Utrecht University

Bachelor Thesis

The four-colour theorem

The basis for a hand-checkable proof

Joris C. Bink (5623138)

Supervisor:
Prof. Dr. Rob H. Bisseling

June 14, 2019

Abstract

The four-colour theorem is easily explained and requires no mathematical knowledge to grasp the concept, yet it took more than a century to prove it. The theorem states that the vertices of every planar graph can be coloured with four colours, such that neighbouring vertices do not have the same colour.

In this thesis we have introduced the theorem. First we limited ourselves to undirected graphs made of triangles. Then we looked at smaller graphs called configurations, and found that we can prove the theorem by finding an unavoidable set of reducible configurations. We have detailed two ways to prove the reducibility of a configuration: D-reducibility, which looks only at the direct surroundings of the configuration, and C-reducibility, which first tries to change the configuration by contracting some of its edges. Proving unavoidability is left as a future project.

With the insight gained during this project we concluded that an increase in computing power will increase the number of reducible configurations that can be found. A recommendation is made for future work to assess the chance of checking unavoidability by hand, as a result of this increase in reducible configurations.

Foreword

I have written this thesis as part of my bachelor in mathematics at Utrecht University. The original goal was to find a hand-checkable proof of the four-colour theorem. This was, however, too ambitious of a project, and due to time constraint I had to limit the scope. I would like to thank my supervisor prof. dr. Rob Bisseling for introducing me to this subject and he has been an immense help during this project.

Contents

1 Introduction

The four-colour theorem concerns the colouring of maps. More precisely, the theorem states that if we were to colour the countries of a map such that neighbouring countries have differing colours, four colours always suffice. This theorem originates from Francis Guthrie in 1852, who came upon it when attempting to colour the counties of a map of England [6].

Over the years many mathematicians attempted to prove the theorem, of note being A.B. Kempe who in 1879 came close to proving the theorem, and H. Heesch, who was the first in proposing many techniques used later. Despite the best efforts, the theorem would not be proved until 1976 by K. Appel and W. Haken [1]. Their proof however, had one flaw: it handled so many cases that it was impossible to check the proof by hand, and a computer had to be used for this purpose. This complicated matters, especially since the theorem has had many faulty proofs in its past. Some were sceptic because hardware and software errors had previously appeared during the development of the proof. In order to convince themselves of the correctness of the proof, N. Robertson, D. Sanders, P. Seymour, and R. Thomas proved the theorem in 1994 with methods similar to those of Appel and Haken [4]. Naturally, the latter proof had the same downside as the former. In 2008 a formal programming proof was given by Georges Gonthier [3], using programs that can be empirically tested on many inputs. As such the case can be considered closed. However the computing power of computers today is much larger than of computers five decades ago, and a completely hand-checkable proof might now be possible.

In this work we will examine a part of the procedure used by Robertson et all., in an attempt to lay the groundwork for such a proof. We will first turn the problem of colouring maps into a problem of colouring graphs. Then we will give some definitions and concepts, followed by the setup for proving reducibility. We will then cover said reducibility in chapter three.

2 Definitions and notation

2.1 From map to graph

The maps considered in the four-colour theorem have certain properties we should mention for completeness: they exist in the Euclidean plane, all countries exist of only one part, and no part of the map belongs to two countries at once.

In order to more easily work with maps, we will make use of graphs. A graph G consists of two sets: $V(G)$, a set of vertices, and $E(G)$, a set of edges that connect two vertices. A vertex represents a country, and an edge connecting two countries represents the fact that those countries are neighbours. Note that the graphs we will consider are undirected, and since a country cannot be its own neighbour, we will not allow loops, which are edges that connect a vertex to itself. We write $e = \{v_1, v_2\}$ for an edge $e \in E(G)$ connecting distinct vertices $v_1, v_2 \in V(G)$. Every vertex has a *degree*, which is the number of vertices it is connected to. For a vertex v in a graph G we write $d(v)$, or $d_G(v)$.

In figure 1 a map is shown with 4 countries, A, B, C, and D, and its corresponding graph. Note that we do not consider countries A and D as neighbours.

Figure 1: a map and its corresponding graph.

Proving the four-colour theorem now means proving that any graph G can be *coloured* with four colours, meaning that we can find a function $\mu : V(G) \to \{A, B, C, D\}$ such that for two distinct vertices $v_1, v_2 \in V(G)$ $\mu(v_1) \neq \mu(v_2)$ if $e = \{v_1, v_2\} \in E(G)$. We will always use the letters A, B, C, D for colours of vertices, and order them lexicographically. We call such a function a *(vertex) four-colouring* of G, or *colouring* for short.

2.2 Further restrictions on our graphs

We can limit the number of graphs we need to consider by adding more restrictions. The first is that the graphs we consider must be *planar*. This means that we can draw the edges in such a way that they do not cross each other. We can see this by taking any map, and drawing any edge connecting A and B through the border they share (figure 2). Since countries do not have multiple parts, edges will only have to cross one border and will therefore not have to cross any other edges. From now on whenever we mention a graph we will assume that it is planar.

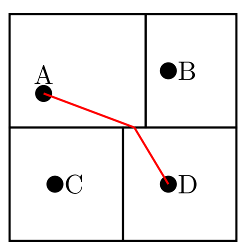

Figure 2: an edge only crossing a relevant border.

A graph G is a *subgraph* of a graph H if $V(G) \subseteq V(H)$ and $E(G) \subseteq E(H)$. G is an *induced subgraph* if, on top of that, for all edges $e = \{v_1, v_2\} \in E(H)$ with $v_1, v_2 \in V(G)$, we have that $e \in E(G)$.

We define a *component* G' of a graph G to be an induced subgraph of G such that there are no $v \in V(G)$ connected to, but not a part of, G'. By connected to G' we mean that there is a $v' \in V(G')$ connected to v by an edge of G. We can assume that all graphs have exactly one component, since the colouring of one component is irrelevant to the colouring of any other component. These components can therefore be treated as separate cases.

A *cycle* C is a sequence of vertices and edges $v_0, e_1, v_1, ..., e_k, v_k$, such that for all $1 \leq n \leq k$ edge $e_n = \{v_{n-1}, v_n\}$, and the only repeating vertex is $v_0 = v_k$. A cycle divides a graph in an inside and an outside, as in figure 3.

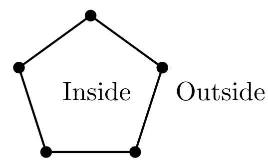

Figure 3: The inside and outside of a cycle.

If a face f of a graph is next to exactly three distinct edges, we refer to it as a *triangle*. A graph G is a *triangulation* if every face of G is a triangle. A *near-triangulation* is a triangulation, of which one face is designated as infinite. The infinite face does not have to be a triangle, see figure 4.

We consider triangulations because of the following:

Lemma 2.1. *Any graph that is not a triangulation can be turned into one by adding edges and vertices.*

Proof. Let G be a graph, not a triangulation. Then it has one or more faces which are incident with either exactly 2, or more than 3 edges. We will show that any such face can be divided into multiple triangles by adding edges and vertices, thereby proving the lemma.

Let f be a face as specified above, C the smallest cycle around f, with $\{v_0, ..., v_{n-1}\}$ the set of its vertices. Place a vertex v inside of f, and connect it to v_k for all $0 \leq k \leq n-1$. The result is that f is divided into n triangles (figure 5). □

It is easy to see that if a graph G has a valid colouring, removing an edge, or vertex along with all edges incident to it, cannot make the colouring of the remaining graph invalid. As such, we need only prove the four-colour theorem for all triangulations.

Figure 4: On the left is a triangulation, and on the right a near-triangulation. Note that the space outside of a graph is considered a face.

Figure 5: Adding a vertex and edges to divide a single face into five triangles.

2.3 Minimal criminal

Let $F(G)$ be the set of all faces of a triangulation G. Then we have by Euler's formula that $|V(G)| - |E(G)| + |F(G)| = 2$. Furthermore, because every face has three edges, we know that $3|F(G)| = 2|E(G)|$, being careful not to count edges twice. These two equations combined give us the following: $|V(G)| = 2 + \frac{1}{3}|E(G)|$. We can now sensibly define a *minimal criminal* as a triangulation G that is not four-colourable such that any triangulation G' with $|V(G')| < |V(G)|$ is four-colourable. It is then our goal to show that G does not exist.

A *short cycle* is a cycle C in a graph G with length $|E(C)| \leq 5$ such that the removal of C from G creates multiple components, at least one on the inside and one on the outside (figure 6). In the case that $|E(C)| = 5$ both the inside and outside must contain at least 2 vertices.

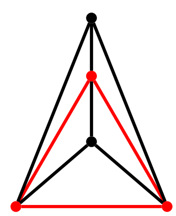

Figure 6: The red cycle is a short cycle. Its removal leaves two vertices, one of which is on the inside of the cycle, the other on the outside.

We then have the following:

Theorem 2.2. *Every minimal criminal is a triangulation T that does not contain a short cycle [2].*

The proof is rather lengthy, and we omit it here.

Note that this means that a minimal criminal cannot have vertices with degree less than 5, because such a vertex would have a cycle of length less than 5 around it. This causes a short cycle unless the other side of the cycle has no vertices, but such a case is trivial to colour.

2.4 The setup

In order to prove the four-colour theorem we will be considering smaller components of graphs, called *configurations*.

Def 2.1. A configuration K is a pair $(G(K), \gamma_K)$, where $G(K)$ is a near-triangulation and the charge function γ_K is a function: $\gamma_K \colon V(G(K)) \to \mathbb{Z}_+$ with the following properties:

i. for every vertex v, $G(K)$ without v has at most two components, and if there are two, then $\gamma_K(v) = d(v) + 2$ (figure 7),

ii. for every vertex v, if v is next to the infinite region $\gamma_K(v) > d(v)$, and otherwise $\gamma_K(v) = d(v)$, and

iii. K has ring-size ≥ 2, where the *ring-size* of K is defined to be $\sum_v (\gamma_K(v) - d(v) - 1)$, summed over all vertices v next to the infinite region such that $G(K) \backslash v$ is connected.

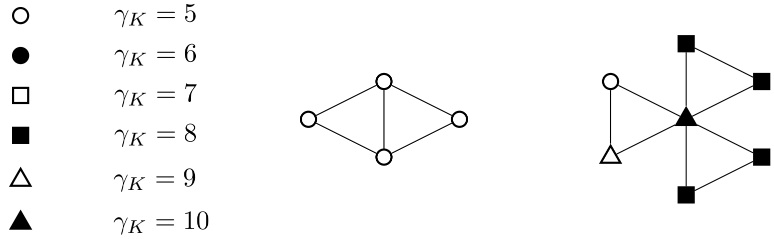

Figure 7: The left figure is a configuration, in the right figure the vertex in the middle breaks the first property.

A configuration K has a *free completion S with ring R* if S is a near-triangulation with the following properties (figure 8):

i. $G(K)$ is an induced subgraph of S,

ii. R is a cycle of S and an induced subgraph of S, such that R and $G(K)$ are disjoint and $V(G(K)) \cup V(R) = V(S)$ and all vertices and edges of R are next to the infinite region of S,

iii. $d_S(v) = \gamma_K(v)$ for every vertex $v \in V(G(K))$.

It is possible to prove that every configuration has a free completion, and that this free completion is unique. We omit that proof here. We will use the free completion to make statements about the surroundings of a configuration.

Lemma 2.3. *The ring-size of a configuration K, with free completion S and ring R, is equal to the length of its ring R.*

Proof. Every vertex $v \in V(G(K))$ has to be connected to $\gamma_K(v) - d_{G(K)}(v)$ vertices of R, because $d_S(v) = \gamma_K(v)$ and $V(R) \cup V(G(K)) = V(S)$. Let $v_1, ..., v_k$ be the vertices of R that are connected to v, in order. Note that v_n is connected to v_{n+1}, and v_1 and v_k are connected to neighbours of v, because S is a triangulation. It is not possible for $v_2, ..., v_{k-1}$ to be connected to any other vertices of $G(K)$, because then some vertices or edges of R would not be next to the infinite region. It is also not possible for R to have vertices which are not at all connected to vertices of $G(K)$, because R is both a cycle and an induced subgraph of S. We must be careful not to count vertices twice, so we count $\gamma_K(v) - d(v) - 1$. The exceptions are vertices w such that $G(K)$ without w has two components, since these share twice as many vertices (figure 9). By the first property of a configuration $\gamma_K(w) = d(w) + 2$, so we need not count them at all, and the lemma follows. □

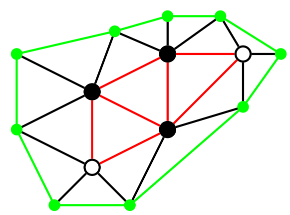

Figure 8: With the vertex labels the same as figure 7, the red edges and their vertices form a configuration K, the green vertices and edges denote the ring R of the free completion, and the graph as a whole is the free completion S.

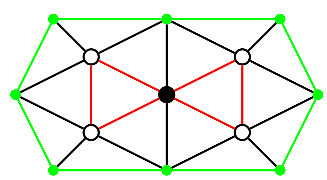

Figure 9: Following the style of figure 8, the vertex w in the middle is connected to two vertices of R, both of which are connected to two other vertices of $G(K)$. Thus these two vertices of R are already counted when we count over all the other vertices of $G(K)$.

A configuration K *appears* in a triangulation T if $G(K)$ is an induced subgraph of T and $d_T(v) = \gamma_K(v)$ for all $v \in V(G(K))$. Since a minimal criminal cannot contain vertices with a degree less than 5, we need only to consider configurations K with $\gamma_K(v) \geq 5$ for all $v \in V(G(K))$.

Def 2.2. A configuration K is *reducible* if every triangulation T in which K appears cannot be a minimal criminal.

Def 2.3. A set of configurations U is *unavoidable* if in every triangulation T there appears at least one configuration $K \in U$ in T.

With these definitions we can give an overview of the proof of the four-colour theorem. We will find a set of reducible configurations U which we will prove to be unavoidable. Then a minimal criminal cannot exist, for if T were a minimal criminal, there would be some configuration $K \in U$ that appears in T. Since K is reducible and T is a triangulation, this leads to a contradiction, and thus T is not a minimal criminal.

There are two reasons as to why the proofs by Appel and Haken, and Robertson et al. were impossible to check. The first is that their unavoidable sets have sizes above 1400 and 600 respectively. Every configuration of these sets need to be checked for reducibility. The second reason is that checking that a set is unavoidable is an enormous task.

3 Reducibility

3.1 Introduction

We will use two kinds of reducibility to prove that a configuration is reducible: D-reducibility and a special case of C-reducibility, both of which we will define later. The definitions and theorems will largely follow [4], however we have changed the order to better explain the two kinds of reducibility. We will not be talking about colouring vertices in this chapter, rather, we will be colouring edges.

An *edge-colouring* of a graph G is a function $\kappa : E(G) \to \{\alpha, \beta, \gamma\}$. We will always use α, β, γ for colours of edges. A triangle t of G is *tri-coloured* by κ if $\{\kappa(a), \kappa(b), \kappa(c)\} = \{\alpha, \beta, \gamma\}$, where a, b, c are the three edges of t. If every triangle of a triangulation T or near-triangulation H is tri-coloured by κ, then κ is a *tri-colouring of* T respectively H. In images concerning the colouring of edges, we will henceforth use green for α, red for β, and black for γ.

In [4] tri-colourings are used instead of four-colourings because those are easier to manipulate in a computer. For this reason we will be using tri-colourings as well, but we first need to show that such an approach will work. To do so we make use of sequences $e_0, f_1, e_1, f_2, ..., e_k$ of edges and faces of a triangulation T with tri-colouring κ, with the following properties:

i. all e_n and all f_n are distinct,

ii. for all $0 \leq n \leq k-1$: e_n is next to f_{n+1} which is next to e_{n+1},

iii. for all $0 \leq n \leq k$: $\kappa(e_n) = \alpha$ or $\kappa(e_n) = \beta$,

iv. $e_0 = e_k$.

We call such a sequence an *$\alpha\beta$-circuit* (figure 10). Analogous is the definition of an $\alpha\gamma$-circuit. Note that any edge e with $\kappa(e) = \alpha$ is part of two circuits; an $\alpha\beta$-circuit and an $\alpha\gamma$-circuit, and any two distinct $\alpha\beta$-circuits are completely disjoint and hence cannot cross each other.

Later on in this chapter we will make use of $\alpha\beta$-ribs in near-triangulations H, which are sequences like $\alpha\beta$-circuits with the difference that $e_0 \neq e_k$, and both e_0 and e_k need to be next to the infinite face of H. Like circuits, ribs cannot cross each other. Note that not all edges of H need to be part of a rib, as some can be part of a circuit instead. Lastly we would like to note that in a given $\alpha\beta$-circuit or rib we can swap the colours of the α-coloured and β-coloured edges to create a new tri-colouring (figure 11).

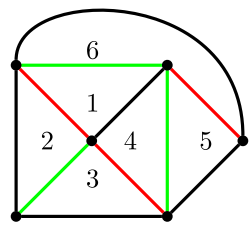

Figure 10: A circuit following the green and red edges and numbered faces.

The following theorem is from [4], where they cite [5] by P.G. Tait as its origin. However, Tait claims only that a four-colouring can be turned into a tri-colouring, and not the other way around. Furthermore, Tait does not provide any proof of this claim, so we provide it here.

Theorem 3.1. *Let T be a triangulation. Then T is four-colourable if and only if it has a tri-colouring.*

Proof. Assume that we have a four-colouring μ of T. We will construct an edge-colouring κ of T as follows: for any edge $e = \{v_1, v_2\} \in E(T)$ where $\mu(v_1) < \mu(v_2)$:

1. $\kappa(e) = \alpha$ if $\mu(v_1) = A, \mu(v_2) = B$ or $\mu(v_1) = C, \mu(v_2) = D$,
2. $\kappa(e) = \beta$ if $\mu(v_1) = A, \mu(v_2) = C$ or $\mu(v_1) = B, \mu(v_2) = D$,
3. $\kappa(e) = \gamma$ if $\mu(v_1) = A, \mu(v_2) = D$ or $\mu(v_1) = B, \mu(v_2) = C$.

We now show that κ is a tri-colouring. Assume the contrary. Then T has a triangle where at least two of its edges have the same colour. Let v_1, v_2, v_3 be the three distinct vertices of this triangle, and let $e_1 = \{v_1, v_2\}, e_2 = \{v_2, v_3\}, e_3 = \{v_3, v_1\}$ be its edges. Assume, without loss of generality, that $\kappa(e_1) = \kappa(e_2) = \alpha$. These edges are both next to vertex v_2. We can assume, again without loss of generality, that $\mu(v_2) = A$. Then we know that $\mu(v_1) = B = \mu(v_3)$, see figure 12. However, v_1 and v_3 are connected by e_3 and μ is a four-colouring, so this is not possible. It follows that κ is a tri-colouring.

Figure 11: The colours of the green and red edges have been swapped in the second figure.

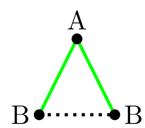

Figure 12: If two edges have the same colour, the third edge connects two vertices with the same colour.

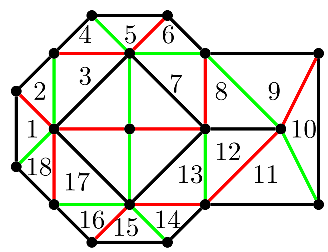

Figure 13: The numbered faces are part of an $\alpha\beta$-circuit, which connects the component formed by γ-edges on the inside with the component formed by γ-edges on the outside.

We now assume we have a tri-colouring κ of T. We divide T into components by removing all the edges of all the $\alpha\beta$-circuits, and label all vertices with either a 1 or a 0. All vertices of the same component get the same label, and if two components were connected by an edge from an $\alpha\beta$-circuit, their vertices do not get the same label.

We can see that this is possible, because a circuit is always next to exactly two components, one on the inside, the other on the outside (figure 13). Then, if a circuit L connects two components with the same label, we can switch the labels of all components inside L to solve the problem, because L is a boundary between the inside and outside.

We then do the same with the $\alpha\gamma$-circuits. Now all vertices of T are labeled in one of four ways: $11, 10, 01, 00$. The first digit corresponds to the label given by the $\alpha\beta$-circuit, and the second digit corresponds to the label given by the $\alpha\gamma$-circuits. Note that if two vertices are connected by an edge with colour β, we know that the first digit in their labels must be different, the second digit must be different if the vertices were connected by a γ-edge, and both digits of the label must be different if the vertices were connected by an α-edge. It follows that if two vertices are connected, they must have different labels, and thus we can transform these labels into a colouring in the obvious way.

We have now shown that we can create a tri-colouring from a four-colouring and a four-colouring from a tri-colouring, which proves the theorem. $\square$

3.2 D-reducibility

In this section we will cover D-reducibility. Let us first define some notation to prevent redundancy: for the rest of this section, T is a triangulation, of which we assume that it is a minimal criminal. K is a configuration that appears in T, and has free completion S with ring R. H is the graph obtained from T by deleting $G(K)$, and all edges connected to it, from T. It is possible to prove that H is a near-triangulation, of which the infinite face is the face which previously contained $G(K)$. We omit that proof here. Furthermore, by lemma 2.1 we can turn H into a triangulation H' . Since we assume T is a minimal criminal, it follows that H' has a tri-colouring, and thus, H has a tri-colouring.

As mentioned previously, R represents the direct surroundings of K. However, if K appears in T we cannot guarantee that S is a subgraph of T, which will cause problems when we try to apply statements about R, to T. We overcome this challenge with the following:

Since H is a near-triangulation, there is a closed walk $v_0, a_1, ..., a_k, v_k$ that traces the infinite face of H, in the natural sense, where v_n are the vertices and a_n are the edges. Assume for now that R is a cycle of length k with edges $b_1, ..., b_k$ in order. Define a function ϕ for $1 \leq n \leq k$ as $\phi(b_n) = a_n$, and let κ be a tri-colouring of H; then $\lambda(b_n) = \kappa(\phi(b_n))$ is an edge-colouring of R. We call λ a *lift* of κ by ϕ, and say that ϕ *wraps* R *around* H (figure 14).

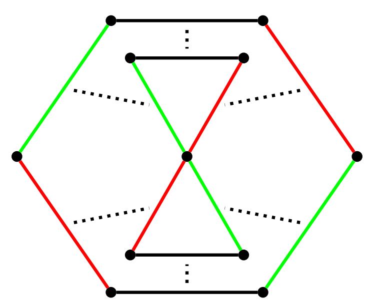

Figure 14: The outer cycle is R, the inner cycle is the walk around the infinite face of H. The dotted line is the function ϕ that wraps R around H.

We just now assumed that R has length k, but this need not be the case. However, in [4] Robertson et al. state as a "folklore theorem" that if a configuration K appears in a triangulation T, there is a function $\psi : V(S) \cup E(S) \cup F(S) \to V(T) \cup E(T) \cup F(T)$ with the following properties:

i. ψ maps $V(S)$ into $V(T)$, $E(S)$ into $E(T)$, $F(S)$ into $F(T)$,

ii. for distinct $v, w \in V(S)$, if $\psi(v) = \psi(w)$ then v and w are both next to the infinite region, likewise for distinct $e, f \in E(S)$, and for all distinct $r, s \in F(S)$ $\psi(r) \neq \psi(s)$,

iii. for $x, y \in S$, if x, y are next to each other in S, then $\psi(x), \psi(y)$ are next to each other in T, and

iv. $\psi(x) = x$ for all $x \in V(G(K)) \cup E(G(K)) \cup F(G(K))$.

No proof of this theorem is given. We call ψ a *corresponding projection of S into T*. We then have from [4] the following theorem:

Theorem 3.2. *Let ψ be a corresponding projection of S into T. Then the restriction of ψ to $E(R)$ wraps R around H.*

For the sake of brevity we omit its proof here. This theorem validates our assumption about R, and allows us to use R to make statements about the surroundings of $G(K)$. From now on ψ is a corresponding projection of S into T, and ϕ is its restriction to R. We can now prove the following lemma:

Lemma 3.3. *Let κ be a tri-colouring of H, and λ a tri-colouring of S. If the lift of κ by ϕ is equal to the restriction of λ to R, then there is a tri-colouring of T.*

Proof. Note that T has been divided into three parts: H, $G(K)$, and a set of edges that connect H and $G(K)$, say, Z.

We create the tri-colouring κ' of T in the following way: for any edge $e \in E(T)$:

i. $\kappa'(e) = \kappa(e)$ if $e \in E(H)$,

ii. $\kappa'(e) = \lambda(e)$ if $e \in E(S)$,

iii. otherwise $e \in Z$, let $\psi(z') = e$, then $\kappa'(e) = \lambda(z')$.

It is easy to see that there will be no problems within either H or $G(K)$. Let $z_1, z_2 \in E(S)$, $f \in F(S)$. If $\psi(z_1)$ and $\psi(z_2)$ are next to the face $\psi(f)$, then z_1 and z_2 are next to f. It follows that $\lambda(z_1) \neq \lambda(z_2)$, hence $\kappa'(\psi(z_1)) \neq \kappa'(\psi(z_2))$. Since for every edge $z \in Z$ there is an edge $s \in E(S)$ such that $\psi(s) = z$, we can conclude that κ' is a tri-colouring of T. $\square$

It is now our goal to prove that, if K is D-reducible, there are tri-colourings κ' of H and λ' of S such that the lift of κ' by ϕ to R is the restriction of λ' to R. To do so we need the following definitions:

A *signed match* of R is a pair (m, u) where $m = \{a, b\} \subseteq R$ is a set of two distinct edges and $u \in \{-1, 1\}$ is the *sign* of the pair. A *signed matching* M of R is a set of signed matches such that for all distinct $(\{a, b\}, u), (\{a', b'\}, u') \in M$:

i. $\{a, b\} \cap \{a', b'\} = \varnothing$, and

ii. edges a and b belong to the same component of the graph obtained by deleting a' and b' from R.

Figure 15 shows a valid and an invalid signed matching. As the image shows, the second property can be seen as: if the cycle is drawn as a regular polygon, straight lines that connect edges of a match do not cross. This property allows us to represent ribs of H using signed matches, using the concept of fitting.

Figure 15: The dashed and dotted lines represent signed matches. The edges connected by the lines are the edges of a match, $u = 1$ if the line is dashed, and $u = 0$ if the line is dotted.

Let M be a signed matching of R, then $E(M) = \{a | (\{a, b\}, u) \in M \text{ for some } b, u\}$ where a and b are edges of R. For $\theta \in \{\alpha, \beta, \gamma\}$, an edge-colouring δ *θ-fits* M if:

i. $a \in E(M)$ if and only if $\delta(a) \neq \theta$, and

ii. for every $(\{a, b\}, u) \in M$: $\delta(a) = \delta(b)$ if and only if $u = 1$.

Figure 16 shows two edge-colourings that γ-fit a matching.

Figure 16: Using the same style as figure 15, the two lines form a signed matching M. If the black edges represent edges with colour γ, then both edge-colourings γ-fit M.

Like we mentioned earlier, we will use signed matches to represent ribs of H. For example, we can represent all the $\alpha\beta$-ribs of a tri-colouring κ with a signed matching M, by adding a signed match whose edges are the ends of a rib, and whose sign follows the second property of θ-fitting. Recall that we can flip the colours of the edges of a rib to create a new tri-colouring of H. We will use these properties to create a set of edge-colourings of R that are the lift of tri-colourings of H. We will explain why afterwards.

Lemma 3.4. *Let κ be a tri-colouring of H, with δ its lift to R by ϕ. Then for all $\theta \in \{\alpha, \beta, \gamma\}$ there is a signed matching M of R such that δ θ-fits M, and all edge-colourings that θ-fit M are the lift of tri-colourings of H by ϕ.*

Proof. Assume without loss of generality that $\theta = \gamma$. Then we first construct a signed matching M such that δ γ-fits M. A signed match $(\{a, b\}, u) \in M$ on the following conditions:

i. There is an $\alpha\beta$-rib $e_0, ..., e_k$ in H such that $\phi(a) = e_0$, $\phi(b) = e_k$,

ii. $u = 1$ if and only if $\delta(a) = \delta(b)$.

Because distinct ribs are disjoint and cannot cross each other, M is a valid signed matching. It is also easy to check that δ γ-fits M by construction.

Now let δ' be an edge-colouring of R that γ-fits M. We create a tri-colouring κ' of H whose lift by ϕ is δ'. We do so from κ by flipping the colours in $\alpha\beta$-ribs $\phi(a), ..., \phi(b)$ if $\delta'(a) \neq \delta(a)$. The result is that $\delta'(a) = \kappa'(\phi(a))$, and $\delta'(b) = \kappa'(\phi(b))$, because $\delta'(a) = \delta(a)$ if and only if $\delta'(b) = \delta(b)$ since both δ and δ' γ-fit M. So we see that $\kappa'(\phi(a)) = \delta'(a)$ for all edges of R, and thus δ' is the lift of κ' by ϕ, as we wanted. $\square$

Let Θ be a set of edge-colourings of R. Θ is *consistent* if for every $\theta \in \{\alpha, \beta, \gamma\}$, and every $\delta \in \Theta$, there is a matching M such that δ θ-fits M and every edge-colouring δ' that θ-fits M is in Θ. Note that the empty set is consistent, and that the union of two consistent sets is itself consistent, so every set of edge-colourings Θ has a maximal consistent subset Θ'.

Note that, by repeatedly applying lemma 3.4, we can create a subset Θ of edge-colourings of R, such that all edge-colourings of Θ are the lift of some tri-colouring of H. We know that H has a tri-colouring κ, so we create Θ by adding all the edge-colourings found by applying lemma 3.4 on κ to it, and then applying the lemma to all edge-colourings already in Θ. We continue to do so until no further edge-colourings are found. By construction Θ is consistent. We then define D-reducibility as follows:

Def 3.1. Let K be a configuration with free completion S with ring R. Let Θ^* be the set of all edge-colourings of R, and $\Theta_1 \subseteq \Theta^*$ be the set of all restrictions to $E(R)$ of tri-colourings of S. Let Θ' be the maximal consistent subset of $\Theta^* - \Theta_1$. If $\Theta' = \varnothing$, then K is *D-reducible.*

The reason we went with such a seemingly roundabout approach is that a computer can easily generate the maximal consistent subset Θ', which allows us to quickly find reducible configurations. We will also make use of this property when defining C-reducibility.

It is now straightforward to prove that D-reducibility implies reducibility:

Theorem 3.5. *Let K be D-reducible. Then K is reducible.*

Proof. To prove that K is reducible we need to prove that T is not a minimal criminal. Recall that we have assumed that T is a minimal criminal, so that H has a tri-colouring κ. Then there is a consistent set Θ of edge-colourings of R, such that every edge-colouring of Θ is the lift of some tri-colouring of H. Note that Θ is nonempty. If $\Theta \subseteq \Theta^* - \Theta_1$, then $\Theta \subseteq \Theta'$, however, K is D-reducible so $\Theta' = \varnothing$. Therefore, $\Theta \cap \Theta_1 \neq \varnothing$, and thus there is some tri-colouring κ' with lift δ' such that there is a tri-colouring of S whose restriction to R is δ'. We then have by lemma 3.3 that there is a tri-colouring of T, which contradicts the assumption that T is a minimal criminal. We can therefore conclude that T cannot be a minimal criminal, and thus D-reducibility implies reducibility. $\square$

3.3 C-reducibility

Not every configuration is D-reducible, so in order to increase the number of possible configurations for our unavoidable reducible set, we can use another form of reducibility: C-reducibility. For a configuration K we will find a set of edges $X \subseteq E(G(K))$ with some special properties defined later. We then create a new configuration K' by contracting the edges in X as follows: for every edge $e = \{v_1, v_2\} \in X$, we remove v_1 and v_2 and all their edges from $G(K)$. We then replace the vertices with a new vertex v, and any edge $\{v_1, v_k\}$ or $\{v_2, v_k\}$ with $\{v, v_k\}$ (except $\{v_1, v_2\}$, to prevent loops). We also need to change the charge-function to ensure that the free completion S' of K' will have the same ring $R' = R$, see figure 17. The smaller number of vertices means that the graph in which K has been replaced by K' can be tri-coloured, and then we can find a colouring of K because of the special properties of X.

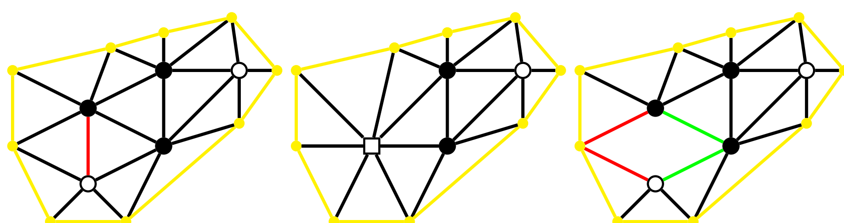

Figure 17: The left graph is a configuration with its ring. The middle graph shows what would happen if we were to contract the red edge, and the right graph shows part of a tri-colouring modulo X.

In order to prevent notational difficulties, we will not speak of actually transforming a configuration into another during the rest of this section. Instead, we will introduce a different kind of edge-colouring.

Let T be a triangulation or a near-triangulation. A set of edges $X \subseteq E(T)$ is *sparse in* T if every face of T is next to at most one edge of X, and in the case of a near-triangulation, the infinite face is next to no edges of X. A *tri-colouring modulo* X of T is a function $\kappa : E(T) - X \to \{\alpha, \beta, \gamma\}$ with the following properties for distinct edges $e_1, e_2, e_3 \in E(T)$ of the same face:

i. if $e_1, e_2, e_3 \in E(T) - X$, then $\{\kappa(e_1), \kappa(e_2), \kappa(e_3)\} = \{\alpha, \beta, \gamma\}$, and

ii. if $e_1 \in X$, then $\kappa(e_2) = \kappa(e_3)$.

The following will allow us to speak about the colouring of the contracted triangulation without explicitly mentioning it. The restriction that there is no cycle C such that $|E(C) - X| = 1$ prevents certain issues, see figure 18.

Theorem 3.6. *Let T be a minimal criminal, and let $X \neq \varnothing$ be sparse in T such that there is no cycle C of T with $|E(C) - X| = 1$. Then T has a tri-colouring modulo X.*

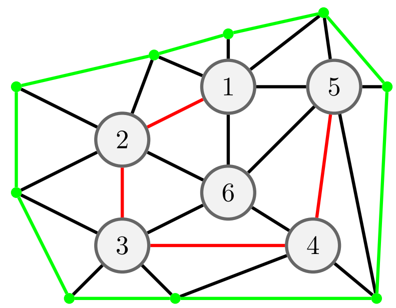

Figure 18: If we were to contract the red edges, the vertices 1 through 5 would be replaced by a single vertex. However, the edge connecting vertices 1 and 5 would not be removed, resulting in an edge that connects a vertex to itself.

Proof. Let $G = (V(T), X)$ be the graph obtained by deleting all edges of $E(T)$ not in X. Then G has multiple components $Z_1, ..., Z_k$. Let Y be the graph with vertices $v_1, ..., v_k$, and edges $E(Y)$, where $e = \{v_n, v_m\} \in E(Y)$ if some $e' \in E(T) - X$ connects Z_n to Z_m. Note that Y is loopless, because there is no cycle C with $|E(C) - X| = 1$. Y is also planar, because it essentially is the contracted version of T. Since $X \neq \varnothing$, Y has less vertices than T, and therefore has a vertex four-colouring. Hence we can find a function $\lambda : V(T) \to \{A, B, C, D\}$ such that $\lambda(v) = \lambda(v')$ if $v, v' \in Z_n$ for all $1 \leq n \leq k$, and $\lambda(v_1) \neq \lambda(v_2)$ for $\{v_1, v_2\} \in E(T) - X$. We then define $\kappa : E(T) - X \to \{\alpha, \beta, \gamma\}$ in the same manner as in theorem 3.1. Let e_1, e_2, e_3 be the three edges of a triangle of T, with vertices v_1, v_2, v_3. If $e_1, e_2, e_3 \in E(T) - X$, then $\{\kappa(e_1), \kappa(e_2), \kappa(e_3)\} = \{\alpha, \beta, \gamma\}$ because $\lambda(v_1)$, $\lambda(v_2)$, and $\lambda(v_3)$ are distinct. If, without loss of generality, $e_1 = \{v_1, v_2\} \in X$, then $\lambda(v_1) = \lambda(v_2)$ and thus $\kappa(e_2) = \kappa(e_3)$. Hence κ is a tri-colouring modulo X of T. $\square$

Let K be a configuration with free completion S and ring R. Let Θ be the set of all edge-colourings of R, $\Theta_2 \subseteq \Theta$ be the set of all restrictions of tri-colourings of S to R, and Θ_3 be the maximal consistent subset of $\Theta - \Theta_2$. Then a *contract* for K is a set $X \neq \varnothing \subseteq E(S) - E(R)$ that is sparse in S, such that no edge-colouring of Θ_2 is the restriction to R of a tri-colouring modulo X of S.

For a vertex v and an edge e we say that e *faces* v if e and v are next to the same triangle, but e is not next to v, see figure 19. A vertex v of S is a *triad* for X if $v \in V(G(K))$, there are at least three distinct vertices of S connected to v that are next to a member of X, and if $\gamma_K(v) = 5$, not every member of X faces v.

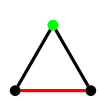

Figure 19: The red edge faces the green vertex.

Theorem 3.7. *Let T be a minimal criminal, and K be a configuration appearing in T with free completion S. Let ϕ be a corresponding projection of S into T, and let X be a contract for K with $|X| \leq 4$ such that if $|X| = 4$, there is a triad for X. Then there is no cycle C of T such that $|E(C) - \phi(X)| \leq 1$.*

Proof. Assume the contrary, let C be such a cycle.

Note that $\phi(X)$ is sparse in T, since if two distinct edges $\phi(x_1), \phi(x_2) \in \phi(X)$ were next to the same face $\phi(f)$, then x_1 and x_2 would be next to f because ϕ is a corresponding projection. This is not possible because X is sparse in S. Let X' be the intersection of the edges of C and $\phi(X)$. Then X' is also sparse in T. This means that the edges of X' are next to distinct triangles inside C. If $|E(C)| \leq 4$ this would mean that C is a short cycle, but T is a minimal criminal hence $|E(C)| \geq 5$. Since $|E(C)| - |X| \leq |E(C)| - |X'| \leq 1$ and $|X| \leq 4$, we know that $|E(C)| = 5$ and $|X| = 4$, and there is a vertex $t \in V(T)$ with $d_T(t) = 5$ inside C. Note that every edge of C faces t.

Because $|X| = 4$, there is a triad $v \in V(G(K))$ for X. Either some edge of X does not face v, or $d_S(v) \geq 6$, so some edge of X' does not face $\phi(v)$, or $d_T(\phi(v)) \geq 6$ because $v = \phi(v)$. In either case $v \neq t$. We know that v is connected to three distinct vertices of S, say v_1, v_2, v_3, in order, that are next to an edge of X'. Naturally these vertices are a part of C, and therefore are connected to t. We now have a short cycle C' of T with vertices v, v_1, t, v_3, with v_2 inside C', see figure 20, which by theorem 2.2 leads to a contradiction. □

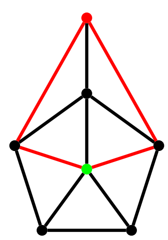

Figure 20: The red vertex is v, the green vertex is t. The red edges form a short cycle.

We can now define C-reducibility and prove that it implies reducibility.

Def 3.2. Let K be a configuration with free completion S. K is *C-reducible* if there is a contract X for K with $1 \leq |X| \leq 4$ such that if $|X| = 4$, there is a triad for X.

Theorem 3.8. *Let K be a C-reducible configuration. Then K is reducible.*

Proof. Let T be a minimal criminal. We show that K cannot appear in T. Assume the contrary, K appears in T. Let S be the free completion of K with ring R, and ϕ a corresponding projection of S into T. Let H be obtained as in theorem 3.2, and let ψ be the restriction of ϕ to R.

Since K is C-reducible, it has a contract X with $|X| \leq 4$, and if $|X| = 4$ there is a triad for X. Then by theorem 3.7 the is no cycle C of T such that $|E(C)-\phi(X)| = 1$, hence by theorem 3.6 there is a tri-colouring modulo X of T, say κ. The restriction of κ to H is a tri-colouring of H, since $E(H)\cap\phi(X) = \varnothing$, and its lift to R via ψ, say λ, belongs to either Θ_2, in which case T is tri-colourable by lemma 3.3, or Θ_3. However, for $e \in E(S)$, let $\kappa'(e) = \kappa(\phi(e))$. Then κ' is a tri-colouring modulo X of S, and λ its restriction to R. Which leads to a contradiction with the definition of C-reducible, hence K cannot appear in T, and is therefore reducible. □

4 Conclusion and discussion

We have explained in detail what D- and C-reducibility are, but have not yet discussed the practical side of things. Both types of reducibility have been defined in terms of consistent sets of edge-colourings, so a computer can easily generate the maximal consistent sets needed to check for D- and C-reducibility. Robertson et al. restricted their search space to accommodate their limited computing power, but we have no need to do so. We can check larger configurations and ring-sizes, resulting in more reducible configurations to chose from for our unavoidable set, and perhaps a set that is easier to check.

We leave the details of proving unavoidability for later, suffice it to say for now that Robertson et al. used a branch-and-bound approach, which results in a large number of cases to be checked. How an increase in availlabe reducible configurations will impact the amount of cases is something to be researched.

Another possible future project might be the implementation of other types of reducibility. So far we have only two types, but this can be expanded upon.

Robertson et al. created an algorithm which can colour any triangulation in $O(|V(T)|^2)$ time, based on their proof of the theorem. The step that requires the most time is finding a short cycle. If this step can be improved, a faster algorithm will be the result.

References

[1] K. Appel and W. Haken. "Every planar map is four colorable. Part I: Discharging". In: *Illinois Journal of Mathematics* 21.3 (Sept. 1977), pp. 429–490. DOI: `10.1215/ijm/1256049011`.

[2] George D. Birkhoff. "The Reducibility of Maps". In: *American Journal of Mathematics* 35.2 (Apr. 1913), p. 115. DOI: `10.2307/2370276`.

[3] Georges Gonthier. "Formal Proof—The Four- Color Theorem The Tale of a Brainteaser". In: 55.11 (2008), pp. 1382–1393. URL: `https://www.ams.org/notices/200811/tx081101382p.pdf`.

[4] Neil Robertson et al. "The Four-Colour Theorem". In: *Journal of Combinatorial Theory, Series B* 70.1 (May 1997), pp. 2–44. DOI: `10.1006/JCTB.1997.1750`.

[5] P. G. Tait. "Note on a Theorem in Geometry of Position." In: *Transactions of the Royal Society of Edinburgh* 29.02 (Jan. 1880), pp. 657–660. DOI: `10.1017/S008045680002617X`.

[6] Robin J. Wilson. *Four colors suffice : how the map problem was solved.* Princeton University Press, 2002, p. 262. ISBN: 0691120234.