Mixing 3-Colourings in Bipartite Graphs

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, Durham University
Science Laboratories, South Road, Durham DH1 3LE, U.K.

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

Abstract

For a 3-colourable graph G, the 3-colour graph of G, denoted $\mathcal{C}_3(G)$, is the graph with node set the proper vertex 3-colourings of G, and two nodes adjacent whenever the corresponding colourings differ on precisely one vertex of G. We consider the following question: given G, how easily can one decide whether or not $\mathcal{C}_3(G)$ is connected? We show that the 3-colour graph of a 3-chromatic graph is never connected, and characterise the bipartite graphs for which $\mathcal{C}_3(G)$ is connected. We also show that the problem of deciding the connectedness of the 3-colour graph of a bipartite graph is coNP-complete, but that restricted to planar bipartite graphs, the question is answerable in polynomial time.

1 Introduction

Throughout this paper a graph $G = (V, E)$ is simple, loopless and finite. Most of our terminology and notation is standard and can be found in any textbook on graph theory such as, for example, [7]. We always regard a k-colouring of a graph G as proper; that is, as a function $\alpha : V \to \{1, 2, \ldots, k\}$ such that $\alpha(u) \neq \alpha(v)$ for any $uv \in E$. 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 joined by an edge in $\mathcal{C}_k(G)$ if they differ in colour on just one vertex of G.

Continuing a theme begun in an earlier paper [4], we investigate the connectedness of $\mathcal{C}_k(G)$ for a given G, this time concentrating on the case $k = 3$. The connectedness of the k-colour graph is an issue of interest when trying to obtain efficient algorithms for almost uniform

[†] Current address: Bayes Forecast, Gran Vía 39, 5ªplanta, 28013 Madrid, Spain. lcereceda@bayesforecast.com.

[‡] Research partially supported by Nuffield grant no. NAL/32772.

sampling of k-colourings of a given graph. 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 further details, see, for example, [9, 10] and references therein.

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 [2] for a discussion of these ideas in the context of cellular phone networks.

We say that G is *k-mixing* if $\mathcal{C}_k(G)$ is connected, and, having defined the colourings as nodes of $\mathcal{C}_k(G)$, the meaning of, for example, the path between two colourings should be clear. Observe that a graph G is k-mixing if and only if every connected component of G is k-mixing, so we will usually take our "argument graph" G to be connected. We assume throughout that $k \geq \chi(G) \geq 2$, where $\chi(G)$ is the chromatic number of G.

In this paper we concentrate on the case $k = 3$. In the next section it will be shown that if G has chromatic number 3, then G is not 3-mixing. We find more interesting behaviour, however, when G has chromatic number less than 3, that is, when G is bipartite. The main results in this paper deal with the following decision problem.

3-Mixing
Instance : A connected bipartite graph G.
Question : Is G 3-mixing?

After proving a characterisation theorem for 3-mixing bipartite graphs, we will prove the following two results :

Theorem 1.1. *The decision problem* 3-Mixing *is* coNP-*complete.*

Theorem 1.2. *Restricted to* planar *bipartite graphs, the decision problem* 3-Mixing *is in* P*.*

We believe that the case $k = 3$ is actually an exceptional case for the more general problem k-Mixing for fixed $k \geq 2$ (where the input graph is not necessarily bipartite). Let us explain the rationale behind that belief.

For two colours, 2-Mixing is trivially in P : If G is a connected bipartite graph with more than one vertex, then $\mathcal{C}_2(G)$ consists of two isolated vertices.

For $k \geq 4$, we do not know the computational complexity of k-Mixing. We do know quite a lot about the related decision problem k-Colour Path, though.

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

Again, the decision problem 2-Colour Path is trivially in P. It is proven by the authors in [5] that 3-Colour Path is in P as well. On the other hand, it is shown in [3] that for all fixed $k \geq 4$, k-Colour Path is PSPACE-complete. Moreover, the computational complexity of k-Colour Path does not change if we restrict the problem to bipartite and/or planar graphs. This strongly suggests that k-Mixing is PSPACE-hard, but we have been unable to prove this. (Note that as k-Colour Path is in PSPACE, it follows that k-Mixing is also in PSPACE.)

We finish this introductory section with some further terminology and notation and an outline of the paper. We use $\alpha, \beta, \ldots$ to denote specific colourings. We use the term *frozen* for a k-colouring of a graph G that forms an isolated node in the k-colour graph. Note that the existence of a frozen k-colouring of a graph immediately implies that the graph is not k-mixing.

If G has a k-colouring α, then we say that we can *recolour* G with β if $\alpha\beta$ is an edge of $\mathcal{C}_k(G)$. If v is the unique vertex on which α and β differ, then we also say that we can *recolour* v.

We denote the cycle on n vertices by C_n, and will often describe a colouring of C_n by just listing the colours as they appear on consecutive vertices.

The remainder of this paper is set out as follows. In the following section we introduce some of our tools and methods, revisiting the short proof (given in [4]) that 3-chromatic graphs are not 3-mixing. Section 3 gives two equivalent characterisations of 3-mixing bipartite graphs. In Section 4 we prove Theorem 1.1, while the final section contains the proof of Theorem 1.2.

2 Preliminaries

In [4] it was shown that if G has chromatic number k for $k = 2, 3$, then G is not k-mixing, but that, on the other hand, for $k \geq 4$, there are k-chromatic graphs that are k-mixing and k-chromatic graphs that are not k-mixing. For completeness, and since several of the ideas are used in later parts of this paper, we include the short proof of the fact that 3-chromatic graphs are not 3-mixing. Let us first give some definitions.

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

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

To *orient* a cycle means to orient each edge on the cycle so that a directed cycle is obtained. If C is a cycle, then by $\overrightarrow{C}$ we denote the cycle with one of the two possible orientations. The *weight* $W(\overrightarrow{C}, \alpha)$ of an oriented cycle $\overrightarrow{C}$ is the sum of the weights of its oriented edges.

Lemma 2.1. *Let α and β be 3-colourings of a graph G that contains a cycle C. Then if α and β are in the same component of $\mathcal{C}_3(G)$, 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 $w(\overrightarrow{v_i v}, \alpha)$ and $w(\overrightarrow{v v_o}, \alpha)$ have opposite sign, hence $w(\overrightarrow{v_i v}, \alpha) + w(\overrightarrow{v v_o}, \alpha) = 0$. Recolouring vertex v will change the signs of the weights of the oriented edges $\overrightarrow{v_i v}$ and $\overrightarrow{v v_o}$, but they will remain opposite. Therefore $w(\overrightarrow{v_i v}, \alpha') + w(\overrightarrow{v v_o}, \alpha') = 0$, and $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C}, \alpha')$.

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

Note that the converse of Lemma 2.1 is not true. For instance the 3-cycle has six 3-colourings. Of these, 1-2-3, 2-3-1 and 3-1-2 give the same weight of the oriented 3-cycle, but they are not connected (in fact, they are all frozen).

Lemma 2.2. *Let α be a 3-colouring of a graph G that contains a cycle C. If $W(\overrightarrow{C}, \alpha) \neq 0$, then $\mathcal{C}_3(G)$ is not connected.*

Proof : Let β be the 3-colouring of G obtained by setting for each vertex v of G :

$$\beta(v) \ = \ \begin{cases} 1, & \text{if } \alpha(v) = 2; \\ 2, & \text{if } \alpha(v) = 1; \\ 3, & \text{if } \alpha(v) = 3. \end{cases}$$

It is easy to check that for each edge e in C, $w(\vec{e}, \alpha) = -w(\vec{e}, \beta)$, which gives $W(\overrightarrow{C}, \alpha) = -W(\overrightarrow{C}, \beta)$. Since $W(\overrightarrow{C}, \alpha) \neq 0$, we must have $W(\overrightarrow{C}, \alpha) \neq W(\overrightarrow{C}, \beta)$, and so, by Lemma 2.1, α and β belong to different components of $\mathcal{C}_3(G)$. $\square$

Theorem 2.3. *Let G be a 3-chromatic graph. Then G is not 3-mixing.*

Proof : As G has chromatic number 3, it contains a cycle C of odd length. Let α be a 3-colouring of G, and note that as the weight of each edge in $\overrightarrow{C}$ is $+1$ or -1, $W(\overrightarrow{C}, \alpha) \neq 0$. We are done by Lemma 2.2. $\square$

3 Characterising 3-mixing bipartite graphs

We have seen that 3-chromatic graphs are not 3-mixing. What can be said for bipartite graphs? Examples of 3-mixing bipartite graphs include trees and C_4, the cycle on 4 vertices. On the other hand, all cycles except C_4 are not 3-mixing; see [4] for details. In Theorem 3.1 we distinguish between 3-mixing and non-3-mixing bipartite graphs in terms of their structure and the possible 3-colourings they may have.

If v and w are vertices of a bipartite graph G at distance two, then a *fold* on v and w is the identification of v and w (together with the removal of any double edges produced). We say that G is *foldable* to a graph H if there exists a sequence of folds that transforms G into H.

Folding of graphs, and its relation to vertex colouring, has been studied before, see for instance [6].

The main result in this section is the following.

Theorem 3.1. *Let G be a connected bipartite graph. The following are equivalent :*

(i) The graph G is not 3-mixing.

(ii) There exists a cycle C in G and a 3-colouring α of G with $W(\overrightarrow{C}, \alpha) \neq 0$.

(iii) The graph G is foldable to the 6-cycle C_6.

To prove Theorem 3.1, we need some definitions and technical lemmas. For the rest of this section, let $G = (V, E)$ denote a connected bipartite graph with vertex bipartition X, Y.

Given a 3-colouring α of G, we define a *height function for α with base X* as a function $h : V \to \mathbb{Z}$ satisfying the following conditions. (See [1, 8] for other, similar height functions.)

H1 For all $v \in X$, $h(v) \equiv 0 \pmod 2$; for all $v \in Y$, $h(v) \equiv 1 \pmod 2$.

H2 For all $uv \in E$, $h(v) - h(u) = w(\overrightarrow{uv}, \alpha)$ ($\in \{-1, +1\}$).

H3 For all $v \in V$, $h(v) \equiv \alpha(v) \pmod 3$.

If $h : V \to \mathbb{Z}$ satisfies conditions H2, H3 and also

H1′ For all $v \in X$, $h(v) \equiv 1 \pmod 2$; while for $v \in Y$, $h(v) \equiv 0 \pmod 2$.

then h is said to be a height function for α with base Y.

Observe that for a particular colouring of a given G, a height function might not exist. An example of this is the 6-cycle C_6 coloured 1-2-3-1-2-3.

Conversely, however, a function $h : V \to \mathbb{Z}$ satisfying conditions H1 and H2 induces a 3-colouring of G: the unique $\alpha : V \to \{1, 2, 3\}$ satisfying condition H3; and h is in fact a height function for this α. Observe also that if h is a height function for α with base X, then so are $h + 6$ and $h - 6$; while $h + 3$ and $h - 3$ are height functions for α with base Y. Because we will be concerned solely with the question of *existence* of height functions, we assume henceforth that for a given G, all height functions have base X. Thus we let $\mathcal{H}_X(G)$ be the set of height functions with base X corresponding to some 3-colouring of G, and define a metric m on $\mathcal{H}_X(G)$ by setting

$$m(h_1, h_2) \;=\; \sum_{v \in V} |h_1(v) - h_2(v)|,$$

for $h_1, h_2 \in \mathcal{H}_X(G)$. Note that condition H1 above implies that $m(h_1, h_2)$ is always even.

For a given height function h, $h(v)$ is said to be a *local maximum* (respectively, *local minimum*) if $h(v)$ is larger than (respectively, smaller than) $h(u)$ for all neighbours u of v. Following [8], we define the following *height transformations* on h.

– An *increasing height transformation* takes a local minimum $h(v)$ of h and transforms h into the height function h' given by $h'(x) = \begin{cases} h(x) + 2, & \text{if } x = v; \\ h(x), & \text{if } x \neq v. \end{cases}$

– A *decreasing height transformation* takes a local maximum $h(v)$ of h and transforms h into the height function h' given by $h'(x) \;=\; \begin{cases} h(x) - 2, & \text{if } x = v; \\ h(x), & \text{if } x \neq v. \end{cases}$

Notice that these height transformations give rise to transformations between the corresponding colourings. Specifically, if we let α' be the 3-colouring corresponding to h', an increasing transformation yields $\alpha'(v) = \alpha(v) - 1$, while a decreasing transformation yields $\alpha'(v) = \alpha(v) + 1$, where addition is modulo 3.

The following lemma, a simple extension of the range of applicability of a similar lemma appearing in [8], shows that colourings with height functions are connected in $\mathcal{C}_3(G)$.

Lemma 3.2 (Goldberg, Martin, and Paterson [8])**.** *Let α, β be two 3-colourings of G with corresponding height functions h_α, h_β. Then there is a path between α and β in $\mathcal{C}_3(G)$.*

Proof: We use induction on $m(h_\alpha, h_\beta)$. The lemma is trivially true when $m(h_\alpha, h_\beta) = 0$, since in this case α and β are identical.

Suppose therefore that $m(h_\alpha, h_\beta) > 0$. We show that there is a height transformation transforming h_α into some height function h with $m(h, h_\beta) = m(h_\alpha, h_\beta) - 2$, from which the lemma follows.

Without loss of generality, let us assume that there is some vertex $v \in V$ with $h_\alpha(v) > h_\beta(v)$, and let us choose v with $h_\alpha(v)$ as large as possible. We show that such a v must be a local maximum of h_α. Let u be any neighbour of v. If $h_\alpha(u) > h_\beta(u)$, then it follows that $h_\alpha(v) > h_\alpha(u)$, since v was chosen with $h_\alpha(v)$ maximum, and $|h_\alpha(v) - h_\alpha(u)| = 1$. If, on the other hand, $h_\alpha(u) \leq h_\beta(u)$, we have $h_\alpha(v) \geq h_\beta(v) + 1 \geq h_\beta(u) \geq h_\alpha(u)$, which in fact means $h_\alpha(v) > h_\alpha(u)$.

Thus $h_\alpha(v) > h_\alpha(u)$ for all neighbours u of v, and we can apply a decreasing height transformation to h_α at v to obtain h. Clearly $m(h, h_\beta) = m(h_\alpha, h_\beta) - 2$. $\square$

The next lemma tells us that for a given 3-colouring, non-zero weight cycles are, in some sense, the obstructing configurations forbidding the existence of a corresponding height function.

Lemma 3.3. *Let α be a 3-colouring of G with no corresponding height function. Then G contains a cycle C for which $W(\overrightarrow{C}, \alpha) \neq 0$.*

Proof: For a path P in G, let $\overrightarrow{P}$ denote one of the two possible directed paths obtainable from P, and let

$$W(\overrightarrow{P}, \alpha) = \sum_{\vec{e} \in E(\overrightarrow{P})} w(\vec{e}, \alpha),$$

where $w(\vec{e}, \alpha)$ takes values as defined in (1).

Notice that if a colouring does have a height function, it is possible to construct one by fixing a vertex $x \in X$, giving x an appropriate height (satisfying properties H1–H3) and then assigning heights to all vertices in V by following a breadth-first ordering from x.

Whenever we attempt to construct a height function h for α in such a fashion, we must come to a stage in the ordering where we attempt to give some vertex v a height $h(v)$ and find ourselves unable to because v has a neighbour u with a previously assigned height $h(u)$ and $|h(u) - h(v)| > 1$. Letting P be a path between u and v formed by vertices that have been assigned a height, and choosing the appropriate orientation of P, we have $w(\overrightarrow{P}, \alpha) = |h(u) - h(v)|$. The lemma now follows by letting C be the cycle formed by P and the edge uv. $\square$

The following lemma is obvious.

Lemma 3.4. *Let u and v be vertices on a cycle C in a graph G, and suppose there is a path P between u and v in G internally disjoint from C. Let α be a 3-colouring of G. Let C' and C'' be the two cycles formed from P and edges of C, and let $\overrightarrow{C'}, \overrightarrow{C''}$ be the orientations of C', C'' induced by an orientation $\overrightarrow{C}$ of C (so the edges of P have opposite orientations in $\overrightarrow{C'}$ and $\overrightarrow{C''}$). Then $W(\overrightarrow{C}, \alpha) = W(\overrightarrow{C'}, \alpha) + W(\overrightarrow{C''}, \alpha)$.*

Note this tells us that $W(\overrightarrow{C}, \alpha) \neq 0$ implies $W(\overrightarrow{C'}, \alpha) \neq 0$ or $W(\overrightarrow{C''}, \alpha) \neq 0$.

Proof of Theorem 3.1: Let G be a connected bipartite graph.

(i) $\implies$ (ii). Suppose $\mathcal{C}_3(G)$ is not connected. Take two 3-colourings of G, α and β, in different components of $\mathcal{C}_3(G)$. By Lemma 3.2 we know at least one of them, say α, has no corresponding height function, and, by Lemma 3.3, there is a cycle C in G with $W(\overrightarrow{C}, \alpha) \neq 0$.

(ii) $\implies$ (iii). Let G contain a cycle C with $W(\overrightarrow{C}, \alpha) \neq 0$ for some 3-colouring α of G. Because $W(\overrightarrow{C_4}, \beta) = 0$ for any 3-colouring β of C_4, it follows that $C = C_n$ for some even $n \geq 6$. If $G = C$, then it is easy to find a sequence of folds that will yield C_6. If G is C plus some chords, then Lemma 3.4 tells us that there is a smaller cycle C' with $W(\overrightarrow{C'}, \alpha) \neq 0$. Thus if $G \neq C$, we can assume that $V(G) \neq V(C)$, and we describe how to fold a pair of vertices so that (ii) remains satisfied (for a specified cycle with G replaced by the graph created by the fold and α replaced by its restriction to that graph; also denoted α); by repetition, we can obtain a graph that is a cycle and, by the previous observations, the implication is proved.

Note that we shall choose vertices coloured alike to fold so that the restriction of α to the graph obtained is well-defined and proper. If C has three consecutive vertices u, v, w with $\alpha(u) = \alpha(w)$, folding u and w yields a graph containing a cycle $C' = C_{n-2}$ with $W(\overrightarrow{C'}, \alpha) = W(\overrightarrow{C}, \alpha)$. Otherwise C is coloured 1-2-3-$\cdots$-1-2-3. We can choose u, v, w to be three consecutive vertices of C, such that there is a vertex $x \notin V(C)$ adjacent to v. Suppose, without loss of generality, that $\alpha(x) = \alpha(u)$, and fold x and u to obtain a graph in which $W(\overrightarrow{C}, \alpha)$ is unchanged.

(iii) $\implies$ (i). Suppose G is foldable to C_6. Take two 3-colourings of C_6 not connected by a path in $\mathcal{C}_3(C_6)$ (1-2-3-1-2-3 and 1-2-1-2-1-2, for example). Considering the appropriate orientation of C_6, note that the first colouring has weight 6 and the second has weight 0. We construct two 3-colourings of G not connected by a path in $\mathcal{C}_3(G)$ as follows. Consider the reverse sequence of folds that gives G from C_6. Following this sequence, for each colouring of C_6, give every pair of new vertices introduced by an "unfolding" the same colour as the vertex from which they originated. In this manner we obtain two 3-colourings of G, α and β, say. Observe that every unfolding maintains a cycle in G which has weight 6 with respect to the colouring induced by the first colouring of C_6 and weight 0 with respect to the second induced colouring. This means G will contain a cycle C for which $W(\overrightarrow{C}, \alpha) = 6$ and $W(\overrightarrow{C}, \beta) = 0$, showing that α and β cannot possibly be in the same connected component of $\mathcal{C}_3(G)$.

This completes the proof of the theorem. □

4 The complexity of 3-mixing for bipartite graphs

Let us now turn our attention to the computational complexity of deciding whether or not a 3-colourable graph G is 3-mixing. From Theorem 2.3 we know that we can restrict our attention to bipartite graphs, so we state the decision problem formally as follows.

3-Mixing
Instance : A connected bipartite graph G.
Question : Is G 3-mixing?

Observing that Theorem 3.1 gives us two polynomial-time verifiable certificates for when G is *not* 3-mixing, we immediately obtain that 3-Mixing is in the complexity class coNP. By the same theorem, the following decision problem is the complement of 3-Mixing.

Foldable-to-C_6
Instance : A connected bipartite graph G.
Question : Is G foldable to C_6?

We will prove the following result, stated in the introduction.

Theorem 1.1. *The decision problem* 3-Mixing *is* coNP*-complete.*

Our proof will in fact show that Foldable-to-C_6 is NP-complete. We will obtain a reduction from the following decision problem.

Retractable-to-C_6
Instance : A connected bipartite graph G with an induced 6-cycle S.
Question : Is G *retractable* to S? That is, does there exist a homomorphism $r : V(G) \to V(S)$ such that $r(v) = v$ for all $v \in V(S)$?

In [11] it is mentioned, without references, that Tomás Feder and Gary MacGillivray proved independently that Retractable-to-C_6 is NP-complete by reduction from 3-Colourability. For completeness we give a sketch of a proof.

Theorem 4.1 (Feder, MacGillivray, see [11])**.** Retractable-to-C_6 *is* NP*-complete.*

Sketch of proof : It is clear that Retractable-to-C_6 is in NP.

Given a graph G, construct a new graph G' as follows: subdivide every edge uv of G by inserting a vertex y_{uv} between u and v. Also add new vertices a, b, c, d, e together with edges za, ab, bc, cd, de, ez, where z is a particular vertex of G (any one will do). The graph G' is clearly connected and bipartite, and the vertices z, a, b, c, d, e induce a 6-cycle S. We will prove that G is 3-colourable if and only if G' retracts to the induced 6-cycle S.

Assume that G is 3-colourable and take a 3-colouring τ of G with $\tau(z) = 1$. From τ we construct a 6-colouring σ of G'. For this, first set $\sigma(x) = \tau(x)$, if $x \in V(G)$. For the new vertices y_{uv} set $\sigma(y_{uv}) = \begin{cases} 4, & \text{if } \tau(u) = 1 \text{ and } \tau(v) = 2, \\ 5, & \text{if } \tau(u) = 2 \text{ and } \tau(v) = 3, \\ 6, & \text{if } \tau(u) = 3 \text{ and } \tau(v) = 1. \end{cases}$ And for the cycle S we take $\sigma(a) = 4$, $\sigma(b) = 2$, $\sigma(c) = 5$, $\sigma(d) = 3$ and $\sigma(e) = 6$. Now define $r : V(G') \to V(S)$ by setting $r(x) = z$, if $\sigma(x) = 1$; $r(x) = a$, if $\sigma(x) = 4$; $r(x) = b$, if $\sigma(x) = 2$; $r(x) = c$, if $\sigma(x) = 5$; $r(x) = d$, if $\sigma(x) = 3$; and $r(x) = e$, if $\sigma(x) = 6$. It is easy to check that r is a retraction of G' to S.

Conversely, suppose G' retracts to S. We can use this retraction to define a 6-colouring of G' in a similar way to that in which we defined r from σ in the preceding paragraph. The restriction of this 6-colouring to G yields a 3-colouring of G, completing the proof. $\square$

Our proof of Theorem 1.1 follows [11], where, as a special case of the main result of that paper, the following problem is proved to be NP-complete.

COMPACTABLE-TO-C_6
Instance: A connected bipartite graph G.
Question: Is G *compactable* to C_6? That is, does there exist an edge-surjective homomorphism $c : V(G) \to V(C_6)$?

In [11] a polynomial reduction from RETRACTABLE-TO-C_k to COMPACTABLE-TO-C_k, with $k \geq 6$ even, is given. We will use exactly the same transformation for $k = 6$ to prove that FOLDABLE-TO-C_6 is NP-complete.

Proof of Theorem 1.1: As mentioned before, we will show that 3-MIXING is coNP-complete by showing that FOLDABLE-TO-C_6 is NP-complete. And we do that by giving a polynomial reduction from RETRACTABLE-TO-C_6 to FOLDABLE-TO-C_6.

So consider an instance of RETRACTABLE-TO-C_6: a connected bipartite graph G and an induced 6-cycle S. From G we construct, in time polynomial in the size of G, an instance G' of FOLDABLE-TO-C_6 such that

$$G \text{ retracts to } S \text{ if and only if } G' \text{ is foldable to } C_6. \tag{$*$}$$

Assume G has vertex bipartition (G_A, G_B). Let $V(S) = S_A \cup S_B$, where $S_A = \{h_0, h_2, h_4\}$ and $S_B = \{h_1, h_3, h_5\}$, and assume $E(S) = \{\, h_0h_1, \ldots, h_4h_5, h_5h_0 \,\}$.

The construction of G' is as follows.

- For every vertex $a \in G_A \setminus S_A$, add to G new vertices $u_1^a, u_2^a, w_1^a, y_1^a, y_2^a$, together with edges $u_1^a h_0, au_2^a, w_1^a h_3, aw_1^a, u_1^a w_1^a, y_1^a h_5, y_2^a h_2, u_1^a y_1^a, w_1^a y_2^a, u_1^a u_2^a, y_1^a y_2^a$.
- For every vertex $b \in G_B \setminus S_B$, add to G new vertices $u_1^b, w_1^b, w_2^b, y_1^b, y_2^b$, together with edges $u_1^b h_0, bu_1^b, w_1^b h_3, bw_2^b, u_1^b w_1^b, y_1^b h_5, y_2^b h_2, u_1^b y_1^b, w_1^b y_2^b, w_1^b w_2^b, y_1^b y_2^b$.
- For every edge $ab \in E(G) \setminus E(S)$, with $a \in G_A \setminus S_A$ and $b \in G_B \setminus S_B$, add two new vertices: x_a^{ab} adjacent to a and u_1^a; and x_b^{ab} adjacent to b, w_1^b and x_a^{ab}.

From the construction it is clear that G' is connected and bipartite. Note that G' contains G as an induced subgraph, and note also that the subgraphs constructed around a vertex $a \in G_A \setminus S_A$ and a vertex $b \in G_B \setminus S_B$ are isomorphic; these are depicted below in Figures 1 and 2.

We will prove ($*$) via a sequence of claims.

Claim 4.1. *Suppose G retracts to S. Then G is foldable to C_6.*

Proof: The fact that G retracts to S means we have a homomorphism $r : V(G) \to V(S)$ such that $r(v) = v$ for all $v \in V(S)$. Define a partition $\{\, R_i \mid i = 0, 1, \ldots, 5 \,\}$ of $V(G)$ by setting $v \in R_i \iff r(v) = h_i$. Because r is a homomorphism, we know any edge $e \in E(G)$

Figure 1 : The subgraph of G' added around a vertex $a \in G_A \setminus S_A$, together with the 6-cycle S.

Figure 2 : The subgraph of G' added around a vertex $b \in G_B \setminus S_B$, together with the 6-cycle S.

has one vertex in R_j and another in R_{j+1}, for some j, where subscript addition is modulo 6. Using this partition of $V(G)$, we show that G is foldable to a 6-cycle (to S, in fact). We describe how to fold a pair of vertices such that the resulting (smaller) graph still has S as an induced subgraph; by repetition, this will eventually yield S. Supposing $V(G) \neq V(S)$ (for else we are done), let $E^- = E(G) \setminus E(S)$. Because G is connected, there must be an edge $uv \in E^-$ with $u \in V(S)$ and $v \in V(G) \setminus V(S)$. Suppose $v \in R_j$, for some $j \in \{0, 1, \ldots, 5\}$. Fold v with h_j, and note that the resulting graph remains bipartite, connected and contains S as an induced subgraph. Denote the resulting graph by G and repeat. □

We now prove the 'only if' part of $(*)$.

Claim 4.2. *Suppose G retracts to S. Then G' is foldable to C_6.*

Proof : By Claim 4.1, G is foldable to C_6. In fact, by the proof of Claim 4.1, we know G is foldable to S. Because G is an induced subgraph of G', we can follow, in G', the sequence of folds that gives S from G. We now show how, after following this sequence of folds, we can choose some further folds that will leave us with S. For a vertex $v \in V(G) \setminus V(S)$, we will fold into S all vertices introduced to G' on account of v, yielding a smaller graph still containing S as an induced subgraph. By repetition, we will eventually end up with just S.

First let us consider where a vertex $a \in G_A \setminus S_A$ with no neighbours in $G_B \setminus S_B$ might have been folded to, and how we could continue folding. There are three possibilities.

1. The vertex a has been folded with h_1. In that case fold y_1^a with h_0, y_2^a with h_1, u_1^a with h_1, u_2^a with h_0, and w_1^a with h_2.
2. The vertex a has been folded with h_3. In that case fold y_1^a with h_4, y_2^a with h_3, u_1^a with h_5, u_2^a with h_4, and w_1^a with h_4.
3. The vertex a has been folded with h_5. In that case fold y_1^a with h_4, y_2^a with h_3, u_1^a with h_5, u_2^a with h_0, and w_1^a with h_4.

Similarly, let us consider where a vertex $b \in G_B \setminus S_B$ with no neighbours in $G_A \setminus S_A$ might have been folded to, and how we could continue folding. Again, there are three possibilities.

1. The vertex b has been folded with h_0. In that case fold y_1^b with h_0, y_2^b with h_1, u_1^b with h_1, w_1^b with h_2, and w_2^b with h_1.
2. The vertex b has been folded with h_2. In that case fold y_1^b with h_0, y_2^b with h_1, u_1^b with h_1, w_1^b with h_2, and w_2^b with h_3.
3. The vertex b has been folded with h_4. In that case fold y_1^b with h_4, y_2^b with h_3, u_1^b with h_5, w_1^b with h_4, and w_2^b with h_3.

Now let us consider the case where a vertex $a \in G_A \setminus S_A$ is adjacent to a vertex $b \in G_B \setminus S_B$. There are six cases to consider, corresponding to the six edges of S to which ab might have been folded. Often there will be a choice of folds; for each case we give just one.

1. The edge ab has been folded to h_1h_2. We can use the previous case analyses to conclude that u_1^a must be folded with h_1 and w_1^b with h_2. Now we must deal with x_a^{ab} and x_b^{ab}. Folding x_a^{ab} with h_2 and x_b^{ab} with h_1 gives us what we require.
2. The edge ab has been folded to h_1h_0. Then we conclude u_1^a must be folded with h_1 and w_1^b with h_2. Now fold x_a^{ab} with h_0 and x_b^{ab} with h_1.
3. The edge ab has been folded to h_3h_4. Then u_1^a must be folded with h_5 and w_1^b with h_4. Now fold x_a^{ab} with h_4 and x_b^{ab} with h_3.
4. The edge ab has been folded to h_3h_2. Then u_1^a must be folded with h_5 and w_1^b with h_2. Now fold x_a^{ab} with h_4 and x_b^{ab} with h_3.
5. The edge ab has been folded to h_5h_0. Then u_1^a must be folded with h_5 and w_1^b with h_2. Now fold x_a^{ab} with h_0 and x_b^{ab} with h_1.
6. The edge ab has been folded to h_5h_4. Then u_1^a must be folded with h_5 and w_1^b with h_4. Now fold x_a^{ab} with h_4 and x_b^{ab} with h_5.

This completes the proof of the claim. □

We must now prove the 'if' part of (∗). We do this via the next three claims.

Claim 4.3. *Suppose G' is foldable to C_6. Then G' is compactable to C_6.*

Proof: The fact that G' is foldable to the 6-cycle $C_6 = k_0k_1k_2k_3k_4k_5k_0$ means there exists a homomorphism $c : V(G') \to V(C_6)$. In order to make this precise, let us define sets P_i, for $i = 0, 1, \ldots, 5$, as follows. Initially, set $P_i = \{k_i\}$. Now let us consider the reverse sequence of "unfoldings" that yields G' from C_6. Following this sequence, suppose a vertex $v \in P_j$ is unfolded. Delete v from P_j and add to P_j the two vertices that were identified to give v in the original fold. Repeat this until G' is obtained, and now define c by setting, for $v \in V(G')$, $c(v) = k_i \iff v \in P_i$. Clearly the sets P_i form a partition of $V(G')$ and so c is well-defined. In addition, by the way the sets P_i have been constructed, it is clear that any edge $uv \in E(G')$ has one vertex in P_j and the other in P_{j+1}, for some $j \in \{0, 1, \dots, 5\}$. This means $c(u)c(v) \in E(C_6)$ and so c is a homomorphism. Moreover, it is edge-surjective: the P_i's are all non-empty and there is at least one edge between every pair P_i, P_{i+1}. □

The proof of the following claim is a specialisation of the proof in [11] that if G' is compactable to C_6, then G' retracts to S.

We need some further notation. As usual, for a set S and a function f, we let $f(S) = \{ f(s) \mid s \in S \}$. For vertices u, v in a graph H, $d_H(u, v)$ denotes the distance between u and v; and for a vertex u and a set of vertices S we have $d_H(S, u) = \min\{ d_H(v, u) \mid v \in S \}$.

Claim 4.4 (Vikas [11])**.** *Suppose G' is foldable to C_6. Then G' retracts to S.*

Proof: By Claim 4.3 we know there exists a compaction $c : V(G') \to V(C_6)$. We prove c is in fact a retraction to S. To do this, we must show that for all $v \in V(S)$, $c(v) = v$. For convenience, we now use the same notation for C_6 and S; that is, we let $V(C_6) = \{h_0, h_1, \dots, h_5\}$ and $E(C_6) = \{ h_0h_1, \dots, h_4h_5, h_5h_0 \}$.

Let $U = \{ u_1^v \mid v \in V(G) \setminus V(S) \} \cup \{h_0, h_1, h_5\}$ and $W = \{ w_1^v \mid v \in V(G) \setminus V(S) \} \cup \{h_2, h_3, h_4\}$. Because both these vertex sets induce subgraphs of diameter 2 in G', $c(U)$ and $c(W)$ must each induce a path of length 1 or 2 in C_6. We prove they each induce a path of length 2.

Suppose that $c(U)$ has only two vertices, adjacent in C_6. Thus we let $c(U) = \{h_0, h_1\}$, with $c(h_0) = h_0$. (Due to the symmetry of C_6, we can, if necessary, redefine c in this way.) Let $U^- = U \setminus \{h_0\}$. Because h_0 is adjacent to every other vertex in U, $c(U^-) = \{h_1\}$. It is easy to check that for any $g \in G'$, $d_{G'}(U^-, g) \leq 2$. But we have $d_{C_6}(c(U^-), h_4) = d_{C_6}(h_1, h_4) = 3$, which means no $g \in G'$ can be mapped to h_4 under c, contradicting the fact that c is a compaction.

Hence $c(U)$ induces a path on three vertices. By a similar argument, the same applies to $c(W)$. By the symmetry of C_6, we can without loss of generality take $c(U) = \{h_1, h_0, h_5\}$. This means that $c(h_0) = h_0$. We now prove that $c(h_3) = h_3$.

Let gg' be an edge of G' that is mapped to h_3h_2 or h_3h_4, with $c(g) = h_3$, and $c(g') = h_2$ or $c(g') = h_4$. Note that h_3 is at distance 2 from $c(U)$ in C_6 while h_2 and h_4 are at distance 1 from $c(U)$ in C_6. This means that $d_{G'}(U, g) \geq 2$ and $d_{G'}(U, g') \geq 1$. Earlier we noted that the distance between U^- and any vertex of G' is at most 2, which means that $d_{G'}(U, g) \leq 2$, so in fact $d_{G'}(U, g) = 2$. Because G' is bipartite, $d_{G'}(U, g') = 1$. Hence g is one of $a, x_b^{ab}, h_3, y_2^a, y_2^b, w_2^b$, and g' is one of $b, x_a^{ab}, u_2^a, h_2, h_4, y_1^a, y_1^b, w_1^a, w_1^b$, for some $a \in G_A \setminus S_A$, $b \in G_B \setminus S_B$. Given that $c(h_0) = h_0$, we cannot have $c(h_3) = h_2$ or $c(h_3) = h_4$. Aiming for a contradiction, let us suppose that $c(h_3) \neq h_3$. Then no edge of G' with h_3 as an endpoint covers h_3h_2 or h_3h_4. Hence gg' must be one of the following: $ax_a^{ab}, ab, au_2^a, aw_1^a$, $x_b^{ab}x_a^{ab}, x_b^{ab}b, x_b^{ab}w_1^b$, $y_2^ay_1^a, y_2^aw_1^a, y_2^ah_2, y_2^by_1^b, y_2^bw_1^b, y_2^bh_2$, $w_2^bw_1^b, w_2^bb$. If ah_2 or ah_4 is an edge of G', then we also need to consider such an edge as a possible candidate for gg'. By previous assumptions, we have $c(h_3) = h_1$ or $c(h_3) = h_5$. We now prove that $c(h_3) \neq h_3$ is impossible as follows. We first assume $c(h_3) = h_1$ and show that no possible edge for gg' covers h_3h_4, and then assume $c(h_3) = h_5$ and show that no possible edge for gg' covers h_3h_2. Thus let us assume $c(h_3) = h_1$.

Let us suppose that for some $v \in V(G) \setminus V(S)$, $y_2^v w_1^v$ covers h_3h_4, so $c(y_2^v) = h_3$ and $c(w_1^v) = h_4$. But $c(h_3) = h_1$, and since h_3 an w_1^v are adjacent, we must have $c(w_1^v) = h_0$ or $c(w_1^v) = h_2$, a contradiction.

By exactly the same argument, we come to the conclusion that none of the edges aw_1^a, $w_2^b w_1^b$, $x_b^{ab} w_1^b$ can cover the edge h_3h_4. A similar argument applies to $y_2^v h_2$.

Suppose that for some $v \in V(G) \setminus V(S)$, $y_2^v y_1^v$ covers h_3h_4, so $c(y_2^v) = h_3$ and $c(y_1^v) = h_4$. Now $c(u_1^v) = h_1$ or $c(u_1^v) = h_5$, but since u_1^v and y_1^v are adjacent we must have $c(u_1^v) = h_5$. Because $c(w_1^v)$ must be adjacent to $c(y_2^v) = h_3$, $c(w_1^v) = h_2$ or $c(w_1^v) = h_4$. But u_1^v is adjacent to w_1^v, so $c(w_1^v) = h_4$. This means $y_2^v w_1^v$ covers h_3h_4, which we have already seen is impossible.

Now suppose that for some $b \in G_B \setminus S_B$, $w_2^b b$ covers h_3h_4, so $c(w_2^b) = h_3$ and $c(b) = h_4$. If $c(b) = h_4$, we must have $c(u_1^b) = h_3$ or $c(u_1^b) = h_5$. But $c(h_0) = h_0$ means $c(u_1^b) = h_1$ or $c(u_1^b) = h_5$, so $c(u_1^b) = h_5$. This implies, since $c(w_1^b) = h_2$ or $c(w_1^b) = h_4$, that $c(w_1^b) = h_4$. But this means that $w_2^b w_1^b$ covers h_3h_4, which we have already excluded as a possibility.

Assume that for some $a \in G_A \setminus S_A$, au_2^a covers h_3h_4, so $c(a) = h_3$ and $c(u_2^a) = h_4$. Because u_1^a and u_2^a are adjacent, $c(u_1^a) = h_3$ or $c(u_1^a) = h_5$, but since u_1^a is adjacent to h_0 and $c(h_0) = h_0$, we have $c(u_1^a) = h_5$. Similarly, $c(w_1^a) = h_2$ or $c(w_1^a) = h_4$, but since w_1^a and u_1^a are adjacent, we have $c(w_1^a) = h_4$. Hence aw_1^a covers h_3h_4, but we have already seen this is impossible.

Now assume that for some $a \in G_A \setminus S_A$, ax_a^{ab} covers h_3h_4, so $c(a) = h_3$ and $c(x_a^{ab}) = h_4$. Now $c(u_1^a) = h_1$ or $c(u_1^a) = h_5$, but since u_1^a and x_a^{ab} are adjacent, we have $c(u_1^a) = h_5$. Because $c(u_2^a)$ must be adjacent to $c(a) = h_3$ as well as $c(u_1^a) = h_5$, we have $c(u_2^a) = h_4$. Hence au_2^a covers h_3h_4, but we have already seen this is impossible.

Suppose that for some $b \in G_B \setminus S_B$, $x_b^{ab} b$ covers h_3h_4, so $c(x_b^{ab}) = h_3$ and $c(b) = h_4$. Now $c(u_1^b) = h_1$ or $c(u_1^b) = h_5$, but since b and u_1^b are adjacent, we must have $c(u_1^b) = h_5$. Because $c(w_1^b)$ must be adjacent to $c(x_b^{ab}) = h_3$, we have $c(w_1^b) = h_2$ or $c(w_1^b) = h_4$. But u_1^b and w_1^b are adjacent, so $c(w_1^b) = h_4$. This means $x_b^{ab} w_1^b$ covers h_3h_4, which we have already ruled out as a possibility.

Now suppose that for some $a \in G_A \setminus S_A$ and some $b \in G_B \setminus S_B$, ab covers h_3h_4, so $c(a) = h_3$ and $c(b) = h_4$. Since u_2^a is adjacent to a and we have seen au_2^a does not cover h_3h_4, we must have $c(u_2^a) = h_2$. Now $c(u_1^a) = h_1$ or $c(u_1^a) = h_5$, but since u_1^a and u_2^a are adjacent, we must have $c(u_1^a) = h_1$. Also, $c(x_a^{ab})$ must be adjacent to $c(u_1^a) = h_1$ and $c(a) = h_3$, so $c(x_a^{ab}) = h_2$. Similarly, $c(x_b^{ab})$ must be adjacent to $c(x_a^{ab}) = h_2$ and $c(b) = h_4$, so $c(x_b^{ab}) = h_3$. But this means $x_b^{ab} b$ covers h_3h_4, which we have already seen is impossible.

Suppose that for some $a \in G_A \setminus S_A$ and some $b \in G_B \setminus S_B$, $x_b^{ab} x_a^{ab}$ covers h_3h_4, so $c(x_b^{ab}) = h_3$ and $c(x_a^{ab}) = h_4$. Since a is adjacent to x_a^{ab} and we have seen ax_a^{ab} does not cover h_3h_4, we must have $c(a) = h_5$. Because $c(b)$ must be adjacent to $c(a) = h_5$ and $c(x_b^{ab}) = h_3$, we have $c(b) = h_4$. But then $x_b^{ab} b$ covers h_3h_4, and we have seen this is impossible.

Lastly, if ah_2 (or ah_4) is an edge of G', assuming $c(a) = h_3$ and $c(h_2) = h_4$ (or $c(a) = h_3$ and $c(h_4) = h_4$) immediately leads us to a contradiction, since $c(h_3) = h_1$.

From all this we obtain that assuming $c(h_3) = h_1$ leads us to the conclusion that no edge of G' covers h_3h_4, contradicting the fact that c is a compaction.

Similarly, one can show that assuming $c(h_3) = h_5$ leads to the conclusion that no edge of G' covers h_2h_3; details are left to the reader.

Hence $c(h_3) = h_3$, which means that $c(W) = \{h_2, h_3, h_4\}$.

Now we show $c(h_1) \neq c(h_5)$. To the contrary, assume $c(h_1) = c(h_5)$. Since $c(h_0) = h_0$, we have $c(h_1), c(h_5) \in \{h_1, h_5\}$. Due to symmetry, we can without loss of generality assume $c(h_1) = c(h_5) = h_1$. Since $c(U) = \{h_1, h_0, h_5\}$, it must be the case that $c(u_1^v) = h_5$ for some $v \in V(G) \setminus V(S)$. Now $c(w_1^v)$ and $c(h_2)$ must both be adjacent to $c(h_3) = h_3$, so $c(w_1^v), c(h_2) \in \{h_2, h_4\}$. Because $c(u_1^v) = h_5$ and u_1^v and w_1^v are adjacent, $c(w_1^v) = h_4$. Similarly, because $c(h_0) = h_0$ and h_1 and h_2 are adjacent, $c(h_2) = h_2$. Now $c(y_2^v)$ must be adjacent to $c(h_2) = h_2$ and $c(w_1^v) = h_4$, so $c(y_2^v) = h_3$. Also, $c(y_1^v)$ must be adjacent to $c(h_5) = h_1$ and $c(u_1^v) = h_5$, so $c(y_1^v) = h_0$. Thus we have that y_1^v and y_2^v are adjacent in G', but $c(y_1^v) = h_0$ and $c(y_2^v) = h_3$ are not adjacent in C_6, a contradiction.

Hence $c(h_1) \neq c(h_5)$. That is, $c(\{h_1, h_5\}) = \{h_1, h_5\}$. Without loss of generality, we can take $c(h_1) = h_1$ and $c(h_5) = h_5$. Since $c(h_3) = h_3$, we have $c(h_2), c(h_4) \in \{h_2, h_4\}$. Because h_1 and h_2 are adjacent in G' and the distance between $c(h_1) = h_1$ and h_4 in C_6 is 3, it must be that $c(h_2) \neq h_4$ and so $c(h_2) = h_2$. Similarly, because h_5 and h_4 are adjacent in G' and the distance between $c(h_5) = h_5$ and h_2 in C_6 is 3, it must be that $c(h_4) \neq h_2$, and so $c(h_4) = h_4$.

Thus $c(h_i) = h_i$ for all $i = 0, 1, \ldots, 5$, and $c : V(G') \to V(C_6)$ is a retraction. ☐

The last claim is a simple observation that completes the proof of (∗) and thus also of Theorem 1.1.

Claim 4.5. *Suppose G' is foldable to C_6. Then G retracts to S.*

Proof : By Claims 4.3 and 4.4 we know there exists a retraction $r : V(G') \to V(S)$. Because S is an induced subgraph of G, and G is an induced subgraph of G', restricting r to G gives us what we need. ☐

5 A polynomial-time algorithm for planar bipartite graphs

In this section, we prove the following.

Theorem 1.2. *Restricted to* planar *bipartite graphs, the decision problem* 3-Mixing *is in the complexity class* P.

To prove the theorem we need two lemmas.

Lemma 5.1. *Let P be a shortest path between distinct vertices u and v in a connected bipartite graph H. Then H is foldable to P.*

Proof : Let P have vertices $u = v_0, v_1, \dots, v_{k-1}, v_k = v$, and let T be a breadth-first spanning tree of H rooted at u that contains P (we can choose T so that it contains P since P is a shortest path). To fold H to P, first fold all vertices at distance one from u in T to v_1. Next fold all vertices at distance two (in T) from u to v_2, and so on until finally all vertices at distance k from u are folded to $v_k = v$. We can then obtain P by making, if necessary, arbitrary folds on the vertices at distance at least $k+1$ from u. □

In the following, when we say some vertices of a graph are *properly precoloured*, we mean that they are assigned colours such that the subgraph induced by these vertices is properly coloured.

Lemma 5.2. *Let H be a bipartite graph, and suppose the vertices of a 4-cycle in H are properly precoloured using colours from* $1, 2, 3$. *Then this 3-colouring can be extended to a proper 3-colouring of H.*

Proof : Since any 3-colouring of a four cycle C_4 has two vertices with the same colour, we can without loss of generality assume the four vertices are coloured 1-2-1-2 or 1-2-1-3. In the first instance, since H is bipartite, we can extend the precolouring to a colouring of H using colours 1 and 2 only. For the second case, we can use the same colouring, except leaving the vertex coloured 3 as it is. □

The sequence of claims that follows outlines an algorithm that, given a connected bipartite planar graph G as input, determines in polynomial time whether or not G is 3-mixing. We first show how we can take the input graph to be 2-connected. Recall that a *block* of a graph is a maximal connected subgraph that has no cut-vertex.

Claim 5.1. *Let G be a connected bipartite planar graph. Then G is 3-mixing if and only if each block of G is 3-mixing.*

Proof : If G is 3-mixing, then clearly so are its blocks. Conversely, if G is not 3-mixing, we know by Theorem 3.1 that there must exist a 3-colouring α of G and a cycle C in G such that $W(\overrightarrow{C}, \alpha) \neq 0$. But because C must lie completely inside a (2-connected) block of G, we know that there is at least one block of G that is not 3-mixing either. □

Let us now consider an embedding of our 2-connected bipartite planar graph G in the plane, and let us identify G with this embedding. (Throughout the rest of this section, we will usually, for ease of reference, identify a planar graph with a given embedding of the graph in the plane.) Given a cycle D in G, denote by $\mathrm{Int}(D)$ and $\mathrm{Ext}(D)$ the sets of vertices inside and outside of D, respectively. Note that the vertices of D itself are not included in $\mathrm{Int}(D)$ nor in $\mathrm{Ext}(D)$. If both $\mathrm{Int}(D)$ and $\mathrm{Ext}(D)$ are non-empty, D is said to be separating and, in this case, we define $G_{\mathrm{Int}}(D) = G - \mathrm{Ext}(D)$ and $G_{\mathrm{Ext}}(D) = G - \mathrm{Int}(D)$. Note that D is part of both these graphs.

We now consider the case where the planar embedding of G has a separating 4-cycle.

Claim 5.2. *Let G be a 2-connected bipartite planar graph, and suppose that G has a planar embedding with a separating 4-cycle D. Then G is 3-mixing if and only if $G_{\mathrm{Int}}(D)$ and $G_{\mathrm{Ext}}(D)$ are both 3-mixing.*

Proof : To prove necessity, we show that if one of $G_{\text{Int}}(D)$ or $G_{\text{Ext}}(D)$ is not 3-mixing, then G is not 3-mixing. Without loss of generality, suppose that $G_{\text{Int}}(D)$ is not 3-mixing, so there exists a 3-colouring α of $G_{\text{Int}}(D)$ and a cycle C in $G_{\text{Int}}(D)$ with $W(\overrightarrow{C},\alpha) \neq 0$. By Lemma 5.2, the 3-colouring of the vertices of the 4-cycle D can be extended to a 3-colouring of $G_{\text{Ext}}(D)$. The combination of the 3-colourings of $G_{\text{Int}}(D)$ and $G_{\text{Ext}}(D)$ gives a 3-colouring of G with a non-zero weight cycle, showing that G is not 3-mixing.

To prove sufficiency, we show that if G is not 3-mixing, then at least one of $G_{\text{Int}}(D)$ and $G_{\text{Ext}}(D)$ must fail to be 3-mixing. Suppose that α is a 3-colouring of G for which there is a cycle C with $W(\overrightarrow{C},\alpha) \neq 0$. If C is contained entirely within $G_{\text{Int}}(D)$ or $G_{\text{Ext}}(D)$ we are done; so let us assume that C has some vertices in $\text{Int}(D)$ and some in $\text{Ext}(D)$. Then applying Lemma 3.4 (repeatedly, if necessary) we can find a cycle C' contained entirely in $G_{\text{Int}}(D)$ or $G_{\text{Ext}}(D)$ for which $W(\overrightarrow{C'},\alpha) \neq 0$, completing the proof. □

We need two further claims to complete the description of our algorithm. We call a face of G with k edges in its boundary a *k-face*, and a face with at least k edges in its boundary a *$\geq k$-face*. The number of ≥ 6-faces in G (which we can now assume has no separating 4-cycle) will in fact determine whether or not G is 3-mixing.

Claim 5.3. *Let G be a 2-connected bipartite planar graph. Suppose that G has a planar embedding with no separating 4-cycle, and suppose that every internal face of the embedding is a 4-face. Then G is 3-mixing.*

Proof : Let α be any 3-colouring of G and let C be any cycle in G. We show $W(\overrightarrow{C},\alpha) = 0$ by induction on the number of faces inside C. If there is just one face inside C, C is in fact a facial 4-cycle and $W(\overrightarrow{C},\alpha) = 0$.

For the inductive step, let C be a cycle with $r \geq 2$ faces in its interior. If, for two consecutive vertices u, v of C, we have vertices $a, b \in \text{Int}(C)$ together with edges ua, ab, bv in G, let C' be the cycle formed from C by the removal of the edge uv and the addition of edges ua, ab, bv. If not, check whether for three consecutive vertices u, v, w of C, there is a vertex $a \in \text{Int}(C)$ with edges ua, aw in G. If so, let C' be the cycle formed from C by the removal of the vertex v and the addition of the edges ua, aw. If neither of the previous two cases apply, we must have, for u, v, w, x four consecutive vertices of C, an edge ux inside C. In such a case, let C' be the cycle formed from C by the removal of vertices v, w and the addition of the edge ux.

In all cases we have that C' has $r-1$ faces in its interior, so, by induction, we can assume $W(\overrightarrow{C'},\alpha) = 0$. By Lemma 3.1, $W(\overrightarrow{C},\alpha) = 0$ as well. □

Claim 5.4. *Let G be a 2-connected bipartite planar graph. Suppose that G has a planar embedding with no separating 4-cycle, and suppose further that the embedding has an internal ≥ 6-face, and that the outer face is a ≥ 6-face. Then G is not 3-mixing.*

Proof : We claim that G, under the given assumptions, is foldable to C_6. Denote the internal ≥ 6-face by f, and the outer face by f_0. We call a cycle D in G *f-separating* if f lies inside D, where we include the possibility that edges on the boundary of f lie on the cycle D. (Note that the cycle bounding f_0 is always an f-separating cycle, and thus an f-separating cycle need not be a separating cycle.) Obviously G contains no f-separating 4-cycle, since such a cycle would constitute a separating 4-cycle. We now claim that if G is not a cycle, then it is possible to find a sequence of one or more folds so that the resulting graph is a smaller planar graph that has an internal ≥ 6-face f', whose outer face is a ≥ 6-face, and without an f'-separating 4-cycle. (Note that bipartiteness is trivially maintained by folding.) Repeating such a sequence of folds will eventually transform G into a cycle of length at least six, proving that G is not 3-mixing.

Let C be the cycle that bounds f: we will initially attempt to fold vertices into C. Let x, y, z be three consecutive vertices of C with y having degree at least 3; if there is no such vertex y, then G is simply a cycle of length at least six and we are done. Let a be a neighbour of y distinct from x and z, such that the edges ya and yz form part of the boundary of a face adjacent to f.

Suppose the result of folding a and z introduces no f-separating 4-cycle. If so, we fold a and z. Note that the resulting graph still contains the internal ≥ 6-face f, and is planar since the edges ya and yz form part of a common face. Note also that the outer face, though it might have decreased in size, remains a ≥ 6-face: if it did not (so the edges ya and yz were originally part of the boundary of f_0, which had length six), then we would have a contradiction to the fact that folding a and z introduced no f-separating 4-cycle. We observe that folding a and z might well introduce a cut-vertex into the graph, but that as long as such a vertex is not included twice on the boundary of the outer face, this is not a problem. (Note that such a situation cannot arise for the internal face f.) If we do find that the boundary of the outer face now includes a vertex v twice, then let us denote by G' the graph resulting from folding a and z. Let us also denote by C'_o and C''_o the two distinct cycles formed by the boundary of the outer face, with $V(C'_o) \cap V(C''_o) = \{v\}$, and where $G'_{\mathrm{Int}}(C'_o)$ is the subgraph of G' containing the internal face f (so C'_o must have length at least six, for otherwise we have introduced an f-separating 4-cycle). Now, considering an edge vw of C''_o, we fold $G'_{\mathrm{Int}}(C''_o)$ to vw (using Lemma 5.1 and the fact that vw is a shortest path between v and w). Using this same sequence of folds in G', followed by folding vw into C'_o, leaves us with a graph with the required invariants, and every vertex on the boundary of the outer face of the resulting graph distinct.

Suppose folding a and z *does* result in the creation of an f-separating 4-cycle. If so, this must be because the path a, y, z forms part of an f-separating 6-cycle D. We now show how we can find alternative folds which do not introduce an f-separating 4-cycle. The fact that D is f-separating means there is a path $P \subseteq D$ of length 4 between a and z. Note that P cannot contain y, for this would contradict the fact that G has no f-separating 4-cycle. Consider the graph $G' = G_{\mathrm{Int}}(D) - \{yz\}$. We claim that the path $P' = P \cup \{y\}$ is a shortest path between y and z in G'. To see this, remember that G is bipartite, so any path between y and z in G has to have odd length. We cannot have another edge $yz \in E(G')$ since G is simple. Now note that any path between y and z in G', together with the edge yz, forms an f-separating cycle in G. Hence a path of length 3 between y and z would contradict the fact that G has no f-separating 4-cycle, and so P' is indeed a shortest path between y and z in G'. Using Lemma 5.1, we see that G' is foldable to P'. Using the same sequence of folds in G will fold

$G_{\text{Int}}(D)$ into D. Note this introduces no separating 4-cycle into the resulting graph, and note also that this graph is planar, since it is a subgraph of G. Moreover, note that the length of the cycle bounding the outer face remains the same, that the vertices of this cycle are all distinct, and that the cycle D now bounds an internal 6-face. It follows that this sequence of folds is a sequence as required by the claim. This completes the proof. $\square$

The sequence of Claims 5.1 – 5.4 can easily be used to obtain a polynomial-time algorithm to check if a given planar bipartite graph G is 3-mixing. This completes the proof of Theorem 1.2.

Acknowledgements

We are indebted to Gary MacGillivray for helpful discussions and for bringing reference [11] to our attention.

References

[1] R.J. Baxter, *Exactly Solved Models in Statistical Mechanics*. Academic Press, New York, 1982.

[2] 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/`.

[3] 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.

[4] L. Cereceda, J. van den Heuvel, and M. Johnson, *Connectedness of the graph of vertex-colourings*. Discrete Math. **308** (2008), 913–919.

[5] L. Cereceda, J. van den Heuvel and M. Johnson, *Finding paths between 3-colourings*. CDAM Research Report LSE-CDAM-2007-31 (2007). Available via `www.cdam.lse.ac.uk/Reports/Abstracts/cdam-2007-31.html`.

[6] C.R. Cook and A.B. Evans, *Graph folding*. In : *Proceedings of the Tenth Southeastern Conference on Combinatorics, Graph Theory and Computing*. Congress. Numer. XXIII–XXIV (1979), 305–314.

[7] R. Diestel, *Graph Theory*, 3rd edition. Springer-Verlag, Heidelberg, New-York, 2005.

[8] L.A. Goldberg, R. Martin, and M. Paterson, *Random sampling of 3-colorings in* $\mathbb{Z}^2$. Random Structures Algorithms **24** (2004), 279–302.

[9] 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.

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

[11] N. Vikas, *Computational complexity of compaction to irreflexive cycles*. J. Comput. Syst. Sci. **68** (2004), 473–496.