Finding Paths Between 3-Colourings

Luis Cereceda [1], Jan van den Heuvel [1] and Matthew Johnson [2] *

[1] Department of Mathematics, London School of Economics
Houghton Street, London WC2A 2AE, U.K.

[2] Department of Computer Science, University of Durham
Science Laboratories, South Road, Durham DH1 3LE, U.K.

email: {luis,jan}@maths.lse.ac.uk, matthew.johnson2@durham.ac.uk

CDAM Research Report LSE-CDAM-2007-31 – December 2007

Abstract

Given a 3-colourable graph G and two proper vertex 3-colourings α and β of G, consider the following question: is it possible to transform α into β by recolouring vertices of G one at a time, making sure that all intermediate colourings are proper 3-colourings?

We prove that this question is answerable in polynomial time. We do so by characterising the instances G, α, β where the transformation is possible; the proof of this characterisation is via an algorithm that either finds a sequence of recolourings between α and β, or exhibits a structure which proves that no such sequence exists. In the case that a sequence of recolourings does exist, the algorithm uses $O(|V(G)|^2)$ recolouring steps and in many cases returns a shortest sequence of recolourings. We also exhibit a class of instances G, α, β that require $\Omega(|V(G)|^2)$ recolouring steps.

1 Introduction

In this paper graphs are finite and do not contain loops or multiple edges unless stated otherwise. We refer the reader to [5] for standard terminology and notation not defined here. A *k-colouring* of a graph $G = (V(G), E(G))$ is a function $\alpha : V(G) \to \{1, 2, \dots k\}$ such that $\alpha(u) \neq \alpha(v)$ for any edge uv. Throughout this paper we will assume that k is large enough to guarantee the existence of k-colourings (i.e., k is at least the *chromatic number of G*).

For a positive integer k and a graph G, we define the *k-colour graph of G*, denoted $\mathcal{C}_k(G)$, as the graph that has the k-colourings of G as its node set, with two k-colourings α and β

* Research partially supported by Nuffield grant no. NAL/32772.

joined by an edge in $\mathcal{C}_k(G)$ if they differ in colour on just one vertex of G. In this case, we shall also say that we can *recolour* G from α to β (and if v is the unique vertex on which α and β differ, then we also say that we can *recolour* v). Note that a path in $\mathcal{C}_k(G)$ can be described by either a sequence of colourings (the vertices of the path) or a sequence of recolourings. In addition, other graph-theoretical notions such as distance and adjacency can now be used for colourings.

The connectedness of the k-colour graph is an issue of interest when trying to obtain efficient algorithms for almost uniform sampling of k-colourings. In particular, $\mathcal{C}_k(G)$ needs to be connected for the single-site Glauber dynamics of G (a Markov chain defined on the k-colour graph of G) to be rapidly mixing. For details, see, for example, [6, 7], and references therein. In this setting, research on the connectedness of colour graphs has concentrated on cases where k is at least the maximum degree, or on cases where G is a highly symmetric graph such as an integer grid.

Properties of the colour graph, and questions regarding the existence of a path between two colourings, also find application in the study of radio channel reassignment. Given that a channel assignment problem can often be modelled as a graph colouring problem, the task of reassigning channels in a network, while avoiding interference and ensuring no connections are lost, can initially be thought of as a graph recolouring problem. See [1] for a discussion of these ideas in the context of cellular phone networks.

In recent work, the present authors have sought to develop a more general theory of the connectedness of colour graphs. In [3], a number of initial observations on properties of colour graphs are made and it is shown that if G has chromatic number $k \in \{2, 3\}$, then $\mathcal{C}_k(G)$ is not connected, but that for $k \geq 4$, there are k-chromatic graphs for which $\mathcal{C}_k(G)$ is not connected and k-chromatic graphs for which $\mathcal{C}_k(G)$ is connected. In [4], a characterisation of bipartite graphs whose 3-colour graph is connected is given and the problem of recognising these graphs is shown to be coNP-complete; while a polynomial algorithm is given for the restriction of the problem to planar graphs. In this paper, we consider the related problem of deciding whether two 3-colourings of a graph G belong to the same component of $\mathcal{C}_3(G)$. Formally, we have the following decision problem.

3-Colour Path
Instance : A connected graph G together with two 3-colourings of G, α and β.
Question : Is there a path between α and β in $\mathcal{C}_3(G)$?

We assume our 3-Colour Path instance graphs G to be connected as it is clear that the problem can be solved component-wise for disconnected graphs : there is a path between 3-colourings α and β of a disconnected graph G, if and only if for every connected component H of G there is a path between the colourings induced by α and β on H.

Our main result is the following.

Theorem 1
The decision problem 3-Colour Path *is in the complexity class* P.

We will prove Theorem 1 by describing a polynomial time algorithm that decides 3-Colour Path. The algorithm stems from the proof of a characterisation of instances G, α, β where α and β belong to the same component of $\mathcal{C}_3(G)$. We will describe and prove this characterisation in Section 3. First, in Section 2, we shall examine what can forbid the existence of a path in $\mathcal{C}_3(G)$ between 3-colourings α and β of a graph G. The proof of the characterisation of connected pairs of 3-colourings is via an algorithm that, given G, α, β, either finds a sequence of recolourings between α and β, or exhibits a structure (described in Section 2) which proves that no such sequence exists. Thus this algorithm also decides 3-Colour Path.

We will see that in the case that α and β belong to the same component of $\mathcal{C}_3(G)$, our algorithm will exhibit a path of length $O(|V(G)|^2)$ between them. This proves the following.

Theorem 2
Let G be a 3-colourable graph with n vertices. Then the diameter of any component of $\mathcal{C}_3(G)$ is $O(n^2)$.

In Section 4 we turn our attention to what else can be said about the distance between a given pair of 3-colourings. We will prove that for many instances our algorithm returns a shortest path between the given 3-colourings. We will also construct a class of instances G, α, β such that α and β are connected and at distance $\Omega(|V(G)^2|)$ in $\mathcal{C}_3(G)$.

The computational complexity of the more general problem k-Colour Path, defined analogously to 3-Colour Path but for k-colourings instead of 3-colourings, is very different. In [2] it is proved that for fixed $k \geq 4$ this problem is PSPACE-complete and that in this case, and in contrast with Theorem 2, the distance between two k-colourings can be superpolynomial in the size of the graph.

Finally, the decision problem 2-Colour Path is more or less trivial: two 2-colourings α, β of a bipartite graph G are connected, if and only if α and β are the same on each non-trivial component of G. And recolouring α to β can only involve changing the colours on isolated vertices, and hence requires at most $O(|V(G)|)$ steps.

2 Obstructions to paths between 3-colourings

In this section we examine what can stop us from being able to find a sequence of recolourings between a pair of 3-colourings α, β of a graph G. Informally, we call a structure in G, α, β forbidding the existence of a path between α and β in $\mathcal{C}_3(G)$ an *obstruction*.

For the remainder of this section we assume we are dealing with some fixed graph G.

The first and most obvious obstruction is given by what we call *fixed vertices*. For a 3-colouring α, we define a vertex v as *fixed* if there is no sequence of recolourings from α which will allow us to recolour v. In other words, a vertex v is fixed if for every colouring β in the same component of $\mathcal{C}_3(G)$ as α we have $\beta(v) = \alpha(v)$. For example, if a cycle with 0 mod 3 vertices is coloured 1-2-3-1-2-3-$\cdots$-1-2-3, then every vertex on the cycle is fixed (as none can be the first to be recoloured). We call such a cycle a *fixed cycle* (as a subgraph of G, and with respect to the 3-colouring α). Similarly, a path coloured $\cdots$3-1-2-3-1-2-3-1-$\cdots$, both whose endvertices lie on fixed cycles, cannot be recoloured and is called a *fixed path.*

Given a graph G with 3-colouring α, we denote the set of fixed vertices of G by F_α. In the next section, we shall prove the following.

Proposition 3
Let α be a 3-colouring of G. Then every $v \in F_\alpha$ belongs to a fixed cycle or to a fixed path.

The next lemma, which illustrates how fixed vertices may act as an obstruction, follows immediately from the definitions.

Lemma 4
Let α and β be two 3-colourings of G. If α and β belong to the same component of $\mathcal{C}_3(G)$, then we must have $F_\alpha = F_\beta$, and $\alpha(v) = \beta(v)$ for each $v \in F_\alpha$.

We proceed to describe two further obstructions that will forbid the existence of a path between a given pair of 3-colourings. For this, we need a few more definitions.

To *orient* a cycle or a path means to orient each edge of the cycle or path to obtain a directed cycle or a directed path. If C is a cycle, then $\overrightarrow{C}$ denotes C with one of the two possible orientations. Similarly, $\overrightarrow{P}$ denotes one of the two possible orientations of a path P.

For a 3-colouring α of G, the *weight* of an edge $e = uv$ oriented from u to v is

$$w(\overrightarrow{uv}, \alpha) \;=\; \begin{cases} +1, & \text{if } \alpha(u)\alpha(v) \in \{12,\, 23,\, 31\}; \\ -1, & \text{if } \alpha(u)\alpha(v) \in \{21,\, 32,\, 13\}. \end{cases}$$

The *weight* $W(\overrightarrow{C}, \alpha)$ of an oriented cycle is the sum of the weights of its oriented edges; the same holds for the weight $W(\overrightarrow{P}, \alpha)$ of an oriented path. The following lemma and its proof are from [3, 4].

Lemma 5
Let α and β be 3-colourings of G, and let C be a cycle in G. If α and β are in the same component of $\mathcal{C}_3(G)$, then we must have $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \beta)$.

Proof Let α and α' be 3-colourings of G that are adjacent in $\mathcal{C}_3(G)$, and suppose the two 3-colourings differ on vertex v. If v is not on C, then we certainly have $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \alpha')$.

If v is a vertex of C, then all its neighbours must have the same colour in α, otherwise we would not be able to recolour v. If we denote the in-neighbour of v on $\overrightarrow{C}$ by v_i and its out-neighbour by v_o, then this means that $w(\overrightarrow{v_iv}, \alpha)$ and $w(\overrightarrow{vv_o}, \alpha)$ have opposite sign, hence $w(\overrightarrow{v_iv}, \alpha) + w(\overrightarrow{vv_o}, \alpha) = 0$. Recolouring vertex v will change the signs of the weights of the oriented edges $\overrightarrow{v_iv}$ and $\overrightarrow{vv_o}$, but they will remain opposite. Therefore $w(\overrightarrow{v_iv}, \alpha') + w(\overrightarrow{vv_o}, \alpha') = 0$, and it follows that $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \alpha')$.

From the above we immediately obtain that the weight of an oriented cycle is constant on all 3-colourings in the same component of $\mathcal{C}_3(G)$. ∎

The following lemma can be proved in the same way.

Lemma 6
Let α and β be 3-colourings of G with $F_\alpha = F_\beta \neq \varnothing$ and $\alpha(v) = \beta(v)$ for all $v \in F_\alpha$. Suppose G contains a path P with endvertices u and v, where $u, v \in F_\alpha$. If α and β are in the same component of $\mathcal{C}_3(G)$, we must have $W(\overrightarrow{P}, \alpha) = W(\overrightarrow{P}, \beta)$.

Lemmas 4, 5 and 6 give necessary conditions for two 3-colourings α and β to belong to the same component of $\mathcal{C}_3(G)$. From Lemmas 4 and 5 we have, respectively:

(C1) $F_\alpha = F_\beta$, and $\alpha(v) = \beta(v)$ for each $v \in F_\alpha$; and

(C2) for every cycle C in G we have $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \beta)$.

If for two 3-colourings α and β of G we take condition (C1) to be satisfied (so they have the same fixed vertices, coloured alike), Lemma 6 gives a third necessary condition for α and β to belong to the same component of $\mathcal{C}_3(G)$:

(C3) for every path P between fixed vertices we have $W(\overrightarrow{P}, \alpha) = W(\overrightarrow{P}, \beta)$.

Bearing in mind that we are only considering condition (C3) if condition (C1) is already satisfied, let us observe that neither conditions (C1) and (C2) taken together, nor conditions (C1) and (C3) taken together, are sufficient to guarantee the existence of a path between 3-colourings α and β.

To see that conditions (C1) and (C2) are not sufficient, consider the graph and two 3-colourings shown in Figure 1. It is easy to check that (C1) and (C2) are satisfied (note that only vertices on the 3-cycles are fixed), but the two colourings are not connected: fix an orientation of the path between the two 3-cycles, and observe that the weight of this oriented path is $+3$ in one colouring and -3 in the other.

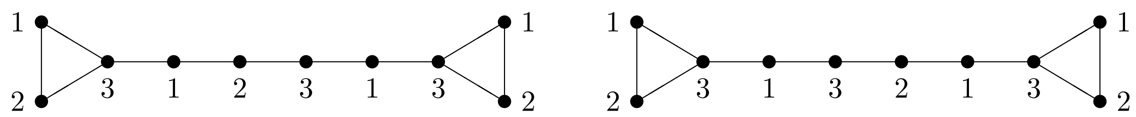

Figure 1: Two 3-colourings of a graph G not connected in $\mathcal{C}_3(G)$.

To see that conditions (C1) and (C3) are not sufficient, consider two 3-colourings α and β of a 5-cycle that differ only in that the colours 1 and 2 are swapped: (C1) and (C3) are satisfied (since $F_\alpha = F_\beta = \varnothing$), but there is no path between the two colourings as the 5-cycle has different weights in the two colourings.

We prove in the next section that if all three conditions are satisfied by a pair of colourings α and β of a graph G, then these colourings are in the same component of $\mathcal{C}_3(G)$.

3 A characterisation of connected pairs of 3-colourings

In this section we prove the following characterisation of connected pairs of 3-colourings. Its proof will yield a polynomial time algorithm for 3-Colour Path, proving Theorem 1. We will also prove Theorem 2 in the process.

Theorem 7
Two 3-colourings α and β of a graph G belong to the same component of $\mathcal{C}_3(G)$, if and only if

(C1) we have $F_\alpha = F_\beta$, and for each $v \in F_\alpha$, $\alpha(v) = \beta(v)$;

(C2) for every cycle C in G we have $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \beta)$; and

(C3) for every path P between fixed vertices we have $W(\overrightarrow{P}, \alpha) = W(\overrightarrow{P}, \beta)$.

The necessity of the three conditions has already been established. We prove that they are sufficient by outlining an algorithm whose input is a graph G and two 3-colourings α and β of G, and whose output is either a path in $\mathcal{C}_3(G)$ from α to β, or an obstruction that shows that (C1), (C2) or (C3) is not satisfied, and hence no such path exists.

The first step of the algorithm is to find F_α and F_β. We claim that the following procedure finds the fixed vertices of a graph G with 3-colouring α.

- Let S_1, S_2, S_3 initially be the three colour classes induced by α.
- For $i \in \{1, 2, 3\}$, and for each vertex $v \in S_i$: if v has no neighbours in one or both of the other two sets, then let $S_i = S_i \setminus \{v\}$.
- Repeat the previous step until no further changes are possible. Return $S = S_1 \cup S_2 \cup S_3$.

Claim 8
The above procedure returns $S = F_\alpha$.

Before proving the claim, let us give some more definitions. Fix a vertex v of G and set $L_0^+ = L_0^- = \{v\}$. For $i = 1, 2, \ldots$, let a vertex u belong to L_i^+ if u has a neighbour $w \in L_{i-1}^+$ and $\alpha(u) \equiv \alpha(w) + 1 \pmod 3$. (So, for example, if v is coloured 3, then L_1^+ contains all its neighbours coloured 1, L_2^+ contains all vertices coloured 2 that have a neighbour in L_1^+, and so on.) For $j = 1, 2, \ldots$, let a vertex u belong to L_j^- if u has a neighbour $w \in L_{j-1}^-$ and $\alpha(u) \equiv \alpha(w) - 1 \pmod 3$. We call these sets the *levels* of v, and the sets are called *positive* or *negative* according to their superscript.

Observe that v lies on a fixed cycle, if and only if there is a vertex $u \in L_i^+ \cap L_j^-$, for some $i, j > 0$. Similarly, v lies on a fixed path with end vertices u and w (each on a fixed cycle), if and only if $u \in L_{i'}^+ \cap L_i^+$ for some $i' > i > 0$ and $w \in L_{j'}^- \cap L_j^-$ for some $j' > j > 0$.

Proof of Claim 8 (and Proposition 3) Suppose the procedure described above is run on G, α, and has terminated. Note that a vertex that lies on a fixed cycle or path is certainly in S. We shall show that for each vertex $v \in V(G)$:

- either v lies on a fixed cycle or path (so is both fixed and in S),
- or v is neither fixed nor in S.

This will prove that $S = F_\alpha$, and also Proposition 3.

Fix a vertex v of G and consider the levels of v. We have observed that if there is a vertex that is in L_i^+, for some $i > 0$, and also in L_j^-, for some $j > 0$, then v lies on a fixed cycle.

Also, if there is a vertex that belongs to L_i^+ and $L_{i'}^+$, for some $i' > i > 0$, and another vertex that belongs to L_j^- and $L_{j'}^-$, for some $j' > j > 0$, then v lies on a fixed path.

If neither of these two properties hold, then either the positive or negative levels (or both) are disjoint and thus only finitely many of them are nonempty. We show that this means we can recolour v, and hence v is not fixed. So assume that $L_t^+ = \varnothing$ or $L_t^- = \varnothing$ for some $t > 0$. Without loss of generality, let us assume $L_t^+ = \varnothing$. Thus each vertex $u \in L_{t-1}^+$ can be recoloured with $\alpha(u) + 1 \pmod 3$. Then each vertex $w \in L_{t-2}^+$ can be recoloured with $\alpha(w) + 1 \pmod 3$, and so on, until v is recoloured. The fact that v can be recoloured implies it is not in S: every vertex in S has a pair of differently coloured neighbours so no vertex in S can be the first to be recoloured. □

Claim 8 allows us to find F_α and F_β. If $F_\alpha \neq F_\beta$, or if there is a vertex $v \in F_\alpha$ such that $\alpha(v) \neq \beta(v)$, then there is no path from α to β. The algorithm outputs F_α, F_β and, if necessary, v.

Henceforth we assume that condition (C1) is satisfied, so $F_\alpha = F_\beta$ and for all $v \in F_\alpha$, $\alpha(v) = \beta(v)$.

If $F_\alpha \neq \varnothing$, we construct, from G, a new graph G^f by identifying, for $i = 1, 2, 3$, all vertices in S_i and denoting the newly created vertex by f_i. In other words:

$$\begin{aligned}
V(G^f) &= (V(G) \setminus F_\alpha) \cup \{f_1, f_2, f_3\}, \quad \text{and} \\
E(G^f) &= \{\, uv \in E(G) \mid u, v \in V(G) \setminus F_\alpha \,\} \\
&\qquad \cup \bigcup_{i=1,2,3} \{\, uv \in E(G) \mid u \in V(G) \setminus F_\alpha,\ v \in S_i \,\} \cup \{f_1f_2, f_1f_3, f_2f_3\}.
\end{aligned}$$

If G has no fixed vertices with respect to α, then we set $G^f = G$.

It is convenient to assume that all edges are retained so that G and G^f have the same edge set. Since S_1, S_2, S_3 are independent sets (they are subsets of the colour classes of the colouring α), this means G^f is a graph with possibly multiple edges, but no loops. Moreover, it is easy to observe that if $F_\alpha \neq \varnothing$,

- f_1, f_2 and f_3 are the only fixed vertices of G^f, and
- f_1, f_2 and f_3 induce a (fixed) 3-cycle in G^f.

Let α^f and β^f be the colourings induced on G^f by α and β. Note that if α and β belong to the same component of $\mathcal{C}_3(G)$, this component is isomorphic to the component of $\mathcal{C}_3(G^f)$ that contains α^f and β^f. Hence we have the following.

Claim 9
There is a path from α to β in $\mathcal{C}_3(G)$ if and only if there is a path from α^f to β^f in $\mathcal{C}_3(G^f)$.

To prove Theorem 7, we shall prove the following claim.

Claim 10
Two 3-colourings α^f and β^f of a graph G^f belong to the same component of $\mathcal{C}_3(G^f)$, if and only if
(C2′) for every cycle C in G^f we have $W(\overrightarrow{C}, \alpha^f) = W(\overrightarrow{C}, \beta^f)$.

Let us first establish that the claim implies the theorem, recalling that we are assuming condition (C1). Let $\overrightarrow{C}$ be an oriented cycle in G. In G^f, the oriented edges of $\overrightarrow{C}$ form a set of edge-disjoint oriented cycles. (Here we use the convention that all edges from G are retained in G^f.) Since these cycles contain the same edges as $\overrightarrow{C}$, similarly oriented, it is easy to see that the sum of the weights of these cycles is equal to $W(\overrightarrow{C}, \alpha)$. Thus if G^f, α^f, β^f satisfy (C2′), then G, α, β satisfy (C2).

Now, let $\overrightarrow{P}$ be an oriented path between fixed vertices in G. If the endvertices of P have the same colour, then the oriented edges of $\overrightarrow{P}$ again form a set of edge-disjoint oriented cycles in G^f, and (C2′) implies that $W(\overrightarrow{P}, \alpha) = W(\overrightarrow{P}, \beta)$. If the endvertices of P have a different colour, then we can suppose, without loss of generality, that the endvertices of P are coloured 1 and 2 and that $\overrightarrow{P}$ is oriented from the endvertex coloured 1 towards the endvertex coloured 2. That means that the union of the oriented edges of $\overrightarrow{P}$ and the edge $\overrightarrow{f_2 f_1}$ forms a set of oriented cycles in G^f. Since we have $w(\overrightarrow{f_2 f_1}, \alpha^f) = w(\overrightarrow{f_2 f_1}, \beta^f)$, (C2′) again implies that $W(\overrightarrow{P}, \alpha) = W(\overrightarrow{P}, \beta)$. We have shown that if G^f, α^f, β^f satisfy (C2′), then G, α, β satisfy (C3).

Conversely, if there is a cycle C in G^f such that $W(\overrightarrow{C}, \alpha^f) \neq W(\overrightarrow{C}, \beta^f)$, then this same cycle can be found in G or, if C intersects $\{f_1, f_2, f_3\}$, then there is a path between fixed vertices in G that has different weights under α and β. This shows that if G^f, α^f, β^f do not satisfy (C2′), then one of (C2) or (C3) fails for G, α, β.

Proof of Claim 10 To prove the claim we describe an algorithm that either finds a path from α^f to β^f in $\mathcal{C}_3(G^f)$, or finds a cycle C in G^f such that $W(\overrightarrow{C}, \alpha^f) \neq W(\overrightarrow{C}, \beta^f)$. The algorithm attempts to find a sequence of recolourings that transforms α^f into β^f. It maintains a set $F \subseteq V(G^f)$ such that the subgraph induced by F is connected and for each $v \in F$, the current colouring of v is $\beta^f(v)$. Initially we let F be the set of fixed vertices of G^f (so $F = \varnothing$ or $F = \{f_1, f_2, f_3\}$) and try to increase the size of F one vertex at a time.

We show how to extend F if $F \neq V(G^f)$. If $F \neq \varnothing$, then choose a $v \notin F$ such that v is adjacent to a vertex $u \in F$. (This is possible by the assumption that G (and thus G^f) is connected.) If $F = \varnothing$, then we choose an arbitrary vertex v, and u does not exist. Suppose the current colouring is α'. If $\alpha'(v) = \beta^f(v)$, we can extend F to include v immediately. Otherwise, let us assume that $\alpha'(v) = 2$ and $\beta^f(v) = 3$. Note that this means that $\alpha'(u) = 1$ (if u exists), since $\alpha'(u) = \beta^f(u)$ and u is adjacent to v.

Now we attempt to find the positive levels of v. This is easily done algorithmically: $L_1^+(v)$ contains those neighbours of v coloured 3; $L_2^+(v)$ contains neighbours of vertices in L_1 coloured 1, and so on. We stop if either

(L1) we reach a level L_i^+ that is empty, or

(L2) we find a level that contains a vertex $w \in F$.

Note that one of (L1) or (L2) must occur. This is because any vertex not in F belongs to at most one level (if a vertex belongs to two levels it is fixed, and all fixed vertices are in F). Hence we eventually reach either a level that contains a vertex $w \in F$, or an empty level. If F is empty, then, of course, (L1) must occur.

If (L1) occurs, then we can recolour each vertex z in L_j^+, $j = i-1, i-2, \ldots, 0$, with $\alpha'(z)+1 \pmod 3$, starting with the highest level and working down. Thus, ultimately, v is recoloured 3 and we can now add v to F. If there are still vertices not in F, we repeat the procedure.

Suppose (L2) occurs. Then there is a path P from u to w coloured 1-2-3-1-2-3-$\cdots$-$\alpha'(w)$. Moreover, no internal vertex of P is in F. As u and w are in F, and F induces a connected subgraph, we can extend P to a cycle C using a path $Q = w \cdots u$ in F. We claim that $W(\overrightarrow{C}, \alpha') \neq W(\overrightarrow{C}, \beta^f)$, and hence the cycle C is an obstruction that shows that α' and β^f do not belong to the same component of $\mathcal{C}_3(G)$. Because α^f and α' belong to the same component of $\mathcal{C}_3(G)$, this cycle is also an obstruction showing that α^f and β^f do not belong to the same component of $\mathcal{C}_3(G)$.

To see that $W(\overrightarrow{C}, \alpha') \neq W(\overrightarrow{C}, \beta^f)$, choose the orientation $\overrightarrow{C}$ so that the edge uv is oriented from u to v. The weight of $\overrightarrow{C}$ is the sum of the weights of $\overrightarrow{P}$ and $\overrightarrow{Q}$ (taking $\overrightarrow{P}$ and $\overrightarrow{Q}$ to have the same orientation as $\overrightarrow{C}$). Let $W(\overrightarrow{Q}, \alpha') = k$. As vertices in F are coloured alike in α' and β^f, $W(\overrightarrow{Q}, \beta^f) = k$. Let p be the number of edges in P. Then $W(\overrightarrow{P}, \alpha') = p$, since each edge has weight $+1$. But $W(\overrightarrow{Q}, \beta^f) < p$, since $w(\overrightarrow{uv}, \beta^f) = -1$. Thus we find $W(\overrightarrow{C}, \beta^f) < k + p = W(\overrightarrow{C}, \alpha')$.

All the above was done under the assumption that $\alpha'(v) = 2$ and $\beta^f(v) = 3$. In the cases $\alpha'(v) = 3$, $\beta^f(v) = 1$ and $\alpha'(v) = 1$, $\beta^f(v) = 2$ we do exactly the same, again using the positive levels $L_i^+(v)$. In the other three cases, we follow the same steps, but now using the negative levels $L_i^-(v)$ of v. This completes the proof of the claim. □

This completes the proof of Theorem 7. ■

Note that if α and β are in the same component of $\mathcal{C}_3(G)$ and G has n vertices, the algorithm in the proof of Claim 10 will use at most $\frac{1}{2}n(n+1)$ recolouring steps: each time a vertex is added to F, we may have to recolour all vertices not in F at most once. This proves Theorem 2. ■

Note also that the procedure which finds the fixed vertices of a given 3-colouring, the construction of G^f from G, and the algorithm in the proof of Claim 10 can clearly be performed in polynomial time. This proves Theorem 1. ■

Using Theorem 7, it is in fact possible to give an alternative proof of Theorem 1. We describe a modification of the algorithm that proves Theorem 7 which, given a graph G together with two 3-colourings α and β as input, decides whether or not α and β belong to the same component of $\mathcal{C}_3(G)$ by simply checking conditions (C1), (C2) and (C3).

As before, we first check whether condition (C1) is satisfied. We proceed by assuming it (else the algorithm terminates), and then transform the instance G, α, β into the instance G^f, α^f, β^f. We have already observed that these operations can be performed in polynomial time.

Having seen that condition (C2′) is equivalent to conditions (C2) and (C3), we now claim that condition (C2′) can be verified in polynomial time. (Note that this is not immediately obvious since the graph G^f may contain an exponential number of cycles.) In order to prove this claim, we need to recall some definitions.

Let H be a connected graph with n vertices and m edges. It is well-known that (the edge sets of) the cycles of H form a vector space over the field $\mathbb{F}_2 = \{0, 1\}$, where addition is symmetric difference. This vector space is known as the *cycle space* of H. Given any spanning tree T of H, adding any of the $m-n+1$ edges $e \in E(H) \setminus E(T)$ to T yields a unique cycle C_e of H. These $m-n+1$ cycles are called the *fundamental cycles* of T, and they form a basis of the cycle space of H known as a *cycle basis*. In fact, it is easy to prove that for every cycle C,

$$C = \sum_{e \in E(C) \setminus E(T)} C_e,$$

where addition is as in the vector space $(\mathbb{F}_2)^m$. We refer the reader to [5, Section 1.9] for further details.

Lemma 11
Let H be a connected graph with n vertices and m edges. Let α be a 3-colouring of H, T a spanning tree of H, and $\{ C_e \mid e \in E(H) \setminus E(T) \}$ the set of fundamental cycles of T. Then for any cycle C in H, $W(\overrightarrow{C}, \alpha)$ is determined by the values of $W(\overrightarrow{C_e}, \alpha)$, for all $e \in E(H) \setminus E(T)$.

Proof Let C be any cycle in H, and write $C = \sum_{e \in E(C) \setminus E(T)} C_e$, with addition as in the vector space $(\mathbb{F}_2)^m$. Choose an orientation $\overrightarrow{C}$ for C. For each $e \in E(C) \setminus E(T)$, orient the fundamental cycle C_e so that e has the same orientation in $\overrightarrow{C}$ and in $\overrightarrow{C_e}$. We claim that

$$W(\overrightarrow{C}, \alpha) = \sum_{e \in E(C) \setminus E(T)} W(\overrightarrow{C_e}, \alpha), \tag{1}$$

where now addition is the normal addition of integers. We prove (1) by counting edge-weight contributions to both sides of the equation.

Let $e = uv$ be an edge of C, with orientation $\overrightarrow{uv}$ on $\overrightarrow{C}$. Clearly, $w(\overrightarrow{uv}, \alpha)$ is counted exactly once on the left-hand side (LHS) of (1). To count the contributions that e makes to the right-hand side (RHS) of (1), we distinguish two cases, according to whether or not e is an edge of T. If $e \notin E(T)$, then the definition of C_e and the choice of the orientation $\overrightarrow{C_e}$ immediately gives that e contributes exactly the weight $w(\overrightarrow{uv}, \alpha)$ to the RHS. If $e = uv \in E(T)$, we claim that it appears oriented as $\overrightarrow{uv}$ exactly one more time than it appears oriented as $\overleftarrow{uv}$ in the cycle expansion of $\overrightarrow{C}$. Note that uv is a cut-edge of T and, as such, its removal splits T into two subtrees T_u and T_v, with $u \in V(T_u)$ and $v \in V(T_v)$. We also have $V(T_u) \cup V(T_v) = V(H)$. Let $f \in E(C) \setminus E(T)$ with $uv \in E(C_f)$. Then, in fact, we can take $f = xy$ with $x \in V(T_u)$ and $y \in V(T_v)$. If f has the orientation $\overrightarrow{xy}$ in $\overrightarrow{C}$, then it has the same orientation in $\overrightarrow{C_f}$, and hence the edge uv has the orientation $\overleftarrow{uv}$ in C_f. The reverse is the case if f has the orientation $\overleftarrow{xy}$ in $\overrightarrow{C}$. Going along the oriented edges of the cycle $\overrightarrow{C}$, we have the same number of edges $\overrightarrow{xy}$ with $x \in V(T_u)$ and $y \in V(T_v)$, as we have edges between $V(T_u)$ and $V(T_v)$ going in the other direction. But since uv is one of the edges of the first count, we get exactly one more edge $xy \neq uv$ of $\overrightarrow{C}$ with $x \in V(T_u)$ and $y \in V(T_v)$ oriented as $\overleftarrow{xy}$ than oriented the other way round. That means that in the sum on the RHS of (1) we have exactly one more contribution of the form $w(\overrightarrow{uv}, \alpha)$ than of the form $w(\overleftarrow{uv}, \alpha)$.

Now suppose that $e = uv$ is not an edge of C. Clearly this edge makes no contribution to the LHS of the equation. Again, to count the contributions of this edge to the RHS of the expression, we distinguish the cases where e is an edge of T and where it is not. If $e = uv \in E(T)$, we can argue as in the preceding paragraph, to see that this time, in the RHS we have exactly the same times a contribution of the form $w(\overrightarrow{uv}, \alpha)$ as of the form $w(\overleftarrow{uv}, \alpha)$. Hence the net contribution to the RHS is zero. Lastly, if $e \notin E(T)$, it makes no contribution either, since the fundamental cycle C_e to which it corresponds does not appear in the cycle expansion of C.

This completes the proof of the lemma. $\square$

Lemma 11 gives an obvious algorithm to check if the (reduced) instance G^f, α^f, β^f satisfies condition (C2′), running in polynomial time.

4 Shortest paths between 3-colourings

Once again, throughout this section we assume that G is some fixed connected graph. We use the notation and terminology from the previous section.

We have seen that if α and β are 3-colourings of G that are in the same component of $\mathcal{C}_3(G)$, then they are at distance $O(|V(G)|^2)$. In this section we show that this bound on the distance between 3-colourings is of the right order. More precisely, we prove that there exists a class of instances G', α, β such that α and β are connected and at distance $\Omega(|V(G')|^2)$ in $\mathcal{C}_3(G')$.

Before doing so, we prove that in the case that α and β are connected and $F_\alpha \neq \varnothing$ (so $F_\beta = F_\alpha$, and for all $v \in F_\alpha$ we have $\alpha(v) = \beta(v)$), the algorithm described in the previous section finds a shortest path from α to β in $\mathcal{C}_3(G)$.

Theorem 12
Let α and β be two 3-colourings of a connected graph G that are in the same component of $\mathcal{C}_3(G)$, and suppose that $F_\alpha \neq \varnothing$. Then the algorithm described in Section 3 finds a shortest path between α and β.

Proof Our algorithm in fact finds a path from α^f to β^f in G^f, but, as we observed earlier, the relevant components of the two colour graphs are isomorphic. For a 3-colouring γ of G^f, denote by $\mathcal{C}_\gamma$ the component of $\mathcal{C}_3(G^f)$ containing γ. Note that, by assumption of connectedness, $\mathcal{C}_{\alpha^f} = \mathcal{C}_{\beta^f}$.

Recall that G^f has exactly three fixed vertices f_1, f_2, f_3 for the colourings α^f and β^f.

Let γ be any 3-colouring in $\mathcal{C}_{\beta^f}$. For any vertex v of G^f, let $\overrightarrow{P}$ be an oriented path from f_1 to v. Then the *height of v in γ* is defined as

$$h(v, \gamma) \;=\; |W(\overrightarrow{P}, \gamma) - W(\overrightarrow{P}, \beta^f)|.$$

We need to prove that this definition is independent of the choice of P. If there are two oriented paths $\overrightarrow{P_1}$ and $\overrightarrow{P_2}$ from f_1 to v, then, noting that their union is a set of oriented cycles and applying Lemma 5, we have $W(\overrightarrow{P_1},\gamma)-W(\overrightarrow{P_2},\gamma) = W(\overrightarrow{P_1},\beta^f)-W(\overrightarrow{P_2},\beta^f)$. Rearranging leads to $|W(\overrightarrow{P_1},\gamma) - W(\overrightarrow{P_1},\beta^f)| = |W(\overrightarrow{P_2},\gamma) - W(\overrightarrow{P_2},\beta^f)|$.

Now let γ and δ be adjacent 3-colourings in $\mathcal{C}_{\beta^f}$ and let w be the unique vertex on which they differ. Note that this means that all neighbours of w are coloured the same as one another, and all these neighbours are coloured the same in both γ and δ. Let $\overrightarrow{P}$ be an oriented path from f_1 to some vertex v and let us consider how the height of v changes as γ is recoloured to δ. If w is not on $\overrightarrow{P}$, then clearly $h(v,\gamma) = h(v,\delta)$. We know $w \neq f_1$, as f_1 is fixed. If w is an internal vertex of $\overrightarrow{P}$, then the sum of the weights of the two edges of $\overrightarrow{P}$ incident with w is zero for both γ and δ, so again $h(v,\gamma) = h(v,\delta)$. If $w = v$, then the sign of the weight of the edge of $\overrightarrow{P}$ incident with v changes as we recolour. So in this last case we have $|h(v,\gamma) - h(v,\delta)| = 2$.

Note that finding a path from α^f to β^f is equivalent to finding a sequence of recolourings that reduces the height of every vertex v from $h(v,\alpha^f)$ to zero. In the previous paragraph we saw that each time we recolour, only the height of the vertex being recoloured changes, and it either increases or decreases by 2. So if we can find a sequence of recolourings that always reduces the height of the vertex being recoloured, we will have found a shortest path. We show that this is indeed what the algorithm of Claim 10 does.

Recall that the algorithm starts with a set $F = \{f_1, f_2, f_3\}$, and then it repeatedly adds vertices v to F, where v has a neighbour $u \in F$. To add v to F, the vertices in either all its positive levels or all its negative levels are recoloured before v itself is recoloured. Assume that we are in the case that to recolour v all positive levels need to be recoloured; the other case is proved in the same way. Let y be a vertex that is about to be recoloured at some stage in this process (this can be v itself, or any of the vertices in the positive levels of v). We must show that its height will be reduced. Let γ and δ be the colourings before and after y is recoloured. Let $\overrightarrow{Q}$ be an oriented path from u to y that contains one vertex from each nonnegative level of v. So if there are k edges in $\overrightarrow{Q}$, then $W(\overrightarrow{Q},\gamma) = k$. Thus $W(\overrightarrow{Q},\delta) = k-2$ since the edge of $\overrightarrow{Q}$ incident with y has its weight changed from 1 to -1 when y is recoloured. Let $\overrightarrow{R}$ be an oriented path from f_1 to u containing only vertices in F, and let $\overrightarrow{P}$ be the union of $\overrightarrow{R}$ and $\overrightarrow{Q}$.

Since the colourings β^f, γ, δ agree on F, we have $W(\overrightarrow{R},\beta^f) = W(\overrightarrow{R},\gamma) = W(\overrightarrow{R},\delta)$. We also know that $w(\overrightarrow{uv},\beta^f) = -1$, and since $\overrightarrow{Q}$ has k edges, this means

$$W(\overrightarrow{Q},\beta^f) \;\le\; k-2 \;=\; W(\overrightarrow{Q},\delta) \;<\; k \;=\; W(\overrightarrow{Q},\gamma).$$

From this we can derive

$$\begin{aligned} h(y,\gamma) \;=\; |W(\overrightarrow{P},\gamma) - W(\overrightarrow{P},\beta^f)| \;&=\; |W(\overrightarrow{Q},\gamma) - W(\overrightarrow{Q},\beta^f)| \\ &=\; W(\overrightarrow{Q},\gamma) - W(\overrightarrow{Q},\beta^f) \;=\; k - W(\overrightarrow{Q},\beta^f) \end{aligned}$$

and, similarly,

$$h(y,\delta) \;=\; k-2-W(\overrightarrow{Q},\beta^f).$$

So indeed, every recolouring according to the lemma, reduces the height of the vertex being recoloured, completing the proof of the theorem. ■

Next let us observe that if there are no fixed vertices, the algorithm may find a much longer path. For example, consider two colourings of a path that differ only on an endvertex v and its neighbour: α = 1-2-3-1-2-3-1-$\cdots$-1-2-3 and β = 2-1-3-1-2-3-1-$\cdots$-1-2-3. The algorithm starts by setting $F = \varnothing$, and then chooses an arbitrary first vertex to start the recolouring. If that first vertex is v, then the algorithm will start by recolouring every vertex on the path. But clearly it is possible to get from α to β via only three recolourings. The reader should check that this shortest number of recolourings would be obtained if the first choice of the algorithm were any vertex other than v.

We believe that the algorithm from Section 3 will also be able to find a shortest path between two 3-colourings without fixed vertices.

Conjecture 13
Let α and β be two 3-colourings of a connected graph G that are in the same component of $\mathcal{C}_3(G)$, and suppose that $F_\alpha = F_\beta = \varnothing$. For $v \in V(G)$, let $T(v)$ be the number of recolourings required by the algorithm in Section 3 when the algorithm starts by adding v to $F = \varnothing$. Then the length of the shortest path between α and β is equal to $\min\limits_{v \in V(G)} T(v)$.

We now proceed to the construction of a class of instances G, α, β where, for each G, α and β are connected and at distance $\Omega(|V(G)|^2)$ in $\mathcal{C}_3(G)$. For $N \in \mathbb{N}$, define the graph G_N as the graph consisting of a 3-cycle with an attached path of length N. More precisely, let

$$\begin{aligned} V(G_N) &= \{f_1, f_2, f_3\} \cup \{v_1, v_2, \ldots, v_N\}, \quad \text{and} \\ E(G_N) &= \{f_1f_2, f_2f_3, f_1f_3\} \cup \{f_3v_1, v_1v_2, v_2v_3, \ldots, v_{N-1}v_N\} \end{aligned}$$

Let α_N be the 3-colouring of G_N given by $\alpha_N(f_i) = i$, for $i = 1, 2, 3$, and where the vertices $v_1, v_2, \ldots, v_N$ are coloured 1-2-3-1-2-3-$\cdots$. Similarly, let β_N be the 3-colouring of G_N given by $\beta_N(t_i) = i$, for $i = 1, 2, 3$, and where the vertices $v_1, v_2, \ldots, v_N$ are coloured 2-1-3-2-1-3-$\cdots$.

Theorem 14
Let $N \in \mathbb{N}$ and let G_N, α_N, β_N be as described above. Then the 3-colourings α_N and β_N of G_N are connected and at distance $\frac{1}{2}\,N\,(N+1) = \Omega(|V(G_N)|^2)$ in $\mathcal{C}_3(G_N)$.

Proof It is clear that G_N, α_N and β_N satisfy conditions (C1), (C2) and (C3). Therefore, by Theorem 7, α_N and β_N are connected in $\mathcal{C}_3(G_N)$.

As in the proof of Theorem 12, we consider heights of vertices. For any vertex v of G_N, let $\overrightarrow{P}$ be an oriented path from f_3 to v, noting that $f_3 \in F_{\alpha_N}$. Define the height of v in α_N as $h(v, \alpha_N) = |W(\overrightarrow{P}, \alpha_N) - W(\overrightarrow{P}, \beta_N)|$.

We have seen in the proof of Theorem 12 that finding a shortest path from α_N to β_N is equivalent to finding a sequence of recolourings that reduces the height of every vertex to zero, and that, with each recolouring, we reduce the height of the recoloured vertex by 2, while the height of all other vertices remains the same. This enables us to calculate the distance between α_N and β_N : we just need to calculate the height of all vertices in α_N.

First observe that $h(f_i, \alpha_N) = 0$, for $i = 1, 2, 3$. For $i = 1, \ldots, N$, let $\overrightarrow{P_i}$ be the oriented path from f_3 to v_i, and observe that $W(\overrightarrow{P_i}, \alpha_N) = i$ while $W(\overrightarrow{P_i}, \beta_N) = -i$. This means $h(v_i, \alpha_N) = |W(\overrightarrow{P_i}, \alpha_N) - W(\overrightarrow{P_i}, \beta_N)| = 2\,i$. We find that the distance between α_N and β_N is equal to $\frac{1}{2} \sum_{i=1}^{N} h(v_i, \alpha_N) = \sum_{i=1}^{N} i = \frac{1}{2}\, N\,(N+1)$. Since G_N has $N+3$ vertices, we obtain that this distance is indeed $\Omega(|V(G_N)|^2)$. ■

References

[1] J. Billingham, R. Leese, and H. Rajaniemi, *Frequency reassignment in cellular phone networks*. Smith Institute Study Group Report (2005). Available via `www.smithinst.ac.uk/Projects/ESGI53/ESGI53-Motorola/Report/`.

[2] P. Bonsma and L. Cereceda, *Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances*. CDAM Research Report LSE-CDAM-2007-12 (2007). Available via `www.cdam.lse.ac.uk/Reports/Abstracts/cdam-2007-12.html`.

Extended abstract in : *Mathematical Foundations of Computer Science 2007. Proceedings of 32nd International Symposium MFCS 2007*. Lect. Notes Comput. Sc. **4708** (2007), 738–749.

[3] L. Cereceda, J. van den Heuvel, and M. Johnson, *Connectedness of the graph of vertex-colourings*. CDAM Research Report LSE-CDAM-2005-11 (2005). Available via `www.cdam.lse.ac.uk/Reports/Abstracts/cdam-2005-11.html`. To appear in Discrete Math.

[4] L. Cereceda, J. van den Heuvel, and M. Johnson, *Mixing 3-colourings in bipartite graphs*. CDAM Research Report LSE-CDAM-2007-06 (2007). Available via `www.cdam.lse.ac.uk/Reports/Abstracts/cdam-2007-06.html`.

[5] R. Diestel, *Graph Theory*. Springer-Verlag, Berlin, 3rd edition, 2005.

[6] M. Jerrum, *A very simple algorithm for estimating the number of k-colourings of a low degree graph*. Random Structures Algorithms **7** (1995), 157–165.

[7] M. Jerrum, *Counting, Sampling and Integrating: Algorithms and Complexity*. Birkhäuser Verlag, Basel, 2003.