The Complexity of Digraph Homomorphisms:
Local Tournaments, Injective Homomorphisms and Polymorphisms

by

Jacobus S. Swarts

BEng, Rand Afrikaans University, 1995
MEng, Rand Afrikaans University, 1999
BSc, Rand Afrikaans University, 2000
MSc, Rand Afrikaans University, 2001

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy

in the Department of Mathematics and Statistics

University of Victoria

The Complexity of Digraph Homomorphisms:
Local Tournaments, Injective Homomorphisms and Polymorphisms

by

Jacobus S. Swarts

BEng, Rand Afrikaans University, 1995
MEng, Rand Afrikaans University, 1999
BSc, Rand Afrikaans University, 2000
MSc, Rand Afrikaans University, 2001

Supervisory Committee

Dr. G. MacGillivray, Supervisor (Department of Mathematics and Statistics)

Dr. C. M. Mynhardt, Member (Department of Mathematics and Statistics)

Dr. U. Stege, Outside Member (Department of Computer Science)

Supervisory Committee

Dr. G. MacGillivray, Supervisor (Department of Mathematics and Statistics)

Dr. C. M. Mynhardt, Member (Department of Mathematics and Statistics)

Dr. U. Stege, Outside Member (Department of Computer Science)

Abstract

In this thesis we examine the computational complexity of certain digraph homomorphism problems. A homomorphism between digraphs, denoted by $f : G \to H$, is a mapping from the vertices of G to the vertices of H such that the arcs of G are preserved. The problem of deciding whether a homomorphism to a fixed digraph H exists is known as the H-colouring problem.

We prove a generalization of a theorem due to Bang-Jensen, Hell and MacGillivray. Their theorem shows that for every semicomplete digraph H, H-colouring exhibits a dichotomy: H-colouring is either polynomial time solvable or it is NP-complete. We show that the class of local tournaments also exhibit a dichotomy. The NP-completeness results are found using direct NP-completeness reductions, indicator and vertex (and arc) sub-indicator constructions. The polynomial cases are handled by appealing to a result of Gutjhar, Woeginger and Welzl: the $\underline{X}$-graft extension. We also provide a new proof of their result that follows directly from the consistency check. An unexpected result is the existence of unicyclic local tournaments with NP-complete homomorphism problems.

During the last decade a new approach to studying the complexity of digraph homomorphism problems has emerged. This approach focuses attention on so-called polymorphisms as a measure of the complexity of a digraph homomorphism problem. For a digraph H, a polymorphism of arity k is a homomorphism $f : H^k \to H$.

Certain special polymorphisms are conjectured to be the key to understanding H-colouring problems. These polymorphisms are known as weak near unanimity functions (WNUFs). A WNUF of arity k is a polymorphism $f : H^k \to H$ such that f is idempotent and $f(y, x, x, \ldots, x) = f(x, y, x, \ldots, x) = f(x, x, y, \ldots, x) = \cdots = f(x, x, x, \ldots, y)$. We prove that a large class of polynomial time H-colouring problems all have a WNUF. Furthermore we also prove some non-existence results for WNUFs on certain digraphs. In proving these results, we develop a vertex (and arc) sub-indicator construction as well as an indicator construction in analogy with the ones developed by Hell and Nešetřil. This is then used to show that all tournaments with at least two cycles do not admit a WNUF_k for $k > 1$. This furnishes a new proof (in the case of tournaments) of the result by Bang-Jensen, Hell and MacGillivray referred

to at the start. These results lend some support to the conjecture that WNUFs are the "right" functions for measuring the complexity of H-colouring problems.

We also study a related notion, namely that of an injective homomorphism. A homomorphism $f : G \to H$ is injective if the restriction of f to the in-neighbours of every vertex in G is an injective mapping. In order to classify the complexity of these problems we develop an indicator construction that is suited to injective homomorphism problems.

For this type of digraph homomorphism problem we consider two cases: reflexive and irreflexive targets. In the case of reflexive targets we are able to classify all injective homomorphism problems as either belonging to the class of polynomial time solvable problems or as being NP-complete. Irreflexive targets pose more of a problem. The problem lies with targets of maximum in-degree equal to two. Targets with maximum in-degree one are polynomial, while targets with in-degree at least three are NP-complete. There is a transformation from (ordinary) graph homomorphism problems to injective, in-degree two, homomorphism problems (a reverse transformation also exists). This transformation provides some explanation as to the difficulty of the in-degree two case. We nonetheless classify all injective homomorphisms to irreflexive tournaments as either being a problem in P or a problem in the class of NP-complete problems. We also discuss some upper bounds on the injective oriented irreflexive (reflexive) chromatic number.

Table of Contents

List of Tables

List of Figures

List of Decision Problems

List of Algorithms

Acknowledgements

I would like to thank the following people for making this thesis possible.

- My supervisor, Gary MacGillivray. If any one person played a larger role in getting me through this, he is the one. His infinite patience with my blank stares and dumb questions, his steady stream of good advice and ideas, and his encouragement in general, turned this from an impossible task into something that I could do. Not to mention the fact that he did all this while being chair. I can only hope that our theorem-to-coffee ratio is close to one.

- The other members of my committee: Kieka Mynhardt, who as a fellow South African, made my move here possible as well as enjoyable, not to mention her role as graduate advisor. Dr. Ulrike Stege and Dr. Richard Nowakowski for their willingness to act as outside member and external member, respectively.

- Jørgen Bang-Jensen of the University of Southern Denmark for suggesting the local tournament problem, getting us off to a good start with it and eventually reading the chapter on local tournaments.

- Ernie Cockayne with whom I spent an enjoyable few months working on graph domination problems — I hope retirement is treating you well.

- The math department at UVic as a whole. From the office staff and the faculty, right through to my fellow graduate students. Everyone made the ride somewhat less traumatic.

- My parents, brother and sister-in-law for the non-stop interest, encouragement and support on many levels through what must have seemed like a never ending life as a student. *Julle alleen weet hoe lank dit my geneem het om hier uit te kom en julle het nooit getwyfel dat ek dit kon doen nie. Baie dankie daarvoor.*

- My wife, Dana, whose love and support made this bearable. You always believed in me and got me to see the brighter side of things. Finally, it's over and I am looking forward to our life together without "the thesis" hanging over our heads.

I never thought this day would come. We have a saying in Afrikaans: *Die agteros kom ook in die kraal.* Roughly translated it means that the ox at the back of the pack will also make it into the pen. This is one ox that's happy to be in the kraal.

Combinatorics, the finite case, is where the genuine, deep insight is. Generalizing, making it infinite is sometimes intricate and sometimes difficult, and I might even be willing to say that it's sometimes deep, but it is nowhere near as fundamental as seeing the finite structure.

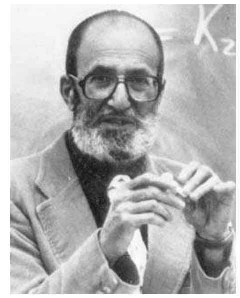

PAUL HALMOS
1916 – 2006

The infinite we shall do right away.
The finite may take a little longer.

STANISLAW ULAM
1909 – 1984

As soon as an Analytical Engine exists, it will necessarily guide the future course of the science. Whenever any result is sought by its aid, the question will then arise — By what course of calculation can these results be arrived at by the machine in the shortest time?

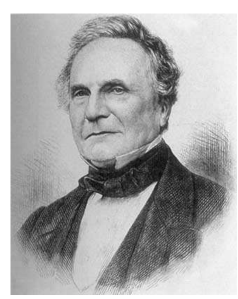

CHARLES BABBAGE
1791 – 1871

1

Introduction

Structure-preserving maps play a central and unifying role in modern mathematics. Eilenberg and MacLane felt so strongly about this, they studied them for their own sake:

> *Frequently in modern mathematics there occur phenomena of "naturality": a "natural" isomorphism between two groups or between two complexes, a "natural" homeomorphism of two spaces and the like. We here propose a precise definition of the "naturality" of such correspondences, as a basis for an appropriate general theory.*
>
> Eilenberg and MacLane [16].

Graph theory is no different. Here we have the concept of a graph homomorphism as a mapping between two graphs that preserves the structure of the first — its edges (or arcs in the case of a digraph).

More precisely, given two graphs (or digraphs) G and H, a *homomorphism* from G to H is a mapping $f : V(G) \to V(H)$ such that $uv \in E(G)$ implies that $f(u)f(v) \in E(H)$.

The simplicity of the concept belies the great wealth of interesting and deep results that have been obtained in the (almost) fifty year history of the subject [32, 33, 37]. A search on MathSciNet for "homomorphism" with `05C` (graph theory) as its primary subject results in about 450 matches. With `05C` as primary or secondary subject, there are about 600 matches.

Our focus in this thesis is on the computational aspects of graph homomorphisms: given two graphs (or digraphs) G and H, is there an efficient procedure for deciding the existence of a homomorphism from G to H?

It is not too surprising that the answer to the question varies depending on the graphs G and H involved. More surprising is the fact that in the case of (undirected) graphs the question has been settled completely by Hell and Nešetřil [36]. In the case of digraphs the search for a complete classification has not met with as much success. Accordingly we focus on the digraph problem for certain restricted classes of digraphs. We will also consider a variation of the basic homomorphism idea, namely that of an injective homomorphism.

For graph theory terms not defined here, the reader may consult [9]. A general reference for directed graphs is [3]. For graph homomorphisms see [37]. For complexity theory see [26].

1.1 A dichotomy for graph homomorphisms

We begin this section by briefly discussing Hell and Nešetřil's result cited above. The existence of a homomorphism from G to H is often abbreviated as $G \to H$. In this case we will also say that G is H-colourable. This terminology stems from the fact that graph colourings may be viewed as homomorphisms to complete graphs. H is also known as the target of the homomorphism problem. The H-colouring problem is defined in Problem 1.1.

Problem 1.1 HOM_H

Instance: A graph or digraph G.

Question: Does there exist a homomorphism $f : G \to H$?

Theorem 1.1.1 (Hell and Nešetřil [36, 37]). *Let H be a graph with loops allowed.*

- *If H is bipartite or contains a loop, then the H-colouring problem has a polynomial time algorithm.*
- *Otherwise the H-colouring problem is NP-complete.*

Another way of viewing this result is that among graph homomorphism problems, there are no problems of intermediate nature. That is, problems that are neither in P nor NP-complete. A result of Ladner's [41] states that the existence of such problems is a real possibility if we assume that P $\neq$ NP (as is widely believed). So for graph homomorphism problems there is a *dichotomy*: every H-colouring problem is either in P or it is NP-complete.

For the past twenty years or so the goal has been to try and extend Hell and Nešetřil's result to digraphs. This seems to be a much harder problem. Various classes of digraphs have been shown to exhibit a dichotomy. That is if H is a member of some well defined family of digraphs (say tournaments), then H-colouring is either in P or it is NP-complete.

A complete classification seems to be out of reach. As an example of this we note a sequence of six (unicyclic) digraphs in [7] where each digraph is obtained from its predecessor by adding a single arc and yet the complexity oscillates between being polynomial and being NP-complete. In addition to this there are also oriented trees

T such that the T-colouring problem is NP-complete [29, 34]. More than that even, the trees in [34] only have one vertex of degree three and $\Delta \leq 3$ (in the underlying undirected tree). The frustration of the homomorphism research community was summarized quite well by Gutjahr, Woeginger and Welzl when they said "Did we really want trees to be NP-complete?" [29].

1.2 A sample of previous results

Undirected graphs may be viewed as directed graphs where each undirected edge corresponds to a pair of symmetric arcs. From now on we will state all of our results and definitions for digraphs only, with the understanding that they also apply to undirected graphs using the above correspondence.

1.2.1 The core of a digraph

From a complexity theoretic viewpoint one would like to consider homomorphism problems that are in some way irreducible. The core of a digraph plays exactly that role.

Two digraphs G and H that are homomorphic to each other are said to be *homomorphically equivalent*.

If H' is a subgraph of H, then a *retraction* of H to H' is a homomorphism $\rho : H \to H'$ such that $\rho(x) = x$ for every $x \in V(H')$. In this case we say that H *retracts to* H' or that H' is a *retract of* H. In this case H and H' are homomorphically equivalent.

A digraph H is said to be a *core* if H does not retract to a proper subgraph.

It turns out that a digraph is a core if and only if it is not homomorphic to a proper subgraph and that every digraph H has a unique retract that is also a core [37]. This retract is called *the core of* H.

The importance of the core of a digraph H lies in the fact that if H' is the core of H, then H and H' have equivalent homomorphism problems: $G \to H$ if and only if $G \to H'$. Thus when considering the complexity of the H-colouring problem, we only need to consider digraphs H that are themselves cores.

1.2.2 Complexity results

The dichotomy conjecture stated below has stimulated a great deal of work in the complexity of digraph homomorphisms [37]

Conjecture 1.2.1. *For each digraph H, the H-colouring problem is polynomial time solvable or NP-complete.*

For relatively simple families of digraphs such as directed paths and directed cycles, polynomial-time algorithms are known [37, 46].

Oriented paths took somewhat longer to settle, but ultimately they were shown to be polynomial-time solvable as well [29].

Oriented cycles are more subtle. The *net length* of an oriented cycle is the number of forward arcs minus the number of backward arcs with respect to some fixed traversal of the cycle. An oriented cycle is *balanced* if its net length is zero and *unbalanced* otherwise.

Theorem 1.2.2 (Hell and Zhu [39], Gutjahr [28])**.** *If C is an unbalanced oriented cycle, then C-colouring is in P. On the other hand there exist balanced oriented cycles C for which C-colouring is NP-complete.*

It was shown later by Feder [19] that the family of all oriented cycles does exhibit a dichotomy.

As stated before there exist oriented trees for which the corresponding homomorphism problem is NP-complete. The first example was constructed in [29] and had 287 vertices. A few years later smaller examples were constructed in [34], the smallest of which has 45 vertices, $\Delta \leq 3$, and exactly one vertex of degree three (in the underlying undirected tree).

One of the first dichotomy results for a “natural” family of digraphs was the result by Bang-Jensen, Hell and MacGillivray [5]. Their result deals with so-called semi-complete digraphs. A *semi-complete* digraph has the property that between every pair of vertices there is at least one arc. Parallel arcs and loops are not allowed, but symmetric arcs may occur. It should be pointed out that the family of semi-complete digraphs contains both undirected complete graphs (every pair of vertices has a pair of symmetric arcs) as well as tournaments (orientations of undirected complete graphs).

Theorem 1.2.3 (Bang-Jensen, Hell and MacGillivray [5]). *Let H be a semi-complete digraph.*

- *If H contains at most one directed cycle, then H-colouring is polynomial time solvable.*
- *Otherwise H-colouring is NP-complete.*

Chapter 2 deals with a generalization of this theorem from tournaments to local tournaments. In a *local tournament* every vertex v has the property that both $N^+(v)$ as well as $N^-(v)$ induce tournaments.

The existence of two cycles in a digraph is often enough for NP-completeness to hold. This was investigated in [4] and [6] and lead to the following conjecture on so-called smooth digraphs. A digraph H is said to be *smooth* if it does not contain any sources or sinks.

Conjecture 1.2.4 (Bang-Jensen and Hell [4]). *Let H be a smooth digraph. If the core of H is a directed cycle, then H-colouring is in P. Otherwise H-colouring is NP-complete.*

This conjecture was only verified recently using techniques from universal algebra [8].

1.3 A connection to constraint satisfaction problems

Large classes of problems in discrete mathematics (and other areas) may be viewed as constraint satisfaction problems. As evidence of this we quote from the appendix to [20].

> *In this paper, we have proposed constraint satisfaction as a unifying framework for problems that can be solved with Datalog, such as Horn clauses and 2-satisfiability; for NP-complete problems, such as 3-coloring and one-in-three satisfiability; for problems from group theory, such as systems of linear equations and labeled graph isomorphism; for linear programming, graph matching, and a family of matroid parity problems; and for network stability and stable matching.*

Feder and Vardi [20] also show that there is a direct connection between constraint satisfaction problems and so called link systems. These systems unify problems in Markov processes and quantum mechanics and also provide an interpretation for relativity and for relational databases.

The importance of constraint satisfaction problems are therefore undeniable.

We now define the constraint satisfaction problem [37]. A *general relational system* S consists of:

- A finite set $V = V(S)$, the *vertices* of S.
- A finite set of relations $R_i(S)$, $i \in I$ (I being the index set). Denote the arity of $R_i(S)$ by k_i.
- The finite set I and the integers k_i form the *pattern* (or *type*) of S.

A *homomorphism* between two general relational systems S and T with the same pattern is a mapping $f : V(S) \to V(T)$ such that $(v_1, v_2, \ldots, v_{k_i}) \in R_i(S)$ implies that $(f(v_1), f(v_2), \ldots, f(v_{k_i})) \in R_i(T)$. The existence of such a homomorphism is denoted by $S \to T$.

A *constraint satisfaction problem* (CSP) is the problem of finding a homomorphism between two general relational systems S and T (with the same pattern). The vertices in S may be thought of as the variables of the problem and the vertices of T as possible values for the variables. A homomorphism $S \to T$ is then an assignment of values to variables satisfying the constraints captured by the relations.

A digraph is a relational system with one binary relation. If one were to consider arcs of different colours, this would correspond to a relational system with several binary relations (one for each colour). So in a sense a relational system is a directed hypergraph with arcs of different colours and a homomorphism between such systems preserves each hyperarc as well as its colour.

If T is a fixed relational system, then T defines a constraint satisfaction problem with respect to T, CSP(T).

Problem 1.2 CSP(T)

Instance: A relational system S with the same pattern as T.

Question: Does there exist a homomorphism $f : S \to T$?

As stated before, Ladner's result [41] guarantees that if P $\neq$ NP, then there are NP-problems that are neither in P nor NP-complete. Feder and Vardi's work was motivated by the question "what is the most general subclass of NP that we can define that may not contain such in-between problems?" [20]. Their work led them to the following conjecture.

Conjecture 1.3.1 (Feder and Vardi [20])**.** *For each general relational system T, the constraint satisfaction problem CSP(T) is polynomial time solvable or NP-complete.*

The importance of CSPs for us follows from the following result of Feder and Vardi [20].

Theorem 1.3.2 (Feder and Vardi [20]). *For every constraint satisfaction problem, CSP(T), there is a corresponding digraph H, such that CSP(T) is polynomially equivalent to H-colouring.*

This result would imply that if the H-colouring problem has a dichotomy, then so will CSP(T). That is, Conjecture 1.2.1 implies Conjecture 1.3.1.

1.4 Indicator and Sub-indicator Constructions

Hell and Nešetřil [36] introduced a number of powerful tools for proving that a given digraph has an NP-complete homomorphism problem. The aim of this section is to introduce these tools and to provide some examples illustrating their use.

1.4.1 The Indicator Construction

Let I be a fixed digraph with two specified vertices i and j. The *indicator construction* (with respect to the *indicator* I, i, j) transforms a digraph H to the digraph H^* as follows. The vertex set of H^* is the same as that of H. Arcs are defined by the following rule: xy is an arc of H^* if and only if there exists a homomorphism from I to H mapping i to x and j to y. We then have the following result.

Lemma 1.4.1 (Hell and Nešetřil [36, 37]). *If the H^*-colouring problem is NP-complete, then the H-colouring problem is also NP-complete.*

1.4.2 The (Vertex) Sub-indicator Construction

Let J be a fixed digraph with specified vertices $k_1, k_2, \ldots, k_t$ and j. The *sub-indicator construction* (with respect to the *sub-indicator* $J, k_1, k_2, \ldots, k_t, j$) transforms a digraph H with specified vertices $x_1, x_2, \ldots, x_t$ to an induced subgraph H^+ defined as follows. Let W be the digraph obtained from a copy H and a copy of J by identifying each k_i with the corresponding x_i for $i = 1, 2, \ldots, t$. Then H^+ is the subgraph of H induced by those vertices u for which some retraction of W to H maps j to u.

Lemma 1.4.2 (Hell and Nešetřil [36, 37]). *Let H be a digraph that is a core. If the H^+-colouring problem is NP-complete, then the H-colouring problem is also NP-complete.*

Often, when using the sub-indicator construction, we take the vertices $k_1, k_2, \ldots, k_t$ above to be a set of isolated vertices in J. This has the effect that the digraph W above is $H \cup (J - \{k_1, k_2, \ldots, k_t\})$. In considering retractions of W to H, we see that we are actually considering homomorphisms of $J - \{k_1, k_2, \ldots, k_t\}$ to H.

1.4.3 The Arc-sub-indicator Construction

Let J be a fixed graph with a specified arc jj' and t specified vertices $k_1, k_2, \ldots, k_t$. The *arc-sub-indicator construction* (with respect to the *arc-sub-indicator* $J, k_1, k_2, \ldots, k_t, jj'$) transforms a digraph H with t specified vertices $x_1, x_2, \ldots, x_t$ into its subgraph H^- determined by the images of the arc jj' under retractions of W (defined as above) to H. This construction is therefore an arc version of the (vertex) sub-indicator outlined above.

Lemma 1.4.3 (Hell and Nešetřil [36]). *Let H be a core. If the H^--colouring problem is NP-complete, then so is the H-colouring problem.*

1.4.4 Examples

We now illustrate the use of these tools.

Example 1

Consider the local tournament H shown in Figure 1.1. This is in fact a round local tournament (see Chapter 2).

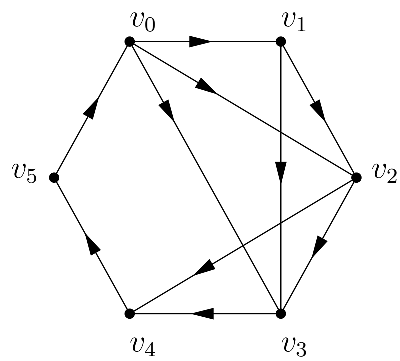

Figure 1.1: A local tournament H for which the H-colouring problem is NP-complete.

Proposition 1.4.4. *The H-colouring problem is NP-complete.*

Proof. To show that this digraph defines an NP-complete problem, we use the indicator construction. Our indicator, I, is a directed path of length seven with i equal to the initial vertex of the path and j equal to the terminal vertex of the path. The images of i and j under homomorphisms from I to H can be found by constructing a table as shown below.

This table shows the kth out-neighbourhood of each vertex in H. Therefore if, for instance, i mapped to v_0, then j can map to any of $\{v_1, v_2, v_3, v_4, v_5\}$. The result of the indicator, H^* is shown in Figure 1.2. Undirected edges (a pair of symmetric arcs) is shown as an edge without any arrows.

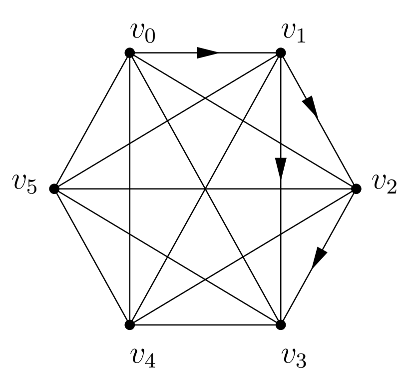

Figure 1.2: The result, H^*, of applying the indicator $I = P_7$ to the local tournament H.

Table 1.1: The kth out-neighbourhood of each vertex of H for $1 \leq k \leq 7$.

v	$N^{+}(v)$	$N^{+2}(v)$	$N^{+3}(v)$	$N^{+4}(v)$	$N^{+5}(v)$	$N^{+6}(v)$	$N^{+7}(v)$
v_0	v_1, v_2, v_3	v_2, v_3, v_4	v_3, v_4, v_5	v_0, v_4, v_5	v_0, v_1, v_2, v_3, v_5	v_0, v_1, v_2, v_3, v_4	v_1, v_2, v_3, v_4, v_5
v_1	v_2, v_3	v_3, v_4	v_4, v_5	v_0, v_5	v_0, v_1, v_2, v_3	v_1, v_2, v_3, v_4	v_2, v_3, v_4, v_5
v_2	v_3, v_4	v_4, v_5	v_0, v_5	v_0, v_1, v_2, v_3	v_1, v_2, v_3, v_4	v_2, v_3, v_4, v_5	v_0, v_3, v_4, v_5
v_3	v_4	v_5	v_0	v_1, v_2, v_3	v_2, v_3, v_4	v_3, v_4, v_5	v_0, v_4, v_5
v_4	v_5	v_0	v_1, v_2, v_3	v_2, v_3, v_4	v_3, v_4, v_5	v_0, v_4, v_5	v_0, v_1, v_2, v_3, v_5
v_5	v_0	v_1, v_2, v_3	v_2, v_3, v_4	v_3, v_4, v_5	v_0, v_4, v_5	v_0, v_1, v_2, v_3, v_5	v_0, v_1, v_2, v_3, v_4

In order to show that H-colouring is NP-complete, we have to examine the new digraph H^*. Here we have two options:

- ❒ H^* is a semi-complete digraph with at least two cycles, so therefore by Theorem 1.2.3, H^*-colouring is NP-complete. This means that by Lemma 1.4.1, H-colouring is NP-complete.

- ❒ Observe that the "undirected portion" of H^* is not bipartite. In order to extract the undirected portion of H^*, we apply, as an arc-sub-indicator, the digraph J with two vertices j and j' and arcs jj' and $j'j$. That is, J is an undirected edge. When $J \to H^*$, j and j' can only map to vertices that have an undirected edge between them. The result, $(H^*)^-$, is therefore the undirected portion of H^*. Since $(H^*)^-$ is not bipartite, Theorem 1.1.1 implies that the $(H^*)^-$-colouring problem is NP-complete. Lemma 1.4.3 then implies that H^*-colouring is NP-complete (since H^* is a core). Now Lemma 1.4.1 finally implies that H-colouring is NP-complete.

∎

Example 2

Our next example is the local tournament D shown in Figure 1.3. This is an example of a round decomposable local tournament with strong components D_0, D_1, D_2 and D_3 (see Chapter 2).

Proposition 1.4.5. *The D-colouring problem is NP-complete.*

Proof. Here we first apply an indicator equal to a directed path of length 2, where the vertices i and j of the indicator are taken to be the initial vertex of the path and the terminal vertex of the path, respectively. This produces the result, D^*, shown in Figure 1.4. Note that the orientation of the C_3 changed, and that there are symmetric

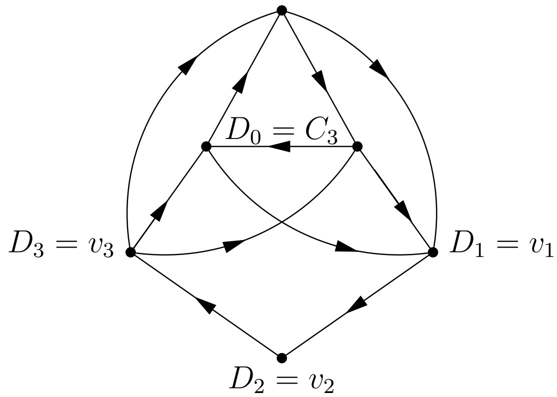

Figure 1.3: An example of a round decomposable local tournament D.

arcs between v_2 and the triangle and between v_1 and v_3. These symmetric arcs are drawn as undirected edges. D^* is also a core.

At this point we apply a sub-indicator, J, that is also equal to a directed path of length 2. We let j be the terminal vertex of the path and k_1 the initial vertex of the path. We also take $x_1 = v_2$. That is we identify k_1 in J with v_2 in D^* and consider all retractions to D^*, keeping track of the possible images of j. The result, D^{*+}, is equal to $D^* - v_3$, the core of which is a wheel with three spokes (or a semi-complete digraph on four vertices with at least two cycles). This sequence of digraphs is shown in Figure 1.4.

Since the wheel-colouring problem is NP-complete (see Section 1.5.1), we find that D^*-colouring is NP-complete, and that ultimately, D-colouring is NP-complete.

∎

1.5 Direct NP-completeness Reductions

Another technique for proving that a given problem P is NP-complete, is to show that there is a polynomial time transformation from some other known NP-complete problem P'. This entails taking an instance I' of P', transforming it (in polynomial time) into an instance I of P and showing that I' is a yes instance of P' if and only

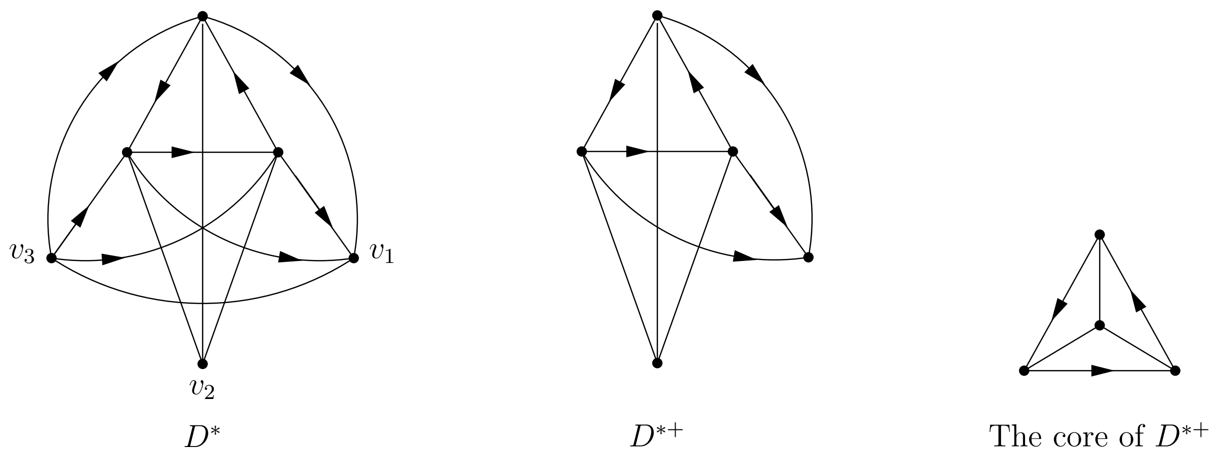

Figure 1.4: The sequence of digraphs obtained in the proof of Proposition 1.4.5.

if I is a yes instance of P.

The known NP-complete problems that we will use are all variations of Boolean satisfiability. These problems all have a formulation as a Boolean constraint satisfaction problem. Schaefer [49] proved the following dichotomy theorem (this formulation is from [35]).

Theorem 1.5.1. *Suppose H is a relational system with $V(H) = \{0, 1\}$ and relations $R_1, R_2, \ldots, R_p$. Then CSP(H) is NP-complete, except in the following polynomial time solvable cases:*

1. *each R_i contains the tuple $(0, 0, \ldots, 0)$,*
2. *each R_i contains the tuple $(1, 1, \ldots, 1)$,*
3. *each R_i is closed under the OR operation,*
4. *each R_i is closed under the AND operation,*
5. *each R_i is closed under the MAJORITY operation, or*
6. *each R_i is closed under the XOR operation.*

The OR operation on two tuples $(a_1, a_2, \ldots, a_s)$ and $(b_1, b_2, \ldots, b_s)$ is the tuple $(z_1, z_2, \ldots, z_s)$ with $z_i = a_i \vee b_i$. The AND operation on two tuples $(a_1, a_2, \ldots, a_s)$ and $(b_1, b_2, \ldots, b_s)$ is the tuple $(z_1, z_2, \ldots, z_s)$ with $z_i = a_i \wedge b_i$. The MAJORITY operation on three tuples $(a_1, a_2, \ldots, a_s), (b_1, b_2, \ldots, b_s)$ and $(c_1, c_2, \ldots, c_s)$ is the tuple $(z_1, z_2, \ldots, z_s)$ with z_i the majority of a_i, b_i, c_i. The XOR (or MINORITY) operation on three tuples $(a_1, a_2, \ldots, a_s), (b_1, b_2, \ldots, b_s)$ and $(c_1, c_2, \ldots, c_s)$ is the tuple $(z_1, z_2, \ldots, z_s)$ with z_i the exclusive-or (or minority) of a_i, b_i, c_i.

The three satisfiability problems that we will need are: ℓ-satisfiability (Problem 1.3), monotone not-all-equal ℓ-satisfiability (Problem 1.4) and monotone one-in-ℓ ℓ-satisfiability (Problem 1.5).

Problem 1.3 ℓ-SAT

Instance: A set U of Boolean variables and a collection C of clauses over U such that each clause has size ℓ.

Question: Is there a truth assignment for U such that each clause in C has at least one true literal?

In the case of ℓ-SAT, $V(H) = \{0, 1\}$ ("False" and "True") and there are 2^ℓ relations

$$\begin{aligned}
R_{00\cdots0} &= \{(r_1, r_2, \ldots, r_\ell) \mid (r_1, r_2, \ldots, r_\ell) \neq (0, 0, \ldots, 0)\}, \\
R_{00\cdots1} &= \{(r_1, r_2, \ldots, r_\ell) \mid (r_1, r_2, \ldots, r_\ell) \neq (0, 0, \ldots, 1)\}, \\
&\vdots \\
R_{11\cdots1} &= \{(r_1, r_2, \ldots, r_\ell) \mid (r_1, r_2, \ldots, r_\ell) \neq (1, 1, \ldots, 1)\},
\end{aligned}$$

where each $r_i = 0, 1$. An instance of ℓ-SAT has $V(U) = \{u_1, u_2, \ldots, u_k\}$ (the variables) and 2^ℓ relations

$$\begin{aligned}
Q_{00\cdots 0} &= \{(u_{i_1}, u_{i_2}, \ldots, u_{i_\ell}) \mid \text{if } u_{i_1} \vee u_{i_2} \vee \cdots \vee u_{i_\ell} \text{ is a clause}\},\\
Q_{00\cdots 1} &= \{(u_{i_1}, u_{i_2}, \ldots, u_{i_\ell}) \mid \text{if } u_{i_1} \vee u_{i_2} \vee \cdots \vee \overline{u_{i_\ell}} \text{ is a clause}\},\\
&\vdots\\
Q_{11\cdots 1} &= \{(u_{i_1}, u_{i_2}, \ldots, u_{i_\ell}) \mid \text{if } \overline{u_{i_1}} \vee \overline{u_{i_2}} \vee \cdots \vee \overline{u_{i_\ell}} \text{ is a clause}\}.
\end{aligned}$$

When $\ell = 2$, we have case 5 above. This is the well-known 2-SAT problem that is polynomial time solvable.

Therefore let $\ell \geq 3$.

It is clear that cases 1 and 2 of Theorem 1.5.1 do not hold. Case 3 does not hold since $(0, 1, 1, \ldots, 1)$ and $(1, 0, 1, \ldots, 1)$ are both in $R_{11\cdots 1}$ whilst their OR, which is equal to $(1, 1, 1, \ldots, 1)$, is not. For case 4, consider the tuples $(1, 0, 0, \ldots, 0)$ and $(0, 1, 0, \ldots, 0)$. They are in $R_{00\cdots 0}$ but their AND is not. For case 5, consider the tuples $(1, 0, 0, \ldots, 0)$, $(0, 1, 0, \ldots, 0)$ and $(0, 0, 1, \ldots, 0)$. These tuples are in $R_{00\cdots 0}$ but the MAJORITY is not. For case 6, consider the tuples $(0, 1, 1, \ldots, 0)$, $(1, 0, 1, \ldots, 1)$ and $(1, 1, 0, \ldots, 1)$. These tuples are in $R_{00\cdots 0}$ but their XOR (or MINORITY) is not. Therefore ℓ-SAT is NP-complete by Theorem 1.5.1.

Problem 1.4 MONOTONE NOT-ALL-EQUAL ℓ-SAT

Instance: A set U of Boolean variables, and a collection C of clauses over U such that each clause has size ℓ and contains only un-negated variables.

Question: Is there a truth assignment for U such that each clause in C has at least one true literal and at least one false literal?

For monotone not-all-equal ℓ-SAT we have $V(H) = \{0, 1\}$ with one relation of arity ℓ

$$R = \{(r_1, r_2, \ldots, r_\ell) \mid r_i = 0, 1\} - \{(0, 0, \ldots, 0), (1, 1, \ldots, 1)\}.$$

An instance of monotone not-all-equal ℓ-SAT has $V(U) = \{u_1, u_2, \ldots, u_k\}$ (the variables) with one relation of arity ℓ, the clauses themselves. As with ℓ-SAT, Theorem 1.5.1 shows that monotone not-all-equal ℓ-SAT is NP-complete (using the same tuples as before).

Problem 1.5 MONOTONE ONE-IN-ℓ ℓ-SAT

Instance: A set U of Boolean variables, and a collection C of clauses over U such that each clause has size ℓ.

Question: Is there a truth assignment for U such that each clause in C has exactly one true literal?

Monotone one-in-ℓ ℓ-SAT has $V(H) = \{0, 1\}$ and one relation of arity ℓ

$$R = \{(1, 0, 0, \ldots, 0), (0, 1, 0, \ldots, 0), (0, 0, 1, \ldots, 0), \ldots, (0, 0, 0, \ldots, 1)\}.$$

An instance is handled in the same way as with monotone not-all-equal ℓ-SAT and again by Theorem 1.5.1, monotone one-in-ℓ ℓ-SAT is NP-complete.

1.5.1 Colouring by Wheels is NP-complete

Let H be the wheel-graph shown below in Figure 1.5. H has vertices $\{0, 1, 2, \ldots, n\}$ and arcs $0i$, $i0$ for $i = 1, 2, \ldots, n$, $j(j + 1)$ for $j = 1, 2, \ldots, n - 1$ and $n1$.

Theorem 1.5.2. *H-colouring is NP-complete, even for acyclic inputs.*

Proof. The proof is via a reduction from not-all-equal 3-SAT without negated variables.

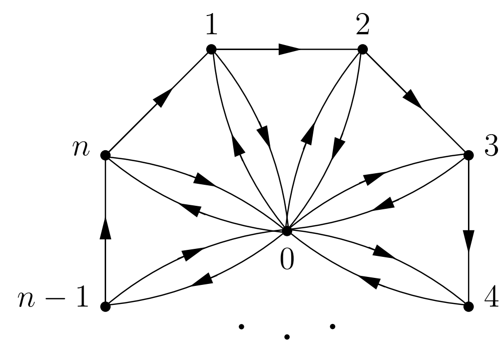

Figure 1.5: The target H.

Throughout let $k \in \{1, 2, \ldots, n\}$ and define

$$k^+ = \begin{cases} k+1 & \text{if } 1 \leq k < n, \\ 1 & \text{if } k = n, \end{cases}$$

and

$$k^- = \begin{cases} k-1 & \text{if } 1 < k \leq n, \\ n & \text{if } k = 1. \end{cases}$$

We have the following transitive triples in H: $0kk^+$, $k0k^+$ and kk^+0. Let F be the digraph shown below in Figure 1.6

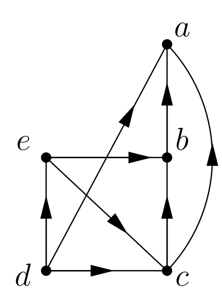

Figure 1.6: The gadget F.

If $F \to H$, then $c \not\mapsto 0$: if $c \mapsto 0$, then $b \mapsto k$ and $a \mapsto k^+$. This implies that $d \mapsto k$ and $e \mapsto k^+$. Therefore the arc eb is mapped to k^+k which is not an arc of H.

On the other hand there is a homomorphism $f : F \to H$ in which $f(c) = k$, where

$k \in \{1, 2, \ldots, n\}$. This homomorphism is given by: $f(a) = 0$, $f(b) = k^+$, $f(c) = k$, $f(d) = k^-$ and $f(e) = 0$.

Let G be the digraph shown in Figure 1.7.

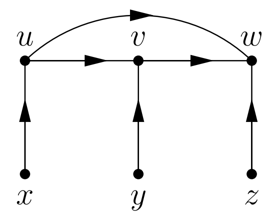

Figure 1.7: The gadget G.

If $G \to H$, then $(x, y, z) \neq (0, 0, 0)$ as this would force a, b and c to all map to nonzero vertices and no transitive triple on nonzero vertices alone exists. Also if $G \to H$, then $(x, y, z) \neq (k, k, k)$. If this was the case then a, b and c are forced to map to $\{0, k^+\}$ and no such transitive triple exists. On the other hand homomorphisms from G to H are shown in Table 1.2 where x, y and z have been pre-coloured with $\{0, k\}$ using a majority of 0's or a majority of k's.

Table 1.2: Homomorphisms of the gadget G to H.

u	v	w	k	k^+	0	k	0	k^+	0	k	k^+
x	y	z	0	0	k	0	k	0	k	0	0
u	v	w	k^+	0	k^{++}	0	k	k^+	k	0	k^+
x	y	z	k	k	0	k	0	k	0	k	k

We are now ready to exhibit the reduction.

Let an instance of not-all-equal 3-SAT without negated variables be given:

$$X = \{x_1, x_2, \ldots, x_\ell\} \text{ — the variables,}$$

$$C_1, C_2, \ldots, C_m \text{ — the clauses.}$$

Each clause $C_i = \{x_{i_1}, x_{i_2}, x_{i_3}\}$ with $x_{i_1}, x_{i_2}, x_{i_3} \in X$ and $i = 1, 2, \ldots, m$.

Construct a digraph D as follows: Take a copy of F and add vertices $x_1, x_2, \ldots, x_m$ to F as well as the arcs cx_i $i = 1, 2, \ldots, m$. For each clause $C_i = \{x_{i_1}, x_{i_2}, x_{i_3}\}$ take a copy of G and identify x, y, z in G with the corresponding $x_{i_1}, x_{i_2}, x_{i_3}$. This is illustrated in Figure 1.8.

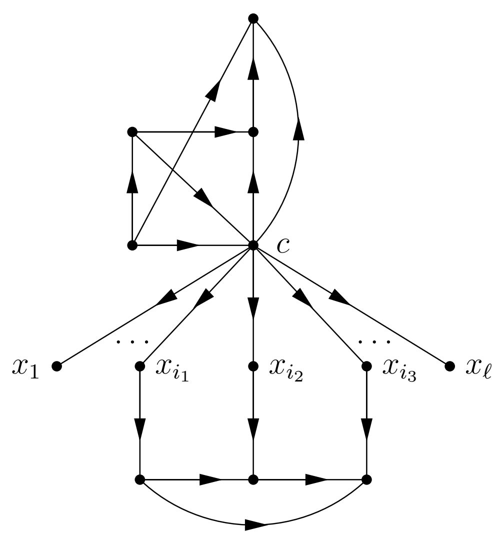

Figure 1.8: The digraph D.

If $D \to H$, then $c \mapsto k \in \{1, 2, \ldots, n\}$ which in turn implies that $x_1, x_2, \ldots, x_\ell \to \{0, k^+\}$. The clause gadget G prevents all of the x's in the same clause from being mapped to the same vertex. This allows one to read off a satisfying truth assignment: $0 =$"False" and $k^+ =$"True."

If there exists a satisfying truth assignment we identify "True" with the vertex 1 in H and "False" with the vertex 0 in H. This produces a pre-colouring on the vertices $x_1, x_2, \ldots, x_\ell$ in D which can be extended to a homomorphism $D \to H$.

Therefore there exists a satisfying truth assignment for not-all-equal 3-SAT without negated variables if and only if $D \to H$. Thus H-colouring is NP-complete (even for acyclic inputs). ∎

1.5.2 The Complexity of Oriented Colouring

An oriented graph is a digraph without 2-cycles.

An oriented colouring of an oriented graph D is an assignment of colours to the vertices of D such that adjacent vertices receive different colours and that between any two colour classes the arcs are all oriented in the same direction [50]. Alternatively, one may view an oriented colouring of D as a homomorphism to some oriented graph H. The smallest number of colours for which D has an oriented colouring is the oriented chromatic number of D, $\chi_o(D)$. Minimizing the number of colours used, is the same as minimizing the order of H in the homomorphism $D \to H$. Since the addition of arcs to H will not affect the existence of a homomorphism $D \to H$, one may assume that the targets in the oriented chromatic problem are always tournaments. The oriented chromatic problem with k colours, OCN_k, is Problem 1.6.

Problem 1.6 OCN_k

Instance: An oriented graph D and a positive integer k.

Question: Does D have an oriented k-colouring ?

Our aim here is to show that the oriented chromatic number problem is NP-complete for every $k \geq 4$, even for acyclic inputs to the problem.

We focus first on the case $k = 4$ and make use of the graphs in Figure 1.9.

The proof that OCN_4 is NP-complete is via a reduction from monotone not-all-equal 3-SAT.

Before proceeding with the proof that OCN_4 is NP-complete, we observe the following facts.

Fact 1.5.3. G_1 has a unique homomorphism to T_4: $a \mapsto 2$, $b \mapsto 1$, $c \mapsto 0$, $d \mapsto 3$, $e \mapsto 2$, $f \mapsto 1$.

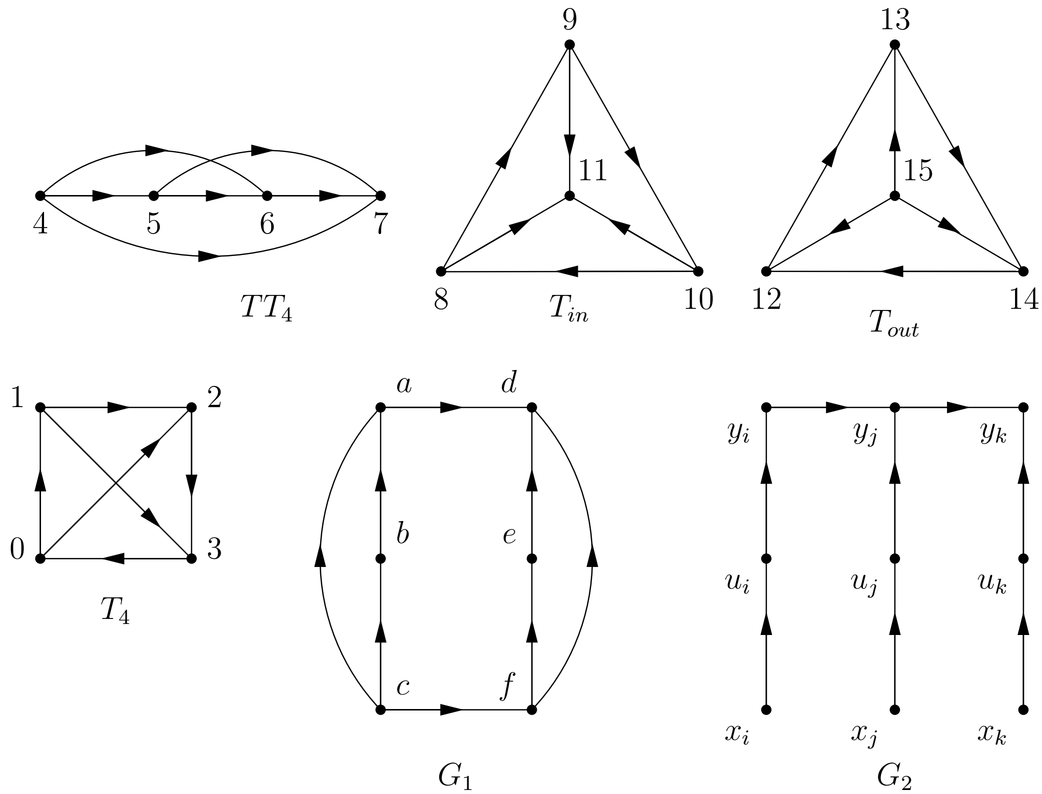

Figure 1.9: The graphs used in proving that OCN_4 is NP-complete.

Proof. Observe that T_4 has only two transitive tournaments on three vertices namely: $0, 1, 2$ and $1, 2, 3$. The two transitive tournaments a, b, c and d, e, f in G_1 have to map to transitive tournaments in T_4. The arcs (a, d) and (c, f) in G_1 force the mapping as shown above. ∎

Fact 1.5.4. Any pre-colouring of x_i, x_j, x_k in G_2 with the colours $\{0, 1\}$ such that at least one 0 and at least 1 is used can be extended to a homomorphism from G_2 to T_4. On the other hand if all of x_i, x_j, x_k are coloured the same (with 0 or 1), it does not extend to a homomorphism to T_4.

Fact 1.5.5. $G_1 \not\to T_{in}$, $G_1 \not\to T_{out}$, $G_2 \not\to TT_4$.

Proof. The only way that $G_1 \to T_{in}$ is if both vertices a and d map to vertex 11, since they're adjacent that is not allowed. Similarly, the only way that $G_1 \to T_{out}$ is if both vertices c and f map to vertex 15, again this is not allowed. Finally $G_2 \nrightarrow TT_4$ since it has a path of length 4: x_i, u_i, y_i, y_j, y_k. ∎

We now describe the reduction to monotone not-all-equal 3-SAT.

Given an instance of monotone not-all-equal 3-SAT: $U = \{x_1, x_2, \ldots, x_n\}$ and $C = \{c_1, c_2, \ldots, c_m\}$, we construct an acyclic, oriented graph D as follows. Take a copy of G_1 and for each variable $x_i \in U$ add a vertex labeled x_i to G_1 and join vertex x_i to vertex a of G_1 by an arc directed in this direction. For each clause $\{y_i, y_j, y_k\} \subseteq U$, construct the corresponding graph G_2 and then identify the vertices x_i, x_j, x_k from this G_2 with the vertices labeled x_i, x_j, x_k from before. Once this has been done for all clauses in C, we have the acyclic, oriented graph D. The construction of D can be carried out in polynomial time. The graph D is shown in Figure 1.10.

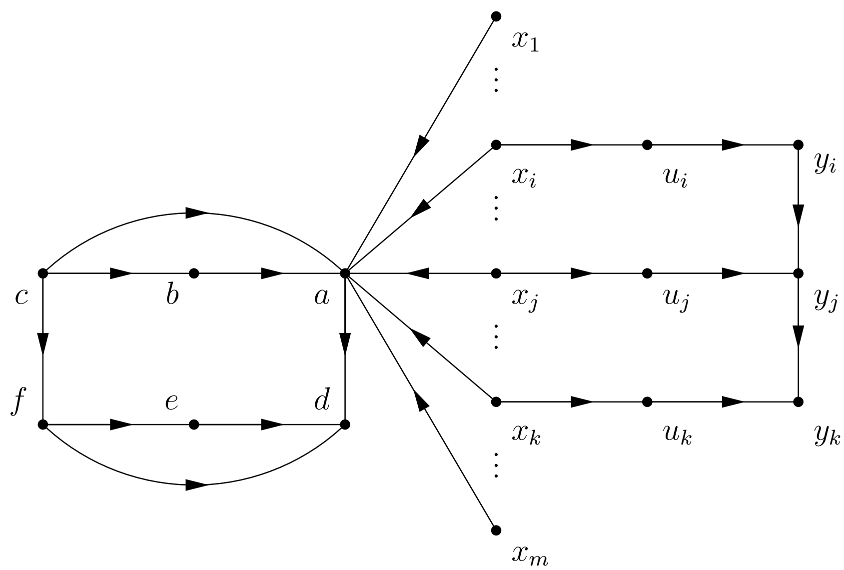

Figure 1.10: The oriented, acyclic graph D.

Lemma 1.5.6. *D has an oriented 4-colouring if and only if there exists a satisfying truth assignment for monotone not-all-equal 3-SAT.*

Proof. $\Rightarrow$:

Assume D is 4-colourable. That implies that D has a homomorphism to one of the four tournaments on four vertices: TT_4, T_{in}, T_{out} or T_4. Since D has both G_1 and G_2 as subgraphs, these prevent D from having a homomorphism to TT_4, T_{in} or T_{out}. Thus $D \rightarrow T_4$. From the facts presented above we know that vertex a will map to vertex 2 of T_4. Therefore the only possible images of the vertices x_i are $\{0, 1\}$. Furthermore they cannot all map to 0 or all map to 1 (by the facts above). The images of the vertices x_i form a satisfying truth assignment for monotone not-all-equal 3-SAT (FALSE = 0 and TRUE = 1).

$\Leftarrow$:

Assume that there exists a satisfying truth assignment for monotone not-all-equal 3-SAT. Use this assignment to colour the vertices x_i in D (FALSE= 0 and TRUE= 1). This pre-colouring can be extended to the whole D to give a homomorphism to T_4 (by the facts above). Therefore D is 4-colourable. ∎

Proposition 1.5.7. *OCN_4 is NP-complete, even when restricted to acyclic, oriented graphs.*

Proof. Since it is easy to check whether a given colouring of an acyclic, oriented graph D is in fact a 4-colouring, OCN_4 is in NP. By the lemma above and since it is known that monotone not-all-equal 3-SAT is NP-complete, we know that OCN_4 is NP-complete. Note that the graph D constructed above is acyclic. ∎

Proposition 1.5.8. *OCN_k is NP-complete for any $k \in \mathbb{Z}^+$, $k \geq 4$, even when restricted to acyclic, oriented graphs.*

Proof. The case $k = 4$ is the theorem above. For any $k \geq 5$, we construct an oriented acyclic graph D' from D by adding $k - 4$ new vertices to D (one at a time) and letting them dominate all the previous vertices. This new graph D' is k-colourable if and only if D is 4-colourable. Therefore k-colourability of D' is equivalent to solving monotone not-all-equal 3-SAT, and so is NP-complete. ∎

1.6 Polymorphisms

An algebraic approach to studying the complexity of the CSP(T) problem (and therefore that of the H-colouring problem) was proposed by Jeavons in [40]. In this framework we associate an algebra to T, and the properties of this algebra will then, hopefully, provide us with useful information on the complexity of CSP(T).

An *algebra* is a pair $\mathbb{A} = (A, F)$ consisting of a set A (the *universe* of $\mathbb{A}$), and a set F of operations on A (the *basic operations* of $\mathbb{A}$). An *operation of rank (or arity)* k (for some natural number k) is a function $f : A^k \to A$.

Let T be a relational system with vertex set $V = V(T)$, index set I and relations $R_i(T) = R_i$ (of arity r_i), $i \in I$. We define the *direct product* T^k (for $k \geq 1$) by $V(T^k) = V^k$ and relations R'_i (of arity r_i), $i \in I$ by

$$\big((x_1^1, x_2^1, \ldots, x_k^1), (x_1^2, x_2^2, \ldots, x_k^2), \ldots, (x_1^{r_i}, x_2^{r_i}, \ldots, x_k^{r_i})\big) \in R'_i \text{ if and only if}$$
$$(x_1^1, x_1^2, \ldots, x_1^{r_i}), (x_2^1, x_2^2, \ldots, x_2^{r_i}), \ldots, (x_k^1, x_k^2, \ldots, x_k^{r_i}) \text{ are all in } R_i.$$

For a digraph $H = (V, A)$ (which is a relational system with one binary relation, A), this definition amounts to: $V(H^k) = V^k$ and arcs $(x_1^1, x_2^1, \ldots, x_k^1) \to (x_1^2, x_2^2, \ldots, x_k^2)$ if and only if $x_i^1 \to x_i^2$ is an arc of H for $1 \leq i \leq k$.

Given a constraint satisfaction problem CSP(T), a *polymorphism of* T is a homomorphism $f : T^k \to T$. That is if

$$\big((x_1^1, x_2^1, \ldots, x_k^1), (x_1^2, x_2^2, \ldots, x_k^2), \ldots, (x_1^{r_i}, x_2^{r_i}, \ldots, x_k^{r_i})\big) \in R'_i,$$

then

$$\big(f(x_1^1, x_2^1, \ldots, x_k^1), f(x_1^2, x_2^2, \ldots, x_k^2), \ldots, f(x_1^{r_i}, x_2^{r_i}, \ldots, x_k^{r_i})\big) \in R_i.$$

For a digraph H this means that $f : H^k \to H$ is a polymorphism if $(x_1^1, x_2^1, \ldots, x_k^1) \to$

$(x_1^2, x_2^2, \ldots, x_k^2)$ implies that $f(x_1^1, x_2^1, \ldots, x_k^1) \rightarrow f(x_1^2, x_2^2, \ldots, x_k^2)$.

Denote the set of polymorphisms of T by $\mathcal{P}ol(T)$. Associate with the constraint satisfaction problem CSP(T), the algebra $(V(T), \mathcal{P}ol(T))$.

Jeavons [40] showed that for each constraint satisfaction problem, $\mathcal{P}ol(T)$ could be classified into one of six categories depending on the polymorphisms present in $\mathcal{P}ol(T)$. The polymorphisms present in $\mathcal{P}ol(T)$ ultimately determine the complexity of CSP(T).

In [13], Bulatov, Jeavons and Krokhin show that we only need to consider idempotent polymorphisms to determine the complexity of CSP(T). An *idempotent* polymorphism $f : T^k \rightarrow T$ has the property that $f(x, x, \ldots, x) = x$ for every $x \in V(T)$. They also show that one only has to consider relational systems, T, that are *cores*: each homomorphism $T \rightarrow T$ is an automorphism. This is analogous to the digraph situation.

An operation $f : A^k \rightarrow A$ is called *essentially unary* if there exists a (nonconstant) unary operation $g : A \rightarrow A$ and an index $i \in \{1, 2, \ldots, k\}$ such that $f(x_1, x_2, \ldots x_k) = g(x_i)$ for all choices of $x_1, x_2, \ldots, x_k$. If g is the identity operation, then f is called a *projection*.

The set of polymorphisms $\mathcal{P}ol(T)$, for a given constraint satisfaction problem CSP(T), always contains the projections. The question then is whether there are other polymorphisms and whether they are of any use. The following theorem shows what happens if one is restricted to essentially unary polymorphisms.

Theorem 1.6.1 (Bulatov, Jeavons and Krokhin [13])**.** *If for a given constraint satisfaction problem CSP(T), $\mathcal{P}ol(T)$ contains only essentially unary polymorphisms, then CSP(T) is NP-complete.*

For an undirected complete graph K_n, with $n \geq 3$, the only idempotent polymorphisms are projections [42]. Therefore Theorem 1.6.1 furnishes a different proof of the result that (undirected) graph k-colouring is NP-complete for $k \geq 3$.

It was conjectured in [13] that the essentially unary polymorphisms are exactly the dividing line between NP-complete CSP problems and CSP problems that are polynomial time solvable.

There is an equivalent formulation of Theorem 1.6.1 due to Larose and Zádori [43]. The equivalence follows from results in Universal Algebra, see [14, 44] for more on this.

An idempotent polymorphism $f : T^k \to T$ is said to be a *Taylor polymorphism* if it satisfies k identities of the form

$$f(x_{i1}, x_{i2}, \ldots, x_{ik}) = f(y_{i1}, y_{i2}, \ldots, y_{ik}), \quad i = 1, 2, \ldots, k,$$

where $x_{ij}, y_{ij} \in \{x, y\}$ for all i, j and $x_{ii} \neq y_{ii}$ for $i = 1, 2, \ldots, k$. Note that a projection is not a Taylor polymorphism.

Theorem 1.6.2 (Larose and Zádori [43])**.** *Let CSP(T) be given. If $\mathcal{P}ol(T)$ does not contain any Taylor polymorphisms, then CSP(T) is NP-complete.*

The conjecture of Bulatov, Jeavons and Krokhin [13] is the following (here we use the alternative formulation from [43]).

Conjecture 1.6.3. *If T admits a Taylor polymorphism, then CSP(T) is polynomial time solvable. Otherwise CSP(T) is NP-complete.*

Thus it is conjectured that the Taylor polymorphisms differ from projections in just the right way to be of use in proving CSP(T) polynomial.

A recent result of Maróti and McKenzie [45] shows (again through Universal Algebra) that the existence of a Taylor polymorphism is equivalent to the existence of a weak near unanimity function of arity $k > 1$. A *weak near unanimity function* *of arity k (*WNUF$_k$*)* is an idempotent polymorphism $f : T^k \to T$ such that

$$\begin{aligned} f(y,x,x,\ldots,x,x) = f(x,y,x,\ldots,x,x) = f(x,x,y,\ldots,x,x) &= \cdots \\ &= f(x,x,x,\ldots,x,y), \end{aligned}$$

for all $x, y \in V(T)$.

The two theorems above (Theorems 1.6.1 and 1.6.2) can now be formulated as follows.

Theorem 1.6.4. *If T does not admit a* WNUF$_k$ *of arity $k > 1$, then CSP(T) is NP-complete.*

Conjecture 1.6.3 now becomes the following.

Conjecture 1.6.5 ([14])**.** *If T admits a* WNUF*, then CSP(T) is polynomial time solvable. Otherwise CSP(T) is NP-complete.*

In Chapter 3 we study weak near unanimity functions in the context of the digraph homomorphism problem. We prove that a wide range of H-colouring problems with known polynomial time algorithms have a WNUF. On the other hand, we prove some non-existence results for WNUFs (including some known NP-complete digraphs). This lends some support to the conjecture above. We also show that there is a close resemblance between traditional NP-completeness proofs and the ones using WNUFs. This is accomplished by proving a version of the vertex (and arc) sub-indicator construction and the indicator construction for WNUFs and then using this result to prove that all tournaments with at least two cycles do not admit a WNUF (using virtually the same proof as in [5]).

1.7 Injective Homomorphisms

Let G and H be digraphs. The homomorphism $f : G \to H$ is said to be *injective* if $f|_{N^-(v)}$ is an injective mapping from $V(G)$ to $V(H)$ for every $v \in V(G)$. Therefore, the in-neighbours of every vertex $v \in V(G)$ have to be mapped to distinct vertices in H while preserving the arcs of G.

One can of course also define a homomorphism $f : G \to H$ to be injective if it is injective on the out-neighbours of the vertices in G. This would be the same as requiring the homomorphism to be injective on the in-neighbours of the converse of G (where of course we are now also mapping to the converse of H).

It is also possible to insist that the homomorphism be injective on $N^+(v) \cup N^-(v)$ for every $v \in V(G)$. This is exactly what happens in the undirected case. If G and H are undirected graphs, then the homomorphism $f : G \to H$ is said to be injective if f is injective on the neighbourhood of every vertex in G.

Injective homomorphisms of undirected graphs are studied in [22, 23, 24, 25, 31]. In some cases they are referred to as *partial covers*, in contrast to full covers which may be viewed as bijective homomorphisms (the mapping is bijective on neighbourhoods). If the target in the homomorphism problem is a complete undirected graph, one obtains the *injective chromatic number* [31]. This is the smallest n such that a given graph G has an injective homomorphism to K_n.

Hahn, Kratochvíl, Širáň and Sotteau [31] were particularly interested in determining the injective chromatic number of the hypercube because of its connections to coding theory. It has to be noted here that in their definition of the injective chromatic number, they allow the possibility that adjacent vertices may receive the same colour. A colouring in their sense may therefore not be a homomorphism to an undirected complete graph, unless we include a loop at every vertex of the complete graph. Since we are mostly interested in complexity results, this will be the only type of result from [31] that we will list here. The injective chromatic number problem (ICN_k) may be stated formally as follows.

Problem 1.7 ICN_k

Instance: A graph G and a natural number k.

Question: Is there an injective k-colouring of G?

Theorem 1.7.1 (Hahn, Kratochvíl, Širáň and Sotteau [31]). *The problem ICN_k is NP-complete for every fixed $k \geq 3$.*

Fiala and Kratochvíl [22, 23, 24] and Fiala, Kratochvíl and Pór [25] consider the more general problem of injective homomorphisms where the target is not a complete graph. In all of these papers it is pointed out that the injective homomorphism problem is connected to so-called $L(2, 1)$ labellings of graphs: adjacent vertices must receive labels that differ by at least two, while vertices at distance two must receive labels that differ by at least one. This has applications in radio frequency assignment. Radio towers (for example cell phone towers) that are close together (where interference is quite possible) need frequencies that are far apart. Towers that are not as close to each other may only need frequencies that differ by a smaller amount.

The complexity results in these papers are mostly centered around a family of graphs which we describe next. Denote by $\Theta(a_1, a_2, \ldots, a_n)$ the (multi)graph that is formed by joining two vertices by n internally disjoint paths of lengths $a_1, a_2, \ldots, a_n$. An abbreviated version of this notation is as follows: $\Theta(a_1^{k_1}, a_2^{k_2}, \ldots, a_n^{k_n})$ is taken to mean that there are k_i paths of length a_i joining the two vertices.

The authors show that the family of graphs defined above exhibit both problems that are polynomial and problems that are NP-complete:

- $\Theta(a^n)$ is polynomial for every fixed a.
- $\Theta(a^i, b^j)$ is polynomial for every odd a, b.

- ❐ $\Theta(a^i, b^j)$ is NP-complete for every a and b of different parity.
- ❐ $\Theta(a, b, c)$ is NP-complete if c is divisible by $a + b$.
- ❐ $\Theta(1, 2, c)$ is NP-complete for every c.
- ❐ $\Theta(a^i, b^j)$ is polynomial when a and b are divisible by the same power of 2 or if $i + j \leq 2$. It is is NP-complete otherwise.
- ❐ $\Theta(1, 2, a)$ and $\Theta(1, 3, b)$ are NP-complete for all positive integers $a > 2$ and $b > 3$.
- ❐ For any three distinct odd positive integers a, b and c, $\Theta(a, b, c)$ is NP-complete.

As with ordinary homomorphisms, the hope is that the injective problems will exhibit a dichotomy. Towards this end Fiala and Kratochvíl [24] considered the list version of the injective homomorphism problem.

The list version of the injective homomorphism problem with target H, is the following problem (INJ-LIST-HOM$_H$).

Problem 1.8 INJ-LIST-HOM$_H$

Instance: A graph G and lists $L(v) \subseteq V(H)$.

Question: Does there exist an injective homomorphism $f : G \to H$ such that $f(v) \in L(v)$ for every $v \in V(G)$?

The lists $L(v)$ are to be thought of as admissible images for the vertex $v \in V(G)$. Fiala and Kratochvíl [24] were able to show that for this problem there is a dichotomy.

Theorem 1.7.2 (Fiala and Kratochvíl [24])**.** *The list, injective, homomorphism problem with target H is solvable in linear time if the graph H contains at most one cycle in each component. It is NP-complete otherwise.*

It is also interesting to note that injective graph homomorphisms are being applied to problems in mathematical biology [11, 17, 18, 21]. Here the interactions between proteins in a given species are modelled as an undirected graph. The problem then is to consider the protein networks of two different species and to determine whether a subgraph of one maps under an injective homomorphism to the network of the other. If it does, it may indicate that the two species share a sub-network that may have been passed along by a common ancestor.

As before, our focus will be on directed graphs. We will consider two versions of the injective homomorphism problem: (i) reflexive targets, where each vertex in the target has a loop and (ii) irreflexive targets, where no vertex in the target has a loop. A dichotomy was found for the reflexive version, while the irreflexive version proved to be much harder to deal with. We show that a polynomial time transformation exists between all digraph homomorphism problems and certain irreflexive injective digraph homomorphism problems. To some degree this explains the difficulty of the irreflexive case.

2

The Complexity of Colouring by Local Tournaments

2.1 Introduction

In this chapter we consider a generalization of the theorem by Bang-Jensen, Hell and MacGillivray [5] on the complexity of colouring by semi-complete digraphs. Recall that a *semi-complete* digraph has the property that between every pair of vertices there is at least one arc; parallel arcs and loops are not allowed, but a pair of symmetric arcs is allowed.

Theorem 2.1.1 (Bang-Jensen, Hell and MacGillivray [5]). *Let H be a semi-complete digraph.*

- *If H contains at most one directed cycle, then H-colouring is polynomial time solvable.*
- *Otherwise H-colouring is NP-complete.*

There is the related notion of a locally semi-complete digraph. A digraph H, is said to be *locally semi-complete* if for every vertex v of H, both $N^+(v)$ and $N^-(v)$

induce semi-complete digraphs. A special case of this is that of a local tournament. A *local tournament*, H, is a digraph such that between every pair of vertices there is at most one arc and that for every vertex v of H both $N^+(v)$ and $N^-(v)$ induce tournaments.

Bang-Jensen introduced the notion of locally semi-complete digraphs in [1] where it was shown that many of the known results on tournaments generalize to this family of digraphs.

It is, of course, the case that Theorem 2.1.1 applies to tournaments since they are a special case of semi-complete digraphs. In this chapter, we aim to generalize this special case of Theorem 2.1.1 to the class of local tournaments.

We will see that Theorem 2.1.1 doesn't generalize in the way one would expect. In particular there are unicyclic local tournaments such that the corresponding colouring problem is NP-complete. This is in contrast with unicyclic tournaments that all define polynomial colouring problems.

Since a classification of the complexity of smooth digraphs has already been obtained (Conjecture 1.2.4 was proved in [8]), this would imply that the complexity of some local tournaments are already known. The result in [8] was obtained by using results from universal algebra though. We give here a completely graph theoretic proof of the complexity of colouring by local tournaments. This is hopefully more appealing to graph theorists. In fact, the graph theoretic proof actually does more than just this. It showcases all of the techniques that have been built up over the last two decades in the study of the complexity of graph homomorphisms.

2.2 Some Tools

In addition to the indicator and sub-indicators mentioned in Chapter 1, we also need a few other tools to fully discuss the complexity of colouring by local tournaments.

2.2.1 The Consistency Check

Suppose that H is a fixed digraph that will act as the target in a homomorphism problem. As input to the problem we have a digraph G. In trying to find a homomorphism $G \to H$, we may start the process by assigning a list $L(v) = V(H)$ to each vertex v of G. These lists record possible images for the vertices of G and initially every vertex of H is a possible image for any given vertex of G. The algorithm we describe in this section processes each list $L(v)$, $v \in V(G)$, by removing any vertices from $L(v)$ that cannot possibly be images of v.

The lists attached to each vertex of G are said to be *consistent* if for any arc uv of G the following two properties hold:

- for any $x \in L(u)$, there exists $y \in L(v)$ such that xy is an arc of H and
- for any $b \in L(v)$, there exists $a \in L(u)$ such that ab is an arc of H.

The goal of the consistency check (Algorithm 2.1) is to reduce the initial lists to ones that are consistent. Our presentation here follows [37].

Algorithm 2.1 The Consistency Check

INPUT: A digraph G with lists $L(v) = V(H)$, $v \in V(G)$.

TASK: Reduce the lists to $L^*(v) \subseteq V(H)$, $v \in V(G)$, that are consistent.

ACTION: Initially set all lists $L^*(v) = L(v)$, and then, as long as changes occur, process each arc uv of G repeatedly as follows: remove from $L^*(u)$ any x for which no element $y \in L(v)$ has xy an arc in H, and remove from $L^*(v)$ any b for which no $a \in L^*(u)$ has ab an arc in H.

The consistency check is often used as a building block in designing polynomial time algorithms for the digraph homomorphism problem [37]. It is also known as the *arc-consistency check* since it checks for consistency across arcs of G. Higher order consistency checks are also possible [37].

2.2.2 The $\underline{X}$-enumeration and the Graft Extension

An enumeration $\{h_1, h_2, \ldots, h_n\}$ of the vertices of a digraph H is called an $\underline{X}$-*enumeration* if the following property holds: if $h_i h_j$ and $h_k h_l$ are arcs of H, then $\min(h_i, h_k)\min(h_j, h_l)$ is also an arc of H, where the minimum is taken with respect to the $\underline{X}$-enumeration.

Theorem 2.2.1 (Gutjahr, Woeginger and Welzl [29]). *Let H be a digraph such that H admits an $\underline{X}$-enumeration. Then the H-colouring problem is solvable in polynomial time.*

This result follows by running the consistency check on the input digraph (this is not the original algorithm presented in [29]). If a list becomes empty at any point during the consistency check, then there is no homomorphism to the target. If, on the other hand, the resulting lists are nonempty the minimum element (with respect to the $\underline{X}$-enumeration) in each list defines a homomorphism to the target [37].

There is one more result from [29] that we need, the so-called graft extension. We will consider a slightly more general problem than the one presented in [29]. Gutjahr, Woeginger and Welzl [29] only considered the H-colouring problem in their paper, HOM_H. We will show here that their result actually applies to a more general problem, namely that of list homomorphisms.

The list homomorphism problem with target H is the following decision problem ($\mathrm{LIST\text{-}HOM}_H$).

Problem 2.1 $\mathrm{LIST\text{-}HOM}_H$

INSTANCE:	$(G, \mathcal{L})$: A digraph G with lists $L(v) \subseteq V(H)$, $v \in V(G)$.
QUESTION:	Does there exist a homomorphism $f : G \to H$ such that $f(v) \in L(v)$ for every $v \in V(G)$?

Let H_1 be a loop-free digraph that has an $\underline{X}$-enumeration, say $\{h_1, h_2, \ldots, h_n\}$.

Let H_2 be a digraph such that H_2-colouring is polynomial. We form a new digraph H by deleting the vertex h_n from H_1 and replacing it by the digraph H_2: every vertex $h_i \in V(H_1)$ that is adjacent to (from) h_n is now adjacent to (from) every vertex in H_2. The digraph H is called the $\underline{X}$-graft(H_1,H_2).

Theorem 2.2.2. *Let H=$\underline{X}$-graft(H_1,H_2) such that LIST-HOM$_{H_2}$ is polynomial. Then LIST-HOM$_H$ is polynomial.*

Proof. Let $(G, \mathcal{L})$ be an instance of LIST-HOM$_H$. Modify H_1 by adding a loop at h_n. Clearly, the existence of a homomorphism of G to H_1 is a necessary condition for the existence of a homomorphism of G to H.

First, alter the lists by replacing any vertices of H_2 by h_n in any list where they occur. Now, apply the arc-consistency check. If it fails, then G is a NO-instance. Hence assume the consistency check succeeds. By the discussion earlier, the mapping $f(x) = \min L(x)$, where the minimum is with respect to the $\underline{\text{X}}$-enumeration, is a list-homomorphism $G \rightarrow H_1$.

Since the minimum element in each list is chosen, the vertices that map to h_n in this list-homomorphism of G to H_1 map to vertices of H_2 in any list-homomorphism of G to H. By the construction of H and the consistency check, there is now a list-homomorphism of G to H if and only if the subgraph of G induced by $f^{-1}(h_n)$ admits a list-homomorphism to H_2, where the lists are the intersections of the initially given lists with $V(H_2)$. As LIST-HOM$_{H_2}$ is polynomial, the result follows. ∎

Corollary 2.2.3 (Gutjahr, Woeginger and Welzl [29]). *Let H=$\underline{X}$-graft(H_1,H_2) such that HOM$_{H_2}$ is polynomial. Then HOM$_H$ is polynomial.*

This result follows since HOM$_H$ is a special case of LIST-HOM$_H$ where each list $L(v) = V(H)$ for each $v \in V(G)$.

2.2.3 The Frobenius-Schur Index

The result discussed in this section is a purely number theoretic result. It will help in choosing the correct lengths for directed paths that are to act as indicators and sub-indicators.

Given a set of relatively prime positive integers $B = \{a_1, a_2, \ldots, a_n\}$, a linear combination of these integers is an expression of the form

$$x_1a_1 + x_2a_2 + \cdots + x_na_n, \tag{2.1}$$

where each $x_i \in \{0, 1, 2, \ldots\}$. A natural question to ask is for the smallest integer ϕ such that each every integer $t \geq \phi$ can be represented in the form of equation 2.1. The existence of such an integer ϕ is guaranteed by a result of Schur (see [12]).

Lemma 2.2.4. *Let S be a nonempty set of positive integers which is closed under addition. Let d be the greatest common divisor of the integers in S. Then there exists a positive integer N such that td is in S for every integer $t \geq N$.*

Since the set of all linear combinations of elements in B is closed under addition, Schur's result guarantees a threshold above which every integer is of the form 2.1.

Frobenius (according to [10]) then posed the problem of finding the smallest integer $\phi = \phi(a_1, a_2, \ldots, a_n)$ such that every integer $t \geq \phi$ is of the form 2.1 (or at least good bounds this number). The integer ϕ is known as the *Frobenius-Schur index* of the set B [12] or as the *conductor* of the set B [51]. Equivalently, one may ask for the largest integer not representable as 2.1. In general, this is a very hard problem with a rich literature [48]. The problem sometimes also goes by the name of *the money changing problem*. Given a fixed set of coins (the set B) what is the largest amount of money that cannot be changed using the coins in B? See [30] or [51] for more on this.

We are typically interested in finding the Frobenius-Schur index of the cycle lengths of a strong round local tournament D. By Lemma 2.3.6 we know that the cycle lengths of D is an interval of integers, $\{\ell, \ell+1, \ldots, n\}$, where ℓ is the girth of D and $n = |V(D)|$. Fortunately, in this case, the Frobenius-Schur index is known exactly.

Lemma 2.2.5 (Brauer [10])**.** *Let ℓ be a positive integer. Then*

$$\phi(\ell, \ell+1, \ldots, n) = \left\lfloor \frac{n-2}{n-\ell} \right\rfloor \ell.$$

We would like to give a short justification of this result as this may aid the reader later on when we apply this result to local tournaments.

Consider the following table. Row k of the table contains numbers that can be written as a linear combination using k numbers from $\{\ell, \ell+1, \ldots, n\}$.

$$\begin{gathered}
\ell,\ \ell+1,\ \ell+2,\ \ldots,\ n = \ell + (n-\ell). \\
2\ell,\ 2\ell+1,\ 2\ell+2,\ \ldots\ 2n = 2\ell + 2(n-\ell). \\
3\ell,\ 3\ell+1,\ 3\ell+2,\ \ldots,\ 3n = 3\ell + 3(n-\ell). \\
\vdots \\
k\ell,\ k\ell+1,\ k\ell+2,\ \ldots,\ kn = k\ell + k(n-\ell). \\
(k+1)\ell,\ (k+1)\ell+1,\ (k+1)\ell+2,\ \ldots,\ (k+1)n = (k+1)\ell + (k+1)(n-\ell). \\
\vdots
\end{gathered}$$

Note that each row is a list of consecutive integers. Therefore each number that cannot be written as a linear combination has to occur somewhere between the rows of the table. Between the last entry of a row and the first entry of the next row there is a "gap" of integers that cannot be written as a linear combination. The size of the gap also decreases as one moves down the table so that it is inevitable that it eventually closes.

We now ask what is the largest integer that occurs somewhere between the rows of the table. This integer has to occur in a gap of size at least two (a gap of size one corresponds to two consecutive integers). So what is the largest integer k such that between rows k and $(k+1)$ there is still a gap of size at least two. That is

$$(k+1)\ell - kn = (k+1)\ell - k(\ell + n - \ell) = \ell - k(n-\ell) \geq 2,$$

implying

$$k = \left\lfloor \frac{\ell - 2}{n - \ell} \right\rfloor.$$

So as soon as we move beyond this row, every integer can be written as a linear combination. Therefore the smallest integer ϕ defined above occurs at $(k+1)\ell$. That is

$$\phi(\ell, \ell+1, \ell+2, \ldots, n) = \left(\left\lfloor \frac{\ell - 2}{n - \ell} \right\rfloor + 1 \right) \ell = \left\lfloor \frac{n - 2}{n - \ell} \right\rfloor \ell.$$

2.3 Local Tournaments

Since we aim to give a graph theoretic proof of the complexity of HOM_H where H is a local tournament, it comes as no surprise that the structure of local tournaments plays a central role in this proof. In this section we state the results on the structure of locally semi-complete digraphs (since local tournaments are a special case of these) that we need. Bang-Jensen and Gutin [3] is a standard reference on these matters.

A digraph D is said to be *strong* or *strongly connected* if for every pair of vertices x and y in D, there is directed path joining x to y and a directed path joining y to x.

A *strong component* of a digraph D is a maximal induced sub-digraph of D that is strong.

The strong components of a digraph D are disjoint and can be labeled as D_1, $D_2, \ldots D_t$ such that there is no arc from D_i to D_j unless $i < j$. We call such an ordering an *acyclic ordering* of the strong components of D.

Theorem 2.3.1 (Guo and Volkmann [3, 27]). *Let D be a connected locally semi-complete digraph that is not strong and let $D_1, D_2, \ldots, D_p$ be the acyclic ordering of the strong components of D. Then D can be decomposed into $r \geq 2$ induced subgraphs $D'_1, D'_2, \ldots, D'_r$ as follows:*

$$D'_1 = D_p, \quad \lambda_1 = p,$$
$$\lambda_{i+1} = \min\{j \mid N^+(D_j) \cap V(D'_i) \neq \emptyset\},$$
$$\textit{and } D'_{i+1} = D\langle V(D_{\lambda_{i+1}}) \cup V(D_{(\lambda_{i+1})+1}) \cup \cdots \cup V(D_{\lambda_i - 1})\rangle.$$

The sub-digraphs $D'_1, D'_2, \ldots, D'_r$ satisfy the properties below:

(a) D'_i consists of some strong components of D and is semi-complete for $i = 1, 2, \ldots, r$.

(b) D'_{i+1} dominates the initial component of D'_i and there exists no arc from D'_i to D'_{i+1} for $i = 1, 2, \ldots, r-1$.

(c) If $r \geq 3$, then there is no arc between D'_i and D'_j for i, j satisfying $|j - i| \geq 2$.

If D is a connected locally semi-complete digraph that is not strong, then the unique sequence $D'_1, D'_2, \ldots, D'_r$ defined in Theorem 2.3.1 is called the *semi-complete decomposition* of D.

A digraph on n vertices is said to be *round* if we can label its vertices $v_1, v_2, \ldots, v_n$ so that for each i, $N^+(v_i) = \{v_{i+1}, \ldots, v_{i+d^+(v_i)}\}$ and $N^-(v_i) = \{v_{i-d^-(v_i)}, \ldots, v_{i-1}\}$, where all subscripts are taken modulo n.

Let D be a digraph with vertex set $\{v_1, v_2, \ldots, v_n\}$ and let $G_1, G_2, \ldots, G_n$ be digraphs which are pairwise vertex disjoint. The *composition* $D[G_1, G_2, \ldots, G_n]$ is the digraph H with vertex set $V(G_1) \cup V(G_2) \cup \ldots \cup V(G_n)$ and arcs $(\cup_{i=1}^{n} A(G_i)) \cup \{g_i g_j \mid g_i \in V(G_i), g_j \in V(G_j), v_i v_j \in A(D)\}$.

A locally semi-complete digraph D is *round decomposable* if there exists a round local tournament R on $r \geq 2$ vertices such that $D = R[S_1, \ldots, S_r]$, where each S_i is a strong semi-complete digraph.

For a strong digraph D a set $S \subseteq V(D)$ is a *separator* or *separating set* if $D - S$ is not strong.

Lemma 2.3.2 (Bang-Jensen, Guo, Gutin and Volkmann [2, 3])**.** *Let D be a strong locally semi-complete digraph which is not semi-complete. Then D is not round decomposable if and only if the following conditions are satisfied:*

(a) There is a minimal separating set S such that $D - S$ is not semi-complete and for each such S, $D\langle S\rangle$ is semi-complete and the semi-complete decomposition of $D - S$ has exactly three components D'_1, D'_2, D'_3;

(b) There are integers α, β, μ, ν with $\lambda_2 \leq \alpha \leq \beta \leq p-1$ and $p+1 \leq \mu \leq \nu \leq p+q$ such that

$$N^-(D_\alpha) \cap V(D_\mu) \neq \emptyset \quad \text{and} \quad N^+(D_\alpha) \cap V(D_\nu) \neq \emptyset,$$
$$\text{or } N^-(D_\mu) \cap V(D_\alpha) \neq \emptyset \quad \text{and} \quad N^+(D_\mu) \cap V(D_\beta) \neq \emptyset,$$

where $D_1, D_2, \ldots, D_p$ and $D_{p+1}, \ldots, D_{p+q}$ are the acyclic orderings of the strong components of $D - S$ and $D\langle S\rangle$, respectively, and D_{λ_2} is the initial component of D'_2.

The structure described in Lemma 2.3.2 is illustrated in Figure 2.1.

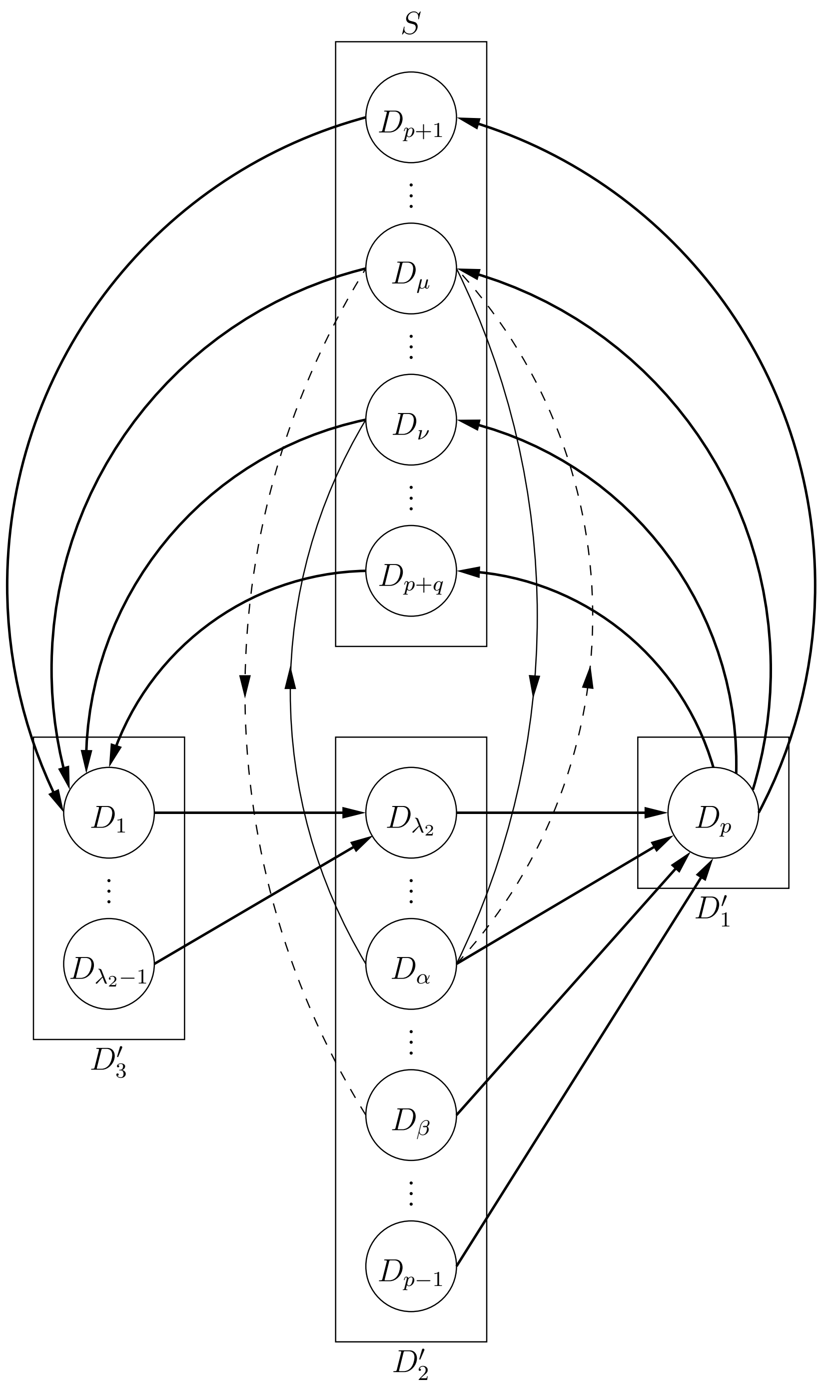

Figure 2.1: The structure of a strong locally semi-complete digraph that is not semi-complete and not round decomposable.

Theorem 2.3.3 (Bang-Jensen, Guo, Gutin and Volkmann [2, 3]). *Let D be a connected locally semi-complete digraph. Then exactly one of the following possibilities holds.*

(a) D is round decomposable with a unique round decomposition given by

$$D = R[D_1, D_2, \ldots, D_\alpha],$$

where R is a round local tournament on $\alpha \geq 2$ vertices and D_i is a strong semi-complete digraph for $i = 1, 2, \ldots, \alpha$;

(b) D is not round decomposable and not semi-complete and it has the structure described in Lemma 2.3.2;

(c) D is a semi-complete digraph which is not round decomposable.

Proposition 2.3.4 (Bang-Jensen, Guo, Gutin and Volkmann [2, 3]). *Let D be a strong non-round decomposable locally semi-complete digraph and let S be a minimal separating set of D such that $D - S$ is not semi-complete. Let $D_1, D_2, \ldots, D_p$ be the acyclic ordering of the strong components of $D - S$ and $D_{p+1}, D_{p+2}, \ldots, D_{p+q}$ be the acyclic ordering of the strong components of $D\langle S\rangle$. Suppose that there is an arc $s \rightarrow v$ from S to D'_2 with $s \in V(D_i)$ and $v \in V(D_j)$, then*

$$D_i \cup D_{i+1} \cup \cdots \cup D_{p+q} \mapsto D'_3 \mapsto D_{\lambda_2} \cup \cdots \cup D_j.$$

Here, $A \mapsto B$, means that A dominates B and there are no arcs from B to A.

Lemma 2.3.5 (Bang-Jensen and Gutin [3]). *Let R be a strong round local tournament and let C be a shortest cycle of R and suppose C has $k \geq 3$ vertices. Then for every round labeling $v_0, v_1, \ldots, v_{n-1}$ of R such that $v_0 \in V(C)$ there exist indices $0 < a_1 < a_2 < \ldots a_{k-1} < n$ so that $C = v_0 v_{a_1} v_{a_2} \cdots v_{a_{k-1}} v_0$.*

Lemma 2.3.6 (Bang-Jensen and Gutin [3])**.** *A strong round local tournament R on r vertices has cycles of length $k, k+1, \ldots, r$, where k is the girth of R.*

Lemma 2.3.7 (Bang-Jensen and Gutin [3])**.** *If a strong round local tournament with r vertices has a cycle of length k through a vertex v, then it has cycles of all lengths $k, k+1, \ldots, r$ through v.*

Strictly speaking a round digraph is also round decomposable (all $|D_i| = 1$). We prefer to distinguish between round and round decomposable digraphs. So when a local tournament is said to be round decomposable then at least one $|D_i| \geq 3$. Furthermore, since a connected locally semi-complete digraph that is not strongly connected is round decomposable (an acyclic ordering of its strong components), a round local tournament (in our sense) is strongly connected and therefore Hamiltonian (by Lemma 2.3.6).

In proving some of our results, we will use the sub-indicator construction (both the vertex and arc versions). A glance at Lemmas 1.4.2 and 1.4.3 confirms the fact that in order to apply a sub-indicator to a digraph H, one has to know that H is a core. We show next that a connected local tournament is indeed a core. The following lemma will be useful in this regard.

Lemma 2.3.8 (Bang-Jensen [1, 3])**.** *A locally semi-complete digraph has a hamilton path if and only if its underlying graph is connected.*

Proposition 2.3.9. *A connected locally semi-complete digraph D is a core.*

Proof. By Lemma 2.3.8, D has a hamilton path, say $v_1, v_2, \ldots, v_n$.

Let $f : D \to D'$ be a retraction of D where D' is a subgraph of D. Let v be a vertex in D'. By the enumeration above $v = v_i$ for some $i \in \{1, 2, \ldots, n\}$ and by the retraction $f(v_i) = v_i$.

Consider the vertices v_{i-1} and v_{i+1} (at least one of these exists), say v_{i-1}. Let $f(v_{i-1}) = v_j$, $j \in \{1, 2, \ldots, n\}$. Clearly $j \neq i$. Since $v_{i-1}v_i \in A(D)$, $f(v_{i-1})f(v_i) =$

$v_j v_i \in A(D)$. That is $v_{i-1}, v_j \in N^-(v_i)$. Therefore by the locally semi-complete property of D, v_{i-1} and v_j are adjacent. If $i-1 \neq j$, the arc(s) between v_{i-1} and v_j are not preserved by f. So the image of v_{i-1} is exactly v_{i-1} under the retraction. If v_{i+1} exists, the same reasoning shows that $f(v_{i+1}) = v_{i+1}$.

Repeating the argument along the hamilton path eventually shows that every vertex has to map to itself. Therefore D is a core. ∎

2.4 Connected vs. Disconnected Local Tournaments

When H is a disconnected local tournament such that each component of H (all of which are local tournaments themselves) is polynomial time solvable, then HOM_H is also polynomial time solvable. NP-completeness results for disconnected local tournaments are much harder to obtain.

The NP-completeness results that follow are all for connected local tournaments since we only give polynomial time transformations from NP-complete problems to HOM_H when H is connected. A natural conjecture would be: if a disconnected local tournament H contains at least one component that is NP-complete, then HOM_H is NP-complete. The difficulty in proving this lies with constructing a polynomial transformation from some NP-complete problem. In general, it is hard to set up the transformation without forcing certain vertices of the transformed instance to map to a specific component of H. This may have the unintended consequence of restricting the images of one or more vertices of the transformed instance too severely.

On the other hand one can easily obtain a polynomial time Turing reduction. Here, an instance of some NP-complete problem Q is transformed into many different instances of HOM_H. The transformation has to run in polynomial time and furthermore an instance I of Q is a yes instance if and only if at least one the transformed instances is a yes instance of HOM_H. This shows that for a disconnected local tournament with at least one NP-complete component, HOM_H is polynomial time solvable if and only if P=NP. Therefore solving HOM_H in polynomial time for a disconnected local tournament H with at least one NP-complete component, is highly unlikely.

Let H be a disconnected local tournament that is also a core and let H' be a component of H such that $\mathrm{HOM}_{H'}$ is NP-complete. The polynomial time Turing reduction is from $\mathrm{HOM}_{H'}$ to HOM_H. Let G be an instance of $\mathrm{HOM}_{H'}$. We now form $|V(H')|$ instances of HOM_H as follows: take $|V(H')|$ copies each of G and H', let $v \in V(G)$ and $V(H') = \{h_1, h_2, \ldots, h_n\}$. Denote by F_i the graph obtained by identifying the vertex v in G with the vertex h_i in H', $1 \leq i \leq n$.

It is now easy to see that $G \to H'$ if and only if there exists at least one $i \in \{1, 2, \ldots, n\}$ such that $F_i \to H$.

From now on we assume that all local tournaments are connected.

2.5 Road map

As an aid to reading our proof that local tournaments exhibit a dichotomy, we provide the following road map.

At its highest level the proof is divided into five cases, as shown below.

1. **Acyclic local tournaments:** Acyclic local tournaments were shown to define polynomial problems by Hell, Zhou and Zhu [38]. Therefore Case 1 has already been dealt with.

2. **Unicyclic local tournaments:** The unicyclic case is treated in Section 2.6 and relies on the graft extension as well as some direct NP-completeness reductions. Unicyclic local tournaments have both problems that are polynomial and problems that are NP-complete. The remaining cases all define NP-complete problems.

3. **Round local tournaments:** Round local tournaments with at least two directed cycles are discussed in Section 2.7. The first lemma we prove deals with the case of a round local tournament with at least two cycles and a unique cycle of shortest length. Such a round local tournament is shown to define an NP-complete problem via an NP-completeness reduction. The rest of the section examines a smallest counterexample (minimum number of vertices and minimum number of arcs). Therefore we assume that there is a round local tournament, T, with at least two cycles, such that T-colouring is not NP-complete. We then derive some structural information about this counterexample showing that it doesn't exist. The structural information is derived using the indicator, vertex sub-indicator and arc sub-indicator constructions. This allows us to conclude that a round local tournament with at least two cycles defines an NP-complete problem.

4. **Round decomposable local tournaments:** The round decomposable case is dealt with in Section 2.8. Here we rely exclusively on the indicator and vertex sub-indicator constructions to show that round decomposable local tournaments with at least two cycles all define NP-complete problems.

5. **Non-round decomposable local tournaments:** In Section 2.9 non-round decomposable local tournaments are shown to define NP-complete problems. Here we again examine a smallest counterexample, derive structural information about it using the indicator and vertex sub-indicator constructions, and finally conclude that it doesn't exist.

This covers all local tournaments and may we therefore conclude that they do indeed exhibit a dichotomy.

2.6 Unicyclic Local Tournaments

Let T be a unicyclic local tournament. Then the cycle, C, in T is an induced cycle. If C has length at least four and T is connected, then $V(T) = V(C)$. Otherwise there is a vertex not on the cycle adjacent to/from a vertex on the cycle and since T is a local tournament, this vertex is adjacent to/from every vertex on the cycle. Hence T is not a local tournament (since C is induced). In this case T-colouring is polynomial. We may therefore assume that the cycle has length 3 and $T = A[D_1, D_2, \ldots, D_l]$. Here A is an acyclic local tournament and furthermore there exists a $j \in \{1, 2, \ldots, l\}$ such that $|D_i| = 1$ for $1 \leq i \leq j-1$, $D_j = C_3$, and $|D_i| = 1$ for $j+1 \leq i \leq l$.

The unicyclic local tournament T may also be viewed as follows. Let S be the set of neighbours (in-and out-neighbours) of the 3-cycle in T. Then $V(T) \setminus S$ is the union of two disjoint sets of vertices: those that come before S in the ordering shown above, call these A, and those that come after S in the ordering above, call these B. Define the following three induced sub-digraphs: $T_1 = T[A]$, $T_2 = T[S]$ and $T_3 = T[B]$. Each T_i is local tournament and T_1 and T_3 are acyclic as well. Note that T_1 or T_3 may be empty. This general structure is illustrated in Figure 2.2 where we have written $D_t = \{y_t\}$ for $t \in \{1, 2 \ldots, l\} \setminus \{j\}$. Note that the arc $y_{j-1}y_{j+1}$ may or may not be present depending on whether T_2 is or isn't a tournament. Furthermore, since T is a local tournament, y_{j-1} dominates the cycle and y_{j+1} is dominated by the cycle.

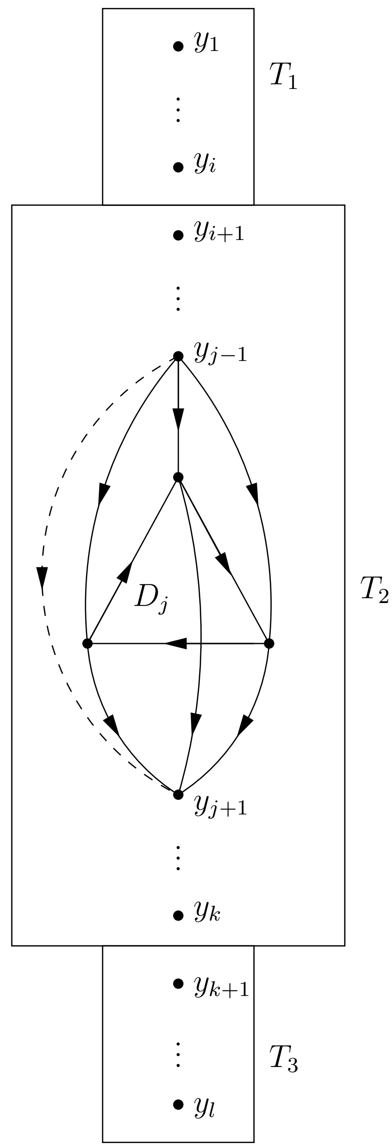

Figure 2.2: The structure of a unicyclic local tournament that's not a directed cycle.

In discussing the complexity of colouring by a unicyclic local tournament where $T_1 = \emptyset$ and $T_3 \neq \emptyset$, or where $T_1 \neq \emptyset$ and $T_3 = \emptyset$, we only need to consider one of these situations since the digraphs are converses of one another.

Lemma 2.6.1. *If T is a unicyclic local tournament in which T_2 is a unicyclic tournament and $T_1 \neq \emptyset$ and $T_3 = \emptyset$, then T-colouring is polynomial.*

Proof. We show that T is an instance of the graft extension.

Let $V(T_1) = \{y_1, y_2, \ldots, y_i\}$ and $V(T_2) = \{y_{i+1}, y_{i+2}, \ldots, y_{j-1}, 0, 1, 2, y_{j+1}, \ldots y_k\}$ (the 3-cycle's vertices are labeled $0, 1$ and 2). Define T' to be the acyclic local tournament obtained by contracting the vertices $\{0, 1, 2, y_{j+1}, \ldots, y_k\}$ to a single vertex (in T) and deleting the resulting loop. Let T'' be the unicyclic tournament induced by the vertices $\{0, 1, 2, y_{j+1}, \ldots, y_k\}$. Since T' is an acyclic local tournament, it has an $\underline{X}$-enumeration. Also, T'' is a unicyclic tournament and so T''-colouring is polynomial. It can now be seen that $T = \underline{X}$-graft(T', T''). ∎

It is worth noting at this point that if T_2 is a unicyclic tournament and both T_1 and T_2 are empty, then T-colouring is polynomial.

In determining the complexity of colouring by a unicyclic local tournament, the local tournament LT_5 shown below in Figure 2.3 plays a special role.

Lemma 2.6.2. *Let T be a unicyclic local tournament. If T contains LT_5 as an induced subgraph, then T-colouring is NP-complete.*

Proof. The proof of NP-completeness is via a reduction from monotone one-in-three 3-SAT (i.e. no negations).

Let $T = A[D_1, D_2, \ldots, D_l]$, with $D_j = C_3$, $j \in \{1, 2, \ldots, l\}$ and $|D_i| = 1$ for $i \in \{1, 2, \ldots, l\}$, $i \neq j$. Label the vertices in D_i with $1 \leq i \leq j-1$ as $t_1, t_2, \ldots t_{j-1}$, the vertices in D_i with $j+1 \leq i \leq l$ as $b_1, b_2, \ldots, b_{l-j}$ and the vertices on the C_3 as $0, 1, 2$. Let $c = \min\{i \mid t_i \text{ dominates the } C_3\}$. Since T is a local tournament and

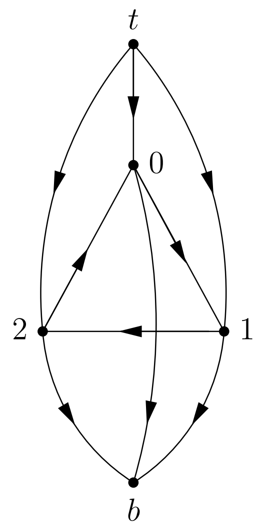

Figure 2.3: The special unicyclic local tournament, LT_5, on five vertices.

LT_5 is an induced subgraph, t_c is not adjacent to any b_i, $1 \leq i \leq l-j$. Let d be equal to the maximum i such that t_i dominates the C_3 and t_i is not adjacent to any b_s, $1 \leq s \leq l-j$. Denote the vertices $\{t_c, t_{c+1}, \ldots, t_d\}$ by F, the vertices $\{t_{d+1}, t_{d+2}, \ldots, t_{j-1}\}$ by M, the vertices $\{0, 1, 2\}$ by C, and the length of the path $t_c t_{c+1} \ldots t_{j-1}$ by k. In a similar way we are able to find an index e, such that the set $B = \{b_1, b_2, \ldots, b_e\}$ is the set of vertices being dominated by the 3-cycle.

The set F will be associated with "False" and $M \cup C$ with "True."

Consider the digraph K shown below in Figure 2.4.

In any homomorphism from K to T, the vertices ℓ_i and ℓ'_i, for $i = 1, 2, 3$, must map into $F \cup M \cup C$. The reason for this is that each ℓ_i (ℓ'_i) dominates a vertex on a 3-cycle, and a 3-cycle in K has to map to the 3-cycle in T.

The vertices v_1, v_4 and v_6 have to map in such a way that they are being dominated by the 3-cycle in T. Therefore each of v_1, v_4 and v_6 map into $C \cup B$. If ℓ_1 maps into F, then v_6 has to map into C. The same is true for ℓ_2 and v_1 and for ℓ_3 and v_4. Therefore if ℓ_1, ℓ_2 and ℓ_3 all map into F, then the oriented 7-cycle on $v_0, v_1, \ldots, v_6$ has to map into the C_3 in T. This is not possible since this oriented cycle has net-length

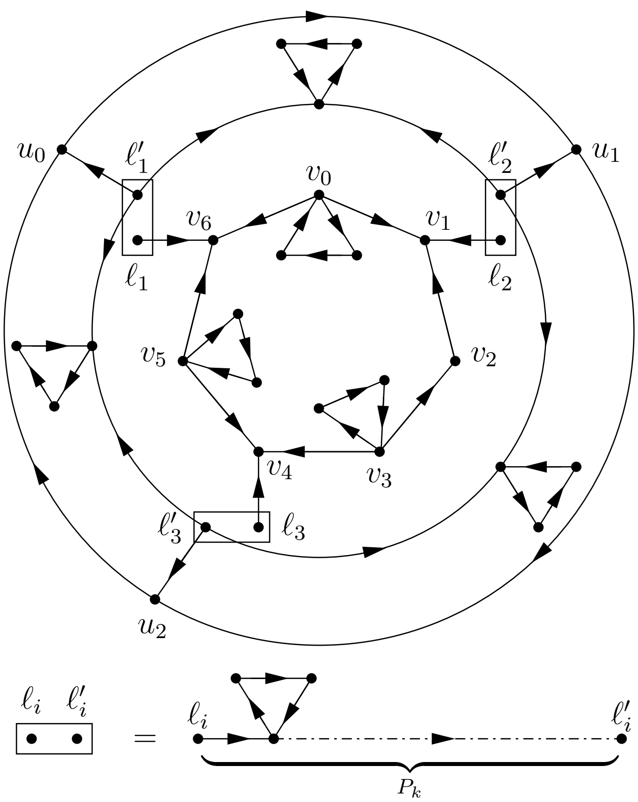

Figure 2.4: The gadget K for unicyclic local tournaments.

one. Thus at most two of ℓ_1, ℓ_2 and ℓ_3 can map into F.

If ℓ_i, $1 \leq i \leq 3$, maps into M, then ℓ_i' has to map into C: ℓ_i' has to map into $M \cup C$ in this case and ℓ_i' cannot map into M since the images of ℓ_i and ℓ_i' are joined by a directed walk of length k (the image of the P_k joining them) in T, the distance between two vertices in M is less than k and the subgraph induced by M is acyclic. On the other hand when ℓ_i, $1 \leq i \leq 3$ maps into C, then again ℓ_i' maps into C. Therefore if $\ell_i \mapsto M \cup C$, then $\ell_i' \mapsto C$ for $i = 1, 2, 3$. Furthermore if any two of $\{\ell_1', \ell_2', \ell_3'\}$ are mapped into C, then in fact these two vertices are mapped to the same vertex in C. The reason for this is that each pair of vertices from $\{\ell_1', \ell_2', \ell_3'\}$ have a common out-neighbour that is on a C_3 in K. If for instance ℓ_1' and ℓ_2' both mapped to vertex $y \in C$, then vertices u_0 and u_1 both have to map to vertex y^+, the successor of y on the C_3 in T. This is clearly not a homomorphism. The other two pairs from $\{\ell_1', \ell_2', \ell_3'\}$ follow in a similar way. Therefore in a homomorphism from K

to T, at most one of $\{\ell'_1, \ell'_2, \ell'_3\}$ can map into C. This implies that at most one of $\{\ell_1, \ell_2, \ell_3\}$ can map into $M \cup C$.

From before we know that at most two of ℓ_1, ℓ_2 and ℓ_3 can map into F. Therefore exactly one of ℓ_1, ℓ_2 and ℓ_3 must map into $M \cup C$ and exactly two of ℓ_1, ℓ_2 and ℓ_3 must map into F.

If there is a homomorphism $f : K \to T$ such that $f(\ell_i) \in M \cup C$, $f(\ell_{i+1}) \in F$ and $f(\ell_{i+3}) \in F$ (subscripts taken mod 3), then $f(\ell'_i) \in C$, so that $f(\ell'_{i+1}), f(\ell'_{i+2}) \notin C$. This implies that $f(\ell_{i+1}) = f(\ell_{i+2}) = t_c$ and that $f(\ell'_{i+1}) = f(\ell'_{i+2}) = t_{j-1}$.

We now show that as long as exactly one of $\{\ell_1, \ell_2, \ell_3\}$ map to $M \cup C$, then there exists a homomorphism from K to T. We denote by y a vertex in $M \cup C$ and by z a vertex in C. These homomorphisms are shown in Table 2.1.

Table 2.1: Homomorphisms from K to T.

ℓ_1	ℓ_2	ℓ_3	ℓ'_1	ℓ'_2	ℓ'_3	v_0	v_1	v_2	v_3	v_4	v_5	v_6	u_0	u_1	u_2
y	t_c	t_c	z	t_{j-1}	t_{j-1}	z	z^+	z	z^{++}	z	z^{++}	b_1	z^+	z^{++}	z
t_c	y	t_c	t_{j-1}	z	t_{j-1}	z	b_1	z^+	z	z^+	z	z^+	z	z^+	z^{++}
t_c	t_c	y	t_{j-1}	t_{j-1}	z	z^+	z^{++}	z^+	z	b_1	z^+	z^{++}	z^{++}	z	z^+

Given an instance of one-in-three 3-SAT with variables $x_1, x_2, \ldots, x_s$ and clauses $C_1, C_2, \ldots, C_t$, we construct a digraph D as follows. For each variable x_i there a corresponding vertex x_i and for each clause $\{\ell_1, \ell_2, \ell_3\}$, we use a copy of K. Identify each vertex x_i with the corresponding vertex in each of the clauses in which it appears. This forms the digraph D.

If there is a homomorphism from D to T, then by what was shown above we see that in a clause $\{\ell_1, \ell_2, \ell_3\}$, exactly one of the ℓ_is map to $M \cup C$, the others have to map to F. We can recover a truth assignment by declaring that a variable is "True" if its corresponding vertex x_i maps to $M \cup C$ declaring it to be "False" if it maps to F.

On the other hand if there is a satisfying truth assignment, then in a clause

$\{\ell_1, \ell_2, \ell_3\}$, at most one ℓ_i is "True." We can set up a homomorphism from K to T by mapping a vertex x_i to vertex 0 in T if the corresponding variable is "True" and mapping it to t_c if the variable is "False." As was shown above these pre-colourings extend to all of D.

Therefore $D \to T$ if and only if there is a satisfying truth assignment. This implies that T-colouring is NP-complete. ∎

Lemma 2.6.2 covers the case where T_2 is not a tournament since the local tournament in Figure 2.3 must be an induced subgraph in this case. Lemma 2.6.1 covers the case where T_2 is a tournament, $T_1 \neq \emptyset$ and $T_3 = \emptyset$ (or $T_1 = \emptyset$ and $T_3 \neq \emptyset$). All that remains therefore is the case where T_2 is a tournament and both T_1 and T_2 are not empty.

Lemma 2.6.3. *Let T be a unicyclic local tournament with T_1, T_2 and T_3 defined as above. If both T_1 and T_3 are non-empty and T_2 is a tournament, then T-colouring is NP-complete.*

Proof. The proof of NP-completeness is by a reduction from not-all-equal 3-SAT without negated variables. We label the vertices of T as in Figure 2.2 (note that the arc $y_{j-1}y_{j+1}$ is present in T since T_2 is a tournament). The vertices on the 3-cycle are labeled as $0, 1$ and 2.

In describing the reduction we use a number of gadgets. The first of which, G_1, is shown in Figure 2.5.

The property of G_1 of interest to us is that in any homomorphism $G_1 \to T$, the vertex labeled v cannot map to any vertex of the three cycle in T. If this happens, the oriented 5-cycle in G_1 is forced to map to the 3-cycle in T which is not possible since the oriented 5-cycle has net-length one.

Consider the directed path $y_{i+1}y_{i+2} \cdots y_{j-1}0$ in T. Suppose that this path has length t. The next gadget, G_2, is constructed from a directed path of length t and

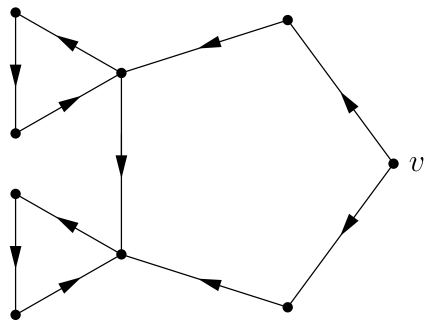

Figure 2.5: The first gadget G_1 for the last unicyclic case.

t copies of G_1. Form G_2 by identifying each vertex on the directed path, except the last vertex, with the vertex v in one copy of G_1. This is shown in Figure 2.6.

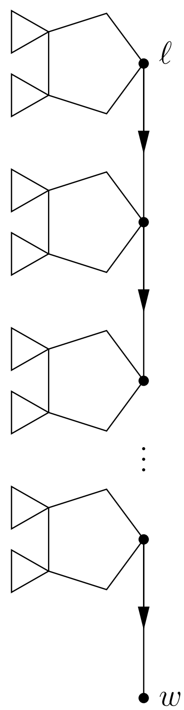

Figure 2.6: The second gadget G_2 for the last unicyclic case.

The third gadget, G_3, is shown in Figure 2.7. It is constructed from three copies of G_2 in which the three respective terminal vertices of the directed paths of length t are identified with particular vertices on an oriented 7-cycle. This oriented 7-cycle has net-length one, and some of its vertices are on 3-cycles. These vertices have to map to the 3-cycle in T under any homomorphism $G_3 \to T$. The gadget G_3 represents half of the clause gadget that will be used in the reduction from not-all-equal 3-SAT.

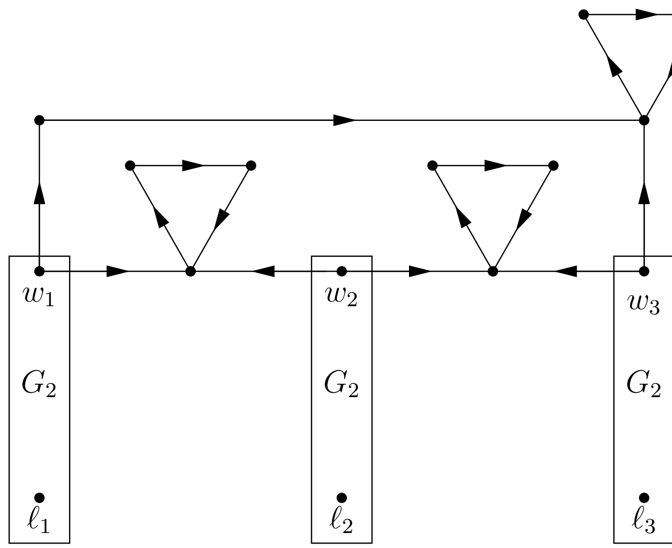

Figure 2.7: The third gadget G_3 for the last unicyclic case.

Let S_1 be the set of vertices in T_1 that are connected to the 3-cycle in T by a path of length two. Consider a homomorphism $G_3 \to T$. Under any such homomorphism ℓ_1, ℓ_2 and ℓ_3 have to map to vertices in T that can reach the 3-cycle in T along a path of length two. This follows from the fact each of ℓ_1, ℓ_2 and ℓ_3 has been identified with the vertex v in G_1. Therefore, the possible images of ℓ_1, ℓ_2 and ℓ_3 are $S_1 \cup \{y_{i+1}, y_{i+2}, \ldots, y_{j-1}\}$. On the other hand, the presence of the directed path of length t in each copy of G_2 in G_3 prevents ℓ_1, ℓ_2 and ℓ_3 from mapping to any of $\{y_{i+2}, y_{i+3}, \ldots, y_{j-1}\}$ (they are "too close" to the 3-cycle). Thus the possible images of ℓ_1, ℓ_2 and ℓ_3 are $S_1 \cup \{y_{i+1}\}$. If $\ell_j \mapsto y_{i+1}$ (for $j \in \{1, 2, 3\}$), then w_j maps to one of $\{0, 1, 2\}$. So if ℓ_1, ℓ_2 and ℓ_3 all mapped to y_{i+1}, then each w_j (for $j = 1, 2, 3$) maps into the 3-cycle. As a consequence of this, the oriented 7-cycle is forced to map to the 3-cycle in T, but this is not possible since the oriented 7-cycle has net-length one. Therefore ℓ_1, ℓ_2 and ℓ_3 cannot all map to y_{i+1}. It is easy to check that if at least one of ℓ_1, ℓ_2 and ℓ_3 map into S_1, then this may be extended to a homomorphism $G_3 \to T$.

Let the length of the path $0y_{j+1}y_{j+2}\dots y_k$ (Figure 2.2) be t'. The other half of the clause gadget, G_4, is formed by taking the converse of G_3, changing the lengths of the paths in G_2 to t', and re-labeling the vertices ℓ_1, ℓ_2 and ℓ_3 as ℓ'_1, ℓ'_2 and ℓ'_3 respectively. Define S_3 to be the vertices in T_3 that are reachable by a path of length two from the 3-cycle in T. It now follows in a similar manner as before that ℓ'_1, ℓ'_2 and ℓ'_3 have to map into the set $S_3 \cup \{y_k\}$ under any homomorphism $G_4 \to T$. As before at least one of ℓ'_1, ℓ'_2 and ℓ'_3 has to map to S_3.

The clause gadget, K, can now (finally) be constructed. K is the disjoint union of G_3 and G_4. The idea behind the clause gadget is that "True" and "False" are to be encoded based on how each pair (ℓ_j, ℓ'_j), $j = 1, 2, 3$, is mapped.

We are now ready to give the reduction. Given an instance of not all equal 3-SAT (without negations) with variables $x_1, x_2, \dots, x_s$ and clauses $C_1, C_2, \dots, C_t$ we construct a digraph D as follows. For each variable x_i, we have a directed path of length two, P^i, with initial vertex (also) labeled x_i and terminal vertex labeled x'_i. For every clause $\{\ell_1, \ell_2, \ell_3\}$ take a copy of K. Each time a variable x_i appears in a clause $\{\ell_1, \ell_2, \ell_3\}$ (say $x_i = \ell_2$), identify the vertex x_i in P^i with the vertex ℓ_2 and the vertex x'_i in P^i with ℓ'_2 in the corresponding clause gadget. This is illustrated in Figure 2.8.

Let $D \to T$ be a homomorphism. Based on the discussion before the possible images for the pair (ℓ_j, ℓ'_j), $j = 1, 2, 3$, is the set $[S_1 \cup \{y_{i+1}\}] \times [S_3 \cup \{y_k\}]$. Since no vertex in S_1 is joined to a vertex in S_3 by a path of length two, we are able to refine the images somewhat. We define two subsets of the set $[S_1 \cup \{y_{i+1}\}] \times [S_3 \cup \{y_k\}]$: $\mathcal{T} = \{(y_{i+1}, y_k) \cup [\{y_{i+1}\} \times S_3]$ and $\mathcal{F} = S_1 \times \{y_k\}$. It can now be seen that the images of each pair (ℓ_j, ℓ'_j), $j = 1, 2, 3$, is the set $\mathcal{T} \cup \mathcal{F}$. A pair is not allowed to map only to $\mathcal{T}$ (as this forces each ℓ_j to map to y_{i+1} which is not possible). Also, no pair maps to $\mathcal{F}$ (as this forces all ℓ'_js to map to y_k which is not possible). Thus each pair (representing a variable) maps in such a way that there is at least one image in $\mathcal{T}$

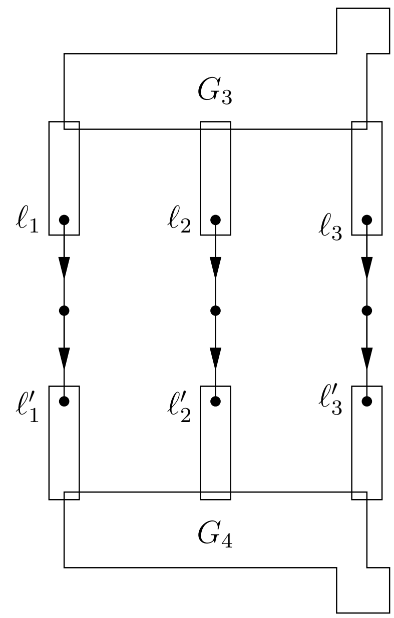

Figure 2.8: A typical clause in the instance D for the last unicyclic case.

and at least one image in $\mathcal{F}$. A satisfying truth assignment can now be recovered by examining the images of the end-vertices of the path P^i, representing the variable x_i: into $\mathcal{T}$ corresponds to “True” and into $\mathcal{F}$ corresponds to “False.”

Given a satisfying truth assignment to not-all-equal 3-SAT (without negations) we construct a homomorphism as follows. Let $y \in S_1$. If the variable x_i is assigned the value “True,” map the end-vertices of the path P^i to the pair (y_{i+1}, y_k). If the variable x_i is assigned the value “False,” map the ends of P^i to the pair (y, y_k). This pre-colouring can be extended to a homomorphism on each clause gadget.

Therefore there is a homomorphism $D \to T$ if and only if there is a satisfying truth assignment to not-all-equal 3-SAT without negations. ∎

For unicyclic local tournaments, we have the following dichotomy.

Theorem 2.6.4. *Let T be a unicyclic local tournament. If T is a directed cycle or T has the structure shown in Figure 2.2 with T_2 a tournament and at least one of T_1 and T_3 is empty, then T-colouring is polynomial. Otherwise T-colouring is NP-complete.*

We close this section by pointing out that if the 3-cycle in all of the above arguments is replaced by a directed 2-cycle, then all of the results carry over. In the polynomial cases, we are still dealing with a graft extension. For the NP-complete cases, we relied on oriented cycles of net-length one not mapping to the 3-cycle. These same oriented cycles do not map to a directed 2-cycle either.

2.7 Round Local Tournaments

Let D be a round local tournament containing at least two cycles and let ℓ denote the length of a shortest cycle in D. Proving NP-completeness is done in the following way:

❒ If D has a unique cycle of length ℓ, D-colouring is NP-complete.

❒ A smallest counterexample, D, (both in the number of vertices and the number of arcs)

- has all vertices on a shortest cycle,
- does not have $\delta^+ \geq \ell$,
- does not have $\Delta^+ \geq \ell$,
- does not have $\phi \geq 2\ell$, where ϕ is the Frobenius-Schur index of the cycle lengths of D.
- This forces $\phi = \ell$ and $\Delta^+ < \ell$ which leads to a contradiction. Therefore the counterexample does not exist and D-colouring is NP-complete.

Lemma 2.7.1. *Let D be a round local tournament containing at least two cycles. If D has a unique cycle of shortest length, then D-colouring is NP-complete.*

Proof. Let $v_0, v_1, \ldots, v_{n-1}$ be a round labeling of $V(D)$, C be a shortest cycle in D and ℓ be the length of C. Since D is strongly connected, by Lemma 2.3.5 there exist indices $0 < a_1 < a_2 < \ldots a_{k-1} < n$ so that $C = v_0 v_{a_1} v_{a_2} \cdots v_{a_{k-1}} v_0$.

Since D has at least two cycles, there must be vertices of D not on C.

The proof is divided into two cases. Firstly, we consider the situation where there are three indices $i, i+1$ and $i+2 \pmod{\ell}$ such that $a_{i+2} - a_{i+1} \geq 2$ and $a_{i+1} - a_i \geq 2$. Therefore there are vertices of D, not on C, between a_i and a_{i+1} and between a_{i+1} and a_{i+2}. Once this has been dealt with, what then remains is the case where these vertices do not exist. This will imply the existence of four indices $i, i+1, i+2$ and $i+3 \pmod{\ell}$ such that $a_{i+3} - a_{i+2} = 1$, $a_{i+2} - a_{i+1} \geq 2$ and that $a_{i+1} - a_i = 1$. That is, $v_{a_{i+3}}$ and $v_{a_{i+2}}$ are consecutive, as are $v_{a_{i+1}}$ and v_{a_i}, and there exists a vertex of D not on C between $v_{a_{i+1}}$ and $v_{a_{i+2}}$.

Case 1: $a_{i+2} - a_{i+1} \geq 2$ and $a_{i+1} - a_i \geq 2$.

Let u be any vertex between a_i and a_{i+1} and let w be the successor of a_{i+1} on the outer n-cycle (w is between a_{i+2} and a_{i+1}). This is illustrated in Figure 2.9.

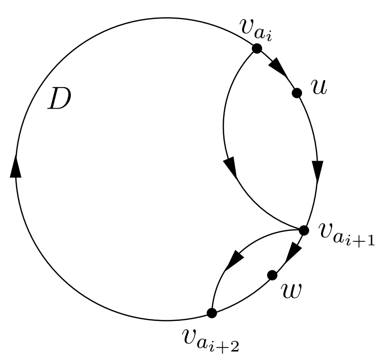

Figure 2.9: The first case where D has a unique cycle of shortest length

We handle this case by assuming the lemma is false and examining a minimum counterexample. That is assume that there exists a round local tournament D with at least two cycles and a unique cycle of shortest length that is of the form shown in Figure 2.9. In addition to this we assume that D has the minimum number of vertices possible. Also, D-colouring is not NP-complete.

We apply the sub-indicator J shown in Figure 2.10 to D. Note that the vertex on the left is identified with vertex a_{i+1} in D and that the vertex j of J is the vertex on the right. The sub-indicator is constructed by starting with a directed path of length $\ell(\ell-1)$. Skip the first two vertices and then attach ℓ directed ℓ-cycles to the next ℓ vertices, skip the next vertex, and then attach ℓ directed ℓ-cycles to the next ℓ vertices ...

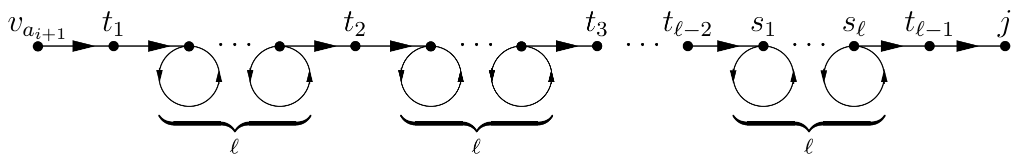

Figure 2.10: The sub-indicator for the first case.

We now consider all retractions of J to D, where the first vertex of J maps to $v_{a_{i+1}}$. First of all, the presence of the ℓ-cycles force the vertices of the path that they are attached to, to map to C. If we map all vertices of J from the first up to s_ℓ onto C, then $s_\ell \mapsto v_{a_{i-1}}$. Therefore $j \mapsto u$ and $j \mapsto v_{a_{i+1}}$ (at least).

We claim that $j \not\mapsto w$. To see why this is the case, consider what happens when $j \mapsto w$. If this was to happen, then $s_\ell \mapsto v_{a_i}$ (if not, we either get a cycle of length less than ℓ or another ℓ cycle different from C). If $s_\ell \mapsto v_{a_i}$, then $s_1 \mapsto v_{a_{i+1}}$. Therefore the image of J between the first vertex and vertex s_1 is a closed directed walk, W, of length $\ell(\ell-2)-1$. Denote by J' the subgraph of J induced by $V(J) - \{s_2, s_3, \ldots, s_\ell, t_{\ell-1}, j\}$.

The walk W can be decomposed into directed cycles. The length of each of these directed cycles is either ℓ or $\ell+1$. To see why this is the case consider any group of $\ell+3$ consecutive vertices on J'. Any such group will include at least $\ell+1$ of the vertices with ℓ cycles attached to them (also keep in mind that the first vertex is mapped onto $v_{a_{i+1}}$). These vertices can only map to $V(C)$ and so the images (of the vertices in J') can not be distinct. This implies that the longest cycle that vertices in J' can map to is at most $\ell+2$. In order for $\ell+2$ vertices in J' to map to an $\ell+2$ cycle in D, one needs $\ell+2$ vertices in J' that can have distinct images. Such a group of $\ell+2$ vertices only occurs between consecutive t_is (see Figure 2.10). In order for the group to map to an $(\ell+2)$-cycle t_i and the vertex following t_{i+1} must have the same image. That means that t_i is mapped to a vertex on $V(C)$ and so the $\ell+2$ vertices do not have distinct images. This shows that the directed cycles in the decomposition of W can only be of length ℓ and $\ell+1$.

Since W can be decomposed into ℓ and $\ell+1$ cycles, we have that $\ell(\ell-2)-1 = k_1\ell + k_2(\ell+1) = (k_1+k_2)\ell + k_2$ where k_1 is the number ℓ-cycles and k_2 the number of $(\ell+1)$-cycles in such a decomposition. Note that $0 \leq k_1, k_2 \leq \ell-2$. This means that $\ell(\ell-1) = (k_1+k_2)\ell+(k_2+1)$ or that ℓ divides (k_2+1), but $1 \leq (k_2+1) \leq \ell-1$, so we have a contradiction. Therefore if the first vertex of J maps to $v_{a_{i+1}}$, then j cannot map to w.

We now also claim that j maps to every vertex of C. From before we already know that $j \mapsto v_{a_{i+1}}$. If we map t_1 to w and then all other vertices of J to C, we find that $j \mapsto v_{a_i}$: if $t_1 \mapsto w$, the vertex preceding t_2 maps to $v_{a_{i+1}}$ and the length of the path that remains in J is $\ell(\ell-2)-1$ — a multiple of ℓ minus 1. In general by mapping $t_1, t_2, \ldots, t_k$ to w and all other vertices (from the first up to the vertex preceding t_{k+1}) to C, we find that the vertex preceding t_k maps to $v_{a_{i+1}}$. The length of the path remaining in J is $\ell(\ell-1)-(k\ell+k) = \ell(\ell-1-k)-k$ (we have used $k\ell+k+1$ vertices or $k\ell+k$ arcs up to this point). If we now map all the remaining vertices to C we see that $j \mapsto v_{a_{i+1-k}}$.

The result of this sub-indicator, D^+, contains (at least) the subgraph induced by $C \cup \{u\}$, but not the vertex w. Therefore D^+ contains at least two cycles and also has a unique shortest cycle of length ℓ. Since D has the minimum number of vertices for a counterexample, D^+-colouring is NP-complete, implying D-colouring is NP-complete — a contradiction.

Case 2: $a_{i+3} - a_{i+2} = 1$, $a_{i+2} - a_{i+1} \geq 2$ and that $a_{i+1} - a_i = 1$.

This case is illustrated in Figure 2.11.

The proof in this case is via a reduction from not all equal ℓ-SAT without nega-

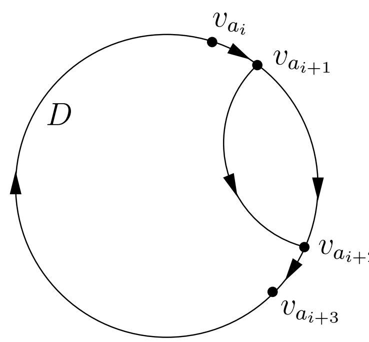

Figure 2.11: The second case where D has a unique cycle of shortest length

tions.

In carrying out the proof we need to construct the gadget Z shown in Figure 2.12. The gadget Z is constructed from a copy of D and two directed paths of length two, say $a'ua$ and $z'vz$. Identify a' and z', attach directed ℓ-cycles to vertices z', v and z and finally identify a with $v_{a_{i+1}}$.

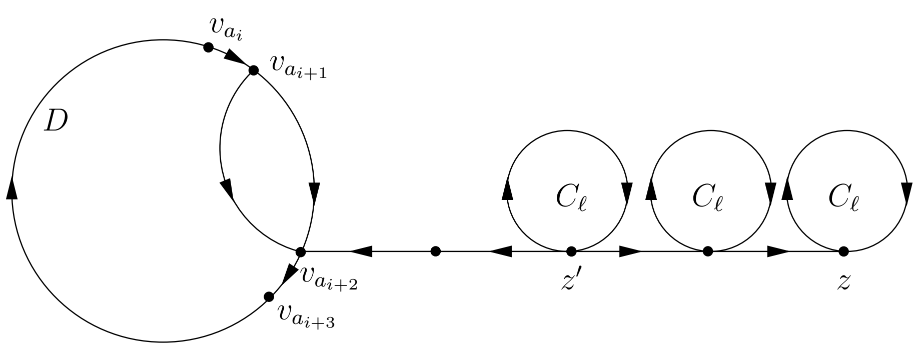

Figure 2.12: The gadget Z

In a retraction of Z to D, the vertex z' can only map to $v_{a_{i+1}}$ or v_{a_i}. This forces the pair of vertices (z', z) to map to the pair $(v_{a_i}, v_{a_{i+2}})$ or the pair $(v_{a_{i+1}}, v_{a_{i+3}})$.

We also need the gadget K shown in Figure 2.13 (it is the same as the sub-indicator used in the first case).

Let $f : K \to D$ be a homorphism from K to D. If $f(r) = v_{a_i}$, then $f(s) \neq v_{a_{i+1}}$. If $f(s) = v_{a_{i+1}}$, then $f(t_{\ell-1}) = v_{a_i}$. This means that the image of K under f corresponds to a closed walk from v_{a_i} to v_{a_i} of length $\ell(\ell-1)-1$. Through a similar technique

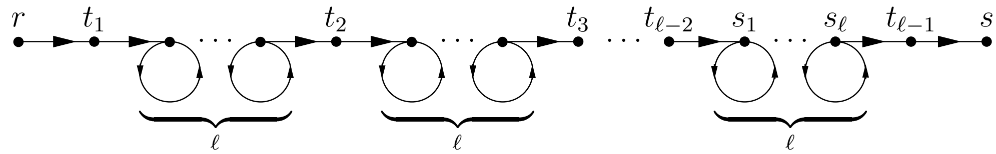

Figure 2.13: The gadget K

as in the previous case, one can show that this is impossible: the walk of length $\ell(\ell-1)-1$ can only be decomposed into ℓ and $\ell+1$ cycles (because of the placement of the ℓ-cycles) and this is impossible. As in case one, it is possible to show that if $f(r) = v_{a_i}$, then s maps to every vertex of C except $v_{a_{i+1}}$. This accomplished by mapping the t_is of K "into" the gap between $v_{a_{i+1}}$ and $v_{a_{i+2}}$. On the other hand if $f(r) = v_{a_{i+1}}$, then s can map to any vertex of C (again the t_is are mapped into the gap between $v_{a_{i+1}}$ and $v_{a_{i+2}}$). In a similar way it is possible to show that if $f(s) = v_{a_{i+3}}$, then $f(r)$ can map to any vertex of C except $v_{a_{i+2}}$ and that if $f(s) = v_{a_{i+2}}$, then r can map to any vertex of C.

Next, we construct a new gadget F from two copies of K, say with end-vertices r_1, s_1 and r_2, s_2 respectively. To construct F identify vertex s_1 with vertex r_2 and call this vertex c. This is shown in Figure 2.14.

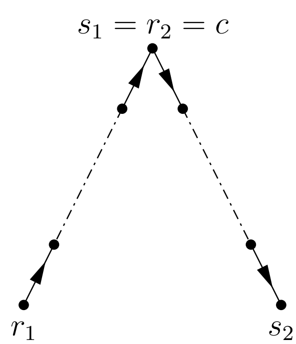

Figure 2.14: The gadget F

The gadget F has the property that if the pair (r_1, s_2) maps to the pair $(v_{a_i}, v_{a_{i+2}})$, under a homomorphism from F to D, then the vertex c can map to any vertex on C, except $v_{a_{i+1}}$. If the pair (r_1, s_2) maps to $(v_{a_{i+1}}, v_{a_{i+3}})$, then c can map to any vertex on C except $v_{a_{i+2}}$.

As mentioned earlier, the proof is via a reduction from not all equal ℓ-SAT without negated variables. So let an instance of not all equal ℓ-SAT without negated variables be given:

$$Z = \{z_1, z_2, \ldots, z_a\} \text{ — the variables },$$

$$K_1, K_2, \ldots, K_b \text{ — the clauses },$$

$$K_j = \{z_{j_1}, z_{j_2}, \ldots, z_{j_\ell}\} \text{ for each } j \in \{1, 2, \ldots, b\}.$$

We now construct an instance of D-colouring. For each variable z_j take a copy of the gadget Z. Identify all copies of D. Label the z and z'-vertices in the copies of Z as $z_1, z_2, \ldots, z_a, z'_1, z'_2, \ldots, z'_a$. These vertices correspond to variables in the instance of not all equal ℓ-SAT without negated variables. This produces the variable gadget.

For each clause $K_j = \{z_{j_0}, z_{j_1}, \ldots, z_{j_{\ell-1}}\}$ take ℓ copies of the gadget F, say $F_{j_1}, F_{j_2}, \ldots, F_{j_\ell}$ with end-vertices $(r_{j_{11}}, s_{j_{21}}), (r_{j_{12}}, s_{j_{22}}), \ldots, (r_{j_{1\ell}}, s_{j_{2\ell}})$ and top-vertices $c_{j_1}, c_{j_2}, \ldots, c_{j_\ell}$, respectively. Now identify $r_{j_{1i}}$ with z'_{j_i} and $s_{j_{2i}}$ with z_{j_i} and form an ℓ-cycle through $c_{j_1}, c_{j_2}, \ldots, c_{j_\ell}$. This produces the instance H of D-colouring shown in Figure 2.15.

We now show that $H \to D$ if and only if there exists a satisfying truth assignment for the instance of not all equal ℓ-SAT without negated variables. Define the pair $(v_{a_i}, v_{a_{i+2}})$ (in D) to be T and the pair $(v_{a_{i+1}}, v_{a_{i+3}})$ to be F.

If $H \to D$, then each pair (z'_i, z_i) maps to T or to F. Furthermore for a given clause $K_j = \{z_{j_1}, z_{j_2}, \ldots, z_{j_\ell}\}$ it is not the case that all pairs (z'_{j_i}, z_{j_i}), $1 \le i \le \ell$, map to T. If this was the case, none of the c_{j_i}s are allowed to map to $v_{a_{i+1}}$ and since there is a unique ℓ-cycle in D, we won't be able to complete the ℓ-cycle through the c_{j_i}s. In a similar way it is possible to show that the pairs (z'_{j_i}, z_{j_i}), $1 \le i \le \ell$ can not all map to F (here we would be missing the vertex $v_{a_{i+2}}$ on the ℓ-cycle). Therefore

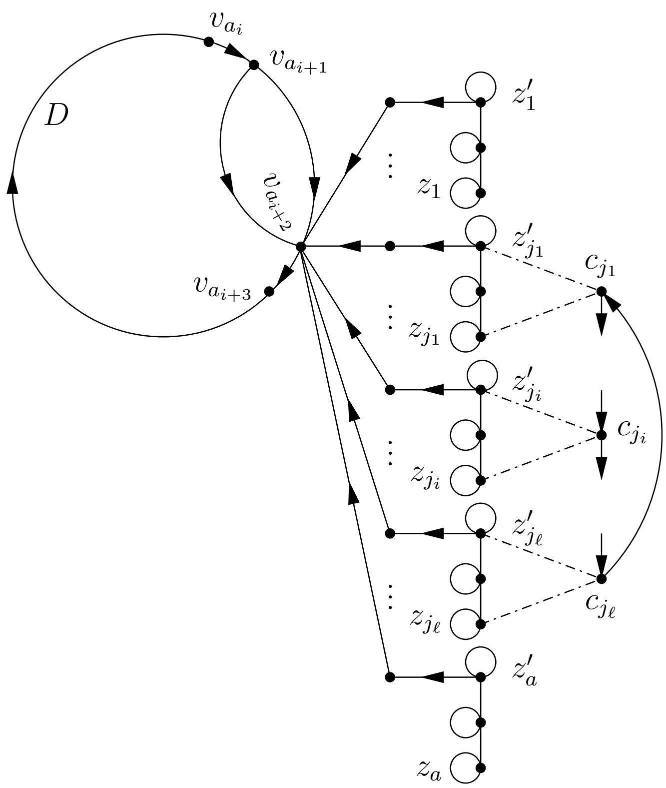

Figure 2.15: The instance H

in every clause there is at least one pair that maps to T and at least one pair that maps to F. A satisfying truth assignment can now be recovered by assigning "True" ("False") to variable z_i if (z'_i, z_i) maps to T (F).

Conversely, let a satisfying truth assignment be given. Pre-colour the pair (z'_i, z_i) by T (F) if z_i is assigned the value "True" ("False"). For the clause gadget corresponding to clause $K_j = \{z_{j_1}, z_{j_2}, \ldots, z_{j_\ell}\}$ we extend the colouring as follows: locate two consecutive vertices c_{j_k} and $c_{j_{k+1}}$ such that (z'_{j_k}, z_{j_k}) maps to F and $(z'_{j_{k+1}}, z_{j_{k+1}})$ maps to T (all subscripts are taken modulo ℓ). Map the corresponding F_{j_k} in such a way that c_{j_k} maps to $v_{a_{i+1}}$ and map $F_{j_{k+1}}$ in such a way that $c_{j_{k+1}}$ maps to $v_{a_{i+2}}$. Now map $F_{j_{k+2}}$ such that $c_{j_{k+2}}$ maps to $v_{a_{i+3}}$, $F_{j_{k+3}}$ such that $c_{j_{k+3}}$ maps to $v_{a_{i+4}}$, $\ldots$, $F_{j_{k-1}}$ such that $c_{j_{k-1}}$ maps to v_{a_i}. By the properties of F and K discussed earlier, this is always possible. The remaining vertices of H map in an obvious way to D.

Therefore $H \to D$. ■

To complete the proof that colouring by a round local tournament is NP-complete, we assume that the result is false and examine a smallest counterexample (minimum number of vertices and minimum number of arcs). Being a counterexample, the corresponding colouring problem is not NP-complete. We derive additional properties of the counterexample to see that it cannot exist and in doing so we can then conclude that the colouring problem is in fact NP-complete.

The first property of this counterexample follows from Lemma 2.7.1.

Lemma 2.7.2. *Every vertex of a smallest counterexample is on a shortest cycle.*

Proof. If D is a smallest counterexample then D-colouring is not NP-complete. This being the case, by Lemma 2.7.1 it must have at least two cycles of shortest length. Let ℓ denote the length of a shortest cycle in D. If, in addition, D has vertices that are not on these shortest cycles, we apply a sub-indicator C_ℓ to D. The result of this, say D', will be the digraph induced by the vertices of D that are on a shortest cycle. Since there are at least two shortest cycles, D' contains at least two cycles and furthermore it has fewer vertices that D. In addition to this it is also a round local tournament because it is an induced subgraph of D. Since it has fewer vertices than D (and D is a smallest counterexample), D'-colouring has to be NP-complete (otherwise it would be smaller than a smallest counterexample). This implies that D-colouring is also NP-complete, a contradiction. ■

The next property deals with the minimum out-degree of the smallest counterexample.

Lemma 2.7.3. *Let D be a smallest counterexample and denote by δ^+ the minimum out-degree of D. Then $\delta^+ < \ell$, where ℓ is the length of a shortest cycle in D.*

Proof. Assume that D has $\delta^+ \geq \ell \geq 3$. To prove the lemma we apply the arc-subindicator shown in Figure 2.16 (a transitive tournament on $\delta^+ + 1$ vertices with the two arcs spanning end-to-end as shown), with respect to the dashed arc, to D. Since every vertex in D has degree at least δ^+, the result of this arc-subindicator is a

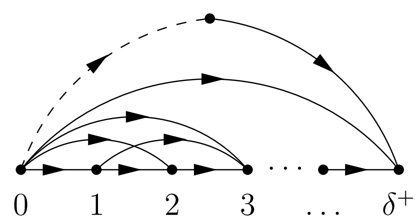

Figure 2.16: The arc-subindicator.

digraph D' in which every vertex has out-degree at least 2 (the two vertices following the given vertex in the round enumeration). Every vertex also loses its outneighbour furthest away from it in the round enumeration. This implies that D' has at least two cycles and is a round local tournament. Furthermore D' has at least one vertex with out-degree less than ℓ (these are the vertices in D that have out-degree exactly δ^+) and so has fewer arcs than D. Since D has the minimum number of arcs possible for a counterexample, D' cannot be a counterexample and so D'-colouring is NP-complete. By the arc-subindicator construction this implies that D-colouring is also NP-complete, a contradiction. ∎

Next, we prove that a similar result holds for Δ^+.

Lemma 2.7.4. *Let D be a smallest counterexample and denote by Δ^+ the maximum out-degree of D. Then $\Delta^+ < \ell$, where ℓ is the length of a shortest cycle in D.*

Proof. Assume that D has $\Delta^+ \geq \ell$. By the previous result this implies that there exist two vertices x and y in D such that $d^+(x) < \ell$ and $d^+(y) \geq \ell$.

Let $V(D) = v_0, v_1, \ldots, v_{n-1}$ be a round enumeration of the vertices of D. We claim that there exists an $i \in \{0, 1, \ldots, n-1\}$ such that $d^+(v_{i-1}) < l$ and $d^+(v_i) \geq \ell$.

To see this, take any vertex u with $d^+(u) \geq \ell$; if the predecessor of u (in the round enumeration) has out-degree less than ℓ we are done. Otherwise, the predecessor has out-degree greater than or equal to ℓ. We then repeat the process of considering the predecessor. Since there is a vertex of out-degree less than ℓ, this process terminates. By shifting the labels of D (mod n) we can assume that $d^+(v_{n-1}) < \ell$ and that $d^+(v_0) \geq \ell$.

Let's say that $d^+(v_0) = m \geq \ell$ and $d^+(v_{n-1}) = a < \ell$. We then use the sub-indicator formed by taking a copy of D and attaching a path of length ℓ to vertex v_0 and letting j be the final vertex on the path of length ℓ. This is shown below in Figure 2.17.

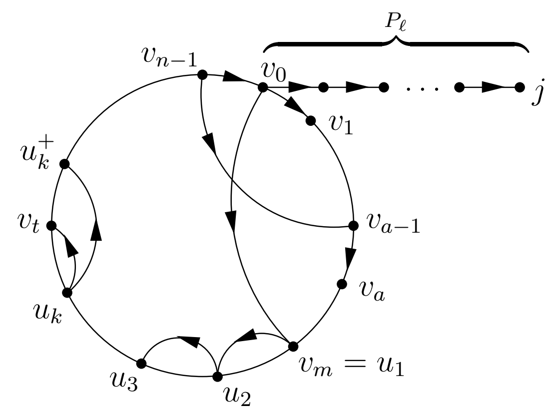

Figure 2.17: The sub-indicator.

Every vertex is on a shortest cycle (an ℓ-cycle). In particular, v_0 is on one and moreover there is a shortest cycle containing v_0 that uses the arc v_0v_m. To see this consider any shortest cycle through v_0, if it is not using v_0v_m then there is at most one vertex of this cycle between v_0 and v_m (otherwise we can find a shorter cycle). Call this vertex u. The vertex u is adjacent to some v_i with $i > m$, otherwise there is a shorter cycle. This also means that v_m is adjacent to v_i. By replacing the arc v_0u with v_0v_m and then following the rest of the ℓ-cycle, we now have a shortest cycle that uses v_0v_m. Let C be such a shortest cycle through v_0. Label the vertices on C (starting with v_0) as $v_0 = u_0, v_m = u_1, \ldots, u_{\ell-1}$.

It is easy to see that vertex j maps to $v_\ell, v_{\ell+1}, \dots, v_m$ by using the outer n-cycle together with the arcs $v_{\ell-1}v_i$ where $i \in \{\ell, \ell+1, \dots, m\}$. Furthermore, j maps to any v_t with $m < t \leq n-1$: for a given v_t, let $s = \max\{x \mid x < t \text{ and } v_x \in C\}$. Since $v_s \in C$, $v_s = u_k$ for some $1 \leq k \leq \ell - 1$. The idea is that u_k is the first vertex on C, not including v_t, that is encountered when moving backwards along C. The successor u_k on C, $u_k^+ = v_i$, with $i \geq t$. This implies that $u_k v_t$ is an arc of D. In order to obtain a walk of length ℓ from $v_0 = u_0$ to v_t we first proceed along the n-cycle from v_0 to $v_{\ell-(k+1)}$ $(0 \leq \ell-(k+1) < m)$. Next, we use the arc $v_{\ell-(k+1)}v_m = v_{\ell-(k+1)}u_1$, then proceed along C to u_k and finally use the arc $u_k v_t$.

By using C, it is clear that j maps to v_0. The resulting digraph derived from using this sub-indicator, contains C. It therefore has at least one cycle.

We now show that vertex j does not map to v_a. Here we have that $N^-(v_a) \subseteq \{v_0, v_1, \dots, v_{a-1}\}$. In order to obtain a contradiction assume that $W = x_0 x_1 \cdots x_{\ell-1}$ is walk of length ℓ from v_0 to v_a. That is $x_0 = v_0, x_1 = v_{j_1}, x_2 = v_{j_2}, \dots, x_{\ell-2} = v_{j_{\ell-2}}, x_{\ell-1} = v_a$. Let $k = \max\{0, j_1, j_2, \dots, j_{\ell-2}, a\}$. Note that $k > a$, otherwise $j_i \leq a$ and W is a walk of length ℓ inside the transitive tournament $D[v_0, v_1, \dots, v_a]$. Furthermore $v_k v_a$ is not an arc of D, for if it was then $d^+(v_{n-1}) \geq \ell$. Therefore $d(v_k, v_a) \geq 2$ and so the length of the sub-walk of W from v_0 to v_k is at most $\ell - 2$. In order to complete W, $v_k v_i \in A(D)$ for some $i \in \{0, 1, \dots, a-1\}$. This implies that $v_k v_0 \in A(D)$ which means that there is a closed walk from v_0 to itself of length at most $\ell - 1$. This is clearly not possible and so j cannot map to v_a.

Now, if $d^+(v_0) > l$, then the result, say D', of this sub-indicator has at least two cycles, but does not contain v_a. Since it is smaller than D it is not a counterexample and so D'-colouring is NP-complete. Therefore D-colouring is NP-complete, a contradiction. Thus $d^+(v_0) = l$.

On the other hand, if there is a vertex $v_i \in \{v_m = v_\ell, v_{\ell+1}, \dots, v_{n-2}\}$ with $d^+(v_i) \geq 2$, then the result of the sub-indicator will again have at least two cycles and as above this leads to D-colouring being NP-complete, a contradiction. This shows that $D[v_m = v_\ell, v_{\ell+1}, \ldots, v_{n-1}]$ is an induced path. Note that, in this case, v_{n-1} is the predecessor of v_0 on the ℓ-cycle C above.

If we now have that $d^+(v_{n-1}) \geq 2$, j will map to v_1 by using C. Therefore, the sub-indicator results in a digraph with at least two cycles and fewer vertices than D (j still does not map to v_a). Once again, we conclude that D-colouring is NP-complete, a contradiction. Hence, $d^+(v_{n-1}) = 1$.

We now apply the sub-indicator shown below in Figure 2.18. This is formed using a copy of D and attaching an oriented path as shown to v_0.

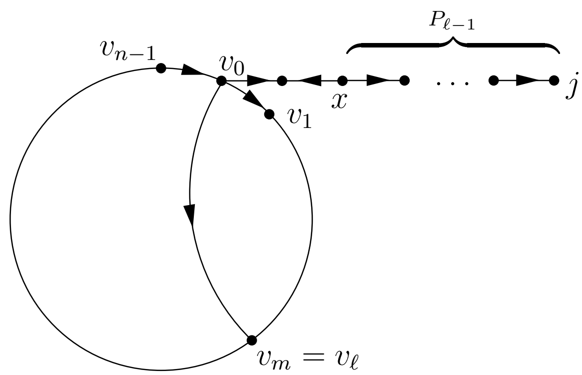

Figure 2.18: The new sub-indicator.

No vertex between v_0 and v_ℓ is adjacent to a vertex v_i with $i > \ell + 1$, otherwise v_ℓ has out-degree at least two. Therefore any vertex between v_0 and v_ℓ is adjacent to $v_{\ell+1}$ so that it can be on a cycle of length ℓ. Also note that the vertex x shown above can map to any of $\{v_0, v_1, \ldots, v_{\ell-1}\}$.

The distance between v_ℓ and v_0 is exactly $\ell - 1$, because of the ℓ-cycle C through v_0. This allows j to map to v_0: first map x to $v_i \in \{v_1, \ldots, v_{\ell-1}\}$, then use $v_i v_{\ell+1}$ together with the path of length $\ell - 2$ from $v_{\ell+1}$ to v_0. It is also possible to map j to $v_i \in \{v_{\ell-1}, v_\ell, \ldots, v_{n-1}\}$: first map x to $v_{i-(\ell-1)}$ ($0 \leq i - (\ell - 1) \leq n - 1 - (\ell - 1) = n - \ell = \ell - 1$), then follow the outer n-cycle from $v_{i-(\ell-1)}$ to v_i.

This sub-indicator therefore guarantees that we have at least the ℓ-cycle C through

v_0 and the vertex $v_{\ell-1}$. Thus the result has at least two cycles.

We now note that j does not map to v_1, otherwise there exists a directed walk of length $\ell - 1$ from $v_i \in \{v_0, v_1, \ldots, v_{\ell-1}\}$ (since x can only map to v_i) to v_1 and such a walk does not exist. Therefore the result has at least one fewer vertex and so its colouring problem is NP-complete, forcing D-colouring to be NP-complete, a contradiction. ∎

The next two results deal with the Frobenius-Schur index of the cycle lengths of a minimum counterexample. The first one may seem somewhat artificial, but it is needed to establish the second.

Lemma 2.7.5. *Let D be a round local tournament on the vertex set $\{0, 1, \ldots, n-1\}$ with $n = 2\ell - 2$, where ℓ is the length of a shortest cycle in D. If for every $v \in \{\ell - 2, \ell - 1, \ldots, n - 1\}$, $N^+(v) = v + 1$ (mod n), then D cannot be a minimum counterexample to the D-colouring problem.*

Proof. Let $\phi(\ell, \ell+1, \ldots, n) = \phi$ be the Frobenius-Schur index of the cycle lengths of D. Recall that ϕ is the smallest integer such that every integer $x \geq \phi$ can be written as a linear combination of $\ell, \ell + 1, \ldots, n$. By the minimality of ϕ, $\phi - 1$ cannot be written as such a linear combination. Therefore D does not possess a closed walk of length $\phi - 1$: If it did, this closed walk can be decomposed into arc disjoint cycles with lengths in the set $\{\ell, \ell+1, \ldots, n\}$. This is equivalent to writing $\phi-1$ as a linear combination of $\ell, \ell + 1, \ldots, n$ which is not possible.

We assume that D is a minimum counterexample and derive a contradiction. Since $n = 2\ell - 2$,

$$\phi = \left\lfloor \frac{n-2}{n-\ell} \right\rfloor \ell = 2\ell.$$

Therefore $\phi-1 = 2\ell-1 = n+1$. Also, since 0 has to be on an ℓ-cycle, $0(\ell-1) \in A(D)$ and by the local tournament property $0a \in A(D)$ where $a \in \{1, 2, \ldots, \ell - 1\}$. To derive the contradiction, we use a sub-indicator construction. This sub-indicator is constructed from a copy of D and by attaching a path of length $\phi - 1$ to vertex 0 and taking the end of this path to be the vertex j. Some of the possible images of j are: $1, 2, \ldots \ell - 1$. This is accomplished by going around the n-cycle once back to 0 and then using the arc $0a$, $a \in \{1, 2, \ldots, \ell - 1\}$. To map j to vertex x, where $\ell \leq x \leq n-1$, we proceed as follows: use an $(n-(x-\ell+1))$-cycle to go from 0 to 0 $(\ell \leq n-(x-\ell+1) \leq n-1)$. Then use the remaining $\phi-1-(n-(x-\ell+1)) = x-\ell+2$ arcs $(2 \leq x-\ell+2 \leq n-\ell+1)$ first to go from 0 to $\ell-1$ and then from $\ell-1$ to x along the n-cycle. Its impossible for j to map to 0 since this would correspond to a closed walk from 0 to 0 of length $\phi - 1$. Therefore the result of applying this sub-indicator is $D - \{0\}$. Furthermore, since $\ell - 2$ is not adjacent to ℓ, 1 is not adjacent to ℓ. In order for 1 to be on an ℓ-cycle, $(n-1)$ is adjacent to 1. Therefore $D - \{0\}$ has at least two cycles. Since D is a minimum counterexample, $(D - \{0\})$-colouring is NP-complete, and so D-colouring is NP-complete, a contradiction. ∎

Lemma 2.7.6. *If D is a minimum counterexample, then the Frobenius-Schur index, ϕ, of the cycle lengths of D satisfies $\phi = \ell$, where ℓ is the length of a shortest cycle in D.*

Proof. Label the vertices of D as $\{0, 1, \ldots n-1\}$. By Lemma 2.2.5,

$$\phi = kl \text{ where } k = \left\lfloor \frac{n-2}{n-\ell} \right\rfloor .$$

We assume now that $\phi \geq 2\ell$, and derive a contradiction.

If $\phi \geq 2\ell$, then

$$k = \left\lfloor \frac{n-2}{n-\ell} \right\rfloor \geq 2,$$

and this happens if and only if $(n-2)/(n-\ell) \geq 2$ which is equivalent to $2\ell \geq n+2$. To obtain the sought after contradiction we employ as an indicator a directed path of length $\phi - 1$. This indicator has the property that its result will not contain loops by the fact that D has no closed walks of length $\phi - 1$.

Note that $\phi-1=(k-1)\ell+(\ell-1)$. This implies that in applying the indicator to D, we obtain an arc from 0 to $\ell-1$ (use $(k-1)$ ℓ-cycles to get from 0 to 0, and then use the remaining $(\ell-1)$ arcs to get from 0 to $\ell-1$). Also, $\phi-1=(k-2)\ell+(\ell+1)+(\ell-2)$. This produces an arc from 0 to $\ell-2$ when applying the indicator.

We now show that it is also possible to have arcs from $\ell-1$ to 0 and from $\ell-2$ to 0. The number of arcs on the outer n-cycle from $\ell-1$ to 0 is $n-\ell+1$. Consider now $\phi-1-(n-\ell+1)=\phi-(n-\ell+2)$. If it is possible to write $\phi-(n-\ell+2)$ as a linear combination of cycle lengths, then there will a closed walk from any vertex back to itself of this length. In particular, there will be a walk of this length from $\ell-1$ to $\ell-1$. A further $n-\ell+1$ arcs will then carry the walk from $\ell-1$ to 0 along the outer n-cycle. Therefore, applying the indicator above will produce an arc from $\ell-1$ to 0 provided that $\phi-(n-\ell+2)$ is a sum of cycle lengths. This is indeed the case: We know that all $x\geq\phi=k\ell$ is a sum of cycle lengths. Furthermore $(k-1)\ell,(k-1)\ell+1,\ldots,(k-1)n$ are also linear combinations (of $k-1$) of the cycle lengths. Let $t=k\ell-(k-1)n$, this represents the length of a "gap" of integers that cannot be written as a linear combination of the cycle lengths. We see that $k(\ell-n)=t-n$ so that

$$k=\frac{t-n}{\ell-n}=\frac{n-t}{n-\ell}=\frac{n-2}{n-\ell}-\frac{t-2}{n-\ell}.$$

Since

$$k=\left\lfloor\frac{n-2}{n-\ell}\right\rfloor,$$

we get that $(t-2)/(n-\ell)<1$. Therefore $t<n-\ell+2$ or $t\leq n-\ell+1$. Further from $2\ell\geq n+2$ it follows that $n-\ell+2\leq\ell$. Thus $\phi-(n-\ell+2)\in\{(k-1)\ell=\phi-\ell,(k-1)\ell+1,\ldots,(k-1)n=\phi-t\}$ and so $\phi-(n-\ell+2)$ is a sum of cycle lengths. To obtain an arc from $\ell-2$ to 0, we are searching for a walk of length $\phi-1$ from $\ell-2$ to 0. The number of arcs on the outer n-cycle from $\ell-2$ to 0 is $n-\ell+2$.

Here $\phi - 1 - (n - \ell + 2) = \phi - (n - \ell + 3)$. Also $\phi - (n - \ell + 3) > \phi - (n - \ell + 2) > t$ and $\phi - (n - \ell + 3) \leq \ell + 1$. If $\phi - (n - \ell + 3) \leq \ell$, then $\phi - (n - \ell + 3) \in \{(k-1)\ell = \phi - \ell, (k-1)\ell + 1, \ldots, (k-1)n = \phi - t\}$. This enables a closed walk from $\ell - 2$ to $\ell - 2$ of length $\phi - (n - \ell + 3)$ to be formed. The remaining $n - \ell + 2$ arcs on the walk of length $\phi - 1$ then goes from $\ell - 2$ to 0. What remains at this point is the case $n - \ell + 3 = \ell + 1$, or $n - \ell + 2 = \ell$. This implies that the number of arcs on the n-cycle from $\ell - 2$ to 0 is ℓ. If there exists a vertex in $\{\ell - 2, \ell - 1, \ldots, n - 2\}$ with out-degree at least two, then there exists a path of length $\ell - 1$ from $\ell - 2$ to 0. This path together with a closed walk of length $\phi - (n - \ell + 2)$ provides us with a walk of length $\phi - 1$ from $\ell - 2$ to 0. If there does not exist a vertex in $\{\ell - 2, \ell - 1, \ldots, n - 2\}$ with out-degree at least two, then $D[\ell - 2, \ell - 1, \ldots, n - 1, 0]$ is an induced path. This is not possible by Lemma 2.7.5.

We now see that using a path of length $\phi - 1$ as an indicator, produces the arcs $0(\ell - 1)$, $(\ell - 1)0$, $(\ell - 2)0$ and $0(\ell - 2)$. The same, of course, applies to any other vertex in D: Starting at vertex a, we have symmetric arcs between a and $a + (\ell - 1)$ as well as between a and $a + (\ell - 2)$, where addition is done modulo n. Again, if it is not possible to obtain a path of length $\ell - 1$ from $a + (\ell - 2)$ to a, then we are in the case dealt with by Lemma 2.7.5. If we now apply as arc-subindicator the digraph with vertices $\{i, j\}$ and arcs ij and ji, with respect to the arc ij, the end result will be only the symmetric arcs or in other words the undirected part of the indicator construction. This undirected part contains the circulant on the vertices $\{0, 1, \ldots, n - 1\}$ and the edges $a(a + (\ell - 1))$ and $a(a + (\ell - 2))$. Such a circulant is not bipartite and so colouring by the undirected portion of the indicator construction (that was obtained through the sub-indicator construction) is NP-complete, and so colouring by the whole result of the indicator construction is NP-complete. This implies that D-colouring is NP-complete, which is a contradiction. ∎

The final contradiction is now obtained by noting that if D is a minimum counterexample to the D-colouring problem, then D has $\Delta^+ < \ell$ and $\phi = \ell$. Here ℓ is the length of a shortest cycle in D and ϕ is the Frobenius-Schur index of the cycle lengths of D. We show next that these two conditions on D are incompatible, and so D does not exist.

Lemma 2.7.7. *Let D be a round local tournament with the Frobenius-Schur index of its cycle lengths $\phi(\ell, \ell+1, \ldots, n) = \phi = \ell$ and $\Delta^+(D) < \ell$, where ℓ is the length of a shortest cycle in D. Then D cannot be a minimum counterexample to the D-colouring problem.*

Proof. We assume that D is a minimum counterexample and then derive a contradiction.

Label the vertices of D as $\{0, 1, \ldots, n-1\}$. Since

$$\phi = \left\lfloor \frac{n-2}{n-\ell} \right\rfloor = \ell,$$

$(n-2)/(n-\ell) < 2$, that is $2\ell \leq n+1$ or $n-\ell+1 \geq \ell$. Also $\Delta^+ < \ell$ and so every vertex has out-degree at most $\ell - 1$.

To obtain the contradiction we use an indicator that is equal to a path of length $\ell-1$ with the two end-vertices of the path as the distinguished vertices of the indicator. We see that the result of this indicator, D^*, has an arc from 0 to $\ell - 1$ (using the n-cycle from 0 to $\ell - 1$). Also, there is an arc from 0 to ℓ in D^*: If there is at least one vertex in $\{0, 1, \ldots, \ell-2\}$ with out-degree at least two, then there exists a path of length $\ell - 1$ from 0 to ℓ. We show that such a vertex exists. Assume that all vertices in $\{0, 1, \ldots, \ell-2\}$ have out-degree one. Then $D[0, 1, \ldots, \ell-1]$ is an induced path of length $\ell - 1$. In order for 0 to be on a cycle of length ℓ, $(\ell-1)0 \in A(D)$. This implies that $d^+(\ell-1) = n - (\ell-1) \geq \ell$, which contradicts $\Delta^+ < \ell$. Therefore, the required vertex exists and there is a path of length $\ell - 1$ from 0 to ℓ.

Next, we show that there are also arcs from $\ell - 1$ to 0 and from ℓ to 0.

The length of the path on the n-cycle from $\ell - 1$ to 0 is $n - (\ell - 1) = n - \ell + 1 \geq \ell$. Since $\Delta^+ < \ell$, $N^-(\ell - 1) \subseteq \{0, 1, \ldots, \ell - 2\}$. Using the ℓ-cycle through vertex $\ell - 1$, we get a path of length $\ell - 1$ from vertex $\ell - 1$ to its predecessor on this ℓ-cycle. Note that this predecessor is in the set $\{0, 1, \ldots, \ell - 2\}$. Let $P = \ell - 1 = x_0, x_1, \ldots, x_{\ell-1}$ be this path of length $\ell - 1$. Set $k = \max\{x_0, x_1, \ldots, x_{\ell-1}\}$. Then $P' = x_0, x_1, \ldots, x_k$ has length at most $\ell - 2$ and $x_{k+1}, x_{k+2}, \ldots, x_{\ell-1} \in \{0, 1, \ldots, \ell - 2\}$. By the local tournament property $P'' = \ell - 1 = x_0, x_1, \ldots, x_k, 0$ is a path from $\ell - 1$ to 0 of length at most $\ell - 1$. Therefore, we have two paths from $\ell - 1$ to 0: P_1 of length at least ℓ that uses the n-cycle, and P_2 of length at most $\ell - 1$ described above. Let t be equal to the difference in the lengths of P_1 and P_2, $t = \ell(P_1) - \ell(P_2) \geq 1$. Then there exist vertices $y_1, y_2, \ldots, y_t$ on P_1 but not on P_2. If $\ell(P_2) = \ell - 1 - s$ and $s = 0$, we have path from vertex $\ell - 1$ to 0 of length $\ell - 1$. If $s \neq 0$ we show how to use $y_1, y_2, \ldots, y_t$ to augment P_2 to obtain a path of length $\ell - 1$ from vertex $\ell - 1$ to 0. First note that $t = \ell(P_1) - \ell(P_2) = \ell(P_1) - \ell + 1 + s \geq \ell - \ell + 1 + s = 1 + s$. Therefore we have enough vertices with which to augment. We use the vertices $y_1, y_2, \ldots, y_s$ along with P_2 as follows: Each $y_i \in \{y_1, y_2, \ldots, y_s\}$ has the property that there exists vertices v_i and v_{i+1} on P_2 such that $v_i < y_i < v_{i+1}$. By the local tournament property we then also have the arcs $v_i y_i$ and $y_i v_{i+1}$. These arcs allow us to incorporate the y_is into P_2 to form a path of length $\ell - 1$ from vertex $\ell - 1$ to vertex 0.

To see that we will obtain an arc from vertex ℓ to vertex 0 we proceed in a similar way as before. Here $N^-(\ell) \subseteq \{1, 2, \ldots, \ell - 1\}$. The length of the path from vertex ℓ to vertex 0 on the n-cycle is $n - \ell \geq \ell - 1$. If $n - \ell = \ell - 1$, we already have a path of length $\ell - 1$ from vertex $\ell - 1$ to vertex 0. So we may assume that $n - \ell \geq \ell$. As before this gives us a path from vertex ℓ to vertex 0 of length at least ℓ. What remains is to see that there is also a path of length at most $\ell - 1$. If $\ell 0 \in A(D)$, this path is a single arc. If $\ell 0 \notin A(D)$, then $N^+(\ell) \subseteq \{\ell + 1, \ell + 2, \ldots, n - 1\}$. In this case we may proceed as above using the ℓ-cycle through vertex ℓ to ultimately construct a path from vertex ℓ to vertex 0 of length at most $\ell - 1$.

The above, of course, also applies to any vertex a: we obtain symmetric arcs between a and $a+(\ell-1)$ and between a and $a+\ell$, where addition is done modulo n. This implies that the result of the indicator construction, D^*, contains the undirected circulant with edges $a(a+\ell-1)$ and $a(a+\ell)$. This circulant is not bipartite. If we now apply the arc sub-indicator with vertices $\{i,j\}$ and arcs ij and ji to D^*, this sub-indicator results in the undirected portion, say D^{**}, of D^* (which contains the circulant). D^{**}-colouring is NP-complete, therefore D^*-colouring is NP-complete. Thus D-colouring is NP-complete, a contradiction. ∎

Theorem 2.7.8. *If D is a round local tournament, containing at least two cycles, then D-colouring is NP-complete.*

Proof. Since a minimum counterexample was shown not to exist, we conclude that the theorem is indeed true. ∎

2.8 Round Decomposable Local Tournaments

Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament with at least two directed cycles. Here R is a round local tournament on $n \geq 2$ vertices and each D_i is a strongly connected tournament. The proof of NP-completeness in this case will proceed as follows:

- ❒ D-colouring is NP-complete if there exists at least one D_i with $|D_i| \geq 4$.
- ❒ Now $|D_i| = 1, 3$. If R is acyclic, D-colouring is NP-complete.
- ❒ If R contains a cycle and $|D_i| = 1, 3$
  - D-colouring is NP-complete if $|D_i| = 3$ for at least two D_is.

- Now $|D_i| = 3$ for exactly one i.
  - D-colouring is NP-complete if $R = \overrightarrow{C}_n$.
  - D-colouring is NP-complete if $R \neq \overrightarrow{C}_n$.

Lemma 2.8.1. *Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament with $|D_i| \geq 4$ for at least one $i \in \{0, 1, \ldots, n-1\}$. Then the D-colouring problem is NP-complete.*

Proof. To prove this result we use the sub-indicator shown below in Figure 2.19. This is constructed using a copy of D_{i-1} and D_{i+1}, provided that both D_{i-1} and D_{i+1} exist. If only one of D_{i-1} and D_{i+1} exists (at least one exists), use only the one that exists. Add to this a vertex j such that $V(D_{i-1})$ dominates j and j dominates $V(D_{i+1})$. We take the vertices $k_1, k_2, \ldots, k_t$ of J to be exactly $V(D_{i-1}) \cup V(D_{i+1})$.

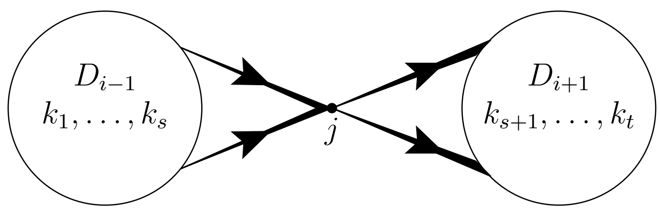

Figure 2.19: The sub-indicator for the first round decomposable case.

Furthermore we take the vertices $x_1, x_2, \ldots, x_t$ of D (required for the sub-indicator construction) also to be $V(D_{i-1}) \cup V(D_{i+1})$. In this way when we perform the sub-indicator construction we force the copy of D_{i-1} (D_{i+1}) in J to map to D_{i-1} (D_{i+1}) in D. Since j retracts to every vertex of D_i, the result of the sub-indicator construction, D^+, is exactly D_i. Since D_i is a strong tournament on at least 4 vertices, D_i-colouring is NP-complete. Therefore D-colouring is NP-complete. ∎

We have now reduced the problem to that of considering round decomposable local tournaments $D = R[D_0, D_1, \ldots, D_{n-1}]$ (with at least two directed cycles) where each $|D_i| = 1, 3$.

Lemma 2.8.2. *Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament containing at least two directed cycles, and with each $|D_i| = 1, 3$. If R is acyclic, then D-colouring is NP-complete.*

Proof. Since D contains at least two directed cycles and R is acyclic, there are at least two D_is with $|D_i| = 3$ (each such D_i is a directed triangle). For a given D_i, let v_i be the corresponding vertex in R. Let $d = \max\{d(v_i, v_j) \mid |D_i| = |D_j| = 3\}$ and define

$$k = \begin{cases} d \text{ if } d \not\equiv 0 \pmod 3, \\ d+1 \text{ if } d \equiv 0 \pmod 3. \end{cases}$$

We use as a sub-indicator two directed C_3s with a path of length k joining the two C_3s. Take as vertex i the first vertex of the path and as vertex j the last vertex of the path. This is shown in Figure 2.20.

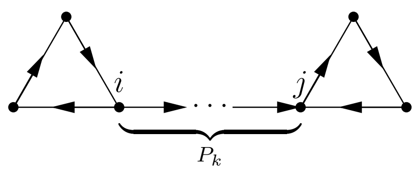

Figure 2.20: The indicator for the second round decomposable case.

Both vertex i and vertex j map to C_3s in D. Furthermore by the choice of k any vertex in a D_i with $|D_i| = 3$ is within a distance k of any other vertex in a D_j with $|D_j| = 3$ where $j \geq i$. Also, i and j cannot map to the same vertex. Each C_3 in D results in a C_3 in D^*, either the C_3 is preserved or its direction is reversed, depending on whether k is congruent to 1 or 2 (mod 3), respectively. Therefore, if D contains $t \geq 2$ C_3s, $D^* = T_t[C_3, C_3, \ldots, C_3]$, where T_t is a transitive tournament on t vertices. Hence D^*-colouring is NP-complete, and so D-colouring is NP-complete. ∎

We now consider the case where $D = R[D_0, D_1, \ldots, D_{n-1}]$ is a round decomposable local tournament where each $|D_i| = 1, 3$ and R contains at least one cycle. The proof is much the same as the case where R is acyclic except that more care is needed to ensure that one does not introduce loops when using the indicator construction.

Lemma 2.8.3. *Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament such that R contains at least one cycle and $|D_i| = |D_j| = 3$ for $i \neq j$ and $i, j \in \{0, 1, \ldots, n-1\}$. Then D-colouring is NP-complete.*

Proof. Let ℓ be the length of a shortest cycle in R. Then for any $u, v \in V(R)$, u and v are within a distance $\lfloor \ell/2 \rfloor + 1$ of each other. This is accomplished by using the arcs of a shortest cycle as well as the local tournament property. Note also that $\ell - 1 \geq \lfloor \ell/2 \rfloor + 1$. Here we use one of the two indicators shown in Figure 2.21, depending on whether $\ell - 1 \not\equiv 0 \pmod 3$ or $\ell - 1 \equiv 0 \pmod 3$.

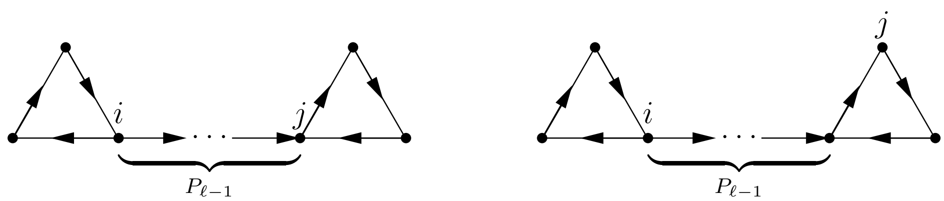

Figure 2.21: The two indicators for the third round decomposable case. The one on the left is used when $\ell - 1 \not\equiv 0 \pmod 3$, the one on the right when $\ell - 1 \equiv 0 \pmod 3$.

The two vertices i and j, in either case, map to a 3-cycle in D. Also, i and j never map to the same vertex: If the whole indicator maps to the same C_3 in D, then clearly i and j have different images (in both cases). Therefore, in order for i and j to possibly map to the same vertex in D, the image of the indicator has to involve vertices in D that are not all restricted to the same C_3. This would imply though that there exists a closed walk in R of length $\ell - 1$, which is not possible.

By mapping the whole indicator (in either case) to the same C_3 in D, we either preserve the C_3, or reverse its orientation. Furthermore any two vertices on two distinct C_3s in D are within distance $\lfloor \ell/2 \rfloor + 1$ of each other. Since $\ell - 1 \geq \lfloor \ell/2 \rfloor + 1$, the result of applying the appropriate indicator, D^*, will be a semi-complete digraph with at least two C_3s. This means that D^*-colouring is NP-complete, and so D-colouring is NP-complete. ∎

At this point we are left with a round decomposable local tournament $D = R[D_0, D_1, \ldots, D_{n-1}]$, where exactly one $|D_i| = 3$. This is dealt with in two ways: first when R is a directed n-cycle, and secondly when R is not a directed n-cycle.

Lemma 2.8.4. *Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament with exactly one $|D_i| = 3$ and $R = C_n$. Then D-colouring is NP-complete.*

Proof. The proof is by induction on n. Without loss of generality we may assume that $|D_0| = 3$. That is, $D_0 = C_3$ and $D_j = K_1$ for $j \in \{1, 2, \ldots, n-1\}$. Also, label the vertices of $R = C_n$ as $v_0, v_1, \ldots, v_{n-1}$, so that v_i corresponds to D_i for $i = 0, 1, \ldots, n-1$. The first base case ($n = 3$) is actually a tournament and is shown below in Figure 2.22. Therefore in this case D-colouring is NP-complete.

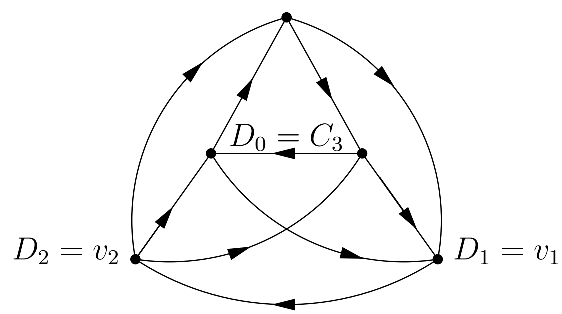

Figure 2.22: The first base case for the fourth round decomposable case.

The second base case ($n = 4$) was dealt with in Proposition 1.4.5.

This completes the two base cases. The rest of the proof proceeds in exactly the same way as the second base case (please refer to Proposition 1.4.5 and its proof). The only exception is that the end-result is not a wheel, but a smaller instance of the same problem.

Let $D = C_n[C_3, K_1, \ldots, K_1]$ and $n \geq 5$. Assume that D'-colouring is NP-complete for every $D' = C_m[C_3, K_1, \ldots, K_1]$ with $m < n$. As with the second base case, we first apply an indicator, I, equal to a path of length 2, with i and j equal to the initial and terminal vertices of the path, respectively. The result, D^*, is shown in Figure

2.23. Note that, as before, the orientation of the C_3 (shown in the figure as the small triangle) changes. We also have the arcs v_0v_1, v_0v_2, $v_{n-2}v_0$, $v_{n-1}v_0$ and $v_{n-1}v_1$. In addition to these we also have the arcs v_iv_{i+2} for $i = 1, 2, \ldots, n-3$.

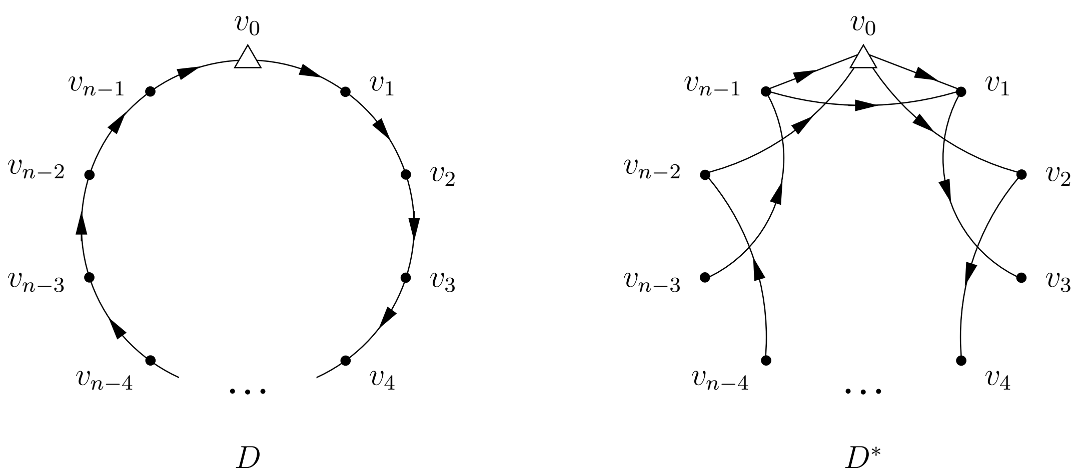

Figure 2.23: Applying the indicator I to D yields D^*.

We now apply the sub-indicator, J, that is equal to a path of length $n-2$, with the terminal vertex equal to j and the initial vertex equal to k_1. We also take $x_1 = v_2$ in D^*. Therefore we identify k_1 and v_2 and consider all retractions to D^*, recording the images of j in the process. We now complete the proof based on the parity of n.

❒ Let n be even. Then v_0 and v_2 are on an $(n/2)$-cycle in D^*: $v_0v_2v_4 \cdots v_{n-4}v_{n-2}$. This means that there exists a path of length $n/2-1$ from v_2 to v_0. Since $n \geq 5$, $n-2 > n/2-1$, and so j maps to every vertex on the C_3 (by varying where we enter the C_3 and then wrapping around it). Furthermore, v_0 is on an $(n/2)$-cycle with all the v_is where i is even and on an $(n/2+1)$-cycle with all the v_js, j odd. Using the first $n/2-1$ arcs of the P_{n-2} sub-indicator, we move from v_2 to v_0, this means that the remaining $(n-2)-(n/2-1) = n/2-1$ arcs can be used to go from v_0 to every v_i, (i even) and from v_0 to every v_j (j odd) except v_{n-1}. This is done by wrapping around the C_3 first, if needed. To see that j cannot map to v_{n-1}, note that the directed distance from v_2 to v_{n-1} is $n-1$. The core of this can be found by wrapping the path $v_0v_1v_3v_5\cdots v_{n-3}$ around the C_3. The result of this sub-indicator, D^{*+}, has a core equal to $C_{n/2}[C_3, K_1, \ldots, K_1]$. By the induction hypothesis, D^{*+}-colouring is NP-complete, so that D^*-colouring is NP-complete, implying that D-colouring is NP-complete in this case.

❒ Let n be odd. Then v_0 and v_2 are on an $((n+1)/2)$-cycle in D^* given by $v_0v_2v_4\cdots v_{n-3}v_{n-1}$. Therefore we have a path of length $(n+1)/2-1 = (n-1)/2$ from v_2 to v_0. As before j maps to every vertex on the C_3. Now v_0 is on an $((n+1)/2)$-cycle with the v_is (i even). Also, v_0 is on an $((n+1)/2)$-cycle with the v_js (j odd). Here we can use the first $(n-1)/2$ arcs of the P_{n-2} sub-indicator to go from v_2 to v_0. Then use the remaining $(n-2)-(n-1)/2 = (n-3)/2$ arcs to map j to $v_1, v_2, \ldots, v_{n-3}$ (again one may have to wrap the path around the C_3 first). It is also possible to map j to v_{n-2}, first go from v_2 to v_{n-1} using the $(n+1)/2-2 = (n-3)/2$ arcs on the cycle through the v_is with i even. Now use the remaining $(n-2)-(n-3)/2 = (n-1)/2$ arcs on the P_{n-2} for the following walk: $v_{n-1}v_1v_3v_5\cdots v_{n-2}$. It is not possible to map j to v_{n-1}, first note that the shortest cycle through v_{n-1} has length $(n+1)/2$ ($v_{n-1}v_0v_2v_4\cdots v_{n-3}$). Secondly, starting at v_2, one is forced to use the arcs $v_{2i}v_{2i+2}$ for $i = 1, 2, \ldots, (n-3)/2$ (a total of $(n-3)/2$ arcs). This takes us to v_{n-1}. At this point, if j is going to map to v_{n-1}, the remaining $(n-2)-(n-3)/2 = (n-1)/2$ arcs will have to form a closed walk from v_{n-1} back to itself, which is not possible. Thus the result of this sub-indicator, D^{*+}, is $D^* - v_{n-1}$. The core of this may be obtained by wrapping the path $v_0v_2v_4\cdots v_{n-3}$ around the C_3. This core is equal to $C_{(n+1)/2}[C_3, K_1, \ldots, K_1]$. Therefore by the induction hypothesis, D^{*+}-colouring is NP-complete, implying that D^*-colouring is NP-complete. It now follows that D-colouring is NP-complete in this case.

This completes the proof. ∎

We have now reached the final case for the round decomposable local tournaments.

Lemma 2.8.5. *Let $D = R[D_0, D_1, \ldots, D_{n-1}]$ be a round decomposable local tournament with exactly one $|D_i| = 3$ and $R \neq C_n$, but R contains at least one cycle. Then D-colouring is NP-complete.*

Proof. Label the vertices of R as $v_0, v_1, \ldots, v_{n-1}$. Then to each vertex v_i of R there corresponds a component D_i of the round decomposition. As in the previous proof, we assume that $D_0 = C_3$ and that $D_i = K_1 = v_i$ for $i = 1, 2, \ldots, n-1$.

Let P be the set of predecessors of v_0 on the shortest cycles through v_0 and define $j = \max\{k \mid v_k \in P\}$. Then v_j is the "closest" predecessor of v_0 on all the shortest cycles through v_0. Denote the length of a shortest cycle through v_0 by ℓ.

If there exists a vertex v_i with $0 \leq i \leq j-1$ such that $d^+(v_i) \geq 2$ (in R), we use the following sub-indicator: a path of length $\ell - 1$ with the terminal vertex equal to j and the initial vertex equal to k_1. Also, let any vertex of $D_0 = C_3$ be x_1. Therefore we identify k_1 with x_1 and consider all retractions of the path to D. The possible images of j are: exactly one vertex of D_0 (by wrapping around the 3-cycle in D_0) and all vertices v_t with $t = 1, 2, \ldots, j$ (by wrapping around D_0, if necessary, and then using one of the shortest cycles through v_0). This sub-indicator is illustrated in Figure 2.24.

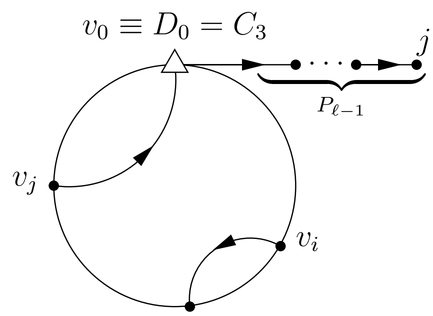

Figure 2.24: An illustration of the first sub-indicator construction in the fifth round decomposable case.

If $0 \leq i \leq j-2$, then since $j \mapsto v_i$, the result of this sub-indicator, D^+, is a round local tournament with at least two cycles. If $i = j - 1$ (that is no vertex v_t with $0 \leq t \leq j-2$ has $d^+(v_t) \geq 2$ and $d^+(v_{j-1}) \geq 2$), then v_{j-1} is the predecessor of v_j on a shortest cycle through v_0 and v_j. This implies that there exists a walk of length $\ell - 2$ from v_0 to v_{j-1}. Also, v_{j-1} is not adjacent to v_0 as this would result in a shorter cycle through v_0. Therefore v_{j-1} has an outneighbour (say v') between v_j and v_0. By following the walk of length $\ell - 2$ by the arc $v_{j-1}v'$, we see that $j \mapsto v'$. This again implies that the result of this sub-indicator, D^+, is a round local tournament with at least two cycles. Therefore D^+-colouring is NP-complete, implying that D-colouring is NP-complete.

We are now left with the case where there does not exist a vertex v_i with $0 \leq i \leq j-1$ such that $d^+(v_i) \geq 2$ (in R). This would imply that $v_0 v_1 \ldots v_j$ is an induced path (in R).

❒ If $v_j = v_{n-1}$, then $d^+(v_j) = d^+(v_{n-1}) \geq 2$ since $R \neq C_n$. Here we apply the same sub-indicator as above and the result is again a round local tournament with at least two cycles.

❒ If $v_j \neq v_{n-1}$, then $v_{j+1} \neq v_0$. Denote the length of a shortest cycle through v_{j+1} by s. Use as a sub-indicator a path of length $s-1$, with the terminal vertex equal to j and the initial vertex equal to k_1. Let x_1 be equal to v_{j+1}, so that we attach the path of length $s-1$ to v_{j+1} and consider all retractions to D. Note that $N^-(v_{j+1}) = \{v_j\}$ and since $v_j v_0 \in A(R)$, we also have $v_{j+1}v_0 \in A(R)$. Therefore the possible images of j are: every vertex in D_0 as well as each v_t with $t = 1, 2, \ldots, v_j$. This sub-indicator is illustrated in Figure 2.25.

So the result of this sub-indicator, $D^+ = C_{j+1}[C_3, K_1, K_1, \ldots, K_1]$. We know that D^+-colouring is NP-complete, therefore D-colouring is NP-complete.

∎

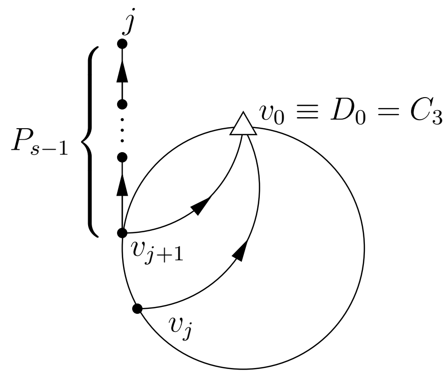

Figure 2.25: An illustration of the second sub-indicator construction in the fifth round decomposable case.

2.9 Non-Round Decomposable Local Tournaments

In discussing the complexity of non-round decomposable local tournaments, its not too surprising that the structural result in Lemma 2.3.2 plays a central role. Note that a non-round decomposable local tournament has at least two directed cycles, since a local tournament with at most one directed cycle is round decomposable. The proof that colouring by a non-round decomposable local tournament is NP-complete proceeds along a familiar line. We assume that the result is false and examine a smallest counterexample (one that has the minimum number of vertices). Such a counterexample is then shown not to exist, and the result follows.

The counterexample will have the structure described in Lemma 2.3.2, since it is non-round decomposable. Let S, D'_1, D'_2 and D'_3 be the subsets of vertices of D defined in Lemma 2.3.2. The first lemma deals with the possibilities for in-and out-neighbours of the vertices in S.

Lemma 2.9.1. *Let D be a non-round decomposable local tournament that is also a minimum counterexample to the D-colouring problem. Then every vertex in S can only have in-neighbours or out-neighbours in D'_2, but not both.*

Proof. Since D is a counterexample, D-colouring is not NP-complete.

D is not round decomposable so it has the structure described in Lemma 2.3.2. S, D'_1, D'_2 and D'_3 are the subsets of vertices of D defined in Lemma 2.3.2.

Let $s \in S$ and $a, a' \in D'_2$ such that $as, sa' \in A(D)$. D does not have symmetric arcs, so $a \neq a'$. By Lemma 2.3.2 there is a vertex $y \in D'_1$ with $ay, a'y, ys \in A(D)$. Lemma 2.3.2 and Proposition 2.3.4 imply that there is a vertex $x \in D'_3$ such that $sx, xa' \in A(D)$. This is illustrated in Figure 2.26.

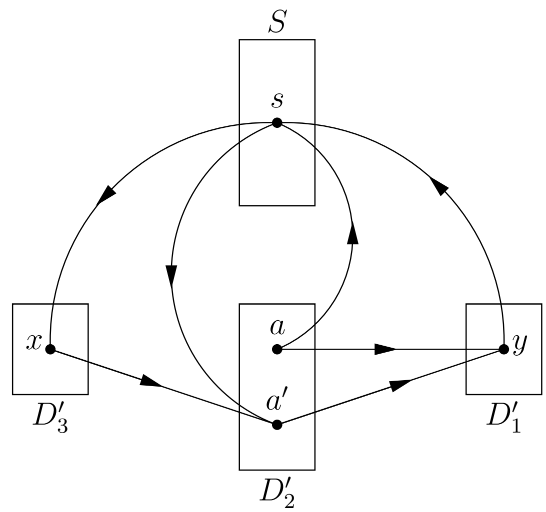

Figure 2.26: Proving that no vertex in S can have both in-and out-neighbours in D'_2.

By Theorem 2.3.1 D'_2 is a tournament. Therefore a and a' are adjacent.

❒ If $aa' \in A(D)$, we use a sub-indicator equal to a directed path of length 2, where the terminal vertex is equal to j and the initial vertex is identified with vertex a in D. That is we consider all retractions of this path of length 2 to D keeping track of the images of j and identifying its initial vertex with vertex a in D. The vertex j maps to (at least) the following vertices: s, a', x, y but not the vertex a. The result of this sub-indicator, D^+, will have fewer vertices, but will still contain at least two cycles ($sxa'y$ and $sa'y$). Since D^+ is smaller than D it cannot be a counterexample, and so D^+-colouring is NP-complete, but this would imply that D-colouring is NP-complete — a contradiction.

❒ If $a'a \in A(D)$, we apply an indicator equal to a directed path of length 2 to D. This will produce the following symmetric arcs: $a's$ ($a'ys$ and sxa'), sa ($sa'a$ and ays), ax (asx and $xa'a$), xy ($xa'y$ and ysx) and ya' (ysa' and $a'ay$). Therefore the result of the indicator, D^*, contains an undirected 5-cycle, $a'saxy$, and so it is not bipartite. The undirected portion of D^* may be extracted using an arc sub-indicator equal to a pair of symmetric arcs. This produces an undirected, non-bipartite graph D^{*+}. D^{*+}-colouring is NP-complete, implying that D^*-colouring is NP-complete, which in turn implies that D-colouring is NP-complete — a contradiction.

∎

Theorem 2.9.2. *Let D be a non-round decomposable local tournament, then D-colouring is NP-complete.*

Proof. Assume that the theorem is false. Among all counterexamples to the theorem, let D be one with the minimum number of vertices.

Therefore D is a non-round decomposable local tournament such that D-colouring is not NP-complete. Furthermore, D has the minimum number of vertices possible.

Hence D has the structure described in Lemma 2.3.2. As before we let S, D'_1, D'_2 and D'_3 be the subsets of vertices of D defined in Lemma 2.3.2.

By Lemma 2.3.2 there are arcs between S and D'_2, oriented in opposite directions and with their vertices in specific locations with respect to one another. That is there are vertices $s, s' \in S$ and $a, a' \in D'_2$ such that $sa, a's' \in A(D)$. By the previous lemma $s \neq s'$, but a may be equal to a'. Furthermore, Lemma 2.3.2 implies the existence of a vertex $y \in D'_1$ such that $ay, a'y, ys, ys' \in A(D)$. Also, by Lemma 2.3.2 and Proposition 2.3.4, there exists a vertex $x \in D'_3$ together with the arcs $sx, s'x, xa, xa'$. We now distinguish two cases: $ss' \in A(D)$ and $s's \in A(D)$ (by Lemma 2.3.2 S is semi-complete).

❒ Let $ss' \in A(D)$. Then $s', a \in N^+(s)$ and since D is a local tournament, s' and a have to be adjacent. The previous lemma forbids the arc $s'a$ and so $as' \in A(D)$. This is illustrated in Figure 2.27.

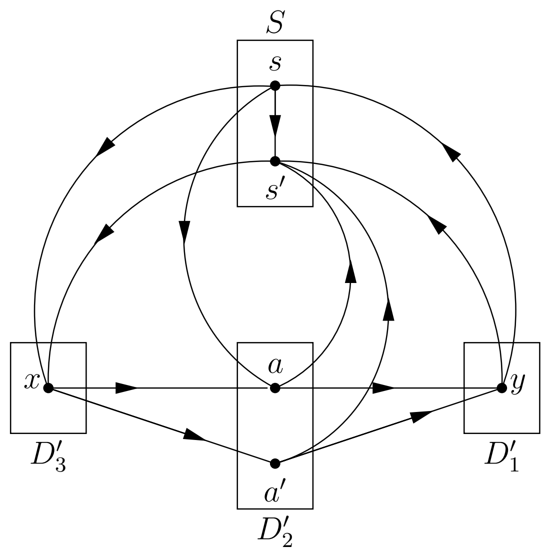

Figure 2.27: The case $ss' \in A(D)$.

Apply a sub-indicator equal to a directed path of length 2, where the terminal vertex is equal to j and the initial vertex is identified with the vertex s in D. We retract the path to D and determine the images of j. The images of j include the following vertices: s', x, y, a but not s. The result of this sub-indicator, D^+, has fewer vertices than D, but still contains at least two cycles ($s'xay$ and $s'xa$). Since its smaller, D^+-colouring is NP-complete, implying that D-colouring is NP-complete — a contradiction.

❒ Let $s's \in A(D)$. This is shown in Figure 2.28.

Here we apply an indicator equal to a directed path of length 2. This produces the following symmetric arcs: $s'x$ ($s'sx$ and $xa's'$), xy (xay and ysx), ys ($ys's$ and say), sa (sxa and ays) and as' (ays' and $s'xa$). The result of this indicator, D^*, will therefore contain an undirected 5-cycle, $s'xysa$. Thus the undirected

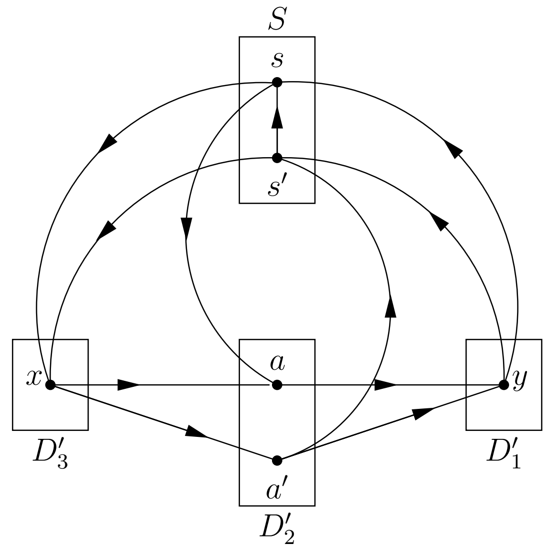

Figure 2.28: The case $s's \in A(D)$.

portion of D^* is not bipartite. As before, we may extract the undirected portion with an arc sub-indicator equal to a pair of symmetric arcs. All of this implies that D^*-colouring is NP-complete which in turn says that D-colouring is NP-complete — a contradiction.

This shows that the counterexample does not exists. Therefore the theorem is true. ∎

2.10 The Dichotomy for Connected Local Tournaments

Summarising the results of the previous sections, we obtain the following dichotomy for connected local tournaments.

Theorem 2.10.1. *Let T be a connected local tournament.*

- *If T is acyclic, then T-colouring is polynomial.*
- *If T is unicyclic and T is a directed cycle or T has the structure shown in Figure 2.2 with T_2 a tournament and at least one of T_1 and T_3 is empty, then T-colouring is polynomial. Otherwise T-colouring is NP-complete.*

❒ *If T contains at least two cycles, then T-colouring is NP-complete.*

We also have the following theorem.

Theorem 2.10.2. *The dichotomy for list homomorphisms to connected local tournaments is exactly the same as the dichotomy for homomorphisms to local tournaments.*

Proof. Let T be a connected local tournament. If HOM_T is NP-c, then LIST-HOM_T will also be NP-complete (set all lists of the input equal to $V(T)$).

On the other hand if HOM_T is polynomial, then LIST-HOM_T is also polynomial since our algorithm in the polynomial case is a list processing algorithm. ∎

3

Weak Near-Unanimity Functions

3.1 Introduction

In this chapter we study the algebraic approach to the CSP(T) problem as outlined in Chapter 1. In particular T is a digraph so that we study the H-colouring problem by examining the polymorphisms of H.

Let H be a digraph. A homomorphism $f : H^k \to H$ is called a *polymorphism of arity* k, where H^k is the k-fold categorical product of H defined as follows.

$$V(H^k) = \underbrace{V(H) \times V(H) \times \cdots \times V(H)}_{k},$$

$$(x_1, x_2, \ldots, x_k) \to (y_1, y_2, \ldots, y_k) \text{ is an arc of } H^k \text{ if and only if}$$

$$x_i \to y_i \text{ is an arc of } H,\ 1 \leq i \leq k.$$

A polymorphism of arity k is called *weak near-unanimity function of arity* k (WNUF$_k$) if it satisfies the following two properties:

❒ $f(x, x, x, \ldots, x) = x$ (idempotent); and,

❒ $f(y, x, x, \ldots, x) = f(x, y, x, \ldots, x) = f(x, x, y, \ldots, x) = \cdots = f(x, x, x, \ldots, y)$ (weakly nearly-unanimous), where x and y are vertices of H.

The collection of all weak near-unanimity functions of arity k on H will be denoted by $\mathrm{WNUF}_k(H)$.

If $f \in \mathrm{WNUF}_k(H)$ and $f(x_1, x_2, \ldots, x_k) \in \{x_1, x_2, \ldots, x_k\}$, then f is termed a *conservative* WNUF_k.

If the WNUF_k $f : H^k \to H$ satisfies $f(y, x, x, \ldots, x) = f(x, y, x, \ldots, x) = f(x, x, y, \ldots, x) = \cdots = f(x, x, x, \ldots, y) = x$, for all $x, y \in V(H)$, then it is called a *near unanimity function of arity k* (NUF_k).

Recall that if H does not have a WNUF_k of any arity $k > 1$, then H-colouring is NP-complete (Theorem 1.6.4). It is conjectured that the existence of a WNUF_k of arity $k > 1$ implies that H-colouring is polynomial time solvable (Conjecture 1.6.5).

In this chapter, we prove that a number of digraphs do not possess a WNUF_k with $k > 1$. Included among these will be all tournaments with at least two cycles. This gives an independent proof of the result in [5]. We also develop WNUF analogs of the vertex (and arc) sub-indicator construction and the indicator construction.

On the other hand, we will show that a number of known polynomial H-colouring problems all have a WNUF_3. This lends support to the conjecture that WNUFs are the right measure for the complexity of H-colouring.

3.2 Boosting the Arity

Often when showing the nonexistence of a WNUF_k we will only consider arities $k \geq 3$. The results in this section show that we don't have to consider an arity of $k = 2$ as this would imply a WNUF_4.

Proposition 3.2.1. *Let $f \in \mathrm{WNUF}_k(H)$ and $g \in \mathrm{WNUF}_l(H)$. Then there exists an $h \in \mathrm{WNUF}_{kl}(H)$.*

Proof. Let $f : H^k \to H$ and $g : H^l \to H$ be weak near-unanimity functions on H of arities k and l respectively. Define $h : H^{kl} \to H$ as follows

$$h(x_1, x_2, x_3, \ldots, x_{kl}) =$$
$$f(g(x_1, x_2, \ldots, x_l), g(x_{l+1}, x_{l+2}, \ldots, x_{2l}), \ldots, g(x_{kl-l+1}, x_{kl-l+2}, \ldots, x_{kl})).$$

If $(x_1, x_2, \ldots, x_{kl})$ is adjacent to $(y_1, y_2, \ldots, y_{kl})$, then $(x_i, x_{i+1}, \ldots, x_{i+l-1})$ is adjacent to $(y_i, y_{i+1}, \ldots, y_{i+l-1})$ for $i = 1, l+1, 2l+1, \ldots, kl-l+1$. Therefore $g(x_i, x_{i+1}, \ldots, x_{i+l-1})$ is adjacent to $g(y_i, y_{i+1}, \ldots, y_{i+l-1})$ for $i = 1, l+1, 2l+1, \ldots, kl-l+1$. Finally $f(g(x_1, x_2, \ldots, x_l), g(x_{l+1}, x_{l+2}, \ldots, x_{2l}), \ldots, g(x_{kl-l+1}, x_{kl-l+2}, \ldots, x_{kl}))$ is adjacent to $f(g(y_1, y_2, \ldots, y_l), g(y_{l+1}, y_{l+2}, \ldots, y_{2l}), \ldots, g(y_{kl-l+1}, y_{kl-l+2}, \ldots, y_{kl}))$. This shows that h is a homomorphism.

It is clear that h is idempotent.

Consider the kl-tuples

$$(y, x, x, \ldots, x),\ (x, y, x, \ldots, x),\ (x, x, y, \ldots, x),\ \ldots,\ (x, x, x, \ldots, y),$$

where x and y are vertices of H. We also have that

$$\begin{aligned} g(y, x, x, \ldots, x) = g(x, y, x, \ldots, x) = g(x, x, y, \ldots, x) = \cdots &= g(x, x, x, \ldots, y) \\ &= a \text{ (say)}. \end{aligned}$$

Note that g is defined on l-tuples. When applying h to the kl-tuples shown above, the result is one of the following

$$f(a, x, x, \ldots, x),\ f(x, a, x, \ldots, x),\ f(x, x, a, \ldots, x),\ \ldots,\ f(x, x, x, \ldots, a),$$

depending on which block of length l y finds itself. Since f is a weak near unanimity function, these are all equal. This shows that h is also weakly nearly unanimous. ∎

Corollary 3.2.2. *Let $f \in \mathrm{WNUF}_k(H)$. Then there exists a $g \in \mathrm{WNUF}_{k^2}(H)$.*

Proof. Take $g = f$ in the proof of Lemma 3.2.1. ∎

3.3 Polynomial Problems With WNUFs

In section we show that a number of polynomial problems that were treated by Gutjahr, Woeginger and Welzl [29] have WNUFs.

We encountered these concepts in Chapter 2. We repeat the necessary definitions and results here.

Recall that an enumeration $\{h_1, h_2, \ldots, h_n\}$ of the vertices of a digraph H is called an $\underline{X}$-*enumeration* if the following property holds: if $h_i h_j$ and $h_k h_l$ are arcs of H, then $\min(h_i, h_k)\min(h_j, h_l)$ is also an arc of H, where the minimum is taken with respect to the $\underline{X}$-enumeration.

Theorem 3.3.1 (Gutjahr, Woeginger and Welzl [29])**.** *Let H be a digraph such that H admits an $\underline{X}$-enumeration. Then the H-colouring problem is solvable in polynomial time.*

Theorem 3.3.2. *If a digraph H has an $\underline{X}$-enumeration, then H has a conservative* WNUF_k, $k > 1$.

Proof. Define a WNUF_k $f : H^k \to H$ as follows

$$f(x_1, x_2, \ldots, x_k) = \min\{x_1, x_2, \ldots, x_k\},$$

where the minimum is with respect to the $\underline{X}$-enumeration. Its easy to see that f is idempotent and weakly nearly unanimous. We show that f is a homomorphism. Consider the arc $(x_1, x_2, \ldots, x_k)(y_1, y_2, \ldots, y_k)$ in H^k. Let's say that $\min\{x_1, x_2, \ldots, x_k\} =$

x_i and $\min\{y_1, y_2, \ldots, y_k\} = y_j$ with $i, j \in \{1, 2, \ldots, k\}$. Then $x_i y_i$ and $x_j y_j$ are arcs in H. By the $\underline{X}$-enumeration, $\min\{x_i, x_j\} \min\{y_i, y_j\} = x_i y_j$ is an arc of H. Therefore f is a homomorphism. ∎

Theorem 3.3.3. *Let C_k be a directed cycle of length k. Then C_k has a conservative* NUF$_t$*, $t \geq 3$.*

Proof. Each vertex of C_k has a unique in-neighbour and a unique out-neighbour. Therefore, if $(x_1, x_2, \ldots, x_t) \in C_k^t$, then $(x_1, x_2, \ldots, x_t)$ has a unique in-neighbour and unique out-neighbour as in C_k^t as well. Define a conservative NUF$_t$ $f : C_k^t \to C_k$ as follows: if $(x_1, x_2, \ldots, x_t)$ is unanimous or nearly unanimous map the 3-tuple accordingly; If $(x_1, x_2, \ldots, x_t)$ is not unanimous nor nearly unanimous, let $f(x_1, x_2, \ldots, x_t) = x_1$. Let $(x_1, x_2, \ldots, x_t)(y_1, y_2, \ldots, y_t)$ be an arc in C_k^t. Due to the uniqueness of in and out neighbours in C_k^t, the degree of unanimity (or lack thereof) in $(x_1, x_2, \ldots, x_t)$ is reflected in $(y_1, y_2, \ldots, y_t)$. Thus f preserves all arcs. ∎

An extension of the $\underline{X}$-enumeration was also discussed by Gutjahr, Woeginger and Welzl [29]. A digraph H is said to have the *C_k-extended $\underline{X}$ property* if the following holds.

❒ There exists a homomorphism $f : H \to C_k$. Label the vertices of C_k as $\{0, 1, \ldots, k-1\}$. Let $V_i = \{v \in V(H) \mid f(v) = i\}$. Therefore for each arc xy of H there exists a unique integer $i \in \{0, 1, \ldots, k-1\}$ such that $x \in V_i$ and $y \in V_{i+1}$, where the subscripts are handled modulo k.

❒ There is an enumeration of the vertices of H such that each subgraph of H induced by $V_i \cup V_{i+1}$, $0 \leq i \leq k-1$, has the $\underline{X}$ property with respect to this enumeration.

Theorem 3.3.4 (Gutjahr, Woeginger and Welzl [29])**.** *If a digraph H has the C_k-extended $\underline{X}$ property, then H-colouring is polynomial time solvable.*

Theorem 3.3.5. *Let H be a digraph that has the C_k-extended $\underline{X}$ property. Then H has a conservative* WNUF_t*, $t \geq 3$.*

Proof. By the definition of the C_k-extended $\underline{X}$ property, H has a homomorphism to C_k. Denote this homomorphism by $f : H \to C_k$. Also C_k has a conservative NUF_t, say $g : C_k^t \to C_k$. Define a WNUF_t on $h : H^t \to H$ as follows: given a t-tuple $(x_1, x_2, \ldots, x_t) \in H^t$, apply f coordinate-wise to the t-tuple to produce $(f(x_1), f(x_2), \ldots, f(x_t))$. Now apply g to this new t-tuple, the result will be $f(x_i)$, $i \in \{1, 2, \ldots, t\}$ since g is conservative. Next, form the intersection $f^{-1}(f(x_i)) \cap \{x_1, x_2, \ldots, x_t\}$, this intersection is non-empty since it contains at least x_i. Finally take the minimum with respect to the $\underline{X}$-enumeration of this set. This is well defined since the C_k-extended $\underline{X}$-enumeration is defined over arcs of C_k and in particular there is a minimum on each set of the form $f^{-1}(u)$, where $u \in V(C_k)$. The mapping therefore is

$$(x_1, x_2, \ldots, x_t) \rightsquigarrow (f(x_1), f(x_2), \ldots, f(x_t)) \rightsquigarrow$$
$$g(f(x_1), f(x_2), \ldots, f(x_t)) = f(x_i) \rightsquigarrow \min\{f^{-1}(f(x_i)) \cap \{x_1, x_2, \ldots, x_t\}\}.$$

It's not too difficult to see that h is idempotent, weakly nearly-unanimous and conservative. What remains to be shown is that h is a homomorphism.

Let $(x_1, x_2, \ldots, x_t)(y_1, y_2, \ldots, y_t)$ be an arc of H^t. Therefore

$$(f(x_1), f(x_2), \ldots, f(x_t))(f(y_1), f(y_2), \ldots, f(y_t))$$

is an arc of C_k^t. Let $g(f(x_1), f(x_2), \ldots, f(x_t)) = a$ and $g(f(y_1), f(y_2), \ldots, f(y_t)) = b$. Since g is a homomorphism, ab is an arc of C_k and, by the extended $\underline{X}$-enumeration, there is an $\underline{X}$-enumeration on the subgraph of H induced by $f^{-1}(a) \cup f^{-1}(b)$. We have $h(x_1, x_2, \ldots, x_t) = \min\{f^{-1}(a) \cap \{x_1, x_2, \ldots, x_t\}\} = x_i$ (say) and $h(y_1, y_2, \ldots, y_t) =$

$\min\{f^{-1}(b) \cap \{y_1, y_2, \ldots, y_t\}\} = y_j$ (say), where $i, j \in \{1, 2, \ldots, j\}$. Consider the arcs $x_i y_i$ and $x_j y_j$ in H. By the homomorphism $f : H \to C_k$, $y_i \in f^{-1}(b)$ and that $x_j \in f^{-1}(a)$. Therefore there is an $\underline{X}$-enumeration on x_i, y_i, x_j, y_j and so $\min\{x_i, x_j\} \min\{y_i, y_j\} = x_i y_j$ is an arc in H. Note that $\min\{x_i, x_j\} = x_i$ since both x_i and x_j are in $f^{-1}(a)$; the same applies to y_i and y_j. ∎

Of course Theorem 3.3.2 is really a corollary of Theorem 3.3.5 since we may take C_k to be a loop in the case of Proposition 3.3.2.

Our next result deals with the $\underline{X}$-graft extension that we encountered in Chapter 2. We repeat the definitions here, albeit in a slightly different (but equivalent) form.

Suppose that $h_1, h_2, \ldots, h_n$ is an $\underline{X}$-enumeration of the vertices of a digraph H_1, where there is a loop at h_n but not at any other vertex. Let H_2 be another digraph. Form a new digraph H by replacing h_n by H_2 and, whenever $h_i h_n$ (respectively $h_n h_i$) is an arc of H_1 (with $i < n$), add arcs joining h_i to (respectively from) all vertices of H_2. The digraph H is called $\underline{X}$-graft(H_1,H_2).

Theorem 3.3.6 (Gutjahr, Woeginger and Welzl [29])**.** *Let H=$\underline{X}$-graft(H_1,H_2) such that HOM_{H_2} is polynomial. Then HOM_H is polynomial.*

We will show that if H=$\underline{X}$-graft(H_1,H_2), where H_2 has a (conservative) WNUF, then H has a (conservative) WNUF. In fact, we will prove a slightly stronger result that does not depend directly on the existence of an $\underline{X}$-enumeration for H_1. Instead, we will replace this by a condition on a WNUF$_k$ from H_1^k to H_1.

Theorem 3.3.7. *Let H_1 be a digraph with exactly one loop, say at the vertex v. Furthermore let $f_1 : H_1^k \to H_1$ be a* WNUF$_k$ *such that $f_1^{-1}(v) = \{(v, v, \ldots, v)\}$. Let H_2 be a digraph with a* WNUF$_k$ *$f_2 : H_2^k \to H_2$. Form a new digraph H by replacing v by H_2 and, whenever hv (respectively vh) is an arc of H_1 (with $v \neq h$), add arcs joining h to (respectively from) all vertices of H_2. Then H has a* WNUF$_k$*. If both f_1 and f_2 have a conservative* WNUF$_k$*, then H has a conservative* WNUF$_k$*.*

Proof. Let H_1, H_2, f_1, f_2 and H be given as in the statement of the theorem. Define a WNUF$_k$ $h : H^k \to H$ by

$$h(x_1, x_2, \ldots, x_k) = \begin{cases} f_1(x_1, x_2, \ldots, x_k) & \text{if } x_1, x_2, \ldots, x_k \in H_1 - \{v\}, \\ f_2(x_1, x_2, \ldots, x_k) & \text{if } x_1, x_2, \ldots, x_k \in H_2, \\ f_1(y_1, y_2, \ldots, y_k) & \text{with } y_i = x_i \text{ if } x_i \in H_1 - \{v\} \text{ and} \\ & y_i = v \text{ if } x_i \in H_2. \end{cases}$$

Note that in the third case above there is at least one $x_i \in H_1 - \{v\}$ and at least one $x_j \in H_2$. This means that $(y_1, y_2, \ldots, y_k)$ is never equal to $(v, v, \ldots, v)$. Furthermore, by the assumption on the inverse image of v, h never maps to v.

We show that h is a homomorphism. Let $(a_1, a_2, \ldots, a_k) \to (b_1, b_2, \ldots, b_k)$ be an arc of H^k. There are nine possibilities to consider because of the three cases present in the definition of h that apply to each k-tuple vertex in the pair forming the arc (throughout keep in mind that vv is a loop in H_1).

❒ If both ends are in H_1^k or in H_2^k, h preserves the arc since the same homomorphism is applied to both (f_1 or f_2).

❒ If one end (either one) is in H_1^k and the other is such that it contains a vertex from $H_1 - \{v\}$ and a vertex from H_2, then h preserves the arc. The reason for this is that ultimately (after certain vertices have been replaced by vs) we apply the same homomorphism to the k-tuples (namely f_1).

❒ If both k-tuples have a vertex in $H_1 - \{v\}$ and a vertex in H_2, then again after the necessary vertices have been replaced by vs we map them under them same homomorphism (f_1), so h preserves the arc.

❒ Consider now the case where each $a_i \in H_1 - \{v\}$ and each $b_j \in H_2$. The image of $(b_1, b_2, \ldots, b_k)$ under h is in H_2. Also $(a_1, a_2, \ldots, a_k) \to (v, v, \ldots, v)$

and this implies that $h(a_1, a_2, \ldots, a_k) = f_1(a_1, a_2, \ldots, a_k) \in N^-(v)$. By the definition of H the arc is preserved. The reverse situation, each $a_i \in H_2$ and each $b_j \in H_1 - \{v\}$ is handled in a similar way.

❒ Finally, consider the case where each $a_i \in H_2$, some $b_s \in H_1 - \{v\}$ and some $b_t \in H_2$. Note that $h(a_1, a_2, \ldots, a_k) = f_2(a_1, a_2, \ldots, a_k) \in H_2$. On the other hand $h(b_1, b_2, \ldots, b_k) = f_1(z_1, z_2, \ldots, z_k)$ where $z_i = b_i$ if $b_i \in H_1 - \{v\}$ and $z_i = v$ if $b_i \in H_2$. Since $(v, v, \ldots, v) \to (z_1, z_2, \ldots, z_k)$ this implies that $f_1(z_1, z_2, \ldots, z_k) \in N^+(v)$. Therefore h preserves the arc. The reverse, some $a_s \in H_1 - \{v\}$, some $a_t \in H_2$ and each $b_i \in H_2$ is handled in a similar way.

The mapping h is idempotent since these are covered by the first two cases in the definition of h. Near unanimity follows easily when $x_1, x_2, \ldots, x_k$ are all in $H_1 - \{v\}$ or when they are all in H_2 (cases one and two in the definition of h, respectively). The other two nearly unanimous cases involve two distinct vertices, say $a \in H_1 - \{v\}$ and $b \in H_2$ (so that we are dealing with case three in the definition of h). There is either exactly one a and $k-1$ bs or $k-1$ as and exactly one b among $\{x_1, x_2, \ldots, x_k\}$. In either case the corresponding tuple of ys that is formed is nearly unanimous so that when f_1 is applied to it, the result is always the same. Therefore h is nearly unanimous.

This shows that h is indeed a WNUF$_k$. ∎

Corollary 3.3.8. *Let H_1 be a digraph with an $\underline{X}$-enumeration $\{h_1, h_2, \ldots, h_n\}$ (i.e. there is a loop at h_n). If H_2 has a (conservative)* WNUF$_k$*, then H=graft(H_1,H_2) has a (conservative)* WNUF$_k$.

Proof. By Theorem 3.3.2, H_1 has a conservative WNUF$_k$, say f_1, such that $f_1^{-1}(h_n) = \{(h_n, h_n, \ldots, h_n)\}$. Therefore by Theorem 3.3.7, if H_2 has a (conservative) WNUF$_k$, then H has a (conservative) WNUF$_k$. ∎

3.4 Some Non-existence Results for WNUFs

In this section we take a different approach to WNUFs and exhibit a number of digraphs that do not possess a WNUF_k of any arity $k > 1$. This immediately implies that the corresponding H-colouring problem is NP-complete.

Let D be a digraph with the following properties:

- ❒ There is a homomorphism $f : P_2 \to D$, where P_2 is a directed path of length two: $V(P_2) = \{a, b, c\}$ and $E(P_2) = \{ab, bc\}$. As a convenience we will write $f(a) = 0, f(b) = 1$ and $f(c) = 2$.
- ❒ $N^+(0) = \{1\}$, $N^+(1) = \{2\}$, $N^-(1) = \{0\}$, $N^-(2) = \{1\}$, $N^-(\{1, 2\}) = \emptyset$ and $N^+(\{0, 1\}) = \emptyset$.
- ❒ The oriented cycles, 3.1 and 3.2, (of net-length one) defined below do not have homomorphisms to D.

Add two vertices t and b to D as well as all arcs tv, $v \in V(D)$, ub, $u \in V(D)$ and lastly either the arc from b to t or no arc between t and b to form a new digraph D_b^t.

We now define oriented cycles in $(D_b^t)^k$ for some integer $k \geq 2$:

$$\begin{gathered}(111\ldots11t) \to (222\ldots221) \leftarrow (111\ldots1t0) \to (222\ldots211) \leftarrow \\ (111\ldots t00) \to \cdots \leftarrow (11t\ldots000) \to (221\ldots111) \leftarrow (1t0\ldots000) \to \\ (211\ldots111) \leftarrow (t00\ldots000) \to (111\ldots111) \to (222\ldots222) \leftarrow (111\ldots11t),\end{gathered} \tag{3.1}$$

$$\begin{aligned}&(111\dots11b) \leftarrow (000\dots001) \rightarrow (111\dots1b2) \leftarrow (000\dots011) \rightarrow\\&(111\dots b22) \leftarrow \cdots \rightarrow (11b\dots222) \leftarrow (001\dots111) \rightarrow (1b2\dots222) \leftarrow\\&(011\dots111) \rightarrow (b22\dots222) \leftarrow (111\dots111) \leftarrow (000\dots000) \rightarrow (111\dots11b).\end{aligned} \tag{3.2}$$

Both cycles have $2k+1$ vertices and net-length one. Furthermore note that the vertices on the cycles are of the form

$$\begin{aligned}&(x_1x_2x_3\dots x_{k-2}x_{k-1}x_k),\\&(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k) \text{ or}\\&(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k),\end{aligned}$$

where each $x_j \in \{0,1,2\}$.

Theorem 3.4.1. *The digraph D_b^t defined above does not have a* WNUF_k *for any $k \geq 2$.*

Proof. We assume that $f : (D_b^t)^k \rightarrow D_b^t$ is a WNUF_k and derive a contradiction. Throughout this proof let $x_i \in \{0,1,2\}$ (defined above).

Since $f(ttt\dots t) = t$ and $(ttt\dots t)$ is adjacent to $(x_1x_2x_3\dots x_k)$, $f(x_1x_2x_3\dots x_k) \in V(D)$. Also, $(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is adjacent to both $(bbb\dots b1b\dots bbb)$ and $(bbb\dots b2b\dots bbb)$.

Also $(000\dots000\dots000)$ is adjacent to $(bbb\dots b1b\dots bbb)$, $(111\dots111\dots111)$ is adjacent to $(bbb\dots b2b\dots bbb)$, $f(000\dots000\dots000) = 0$ and $f(111\dots111\dots111) = 1$, it follows that $f(bbb\dots b1b\dots bbb) \in \{1,b\}$ and $f(bbb\dots b2b\dots bbb) \in \{2,b\}$.

If either one of $(bbb\dots b1b\dots bbb)$ or $(bbb\dots b2b\dots bbb)$ is mapped to b, then $(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ must be mapped into $V(D)$ (for any choice of the x_js). This would force the first oriented cycle, 3.1, shown above to be mapped into

$V(D)$ which is not possible. Therefore

$$f(bbb\dots b1b\dots bbb) = 1 \text{ and } f(bbb\dots b2b\dots bbb) = 2.$$

The vertex $(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is an in-neighbour of the vertices $(bbb\dots b1b\dots bbb)$ and $(bbb\dots b2b\dots bbb)$, and so $f(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is a common in-neighbour of vertices 1 and 2. Therefore

$$f(x_1x_2x_3\dots x_itx_{i+1}\dots x_{k-2}x_{k-1}x_k) = t.$$

It is also the case that $(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is adjacent from both $(ttt\dots t0t\dots ttt)$ and $(ttt\dots t1t\dots ttt)$. Furthermore since $(111\dots 111\dots 111)$ is adjacent from $(ttt\dots t0t\dots ttt)$, $(222\dots 222\dots 222)$ is adjacent from $(ttt\dots t1t\dots ttt)$, $f(111\dots 111\dots 111) = 1$ and $f(222\dots 222\dots 222) = 2$, it follows that

$$f(ttt\dots t0t\dots ttt) \in \{t, 0\} \text{ and } f(ttt\dots t1t\dots ttt) \in \{t, 1\}.$$

Again if either one of $(ttt\dots t0t\dots ttt)$ or $(ttt\dots t2t\dots ttt)$ is mapped to t, then $(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ must be mapped into $V(D)$. This would force the second oriented cycle, 3.2, shown above to be mapped into $V(D)$ which is not possible. Therefore, $f(ttt\dots t0t\dots ttt) = 0$ and $f(ttt\dots t1t\dots ttt) = 1$.

The vertex $(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is an out-neighbour of the vertices $(ttt\dots t0t\dots ttt)$ and $(ttt\dots t1t\dots ttt)$, and so $f(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k)$ is a common out-neighbour of vertices 0 and 1. Therefore

$$f(x_1x_2x_3\dots x_ibx_{i+1}\dots x_{k-2}x_{k-1}x_k) = b.$$

We now observe that the vertex $(t00\dots 0)$ is adjacent to the vertex $(1b1\dots 1)$. Also

$f(t00\dots0) = t$ and $f(1b1\dots1) = b$, but t and b are not adjacent or the adjacency is from b to t. Therefore f is not a homomorphism. ∎

Corollary 3.4.2. *If D is a directed cycle or a directed path of length at least two, then D_b^t does not have a* WNUF_k *for any $k \geq 2$.*

The following Corollary was actually the starting point for many of the WNUF results in this chapter.

Corollary 3.4.3. *The local tournament shown in Figure 2.3 does not have a* WNUF_k *for any $k \geq 2$.*

It is interesting to note that D may be a directed cycle of length two, in which case $0 = 2$ in all of the above.

Let D_1 and D_2 be two digraphs with the following properties:

- ❐ There exist homomorphisms $f : P_3 \to D_1$ and $g : P_2 \to D_2$, where $P_3 = wxyz$ is a directed path of length three and $P_2 = tuv$ is a directed path of length two. Let $f(w) = 0$, $f(x) = 1$, $f(y) = 2$, $f(z) = 3$, $g(t) = a$, $g(u) = b$ and $g(v) = c$.
- ❐ $N^+(a) = \{b\}$, $N^-(b) = \{a\}$ and $N^-(\{a, b\}) = \emptyset$ (that is a and b do not have a common in-neighbour).
- ❐ The oriented cycle, 3.3, defined below does not have a homomorphism to D_1.

Form a new digraph $D_1 \rightharpoonup D_2$ by taking a copy of D_1, a copy of D_2 and adding the following arcs: $0a, 1a, 0b, 1b, 2b, 1c, 2c, 3c$.

Consider the following oriented cycle in $(D_1 \rightharpoonup D_2)^k$ for some integer $k \geq 2$:

$$\begin{gathered}(111\dots11a) \leftarrow (000\dots001) \rightarrow (111\dots1a2) \leftarrow (000\dots011) \rightarrow \\ (111\dots a22) \leftarrow \cdots \rightarrow (11a\dots222) \leftarrow (001\dots111) \rightarrow (1a2\dots222) \leftarrow \\ (011\dots111) \rightarrow (a22\dots222) \leftarrow (111\dots111) \leftarrow (000\dots000) \rightarrow (111\dots11a).\end{gathered} \tag{3.3}$$

The vertices on the oriented cycle are of the form

$$(111\dots 1a2\dots 222),$$
$$(000\dots 011\dots 111),$$
$$(000\dots 000),\ \text{or}$$
$$(111\dots 111).$$

The cycle has $2k+1$ vertices and net-length one.

Theorem 3.4.4. *The digraph* $D_1 \rightharpoonup D_2$ *does not have a* WNUF_k *for any* $k \geq 2$.

Proof. We assume that $f : (D_1 \rightharpoonup D_2)^k \to (D_1 \rightharpoonup D_2)$ is a WNUF_k and derive a contradiction.

The vertices $(000\dots 011\dots 111)$, $(000\dots 000)$ and $(111\dots 111)$ are in-neighbours of the vertices $(aaa\dots aaa)$ and $(bbb\dots bbb)$.

Since $f(aaa\dots aaa) = a$ and $f(bbb\dots bbb) = b$, we see that $(000\dots 011\dots 111)$, $(000\dots 000)$ and $(111\dots 111)$ all map into $V(D_1)$ under f.

Furthermore $(111\dots 1a2\dots 222)$ is adjacent to $(bbb\dots bbb)$. Therefore

$$f(111\dots 1a2\dots 222) \in \{a\} \cup V(D_1).$$

If $f(111\dots 1a2\dots 222) = a$, then since $(222\dots 2b3\dots 333)$ and $(ccc\dots cbc\dots ccc)$ are both out-neighbours of $(111\dots 1a2\dots 222)$, we see that $f(222\dots 2b3\dots 333) = b$ and $f(ccc\dots cbc\dots ccc) = b$. There is also an arc from $(222\dots 2b3\dots 333)$ to $(bcc\dots ccc\dots ccc)$. By near-unanimity the last arc also has to map to b, which is not possible. This shows that $f(111\dots 1a2\dots 222) \in V(D_1)$.

All of the above now implies that the oriented cycle, 3.3, maps to D_1 which is not possible. ∎

Corollary 3.4.5. *Let C_k be a directed cycle of length k and P_l be a directed path of length l. The following digraphs do not have a* WNUF_k *for any $k \geq 2$.*

- ❒ $(C_{k_1} \rightharpoonup C_{k_2})$,
- ❒ $(C_k \rightharpoonup P_l)$ *as long as* $l \geq 2$,
- ❒ $(P_{l_1} \rightharpoonup P_{l_2})$ *as long as* $l_1 \geq 3$ *and* $l_2 \geq 2$,
- ❒ $(P_l \rightharpoonup C_k)$ *as long as* $l \geq 3$.

It is worth noting that more arcs may be added between the digraphs D_1 and D_2 above. The list of arcs between D_1 and D_2, given above, represents a minimum that is needed to prove the non-existence of a WNUF.

3.5 Indicators and Sub-indicators

In this section we develop the WNUF analogs to the indicator construction (Lemma 1.4.1) and the vertex (and arc) sub-indicator construction (Lemmas 1.4.2 and 1.4.3).

These results have a very interesting implication. Suppose that HOM_H is shown to be NP-complete for a family of digraphs H. If the proof of this follows from direct NP-completeness reductions for the bases cases $B_1, B_2, \ldots, B_t$ and by treating the remaining cases through a combination of the indicator and vertex (arc) sub-indicators, then there exists a proof that the family of digraphs H do not admit a WNUF_k for any $k > 1$, provided that one can show the base cases $B_1, B_2, \ldots, B_t$ do not admit a WNUF_k for any $k > 1$. In fact, the two proofs will look exactly the same once the bases cases have been dealt with. The reason for this is the correspondence between the constructions in both cases: "NP-complete" can be replaced by "no WNUF."

3.5.1 Indicator Construction

Let I be a fixed digraph with two specified vertices i and j. The *indicator construction* (with respect to the indicator I, i, j) transforms a digraph H into the digraph H^* defined as follows. $V(H^*) = V(H)$ and arcs xy if there exists a homomorphism $f : I \to H$ such that $f(i) = x$ and $f(j) = y$.

Lemma 3.5.1. *If H^* does not admit a* WNUF_k *of any arity $k > 1$, then H does not admit a* WNUF *of any arity greater than one.*

Proof. Let I, H and H^* be defined as above. We will assume that H has a WNUF and derive a contradiction.

Assume that $f : H^k \to H$ is a WNUF_k. Since H and H^* have the same vertex set, f is also a mapping of the vertices of $(H^*)^k$ to the vertices of H^*. In fact, we will show that f is also a WNUF_k of H^*. To do this all we have to show is that f is a homomorphism from $(H^*)^k$ to H^* since all other WNUF properties are inherited from $f : H^k \to H$.

Let $(x_1, x_2, \ldots, x_k) \to (y_1, y_2, \ldots, x_k)$ be an arc of $(H^*)^k$. Then by the definition of H^* there exist k homomorphisms $g_t : I \to H$ such that $g_t(i) = x_t$ and $g_t(j) = y_t$, $1 \leq t \leq k$. Define $H' = g_1(I) \times g_2(I) \times \cdots \times g_k(I) \subseteq H^k$.

Define a homomorphism $h : I \to H'$ by $h(w) = (g_1(w), g_2(w), \ldots, g_k(w))$. The composition $f \circ h : I \to H$ is a homomorphism from I to H. We now note that $f(x_1, x_2, \ldots, x_k) = f \circ h(i)$ and that $f(y_1, y_2, \ldots, y_k) = f \circ h(j)$. By the definition of the arcs of H^*, $f(x_1, x_2, \ldots, x_k) \to f(y_1, y_2, \ldots, y_k)$ is an arc of H^*. ∎

3.5.2 Vertex Sub-indicator

Let J be a digraph with distinguished vertices u and j. The sub-indicator construction with respect to J, u, and j transforms a digraph H with a distinguished vertex v into the induced subgraph H^+ of H defined as follows: $V(H^+) = \{x \in$

$V(H) \mid$ there exists $f : J \to H$ with $f(u) = v$ and $f(j) = x\}$.

Lemma 3.5.2. *Let H be a digraph. If H^+ does not have a* WNUF$_k$ *for $k > 1$, then H does not have a* WNUF *of arity greater than one.*

Proof. The proof is by contradiction. Assume that $g : H^k \to H$ is a WNUF$_k$ and that H^+ does not have a WNUF. We show that $g|_{H^+} = g' : (H^+)^k \to H^+$ is in fact a WNUF — a contradiction.

Let $(x_1, x_2, \ldots, x_k) \in (V(H^+))^k$. Then, there exist homomorphisms $f_i : J \to H$ such that $f_i(u) = v$ and $f_i(j) = x_i$ for $1 \leq i \leq k$. Set $H' = f_1(J) \times f_2(J) \times \cdots \times f_k(J)$. Define $h : V(J) \to V(H')$ by $h(w) = (f_1(w), f_2(w), \ldots, f_k(w))$ for each $w \in V(J)$. Then it is easy to see that h is a homomorphism $J \to H'$. That is, H' contains a homomorphic image of J.

Now $g'(x_1, x_2, \ldots, x_k) = g'(f_1(j), f_2(j), \ldots, f_k(j)) = g' \circ h(j)$ and $g' \circ h(u) = g'(f_1(u), f_2(u), \ldots, f_k(u)) = g'(v, v, \ldots, v) = v$. Therefore, $(x_1, x_2, \ldots, x_k)$ is the homomorphic image of vertex j under the homomorphism $g' \circ h : J \to H$ where $g' \circ h(u) = v$. By the definition of H^+, $g'(x_1, x_2, \ldots, x_k) \in V(H^+)$. That is, g' maps into H^+.

Since g' inherits all of g's WNUF properties, g' is a WNUF$_k$ on H^+. ∎

3.5.3 Arc Sub-indicator

Let J be a fixed digraph with a distinguished arc jj' and with a distinguished vertex u. The arc-subindicator construction transforms a given digraph H with a distinguished vertex v into its subgraph H^- induced by the images of the arc jj' under homomorphisms $f : J \to H$, where $f(u) = v$.

Lemma 3.5.3. *If H^- does not admit a* WNUF$_k$ *of arity $k > 1$, then H does not admit a* WNUF *of any arity greater than one.*

Proof. Let H, H^- and J be given as above. We assume that H has a WNUF, and then show that this implies that H^- also has a WNUF_k — a contradiction.

Assume that $f : H^k \to H$ is a WNUF_k. We will show that the restriction of f to (H^-) (suitably modified) defines a WNUF_k on (H^-). The first concern is that some vertices of $(H^-)^k$ may be isolated. Choose a vertex of H^-, say z, and use z as the image of all isolated k-tuples in $(H^-)^k$. This partial mapping is a homomorphism from $(H^-)^k$ to H^- (since no arcs are involved yet). If an isolated tuple is nearly unanimous, any tuple formed by rearranging its individual vertices will also be isolated, and so they will all be mapped in the same way (to z). No unanimous tuple $(x, x, \ldots, x)$ in $(H^-)^k$ is isolated since no x in (H^-) is an isolated vertex (by the definition of H^-).

We now extend the partial homomorphism to all of $(H^-)^k$. By the above we may assume that $(H^-)^k$ does not have isolated vertices. Let $(x_1, x_2, \ldots, x_k) \in (H^-)^k$. Since $(x_1, x_2, \ldots, x_k)$ is not isolated we assume without loss of generality that there is another k-tuple $(y_1, y_2, \ldots, y_k) \in (H^-)^k$ such that $(x_1, x_2, \ldots, x_k) \to (y_1, y_2, \ldots, y_k)$ is an arc of $(H^-)^k$. For these non-isolated vertices we restrict the WNUF_k $f : H^k \to H$ to $(H^-)^k$. We only have to show that $f|_{(H^-)^k}$ actually maps into H^-, the other WNUF properties are inherited from f.

Since $(x_1, x_2, \ldots, x_k) \to (y_1, y_2, \ldots, y_k)$ is an arc of $(H^-)^k$, $x_i \to y_i$ is an arc of H^-, $1 \le i \le k$. Furthermore, there exist homomorphisms $g_i : J \to H$ such that $g_i(j) = x_i$, $g_i(j') = y_i$ and $g_i(u) = v$, $1 \le i \le k$. Let $H' = g_1(J) \times g_2(J) \times \cdots \times g_k(J) \subseteq H^k$. Define a homomorphism $h : J \to H'$ by $h(w) = (g_1(w), g_2(w), \ldots, g_k(w))$. Then $f(x_1, x_2, \ldots, x_k) = f \circ h(j)$, $f(y_1, y_2, \ldots, y_k) = f \circ h(j')$ and $f \circ h(u) = f(g_1(u), g_2(u), \ldots, g_k(u)) = f(v, v, \ldots, v) = v$. Therefore the arc $(x_1, x_2, \ldots, x_k) \to (y_1, y_2, \ldots, y_k)$ is the image of the arc jj' under the homomorphism $f \circ h : J \to H$ such that $f \circ h(u) = v$. By the definition of H^-, $(x_1, x_2, \ldots, x_k)$ is mapped into H^-. ∎

3.6 Tournaments and WNUFs

The goal of this section is to prove that WNUFs follow the dichotomy laid out for tournaments by Bang-Jensen, Hell and MacGillivray in [5] exactly. Recall that the dichotomy says that tournaments with at most one cycle define a polynomial problem, while tournaments with at least two cycles define an NP-complete problem. We will show that tournaments with at most one cycle all have a conservative WNUF$_3$. No tournament with at least two cycles admits a WNUF$_k$, $k > 1$.

Furthermore this section is an illustration of the correspondence between NP-completeness proofs and no WNUF proofs alluded to at the start of Section 3.5. All of the results here are derived without assuming P $\neq$ NP.

3.6.1 Acyclic Tournaments and WNUFs

Since acyclic tournaments have an $\underline{X}$-enumeration (the acyclic order of its vertices $\{v_1, v_2, \ldots, v_n\}$ where v_i is adjacent to v_j if and only if $i < j$) it has a conservative WNUF$_3$ by Proposition 3.3.2.

3.6.2 Unicyclic Tournaments and WNUFs

We start by showing that most unicyclic tournaments do not have a NUF$_k$ for any $k \geq 3$ even although the corresponding homomorphism problem is polynomial time solvable. On the other hand they all have a conservative WNUF$_3$. This is further evidence that WNUFs are "the right way to go."

Let T be a unicyclic tournament. Then the cycle in T has length three. It is well known that T may be constructed by starting with a 3-cycle and then adding sources/sinks recursively.

Theorem 3.6.1. *Let T be a unicyclic tournament in which there are at least two vertices that dominate the cycle or in which there are at least two vertices dominated* *by the cycle. Then T does not have a* NUF_k *for any* $k \geq 3$*.*

Proof. We only focus on the first case: there are at least two vertices dominating the cycle. The second case is the converse of the first.

Label the vertices on the cycle with $\{0, 1, 2\}$. If there are vertices dominating the cycle, let t and t' be the first two source vertices that were added in forming T. Therefore they both dominate the cycle, t' is adjacent to t and there are no other vertices "between" t and t' and the cycle. Assume that $f : T^k \to T$ is a NUF_k.

Then $f(ttt \ldots t) = t$ and $f(ttt \ldots tt't \ldots t) = t$. Furthermore $(ttt \ldots t)$ is adjacent to $(x_1x_2x_3 \ldots x_k)$ where $x_j \in \{0, 1, 2\}$. Therefore $f(x_1x_2x_3 \ldots x_k) \in N^+(t)$. If there are no vertices dominated by the cycle, this would imply that $f(x_1x_2x_3 \ldots x_k) \in \{0, 1, 2\}$. If, on the other hand, there is a vertex dominated by the cycle, let b be the first sink vertex that was added in forming T. Then $f(bbb \ldots b) = b$ and $(x_1x_2x_3 \ldots x_k)$ is adjacent to $(bbb \ldots b)$. Therefore $f(x_1x_2x_3 \ldots x_k)$ has to be an out-neighbour of t and an in-neighbour of b implying that again $f(x_1x_2x_3 \ldots x_k) \in \{0, 1, 2\}$.

In the same way, it may be shown that $f(x_1x_2x_3 \ldots x_itx_{i+1} \ldots x_k) \in \{0, 1, 2\}$ since $(ttt \ldots tt't \ldots t)$ is adjacent to $(x_1x_2x_3 \ldots x_itx_{i+1} \ldots x_k)$ which in turn is adjacent to $(bbb \ldots b)$ (if b exists).

The above implies that the oriented cycle, 3.1 (which also exists in T^k), maps to the 3-cycle in T. This is not possible since the oriented cycle has net-length one. ∎

Corollary 3.6.2. *Let T be a unicyclic tournament. Then T has a conservative* WNUF_3*.*

Proof. Let T be a unicyclic tournament on n vertices in which its unique 3-cycle is being dominated by all other vertices in T. Then there exists a transitive tournament, H_1, on $n-2$ vertices and a directed 3-cycle, H_2, such that T=graft(H_1,H_2). Therefore T has a conservative WNUF_3, say f, by Corollary 3.3.8.

If T' is a unicyclic tournament in which its unique 3-cycle is dominating all other vertices in T', then T' is the converse of T above. In this case f from above is a conservative WNUF_3 for T'.

Lastly, let T'' be a unicyclic tournament in which its 3-cycle is dominating and is dominated by other vertices in T''. Let $A \subseteq V(T'')$ be the vertices dominating the 3-cycle, $B \subseteq V(T'')$ be the vertices dominated by the 3-cycle and C the vertices on the 3-cycle. Then T'' =graft(G_1,G_2) where G_1 is a transitive tournament on $|A| + 1$ vertices and G_2 is equal to the sub-tournament of T'' induced by the vertices in $B \cup C$. G_1 has an $\underline{X}$-enumeration and G_2 has conservative WNUF_3 (it resembles the T's from above). Corollary 3.3.8 now implies that T'' has a conservative WNUF_3. ∎

3.6.3 Tournaments With At Least Two Cycles and WNUFs

In this section we show that no tournament with at least two cycles has a WNUF_k for any $k \geq 3$.

We start by showing that the strong tournament on four vertices shown below in Figure 3.1 has no WNUF of any arity $k \geq 3$.

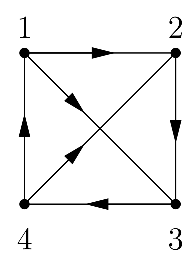

Figure 3.1: The strong tournament on four vertices T_4.

Lemma 3.6.3. *The strong tournament on four vertices does not have a* WNUF_k *for any* $k \geq 3$.

Proof. Label the vertices of the tournament as shown in Figure 3.1.

Let $f : T_4^k \to T$ be a WNUF_k. We define a family of subgraphs, G_t, of T_4^k as shown in Figure 3.2. The parameter identifying different members of the family is the number, t, of 2s in the tuple $(44\ldots422\ldots2)$

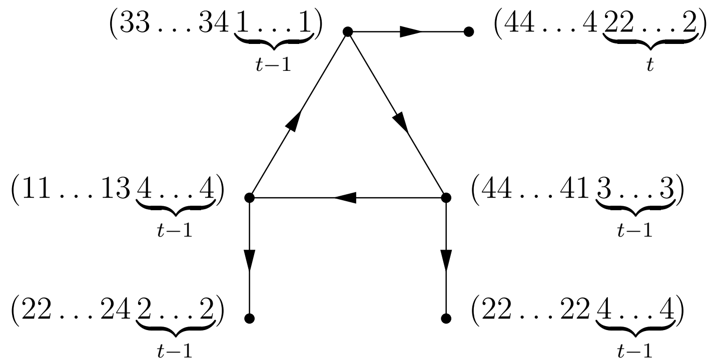

Figure 3.2: The family of subgraphs G_t of T_4^k.

Consider the subgraph G_2 of T_4^k. The vertices on the 3-cycle map to one of the three cycles in T_4. Furthermore since f is a WNUF, $(22\ldots242)$ and $(22\ldots224)$ map to the same vertex in T_4. Among the two 3-cycles in T_4 (234 and 134) only one of these is such that two consecutive vertices on it have a common out-neighbour, namely 134.

This implies that

$$\begin{aligned} f(11\ldots134) &= 1, \quad f(44\ldots413) = 4, \\ f(33\ldots341) &= 3, \quad f(22\ldots242) = 2, \\ f(22\ldots224) &= 2, \quad f(44\ldots422) = 4. \end{aligned} \tag{3.4}$$

Since f is a WNUF, any tuple with $k-1$ 2s and exactly one 4 has to map to 2. In particular $f(22\ldots24\underbrace{2\ldots2}_{t-1}) = 2$

Next consider the family of 4-cycles, H_j, in T_4^k shown in Figure 3.3. This family is parametrized by the number, j, of 2s in the tuple $(44\ldots422\ldots2)$.

By (3.4), $f(44\ldots422) = 4$. If we apply this to H_2 and note that the 4-cycle maps to the 4-cycle in T_4, it follows that $f(22\ldots244) = 2$. Using this on G_3, together with the fact that (as before) there is only one 3-cycle in T_4 such that two consecutive

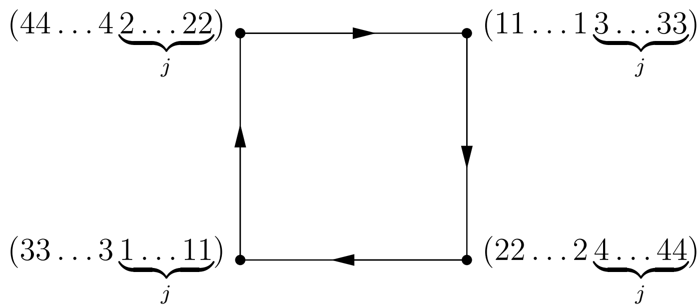

Figure 3.3: The family of 4-cycles H_j in T_4^k.

vertices have a common out-neighbour, we see that $f(44\ldots4222) = 4$. This can now be used on H_3 to find that $f(22\ldots2444) = 2$, which in G_4 will imply that $f(44\ldots42222) = 4$. Alternating in this way between H_i and G_{i+1} we conclude (when we reach G_{k-1}) that $f(422\ldots22) = 4$. This contradicts $f(22\ldots224) = 2$ from (3.4), since f is weakly nearly unanimous.

Therefore f does not exist. ∎

By Corollary 3.4.5 and the remark following it, the six-vertex tournament defined as two 3-cycles, with one cycle completely dominating the other, does not have a WNUF.

We are now ready to show that every tournament with at least two directed cycles does not possess a WNUF_k for any $k \geq 3$. The proof is carried out be examining a minimal counterexample and showing that it cannot exist. Our proof follows (almost exactly) the proof for tournaments in [5].

Theorem 3.6.4. *If T is a tournament with at least two cycles, then T does not possess a* WNUF_k *for any $k \geq 3$.*

Proof. Assume that the theorem is false. That is, there exists a tournament with at least two cycles that has a WNUF. Among all such tournaments, let T be one with the minimum number of vertices and let $f : T^k \to T$ be a WNUF_k. No proper induced sub-tournament of T, with at least two cycles, has a WNUF_t for any $t \geq 3$. By Lemma 3.6.3, we may assume that T has at least five vertices.

1. For every vertex v in T, $N^{+}(v)$, $N^{+2}(v)$, $N^{-}(v)$ and $N^{-2}(v)$ induce sub-tournaments with at most one cycle.

   Proof. If $N^{+}(v)$ ($N^{+2}(v)$, $N^{-}(v)$, $N^{-2}(v)$) contains at least two cycles, we could apply the sub-indicator from Lemma 3.5.2 with J equal to the path $u \to j$ ($u \to a \to j$, $u \leftarrow j$, $u \leftarrow a \leftarrow j$). The proper sub-tournament (none of these paths ever reach v), T^{+}, would then contain at least two cycles, and T^{+} would not possess a WNUF forcing T not have a WNUF — a contradiction. ∎

2. Consider the case where T is not strong. Let $T_1, T_2, \ldots, T_k$, $k \geq 2$ be the acyclic ordering of the strong components of T. By point number 1, $|T_i| \leq 3$ for $1 \leq i \leq k$, since a strong tournament is pancyclic. By Corollary 3.4.5 (and the remark following it), $k \geq 3$. Point number 1 now implies that $|T_i| = 1$ for $2 \leq i \leq k-1$. Therefore T_1 and T_k are 3-cycles and none of the intermediate vertices lies on a cycle. (At this point our proof differs slightly from that in [5].) Denote by T' the proper sub-tournament induced by $T_1 \cup T_2$. Let $g = f|_{T'}$. We claim that g is in fact a WNUF_k on T', which is impossible since T' is a proper sub-tournament. All that we have to show is that g maps into T', all other properties are inherited from f. Let $(x_1, x_2, \ldots, x_k) \in V(T')$. Then

   $$(x_1, x_2, \ldots, x_k) \to (x_1^{+}, x_2^{+}, \ldots, x_k^{+}) \to (x_1^{++}, x_2^{++}, \ldots, x_k^{++}) \to (x_1, x_2, \ldots, x_k),$$

   is a 3-cycle in $(T')^k$. Since T_1 and T_k are exactly the 3-cycles of T, g maps $(x_1, x_2, \ldots, x_k)$ into T'.

   This completes the case where T is not strong. We may therefore assume that T is strong.

3. We now show that T is one of the exceptional tournaments shown in Figure 3.4.

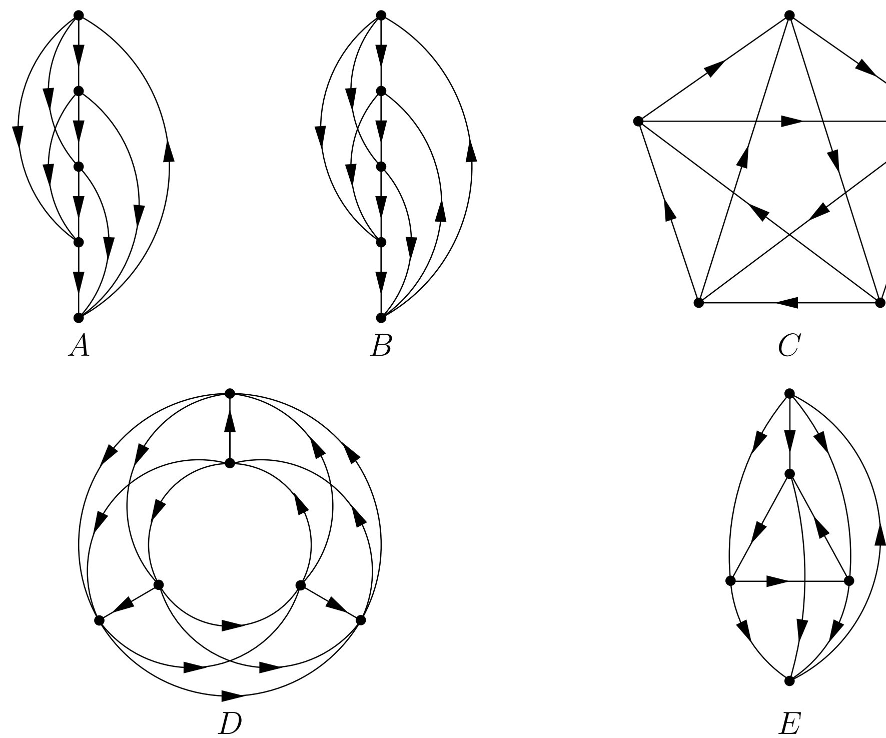

Figure 3.4: The exceptional tournaments.

❒ Consider first the case where T is not two-connected. Let x be a vertex such that $T-x$ is not strongly connected and let $T_1, T_2, \ldots, T_k$, $k \geq 2$, be the acyclic ordering of the strong components of $T-x$. As in point number 2, $|T_i| \leq 3$ for $i = 1, 2, \ldots, k$. The strong connectivity of T implies that x is dominated by some vertex in T_k and that x dominates some vertex in T_1.

- Suppose now that $|T_i| = |T_j| = 3$ for some $i < j$. By the same method as in point 2, we find that $|T_1| = |T_k| = 3$ and $|T_i| = 1$ for $2 \leq i \leq k-1$. Denote the vertices of T_1 (in their cyclic order) as a, b, c and the vertices of T_k (in their cyclic order) as u, v, w. Without of loss of generality, $u \rightarrow x \rightarrow a$. Hence $\{x, u, v, w\} \subseteq N^{+2}(a)$. Point number 1 now implies that $v, w \rightarrow x$ and that $x \rightarrow b, c$. Hence $\{x, a, b, c, w\} \subseteq N^{+2}(u)$. These vertices induce more than one directed cycle, contradicting point 1.

Therefore at most one of the strong components of $T - x$ is a 3-cycle.

- Suppose that T_1 is a 3-cycle, abc, and let $T_k = \{y\}$. As before, $x \to a, b, c$ since $\{x, a, b, c\} \subseteq N^{-2}(y)$. If $|T| = 5$, then T is E in Figure 3.4. Therefore we may assume that $|T| \geq 6$, so that $k \geq 3$. Let $T_2 = \{u\}$. Then $\{c, x, y, u\} \subseteq N^{+2}(a)$ and these vertices induce more than one cycle, contradicting point 1. The case where $|T_k| = 3$ and $|T_1| = 1$ is similar and also leads to E in Figure 3.4.
- Suppose that T_i is a 3-cycle, abc, for some $2 \leq i \leq k-1$. Let $T_1 = \{u\}$ and $T_k = \{y\}$. By the strong connectivity of T, $y \to x \to u$. Since $\{x, a, b, c\} \subseteq N^{-2}(y)$, either $x \to a, b, c$ or $a, b, c \to x$. In the first instance, $\{a, x, y, u\} \subseteq N^{-2}(c)$ and these vertices induce two directed cycles. In the second instance $\{c, x, y, u\} \subseteq N^{+2}(a)$ and these vertices also induce two directed cycles. Neither of these cases is possible by point number 1.

Therefore $T - x$ is acyclic. Let $T_i = \{v_i\}$, $1 \leq i \leq k$, as before $v_k \to x \to v_1$ since T is strong.

- Suppose $|T| \geq 6$. If $v_2 \to x$, then $\{v_1, v_{k-1}, v_k, x\} \subseteq N^{+2}(v_2)$ and these vertices would induce two directed cycles. Therefore $x \to v_2$. If $v_3 \to x$, then $\{v_1, v_2, v_k, x\} \subseteq N^{+2}(v_3)$ and these vertices would induce two directed cycles. So, $x \to v_3$. Now, $\{v_3, v_4, \ldots, v_k, x\} \subseteq N^{+2}(v_1)$. These vertices induce at least two directed cycles.

Therefore $|T| = 5$. If $x \to v_2, v_3$, then T is A in Figure 3.4, where x is the top vertex of A. If $v_2, v_3 \to x$, then T is also A in Figure 3.4, with x the bottom vertex of A. If $v_3 \to x \to v_2$, then T is B in Figure 3.4, where x is the bottom vertex of B. It is impossible for $v_2 \to x \to v_3$, otherwise

$\{x, v_1, v_3, v_4\} \subseteq N^{+2}(v_2)$ and these vertices induce two cycles.

❒ We now consider the case where T is 2-connected. If $|T| = 5$, then T is C in the list as C is the only 2-connected tournament on five vertices. Therefore we may assume that $|T| \geq 6$. Let v be a vertex of maximum out-degree in T and write $O = N^+(v)$ and $I = N^-(v)$. By examining the average out-degree of T, we see that $|O| \geq 3$. Also, $I \subseteq N^{+2}(v)$ since v is a king. Since T is 2-connected, every vertex has at least two in-neighbours and at least two out-neighbours. Therefore $|I| \geq 2$. Let $O_1, O_2, \ldots, O_q$ and $I_1, I_2, \ldots, I_k$ be the acyclic ordering of the strong components of $T[O]$ and $T[I]$ respectively.

- If $|O_1| \geq 3$. Then $N^{+2}(v) = T - v$ which is strongly connected and so has at least two cycles.

Therefore $|O_1| = 1$. Let $O_1 = \{o_1\}$.

- Suppose that there exists a vertex $u \in I_1$ and a vertex $r \in O_q$ such that $u \to r$. (Here our proof also differs from that in [5].) Then $I \subseteq N^{-2}(r)$. Furthermore since v is a king, u has an in-neighbour from O, say x. Therefore $x \in N^{-2}(r)$. Also, $v \to o_1 \to r$ since o_1 and r cannot be in the same strong component of O. If $|I_1| \geq 3$ choose w to be in $N^+(u) \cap I_1$. Otherwise $|I_1| = 1$ and u dominates all of $I - u$; in this case choose w to be any vertex in I different from u. Therefore, $\{u, w, x, v\} \subseteq N^{-2}(r)$ and these vertices induce two directed cycles.

Thus O_q completely dominates I_1. Let $r \in O_q$ and $u \in I_1$. Note that $N^{+2}(v) = T - \{v, o_1\}$ (use o_1 to get every vertex in $O - o_1$ and use that fact that v is a king). Therefore $T - \{v, o_1\}$ cannot be a strong tournament as this would force at least two directed cycles since $|T - \{v, o_1\}| \geq 4$. This in turn implies that O_2 completely dominates I_k, otherwise we can find a

Hamilton cycle in $T - \{v, o_1\}$.

- Suppose that $o_1 \to z$ for some $z \in I$. Then $\{r, u, v\} \cup I_k \subseteq N^{+2}(o_1)$. Since o_1 and r are not in the same strong component, $o_1 \to r \to u$, $o_1 \to z \to v$ and O_2 dominates I_k (as well as $|I| \geq 2$), respectively. These vertices induce at least two directed cycles.

Therefore I completely dominates o_1.

- Suppose $|O| \geq 4$. We claim that there exists a vertex x in O_2 such that x is joined to r by a directed path of length two. If $q = 2$, then $r \in O_2$ and $|O_2| \geq 3$. Since O_2 is strong, its easy to find the vertex x by using a Hamilton cycle. Otherwise $q \geq 3$. If $q > 3$, there exists a strong component between O_2 and O_q and x can then be any vertex in O_2. Hence $q = 3$. Since $|O| \geq 4$, there either are at least two vertices in O_2 and so an arc in O_2 and x is easy to find, or at least two vertices in O_q. O_q is strong and so r has an in-neighbour in O_q and x is any vertex in O_2.

  - If x does not have a directed path of length two to u, then x and r are both in O_2, $r \to x$ and $u \to f$ where $x \to f \to r$ is an (x, r)-path from above. In this case $\{x, f, r, u\} \subseteq N^{+2}(v)$ (use o_1 to get x, f, r and $v \to r \to u$ to get u). These vertices induce two cycles.

  Therefore there is a directed path of length two from x to u. Now since $x \in O_2$ dominates I_k, and I dominates o_1, there is a path of length two to r from above. We also have one to u, and x dominates I_k which are all in-neighbours of v. Thus $\{o_1, r, u, v\} \subseteq N^{+2}(x)$. These vertices induce two cycles.

Therefore $|O| = 3$.

- If $|I| \geq 4$, there exists a vertex a in I with at least two out-neighbours in I. This vertex also has o_1 and v as out-neighbours. Hence a has at least four out-neighbours, which is impossible since v is a vertex of maximum out-degree and we have now determined that v has degree three.

Thus $|I| \leq 3$. Also, since O is not strong (because of O_1) we have that $T[O]$ is a transitive triple. Write $O_i = \{o_i\}$ for $i = 1, 2, 3$. Hence $|I_k| = 1$, otherwise the out-degree of o_2 is too big (since it dominates I_k). This implies that I is also transitive. Write $I_j = \{i_j\}$ for $1 \leq j \leq k$.

- If $k = 3$, then $\{i_1, v, o_1, o_3\} \subseteq N^{-2}(i_3)$ ($i_1 \to i_2 \to i_3$, $v \to o_2 \to i_3$, $o_1 \to o_2 \to i_3$, $o_3 \to i_1 \to i_3$) and these vertices have two directed cycles.

Therefore $k = 2$, as $k \neq 1$ since $|T| \geq 6$. Then $o_3 \to i_2$, so that $d^+(o_3) \geq 2$, by the 2-connectivity of T. Also, $o_2 \to i_1$, as i_1 does not have any other out-neighbours because of a maximum out-degree of three. This now means that T is in fact D in Figure 3.4.

The tournament E does not have a WNUF by Corollary 3.4.2. Tournaments A through D can be handled by an appropriate sub-indicator; see Figure 3.5.

For tournaments B, C and D, let $J = P_3$ (a directed path of length 3) with u the initial vertex and j the terminal vertex of the path. Map u to the white vertices in Figure 3.5. In the case of tournament A use $J = P_5$ (directed path of length 5) with u the initial vertex and j the terminal vertex. The images of j in all cases are indicated by the check marks (✓) in Figure 3.5. In each case the resulting tournament T^+ is the strong tournament on four vertices. Hence T does not have a WNUF — a contradiction. ∎

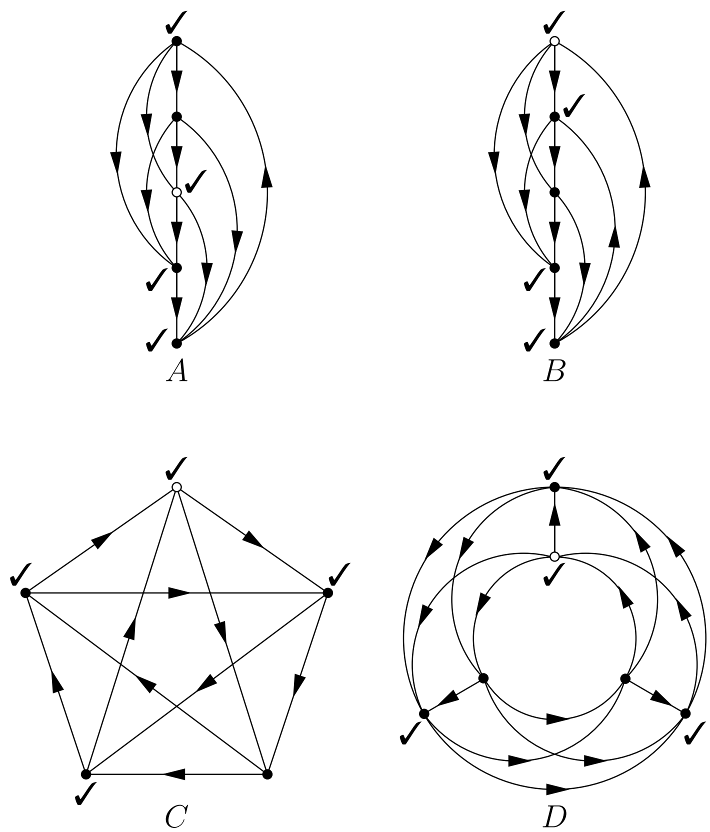

Figure 3.5: The exceptional tournaments, A–D, and their sub-indicators.

4

Injective Homomorphisms

4.1 Introduction

A homomorphism $f : G \to H$ is *injective* if the restriction of f to $N^{-}(v)$ is an injective mapping, for each $v \in V(G)$. Alternatively one may think of each vertex in $N^{-}(v)$ as receiving a different "colour" in H. The existence of an injective homomorphism between G and H is denoted by $G \xrightarrow{\text{inj}} H$.

The complexity results in this chapter are derived through the repeated application of an indicator construction. This construction is similar in spirit to the one in Lemma 1.4.1, except that it is tailored to injective homomorphism problems.

We classify all reflexive injective homomorphism problems (where the target graph H is reflexive — H has a loop at every vertex). Irreflexive targets (no loops) pose more of a problem. The difficulty lies with targets where the maximum in-degree is equal to two. All targets with in-degree at least three are NP-complete and those with in-degree one are polynomial. The in-degree two case exhibits both polynomial problems and NP-complete problems. We also describe two transformations, the one transforming any (ordinary) H-colouring problem into an injective, in-degree two problem and the other performing a reverse transformation of the above (unfortunately they are not inverses).

Finally we discuss some upper bounds on related colouring problems.

4.2 Injective Indicator Construction

H is a digraph and is used as the target in the injective H-colouring problem. H may be reflexive or irreflexive.

Let I be a digraph with two distinguished vertices i and j such that $N^-(i) = N^-(j) = \emptyset$, $N^+(i) \neq \emptyset$ and $N^+(j) \neq \emptyset$.

We define a new digraph H^* with the same vertex set as H and with arcs given by $xy \in A(H^*)$ if and only of there exists an injective homomorphism $f : I \stackrel{\text{inj}}{\longrightarrow} H$ such that $f(i) = x$ and $f(j) = y$.

Lemma 4.2.1. *H^*-colouring polynomially reduces to injective H-colouring.*

Proof. Let G be an instance of the H^*-colouring problem. Form a new digraph ${}^*G$ by replacing each arc xy in G by a copy of I, identifying x with i and y with j.

Then $G \rightarrow H^*$ if and only if ${}^*G \stackrel{\text{inj}}{\longrightarrow} H$:

$\Rightarrow$:

Let $f : G \rightarrow H^*$ be a homomorphism. If $xy \in A(G)$, then $f(x)f(y) \in A(H^*)$. Therefore there exists and injective homomorphism $f_{xy} : I \stackrel{\text{inj}}{\longrightarrow} H$ such that $f_{xy}(x) = f(x)$ and $f_{xy}(y) = f(y)$. In this way f may be extended to an injective homomorphism ${}^*G \stackrel{\text{inj}}{\longrightarrow} H$. This is possible since the vertices of G in ${}^*G$ only have out-neighbours.

$\Leftarrow$:

Let $g :{}^*G \stackrel{\text{inj}}{\longrightarrow} H$ be an injective homomorphism. If $xy \in A(G)$, let I_{xy} be the copy of I that was used to replace xy in forming ${}^*G$. Then $g|_{I_{xy}} : I_{xy} \stackrel{\text{inj}}{\longrightarrow} H$ is an injective homomorphism from I_{xy} to H. Therefore $g|_{I_{xy}}(x)g|_{I_{xy}}(y) \in A(H^*)$ by the definition of H^*. We can therefore find a homomorphism $G \to H^*$ by restricting g to $V(G)$. ∎

Note that if I is chosen such that $f : I \xrightarrow{\text{inj}} H$ is an injective homomorphism with $f(i) = x$ and $f(j) = y$ implies that there also exists an injective homomorphism $g : I \xrightarrow{\text{inj}} H$ with $g(i) = y$ and $g(j) = x$, for all $x, y \in V(H)$, then H^* above may be viewed as an undirected graph. Furthermore if the input to the H^*-colouring problem contains an undirected edge, it is sufficient to replace this edge with one copy of I using any orientation of I (i.e., i can be identified with x or y and j with y or x, respectively).

4.3 Reflexive Targets

Lemma 4.3.1. *If H is a reflexive digraph and $\delta^-(H) \geq 3$ (H has a vertex with at least two distinct in-neighbours, other than itself), then injective H-colouring is NP-complete.*

Proof. Let H be a reflexive digraph with $\delta^-(H) \geq 3$ and let a be a vertex of H with at least two distinct in-neighbours, say b and c. For this digraph we use the indicator I with vertices $\{i, x, j\}$ and arcs ix and jx. If uv is an arc of H, then the following mappings are injective homomorphisms from I to H: $i \mapsto u$, $x \mapsto v$, $j \mapsto v$ and $i \mapsto v$, $x \mapsto v$, $j \mapsto u$. This produces the undirected edge uv in H^*. In particular, ab and ac are undirected edges. Furthermore the edge bc is in H^* since $i \mapsto b$, $x \mapsto a$, $j \mapsto c$ and $i \mapsto c$, $x \mapsto a$, $j \mapsto b$ are injective homomorphisms from I to H. Therefore, H^* contains the triangle abc and so is not bipartite. Since H^*-colouring is NP-complete, injective H-colouring is NP-complete. ∎

Lemma 4.3.2. *If P_n is a directed, reflexive path of length $n \geq 2$, then injective P_n-colouring is NP-complete.*

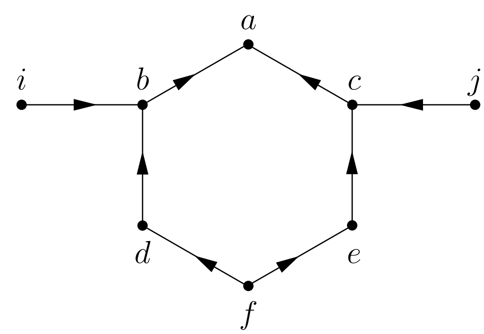

Figure 4.1: The indicator I for directed, reflexive paths.

Proof. To prove this result we use the indicator I shown below in Figure 4.1.

Note that I has an automorphism that exchanges i and j. Hence the digraph obtained by applying I to P_n has undirected edges.

Let P_n be a directed, reflexive path of length $n \geq 2$. Let $f : I \overset{\text{inj}}{\longrightarrow} P_n$ be an injective homomorphism. Note that $f(b) \neq f(c)$, $f(d) \neq f(i)$ and $f(e) \neq f(j)$. Also, $|f(d) - f(e)| = 0, 1$. If x is a vertex of P_n we denote by x^- and x^+ the predecessor and successor (respectively) of x, when they exist. We now observe that we cannot have $f(i) = f(j) = x$, if this was the case then $d \mapsto x^-, x^+$ and $e \mapsto x^-, x^+$. Since $|f(d) - f(e)| = 0, 1$ we must have $f(d) = f(e) = x^+$ or $f(d) = f(e) = x^-$. This forces $f(b) = f(c) = x^+$ or $f(b) = f(c) = x$ (respectively) which is not possible. Therefore there are no loops when we apply this indicator.

As an example we have the injective homomorphisms from I to the first three vertices of P_n (labeled 0, 1 and 2 respectively) shown in Table 4.1.

Keeping in mind that we obtain undirected edges in P_n^*, we see that P_n^* contains an undirected triangle on the vertices 0, 1, and 2 and so is not bipartite. Therefore P_n^*-colouring is NP-complete implying that injective P_n-colouring is NP-complete. ∎

Lemma 4.3.3. *If C is a reflexive, directed cycle of length at least three, then injective C-colouring is NP-complete.*

Proof. Let C be a directed, reflexive n-cycle and $x \in V(C)$. Denote by x^- and x^+

Table 4.1: Some injective homomorphisms from I to the first three vertices of P_n.

		a					2					2		
$\boldsymbol{i}$	b		c	$\boldsymbol{j}$	**0**	1		2	**2**	**0**	1		2	**1**
	d		e			1		1			1		2	
		f					1					1		
		a					2					2		
$\boldsymbol{i}$	b		c	$\boldsymbol{j}$	**1**	1		2	**2**	**1**	1		2	**1**
	d		e			0		1			0		2	
		f					0					$\nexists$		

the predecessor and successor (respectively) of x on C. Here we use the indicator I shown below in Figure 4.2.

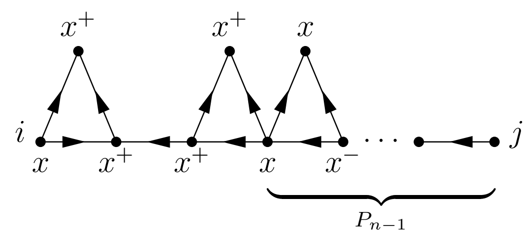

Figure 4.2: The indicator for directed, reflexive cycles.

Also shown in the figure is a partial injective homomorphism where the vertex i (of I) has been mapped to the vertex x in C. Note that vertex j can be mapped to any vertex on C except the vertex x. This is accomplished by using the loops and the directed path from x^- to x^+ on C. It is impossible to map j to x since the length of the path in I from j to the vertex mapped to x^- is only $n-1$. The result of this indicator, C^*, is a complete graph on n vertices and so the NP-completeness of injective C-colouring follows from the NP-completeness of C^*-colouring. ∎

Lemma 4.3.4. *If H is a reflexive, directed cycle of length two or H is a reflexive, directed path of length one, then injective H-colouring is polynomial.*

Proof. Let be a H reflexive, directed cycle of length two: $V(H) = \{0, 1\}$ and $A(H) = \{00, 11, 01, 10\}$. We show that injective H-colouring is polynomial by reducing the problem to an instance of 2-SAT.

Let G be the input digraph to the injective H-colouring problem. If $\Delta^-(G) \geq 3$, then G does not map injectively to H. We therefore restrict our attention to inputs with $\Delta^-(G) \leq 2$.

For each vertex v of G there are two variables: v_0 and v_1 (v_0 is true iff $v \mapsto 0$ and v_1 is true iff $v \mapsto 1$). Since v may only map to exactly one of 0 and 1, we form the clause $v_0 \veebar v_1 = (v_0 \vee v_1) \wedge (\overline{v_0} \vee \overline{v_1})$ for each vertex $v \in V(G)$. For each vertex $x \in V(G)$ with $N^-(x) = \{u, v\}$, $u \neq v$ add the clauses: $u_0 \veebar v_0$ and $u_1 \veebar v_1$. It is now a simple matter to check that $G \xrightarrow{\text{inj}} H$ if and only if there exists a satisfying truth assignment.

We now deal with the case where H is a reflexive directed path of length one, with $V(H) = \{0, 1\}$ and $A(H) = \{00, 11, 01\}$. This is handled in exactly the same way as the reflexive 2-cycle except that we add one more set of clauses: for each arc uv in the input digraph G we form the clause $u_1 \rightarrow v_1$. Again it's easy to check that $G \xrightarrow{\text{inj}} H$ if and only if there exists a satisfying truth assignment. ∎

Theorem 4.3.5. *If H is a loop, a reflexive arc or a reflexive 2-cycle, then injective H-colouring is polynomial. Otherwise injective H-colouring is NP-complete.*

Proof. The cases where H is a reflexive arc or a reflexive 2-cycle are covered by Lemma 4.3.4. If H is a loop, then the input to the problem has in-degree at most one in order to have an injective homomorphism to H. In this case, all vertices can be coloured with the same colour.

Let H be a reflexive digraph that is not a loop, a reflexive arc or a reflexive 2-cycle. In addition, we may assume that H is a core with respect to injective homomorphisms. The proof that injective H-colouring is NP-complete in this case proceeds according to the maximum in-degree of H. Note that $\Delta^-(H) \neq 1$ as this would imply that H is a single loop.

❐ $\Delta^-(H) \geq 3$. This case is covered by Lemma 4.3.1.

❒ $\Delta^{-}(H) = 2$.

- We first assume that H is acyclic (there are no directed cycles of length greater than one; loops are not considered to be cycles). Then there exists a vertex v in H with no in-neighbours other than v itself. The vertex v has to have an out-neighbour other than v itself, otherwise the component containing v is just a single loop and since $\Delta^{-}(H) = 2$, H will not be a core. If v has more than one out-neighbour different from v, then H is a reflexive, oriented tree rooted at v with all arcs directed away from v. In this case, H is not a core. Therefore v has exactly one out-neighbour, say $N^{+}[v] = \{v, w\}$. The same idea may be applied to w to conclude that w has exactly one out-neighbour. Continuing in this way find that the component containing v is a reflexive, directed path. In fact, if H has more than one component, then each component of H is a reflexive, directed path. In this case though, H is not a core. Therefore, H has exactly one component and this component is a reflexive, directed path of length at least two. By Lemma 4.3.2, H-colouring is NP-complete.

- If H contains a cycle (of length at least two), this cycle is an induced cycle since $\Delta^{-}(H) = 2$. Furthermore, every component of H contains a cycle, because an acyclic component is a reflexive, directed path which can easily be mapped injectively to a directed cycle in H, implying that H is not a core. Consider any component of H and let v be a vertex on a cycle in this component. The vertex v only has it predecessor on the cycle as an in-neighbour (different from itself). As for out-neighbours of v (different from v), v obviously has its successor on the cycle as an out-neighbour. If it had more than one out-neighbour different from v, these out-neighbours are not on the cycle. As in the acyclic case above we can form an oriented tree rooted at v with all arcs directed away from v. This tree is equivalent to a reflexive, directed path that can be mapped injectively to the cycle in the component under consideration. This would show that H is not a core. Therefore each component of H is a reflexive, directed cycle and these cycles have relatively prime lengths (in order for H to be a core).

Let ℓ be the length of a longest cycle in H.

If $\ell = 2$, then H is a reflexive, directed 2-cycle. Therefore $\ell \geq 3$. Let t be the length of a shortest cycle, not equal to two, in H. We now show that there is a polynomial reduction from injective C_t-colouring to injective H-colouring. Since injective C_t-colouring is NP-complete for $t \geq 3$ by Lemma 4.3.3, we will have that injective H-colouring is NP-complete.

Let G be an instance of injective C_t-colouring. We form an instance G' of injective H-colouring as shown Figure 4.3 below.

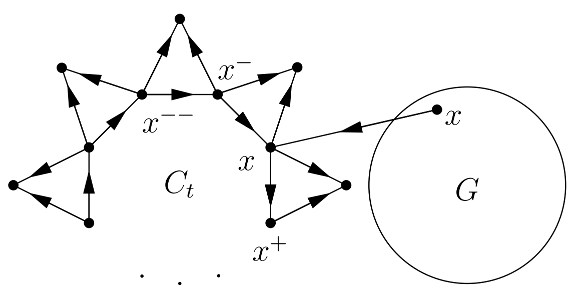

Figure 4.3: A reduction from injective C_t-colouring.

If t is even, then H does not contain a 2-cycle. This follows from the fact that an even cycle would map injectively to a two cycle showing that H is not a core. If $G \xrightarrow{\text{inj}} C_t$, then its easy to see that $G' \xrightarrow{\text{inj}} H$ as shown in Figure 4.3. Conversely if $G' \xrightarrow{\text{inj}} H$, then the copy of C_t in G' can only map to the copy of C_t in H: the C_t in G' has to map onto one of the cycles in H because the ends of each arc are also common in-neighbours of a vertex of in-degree two. Furthermore this C_t cannot map to a smaller cycle in H (if t is even there is no 2-cycle to map to and if t is odd it cannot map to a 2-cycle if one is present in H). This forces the copy of G in G' to map injectively to C_t as well.

∎

In the special case of reflexive, injective, oriented colourings, the target is always a reflexive tournament. By Lemma 4.3.4 it is possible to decide in polynomial time whether the reflexive, injective oriented chromatic number of an input digraph is at most two.

Corollary 4.3.6. *It is an NP-complete problem to decide whether the reflexive, injective oriented chromatic number of an input digraph is at most k for $k \geq 3$.*

Proof. Let T be a reflexive tournament on $n \geq 3$ vertices. If $n = 3$, then T is either a reflexive, directed cycle or T is a reflexive, transitive tournament. By Lemmas 4.3.3 and 4.3.1 respectively, injective T-colouring is NP-complete. Therefore we may assume that $n \geq 4$. The sum of the in-degrees of T is equal to the number of arcs which is $n(n-1)/2$ plus n (for the loops). Therefore the average in-degree of T is $(n-1)/2+1 = (n+1)/2 \geq 5/2$. Since T has a vertex with in-degree greater than or equal to its average in-degree, T has a vertex of in-degree at least 3. So, by Lemma 4.3.1 injective T-colouring is NP-complete. ∎

4.4 Irreflexive Targets

Lemma 4.4.1. *If the irreflexive digraph H has $\Delta^-(H) \geq 3$, then injective H-colouring is NP-complete.*

Proof. This follows in almost exactly the same way as for the reflexive case. Let a be a vertex with at least three distinct in-neighbours b, c, d. By applying the indicator I with vertices $\{i, x, j\}$ and arcs ix and jx to H we obtain an undirected 3-cycle: bcd.

Since the undirected portion of H^* is not bipartite, H^*-colouring is NP-complete, implying that injective H-colouring is NP-complete. ∎

We note at this point that if the irreflexive digraph H has $\Delta^-(H) = 1$, then H is a union of paths and cycles. The core of H (under injective homomorphisms) is either a union of directed cycles (of relative prime lengths) or a directed path. In each of these cases the input has a maximum in-degree of at most 1 in order for it to map to H. In fact an injective homomorphism is exactly an ordinary homomorphism in this case. Therefore the injective H-colouring problem is polynomial when $\Delta^-(H) = 1$.

Let $\mathscr{D}$ be the set of all irreflexive digraphs. We now define a new set of irreflexive digraphs with maximum in-degree two, denoted by $\mathscr{I}$. If $H \in \mathscr{D}$, define $\widehat{H} \in \mathscr{I}$ as follows: replace each arc uv of H with an oriented path on four vertices P_{uv} given by $V(P_{uv}) = \{u, x_{uv}, y_{uv}, v\}$ and arcs ux_{uv}, $x_{uv}y_{uv}$, vy_{uv}. This replacement operation is shown in Figure 4.4.

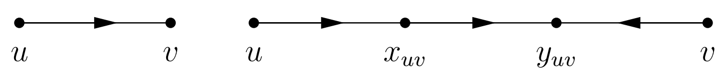

Figure 4.4: The replacement operation.

Lemma 4.4.2. *If $H \in \mathscr{D}$, then H-colouring polynomially transforms to injective $\widehat{H}$-colouring.*

Proof. We need to prove the following. If $G, H \in \mathscr{D}$, then $G \to H$ if and only if $\widehat{G} \xrightarrow{\text{inj}} \widehat{H}$.

Let $f : G \to H$ be a homomorphism. If uv is an arc of G, then $f(u)f(v)$ is an arc of H. There exists an oriented path on the vertices u, x_{uv}, y_{uv}, v together with the arcs ux_{uv}, $x_{uv}y_{uv}$ and vy_{uv} in $\widehat{G}$. Similarly, in $\widehat{H}$ there is an oriented path on the vertices $f(u), x_{f(u)f(v)}, y_{f(u)f(v)}, f(v)$ together with the arcs $f(u)x_{f(u)f(v)}$, $x_{f(u)f(v)}y_{f(u)f(v)}$ and

$f(v)y_{f(u)f(v)}$. We define an injective homomorphism $\widehat{f} : \widehat{G} \xrightarrow{\text{inj}} \widehat{H}$ as follows:

$$\begin{aligned} \widehat{f}(u) &= f(u), \\ \widehat{f}(v) &= f(v), \\ \widehat{f}(x_{uv}) &= x_{f(u)f(v)}, \\ \widehat{f}(y_{uv}) &= y_{f(u)f(v)}. \end{aligned}$$

Conversely, let $\widehat{f} : \widehat{G} \xrightarrow{\text{inj}} \widehat{H}$ be an injective homomorphism. Note that $V(G) \subset V(\widehat{G})$. We define a homomorphism $f : G \to H$ by setting $f(u) = \widehat{f}(u)$ for all $u \in V(G)$. Let uv be an arc of G and P_{uv} the corresponding oriented path in $\widehat{G}$. P_{uv} is a core under injective homomorphisms and so P_{uv} maps onto a similar oriented path in $\widehat{H}$ under $\widehat{f}$, say $P_{\widehat{f}(u)\widehat{f}(v)}$. The existence of $P_{\widehat{f}(u)\widehat{f}(v)}$ in $\widehat{H}$ implies that $\widehat{f}(u)\widehat{f}(v) = f(u)f(v)$ is an arc of H. ∎

Corollary 4.4.3. *If $H \in \mathscr{D}$ and H-colouring is NP-complete, then injective $\widehat{H}$-colouring is NP-complete. On the other hand if injective $\widehat{H}$-colouring is polynomial, then H-colouring is polynomial.*

This corollary in particular implies that there are oriented trees with in-degree two for which the corresponding injective colouring problem is NP-complete. The problems in $\mathscr{I}$ therefore display a richness comparable to that of the problems in $\mathscr{D}$.

We now show that there is also a polynomial-time reduction from injective colouring problems in $\mathscr{I}$ to colouring problems in $\mathscr{D}$. In order to describe this reduction we perform the following replacement on a digraph $H \in \mathscr{I}$: for each $x \in V(H)$ let $x_1x_2x_3$ be a 3-cycle on three new vertices, identify x with x_1. Next, if y is a vertex of in-degree two in H and $N^-(y) = \{x, z\}$, add two new vertices u_y and v_y as well as the arcs xu_y, u_yz, zv_y and v_yx. This replacement is shown in Figure 4.5.

If $H \in \mathscr{I}$ we denote the result of this replacement by $H^\circ \in \mathscr{D}$. Note that $V(H) \subseteq V(H^\circ)$ and that each of the original vertices of H is now on a 3-cycle in H°.

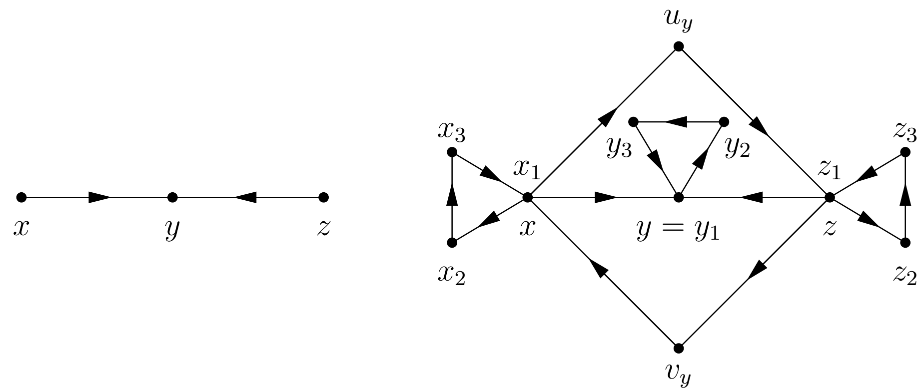

Figure 4.5: The second replacement operation.

Let G be an instance of injective H-colouring, where $H \in \mathscr{I}$. If G contains a directed cycle, or a directed path of length at least three or a vertex of indegree at least three, then G is a NO instance of injective H-colouring. The reason for this is that H is acyclic, the length of a longest path in H is at most two and that $\Delta^-(H) = 2$. Furthermore if $x, z \in N^-(y)$, then x and z are not adjacent, since H does not contain any transitive triples.

We may therefore assume that G is acyclic, that the length of a longest path in G is at most two, that $\Delta^-(G) \leq 2$ and that any common in-neighbours of a vertex are independent.

The transformed instance G° of H°-colouring is obtained in the same way as H°: attach 3-cycles to each vertex of G and if y is a vertex of in-degree two in G and $N^-(y) = \{x, z\}$, add two new vertices u_y and v_y as well as the arcs xu_y, u_yz, zv_y and v_yx.

Lemma 4.4.4. *If $H \in \mathscr{I}$, then injective H-colouring polynomially transforms to H°-colouring.*

Proof. Let $f : G \xrightarrow{\text{inj}} H$ be an injective homomorphism. We now extend f to a homomorphism $g : G^\circ \to H^\circ$. Note that $V(G) \subseteq V(G^\circ)$. If y is a vertex of indegree two in G, then $f(y)$ is a vertex of indegree two in H, by injectivity.

$$g(x) = \begin{cases} f(x) & \text{if } x \in V(G), \\ v_{f(y)} & \text{if } x = v_y \text{ for some } y \in V(G), \\ u_{f(y)} & \text{if } x = u_y \text{ for some } y \in V(G), \\ f(z)_i & \text{if } x = z_i \text{ for some } z \in V(G) \text{ and } i = 2, 3. \end{cases}$$

It is easy to check that g is a homomorphism from G° to H°.

Suppose that $f : G^\circ \to H^\circ$ is a homomorphism. We show that $g = f|_{V(G)}$ is an injective homomorphism, $g : G \xrightarrow{\text{inj}} H$. Let $x, y \in V(G) \subseteq V(G^\circ)$. Then x and y are both on 3-cycles. The homomorphism, f, maps x and y in such a way that $f(x)$ and $f(y)$ are also on 3-cycles. The only 3-cycles in H° are the ones incident with vertices of H (recall that H itself is acyclic). Therefore g maps into $V(H)$. If xy is an arc of G, then xy is also an arc of G°, therefore $f(x)f(y)$ is an arc of H. Let y be a vertex of indegree two in G, with $N^-(y) = \{x, z\}$. Then x and z are on the 4-cycle xu_yzv_y. Since there are no 2-cycles in H°, f must map x and z to distinct vertices. Therefore g is an injective homomorphism from G to H. ∎

Corollary 4.4.5. *If $H \in \mathscr{I}$ and injective H-colouring is NP-complete, then H°-colouring is NP-complete. On the other hand if H°-colouring is polynomial, then injective H-colouring is polynomial.*

4.4.1 The Complexity of Injective Homomorphisms to Irreflexive Tournaments

Lemma 4.4.6. *If T is a tournament on at least 5 vertices, then injective T-colouring is NP-complete.*

Proof. The average in-degree of a tournament on n vertices is $(n-1)/2$. If T is a tournament on $n \geq 6$ vertices, the average in-degree is at least $5/2$. This implies that T has a vertex of in-degree at least 3. We now apply the indicator with vertices

i, x, j and arcs ix and jx to T. The result, T^*, will contain an undirected 3-cycle. Therefore T^*-colouring is NP-complete and so injective T-colouring is NP-complete.

If T is a tournament on 5 vertices then T has a vertex of in-degree at least 3, except in one case: the quadratic residue tournament on 5 vertices, shown below in Figure 4.6.

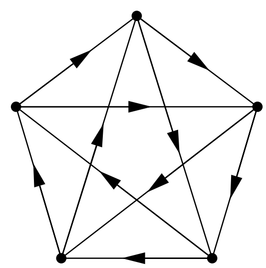

Figure 4.6: The quadratic residue tournament on 5 vertices.

On the other hand this tournament has the property that every vertex has in-degree equal to 2. If we apply the same indicator as above, the result is an undirected 5-cycle. Again, this leads to injective T-colouring being NP-complete. ∎

We now examine the remaining tournaments individually.

The strong tournament on four vertices, T_4, is shown below in Figure 4.7.

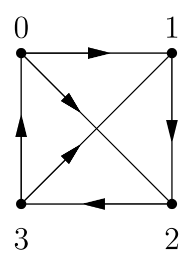

Figure 4.7: The strong tournament on four vertices T_4.

Lemma 4.4.7. *Injective T_4-colouring is NP-complete.*

Proof. The proof is via a reduction from 3-SAT.

There are two 3-cycles in T_4: 123 and 023. There are two transitive triples in T_4: 301 and 012.

Consider the digraph G shown in Figure 4.8.

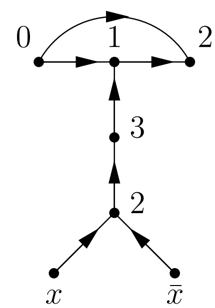

Figure 4.8: The variable gadget, G_x, for variable x.

In an injective homomorphism from G to T_4, the vertices map as shown. In addition to this, either $x \mapsto 0$ and $\overline{x} \mapsto 1$, or $x \mapsto 1$ and $\overline{x} \mapsto 0$. Consider also a directed path of length 4, $abcde$. If we pre-colour vertex e with the vertices $0, 1, 2, 3$ of T_4, we find the injective homomorphisms from the path to T_4 shown in Table 4.2

Table 4.2: Injective homomorphisms from $P_4 = abcde$ to T_4.

e	d	c	b	a
0	3	2	$0, 1$	$0, 3$
1	$0, 3$	$2, 3$	$0, 1, 2$	$0, 1, 3$
2	$0, 1$	$0, 3$	$2, 3$	$0, 1, 2$
3	2	$0, 1$	$0, 3$	$2, 3$

Next, consider the digraph K shown below in Figure 4.9. It consists of a directed 3-cycle, with paths of length 4 attached to each vertex of the 3-cycle and directed away from the 3-cycle. To this we add three copies of the oriented path with vertices x, y, z and arcs xy zy, each attached to a path of length 4.

Note if ℓ_1, ℓ_2 and ℓ_3 are all pre-coloured 1, this does not extend to an injective homomorphism to T_4: pre-colouring in this way forces one to use a 0 on the terminal vertices of each P_4, which in turn forces one to use only 0 and 3 on the 3-cycle. On

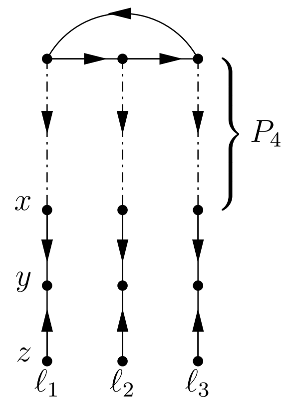

Figure 4.9: The clause gadget, K_{ℓ_1,ℓ_2,ℓ_3}, corresponding to the clause $\{\ell_1, \ell_2, \ell_3\}$.

the other hand using only $\{0, 1\}$ on ℓ_1, ℓ_2 and ℓ_3 with the requirement that there be at least one 0, extends to an injective homomorphism to T_4.

Given an instance of 3-SAT we construct the following digraph D: for each variable x, there is a corresponding variable gadget G_x. For each clause $\{\ell_1, \ell_2, \ell_3\}$ (where each literal ℓ_i is either equal to a variable x or its negation $\overline{x}$) there is a corresponding clause gadget K_{ℓ_1,ℓ_2,ℓ_3}. D is formed by identifying each variable vertex x (or $\overline{x}$) with the corresponding variable vertex in each clause gadget in which it occurs.

If there is a satisfying truth assignment, we can form an injective homomorphism from D to T_4: whenever a variable is assigned the value "True", map the corresponding vertex to vertex 0 of T_4, if its assigned the value "False" map the corresponding vertex to vertex 1 of T_4. This pre-colouring extends to an injective homomorphism, as shown above.

On the other hand if there is an injective homomorphism from D to T_4, we can form a satisfying truth assignment. The vertices x and $\overline{x}$ can only map to 0 and 1 respectively (or 1 and 0 respectively). As shown above the clause gadget ensures that at least one literal vertex per clause has to map to 0. Now by assigning the value "True" to x whenever the variable vertex x has been mapped to 0 and "False" when it maps to 1, we have a satisfying truth assignment. ∎

Lemma 4.4.8. *If T is a transitive tournament on 4 vertices or T is equal to a 3-cycle with a sink vertex, then injective T-colouring is NP-complete.*

Proof. In each case there is a vertex of in-degree 3. We can therefore proceed as above. ∎

The only remaining four-vertex case is a 3-cycle that is dominated by a source vertex. We will prove that this problem is polynomial by proving a slightly stronger result. First we prove the following lemma.

Lemma 4.4.9. *Let T_0 be the digraph with vertices t and c, and arcs tc and cc, then injective T_0-colouring is polynomial.*

Proof. The digraph T_0 is shown below.

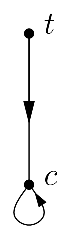

Figure 4.10: The digraph T_0.

We prove that injective T_0-colouring is polynomial by showing how the problem may be reduced to 2-SAT.

Let D be an instance of T_0-colouring. With each vertex of D we associate a variable v_t. The interpretation of v_t is that v_t is "True" if and only if there is an injective homomorphism from D to T_0 that maps v to t. For each arc xy in D we have the following clauses:

$$\begin{aligned} \overline{y_t} &= \overline{y_t} \vee \overline{y_t}, \\ x_t \rightarrow \overline{y_t} &= \overline{x_t} \vee \overline{y_t}, \\ \overline{x_t} \rightarrow \overline{y_t} &= x_t \vee \overline{y_t}. \end{aligned}$$

If u is a vertex of in-degree at least two in D, then for each $\{x, y\} \subseteq N^-(u)$ we have the following clause: $x_t \veebar y_t = (x_t \vee y_t) \wedge (\overline{x_t} \vee \overline{y_t})$. It is now easy to see that $D \xrightarrow{\text{inj}} T_0$ if and only if there is a satisfying truth assignment. ∎

Lemma 4.4.10. *Let H be the digraph formed by taking a copy of an n-cycle, C_n, and adding to this a new vertex, t, as well as all arcs from t to the n-cycle. Then injective H-colouring is polynomial.*

Proof. The digraph H is shown below.

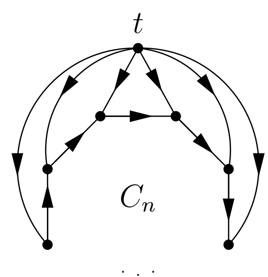

Figure 4.11: The digraph H.

Let D be an instance of injective H-colouring. Note first of all that if $D \xrightarrow{\text{inj}} H$, then $D \xrightarrow{\text{inj}} T_0$. We therefore start by testing whether D has an injective homomorphism to T_0 and if the result of this is negative, we may declare that D does not have an injective homomorphism to H either.

Assume now that $f : D \xrightarrow{\text{inj}} T_o$ is an injective homomorphism. Let $D_c = \{v \in V(D) \mid f(v) = c\}$, $D' = D\langle D_c \rangle$ and $D_t = \{v \in V(D) \mid f(v) = t\}$. We now test whether $D' \to C_n$. We only need to consider ordinary homomorphisms here since $\Delta^-(D') \leq 1$. In fact $D \xrightarrow{\text{inj}} H$ if and only if $D' \to C_n$: let $g : D' \to C_n$ be a homomorphism. Then the map h given by $h(v) = g(v)$ if $v \in D_c$ and $h(v) = t$ if $v \in D_t$ is an injective homomorphism. Conversely, assume that $D \xrightarrow{\text{inj}} H$, but that $D' \not\to C_n$. Then there exists a closed walk in D' of net-length not congruent to zero (mod n). Since $\Delta^-(D') = 1$ this closed walk is in fact a directed cycle. Thus D does not have an injective homomorphism to H, contrary to the assumption.

Therefore, once the test $D' \to C_n$ has been performed, we may declare that $D \stackrel{\text{inj}}{\longrightarrow} H$ if the test succeeds or declare that D does not have an injective homomorphism to H if it fails. ∎

Corollary 4.4.11. *If T is the four-vertex tournament consisting of a 3-cycle dominated by a source, then injective T-colouring is polynomial.*

There are only two three-vertex tournaments: the transitive triple and the 3-cycle.

Lemma 4.4.12. *If T is a three-vertex tournament, then injective T-colouring is polynomial.*

Proof. Let T be a directed 3-cycle and D an instance of injective T-colouring. We start by checking the maximum in-degree of D. If it exceeds one we may declare that D does not have an injective homomorphism to T. If $\Delta^-(D) \leq 1$, the sought after injective homomorphism is in fact an ordinary homomorphism and we may proceed by testing whether $D \to C_3$.

Let T be the transitive triple, $V(T) = \{0, 1, 2\}$ and $A(D) = \{01, 12, 02\}$, and D an instance of injective T-colouring. We show that injective T-colouring is polynomial by providing a reduction to 2-SAT.

To each $u \in V(D)$ we associate three variables: u_0, u_1 and u_2. We interpret these to mean that u_i is "True" if and only if there exists an injective homomorphism from D to T in which u maps to i for $i = 0, 1, 2$. For each arc $xy \in A(D)$ we form the following clauses:

$$\begin{aligned}
\overline{x_2} &= \overline{x_2} \vee \overline{x_2}, \\
\overline{y_0} &= \overline{y_0} \vee \overline{y_0}, \\
x_1 \to y_2 &= \overline{x_1} \vee y_2, \\
y_1 \veebar y_2 &= (y_1 \vee y_2) \wedge (\overline{y_1} \vee \overline{y_2}), \\
x_0 \veebar x_1 &= (x_0 \vee x_1) \wedge (\overline{x_0} \vee \overline{x_1}).
\end{aligned}$$

For each $u \in V(D)$ with $d^-(u) \geq 2$ and for every $\{x, y\} \subseteq N^-(u)$ we form the following clauses:

$$\begin{aligned} u_2 &= u_2 \vee u_2, \\ x_0 \veebar y_0 &= (x_0 \vee y_0) \wedge (\overline{x_0} \vee \overline{y_0}). \end{aligned}$$

If there exists an injective homomorphism $D \xrightarrow{\text{inj}} H$, then it is clear that there exists a satisfying truth assignment. Conversely if there is a satisfying truth assignment, we note first of all that for each $u \in V(D)$ exactly one of u_0, u_1 and u_2 is "True:" if u has an out-neighbour, then u_2 is "False" and the clause $u_0 \veebar u_1$ ensures that exactly one of u_0 and u_1 is "True." Similarly if u has an in-neighbour, u_0 is "False" and the clause $u_1 \veebar u_2$ ensures that exactly one of u_1 and u_2 is "True." If u has both an in-neighbour and an out-neighbour, then u_2 and u_0 are both "False" and the clauses $(u_0 \veebar u_1) \wedge (u_1 \veebar u_2)$ ensure that u_1 is "True." Therefore we may define a mapping $f : V(D) \to V(T)$ by $f(v) = i \in \{0, 1, 2\}$ if v_i is "True." It is now easy to see that f is in fact an injective homomorphism by the way the clauses were defined. ∎

The case where T is a one or two-vertex tournament is handled in a similar way to the 3-cycle case above: first check the in-degree of the input, then proceed by testing for the existence of an ordinary homomorphism.

In summary, we have the following result.

Theorem 4.4.13. *Let T be a tournament. Then the injective T-colouring problem is NP-complete, except when T is a tournament on at most three vertices or when T consists of 3-cycle dominated by a source.*

4.5 Bounds on the Injective Oriented Chromatic Number

Let D be an oriented graph, that is D does not have any 2-cycles. The reflexive (irreflexive) injective oriented chromatic number of D, $\chi_{io}^r(D)$ ($\chi_{io}^i(D)$), is defined to be the least k such that D has an injective homomorphism to some reflexive (irreflexive) tournament on k vertices. One may also think of the reflexive (irreflexive) injective oriented chromatic number as the least amount of colours that can be used in an improper (proper) injective oriented colouring of the vertices of D. In an oriented colouring of D, the arcs between two colour classes must all be oriented in the same direction. In an improper (proper) injective oriented colouring of D one adds the restriction that the colouring be injective on the in-neighbourhood of every vertex with the possibility that the same colour may (may not) appear on the ends of an arc.

Since injective colourings are defined in terms of in-neighbourhoods, one may suspect that there may be a Brooks-type theorem giving upper bounds on $\chi_{io}^r(D)$ and $\chi_{io}^i(D)$ in terms of $\Delta^-(D)$. This turns out to be false, as the next proposition shows. Note that $\Delta^-(D)$ ($\Delta^-(D)+1$) is always a lower bound for $\chi_{io}^r(D)$ ($\chi_{io}^i(D)$).

Proposition 4.5.1. *It is impossible find an upper bound on $\chi_{io}^r(D)$ and $\chi_{io}^i(D)$ in terms of $\Delta^-(D)$.*

Proof. To show that it is impossible to bound $\chi_{io}^r(D)$ and $\chi_{io}^i(D)$ in terms of $\Delta^-(D)$, we have to show that there exist oriented graphs D with a fixed $\Delta^-(D)$, but with arbitrarily high $\chi_{io}^r(D)$ and $\chi_{io}^i(D)$.

Let Δ^- be given and let $k \geq \Delta^-$ be a fixed positive integer.

Define two sets of vertices $L = \{v_1, v_2, \ldots, v_k\}$ and $R = \{u_1, u_2, \ldots, u_t\}$ where

$$t = \binom{k}{\Delta^-}.$$

For every Δ^--subset of L there is one vertex in R. For a fixed set of Δ^- elements of L, say $L_j = \{v_{j_1}, v_{j_2}, \ldots, v_{j_{\Delta^-}}\}$ join every member of L_j to its corresponding vertex u_j in R.

Every vertex in R has in-degree Δ^- and any two vertices in L are in-neighbours of some vertex in R. In colouring this oriented graph injectively, every vertex in L has to receive a different colour. Therefore we need at least k colours to colour this graph injectively (whether proper or improper). ∎

It turns out that one is able to bound the injective oriented chromatic number if one takes into account the maximum in-degree, maximum out-degree and the maximum degree of the underlying undirected graph. In general though, the bound has to be exponential as the following example shows.

Fix a positive integer Δ. We construct a digraph D as follows: take the disjoint union of all tournaments on Δ vertices, say $T_1 \cup T_2 \cup \cdots \cup T_m$. Now add m new vertices, say $x_1, x_2, \ldots, x_m$ and let T_i dominate x_i. At this point we have to find a tournament T so that $D \xrightarrow{\text{inj}} T$. Since the vertices of every T_i are the in-neighbours of x_i, we see that each T_i has to map injectively to T (whether T is reflexive or irreflexive). Therefore T has to contain each tournament on Δ vertices as a subgraph. Such a tournament is known as a Δ-*universal tournament* and lower bounds on its size are known [47]: $|V(T)| \geq 2^{(\Delta-1)/2}$. This shows that $\chi^r_{io}(D), \chi^i_{io}(D) \geq 2^{(\Delta-1)/2}$.

Next, we present an upper bound on $\chi^r_{io}(D)$ and $\chi^i_{io}(D)$. We treat the irreflexive case (different colours have to appear on the ends of an arc) first and then show how this can be modified to handle the reflexive case. The bound is derived by constructing a suitable undirected graph, colouring this undirected graph, using Brooks' Theorem to bound the chromatic number of this graph and then converting the (undirected) colouring into an irreflexive injective oriented colouring.

Theorem 4.5.2. *If D is an oriented graph, then*

$$\chi_{io}^{i}(D) \leq 2^{(\Delta+(\Delta^{-}-1)\Delta^{+}+(\Delta^{+})^{2}+1)} - 1,$$

where Δ is the maximum degree of the underlying graph of D, Δ^{-} is the maximum in-degree of D and Δ^{+} is the maximum out-degree of D.

Proof. Given an oriented graph D, we define an undirected graph, $D^{\bullet}$, that will capture the properties of an irreflexive injective oriented colouring that we are interested in:

- ❒ If $x, y \in V(D)$ and $d(x,y) = 2$, then x and y must receive different colours. This follows from the fact that our target is an irreflexive tournament and so does not contain any 2-cycles.

- ❒ If $x, y \in N^{-}(v)$ for some vertex v of D, then x and y must receive different colours.

- ❒ If x and y are adjacent in D, they must receive different colours.

The graph $D^{\bullet}$ is defined as follows: $V(D^{\bullet}) = V(D)$, $E(D^{\bullet}) = A^{*}(D) \cup \{xy \mid x, y \in V(D) \text{ and } d(x,y) = 2\} \cup \{xy \mid x, y \in N^{-}(v) \text{ for some } v \in V(D)\}$, where $A^{*}(D)$ denotes the edges of of the underlying graph of D.

Each vertex in $D^{\bullet}$ has at most Δ neighbours from the underlying graph of D, at most $(\Delta^{-}-1)\Delta^{+}$ neighbours that are derived from joining common in-neighbours by an edge and at most $(\Delta^{+})^{2}$ vertices that are at a distance two from it. Therefore the maximum degree of $D^{\bullet}$ is at most $\Delta + (\Delta^{-}-1)\Delta^{+} + (\Delta^{+})^{2}$. We now colour $D^{\bullet}$ and by Brooks' Theorem we need at most $\Delta + (\Delta^{-}-1)\Delta^{+} + (\Delta^{+})^{2} + 1$ colours.

Let the colour classes of this colouring be $C_1, C_2, \ldots, C_k$ where $k = \Delta + (\Delta^{-}-1)\Delta^{+} + (\Delta^{+})^{2} + 1$ and colour every vertex of D with the same colour it received in $D^{\bullet}$. This colouring of D is certainly irreflexive (because of $A^{*}(D)$) and it is also injective because common in-neighbours were joined by an edge. At this point it may not be an oriented colouring. We now recolour the vertices of D in order to ensure an oriented colouring.

Let the colour of $v \in V(D)$ be j. For any colour $i \neq j$, v has at most one in-neighbour in colour class i (the in-neighbours of v are adjacent in $D^{\bullet}$ and so receive different colours). Furthermore, v does not have both an in-neighbour and an out-neighbour in colour class i (vertices at distance two received different colours since they are adjacent). Therefore v either has exactly one in-neighbour in colour class i or only out-neighbours in colour class i. Define the *signature* of v to be a k-tuple where the jth entry is a "$\cdot$" and the ith entry is either "+" or "−" depending on whether v has out-neighbours or in-neighbours respectively in colour class i, $1 \leq i \leq k$, $i \neq j$.

We now use the information encoded in the signatures of each vertex to decide how to recolour it so that we have an oriented colouring. Since the vertices in C_1 are independent, we assign them all the same colour. Next, consider the possible arcs between C_1 and C_2: every vertex in C_2 either has an out-neighbour in C_1 (so its signature starts with $(+\ \cdot\ \ldots)$) or an in-neighbour in C_1 (their signatures start with $(-\ \cdot\ \ldots)$). These vertices have to receive different colours in an oriented colouring. If we now consider the arcs between C_3 and C_i, $i = 1, 2$, we see that the signatures start in one of four ways: $(-\ -\ \cdot\ \ldots)$, $(-\ +\ \cdot\ \ldots)$, $(+\ -\ \cdot\ \ldots)$ or $(+\ +\ \cdot\ \ldots)$. The vertices in C_3 with different signatures on the first two coordinates need different colours in an oriented colouring. In general, the vertices in colour class i can be partitioned into 2^{i-1} sets depending on their signatures on the first $i-1$ coordinates. As before, vertices in C_i in different blocks of the partition need different colours in an oriented colouring. In total we therefore need $1 + 2 + 2^2 + \cdots + 2^{k-1} = 2^k - 1$ colours to recolour the vertices of D. ∎

Theorem 4.5.3. *If D is an oriented graph, then*

$$\chi_{io}^r(D) \leq 2^{((\Delta^- - 1)\Delta^+ + (\Delta^+)^2 + 1)} - 1,$$

where Δ^- is the maximum in-degree of D and Δ^+ is the maximum out-degree of D.

Proof. The proof here is exactly the same as above except for $E(D^\bullet) = \{xy \mid x, y \in V(D)$ and $d(x, y) = 2\} \cup \{xy \mid x, y \in N^-(v)$ for some $v \in V(D)\}$. Here we don't need the edges of the underlying graph of D since the same colours may appear on the ends of an arc. ∎

4.6 Obstructions to Injective Oriented Reflexive Colouring

In the case of injective oriented reflexive colourings, there are three polynomial cases: a target digraph that is a loop, a reflexive arc or a reflexive 2-cycle. These correspond to $\chi_{io}^r \leq 2$.

In this section we are interested in studying obstructions to a digraph D having $\chi_{io}^r(D) \leq 2$. An *obstruction* is a digraph F such that $\chi_{io}^r(D) \neq k$ if and only if $F \xrightarrow{\text{inj}} D$, for $k = 1$ or 2. Note that if a digraph D has an injective homomorphism to the reflexive arc, then it has an injective homomorphism to the reflexive 2-cycle. So, for the purposes of studying injective oriented reflexive two-colourings, we may assume that the target (in this case) is the reflexive 2-cycle.

Let J_1 be the digraph with vertices $\{a, b, c\}$ and arcs ab and cb. J_1 is the obstruction to a digraph D having $\chi_{io}^r(D) = 1$.

Theorem 4.6.1. *A digraph D has $\chi_{io}^r(D) \neq 1$ if and only if $J_1 \xrightarrow{\text{inj}} D$.*

Proof. If $\Delta^-(D) \leq 1$, all vertices can be coloured with one colour. So if $\chi_{io}^r(D) \neq 1$, then D must have a vertex of in-degree at least two, and so $J_1 \xrightarrow{\text{inj}} D$.

On the other hand if $J_1 \xrightarrow{\text{inj}} D$, the image of J_1 in D needs at least two colours. So $\chi^r_{io}(D) \neq 1$. ∎

If a digraph D has $\Delta^-(D) \geq 3$, then $\chi^r_{io}(D) \neq 2$. In this case the digraph J_2 with $V(J_2) = \{a, b, c, d\}$ and $A(D) = \{ba, ca, da\}$ is an obstruction. It is possible for a digraph D to have $\Delta^-(D) = 2$, but $\chi^r_{io}(D) \neq 2$. We now define a family obstructions for this case. Let $\mathcal{C}$ be the family of all odd undirected cycles and define $\mathscr{C}$ to be the family of digraphs obtained from a graph in $\mathcal{C}$ be replacing each edge with a copy of J_1.

Theorem 4.6.2. *A digraph D has $\chi^r_{io}(D) > 2$ if and only if either $J_2 \xrightarrow{inj} D$ or there exists a digraph $C \in \mathscr{C}$ such that $C \xrightarrow{inj} D$.*

Proof. Since $\chi^r_{io}(J_2) = 3$ and since every $C \in \mathscr{C}$ has $\chi^r_{io}(C) = 3$ it follows immediately that if either $J_2 \xrightarrow{\text{inj}} D$ or $C \xrightarrow{\text{inj}} D$, then $\chi^r_{io}(D) > 2$.

Conversely let $\chi^r_{io}(D) > 2$. If $\Delta^-(D) \geq 3$, then $J_2 \xrightarrow{\text{inj}} D$. Therefore we may assume that $\Delta^-(D) \leq 2$. Define an auxiliary (undirected) graph $D^\bullet$ by

$$V(D^\bullet) = V(D),\ E(D^\bullet) = \{uv \mid u, v \in N^-(x) \text{ for some } x \in V(D)\}.$$

We now claim that $D^\bullet$ is 2-colourable (bipartite) if and only if $D \xrightarrow{\text{inj}} T$, where T is the reflexive 2-cycle: $V(T) = \{0, 1\}$, $A(T) = \{00, 01, 10, 11\}$. A 2-colouring of $D^\bullet$ (with the colours 0 and 1) is immediately an injective homomorphism $D \xrightarrow{\text{inj}} T$ and conversely if $D \xrightarrow{\text{inj}} T$, we have a 2-colouring of $D^\bullet$.

Since $\chi^r_{io}(D) > 2$, it must be the case that $D^\bullet$ is non-bipartite and so $D^\bullet$ contains an odd cycle. From the definition of $D^\bullet$ it is now clear that some member of $\mathscr{C}$ maps injectively to D. ∎

The 2-colouring of the auxiliary graph above provides another way, apart from the 2-SAT formulations for the reflexive arc and the reflexive 2-cycle, of showing that injective oriented reflexive 2-colouring is polynomial time solvable. If the input has $\Delta^-(D) > 2$, then D is a NO instance. Otherwise, form $D^\bullet$ and test for (undirected) 2-colourability. If it succeeds, we have an oriented reflexive 2-colouring, if it fails $\chi_{io}^r(D) > 2$. Of course, here we are assuming that the target is the reflexive 2-cycle.

4.7 Injective Oriented Colourings on Trees

Let T be an oriented tree. In this section we give an algorithm that tests for a fixed irreflexive (reflexive) digraph H, whether $T \xrightarrow{\text{inj}} H$ and finds an injective irreflexive (reflexive) homomorphism if one exists.

If T is an oriented tree rooted at a vertex x, we may arrange the vertices of T into its level sets based on the distance from x (in the underlying undirected tree T'): $V_i = \{v \mid d_{T'}(v, x) = i\}$, $0 \leq i \leq \ell$, where ℓ is the eccentricity of x. V_ℓ is considered to be the bottom of the tree and V_0 the top of the tree.

The algorithm first assigns lists, $L(v) \subseteq V(H)$, $v \in V(T)$, to the vertices of T. These lists are to be thought of as possible images (or colours) for the vertices of T. The assignment is based on the fact that if $f : T \xrightarrow{\text{inj}} H$ is an injective homomorphism, then $d_H^-(f(v)) \geq d_T^-(v)$ since the image of v has to accommodate the in-neighbours of v in H. The lists are then processed from the bottom-up and the eventual colouring is from the top-down.

Since the colouring will be from the top-down, the processing of the lists has to ensure that once a possible image for a vertex has been decided upon (i.e. we've made a choice from $L(v)$), this choice can be extended downwards. "Extending downwards" means that we can make choices for vertices lower down in the tree (from their lists) that will preserve arcs as well as respect injectivity on in-neighbours.

To preserve arcs we essentially do a one-sided consistency check. That is if $u \in V_i$ and $v \in V_{i+1}$ such that uv (vu) is an arc of T, we remove $a \in L(u)$ if there does not exist a $b \in L(v)$ such that ab (ba) is an arc of H. This is done for all $u \in V_i$, for i

from $\ell - 1$ to 0. In this way if we reach the root, x, and $L(x) \neq \emptyset$, we can make a choice for the image of x and extend all the way down preserving arcs in the process.

Let $A = \{v \in V(T) \mid |N^-(v)| \geq 2\}$. If $u \in A$, with $N^-(u) = \{u_1, u_2, \ldots, u_k\}$ and $y \in L(u)$, then there exists an injective mapping of $N^-(u)$ with $u \mapsto y$, if and only if there exists a system of distinct representatives (SDR) for the sets

$$L(u_1) \cap N^-(y),\ L(u_2) \cap N^-(y),\ \ldots,\ L(u_k) \cap N^-(y).$$

Therefore to process $L(u)$ we remove y from $L(u)$ if there does not exist an SDR as shown above. Note that the lists have to be intersected with $N^-(y)$ in order to "localize" them around y in H.

If $u \in A \cap V_i$, for some $0 \leq i \leq \ell - 1$, is such that $N^-(u) \subseteq V_{i+1}$, then the processing above is sufficient since a top-down colouring will colour u first (make a choice from $L(u)$) and if the SDR checks succeeded we will then be able to extend downwards (injectively). On the other hand it could happen that $N^-(v) \cap V_{i-1} \neq \emptyset$ for some $v \in A \cap V_i$. In this case $|N^-(v) \cap V_{i-1}| = 1$. Let u be the vertex in $|N^-(v) \cap V_{i-1}|$. Here u will be coloured (in a top down colouring) before any of $\{v\} \cup (N^-(v) - \{u\})$ have been coloured. For such a vertex u we have to ensure that a choice $a \in L(u)$ can be extended downwards. Of course, u may be an in-neighbour of more than one vertex in V_i. Let $A^* = \{v \in V(T) \mid v \in V_i \text{ and } N^+(v) \cap V_{i+1} \neq \emptyset,\ 0 \leq i \leq \ell - 1\}$. Suppose that $u \in A^* \cap V_i$. That is, u is an in-neighbour of some vertices $v_1, v_2, \ldots, v_k \in V_{i+1}$. We remove $a \in L(u)$ if there exists $v_j \in \{v_1, v_2, \ldots, v_k\} \cap A$, such that for every $z \in L(v_j)$ there does not exist an SDR for the sets

$$\{a\} \cap N^-(z),\ L(u_{j_1}) \cap N^-(z),\ L(u_{j_2}) \cap N^-(z),\ \ldots,\ L(u_{j_k}) \cap N^-(z),$$

where $(N^-(v_j) - \{u\}) = \{u_{j_1}, u_{j_2}, \ldots, u_{j_k}\}$. In essence, we are checking here whether $u \mapsto a$ can be extended to an injective mapping on all the in-neighbourhoods of vertices in $\{v_1, v_2, \ldots, v_k\} \cap A$ — we only need to consider vertices with more than one in-neighbour, hence the intersection with A.

The algorithm is shown below (Algorithm 4.1). The algorithm is for irreflexive, injective homomorphisms $T \xrightarrow{\text{inj}} H$. To obtain the corresponding algorithm for the reflexive problem, we change all intersections of the form $L(u_j) \cap N^-(y)$ to $L(u_j) \cap N^-[y]$ since we are now allowed to re-use the colour (y) of the "centre" vertex on the in-neighbours.

Algorithm 4.1 Irreflexive Injective Hom $T \xrightarrow{\text{inj}} H$

INPUT: An oriented tree T rooted at a vertex x.

Level sets $V_0 = \{x\}, V_1, \ldots, V_\ell$.

The sets A and A^* as defined as above.

A target digraph H.

TASK: Find an injective homomorphism $T \xrightarrow{\text{inj}} H$ if one exists.

ACTION: Assign lists to $V(T)$ as follows: $L(v) = \{y \in V(H) \mid d_H^-(y) \geq d_T^-(v)\}$.

For $i = \ell - 1$ to 0 perform the following for all $u \in V_i$.

For each arc uv (vu) with $v \in V_{i+1}$, remove $a \in L(u)$ if there does not exist a $b \in L(v)$ with ab (ba) an arc of H.

If $u \in A$, remove $y \in L(u)$ if there does not exist an SDR for the sets: $L(u_1) \cap N^-(y)$, $L(u_2) \cap N^-(y)$, $\ldots$, $L(u_k) \cap N^-(y)$, where $N^-(u) = \{u_1, u_2, \ldots, u_k\}$.

If $u \in A^*$ and $N^+(u) \cap V_{i+1} = \{v_1, v_2, \ldots, v_k\}$, remove $a \in L(u)$ if there exists a $v_j \in \{v_1, v_2, \ldots, v_k\} \cap A$, such that for every $z \in L(v_j)$ there does not exist an SDR for the sets: $\{a\} \cap N^-(z)$, $L(u_{j_1}) \cap N^-(z)$, $L(u_{j_2}) \cap N^-(z)$, $\ldots$, $L(u_{j_k}) \cap N^-(z)$, where $(N^-(v_j) - \{u\}) = \{u_{j_1}, u_{j_2}, \ldots, u_{j_k}\}$.

Let T be an oriented tree rooted at a vertex x together with the corresponding level sets $V_0 = \{x\}, V_1, \ldots, V_\ell$. For $v \in V_i$, consider the forest induced by $V(T) - (V_0 \cup V_1 \cup \cdots \cup V_{i-1})$. Let T_v be the sub-tree of T that contains v in the aforementioned forest. This is the sub-tree of T rooted at v, relative to x. The height of T_v is at most $\ell - i$.

Theorem 4.7.1. *Let H be a digraph and T an oriented tree rooted at a vertex x. If Algorithm 4.1 terminates with $L(x) \neq \emptyset$, then for every $v \in V(T)$ and for every $a \in L(v)$, there exists an injective homomorphism $f : T_v \xrightarrow{inj} H$ such that $f(v) = a$.*

Proof. If $L(x) \neq \emptyset$, then $L(v) \neq \emptyset$ for every $v \in V(T)$ (an empty list in V_i would lead to empty lists in V_j, $j \leq i$, in particular $V_0 = \{x\}$ would have an empty list).

Let $v \in V_i$. The proof is by induction on the height of T_v: $h = \ell - i$, $0 \leq i \leq \ell$.

When $h = 0$, $i = \ell$, v is a leaf and $T_v = \{v\}$. Therefore any element in $L(v)$ defines an injective homomorphism $f : T_v \xrightarrow{\text{inj}} H$.

Assume that the statement is true for all $0 \leq h < k \leq \ell$. That is, the statement is true for all $v \in V_j$ with $\ell - k + 1 \leq j \leq \ell$. Let $v \in V_{\ell-k}$ so that T_v has height $h = k$. Since $L(v) \neq \emptyset$, there exists an $a \in L(v)$ such that we can define an injective homomorphism $f : T[(N(v) \cap V_{\ell-k+1}) \cup \{v\}] \xrightarrow{\text{inj}} H$ with $f(v) = a$ (we might have to re-compute some of the SDRs to do this). Let $u \in N(v) \cap V_{\ell-k+1}$, with $f(u) = b \in L(u)$. The height of $(T_v)_u$ is $k - 1$ and by the induction hypothesis there exists an injective homomorphism $f_u : (T_v)_u \xrightarrow{\text{inj}} H$ such that $f_u(u) = f(u) = b$. Since this applies to every $u \in N(v) \cap V_{k-1}$ we can extend f to all of T_v. ∎

Corollary 4.7.2. *Let H be a digraph and T an oriented tree rooted at a vertex x. Then there exists an injective homomorphism $f : T \xrightarrow{inj} H$ if and only if Algorithm 4.1 terminates with $L(x) \neq \emptyset$.*

Proof. If $f : T \xrightarrow{\text{inj}} H$ is an injective homomorphism, then $d_H^-(f(v)) \geq d_T^-(v)$ and so $f(v) \in L(v)$. Furthermore, $f(v)$ is never removed from $L(v)$ during the execution of the algorithm. In particular $f(x) \in L(x)$, and so $L(x) \neq \emptyset$.

The converse follows from Theorem 4.7.1. ∎

We mention in closing that the algorithm has a running time that is proportional to the number of vertices in T. This follows from the fact that each vertex of T is processed only once and the processing (one-sided consistency check and SDR computation) at each vertex is a function of $|V(H)|$, which is fixed.

5

Conclusions and Future Work

In this chapter we look back at some of the results in this thesis, and we also look forward to possible future work.

5.1 Local Tournaments

By proving the dichotomy theorem for local tournaments, we have extended the class of digraphs for which a dichotomy is known. The most surprising result is the existence of unicyclic tournaments that define NP-complete problems. This together with the result on smooth digraphs by Barto, Kozik and Niven [8] highlights the fact that sources and sinks pose a significant problem to determining the complexity of digraph homomorphisms.

Possible future work in this area may include trying to extend the theorem to all locally semicomplete digraphs. New ideas may be needed here since our sub-indicators (for local tournaments) were often paths of length two. The presence of 2-cycles in locally semicomplete digraphs may negate the usefulness of these sub-indicators.

5.2 Complexity with Acyclic Inputs

The results in Section 1.5.2 on the complexity of oriented colourings with acyclic inputs can possibly be extended to tournaments (or even semicomplete digraphs) in general. That is, one would like to prove that colouring by a tournament (or maybe even a semicomplete digraph) with at least two cycles, is still NP-complete even if the input is an acyclic digraph. All tournaments with at most six vertices have been examined and this result holds for these tournaments.

5.3 Weak Near Unanimity Functions

We know of at least one homomorphism problem that has a polynomial time algorithm and has a NUF: directed C_k-colouring. Most unicyclic tournaments do not have a NUF, but they do have a WNUF; they are also polynomial problems.

- ❒ What is the difference (if any) between polynomial problems with a NUF and those with no NUF, but a WNUF?
- ❒ Is it possible to find more parallels between polynomial problems and WNUFs?
- ❒ Is there a polynomial digraph homomorphism problem that does not have a WNUF_3, but has one of higher arity?
- ❒ Extend the no-WNUF result for tournaments to semi-complete (local tournaments?) digraphs. A result like this will be embedded in the smooth digraph result of Barto, Kozik and Niven [8]. It may nonetheless be instructive to find an alternative proof using sub-indicators.
- ❒ Let G be an undirected, non-bipartite graph. Bulatov [15] has shown that all polymorphisms of G are essentially unary (re-proving Hell and Nešetřil's

result [36] using polymorphisms). Through the universal algebraic equivalences outlined in Chapter 1, this must mean that G does not have a WNUF. Give a direct proof of this.

❒ What can be said about oriented cycles and and their WNUFs?

5.4 Injective Homomorphisms

As far as irreflexive injective homomorphisms are concerned, our results have shown that targets with in-degree two exhibit a richness in complexity comparable to that of ordinary digraph homomorphisms. It would be beneficial to try and identify more polynomial cases of in-degree two.

❒ In particular, what can be said about oriented paths? Are they always polynomial, NP-complete, or do they exhibit a dichotomy (for injective homomorphisms)?

❒ Is it possible to extend results that hold for ordinary homomorphisms to injective homomorphisms?

❒ Is there a product of digraphs that's relevant to injective homomorphisms?

❒ Can the injective homomorphism problem be recast as a constraint satisfaction problem? When this is achieved, the corresponding polymorphisms may be interesting.

References

[1] Jørgen Bang-Jensen, *Locally semicomplete digraphs: a generalization of tournaments*, J. Graph Theory **14** (1990), no. 3, 371–390. MR MR1060865 (91g:05055)

[2] Jørgen Bang-Jensen, Yubao Guo, Gregory Gutin, and Lutz Volkmann, *A classification of locally semicomplete digraphs*, Discrete Math. **167/168** (1997), 101–114, 15th British Combinatorial Conference (Stirling, 1995). MR MR1446736 (98a:05121)

[3] Jørgen Bang-Jensen and Gregory Gutin, *Digraphs: Theory, Algorithms and Applications*, Springer Monographs in Mathematics, Springer-Verlag London Ltd., London, 2001. MR MR1798170 (2002e:05002)

[4] Jørgen Bang-Jensen and Pavol Hell, *The effect of two cycles on the complexity of colourings by directed graphs*, Discrete Appl. Math. **26** (1990), no. 1, 1–23. MR MR1028872 (91c:05072)

[5] Jørgen Bang-Jensen, Pavol Hell, and Gary MacGillivray, *The complexity of colouring by semicomplete digraphs*, SIAM J. Discrete Math. **1** (1988), no. 3, 281–298. MR MR955645 (89e:05095)

[6] Jørgen Bang-Jensen and Gary MacGillivray, *Further effects of two directed cycles on the complexity of H-colouring*, J. Combin. Math. Combin. Comput. **10** (1991), 33–50. MR MR1137820 (93a:05116)

[7] ______, *On the complexity of colouring by digraphs with at most one directed cycle*, Ars Combin. **35** (1993), no. A, 175–186. MR MR1275235 (94m:05075)

[8] Libor Barto, Marcin Kozik, and Todd Niven, *The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell)*, To appear in SIAM J. Computing.

[9] J. A. Bondy and U. S. R. Murty, *Graph theory*, Graduate Texts in Mathematics, vol. 244, Springer, New York, 2008. MR MR2368647

[10] Alfred Brauer, *On a problem of partitions*, Amer. J. Math. **64** (1942), 299–312. MR MR0006196 (3,270d)

[11] Gaëlle Brevier, Romeo Rizzi, and Stéphane Vialette, *Pattern matching in protein-protein interaction graphs*, Fundamentals of Computation Theory, Lecture Notes in Comput. Sci., vol. 4639, Springer, Berlin, 2007, pp. 137–148.

[12] Richard A. Brualdi and Herbert J. Ryser, *Combinatorial matrix theory*, Encyclopedia of Mathematics and its Applications, vol. 39, Cambridge University Press, Cambridge, 1991. MR MR1130611 (93a:05087)

[13] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin, *Classifying the complexity of constraints using finite algebras*, SIAM J. Comput. **34** (2005), no. 3, 720–742 (electronic). MR MR2137072 (2005k:68181)

[14] Andrei Bulatov and Matthew Valeriote, *Recent results on the algebraic approach to the CSP*, Submitted. Available at `http://www.cs.sfu.ca/~abulatov/jpapers.html`.

[15] Andrei A. Bulatov, *H-coloring dichotomy revisited*, Theoret. Comput. Sci. **349** (2005), no. 1, 31–39. MR MR2183466 (2006i:05057)

[16] Samuel Eilenberg and Saunders MacLane, *Natural isomorphisms in group theory*, Proc. Nat. Acad. Sci. U. S. A. **28** (1942), 537–543. MR MR0007421 (4,134d)

[17] Isabelle Fagnot, Gaëlle Lelandais, and Stéphane Vialette, *Bounded list injective homomorphism for comparative analysis of protein-protein interaction graphs*, CompBioNets 2004: Algorithms & computational methods for biochemical and evolutionary networks, Texts Algorithmics, vol. 3, King's Coll. Publ., London, 2004, pp. 45–70. MR MR2209695 (2006j:92036)

[18] ———, *Bounded list injective homomorphism for comparative analysis of protein-protein interaction graphs*, J. Discrete Algorithms **6** (2008), no. 2, 178–191. MR MR2418976

[19] Tomás Feder, *Classification of homomorphisms to oriented cycles and of k-partite satisfiability*, SIAM J. Discrete Math. **14** (2001), no. 4, 471–480 (electronic). MR MR1861785 (2002j:05063)

[20] Tomás Feder and Moshe Y. Vardi, *The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory*, SIAM J. Comput. **28** (1999), no. 1, 57–104 (electronic). MR MR1630445 (2000e:68063)

[21] Guillaume Fertin, Romeo Rizzi, and Stéphane Vialette, *Finding exact and maximum occurrences of protein complexes in protein-protein interaction graphs*, Mathematical foundations of computer science 2005, Lecture Notes in Comput. Sci., vol. 3618, Springer, Berlin, 2005, pp. 328–339. MR MR2237379 (2007b:92025)

[22] Jiří Fiala and Jan Kratochvíl, *Complexity of partial covers of graphs*, Algorithms and computation (Christchurch, 2001), Lecture Notes in Comput. Sci., vol. 2223, Springer, Berlin, 2001, pp. 537–549. MR MR1917771 (2003e:68108)

[23] ______, *Partial covers of graphs*, Discuss. Math. Graph Theory **22** (2002), no. 1, 89–99, Conference on Graph Theory (Elgersburg, 2000). MR MR1936227 (2003i:05125)

[24] ______, *Locally injective graph homomorphism: lists guarantee dichotomy*, Graph-theoretic concepts in computer science, Lecture Notes in Comput. Sci., vol. 4271, Springer, Berlin, 2006, pp. 15–26. MR MR2290714 (2007i:05063)

[25] Jiří Fiala, Jan Kratochvíl, and Attila Pór, *On the computational complexity of partial covers of theta graphs*, Discrete Appl. Math. **156** (2008), no. 7, 1143–1149. MR MR2404227

[26] Michael R. Garey and David S. Johnson, *Computers and intractability*, W. H. Freeman and Co., San Francisco, Calif., 1979, A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences. MR MR519066 (80g:68056)

[27] Yubao Guo and Lutz Volkmann, *On complementary cycles in locally semicomplete digraphs*, Discrete Math. **135** (1994), no. 1-3, 121–127. MR MR1310875 (96e:05094)

[28] Wolfgang Gutjahr, *Graph colourings*, Ph.D. thesis, Freie Universität Berlin, 1991.

[29] Wolfgang Gutjahr, Emo Welzl, and Gerhard Woeginger, *Polynomial graph-colorings*, Discrete Appl. Math. **35** (1992), no. 1, 29–45. MR MR1138082 (92m:05081)

[30] Richard K. Guy, *Unsolved problems in number theory*, third ed., Problem Books in Mathematics, Springer-Verlag, New York, 2004. MR MR2076335 (2005h:11003)

[31] Geňa Hahn, Jan Kratochvíl, Jozef Širáň, and Dominique Sotteau, *On the injective chromatic number of graphs*, Discrete Math. **256** (2002), no. 1-2, 179–192. MR MR1927065 (2003h:05081)

[32] Geňa Hahn and Gary MacGillivray, *Graph homomorphisms: computational aspects and infinite graphs*, Unpublished document. http://www.math.uvic.ca/faculty/gmacgill/homomorphisms/.

[33] Geňa Hahn and Claude Tardif, *Graph homomorphisms: structure and symmetry*, Graph symmetry (Montreal, PQ, 1996), NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., vol. 497, Kluwer Acad. Publ., Dordrecht, 1997, pp. 107–166. MR MR1468789 (99c:05091)

[34] P. Hell, J. Nešetřil, and X. Zhu, *Complexity of tree homomorphisms*, Discrete Appl. Math. **70** (1996), no. 1, 23–36. MR MR1402736 (97e:68049)

[35] Pavol Hell, *From graph colouring to constraint satisfaction: there and back again*, Topics in discrete mathematics, Algorithms Combin., vol. 26, Springer, Berlin, 2006, pp. 407–432. MR MR2249279 (2007f:68081)

[36] Pavol Hell and Jaroslav Nešetřil, *On the complexity of H-coloring*, J. Combin. Theory Ser. B **48** (1990), no. 1, 92–110. MR MR1047555 (91m:68082)

[37] ———, *Graphs and homomorphisms*, Oxford Lecture Series in Mathematics and its Applications, vol. 28, Oxford University Press, Oxford, 2004. MR MR2089014 (2005k:05002)

[38] Pavol Hell, Hui Shan Zhou, and Xu Ding Zhu, *On homomorphisms to acyclic local tournaments*, J. Graph Theory **20** (1995), no. 4, 467–471. MR MR1358537 (96f:05085)

[39] Pavol Hell and Xu Ding Zhu, *The existence of homomorphisms to oriented cycles*, SIAM J. Discrete Math. **8** (1995), no. 2, 208–222. MR MR1329507 (96i:05096)

[40] Peter Jeavons, *On the algebraic structure of combinatorial problems*, Theoret. Comput. Sci. **200** (1998), no. 1-2, 185–204. MR MR1625523 (99d:68085)

[41] Richard E. Ladner, *On the structure of polynomial time reducibility*, J. Assoc. Comput. Mach. **22** (1975), 155–171. MR MR0464698 (57 #4623)

[42] Benoit Larose and Claude Tardif, *Projectivity and independent sets in powers of graphs*, J. Graph Theory **40** (2002), no. 3, 162–171. MR MR1905132 (2003d:05159)

[43] Benoit Larose and László Zádori, *The complexity of the extendibility problem for finite posets*, SIAM J. Discrete Math. **17** (2003), no. 1, 114–121 (electronic). MR MR2033309 (2004h:06007)

[44] ______, *Taylor terms, constraint satisfaction and the complexity of polynomial equations over finite algebras*, Internat. J. Algebra Comput. **16** (2006), no. 3, 563–581. MR MR2241624 (2007i:08004)

[45] Miklós Maróti and Ralph McKenzie, *Existence theorems for weakly symmetric operations*, To appear in Algebra Universalis. `http://www.math.u-szeged.hu/~mmaroti/publications.html`.

[46] H. A Maurer, J. H. Sudborough, and E. Welzl, *On the complexity of the general coloring problem*, Inform. and Control **51** (1981), no. 2, 128–145. MR MR686834 (84i:68073)

[47] John W. Moon, *Topics on tournaments*, Holt, Rinehart and Winston, New York, 1968. MR MR0256919 (41 #1574)

[48] J. L. Ramírez Alfonsín, *The Diophantine Frobenius problem*, Oxford Lecture Series in Mathematics and its Applications, vol. 30, Oxford University Press, Oxford, 2005. MR MR2260521 (2007i:11052)

[49] Thomas J. Schaefer, *The complexity of satisfiability problems*, Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978), ACM, New York, 1978, pp. 216–226. MR MR521057 (80d:68058)

[50] Eric Sopena, *Oriented graph coloring*, Discrete Math. **229** (2001), no. 1-3, 359–369, Combinatorics, graph theory, algorithms and applications. MR MR1815613 (2002g:05088)

[51] Herbert S. Wilf, *generatingfunctionology*, second ed., Academic Press Inc., Boston, MA, 1994. MR MR1277813 (95a:05002)