the T -adjacency and $d, e \in \tau$. Each dashed line is a chord with one end in $S \cup T$ and the other end $y \notin S \cup T$. Chords (y, c) and (y, f) separate S and τ ; then M_2 is the side containing τ . Chords $(y, g), (y, a), (y, b)$ do not separate S and τ ; then M_2 is the side that does not contain τ . Note that if (y, g) is present, then neither (y, a) nor (y, b) can have minimal M_2 . On the other hand, if (y, f) is present, then (y, c) does not have minimal M_2	46
3.3	Example depicting the reduction of a map M that satisfies the 3-1 hypothesis Case 1 and does not contain any chords with one end in $S \cup T$, or $(S \cup T)$ - $(T \cup N_C(\tau))$ 2-paths.	50
3.4	Reducing separating 4-cycles in M	54

Chapter 1

Introduction

According to [13], in 1852, the English mathematician and botanist Francis Guthrie posed what is now called the Four Colour Problem. He noticed that at most four colours were needed to colour the map of the counties of England, such that any two counties sharing a common boundary segment (but not just a point) did not receive the same colour. He wondered if this was a coincidence. In 1854, the problem was published in *Athenaeum* [7], a famous literary journal of London, and caught the attention of other mathematicians. A formally stated version of the problem became known as the 4-Colour Theorem.

4-Colour Theorem: Every planar graph has a 4-colouring.

The 4-Colour Theorem was extensively studied in its original form as well as in its dual version, the 4-Flow Conjecture (see, e.g. [14]) for planar graphs. While attempting to prove it, several other significant and closely related results were produced. One of them established that maps satisfying certain constraints could be coloured with at most three colours. This result was given by the renowned German mathematician Herbert Grötzsch in 1958 and is now known as Grötzsch's Theorem [8]. Its equivalent, dual version is the 3-Flow Theorem [16].

This essay is a self-contained discourse on Grötzsch's Theorem and its dual, the 3-Flow Theorem. Chapter 1 introduces the notions of flows in graphs, graph colourings and the flow-colouring duality in Sections 1.1, 1.2 and 1.3 respectively. Sections 1.4 and 1.5 briefly discuss the history and motivation of Grötzsch's Theorem. Chapter 2 provides a detailed study of the proof of 3-Flow Theorem given by Steinberg and Younger [17] in 1989. This is the first correct proof in English of Grötzsch's Theorem (in dual form), and it extends the result to include projective planar graphs. Chapter 3 gives a direct proof of Grötzsch's

Theorem that is inspired by Carsten Thomassen’s work [19], but is due to C. Nunes da Silva, R.B. Richter and D. Younger (unpublished).

1.1 Flows on Graphs

In this section, we introduce flows in graphs and discuss some fundamental properties and results about them.

Let $G = (V, E)$ be an undirected graph. An orientation of G is an assignment to each edge e of G a direction from its root $\sigma(e)$ to its terminal $\tau(e)$, so $\sigma(e)$ and $\tau(e)$ are the two vertices of G incident with e . An integer flow on G is a pair (D, f) , where D is an orientation of G and f is an integer valued function $f : E \rightarrow \mathbb{Z}$ such that, for any $v \in V$,

$$\sum_{e: v=\sigma(e)} f(e) = \sum_{e: v=\tau(e)} f(e). \quad (1.1)$$

Equation (1.1) is called the Flow Conservation Property and it says that, for any vertex in G , the sum of flows on all its incident edges must equal zero.

Informally, a flow on an undirected graph G consists of an assignment D of directions to each edge in G , and a function f which assigns a numeric value to each directed edge such that, for any vertex v , the total ‘incoming flow’ (that is, the sum of the flows on edges directed into v) equals the total ‘outgoing flow’ (that is, the sum of the flows on edges directed out of v).

The net-flow at a set S of vertices is defined to be
$$\sum_{\sigma(e) \in S, \tau(e) \notin S} f(e) - \sum_{\sigma(e) \notin S, \tau(e) \in S} f(e).$$

The net-flow along a set F of edges is defined to be
$$\sum_{e \in F} f(e).$$

In this essay, we consider very specific kinds of flows, called nowhere-zero k -flows. For an undirected graph $G = (V, E)$, a nowhere-zero (NWZ) flow is a flow (D, f) where $f : E \rightarrow \mathbb{Z} \setminus \{0\}$. A NWZ k -flow is a flow (D, f) where $f : E \rightarrow \{\pm 1, \pm 2, \dots, \pm(k - 1)\}$. Figure 1.1 shows an example of a NWZ 5-flow on a graph.

An isthmus, also called a 1-cut, is an edge whose deletion increases the number of connected components of a graph. It is known that any graph containing an isthmus cannot have a NWZ k -flow for any k , so G must be 2-edge-connected in order to have a NWZ flow. Moreover, any 2-edge-connected graph has a k -flow for at least one value of k . An interesting graph theoretic problem is to find the smallest positive integer k , for an




Figure 1.1: A graph with a NWZ 5-flow.

arbitrary undirected graph $G = (V, E)$, such that G has a NWZ k -flow (D, f) . This is equivalent to finding a NWZ flow (D', f') on G , where $f' : E \rightarrow \{1, 2, 3, \dots, k-1\}$, because the orientation D can be adjusted by reversing the direction of any edge that is assigned a negative flow to obtain D' . This problem is commonly referred to as ‘The k -Flow Problem’. In fact, according to [3], the problem of deciding whether a graph has a k -flow for any given integer k is in NP^1 . For $k = 3$ in particular, this problem is NP-complete^2 , even when G is planar. However, several results have been published over the years for specific values of k . They illustrate the relationship between a graph’s edge-connectivity and the existence of NWZ k -flows for various values of k . We list some of the famous results below to provide a historical context.

For $k = 2$ and 3 , the following are two of the earliest known results (see, e.g. [5]).

2-Flows: A graph has a NWZ 2-flow if and only if every vertex has even degree.

3-Flows: A 2-edge-connected, cubic³ graph has a NWZ 3-flow if and only if it is bipartite.

For $k = 3, 4$ and 5 , Tutte proposed the following well-known conjectures. They are all still open. The first one is equivalent to Grötzsch’s Theorem without the assumption of planarity. Its first appearance in print seems to be via Tutte’s student Steinberg [16].

Tutte’s 3-Flow Conjecture: Every 4-edge-connected graph has a NWZ 3-flow.

Tutte’s 4-Flow Conjecture [20]: Every 2-edge-connected graph without a Peterson minor has a NWZ 4-flow.

¹ NP is the class of decision problems such that, if the answer to the problem is ‘yes’, then there exists a proof of this fact and the proof can be verified in polynomial time.

²A decision problem is NP-complete if it is in NP , and any problem in NP can be reduced to it.

³A cubic graph is a graph in which every vertex has degree 3.

Tutte’s 5-Flow Conjecture [20]: Every 2-edge-connected graph has a NWZ 5-flow.

In 1979, François Jaeger gave two nice results regarding NWZ 4- and 8-flows.

Jaeger’s 4-Flow Theorem [11]: Every 4-edge-connected graph has a NWZ 4-flow.

Jaeger’s 8-Flow Theorem [11]: Every 2-edge-connected graph has a NWZ 8-flow.

The latter was improved by Paul Seymour in 1995.

Seymour’s 6-Flow Theorem [15]: Every 2-edge-connected graph has a NWZ 6-flow.

Recently, Thomassen has given the following result regarding NWZ 3-flows.

Thomassen’s 3-Flow Theorem [18]: Every 8-edge-connected graph has a NWZ 3-flow.

In 1950, Tutte provided a very useful tool to prove the existence of a NWZ k -flow on a graph. He introduced the idea of the ‘NWZ (mod k)-flow’. A NWZ (mod k)-flow f is similar to a NWZ k -flow, except that

$$\sum_{e: v=\sigma(e)} f(e) \equiv \sum_{e: v=\tau(e)} f(e) \pmod{k} \text{ for any } v \in V.$$

The following is a not very difficult theorem due to Tutte.

Tutte’s Flow Theorem (see, e.g. [5]): A graph has a NWZ k -flow if and only if it has a NWZ (mod k)-flow.

Since this essay is about Grötzsch’s Theorem, the focus is on NWZ 3-flows and NWZ (mod 3)-flows for planar graphs. Note that in a NWZ (mod 3)-flow, the edges are assigned flow-values from the set $\{-2, -1, 1, 2\}$. However, $-2 \equiv 1 \pmod{3}$ and $2 \equiv -1 \pmod{3}$, so we can change -2 and 2 to 1 and -1 respectively. Also, a -1 on an edge can be changed to a 1 by reversing the direction of the edge. Thus, if there is a NWZ (mod 3)-flow, there is one in which all the flow-values are 1. In this sense, the flow-values can be completely captured by the edge directions. More specifically, the edge directions record all the information: the number of incoming edges at a vertex is the same as the number of outgoing edges at that vertex modulo 3. Hereon, we refer to NWZ (mod 3)-flows as ‘(mod 3)-orientations’, or $m3$ -orientations in short.

Let D be an orientation of a graph G . The directional dual of D is the orientation \overleftarrow{D} (or $-D$) obtained from D by reversing the direction of every edge of G . It is a simple matter to verify that D is an $m3$ -orientation if and only if \overleftarrow{D} is an $m3$ -orientation.

The following lemma contains some key concepts from the theory of flows that will be used in proofs throughout this essay. For that, we require the notions of coboundaries and

cuts. A set c of edges in G is a coboundary in G if there is a set $X \subseteq V(G)$ such that $c = \delta X$, where δX is the set of all edges that have exactly one end in X . A cut in G is a minimal non-null coboundary. A k -cut is a cut of size k .

Lemma 1.1.1: Let $G = (V, E)$ be an undirected graph, and let $F = (D, f)$ be a NWZ k -flow on G . Then the following statements are true:

- (i) Vertex identifications preserve the Flow Conservation Property at every vertex.
- (ii) In F , the net-flow along the edges of any cut is zero.
- (iii) For any function $g : E \rightarrow \{1, 2, \dots, k - 1\}$, if (1.1) holds for all vertices except one vertex u , then it holds for u as well.
- (iv) If G contains a 1-cut, no NWZ k -flow exists on G for any value of k .

Proof:

- (i) Let $u, v \in V$. Obtain a new graph $G' = (V', E')$ by identifying u and v to a new vertex u^* . In G' , the net-flow at each vertex except u^* is the same as that in G (that is, zero), and the net-flow at u^* is the sum of the flows on edges incident to u and to v in G . If $e = (u, v) \notin E$, this sum is equal to the sum of net-flows at u and at v , so it is zero. If $e \in E$, e can be thought of as a loop at u^* in G' ; it contributes equally to the incoming flow and outgoing flow at u^* . Hence e 's absence in G' means that the net-flow at u^* is still the sum of net-flows at u and at v in G , so it is zero.
- (ii) Identify one side of the cut c to a single vertex u^* . By (i), the net-flow at u^* , that equals the net-flow along the edges of c , is zero.
- (iii) Identify $G - u$ to a single vertex v . Since (1.1) holds for each vertex in $G - u$, it holds for v by (i), so the net-flow at v is zero. This implies that the net-flow at u is zero.
- (iv) Let e be the edge in the 1-cut and let $F = (D, f)$ be a NWZ k -flow on G for some value of k . By (ii), $f(e) = 0$. This contradicts the fact that F is nowhere-zero. Therefore, F does not exist for any value of k . □

1.2 Graph Colourings

In this section, we introduce graph colourings and discuss some of their important properties.

The subject of graph colouring was born with the Four Colour Problem [7] in 1852. The problem was to determine whether the countries in any map could be coloured with at most four colours, provided that no two countries sharing a common boundary segment (and not

just a single point) received the same colour. The problem is translated to graph theory by representing the map with a graph as follows. Each vertex in the graph represents a country and an edge between two vertices represents the common border between the two corresponding countries. This gives us an instance of what may be called the ‘4-Vertex-Colouring Problem’: the goal is to assign colours to vertices such that no two adjacent vertices receive the same colour and no more than four colours are used.

A vertex-colouring (or a colouring, in short) of a graph $G = (V, E)$ is a function $c : V \rightarrow \mathbb{N}$ such that $c(u) \neq c(v)$ for all $(u, v) \in E$. Simply put, it is an assignment of colours (or labels) to the vertices of G such that no two adjacent vertices are assigned the same colour. A k -colouring of G is a (vertex) colouring of G with at most k colours: it is a function $c : V \rightarrow \{0, 1, 2, \dots, k - 1\}$ such that $c(u) \neq c(v)$ for all $(u, v) \in E$. If G has a k -colouring, then G is k -colourable. There is an analogous notion of face-colouring in planar graphs. The faces of a planar graph embedded in the plane are domains obtained by deleting from the plane the arcs and points representing the edges and vertices respectively of the graph. A face-colouring of an embedded graph G is an assignment of colours to the faces of G such that no two faces containing a common edge in their boundaries are assigned the same colour. A k -face-colouring of G is a face-colouring of G with at most k colours. If G has a k -face-colouring, then G is k -face-colourable. There also exists a notion of edge-colouring. An edge-colouring of a graph G is an assignment of colours to the edges of G such that no two edges with a common end-vertex are assigned the same colour. In this essay, we restrict ourselves to the notions of vertex-colourings and face-colourings only.

There is a ‘duality’ between vertex-colouring and face-colouring, and for that, we need to understand the notion of the dual of a graph. Let $G = (V, E)$ be a planar graph embedded in the plane. The dual graph $G' = (V', E')$ of G has a vertex v_f for each face of G , and an edge e' for each edge e of G , so that, if e is incident with the (possibly identical) faces f_1 and f_2 , then e' is incident with v_{f_1} and v_{f_2} . Observe that a k -vertex-colouring of a planar graph is a k -face-colouring of its dual graph; this is the duality between vertex-colouring and face-colouring. A loop in G is a 1-cut in G' and vice versa. Also, a cycle⁴ in a graph corresponds to a cut in the dual graph.

One of the main goals of colouring theory is to determine whether an arbitrary graph is k -colourable for a given positive integer k . It is worthwhile to note straight away that if G has a loop, it does not have any colouring because the end-points of the loop are always assigned the same colour. Equivalently, in the dual version, a graph containing a 1-cut does not have any face-colouring. It is easy to prove that all planar graphs are 6-colourable (see, e.g. [4]). Heawood proved in 1890 that, in fact, they are 5-colourable [10]. The most

⁴A cycle is a connected non-null graph in which each vertex has degree 2.

challenging conjecture, however, was the 4-Colour Theorem.

4-Colour Theorem (see, e.g. [13]): Every planar graph has a 4-colouring.

After various unsuccessful attempts at its proof by renowned mathematicians like Kempe and Heawood, the proof was finally obtained in the 1970s [2] with substantial help from computers. This was the first proof of its kind.




Figure 1.2: A 3-colourable graph and a list-colouring of a graph.

A related concept in graph theory is the list-colouring of a graph, which was introduced by Erdős, Rubin and Taylor [6]. It is a type of (vertex) colouring where each vertex is assigned a list of colours and a vertex can receive a colour only from its own list. More formally, assign to each vertex $v \in G$ a set $L(v) \subseteq \mathbb{N}$. An L -colouring of G is a colouring $c : G \rightarrow \mathbb{N}$ such that $c(v) \in L(v)$ for all $v \in G$. If G has an L -colouring, then G is L -colourable. A graph G is k -list-colourable, or k -choosable, if, for every list assignment L with, for every vertex v , $|L(v)| \geq k$, there is an L -colouring of G . It is easy to see that if G is k -list-colourable, it is k -colourable (one simply assigns the same list of k colours to each vertex). Figure 1.2 illustrates a 3-colourable graph and a list colouring of a graph.

1.3 The Flow-Colouring Duality

In this section, we describe the duality between NWZ flows on graphs and graph colourings. This duality was established by Tutte in 1954, in one of his seminal works.

Tutte Flow-Colouring Duality Theorem [20]: Let G be a 2-edge-connected graph embedded in the plane. Then G has a NWZ k -flow if and only if it is k -face-colourable (that is, its dual is k -colourable).

Proof: Let (D, f) be a k -flow on $G = (V, E)$. We need to construct a function $c : \mathcal{F} \rightarrow \{0, 1, \dots, k - 1\}$, where $\mathcal{F} = \text{faces}(G)$, such that $c(F_1) \neq c(F_2)$ for any two adjacent faces F_1 and F_2 . Assign an arbitrary colour, say 0, to the infinite face. Let F be some face in G , other than the infinite face. We assign a colour $c(F)$ to it as follows. Choose a point a in the infinite face and a point b in F . Consider any directed, open curve⁵ P from a to b that does not pass through any vertex of G and so that P meets each edge at most once. Let $c_P(F)$ be the sum of flows at the edges crossed by P such that, while moving along the direction of P , an edge crossing P from left to its right, relative to the orientation of P , contributes $f(e)$ to the sum and an edge crossing P from right to its left contributes $-f(e)$. Choose $c_P(F)$ such that it is the least non-negative integer modulo k . Finally, let $c(F) = c_P(F)$.

This method of face-colouring is well defined and not dependent on the choice of P . If Q is any other directed, open curve from a to b that does not pass through any vertex of G and $c_Q(F)$ is the sum of flows at the edges crossed by Q , then P and Q' (that is, the reverse of the path Q) together form a directed, closed curve⁶ S . Let $d_S(F)$ be the sum of flows at the edges crossed by S . Then $d_S(F) = c_P(F) + c_{Q'}(F) = c_P(F) - c_Q(F)$. For each region R enclosed by S , let V_R be the set of vertices in G that are in R . There are three kinds of edges that cross the boundary of R (and hence they cross S).

- (i). If $V_R = \emptyset$, an edge that enters R must also leave R . Thus, each such edge crosses the boundary of R twice.
- (ii). If $V_R \neq \emptyset$, there may be edges that have both endpoints in R , yet they cross the boundary of R . Each such edge crosses R 's boundary twice.
- (iii). If $V_R \neq \emptyset$, each edge of the cut separating R and $V - R$ crosses the boundary of R once.

Observe that $d_S(F)$ equals the sum of contributions made by edges of type (i), (ii) and (iii). Edges of type (i) and (ii) contribute zero to $d_S(F)$ because the incoming flow on

⁵An open curve is a continuous image of the closed interval $[0,1]$, with 0 and 1 mapping to distinct points.

⁶A closed curve is a continuous image of the unit circle.

each such edge balances the outgoing flow on that edge. Consider type (iii) edges now. The net-flow N_R along the edges of the cut separating R from $V - R$ is zero, as shown in Section 1.1. So, type (iii) edges contribute zero to $d_S(F)$. Thus, $d_S(F) = 0$, which implies $c_P(F) = c_Q(F)$. So, the colouring is well-defined.

We show that the colouring of G 's faces obtained this way is a k -face-colouring as follows. By construction, $0 \leq c(F) < k$ for any face F . Let F and F' be two adjacent faces of G , sharing an edge e . Let A be a point in the infinite face, B a point in F and B' a point in F' . Let P be a directed open curve from A to B that does not cross e and let P' be a directed open curve from A to B' that extends P to meet B' such that e is the only additional edge crossed by P' . So, $c_{P'}(F') = c_P(F) \pm f(e)$. Since $f(e) \neq 0$, $c_{P'}(F') \neq c_P(F)$, and hence F and F' are assigned distinct colours.

Conversely, let $c : \mathcal{F} \rightarrow \{0, 1, \dots, k - 1\}$ be a k -face-colouring of G . We construct a function $f : E \rightarrow \{\pm 0, \pm 1, \dots, \pm(k - 1)\}$ as follows. Assign a direction to an edge e such that the colour value c_L of the face on e 's left, relative to e 's orientation, is greater than the colour value c_R of the face on its right. Let the flow-value at e be $f(e) = c_L - c_R$. We show that this procedure gives a NWZ k -flow on G by showing that flow is conserved at each vertex. Since the face-colouring is proper, $f(e) \neq 0$ and $f(e) < k$ for any edge e . Let $d = \text{degree}(v)$ for some $v \in V(G)$ and let the colours of the faces incident to v be c_1, c_2, \dots, c_d in anti-clockwise order. For each $i \pmod d$, the flow in to v on the edge e_i incident to c_i and c_{i+1} is $c_i - c_{i+1}$ (if this value is positive, the edge is directed into v , otherwise it is directed out of v). The net-flow at v equals $\sum_{i=0}^{d-1} (c_i - c_{i+1})$, where $c_0 = c_d$. This sum equals zero. \square

1.4 History of Grötzsch's Theorem

We now give a history of Grötzsch's Theorem and a list of its known proofs.

Grötzsch's Theorem was given by the renowned German mathematician Herbert Grötzsch in 1958.

Grötzsch's Theorem [8]: Every loop-free and triangle⁷-free planar graph is 3-colourable.

Since K_4 is planar and is not 3-colourable, the girth constraint (at least 4) in Grötzsch's Theorem is the best possible. There exist infinitely many planar graphs that are not 3-colourable because they contain K_4 as a subgraph. Grötzsch's proof was very complex and in 1960, Berge published a simpler proof which turned out to be erroneous. In 1963,

⁷A triangle is a cycle of length 3.

Branko Grünbaum presented an erroneous proof of the following extension of Grötzsch's Theorem, which was corrected by Vasili Aksionov in 1974.

Grünbaum-Aksionov [9], [1]: Every planar graph with at most three triangles is 3-colourable.

Their arguments were largely based on those given by Grötzsch in the original proof. In 1989, Steinberg and Younger presented the first correct proof in English of Grötzsch's Theorem; they proved the dual form of Grünbaum-Aksionov Theorem and extended it to the projective plane.

Steinberg-Younger's 3-Flow Theorem [17]: Every planar graph without 1-cuts and with at most three 3-cuts has a NWZ 3-flow. Moreover, every graph without 1-cuts and with at most one 3-cut that can be embedded in the projective plane has a NWZ 3-flow.

Fourteen years later, Carsten Thomassen presented a simpler proof of Grötzsch's Theorem using list-colourings. His approach was complicated, but ingenious.

Thomassen [19]: Every planar graph with girth at least five is 3-list-colourable. Moreover, every triangle-free planar graph is 3-colourable.

It should be noted that Grötzsch's Theorem cannot be extended from a 3-colourable graph to 3-list-colourable graph. As a counterexample, Margit Voigt [21] presented the construction of a planar graph with girth 4, that is not 3-list-colourable. We present Voigt's construction below.

We describe the construction of a planar, triangle-free graph G , with a list-assignment L , in several steps and then show that it is not L -colourable. The symbols $1, 2, 3, \dots$ and a, b, c, \dots are used to denote colours.

We start with three copies G_1, G_2, G_3 of the graph in Figure 1.3. Except for the vertex P , the list assignments on the vertices of G_a are given in the figure. For each of the four 4-cycles in G_a that contain a copy of PE , we shall attach a copy of the graph in Figure 1.4 (a), with the 4-cycle in G_a being identified with the 4-cycle in the figure. The list assignments given to the new vertices in Figure 1.4 (a) are determined by b and c , which are chosen as follows:

1. for the cycles containing PEF , $b = 6$; the one with H has $c = 4$, while the one with K has $c = 5$;
2. for the cycles containing PEC , $b = 7$; the one with D has $c = 4$, while the one with B has $c = 5$.




Figure 1.3: An auxiliary graph used in the construction of G . It contains 8 vertices; every vertex other than P has a 3-list available to it.

The resulting graph is G_a^* .


Finally the graph G is obtained from G_1^* , G_2^* , and G_3^* by identifying the three copies of P and giving P the list $\{1, 2, 3\}$.

Claim 1.4.1: G is planar, triangle-free and not L -colourable.

Proof: Suppose by way of contradiction that G has an L -colouring. Clearly, G is planar and triangle-free by construction. By symmetry, we may suppose P is coloured 1. We now restrict our attention to G_1^* .

There is a left-right symmetry in G_1 (interchanging 6 and 7), so we may assume E is coloured 6. If F is coloured 4, then H is coloured 6, while if F is coloured 5, then K is coloured 6. These two cases are symmetric, so we assume F is coloured 4 and H is coloured 6. We now look at the subgraph of G_1^* obtained from identifying the cycle $PEFH$ with the outer cycle in the graph in Figure 1.4 (a). We know $a = 1$, $b = 6$, and $c = 4$, and the outer cycle has P coloured 1, E coloured 6, X coloured 4, and Y coloured 6.

For both the 5-cycles in G_1^* , the three vertices joined to the outer 4-cycle containing P are coloured either 8 or 9; it makes no difference which of the two choices is made for one. In either case, one of the remaining two vertices on each 5-cycle is coloured either 4 or 5. Now, one of the following two cases must occur: either both 4's are used (which is equivalent to using both 5's) or one 4 and one 5 is used. The latter gives a contradiction at their common neighbour. The former requires the two remaining uncoloured, adjacent



(a) A copy of the graph that is attached to each of the four 4-cycles in G_a that contain a copy of PE , with the 4-cycle in G_a being identified with the 4-cycle in the figure.

(b) A colour assignment that does not work out for $a = 1, b = 6, c = 4$. The vertex labelled ‘?’ can not be assigned any colour.

Figure 1.4: The interior of one 4-cycle of G_a containing a copy of PE , and an example of a invalid list-colouring.

vertices in G_1^* to be coloured 5. So, no matter how we assign colours to the vertices of G_1^* , one of the two 5-cycles is not L -colourable. Thus, G_1^* is not L -colourable and neither is G . Figure 1.4(b) depicts an example of a colour assignment that does not work out. \square

Thus, we have established that the idea of 3-list-colourability can be used to prove Grötzsch’s Theorem, but Grötzsch’s Theorem cannot be generalized to include 3-list-colourability.

Chapter 2

Steinberg-Younger Proof of Grötzsch's Theorem

In this chapter, we present the proof of the dual of Grötzsch's Theorem given by Steinberg and Younger [17]. Their original paper extends the result to include graphs that embed in the real projective plane. We omit this part.

Steinberg-Younger's 3-Flow Theorem: Every planar graph without 1-cuts and with at most three 3-cuts has a NWZ 3-flow.

In this chapter, the goal is to present Steinberg and Younger's proof in complete detail by filling in the gaps that the authors may have left in their arguments intentionally or otherwise. Also, special emphasis is placed on clarity of exposition and improved presentation of arguments. Section 2.1 contains some preliminaries and states the form of the Steinberg-Younger Theorem that was originally proved. Section 2.2 gives a brief sketch of the proof. Section 2.3 gives the complete proof of the theorem.

2.1 Preliminaries & Main Theorem

In this section, we present the form of the Steinberg-Younger Theorem as it is actually proved. We begin with the few definitions and remarks required to do so.

Edges whose deletion disconnects a connected graph play a large role in the theory of flows. Let $G = (V, E)$ be a graph. Recall the definitions of coboundary and cut from Chapter 1. For a subset S of V , a cut $c = \delta S$ is peripheral if either $|S| = 1$ or $|V \setminus S| = 1$ and is nonperipheral otherwise.

We have already seen in Chapter 1 that an isthmus is a 1-cut. It is worthwhile to explore the motivation for excluding 1-cuts and more than three 3-cuts from the premise of the theorem. As explained in Section 1.1, a NWZ 3-flow (or equivalently, a 3-face-colouring) does not exist for graphs (or embedding of graphs) containing a 1-cut. To figure out the extent to which 3-cuts can be excluded, we look for the smallest planar, embedded graph without 1-cuts and 3-face-colourings. The number of 3-cuts in it potentially leads us to the answer. The smallest planar, embedded graph without 1-cuts and 3-face-colourings is K_4 . It has exactly four 3-cuts and this suggests that up to three 3-cuts are allowed in an embedded graph, if we want to obtain its 3-face-colouring.

A vertex incident with precisely k edges is both k -valent and a k -vertex. A distinguished vertex is a 4- or 5-vertex d , not a cut-vertex, at which an $m3$ -orientation is specified; that is, the edges at d are directed so that the net outflow from d is congruent to zero modulo 3. If d is a 5-vertex, one incident edge must oppose the other four in direction; this edge is called the minority edge.

The following is the form of the Steinberg-Younger Theorem that we shall prove.

Theorem 2.1.1: Let G be a planar graph such that either

- i. G has at most three 3-cuts, or
- ii. G has at most one 3-cut and a distinguished vertex d so that, if d is 5-valent, then the minority edge does not lie in a 3-cut.

Then G has an $m3$ -orientation D ; in Case (ii), D extends the $m3$ -orientation at d .

2.2 Proof Outline

This section contains a brief outline of the proof of the Steinberg-Younger Theorem (Theorem 2.1.1).

The proof uses induction on the total number of edges in the graph. The main strategy is to show, through a series of seven lemmas, that, in case of existence of each of the seven configurations listed below, we can contract/delete edges to obtain a smaller graph. The inductive hypothesis gives an $m3$ -orientation of the smaller graph. It is then shown that this $m3$ -orientation can be extended to an $m3$ -orientation of the original graph. The seven configurations are:

1. a cut-vertex;
2. a 2-vertex;
3. a nonperipheral k -cut, $k = 2, 3, 4, 5$;

4. a digon (that is, a cycle of length 2);
5. an m -vertex, $m = 4$ or $m \geq 6$;
6. a triangle containing a 3-vertex in a planar embedding of the graph; and
7. a 6-cut that contains a zigzag.

Extending the $m3$ -orientation of the smaller graph to that of the original graph is straightforward for configurations 1, 2 and 4. The other cases use the idea of segments of the graph, which will be defined later.

If the graph does not contain any of the above configurations, it is shown through a lemma that there must exist an eighth configuration, called a ‘Grötzsch Configuration’, in a planar embedding of the graph. The proof of the lemma makes use of Euler’s Formula. Given a Grötzsch Configuration in the embedded graph, we shrink the graph and apply the inductive hypothesis to obtain an $m3$ -orientation of the smaller graph, which is then shown to extend to an $m3$ -orientation of the original graph. This completes the proof of the Theorem 2.1.1. In particular, it proves the Grünbaum-Aksionov Extension.

2.3 Complete Proof of Theorem 2.1.1

We proceed by induction on $|E(G)|$, the total number of edges in G .

Base case: If G has no edges, then G consists of isolated vertices. Therefore, it satisfies hypothesis (i) and has the trivial $m3$ -orientation. If G has exactly one edge e , then e is not a 1-cut, so e is a loop. Thus, any orientation of G is an $m3$ -orientation.

Inductive Step: Let G be a graph satisfying one of the hypotheses and having more than one edge. It is a triviality that G has an $m3$ -orientation if and only if each component of G has an $m3$ -orientation. Therefore we may assume G is connected. We now present eight lemmas; each of them tackles a configuration that gives rise to an $m3$ -orientation of G . The lemmas are stated in order of priority of reduction of the different configurations. Lemmas 2.3.1 to 2.3.7 below are all in the context of the induction; an unstated hypothesis is that the theorem is true for all planar graphs with fewer edges than G .

Lemma 2.3.1: If G has a cut-vertex, then G has an $m3$ -orientation.

Proof: Let v be a cut-vertex in G . Then there exist two subgraphs H and J of G , each containing at least one edge, such that $H \cap J = \{v\}$. Note that any cut of G is a cut of either H or J , and vice-versa. Thus, each of H and J contains at most the number of 3-cuts in G . This implies that if G satisfies hypothesis (i), so do H and J . If G satisfies hypothesis (ii), then the distinguished vertex is not the cut-vertex (by definition), so exactly

one of H and J satisfies hypothesis (i) and the other satisfies hypothesis (ii). Moreover, each of H and J contains fewer edges than G , so it has an $m3$ -orientation by the inductive hypothesis. The union of these two $m3$ -orientations gives an $m3$ -orientation of G . \square

Lemma 2.3.2: If G has a 2-vertex, then G has an $m3$ -orientation.

Proof: Let v be a 2-vertex in G . Then v is not the distinguished vertex, since it is not a 4- or 5-vertex. If v has exactly one neighbour u , then u is a cut-vertex, so G has an $m3$ -orientation by Lemma 2.3.1. Therefore, we may assume that the two neighbours of v are distinct.

Let α be an edge incident with v whose other end is not d . Contract α to obtain a smaller graph G' (Figure 2.1) that satisfies the same hypothesis as G (contracting α does not increase the number of 3-cuts, and does not disturb d because α is not incident with d). By the inductive hypothesis, G' has an $m3$ -orientation. Transfer this orientation to G . Only α remains without orientation now; the other edge incident to v has a direction. Assign a direction to α which gives net-flow zero at v . This automatically ensures that the other end of α has net-flow zero (if not, we would have a non-zero net-flow at exactly one vertex of the graph, which is not possible). This completes an $m3$ -orientation of G . \square




Figure 2.1: Reducing a 2-vertex.

We give some more definitions. Let G be a graph.

For a cut $c = \delta S$, S is a side of c . In a connected graph, c has just one other side, $V \setminus S$. A side graph of c corresponding to side S is the induced subgraph $G[S]$. Observe that each cut has two side graphs, $G[S]$ and $G[V \setminus S]$. Both the side graphs of a planar graph are planar.

The segment of G by cut c corresponding to side S is obtained from G by contracting $G[V \setminus S]$ to a vertex. Thus, each cut has two segments. The cuts of each segment are the cuts

of G whose edges all belong to that segment. A planar embedding of a graph G determines a set of rotations, where each rotation is a permutation of the edges incident with a vertex v in G . A loop at v occurs twice in the rotation at v . Each rotation can be adjusted to a cyclic permutation. Two edges incident with v are called consecutive edges at v if they are adjacent in the rotation at v . For example, α and β are consecutive at $v = (\alpha, \beta, \gamma, \dots)$.

As a general rule, we can use the segments of a small cut to find an $m3$ -orientation of G by using induction to first obtain an orientation of one segment, and then to obtain an orientation of the other. This is the point of the following lemma.

Lemma 2.3.3: If G has a nonperipheral k -cut for $k = 2, 3, 4, 5$, then G has an $m3$ -orientation.

Proof: Choose the labelling of the two segments G^1 and G^2 of G by c , so that G^2 has at most one 3-cut, and if there is a distinguished vertex d , then d is in G^1 . The main strategy for each value of k is to obtain $m3$ -orientations of G^1 and G^2 , and take their union to get an $m3$ -orientation of G .

Case 1: Nonperipheral 2-, 3-cuts:

Let c be a nonperipheral 2- or 3-cut in G . Then each segment of G by c satisfies one of the two hypotheses: G^1 satisfies the same hypothesis as G (it either contains at most three 3-cuts of G , or the distinguished vertex d and at most one 3-cut of G) and G^2 satisfies hypothesis (i) (it contains at most one 3-cut of G). Since each of G^1 and G^2 has fewer edges than G , it has an $m3$ -orientation by the inductive hypothesis. If D^1 and D^2 are the $m3$ -orientations of G^1 and G^2 respectively, then either they agree on the edges of c , or there is agreement for D^1 and $\overleftarrow{D^2}$. The union of the two agreeing orientations gives an $m3$ -orientation of G .

Case 2: Nonperipheral 4-cuts and nonperipheral 5-cuts with no edge in the 3-cut of G^2 :

Note: We separate the case of nonperipheral 5-cuts into two subcases: (a) no edge in the 5-cut lies in the 3-cut of G^2 , and (b) at least one edge in the 5-cut lies in the 3-cut of G^2 . The motivation to make this bifurcation is that we wish to use the contracted vertex u of G^2 as a distinguished vertex when orienting the edges of G^2 . This is simple to accomplish in (a). In (b), however, we must make sure that the minority edge at u does not lie in the 3-cut in G^2 . This makes the analysis of (a) and (b) different.

Let c be either a nonperipheral 4-cut or a nonperipheral 5-cut with no edge in the 3-cut of G^2 . Observe that G^1 either contains d and at most one 3-cut, or at most three 3-cuts. Therefore, G^1 satisfies the same hypothesis as G and has an $m3$ -orientation by the inductive hypothesis. Transfer the directions assigned to the edges of c to their corresponding edges incident to the contracted vertex u in G^2 . Now u behaves as a distinguished vertex in G^2

and has at most one 3-cut by the definition of G^2 . Moreover, none of its edges is in a 3-cut by assumption. Thus, G^2 satisfies hypothesis (ii). By the inductive hypothesis, G^2 has an $m3$ -orientation that agrees with the $m3$ -orientation of G^1 on the common edges c . The union of these two orientations gives an $m3$ -orientation of G .

Case 3: Nonperipheral 5-cuts with at least one edge in the 3-cut of G^2 :

Let c be one such 5-cut such that its side S^1c , that contains the (perhaps not strict) majority of 3-cuts if G satisfies hypothesis (i), and the distinguished vertex d if G satisfies hypothesis (ii), is minimal. We distinguish two cases.

Subcase 1: $\delta\{d\} \cap c$ contains the minority edge α and a second edge β not in the 3-cut of G^2 . In this case, G^1 has an $m3$ -orientation by the inductive hypothesis (it satisfies the same hypothesis as G). Transfer the directions assigned to the edges of c to their corresponding edges incident to the contracted vertex u in G^2 . Now u functions as the distinguished vertex in G^2 . Since α is the minority edge at d , it is oppositely directed to β , hence one of α and β is the minority edge at u . Note that α does not lie in the 3-cut of G^2 because it is the minority edge at d . Neither does β , by assumption. Therefore, the minority edge at u does not lie in the 3-cut of G^2 . Now, G^2 satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The union of these two orientations gives the $m3$ -orientation of G .

Subcase 2: the remaining case. In this case, let w be the contracted vertex of G^1 . We may assume at this point that 3-cuts are peripheral and there is a 3-cut c' in G^2 . We start with the discussion of c' relative to c . Since c' is peripheral, it is incident with a vertex x in the side S^2c of c that contains at most one 3-cut. Since c is nonperipheral, there is another vertex in S^2c . Since c is a cut, S^2c is connected. Therefore, x has degree at least one in S^2c , so there are at most two edges in $c' \cap c$. Since w has degree 5 in G^1 , in any embedding of G^1 , there are consecutive edges α and β incident with w that are not in c' . We now show that the edges in $c' \cap c$ occur consecutively in any embedding of G^1 . If there is only one edge in $c' \cap c$, this is trivial. If there are two, then x has degree 1 in S^2c . As we contract S^2c to w , think of the last contraction being the one edge in S^2c incident with x . Then the three edges in c from $S^2c - x$ to S^1c are going to occur in the part of the rotation at w (which we may think of as x), in the same place the one edge at x was, so the other two edges (that is, those in $c' \cap c$) are consecutive, as claimed.

Our next aim is to find two w -consecutive edges to pull off of w . We consider two cases. In the first case, the minority edge at d is in c . Then it is the only edge in $\delta\{d\} \cap c$, otherwise we are in Subcase 1. In this case, we choose α to be the minority edge and β to be a w -consecutive edge to α that is not in c' . In the other case, since $c' \cap c$ has at most two edges, there are three w -consecutive edges e_1, e_2, e_3 that are all not in c' . If all three

were incident with the 5-vertex z of S^1c , then moving z out of S^1c produces a 4-cut (by Lemma 2.3.1 there is no loop incident with z) in G . This 4-cut is not peripheral: obviously $S^2c \cup \{z\}$ has at least two vertices, while $S^1c \setminus \{z\}$ has the 5-vertex d and, therefore, some neighbour of d . But now G has an $m3$ -orientation by Case 2. It follows that two consecutive ones of e_1, e_2, e_3 are not both incident with the same 5-vertex of S^1c ; let α and β be two such edges.

Form G^{11} from G^1 (Figure 2.2) by splitting w into a 2-vertex $w' = \{\alpha, \beta\}$ and a three vertex w'' . We proceed to find an $m3$ -orientation of G^{11} and subsequently that of G^1 as follows.




Figure 2.2: Forming G^{11} from G^1 (Case 2). The squiggly edges lie in the 3-cut of G^2 . This example shows exactly two edges in the 3-cut of G^2 .

There is exactly one 3-cut of G not in G^{11} , that is, the 3-cut of G^2 . However, this is balanced by a new 3-cut $\delta\{w''\}$ in G^{11} . In addition, G^{11} may contain some new nonperipheral 3-cuts; each of these must separate w' from w'' , otherwise it would not be new. If no such new 3-cuts are created, G^{11} has exactly the same number of 3-cuts as G (and contains d , if G does). Then G^{11} satisfies the same hypothesis as G and has an $m3$ -orientation by the inductive hypothesis.

On the other hand, if G^{11} contains a nonperipheral 3-cut $\{x, y, z\}$ that separates w' from w'' , it must be the image of the nonperipheral 5-cut $c'' = \{\alpha, \beta, x, y, z\}$ in G^1 . Let S^1c'' be the side of c'' that contains the majority of 3-cuts, in case hypothesis (i) holds for G , and d , if hypothesis (ii) holds for G . Let the other side be S^2c'' . Since S^1c is minimal,

$w \in S^1 c''$. We may assume that this 5-cut contains at least one edge that lies in the peripheral 3-cut of $S^2 c''$, otherwise G has an $m3$ -orientation by Case 1. This implies that $S^1 c$ contains a 3-vertex u of G . So G has at least two 3-cuts, $\delta\{u\}$ and c' , and hence it satisfies hypothesis (i).

Now, G^{11} possibly has more than three nonperipheral 3-cuts, so to obtain an $m3$ -orientation of G^{11} , segment it iteratively by each nonperipheral 3-cut that separates w' from w'' . More specifically, choose such a nonperipheral 3-cut c_1 so that one of its sides is minimal (that is, it does not contain any nonperipheral 3-cuts). Segment G^{11} by c_1 and obtain an $m3$ -orientation of the minimal side graph by the inductive hypothesis, since it satisfies hypothesis (i). Once c_1, \dots, c_i are chosen, choose c_{i+1} so that one of its sides contains c_1, \dots, c_i and no other nonperipheral 3-cuts. The unoriented subgraph of this side graph satisfies hypothesis (i); obtain an $m3$ -orientation of this subgraph that agrees with the oriented edges of this side. Stop when there are no more nonperipheral 3-cuts to choose. This completes an $m3$ -orientation of G^{11} .


The $m3$ -orientation of G^{11} directly gives the $m3$ -orientation of G^1 by restoring vertex w . As before, the directions of edges in c are transferred to their corresponding edges in G^2 , whose contracted vertex functions as the distinguished vertex. Its minority edge is either α or β , neither of which lies in the 3-cut in G^2 , so the minority edge condition is satisfied. Thus, G^2 satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The union of the $m3$ -orientations of G^1 and G^2 gives the $m3$ -orientation of G . \square

Lemma 2.3.4: If G has a digon, then G has an $m3$ -orientation.


Proof: Let u, v be the vertices of a digon (which we may assume is facial in the embedding of G) and α, β its edges. We distinguish two cases (Figure 2.3).

Case 1: neither u nor v is the distinguished vertex. In this case, contract α to get a smaller graph G' . This does not perturb the distinguished vertex, if it exists. Also, contracting an edge does not increase the number of 3-cuts. Therefore G' satisfies the same hypothesis as G . By the inductive hypothesis, G' has an $m3$ -orientation. Transfer this orientation to G . The net-flow is zero modulo 3 at every vertex except maybe u and v . If both u and v are imbalanced, we try to rectify the balance on one of them, say u . At u , the balance can be off by 1 or -1. Giving α the appropriate direction so as to create a balance (modulo 3) at u automatically balances v . On the other hand, if u and v are already balanced, reverse the direction of β to make u imbalanced. This balance is off by either 1 or -1; giving α the appropriate direction balances u and subsequently v too. This completes an $m3$ -orientation of G .


Case 2: u is the distinguished vertex d . In this case, G has at most one 3-cut. We distinguish two subcases.



(a) Case 1: Neither u nor v is the distinguished vertex in G . Contract α to obtain G' .



(b) Case 2, Subcase 1: u is the distinguished vertex d , with α and β oppositely directed at d in G . Delete α, β to obtain G' .



(c) Case 2, Subcase 2: u is the distinguished vertex d , with α and β similarly directed at d in G . Delete β and reverse the direction of α to obtain G' .

Figure 2.3: Different cases in the reduction of a digon.

Subcase 1: α and β are oppositely directed at d . In this subcase, we start by showing that v is not a 3-vertex. If v were a 3-vertex, then recall that d is either a 4-vertex or a 5-vertex. In the former case, we have the contradiction that $\delta\{d, v\}$ is a second 3-cut in G . In the latter case, we have the contradiction that the minority edge (that is, either α or β) is in a 3-cut. Thus, v is not a 3-vertex.

Delete α, β to obtain a smaller graph G' . The 3-vertex d is not the distinguished vertex in G' . Also, G' can have at most three 3-cuts: one that it possibly inherited from G , and either one or both of $\delta\{d\}$ and $\delta\{v\}$. We may assume that there are no nonperipheral 1- or 3-cuts in G' since each would be, with α and β , a nonperipheral 3- or 5-cut in G , in which

case an $m3$ -orientation of G can be obtained by Lemma 2.3.3. Therefore, G' satisfies hypothesis (i) and has an $m3$ -orientation by the inductive hypothesis. This orientation agrees with the required directions at d either directly or by reversing the directions of all the edges in G' . Together with the assigned directions on α and β , this orientation gives an $m3$ -orientation of G .

Subcase 2: α and β are similarly directed at d . In this subcase, delete β and reverse the direction of α to obtain a smaller graph G' . This decreases the degree of each of d and v by one. We can assume that no new nonperipheral 3-cut is created, otherwise the edges in this 3-cut, together with β , give a nonperipheral 4-cut in G and G has an $m3$ -orientation by Lemma 2.3.3. If d becomes a 3-vertex, it is not a distinguished vertex in G' : this shifts the hypothesis from (ii) to (i) and also accomodates a new 3-cut that arises if v becomes a 3-vertex in G' .

On the other hand, if d becomes a 4-vertex, there are two possibilities. One possibility is that v becomes a 3-vertex. Then $\delta\{d, v\}$ is a cut of size at most 5 in G' , and therefore in G too. By Lemma 2.3.3, we may assume $\delta\{d, v\}$ is peripheral. Thus G has exactly three vertices. All the edges in G that are incident with d go to v and a third vertex w . Since d is not a cut-vertex, there is at least one edge between v and w . Thus, at least two and at most three edges join d and v in G . Since d is 5-valent, there are also at least two and at most three edges between d and w . Precisely one of the digons (d, v, d) and (d, w, d) must consist of an oppositely oriented pair, which is Subcase 1.

In the remaining case, v does not become a 3-vertex. Then d remains intact as the distinguished vertex and G' has the same number of 3-cuts as G . So G' satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. This is converted to an $m3$ -orientation of G by restoring β and the direction of α . \square

Lemma 2.3.5: If G has an m -vertex where $m = 4$ or $m \geq 6$, then G has an $m3$ -orientation.

Proof: Let the other end of the minority edge at a distinguished 5-vertex d be called the *distinguished neighbour* d^* . Let G be embedded in the plane. We identify two cases.

Case 1: v is an m -vertex other than d . If $v \neq d^*$, then choose α and β to be any consecutive edges incident with v in G . If $v = d^*$, then Lemma 2.3.4 implies we may assume G has no digon, while Lemma 2.3.3 implies there is no nonperipheral 3-cut. Therefore, at most one edge incident with d^* is in a 3-cut in G . It follows that we may choose α to be the minority edge and β to be an edge consecutive to α at d^* and not in a 3-cut.

Split v into a 2-vertex $v' = (\alpha, \beta)$ and v'' to get a new embedded graph G' . We may assume that G' has no new 1-cuts or 3-cuts. This is because such new cuts must separate v' from v'' (otherwise they would not be new); together with α and β , they correspond to 3-

or 5-cuts respectively in G . If these 3- or 5-cuts are nonperipheral, G has an $m3$ -orientation by Lemma 2.3.3. Otherwise, they are peripheral; in this case G contains a digon and hence, it has an $m3$ -orientation by Lemma 2.3.4. If d is a cut-vertex in G' , then we show that G has an $m3$ -orientation. To see this, its removal must separate v' from v'' , (otherwise it would be a cut-vertex in G , a contradiction). Note there are subgraphs H and J of G' containing v' and v'' , respectively, so that $H \cup J = G'$ and $H \cap J = \{d\}$. Both H and J must contain at least 2 edges incident to d , otherwise G' has a 1-cut, a contradiction.

As the first of two cases (Figure 2.4), suppose d is not a neighbour of v' . Then $\delta\{d\} \cap H$ together with α and β gives a nonperipheral 4- or 5- cut in G , so G has an $m3$ -orientation by Lemma 2.3.3. In the other case, d is a neighbour of v' . Now $\delta\{d\} \cap J$, together with α and β , gives a nonperipheral 3- or 4-cut in G , so G has an $m3$ -orientation by Lemma 2.3.3.

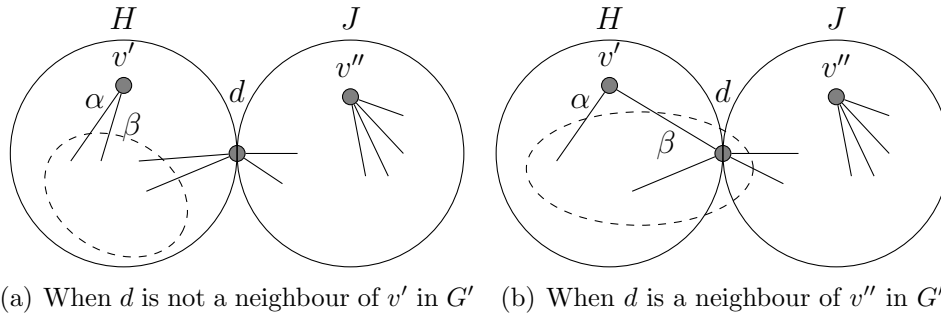


Figure 2.4: Reducible scenarios in Case 1 when a 4-vertex d is a cut-vertex in G' . H, J are subgraphs in G' such that $H \cap J = \{d\}$ with $v' \in H$ and $v'' \in J$. In both cases, the dashed line shows a nonperipheral 4-cut contained in G .

Thus, we may assume d is not a cut vertex in G' and therefore, G' satisfies the same hypothesis as G . Since G' contains a 2-vertex v' , it has an $m3$ -orientation by Lemma 2.3.2; this can be converted to an $m3$ -orientation of G by restoring v', v'' to v .

Case 2: $v = d$. In this case, v has degree 4. By Case 1, we may assume that every other vertex of G has degree either 3 or 5. If d has a neighbour w that is a 3-vertex, then $c = \delta\{d, w\}$ is a 5-cut. If c is nonperipheral, then G has an $m3$ -orientation by Lemma 2.3.3, so assume c is peripheral. Then G contains a digon, and has an $m3$ -orientation by Lemma 2.3.4. Therefore, we may assume that each neighbour of d is a 5-vertex. Form a new graph G' by splitting one such neighbour w of d into $w' = (\alpha, \beta)$ and w'' , where α and β are consecutive edges at w and neither of them is incident with d (Figure 2.5).

Consider the 3-cuts that separate w' from w'' . One is $\delta\{w''\}$. Any other such 3-cut, together with α and β , is a 5-cut in G . If this 5-cut is nonperipheral, G has an $m3$ -




Figure 2.5: Forming G' from G to reduce a 4-valent distinguished vertex d .

orientation by Lemma 2.3.3. So assume that it is peripheral. This implies that α and β make a digon in G , so G has an $m3$ -orientation by Lemma 2.3.4. Therefore, we may assume that $\delta\{w''\}$ is the only 3-cut that separates w' from w'' .

Now $\delta\{d, w''\}$ is a 5-cut in G' . The edges incident with w'' have an enforced orientation, leaving the minority edge of this 5-cut being an edge incident with d . This minority edge is not in any 3-cut because d has no 3-vertex as its neighbour. Segment G' by the 5-cut $\delta\{d, w''\}$. The segment that contains d has d as the distinguished vertex and has exactly one 3-cut: $\delta\{w''\}$. So, it satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The other segment contains the contracted vertex as the distinguished vertex and has at most one 3-cut, which is the one inherited from G . Thus, this segment also satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. Their union is an $m3$ -orientation of G' that is converted to an $m3$ -orientation of G by restoring vertex w . \square

Lemma 2.3.6: If G has a triangle containing a 3-vertex, then G has an $m3$ -orientation. Furthermore, either G has an $m3$ -orientation or, for any triangle u, v, w in G , $\delta(\{u, v, w\})$ is a 9-cut.

Proof: First, we show that such a triangle must be facial in any planar embedding of G . Let u, v, w be the vertices of the triangle, labelled so that u has degree 3. Each of v and w has degree 3 or 5, so $|\delta(\{u, v, w\})|$ is one of 3, 5 and 7. It cannot be 3, since then there are three 3-vertices and a 3-cut, which is more 3-cuts than is allowed. If it is 5, then either it is a non-peripheral 5 cut, in which case, G has an $m3$ -orientation by Lemma 2.3.3, or it is

peripheral; in that case, there is a digon incident with the 5-vertex on the other side and G has an $m3$ -orientation by Lemma 2.3.4.

So $|\delta(\{u, v, w\})| = 7$. Clearly, the side of $\delta(\{u, v, w\})$ that contains u, v, w is connected. If the other side is not connected, then it consists of disjoint subgraphs. Any one of these subgraphs is not incident to exactly 1 or 2 of the 7 edges, because these two situations respectively imply that G has a 1-cut or a 2-cut. Thus, there are exactly two disjoint subgraphs H_1 and H_2 , so that H_1 is incident to 3 of the 7 edges and H_2 is incident to 4. If any of these cuts is nonperipheral, then G has an $m3$ -orientation by Lemma 2.3.3, so assume they are both peripheral. But then H_2 consists of a 4-vertex that is incident with a digon at v or w . So G has an $m3$ -orientation by Lemma 2.3.4.

Therefore, both sides of $\delta(\{u, v, w\})$ are connected. It follows that, in any planar embedding, this triangle is facial.





Figure 2.6: Reducing a triangle that contains exactly two 3-vertices.


We now show that if G has a triangle containing a 3-vertex, then it has an $m3$ -orientation. We distinguish two cases.

Case 1: the edge (w, v) is not the minority edge. In this case, exchange the labels on v and w , if needed, so that the minority edge is not incident with v . Split v into $v' = ((v, w), (u, v))$ and v'' . Call the new embedded graph G' (Figure 2.7). This process may create new cuts in the graph; each of these must separate v' from v'' in G' , otherwise they would not be new.

Any new 1-cut, together with the edges (v, w) and (u, v) , is a 3-cut in G . If this cut is peripheral, u and w are on one side and v is on the other. So, it must be the v side that is just one vertex; but v is 5-valent, a contradiction. Thus, this cut is nonperipheral and G has an $m3$ -orientation by Lemma 2.3.3. So, we may assume that there are no new 1-cuts.



(a) when the minority edge at d is not incident with v . The figure on left shows a facial triangle containing a 3-vertex u in G . The figure on right shows the 4-cut $\delta\{u, v', w\}$ in G' after splitting v .



(b) when (w, v) is the minority edge at $w = d$. The figure on left shows a facial triangle containing a 3-vertex u and a distinguished vertex d with its five oriented edges in G . The figure on right shows the reduced graph G' after deleting edge (u, v) and reversing the directions of (u, w) and (v, w) .

Figure 2.7: Reducing a facial triangle containing exactly one 3-vertex.

If there is a new 3-cut other than $\delta\{v''\}$, then it, together with (v, w) and (u, v) , is a 5-cut in G . This 5-cut cannot be peripheral: in G' the 3-cut cannot have u, w, v'' on one side and just v' on the other. So, G has a nonperipheral 5-cut and an $m3$ -orientation by Lemma 2.3.3. We may, therefore, assume that the only new 3-cut is $\delta\{v''\}$.

Note that $\delta\{u, v', w\}$ is a 4-cut in G' . The segment containing u, v' and w has one 3-cut $\delta\{u\}$. If either G has no distinguished vertex or the distinguished vertex is in the other segment, then orient the other segment first; that segment suppresses the 3-cut $\delta\{u\}$, so it has the same number of 3-cuts as G . Thus, it satisfies the same hypothesis as G and

has an $m3$ -orientation by the inductive hypothesis. Now the segment containing u, v' and w has one 3-cut and the contracted vertex functions as the distinguished vertex, so this segment satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The union of these orientations is an $m3$ -orientation of G' .

In the remaining case, w is the distinguished vertex in G . The segment of G' containing u, v' and w satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The other segment has only one 3-cut, namely $\delta\{v''\}$, and the contracted vertex functions as the distinguished vertex. So, this segment satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. The union of these orientations is an $m3$ -orientation of G' .

Restoring v to G' gives an $m3$ -orientation of G .

Case 2: (v, w) is the minority edge. We choose the labelling so that $w = d$. Delete edge (u, v) to get a smaller embedded graph G' . Reverse the directions of the edges (d, v) and (d, u) ; the edge (d, u) is now the minority edge at d in G' . Suppose c is a 3-cut in G' . If c does not separate u from v , then c is a 3-cut in G . Since the only 3-cut in G is $\delta\{u\}$, we have that $c = \delta\{u\}$. But $\delta\{u\}$ is not a 3-cut in G' , a contradiction. So it must be that c separates u from v . Then $c \cup \{(u, v)\}$ is a 4-cut in G , and consequently, G has an $m3$ -orientation, either by Lemma 2.3.3 (if the cut is nonperipheral) or Lemma 2.3.5 (if the cut is peripheral).

Therefore, G' satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis. In this orientation, one of the two edges at u has u as origin and the other has u as terminal. So in G , both edges have u as either the origin or the terminal (after we unreverse the directions of the edges (d, u) and (d, v)). This determines a unique direction for edge (u, v) and completes an $m3$ -orientation of G .

Finally, we show that if none of the triangles in G contain a 3-vertex, then, for any triangle u, v, w in G , $\delta(\{u, v, w\})$ is a 9-cut. To see this, first observe that each of u, v, w is a 5-vertex. If the side of the 9-coboundary $\delta(\{u, v, w\})$ that does not contain u, v, w , is not connected, then the 9-coboundary must be split in a 3-cut and a 6-coboundary or a 4-cut and a 5-cut. The latter either yields a nonperipheral 4- or 5-cut or a digon, both previously treated in Lemmas 2.3.3 and 2.3.4. The former yields either a nonperipheral 3-cut, which has been treated in Lemma 2.3.3, or a 3-vertex x . Now x is either incident to a digon at u, v or w , in which case, Lemma 2.3.4 yields an $m3$ -orientation of G , or it is adjacent to all of u, v, w . But then x is a 3-vertex in the triangle u, v, x , and G has an $m3$ -orientation by the previous discussion.

It follows that both sides of the 9-coboundary $\delta(\{u, v, w\})$ are connected, implying it is a 9-cut. \square

Our last preliminary reduction requires the introduction of a new notion. A path is a connected graph with two vertices of degree 1 and remaining vertices of degree 2. The vertices of degree 2 are called internal vertices of the path. A zigzag in G is a path Z with three edges so that, in some planar embedding of G , not all three edges are in the boundary of a common face, but with each adjacent pair consecutive at their common vertex. We further require that each internal vertex of Z is 5-valent and distinct from d and d^* (Figure 2.8).




Figure 2.8: A zigzag in a planar embedding of G . It consists of edges $(t, u), (u, v)$ and (v, w) . Vertices u, v are distinct from d and d^* .

Lemma 2.3.7: If G has a zigzag that is contained in a 6-cut, then G has an $m3$ -orientation.

Proof: Let c be a 6-cut in G that contains the three edges of a zigzag Z . Let u, v be the internal vertices of Z and $\alpha = (u, v)$ the middle edge of Z . We choose the labelling of u and v so that v lies in the segment G^2 of G by c , that contains at most one 3-cut if G satisfies hypothesis (i) and does not contain the distinguished vertex if G satisfies hypothesis (ii). We form a new graph G' as follows.

1. If there is no 3-cut in G^2 , let β be the other edge of Z incident with v . Since Z is contained in c , β is also in c . Split v into $v' = (\alpha, \beta)$ and v'' .
2. If there exists a 3-cut in G^2 , let β be the other edge of Z incident with u . Again, β is in c . Split u into $u' = (\alpha, \beta)$ and u'' .

In each case, $c' = c - \{\alpha, \beta\}$ is a nonperipheral 4-cut in G' (Figure 2.9). In the first case, v'' is on the same side that v was for c , while v' is on the other side. In the second case, u'' is on the same side that u was for c , while u' is on the other side. Consider the segment G'^1 of G' by c' that contains the majority (not necessarily strict) of 3-cuts, if G satisfies hypothesis (i), and d , if G satisfies hypothesis (ii). This segment contains at most as many 3-cuts as G does, so it satisfies the same hypothesis as G , and has an $m3$ -orientation by the

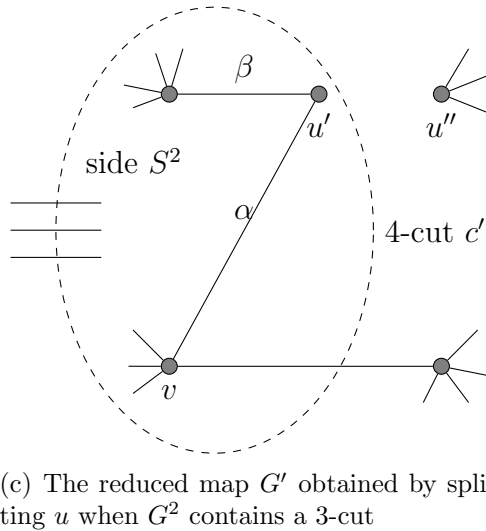
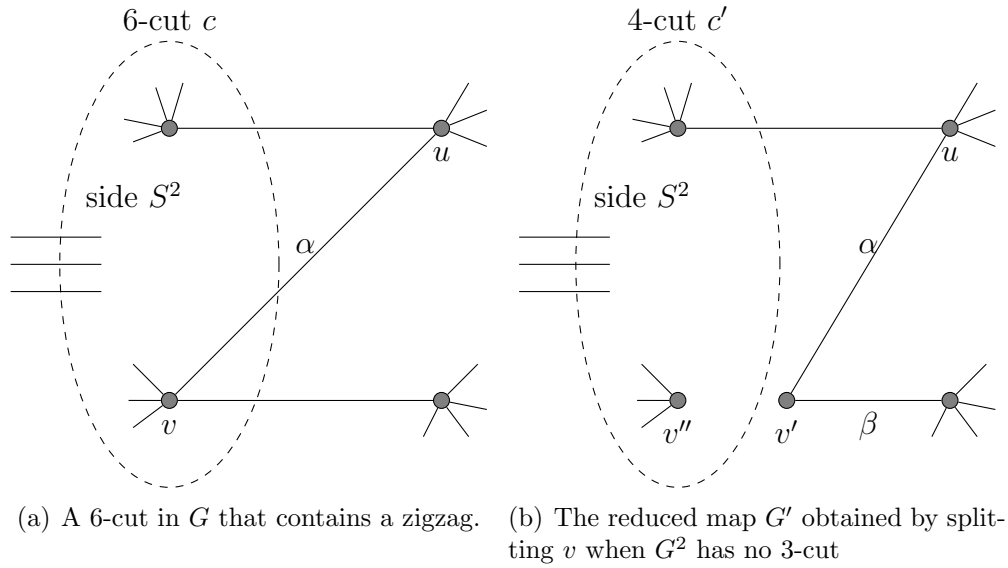


Figure 2.9: Reducing a 6-cut that contains a zigzag. The side S^2 corresponds to the segment G^2 of G by the 6-cut c .

inductive hypothesis. The contracted vertex of G'^1 functions as the distinguished vertex for the other segment G'^2 . Note that G'^2 has exactly one 3-cut (which is $\delta\{v''\}$ in the first case and the one inherited from G in the second case). So, G'^2 satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis, which agrees with that of G'^1 on their

common edges. Taking their union gives an $m3$ -orientation of G' , which is converted to an $m3$ -orientation of G by restoring the vertex we split. \square

We now prove that if a graph satisfying hypothesis (i) or (ii) does not contain any of the seven configurations considered in Lemmas 2.3.1 to 2.3.7, it must contain a ‘Grötzsch Configuration’. The existence lemma and its proof are followed by a method to reduce this configuration. We begin by defining a Grötzsch Configuration.

Grötzsch Configuration

Let R be a 5-wheel with one rim edge deleted. Choose the labelling so that the rim path of R is p, q, r, s, t , while v is the central 5-vertex.

Let G be a graph. A Grötzsch Configuration in G is a subgraph H of G isomorphic to R so that every vertex of H has degree 5 in G , d is not in H , and, furthermore, with the labelling as in the preceding paragraph, d^* is neither q nor s .

Lemma 2.3.6 shows we may assume every triangle in G is facial in any embedding of G . Thus, if G is embedded in the plane, then the embedding of a Grötzsch Configuration is precisely as depicted in Figure 2.10 (a).

To the edges that have exactly one end in $\{v, p, q, r, s, t\}$, we assign labels as follows (these labels are shown in Figure 2.10 (a) as well).

- $\alpha_1, \alpha_2, \alpha_3$ consecutive at p , with α_1 consecutive to the edge (p, v) and α_3 consecutive to the edge (p, q) .
- β_1, β_2 consecutive at q , with β_1 consecutive to the edge (q, p) and β_2 consecutive to the edge (q, r) .
- γ_1, γ_2 consecutive at r , with γ_1 consecutive to the edge (r, q) and γ_2 consecutive to the edge (r, s) .
- δ_1, δ_2 consecutive at s , with δ_1 consecutive to the edge (s, r) and δ_2 consecutive to the edge (s, t) .
- $\epsilon_1, \epsilon_2, \epsilon_3$ consecutive at t , with ϵ_1 consecutive to the edge (t, s) and ϵ_3 consecutive to the edge (t, v) .

Since the embedding of G is planar, only ϵ_3 and α_1 may be equal. All other edges labelled above may be assumed to be distinct. To see this, observe that if one of these edges (other than ϵ_3 and α_1) is incident with two vertices u and v on the rim of R , then there are two possibilities. Either u and v are already adjacent in R , so G has a digon, in which case we are done by Lemma 2.3.4, or they are not adjacent in R , in which case, these two vertices, together with v , give us a non-facial triangle, and we are done by Lemma 2.3.6.

Existence Lemma: Let G be a planar graph with no loop or digon, all vertices of degree 5, except at most three 3-vertices, and no 3-vertex is on a triangle. Then G has a Grötzsch Configuration.

Proof: We first establish the existence of several 5-vertices in G , each incident with four triangle-faces, and then show that at least one of them qualifies to be the vertex v of a Grötzsch configuration. By assumption, the neighbours of each such 5-vertex must be distinct 5-vertices.

Let the number of edges incident with a face f be denoted by $\underline{\text{degree}(f)}$ ¹. Since each vertex of G is either a 3- or 5-vertex,

$$|V| = |V_3| + |V_5| \text{ and } 2|E| = 3|V_3| + 5|V_5|.$$

A k -face is a face of degree k .

Following Lebesgue [12], assign a weight $w(u)$ to each vertex u , where

$$w(u) = \sum_{f \in F_u} \frac{1}{\text{degree}(f)},$$

and F_u is the set of faces incident with u . Note that u may occur several times in the boundary walk of a face f , and there is a summand $\frac{1}{\text{degree}(f)}$ in $w(u)$ for every such occurrence. Then

$$|F| = \sum_{u \in V} w(u) = \sum_{u \in V_3} w(u) + \sum_{u \in V_5} w(u),$$

where F is the set of faces in G . Substituting these expressions for $|V|$, $|E|$ and $|F|$ into Euler's Formula $|V| - |E| + |F| = 2$, we get

$$|V_3| + |V_5| - \frac{3}{2}|V_3| - \frac{5}{2}|V_5| + \sum_{u \in V_3} w(u) + \sum_{u \in V_5} w(u) = 2.$$

Collecting terms yields

$$-\frac{1}{2}|V_3| - \frac{3}{2}|V_5| + \sum_{u \in V_3} w(u) + \sum_{u \in V_5} w(u) = 2.$$

¹The face f has a boundary walk and $\text{degree}(f)$ is the number of edges in the boundary walk. In general, this number may be larger than the number of edges incident with f , as an edge can occur twice in a single boundary walk. However, G is embedded in the plane, so such an edge is necessarily a 1-cut. Since G has no 1-cuts, each edge separates two faces and so the number of edges in the boundary walk is equal to the number of edges incident with the face.

Now replacing $|V_3|$ by $\sum_{u \in V_3} 1$ and $|V_5|$ by $\sum_{u \in V_5} 1$, and then collecting terms gives

$$\sum_{u \in V_3} (w(u) - \frac{1}{2}) + \sum_{u \in V_5} (w(u) - \frac{3}{2}) = 2.$$

Since G has no digons and triangles containing a 3-vertex, each face incident with a 3-vertex has degree at least 4. So for each $u \in V_3$, $w(u) \leq 3/4$. Thus

$$\sum_{u \in V_3} (w(u) - \frac{1}{2}) \leq \sum_{u \in V_3} \frac{1}{4}.$$

If G satisfies hypothesis (i), the expression above has maximum value $3/4$ because there are at most three 3-vertices. If G satisfies hypothesis (ii), it has at most one 3-vertex so the maximum value is $1/4$. Thus

$$\sum_{u \in V_5} (w(u) - \frac{3}{2}) \geq \begin{cases} 2 - \frac{3}{4} = \frac{5}{4}, & \text{if } G \text{ satisfies hypothesis (i),} \\ 2 - \frac{1}{4} = \frac{7}{4}, & \text{if } G \text{ satisfies hypothesis (ii).} \end{cases}$$

If at most three 3-faces are incident to a 5-vertex v , then v contributes at most zero to the sum. So there must exist a 5-vertex with at least four incident 3-faces. If G satisfies hypothesis (i), it is enough to establish the existence of one such vertex, so we are done. If G satisfies hypothesis (ii), vertex v must be distinct from the distinguished vertex, its five neighbours and the two neighbours of the distinguished neighbour d^* whose incident faces are not incident with d (the last condition ensures that d^* is neither q nor s). The total number of excluded vertices is eight. The maximum that a 5-vertex v can contribute to the sum (which is at least $\frac{7}{4}$) is when all its incident faces are 3-faces, in which case $w(v) - \frac{3}{2} = \frac{1}{6}$. Thus, there are at least $\frac{7}{4} \div \frac{1}{6} = 10.5$ vertices that have at least four 3-faces incident to it. Since this number is greater than 8, there exists at least one vertex which is sufficiently far from d and d^* . This completes the proof. \square

We now show that the embedded graph G with a Grötzsch Configuration has an $m3$ -orientation, by forming a reduced graph G' and a contracted graph G^* and inductively obtaining their $m3$ -orientations first (Figure 2.10).

Lemma 2.3.8: If G has a Grötzsch Configuration, then G has an $m3$ -orientation.

Proof: Form a reduced graph G' from G as follows. Split vertex q into a 2-vertex $q' = (\beta_1, \beta_2)$ and a 3-vertex q'' . Delete edges (r, s) , (s, v) and (v, t) , that is, the edges of a zigzag. Add two new vertices z' and z'' . Join z' to s with two parallel edges and z' to t with a single

edge. Join z'' to v with two parallel edges and z'' to r with a single edge. The graph G' has three 3-vertices not inherited from G : q' , z' and z'' . Form a contracted graph G^* from G' by contracting $Y = \{p, q'', r, v, z''\}$ and $Z = \{s, t, z'\}$ to single vertices y and z respectively. Observe that nonperipheral 3-cuts in G have been taken care of, so deleting these three edges (and adding new vertices and edges) of G does not create a disconnection. Thus, G' is connected, and hence, so is G^* . We first establish that G^* has an $m3$ -orientation, extend it to one of G' , and then finally to an $m3$ -orientation of G .

The cuts of G^* that are not inherited from G are those that separate the set $\{q', y, z\}$ (that is, each side of the cut contains at least one vertex in this set). The cuts with q' on one side and y, z on the other are left cuts, the ones with z on one side and q', y on the other are right cuts and those with y on one side and q', z on the other are central cuts.

Claim: If G^* has a 1-cut or a nonperipheral 3-cut, then G has an $m3$ -orientation.

Proof: Suppose G^* contains a left 1- or 3-cut c . Replace y and z with $(Y \setminus \{q''\}) \cup \{q\}$ and Z . At most the edges β_1 and β_2 are added to c . Call this cut c' . Observe that not both β_1 and β_2 can be in c , because the side graphs of every cut are connected. This shows that c is nonperipheral in G^* , and hence c' is nonperipheral in G' . If the corresponding cut (of size at most 5) in G is peripheral, then G has a digon containing the vertex q , in which case Lemma 2.3.4 can be employed to obtain an $m3$ -orientation of G . Otherwise, G has an $m3$ -orientation by Lemma 2.3.3. Therefore, we may assume that there are no left 1- or 3-cuts in G^* .

If G^* contains a right 1-cut, it is the image of a nonperipheral 4-cut in G that contains edges (r, s) , (s, v) and (v, t) of the zigzag. If G^* has a right 3-cut, it is the image of a 6-cut in G that contains the edges of the zigzag. In these two cases, G has an $m3$ -orientation by Lemmas 2.3.3 and 2.3.7 respectively. Therefore, we may assume that there are no right 1- or 3-cuts in G^* .

Suppose G^* contains a central 1-cut c . In G' , at most the edges $\{\beta_1, \beta_2, (r, s), (s, v), (v, t)\}$ are added to c . This new cut c' is clearly nonperipheral in G' and has at most 5 more edges than c . Thus, the corresponding cut in G is also nonperipheral and has size at most 6. If its size is at most 5, we invoke Lemma 2.3.3, while if it is 6, we invoke Lemma 2.3.7, to obtain an $m3$ -orientation of G . Therefore, we may assume that there are no central 1-cuts in G^* .

This leaves the possibility of existence of central 3-cuts in G^* . We show below that if a central 3-cut exists in G^* , then the embedding of G is not planar, which is a contradiction.

Let c be a central 3-cut in G^* of the form $c = \delta X$, where $y \in X$. The coboundary $\delta X'$ in G' , where $X' = (X - \{y\}) \cup ((Y \setminus \{q''\}) \cup \{q\})$ (note that $v \in X'$), is two larger than




Figure 2.10: Graph G with a Grötzsch Configuration, and the reduced and contracted graphs G', G^* respectively.

c because it contains edges β_1, β_2 . The corresponding coboundary in G contains the edges of the zigzag as well, so it has size 8.

If K is a connected component of $G - \delta(X')$, then $\delta(K)$ has at least three edges in it. Thus, each of $G[X']$ and $G[V(G) \setminus X']$ has at most two components. If either has precisely two components, then the 8 edges of $\delta(X')$ are divided either 4 and 4 or 3 and 5 between

them. By Lemmas 2.3.3 and 2.3.5, we may assume there are no 4-cuts at all and the 3- and 5-cuts must be peripheral. Thus, a disconnected side consists of a 3-vertex and a 5-vertex. We note that X' has at least the vertices r, v, p , so we conclude $G[X']$ is connected. On the other hand, $V(G) \setminus X'$ contains the adjacent vertices s and t , so $G[V(G) \setminus X']$ is also connected.

Now, $G[X']$ contains a path with ends p and r . The union of this path and the path (p, v, r) is a cycle C in G . Similarly, $G[V(G) \setminus X']$ contains a path A with ends q and s . The vertices q and s are separated in the rotation at v by the vertices p and r . Therefore, q and s are on different sides of C in the plane, yet the path A joins them and is disjoint from K , contradicting the Jordan Curve Theorem. Thus, G^* does not contain any central 3-cuts. \square

Since G^* does not contain any new 3-cuts and, if it exists, the distinguished vertex of G is not disturbed in G^* , G^* satisfies the same hypothesis as G and has an $m3$ -orientation by the inductive hypothesis.

We now extend G^* 's $m3$ -orientation to that of G' . Consider the segment of the 5-cut δZ in G' that corresponds to side Z . The contracted vertex of this segment is incident to edges $\delta_1, \delta_2, \epsilon_1, \epsilon_2$ and ϵ_3 , all of which have already been assigned orientations. So, this segment has the contracted vertex as a distinguished vertex and has one 3-cut $\delta\{z'\}$. Thus, it satisfies hypothesis (ii) and has an $m3$ -orientation by the inductive hypothesis, which agrees with that of G^* on the edges at z .

It remains to orient the edges in Y . Note that the $m3$ -orientation of the edges inside Z obtained above determines a set of directions in G' on the edges at the 3-vertex z' . Transfer these directions, either all in or all out, to the edges at z'' . Assume, for the sake of definiteness and without loss of generality, that the three edges at z'' are directed away from z'' . Consider the segment of G' by $\delta Y = \{\gamma_1, \gamma_2, \alpha_1, \alpha_2, \alpha_3\}$ which contains vertex v . The contracted vertex is a distinguished vertex, and $\delta\{z''\}, \delta\{q''\}$ are the two 3-cuts. Clearly this segment does not satisfy hypothesis (i) or (ii), but by exhaustive searching, it can be seen that it has an $m3$ -orientation that extends any directions on δY and agrees with the directions at z'' . These orientations are shown in Figure 2.11 (there are three cases, depending on the direction of the minority edge e of vertex y in G^* . Figure 2.11(a) depicts the situation when e lies at r in G' ; the three oriented edges at p are all pointing in the same direction, and it does not matter which. Figures 2.11(b) and 2.11(c) depict the situation when e lies at p in G' , with one and two incoming edges at p respectively).

This completes an $m3$ -orientation of G' , which extends to that of G in a straightforward way by restoring vertex q and the edges of the zigzag: the edge (v, t) gets the direction of the edge (z', t) , and the edge (r, s) gets the directions of the edge (r, z'') . This determines




Figure 2.11: m_3 -orientations of the segment of G' by δY that contains vertex v , for all possible sets of directions on δY . The directions inherited from G^* and z' are shown with solid arrows.

a unique direction for (v, s) . □

This completes the inductive step and the proof of Theorem 2.1.1. □

Chapter 3

A new proof of Grötzsch's Theorem

In this chapter, we give a proof of Grötzsch's Theorem that is inspired by Thomassen's [19], but is due to the unpublished work of C. Nunes da Silva, R.B. Richter and D. Younger. We state the theorem again for the reader's ease.

Grötzsch's Theorem: Every loop-free and triangle-free planar graph has a 3-colouring.

Before we begin, it is worth reiterating that a triangle-free planar graph G must be loop-free for it to have a 3-colouring, as described in Section 1.1. Moreover, the existence of multiple edges in G does not change the way its vertices can be coloured; as long as there is one edge between two vertices u and v , they receive different colours and having more edges between u and v does not require a change in colours. Therefore, it is acceptable to assume that G is free of loops and multiple edges, and we concern ourselves only with planar graphs having girth at least 4.

Section 3.1 presents some preliminaries. Section 3.2 gives a brief sketch of the proof, followed by the complete proof.

3.1 Preliminaries

Let M be a planar map of a 2-connected planar graph $G = (V, E)$, and let \mathcal{C} be the cycle in G that bounds the infinite face in M . We present some terminology that will be used throughout the chapter.

Let $X, Y \subseteq V(G)$. If $v \in X$, $w \in Y$ and $(v, w) \in E(G)$, then v is called an X -vertex and w is called a Y -neighbour of v . An edge between two X -vertices is called an

X -adjacency, and an edge between an X -vertex and a Y -vertex is called an X - Y adjacency. A uX -adjacency refers to an edge between a vertex $u \in V(G)$ and an X -vertex. We use $N(v)$ to denote the set of neighbours of a vertex v in G , and $N(X)$ denotes the set of vertices adjacent to at least one X -vertex.

Paths play an important role in the arguments we present in the next section. In particular, we use the notion of the length of a path to construct inductive arguments at various stages of our proof. A path that has an end-vertex in $X \subseteq V(G)$ is said to have an X -end. For $X, Y \subseteq V(G)$ and any positive integer k , an X - Y path is a path with one end-vertex in X and the other in Y . An x - y path refers to a path between $x, y \in V(G)$. A k -path is a path of length k (that is, it contains k edges) with its end-vertices in the outer cycle \mathcal{C} of M and the interior vertices not in \mathcal{C} . For simplicity, we refer to a 1-path as a chord. A k -path P separates $X \subseteq V(G)$ from $Y \subseteq V(G)$ if $G - P$ contains no path from an X -vertex to a Y -vertex.

In 2-connected maps, cycles and paths are fundamentally related. A k -cycle refers to a cycle of length k . It is easy to see that any k -path P in M induces two submaps M_1 and M_2 such that $M_1 \cap M_2 = P$. These submaps are called the two sides of P . Observe that the infinite faces of M_1 and M_2 are bounded by the cycles C_1 and C_2 respectively, and these are the two cycles in $\mathcal{C} \cup P$ that contain P . We call these cycles the inner i - and j -cycles, where $|E(C_1)| = i$ and $|E(C_2)| = j$.

Cycles can be divided into two broad categories. A cycle C in M is a facial cycle if it bounds a finite face (note that \mathcal{C} is facial if and only if $M = \mathcal{C}$). Otherwise, it is called a separating cycle. We denote by $Int(C)$ the submap induced by the vertices lying on the bounded side of C in the map and not in C ; $Ext(C)$ denotes the submap induced by $V(M - (Int(C) \cup C))$.

For a graph that is not 2-connected, we use the notion of blocks. Let G' be a connected subgraph of G . A block in G' is a maximally connected subgraph with no cut-vertex (that is, a maximal 2-connected subgraph). Any two blocks of G' intersect in at most one cut-vertex, and every cut-vertex is in at least two blocks of G' . It is not hard to see that two blocks are adjacent if they share a cut-vertex of G' , and that each edge of G' is in exactly one block. The block-cutpoint tree of G' is the bipartite tree whose vertices are the blocks and cut-vertices of G' , and a cut-vertex v is adjacent to a block B if $v \in V(B)$.

Lastly, for $v \in V(G)$, $L(v)$ denotes the list of colours available to v in a list assignment L .

3.2 Proof of Grötzsch's Theorem

Grötzsch's Theorem is about a planar graph G ; thus, we may assume G is the graph of a planar map M . It is not hard to see that it suffices to assume G is 2-connected, so that every face of M is bounded by a cycle; we will typically let \mathcal{C} be the cycle bounding the infinite face of M . The inductive proof will reduce 4-cycles, leaving us to prove the result for 2-connected planar maps of girth 5. Thomassen proved a variant of Theorem 3.2.1 below; its proof is the main point of this chapter.

Proof Sketch: We begin by proving Theorem 3.2.1, which asserts 3-list-colourability of a special class of maps: 2-connected with girth ≥ 5 , having some specific constraints on the colours of a few vertices on the cycle bounding the infinite face and the sizes of colour-lists of vertices. We then prove Grötzsch's Theorem for a graph G by embedding G in the plane to obtain a map M and reducing all 4-cycles in M , except perhaps the one bounding the infinite face. Then, M is transformed to satisfy the conditions of Theorem 3.2.1. The theorem asserts that M is 3-list-colourable. Thus, we can obtain a 3-colouring of M and, consequently, a 3-colouring of G .

Theorem 3.2.1: Let M be a 2-connected, planar map with the cycle \mathcal{C} bounding the infinite face. Let \mathcal{C} have length at least 4, while all other cycles of M have length at least 5. Let S and T be disjoint subsets of $V(\mathcal{C})$ and let L be a list assignment to M so that:

- (a) S consists of a set of consecutive vertices on \mathcal{C} such that, for each $s \in S$, $|L(s)| = 1$;
- (b) if $t \in T$, then $|L(t)| = 2$;
- (c) there is no S - T adjacency; and
- (d) for every vertex v of M not in $S \cup T$, $|L(v)| = 3$.

Suppose M satisfies one of the following hypotheses:

- i. $4 \leq |V(\mathcal{C})| \leq 7$ and $S = V(\mathcal{C})$;
- ii. $|S| \leq 5$ and no two T -vertices are adjacent (call this the $|S|$ -0 hypothesis for $|S| \leq 5$);
or
- iii. $|S| \leq 3$ and exactly one pair of T -vertices is adjacent (call this the $|S|$ -1 hypothesis for $|S| \leq 3$).

For $s \in S$, let $c(s)$ be the colour in $L(s)$. If $c(S)$ is an L -colouring of $M[S]$ (*this is the submap of M induced by S*), then c extends to an L -colouring of M .

Proof: We proceed by induction on the number of vertices of M .

Base Case: The smallest 2-connected, planar map M satisfying the conditions given above is a facial 4-cycle \mathcal{C} . If $S = V(\mathcal{C})$, we are done. Otherwise, $S \neq V(\mathcal{C})$, and there is a vertex v of \mathcal{C} not in $S \cup T$ so that $|L(v)| = 3$. Working outwards from S , we can L -colour

the path $\mathcal{C} - v$, because any vertex that is not yet coloured has at least two available colours. Once $\mathcal{C} - v$ is coloured, we can colour v , since it has only two neighbours and three available colours. This completes an L -colouring of M .

Inductive Step: First of all, if M contains a T -adjacency as a chord (x, y) , then an L -colouring of M can be obtained as follows. Let M_1 be the side of the chord that contains all the S -vertices, if there are any. M_1 satisfies the same hypothesis as M and has an L -colouring by the inductive hypothesis. This colours x, y on the other side, M_2 . Since x, y have no $(T \cap M_2)$ -neighbours, x and y are the only coloured vertices in M_2 . Therefore M_2 satisfies the 2-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M . Thus, for the remainder of the proof, we may assume that the T -adjacency in M is contained in \mathcal{C} .

As part of the induction, we will often be deleting up to three vertices that are consecutive on \mathcal{C} to produce a submap M' , which may or may not be connected. They will have been properly coloured before deletion (either because they were in S or for some other reason). The colour lists of their neighbours will be changed by removing the colour of the deleted vertex (by girth considerations, no vertex has two deleted vertices as neighbours). Thus, we may be reducing the size of S and we will almost certainly be changing T . Note that M' may not be connected or 2-connected, so it may satisfy one of the three hypotheses without the assumption of 2-connectedness. The following lemma takes care of this scenario. Let S' be the vertices of M' from S and let T' be the vertices of M' that now have lists of size 2. Let L' be the list colouring on M' inherited in this way.

Lemma 3.2.2: If either no two vertices of T' are adjacent or there is only one T' -adjacency and $|S'| \leq 3$, then M has an L -colouring.

Proof: It suffices to show how to use the induction to get an L' -colouring of M' .

Because the deleted vertices are consecutive on \mathcal{C} , $\mathcal{C} \cap M'$ is a path and, therefore, is contained in a component of M' . Thus, all the vertices in S' are contained in the same component K_1 of M' .

Let B_1, B_2, \dots, B_k be the blocks of M' , labelled so that $|S' \cap V(B_1)|$ is at least as large as any other $|S' \cap V(B_i)|$. Furthermore, if any of the blocks with maximum $|S' \cap V(B_i)|$ also contains the T -adjacency, then this block must be B_1 .

We begin our L' -colouring of M' by using the induction to colour B_1 . In any component of M' other than the one containing B_1 , we begin by colouring inductively any of its blocks.

Claim: Any uncoloured block adjacent to an L' -coloured block has an L' -colouring.

Proof: Let D be an uncoloured block adjacent to a coloured block B in M' and let $u \in$

$V(D) \cap V(B)$. The block D is not adjacent to any other coloured block of M' , otherwise the block-cutpoint tree of M' contains a cycle, a contradiction.

If u was originally coloured in M' , then D may or may not contain other coloured vertices. In this case, D satisfies the $|S' \cap V(D)|=0$ or $|S' \cap V(D)|=1$ hypothesis, where $|S' \cap V(D)|$ is at most the number of originally coloured vertices in M' , and has an L' -colouring by the inductive hypothesis. On the other hand, if u was originally uncoloured in M' , then u is the only coloured vertex in D . If there is no uT' -adjacency, then D satisfies the 1-0 or 1-1 hypothesis and has an L' -colouring by the inductive hypothesis. So suppose there exists at least one uT' -adjacency in D .

The uT' -adjacencies in D partition D into minimal uT' -chord-free 2-connected submaps. In each of these submaps, at most u and any u -neighbour that is in the T' -adjacency can have T' -neighbours and these are in the boundary. Assign a colour, different from $c(u)$, to each T' -neighbour of u . If this colours one vertex of the T' -adjacency, colour the other vertex of the T' -adjacency as well. So each submap satisfies either the $k=0$ hypothesis for $k \leq 4$, or the $k=1$ hypothesis for $k \leq 3$. Obtain an L' -colouring of each submap by the inductive hypothesis. The union of these L' -colourings gives an L' -colouring of D . \square

A trivial induction completes the proof. \square

Now we proceed to treat the case where M is 2-connected and satisfies Hypothesis i, ii or iii. We treat each of the three hypotheses separately via three lemmas, and prove that M has an L -colouring in each case. If $S \neq V(\mathcal{C})$, we denote by \underline{S} -path the path in \mathcal{C} consisting of all the vertices in S . For any connected submap H of M with list assignment L' , S_H is the set of all $s \in V(H)$ such that $|L'(s)| = 1$ and T_H is the set of all $t \in V(H)$ such that $|L'(t)| = 2$. If H is 2-connected, $\underline{\mathcal{C}}_H$ denotes the cycle bounding the infinite face in H .

Lemma 3.2.3: If $4 \leq |V(\mathcal{C})| \leq 7$ and $S = V(\mathcal{C})$ in M , then M has an L -colouring.

Proof: If \mathcal{C} is a facial 4-cycle, then $M = \mathcal{C}$ is already coloured. So we may assume that \mathcal{C} is not a facial 4-cycle.

Claim 3.2.4: There are no chords in M . If M has a 2-path, then M is L -colourable.

Proof: Let P be a chord or 2-path in M . Then P induces two inner cycles C_1 and C_2 in M , such that

$$\text{length}(C_1) + \text{length}(C_2) = \text{length}(\mathcal{C}) + 2 \times \text{length}(P). \quad (3.1)$$

Case 1: P is a chord. In this case,

$$6 = 4 + 2(1) \leq \text{length}(C_1) + \text{length}(C_2) \leq 7 + 2(1) = 9$$

by (3.1). The inequality implies that one of C_1 and C_2 has length at most 4, a contradiction. So M has no chords.

Case 2: P is a 2-path. In this case, let $P = (x, y, z)$ where $x, z \in V(\mathcal{C})$. By (3.1),

$$8 = 4 + 2(2) \leq \text{length}(C_1) + \text{length}(C_2) \leq 7 + 2(2) = 11.$$

By assumption, both C_1 and C_2 have length at least 5; we conclude from the inequality that one has length 5 and the other has length 5 or 6. Since $y \in \text{Int}(\mathcal{C})$, $|L(y)| = 3$. Assign a colour to y different from $c(x)$ and $c(z)$ to make C_1 and C_2 fully coloured. Now each of $\text{Int}(C_1) \cup C_1$ and $\text{Int}(C_2) \cup C_2$ has an L -colouring by the inductive hypothesis. These L -colourings agree on their colours on $V(P)$ and their union gives an L -colouring of M . \square

Assume now that M does not contain chords and 2-paths. Obtain a smaller map M' by deleting at most two adjacent vertices of \mathcal{C} and removing their colours from the colour-lists of their neighbours, so that $2 \leq |S_{M'}| \leq 5$. Observe that these neighbours, whose colour-lists decrease by 1, are now in $T_{M'}$. No two $T_{M'}$ -vertices are adjacent, otherwise M contains a triangle or an inner 4-cycle, a contradiction. Also, there is no $S_{M'}-T_{M'}$ adjacency, otherwise M contains a 2-path, a contradiction. Therefore M' is a submap of M that satisfies the $|S_{M'}|-0$ hypothesis for $2 \leq |S_{M'}| \leq 5$. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to an L -colouring of M by restoring the deleted vertices. \square

We now move on to Hypothesis ii.

Lemma 3.2.5: If M satisfies the $|S|-0$ hypothesis for $|S| \leq 5$, then M has an L -colouring.

Proof: If M satisfies the 0-0 hypothesis, we transform M to satisfy the 1-0 hypothesis as follows. If possible, choose $v \in T$; otherwise, choose $v \in V(\mathcal{C})$. Assign to v a colour from $L(v)$. Since v has no T -neighbours, this shifts the hypothesis to 1-0.

Thus, we may assume $S \neq \emptyset$. We next show how to reduce $(S \cup T)-N(S)$ chords and $S-(S \cup T)$ 2-paths in M , and then prove the lemma for the reduced map.

Claim 3.2.6: If M contains a $(S \cup T)-N(S)$ chord, then M has an L -colouring.

Proof: For any $T-N(S)$ chord, choose the labelling of the two sides M_1 and M_2 so that S is contained in M_1 . We pick a chord (x, y) , where $x \in T$ and $y \in N(S)$, so that M_2 is minimal. Observe that M_1 does not contain a T -adjacency because M does not. Therefore, it satisfies the same hypothesis as M and has an L -colouring by the inductive hypothesis. This colours x and y in M_2 . Observe that x cannot have a T -neighbour because $x \in T$. If y has a T -neighbour in M_2 , it must be via an edge in \mathcal{C}_{M_2} by the minimality of M_2 ; assign to it a colour different from $c(y)$. Therefore, M_2 satisfies the 2-0 or 3-0 hypothesis and has

an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Now assume M does not contain T - $N(S)$ chords. Let (x, y) be a chord in M where $x \in S$ and $y \in N(S)$. Because of the girth assumption on M , there is a unique x - y subpath P of \mathcal{C} containing 4 or 5 S -vertices. Let M_1 be the submap of M bounded by $P + (x, y)$. Since $|L(y)| = 3$, assign to y a colour different from $c(x)$ and its other S -neighbour in M_1 . Now M_1 has a fully-coloured outer boundary \mathcal{C}_{M_1} of length 5 or 6, so it has an L -colouring by the inductive hypothesis. This colours x and y in the other side of P , M_2 . If y has a T -neighbour in M_2 , it must be via an edge in \mathcal{C}_{M_2} ; assign to it a colour different from $c(y)$. Now S_{M_2} consists of x, y , possibly a T -neighbour of y and at most one other S -vertex. Therefore M_2 satisfies the 2-0, 3-0 or 4-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M . \square

Claim 3.2.7: If M contains an S - $(S \cup T)$ 2-path, then M has an L -colouring.

Proof: For any S - T 2-path in M , choose a labelling of its two sides M_1 and M_2 so that $|V(M_2) \cap S| \leq |V(M_1) \cap S|$. Let (x, y, z) be an S - T 2-path in M , where $x \in T$ and $z \in S$, such that M_2 is minimal. M_1 contains at most as many coloured vertices as M . So, M_1 satisfies the $|S_{M_1}|$ -0 hypothesis for $|S_{M_1}| \leq |S|$ and hence, it has an L -colouring by the inductive hypothesis. This colours x and y in M_2 . Since $x \in T$, x does not have any T -neighbours. Moreover, y does not have any T -neighbours in M_2 by the minimality of M_2 . So S_{M_2} consists of x, y, z and at most two S -vertices. Thus, M_2 satisfies the 3-0, 4-0 or 5-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Now assume M does not contain S - T 2-paths. Let (x, y, z) be an S - S 2-path in M . Since there are no 4-cycles in $\text{Int}(\mathcal{C})$, there are either two or three S -vertices between x and z on the S -path; this gives an inner 5- or 6-cycle C in M containing x, y and z . Observe that C is fully coloured, except for y . Since $y \in \text{Int}(\mathcal{C})$, $L(y) = 3$. Assign a colour to y different from $c(x)$ and $c(z)$. Now C is a fully-coloured, inner 5- or 6-cycle, so $C \cup \text{Int}(\mathcal{C})$ has an L -colouring by the inductive hypothesis. The submap $M' = \text{Ext}(C) \cup \{x, y, z\}$ is planar and 2-connected. Moreover, y has no T -neighbours in M' because M has no S - T 2-paths. Now $S_{M'}$ consists of x, y, z and possibly another S -vertex. Thus, M' satisfies the 3-0 or 4-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M . \square

We may now assume that M has no $(S \cup T)$ - $N(S)$ chords, and no S - $(S \cup T)$ 2-paths. If M satisfies the 1-0 hypothesis, it has exactly one S -vertex v ; colour one of its neighbours $u \in \mathcal{C}$ with a colour different from $c(v)$. If u has a T -neighbour w , it must be via an edge in \mathcal{C} because M has no T - $N(S)$ chords; assign to w a colour different from $c(u)$. This shifts

the hypothesis to 2-0 or 3-0.

Thus, we may assume that $|S| \geq 2$. Obtain a reduced map M' from M as follows (Figure 3.1). If M satisfies the 2-0 hypothesis, then delete one vertex at the end of the S -path; otherwise, delete two vertices at one end of the S -path. Remove the colours of deleted vertices from the colour-lists of their neighbours, adding these neighbours to T , creating the new set $T_{M'}$. Observe that $1 \leq |S_{M'}| \leq 3$. There is no $S_{M'}-T_{M'}$ adjacency; otherwise, M contains an $S-N(S)$ chord or an $S-S$ 2-path, a contradiction.

We now consider possible $T_{M'}$ -adjacencies. If two vertices of $T_{M'} \setminus T$ are adjacent, then M has a 3- or 4-cycle, which is impossible. If a vertex of $T_{M'} \setminus T$ in $Int(\mathcal{C})$ is adjacent to a T -vertex, then M has an $S-T$ 2-path, another contradiction.

In the remaining possibility, a vertex v of $T_{M'} \setminus T$ in \mathcal{C} is adjacent to a T -vertex w . If $(v, w) \notin E(\mathcal{C})$, then (v, w) is a $T-N(S)$ chord in M . Since no such chords exist, $(v, w) \in E(\mathcal{C})$. Since we retain at least one vertex of S in M' , this can happen for v , the \mathcal{C} -neighbour of the deleted end of the S -path. Therefore, there is at most one $T_{M'}$ -adjacency in M' . Thus, M' satisfies the $|S_{M'}|-0$ or $|S_{M'}|-1$ hypothesis for $1 \leq |S_{M'}| \leq 3$. Lemma 3.2.2 implies M' has an L -colouring; restoring the deleted vertices provides an L -colouring of M . \square

Finally, we consider Hypothesis iii.

Lemma 3.2.8: If M satisfies the $|S|-1$ hypothesis for $|S| \leq 3$, then M has an L -colouring.


Proof: Let τ denote the set of vertices in the T -adjacency contained in M . The set $N_{\mathcal{C}}(\tau)$ denotes the \mathcal{C} -neighbours of the τ -vertices via edges in \mathcal{C} .

If M satisfies the 0-1 hypothesis, we colour M as follows. Pick a vertex $t_1 \in \tau$ and assign to it a colour from $L(t_1)$. Assign to the other τ -vertex a colour different from $c(t_1)$. This shifts the hypothesis to 2-0, in which case, Lemma 3.2.5 implies M has an L -colouring. Thus, we may assume that M satisfies the $|S|-1$ hypothesis for $1 \leq |S| \leq 3$.


Now we reduce some chords and 2-paths in M and then prove the lemma for the reduced map.

Claim 3.2.9: If M contains a chord with one end in $S \cup T$, then M has an L -colouring.


Proof: Suppose first that both ends a and b of the chord are in $S \cup T$. Since there are no $S-T$ adjacencies, both a and b are in the same one of S and T . The girth constraint implies they are not both in S , so they are both in T . The side that contains all the S -vertices satisfies the $|S|-0$ or $|S|-1$ hypothesis and has an L -colouring by the inductive hypothesis. This colours a and b in the other side. In this side, a and b do not have any T -neighbours, except maybe one τ -vertex. Thus, this side satisfies the 2-0 or 3-0 hypothesis and has




(a) M satisfies the 5-0 hypothesis and $a, b, c, d, e \in S$, $g \in T$.



(b) Obtain M' by deleting d, e . This puts f in $T_{M'}$ (in this example). Also, $h, i, j \in N(\{d, e\}) \cap T_{M'}$.



(c) The $S_{M'}-T_{M'}$ adjacencies (b, h) and (a, f) (depicted by dashed edges) imply an $S-S$ 2-path (b, h, d) and an $S-N(S)$ chord (a, f) in M respectively.



(d) The $T_{M'}$ -adjacencies (h, j) , (j, k) and (f, k) imply a 4-cycle (d, h, j, e, d) , an $S-T$ 2-path (e, j, k) and a $T-N(S)$ chord (f, k) in M respectively.

Figure 3.1: Example depicting the reduction of a map M that satisfies the 5-0 hypothesis and does not contain any $(S \cup T)-N(S)$ chords or $S-(S \cup T)$ 2-paths.

an L -colouring by the inductive hypothesis. Thus, we may assume each such chord has precisely one end in $S \cup T$.

For each such chord in M , choose the labelling of its two sides M_1 and M_2 so that $|V(M_2) \cap S| \leq |V(M_1) \cap S|$. In case of equality, choose the labelling so that $\tau \subseteq V(M_1)$. Pick a chord (x, y) , where $x \in S \cup T$ and $y \notin S \cup T$, so that M_2 is minimal (Figure 3.2).




Figure 3.2: Reducing chords with one end in $S \cup T$ in a map M that satisfies the 3-1 hypothesis. Here, $a, b, c \in S$, (d, e) is the T -adjacency and $d, e \in \tau$. Each dashed line is a chord with one end in $S \cup T$ and the other end $y \notin S \cup T$. Chords (y, c) and (y, f) separate S and τ ; then M_2 is the side containing τ . Chords (y, g) , (y, a) , (y, b) do not separate S and τ ; then M_2 is the side that does not contain τ . Note that if (y, g) is present, then neither (y, a) nor (y, b) can have minimal M_2 . On the other hand, if (y, f) is present, then (y, c) does not have minimal M_2 .

Case 1: $x \in T$. In this case, which side of the chord is M_1 is determined by the fact that it contains all the S -vertices. Then M_1 satisfies the $|S|-j$ hypothesis for $j \in \{0, 1\}$ and has an L -colouring by the inductive hypothesis. This colours x and y in M_2 . Observe that $j = 0$ if (x, y) separates S from τ , and $j = 1$ otherwise. If y has a T -neighbour t_1 in M_2 , then minimality of M_2 implies that $(y, t_1) \in E(\mathcal{C}_{M_2})$; assign to t_1 a colour different from $c(y)$, and if $t_1 \in \tau$, assign to the second τ -vertex t_2 a colour different from $c(t_1)$. On the other hand, if $x \in \tau$, assign to the other τ -vertex a colour different from $c(x)$. Now S_{M_2} consists of x, y , and at most two other vertices. If $j = 0$, only y may have a T -neighbour in M_2 , so $|S_{M_2}| \leq 3$. If $j = 1$, $|S_{M_2}| \leq 4$. Thus, M_2 satisfies the $|S_{M_2}|-(1-j)$ hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Case 2: $x \in S$. In this case, we distinguish two subcases.

Subcase 1: (x, y) separates $S \setminus \{x\}$ from τ . In particular, S has either 2 or 3 vertices. In this case, $\tau \subset V(M_2)$ and $S \subset V(M_1)$. Every S -vertex in M is contained in M_1 , so M_1 satisfies the $|S|-0$ hypothesis and has an L -colouring by the inductive hypothesis. This

colours x and y in M_2 . If y is adjacent to a T -vertex $t_1 \in V(M_2)$, then $(y, t_1) \in E(\mathcal{C}_{M_2})$, otherwise we are back in Case 1. Assign to t_1 a colour different from $c(y)$, and if $t_1 \in \tau$, assign to the second τ -vertex t_2 a colour different from $c(t_1)$. Now S_{M_2} consists of x, y , possibly a T -neighbour t_1 of y , and possibly the τ -neighbour t_2 of t_1 . Thus, $|S_{M_2}| \leq 4$ and, if $|S_{M_2}| = 4$, then the vertices of τ are contained in M_2 . It follows that M_2 satisfies either the $|S_{M_2}|-1$ hypothesis for $|S_{M_2}| \leq 3$ or the 4-0 hypothesis. So M_2 has an L -colouring by the inductive hypothesis. The union of the L -colourings of M_1 and M_2 gives an L -colouring of M .

Subcase 2: (x, y) does not separate $S \setminus \{x\}$ from τ . In this case, M_1 contains τ and most of the S -vertices. Hence it satisfies the $|S_{M_1}|-1$ hypothesis for $|S_{M_1}| \leq |S|$ and has an L -colouring by the inductive hypothesis. This colours x and y in M_2 . If y has a T -neighbour in M_2 , it must be via an edge in \mathcal{C}_{M_2} ; assign to it a colour different from $c(y)$. Now S_{M_2} consists of x, y , possibly a neighbour of y and at most one other S -vertex. Therefore M_2 satisfies the $|S_{M_2}|-0$ hypothesis for $2 \leq |S_{M_2}| \leq 4$ and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M . \square

Now assume that M has no chords with one end in $S \cup T$.

Claim 3.2.10: If M contains a $(T \cup S)$ - $(T \cup N_{\mathcal{C}}(\tau))$ 2-path, then M has an L -colouring.

Proof: First, consider T - $(T \cup N_{\mathcal{C}}(\tau))$ 2-paths in M .

Let (x, y, z) be a 2-path in M , where $x \in T$ and $z \in T \cup N_{\mathcal{C}}(\tau)$, such that the side containing no S -vertices is minimal. Call this side M_2 and the other side M_1 . We distinguish two cases.

Case 1: $z \in (T - \tau)$. In this case, M_1 contains all the S -vertices and possibly the T -adjacency as well. Thus, it satisfies the $|S|-j$ hypothesis for $j \in \{0, 1\}$ and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . Observe that y has no T -neighbours in M_2 by the minimality of M_2 . Neither does z , by assumption. However, x may have a T -neighbour t in M_2 if $x \in \tau$. If this is the case, assign a colour to t different from $c(x)$ so that M_2 satisfies the 4-0 hypothesis, and obtain M_2 's L -colouring by the inductive hypothesis. Otherwise, $S_{M_2} = \{x, y, z\}$. Then M_2 satisfies the $3-(1-j)$ hypothesis and has an L -colouring by the inductive hypothesis. The union of the L -colourings of M_1 and M_2 gives an L -colouring of M .

Case 2: $z \in \tau \cup N_{\mathcal{C}}(\tau)$. We further distinguish two subcases.

Subcase 1: (x, y, z) separates S from τ . In this case, $\tau \subset V(M_2)$ and $S \subset V(M_1)$. So M_1 satisfies the $|S|-0$ hypothesis and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . By the minimality of M_2 , y has no T -neighbours in M_2 . Moreover, z has no T -neighbour in M_2 except via an edge in \mathcal{C}_{M_2} . Since one τ -vertex t_1 is either equal

or adjacent to z in M_2 , assign to t_1 a colour different from $c(z)$ if it is uncoloured and assign to the second τ -vertex t_2 a colour different from $c(t_1)$. Now $S_{M_2} = \{x, y, z, t_1, t_2\}$, with the possibility that $z = t_1$. So M_2 satisfies the 4-0 or 5-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of the L -colourings of M_1 and M_2 gives an L -colouring of M .

Subcase 2: (x, y, z) does not separate S from τ . In this case, $S, \tau \subset V(M_1)$, therefore M_1 satisfies the same hypothesis as M and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . If $z \in \tau$, neither x nor z has a T -neighbour in M_2 ; then M_2 satisfies the 3-0 hypothesis. If $z \in N_{\mathcal{C}}(\tau)$ and z has a T -neighbour $t \in V(M_2)$, then $(z, t) \in E(\mathcal{C}_{M_2})$. Assign to t a colour different from $c(z)$. Now, S_{M_2} consists of $\{x, y, z\}$ and possibly t . Thus, M_2 satisfies the 3-0 or 4-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Now consider S - $(T \cup N_{\mathcal{C}}(\tau))$ 2-paths in M . For each S - $(T \cup N_{\mathcal{C}}(\tau))$ 2-path in M , choose the labelling of its sides M_1 and M_2 so that $|V(M_2) \cap S| \leq |V(M_1) \cap S|$. In case of equality, choose the labelling so that $\tau \subset V(M_1)$. Pick a 2-path (x, y, z) , where $x \in S$ and $z \in T \cup N_{\mathcal{C}}(\tau)$, so that M_2 is minimal. We distinguish two cases.

Case 1: $z \in (T - \tau)$. We further distinguish two subcases.

Subcase 1: (x, y, z) separates $S \setminus \{x\}$ from τ . In this case, $\tau \subset V(M_2)$ and $S \subset V(M_1)$. So M_1 satisfies the $|S|-0$ hypothesis and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . Both x and z have no T -neighbours. Also, y has no T -neighbours in M_2 by the minimality of M_2 , so $S_{M_2} = \{x, y, z\}$. Thus, M_2 satisfies the 3-1 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Subcase 2: (x, y, z) does not separate $S \setminus \{x\}$ from τ . In this case, M_1 contains τ and at most as many coloured vertices as M . So M_1 satisfies the $|S_{M_1}|-1$ hypothesis for $|S_{M_1}| \leq |S|$ and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . Both x and z have no T -neighbours. Also, y has no T -neighbours in M_2 by the minimality of M_2 , so S_{M_2} consists of x, y, z and at most one other S -vertex. Thus, M_2 satisfies the 3-0 or 4-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Case 2: $z \in \tau \cup N_{\mathcal{C}}(\tau)$. We further distinguish two subcases.

Subcase 1: (x, y, z) separates $S \setminus \{x\}$ from τ . In this case, $\tau \subset V(M_2)$ and $S \subset V(M_1)$. So M_1 satisfies the $|S|-0$ hypothesis and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . Also, z has no T -neighbours in M_2 , except via an edge in \mathcal{C} . Since one τ -vertex t_1 is either equal or adjacent to z in M_2 , assign to it a colour different

from $c(z)$ if it is uncoloured and assign to the second τ -vertex t_2 a colour different from $c(t_1)$. Now $S_{M_2} = \{x, y, z, t_1, t_2\}$, with the possibility that $z = t_1$. Thus, M_2 satisfies the 4-0 or 5-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M .

Subcase 2: (x, y, z) does not separate $S \setminus \{x\}$ from τ . In this case, M_1 contains τ and at most as many coloured vertices as M . So M_1 satisfies the $|S_{M_1}|-1$ hypothesis for $|S_{M_1}| \leq |S|$ and has an L -colouring by the inductive hypothesis. This colours x, y and z in M_2 . If $z \in N_{\mathcal{C}}(\tau)$ and z has a T -neighbour $t \in V(M_2)$, assign a colour to t different from $c(z)$. Now S_{M_2} consists of x, y, z , possibly t and at most one other S -vertex. So, M_2 satisfies the 3-0, 4-0 or 5-0 hypothesis and has an L -colouring by the inductive hypothesis. The union of these two L -colourings gives an L -colouring of M . \square


Assume now that M contains no chords with one end in $S \cup T$, and no $(S \cup T)$ - $(T \cup N(\tau))$ 2-paths. If M satisfies the 1-1 hypothesis, it has exactly one S -vertex v ; colour one of its neighbours $u \in \mathcal{C}$ with a colour different from $c(v)$. If u has a T -neighbour t_1 , then $(u, t_1) \in E(\mathcal{C})$ because M does not contain chords with one end in T . Assign to t_1 a colour different from $c(u)$. Moreover, if t_1 is a τ -vertex, assign a colour to the second τ -vertex different from $c(t_1)$. This shifts the hypothesis to 4-0 (in which case M has an L -colouring by Lemma 3.2.5), 2-1 or 3-1. We now describe how to obtain an L -colouring of M .

Let u_1, u_2, u_3, u_4, u_5 and u_6 be 6 consecutive vertices on \mathcal{C} , where $u_3, u_4 \in \tau$. We distinguish two cases.

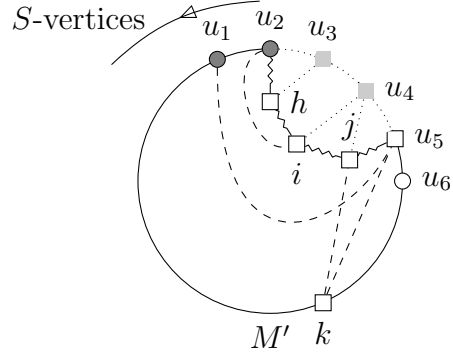
Case 1: $u_6 \notin S \cup T$ (Figure 3.3). We distinguish two subcases.

Subcase 1: $u_1 \in S$. In this subcase, assign a colour to u_2 different from $c(u_1)$, a colour to u_3 different from $c(u_2)$, and a colour to u_4 different from $c(u_3)$. Obtain a reduced map M' from M by deleting u_3 and u_4 and removing their colours from the colour-lists of their neighbours. Observe that $S_{M'} = S \cup \{u_2\}$. Since M does not contain S - $N_{\mathcal{C}}(\tau)$ chords or S - τ 2-paths, no $(T_{M'} \setminus T)$ -vertex is adjacent to an $S_{M'}$ -vertex. Also, no two $T_{M'}$ -vertices are adjacent: an edge between a $(T_{M'} \setminus T)$ -vertex and a T -vertex implies a T - $N_{\mathcal{C}}(\tau)$ chord or a T - τ 2-path in M , and an edge between two $(T_{M'} \setminus T)$ -vertices implies a triangle or an inner 4-cycle in M , a contradiction. Hence M' is a connected, but possibly not 2-connected, submap of M that satisfies the $|S_{M'}|-0$ hypothesis for $|S_{M'}| = |S| + 1$. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to that of M by restoring the deleted vertices.


Subcase 2: $u_1 \notin S$. In this subcase, assign any colour to u_3 from $L(u_3)$ and a colour to u_4 different from $c(u_3)$. Obtain a reduced map M' from M by deleting u_3, u_4 and removing their colours from the colour-lists of their neighbours. Since M does not contain S - $N_{\mathcal{C}}(\tau)$




(a) Subcase 1: $u_1 \in S$ in M .



(b) Subcase 1: Colour u_2, u_3, u_4 . Delete u_3, u_4 to obtain M' . This puts u_5 in $T_{M'}$ (in this example); $h, i, j \in N(\{u_3, u_4\}) \cap T_{M'}$. The dashed edges (u_1, u_5) , (k, u_5) , (u_2, i) , (j, k) correspond to $S-N_C(\tau)$ chord, $T-N_C(\tau)$ chord, 4-cycle & $T-\tau$ 2-path in M respectively.



(c) Subcase 2: $u_1 \notin S$ in M . In this example, $u_1 \in T$. $u, v, w \in S$.



(d) Subcase 2: Colour u_3, u_4 . Delete them to obtain M' . This puts u_2, u_5 in $T_{M'}$ (in this example); $h, i, j \in N(\{u_3, u_4\}) \cap T_{M'}$. The dashed edges (w, u_5) , (k, u_5) , (u_2, i) , (j, k) , (v, h) correspond to $S-N_C(\tau)$ chord, $T-N_C(\tau)$ chord, 4-cycle, $T-\tau$ 2-path & $S-\tau$ 2-path in M respectively.

Figure 3.3: Example depicting the reduction of a map M that satisfies the 3-1 hypothesis Case 1 and does not contain any chords with one end in $S \cup T$, or $(S \cup T)-(T \cup N_C(\tau))$ 2-paths.

chords and S - τ 2-paths, M' does not contain any $S_{M'}\text{-}T_{M'}$ adjacency. No two $(T_{M'}\setminus T)$ -vertices are adjacent; otherwise, M contains a triangle or an inner 4-cycle, a contradiction. Finally, except possibly (u_1, u_2) , there is no $(T_{M'}\setminus T)\text{-}T$ adjacency; otherwise, M contains a $T\text{-}N_{\mathcal{C}}(\tau)$ chord or a $T\text{-}\tau$ 2-path, a contradiction. Thus, M' is a submap of M that satisfies the $|S|\text{-}0$ or $|S|\text{-}1$ hypothesis. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to that of M by restoring the deleted vertices.

Case 2: $u_6 \in S \cup T$. Case 1 and symmetry allow us to assume $u_1 \in S \cup T$. We distinguish three subcases.

Subcase 1: $u_1, u_6 \in S$. In this subcase, either u_1 and u_6 are adjacent or there exists a vertex $w \in S$ on \mathcal{C} such that both u_1 and u_6 are adjacent to w . Assign a colour to u_2 different from $c(u_1)$, a colour to u_3 different from $c(u_2)$, and a colour to u_4 different from $c(u_3)$. Since $L(u_5) = 3$, assign a colour to u_5 different from $c(u_4)$ and $c(u_6)$. Now \mathcal{C} is a fully-coloured 6- or 7-cycle, so M has an L -colouring by Lemma 3.2.3.

Subcase 2: $u_1 \in S, u_6 \in T$. There are two possibilities.

(i). There is no 3-path between u_3 and u_5 . Assign a color to u_5 that is not available to u_6 , a color to u_4 different from $c(u_5)$, a color to u_3 different from $c(u_4)$ and a color to u_2 different from $c(u_1), c(u_3)$. Obtain a smaller map M' by deleting u_3, u_4, u_5 . Observe that $S_{M'} = S \cup \{u_2\}$. There is no $S_{M'}\text{-}(T_{M'} \setminus T)$ adjacency; otherwise, M contains a $S\text{-}(\tau \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. There is no $(T_{M'} \setminus T)\text{-}T$ adjacency; otherwise, M contains a $T\text{-}(\tau \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. Finally, no two $(T_{M'} \setminus T)$ -vertices are adjacent; otherwise, M contains a triangle, an inner 4-cycle or a 3-path between u_3 and u_5 , a contradiction. So M' is a submap of M that satisfies the $|S_{M'}|\text{-}0$ hypothesis for $|S_{M'}| = |S| + 1$. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to that of M by restoring the deleted vertices.

(ii). M contains a 3-path from u_3 to u_5 . Let a and b be the vertices of this 3-path in $\text{Int}(\mathcal{C})$. If the 5-cycle $C = (u_3, a, b, u_5, u_4, u_3)$ is a separating cycle, we can obtain an L -coloring of M as follows. Obtain a new map M' by deleting $\text{Int}(C)$. M' satisfies the same hypothesis as M and has an L -coloring by the inductive hypothesis. This colors C and an L -coloring of $C \cup \text{Int}(C)$ can be obtained by Lemma 3.2.3. The union of the L -colorings of M' and $C \cup \text{Int}(C)$ gives an L -coloring of M .

In the remaining case, $C = (u_3, a, b, u_5, u_4, u_3)$ is a facial 5-cycle. Assign a color to u_5 that is not available to u_4 , and a color to u_6 different from $c(u_5)$. Obtain a smaller map M' by deleting u_5, u_6 . Observe that $|S_{M'}| = |S|$, and that the degree of u_4 is 1. There is no $(T_{M'} \setminus T)\text{-}S_{M'}$ adjacency; otherwise, M has a chord with an S -end, or an $S\text{-}(T \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. No $(T_{M'} \setminus T)$ -vertex is adjacent to a T -vertex except perhaps the neighbour of u_6 on \mathcal{C} with its other neighbour on \mathcal{C} ; otherwise, M has a chord with a T -end,

or a T - $(T \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. Finally, no two $(T_{M'} \setminus T)$ -vertices are adjacent; otherwise, M contains a triangle or an inner 4-cycle, a contradiction. So M' is a submap of M that satisfies the $|S_{M'}|-0$ or $|S_{M'}|-1$ hypothesis for $|S_{M'}| = |S|$. Obtain its L -coloring by Lemma 3.2.2. This coloring is extended to that of M by restoring the deleted vertices.

Subcase 3: $u_1, u_6 \in T$. If (u_3, x, y, u_5) (or (u_2, x, y, u_4) , by symmetry) is a 3-path, then obtain an L -coloring of M the same way as in subcase 2, (ii). In the remaining case, there are no 3-paths connecting u_2 and u_4 , and u_3 and u_5 . There are two possibilities.

(i). The degree of u_4 is 2. Then assign a colour to u_5 that is not available to u_4 , and a colour to u_6 different from $c(u_5)$. Obtain a smaller map M' by deleting u_5, u_6 and removing their colours from the colour-lists of their neighbours. From M' , obtain a smaller map M'' by deleting u_4 . Observe that $|S_{M''}| = |S|$. No $(T_{M''} \setminus T)$ -vertex is adjacent to an $S_{M''}$ -vertex; otherwise, M has a chord with an S -end, or an S - $(T \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. There is no $(T_{M''} \setminus T)$ - T adjacency, except perhaps the neighbour of u_6 on \mathcal{C} with its other neighbour on \mathcal{C} ; otherwise, M has a chord with a T -end, or a T - $(T \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. Finally, no two $(T_{M''} \setminus T)$ -vertices are adjacent; otherwise, M contains a triangle or an inner 4-cycle, a contradiction. So, M'' is a submap of M that satisfies the $|S|-0$ or $|S|-1$ hypothesis. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to that of M' by restoring u_4 and giving it a colour different from $c(u_3)$ (observe that this possible because u_4 has degree 2 and $c(u_5)$ is not in the list of u_4). Finally, restore u_5 and u_6 to obtain an L -colouring of M .

(ii). The degree of u_4 is greater than 2. Then assign a colour to u_5 that is not available to u_6 , a colour to u_4 different from $c(u_5)$, and a colour to u_3 different from $c(u_4)$. Obtain a smaller map M' by deleting u_3, u_4, u_5 and removing their colours from the colour-lists of their neighbours. Observe that $|S_{M'}| = |S|$. No $(T_{M'} \setminus T)$ -vertex is adjacent to an $S_{M'}$ -vertex; otherwise, M contains an S - $N_{\mathcal{C}}(\tau)$ chord, or an S - $(\tau \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction. There is no $(T_{M'} \setminus T)$ - T adjacency, except perhaps (u_1, u_2) ; otherwise, M contains a chord with a T -end, or a T - $(\tau \cup N_{\mathcal{C}}(\tau))$ 2-path, a contradiction.

Finally, consider adjacencies between two $(T_{M'} \setminus T)$ -vertices. First of all, if u_2 is adjacent to a $(T_{M'} \setminus T)$ -vertex x , then x cannot be adjacent to u_3 or u_4 (because we would then have a triangle or an inner 4-cycle), so x must be adjacent to u_5 . Then $C = (u_2, x, u_5, u_4, u_3, u_2)$ is a 5-cycle. Since u_4 has degree greater than 2, C must be a separating cycle; in that case, we can obtain an L -colouring of M in the same way as in subcase 2, (ii). So, u_2 is not adjacent to any $(T_{M'} \setminus T)$ -vertex. Secondly, no other $(T_{M'} \setminus T)$ -vertices are adjacent to one another; otherwise, M contains a triangle, an inner 4-cycle, or a 3-path between u_3 and u_5 , a contradiction. So, M' is a submap of M that satisfies the $|S|-0$ or $|S|-1$ hypothesis. Obtain its L -colouring by Lemma 3.2.2. This L -colouring is extended to that of M by

restoring the deleted vertices.

This completes the proof of Theorem 3.2.1. \square

Having proved Theorem 3.2.1, we are now in a position to prove Grötzsch's Theorem for a graph G . We proceed by induction. First, by considering blocks of G , we may assume G is 2-connected. Now let M be a map of G in the plane; we separately reduce facial and separating 4-cycles in M . Ultimately, we apply Theorem 3.2.1 to obtain a 3-colouring of G from a particular list assignment L on M .

Theorem 3.2.11 (Grötzsch's Theorem): Every loop-free and triangle-free planar graph has a 3-colouring.

Proof: Let G be a loop-free and triangle-free planar graph. As explained in the beginning of the chapter, we may assume that G does not contain multi-edges. If G is planar graph with no 3-cycles, then this is also true for each block of G . If each block of G is 3-colourable, then so is G , so it suffices to work with each block of G . Now consider M to be a map of the block and let \mathcal{C} be the cycle bounding the infinite face of M . We proceed by induction on the number of vertices in M .

Base Case: For $|V(M)| = 1, 2$ and 3 , a 3-colouring can be trivially obtained by assigning a different colour to each vertex.

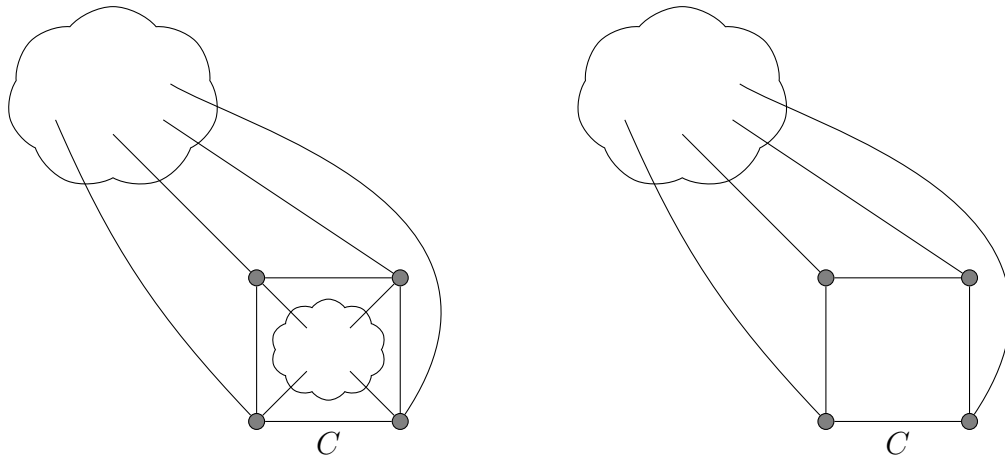
Inductive Step: First we reduce facial and separating 4-cycles in M that do not bound the infinite face in M .

Claim 3.2.12: If M contains a facial 4-cycle, then M has 3-colouring.

Proof: Let (a, b, c, d, a) be a 4-cycle in M . Obtain a smaller map M' from M by collapsing a and c into a single vertex a^* . Obtain another smaller map M'' from M by collapsing b and d into a single vertex b^* . If neither of M' and M'' is triangle-free, there must exist:

- 1) a triangle in M' containing a^* . The other two vertices w and x must be distinct from b and d , otherwise, if $w = b$ for example, then M contains at least one of the two triangles (a, b, x, a) and (c, b, x, c) , a contradiction; and
- 2) a triangle in M'' containing b^* . The other two vertices y and z must be distinct from a and c .

These triangles translate to 5-cycles $C_1 = (a, w, x, c, d, a)$ and $C_2 = (b, y, z, d, c, b)$ respectively in M . Since M is planar, C_1 and C_2 can exist together in M if and only if they have at least one vertex, other than c and d , in common. However, the equalities $w = y, w = z, x = y, x = z$ give rise to a triangles $(a, b, y, a), (w, a, d, w), (y, b, c, y), (x, c, d, x)$ respectively in M , a contradiction. Therefore at least one of M' and M'' is triangle-free.



(a) C is a separating 4-cycle in M such that $\text{Int}(C)$ is minimal. (b) M' is obtained from M by deleting $\text{Int}(C)$.

Figure 3.4: Reducing separating 4-cycles in M .

Without loss of generality, let the triangle-free map be M' . Obtain a 3-colouring of M' by the inductive hypothesis. This 3-colouring is extended to that of M by restoring vertices a, c and giving them the same colour as a^* . \square

Claim 3.2.11: If M contains a separating 4-cycle that does not bound the infinite face, then M has a 3-colouring.

Proof: In view of Claim 3.2.12, we may assume no interior face of M has length 4. Let C be a separating 4-cycle in M that does not bound the infinite face, such that $\text{Int}(C)$ is minimal. Obtain a reduced map M' from M by deleting $\text{Int}(C)$ (Figure 3.4). This does not create any triangles in M' and M' is planar, so M' has a 3-colouring by the inductive hypothesis. This colours $V(C)$ in M ; we may assume the colours are from the set $\{1, 2, 3\}$.

Now consider the submap $\text{Int}(C) \cup C$ in M . We make a list assignment L for $\text{Int}(C) \cup C$ by assigning to each vertex of C the colour it has been assigned in M' and to each vertex of $\text{Int}(C)$ the list $\{1, 2, 3\}$. From Claim 3.2.12, we may assume that no interior face of $\text{Int}(C) \cup C$ is a 4-cycle, and minimality of $\text{Int}(C)$ implies $\text{Int}(C) \cup C$ has no separating 4-cycles. Theorem 3.2.1 implies $\text{Int}(C) \cup C$ has an L -colouring; evidently, this extends the 3-colouring of M' to a 3-colouring of M . \square

By Claims 3.2.12 and 3.2.13, we can assume M has girth at least 4, where C is the only cycle allowed to have length 4. We make a list assignment L for M as follows. To each uncoloured vertex, assign the list $\{1, 2, 3\}$. If C has length 4, colour it completely using

at most three colours from the set $\{1, 2, 3\}$. Now, M satisfies either the 0-0 hypothesis of Theorem 3.2.1, or it has a fully coloured outer boundary of length 4. Use Theorem 3.2.1 to obtain an L -colouring of M , which implies a 3-colouring of M , and hence, that of G . This completes the proof of Grötzsch's Theorem. \square

References

- [1] V.A. Aksionov. On an extension of a 3-coloring on planar graphs (in Russian). *Diskret. Analiz.*, 16:3–19, 1974.
- [2] K. Appel, W. Haken, and J. Koch. Every planar map is four colorable. *Illinois J. Math.*, 21:439–567, December 1977.
- [3] J. A. Bondy and U.S.R. Murty. *Graph Theory*. Springer, New York, 2008.
- [4] J.B. Conway. *Mathematical Connections. A Capstone Course*. American Mathematical Society, 2010.
- [5] R. Diestel. *Graph Theory*. Springer-Verlag, Berlin, third edition, 2005.
- [6] P. Erdős, A.L. Rubin, and H. Taylor. Choosability in graphs. *Proceedings of the West Coast Conference on Combinatorics, Graph Theory and Computing, Congress. Numer.*, 26:125–157, 1979.
- [7] F.G. Tinting maps. *The Athenaeum*, 1389:726, June 10, 1854.
- [8] H. Grötzsch. Ein dreifarbensatz für dreikreisfreie netze auf der kugel. *Math.-Naturwiss.*, 8:109–120, 1958.
- [9] B. Grünbaum. Grötzsch’s theorem on 3-colorings. *Michigan Math. J.*, 10:303–310, 1963.
- [10] P.J. Heawood. Map-colour theorems. *Quart. J. Math. Oxford*, 24:332–338, 1890.
- [11] F. Jaeger. Flows and generalized coloring theorems in graphs. *J. Combin. Theory Ser. B*, 26:205–216, 1979.
- [12] H. Lebesgue. Quelques conséquences simples de la formule d’Euler (French). *J. Math. Pures Appl.*, 19:27–43, 1940.

- [13] J. Mitchem. On the history and solution of the four-color map problem. *College Math. J.*, 12(2):108–116, 1981.
- [14] N. Robertson, P. Seymour, and R. Thomas. Tutte’s edge-colouring conjecture. *J. Combin. Theory Ser. B*, 70(1):166–183, 1997.
- [15] P. Seymour. Nowhere-zero 6-flows. *J. Combin. Theory Ser. B*, 30:130–135, 1981.
- [16] R. Steinberg. *Grötzsch’s Theorem Dualized*. M.Math. Thesis, University of Waterloo, Ontario, Canada, 1976.
- [17] R. Steinberg and D.H. Younger. Grötzsch’s theorem for the projective plane. *Ars Combin.*, 28:15–31, December 1989.
- [18] C. Thomassen. The weak 3-flow conjecture and the weak circular flow conjecture. *J. Combin. Theory Ser. B*, 102(2):521–529, 2012.
- [19] C. Thomassen. A short list color proof of Grötzsch’s theorem. *J. Combin. Theory Ser. B*, 88(1):189–192, May 2003.
- [20] W.T. Tutte. A contribution to the theory of chromatic polynomials. *Canadian J. Math.*, 6:80–91, 1954.
- [21] M. Voigt. A not 3-choosable planar graph without 3-cycles. *Discrete Math.*, 146:325–328, 1995.