RECONFIGURATION OF HOMOMORPHISMS TO REFLEXIVE DIGRAPH CYCLES

RICHARD C. BREWSTER, JAE BAEK LEE, AND MARK SIGGERS

ABSTRACT. Let G be a digraph and B be a fixed reflexive digraph cycle. Given two homomorphisms $\phi, \phi' : G \to B$, a walk from ϕ to ϕ' in the Hom-graph $\mathrm{Hom}(G, B)$ corresponds to what is often called a reconfiguration sequence of the homomorphisms. Except in the case that B contains a 4-cycle, containing an oriented cycle of algebraic girth 0, we give a polynomial time algorithm that either finds a path between two given homomorphisms or discovers an obstruction certifying the non-existence of such a path.

1. INTRODUCTION

The problem RECOL(H) of *graph recolouring* has seen recent interest in such papers as [1, 3, 4, 5, 6, 10]. The problem is the following. Given a given pair of homomorphisms from a graph G to a graph H, can one find a sequence of homomorphisms starting at the first and ending at the second in which consecutive homomorphisms differ on only one vertex? This problem can be viewed as a problem of connectivity in the so called Hom-graph, and from this point of view, there are two natural extensions of the problem to the more general setting of digraphs (which we describe below). In [2] we considered one of these extensions, here we consider the other.

For digraphs G and H in which loops are allowed, $\mathrm{Hom}(G, H)$ is the set of homomorphisms from G to H. It causes no confusion to let $\mathrm{Hom}(G, H)$ also denote the *Hom-graph*– the digraph on the vertex set $\mathrm{Hom}(G, H)$ in which (ϕ, ϕ') is an arc if $(\phi(u), \phi'(v))$ is an arc of H whenever (u, v) is an arc of G. If $\phi(v) \neq \phi'(v)$, then we say that v is *recoloured* (under the arc (ϕ, ϕ')) or that we *move* v to $\phi'(v)$.

The *weight* of an arc (ϕ, ϕ') of $\mathrm{Hom}(G, H)$ is the number of vertices on which ϕ and ϕ' differ, and $\mathrm{Hom}_i(G, H)$ is the spanning subgraph of $\mathrm{Hom}(G, H)$ consisting of arcs of weight at most i.

It is not hard to see, that for simple, undirected graphs, the problem of RECOL(H) is simply the problem of deciding, for a given graph G and given pair of maps ϕ and ψ in $\mathrm{Hom}(G, H)$, if there is a walk between ϕ and ψ in $\mathrm{Hom}_1(G, H)$. However, results from [4, 8] show that for undirected graphs (symmetric digraphs) H, there is a walk between two maps in $\mathrm{Hom}_1(G, H)$ if and only if there is a walk between them in $\mathrm{Hom}(G, H)$. Thus, for undirected graphs, RECOL(H) can equivalently be defined as the problem of deciding if there is a walk between maps in $\mathrm{Hom}(G, H)$.

In a sense made explicit in [7], the language of digraph homorphisms is more expressive than that of graph homomorphisms. Hence, we expect to see more varied behaviour in digraph reconfiguration problems than in the graph setting,

Author was supported by NSERC Discovery Grant Program RGPIN-2014-04760.
Author supported by Kyungpook National University Research Fund .

and indeed we do. The problem of deciding if there is a path between a pair of maps ϕ and ψ in $\mathrm{Hom}(G,H)$, can depend on whether we are considering paths $\mathrm{Hom}_1(G,H)$, or in $\mathrm{Hom}(G,H)$. Specifically, it was observed in [9] that when H is a general digraph there can be a walk between vertices in $\mathrm{Hom}(G,H)$ even if there are none in $\mathrm{Hom}_1(G,H)$. Further, in the reflexive setting, 'recolouring a single vertex' and 'adjacency in $\mathrm{Hom}_1(G,H)$' are no longer equivalent.

In [2] we considered the problem of finding a walk in $\mathrm{Hom}_1(G,H)$ between given vertices. In this paper we focus on the more difficult, but in terms of the Hom-graph more natural, problem of finding a walk between given vertices in $\mathrm{Hom}(G,H)$. As in [2], we consider the problem for H being a reflexive digraph cycle. We begin with some definitions.

2. Results

A digraph is *reflexive* (resp. *irreflexive*) if all vertices (resp. no vertices) have a loop. If (u,v) is an arc of a digraph G we write $u \to v$. Varying from standard notation, we refer to an **ordered** pair uv of vertices of G as an *edge*, and write $u \sim v$ if either $u \to v$ or $u \leftarrow v$ (or both). So $u \sim v$ if $v \sim u$. When both $u \to v$ and $u \leftarrow v$ hold, we write $u \leftrightarrow v$. A *walk* W from w_1 to w_ℓ is a sequence $w_1 \sim w_2 \sim \cdots \sim w_\ell$ of consecutively adjacent vertices. It is *directed* if $w_1 \to w_2 \to \cdots \to w_\ell$ and *symmetric.* if $w_1 \leftrightarrow w_2 \leftrightarrow \cdots \leftrightarrow w_\ell$.

A *cycle* $C = v_1 \sim v_2 \sim \cdots \sim v_g \sim v_1$ is a closed walk in which each vertex is distinct. An edge v_iv_{i+1} of C is *forward* if $v_i \to v_{i+1}$ and $v_i \not\leftarrow v_{i+1}$, *backward* if $v_i \leftarrow v_{i+1}$ and $v_i \not\to v_{i+1}$, and *symmetric* if $v_i \leftrightarrow v_{i+1}$. A cycle is *symmetric* if every edge is symmetric, *oriented* if no edge is symmetric, and *directed* if oriented and every edge is forward or every edge is backward.

The *algebraic girth* of an oriented cycle is the absolute value of the number of forward edges less the number of backward edges. Note, for example, that any 4-cycle with at least two symmetric edges contains a 4-cycle of algebraic girth 0.

The *homomorphism reconfiguration problem* is defined for a fixed digraph H as follows.

Recon(H)
Instance: (G,ϕ,ψ): a digraph G, and two homomorphisms $\phi,\psi \in \mathrm{Hom}(G,H)$.
Question: Is there a walk from ϕ to ψ in $\mathrm{Hom}(G,H)$?

Our results are for reflexive digraph targets H. Though an instance G of Recon(H) need not be reflexive, it is natural to consider this restriction. Indeed our main results are simpler to state (and prove) with this restriction added. Formally, we define the following variation on the above problem. Let H be a fixed, reflexive digraph.

Recon^{r}(H)
Instance: (G,ϕ,ψ): a reflexive digraph G, and two homomorphisms $\phi,\psi \in \mathrm{Hom}(G,H)$.
Question: Is there a walk from ϕ to ψ in $\mathrm{Hom}(G,H)$?

In this paper, we prove the following results. The analogous results were shown with Recol in place of Recon in [2]. The proofs here require several new tools generalizing those in [2].

Theorem 2.1. *Let B be a fixed reflexive cycle. Then the following problems are polynomial time solvable.*

(i) $\textsc{Recon}^{\mathrm{r}}(B)$ *and,*
(ii) $\textsc{Recon}(B)$ *provided* B *does not contain a* 4*-cycle of algebraic girth* 0*.*

It is known from [10] that $\textsc{Recol}(B)$ and so $\textsc{Recon}(B)$ is PSPACE-complete when B is a reflexive symmetric 4-cycle. The complexity of both $\textsc{Recol}(B)$ and $\textsc{Recon}(B)$ is unknown for other reflexive 4-cycles containing a 4-cycle of algebraic girth 0.

Of course, one can ask about finding walks between maps in $\mathrm{Hom}_k(G,B)$ for any k. Let $\textsc{Recon}_{\mathrm{k}}(B)$ be the problem of deciding for an instance (G,ϕ,ψ) of $\textsc{Recon}(B)$ if there is a walk from ϕ to ψ in $\mathrm{Hom}_k(G,B)$. As it adds almost no difficulty to our proofs, we prove the following result, which gives the main result of [2] by taking $k = 1$.

Theorem 2.2. *Let* B *be a fixed reflexive cycle. Then for any* $k \geq 1$*, the following problems are polynomial time solvable.*

(i) $\textsc{Recon}^{\mathrm{r}}_{\mathrm{k}}(B)$*, and*
(ii) $\textsc{Recon}_{\mathrm{k}}(B)$*, provided* B *does not contain a reflexive* 4*-cycle of algebraic girth* 0*.*

After proving Theorem 2.2, we will observe that the algorithm we give has polynomial running time independent of k, and so yields a proof of Theorem 2.1 also.

3. Outline of the proof and Initial reductions

In this section we outline the proof, introducing definitions and results from [2] that we will need. Henceforth, $B = 1 \sim 2 \sim \cdots \sim b \sim 1$ shall be a reflexive cycle on the integers modulo b. This gives an *underlying orientation* to B that is independent of the orientation of its arcs. Occasionally we will 'reverse the underlying orientation' of B in definitions or arguments, to avoid repeating them. By this we mean that we relabel the vertex i with $b - i + 1$. Further, G will henceforth be a digraph and (G,ϕ,ψ) will be an instance of $\textsc{Recon}_{\mathrm{k}}(B)$ or $\textsc{Recon}^{\mathrm{r}}_{\mathrm{k}}(B)$. We may clearly assume that G is connected. We start with reductions that allow us to focus on the main case of Theorem 2.2.

In a reflexive digraph H, a vertex u *dismantles forward* to a vertex v if the in-neighbourhood of v contains the in-neighbourhood of u; as H is reflexive, this impies that $u \to v$. It *dismantles backward* to v if the out-neighbourhood of v contains that of u. Observe that if u dismantles forward to v then for any $\phi \in \mathrm{Hom}(G,B)$ we get that $\phi \to \phi_{v\mapsto u}$ where

$$\phi_{v\mapsto u}(x) = \begin{cases} v & \text{if } \phi(x) = u \\ \phi(x) & \text{otherwise.} \end{cases}$$

Thus we can decide if there is a walk from ϕ to ψ in $\mathrm{Hom}(G,H)$ by deciding if there is a walk from $\phi_{v\mapsto u}$ to $\psi_{v\mapsto u}$ in $\mathrm{Hom}(G,H)$.

Fact 3.1. *Let* B *contains vertices* u *and* v *such that* u *dismantles forward to* v*. The problem* $\textsc{Recon}(B)$ *polynomial time solvable if* $\textsc{Recon}(B \setminus \{u\})$ *is. Moreover, if* u *also dismantles backward to* v *then* $\textsc{Recon}_{\mathrm{k}}(B)$ *is polynomial time solvable if* $\textsc{Recon}_{\mathrm{k}}(B \setminus \{u\})$ *is.*

As $\mathrm{Hom}_k(G,C_1)$ is clearly connected for the reflexive 1-cycle C_1 this shows that $\mathrm{Hom}_k(G,C_2)$ is also connected for the reflexive 2-cycle C_2; thus Theorem 2.2 holds when B is a 1-cycle or a 2-cycle. The above argument is basically from [2] where

we used it to claim also that $\textsc{Recon}_k(B)$ and $\textsc{Recon}^r_k(B)$ are trivial when B is a transitive tournament or contains a transitive triangle. Our proof was mistaken, it only showed that $\textsc{Recon}(B)$ and $\textsc{Recon}^r(B)$ are trivial in these cases. We did not notice that the edge $\phi \to \phi'$ we get from the forward-dismantling might have weight more than k. In fact $\textsc{Recon}_k(B)$ and $\textsc{Recon}^r_k(B)$ are not trivial for such B, but they are linear time solvable. We rectify this now.

Lemma 3.2. *Let B be a reflexive triangle containing the transitive triangle, then Theorem 2.2 holds.*

Proof. We first observe that $\textsc{Recon}_k(T_n)$ is linear-time solvable for the reflexive transitive tournament T_n on the vertex set $[n]$ in which $i \to j$ if $i \le j$. Indeed, let G be a connected instance, and order its strong components by setting $C \le C'$ if there is some arc from C to C'. It is clear that a vertex map ϕ of G to T_n is a homomorphism if and only if

(i) ϕ is constant on strong components of G, (we let $\phi(C)$ denote the image $\phi(v)$ of all $v \in V(C)$), and
(ii) it preserves the order of strong components: if $C \le C'$, then $\phi(C) \le \phi(C')$.

Furthermore, for two such homomorphisms ϕ and ϕ', it is clear that $\phi \to \phi'$ in $\mathrm{Hom}_k(G, T_n)$ if and only if the following hold:

(i) ϕ and ϕ' differ on at most k vertices, and
(ii) for each recoloured vertex v, $\phi(u) \le \phi'(v)$ for all $u \in N^-(v)$.

Note the second condition is satisfied when $\phi(v) \le \phi'(v)$ for all v. In particular, strong components C of G of order greater than k are *fixed*, which means that for any (ϕ, ϕ') in $\mathrm{Hom}_k(G, T_n)$ we have $\phi(C) = \phi'(C)$. Every component of $\mathrm{Hom}_k(G, T_n)$ thus has a canonical sink that we can reach from any map ϕ in the component by mapping any component of G of size at most k forward to $\phi(C')$ where C' is the least fixed strong component of G greater than C in the component order. In the absence of such a C' we set $\phi(C) = n$. (The order in which we move components is selected by topologically sorting the components that need to move.) There is a path between ϕ and ϕ' if and only if they have this same canonical sink. If follows that there is a path between ϕ and ϕ' if and only if they map fixed strong components to the same vertices. Strong components can be found by a tree search, so this can be determined in linear time.

Now, let B be a triangle properly containing a T_3, so having at least one symmetric edge. If the edge 12 is symmetric, the the vertex 1 dismantles forward and backward to 2, so by Fact 3.1 the problem reduces to that for $B \setminus \{1\}$, which is either T_2 or K_2, so polynomial time solvable A similar argument works if 23 is symmetric in B. So we may assume that neither 12 or 23 are symmetric, and so the only symmetric edge is 13. To see that $\textsc{Recon}_k(B)$ is polynomial in this case, we observe that two maps are in the same component if the set of strong components of size greater than k that map to the vertex 2 are the same. (As these components cannot move from 2, this is necessary. A component C that sends an edge to a fixed component C' with $\phi(C') = 2$ must have each of $\phi(C)$ equal to either 1 or 2. If C is a non-fixed strong component, then it can move to 2. If it is in a fixed strong component, then it must stay where it is mapped. Finally, all other vertices can be moved one at a time, forward to 3 even from strong components of size exceeding k. This defines a unique canonical form that depends only on the fixed components mapping to 2.) □

We may therefore restrict ourselves to the following situation.

Definition 3.3. The pair (G, B) is a $\mathcal{WP}$-pair if

- B is a reflexive cycle of girth at least 3 not containing a transitive triangle, and G is reflexive; or
- B is a reflexive cycle of girth at least 4 not containing a 4-cycle of algebraic girth 0.

Theorem 2.2 is thus an immediate corollary of the following theorem, the proof of which is the focus of the remainder of the paper.

Theorem 3.4. *Let (G, B) be a $\mathcal{WP}$-pair. Let $\phi, \psi \in \mathrm{Hom}(G, B)$. Then there is a polynomial time algorithm to decide if there is a walk in $\mathrm{Hom}_k(G, B)$ between ϕ and ψ.*

The property that differentiates $\mathcal{WP}$-pairs from the trivial cases mentioned above, and from the case where G is not reflexive and B contains a C_4 of algebraic girth 0, is that for $\mathcal{WP}$-pairs we have a topological notion of preservation of wind. (Hence the notation, we think of them as $\mathcal{W}$ind $\mathcal{P}$reserving pairs.) What this means is that for a closed walk C in G, a homomorphism ϕ in $\mathrm{Hom}(G, B)$ winds C around B some well defined number of times. For $\phi' \sim \phi$ in $\mathrm{Hom}(G, B)$, this wind is preserved, and so the wind of cycles is be preserved over components of $\mathrm{Hom}_k(G, B)$.

Let us define this concept more formally. The *net-length* of an edge uv of B is $\mathrm{nl}(uv) = v - u$. Note the orientation of uv in B does not affect net-length, though $\mathrm{nl}(uv) = -\mathrm{nl}(vu)$. We reduce modulo b so that $\mathrm{nl}(uv)$ is $1, 0$, or -1. The *net-length* of a walk $W = w_0, w_1, \ldots, w_n$ of B is defined as $\sum_{i=1}^{n} \mathrm{nl}(w_{i-1}w_i)$. For a closed walk C of B we clearly have that $\mathrm{nl}(C) = b \cdot w(C)$ for some unique integer $w(C)$ called the *wind* of C. Given any closed walk C and a homomorphism $\phi : C \to B$, $\phi(C)$ is a closed walk in B; we let $w_\phi(C) := w(\phi(C))$. We showed in [2] that if there is a walk between ϕ and ψ in $\mathrm{Hom}_k(G, B)$, for a $\mathcal{WP}$-pair (G, B), then $w_\phi(C) = w_\psi(C)$ for every cycle C in G.

A key step in our algorithm is to determine whether or not $w_\phi(C) = w_\psi(C)$ for every cycle C in G. The following ideas are used in the algorithm. For a walk

$$\mathcal{W} = \phi_1 \sim \phi_2 \sim \cdots \sim \phi_\ell$$

from $\phi = \phi_1$ to $\psi = \phi_\ell$ in $\mathrm{Hom}_k(G, B)$ and a vertex v in G, let $\mathcal{W}(v)$ be the sequence $\phi_1(v), \phi_2(v), \ldots, \phi_\ell(v)$. We would like the sequence of vertices to unambiguously define a walk in B, however images of v may jump between non-adjacent vertices in B. Formally, a vertex v of G *jumps* with respect to (or under) an edge $\gamma\gamma'$ of $\mathrm{Hom}_k(G, B)$ if $\gamma'(v) \notin \{\gamma(v) - 1, \gamma(v), \gamma(v) + 1\}$. (We slightly generalise the definition of 'jump' later in the case that B has girth 3.) The edge is *non-jumping* if no vertex jumps with respect to it. In Section 4 we prove Lemma 4.3 which allows us to consider only edges of $\mathrm{Hom}_k(G, B)$ that are non-jumping. Consequently, $\mathcal{W}(v)$ defines the walk

$$\phi_1(v) \sim \phi_2(v) \sim \cdots \sim \phi_\ell(v)$$

in B. Let $\delta_\mathcal{W} : V(G) \to \mathbb{Z}$ be defined by $\delta_\mathcal{W}(v) = \mathrm{nl}(\mathcal{W}(v))$. The value $\delta_\mathcal{W}(v)$ determines how much the image of v moves under the sequence of homomorphisms from ϕ_1 to ϕ_ℓ. When $\delta_\mathcal{W}(v) > 0$ ($\delta_\mathcal{W}(v) < 0$) we say v moves up (down).

It is easy to see, and is shown in [2], that for a walk $\mathcal{W}$ from ϕ to ψ, $\delta_\mathcal{W}$ satisfies the following conditions, so is a topologically valid difference function for ϕ and ψ.

Definition 3.5. A function $\delta : V(G) \to \mathbb{Z}$ is is a *difference function* for ϕ and ψ in $\mathrm{Hom}_k(G, B)$ if for all $v \in V(G)$,

$$\delta(v) \equiv \psi(v) - \phi(v) \mod b$$

where b is the girth of B. It is *topologically valid* if for any two vertices $v, w \in V(G)$, and any walk W from v to w, we have

$$\delta(w) = \mathrm{nl}_\psi(W) + \delta(v) - \mathrm{nl}_\phi(W). \tag{1}$$

It is shown in [2] that there is a topologically valid difference function for ϕ and ψ if and only if $w_\phi(C) = w_\psi(C)$ for every cycle in C. It is also shown that the existence of such a function can be determined in polynomial time. The following comes immediately from Corollary 2.4 and Lemma 4.3 of [2].

Lemma 3.6. [2] *For a $\mathcal{WP}$-pair (G, B), if there is a walk between ϕ and ψ in $\mathrm{Hom}_k(G, B)$ then there is a topologically valid difference function δ for ϕ and ψ. This δ is unique up to a constant shift by a multiple of $b = |V(B)|$, and its existence can be determined in time polynomial in the size of G.*

That we have a topologically valid difference function is not enough to assure that there is a walk from ϕ to ψ in $\mathrm{Hom}(G, B)$. One simple example of this is when G and B are both the symmetric reflexive 5-cycle C_5, and ϕ and ψ are different wind one maps, e.g. $\phi(i) = i$ and $\psi(i) = i + 1 \bmod 5$. It is easy to see that both ϕ and ψ are isolated vertices in $\mathrm{Hom}_k(G, B)$.

Definition 3.7. A topologically valid difference function δ for ϕ and ψ is *realisable* if there is a walk $\mathcal{W}$ from ϕ to ψ in $\mathrm{Hom}_k(G, B)$ with $\delta_\mathcal{W} = \delta$.

So our goal is to determine the existence of a realisable topologically valid difference function for ϕ and ψ.

In Sections 5, 6, and 7, we develop ideas necessary to prove the correctness of an algorithm that will test if a given difference function is realisable. With these ideas we also prove Proposition 7.8 in Section 7, which says that if there is a topologically valid difference function for ϕ and ψ, then there is one such that the maximum value of $|\delta(v)|$ is at most $2n + b$.

In Section 8 we give Algorithm 8.1 which decides, in time polynomial in $n = |V(G)|$, if a given $(2n+b)$-bounded function δ which is topologically valid for ϕ and ψ, is realisable. The difference function δ tells us which vertices should move up, and which should move down. The main idea behind the Algorithm 8.1 is that if δ is realisable, then we can find a path realising it by greedily moving up vertices that it tells us should be moved up, and then moving down vertices that it tells us should be moved down.

Putting together Lemma 3.6, Proposition 7.8, and Proposition 8.2 in the obvious way, we then finish off Section 8 with the formal proof Theorem 3.4, and so get Theorem 2.2 also.

In so much as we have outlined it here, our proof is very similar to the proof in [2] which in turn borrows many ideas from [10]. The details however are very different. Analogues of Sections 5 and 6 are unnecessary in [2], and the analogous results to those in Section 4 are much easier to prove in the previous work. On the other hand, our main algorithm is conceptually simplified here by allowing a slightly longer running time.

4. Proof of the No-jumping Lemma

In this section, we prove the No-jumping Lemma. Before we do this, we look in more detail at the notion of jumping.

Earlier, we defined a vertex v as jumping under an edge $\phi\phi'$ of $\mathrm{Hom}_k(G,B)$ if $\phi'(v) \notin \{\phi(v)-1, \phi(v), \phi(v)+1\}$. However, we will need to distinguish 'up' jumps and 'down' jumps. Observe that if v has a loop, then $\phi(v) \sim \phi'(v)$, so no jumping vertex has a loop. Moreover, by the assumption that G is connected, every jumping vertex v must have a neighbour.

Consider a vertex v and an arc (ϕ,ϕ') of $\mathrm{Hom}_k(G,B)$. For any in-neighbour u, $u \to v$ implies that

$$\phi(v) \leftarrow \phi(u) \to \phi'(v) \tag{2}$$

and for any out-neighbour w, $v \to w$ implies that

$$\phi(v) \to \phi'(w) \leftarrow \phi'(v). \tag{3}$$

When B has girth at least 5, we therefore get that v jumps under an arc (ϕ,ϕ') exactly if $\phi'(v) = \phi(v) \pm 2$. We will say it *jumps up* if $\phi'(v) = \phi(v)+2$, and *jumps down* if $\phi'(v) = \phi(v)-2$. When B has smaller girth, the definition of up and down jumps is not so clear. In fact, by our definition of 'jump' there are no jumping vertices when B has girth 3. It may appear we may need to distinguish jumping up two versus moving down one in this case, but we will show that when the girth is 3, there are no jumping vertices in the sense defined below. In our main algorithm we will have a vertex v mapped to B that must move up, and we establish that it can move up one vertex at time, rather than jumping up two vertices. The definition we want for 'jumps up' is clear in the case that v is the only vertex moving under (ϕ,ϕ'): if v jumps up, then any neighbour is above it before the jump and is below it after the jump. (It jumps over the fixed neighbour.) The following is the extension of this idea to the present situation where several vertices of G may move under one arc of $\mathrm{Hom}_k(G,B)$.

Definition 4.1. Let $\phi \to \phi'$ in $\mathrm{Hom}_k(G,B)$. A vertex v *jumps up* under $\phi\phi'$ if all of the following hold:

(a) $\phi'(v) = \phi(v)+2$.
(b) For all u with $u \to v$, $\phi(u) = \phi(v)+1$.
(c) For all v with $v \to w$, $\phi'(w) = \phi'(v)-1$.

A vertex *jumps down* under $\phi\phi'$ if it jumps up when reversing the underlying orientation of B. For $\phi \leftarrow \phi'$ in $\mathrm{Hom}_k(G,B)$, v *jumps up* under $\phi\phi'$ if it jumps down under $\phi'\phi$, and so we reverse the arrows in (b) and (c). A vertex *jumps up* under a symmetric edge $\phi\phi'$ if it satisfies either (and so both) of the above conditions. A vertex *jumps* under an edge if it jumps up or down under the edge.

Now, observe that if B has girth 3 and a vertex v with an in-neighbour u jumps up, under an edge $\phi\phi'$ with $\phi \to \phi'$, then assuming without loss of generality that $\phi(v) = 1$, we would have

$$1 = \phi(v) \leftarrow \phi(u) = 2 \to \phi'(v) = 3,$$

and so B would contain a transitive triangle. A similar argument shows that a vertex v with an out-neighbour w cannot jump up when B has girth 3. It follows that when B has girth 3, vertices do not jump. So when G is a connected graph

with at least two vertices, this definition is the same as our our original definition of 'jump'. Also, clearly a vertex cannot jump both up and down.

In our arguments, we will often want to assert that a vertex jumps up. When B has girth at least 5, verifying Def. 4.1(a) is enough. But when B has girth 4 it would seem that we have to check Def. 4.1(b) and (c). The following shows that we do not.

Fact 4.2. *Let v be a jumping vertex under (ϕ, ϕ') for which Def. 4.1(a) holds. If there is some u with $u \to v$, such that $\phi(u) = \phi(v) + 1$, or there is some w with $v \to w$, such that $\phi'(w) = \phi'(v) - 1$, then v jumps up.*

Proof. In the case that B has girth greater than 4, the result is trivial as Def. 4.1(a), with equations (2) and (3), implies Def. 4.1(b) and (c).

When B has girth 4 however, it is possible that some neighbours map to $\phi(v)+1$ and some to $\phi(v)-1$. If v has two neighbours (both in or both out or one of each) and one maps to $\phi(v)+1$ while the other maps to $\phi(v)-1$, their image in B makes a 4-cycle of algebraic girth 0. This contradicts our assumptions. Thus all such neighbours map to the same place. If one of them satisfies Def. 4.1(b) or (c) then they all do, and so v jumps up. □

We are now ready to state and prove the main result of this section.

Lemma 4.3 (No-jumping Lemma)**.** *For a $\mathcal{WP}$ pair (G, B) such that $\phi \to \phi'$ in $\mathrm{Hom}_k(G, B)$, there is a directed path from ϕ to ϕ' in $\mathrm{Hom}_k(G, B)$ in which no vertex jumps under any of the arcs in the path.*

Proof. We say that (ϕ, ϕ') has a *good decomposition* if there is a path from ϕ to ϕ' in $\mathrm{Hom}_k(G, B)$ such that each arc in the path has fewer jumping vertices than (ϕ, ϕ'). We therefore assume that (ϕ, ϕ') has jumping vertices, and show that it has a good decomposition. We start with a series of assumptions for which the following notation will be useful. For a vertex v of G, let ϕ_v be the map we get from ϕ by increasing the image of v by 1, and let ϕ'_{-v} be the map we get from ϕ' by reducing the image of v by 1.

Let v_0 be one of the jumping vertices. Clearly we may assume, possibly by reversing the underlying orientation of B, that v_0 jumps up. This next claim allows us to assume that any vertex that jumps up has an in-neighbour.

Claim 4.4. *If v jumps up, then we may assume that $\phi(v) \leftarrow \phi(v) + 1$.*

Proof. Suppose v jumps up. If $\phi(v) \not\leftarrow \phi(v) + 1$, then by Def. 4.1(b), v has no in-neighbours. Hence ϕ' maps all out-neighbours of v, and so all neighbours of v, to $\phi'(v) - 1$ by Def. 4.1(c). It is clear that the map ϕ'_{-v} is a homomorphism. This implies that $\phi \to \phi'_{-v} \to \phi'$ is a good decomposition as required. ⋄

Claim 4.5. *If v jumps up, then we may assume that it has an out-neighbour w with $\phi(w) = \phi(v) - 1$.*

Proof. Indeed, if v jumps up and this is not true, then for any out-neighbour w of v (if it exists) we have $\phi(w) \in \{\phi(v), \phi(v) + 1\}$. As v jumps up, any in-neighbour u has $\phi(u) = \phi(v) + 1$ by Def. 4.1(b). So the only neighbours of v not mapped to $\phi(v) + 1$ by ϕ (and so ϕ_v), are out-neighbours mapped to $\phi(v)$. By Claim 4.4, ϕ_v is a homomorphism. To see that $\phi \to \phi_v$ it is enough to check that $\phi(u) \to \phi_v(v)$ for any u with $u \to v$. But this is clear as $\phi(u) = \phi(v) + 1 = \phi_v(v)$ in this case.

As $\phi \to \phi'$, and ϕ and ϕ_v only differ on v, to see that $\phi_v \to \phi'$ holds, it is enough to check that $\phi(v)+1 = \phi_v(v) \to \phi'(w)$ holds for any out-neighbour w of v. But by Def. 4.1(c), $\phi'(w) = \phi'(v) - 1 = \phi(v) + 1$ as v jumps up, so this holds. Again $\phi \to \phi_v \to \phi'$ is a good decomposition as required. ⋄

We have reduced the problem to allow us the following assumptions. There is some vertex v_0 that jumps up. For any vertex v that jumps up, $\phi(v) \leftarrow \phi(v)+1$ and there is an out-neighbour w with $\phi(w) = \phi(v) - 1$. By Def. 4.1(c) we have that $\phi'(w) = \phi'(v) - 1$ and so by Fact 4.2, w also jumps up. Let J' be the set of vertices of G that jump up, and let A be the graph on J' we get by setting $v \to w$ in A if $v \to w$ in G and $\phi(w) = \phi(v) - 1$. By Claim 4.5, A has no sinks, so contains some *terminal set* J: a strong component with no out-neighbours. Let ϕ_J be the map we get from ϕ by increasing the image of every vertex in J by 1.

The following claim finishes the proof of the lemma, where ϕ_J is the map we get from ϕ by moving all vertices in J up by one.

Claim 4.6. *$\phi \to \phi_J \to \phi'$ is a good decomposition.*

Proof. First we show that B is a reflexive symmetric cycle. Indeed, as $v \to w$ in A implies that $\phi(w) = \phi(v) - 1$, the fact that J is a non-trivial strongly connected component of A means that ϕ restricted to J is onto B. For a vertex i of B, we have a vertex v such that $\phi(v) = i$. By the definition of J, there is a vertex w such that $v \to w$ and $i = \phi(v) \to \phi(w) = i - 1$. Since v jumps up, $i = \phi(v) \to \phi'(w) = i + 1$.

We now show that ϕ_J is a homomorphism. Indeed, let $u \to v$ in G. If u and v are both in J, or both not in J, then that $\phi_J(u) \to \phi_J(v)$ is trivial by the facts that ϕ is a homomorphism and B is symmetric. If u is in J but v is not, then $u \not\to v$ in A, and so by the definition of A, we have that $\phi_J(v) = \phi(v) \in \{\phi(u), \phi(u)+1\}$. As B is reflexive and symmetric $\phi_J(u) = \phi(u)+1$ has an arc to both of these. If v is in J but u is not, then $\phi_J(u) = \phi(u) = \phi(v) + 1$ as v jumps up, and $\phi_J(v) = \phi(v) + 1$, so we are done as B is reflexive.

To see that $\phi \to \phi_J$, it is enough to check, when $u \to v$ and $v \in J$, that $\phi(u) \to \phi_J(v)$. But $\phi(u) = \phi(v) + 1 = \phi_J(v)$.

To see that $\phi_J \to \phi'$, it is enough to check, when $v \to w$ and $v \in J$, that $\phi_J(v) \to \phi'(w)$. But both of these are $\phi(v) + 1$. ⋄

This completes the proof of the Lemma. □

Remark 4.7. As we are only interested in the existence of walks between maps, we will henceforth apply Lemma 4.3 implicitly, and assume that $\mathrm{Hom}_k(G, B)$ contains only non-jumping arcs. In particular given a walk $\mathcal{W}$ in $\mathrm{Hom}_k(G, B)$, we implicitly assume that it is a walk of non-jumping edges, and so $\delta_{\mathcal{W}}$, as defined in Section 3, is well-defined.

5. Indecomposable arcs

An arc $(\phi, \psi) \in \mathrm{Hom}_k(G, H)$ is *decomposable* if there are arcs (ϕ, γ) and (γ, ψ) of strictly smaller weight. It is *indecomposable* if is not decomposable. It is immediate from these definitions that there is a walk between ϕ and ψ in $\mathrm{Hom}_k(G, H)$ if and only if there is a walk of indecomposable arcs between them; so our task becomes finding a walk of indecomposable arcs between ϕ and ψ.

The following is useful in determining how an arc can be decomposed.

Lemma 5.1. *For digraphs G and H, let $\phi \to \phi'$ in $\mathrm{Hom}_k(G,H)$ and let X be the set of vertices x of G for which $\phi(x) \neq \phi'(x)$. Let $Y \subset X$ and let $\phi_Y : V(G) \to V(H)$ be defined by*

$$\phi_Y(y) = \begin{cases} \phi'(y) & \text{if } y \in Y \\ \phi(y) & \text{otherwise.} \end{cases}$$

Then $\phi \to \phi_Y \to \phi'$ in $\mathrm{Hom}(G,H)$ if and only if $\phi_Y \in \mathrm{Hom}(G,H)$.

Proof. It is clear that ϕ_Y must be in $\mathrm{Hom}(G,H)$ to have $\phi \to \phi_Y \to \phi'$ in $\mathrm{Hom}(G,H)$. For the other direction, assume that ϕ_Y is in $\mathrm{Hom}(G,H)$. We show that $\phi_Y \to \phi'$. The proof that $\phi \to \phi_Y$ is similar.

Let $u \to v$ in G, we must show that $\phi_Y(u) \to \phi'(v)$. If $u \in Y$, then $\phi_Y(u) = \phi'(u)$, and so $\phi_Y(u) \to \phi'(v)$ as ϕ' is a homomorphism. If $u \notin Y$ then $\phi_Y(u) = \phi(u)$, and so it holds as $\phi \to \phi'$. □

A vertex v of G *moves up* with respect to an arc (γ, γ') of $\mathrm{Hom}_k(G,B)$ if $\mathrm{nl}(\gamma(v), \gamma'(v)) = 1$. Given a single map γ in $\mathrm{Hom}(G,B)$, a vertex v *moves up* (from γ) if there is an arc (γ, γ') with respect to which v moves up. We similarly define moving down.

The following lemma allows us to define some useful notation.

Lemma 5.2. *Let (G,B) be a $\mathcal{WP}$ pair and $\phi \to \phi'$ be an indecomposable arc in $\mathrm{Hom}(G,B)$. Then for all vertices that move, they either all move up, or they all move down.*

Proof. If some vertices U move up and some others D move down with respect to an arc (ϕ, ϕ'), we show that ϕ_U, as defined in Lemma 5.1 is a homomorphism. It then follows from that lemma that (ϕ, ϕ') can be decomposed.

Now we verify ϕ_U is a homomorphism: for (u,v) in G we clearly have $\phi_U(u) \to \phi_U(v)$ if u and v are either both in U or both not in U. If $u \notin U$ but $v \in U$ then

$$\phi_U(u) = \phi(u) \to \phi'(v) = \phi_U(v).$$

So assume that $u \in U$ but $v \notin U$. So $\phi_U(u) = \phi'(u)$ and $\phi_U(v) = \phi(v)$. We must show that $\phi'(u) \to \phi(v)$. If $v \in D$, then we reason as follows. By our hypothesis we have $\phi(u) \to \phi'(v)$ and $\phi'(u) \to \phi'(v)$. If the three images in B are distinct, then there is a transitive triangle in B. Similarly, if $\phi'(u) = \phi'(v)$, then $\phi(u), \phi(v), \phi'(u)$ is a transitive triangle. As B has no transitive triangles, we conclude $\phi(u) = \phi'(v)$ which implies $\phi'(u) = \phi(v)$, and we are done. If $v \notin D \cup U$, then $\phi(v) = \phi'(v) \leftarrow \phi'(u)$. □

Because of this lemma, we can represent any indecomposable arc (γ, γ') by $\uparrow T$ if we get γ' from γ by moving the vertices of the set T on which they differ up, or by $\downarrow T$ if we move them down. So a walk of indecomposable arcs can be represented by a sequence of symbols

$$\square V_1, \square V_2, \ldots, \square V_\ell$$

where each $\square$ is either $\uparrow$ or $\downarrow$. If this walk is a walk from ϕ to ψ, we can represent this by writing

$$\psi = \phi + \square V_1 + \square V_2 + \cdots + \square V_\ell.$$

6. Minimal up-sets and up obstructions

Given a homomorphism γ in $\mathrm{Hom}(G,B)$, a set $T \subseteq V(G)$ *moves up* for γ if $\gamma \sim (\gamma + \uparrow T)$ in $\mathrm{Hom}_k(G,B)$; it is a *minimal up-set* if no proper subset $T' \subset T$ moves up. The following is immediate from Lemma 5.1.

Fact 6.1. *If $\gamma \sim \gamma + \uparrow T$ is indecomposable, then T is a minimal up-set for γ.*

In this section we characterise minimal up-sets and show that for a given homomorphism γ one can find all minimal up-sets in polynomial time. Then in the next section, we will show the converse of Fact 6.1, characterising indecomposable arcs.

We begin by recalling a construction from [2]. Let γ be a map in $\mathrm{Hom}(G,B)$, and $\gamma|_{uv}$ be its restriction to an edge uv of G. Consider the situation that the map

$$\gamma|_{uv} + \uparrow u \text{ is not a homomorphism.} \tag{4}$$

Clearly in this situation, if u is in a set T that moves up from γ, then v must also be in T. This suggests the following definition from [2].

Definition 6.2. For $\gamma \in \mathrm{Hom}(G,B)$ let the *auxillary graph* $\mathbb{A}^+_\gamma$ have vertex set $V(G)$ and for an edge uv of G let (u,v) be in $\mathbb{A}^+_\gamma$ if u and v satisfy (4).

The arc (u,v) in $\mathbb{A}^+_\gamma$ means u cannot move up unless v moves up (first or simultaneously). In [2] we used the clear fact that $\gamma \sim \gamma + \uparrow u$ is an edge of $\mathrm{Hom}_1(G,B)$ if and only if u it is a sink of $\mathbb{A}^+_\gamma$. To characterise the sets T that move up, the construction $\mathbb{A}^+_\gamma$ remains useful, but we must look a little deeper. For $T \subseteq V(G)$, we denote by $G[T]$ the subgraph of G induced by T.

Definition 6.3. A set $T \subseteq V(G)$ is an *up-obstruction* for γ if $\gamma|_T \sim (\gamma|_T + \uparrow T)$ is not an edge of $\mathrm{Hom}_k(G[T],B)$.

It is important that T being an up-obstruction is internal to T, that is, it is different than saying that $\gamma \sim (\gamma + \uparrow T)$ is not an edge of $\mathrm{Hom}_k(G,B)$. But if T is a terminal strong component, these definitions are the same, as the following shows.

Lemma 6.4. *A set $T \subset V(G)$ moves up for γ if and only if it is a terminal set of $\mathbb{A}^+_\gamma$, but not an up-obstruction.*

Proof. It is clear that if T moves up for γ then it must be a terminal set of $\mathbb{A}^+_\gamma$. Indeed, if some $u \in T$ has an arc in $\mathbb{A}^+_\gamma$ to some $v \notin T$, then there is an edge in G between u and v. The restriction of $\gamma + \uparrow T$ to this edge is $\gamma|_{uv} + \uparrow u$, which by the definition of $\mathbb{A}^+_\gamma$ is not a homomorphism. It is also clear that the restriction of γ and $\gamma + \uparrow T$ to T is also a homomorphism, and so T is not an up-obstruction.

On the other hand, let T be a terminal set of $\mathbb{A}^+_\gamma$, and let $\gamma|_T \to \gamma|_T + \uparrow T$ in $\mathrm{Hom}_k(G[T],B)$. Let $\gamma' = \gamma + \uparrow T$. We must check for any arc (u,v) of G that

(i) $\gamma(u) \to \gamma'(v)$ and
(ii) $\gamma'(u) \to \gamma'(v)$.

This clearly holds if $u, v \notin T$, since then $\gamma' = \gamma$ on u and v, and γ is a homomorphism. If $u, v \in T$ we get (i) from the fact that $\gamma|_T \to (\gamma|_T + \uparrow T)$ and (ii) from the fact that $\gamma|_T + \uparrow T$ is a homomorphism. If $u \in T$ and $v \notin T$, then (i) is $\gamma(u) \to \gamma(v)$ which is true as γ is a homomorphism, and (ii) is $\gamma(u) + 1 \to \gamma(v)$ which is true as the arc (u,v) of G is not in $\mathbb{A}^+_\gamma$. If $u \notin T$ and $v \in T$, then (i) is $\gamma(v) + 1 \leftarrow \gamma(u)$ which is true as $v \not\to u$ in $\mathbb{A}^+_\gamma$, and (ii) follows from (i). □

As a terminal set of vertices in a digraph is minimally terminal if and only if it is a terminal strong component, we immediately get the the following.

Proposition 6.5. *A set $T \subset V(G)$ is a minimal up-set for γ if and only if it is a terminal strong component of $\mathbb{A}^+_\gamma$ and is not an up-obstruction.*

Remark 6.6. Clearly a terminal strong component of size greater than k is an up-obstruction, however, there are other up-obstructions too. One can give an explicit description of them but it requires more definitions, and the proof is tedious. Luckily, we were able to remove the necessity of an explicit description of up-obstructions from our proof. This allows an easier 'removal' of the parameter k at the end of the paper. Moreover, it allows us, at the end of the paper, to easily settle the problem of finding a directed path in $\mathrm{Hom}(G, B)$ between two homomorphisms.

Clearly one can construct $\mathbb{A}^+_\gamma$ from γ in polynomial time, one can find all its terminal strong components T in polynomial time, and one can check, in polynomial time, whether or not any of them are up-sets by checking if $\gamma|_T \to (\gamma|_T + \uparrow T)$. This gives the following.

Corollary 6.7. *There is a polynomial function $f(n)$, independent of k, where $n = |V(G)|$, such that for a given a map γ in $\mathrm{Hom}_k(G, B)$, one can find all minimal up-sets of γ in time $O(f(n))$.*

Remark 6.8. For a homomorphism $\gamma : G \to B$, defining $\mathbb{A}^-_\gamma$ as the graph $\mathbb{A}^+_\gamma$ that we get by reversing the underlying orientation of B. We get 'negative' or 'down-set' analogues of Proposition 6.5 and Corollary 6.7.

7. Standard Form of a walk and a bound on shift

Having found a difference function δ that is topologically valid for two maps ϕ and ψ in $\mathrm{Hom}_k(G, B)$, there are two steps remaining to decide if there is a walk in $\mathrm{Hom}_k(G, B)$ from ϕ to ψ. The first step is to decide if δ is realisable. If it is, we are done. If it is not, it is possible that there are other realisable difference functions that are topologically valid for ϕ and ψ. By Lemma 3.6 we have that the only other topologically valid difference functions are the shifts of δ by constant multiples of $b = |V(B)|$, but there are infinitly many of these. The second task therefore, is to decide if any of these are realisable. To do this, we show that if δ is not realisable, but some shift $\delta \pm s \cdot b$ is, then s cannot be too big.

Extending ideas from [2], we will prove three results that will allow us, in the next section, to accomplish the above steps. First, we prove Lemma 7.1 which tells us that if there is a walk realising a topologically valid difference function δ, then we can assume it has some nice properties. This will allow us to use a greedy algorithm to decide if there is a walk realising δ, and to find it if it exists. It also yields our second result, Lemma 7.5 which, in the case that the greedy algorithm does not yield a walk realising δ, gives us an obstruction to the existence of such a walk. Using these two lemmas, we prove the third main result of the section, Proposition 7.8, which says that if there is a walk realising some shift of δ, then there is one realising a shift of δ that is (linearly) bounded by the size of G.

Recall that a walk of indecomposable arcs between γ and γ' can be represented by $\square V_1, \square V_2, \ldots, \square V_\ell$ where each $\square$ is either $\uparrow$ or $\downarrow$ and each V_i is a minimal up-set (or down-set) with respect to $\gamma + \square V_1 + \square V_2 + \cdots + \square V_{i-1}$. A walk $\mathcal{W}$ is *in standard form* if the following all hold.

(i) $\mathcal{W}$ consists of indecomposable edges.
(ii) $\mathcal{W}$ is *unimodal*: there is some j such that $\square V_i = \uparrow V_i$ if and only if $i \le j$.
(iii) *Vertices are monotone*: no vertex v is in V_i for $i \le j$ and in V_k for $j < k$.

Lemma 7.1. *If there is a walk from ϕ to ψ in $\mathrm{Hom}_k(G,B)$ then there is a walk in standard form from ϕ to ψ.*

Before we prove this, we prove two useful claims.

Claim 7.2. *If $\uparrow U, \downarrow D$ is a walk of indecomposable arcs between some γ and γ', then either $U \cap D = \varnothing$ or $U = D$.*

Proof. We assume that $U \cap D \neq \varnothing$ and show that $U \subseteq D$. From this we can assume that $U = D$ as otherwise by Lemma 5.1 the indecomposible arc $\downarrow D$ would decompose into the walk $\downarrow U, \downarrow(D \setminus U)$.

So assume that $v \in U \cap D$. As $\uparrow U$ is indecomposable, U is a minimal up-set. By Proposition 6.5, U is a terminal strong component of $\mathbb{A}^+_\gamma$. Let $\gamma^U = \gamma + \uparrow U$. We show that U is also strongly connected in $\mathbb{A}^-_{\gamma^U}$, thus by the negative analogue (see Remark 6.8) of Proposition 6.5, we get that D contains all of U.

It is enough to show that if $v \to u$ in $\mathbb{A}^+_\gamma$, then $u \to v$ in $\mathbb{A}^-_{\gamma^U}$. So assume $v \to u$ in $\mathbb{A}^+_\gamma$. Both u and v are in U, as U is terminal and $v \in U$. Using the definition of $\mathbb{A}^+_\gamma$ we get, observing that the maps in the second and third line are the same,

$$\begin{aligned} v \to u \text{ in } \mathbb{A}^+_\gamma \\ &\Rightarrow \gamma|_{uv} + \uparrow v \text{ is not a homomorphism of } \{u,v\} \text{ to } B \\ &\Rightarrow \gamma^U_{uv} + \downarrow u \text{ is not a homomorphism of } \{u,v\} \text{ to } B \\ &\Rightarrow u \to v \text{ in } \mathbb{A}^-_{\gamma^U}, \end{aligned}$$

as needed. ◇

The following shows that up moves and down moves commute.

Claim 7.3. *If $\downarrow D, \uparrow U$ is a walk of indecomposable arcs between γ and γ', then $\uparrow U, \downarrow D$ is also a walk between them.*

Proof. Our proof is by contradiction. Towards this, we assume that $\gamma \sim \gamma_D \sim \gamma^U_D$ is a walk but $\gamma \sim \gamma^U \sim \gamma^U_D$ is not, where $\gamma_D = \gamma + \downarrow D$, $\gamma^U = \gamma + \uparrow U$ and $\gamma^U_D = \gamma'$. By Claim 7.2 we have that U and D are disjoint. (In the case that $U = D$, $\gamma = \gamma^U_D$ and we can simply shorten the walk in $\mathrm{Hom}(G,B)$ by the two arcs.) We show U is an up-set for γ. The proof that D is a down-set for γ^U is similar, and the claim follows.

To show U is an up-set for γ we show that it is a terminal set in $\mathbb{A}^+_\gamma$, and that it is not an up-obstruction: $\gamma|_U \sim \gamma^U|_U$ is an edge of $\mathrm{Hom}_k(G[U],B)$. That it is not an up-obstruction is immediate from the fact that $\gamma_D|_U \sim \gamma^U_D|_U$ and the fact that U and D are disjoint. So we have to show that U is terminal in $\mathbb{A}^+_\gamma$.

Assume that $u \to d$ in $\mathbb{A}^+_\gamma$ for some $u \in U$ and $d \notin U$. As $u \not\to d$ in $\mathbb{A}^+_{\gamma_D}$ we have that d must be in D. Now, ud is an edge of G; and γ^U_D, which maps d to $\gamma(d)-1$ and u to $\gamma(u)+1$, is a homomorphism. Hence, either $\gamma(u) = \gamma(d)-1$, or $b = 3$. The former contradicts the assumption that $u \to d$ in $\mathbb{A}^+_\gamma$, so we have $b = 3$. As above, $\gamma(u)+1 = \gamma(d)$ cannot give rise to $u \to d$ in $\mathbb{A}^+_\gamma$, so $\gamma(u) = \gamma(d)$ or $\gamma(u) = \gamma(d)-1$. We show the former cannot happen; the proof the latter cannot occur is similar. Towards contradiction, assume, without loss of generality that $u \to d$ in G, and

$\gamma(u) = \gamma(d) = 2$. (The case $d \to u$ is analogous.) As $u \to d$ in $\mathbb{A}^+_\gamma$ we have $3 \not\to 2$ in B, and thus $2 \to 3$. As γ_D is a homomorphism we have $2 = \gamma_D(u) \to \gamma_D(d) = 1$, and as γ^U_D is a homomorphism, we have $3 = \gamma^U_D(u) \to \gamma^U_D(d) = 1$. But then B contains a transitive triangle, which is false by the fact that (G, B) is a $\mathcal{WP}$-pair. This contradiction shows that $u \to d$ cannot be in $\mathbb{A}^+_\gamma$ and so U is terminal in $\mathbb{A}^+_\gamma$, as needed. ⋄

We are now ready to prove Lemma 7.1

Proof. Let there be a walk $\mathcal{W}$ from ϕ to ψ. As we have observed before, we may assume that $\mathcal{W}$ consists only of indecomposable edges. Throughout the proof, we will frequently alter $\mathcal{W}$. Whenever we do, we may introduce decomposable edges, but at each step, we implicitly decompose these into indecomposable edges. This allows us to use the claims above, and gives us item (i).

Using Claim 7.3 we can make $\mathcal{W}$ unimodal giving item (ii).

Assume that v moves up in some $\uparrow V_i$ and down in some $\downarrow V_k$. Using Claim 7.3 we can commute moves to get a walk in which all down moves come first. Then we can commute moves again so that all up moves come first. In doing so we had, at some point, $\uparrow V_i$ followed immediately by $\downarrow V_k$. But then by Claim 7.2 V_i and V_k are the same, so we can cancel $\uparrow V_i$ and $\downarrow V_k$ out of the walk. This establishes item (iii). □

Using this, we now define a certificate to the non-existence of a realisation of given topologically valid difference function δ: a walk $\mathcal{W}$ and an up-obstruction T showing that a vertex v cannot move up as far as $\delta(v)$ needs it too.

Definition 7.4. For a homomorphism $\phi \in \mathrm{Hom}(G, B)$ and a vertex v in $V(G)$, let

$$\max_\phi(v) = \max_{\mathcal{W}} \delta_{\mathcal{W}}(v)$$

where the $\mathcal{W}$ runs over all walks $\mathcal{W}$ in $\mathrm{Hom}_k(G, B)$ that start at ϕ. If there is no maximum, let $\max_\phi(v) = \infty$.

Lemma 7.5. *Let $\mathcal{W}$ be a walk from ϕ to γ in $\mathrm{Hom}_k(G, B)$ and let T be an up-obstruction of γ. For all v in T, $\max_\phi(v) = \delta_{\mathcal{W}}(v)$.*

Proof. By definition we have that $\max_\phi(v) \geq \delta_{\mathcal{W}}(v)$ so we just have to show that $\max_\phi(v) \leq \delta_{\mathcal{W}}(v)$ for all v in the up-obstruction T of γ. Assume, towards contradiction, that there is a walk $\mathcal{W}'$ from ϕ to some γ' in $\mathrm{Hom}_k(G, B)$ such that $\delta_{\mathcal{W}'}(v) > \delta_{\mathcal{W}}(v)$ for some $v \in T$. Then restricting from $\mathrm{Hom}(G, B)$ to $\mathrm{Hom}(G[T], B)$ we have that $-\mathcal{W} + \mathcal{W}'$ restricts to a walk $\mathcal{U}$ from $\gamma|_T$ to $\gamma'|_T$, and $\delta_{\mathcal{U}}$ is positive on v. Consequently, Lemma 7.1 allows us to assume this walk begins by moving some up-set of $\gamma|_T$ up. However, T is an up-obstruction for γ, so it is an up-obstruction of $\gamma|_T$. Thus, $\gamma|_T$ has no up-sets. □

We will need the following from [2].

Lemma 7.6. *Let δ be a topologically valid difference function for ϕ and ψ in $\mathrm{Hom}(G, B)$. If $\delta(u) > 0$ and $u \to v$ in $\mathbb{A}^+_\phi$, then $\delta(v) > 0$.*

This allows us to prove the following.

Lemma 7.7. *Let δ be a non-negative topologically valid difference function for ϕ and ψ in $\mathrm{Hom}(G, B)$. If δ is not realisable for ϕ and ψ, then neither is $\delta + b$.*

Proof. Assume that non-negative δ is not realisable for ϕ and ψ. Let $\gamma : V(G) \to \mathbb{Z}$ be a function such that there is a walk in $\mathrm{Hom}_k(G,B)$ from ϕ to $\phi+\gamma$ and that γ is a maximal such function for which $0 \le \gamma \le \delta$ holds. Such a γ exists as the 0 function is realisable for ϕ and $\phi+0$, and δ is not realisable for ϕ and $\psi = \phi+\delta$. Let P be the set of vertices v for which $\gamma(v) < \delta(v)$. The set P is non-empty as δ is not realisable. By the maximality of γ no vertices of P are sinks of $\mathbb{A}^+_{\phi+\gamma}$.

As δ is topologically valid for ϕ and ψ, and γ is topologically valid for ϕ and $\phi+\gamma$, $\delta-\gamma$ is topologically valid for $\phi+\gamma$ and ψ, and is positive on P. As no vertices of P are sinks of $\mathbb{A}^+_{\phi+\gamma}$, every vertex of P has an out-neighbour. By Lemma 7.6 (using $\delta-\gamma$), P is closed under out-neighbours. Thus P contains a terminal strong-component T of $\mathbb{A}^+_{\phi+\gamma}$. By the maximality of γ, T cannot be a minimal up-set, so by Proposition 6.5 it is an up-obstruction to $\phi+\gamma$. But then by Lemma 7.5 we have $\max_\phi(v) = \delta_W(v) < \delta(v) < (\delta+b)(v)$ for each $v \in T$, so $\delta+b$ is certainly not realisable for ϕ. □

A difference function δ is *N-bounded* if $|\delta(v)| < N$ for all v.

Proposition 7.8. *Let δ be a topologically valid difference function for ϕ and ψ. If some shift $\delta+cb$ of δ is realisable for ϕ and ψ, then some $(2n-2+b)$-bounded shift is realisable.*

Proof. Recall that if $u \sim v$, then $\mathrm{nl}_\phi(uv) \in \{-1,0,1\}$. This observation together with (1) from Definition 3.5, gives for the topologically valid function δ and for any two vertices v and w,

$$\delta(v) - 2d(v,w) \le \delta(w) \le \delta(v) + 2d(v,w),$$

where $d(v,w)$ is the distance in the underlying undirected graph. Note that we can shift δ by tb so that the minimum value of $\delta+tb$ is between 0 and $b-1$. Then the maximum value is at most $b-1+2(n-1) = b+2n-3$. Hence $0 \le \delta+tb < b+2n-2$. That is, $\delta+tb$ is $(2n-2+b)$-bounded. If no $(2n-2+b)$-bounded shift of δ is realisable, then in particular this holds for the minimum shift of δ that is nonnegative (as described above) and the maximum shift that is nonpositive. By induction, using Lemma 7.7 and its negative analogue, there is therefore no realisable difference function for ϕ and ψ. □

8. Realising topologically valid difference functions

Let (G,B) be a $\mathcal{WP}$-pair. In this section we give a polynomial time algorithm to decide if there is a realisation in $\mathrm{Hom}_k(G,B)$ of a given difference function δ that is topologically valid for two maps ϕ and ψ in $\mathrm{Hom}_k(G,B)$.

Consider the following greedy algorithm for determining whether or not a difference function δ, which is topologically valid for ϕ and ψ in $\mathrm{Hom}_k(G,B)$, is realisable.

Algorithm 8.1.

Input: A difference function δ that is topologically valid for ϕ and ψ in $\mathrm{Hom}_k(G,B)$.

Output: 'Yes' if δ is realisable for ϕ and ψ; 'No' otherwise.

Steps: Initialise by setting $i = 0$, letting $\delta^* = \delta$, and letting $\mathcal{W}^*$ be the length 0 walk at the vertex ϕ.

(I) While there is a vertex v with $\delta^*(v) > 0$, let γ be the last map in $\mathcal{W}^*$ and do the following.

(i) Construct $\mathbb{A}^+_\gamma$ and let P be the set of vertices of G on which δ^* is positive.
(ii) Check if γ has any minimal up-sets U in P.
(iii) If so, append $\gamma + \uparrow U$ to the end of $\mathcal{W}^*$ and reduce δ^* by one on all of U; if not return 'No' and quit.

(II) While there is a vertex v with $\delta^*(v) < 0$ do the same only moving down minimal down-sets of γ on on which δ^* is negative, and increasing δ^*.

(III) Now $\delta^* = 0$, return 'Yes'.

The following verifies that Algorithm 8.1 finishes correctly in polynomial time.

Proposition 8.2. *Algorithm 8.1 correctly determines, in time $O(N \cdot f(n))$, where $f(n)$ is a polynomial function independent of k, whether or not an N-bounded difference function δ, which is topologically valid for ϕ and ψ in $\mathrm{Hom}_k(G, B)$, is realisable.*

Proof. For the running time, observe that in each iteration of part (I) or (II) the sum of $|\delta^*(v)|$ is reduced. As δ is N-bounded, there can be at most $N \cdot n$ itereations. That each iteration completes in polynomial time comes from the observation that $\mathbb{A}^+_\gamma$ can be constructed in quadratic time, and from Corollary 6.7 which tells us we can find all minimal up-sets in P in polynomial time.

If the algorithm completes with $\delta^* = 0$ then we have clearly found a realisation, so what is left to show for correctness is that if there is realisation of δ from ϕ to ψ, then the algorithm finds one. For this it is enough to show that if δ^* is not zero, so without loss of generality δ has some positive values, necessarily on P, then there is a minimal up-set U in P for step (I)(ii).

So assume that there is a walk $\mathcal{W}$ from ϕ to ψ and we have constructed $\mathcal{W}^*$ from ϕ to γ. As $-\mathcal{W}^* + \mathcal{W}$ is a walk from γ to ψ we have by Lemma 7.1 that there is a standard form walk $\mathcal{U}$ from γ to ψ, and clearly $\delta_\mathcal{U} = \delta^*$. As δ^* has positive values on P, $\mathcal{U}$ starts with an indecomposable arc $\uparrow U$ for $U \subset P$, so this set U is a minimal up-set for γ, as needed. □

Remark 8.3. If Algorithm 8.1 returns 'Yes' then the final walk $\mathcal{W}^*$ is a walk from ϕ to ψ. When it returns 'No' we also have a certificate that such a walk cannot exist: indeed as we terminated with 'No' we have a walk $\mathcal{W}^*$ to γ for which the set P is an up-obstruction, so by Lemma 7.5, $\max_\phi(v) = \delta_{\mathcal{W}^*}(v) < \delta(v)$ for any v in P. This is a certificate that δ is not realisable for ϕ.

The proof of Theorem 3.4 is now immediate, and so, as observed before the statement of Theorem 3.4, the proof Theorem 2.2 is too

Proof of Theorem 3.4. Let (G, B) be a $\mathcal{WP}$-pair, and (G, ϕ, ψ) be an instance of Recon(B). By Lemma 3.6 we can determine in time polynomial in n if there is a topolgically valid difference function for ϕ and ψ. If not, we have by Lemma 3.6 that there is no walk in $\mathrm{Hom}_k(G, B)$ from ϕ to ψ, and we are done. So we may assume that there is such a function δ, and so also $\delta + cb$ for any integer c. If any

of these are realisable, then by Proposition 7.8 one of the at most $4n/b$ different $(2n-2+b)$-bounded shifts $\delta + cb$ are realisable. By Proposition 8.2 we can run Algorithm 8.1 on all of these in time polynomial in n to determine if any of these are realisable. If any are, then there is a walk in $\mathrm{Hom}_k(G,B)$ from ϕ to ψ, if none are, then there is no such walk. □

Observe that we have also proved Theorem 2.1 as the determination that there is a topologically valid difference function is completley independent of k, and by Proposition 8.2 the rest of our algorithm runs in time that is polynomial, independent of k. For any size n instance of $\textsc{Recon}(B)$ or $\textsc{Recon}^{\mathrm{r}}(B)$ we can simply solve $\textsc{Recon}_n(B)$ or $\textsc{Recon}^{\mathrm{r}}{}_n(B)$.

Remark 8.4. In [2] we used a more complicated algorithm that does not require linearly many applications of Algorithm 8.1. It shows that if there is no walk between ϕ and ψ in $\mathrm{Hom}_k(G,B)$, then there is some cycle C in G such that there is no walk between $\phi|_C$ and ψ_C in $\mathrm{Hom}_k(C,B)$. We could do the same here, but opted for a more transparent algorithm.

9. Discussion

In [2] we observed that once one has found a walk between maps in $\mathrm{Hom}_1(G,B)$, then it is a simple task to decide if there is a directed walk or a symmetric walk between them. For reconfiguration the situation is a little different, but we can replace 'walk' in Theorem 3.4 with either 'symmetric walk' or 'directed walk'.

9.1. **Symmetric Walks.** The case of symmetric walks is almost trivial, indeed the following observation, which generalises the earlier mentioned fact that for symmetric G, there is a path between two vertices in $\mathrm{Hom}(G,H)$ if and only if there is a path between them in $\mathrm{Hom}_1(G,H)$, reduces the problem to the finding a symmetric path in $\mathrm{Hom}_1(G,B)$, so is done in [2].

Fact 9.1. *Let G and H be digraphs. If (ϕ,ψ) is a symmetric arc in $\mathrm{Hom}(G,H)$ then there is a path of symmtric arcs between ϕ and ψ in $\mathrm{Hom}_1(G,H)$.*

Proof. Let s be any vertex that has different images under ϕ and ψ, and let ϕ' be the same as ϕ except that $\phi'(s) = \psi(s)$. Clearly it is enough to verify that ϕ' is a homomorphism, and that

$$\phi \leftrightarrow \phi' \leftrightarrow \psi.$$

To see that all ϕ' is a homomorphism, it is enough to check, when $u \to s \to w$, that $\phi'(u) \to \phi'(s) \to \phi'(w)$. But we have $\phi(u) \to \psi(s) \to \phi(w)$ by assumption and we have $\phi'(u) = \phi(u)$ and $\phi'(w) = \phi(w)$ by definition of ϕ', so ϕ' is a homomorphism.

To see that $\phi \leftarrow \phi' \to \psi$, it is enough to check, for $u \to s$ and $s \to w$, that

$$\phi(s) \leftarrow \phi'(u) \to \psi(s) \text{ and } \phi(w) \leftarrow \phi'(s) \to \psi(w).$$

Since $\phi'(u) = \phi(u)$, the former holds because ϕ is a homomorphism and $\phi \to \psi$. Since $\phi'(s) = \psi(s)$, the later holds, because ψ is a homomorphism and $\psi \to \phi$. Similarly we can prove $\phi \to \phi' \leftarrow \psi$, and so get $\phi \leftrightarrow \phi' \leftrightarrow \psi$. □

9.2. **Directed Walks.** For the task of deciding if there is a directed walk between maps in $\mathrm{Hom}(G,B)$ or $\mathrm{Hom}_k(G,B)$ we cannot find a post-hoc argument. However, it is not too hard to see that our main proof works in this situation to give a version of Theorem 3.4 in which we replace 'walk' with 'directed walk'. In fact, we have written it so that one need make only minimal changes. (Though integrating these in the proof was distracting.) One will notice that the defintion and properties of $\mathbb{A}^+_\phi$ are unchanged as they do not consider the arcs of $\mathrm{Hom}_1(G,B)$, only that the 'moved' vertices must still yield a homomorphism. Starting from Section 6 we must make some small and obvious changes to decide if there is a directed path from ϕ to ψ:

- At the beginning of Section 6 we must say T *moves up* if $\gamma \to (\gamma + \uparrow T)$ in $\mathrm{Hom}_k(G,B)$, and adjust the definition of minimal up-set accordingly.
- In Definition 6.3 we say T is a **forward up-obstruction** if $\gamma|_T \not\to (\gamma|_T + \uparrow T)$ in $\mathrm{Hom}_k(G[T],B)$.
- In Lemma 6.4 the same proof shows that $T \subset V(G)$ moves up (now via a forward arc) if and only if it is a terminal set but not a forward up-obstruction.
- The rest of Section 6 follows making only these same replacements.
- In Section 7 the only place where one must be careful is in Claim 7.3. One must check in the proof that if these two edges, $\uparrow U$ and $\downarrow D$ are both forward arcs, then the remain forward arcs when they commute. But this is clear from the second paragraph of the proof.

With these changes, everything goes through as it does in the main case of not necessarily directed walks.

10. Concluding Remarks

The proof of our result shares many features with the result of Wrochna from [10] which says that for C_4-free graphs H the problem $\textsc{Recol}(H)$ is polynomial time solvable. It appears that Wrochna's results have spawned several generalisations, but there do not appear to be any that deal with digraphs. One wonders if one could generalise our result to a result of the same breadth as Wrochna's, showing that $\textsc{Recon}(D)$ is polynomial time solvable for any C_4-free reflexive digraph.

11. Acknowledgements

The first author is supported by the Natural Sciences and Engineering Research Council of Canada Grant RGPIN-2014-04760. The third author is supported by Korean NRF Basic Science Research Program (NRF-2018R1D1A1A09083741) funded by the Korean government (MEST) and the Kyungpook National University Research Fund.

References

[1] P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoret. Comput. Sci.*, 410(2009) 5215–5226.

[2] R. C. Brewster, J. Lee, and M. Siggers. Recolouring Reflexive Digraphs Discrete Math. 341(6), 1708-1721.

[3] R. C. Brewster, S. McGuinness, B. Moore, and J. A. Noel. On the complexity of the reconfiguration problem for graph homomorphisms. *J. Theoret. Comput. Sci.*, 639(2016) 1–13.

[4] R. C. Brewster and J. A. Noel. Mixing homomorphisms, recolorings, and extending circular precolorings. *J. Graph Theory*, 80(2015) 173–198.

[5] L. Cereceda. *Mixing Graph Colourings*. PhD thesis, The London School of Economics and Political Science, 2007.
[6] L. Cereceda, J. van den Heuvel, and M. Johnson. Connectedness of the graph of vertex-colourings. *Discrete Math.*, 308(2008) 913–919.
[7] T. Feder and M. Y. Vardi. *The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory.* SIAM J. Comput., 28/1 (1998), pp. 57–104.
[8] B. Larose, C. Loten, and L. Zádori, A polynomial-time algorithm to recognize near-unanimity graphs, *J. Algorithms*, 55(2005) 177–191.
[9] M. Maróti, and L. Zádori *Reflexive digraphs with Near Unanimity Operations Discrete Math.*, 12(2012) 2316–2328.
[10] M. Wrochna. Homomorphism reconfiguration via homotopy. In *32nd International Symposium on Theoretical Aspects of Computer Science*, volume 30 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 730–742. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2015.

Department of Mathematics, Thompson Rivers Univesity, Kamloops, BC, Canada
Email address: rbrewster@tru.ca

College of Natural Sciences, Kyungpook National University, Daegu 702-701, South Korea
Email address: dlwoqor0923@gmail.com

College of Natural Sciences, Kyungpook National University, Daegu 702-701, South Korea
Email address: mhsiggers@knu.ac.kr