Surjective H-Colouring over Reflexive Digraphs*

BENOÎT LAROSE, LACIM, Université du Québec a Montréal, Canada
BARNABY MARTIN, Department of Computer Science, Durham University, U.K.
DANIËL PAULUSMA, Department of Computer Science, Durham University, U.K.

The Surjective H-Colouring problem is to test if a given graph allows a vertex-surjective homomorphism to a fixed graph H. The complexity of this problem has been well studied for undirected (partially) reflexive graphs. We introduce *endo-triviality*, the property of a structure that all of its endomorphisms that do not have range of size 1 are automorphisms, as a means to obtain complexity-theoretic classifications of Surjective H-Colouring in the case of reflexive *digraphs*. Chen [2014] proved, in the setting of constraint satisfaction problems, that Surjective H-Colouring is NP-complete if H has the property that all of its polymorphisms are essentially unary. We give the first concrete application of his result by showing that every endo-trivial reflexive digraph H has this property. We then use the concept of endo-triviality to prove, as our main result, a dichotomy for Surjective H-Colouring when H is a reflexive tournament: if H is transitive, then Surjective H-Colouring is in NL, otherwise it is NP-complete. By combining this result with some known and new results we obtain a complexity classification for Surjective H-Colouring when H is a partially reflexive digraph of size at most 3.

CCS Concepts: • **Mathematics of computing → Graph algorithms**; • **Theory of computation → Problems, reductions and completeness**;

Additional Key Words and Phrases: Surjective H-Coloring, Computational Complexity, Algorithmic Graph Theory, Universal Algebra, Constraint Satisfaction

ACM Reference Format:
Benoît Larose, Barnaby Martin, and Daniël Paulusma. 2010. Surjective H-Colouring over Reflexive Digraphs. *ACM Trans. Comput. Theory* 9, 4, Article 39 (March 2010), 21 pages. https://doi.org/0000001.0000001

1 INTRODUCTION

The classical *homomorphism problem*, also known as H-Colouring, involves a fixed structure H, with input another structure G, of the same signature, invoking the question as to whether there is a function from the domain of G to the domain of H that is a homomorphism from G to H. The H-Colouring problem is an intensively studied problem, which has additionally attracted attention in its guise of the *constraint satisfaction problem* (CSP), especially since the seminal paper of Feder and Vardi 1998. Their well-known conjecture, recently proved by Bulatov 2017 and Zhuk 2017, stated that every CSP(H) has complexity either in P or NP-complete, omitting any Ladner-like complexities in between.

*An extended abstract of this article has appeared at the 35th International Symposium on Theoretical Aspects of Computer Science (STACS 2018) [Larose et al. 2018]. The first author was supported by NSERC and FRQNT and the third author was supported by The Leverhulme Trust (RPG-2016-258).

Authors' addresses: Benoît Larose, LACIM, Université du Québec a Montréal, Canada; Barnaby Martin, Department of Computer Science, Durham University, U.K.; Daniël Paulusma, Department of Computer Science, Durham University, U.K..

List H-Colouring → H-Retraction → H-Compaction → Surj H-Colouring → H-Colouring

Fig. 1. Relations between Surjective H-Colouring and its variants (from [Golovach et al. 2017]). An arrow from one problem to another indicates that the latter problem is polynomial-time solvable for a graph H whenever the former is polynomial-time solvable for H. Reverse arrows do not hold for the leftmost and rightmost arrows, as witnessed by the reflexive 4-vertex cycle for the rightmost arrow and by any reflexive tree that is not a reflexive interval graph for the leftmost arrow (Feder, Hell and Huang 2003 showed that the only reflexive bi-arc graphs are reflexive interval graphs). It is not known whether the reverse direction holds for the two middle arrows.

This paper concerns the computational complexity of the *surjective homomorphism* problem, which is also known as Surjective H-Colouring [Golovach et al. 2017, 2012] and H-Vertex-Compaction [Vikas 2013] and which is to test whether a given structure G admits a *surjective* homomorphism to a fixed structure H. The surjective homomorphism problem is a cousin of the *list homomorphism* problem and is even more closely related to the *retraction* and *compaction* problems. Indeed, the H-Compaction problem, hitherto defined only for graphs H, takes as input a graph G and asks if there exists a function f from $V(\mathrm{G})$ to $V(\mathrm{H})$ so that for each non-loop edge $(x, y) \in E(\mathrm{H})$ (i.e. with $x \neq y$), there exists $u, v \in V(\mathrm{G})$ so that $f(u) = x$ and $f(v) = y$. Thus, compaction can be seen as the *edge-surjective homomorphism* problem.[1] The problem H-Retraction takes as input a superstructure G of H and asks whether there is a homomorphism from G to H that is the identity on H. The H-Retraction problem is polynomially equivalent with a special type of CSP, CSP(H′), where H′ is H decorated with constants naming the elements of its domain. Feder and Vardi 1998 showed that the task of classifying the complexities of the retraction problems is equivalent to that for the CSPs. Hence, owing to [Bulatov 2017; Zhuk 2017], H-Retraction has now been fully classified.

The list homomorphism problem, List H-colouring, allows one to express restricted lists for each of the input structure's elements, that are the only domain elements permitted in a solution homomorphism. List H-colouring is also a special type of CSP, CSP(H′), where H′ is H replete with all possible unary relations over the domain of H. Historically, the complexities of List H-colouring were first settled in the case where H is a graph [Feder and Hell 1998; Feder et al. 1999, 2003], followed by a general proof of dichotomy for all structures by Bulatov 2011. For a thorough treatment of graphs and homomorphisms we refer to the book [Hell and Nešetřil 2004].

In contrast to the situation for H-Colouring, List H-Colouring and H-Retraction, the complexity classifications for H-Compaction and Surjective H-Colouring are far from settled, and there are concrete open cases (see 3-No-Rainbow-Colouring in the survey [Bodirsky et al. 2012]). Obtaining NP-hardness for compaction and surjective homomorphism problems appears to be especially challenging. The complexity-theoretic relationship between these various problems is drawn in Figure 1. At present it is not known whether there is a graph H so that H-Retraction, H-Compaction and Surjective H-Colouring do not have the same complexity up to polynomial time reduction (see [Golovach et al. 2017; Vikas 2005]).

Nevertheless, classification results for Surjective H-Colouring have tried to keep pace with similar ones for H-Retraction. In [Feder et al. 2010] it is proved, among partially reflexive pseudoforests H, where the problem H-Retraction is either in P or is NP-complete. A similar classification for Surjective H-Colouring over partially reflexive forests can be inferred from the classification for partially reflexive trees in [Golovach et al. 2012]. The quest for a classification for H-Compaction and Surjective H-Colouring over pseudoforests is ongoing, but for both

[1] Except for the treatment of self-loops, which appears to be an idiosyncrasy that plays no vital role in computational complexity. For some history of the definition see [Vikas 2002].

problems already the reflexive 4-cycle took some time to classify [Martin and Paulusma 2015; Vikas 2002], as well as the irreflexive 6-cycle [Vikas 2004, 2017].

The above results are for undirected graphs, whereas we focus on digraphs. A known classification for H-Retraction comes for irreflexive semicomplete digraphs H. Bang-Jensen, Hell, and MacGillivray 1988 proved that H-Colouring is always in P or is NP-complete if H is irreflexive semicomplete. This is a fortiori a classification for H-Retraction since semicomplete digraphs are *cores* (all endomorphisms are automorphisms), which ensures that H-Colouring and H-Retraction are polynomially equivalent. For irreflexive semicomplete digraphs H, the classification for Surjective H-Colouring can be read trivially from that for H-Colouring, and they are the same. An obvious next place to look is at the situation if H is *reflexive* semicomplete, where surely the classifications will not be the same as H-Colouring is trivial in this case.

Reflexive tournaments form an important subclass of the class of reflexive semicomplete graphs and are well-understood algebraically [Larose 2006]. In particular, the classification for H-Retraction where H is a reflexive tournament can be inferred from the algebraic characterisation from [Larose 2006]: for a reflexive tournament H, the H-Retraction problem is in NL if H is transitive, and it is NP-complete otherwise. This raises the question whether the same holds for Surjective H-Colouring and whether we can develop algebraic methods further to prove this. In fact, the algebraic method is by now well known for CSPs and their relatives, including its use with digraphs; see the recent survey [Larose 2017]. However, the algebraic method is not so far advanced for surjective homomorphism problems. So far it only exists in the work of Chen 2014, who proved that Surjective H-Colouring is NP-complete if H has the property that all of its polymorphisms depend only on one variable, that is, are essentially unary. Chen's result has not yet been put to work (even on toy open problems) and a key driver for our research has been to find, in the wild, a place for its application.

Our Results

We give, for the first time, complexity classifications for Surjective H-Colouring for digraphs instead of undirected graphs. To prove our results, we further develop algebraic machinery to tackle surjective homomorphism problems. That is, in Section 2 we introduce, after giving the necessary terminology, the concept of endo-triviality. We show how this concept is closely related to some known algebraic concepts and explore its algorithmic consequences in the remainder of our paper. We also exhibit an infinite family of reflexive tournaments that are endo-trivial, in order to evidence the significance of this class.

Firstly, in Section 3, we prove that a reflexive digraph H that is endo-trivial has the property that all of its polymorphisms are essentially unary. Combining this result with the aforementioned result of Chen 2014 immediately yields that Surjective H-Colouring is NP-complete for any such digraph H. This is the first concrete application of Chen's result to settle a problem of open complexity; it shows, for instance, that Surjective H-Colouring is NP-complete if H is a reflexive directed cycle on $k \geq 3$ vertices. As the case $k \leq 2$ is trivial, this gives a classification of Surjective H-Colouring for reflexive directed cycles, which we believe form a natural class of digraphs to consider given the results in [Martin and Paulusma 2015; Vikas 2017].

Secondly, in Section 4 we give a complexity classification for Surjective H-Colouring, when H is a reflexive tournament. We use endo-triviality in an elaborate and recursive encoding of an NP-hard retraction problem within Surjective H-Colouring. In doing this, we show that on this class, the complexities of Surjective H-Colouring and H-Retraction coincide.

Finally, our results enable us to give, in Section 5, a complexity classification for Surjective H-Colouring when H is a partially reflexive digraph of size at most 3. In doing this, we show that on this class, the complexities of Surjective H-Colouring and H-Retraction coincide. We are not aware of an existing classification for H-Retraction on this class, but we do build on one existing for List H-Colouring from [Feder et al. 2006].

2 PRELIMINARIES

2.1 Directed graphs

Let $[n] := \{1, \ldots, n\}$. For a k-tuple $\bar{t}$ and $i \in [k]$, let $\bar{t}[i]$ be the ith entry in $\bar{t}$. In a digraph G, a forward- (resp., backward-) *neighbour* (or *adjacent*) to a vertex $u \in V(\mathrm{G})$ is another vertex $v \in V(\mathrm{G})$ so that $(u, v) \in E(\mathrm{G})$ (resp., $(v, u) \in E(\mathrm{G})$). The *out-degree* and *in-degree* of a vertex are the number of its forward-neighbours and backward-neighbours, respectively. A vertex with out-degree and in-degree both 0 is said to be *isolated*. A vertex with a self-loop is *reflexive* and otherwise it is *irreflexive*. A digraph is *(ir)reflexive* if all its vertices are (ir)reflexive.

The *directed path* on k vertices is the digraph with vertices $u_0, \ldots, u_{k-1}$ and edges (u_i, u_{i+1}) for $i = 0, \ldots, k-2$. The *directed cycle* on k vertices is obtained from the directed path on k vertices after adding the edge (u_{k-1}, u_0). A digraph G is *strongly connected* if for all $u, v \in V(\mathrm{G})$ there is a directed path in $E(\mathrm{G})$ from u to v (note that we take this to include the situation $u = v$, but for reflexive graphs the distinction is moot). A digraph is *weakly connected* if its symmetric closure (underlying undirected graph) is connected. A *double edge* (or *digon*) in a digraph G consists in a pair of distinct vertices $u, v \in V(\mathrm{G})$, so that $(u, v), (v, u) \in E(\mathrm{G})$. A digraph G is *semicomplete* if for every two distinct vertices u and v, at least one of (u, v), (v, u) belongs to $E(\mathrm{G})$.

A digraph G is a *tournament* if for every two distinct vertices u and v, exactly one of (u, v), (v, u) belongs to $E(\mathrm{G})$. Some (reflexive) tournaments are drawn in Figures 2 and 3. We demand our tournaments have more than one vertex (to rule out certain trivial cases in proofs). A reflexive tournament G is *transitive* if for every triple of vertices u, v, w with $(u, v), (v, w) \in E(\mathrm{G})$, also (u, w) belongs to $E(\mathrm{G})$. A digraph F is a *subgraph* of a digraph G if $V(\mathrm{F}) \subseteq V(\mathrm{G})$ and $E(\mathrm{F}) \subseteq E(\mathrm{G})$. It is *induced* if $E(\mathrm{F})$ coincides with $E(\mathrm{G})$ restricted to pairs containing only vertices of $V(\mathrm{F})$. A *subtournament* is an induced subgraph of a tournament (note that this is a fortiori a tournament). All subgraphs we consider in this paper will be induced.

2.2 Homomorphisms and Algebra

A *homomorphism* from a digraph G to a digraph H is a function $f : V(\mathrm{G}) \to V(\mathrm{H})$ so that for all $u, v \in V(\mathrm{G})$ with $(u, v) \in E(\mathrm{G})$ we have $(f(u), f(v)) \in E(\mathrm{H})$. We say that f is *(vertex)-surjective* if for every vertex $x \in V(\mathrm{H})$ there exists a vertex $u \in V(\mathrm{G})$ with $f(u) = x$. Let H be a digraph. A *homomorphic image* of H is a digraph H′ so that there is a surjective homomorphism $h : \mathrm{H} \to \mathrm{H}'$ in which, for all $(x', y') \in E(\mathrm{H}')$ there exists $(x, y) \in E(\mathrm{H})$ so that $x' = h(x)$ and $y' = h(y)$. That is, h is vertex- and edge-surjective.

The *direct product* of two digraphs G and H, denoted $\mathrm{G} \times \mathrm{H}$, has vertex set $V(\mathrm{G}) \times V(\mathrm{H})$ and edges $((x, y), (x', y'))$ exactly when $(x, x') \in E(\mathrm{G})$ and $(y, y') \in E(\mathrm{H})$. This product is associative and commutative, up to isomorphism, and spawns a natural power. A k-ary *polymorphism* of G is a function $f : \mathrm{G}^k \to \mathrm{G}$ so that when $(x_1, y_1), \ldots, (x_k, y_k) \in E(\mathrm{G})$ then $(f(x_1, \ldots, x_k), f(y_1, \ldots, y_k)) \in E(\mathrm{G})$. A polymorphism of G can be seen as a homomorphism from the kth (direct) power of G, G^k, to G. A polymorphism f is *idempotent* if for all $x \in V(\mathrm{G})$, $f(x, \ldots, x) = x$. The k-ary ith *projection*, for $i \in [k]$, is the polymorphism π^i_k given by $\pi^i_k(x_1, \ldots, x_k) = x_i$. A k-ary operation f is called *essentially unary* if there exists a unary operation g and $i \in [k]$ so that $f(x_1, \ldots, x_k) = g(x_i)$ for all $(x_1, \ldots, x_k) \in \mathrm{G}^k$.

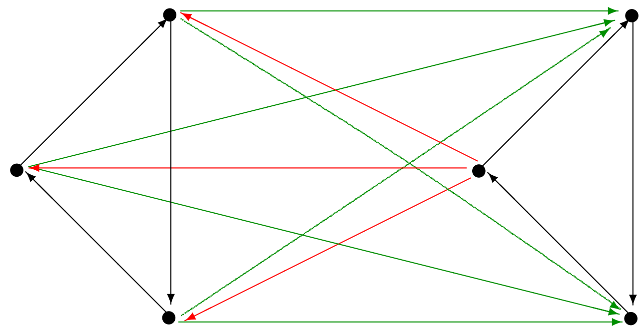

Fig. 2. A tournament on six vertices (self-loops are not drawn), which retracts to the directed 3-cycle (in black) on the right-hand side, but not to the one on the left-hand side (in black as well). However, there is no endomorphism that maps the left-hand one isomorphically to the right. We can use this tournament to build a structure that is a counterexample to the generalisation of Lemma 2.5 stating that the notions of endo-triviality and retract-triviality coincide. Let us label the vertices in the tournament: α, β, γ (left-hand DC_3^*, clockwise from bottom) and $0, 1, 2$ (right-hand DC_3^*, clockwise from bottom). Let us build a structure B by augmenting a new 6-ary relation with tuples in $\{(\alpha, \beta, \gamma, 0, 1, 2), (\alpha, \alpha, \alpha, \alpha, \beta, \gamma), (\alpha, \alpha, \alpha, \alpha, \alpha, \alpha)\}$. The structure B is retract-trivial but is not endo-trivial, since it has an interesting endomorphism that takes $(\alpha, \beta, \gamma, 0, 1, 2)$ to $(\alpha, \alpha, \alpha, \alpha, \beta, \gamma)$.

Let G be a digraph. A *self-map* of G is a function $f : V(\mathrm{G}) \to V(\mathrm{G})$. The *self-map digraph* G^G has as its vertices the self-maps of G and there is an edge $(f, g) \in E(\mathrm{G}^\mathrm{G})$ between self-maps f and g if and only if for every edge $(x, y) \in E(\mathrm{G})$, we have that $(f(x), g(y)) \in E(\mathrm{G})$. An *endomorphism* of G is a homomorphism from G to itself. The *endomorphism digraph* $\widehat{\mathrm{G}^\mathrm{G}}$ is the restriction of the self-map digraph G^G to the vertices induced by endomorphisms of G. Note that the self-loops of G^G are precisely the endomorphisms of G, so $\widehat{\mathrm{G}^\mathrm{G}}$ is reflexive when G is reflexive.

We now make two more observations. The first one follows directly from the definition of $\widehat{\mathrm{G}^\mathrm{G}}$. The second one can, for example, be found in Section 5.2 of [Larose et al. 2007].

Lemma 2.1. *Let* G *be a digraph. If* $(f_1, g_1) \in E(\widehat{\mathrm{G}^\mathrm{G}})$ *and* $(f_2, g_2) \in E(\widehat{\mathrm{G}^\mathrm{G}})$*, then* $((f_1 \circ f_2), (g_1 \circ g_2)) \in E(\widehat{\mathrm{G}^\mathrm{G}})$.

Lemma 2.2. *Let* G *and* H *be two digraphs. Let* φ *be a homomorphism from* $\mathrm{H} \times \mathrm{G}$ *to* G*. Then the function* ψ *defined by* $\psi(x)(u) = \varphi(x, u)$ *for all* $x \in V(\mathrm{H})$, $u \in V(\mathrm{G})$ *is a homomorphism from* H *to* $\widehat{\mathrm{G}^\mathrm{G}}$.

An endomorphism e of G is a *constant map* if there exists a vertex $v \in V(\mathrm{G})$ such that $e(u) = v$ for all $u \in V(\mathrm{G})$. A bijective endomorphism whose inverse is a homomorphism is an *automorphism*. An endomorphism is *non-trivial* if it is neither an automorphism nor a constant map. A digraph, all of whose endomorphisms are automorphisms, is termed a *core*. An endomorphism e of a digraph H *fixes* a subset $S \subseteq V(\mathrm{H})$ if $e(S) = S$, that is, $e(x) \in S$ for all $x \in S$, and it fixes a subgraph F of H if $e(\mathrm{F}) = \mathrm{F}$. It fixes an induced subgraph F *up to automorphism* if $e(\mathrm{F})$ is an automorphic copy of F (this is a stronger condition than $e(\mathrm{F})$ being isomorphic to F). An endomorphism r of G is a *retraction* of G if r is the identity on the image $r(\mathrm{G})$ (thus a retraction must have at least one fixed point).

2.3 Endo-triviality and Retract-triviality

We now define the key concept of endo-triviality and the closely related concept of retract-triviality.

Definition 2.3. A digraph is *endo-trivial* if all of its endomorphisms are automorphisms or constant maps.

The concept of endo-triviality also arises from the perspective of the algebra of polymorphisms. An algebra is called minimal if its unary polynomials are either constants or permutations (see Definition 2.14 in [Hobby and McKenzie 1988]).

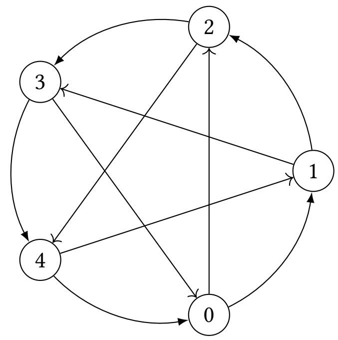

Fig. 3. The digraph G_2 (self-loops are not drawn).

For reflexive digraphs, polynomials and polymorphisms coincide. In other words, a reflexive digraph is endo-trivial if and only if its associated algebra of polymorphisms is minimal.

We will also need the following closely related concept.

Definition 2.4. A digraph is *retract-trivial* if all of its retractions are the identity or constant maps.

The concept of retract-triviality also appears in the algebraic theory but has, as far as we are aware, not been studied in a combinatorial setting. An algebra is term-minimal if the only retractions in its clone of terms are the identity and constants (see [Szendrei 1994]). A reflexive digraph is retract-minimal if its associated algebra of polymorphisms is term-minimal. It follows that on reflexive digraphs, the concepts of retract-minimality and retract-triviality coincide.

We note that every endo-trivial structure is also retract-trivial. However, the reverse implication is not necessarily true: in Figure 2 we give an example of a structure that is retract-trivial but not endo-trivial. This example is based on a digraph but is not itself a digraph. It is also possible to construct a retract-trivial digraph that is not endo-trivial [Siggers 2017], but on reflexive tournaments both concepts do coincide.

Lemma 2.5. *A reflexive tournament is endo-trivial if and only if it is retract-trivial.*

Proof. (Forwards.) Trivial. (Backwards.) By contraposition, suppose e is a non-trivial endomorphism of a reflexive tournament H. Consider $e(\mathrm{H})$ and build some function e^{-1} from $e(\mathrm{H})$ to H by choosing $e^{-1}(y) = x$ if $e(x) = y$ arbitrarily. Since H is a (reflexive) tournament, e^{-1} is an isomorphism, whereupon $e^{-1} \circ e$ is the identity automorphism when restricted to some subtournament H_0 of H. Hence $e^{-1} \circ e$ is a non-trivial retraction of H (to H_0). ◻

The class of endo-trivial digraphs is clearly infinite, since all cores are endo-trivial. However, this observation is not interesting. Of more significance is that there is even an infinite family of endo-trivial reflexive tournaments. Let $n \geq 2$ be an integer. The digraph G_n has set of vertices $V(\mathrm{G}_n) = [0, 2n] = \{0, 1, \ldots, 2n\}$ and edges $(i, i + j) \in E(\mathrm{G}_n)$ for $i \in [0, 2n]$ and $j \in [0, n]$ (where the sum is taken modulo $2n + 1$.) By way of example, G_2 is depicted in Figure 3. Clearly each digraph G_n is a reflexive tournament.

Lemma 2.6. *For every $n \geq 2$ the digraph G_n is endo-trivial.*

Proof. Owing to Lemma 2.5, it suffices to prove that it is retract-trivial. Let f be a retraction of G, i.e. a homomorphism $f : \mathrm{G} \to \mathrm{G}$ such that $f \circ f = f$.

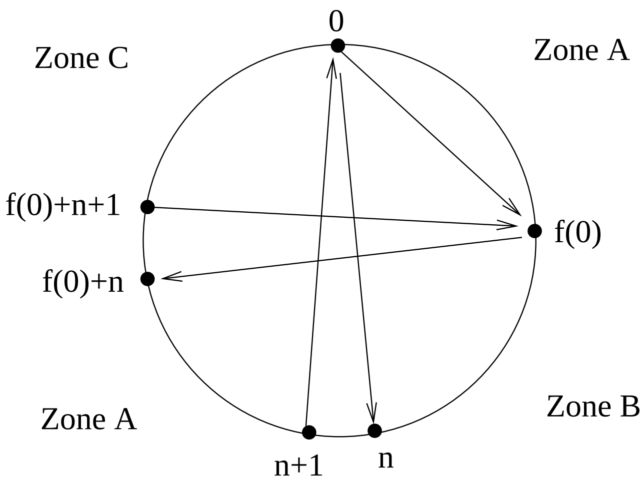

Fig. 4. Illustration for the proof of Lemma 2.6 (self-loops are not drawn).

Claim. *If* $(s, w), (w, f(s)) \in E(\mathrm{G}_n)$ *or* $(f(s), w), (w, s) \in E(\mathrm{G}_n)$, *then* $f(s) = f(w)$.

Proof of Claim. The proof of both statements is identical, so suppose $(s, w), (w, f(s)) \in E(\mathrm{G}_n)$: apply f to obtain $(f(s), f(w)), (f(w), f(s)) \in E(\mathrm{G}_n)$ (since f is a retraction) and thus $f(s) = f(w)$.

Assume that f is not the identity, i.e. $f(s) \neq s$ for some $s \in V(\mathrm{G}_n)$. We prove that f is a constant map. By symmetry, and without loss of generality, we may assume that $0 < f(0) \leq n$ (i.e $s = 0$). This is because otherwise we could make an identical argument but with the edges reversed.

Let us consider separately u in each of the three zones described in Figure 4. (A.) By the claim, we immediately get that $f(u) = f(0)$ for all $0 \leq u \leq f(0)$ and all $n + 1 \leq u \leq f(0) + n$ (this latter implies $(f(0), u), (u, 0) \in E(\mathrm{G}_n)$). (B.) Now let $f(0) + 1 \leq u \leq n$; we then have that $(f(0), u), (u, f(0) + n) \in E(\mathrm{G}_n)$, and since $f(f(0) + n) = f(0)$ we may apply the claim again and conclude that $f(u) = f(0)$. (C.) Finally let $f(0) + (n + 1) \leq u \leq 2n$; then $(f(0) + n, u), (u, f(0)) \in E(\mathrm{G}_n)$ and by the claim once more we conclude that $f(u) = f(0)$. Thus f is the constant map with value $f(0)$. ◻

3 ESSENTIAL UNARITY AND A DICHOTOMY FOR REFLEXIVE DIRECTED CYCLES

In this section we give the first concrete application, of which we are aware, of the aforementioned result of Chen, formally stated below.

Theorem 3.1 (Corollary 3.5 in [Chen 2014]). *Let* H *be a finite structure whose universe* $V(\mathrm{H})$ *has size strictly greater than* 1. *If each polymorphism of* H *is essentially unary, then* Surjective H-Colouring *is* NP*-complete.*

In order to prove this, we make use of the endomorphism graph and a result from Mároti and Zádori 2012. Let id_{H} denote the identity map on a digraph H.

Lemma 3.2 (Lemma 2.2 in [Maróti and Zádori 2012]). *Let* H *be a reflexive digraph. If* $(id_{\mathrm{H}}, f) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$, *where* f *is different from* id_{H}, *then* H *has a non-surjective retraction* r *such that* $(id_{\mathrm{H}}, r) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$.

The following lemma is crucial and will be of use in the next section as well.

Lemma 3.3. *Let* H *be a retract-trivial reflexive digraph with at least three vertices. Then*

(1) H *has no double edge;*
(2) H *is strongly connected; and*
(3) *the automorphisms of* H *are isolated vertices in* $\widehat{\mathrm{H}^{\mathrm{H}}}$.

Proof. (1) As any reflexive digraph can be retracted onto a double edge, the result follows.

(2) A reflexive digraph that is not weakly connected may be retracted onto a 2-vertex digraph. So we may safely assume H is weakly connected. If H is not strongly connected, then H has an edge (a, b) such that a and b are in different strong components, that is, there is no directed path from b to a. We define a retraction of H onto the subgraph induced by $\{a, b\}$ as follows: let $r(x) = a$ if there exists a directed path from x to a and $r(x) = b$ otherwise. As $r(a) = a$ and $r(b) = b$, it remains to check if r is an endomorphism. Let (x, y) be an edge. If $r(y) = a$, then there exists a directed path from y to a, and thus a directed path from x to a implying that $r(a) = a$. If $r(y) = b$, then $r(x) = a$ and $r(x) = b$ are both allowed.

(3) We first prove that no constant map is adjacent to an automorphism in $\widehat{\mathrm{H}^{\mathrm{H}}}$. Suppose for a contradiction that there exists an automorphism σ backwards-adjacent to some constant (the case of forwards-adjacent is dual). Thus without loss of generality we may assume that $(\sigma, c) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$ for some constant map c, say for all $u \in V(\mathrm{H})$, $c(u) = v$ for some $v \in V(\mathrm{H})$. By composing sufficiently many times via Lemma 2.1, we obtain that $(id_{\mathrm{H}}, c) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$. Since H is strongly connected, every vertex $u \in V(\mathrm{H})$ has out-degree at least 1. Hence $(u, v) \in E(\mathrm{H})$ for all $u \in V(\mathrm{H})$. Since v has out-degree at least 1, this means that H contains a double edge, contradicting statement (1).

Now suppose that there exists an automorphism σ that is backward-adjacent to some endomorphism $f \neq \sigma$ (the case of forwards-adjacent is dual). Thus, without loss of generality we have $(\sigma, f) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$. Apply σ^{-1} to both sides of the edge to obtain that $(id_{\mathrm{H}}, g) \in E(\widehat{\mathrm{H}^{\mathrm{H}}})$ for some $g \neq id_{\mathrm{H}}$. By the preceding lemma, and the fact that H is retract-trivial, this means g is a constant map, contradicting the claim above. □

We use Lemma 3.3 to obtain the following structural result.

Theorem 3.4. *Let* H *be an endo-trivial reflexive digraph with at least three vertices. Then every polymorphism of* H *is essentially unary.*

Proof. Since H is endo-trivial, H is retract-trivial. Hence, by Lemma 3.3, H is strongly connected, and furthermore the automorphisms of H are isolated vertices of $\widehat{\mathrm{H}^{\mathrm{H}}}$. As H is endo-trivial, this means that $\widehat{\mathrm{H}^{\mathrm{H}}}$ is the disjoint union of a copy of H that corresponds to the constant maps and a set of isolated vertices, one for each automorphism of H. Suppose for a contradiction that there exists an an n-ary polymorphism f of H which is not essentially unary. We may without loss of generality assume that f depends on all of its n variables, where $n \geq 2$. By Lemma 2.2, the mapping $F : \mathrm{H}^{n-1} \to \widehat{\mathrm{H}^{\mathrm{H}}}$ defined by $F(x_1, \ldots, x_{n-1})(y) = f(x_1, \ldots, x_{n-1}, y)$ is a homomorphism. Since H is strongly connected, so is H^{n-1}, and hence so is the image of F in $\widehat{\mathrm{H}^{\mathrm{H}}}$. Thus this image is either contained in the component of constants, in which case f does not depend on its last variable, else it is a singleton, in which case f does not depend on any of its first $n - 1$ variables. □

Combining Theorems 3.1 and 3.4 yields the main result of this section.

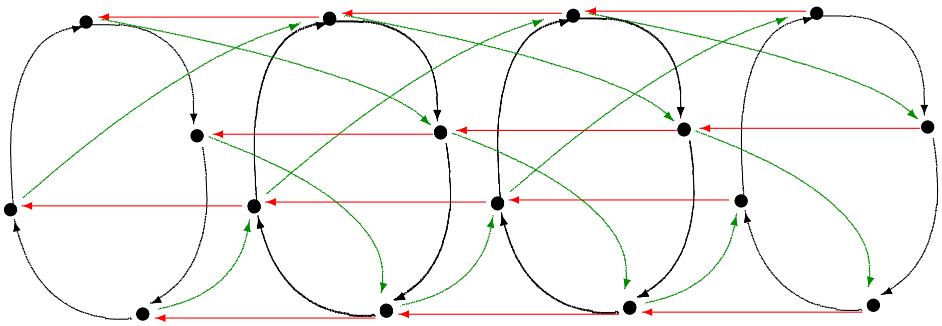

Fig. 5. The gadget Cyl^*_m in the case $m := 4$ (self-loops are not drawn). We usually visualise the right-hand copy of DC^*_4 as the "bottom" copy and then we talk about vertices "above" and "below" according to the red arrows.

Corollary 3.5. *If* H *is an endo-trivial reflexive digraph on at least three vertices, then* Surjective H-Colouring *is* NP*-complete.*

Let DC^*_k denote the reflexive directed cycle on k vertices, which is readily seen to be endo-trivial. Corollary 3.5 yields the following dichotomy for reflexive directed cycles after noting that Surjective DC^*_k-Colouring is trivial for $k \leq 2$.

Corollary 3.6. Surjective DC^*_k-Colouring *is in* L *if* $k \leq 2$ *and* NP*-complete if* $k \geq 3$.

In the next section though we give a combinatorial NP-hardness proof for Surjective H-Colouring whenever H is *any* non-transitive reflexive tournament. As DC^*_3 is such a digraph, this proof also can be used for the case $\mathrm{H} = \mathrm{DC}^*_3$. However, it does not extend to Surjective DC^*_k-Colouring for $k \geq 4$.

4 A DICHOTOMY FOR REFLEXIVE TOURNAMENTS

In this section we prove our main result, namely a dichotomy of Surjective H-Colouring for reflexive tournaments H by showing that transitivity is the crucial property for tractability. In the next subsections we prove that Surjective H-Colouring is NP-complete when H is a non-transitive tournament.

4.1 Two Elementary Lemmas

It is well-known that every strongly connected tournament has a directed Hamilton cycle [Camion 1959]. Hence we derive the following corollary to Lemmas 2.5 and 3.3 Part 2.

Lemma 4.1. *If* H *is a reflexive tournament that is endo-trivial, then* H *contains a directed Hamilton cycle.*

We will also need the following lemma.

Lemma 4.2. *If* H *is a reflexive tournament that is endo-trivial, then any homomorphic image of* H *of size* $1 < n < |V(\mathrm{H})|$ *possesses a double edge.*

Proof. Suppose H has a homomorphic image of size $1 < n < |V(\mathrm{H})|$ without a double edge. By looking at the equivalence classes of vertices identified in the homomorphic image, we can deduce a non-trivial retraction, namely by mapping each of the vertices in an equivalence class to any particular one of them. □

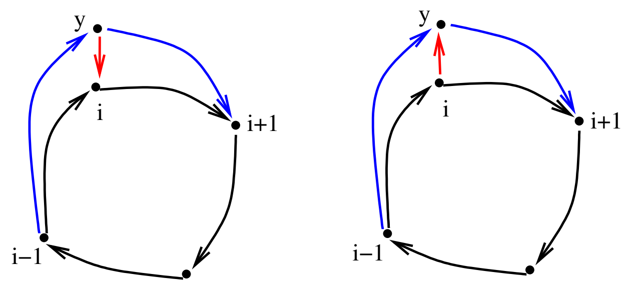

Fig. 6. A diagram of the two cases of Lemma 4.4 (self-loops are not drawn).

4.2 The NP-Hardness Gadget

We now introduce the gadget Cyl^*_m drawn in Figure 5. We take m disjoint copies of the directed m-cycle DC^*_m arranged in a cylindrical fashion so that there is an edge from i in the jth copy to i in the $j+1$th copy (drawn in red), and an edge from i in the $j+1$th copy to $i+1$ in the jth copy (drawn in green). We consider DC^*_m to have vertices $\{1,\dots,m\}$. A key role will be played by Hamilton cycles HC_m in a strongly connected reflexive tournament on m vertices. We consider this cycle also labelled $\{1,\dots,m\}$, in order to attach it to the gadget Cyl^*_m. The gadget Cyl^*_m is an alteration of a gadget that appears in [Feder and Hell 1998] for proving that List H-Colouring is NP-complete when H is an undirected cycle on at least four vertices, but our proof is very different.

The following lemma follows from induction on the copies of DC^*_m, since a reflexive tournament has no double edges.

Lemma 4.3. *In any homomorphism h from Cyl^*_m, with bottom cycle DC^*_m, to a reflexive tournament, if $|h(\mathrm{DC}^*_m)| = 1$, then $|h(\mathrm{Cyl}^*_m)| = 1$.*

We will use another property, denoted (†), of Cyl^*_m, which is that the retractions from Cyl^*_m to its bottom copy of DC^*_m, once propagated through the intermediate copies, induce on the top copy precisely the set of automorphisms of DC^*_m. That is, the top copy of DC^*_m is mapped isomorphically to the bottom copy, and all such isomorphisms may be realised. The reason is that in such a retraction, the $(j+1)$th copy may either map under the identity to the jth copy, or rotate one edge of the cycle clockwise, and Cyl^*_m consists of sufficiently many (namely m) copies of DC^*_m.

Now let H be a reflexive tournament that contains a subtournament H_0 on m vertices that is endo-trivial. By Lemma 4.1, we find that H_0 contains at least one directed Hamilton cycle HC_0. Define $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0,\mathrm{HC}_0])$ as follows. Begin with H and add a copy of the gadget Cyl^*_m, where the bottom copy of DC^*_m is identified with HC_0, to build a digraph $\mathrm{F}(\mathrm{H}_0,\mathrm{HC}_0)$. Now ask, for some $y \in V(\mathrm{H})$ whether there is a retraction r of $\mathrm{F}(\mathrm{H}_0,\mathrm{HC}_0)$ to H so that some vertex x in the top copy of DC^*_m in Cyl^*_m is such that $r(x) = y$. Such vertices y comprise the set $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0,\mathrm{HC}_0])$.

Remark 1. If x belongs to some copy of DC^*_m that is not the top copy, we can find a vertex x' in the top copy of DC^*_m and a retraction r' from $\mathrm{F}(\mathrm{H}_0,\mathrm{HC}_0)$ to H with $r'(x') = r(x) = y$, namely by letting r' map the vertices of higher copies of DC^*_m to the image of their corresponding vertex in the copy that contains x. In particular this implies that $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0,\mathrm{HC}_0])$ contains $V(\mathrm{H}_0)$.

We note that the set $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0,\mathrm{HC}_0])$ is potentially dependent on which Hamilton cycle in H_0 is chosen. We now prove that $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0,\mathrm{HC}_0]) = V(\mathrm{H})$ if H retracts to H_0.

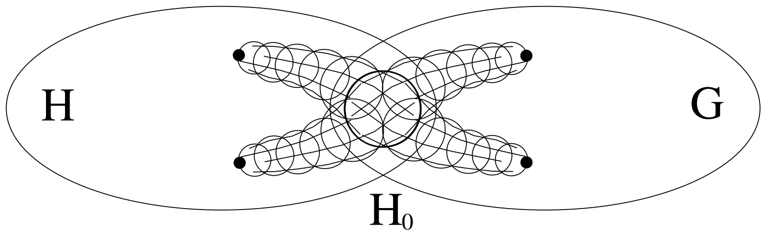

Fig. 7. A stylised depiction of the construction in Base Case I. The central circle is the Hamilton cycle and the eccentric circles emanating thereout are the gadgets Cyl^*_m.

Lemma 4.4. *If* H *is a reflexive tournament that retracts to a subtournament* H_0 *with Hamilton cycle* HC_0*, then* $\mathrm{Spill}_m(H[H_0, HC_0]) = V(H)$.

Proof. let $y \in V(H) \setminus V(H_0)$. We need to prove that there exists a retraction r from $F(H_0, HC_0)$ to H with $r(x) = y$ for some vertex x in the top copy of DC^*_m in Cyl^*_m. Let h be a retraction of H to H_0. Suppose $h(y) = i$. We observe that both $(i-1, y)$ and $(y, i+1)$ are edges of $E(H)$. However, we might have either (y, i) or (i, y) and distinguish between these two cases; see also Figure 6.

First suppose that $(y, i) \in E(H)$. Then we retract the gadget Cyl^*_m associated with y in the following fashion. Using property (†) we turn successive copies of DC^*_m in such a way as necessary to ensure that the vertex directly below y is at position i; for a diagrammatic description of what means "below", see Figure 5. Ih the top copy of DC^*_m we map the ith vertex to y and the jth vertex ($j \neq i$) to j in H_0.

Now suppose that $(i, y) \in E(H)$.) Then we retract the gadget Cyl^*_m associated with y in the following fashion. Using property (†) we turn successive copies of DC^*_m in such a way as necessary to ensure that the vertex directly below y is at position $i-1$; wherein with the last copy of DC^*_m we map the ith vertex to y and the jth vertex ($j \neq i$) to j in H_0.

Note that, in both cases, all of the vertices of Cyl^*_m, except one, are mapped to H_0. □

4.3 Two Base Cases

Recall that if H is an endo-trivial tournament, then Surjective H-Colouring is NP-complete due to Corollary 3.5. However H may not be endo-trivial. We will now show how to deal with the case where H is not endo-trivial but retracts to an endo-trivial subtournament. For doing this we use the above gadget, but we need to distinguish between two different cases.

Lemma 4.5 (Base Case I.). *Let* H *be a reflexive tournament that retracts to an endo-trivial subtournament* H_0 *with Hamilton cycle* HC_0*. Assume that* H *retracts to* H'_0 *for every isomorphic copy* $H'_0 = i(H_0)$ *of* H_0 *in* H *with* $\mathrm{Spill}_m(H[H'_0, i(HC_0)]) = V(H)$*. Then* H_0*-Retraction can be polynomially reduced to* Surjective H-Colouring.

Proof. Let G be an instance of H_0-Retraction. We build an instance G'' of Surjective H-Colouring in the following fashion. First, take a copy of H together with G and build G' by identifying these on the copy of H_0 that they both possess as a subgraph. Let m be the size of H_0 and consider its Hamilton cycle HC_0. We build G'' from G' by augmenting a new copy of Cyl^*_m for every vertex $v \in V(G') \setminus V(H_0)$. Vertex v is to be identified with any vertex in the top copy of DC^*_m in Cyl^*_m and the bottom copy of DC^*_m is to be identified with HC_0 in H_0 according to the identity function. See Figure 7 for an example. We claim that G retracts to H_0 if and only if there exists a surjective homomorphism from G'' to H.

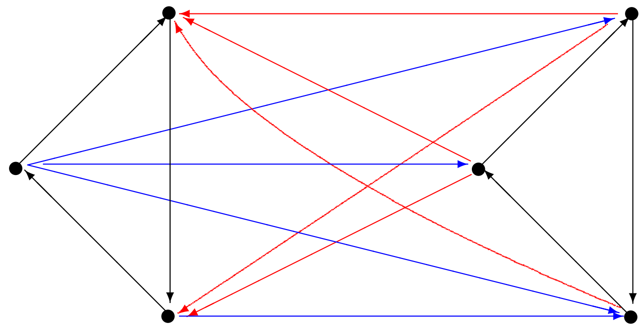

Fig. 8. An interesting tournament H on six vertices (self-loops are not drawn). This tournament does not retract to the DC_3^* on the left-hand side, yet $\mathrm{Spill}_3(\mathrm{H}[\mathrm{DC}_3^*, \mathrm{DC}_3]) = V(\mathrm{H})$.

First suppose that G retracts to H_0. Let h be a retraction from G to H_0. We extend h as follows. First we map the copy of H in G'' to itself in H by the identity. This will ensure surjectivity. We then map the various copies of Cyl_m^* in G''. This is always possible: because H retracts to H_0, we have $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0, \mathrm{HC}_0]) = V(\mathrm{H})$ due to Lemma 4.4. Hence, if $h(x) = y$ for two vertices $x \in V(\mathrm{G}') \setminus V(\mathrm{H}_0)$ and $y \in V(\mathrm{H})$, we can always find a retraction of the graph $\mathrm{F}(\mathrm{H}_0, \mathrm{HC}_0)$ to H that maps x to y, and we mimic this retraction on the corresponding subgraph in G''. The crucial observation is that this can be done independently for each vertex in $V(\mathrm{G}') \setminus V(\mathrm{H}_0)$, as two vertices of different copies of Cyl_m^* are only adjacent if they both belong to G'. This leads to a surjective homomorphism from G'' to H.

Now suppose there exists a surjective homomorphism h from G'' to H. If $|h(\mathrm{H}_0)| = 1$, then by Lemma 4.3, $|h(\mathrm{Cyl}_m^*)| = 1$ for all copies of Cyl_m^* in G''. Hence $|h(\mathrm{G}'')| = 1$ and h is not surjective, a contradiction. As $1 < |h(\mathrm{H}_0)| < m$ is not possible either due to Lemma 4.2, we find that $|h(\mathrm{H}_0)| = m$, and indeed h maps H_0 to a copy of itself in H which we will call $\mathrm{H}_0' = i(\mathrm{H}_0)$ for some isomorphism i.

We claim that $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0', i(\mathrm{HC}_0)]) = V(\mathrm{H})$. In order to see this, consider a vertex $y \in V(\mathrm{H})$. As h is surjective, there exists a vertex $x \in V(\mathrm{G}'')$ with $h(x) = y$. By construction, x belongs to some copy of DC_m^*, and thus also belongs to some copy of DC_m^* in $\mathrm{F}(\mathrm{H}_0, \mathrm{HC}_0)$. We can extend i^{-1} to an isomorphism from the copy of Cyl_m^* (which has $i(\mathrm{HC}_0)$ as its bottom cycle) in the graph $\mathrm{F}(\mathrm{H}_0', i(\mathrm{HC}_0))$ to the copy of Cyl_m^* (which has HC_0 as its bottom cycle) in the graph $\mathrm{F}(\mathrm{H}_0, \mathrm{HC}_0)$. We define a mapping r^* from $\mathrm{F}(\mathrm{H}_0', i(\mathrm{HC}_0))$ to H by $r^*(u) = h \circ i^{-1}(u)$ if u is on the copy of Cyl_m^* in $\mathrm{F}(\mathrm{H}_0', i(\mathrm{HC}_0))$ and $r^*(u) = u$ otherwise. We observe that $r^*(u) = u$ if $u \in V(\mathrm{H}_0')$ as h coincides with i on H_0. As H_0 separates the other vertices of the copy of Cyl_m^* from $V(\mathrm{H}) \setminus V(\mathrm{H}_0)$, in the sense that removing H_0 would disconnect them, this means that r^* is a retraction from $\mathrm{F}(\mathrm{H}_0', i(\mathrm{HC}_0))$ to H. We find that r^* maps $i(x)$ to $h \circ i^{-1}(i(x)) = h(x) = y$. Moreover, as x is in some copy of DC_m^* in $\mathrm{F}(\mathrm{H}_0, \mathrm{HC}_0)$, we have that $i(x)$ is in some copy of DC_m^* in $\mathrm{F}(\mathrm{H}_0', i(\mathrm{HC}_0))$. We may assume without loss of generality that $i(x)$ belongs to the top copy (cf. Remark 1). We conclude that y always belongs to $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0', i(\mathrm{HC}_0)])$ (cf. Remark 1).

As $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0', i(\mathrm{HC}_0)]) = V(\mathrm{H})$, we find, by assumption of the lemma, that there exists a retraction r from H to H_0'. Now $i^{-1} \circ r \circ h$ ia the desired retraction of G to H_0. □

We now need to deal with the situation in which we have an isomorphic copy $\mathrm{H}_0' = i(\mathrm{H}_0)$ of H_0 in H with $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}_0', i(\mathrm{HC}_0)]) = V(\mathrm{H})$, such that H does not retract to H_0' (see Figure 8 for an example). We cannot deal with this case in a direct matter and first show another base case. For this we need the following lemma and an extension of endo-triviality that we discuss afterwards.

Lemma 4.6. *Let* H *be a reflexive tournament, containing a subtournament* H_0 *so that any endomorphism of* H *that fixes* H_0 *is an automorphism. Then any endomorphism of* H *that maps* H_0 *to an isomorphic copy* $H_0' = i(H_0)$ *of itself is an automorphism of* H.

Proof. For contradiction, suppose there is an endomorphism h that maps H_0 to an isomorphic copy $H_0' = i(H_0)$ of itself that is not an automorphism of H. In particular, $|h(H)| < |V(H)|$. Choose h^{-1} in the following fashion. We let h^{-1} of $h(H_0)$ be the natural isomorphism of $h(H_0)$ to H_0 (that inverts the isomorphism given by h from H_0 to H_0'). Otherwise we choose h^{-1} arbitrarily, such that $h^{-1}(y) = x$ only if $h(x) = y$. Since H is a reflexive tournament, containing precisely one edge between distinct vertices, h^{-1} is an isomorphism. Moreover, $h^{-1} \circ h$ is an endomorphism of H that fixes H_0 and that is not an automorphism, a contradiction. ◻

Let H_0 be an induced subgraph of a digraph H. We say that the pair (H, H_0) is *endo-trivial* if all endomorphisms of H that fix H_0 are automorphisms.

Lemma 4.7 (Base Case II). *Let* H *be a reflexive tournament with a subtournament* H_0 *with Hamilton cycle* HC_0 *so that* (H, H_0) *and* H_0 *are endo-trivial and* $\text{Spill}_m(H[H_0, HC_0]) = V(H)$. *Then* H-Retraction *can be polynomially reduced to* Surjective H-Colouring.

Proof. Let G be an instance of H-Retraction. We build an instance G′ of Surjective H-Colouring in the following fashion. First we build G′ from G by augmenting a new copy of Cyl^*_m for every vertex $v \in V(G) \setminus V(H_0)$. Vertex v is to be identified with any vertex in the top copy of DC^*_m in Cyl^*_m and the bottom copy of DC^*_m is to be identified with HC_0 in H_0 according to the identity function. We claim that G retracts to H if and only if there exists a surjective homomorphism from G′ to H.

First suppose G retracts to H. Let r be a retraction from G to H. Then any extension of r from G to G′ is surjective. As $\text{Spill}_m(H[H_0, HC_0]) = V(H)$ and two vertices in different copies of Cyl^*_m are only adjacent if both of them are in G, we can in fact extend r to a surjective homomorphism from G′ to H.

Now suppose there exists a surjective homomorphism h from G′ to H. If $|h(H_0)| = 1$, then Lemma 4.3 tells us that $|h(\text{Cyl}^*_m)| = 1$ for all copies of Cyl^*_m in G′, and then we derive $|h(G')| = 1$, contradicting the surjectivity of h. Moreover, $1 < |h(H_0)| < m$ is not possible either due to Lemma 4.2. Thus, $|h(H_0)| = m$ and h maps H_0 to a copy of itself. As (H, H_0) is endo-trivial, Lemma 4.6 tells us that the restriction of h to H is an automorphism of H, which we call α. The required retraction from G to H is now given by $\alpha^{-1} \circ h$. ◻

4.4 Generalising the Base Cases

We now generalise the two base cases to more general cases via some recursive procedure. Afterwards we will show how to combine these two cases to complete our proof. We will first need a slightly generalised version of Lemma 4.6, which nonetheless has virtually the same proof.

Lemma 4.8. *Let* $H_2 \supset H_1 \supset H_0$ *be a sequence of strongly connected reflexive tournaments, each one a subtournament of the one before. Suppose that any endomorphism of* H_1 *that fixes* H_0 *is an automorphism. Then any endomorphism* h *of* H_2 *that maps* H_0 *to an isomorphic copy* $H_0' = i(H_0)$ *of itself also gives an isomorphic copy of* H_1 *in* $h(H_1)$.

Proof. For contradiction, suppose there is an endomorphism h of H_2 that maps H_0 to an isomorphic copy $H_0' = i(H_0)$ of itself that does not yield an isomorphic copy of H_1. In particular, $|h(H_1)| < |V(H_1)|$. We proceed as in the proof of the Lemma 4.6. Choose h^{-1} in the following fashion. We let h^{-1} of $h(H_0)$ be the natural isomorphism of $h(H_0)$ to H_0

(that inverts the isomorphism given by h from H_0 to H'_0). Otherwise we choose h^{-1} arbitrarily, such that $h^{-1}(y) = x$ only if $h(x) = y$. Since H_2 is a reflexive tournament, h^{-1} is an isomorphism. And $h^{-1} \circ h$ is an endomorphism of H_2 that fixes H_0 that does not yield an isomorphic copy of H_1 in $h(H_1)$, a contradiction. □

The following two lemmas generalize Lemmas 4.5 and 4.7.

Lemma 4.9 (General Case I). *Let* $H_0, H_1, \dots, H_k, H_{k+1}$ *be reflexive tournaments, the first* k *of which have Hamilton cycles* $HC_0, HC_1, \dots, HC_k$, *respectively, so that* $H_0 \subseteq H_1 \subseteq \cdots \subseteq H_k \subseteq H_{k+1}$. *Assume that* $H_0, (H_1, H_0), \dots, (H_k, H_{k-1})$ *are endo-trivial and that*

$$\begin{array}{lcl} \mathrm{Spill}_{a_0}(H_1[H_0, HC_0]) & = & V(H_1) \\ \mathrm{Spill}_{a_1}(H_2[H_1, HC_1]) & = & V(H_2) \\ \vdots & \vdots & \vdots \\ \mathrm{Spill}_{a_{k-1}}(H_k[H_{k-1}, HC_{k-1}]) & = & V(H_k). \end{array}$$

Moreover, assume that H_{k+1} *retracts to* H_k *and also to every isomorphic copy* $H'_k = i(H_k)$ *of* H_k *in* H_{k+1} *with* $\mathrm{Spill}_{a_k}(H_{k+1}[H'_k, i(HC_k)]) = V(H_{k+1})$. *Then* H_k-Retraction *can be polynomially reduced to* Surjective H_{k+1}-Colouring.

Proof. Let G be an instance of H_k-Retraction. We will build an instance G'' of Surjective H_{k+1}-Colouring in the following fashion. First, take a copy of H_{k+1} together with G and build G' by identifying these on the copy of H_k that they both possess as a subgraph. We now build G'' as follows. First we augment G' with a new copy of $\mathrm{Cyl}^*_{a_k}$ for every vertex $v \in V(G') \setminus V(H_k)$. Vertex v is to be identified with any vertex in the top copy of $\mathrm{DC}^*_{a_k}$ in $\mathrm{Cyl}^*_{a_k}$, and the bottom copy of $\mathrm{DC}^*_{a_k}$ is to be identified with HC_k according to the identity function. Then, for each $i \in [k+1]$, and $v \in V(H_i) \setminus V(H_{i-1})$, add a copy of $\mathrm{Cyl}^*_{a_{i-1}}$, where v is identified with any vertex in the top copy of $\mathrm{DC}^*_{a_{i-1}}$ in $\mathrm{Cyl}^*_{a_{i-1}}$ and the bottom copy of DC^*_{i-1} is to be identified with H_{i-1} according to the identity map of $\mathrm{DC}^*_{a_{i-1}}$ to HC_{i-1}. We claim that G retracts to H_k if and only if there exists a surjective homomorphism from G'' to H_{k+1}.

First suppose that G retracts to H_k. Let h be a retraction from G to H_k. Extend h mapping H_{k+1} according to the identity to ensure the final mapping is surjective. Finally, we map the various copies of $\mathrm{Cyl}^*_{a_{i-1}}$ in G'' in any suitable fashion, which will always exist due to our assumptions and the fact that $\mathrm{Spill}_{a_k}(H_{k+1}[H_k, HC_k]) = V(H_{k+1})$, which follows from our assumption that H_{k+1} retracts to H_k and Lemma 4.4.

Now suppose that if there exists a surjective homomorphism h from G'' to H_{k+1}. Suppose that $|h(H_0)| = 1$. Then $|h(\mathrm{Cyl}^*_{a_0})| = 1$ by Lemma 4.3. Now we follow the chain of spills to deduce that $|h(H_{k+1})| = 1$, which is not possible. Now, $1 < |h(H_0)| < a_0$ is not possible either due to Lemma 4.2. Thus, $|h(H_0)| = |V(H_0)|$ and indeed h maps H_0 to a copy of itself in H_{k+1} which we will call $H'_0 = i(H_0)$. We now apply Lemma 4.8 as well as our assumed endo-trivialities to derive that h in fact maps H_k by the isomorphism i to a copy of itself in H_{k+1} which we will call H'_k. Since h is surjective, we can deduce that $\mathrm{Spill}_{a_k}(H_{k+1}[H'_k, i(HC_k)]) = V(H_{k+1})$ in the same way as in the proof of Lemma 4.5. and so there exists a retraction r from H_{k+1} to H'_k. Now $i^{-1} \circ r \circ h$ gives the desired retraction of G'' to H_k. □

Lemma 4.10 (General Case II). *Let* $H_0, H_1, \dots, H_k, H_{k+1}$ *be reflexive tournaments, the first* $k+1$ *of which have Hamilton cycles* $HC_0, HC_1, \dots, HC_k$, *respectively, so that* $H_0 \subseteq H_1 \subseteq \cdots \subseteq H_k \subseteq H_{k+1}$. *Suppose that* $H_0, (H_1, H_0), \dots,$

$(H_k, H_{k-1}), (H_{k+1}, H_k)$ *are endo-trivial and that*

$$\begin{array}{rcl} \mathrm{Spill}_{a_0}(H_1[H_0, HC_0]) & = & V(H_1) \\ \mathrm{Spill}_{a_1}(H_2[H_1, HC_1]) & = & V(H_2) \\ \vdots & \vdots & \vdots \\ \mathrm{Spill}_{a_{k-1}}(H_k[H_{k-1}, HC_{k-1}]) & = & V(H_k) \\ \mathrm{Spill}_{a_k}(H_{k+1}[H_k, HC_k]) & = & V(H_{k+1}) \end{array}$$

Then H_{k+1}-Retraction *can be polynomially reduced to* Surjective H_{k+1}-Colouring.

Proof. Let G be an instance of H_{k+1}-Retraction. We build an instance G′ of Surjective H_{k+1}-Colouring in the following fashion. Build G′ by, for each $i \in [k+1]$, and $v \in V(H_i) \setminus V(H_{i-1})$, adding a copy of $\mathrm{Cyl}^*_{a_{i-1}}$, where v is identified with any vertex in the top copy of $\mathrm{DC}^*_{a_{i-1}}$ in $\mathrm{Cyl}^*_{a_{i-1}}$ and the bottom copy of DC^*_{i-1} is to be identified with H_{i-1} according to the identity map of $\mathrm{DC}^*_{a_{i-1}}$ to HC_{i-1}. We claim that G retracts to H_{k+1} if and only if there exists a surjective homomorphism from G′ to H_{k+1}.

First suppose that G retracts to H_{k+1}. Let h be a retraction from G to H_{k+1}. Then we can extend h by mapping G′ in some suitable fashion, which is possible due to the spill assumptions.

Now suppose that there exists a surjective homomorphism h from G′ to H_{k+1}. Suppose that $|h(H_0)| = 1$. Then $|h(\mathrm{Cyl}^*_{a_0})| = 1$ by Lemma 4.3. Now we follow the chain of spills to deduce that $|h(H_{k+1})| = 1$, a contradiction. Now, $1 < |h(H_0)| < a_0$ is not possible either due to Lemma 4.2. Thus, $|h(H_0)| = |V(H_0)|$ and indeed h maps H_0 to a copy of itself in H_{k+1} which we will call $H'_0 = i(H_0)$. We now apply Lemma 4.8 as well as our assumed endo-trivialities to derive that h in fact maps H_k by the isomorphism i to a copy of itself in H_{k+1}, which we will call H'_k. Now we can deduce, via Lemma 4.6, that $h(H_{k+1})$ is an automorphism of H_{k+1}, which we call α. The required retraction from G′ to H_{k+1} is now given by $\alpha^{-1} \circ h$. □

4.5 Final Steps for Hardness for Non-Transitive Reflexive Tournaments

We first prove, by using the lemmas from Section 4.4, that Surjective H-Colouring is NP-complete if H is a non-transitive reflexive tournament that is strongly connected. For our discourse it is not necessary to know precisely what is a Taylor operation, but we will use the following result.

Theorem 4.11 ([Bulatov et al. 2005; Larose and Zádori 2005]). *Let* H *be a finite structure so that the idempotent polymorphisms of* H *do not contain any Taylor operations. Then* H-Retraction *is* NP-*complete.*

Corollary 4.12. *Let* H *be a strongly connected reflexive tournament. Then* Surjective H-Colouring *is* NP-*complete.*

Proof. As H is is a strongly connected reflexive tournament, which has more than one vertex by our definition, H is not transitive. Note that H-Retraction is NP-complete, since non-transitive reflexive tournaments do not have any Taylor polymorphisms [Larose 2006], following Theorem 4.11. Thus, if H is endo-trivial, the result follows from Lemma 4.5 (note that we could also have used Corollary 3.5).

Suppose H is not endo-trivial. Then, by Lemma 2.5, H is not retract-trivial either. This means that H has a non-trivial retraction to some subtournament H_0. We may assume that H_0 is endo-trivial, as otherwise we will repeat the argument until we find a retraction from H to an endo-trivial subtournament.

Suppose that H retracts to all isomorphic copies $H'_0 = i(H_0)$ of H_0 within it, except possibly those for which $\mathrm{Spill}_m(H[H'_0, i(HC_0)]) \neq V(H)$. Then the result follows from Lemma 4.5. So there is a copy $H'_0 = i(H'_0)$ to which H does not retract for which $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}'_0, i(\mathrm{HC}_0)]) = V(\mathrm{H})$. If $(\mathrm{H}, \mathrm{H}'_0)$ is endo-trivial, the result follows from Lemma 4.7. Thus we assume $(\mathrm{H}, \mathrm{H}'_0)$ is not endo-trivial and we deduce the existence of $\mathrm{H}'_0 \subset \mathrm{H}_1 \subset \mathrm{H}$ (H_1 is strictly between H and H'_0) so that $(\mathrm{H}_1, \mathrm{H}'_0)$ is endo-trivial and H retracts to H_1. Now we are ready to break out. Either H retracts to all isomorphic copies of $\mathrm{H}'_1 = i(\mathrm{H}_1)$ in H, except possibly for those so that $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}'_1, i(\mathrm{HC}_1)]) \neq V(\mathrm{H})$, and we apply Lemma 4.9, or there exists a copy H'_1, with $\mathrm{Spill}_m(\mathrm{H}[\mathrm{H}'_1, i(\mathrm{HC}_1)]) = V(\mathrm{H})$, to which it does not retract. Then H'_1 contains H''_0 a copy of H''_0 so that $(\mathrm{H}'_1, \mathrm{H}''_0)$ and H''_0 are endo-trivial. We now continue iterating this method, which will terminate because our structures are getting strictly larger. □

In order to deal with reflexive tournaments that are not strongly connected we need the following strengthened version of Corollary 4.12.

Corollary 4.13. *Let* H *be a strongly connected reflexive tournament. Then* Surjective H-Colouring *is* NP*-complete even for strongly connected digraphs.*

Proof. We need to argue that the instances of Surjective H-Colouring that we have constructed before can be assumed to be strongly connected. Noting that H and the gadgets Cyl^*_m are strongly connected, this is clear once we can assume the inputs to our Retraction problems are strongly connected. For H′-Retraction, where H′ is a strongly connected reflexive tournament, we can surely assume our inputs are strongly connected. If they were not, then we add individual directed paths of length $|V(\mathrm{H}')|$ between the relevant vertices. This will not affect the truth of an instance, in the sense that the input without the directed paths is a yes-instance if and only if the input with the directed paths is a yes-instance, and the result follows. □

When G is a reflexive tournament, we may break it up into strongly connected components $\mathrm{G}(1), \dots, \mathrm{G}(k)$ so that, for all $i < j \in [k]$, for all $x \in \mathrm{G}(i)$ and for all $y \in \mathrm{G}(j)$, $(x, y) \in E(\mathrm{G})$. This is the standard decomposition, inducing a standard order on the connected components, that we will use.

We now prove our main hardness result.

Theorem 4.14. *Let* H *be a non-transitive reflexive tournament. Then* Surjective H-Colouring *is* NP*-complete.*

Proof. For strongly connected tournaments, the result follows from Corollary 4.12. Let H instead have $k > 1$ strongly connected components $\mathrm{H}(1), \dots, \mathrm{H}(k)$. Since H is not transitive, one of these strongly connected components, $\mathrm{H}(i)$, must be of size greater than 1, whereupon we know from Corollary 4.13 that Surjective $\mathrm{H}(i)$-Colouring is NP-complete, even when restricted to strongly connected inputs.

Let us reduce Surjective $\mathrm{H}(i)$-Colouring to Surjective H-Colouring by taking a strongly connected input G for the former and building G′ by adding a copy of H restricted to $V(\mathrm{H}(1)), \dots, V(\mathrm{H}(i-1))$, where every vertex here has an edge to every vertex of G, and adding a copy of H restricted to $V(\mathrm{H}(i+1)), \dots, V(\mathrm{H}(k))$, where every vertex there has an edge from every vertex of G. Note that G′ has k strongly connected components $\mathrm{G}'(1), \dots, \mathrm{G}'(k)$, where $\mathrm{G}'(h)$ is isomorphic to $\mathrm{H}'(h)$ for $h = 1, \dots, k$, $h \neq i$. We claim that there exists a surjective homomorphism from G to $\mathrm{H}(i)$ if and only if there exists a surjective homomorphism from G′ to H.

(Forwards.) Map the additional vertices in G′ in the obvious fashion (by the "identity") to extend a surjective homomorphism from G to H_i so that it is surjective from G′ to H.

(Backwards.) In any surjective homomorphism s from G′ to H, we have that $s(V(\mathrm{G}'(i))) \subseteq V(\mathrm{H}(i))$ for $i = 1, \dots k$. □

Corollary 4.15. *Let* H *be a reflexive tournament. If* H *is transitive, then* Surjective H-Colouring *is in* NL*; otherwise it is* NP*-complete.*

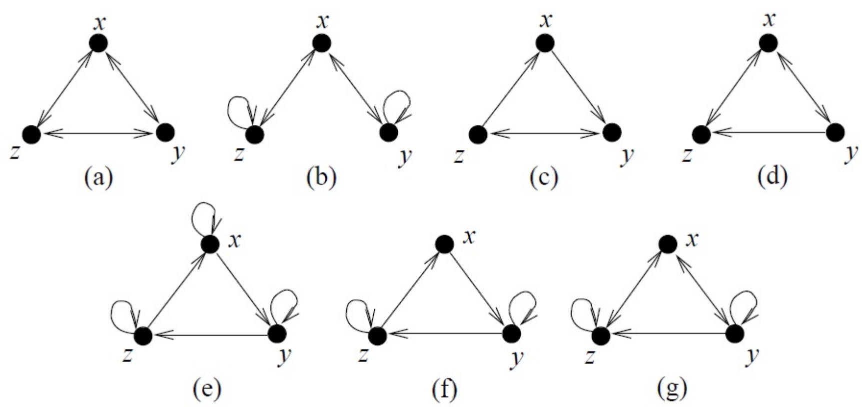

Fig. 9. The sporadic cases from Figure 1 in [Feder et al. 2006]. Cases (a)–(f) are NP-complete cases, whereas Case (g) is polynomial-time solvable (see also Theorem 5.1).

Proof. For the transitive case we can say that H-Retraction is in NL from [Dalmau and Krokhin 2008], since H enjoys the ternary median operation as a polymorphism (this has been observed, inter alia, in [Larose 2006]). It follows of course that Surjective H-Colouring is in NL also. The non-transitive case follows from Theorem 4.14. □

5 DIGRAPHS WITH A MOST THREE VERTICES

In this section we prove the following result.

Theorem 5.1. *Let* H *be a partially reflexive digraph of size at most* 3. *Then* Surjective H-Colouring *is polynomially equivalent to* H-Retraction. *In particular, it is always in* P *or is* NP*-complete, with the* NP*-complete cases being precisely those drawn in Figure 9 (a)–(f).*

We are not aware of a published classification for H-Retraction, when H is a partially reflexive digraph of size at most 3, though we know of one for List H-Colouring from [Feder et al. 2006]. Our starting point is therefore Theorem 3.1 from [Feder et al. 2006], and in particular the sporadic digraphs drawn in Figure 9 that are precisely those for which List H-Colouring is not in P. Bearing in mind that membership in P for List H-Colouring gives this a fortiori for H-Retraction and Surjective H-Colouring, these sporadic digraphs are the only ones we need to consider. Note that the principal objects of study in [Feder et al. 2006] are *trigraphs* and the reference to *complement* in that paper's Theorem 3.1 is to trigraph complement, which is different from the various notions of (di)graph complement.

We will need two lemmas to deal with two of the cases from Figure 9.

Lemma 5.2. *Let* H *be the digraph from Figure 9 (f). All polymorphisms of* H *are essentially unary.*

Proof. Recall the self-map digraph H^H, with vertices the self-maps of H, and an edge $(f, g) \in E(H^H)$ between self-maps f and g if and only if for every edge $(x, y) \in E(H)$, we have that $(f(x), g(y)) \in E(H)$. In the self-map digraph, the endomorphisms are precisely the looped vertices. Without fear of confusion, we denote the constant maps by 0, 1 and 2. We denote the identity map by id.

Consider the homomorphism $\varphi : \mathrm{H}^n \to \mathrm{H}^{\mathrm{H}}$ induced by an $(n+1)$-ary polymorphism of H where $n \geq 1$; suppose for a contradiction that the polymorphism depends on all its variables. Then clearly the image of φ contains at least two elements, and at least one of these elements is not a constant map.

Claim 0. *The only loops in* H^{H} *are the constant maps* 1 *and* 2 *and the identity.*

Proof of Claim 0: let f be an endomorphism of H. Then f maps loops to loops. Since $(1,2) \in E(\mathrm{H})$, f can either map both to 1, both to 2 , or fix both. Clearly if it fixes both f is the identity. Otherwise, if $f(1) = f(2) = i$, then $(f(0), i), (i, f(0)) \in E(\mathrm{H})$ so $f(0) = i$.

Claim 1. *(a) There are no edges in* H^{H} *between the identity and the constant maps; (b) if* $(2, f), (f, 1) \in E(\mathrm{H}^{\mathrm{H}})$*, then* $f = 0$*; (c) if* c *is a constant map such that* $(f, c), (c, f) \in E(\mathrm{H}^{\mathrm{H}})$ *then* $f = c$.

Proof of Claim 1: $(2, f) \in E(\mathrm{H}^{\mathrm{H}})$ implies $(2, f(i)) \in E(\mathrm{H})$ for all i, so $f(i) \in \{0, 2\}$ for all i. Similarly, $(f, 1) \in E(\mathrm{H}^{\mathrm{H}})$ implies $f(i) \in \{0, 1\}$ for all i so (b) follows immediately. To prove (a), observe that any in- or out-neighbour of a constant cannot be surjective since no constant has in- or out-neighbourhood of size 3. For (c), argue again as we just did: $(1, f) \in E(\mathrm{H}^{\mathrm{H}})$ means $f(i) \in \{1, 2\}$ for all i, and $(f, 1) \in E(\mathrm{H}^{\mathrm{H}})$ implies $f(i) \in \{0, 1\}$ for all i and hence $f = 1$; the proof for $c = 2$ is identical.

Claim 2. *If* $(f, id), (id, f) \in E(\mathrm{H}^{\mathrm{H}})$ *then* $f = id$.

Proof of Claim 2: Assume that $(f, id), (id, f) \in E(\mathrm{H}^{\mathrm{H}})$. If i is a loop then $(i, i) \in E(\mathrm{H})$ implies $(f(i), i), (i, f(i)) \in E(\mathrm{H})$ thus $f(i) = i$. Now $(0, 1) \in E(\mathrm{H})$ implies $(f(0), 1) \in E(\mathrm{H})$, and $(2, 0) \in E(\mathrm{H})$ implies $(2, f(0)) \in E(\mathrm{H})$ so $f(0) = 0$.

Obviously the loops of H^n are precisely the tuples whose coordinates belong to $\{1, 2\}$. Call the subdigraph induced by these vertices W, and notice it is weakly connected. Then the homomorphism φ must map W onto a weakly connected digraph consisting of loops only. By Claim 1 (a), either (i) W is entirely mapped to $\{id\}$ or (ii) W is mapped by φ to the set $\{1, 2\}$.

Choose any tuple X of H^n containing a coordinate equal to 0. Then there exist $Y, Z \in W$ such that $(X, Z), (Z, Y), (Y, X) \in E(\mathrm{H}^n)$: indeed, let Z be obtained from X by replacing each 0 entry by 1, fixing all other coordinates, and let Y be obtained from X by replacing each 0 entry by 2, fixing all other coordinates. Thus in case (i) we get that $id = \varphi(Y), (\varphi(Y), \varphi(X)), (\varphi(X), \varphi(Z)) \in E(\mathrm{H}^{\mathrm{H}})$ and $\varphi(Z) = id$ so by Claim 2 φ maps all of H^n to id. In case (ii), we get that $\varphi(Y)$ and $\varphi(Z)$ belong to $\{1, 2\}$; since $(\varphi(Z), \varphi(Y)) \in E(\mathrm{H}^{\mathrm{H}})$ we must in fact have that $(\varphi(Z), \varphi(Y))$ belongs to $\{(1, 1), (1, 2), (2, 2)\}$; thus we get one of the following cases, namely $(1, \varphi(X)), (\varphi(X), 1) \in E(\mathrm{H}^{\mathrm{H}})$ or $(1, \varphi(X)), (\varphi(X), 2) \in E(\mathrm{H}^{\mathrm{H}})$ or $(2, \varphi(X)), (\varphi(X), 2) \in E(\mathrm{H}^{\mathrm{H}})$. By Claim 1 (b) and (c), $\varphi(X)$ must be one of the constant maps 0, 1 or 2.

Thus we conclude that φ either maps all of H^n to $\{id\}$, in which case our polymorphism cannot depend on its first n variables; or otherwise φ maps all of H^n to constant maps, and then the polymorphism does not depend on its last variable, also a contradiction.

□

An operation $t : D^k \to D$ is a *weak near-unanimity* operation if t satisfies

$$t(x, \ldots, x) = x, \text{ and}$$
$$t(y, x, \ldots, x) = t(x, y, x, \ldots, x) = \cdots = t(x, \ldots, x, y).$$

Lemma 5.3. *The digraph* H *from Figure 9 (g) admits a weak near-unanimity polymorphism.*

Proof. Let H be given over vertex-set $\{0, 1, 2\}$ with edge-set

$$\{(0, 0), (2, 2), (0, 1), (1, 0), (0, 2), (2, 0), (2, 1)\}.$$

We have found by computer the following ternary weak near-unanimity polymorphism.

$000 = 0$	$001 = 0$	$002 = 0$	$010 = 0$	$011 = 0$	$012 = 0$	$020 = 0$	$021 = 0$	$022 = 0$
$100 = 0$	$101 = 0$	$102 = 0$	$110 = 0$	$111 = 1$	$112 = 0$	$120 = 0$	$121 = 0$	$122 = 0$
$200 = 0$	$201 = 0$	$202 = 0$	$210 = 0$	$211 = 0$	$212 = 0$	$220 = 0$	$221 = 0$	$222 = 2$

We have used the excellent program of Miklós Maróti to find this polymorphism.[2] □

For the benefit of the interested reader, we note that a finite core possesses a weak near-unanimity polymorphism if and only if it possesses a Taylor polymorphism.

Proof of Theorem 5.1. As discussed, we only need to settle the complexity of H-Retraction and Surjective H-Colouring for the digraphs depicted in Figure 9, as for all other partially reflexive digraphs of size at most 3, List H-Colouring is in P [Feder et al. 2006].

When H is as depicted in Figure 9 (a), (c) or (d), then H is an irreflexive semicomplete digraph with more than one cycle and it is known from [Bang-Jensen et al. 1988] that H-Colouring is NP-hard, thus the same can be said for both H-Retraction and Surjective H-Colouring.

When H is the partially reflexive tree of Figure 9 (b), H-Retraction is known to be NP-hard from [Feder et al. 2010] and Surjective H-Colouring is known to be NP-hard from [Golovach et al. 2012].

When H is the reflexive tournament from Figure 9 (e), the result follows from Corollary 3.5 or Lemma 4.5. For the related case (f), we appeal to Lemma 5.2 (finishing via Corollary 3.5).

When H is as depicted in Figure 9 (g), membership in P for H-Retraction, and therefore Surjective H-Colouring, follows from Lemma 5.3, in light of the recent groundbreaking proofs of [Bulatov 2017; Zhuk 2017]. □

6 CONCLUSION

We have given the first significant classification results for Surjective H-Colouring where H comes from a class of digraphs (that are not graphs). To do this, we have developed both a novel algebraic method and a novel recursive combinatorial method. Below we discuss some directions for future research.

Let 3NRC be the hypergraph with vertex-set $\{r, g, b\}$ and hyperedge-set

$$\{r, g, b\}^3 \setminus \{(r, g, b), (r, b, g), (g, b, r), (g, r, b), (b, r, g), (b, g, r)\}.$$

Then 3-No-Rainbow-Colouring is the problem Surjective 3NRC-Colouring, in which one looks for a surjective colouring of the vertices, such that no hyperedge is rainbow-coloured (i.e. uses all colours). We recall that the complexity of this problem is open since it arose (under a different name) in [Král et al. 2006], see also Question 3 in [Bodirsky et al. 2012]. The Surjective DC_3^*-Colouring problem is the digraph problem most closely related to 3-No-Rainbow-Colouring. To explain this, when looking for digraphs with a similar character to 3NRC, we would insist at least that the automorphism group is transitive. This leaves just the reflexive and irreflexive directed 3-cycles and the reflexive and irreflexive 3-cliques, that is, 3-cycles with a double edge between every pair of vertices (admittedly, the cycles have only some of the automorphisms of the cliques). If H is the reflexive 3-clique, then H-Retraction and Surjective

[2]See: `http://www.math.u-szeged.hu~maroti/applets/GraphPoly.html`.

H-Colouring are trivial. If H is the irreflexive directed 3-cycle, then H has a majority polymorphism, which shows that H-Retraction, and thus Surjective H-Colouring (see Figure 1), can be solved in polynomial time [Bang-Jensen et al. 1988]. If H is the irreflexive 3-clique, then Surjective H-Colouring is NP-complete, as there exists a straightforward reduction from 3-Colouring. Hence $H = DC_3^*$ was indeed the only case for which determining the complexity of Surjective H-Colouring was not immediately obvious.

It would be great to extend our results to larger reflexive digraph classes. Reflexive digraphs with a double edge are not endo-trivial and further fail to be endo-trivial in the worse way, since Surjective DC_2^*-Colouring is nearly trivial. Thus, our methods are likely only to be applicable to reflexive oriented digraphs, that is, those without a double edge. On the way, a natural question arising is exactly which reflexive digraphs are endo-trivial?

Finally, there is the question as to whether the assumption of endo-triviality can be weakened to that of retract-triviality in Theorem 3.4. Endo-triviality is used right at the beginning of the proof to show that G^G is the disjoint union of a copy of G (the constant maps) and isolated automorphisms. We do not know if retract-triviality is here sufficient.

Acknowledgements

We thank Jan Bok for fruitful discussions and several reviewers for useful comments on drafts of this manuscript.

REFERENCES

Jørgen Bang-Jensen, Pavol Hell, and Gary MacGillivray. 1988. The Complexity of Colouring by Semicomplete Digraphs. *SIAM Journal on Discrete Mathematics* 1, 3 (1988), 281–298.

Manuel Bodirsky, Jan Kára, and Barnaby Martin. 2012. The Complexity of Surjective Homomorphism Problems - a Survey. *Discrete Applied Mathematics* 160, 12 (2012), 1680–1690.

Andrei A. Bulatov. 2011. Complexity of Conservative Constraint Satisfaction Problems. *ACM Transactions on Computational Logic* 12, 4 (2011), 24:1–24:66.

Andrei A. Bulatov. 2017. A Dichotomy Theorem for Nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*. 319–330.

Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. 2005. Classifying the Complexity of Constraints Using Finite Algebras. *SIAM J. Comput.* 34, 3 (2005), 720–742.

Paul Camion. 1959. Chemins et Circuits Hamiltoniens de Graphes Complets. *Comptes Rendus de l'Académie des Sciences Paris* 249 (1959), 2151–2152.

Hubie Chen. 2014. An Algebraic Hardness Criterion for Surjective Constraint Satisfaction. *Algebra Universalis* 72, 4 (2014), 393–401.

Víctor Dalmau and Andrei A. Krokhin. 2008. Majority Constraints have Bounded Pathwidth Duality. *European Journal of Combinatorics* 29, 4 (2008), 821–837.

Tomás Feder and Pavol Hell. 1998. List Homomorphisms to Reflexive Graphs. *Journal of Combinatorial Theory, Series B* 72, 2 (1998), 236–250.

Tomás Feder, Pavol Hell, and Jing Huang. 1999. List Homomorphisms and Circular Arc Graphs. *Combinatorica* 19, 4 (1999), 487–505.

Tomás Feder, Pavol Hell, and Jing Huang. 2003. Bi-Arc Graphs and the Complexity of List Homomorphisms. *Journal of Graph Theory* 42, 1 (2003), 61–80.

Tomás Feder, Pavol Hell, Peter Jonsson, Andrei A. Krokhin, and Gustav Nordh. 2010. Retractions to Pseudoforests. *SIAM Journal on Discrete Mathematics* 24, 1 (2010), 101–112.

Tomás Feder, Pavol Hell, and Kim Tucker-Nally. 2006. Digraph Matrix Partitions and Trigraph Homomorphisms. *Discrete Applied Mathematics* 154, 17 (2006), 2458–2469.

Tomás Feder and Moshe Y. Vardi. 1998. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM J. Comput.* 28, 1 (1998), 57–104.

Petr A. Golovach, Matthew Johnson, Barnaby Martin, Daniël Paulusma, and Anthony Stewart. 2017. Surjective H-Colouring: New Hardness Results. *In 13th Conference on Computability in Europe (CiE) 2017, Lecture Notes in Computer Science* 10307 (2017), 270–281.

Petr A. Golovach, Daniël Paulusma, and Jian Song. 2012. Computing Vertex-Surjective Homomorphisms to Partially Reflexive Trees. *Theoretical Computer Science* 457 (2012), 86–100.

Pavol Hell and Jaroslav Nešetřil. 2004. *Graphs and Homomorphisms*. Oxford University Press.

David Charles Hobby and Ralph McKenzie. 1988. *The Structure of Finite Algebras, Contemporary Mathematics*. American Mathematical Society.

Daniel Král, Jan Kratochvíl, Andrzej Proskurowski, and Heinz-Jürgen Voss. 2006. Coloring Mixed Hypertrees. *Discrete Applied Mathematics* 154, 4 (2006), 660–672.

Benoit Larose. 2006. Taylor Operations on Finite Reflexive Structures. *International Journal of Mathematics and Computer Science* 1, 1 (2006), 1–26.

Benoit Larose. 2017. Algebra and the Complexity of Digraph CSPs: a Survey. In *The Constraint Satisfaction Problem: Complexity and Approximability*. 267–285.

Benoit Larose, Cynthia Loten, and Claude Tardif. 2007. A Characterisation of First-Order Constraint Satisfaction Problems. *Logical Methods in Computer Science* 3, 4 (2007).

Benoit Larose, Barnaby Martin, and Daniël Paulusma. 2018. Surjective H-Colouring over Reflexive Digraphs. In *35th Symposium on Theoretical Aspects of Computer Science, STACS 2018, February 28 to March 3, 2018, Caen, France.* 49:1–49:14.

Benoit Larose and László Zádori. 2005. Finite Posets and Topological Spaces in Locally Finite Varieties. *Algebra Universalis* 52, 2 (2005), 119–136.

Miklós Maróti and László Zádori. 2012. Reflexive Digraphs with Near Unanimity Polymorphisms. *Discrete Mathematics* 312, 15 (2012), 2316–2328.

Barnaby Martin and Daniël Paulusma. 2015. The Computational Complexity of Disconnected Cut and 2K2-partition. *Journal of Combinatorial Theory, Series B* 111 (2015), 17–37.

Mark Siggers. 2017. personal communication. (2017).

Agnes Szendrei. 1994. Term Minimal Algebras. *Algebra Universalis* 32, 4 (1994), 439–477.

Narayan Vikas. 2002. Computational Complexity of Compaction to Reflexive Cycles. *SIAM J. Comput.* 32, 1 (2002), 253–280.

Narayan Vikas. 2004. Computational Complexity of Compaction to Irreflexive Cycles. *J. Comput. System Sci.* 68, 3 (2004), 473–496.

Narayan Vikas. 2005. A Complete and Equal Computational Complexity Classification of Compaction and Retraction to All Graphs with at most Four Vertices and Some General Results. *J. Comput. System Sci.* 71, 4 (2005), 406–439.

Narayan Vikas. 2013. Algorithms for Partition of Some Class of Graphs under Compaction and Vertex-Compaction. *Algorithmica* 67, 2 (2013), 180–206.

Narayan Vikas. 2017. Computational Complexity of Graph Partition under Vertex-Compaction to an Irreflexive Hexagon. In *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21-25, 2017 - Aalborg, Denmark.* 69:1–69:14.

Dmitriy Zhuk. 2017. A Proof of CSP Dichotomy Conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017.* 331–342.