CYCLE STRUCTURE AND COLORINGS OF DIRECTED GRAPHS

vorgelegt von M. Sc.

Raphael Mario Steiner

an der Fakultät II – Mathematik und Naturwissenschaften der Technischen Universität Berlin zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften

– Dr. rer. nat. –

genehmigte Dissertation

Promotionsausschuss

Vorsitzender: Prof. Dr. Boris Springborn Gutachter: Prof. Dr. Stefan Felsner Prof. Dr. Bojan Mohar Prof. Dr. Tibor Szabó

Tag der wissenschaftlichen Aussprache: 18. $08.\ 2021$

Berlin 2021

"Mathematics is the music of reason."

James Joseph Sylvester

Zusammenfassung

Diese Dissertation beschäftigt sich mit Problemstellungen aus der Theorie der endlichen gerichteten Graphen. Ein (endlicher) gerichteter Graph ist eine binäre Relation, deren Quellmenge endliche Größe besitzt. Gerichtete Graphen stellen damit eine sehr allgemeine Art und Weise dar, möglicherweise asymmetrische Beziehungen zwischen einer endlichen Menge von Objekten zu codieren. Selbstverständlich erlaubt es eine solche Allgemeinheit, viele Probleme durch gerichtete Graphen zu abstrahieren, besonders dann, wenn sich wichtige Eigenschaften durch Beziehungen oder Verbindungen zwischen Objekten ausdrücken lassen. Als ausgewählte Beispiele seien hier Straßennetzwerke, Funknetze, Gasnetzwerke, das Internet, Schaltkreise in elektronischen Geräten, sowie neuronale Netzwerke genannt.

Ein Schwerpunkt der vorliegenden Arbeit liegt auf der Untersuchung von Grapheneigenschaften im Zusammenhang mit einem der wohl fundamentalsten Objekte der Graphentheorie, dem sogenannten *Kreis*. Ein Kreis in einem Graphen wird beschrieben durch eine geschlossene Folge von zyklisch benachbarten Knoten und Kanten ohne auftretende Wiederholungen.

In einem Graphen endlicher Größe kann man typischerweise erwarten, eine ganze Reihe verschiedenster Kreise zu finden. Aufgrunddessen sind Kreise ein wichtiges und wiederkehrendes Motiv in fast allen Zweigen der Graphentheorie und treten beispielsweise in der strukturellen Graphentheorie auf, in der Theorie von Flüssen auf gerichteten Graphen, in theoretischen Charakterisierungen von Graphenklassen und in der Theorie der Färbung von Graphen. Zudem spielen Kreise in etlichen Verfahren zur Lösung von algorithmischen Problemen auf Graphen eine entscheidende Rolle, wie beispielsweise dem Problem des Handlungsreisenden, Algorithmen zur Berechnung eines größten Matchings oder dem Maximum-Flow-Problem und auch in Subprozeduren wie beispielweise Kruskals Algorithmus zum Finden eines leichtesten Spannbaums. Aus diesen Gründen hat sich eine beachtliche Menge an Resultaten der Graphentheorie darauf spezialisiert, die Kreisstruktur in Graphen selbst als Untersuchungsobjekt zu betrachten.

Im ersten Teil dieser Arbeit befassen wir uns mit Kreisen, die in gerichteten Graphen auftreten, beweisen verschiedene hinreichende und notwendige Bedingungen für die Existenz solcher Kreise und betrachen algorithmische Fragestellungen in diesem Zusammenhang. Im zweiten Teil der Arbeit beschäftigen wir uns mit dem Problem der gerichteten Graphenfärbung, das auch mit der Existenz gewisser Kreise zu tun hat. Hierbei handelt es sich um ein Optimierungsproblem, bei dem man die Knoten des Graphen mit möglichst wenigen Farben so zu färben versucht, dass *keine* gerichteten Kreise innerhalb einer Farbklasse auftreten. Dieses Problem wurde vor vier Jahrzehnten von Paul Erdős und Victor Neumann-Lara aufgeworfen und seitdem in der Literatur studiert. In dieser Arbeit tragen wir neue Resultate bei, die einerseits neue Schranken an die Zahl der benötigten Farben liefern und andererseits die algorithmische Komplexität des Problems beleuchten.

Der dritte Teil der Arbeit knüpft an den zweiten Teil an und liefert neue Resultate für diverse weitere Arten, einen gerichteten Graphen zu färben oder in möglichst wenige strukuriertere Subgraphen zu zerlegen.

Abstract

This thesis deals with problems from the theory of *finite directed graphs*. A directed graph (digraph for short) is a binary relation whose domain has finite size. With that digraphs can be seen as a very general way of representing (possibly asymmetric) relations between pairs from a finite set of objects. Undoubtedly, such a generality allows to encode many structures by digraphs. This works particularly well if important properties of the structure at hand can be expressed as relations or connections between objects. To give some selected examples, let us mention road networks, electricity networks, radio networks, the world wide web, circuits in electronic devices, or neural networks.

A main focus of the thesis at hand is the investigation of properties of one of the most fundamental objects all over graph theory, the so-called *cycle* (sometimes also called *circuit*). A cycle in a graph is determined by a closed alternating sequence of cyclically connected vertices and edges.

In a graph of finite size one will typically see loads of distinct cycles of various types. Therefore cycles constitute an important and recurring motive in almost all branches of graph theory, for instance, they play important roles in structural graph theory, in the theory of flows on directed networks, in theoretical characterizations of graph classes, as well as in the theory of graph colorings. Additionally, cycles play a decisive role in numerous algorithmic problems and their solutions, such as in the Traveling Salesman Problem, algorithms for finding a largest matching in a given graph, in the max-flow problem, and also in subprocedures such as Kruskal's algorithm for finding a minimum weight spanning tree. For those reasons, a substantial amount of research in graph theory has specialised on the structure of cycles in graphs.

In the first major part of this thesis we deal with cycles which occur in directed graphs, and prove several necessary and sufficient theoretical conditions for the existence of cycles of certain types. Additionally, we deal with algorithmic problems related to cycles in directed graphs.

In the second part we deal with the problem of *acyclic colorings* of directed graphs, which also relates to the (non-)existence of certain cycles. The *dichromatic number* represents an optimization problem in which we seek to color the vertices of a given digraph with the fewest number of colors while avoiding monochromatic directed cycles. This topic was introduced 40 years ago by Paul Erdős and Victor Neumann-Lara and since then, particularly in the last two decades, has been considered by many researchers. In this thesis we contribute new results that on the one hand establish new theoretical bounds on the dichromatic number and on the other hand shed more light on the computational complexity of this problem.

The third and last major part of this thesis carries on with the topic of digraph colorings and presents new results for various further notions of digraph coloring and ways of decomposing a given digraph into the fewest number of simpler subdigraphs.

Acknowledgements

The research during my PhD was funded by a scholarship of the research training group GRK 2434 Facets of Complexity. I am grateful to Deutsche Forschungsgemeinschaft and to the faculty of the GRK for giving me the opportunity of coming to Berlin, where I was able to conduct my own research in an inspiring and active research environment both at Technische and Freie Universität Berlin and to make many new contacts.

My sincerest thanks go to my supervisor Stefan Felsner for teaching me how to think and write about maths, for giving me lots of free space to develop as an independent researcher, for giving me extremely useful academic and personal guidance whenever it was required, and for defending me in a very difficult situation.

I would also like to thank Tibor Szabó and Winfried Hochstättler for additional guidance and advise through the past years.

I am highly indebted to the referees Stefan Felsner, Bojan Mohar, and Tibor Szabó for reading and evaluating this thesis, and to Boris Springborn for accepting the chairmanship. Many thanks to David, Manfred and Sebastian for proof-reading parts of this thesis.

A big thank you to all my colleagues and coauthors, and in particular Felix, Hendrik, Lior, Manfred, Max, Sebastian and Tamás, for sharing many exciting moments on and off the white-board or IDroo with me, which I will keep as beautiful memories.

Most of all I would like to thank my parents and my two sisters for always being there for me, no matter what.

Credits

Contents of the following research projects are represented fully or partially in this thesis:

Oriented cycles in digraphs with large out-degree with *Lior Gishboliner* and *Tibor Szabó* [GSS20b].

Content of this project appears in Chapter 2.

Disjoint cycles with length constraints in digraphs of large connectivity or minimum degree [Ste20].

Content of this project appears in the Chapters 2 and 3.

Flip distances between graph orientations with Oswin Aichholzer, Jean Cardinal, Tony Huynh, Kolja Knauer, and Birgit Vogtenhuber [ACH⁺21]. Content of this project appears in Chapter 4.

Even circuits in oriented matroids with *Karl Heuer* and *Sebastian Wiederrecht* [HSW20]. Content of this project appears in Chapter 5.

Complete minors in digraphs with given dichromatic number with *Tamás Mészáros* [MS21a]. Content of this project appears in the Chapters 6, 7, and 8.

Dichromatic number and forced subdivisions with *Lior Gishboliner* and *Tibor Sz-abó* [GSS20a].

Content of this project appears in Chapter 7.

Colouring non-even digraphs with *Marcelo Millani* and *Sebastian Wiederrecht* [MSW19]. Content of this project appears in the Chapters 8, 10, and 14.

Coloring digraphs with forbidden induced subgraphs [Ste21]. Content of this project appears in Chapter 9.

Parametrized algorithms for directed modular width with *Sebastian Wiederrecht* [SW19, SW20].

Content of this project appears in Chapter 10.

Complete acyclic colorings with *Stefan Felsner*, *Winfried Hochstättler*, and *Kolja Knauer* [FHKS20].

Content of this project appears in Chapter 11.

Majority colorings of sparse digraphs with *Michael Anastos, Ander Lamaison* and *Tibor Szabó* [ALSS19].

Content of this project appears in Chapter 12.

On the complexity of digraph colourings and vertex arborcity with *Winfried Hochstättler* and *Felix Schröder* [HSS20].

Content of this project appears in Chapter 13.

The star dichromatic number with *Winfried Hochstättler* [HS19]. Content of this project appears in Chapter 14.

Contents

Zι	usammenfassung	5				
\mathbf{A}	Abstract					
A	Acknowledgements					
Co	Contents					
1	Preliminaries	31				
Ι	Existence and Structure of Oriented Cycles	47				
2	Oriented Cycles in Digraphs of Large Out-Degree	49				
	2.1 Introduction	49				
	2.2 Subdivisions of Oriented Cycles	53				
	2.2.1 The Gadgets	54				
	2.2.2 Gadget Chains	57				
	2.2.3 Embedding Gadgets	61				
	2.2.4 Putting It All Together	64				
	2.3 Disjoint Oriented Cycles	68				
	2.4 Oriented Cycles with Two Blocks	71				
	2.5 Subdivisions of $K_3 - e$	73				
	2.6 Subdivisions and Arc-Connectivity	75				
	2.7 Conclusion	77				
3	Disjoint Directed Cycles with Length Constraints	79				
	3.1 Introduction	79				
	3.2 Preliminaries	82				
	3.3 Finding Disjoint Dicycles of Distinct Lengths	83				
	3.3.1 Digraphs of Bounded Directed Tree-Width	84				
	3.3.2 Digraphs Containing a Large Complete Minor	85				
	3.3.3 Digraphs Containing a Flat Wall	86				
	3.3.4 Proofs of Theorem 3.1 and Theorem 3.3	91				
	3.4 Equicardinal Disjoint Directed Cycles	91				
	3.5 Conclusion	92				
4	Cycle-Reversion Distance between Graph Orientations	95				
	4.1 Introduction	95				
	4.2 Problems and Main Results	97				

	4.3	Flip Distance between Perfect Matchings and between $\alpha\text{-}\mathrm{Orientations}$. 106
	4.4	Vertex Flip Distance between <i>c</i> -Orientations	. 108
	4.5	Flip Distance with Larger Cut Sets	. 113
	4.6	Conclusion	. 114
5	Eve	en Circuits in Regular Oriented Matroids	117
	5.1	Introduction	. 117
		5.1.1 The Even Directed Circuit Problem in Oriented Matroids \ldots .	. 118
	5.2	Non-Evenness and GB-minors	. 119
	5.3	The Complexity of the Even Directed Circuit Problem	. 121
		5.3.1 Proof of Theorem 5.3 \ldots \ldots \ldots \ldots \ldots \ldots	. 123
	5.4	Digraphs Admitting an Odd Dijoin	. 129
		5.4.1 Forbidden Cut Minors for Digraphs with an Odd Dijoin	. 130
	5.5	Conclusion	. 137
II	D	ichromatic Number	139
6	Exc	cluding Strong Minors	143
	6.1	Introduction	. 143
	6.2	Proof of Theorem 6.2	. 144
	6.3	Conclusion	. 145
7	Exc	eluding Topological Minors	147
	7.1	Introduction	. 147
		7.1.1 Our Results	. 149
	7.2	Preliminaries	. 151
	7.3	Oriented Cacti and Bioriented Forests	. 152
	7.4	Tournaments of Order 4	. 156
	7.5	Subcubic Digraphs	. 165
	7.6	Mader-Perfect Digraphs and Conclusion	. 166
8	Exc	cluding Butterfly Minors	171
	8.1	Introduction	. 171
	8.2	2-Colorings of Non-Even Digraphs	. 177
	8.3		. 183
	8.4	List Colorings of Non-Even Digraphs	. 184
	8.5	Non-Bipartite Graphs	. 186
	8.6	Conclusion	. 190
9	Exc	cluding Induced Subdigraphs	193
	9.1	Introduction	. 193
	9.2	$\{S_2^+, W_3^+\}$ -Free Oriented Graphs	. 198
	9.3	$\{S_2^+, W_3^-\}$ -Free Oriented Graphs	. 205
	9.4	Adding a Dominating Sink to a Hero	. 207
	9.5	Oriented 4-Vertex-Paths	. 209
	9.6	Conclusion	. 210

10 Complexity Results	213
10.1 Introduction	213
10.2 Parametrization by Feedback Vertex Set	217
10.3 Parametrization by Modular Width	220

III Other Coloring Concepts for Digraphs

11	Complete Colorings	229
	11.1 Introduction	229
	11.2 First Observations and Results	231
	11.2.1 Relations between Parameters	231
	11.2.2 Behavior with Respect to Graph Operations	235
	11.3 Upper Bounds for Minimum Feedback Vertex Sets	237
	11.3.1 Interplay of ava and adi \ldots	241
	11.4 Conclusion	242
12	Majority Colorings	245
	12.1 Introduction	245
	12.1.1 Our Results	246
	12.2 Digraphs without Odd Directed Cycles	249
	12.3 Majority 3-Colorings of Sparse Digraphs	251
	12.4 Fractional Majority Colorings	253
	12.5 Conclusion	256
13	Complexity of Fractional and Circular Colorings	259
	13.1 Introduction	259
	13.2 Complexity of the Star Dichromatic Number	264
	13.3 Complexity of the Fractional Dichromatic Number	268
	13.4 Complexity of the Circular Vertex Arboricity	272
	13.5 Conclusion	274
14	Fractional Colorings of Planar Digraphs	277
	14.1 Introduction	277
	14.2 Proofs	279
\mathbf{A}	Open Problems Collection	283
Bi	bliography	287

CONTENTS

Introduction

This dissertation is concerned with problems from the theory of *finite directed graphs*. To start with, a (finite) directed graph is a binary relation whose domain has finite size. Directed graphs can be seen as a very general way of representing possibly asymmetric relations between a finite set of objects. In the language of directed graphs, the elements of the domain of the relation are called *vertices*, while the relations between the objects are called *arcs* or *directed edges*. As a matter of course, the generality of the definition of directed graphs (digraphs for short) allows to encode and abstract many real-world problems through digraphs. Such an encoding is particularly successful if important properties of the problem can be expressed by connections between certain objects. To give some selected examples, let us mention here road networks, electricity networks, radio networks, the world wide web, circuits in electronic devices, or neural networks.

The definition of directed graphs stands in contrast to the definition of an undirected graph, whose representing binary relation is required to be symmetric. The unordered pairs of an undirected graph in relation are called *edges*, the corresponding vertices *adjacent*. The symmetry condition required for undirected graphs shows that directed graphs constitute a far-reaching generalization of graphs, as the relation of an unordered pair of vertices in a simple digraph may assume up to four different states, rather than two as in simple undirected graphs. Accordingly, it is not very surprising that directed graphs exhibit more complex structures than the ones found in undirected graphs. For that reason, a phenomenon that we will witness repeatedly in the course of this thesis is that many an inconspicuous problem formulated on directed graphs, compared to its undirected analogue, will turn out quite difficult or will have a very surprising answer.

A main focus of the thesis at hand is the study of properties of directed graphs in connection with certainly one of the most fundamental objects in graph theory, the so-called *cycle* (less commonly also called circuit). A cycle in an undirected graph is a cyclical repetition-free alternating sequence of vertices and edges such that edges connect their neighboring vertices in the sequence. In digraphs, we differentiate between *oriented cycles* and *directed cycles*. While an oriented cycle is a cyclical repetition-free alternating sequence of vertices acyclical repetition-free alternating sequence of vertices acyclical repetition-free alternating sequence of vertices and acycle is a cyclical repetition-free alternating sequence, a *directed cycle* is a special kind of oriented cycle in which all arcs are oriented consistently in the direction of circulation.

A finite graph can typically be expected to contain a load of distinct cycles, since every walk along the edges of the graph will either have to stop or return to an already visited vertex of the graph after finitely many steps, and thereby create a cycle. Graphs containing no cycles, known as forests, are very sparse and exceptional. Therefore cycles constitute an important and recurring motive in almost all branches of graph theory. For instance, cycles play a fundamental role in structural graph theory, in the theory of flows on directed networks, in theoretical characterizations of graph classes, and in the theory of graph coloring. Moreover, cycles play a decisive role in numerous procedures for solving algorithmic problems on graphs, as exemplified by the famous Traveling Salesman Prob-



Figure 1: An undirected cycle, an oriented cycle, and a directed cycle, each containing four vertices. Vertices are represented by points, arcs (u, v) between vertices u and v are represented by an arrow pointing from u to v.

lem, maximum matching- and maximum flow-algorithms, as well as Kruskal's algorithm for finding a minimum-weight spanning tree.

For those reasons a substantial amount of research in graph theory has specialized on studying the cycle structure of undirected and directed graphs for its own sake. A wide range of questions has been studied, but many remain open. Accordingly, the first main goal of the thesis at hand is to contribute new results which enhance the theory of oriented and directed cycles in digraphs.

The second main topic of this thesis are colorings of directed graphs. Colorings of graphs in general can be seen as optimization problems which seek to partition the vertexset (or sometimes the edge-set) of a given graph into the fewest number of parts which induce subgraphs with a certain *simpler* structure. What exactly is meant here by *simpler* varies depending on the definition of the coloring concept. However one can hardly disagree that a graph consisting of a set of isolated vertices (with an empty edge-set) is probably as simple as it gets. So if we say that an edgeless graph is what we would like to call simple, then the coloring concept we end up with is exactly the well-known *chromatic* number $\chi(G)$ of an undirected graph G, i.e., the smallest number of independent sets in G partitioning the vertex set. Instead of a partition, we may also think of a coloring of a graph as a mapping c which assigns to every vertex v of the graph an element c(v) of a finite (color) set, in such a way that the *color classes* (i.e., the sets of vertices sharing the same color) induce subgraphs with a simpler structure. The goal is then to find a coloring minimizing the size of the color set which is used for the coloring. Therefore in its standard definition, $\chi(G)$ is the smallest number of colors which can be used in a proper *coloring*, i.e., a vertex-coloring such that adjacent vertices receive distinct colors.

Graph coloring and in particular the study of the chromatic number of graphs constitutes one of the most active and biggest parts of graph theory up until now. Moreover, graph coloring is also one of the earliest and most influential branches of graph theory, whose first occurences date back more than 150 years. Maybe the most celebrated problem in graph theory is the 4-Color-Problem, asking whether the bounded regions of every planar map can be colored using 4 colors such that regions sharing a common border receive different colors. This problem was finally resolved in the positive in 1976 when Appel and Haken [AH76, AHK77] presented a computer-assisted proof of their famous 4-Color-Theorem, which formally states that the chromatic number of every planar graph is at most four.

While the chromatic number is the most popular graph coloring parameter for undirected graphs, it is not very well-suited for directed graphs, as the independent sets in a digraph do not depend upon the orientation of the directed edges. Hence, for a more meaningful coloring parameter for directed graphs, the definition of what makes a digraph simple should be dependent not only on the underlying graph. A natural choice for such a coloring concept was made by Erdős and Neumann-Lara [Erd80, NL82], who defined a coloring of a digraph (subsequently and throughout the thesis called *acyclic coloring*) as a partition into sets inducing acyclic subdigraphs, or equivalently, a vertex-coloring in which no directed cycle is monochromatic. The corresponding coloring parameter of a digraph D was named dichromatic number by Neumann-Lara and is denoted by $\vec{\chi}(D)$. Indeed, acyclic digraphs exhibit several properties which make them very special and a lot more structured than general digraphs. Since its introduction in the 1980s, the dichromatic number has grown in importance and has now been an active object of research for two decades. It may well be considered the most popular and well-studied coloring concept for directed graphs up to date. Motivated by this development, a main part of this thesis focuses on deriving new insights concerning the dichromatic number, and in particular finding structural conditions which enforce a digraph to have small or bounded dichromatic number. These structural conditions usually mean that we consider digraphs which do not contain certain types of substructures, such as minors, subdivisions or induced subgraphs, and try to find the best possible upper and lower bounds on the dichromatic number of digraphs in these classes. In another major part of this thesis, we will also turn our attention to other coloring parameters for directed graphs which were introduced more recently, and contribute some new results for these parameters.

The thesis at hand starts off with some fundamental definitions as well as notation used repeatedly throughout the thesis. The main content of the thesis is divided into three main parts, covering the cycle structure of digraphs, the dichromatic number, and other coloring concepts for digraphs, respectively. Every one of the three parts is again subdivided into several chapters, which deal with different problems to which we have contributed new results. Every chapter has a more specialized introduction on its own. In the following let us give a broad outline of the three different parts and their chapters. For detailed definitions we refer to the Preliminaries (Chapter 1) or the respective chapters.

The Chapters of Part I

Part I, entitled *Existence and Structure of Oriented Cycles* focuses on both theoretical conditions for digraphs that enforce the existence of oriented and directed cycles of certain types and lengths, as well as algorithmic problems which are related to oriented cycles in digraphs.

In Chapter 2 we solve problems raised in a paper by Aboulker, Cohen, Havet, Lochet, Moura, and Thomassé [ACH⁺19] concerning subdivisions in digraphs of large minimum out-degree. A subdivision of a (di)graph is any (di)graph that can be obtained by replacing the edges (arcs) of the (di)graph by internally vertex-disjoint (directed) paths connecting the endpoints of the original edges (arcs). In the case of directed graphs, we require that the directed path replacing the arc (u, v) is oriented consistently from u towards v. Subdivision is a natural concept that appears in many fundamental graph theory results, such as Kuratowski's characterization of planar graphs. It is therefore desirable to find means which can force any graph to contain a subdivision of a fixed graph, such as K_5 , as a subgraph. For undirected graphs, a classical result by Mader [Mad67] asserts that graphs of sufficiently large minimum degree contain a subdivision of any given graph. Surprisingly, the analogous result for directed graphs does not hold true in a strong sense: It follows from a result by Thomassen [Tho85b] that every so-called even digraph¹, including the

¹A digraph F is called even, if every subdivision of F contains an even directed cycle.

complete digraph \vec{K}_3 consisting of three vertices and all 6 possible connecting arcs², cannot be forced as a subdivision by means of large minimum out- and in-degree. However, every even digraph contains a directed cycle. In 1985, Mader [Mad85] conjectured that maybe every *acyclic digraph* containing no directed cycles can be forced as a subdivision by large minimum out-degree. Despite its popularity, this conjecture so far has not been resolved, and remains open for acyclic digraphs on only five vertices. Aboulker et al. [ACH⁺19] worked on special cases of Mader's problem and raised the conjecture that every oriented cycle can be forced as a subdivision by large minimum out-degree. As the main contribution of Chapter 2 we fully resolve this conjecture. We also show the stronger result that digraphs of large minimum out-degree contain any given disjoint union of oriented cycles as a subdivision. Previously this was known only for disjoint unions of directed cycles (cf. [Tho83, Alo96]). Finally, we provide answers to related open problems raised in the paper by Aboulker et al.

As mentioned above, it is known that there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that every digraph of minimum out-degree at least f(k) contains k vertex-disjoint directed cycles. This was proved in 1983 by Thomassen [Tho83], who showed that we can take f(k) = (k + 1)! for every k. Several improvements on this bound have been made, but a famous conjecture by Bermond and Thomassen, claiming that f(k) = 2k - 1, still remains widely open. The proofs for the existence of f known so far, however, do not give any information concerning the lengths of the disjoint cycles.

In Chapter 3 we are motivated by a conjecture of Lichiardopol [Lic14] from 2014, which states the existence of a function $g: \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathbb{N}$ every digraph whose minimum out-degree is at least g(k) contains k vertex-disjoint directed cycles whose lengths are pairwise different. Note that this is a qualitative strengthening of Thomassen's result. Our main contribution of Chapter 3 is to show that Lichiardopol's conjecture holds true for digraphs of sufficiently large connectivity. More precisely, we show that there exists a function $s: \mathbb{N} \to \mathbb{N}$ such that every strongly s(k)-connected digraph contains k vertexdisjoint directed cycles of pairwise different lengths. In contrast, we give a construction showing that there are digraphs of arbitrarily large strong connectivity not containing two disjoint directed cycles of equal length, which strengthens an earlier construction by Alon [Alo96]. We further verify Lichiardopol's conjecture for digraphs of bounded *directed tree-width*. Directed tree-width is an important structural digraph parameter and will be used several times in this thesis as a technical tool for some of our proofs.

In Chapter 4 we consider an algorithmic problem involving directed cycles in oriented graphs. Given an undirected graph G and some assignment $\alpha : V(G) \to \mathbb{N}$, an α -orientation of G is an orientation of G such that every vertex $v \in V(G)$ has out-degree exactly $\alpha(v)$. Orientations with a prescribed out-degree pattern have received quite some attention in the literature previously, and have had applications in enumerative combinatorics, discrete geometry and order theory. Sometimes it is possible to find non-trivial bijections between a set of combinatorial objects and a set of α -orientations of a particular planar graph with some function α , and in this case the combinatorial objects automatically inherit a beautiful distributive lattice structure from the set of α -orientations of this graph [Fel04, dFOdM01].

Flip graphs are a ubiquitous class of graphs, which encode relations on a set of combinatorial objects by elementary, local changes. Skeletons of associahedra, for instance, are the graphs induced by quadrilateral flips in triangulations of a convex polygon.

²In general, for an integer $t \ge 1$ we denote by \vec{K}_t the complete digraph having t vertices and all possible t(t-1) ordered pairs of distinct vertices as its arcs.

Also for the set of α -orientations of a given graph G, such a flip graph exists, whose vertices are the different α -orientations of G and where two orientations are adjacent if one can be obtained from the other by reversing all arcs of some directed cycle. It is not hard to show that the so-defined flip graph is connected.

Given some definition of a connected flip graph, a natural computational problem to consider is the flip distance: Given two objects, what is the minimum number of flips needed to transform one into the other? In Chapter 4 we consider this problem for flip graphs of α -orientations and prove that deciding whether the flip distance between two α -orientations of a planar graph G is at most two is NP-complete. This also holds in the special case of perfect matchings, where flips involve alternating cycles. This problem amounts to finding geodesics on the common base polytope of two partition matroids, or, alternatively, on an alcoved polytope. It therefore provides an interesting example of a flip distance question that is computationally intractable despite having a natural interpretation as a geodesic on a nicely structured combinatorial polytope. We also consider the dual question of the flip distance between graph orientations in which every cycle has a specified number of forward edges, and a flip is the reversal of all edges in a minimal directed cut. In general, the problem remains hard. However, if we restrict to flips that only change sinks into sources, or vice-versa, then the problem can be solved in polynomial time. Here we exploit the fact that the flip graph is the cover graph of a distributive lattice. This generalizes a recent result by Zhang et al. [ZQZ19].

Part I concludes with Chapter 5. Here we are concerned with yet another algorithmic problem involving directed cycles. We study directed circuits of even size in a structure which generalizes directed graphs, known as *oriented matroids*. A *matroid* is a dependency structure which can be described by a finite ground set of *elements* and a collection of subsets of the ground set which are called *circuits*. Matroids can arise both from graphs, in which the element set consists of the edges of the graph, as well as from matrices over fields, in which the elements correspond to the column vectors. In the graphic case, the circuits correspond exactly to the edge-sets of the cycles, while in the column-matroid generated by a matrix, the circuits form the minimal linearly dependent sets of column vectors of the matrix.

In a similar vein, oriented matroids are an oriented version of matroids described by a finite element set and a collection of *oriented* circuits, which are special subsets of the ground set whose elements are equipped with either a plus- or a minus-sign. Again, every digraph gives rise to an oriented matroid, in which the oriented circuits are retrieved from the oriented cycles in the digraph. In this sense the oriented cycle structure of a digraph is also represented in its associated oriented matroid. Given a problem related to cycles in graphs or digraphs, it is therefore natural to consider its generalization to matroids or oriented matroids. In Chapter 5 we generalize the so-called *even cycle problem* to oriented matroids and study its generalization, the *even circuit problem*.

Until its resolution in 1999 by Robertson, Seymour, and Thomas [RST99], the even cycle problem was a popular open algorithmic problem for directed graphs, which asked whether there exists a polynomially bounded algorithm which can check whether a given digraph contains an even directed cycle or not. Motivated by this problem, Seymour and Thomassen [ST87] proved in 1987 that the even cycle problem is polynomially equivalent to the algorithmic problem of recognizing whether a given digraph is *even*. As mentioned before, a digraph is called *even* if every subdivision of it contains a directed cycle of even length. Alternatively, one may say that a digraph is even if for every labelling of its arcs with either 0 or 1, there is a directed cycle such that if we sum up the labels on the arcs of the cycle, we obtain an even number. The equivalence proof of Seymour and

Thomassen showed that in order to solve the even cycle problem, one may as well first try to obtain a structural understanding of the class of even digraphs, or alternatively, its complement, the *non-even digraphs*. Seymour and Thomassen made the first step in this direction by showing that a digraph is non-even if and only if it excludes all bioriented cycles of odd length as butterfly-minors. Here, a *butterfly-minor* is a concept of minors for directed graphs which is nowadays used frequently in structural digraph theory [JRST01, GT11, KK15]. Later on, Robertson, Seymour and Thomas built on the characterization by Seymour and Thomas to prove a generation theorem for the class of non-even digraphs. This eventually led to a recognition algorithm and hence a resolution of the even cycle problem.

The even circuit problem is the algorithmic problem that asks whether a given oriented matroid contains a directed circuit of even size. Finding a sensible definition of an algorithmic problem for matroids or oriented matroids in general, however, can be a difficult task, since matroids are lacking a compact encoding scheme, which would allow for a sensible measure of efficiency of an algorithm solving the problem. As a consequence, we argue in Chapter 5 that in order to formally treat the even circuit problem, one should restrict the problem to the class of *regular* oriented matroids. These are oriented matroids which allow for compact representations by totally unimodular matrices, a special class of matrices which is closely related to directed graphs. Regular oriented matroids are still a general class of matroids and extend the classes of oriented graphic matroids and oriented bond matroids, which both arise from directed graphs. In comparison to general matroids, they have the advantage that they allow for the usage of additional tools such as *circuit* bases, which is not possible in general oriented matroids. In the rest of the chapter, we then focus on the even circuit problem for regular oriented matroids and prove generalizations and supplements of the results of Seymour and Thomassen in this setting. We define non-even oriented matroids generalizing non-even digraphs and prove that the even circuit problem for regular matroids is polynomially equivalent to the problem of recognizing whether a given regular oriented matroid is non-even. Our main result is a precise characterization of the class of non-even oriented bond matroids in terms of forbidden minors, which complements the result of Seymour and Thomassen [ST87] that amounts to a characterization of the non-even oriented graphic matroids.

The Chapters of Part II

Part II, entitled *Dichromatic Number*, focuses on the study of acyclic colorings of digraphs. In particular, we prove conditions that exhibit well-structured classes of digraphs with bounded dichromatic number. In addition, motivated by earlier hardness results we revisit the algorithmic problem of determining the dichromatic number of a digraph and prove refined hardness results as well as a positive parametrized complexity result.

In Chapter 6 we start off by studying bounds on the dichromatic number of digraphs excluding the complete digraph K_t for some integer $t \ge 2$ as a strong minor. Strong minors in a very natural way generalize the ordinary undirected graph minors to directed graphs and have been introduced by Jagger [Jag96] in 1996. In 2015 Kim and Seymour [KS15] reconsidered the notion and proved that the set of semi-complete digraphs (digraphs in which every pair of vertices is adjacent) are well-quasi-ordered by the strong minor relation.

Colorings of minor-closed undirected graph classes such as the planar graphs, and in particular their chromatic number, have received widespread attention in the history of graph theory. Maybe the most intriguing and challenging open problem all over graph theory is Hadwiger's Conjecture, which was formulated back in 1943 by Hugo Hadwiger as a generalization of the famous 4-Color-Conjecture.

Conjecture 0.1 (Hadwiger [Had43], 1943). If $t \in \mathbb{N}$ and G is a graph such that $\chi(G) \ge t$, then G contains K_t as a minor.

Hadwiger's conjecture admits an elementary proof for all $t \leq 4$. From a result of Wagner [Wag37] it was known that the case t = 5 is equivalent to the Four-Color-Conjecture. After its resolution in 1976 by Appel and Haken, the next case t = 6 was solved in yet another breakthrough-result by Robertson, Seymour, and Thomas [RST93] in 1993. All the cases $t \geq 7$ remain open as of today.

Inspired by Hadwiger's conjecture, Axenovich, Girão, Snyder, and Weber [AGSW20] recently studied the existence of strong complete minors in digraphs with a given dichromatic number or given minimum degree. On the negative side, they showed that (in contrast to undirected graphs) digraphs of large minimum out- and in-degree might not have a strong \vec{K}_3 -minor. On the positive side they established that for every integer $t \geq 1$ there exists a smallest natural number $sm_{\vec{\chi}}(t)$ such that every digraph D with $\vec{\chi}(D) \geq sm_{\vec{\chi}}(t)$ contains a strong \vec{K}_t -minor, and proved the bounds

$$t+1 \le sm_{\vec{\chi}}(t) \le t4^t$$

for every $t \ge 2$. They then raised the problem of improving in particular the upper, exponential, bound. In Chapter 6, we solve this problem by showing the almost linear upper bound $sm_{\vec{\chi}}(t) = O(t(\log \log t)^6)$.

In Chapter 7 we study bounds on the dichromatic number of digraphs excluding a fixed digraph F as a *topological minor*. Here we say that a digraph D contains F as a topological minor if D contains a subdigraph isomorphic to a subdivision of F.

The corresponding problem for undirected graphs has been studied in particular when the excluded graph is the *t*-vertex-clique K_t with some $t \in \mathbb{N}$, but recently there have also been results for sparser graphs F, such as grids [HKL20]. In this context a famous problem known as Hajós' Conjecture stated that every graph G excluding K_t as a topological minor satisfies $\chi(G) \leq t-1$. Note that if true, this would be a strengthening of Hadwiger's conjecture (as topological minors specialize ordinary graph minors). Hajós' conjecture follows from Hadwiger's Conjecture for all $t \leq 4$, where the case t = 4 is a classical result by Dirac [Dir52]. In 1979, Catlin [Cat79] disproved Hajós' Conjecture for all values $t \ge 7$, by constructing t-chromatic graphs not containing a K_t -subdivision. The cases t = 5, 6 of the conjecture still remain open. The counterexamples of Catlin were later on strengthened by Erdős and Fajtlowicz [EF81] in 1981, who used the probabilistic method to show that large random graphs are counterexamples to Hajós conjecture with probability tending to 1. Their method also yields that for some positive constant c > 0 and every $t \ge 2$ there exists a graph G_t excluding K_t as a topological minor and such that $\chi(G_t) \ge c \frac{t^2}{\log t}$. On the positive side, Bollobás and Thomason [BT98] and independently Komlós and Szemerédi [KS96] proved that there exists a constant C > 0 such that every graph with minimum degree at least Ct^2 contains a K_t -subdivision. Using degeneracy this implies that every graph G excluding K_t as a subdivision satisfies $\chi(G) \leq Ct^2$. The precise asymptotics of the maximum possible chromatic number of a graph exluding K_t as a topological minor still remain open, but the truth was conjectured to lie with the lower bound, i.e. $\Theta\left(\frac{t^2}{\log t}\right)$, by Fox, Lee, and Sudakov [FLS13].

While for undirected graphs F it follows from the previous discussion that every graph with sufficiently large chromatic number contains an F-subdivision, the analogous statement for the dichromatic number of digraphs appears more difficult, since it is no longer true that digraphs D excluding a fixed digraph F must contain vertices of small degree. As mentioned above, there are digraphs D of arbitrarily large minimum out- and in-degree not containing a \vec{K}_3 -subdivision.

Nevertheless, Aboulker et al. [ACH⁺19] investigated this problem and managed to find a beautiful proof that there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that every digraph Dwith $\vec{\chi}(D) \geq f(t)$ contains a subdivision of the complete digraph \vec{K}_t , and hence of any other digraph on at most t vertices. This result is truly remarkable, since the dichromatic number thus far is the only natural digraph parameter one knows of that is capable of forcing a subdivision of *every* digraph F. However, the upper bound given by Aboulker et al. is $f(t) \leq 4^{t^2-2t+1}(t-1) + 1$, which is probably quite far from the truth. In their paper, Aboulker et al. also started a systematic study of the following problem:

For a given (not necessarily complete) digraph F, what is the smallest integer k such that every digraph D with $\vec{\chi}(D) \ge k$ contains an F-subdivision? This integer k is called the *(Dichromatic) Mader-number* of the digraph F and denoted by mader $_{\vec{\chi}}(F)$. Determining this value precisely turns out to be quite difficult, even if F is a rather sparse digraph. In this context, Aboulker et al. posed several open problems. They proved that mader $_{\vec{\chi}}(C) \le 2\ell - 3$ for every oriented cycle C of length $\ell \ge 3$ and conjectured that mader $_{\vec{\chi}}(C) = \ell$ for every such cycle C.

As the first main result of Chapter 7, we prove this conjecture by Aboulker et al. We further generalize this result and show that $\operatorname{mader}_{\vec{\chi}}(F)$ equals the number of vertices of F for all digraphs in an inductively defined class which we call *octus-digraphs*. These digraphs, among others, contain all bioriented forests and all orientations of cactus graphs.

Our second main result extends the classical result by Dirac [Dir52] that every 4chromatic graph contains a subdivision of K_4 to digraphs as follows: For every orientation F of K_4 , every digraph D with dichromatic number at least 4 contains a subdivision of F.

Our third and last main result proves an asymptotically tight (linear) upper bound on $\operatorname{mader}_{\vec{\chi}}(F)$ in terms of the number of vertices of F for a large class of (relatively sparse) digraphs F which we call *subcubic*. Here a digraph F is called subcubic if it has maximum degree at most 3, and maximum in- and out-degree at most 2.

The chapter concludes with a discussion of certain digraphs F which we call Maderperfect. We say that F is Mader-perfect if it holds that mader $_{\vec{\chi}}(F')$ equals the order of F'for every subdigraph $F' \subseteq F$. Our first two main results exhibit several digraphs which are Mader-perfect, and it is natural to ask for a structural characterization of such directed graphs. We obtain some preliminary results concerning digraphs in this class and pose several open problems.

In Chapter 8 we study bounds on the dichromatic number of digraphs excluding socalled *butterfly-minors*, a minor concept for digraphs that is used in structural digraph theory. We say that a digraph D contains another digraph F as a butterfly minor if F can be obtained from a subdigraph of D by contracting certain *butterfly-contractible arcs*. These contractions have the property that they leave the directed cycle structure of the digraph unaffected, and are therefore particularly suited for investigations related to directed cycles. The notion of butterfly-minors generalizes the topological minors discussed in the previous section, and hence, given a digraph D of sufficiently large dichromatic number, it will contain any fixed digraph F as a topological minor and hence also as a butterfly-minor. However, since butterfly-minors are more general than topological minors, the concrete bounds for the dichromatic number may improve by a lot if we consider excluded butterfly-minors instead of topological minors.

Butterfly-minors up to some extent resemble the ordinary minors used in undirected graph theory, where one is allowed to contract arbitrary edges (or connected subgraphs).

For $t \in \mathbb{N}$ let us denote by \mathcal{G}_t the inclusion-wise largest minor-closed class of t-colorable graphs. This class certainly does not contain K_{t+1} , and Hadwiger's conjecture states that \mathcal{G}_t is exactly the class of graphs excluding (only) K_{t+1} as a minor. In Chapter 8 we are motivated by the analogous question for directed graphs, where we replace the chromatic number by the dichromatic number.

Question 0.1. What is the largest butterfly-minor closed class \mathcal{D}_t of t-colorable digraphs?

While \mathcal{D}_1 is the class of acyclic digraphs, the first main contribution of Chapter 8 is to provide a precise characterization of \mathcal{D}_2 , showing that \mathcal{D}_2 constitutes exactly the class of *non-even digraphs*, which we have encountered earlier. We also consider the recently studied notion of *list colorings* of digraphs [BHL18] and show that the bestpossible upper bound on the list coloring number of the non-even digraphs is 3. Together with the characterization of non-even digraphs by Seymour and Thomassen [ST87], our main result states that every digraph D with $\vec{\chi}(D) \geq 3$ contains a bioriented cycle of odd length as a butterfly-minor. In contrast to graph minor containment, which induces a well-quasi-ordering on the set of graphs [RS04], the same is not true for directed graphs, as for instance bioriented cycles cannot be reduced to one another via butterfly minors.

In the chapter we move on by pointing out that Question 0.1 is closely related to the following problem: For $t \ge 1$, what is the smallest integer $bm_{\vec{\chi}}(t)$ such that every digraph D satisfying $\vec{\chi}(D) \ge bm_{\vec{\chi}}(t)$ contains the complete digraph \vec{K}_t as a butterfly-minor? As the second main result of Chapter 8 we give an asymptotically almost tight answer to this question by showing that $t + 1 \le bm_{\vec{\chi}}(t) = O(t(\log \log t)^6)$.

Interestingly, there is a bijection between directed graphs and bipartite graphs equipped with perfect matchings which plays a crucial role in our characterization of \mathcal{D}_2 . Using this bijection, butterfly minors can be generalized to another minor concept called *matching minors* which is used in structural matching theory. Similarly, the dichromatic number has an analogous concept in graphs with perfect matchings which we call the *matching chromatic number* $\chi(G, M)$ of a pair (G, M) consisting of an undirected graph G and a perfect matching M. In this setting, our result concerning \mathcal{D}_2 has a much nicer formulation: If $\chi(G, M) \geq 3$ for a bipartite graph G, then G contains $K_{3,3}$ as a matching minor. It is remarkable that in this setting we need to exclude only one single minor $K_{3,3}$ instead of the infinite anti-chain of bioriented cycles to describe a result of equal strength. This inspires us to formulate a conjecture very similar to Hadwiger's conjecture:

Conjecture 0.2 (Matching-Hadwiger). If $t \in \mathbb{N}$, G is a bipartite graph and M a perfect matching such that $\chi(G, M) \geq t$, then G contains $K_{t,t}$ as a matching minor.

The conjecture remains open for $t \ge 4$. Some evidence for the conjecture can be retrieved from our result concerning butterfly-minors: If $\chi(G, M) \ge O(t(\log \log t)^6)$, for a bipartite graph G and a perfect matching M, then G contains $K_{t,t}$ as a matching minor.

The chapter contains further results concerning the matching chromatic number of non-bipartite graphs, and an application of our results to so-called *forcing matchings* in *Pfaffian graphs*, an important graph class which is related to the problem of efficiently counting perfect matchings in graphs.

In Chapter 9 we aim at bounding the dichromatic number of digraphs which exclude a finite list of digraphs as *induced subdigraphs*. A subdigraph D' of a digraph D is called *induced* if there exists $X \subseteq V(D)$ such that D' consists of the vertex-set X and all arcs of D with both endpoints in X. Given a set \mathcal{F} of (di)graphs, we denote by $\operatorname{Forb}_{ind}(\mathcal{F})$ the set of (di)graphs containing no induced sub(di)graphs isomorphic to a member of \mathcal{F} . We are motivated by the following natural question posed recently by Aboulker, Charbit, and Naserasr [ACN20]: **Question 0.2.** For which finite sets \mathcal{F} of digraphs is the dichromatic number of digraphs in Forb_{ind}(\mathcal{F}) bounded?

If Forb_{ind}(\mathcal{F}) has bounded dichromatic number, then \mathcal{F} is called *heroic*.

For undirected graph coloring, the analogous question has led to the following famous Gyárfás-Sumner Conjecture, which remains open in general:

Conjecture 0.3 (Gyárfás-Sumner Conjecture, cf. [Gyá75, Sum81]). For every forest F and every $k \in \mathbb{N}$ there exists c(F, k) such that every graph G without induced subgraphs isomorphic to F or K_k satisfies $\chi(G) \leq c(F, k)$.

Aboulker, Charbit, and Naserasr recently investigated Question 0.2 in [ACN20] and made the following challenging conjecture, which resembles the Gyárfás-Sumner Conjecture in the setting of directed graphs.

Conjecture 0.4 (Directed Gyárfás-Sumner Conjecture, cf. [ACN20]).

- For every oriented star forest F, and every hero H, the set $\{ \vec{K}_2, F, H \}$ is heroic.
- For every oriented forest F, and every $k \in \mathbb{N}$, the set $\{\vec{K}_2, F, \vec{K}_k\}$ is heroic.

Here, \vec{K}_k denotes the transitive tournament on k vertices, and a *hero* is a special type of tournament introduced by Berger, Choromanski, Chudnovsky, Fox, Loebl, Scott, Seymour, and Thomassé in [BCC⁺13].

In Chapter 9 we present several new results verifying special cases of Conjecture 0.4. The smallest hero H for which the first part of Conjecture 0.4 gets interesting is the directed triangle. For this case Aboulker et al. conjectured that every directed triangle-free oriented graph in which the out-neighborhood of every vertex induces a tournament, is acyclically 2-colorable, without being able to prove any finite bound on the dichromatic number. Among other results, in Chapter 9 we prove this conjecture.

Chapter 10 deals with the computational complexity of acyclic colorings of digraphs. It was shown first by Bokal, Fijavz, Juvan, Kayll, and Mohar [BFJ⁺04] that deciding whether a given digraph D satisfies $\vec{\chi}(D) \leq 2$ is an NP-complete problem, and therefore this problem is expected not to admit a polynomial time solution.

In the first part of Chapter 10, we investigate refined hardness results for computing the dichromatic number. Strengthening the hardness result by Bokal et al., we present a reduction showing that for every $k \ge 2$ deciding whether a given digraph D is acyclically k-colorable remains NP-hard even if we restrict to input digraphs D which possess a bounded-size *feedback vertex set*, i.e., in which we can hit every directed cycle with a small set of vertices. We conclude that this hardness result rules out parametrized algorithms computing the dichromatic number with most well-known width measures for directed graphs, such as directed tree-width.

In the second part of Chapter 10 we present a positive algorithmic result, by presenting an FPT-algorithm that computes the dichromatic number of a given digraph in polynomial time, provided that its *directed modular width* is bounded by a constant. Directed modular width is a width measure for directed graphs introduced in [SW20], which behaves quite differently from other width measures in structural digraph theory in the sense that it may be small on dense networks but quite large on certain sparse digraphs.

The results of Chapter 10 were recently taken up by other researchers, whose results yield strengthenings of both our results. The first result is improved by Harutyunyan, Lampis, and Melissinos [HLM20], and the second by Gurski, Komander, and Rehs [GKR20]. In the chapter we give a short summary of these new results.

The Chapters of Part III

Part III, entitled *Other Coloring Concepts for Digraphs*, focuses on the study of colorings of digraphs which are different from, but sometimes related to, the acyclic colorings of digraphs treated in Part II.

In Chapter 11 we start off with the notion of *complete acyclic colorings* of undirected and directed graphs. This is inspired by the notion of the achromatic number of graphs, which is a variant of the chromatic number that was introduced in [HHP67] and further investigated in several works, we refer to [Edw97, HM97] for survey articles. To illustrate this parameter, let us consider the following naive (polynomial time) coloring algorithm which outputs a proper graph coloring of a given graph G as follows:

We start by assigning to every vertex in the graph a different color. Next, we repeatedly go through all different pairs $\{i, j\}$ of two distinct colors, and check whether there is an edge in the graph whose endpoints are colored by i and j. If we find a pair for which such an edge does not exist, we merge those two colors into one common color, preserving a proper graph coloring, and repeat. Once we do not find a mergeable pair any more, we return the current coloring.

Clearly, this is not a very sophisticated coloring algorithm, but as graph coloring is an NP-hard problem, we should not expect a polynomial algorithm to find an optimal coloring. Yet it is natural to ask *how* bad a coloring produced by the above procedure could be in the worst case.

To measure the performance of the algorithm, we may define the *achromatic number* $\Psi(G)$ of a graph as the worst-case number of colors used by a coloring generated with the above algorithm. It is not hard to see that a coloring $c : V(G) \to \{1, \ldots, k\}$ can be produced by the above algorithm if and only if it is a *complete* coloring, i.e., for every pair $i \neq j \in \{1, \ldots, k\}$, there exists an edge $e \in E(G)$ whose endpoints are colored i and j. Hence, the achromatic number is the *largest* number of colors that can be used in a complete coloring of G.

Clearly, the same coloring procedure can be applied to any other definition of a conflict free-coloring. Considering the dichromatic number, we may define in the same spirit a *complete acyclic coloring* of a digraph D as an acyclic coloring with the property that the subdigraph induced by the union of any two color classes contains a directed cycle, and the *adichromatic number* adi(D) of a digraph D as the maximum possible number of colors used by a complete acyclic coloring.

A coloring concept for undirected graphs that is directly analogous to the dichromatic number of a digraph is the *vertex-arboricity*. An *arboreal coloring* of an undirected graph is a vertex-coloring which avoids monochromatic cycles. Hence, the color classes in an arboreal coloring induce forests. The smallest number of colors that can be used for an arboreal coloring of a graph G is called the *vertex-arboricity* va(G).

Applying the above greedy-coloring scheme also to the arboreal colorings, we define in the same way a *complete arboreal coloring* of a graph G as an arboreal coloring such that the union of any two color classes induces a subgraph containing a cycle, and the *a-vertex arboricity* ava(G) as the largest number of colors used by a complete arboreal coloring.

After observing basic properties of the new parameters, in Chapter 11 we first establish so-called *interpolation theorems* for the adichromatic number and the a-vertex arboricity, which have been proved previously for the achromatic number. We then consider the relation between the parameters $\operatorname{adi}(D)$ and $\operatorname{ava}(G)$ and the *feedback vertex numbers* $\tau(D)$ and $\tau(G)$, which denote, respectively, the smallest size of a vertex-set in the (di)graph hitting all (directed) cycles. We show that the parameters adi and ava are bounded from above in terms of τ and consider the natural question whether an inverse relationship between the feedback-parameters and the coloring parameters also holds. While we construct digraphs with bounded adichromatic number and unbounded feedback vertex number, we prove as the main result of the chapter that there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that $\tau(G) \leq f(\operatorname{ava}(G))$ for every simple graph G, showing that the parameters $\operatorname{ava}(G)$ and $\tau(G)$ are tied to one another. As a consequence, we also show that a graph has large a-vertex-arboricity if and only if it admits an orientation whose adichromatic number is large. It is interesting that a corresponding relation between vertex-arboricity and dichromatic number is not known. In fact, it is subject of an open conjecture by Erdős and Neumann-Lara [Erd80] from 1980 that a graph has large chromatic number (or large vertex-arboricity³) if and only if it admits an orientation with large dichromatic number.

Chapter 12 deals with majority colorings of digraphs, a coloring concept introduced rather recently by Kreutzer, Oum, Seymour, van der Zypen, and Wood in [KOS⁺17], which is quite different from the coloring notions considered in the previous chapters. Given a digraph D, a majority k-coloring of D is a vertex-coloring $c: V(D) \to \{1, \ldots, k\}$ such that for every vertex $x \in V(D)$, at most half of the out-neighbors of x have the same color as x. This notion of coloring was originally inspired by a problem for neural networks [vdZ19]. While the definition in itself might seem a bit artificial, it is not uncommon in graph theory to seek partitions of the vertex-set of a given graph that satisfy certain degree restrictions, as they can have useful applications, for instance in order to find disjoint substructures in graphs. We refer to the Chapters 2 and 3, as well as to the overview article of Alon [Alo06] with further evidence for this claim.

Kreutzer et al. [KOS⁺17] proved as their main result that every digraph admits a majority 4-coloring. While there are digraphs requiring at least 3 colors in any majority coloring (for instance the directed triangle), Kreutzer et al. conjectured that in fact every digraph should admit a majority 3-coloring. They provided certain evidence for their claim by proving it for very dense digraphs, whose out- and in-degrees are large and relatively balanced, by using a random coloring approach. However, probabilistic methods seem incapable of proving majority 3-colorability for digraphs which are comparatively unbalanced, and admit vertices of both very large and very small degrees. In Chapter 12, we verify the majority 3-coloring conjecture by Kreutzer et al. for several structurally nice classes of digraphs, which can be both sparse and unbalanced, such as orientations of properly 6-colorable graphs. Instead of a probabilistic approach, our proofs are based on deterministic list coloring arguments. In addition, answering a question of Kreutzer et al., we prove that there is $\varepsilon > 0$ such that every digraph can be fractionally $(4 - \varepsilon)$ -colored.

In the last two Chapters of Part III, we deal with variants of acyclic colorings of digraphs called *fractional* and *circular* colorings. While the *fractional dichromatic number* of a digraph arises as a linear relaxation of an integer program defining the dichromatic number, we also consider the *star dichromatic number* of digraphs and the *circular vertex arboricity* of undirected graphs introduced in [HS19] and [WZLW11] respectively, which are rational refinements of the dichromatic number and the vertex arboricity, respectively.

In Chapter 13 we start by displaying basic properties of these parameters and then go on to study the complexity of computing these parameters. The results are unfortunately negative, as we prove strong NP-hardness results for all three parameters. Our results answer an open question from [HS19] as well as questions in the context of the work by Wang et al. [WZLW11]. The results resemble the hardness results achieved by Feder, Hell, and Mohar [FHM03] for the circular dichromatic number, which is another fractional coloring parameter for digraphs introduced by Bokal et al. [BFJ⁺04].

³Since every forest is properly 2-colorable, we have $va(G) \leq \chi(G) \leq 2va(G)$ for every graph G.

Our NP-hardness reductions make use of the concept of graph homomorphisms which are certain types of mappings between undirected graphs generalizing proper graph colorings. In order to show hardness results for digraphs, we introduce so called *circular homomorphisms*. These are mappings acting between digraphs which are appropriate for the analysis of acyclic colorings of digraphs and the star dichromatic number, and which generalize undirected graph homomorphisms in a natural way. We think that this concept could be of independent interest.

In the last Chapter 14 we study circular and fractional acyclic colorings of *planar* digraphs, i.e., digraphs whose underlying graphs are planar. Motivated by a famous conjecture of Jaeger [Jae84] from 1984 concerning circular colorings of planar graphs without short cycles as well as by a conjecture of Woodall [Woo78] concerning feedback arc-sets in planar digraphs, we investigate bounds on the fractional dichromatic number of planar digraphs without short directed cycles. Using a classical min-max packing result by Lucchesi and Younger [LY78] and basic results from the theory of certain set families called *clutters* we can show that (1) for every $\varepsilon > 0$ the fractional dichromatic number of planar digraphs of sufficiently large digirth is at most $1 + \varepsilon$ and (2) the fractional dichromatic number of a strongly planar digraph equals g/(g-1), where g denotes the length of its shortest directed cycle. Here, strongly planar digraphs refer to the subclass of planar digraphs which can be embedded in the plane without crossings such that sets of out-arcs and in-arcs of any vertex form two intervals in the cyclical ordering of incident arcs around the vertex. Interestingly, (2) means that the fractional dichromatic number of strongly planar digraphs can be computed in polynomial time, contrasting hardness results from Chapter 13 for general planar digraphs.

Chapter 1

Preliminaries

In this chapter we introduce basic notation and fundamental concepts used in the thesis. The reader very familiar with graph and matroid theory may skip large parts of this chapter and revisit it to look for special notations or definitions more targetedly.

General Mathematical Notation. Given a subset $X \subseteq S$ of a ground-set S (given by context), we will use the notation \overline{X} to denote the complement $S \setminus X$ of X. By 2^S we denote the power-set of S, and for every integer $k \ge 0$ we denote by $\binom{S}{k}$ the set of all k-element subsets of S. For two sets X, Y we denote their symmetric difference by $X+Y := (X \cup Y) \setminus (X \cap Y)$. The smallest natural number is 1, and we write $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For $n \in \mathbb{N}$ we denote $[n] := \{1, 2, \ldots, n\}$. For $k \in \mathbb{N}$ we use \mathbb{Z}_k to denote the cyclic group $(\mathbb{Z}/k\mathbb{Z}, +)$ and identify its elements with the numbers $\{0, 1, 2, \ldots, k-1\}$.

We use standard Landau notation such as $f = O(g), f = \Omega(g), f = o(g), f = \omega(g)$ to express relationships between the asymptotic growth of real-valued functions f and g. For $t \in \mathbb{R}, t > 0$ we denote by log t the *natural* logarithm of t.

Digraphs. For basic notation and facts about graphs and digraphs not included in the following we refer the reader to the standard reading [Die17, BJG08].

A graph is a tuple G = (V, E) consisting of a finite set V of vertices and a multi-set $E \subseteq \{\{u, v\} | u, v \in V\}$ of edges, which consist of unordered pairs of vertices or single vertices. Given a graph G = (V, E), we denote by V(G) := V the vertex-set of G, by E(G) := E the edge-(multi-)set of G, and by v(G) := |V(G)|, e(G) := |E(G)| their respective sizes. The value of v(G) is called the order of G, while e(G) is its size. Two vertices $u, v \in V(G)$ are called adjacent in G if $\{u, v\} \in E(G)$, a vertex $v \in V(G)$ and an edge $e \in E(G)$ are called incident if $v \in e$. For an edge $e = \{u, v\}$ in a graph, we also use the short notation e = uv and e = vu (symmetrically). For a vertex $v \in V(G)$, we denote by $E_G(v) \subseteq E(G)$ the set of edges in G incident with v, and by $N_G(v)$ the neighborhood of v, i.e., the set of vertices in G adjacent to v.

An edge e is called a *loop* if |e| = 1. Two edges $e_1 \neq e_2$ in G are called *parallel* if they are not loops and have the same endpoints. A graph G is called *loopless* if it contains no loops and *simple* if it contains neither loops nor pairs of parallel edges.

The degree of a vertex v in a graph G is defined as $d_G(v) := |E_G(v)| + |L_G(v)|$, where $L_G(v)$ is the set of loops incident with v. If G is a simple graph, then $d_G(v) = |N_G(v)|$. While using these notations, we sometimes drop the subscript G if it is clear by context. The largest respectively smallest degree achieved by a vertex in G is the maximum degree $\Delta(G)$ or the minimum degree $\delta(G)$ of G. The average degree $\overline{d}(G) := \frac{2e(G)}{v(G)}$ is the arithmetic mean of the vertex-degrees.

sym. difference

 \mathbf{graph}

vertex-set edge-set order size adjacent incident

neighborhood

loop parallel loopless simple degree

extremal degrees average degree A directed graph is a tuple D = (V, A) consisting of a finite set V of vertices of the digraph, and a multi-set $A \subseteq V \times V$ of arcs (also called directed edges or sometimes just edges) of the digraph D. Given a digraph D = (V, A), we denote by V(D) := V its vertex-set and by A(D) := A its arc-set (sometimes also called edge-set). We denote v(D) := |V(D)|, a(D) := |A(D)| and call v(D) the order and a(D) the size of D. Given an arc $e = (u, v) \in A(D)$, the vertex u is called the tail of e, while v is the head of e, denoted by tail(e), head(e). Given two vertices $u, v \in V(D)$ such that $(u, v) \in A(D)$, we say that u sees v and v is seen by u in D. Given two digraphs D', D we say that D' is a subdigraph of D and write $D' \subseteq D$, if $V(D') \subseteq V(D)$ and $A(D') \subseteq A(D)$. If in addition V(D') = V(D), then we say that D' is a spanning subdigraph of D.

The underlying graph U(D) of a digraph D = (V, A) is an undirected graph defined as U(D) = (V, E) where $E := \{\{u, v\} | (u, v) \in A\}$ is the multi-set obtained from the arc-multi-set of D by replacing every ordered pair of vertices with the corresponding unordered pair. Given an undirected graph G, an orientation of G is any digraph D satisfying U(D) = G.

Two vertices of a digraph D are called *adjacent* if they are adjacent in U(D), and a vertex $x \in V(D)$ and an arc $e \in A(D)$ are *incident* in D if $x \in \{\text{tail}(e), \text{head}(e)\}$. We say that e *leaves* x or *emanates* from x or that e is *outgoing* at x if x = tail(e), and that e *enters* x or that e is *incoming* at x if x = head(e). An arc $e \in A(D)$ is called a *loop* if tail(e) = head(e). Two distinct arcs $e_1, e_2 \in A(D)$ are called *parallel* if $\text{tail}(e_1) = \text{tail}(e_2) \neq \text{head}(e_1) = \text{head}(e_2)$, and *anti-parallel* if $\text{tail}(e_1) = \text{head}(e_2) \neq$ $\text{head}(e_1) = \text{tail}(e_2)$. Given a graph G, its *biorientation* or *bidirection* G is defined as the digraph obtained from G by replacing every edge with an anti-parallel pair of arcs, i.e., $A(G) := \{(u, v), (v, u) \mid \{u, v\} \in E(G)\}$. A digraph D is called *simple* if it does not contain loops nor parallel arcs. If in addition, D does not contain anti-parallel pairs of arcs (i.e., if U(D) is a simple graph), then we call D oriented or oriented graph.

During this thesis, we will mostly work with simple digraphs, there are exceptions however. To make these cases clear, at the beginning of every chapter we emphasize which kinds of graphs (multi-, simple) or digraphs (multi-, simple, oriented) we deal with.

For a vertex $v \in V(D)$, we denote by $E_D(v) := E_{U(D)}(v)$ and $N_D(v) := N_{U(D)}(v)$ the set of incident edges respectively the *neighborhood* of v in D and define $d_D(v) := |E_D(v)| = d_{U(D)}(v)$ as the *(total)* degree of v in D. We furthermore use the notation $E_D^+(v) := \{e \in A(D) | \text{tail}(e) = v\}$ and $E_D^-(v) := \{e \in A(D) | \text{head}(e) = v\}$; and denote by $d_D^+(v) := |E_D^+(v)|$ the out-degree and by $d_D^-(v) := |E_D^-(v)|$ the *in-degree* of v in D. We further let $N^+(v) := \{u \in V(D) | (v, u) \in A(D)\}$, $N^-(v) := \{u \in V(D) | (u, v) \in A(D)\}$ denote the out- and in-neighborhood of v in D. More generally, for a non-empty vertex-set U in a digraph D, we denote $N_D^+(U) := \bigcup_{u \in U} N_D^+(u) \setminus U$ and $N_D^-(U) := \bigcup_{u \in U} N_D^-(u) \setminus U$. Again, we may drop the subscript D while using these notations if it is clear from context. If D is a simple digraph, then $d_D^+(v) = |N_D^+(v)|, d_D^-(v) = |N_D^-(v)|$, and if D is oriented, then $d_D(v) = |N_D(v)| = d_D^+(v) + d_D^-(v)$. A vertex $v \in V(D)$ is called a sink if $d_D^+(v) = 0$ and a source if $d_D^-(v) = 0$. We say that a digraph D is Eulerian if $d_D^+(x) = d_D^-(x)$ for every $x \in V(D)$ and r-regular for an integer $r \ge 1$ if $d_D^+(x) = d_D^-(x) = r$ for every $x \in V(D)$.

The extremal out- and in-degrees of vertices in D are denoted as follows: $\delta(D), \Delta(D)$ are the minimum and maximum (total) degrees of D, $\delta^+(D)$ and $\delta^-(D)$ are the minimum out- and the minimum in-degree of D, and finally $\Delta^+(D)$ and $\Delta^-(D)$ denote the maximum out- and the maximum in-degree of D.

For a vertex-set $X \subseteq V(D)$ in a digraph D, we denote by D[X] the subdigraph of Dinduced by X, that is $D[X] := (X, A_X)$ with $A_X := \{e \in A(D) | tail(e), head(e) \in X\}$. For a set X of vertices or arcs in D, we denote by D - X the subdigraph obtained by deleting the objects in X from D. We write $D - v := D - \{v\}$ and $D - e := D - \{e\}$ for all

digraph

vertex-set arc-set order size tail and head subdigraph spanning

underlving

orientation

adjacent incident outgoing incoming

loop (anti-)parallel

biorientation

simple

oriented

neighborhood total degree

out-degree in-degree

out- and inneighborhood

sink and source Eulerian regular

extremal degrees

induced subdigraph deletion $v \in V(D)$ and $e \in A(D)$. If $A \subseteq A(D)$, we write D[A] for the subdigraph (V(D), A) of D induced by A.

A pair of (di)graphs is said to be *isomorphic* if there is a bijection between the vertexsets which also induces a bijective map between the edge- or arc-sets, respectively. For most of our work, it will not matter much whether two graphs are equal or isomorphic. In this sense, we use a slightly sloppy notation as follows: For $k \in \mathbb{N}$ we use K_k to denote any simple graph of order k in which any pair of distinct vertices is adjacent and call such a graph a (k-)clique or a complete graph of order k. By P_k we denote a path of order k, i.e. a graph isomorphic to $([k], \{\{i, i+1\} | 1 \leq i < k\})$. Furthermore, we denote by C_k for $k \geq 2$ any graph isomorphic to the graph $([k], \{\{i, i+1\} | 1 \leq i < k\} \cup \{1, k\})$ and call it a k-cycle or cycle of length k. By S_k we denote the graphs which are isomorphic to $([k+1], \{\{1, i\}, i = 2, \ldots, k+1\})$ and call them k-stars. By W_k we denote the graphs isomorphic to $([k+1], \{\{1, i\}, i = 2, \ldots, k+1\} \cup \{\{i, i+1\} | i = 2, \ldots, k\} \cup \{2, k+1\})$, and call them k-wheels. Finally, for $m, n \in \mathbb{N}$ we denote by $K_{m,n}$ any complete bipartite graph isomorphic to $([m+n], \{\{i, j\} | 1 \leq i \leq m < j \leq m + n\})$.

Similar definitions apply to digraphs: For $k \in \mathbb{N}$ we denote by $\vec{K}_k, \vec{P}_k, \vec{C}_k$ any digraphs isomorphic to a biorientation of a k-clique, a k-path or a k-cycle, and call them bioriented k-cliques or complete digraphs of order k, bioriented (bidirected) k-path and bioriented (bidirected) k-cycles or k-bicycles, respectively. Further important digraphs are the directed k-paths \vec{P}_k which are isomorphic to $([k], \{(i, i+1)|1 \leq i < k\})$, the directed k-cycles \vec{C}_k isomorphic to $([k], \{(i, i+1)|1 \leq i < k\} \cup \{(k, 1)\})$, as well as \vec{K}_k , the transitive tournaments of order k isomorphic to $([k], \{(i, j)|1 \leq i < j \leq k\})$. We denote by S_k^+, S_k^- the orientations of the graphs S_k where all arcs incident to the central vertex (with label 1) are oriented outwards (respectively inwards). Similarly we denote by W_k^+, W_k^- the orientations of the graphs W_k where all arcs incident to the central vertex (with label 1) are oriented outwards (respectively inwards) and where the cycle spanned by the remaining vertices is oriented cyclically. Lastly, for $m, n \in \mathbb{N}$ we denote by $\vec{K}_{m,n}$ an orientation of $K_{m,n}$ isomorphic to $([m + n], \{(i, j)|1 \leq i \leq m < j \leq m + n\})$. These definitions are illustrated in Figure 1.1.

A walk W in a digraph is an alternating sequence $W = v_1, e_1, v_2, \ldots, v_{k-1}, e_{k-1}, v_k$ of vertices and arcs such that $e_i = (v_i, v_{i+1})$ or $e_i = (v_{i+1}, v_i)$ for all $1 \le i \le k - 1$. We will call v_1 the first vertex of W, v_2 the second vertex of W, v_k the last vertex of W, etc. The walk is called *closed* if $v_k = v_1$. Given a walk W in a digraph, we will often use the short notation $W = v_1, \ldots, v_k$ instead of $W = v_1, e_1, \ldots, e_{k-1}, v_k$ to express that the walk W visits the vertices v_1, \ldots, v_k in this order. If a walk does not repeat arcs, then it is called a *trail*. If it does not repeat vertices, then it is called *oriented path*. If in a walk $v_1, e_1, v_2, \ldots, v_{k-1}, e_{k-1}, v_k$ the vertices v_1, \ldots, v_{k-1} are pairwise distinct and $v_k = v_1$, then we speak of an *oriented cycle* in the digraph.

A directed walk in a digraph is a walk $v_1, e_1, v_2, \ldots, v_{k-1}, e_{k-1}, v_k$ such that $tail(e_i) = v_i$, head $(e_i) = v_{i+1}$ for all $1 \le i \le k-1$. We have a closed directed walk if $v_k = v_1$. If a directed walk does not repeat arcs, then it is called a *directed trail*. If it does not repeat vertices, then it is called a *directed path*. If in a directed walk $v_1, e_1, v_2, \ldots, v_{k-1}, e_{k-1}, v_k$ we have $k \ge 2$, the vertices v_1, \ldots, v_{k-1} are pairwise distinct and $v_k = v_1$, then we speak of a *directed cycle* in the digraph.

It is clear that oriented paths and oriented cycles in a digraph D describe certain subdigraphs of D isomorphic to an orientation of a path or an orientation of a cycle. In this sense, we will not distinguish between the vertex-arc-list and the subgraph-definition of oriented paths and cycles. We may therefore use notations like V(P), V(C), A(P), A(C)

33

isomorphic

clique path graph

 $_{\rm star}$

wheel $K_{m,n}$

cycle graph

biclique bioriented path bicycle directed path directed cycle transitive tournament out/in-stars out/in-wheels

walk

closed walk trail oriented path oriented cycle directed walk

first/last vertex

directed trail directed path

directed cycle



Figure 1.1: From top left to bottom right going in row-wise order the following digraphs are depicted: The complete digraph of order 4, the bioriented 4-path, the 4-bicycle, S_3^+ , W_3^+ , the transitive tournament of order 4, the directed 4-path, the directed 4-cycle, S_3^- , W_3^- , as well as the complete bipartite graph $K_{3,3}$ and the one-direction $\vec{K}_{3,3}$.

to refer to the vertex- or arc-sets of an oriented path or an oriented cycle P or C in D. We will also often refer to oriented paths and cycles in digraphs simply as paths and cycles.

Given a directed path $P = v_1, e_1, v_2, \ldots, v_{k-1}, e_{k-1}, v_k$, we say that P is a *dipath* from v_1 to v_k (v_1 - v_k -dipath for short). We denote by |P| = k - 1 the length of P (i.e. its number of arcs). Similarly, for a directed cycle C we define its length as |C| := a(C) = v(C). A directed cycle of length 2 in a digraph corresponds to an anti-parallel pair of arcs and is called a *digon*. Given two vertices x, y on a path P, we denote by P[x, y] = P[y, x] the subpath of P with endpoints x and y. A vertex v in a digraph D is said to be reachable from a vertex u if there exists a u-v-dipath. In this case, the distance from u to v in D is defined as the length of a shortest u-v-dipath. For a pair of walks P, Q such that the first vertex of Q is the last vertex of P, we denote by $P \circ Q$ the concatenation of P and Q, i.e. the walk obtained by first traversing P and then traversing Q. When P (resp. Q) consists of a single arc (x, y), we will sometimes write $(x, y) \circ Q$ (resp. $P \circ (x, y)$) instead of $P \circ Q$. The directed girth or digirth¹ of D, i.e. the minimum length of a directed cycle in D, is denoted by $\vec{q}(D)$. For a *directed* cycle C and two vertices $x, y \in V(C)$, we denote by C[x, y] the segment of C which forms a dipath from x to y. Two paths or cycles in a digraph are called *vertex-disjoint* if they do not share a vertex and two paths are *internally vertex-disjoint* if they only intersect at common endpoints.

A vertex-set $X \subseteq V(G)$ in a graph G is called *independent* if the vertices in X are pairwise non-adjacent. Similarly a vertex-set $X \subseteq V(D)$ in a digraph D is *independent* if it is independent in U(D). By $\alpha(G)$ (resp. $\alpha(D)$) we denote the largest size of an independent set in G (resp. D) and call it the *independence number*.

A graph G is called a *forest* or *acyclic* if it contains no cycles, and a *tree* if it is also connected. An *out-(in-)arborescence* is an orientation of a rooted tree in which all arcs

length

dipath

digon reachable distance

concatenation

directed girth

internally disjoint independent set

independence number forest, tree arborescence

¹If D does not contain directed cycles, then $\vec{q}(D) := \infty$.

are directed away from (towards) its root. A (di)graph is *bipartite* if its vertex-set can be bipartitioned into at most two independent sets.

We call a vertex-set $X \subseteq V(G)$ acyclic, if G[X] is a forest. A feedback vertex-(edge-)set of G is a set X of vertices (edges) such that G - X is acyclic. By $\tau(G)$ we denote the smallest size of a feedback vertex set in G.

Analogously, a digraph is called *acyclic* if it does not contain directed cycles. A digraph is acyclic if and only if it admits a *topological vertex-ordering*, i.e., a linear ordering v_1, \ldots, v_n of V(D) such that $(v_i, v_j) \in A(D) \Rightarrow i \leq j$ for all $i, j \in \{1, \ldots, n\}$.

We call a vertex-set $X \subseteq V(D)$ acyclic, if D[X] is acyclic. A feedback vertex-(arc-)set of D is a set X of vertices (arcs) such that D - X is acyclic. By $\tau(D)$ we denote the smallest size of a feedback vertex set in D.

A digraph D is called *weakly connected* (or just *connected*) if every two vertices can be connected by a path (i.e., if U(D) is connected), and is called *strongly connected* if for every ordered pair of vertices $(x, y) \in V(D) \times V(D)$, x can reach y in D. The (vertex-sets of the) maximal strongly connected subgraphs of a digraph D are called *(strong) components* and induce a partition of V(D).

For a natural number $k \in \mathbb{N}$, a graph G is called k-vertex (edge)-connected if it has at least k + 1 vertices (edges) and for every set K of at most k - 1 vertices (edges) of G, the digraph G - K is connected. Analogously, a digraph D is called *strongly k-vertex* (arc)-connected if it has at least k + 1 vertices (arcs) and for every set K of at most k - 1 vertices (arcs) of D, the digraph D - K is strongly connected. The smallest integer $k \ge 0$ such that a digraph D is strongly k-vertex (arc)-connected is called the *strong vertex* (arc)-connectivity of D, denoted by $\kappa(D)$ (resp. $\kappa'(D)$).

Next we quickly recall Menger's Theorem, which is a fundamental result relating connectivity in digraphs with the existence of disjoint directed paths. Menger's Theorem will be used repeatedly in the course of this thesis.

Theorem 1.1 (cf. [Men27]). Let D be a digraph and $u, v \in V(D)$ be distinct vertices such that $(u, v) \notin A(D)$. Then for every $k \in \mathbb{N}$, either there are k internally vertex-disjoint u-v-dipaths in D, or there is a set $K \subseteq V(D) \setminus \{u, v\}$ such that |K| < k and D - K contains no u-v-dipath.

Given a digraph D and two (not necessarily disjoint) subsets $A, B \subseteq V(D)$, an A-Bdipath is a directed path in D which starts at a vertex of A, ends at a vertex of B, and is internally vertex-disjoint from $A \cup B$ (an A-B-dipath is allowed to consist of a single vertex in $A \cap B$). If A or B are of size one, say $A = \{u\}$ or $B = \{u\}$, then we will simply write "u-B-dipath" or "A-u-dipath", respectively.

We will frequently use the following well-known direct consequences of Menger's Theorem for directed graphs.

Theorem 1.2 (cf. [Men27] and [Gör00]). Let D be a digraph and $k \in \mathbb{N}$.

- (i) If D is strongly k-vertex-connected, then for any two subsets $A, B \subseteq V(D)$ such that $|A|, |B| \ge k$, there are k pairwise vertex-disjoint A-B-dipaths.
- (ii) If $v \in V(D)$ and $A \subseteq V(D) \setminus \{v\}$, then either there are k different v-A-dipaths (resp. A-v-dipaths) which pairwise only intersect at v, or there is $K \subseteq V(D) \setminus \{v\}$ such that |K| < k and such that in D - K there is no dipath starting at v and ending in A (resp. starting in A and ending at v).

Proof. For (i), consider an auxiliary digraph H obtained from D by adding two distinct artificial vertices $v_A, v_B \notin V(D)$, and adding all the arcs (v_A, y) with $y \in A$ and (y, v_B)

bipartite

acyclic set feedback set

acyclic topological ordering acyclic set

connected strongly connected strong component

(strongly) k-connected

strong connectivity

Menger

with $y \in B$. We claim that no $K \subseteq V(D) = V(H) \setminus \{v_A, v_B\}$ of size |K| < k can separate v_B from v_A . Indeed, for each such K we have $A \setminus K, B \setminus K \neq \emptyset$ (because $|A|, |B| \ge k$), and since D was assumed to be strongly k-vertex-connected, there is an A-B-dipath in D - K, and hence a v_A - v_B -dipath in H - K. We can now apply Theorem 1.1 to the vertices v_A and v_B to conclude that there must be a collection \mathcal{P} of k internally vertex-disjoint v_A - v_B -dipaths in H. For each path $P \in \mathcal{P}$, let $v_{P,A} \in A$ be the last vertex in A we meet when traversing P from v_A towards v_B , and let $v_{P,B}$ be the first vertex in B we meet when traversing $P[v_{P,A}, v_B]$ starting from $v_{P,A}$. Now the family $\mathcal{P}' := \{P[v_{P,A}, v_{P,B}] | P \in \mathcal{P}\}$ forms a collection of k vertex-disjoint A-B-dipaths in D, proving the claim.

Let us now derive (ii) for the case of v-A-dipaths (the proof for A-v-dipaths is symmetrical). Consider an auxiliary digraph H obtained from D by adding an artificial vertex $v_A \notin V(D)$ and adding the arc (y, v_A) for every $y \in A$. The claim now follows by applying Theorem 1.1 to the vertices v and v_A in H: either there are k internally vertex-disjoint v- v_A -dipaths in H, and by cutting these dipaths as soon as they hit A (similarly as we do in the proof of the first item) we obtain k distinct v-A-dipaths in D which pairwise only share the vertex v; or we can hit all v- v_A -dipaths in H with a subset $K \subseteq V(H) \setminus \{v, v_A\} = V(D) \setminus \{v\}$ of size |K| < k which means that there are no dipaths in D - K starting at v and ending in A. This proves the second item.

Given a directed graph D and a subset $U \subseteq V(D)$ of vertices, we denote by $\partial(U)$ the set of all arcs in A(D) having one end in U and the other in $\overline{U} := V(D) \setminus U$. A *cut* of a edge-cut digraph is a non-empty arc set of the form $\partial(U)$ for some $U \subseteq V(D)$. A cut of D is called minimal or equivalently bond, if there is no other cut of D properly contained in it. By bond $\partial^+(U)$, we denote the set of all arcs in A(D) directed from U to \overline{U} . If $\partial(U) = \partial^+(U)$, we call $S = \partial(U)$ a directed cut or dicut induced by U, and U is referred to as a cut set of S. directed cut In the case that D is weakly connected, the cut set is uniquely determined by the dicut. A minimal directed cut or directed bond is a directed cut which is minimal. A dijoin in a directed bond dijoin digraph is a set of arcs intersecting every directed cut (resp. every directed bond). Minor Concepts for Graphs and Digraphs. Given a graph H, an H-minor model minor model is any graph G whose vertex-set admits a partition into non-empty sets $X_h, h \in V(H)$ with the following properties: For every $h \in V(H)$, the graph $G[X_h]$ is connected and for every edge $e = \{h_1, h_2\} \in E(H)$, there exists a representing edge $f = \{u, v\} \in E(G)$ such that $u \in X_{h_1}$ and $v \in X_{h_2}$. Note that in case that H is a multi-graph, we require that the representing edges of distinct parallel edges or loops in E(H) are also pairwise distinct. Given two undirected graphs G and H, we say that G contains H as a minor and write minor $G \succeq H$ if G contains a subgraph which is an H-minor model. An alternative definition of graph minors uses *edge contractions*: Given a graph G and an edge $e = \{u, v\} \in E(G)$, the contraction

> graph G/e obtained from G by contracting e is defined by $V(G/e) := (V(G) \setminus \{u, v\}) \cup \{x_e\}$, where $x_e \notin V(G)$ is a newly created vertex representing the contracted edge e, and

$$E(G/e) = E(G - \{u, v\}) \cup \{\{x, x_e\} | e = \{x, y\} \in E(G), x \notin \{u, v\}, y \in \{u, v\}\}.$$

When working in settings where parallel edges or loops are allowed, we will interpret G/e as a multi-graph. In contrast and with a slight abuse of notation, when working in the setting of simple graphs only, we will use the same symbol G/e to denote the simple graph obtained from G by identifying the endpoints of e, and then deleting all loops and identifying all parallel edges that have arisen by performing the contraction at e. Which of the two definitions is used will be clear from context. Using the notion of edge contractions, we may give the equivalent definition of graph minors as follows: A
inite seque

graph G contains another graph H as a minor if and only if there exists a finite sequence $G = G_0, G_1, \ldots, G_k$ of graphs such that for every $1 \leq i \leq k$ we can obtain G_i from G_{i-1} by a vertex-deletion, an edge-deletion or an edge-contraction, and such that G_k is isomorphic to H. A class \mathcal{G} of graphs is *minor-closed* if for every graph $G \in \mathcal{G}$ and every graph H such that $G \succeq H$ we have $H \in \mathcal{G}$. Graph minors are certainly one of the most important tools in graph theory and have been heavily influential to the development of structural graph theory. One of the main reason for the popularity and importance of graph minors may be that many important graph parameters, such as *tree-width*, as well as many important topologically defined graph classes, such as the planar graphs or the graphs embeddable on a fixed surface, are not only closed under the subgraph operation, but are also well-behaved with respect to taking minors. It is not hard to see that minor containment defines a quasi-ordering (a reflexive and transitive binary relation) on the set of graphs. In contrast, it is not an easy question whether this ordering is also a so-called well-quasi-order, i.e., if or not there exist infinite sequences of graphs G_1, G_2, \ldots with the property that G_i is not a minor of G_j for any ordered pair $(i, j) \in \mathbb{N}^2$ (so-called antichains). A famous conjecture attributed to Wagner stated that the minor relation is indeed a wellquasi-ordering. This fundamental question in graph theory has inspired much research in the late 20th and the early 21th century, culminating in the famous Graph Minors Project by Robertson and Seymour [RS04], a sequence of 20 papers spanning from 1984 to 2004 which resulted in a proof of Wagner's conjecture, and, on its way, produced a whole structural theory of minor-closed graph classes. These characterizations still find applications in present research.

Another important notion of minors in graphs called *topological minors* also originally stems from topological graph theory: Given a graph drawn in the plane, the topological properties of this planar embedding, and in particular whether or not the drawing is crossing-free, do not change by subdividing the curves representing these edges with further points (vertices) placed on the interior of these curves. Combinatorially, we define a *subdivision* of a graph G to be any graph obtainable from G by (repeatedly) replacing an edge e = uv of G by a path which is internally vertex-disjoint from G but intersects G at its endpoints, u and v. A *topological minor* of a graph G is defined as any graph H such that G contains as a subgraph a subdivision of a graph isomorphic to H. The fact that planarity is preserved by topological minors allows for beautiful combinatorial characterizations of topologically defined graph classes by excluded topological minors, such as Kuratowski's characterization of planar graphs [Kur30].

Given the tremendous success of undirected graph minors in history, it is only natural to try and establish an analogous structural theory of digraphs based on a suitable concept of *digraph minors*. However, it has turned out that there is not only one natural generalization of minor concepts from undirected graphs to digraphs, and several concurring notions have been studied in the literature. The maybe least debatable minor concept of minors for digraphs are *topological minors*, whose definition is basically the same as for undirected graphs.

Definition 1.1 (Topological minor). Given a digraph D, a subdivision of D is any digraph D' obtained from D by replacing every arc $e = (u, v) \in A(D)$ by a directed path P_e starting in u and ending in v, such that the paths $P_e, e \in A(D)$ are pairwise internally vertexdisjoint and intersect V(D) only in their endpoints. Given a subdivision D', the vertices in D' originally contained in D are called branch vertices, while the vertices internal to a path $P_e, e \in A(D)$ are called subdivision vertices. Let D_1, D_2 be digraphs. We say that D_2 is a topological minor of D_1 and write $D_1 \succeq_t D_2$ if D_1 contains a subdivision of a digraph isomorphic to D_2 as a subdigraph.

subdivision

topological minor

digraph subdivision

branch vertex subdivision vertex topological digraph minor

tree-width

minor-closed

arc contraction

Given a digraph D and an arc $e = (u, v) \in A(D)$, the digraph D/e obtained from Dby contracting e is defined by $V(D/e) := (V(D) \setminus \{u, v\}) \cup \{x_e\}$, where $x_e \notin V(G)$ is a newly created vertex representing the contracted arc e, and

$$A(D/e) = A(D - \{u, v\}) \cup \{(x, x_e) | e = (x, y) \in A(D), x \notin \{u, v\}, y \in \{u, v\}\}$$
$$\cup \{(x_e, y) | e = (x, y) \in A(D), x \in \{u, v\}, y \notin \{u, v\}\}.$$

More generally, for a set $A \subseteq A(D)$ of arcs we denote by D/A the digraph obtained from D by successively² contracting the arcs in A. Again, depending on whether we work in the realm of simple digraphs or multi-digraphs, we will identify or keep parallel arcs or loops arising by arc contractions.

The second minor notion which we discuss in this thesis are the so-called *strong minors*, which aim to generalize the definition of undirected graph minors by replacing connectivity in graphs with strong connectivity in directed ones. Strong minors have been introduced in 1996 by Jagger [Jag96] and further studied by Kim and Seymour [KS15], who proved that the strong minor relation induces a well-quasi ordering on the set of *semicomplete digraphs* (digraphs in which every pair of vertices is adjacent). The notion has also been studied more recently by Axenovich, Girão, Snyder, and Weber [AGSW20].

strong minor model

strong minor

Definition 1.2 (Strong minor). Let H be a digraph. A strong H-minor-model is a digraph D equipped with a partition $V(D) = \bigcup_{h \in V(H)} V_h$ of its vertex-set such that $D[V_h]$ is strongly connected for every $h \in V(H)$, and such that for every $\operatorname{arc}(h_1, h_2) \in A(H)$ there exists $e \in A(D)$ with $\operatorname{tail}(e) \in V_{h_1}$ and $\operatorname{head}(e) \in V_{h_2}$. We say that a digraph D contains a strong H-minor and write $D \succeq_s H$ if D contains a strong H-minor-model as a subdigraph. Less formally, one may define a strong minor of a digraph D to be any digraph obtainable from D by repeated application of the followings steps:

- Deleting vertices or arcs, and
- Contracting the arc-set of a strongly connected subdigraph.

The last notion of digraph minors which plays an important role in this thesis are the *butterfly-minors*, which is the minor notion for digraphs used most frequently in structural digraph theory. One of the first appearances of this notion is in [JRST01], and since then it has reoccurred frequently in digraph literature; let us mention the papers [GT11], [KK15],[GKKK20] for a small selection of results based on butterfly-minors.

The basic idea of butterfly-minors is to only allow for contractions of certain special *contractible arcs*. The reason for that is that when focussing on structural digraph theory one is mainly concerned with strong connectivity and directed substructures such as directed cycles and directed paths. However, contracting an arbitrary arc in a digraph could possibly hide information, in the sense that new directed cycles might be created that did not exist in the digraph before the contraction. In order to be able to keep track of this directed information, one insists on contracting only such arcs which have the property that any directed path or cycle in the digraph.

More precisely, for a butterfly-contractible edge e within a digraph D every directed cycle in D/e either equals one in D or induces one in D by incorporating e, a property which does not necessarily hold if arbitrary edges are contracted.

Definition 1.3 (Butterfly minor). Let D be a digraph. An arc $e = (u, v) \in A(D)$ is called contractible arc contractible if $d_D^+(u) = 1$ or $d_D^-(v) = 1$, i.e., if e is the only arc leaving u or the only arc

 $^{^{2}}$ Up to isomorphism, the outcome of this process does not depend upon the order in which the arcs are contracted.

entering v. We say that a digraph D_1 contains another digraph D_2 as a butterfly minor butterfly minor and write $D_1 \succeq_b D_2$, if D_2 is isomorphic to a digraph which can be obtained from D_1 via a finite sequence of arc-deletions, vertex-deletions, and contractions of contractible arcs.

We note that it can be proved by induction that if F is a *subcubic* digraph, that is, a digraph satisfying $\Delta(F) \leq 3, \Delta^+(F), \Delta^-(F) \leq 2$, then $D \succeq_b F$ if and only if $D \succeq_t F$.

Let us also note at this point that $\succeq_t, \succeq_s, \succeq_b$ are transitive and therefore define quasiorders on the set of digraphs.

Directed tree-width. The *tree-width* tw(G) of an undirected graph G is a positive integer measuring the structural similarity between G and a tree or a forest. Intuitively, graphs with small tree-width are such graphs which are nowhere well-connected in the sense that they have bounded-size vertex-separators all over the place. For example in forests, which constitute exactly the graphs with $tw(G) \leq 1$, the deletion of any non-leaf vertex splits the graph into further connected components.

Tree-width has been a very successful and highly popular graph parameter both in theoretical computer science and structural graph theory. For computer scientists, treewidth is and has been interesting because a famous result by Courcelle [Cou90] showed that many natural algorithmic problems on graphs which are known to be NP-complete on graphs in general, such as the problem of detecting a Hamiltonian cycle in a given graph, become tractable and in fact admit linear-time algorithms for input graphs of tree-width at most k, for any fixed value k. Such algorithms which perform in polynomial time given that some parameter is bounded are known as *parametrized algorithms*.

In structural graph theory, tree-width has been important mainly because of the famous Grid Theorem proved by Robertson and Seymour [RS86], which has later turned out to be one of the major cornerstones for understanding minor-closed graph classes. Indeed, tree-width and graph minors act nicely together in the sense that tree-width is minormonotone: $\operatorname{tw}(G) \geq \operatorname{tw}(H)$ for all graphs G, H such that $G \geq H$. The Grid Theorem states that qualitatively, what makes the tree-width of a graph large, is its containment of a large grid minor. The precise statement is the following: For every $k \in \mathbb{N}$, there exists $t(k) \in \mathbb{N}$ such that every graph G either contains a $k \times k$ -grid graph as a minor³, or G has tree-width at most t(k).

Motivated by the success of undirected tree-width, there have been several attempts to define a structural digraph parameter which unifies nice structural properties such as monotonicity under taking digraph minors, as well as algorithmic applications for the design of parametrized algorithms. We refer the interested reader to $[GHK^{+}10]$ for an interesting theoretical discussion of the possible success of such attempts. In this thesis we only deal with the oldest and most important of these generalizations of tree-width to directed graphs, the *directed tree-width*. This parameter was introduced by Johnson, Robertson, Seymour, and Thomas in 2001 [JRST01]. The basic idea of directed tree-width is to measure the richness of the directed cycle structure of a digraph, or how different the digraph is structurally from an acyclic digraph. In particular, digraphs of directed tree-width 0 are exactly the acyclic digraphs. We now give the formal definition.

For two distinct vertices $r, r' \in V(T)$ in an out-arborescence T we write r < r' if r'is reachable from r in T. We write $r \leq r'$ to express that r = r' or r < r'. If $e \in A(T)$ whose head is r, then we write e < r' if $r \leq r'$. For a digraph D and a subset $Z \subseteq V(D)$, we call $S \subseteq V(D) \setminus Z$ a Z-normal set if the vertex set of every directed walk in D-Z

subcubic

Grid Theorem

³The $k \times k$ -grid is the graph with vertex set $[k] \times [k]$ and where $(i, j), (i', j') \in [k] \times [k]$ are adjacent iff |i - i'| + |j - j'| = 1.

which starts and ends in S is completely contained in S. Note that every Z-normal set is a union of certain strong components of D - Z.

Definition 1.4 (cf. [JRST01]). A directed tree decomposition of a digraph D is a triple (T, β, γ) , where T is an out-arborescence, $\beta : V(T) \rightarrow 2^{V(D)}$, $\gamma : A(T) \rightarrow 2^{V(D)}$ are functions such that

(i) $\{\beta(t)|t \in V(T)\}$ is a partition of V(D) into non-empty sets and

(ii) if $e \in A(T)$, then $\bigcup \{\beta(t) | t \in V(T), e < t\}$ is $\gamma(e)$ -normal.

For every $t \in V(T)$ we denote

$$\Gamma(t) := \beta(t) \cup []\{\gamma(e) | e \in A(T), e \sim t \},$$

where $e \sim t$ expresses that e is incident with t in T. Then the width of the treedecomposition (T, β, γ) is defined as $w = \max_{t \in V(T)} |\Gamma(t)| - 1$. The directed tree-width of D, denoted by dtw(D), is the smallest possible width of a directed tree-decomposition.

The sets $\beta(t)$ are called *bags*, while the sets $\gamma(e)$ are the *guards* of the tree-decomposition. As for undirected tree-width, there exists a duality theorem for directed tree-width and the containment of large grid-like substructures as butterfly-minors, known as the *Directed Grid Theorem*. This result was conjectured by Johnson et al. in [JRST01] and proved in a breakthrough-result by Kawarabayashi and Kreutzer [KK15]. Let us specify what we mean by "grid-like substructures" in digraphs.

cylindrical grid **Definition 1.5** (Cylindrical Grid, cf. [GKKK20]). A cylindrical grid of order k for $k \in \mathbb{N}$, $k \geq 1$ is a digraph G_k consisting of k directed cycles C_1, \ldots, C_k , pairwise vertex-disjoint, together with a set of 2k pairwise vertex-disjoint directed paths P_1, \ldots, P_{2k} , satisfying the following properties:

- (i) P_i has length k-1 for $0 \le i \le 2k$, C_j has length 2k for $1 \le j \le k$.
- (ii) each path P_i has exactly one vertex in common with each cycle C_j and both endpoints of P_i are in $V(C_1) \cup V(C_k)$,
- (iii) the paths P_1, \ldots, P_{2k} appear on C_i in this order, and
- (iv) for odd i the cycles C_1, \ldots, C_k occur on P_i in this order and for even i they appear in reverse order C_k, \ldots, C_1 .

Definition 1.6 (Cylindrical Wall, cf. [GKKK20]). An elementary cylindrical wall of order k for $k \in \mathbb{N}$ is the planar digraph W_k obtained from the cylindrical grid G_k of order k by replacing every vertex $v \in V(G_k)$ by two new vertices v_s, v_t connected by an arc (v_s, v_t) such that for every arc $(v, w) \in A(G_k)$, we have the corresponding arc $(v_t, w_s) \in A(W_k)$. A cylindrical wall of order k is any digraph isomorphic to a subdivision of W_k .

We remark that G_k is a butterfly-minor of W_k obtained by contracting all the splitedges $(v_s, v_t), v \in V(G_k)$, while it can be observed that $W_k \subseteq G_{2k}$ for every $k \in \mathbb{N}$. It is convenient to imagine grids, elementary cylindrical walls and their subdivisions as embedded in the plane, as depicted in Figure 1.2. Every cylindrical wall W of order kcontains in this canonical embedding k pairwise vertex-disjoint "vertical" directed cycles Q_1, \ldots, Q_k , as well as 2k pairwise vertex-disjoint directed "horizontal" directed paths $P_i^1, P_i^2, i = 1, \ldots, k$ which are alternately directed from left to right or from right to left. To make references to specific vertices in a cylindrical wall easier, we assign coordinates

40

directed tree-width

directed tree decomposition

bags guards

elementary wall

cylindrical wall

to the branch vertices of a cylindrical wall of order k based on its canonical embedding as follows: For every $1 \le i \le k$, and $1 \le j \le 2k$, the *j*-th branch vertex on the directed path P_i^1 receives coordinates (j, 2i - 1), while the *j*-th branch vertex on P_i^2 receives coordinates (2k + 1 - j, 2i).



Figure 1.2: Left: The cylindrical grid G_4 of order 4. Right: The elementary cylindrical wall W_4 in its canonical depiction with the cycles Q_1, \ldots, Q_4 (marked red) and the horizontal paths $P_1^1, P_1^2, \ldots, P_4^1, P_4^2$. The red half-edges at the top and the bottom of the wall indicate arcs connecting the paths P_1^1 and P_4^2 .

We are now ready to state the Directed Grid Theorem.

Theorem 1.3 (cf. [KK15]). For every $k \in \mathbb{N}$ there exists an integer d(k) such that every Diministry digraph D with $dtw(D) \ge d(k)$ contains the cylindrical grid of order k as a butterfly-minor.

Since $W_k \subseteq G_{2k}$, every digraph containing a G_{2k} -butterfly-minor also contains a W_k minor and hence a subdivision of W_k , since W_k is subcubic. We conclude that every digraph with sufficiently high directed tree-width (at least d(2k)) contains a wall of order k as a subdigraph.

Matroids. Matroids can be used to represent several algebraic and combinatorial structures of dependencies. For missing terminology and basic facts from matroid theory not mentioned or mentioned without proof in the following, please consult the standard reading [Ox111, Wel76].

In this thesis we will mostly use the representation of a matroid as a tuple M = (E, C)consisting of a finite ground set E(M) := E containing the elements of M and the family $C \subseteq 2^E$ of circuits of M, which satisfy certain so-called circuit axioms [Oxl11, Wel76]. A subset of E(M) is called independent in M if it does not contain a circuit as a subset, and a basis if it is inclusion-wise maximally independent. Bases of a matroid satisfy the following important so-called exchange property: If B_1, B_2 are bases of M, then for any $e \in B_1 \setminus B_2$ there exists $f \in B_2 \setminus B_1$ such that $(B_1 \setminus \{e\}) \cup \{f\}$ is a basis of M such that $X \cap B = \emptyset$. Instead of by their circuits, matroids may as well be represented uniquely by their collections of independent sets/bases/coindependent sets.

Two matroids $M_1 = (E_1, \mathcal{C}_1), M_2 = (E_2, \mathcal{C}_2)$ are called *isomorphic*, in symbols $M_1 \simeq M_2$, if there exists a bijection $f: E_1 \to E_2$ such that $C \in \mathcal{C}_1 \Leftrightarrow f(C) \in \mathcal{C}_2$ for all $C \subseteq E_1$.

Important examples of matroids are the so-called *linear or representable* matroids

ground set

circuits

isomorphic

induced by vector configurations in linear spaces. Let $V = \mathbb{F}^n$ be a vector-space over a field \mathbb{F} and let $X = \{x_1, \ldots, x_k\} \subseteq V$ for some $k \in \mathbb{N}$. Let A be the $n \times k$ -matrix over \mathbb{F} whose columns are x_1, \ldots, x_k . Then we define the *column matroid induced by* Aas $M[A] := (\{x_1, \ldots, x_k\}, \mathcal{C}_A)$, where its set of circuits \mathcal{C}_A consists of the inclusion-wise minimal collections of linearly dependent vectors from $\{x_1, \ldots, x_k\}$.

It is a well-known fact that M[A] is indeed a matroid for any choice of a matrix A. A matroid M is called \mathbb{F} -linear or representable over the field \mathbb{F} if there is a matrix A with entries in \mathbb{F} such that $M \simeq M[A]$.

Classical examples of matroids can further be derived from undirected graphs. Let G = (V, E) be a graph. The graphical matroid of G, denoted by M(G), is the matroid (E, \mathcal{C}) where the set \mathcal{C} of circuits consists of all edge-sets of the cycles of G. Analogously, the bond matroid of G is $M^*(G) = (E, \mathcal{S})$ where \mathcal{S} is the set of bonds (or minimal nonempty edge cuts) of G. Note that M(G) and $M^*(G)$ are the dual matroids of one another, that is, the bases of M(G) are complementary to the bases of $M^*(G)$.

A matroid is called a *graphic matroid* (respectively a *bond matroid* or *cographic*) if it is, respectively, isomorphic to the graphical or the bond matroid of some graph.

Graphic matroids and bond matroids form part of a larger matroid class, the so-called *regular matroids*. A matroid M is called regular if it is \mathbb{F} -linear for every field \mathbb{F} . The following equivalent characterization of regular matroids shows that they possess a very special structure which is not shared by general linear matroids. A matrix with entries in \mathbb{R} is called *totally unimodular* if every square submatrix has determinant -1, 0 or 1.

Theorem 1.4 ([Tut58]). Let M be a matroid. Then M is regular if and only if $M \simeq M[A]$ for a totally unimodular real-valued matrix A. Furthermore, for any field \mathbb{F} , reinterpreting the $\{-1, 0, 1\}$ -entries of A as elements of \mathbb{F} , we obtain an \mathbb{F} -linear representation of M.

Every graphic and every bond matroid is regular, but not vice-versa. Regular matroids are in turn generalized by the *binary matroids*, which are the \mathbb{F}_2 -linear matroids.

We conclude this paragraph with the important notion of matroid minors, which generalizes the concept of minors in graph theory. Given a matroid M and an element $e \in E(M)$, we denote by M - e and M/e the matroids obtained from M by deleting and contracting e. The circuits of M - e are defined as the circuits of M not containing e, while the circuits of M/e are the inclusion-wise minimal sets of the form $C \setminus \{e\} \neq \emptyset$ where C is a circuit of M. These operations are consistent with deletions and contractions in graph theory in the following sense: If G is a (multi-)graph and $e \in E(G)$, then it holds that $M(G/e) \simeq M(G)/e$, $M(G-e) \simeq M(G) - e$, $M^*(G-e) = M^*(G)/e$, and finally $M^*(G/e) \simeq M^*(G) - e$. A minor of M is any matroid isomorphic to a matroid obtained from M by a sequence of element deletions and contractions.

Oriented Matroids. For missing terminology and basic facts from the theory of oriented matroids not mentioned or mentioned without proof in the following, please consult the standard reading [BLVS⁺99].

An oriented matroid \overline{M} is a tuple (E, \mathcal{C}) consisting of a ground set E of elements and a collection \mathcal{C} of signed subsets of E, i.e., ordered partitions (C^+, C^-) of subsets C of Einto positive and negative parts such that the following four axioms are satisfied:

- (i) $(\emptyset, \emptyset) \notin C$
- (ii) If $(C^+, C^-) \in \mathcal{C}$, then $(C^-, C^+) \in \mathcal{C}$.
- (iii) If $(C_1^+, C_1^-), (C_2^+, C_2^-) \in \mathcal{C}$ such that $C_1^+ \cup C_1^- \subseteq C_2^+ \cup C_2^-$, then one of the equations $(C_1^+, C_1^-) = (C_2^+, C_2^-)$ or $(C_1^+, C_1^-) = (C_2^-, C_2^+)$ holds.

column matroid

representable

graphic matroid bond matroid

regular matroid

totally unimodular

binary matroid matroid minors

deletion and contraction

matroid minor

oriented matroid

signed set

(iv) Let $(C_1^+, C_1^-), (C_2^+, C_2^-) \in \mathcal{C}$ such that $(C_1^+, C_1^-) \neq (C_2^-, C_2^+)$, and let $e \in C_1^+ \cap C_2^-$. Then there is a $(C^+, C^-) \in \mathcal{C}$ with $C^+ \subseteq (C_1^+ \cup C_2^+) \setminus \{e\}$ and $C^- \subseteq (C_1^- \cup C_2^-) \setminus \{e\}$.

In case these axioms are satisfied, the elements of C are called *signed circuits* of \vec{M} .

Two oriented matroids $\vec{M_1} = (E_1, C_1)$ and $\vec{M_2} = (E_2, C_2)$ are called *isomorphic* if there exists a bijection $\sigma : E_1 \to E_2$ such that $\{(\sigma(C^+), \sigma(C^-)) \mid (C^+, C^-) \in C_1\} = C_2$. For every oriented matroid $\vec{M} = (E, C)$ and a signed circuit $X = (C^+, C^-) \in C$, we denote by $\underline{X} := C^+ \cup C^-$ the so-called *support* of X. From the axioms for signed circuits it follows that the set family $\underline{C} := \{\underline{X} \mid X \in C\}$ over the ground set E defines a matroid $M = (E, \underline{C})$, which we refer to as the *underlying matroid* of \vec{M} , and vice versa, \vec{M} is called an *orientation* of M. A matroid is called *orientable* if it admits at least one orientation. A signed circuit (C^+, C^-) is called *directed* if either $C^+ = \emptyset$ or $C^- = \emptyset$. We use this definition also for the circuits of the underlying matroid M, i.e., a circuit of M is *directed* in \vec{M} if (C, \emptyset) (or equivalently (\emptyset, C)) is a directed signed circuit, and *acyclic* if there exists no directed circuit.

Classical examples of oriented matroids can be derived from vector configurations in real-valued vector spaces and, most importantly for the investigations in this thesis, from directed graphs.

Given a configuration $\{x_1, x_2, \ldots, x_k\} \in \mathbb{R}^n$ of vectors for some $k \in \mathbb{N}$, consider the matroid M[A] with $A = (x_1, x_2, \ldots, x_k) \in \mathbb{R}^{n \times k}$. For a circuit $C = \{x_{i_1}, x_{i_2}, \ldots, x_{i_\ell}\} \in \mathcal{C}_A$, there are scalars $\alpha_1, \ldots, \alpha_\ell \in \mathbb{R} \setminus \{0\}$ such that $\sum_{j=1}^{\ell} \alpha_j x_{i_j} = 0$, and the coefficients α_j are determined up to multiplication with a common scalar. It is therefore natural to assign two signed sets to the circuit as follows: $X(C) := (C^+, C^-)$ and $-X(C) := (C^-, C^+)$, where $C^+ := \{x_{i_j} \mid \alpha_j > 0\}$ and $C^- := \{x_{i_j} \mid \alpha_j < 0\}$. The oriented column matroid induced by A is then defined as $\vec{M}[A] = (\{x_1, \ldots, x_k\}, \{X(C), -X(C) \mid C \in \mathcal{C}_A\})$.

Given a digraph D we can, as in the undirected case, associate with it two different kinds of oriented matroids with ground set A(D). Unsurprisingly, their underlying matroids are exactly the graphical respectively the bond matroid of U(D).

Definition 1.7. Let D be a digraph.

• For every (oriented) cycle C in D, let $(C^+, C^-), (C^-, C^+)$ be the two tuples describing a partition of A(C) into sets of forward and backward edges, according to some choice of cyclical traversal of C. Then

 $\{(C^+, C^-), (C^-, C^+) \mid C \text{ cycle in } D\}$

forms the set of signed circuits of an orientation M(D) of M(U(D)), called the oriented graphic matroid induced by D.

For every bond S = ∂(X) in D, let S⁺ = ∂⁺(X) be the set of edges in S with tail in X and head in Y := V(D) \ X, and let S⁻ := ∂⁺(Y). Then

$$\{(S^+, S^-), (S^-, S^+) \mid S \text{ is a bond in } D\}$$

forms the set of signed circuits of an orientation $M^*(D)$ of $M^*(U(D))$, called the oriented bond matroid induced by D.

Note that the directed circuits of an oriented graphic matroid are exactly the arc-sets of the directed cycles of the corresponding digraph D. Similarly, the directed circuits in an oriented bond-matroid are the arc-sets of the directed bonds in the corresponding digraph. Given an oriented matroid $\vec{M} = (E, C)$ and an element $e \in E(\vec{M})$, we denote by oriented column matroid

oriented graphic

matroid

oriented bond matroid

signed circuit isomorphic support underlying orientation orientable directed circuit totally cyclic acyclic deletion and contraction

 $\vec{M} - e$ and \vec{M}/e the oriented matroids obtained from \vec{M} by *deleting* and *contracting* e, respectively. The signed circuits of these matroids are defined as follows:

$$\mathcal{C}(\vec{M} - e) := \{ (C^+, C^-) \in \mathcal{C} \mid e \notin C^+ \cup C^- \},\$$

and the signed circuits of \vec{M}/e are the (inclusion-wise) support-minimal members of

 $\{(C^+ \setminus \{e\}, C^- \setminus \{e\}) \mid (C^+, C^-) \in \mathcal{C}\} \setminus \{(\emptyset, \emptyset)\}.$

These definitions generalize to subsets $Z \subseteq E(\vec{M})$, here we denote by $\vec{M} - Z$ and \vec{M}/Z , respectively, the oriented matroids obtained from \vec{M} by successively deleting (resp. contracting) all elements of Z (in arbitrary order⁴). We also denote $\vec{M}[Z] = \vec{M} - (E(\vec{M} - Z))$.

Again, in the case of graphic and cographic oriented matroids, the deletion and contraction operations resemble the same operations in directed graphs:

If D is a digraph and $e \in A(D)$, then $M(D) - e \simeq M(D-e)$, $M(D)/e \simeq M(D/e)$ and $M^*(D) - e \simeq M^*(D/e)$, $M^*(D)/e \simeq M^*(D-e)$.

For an oriented matroid \vec{M} with a collection \mathcal{C} of signed circuits, let \mathcal{S} be defined as the set of signed vectors (S^+, S^-) satisfying the following *orthogonality property* for every signed circuit $C = (C^+, C^-) \in \mathcal{C}$:

$$(S^+ \cap C^+) \cup (S^- \cap C^-) \neq \emptyset \iff (S^+ \cap C^-) \cup (S^- \cap C^+) \neq \emptyset.$$
(*)

Then S is called the set of *signed cocircuits* of \overline{M} . The supports of the signed cocircuits form exactly the cocircuits of the underlying matroid M. A signed cocircuit (S^+, S^-) is called *directed* if $S^+ = \emptyset$ or $S^- = \emptyset$. If the underlying matroid M of \overline{M} is regular, then the following stronger orthogonality holds for every signed circuit $(C^+, C^-) \in \mathcal{C}$, and every signed cocircuit $(S^+, S^-) \in \mathcal{S}$:

orthogonality

Farkas' Lemma

cover relation

Hasse diagram

cover graph

poset

strong

orthogonality

$$|S^{+} \cap C^{+}| + |S^{-} \cap C^{-}| = |S^{+} \cap C^{-}| + |S^{-} \cap C^{+}|.$$
(**)

For any digraph D the signed cocircuits of M(D) are the same as the signed circuits of $M^*(D)$, while the signed cocircuits of $M^*(D)$ are exactly the signed circuits of M(D).

We conclude this first part of the preliminary section by stating a couple of important facts concerning orientations of (regular) matroids from the literature.

Theorem 1.5 ([BLVS⁺99]). Let \vec{M} be an orientation of a regular matroid M. Then there exists a totally unimodular matrix A such that $\vec{M} \simeq \vec{M}[A]$ and $M \simeq M[A]$.

We will also need the following matroidal version of the famous Farkas' Lemma:

Theorem 1.6 ([BLVS⁺99]). Let \vec{M} be an oriented matroid and $e \in E(M)$. Then e is contained in a directed circuit of \vec{M} if and only if it is not contained in a directed cocircuit.

Posets and lattices. A partially ordered set, or poset for short, is a pair (P, \preceq) , where P is a set and \preceq is a reflexive, antisymmetric, and transitive binary relation on P. We write $x \prec y$ for two elements of P if $x \preceq y$ and $x \neq y$.

Posets can be represented more compactly by their minimal comparabilities: We say that $x \prec y$ is a cover relation, or y covers x, if there is no z in the poset with $x \prec z \prec y$. This defines the cover graph of P, which has the elements of P as vertices, and an edge for every cover relation. A Hasse diagram of P is a drawing of the cover graph in the plane, where vertices are represented by distinct points and for every cover relation $x \prec y$, the edge between x and y is drawn as a straight line going upwards from x to y. A poset P

⁴It is well-known that the order in which elements are deleted and contracted, respectively, does not affect the outcome of the process, up to isomorphism.

is called a *lattice*, if for any two elements x and y in P there is a unique smallest element z such that $x \leq z$ and $y \leq z$, and a unique largest element z such that $z \leq x$ and $z \leq y$. These elements are called the *join* and the *meet* of x and y, respectively. A lattice is called

distributive, if the join and meet operations distribute over each other. The downset (resp. *upset*) of an element y in P is the set of all x with $x \leq y$ (resp. $y \leq x$).

lattice

join and meet distributive down/up-set

Part I

Existence and Structure of Oriented Cycles

Chapter 2

Oriented Cycles in Digraphs of Large Out-Degree

2.1 Introduction

The graphs and digraphs considered in this chapter are always simple.

Recall that a subdivision of a graph F is a graph obtained from F by replacing its edges with internally vertex-disjoint paths. This notion appears in some of the most fundamental results of graph theory, such as Kuratowski's characterization of planar graphs, as well as many classical results in the structure theory of sparse graphs. Because of these applications, it is desirable to understand by which means a given graph G can be forced to contain a subdivision of a fixed graph F. One such direction of study that has received a great amount of attention in the literature is the question of how "dense" G should be to guarantee a subdivided F. For undirected graphs, this problem has been solved with great precision. Mader [Mad67] was the first to prove that for every fixed $k \in \mathbb{N}$, every graph of sufficiently large average degree contains a subdivision of K_k , and hence also of any other graph on at most k vertices. The precise asymptotic dependence of the average degree on k, that is required to force K_k as a subdivision, was independently determined by Bollobás and Thomason [BT98] and by Komlós and Szemerédi [KS96].

Theorem 2.1 ([BT98, KS96]). There is an absolute constant C > 0 such that every graph with average degree at least Ck^2 contains a subdivision of K_k . This bound is best-possible up to the value of C.

In view of the analogous definition of subdivisions in the directed setting (cf. Definition 1.1, Chapter 1), it is natural to ask to what extent the above phenomenon, that every "sufficiently dense" graph contains a subdivision of a fixed graph F, extends to digraphs.

Aboulker et al. [ACH⁺19] introduced the following handy terminology for the study of forcing subdivisions of digraphs through various digraph parameters. Given a digraph parameter γ ranging in \mathbb{N} , a digraph F is called γ -maderian if there exists a (smallest) number mader_{γ}(F) $\in \mathbb{N}$ such that every digraph D with $\gamma(D) \geq \text{mader}_{\gamma}(F)$ contains a subdivision of F as a subdigraph. We call mader_{γ}(F) the Mader number of F (with respect to γ). For example, using the natural analogue of these notions for undirected graphs, Theorem 2.1 states that the Mader number of K_k with respect to the graph parameter \overline{d} , namely the average degree, is quadratic in k, and in particular every graph F is \overline{d} -maderian.

The average out-degree (or, equivalently, average in-degree) of a digraph D is defined by $\overline{d}(D) := \frac{a(D)}{v(D)}$. As the transitive tournament is a digraph of very high average out-degree which does not even contain a subdivision of any directed cycle, it should be clear that an analogue of Theorem 2.1 for digraphs cannot hold in its full generality. It turns out that the family of \overline{d} -maderian digraphs is limited to the so-called *anti-directed forests*: forests in which every vertex is a sink or a source. The positive direction of this result is the consequence of a theorem of Burr [Bur82], who proved that every digraph of sufficiently large average degree contains every anti-directed forest as a subgraph (and hence also as a subdivision). The negative direction, as pointed out by Aboulker et al. [ACH⁺19], follows by considering bipartite graphs of large average degree and girth and orienting all their edges from one side of the bipartition to the other.

The above constructions of dense digraphs without certain subdivisions all contain sinks (i.e. vertices of out-degree zero); this motivates the study of subdivisions in digraphs with large minimum out-degree.

Since $\delta^+ \leq \overline{d}$, every \overline{d} -maderian digraph is obviously also δ^+ -maderian. However, a characterization of δ^+ -maderian digraphs is still widely unknown. Thomassen [Tho85b], answering a question of Seymour in the negative, constructed digraphs of arbitrarily large minimum out-degree not containing directed cycles of even length. As a consequence, if a digraph F has the property that each of its subdivisions contains a directed cycle of even length, then F is not δ^+ -maderian. As mentioned in the introduction, digraphs with this property are known in the literature as *even digraphs*, and have been extensively studied due to their relation to the so-called *even cycle problem*. For further information on this topic we refer the reader to Chapters 5 and 8, and to [RST99, ST87, Tho85b, Tho86, Tho92] for a selection of relevant literature. As can easily be verified by hand, the smallest even digraph is the bioriented clique \vec{K}_3 of order 3. This is also the smallest non- δ^+ -maderian digraph; indeed, the following theorem, which constitutes the first new result of this chapter, states that $\vec{K}_3 - e$, the digraph obtained from \vec{K}_3 by removing a single arc, is δ^+ -maderian. The proof of this theorem appears in Section 2.5.

Theorem 2.2. Every digraph D with $\delta^+(D) \ge 2$ contains a subdivision of $\overrightarrow{K}_3 - e$.

Observe that for every digraph F it holds that $\operatorname{mader}_{\delta^+}(F) \ge v(F) - 1$, since the bioriented clique on v(F) - 1 vertices has minimum out-degree v(F) - 2 but no subdivision of F. Hence, the bound in Theorem 2.2 is optimal.

Theorem 2.2 strengthens an earlier result by Thomassen (cf. [Tho85a], Theorem 6.2), who proved that every digraph of minimum out-degree 2 contains two directed cycles sharing precisely one vertex (this configuration is present in every subdivision of $\vec{K}_3 - e$). On the negative side, another construction by Thomassen [Tho85b] shows that there are digraphs of arbitrarily high minimum out-degree having no three directed cycles which share exactly one common vertex (and are otherwise disjoint). In other words, the bioriented 3-star \vec{S}_3 is not δ^+ -maderian. This result is somewhat surprising when compared to another positive result of Thomassen [Tho83], which shows that for every $k \in \mathbb{N}$ the digraph $k\vec{K}_2$ (i.e., the disjoint union of k digons) is δ^+ -maderian. More concretely, Thomassen proved that for every $k \in \mathbb{N}$ we have mader $_{\delta^+}(k\vec{K}_2) \leq (k+1)!$. The first linear bound on mader $_{\delta^+}(k\vec{K}_2)$ was proven by Alon [Alo96], and then further improved by Bucić [Buc18]. The famous Bermond-Thomassen-Conjecture [BT81] states that in fact mader $_{\delta^+}(k\vec{K}_2) = 2k - 1$, but this remains widely open.

A further negative result was established by DeVos et al. [DMMS12]. Building on previous work of Mader [Mad85], they constructed digraphs of arbitrarily high minimum out-degree having no pair of vertices x, y with two arc-disjoint dipaths from x to y as well as two from y to x (see [DMMS12, Observation 8]). This result shows that every δ^+ -maderian digraph F has arc-connectivity $\kappa'(F) \leq 1$. Yet another restriction follows

2.1. INTRODUCTION

from a construction we will present in Remark 2.20 in Section 2.3: Every δ^+ -maderian digraph F had directed tree-width at most one.

On the positive side, Aboulker et al. [ACH⁺19] proved that if F is a digraph consisting of two vertices x and y and three internally vertex-disjoint dipaths between x and y – two from x to y and one from y to x – then F is δ^+ -maderian.

The negative results above indicate that digraphs F with a sufficiently rich directed cycle structure are not δ^+ -maderian. However, to this date, no *acyclic* digraph is known that is not δ^+ -maderian. This led Mader [Mad85] to the following intriguing conjecture.

Conjecture 2.1 (Mader, 1985). Every acyclic digraph is δ^+ -maderian.

Clearly, it would suffice to prove Mader's conjecture for the transitive tournaments \vec{K}_k . Mader [Mad96] proved that mader_{\delta^+}(\vec{K}_4) = 3, but the existence of mader_{\delta^+}(\vec{K}_k) remains unknown for any $k \geq 5$. In view of the apparent difficulty of Mader's question, it is natural to try and verify Mader's conjecture for subclasses of acyclic digraphs. Mader himself [Mad95] considered the digraph consisting of two vertices x and y and k dipaths of length two from x to y, and showed that it is δ^+ -maderian for all $k \in \mathbb{N}$. Aboulker et al. [ACH⁺19] proposed to study the following two special cases of Mader's conjecture:

Conjecture 2.2 ([ACH⁺19]). Every orientation of a forest is δ^+ -maderian.

Conjecture 2.3 ([ACH⁺19]). Every orientation of a cycle is δ^+ -maderian.

Aboulker et al. [ACH⁺19] proved two special cases of Conjecture 2.2, showing that every orientation of a path and every in-arborescence is δ^+ -maderian¹. They also proved Conjecture 2.3 for oriented cycles consisting of two *blocks*², i.e., oriented cycles having exactly one source and one sink.

Our main contribution in this chapter is to verify Conjecture 2.3 in its full generality. Moreover, we show that the Mader number mader_{δ^+} of an oriented cycle grows (only) polynomially with the cycle length. Let C_{ℓ} denote the undirected cycle of length ℓ .

Theorem 2.3. There exists a polynomial function $K : \mathbb{N} \to \mathbb{N}$ such that for every $\ell \geq 2$, every digraph D with $\delta^+(D) \geq K(\ell)$ contains a subdivision of every orientation of C_{ℓ} .

The proof of Theorem 2.3 is presented in Section 2.2.

Disjoint union is a basic graph operation under which one might naturally anticipate the δ^+ -maderian property to be preserved. Aboulker et al. formulated this as a conjecture.

Conjecture 2.4 (cf. [ACH⁺19], Conjecture 7). If F_1 and F_2 are δ^+ -maderian, then also the disjoint union $F_1 \cup F_2$ of F_1 and F_2 is δ^+ -maderian.

Yet, despite quite a bit of effort, Conjecture 2.4 is only known to hold in a few special cases. Thomassen's Theorem for example states that the disjoint union of k digons is δ^+ -maderian for all k. As noted by Aboulker et al., Conjecture 2.4 would follow from a positive resolution of the following well-known problem raised by Stiebitz and Alon:

Problem 2.1 (cf. [Alo96, Alo06]). Given an integer $k \ge 1$, does there exist an integer f(k) such that every digraph D with $\delta^+(D) \ge f(k)$ contains (non-empty) disjoint sets $A, B \subseteq V(D)$ satisfying $\delta^+(D[A]), \delta^+(D[B]) \ge k$?

¹It is an easy exercise to show that every out-arborescence is δ^+ -maderian as well.

²By a *block* in an oriented cycle we mean a maximal directed subpath.

The difficulty of this problem is remarkable. While f(1) = 3 follows from the results of Thomassen [Tho83] on disjoint dicycles, even the existence of f(2) remains unknown.

As the second main result of this chapter we prove Conjecture 2.4 for a large class of digraphs. We say that a digraph F is a *topological grid minor* if there exists an elementary wall (cf. Section 1) containing F as a topological minor. The proof circumvents Problem 2.1 by using directed tree-width as well as the Directed Grid Theorem 1.3.

Theorem 2.4. If F_1, F_2, \ldots, F_r are topological grid minors and each of them is δ^+ -maderian, then also their disjoint union $F_1 \cup F_2 \ldots \cup F_r$ is δ^+ -maderian.

It is an easy observation that in particular, every orientation of a cycle is a topological grid minor. Combining Theorem 2.3 and Theorem 2.4 we can therefore generalize Thomassen's result that disjoint unions of digons are δ^+ -maderian to arbitrary oriented cycles as follows.

Corollary 2.5. Any disjoint union of (arbitrarily) oriented cycles is δ^+ -maderian.

The proof of Theorem 2.4 is given in Section 2.3.

Let $k_1, k_2 \in \mathbb{N}$. Following the notation in [ACH⁺19], we denote by $C(k_1, k_2)$ the twoblock cycle consisting of two vertices x, y and two internally vertex-disjoint dipaths from x to y of length k_1 and k_2 , respectively. As mentioned above, we have the trivial lower bound mader_{δ^+}($C(k_1, k_2)$) $\geq k_1 + k_2 - 1$. Aboulker et al. (see [ACH⁺19, Theorem 24]) proved the upper bound mader_{δ^+}($C(k_1, k_2)$) $\leq 2(k_1 + k_2) - 1$. They also observed that the trivial lower bound gives the truth if $k_2 = 1$, showing that mader_{δ^+}($C(k_1, k_2)$) = k for every $k \geq 1$. They then asked whether or not their bound on mader_{δ^+}($C(k_1, k_2)$) is tight.

Problem 2.2 ([ACH⁺19], Problem 25). What is the value of $mader_{\delta^+}(C(k_1, k_2))$?

Our next result improves upon the bound given by Aboulker et al. [ACH⁺19].

Theorem 2.6. Let $k_1 \ge k_2 \ge 2$ be integers. Then $mader_{\delta^+}(C(k_1, k_2)) \le k_1 + 3k_2 - 5$.

Theorem 2.6 improves upon the result of $[ACH^+19]$ for all values of $k_1, k_2 \ge 2$, and is asymptotically better if $k_1 \gg k_2$. Furthermore, if $k_2 = 2$ then the bound in Theorem 2.6 is optimal, as it matches the aforementioned trivial lower bound, thus showing that $mader_{\delta^+}(C(k,2)) = k+1$ for every $k \ge 1$. The proof of Theorem 2.6 appears in Section 2.4.

To conclude, let us mention that in contrast to the aforementioned negative results for general directed graphs, if we restrict our attention to the class of tournaments, which have an inherent density property, then it can be proved that every digraph is forcible as a subdivision by means of large minimum out-degree. This is a recent result by Girão, Popielarz, and Snyder [GPS21], which in addition gives a best-possible asymptotic bound of Ck^2 on the minimum out-degree of a tournament required to guarantee the existence of a subdivision of the bioriented k-clique.

As the family of δ^+ -maderian digraphs is still somewhat limited, Aboulker et al. [ACH⁺19] initiated the study of the effect of even stronger density conditions, involving the *strong* vertex-connectivity κ , and the strong arc-connectivity κ' of digraphs. Since $\kappa \leq \kappa' \leq \delta^+$, every δ^+ -maderian digraph is obviously κ' - and κ -maderian. Not much is known however concerning how much richer the families of κ - and κ' -maderian digraphs are. The following interesting questions were posed in [ACH⁺19]:

Problem 2.3 ([ACH⁺19], Problem 16). Is every digraph κ -maderian? Is every digraph κ' -maderian?

While the first question remains open, we can resolve the second question in the negative by proving that neither the bioriented 4-clique \vec{K}_4 nor the bioriented 4-star \vec{S}_4 is κ' -maderian:

Proposition 2.7. For every $k \in \mathbb{N}$, there exists a digraph G_k with $\kappa'(G_k) \ge k$ such that G_k contains no subdivision of \vec{K}_4 .

Proposition 2.8. For every $k \in \mathbb{N}$, there exists a digraph H_k with $\kappa'(H_k) \ge k$ such that H_k contains no subdivision of \vec{S}_4 .

The proofs of Propositions 2.7 and 2.8 are presented in Section 2.6.

We note that a main difficulty arising when studying subdivisions in digraphs (as opposed to undirected graphs) is that digraphs of large (strong) vertex-connectivity may not be linked. A digraph is called *k*-linked if for every 2*k*-tuple of distinct vertices $x_1, \ldots, x_k, y_1, \ldots, y_k$, there are vertex-disjoint dipaths P_1, \ldots, P_k such that P_i goes from x_i to y_i . In undirected graphs, it is known that a graph with sufficiently large vertex-connectivity is *k*-linked (see [BT96]), and linkedness has proven very useful for embedding subdivisions. In stark contrast, a construction of Thomassen [Tho91] shows that for every $k \in \mathbb{N}$ there is a strongly *k*-vertex-connected digraph which is not 2-linked. This makes subdivision questions for digraphs significantly more challenging.

2.2 Subdivisions of Oriented Cycles

In this section, we prove Theorem 2.3, which we restate here for convenience.

Theorem 2.9. For every $\ell \geq 2$ there is a polynomially bounded $K = K(\ell)$ such that every digraph D with $\delta^+(D) \geq K$ contains a subdivision of every oriented cycle of length ℓ .

It is well-known and easy to show that every digraph with minimum out-degree k contains a directed cycle of length at least k + 1. Thus, in what follows we restrict our attention to acyclic oriented cycles. For integers $a, b \ge 1$, let $C_{a,b}$ be the oriented cycle consisting of 2a vertices $s_1, \ldots, s_a, t_1, \ldots, t_a$ and 2a internally-disjoint length-b dipaths: one from s_i to t_i and one from s_{i+1} to t_i for each $1 \le i \le a$ (with indices taken modulo a). See Figure 1 for an illustration of $C_{2,3}$. It is easy to see that for every acyclic oriented cycle C, there are $a, b \ge 1$ such that every subdivision of $C_{a,b}$ is also a subdivision of C (specifically, a is the number of sources (or, equivalently, sinks) in C, and b is the largest length of a dipath contained in C). Therefore, it is sufficient to show that digraphs with minimum out-degree at least k(a, b) contain a subdivision of $C_{a,b}$ (for some suitable choice of k(a, b) = poly(a, b)). For a = 1, this statement was proven in [ACH⁺19], and we also give a new proof in Section 2.4. Consequently, it is sufficient to consider the case $a \ge 2$ (and, in fact, the assumption $a \ge 2$ is required by our method).

Dellamonica, Koubek, Martin, and Rödl [DKMR11] proved that for every $k \ge 1$ and $g \ge 3$ there exists K = K(k,g) such that every digraph D with $\delta^+(D) \ge K$ contains a subdigraph D' with $\delta^+(D') \ge k$ and with directed girth $\vec{g}(D')$ at least g. Thus, in order to prove Theorem 2.9, it suffices to establish the following:

Theorem 2.10. There is a constant C > 0 such that for any pair of integers $a \ge 2$, $b \ge 1$, every digraph D with $\delta^+(D) \ge Cab^7$ and $\vec{g}(D) \ge 4b^2$ contains a subdivision of $C_{a,b}$.



Figure 2.1: The oriented cycle $C_{2,3}$

A quantitative version³ of the aforementioned result of [DKMR11] is that $K(k,g) \leq O(kg^2 \log g)$. It follows that minimum out-degree at least $a \cdot \text{poly}(b)$ is enough to force a subdivision of $C_{a,b}$, and that the conclusion of Theorem 2.9 holds with $K(\ell) = \text{poly}(\ell)$.

For the rest of this section we set $g := 4b^2$. Our proof of Theorem 2.10 will use a certain structure we call a *chain*, that will consist of some carefully chosen *gadgets*. This structure will enable us to embed subdivisions of $C_{a,b}$ in the given digraph. We start by presenting these key definitions.

2.2.1 The Gadgets

We will use three types of gadgets. Each of the gadgets will have a special pair of vertices p, q with an arc from p to q. The gadgets are defined as follows:

- (I) A gadget of type I is a directed cycle of length at least g through the arc (p,q).
- (II) A basic gadget of type II is a digraph consisting of vertices p, q, r and a dipath P_1 from r to p, such that P_1 has length at least $2b^2 + b 2$, $q \notin V(P)$, and every vertex of P_1 has an arc to q (so in particular, (p, q) is an arc). An extended gadget of type II consists of a basic gadget of type II, comprised of vertices p, q, r and a dipath P_1 as above, as well as an additional dipath P_2 of length at least b having the following properties:
 - (a) The last vertex of P_2 is $r, V(P_1) \cap V(P_2) = \{r\}$, and $q \notin V(P_2)$.
 - (b) Either there is an arc from the first vertex of P_2 to the second vertex of P_1 , or there is an arc from some vertex in $V(P_1) \setminus \{r\}$ to the first vertex of P_2 .

For an extended type-II gadget G, the *basic part* of G is the corresponding basic type-II gadget, namely the subgraph of G induced by $V(P_1) \cup \{q\}$.

(III) A gadget of type III is a digraph consisting of vertices p, q, r, the arc (p, q), and two internally disjoint dipaths P_1, P_2 from p and q, respectively, to r, such that P_1 and P_2 have length at least 2b - 1 each.

³In fact, the bound on K(k,g) appearing in [DKMR11] was slightly weaker — in that the logarithmic factor depended on k — but it is easy to see that by using the argument from [DKMR11] and replacing a union bound used there with a tighter concentration inequality (say, Chernoff's bound), one obtains the stronger estimate stated here.



Figure 2.2: A basic gadget of type II (left) and a gadget of type III (right)



Figure 2.3: The two options for an extended gadget of type II: either there is an arc from the first vertex of P_2 to the second vertex of P_1 (left), or there is an arc from some vertex in $V(P_1) \setminus \{r\}$ to the first vertex of P_2 (right).

The various types of gadgets are depicted in Figures 2-3. For convenience, we also introduce the notion of a *trivial gadget*: a trivial gadget simply consists of vertices p, q and the arc (p, q) (and no other vertices).

We now introduce another useful definition. For integers $a, b \ge 1$, an (a, b)-alternatingpath is an oriented path R consisting of vertices $s_1, \ldots, s_a, t_1, \ldots, t_a$ and pairwise internallydisjoint dipaths $Q_1, \ldots, Q_a, Q'_1, \ldots, Q'_{a-1}$, such that Q_i is a dipath from s_i to t_i (for each $1 \le i \le a$), Q'_i is a dipath from s_{i+1} to t_i (for each $1 \le i \le a - 1$), and Q_2, \ldots, Q_{a-1} , Q'_1, \ldots, Q'_{a-1} have length at least b each. We note that Q_1 or Q_a may have length zero (in which case $s_1 = t_1$ or $s_a = t_a$, respectively). In particular, for vertices u, v, any dipath from u to v is a (1, b)-alternating-path with $s_1 = u$ and $t_1 = v$ (for any value of b); and any dipath of length at least b from u to v is a (2, b)-alternating-path with $s_2 = t_2 = u$ and $s_1 = t_1 = v$. The path R is called strong if Q_1 and Q_a also have length at least b. When several paths are considered at the same time, we will write $s_i(R), t_i(R), Q_i(R), Q'_i(R)$ (instead of s_i, t_i, Q_i, Q'_i) so as to prevent confusion. The following observation follows immediately from the definitions of $C_{a,b}$ and alternating-paths. **Observation 2.11.** Let $a_1, a_2, b \ge 1$ be integers, and for each i = 1, 2, let R_i be a strong (a_i, b) -alternating path. Suppose that $s_1(R_1) = t_{a_2}(R_2)$, $s_1(R_2) = t_{a_1}(R_1)$ and that R_1 and R_2 do not share any other vertices. Then $R_1 \cup R_2$ spans a subdivision of $C_{a_1+a_2-2.b}$.

Let us now prove some simple facts about type I and type II gadgets.

Lemma 2.12. Let G be a gadget of type I or II (either basic or extended). Then:

- 1. G contains a (2,b)-alternating path R_0 with $s_1(R_0) = t_1(R_0) = p$ and $t_2(R_0) = q$.
- 2. $\{p,q\}$ is reachable from every vertex of G.

Proof. Item 2 follows immediately from the definitions of these gadgets. Let us prove Item 1. If G is of type I, i.e. a directed cycle of length at least g > b through (p,q), then define R_0 by letting $s_1(R_0) = t_1(R_0) = p$, $s_2(R_0) = t_2(R_0) = q$ and $Q'_1(R_0) = G[q, p]$ (i.e., $Q'_1(R_0)$ is simply the q-p-dipath obtained from the cycle by removing the arc (p,q)). If G is of type II then define R_0 by letting $s_1(R_0) = t_1(R_0) = p$, $s_2(R_0) = t_1(R_0) = p$, $s_2(R_0) = r$, $t_2(R_0) = q$, $Q'_1(R_0) = P_1$ and $Q_2(R_0) = (r,q)$.

Lemma 2.13. Let G be an extended gadget of type II, and let p, q, r and P_1, P_2 be as in the definition of such a gadget. Then:

- 1. For every $x \in V(G) \setminus \{p,q\}$, there exists $1 \le a \le 2$ and an (a,b)-alternating-path R with $t_a(R) = x$, $s_1(R) \in \{p,q\}$ and $|V(R) \cap \{p,q\}| = 1$.
- 2. For every set $\emptyset \neq X \subseteq V(P_1) \setminus \{p, r\}$, there exists $1 \leq a \leq 2$ and an (a, b)-alternatingpath R with $t_a(R) \in X$, $s_1(R) \in \{p, q\}$ and $|V(R) \cap \{p, q\}| = |V(R) \cap X| = 1$.

Proof. We start by proving Item 2, from which Item 1 will then follow. So let $\emptyset \neq X \subseteq V(P_1) \setminus \{p, r\}$. Denote by z the first vertex of P_2 , and by y the second vertex of P_1 . By the definition of an extended type II gadget, either $(z, y) \in A(G)$ or there is some $w \in V(P_1) \setminus \{r\}$ such that $(w, z) \in A(G)$. Suppose first that $(z, y) \in A(G)$. Traverse the dipath P_1 starting from y until the first vertex of X is reached, and denote this vertex by x. Evidently, we have $X \cap V(P_1[y, x]) = \{x\}$. Now define R by setting $s_1(R) = t_1(R) = q$, $s_2(R) = z, t_2(R) = x, Q'_1(R) = P_2 \circ (r, q)$ and $Q_2(R) = (z, y) \circ P_1[y, x]$. Then R is indeed a (2, b)-alternating-path (since $|P_2| \ge b$), and we have $V(R) \cap \{p, q\} = \{q\}$ (since $x \neq p$) and $V(R) \cap X = \{x\}$, as required.

Suppose now that there is $w \in V(P_1) \setminus \{r\}$ such that $(w, z) \in A(G)$. If w = p then, as before, we let x be the first vertex of X reached when traversing $P_1[y, p]$. Observe that $(w, z) \circ P_2 \circ P_1[r, x]$ is a dipath from p = w to x, and thus also a (1, b)-alternating-path R with $s_1(R) = p$ and $t_1(R) = x$. Moreover, our choice of x implies that $V(R) \cap X = \{x\}$, as required.

So from now on we assume that $w \neq p$. In this case, choose an element $x' \in X$, which is closest to w in the *undirected* path underlying P_1 . In other words, we choose $x' \in X$ such that the subpath of P_1 between w and x' contains no vertex of X other than x' itself. We consider two cases, based on the relative position of x' and w along P_1 . Assume first that when traversing the dipath P_1 (starting from r), w is reached before x' is (here we allow w = x'). In this case, define R by setting $s_1(R) = t_1(R) = q$, $s_2(R) = w$, $t_2(R) = x'$, $Q'_1(R) = (w, z) \circ P_2 \circ (r, q)$ and $Q_2(R) = P_1[w, x']$. Assume now that x' is reached before w when traversing P_1 . In this case, define a (2, b)-alternating-path R by setting $s_1(R) =$ $t_1(R) = q$, $s_2(R) = t_2(R) = x'$ and $Q'_1(R) = P_1[x', w] \circ (w, z) \circ P_2 \circ (r, q)$. Observe that in both cases, R is indeed a (2, b)-alternating-path (because $|P_2| \ge b$), $V(R) \cap \{p,q\} = \{q\}$ (because $w, x' \neq p$), and $V(R) \cap X = \{x'\}$ (by our choice of x'). This concludes the proof of Item 2.

It remains to prove Item 1. So let $x \in V(G) \setminus \{p,q\}$. If $x \in V(P_2)$, then we define a (2,b)-alternating-path R by setting $s_1(R) = t_1(R) = p$, $s_2(R) = t_2(R) = x$ as well as $Q'_1(R) = P_2[x,r] \circ P_1$. This path is indeed (2,b)-alternating since $|Q'_1(R)| \ge |P_1| \ge 2b^2 + b - 2 \ge b$. And if $x \in V(G) \setminus (V(P_2) \cup \{p,q\}) = V(P_1) \setminus \{p,r\}$, then we obtain the required alternating-path R by applying Item 2 with $X := \{x\}$. It is easy to see that in both cases, R satisfies the assertion of Item 1. This completes the proof of the lemma. \Box

2.2.2 Gadget Chains

We now define the notion of a *chain* of gadgets, a structure which will be instrumental to our proof of Theorem 2.10. In what follows, for a gadget G, we will denote by p(G) and q(G) the designated vertices p and q of G, respectively.

Definition 2.1. A chain C consists of a directed path $P = v_0, \ldots, v_m$, a partition $A_1 \cup A_2 = A(P) = \{(v_0, v_1), \ldots, (v_{m-1}, v_m)\}$ of the arc-set of P, and a collection of (non-trivial) gadgets ($G_e : e \in A_2$) having the following four properties:

- 1. For every $e \in A_2$, the gadget G_e is either of type I, basic type II, or type III.
- 2. For every $e = (v_i, v_{i+1}) \in A_2$, $p(G_e) = v_i$ and $q(G_e) = v_{i+1}$.
- 3. $V(G_{(v_i,v_{i+1})}) \cap \{v_0,\ldots,v_m\} = \{v_i,v_{i+1}\}$ for every $(v_i,v_{i+1}) \in A_2$.
- 4. $V(G_e) \cap V(G_f) \subseteq \{v_0, \ldots, v_m\}$ for every pair of distinct $e, f \in A_2$.

We will use the following terminology and notation in relation with gadget chains:

- With a slight abuse of notation, we identify the chain C and the digraph consisting of the union of P and the gadgets $G_e, e \in A_2$.
- For convenience, for $(v_i, v_{i+1}) \in A_1$ we denote by $G_{(v_i, v_{i+1})}$ the trivial gadget with vertices v_i, v_{i+1} and arc (v_i, v_{i+1}) .
- In cases where several chains are considered at the same time, we will write $A_1(\mathcal{C})$, $A_2(\mathcal{C})$ and $G_e(\mathcal{C})$ to indicate that we are considering the chain \mathcal{C} .
- The dipath P is called the *spine* of the chain, and |P| = m is the *length* of the chain.
- The vertex set of \mathcal{C} , denoted $V(\mathcal{C})$, is defined as $V(\mathcal{C}) = V(P) \cup \bigcup_{e \in A_2} V(G_e)$.
- For integers $0 \leq i < j \leq m$, we denote by $\mathcal{C}[v_i, v_j]$ the subchain of \mathcal{C} whose spine is $P[v_i, v_j] = v_i, v_{i+1}, \ldots, v_j$; so $A_{\ell}(\mathcal{C}[v_i, v_j]) = A_{\ell}(\mathcal{C}) \cap A(P[v_i, v_j])$ for $\ell = 1, 2$, and $\mathcal{C}[v_i, v_j]$ inherits the gadgets of \mathcal{C} .

The next sequence of lemmas is concerned with embedding subdivisions of $C_{a,b}$ using gadget chains. The following lemma asserts that such chains can be used to find (a, b)-alternating-paths.

Lemma 2.14. Let $a, b \ge 1$ be integers. Let C be a chain, let $P = v_0, \ldots, v_m$ and A_1, A_2 be as in Definition 2.1, and suppose that $|A_2| \ge a(b+1) - 1$. Then C contains a strong (a, b)-alternating-path R with $s_1(R) = v_0$ and $t_a(R) = v_m$.

Proof. The proof is by induction on a. In the base case a = 1, the condition in the lemma states that $|A_2| \ge b$. This implies that $m = |P| \ge b$, meaning that P is a dipath of length at least b from v_0 to v_m , and hence also a strong (1, b)-alternating-path with $s_1(P) = v_0$ and $t_1(P) = v_m$.

We now move on to the induction step. So let $a \ge 2$. Let j be the largest integer in the set $\{0, \ldots, m-b-1\}$ satisfying $(v_j, v_{j+1}) \in A_2$. Set $\mathcal{C}' := \mathcal{C}[v_0, v_j]$. Then $|A_2(\mathcal{C}')| \ge$ $|A_2(\mathcal{C})| - (b+1) \ge (a-1)(b+1) - 1$. By the induction hypothesis, \mathcal{C}' contains a strong (a-1,b)-alternating-path R' with $s_1(R') = v_0$ and $t_{a-1}(R') = v_j$.

Setting $e := (v_j, v_{j+1})$, suppose first that G_e is either of type I or a basic gadget of type II. By Item 1 of Lemma 2.12, G_e contains a (2, b)-alternating path R_0 with $s_1(R_0) = t_1(R_0) = v_j$ and $t_2(R_0) = v_{j+1}$. Now let R be the (a, b)-alternating-path obtained by attaching to R' the dipaths $Q'_1(R_0)$ and $Q_2(R_0) \circ P[v_{j+1}, v_m]$. Formally, R is defined by setting $s_i(R) = s_i(R')$ and $t_i(R) = t_i(R')$ for every $1 \le i \le a-1$ (so in particular, $s_1(R) = v_0$), $s_a(R) = s_2(R_0)$, $t_a(R) = v_m$, $Q_i(R) = Q_i(R')$ for every $1 \le i \le a-1$, $Q'_i(R) = Q'_i(R')$ for every $1 \le i \le a-2$, $Q'_{a-1}(R) = Q'_1(R_0)$ and $Q_a(R) = Q_2(R_0) \circ P[v_{j+1}, v_m]$. Note that $|Q_a(R)| \ge b$ because $j \le m-b-1$. It follows that R is indeed a strong (a, b)-alternating-path, as required.

Suppose now that G_e is of type III. Then G_e consists of the vertices v_j, v_{j+1} , a vertex r, and two internally vertex-disjoint dipaths P_1, P_2 from v_j and v_{j+1} , respectively, to r, such that P_1 and P_2 have length at least $2b-1 \ge b$ each. Now let R be the (a, b)-alternating-path obtained by attaching to R' the dipaths P_1, P_2 and $P[v_{j+1}, v_m]$. Formally, R is defined by setting $s_i(R) = s_i(R')$ for every $1 \le i \le a-1$ (so in particular, $s_1(R) = v_0$), $t_i(R) = t_i(R')$ for every $1 \le i \le a-2$, $t_{a-1}(R) = r$, $s_a(R) = v_{j+1}$, $t_a(R) = v_m$, $Q_i(R) = Q_i(R')$ and $Q'_i(R) = Q'_i(R')$ for every $1 \le i \le a-2$, $Q_{a-1}(R) = Q_{a-1}(R') \circ P_1$, $Q'_{a-1}(R) = P_2$ and $Q_a(R) = P[v_{j+1}, v_m]$. Again, it is easy to check that R is a strong (a, b)-alternating-path, as required.

Lemma 2.15. Let G, G^* be gadgets such that $V(G) \cap V(G^*) \neq \emptyset$, $p(G^*), q(G^*) \notin V(G)$, and G^* is an extended gadget of type II. Then there exists $1 \le a \le 3$ such that $G \cup G^*$ contains an (a, b)-alternating-path R with $t_a(R) \in \{p(G), q(G)\}, s_1(R) \in \{p(G^*), q(G^*)\}$ and $|V(R) \cap \{p(G), q(G)\}| = |V(R) \cap \{p(G^*), q(G^*)\}| = 1$.

Proof. For convenience, let us put p := p(G), q := q(G), $p^* := p(G^*)$ and $q^* := q(G^*)$. The assumption $p^*, q^* \notin V(G)$ will be used implicitly throughout the proof. We proceed by a case analysis over the type of G.

Case 1. *G* is trivial, a gadget of type I, or a gadget of type II. Recall that by assumption, $V(G) \cap V(G^*) \neq \emptyset$. By Item 2 of Lemma 2.12, $\{p,q\}$ is reachable from every vertex of V(G) via a dipath inside *G* (this is evident if *G* is trivial). In particular, *G* contains a dipath from $V(G) \cap V(G^*)$ to $\{p,q\}$. Fix a shortest such dipath $P \subseteq G$, and let $x \in V(G) \cap V(G^*)$ be the first vertex of *P*. The minimality of *P* implies that $V(P) \cap V(G^*) = \{x\}$ and $|V(P) \cap \{p,q\}| = 1$. By Item 1 of Lemma 2.13, there is $1 \le a \le 2$ such that G^* contains an (a, b)-alternating-path R^* with $t_a(R^*) = x$, $s_1(R^*) \in \{p^*,q^*\}$ and $|V(R^*) \cap \{p^*,q^*\}| = 1$. (The condition $x \notin \{p^*,q^*\}$ appearing in Item 1 of Lemma 2.13 is satisfied here because $x \in V(G) \cap V(G^*) = \{x\}$ and $V(R^*) \subseteq V(G^*)$. Now it is easy to see that by combining *P* and R^* we obtain an (a, b)-alternating-path *R* with $t_a(R) \in \{p,q\}$, $s_1(R) \in \{p^*,q^*\}$ and $|V(R) \cap \{p,q\}| = |V(R) \cap \{p^*,q^*\}| = 1$. Formally, *R* is defined by setting $Q_a(R) := Q_a(R^*) \circ P$ (so $t_a(R) \in \{p,q\}$ is the last vertex of *P*); $s_i(R) = s_i(R^*)$ for every $1 \le i \le a$;

and $t_i(R) = t_i(R^*)$, $Q_i(R) = Q_i(R^*)$ and $Q'_i(R) = Q'_i(R^*)$ for every $1 \le i \le a - 1$. This completes the proof in Case 1.

Case 2. G is a gadget of type III. In this case G consists of the arc (p,q), a vertex r, and two internally vertex-disjoint dipaths P_1, P_2 from p and q, respectively, to r, such that P_1 and P_2 have length at least 2b - 1 each.

As G^* is an extended gadget of type II, it consists of the vertices p^*, q^* , a vertex r^* and dipaths P_1^*, P_2^* , all satisfying the properties stated in the definition of a type II gadget. We start by handling the case that there is some $x \in V(P_1^*) \cap V(G)$ such that the distance from $\{p,q\}$ to x in G is at least b-1. Since $V(G) = V(P_1) \cup V(P_2)$, we have either $x \in V(P_1)$ or $x \in V(P_2)$. Suppose without loss of generality that $x \in V(P_1)$ (the case $x \in V(P_2)$ is symmetric). Our assumption on x then means that $|P_1[p,x]| \geq b-1$. Now, $R := P_1[p,x] \circ (x,q^*)$ is a dipath of length at least b from $\{p,q\}$ to q^* (note that we have $(x,q^*) \in A(G^*)$ by the definition of a type II gadget). Hence, R constitutes a (2,b)-alternating-path with $s_1(R) = t_1(R) = q^*$ and $s_2(R) = t_2(R) \in \{p,q\}$. Moreover, $|V(R) \cap \{p,q\}| = 1$ and $V(R) \cap \{p^*,q^*\} = \{q^*\}$ (since $p^* \notin V(G)$), as required.

So from now on we assume that every $x \in V(P_1^*) \cap V(G)$ is at distance at most b-2 from $\{p,q\}$ in G (in particular, if b=1 then $V(P_1^*) \cap V(G) = \emptyset$). It follows that $|V(G) \cap V(P_1^*)| \leq 2(b-1)$. Moving forward, we will consider two cases, based on the intersection of V(G) with $V(P_2^*)$.

Case 2.1. $V(G) \cap V(P_2^*) = \emptyset$. Set $X := V(G) \cap V(G^*)$, noting that $X \neq \emptyset$ by assumption. From $V(G) \cap V(P_2^*) = \emptyset$ and $p^*, q^* \notin V(G)$, we conclude $X \subseteq V(G^*) \setminus (V(P_2^*) \cup \{p^*, q^*\}) = \emptyset$ $V(P_1^*) \setminus \{p^*, r^*\}$. By Item 2 of Lemma 2.13, there exists $1 \le a \le 2$ and an (a, b)-alternatingpath R^* contained in G^* , such that $t_a(R^*) \in X$, $s_1(R^*) \in \{p^*, q^*\}$ and $|V(R^*) \cap \{p^*, q^*\}| =$ $|V(R^*) \cap X| = 1$. For convenience, put $x := t_a(R^*)$. Note that $V(R^*) \cap V(G) = \{x\}$ by our choice of X and R^* . We now see that if $x \in \{p, q\}$, then $R := R^*$ satisfies the requirements of the lemma. Suppose then that $x \notin \{p,q\}$. Since $x \in V(G)$, we have either $x \in V(P_1)$ or $x \in V(P_2)$. Without loss of generality, we assume that $x \in V(P_1)$ (the case that $x \in V(P_2)$ is symmetric). Recall that by our assumption, x is at distance at most b-2from $\{p,q\}$ in G; in other words, the length of the dipath $P_1[p,x]$ is at most b-2. As $|P_1| \ge 2b-1 \ge 2b-2$, we get that $|P_1[x, r]| = |P_1| - |P_1[p, x]| \ge b$. Now let R be the (a+1, b)alternating-path obtained by combining R^* with the dipaths $P_1[x, r]$ and P_2 . Formally, R is defined by setting $s_i(R) = s_i(R^*)$ for every $1 \le i \le a$; $t_i(R) = t_i(R^*)$, $Q_i(R) = Q_i(R^*)$ and $Q'_i(R) = Q'_i(R^*)$ for every $1 \le i \le a-1$; $t_a(R) = r$; $s_{a+1}(R) = t_{a+1}(R) = q$; $Q_a(R) = Q_a(R^*) \circ P_1[x, r]$; and $Q'_a(R) = P_2$. Note that R is indeed an (a+1, b)-alternatingpath; this follows from our choice of R^* , the fact that $V(R^*) \cap V(G) = \{x\}$, and the bounds $|Q'_a(R)| = |P_2| \ge 2b - 1 \ge b$ and $|Q_a(R)| \ge |P_1[x, r]| \ge b$. We also have $|V(R) \cap \{p^*, q^*\}| = 1$ $(as V(R) \cap \{p^*, q^*\} = V(R^*) \cap \{p^*, q^*\})$ and $V(R) \cap \{p, q\} = \{q\}$ (by our definition of R and as $x \notin \{p,q\}$). Since $a + 1 \leq 3$, we see that the assertion of the lemma holds in Case 2.1.

Case 2.2. $V(G) \cap V(P_2^*) \neq \emptyset$. In this case, we traverse the dipath P_2^* backwards (i.e., starting from its last vertex, r^*), until the first time a vertex of V(G) is reached, and denote this vertex by w. Evidently, $V(G) \cap V(P_2^*[w, r^*]) = \{w\}$. For convenience, let us set $P^* := P_2^*[w, r^*] \circ P_1^*$, noting that P^* starts at w, ends at p^* , and has length at least $|P_1^*| \ge 2b^2 + b - 2$ (by the definition of a type II gadget). For every $u \in V(G) \cap (V(P^*) \setminus \{w\})$, denote by e_u the (unique) arc of P^* whose head is u. Let R_1, \ldots, R_m be the connected components of the digraph obtained from the path P^* by deleting the arcs e_u for every

 $u \in V(G) \cap (V(P^*) \setminus \{w\})$ (this digraph is a dipath forest). Then for each $1 \leq i \leq m, R_i$ is a dipath whose first vertex is in V(G) and all of whose other vertices are not in V(G). Recall that by our assumption, $|V(G) \cap V(P_1^*)| \leq 2(b-1)$. Now, our choice of w implies that $V(G) \cap V(P^*) = (V(G) \cap V(P_1^*)) \cup \{w\}$. The number of edges we deleted from P^* to obtain R_1, \ldots, R_m is, one the one hand, equal to m-1, and on the other hand equal to $|V(G) \cap (V(P^*) \setminus \{w\})| \leq |V(G) \cap V(P_1^*)| \leq 2(b-1)$. It follows that $m \leq 2(b-1)+1 = 2b-1$ and $|R_1| + \cdots + |R_m| \geq |P^*| - 2(b-1) \geq 2b^2 + b - 2 - 2(b-1) = 2b^2 - b$. By averaging, there is some $1 \leq i \leq m$ such that $|R_i| \geq \frac{2b^2-b}{m} \geq \frac{2b^2-b}{2b-1} \geq b$. Let u (resp. v) be the first (resp. last) vertex of R_i . Note that $v \in V(P_1^*)$ due to our

Let u (resp. v) be the first (resp. last) vertex of R_i . Note that $v \in V(P_1^*)$ due to our choice of w. Define a dipath R^* as follows: if $v = p^*$ then set $R^* := R_i$, and otherwise set $R^* := R_i \circ (v, q^*)$. (That $(v, q^*) \in A(G^*)$ follows from the definition of a type II gadget and the fact that $v \in V(P_1^*)$.) Then $|V(R^*) \cap \{p^*, q^*\}| = 1$ and $V(G) \cap V(R^*) = \{u\}$ (because $V(G) \cap V(R_i) = \{u\}$ and $q^* \notin V(G)$). In particular, $u \in V(G) = V(P_1) \cup V(P_2)$. Suppose, without loss of generality, that $u \in V(P_1)$ (the case $u \in V(P_2)$ is symmetric). Now define a (2, b)-alternating-path R as follows: $s_1(R) = t_1(R) = q^*$, $s_2(R) = t_2(R) = p$, $Q'_1(R) = P_1[p, u] \circ R^*$. Note that $Q'_1(R)$ is indeed a dipath (because $V(P_1) \cap V(R^*) \subseteq V(G) \cap V(R^*) = \{u\}$), and that $|Q'_1(R)| \ge |V(R^*)| \ge |V(R_i)| \ge b$. Furthermore, $q \notin V(R)$ and $V(R) \cap \{p^*, q^*\} = V(R^*) \cap \{p^*, q^*\}$. Thus, R satisfies the requirements of the lemma. This completes the proof.

Note that the gadget G in Lemma 2.15 is allowed to be trivial.

In the following lemma we show that if a sufficiently "rich" chain (i.e., a chain C with $|A_2(C)|$ large enough) "self-intersects" in some well-defined way, then it contains a subdivision of $C_{a,b}$. Roughly speaking, this can be thought of as closing the alternating-path obtained from Lemma 2.14 to form a cycle (i.e., a $C_{a,b}$ -subdivision). The purpose of Lemma 2.15 is to achieve this closure.

Lemma 2.16. Let $a \ge 2$ and $b \ge 1$ be integers, let C be a chain contained in a digraph D and let $P = z_0, \ldots, z_\ell$ and A_1, A_2 be as in Definition 2.1. Suppose $|A_2| \ge (a+3)(b+1)-2$ and that at least one of the following two conditions is satisfied:

- 1. There exists $x \in V(G_{(z_0,z_1)})$ such that $(z_{\ell}, x) \in A(D)$.
- 2. There exists a vertex $z^* \in V(D) \setminus V(\mathcal{C})$ such that $(z_{\ell}, z^*) \in A(D)$, and there exists an extended type II gadget G^* such that $p(G^*) = z_{\ell}$, $q(G^*) = z^*$, $V(G_{(z_0, z_1)}) \cap V(G^*) \neq \emptyset$ and $V(\mathcal{C}) \cap V(G^*) \subseteq V(G_{(z_0, z_1)}) \cup \{z_{\ell}\}.$

Then D contains a subdivision of $C_{a,b}$.

Proof. For convenience, put $G := G_{(z_0,z_1)}$. We start by showing that for some $1 \le a_1 \le 3$, $G \cup G^*$ contains an (a_1, b) -alternating-path R^* with the properties $t_{a_1}(R^*) \in \{z_0, z_1\}$, $s_1(R^*) \in \{z_\ell, z^*\}$ and $|V(R^*) \cap \{z_0, z_1\}| = |V(R^*) \cap \{z_\ell, z^*\}| = 1$. If Condition 2 in the lemma holds, then this assertion follows immediately from Lemma 2.15. Note that the conditions of Lemma 2.15 are indeed satisfied in our setting: we have $V(G) \cap V(G^*) \neq \emptyset$ and $z^* \notin V(G)$ by assumption, and $z_\ell \notin V(G)$ by the definition of a chain and as $m \ge |A_2| \ge 2$.

Suppose now that Condition 1 in the lemma holds. Let $x \in V(G)$ be such that $(z_{\ell}, x) \in A(D)$. If $x \in \{z_0, z_1\}$ then the arc (z_{ℓ}, x) itself constitutes a (1, b)-alternating-path R^* with the required properties. Suppose from now on that $x \notin \{z_0, z_1\}$. So in particular, G is not a trivial gadget. Assume first that G is of type I or II. By Item 2 of Lemma 2.12, $\{z_0, z_1\}$ is reachable from x inside G. Fix a shortest path P_0 from x to $\{z_0, z_1\}$ contained in G. Then $|V(P_0) \cap \{z_0, z_1\}| = 1$. Now $R^* := (z_{\ell}, x) \circ P_0$ is a dipath from z_{ℓ} to $\{z_0, z_1\}$, and hence also a (1, b)-alternating-path with $t_1(R^*) \in \{z_0, z_1\}$ and $s_1(R^*) = z_{\ell}$,

as required. Assume now that G is of type III. Let $r \in V(G)$ and $P_1, P_2 \subseteq V(G)$ be as in the definition of a type III gadget (so P_1, P_2 are dipaths from z_0, z_1 , respectively, to r, each having length at least 2b - 1). Suppose without loss of generality that $x \in V(P_1)$ (the case that $x \in V(P_2)$ is symmetric). Now define a (2, b)-alternating-path R^* by setting $s_1(R^*) = z_\ell$, $t_1(R^*) = r$, $s_2(R^*) = t_2(R^*) = z_1$, $Q_1(R^*) = (z_\ell, x) \circ P_1[x, r]$ and $Q'_1(R^*) = P_2$, noting that $|Q'_1(R^*)| = |P_2| \ge 2b - 1 \ge b$ by the definition of a type III gadget. Note that $z_0 \notin V(R^*)$ because $x \notin \{z_0, z_1\}$ by assumption. Thus, the oriented path R^* satisfies our requirements.

We have thus shown that for some $1 \le a_1 \le 3$, $G \cup G^*$ contains an (a_1, b) -alternatingpath R^* satisfying $t_{a_1}(R^*) \in \{z_0, z_1\}$, $s_1(R^*) \in \{z_\ell, z^*\}$ as well as $|V(R^*) \cap \{z_0, z_1\}| =$ $|V(R^*) \cap \{z_\ell, z^*\}| = 1$. Let R_1 be the (a_1, b) -alternating-path obtained by combining R^* with the dipaths $P[t_{a_1}(R^*), z_{b+1}]$ and $(P \circ (z_\ell, z^*))[z_{\ell-b}, s_1(R^*)]$. Formally, we set

$$s_1(R_1) = z_{\ell-b}, s_i(R_1) = s_i(R^*), 2 \le i \le a_1,$$

$$t_i(R_1) = t_i(R^*), 1 \le i \le a_1 - 1, t_{a_1}(R_1) = z_{b+1}$$

and

$$Q_1(R_1) = (P \circ (z_{\ell}, z^*))[z_{\ell-b}, s_1(R^*)] \circ Q_1(R^*), Q_{a_1}(R_1) = Q_{a_1}(R^*) \circ P[t_{a_1}(R^*), z_{b+1}],$$
$$Q_i(R_1) = Q_i(R^*), 2 \le i \le a_1 - 1,$$
$$Q'_i(R_1) = Q'_i(R^*), 1 \le i \le a_1 - 1.$$

Note that R_1 is strong, i.e., that $Q_1(R_1)$ and $Q_{a_1}(R_1)$ have length at least b each.

Put $a_2 := a + 2 - a_1$. Then $1 \le a_2 \le a + 1$ because $a \ge 2$ and $1 \le a_1 \le 3$. Now set $\mathcal{C}' := \mathcal{C}[z_{b+1}, z_{\ell-b}]$, noting that $|A_2(\mathcal{C}')| \ge |A_2(\mathcal{C})| - (2b+1) \ge (a+1)(b+1) - 1 \ge a_2(b+1) - 1$. By Lemma 2.14, applied with parameter a_2 , the chain \mathcal{C}' contains a strong (a_2, b) -alternating-path R_2 with $s_1(R_2) = z_{b+1} = t_{a_1}(R_1)$ and $t_{a_2}(R_2) = z_{\ell-b} = s_1(R_1)$. Let us note that $V(R_1) \cap V(R_2) = \{z_{b+1}, z_{\ell-b}\}$, since $V(R_1) \subseteq V(G) \cup V(G^*) \cup \{z_0, \ldots, z_{b+1}\} \cup \{z_{\ell-b}, \ldots, z_{\ell}\}$, $V(R_2) \subseteq V(\mathcal{C}') = V(\mathcal{C}[z_{b+1}, z_{\ell-b}])$ and $V(\mathcal{C}) \cap V(G^*) \subseteq V(G) \cup \{z_\ell\}$, and by the definition of a chain (see Items 3-4 in Definition 2.1). By Observation 2.11, $R_1 \cup R_2$ spans a subdivision of $C_{a,b}$, as required.

2.2.3 Embedding Gadgets

In this section we prove two lemmas, each asserting that one can find certain gadgets in digraphs D possessing some suitable properties. Recall that g is chosen as $g = 4b^2$.

Lemma 2.17. Let D be a digraph of directed girth at least g, and assume that for every $(x,y) \in A(D)$, either D contains a directed cycle of length exactly g through (x,y), or there is $z \in V(D) \setminus \{x,y\}$ such that $(z,x), (z,y) \in A(D)$. Then for every $(p,q) \in A(D)$, there is a type I or extended type II gadget G contained in D such that p(G) = p, q(G) = q and $v(G) \leq 2g$.

Proof. Let $(p,q) \in A(D)$. We inductively define a sequence of vertices r_i , $i \ge 0$, with the property that $(r_i,q) \in A(D)$ for every $i \ge 0$ and $(r_i,r_{i-1}) \in A(D)$ for every $i \ge 1$. Set $r_0 := p$. Let $i \ge 1$, and suppose we have already defined r_0, \ldots, r_{i-1} . By assumption, either D contains a directed cycle C of length exactly g through (r_{i-1},q) , or there is $z \in V(D) \setminus \{r_{i-1},q\}$ such that $(z,r_{i-1}), (z,q) \in A(D)$. In the latter case, we set $r_i := z$. In the former case, we stop, noting that $r_{i-1}, \ldots, r_1, r_0 = p, q, C[q, r_{i-1}]$ is a closed directed walk of length i + (|C| - 1) = g + i - 1 containing the arc (p,q). It follows that if we

stop at step *i*, then there is a directed cycle of length at most g + i - 1 containing (p, q). Therefore, if the process stopped at step *i* for some $0 \le i \le 2b^2 + b - 2$, then *D* must contain a directed cycle of length at most $g + 2b^2 + b - 3 \le 2g$ through (p, q). Moreover, this cycle must have length at least *g* since the directed girth of *D* is at least *g*. So we see that in this case, *D* contains a gadget *G* of type I with p(G) = p, q(G) = q and $v(G) \le 2g$, as required by our claim.

Suppose then that the process carried through to step $2b^2 + b - 2$ (inclusive), and let $r_0, r_1, \ldots, r_{2b^2+b-2}$ be the vertices produced by the process. The vertices $r_{2b^2+b-2}, \ldots, r_1, p$ describe a directed walk in D, all of whose vertices have an arc to q. Let $r := r_{2b^2+b-2}$ and $u := r_{2b^2+b-3}$ denote the first and second vertex of this directed walk, respectively. We now inductively define a sequence of vertices $w_i, i \ge 0$, with the property that $(w_i, u) \in A(D)$ for every $i \ge 0$ and $(w_i, w_{i-1}) \in A(D)$ for every $i \ge 1$. Set $w_0 := r$. Let $i \ge 1$, and suppose we have already defined w_0, \ldots, w_{i-1} . By assumption, either D contains a directed cycle C of length exactly g through (w_{i-1}, u) , or there is $z \in V(D) \setminus \{w_{i-1}, u\}$ such that $(z, w_{i-1}), (z, u) \in A(D)$. In the latter case, we set $w_i := z$. In the former case, we stop and output the directed cycle C.

Suppose first that the process carried through to step b, and let $w_0 = r, w_1, \ldots, w_b$ be the vertices produced by the process. Note that

$$P := (w_b, w_{b-1}, \dots, w_0 = r = r_{2b^2+b-2}, r_{2b^2+b-3}, \dots, r_1, r_0 = p, q)$$

is a directed walk of length at most $2b^2 + 2b - 1$ in *D*. Since $\vec{g}(D) \ge g = 4b^2 > 2b^2 + 2b - 1$, the vertices of *P* must be pairwise distinct. Now set $P_1 := (r = r_{2b^2+b-2}, \ldots, r_1, r_0 = p)$ and $P_2 := (w_b, \ldots, w_1, w_0 = r)$, and observe that P_1 and P_2 satisfy all the requirements in the definition of an extended type II gadget (note that there is an arc from the first vertex of P_2 , namely w_b , to the second vertex of P_1 , namely u, by our choice of the vertices w_i , $i \ge 0$). Moreover, the resulting gadget has $2b^2 + 2b \le 2g$ vertices, as required.

Suppose now that the process stopped at step *i* for some $1 \leq i \leq b$, and let *C* be the outputted directed cycle of length (exactly) *g* through the arc (w_{i-1}, u) . As before, the $2b^2 + b + i - 1$ vertices $w_{i-1}, \ldots, w_0 = r = r_{2b^2+b-2}, \ldots, r_1, r_0 = p, q$ are pairwise distinct because $\vec{g}(D) \geq g = 4b^2 > 2b^2 + b + i - 1$ (as $i \leq b$). Traverse the directed cycle *C* backwards, starting from w_{i-1} , until the first time a vertex *v* in the set $V := \{w_{i-2}, \ldots, w_0 = r_{2b^2+b-2}, \ldots, r_1, p, q\}$ is hit (this will surely happen because $u = r_{2b^2+b-3} \in V(C) \cap V$). By our choice of *v* we have $V(C[v, w_{i-1}]) \cap V = \{v\}$. We now rule out the possibility that $v \in \{w_0, \ldots, w_{i-2}\}$. To this end, suppose by contradiction that $v = w_j$ for some $0 \leq j \leq i-2$. Since $w_{i-1}, w_{i-2}, \ldots, w_j = v, C[v, w_{i-1}], w_{i-1}$ is a directed cycle and $\vec{g}(D) \geq g$, it must be the case that $|C[v, w_{i-1}]| \geq g - (i - 1 - j) \geq g - b + 1$. Now, as *C* consists of the arc (w_{i-1}, u) and the dipaths $C[v, w_{i-1}]$ and C[u, v], we have $|C[u, v]| = |C| - 1 - |C[v, w_{i-1}]| = g - 1 - |C[v, w_{i-1}]| \leq b - 2$. Finally, we get that $v = w_j, (w_j, u), u, C[u, v], v = w_j$ is a (non-trivial) directed closed walk of length at most b - 1 < g, a contradiction.

We have thus shown that $v \notin \{w_0, \ldots, w_{i-2}\}$. If v = q then the vertex-sequence $w_{i-1}, u = r_{2b^2+b-3}, \ldots, r_1, r_0 = p, q = v, C[v, w_{i-1}]$ describes a directed cycle which goes through the arc (p,q) and has length at most $g + 2b^2 + b - 1 \leq 2g$ and at least g. Hence, in this case D contains a gadget of type I with the required properties. It remains to handle the case that $v \in \{u = r_{2b^2+b-3}, \ldots, r_1, r_0 = p\}$. In this case, let s be the second vertex of $C[v, w_{i-1}]$ (so in particular, $(v, s) \in A(D)$). Now define the dipaths $P_1 := (r_{2b^2+b-2} = r, \ldots, r_1, r_0 = p)$ and $P_2 := C[s, w_{i-1}] \circ (w_{i-1}, \ldots, w_1, w_0 = r)$. We claim that $|P_2| \geq b$. Indeed, since $(v, s) \circ P_2 \circ P_1[r, v]$ is a directed cycle in D and $\vec{g}(D) \geq g$, it must be the case that $|P_2| \geq g - 1 - |P_1| = g - 1 - (2b^2 + b - 2) \geq b$. We now see that

all of the requirements in the definition of an extended type II gadget are met (note that the vertex $v \in V(P_1) \setminus \{r\}$ has an arc to the first vertex of P_2 , namely s). Finally, observe that the resulting type II gadget has $2b^2 + b + (i-1) + |C[s, w_{i-1}]| \le 2b^2 + 2b + g \le 2g$ vertices, as required.

Lemma 2.18. Let $b, h, d \ge 1$ be integers, let D' be a digraph and let $v \in V(D')$. Suppose that the following two conditions hold.

- 1. Every vertex of D' reachable from v has out-degree at least $(h+1) \cdot (d(2b-2)+1) + d$;
- 2. There are less than d^h vertices in D' at distance at most (h+1)(2b-1) from v.

Then D' contains a type III gadget G and a dipath P_0 from v to p(G) with the properties $V(P_0) \cap V(G) = \{p(G)\}, v(G) \leq (2h+2)(2b-1) \text{ and } v(P_0) \leq h(2b-1).$

Proof. We describe a process for producing a (specific) out-arborescence $T \subseteq D$ with root v. The idea is as follows: going level by level (in a breadth-first manner), we will try to attach to each vertex u of the (current) lowest level a collection of d dipaths of length 2b-1 each, which intersect only at u and do not intersect the (current) tree in any other vertex. In this manner, we will construct a (2b-1)-subdivision of a d-ary out-arborescence, where an s-subdivision of a digraph F is a subdivision of F in which every arc is replaced with a dipath of length (exactly) s, and a d-ary out-arborescence is an out-arborescence in which every non-leaf vertex has exactly d children. We will then use Item 2 to argue that rather soon in this process, intersections of branches must occur. Such an intersection will give rise to the desired type III gadget. The details follow.

Throughout the process, we will maintain and update an out-arborescence T and sets $L_i, i \ge 0$. We start by setting $L_0 = \{v\}$ and initializing T to be the one-vertex tree with root v. Let $i \ge 0$, and suppose that we have already defined L_0, \ldots, L_i . If $L_i = \emptyset$ then we stop and say that the process terminated at step i. Otherwise, initialize L_{i+1} to be the empty set and proceed as follows. Let u_1, \ldots, u_t be an enumeration of the vertices in L_i . Going over $j = 1, \ldots, t$ in increasing order, we let $\mathcal{P}(u_j)$ be the set of all dipaths of length 2b - 1 which start at u_j and are otherwise disjoint from V(T). If $\mathcal{P}(u_j)$ contains d dipaths Q_1, \ldots, Q_d with $V(Q_k) \cap V(Q_\ell) = \{u_j\}$ for all $1 \le k < \ell \le d$, then attach these dipaths to T and add their endpoints to L_{i+1} . Otherwise, i.e. if \mathcal{P}_j does not contain d dipaths Q_1, \ldots, Q_d which pairwise intersect only at u_j , then do nothing; in this case u_j will remain a leaf of T throughout the process.

Consider the out-arborescence T at the end of the process. It is easy to see that T is indeed the (2b - 1)-subdivision of some d-ary out-arborescence T_0 , and that the branch vertices of this subdivision are precisely the elements of $\bigcup_{i\geq 0} L_i$. It follows that $|L_i| \leq d^i$ for every $i \geq 0$.

We claim that there is $0 \le i \le h$ such that L_i contains a leaf of T. Indeed, suppose by contradiction that for every $0 \le i \le h$, no vertex of L_i is a leaf. Then $|L_i| = d^i$ for every $0 \le i \le h + 1$. Observe that the number of vertices of T which are at distance at most (h+1)(2b-1) from v (in T) is exactly $|L_{h+1}| + \sum_{i=0}^{h} (d(2b-2)+1) \cdot |L_i|$. Hence, the number of such vertices is at least

$$\sum_{i=0}^{h} \left(d(2b-2) + 1 \right) \cdot |L_i| = \left(d(2b-2) + 1 \right) \cdot \sum_{i=0}^{h} d^i \ge d^h,$$

in contradiction to the assumption in Item 2 of the lemma.

So let $0 \le i \le h$ be such that L_i contains a leaf of T, and let $u \in L_i$ be such a leaf. Let X be the set of vertices of T which are at distance at most $(i + 1) \cdot (2b - 1)$ from the root v. In other words, X consists of the sets L_0, \ldots, L_{i+1} and the (vertices of the) subdivision dipaths connecting L_j to L_{j+1} for $j = 0, \ldots, i$. Let us say that a d-tuple of dipaths (Q_1, \ldots, Q_d) is good if

- (a) For every k = 1, ..., d, it holds that $|Q_k| \le 2b 1$, $V(Q_k) \cap X = \{u\}$, and u is the first vertex of Q_k .
- (b) $V(Q_k) \cap V(Q_\ell) = \{u\}$ for all $1 \le k < \ell \le d$.

Choose (Q_1, \ldots, Q_d) among all good *d*-tuples of dipaths such that $|Q_1| + \cdots + |Q_d|$ is maximized (note that taking Q_1, \ldots, Q_d to be empty dipaths (starting at *u*) gives a good *d*-tuple, so the set of good *d*-tuples is non-empty). If we had $|Q_1| = \cdots = |Q_d| = 2b-1$, then the algorithm would have attached Q_1, \ldots, Q_d to *T*, in contradiction to our assumption that *u* is a leaf. Thus, there must be some $1 \le k \le d$ such that $|Q_k| \le 2b-2$. Suppose without loss of generality that $|Q_1| \le 2b-2$, and let *w* be the last vertex of Q_1 . Evidently, *w* is reachable from *u* and hence also from *v*, implying that $d^+(w) \ge (h+1) \cdot (d(2b-2)+1)+d$ by Item 1. If *w* had an out-neighbour in $V(D') \setminus (X \cup V(Q_1) \cup \cdots \cup V(Q_d))$, then we could extend Q_1 and thus obtain a longer good *d*-tuple of dipaths, in contradiction to the maximality of (Q_1, \ldots, Q_d) . Thus, $N^+(w) \subseteq X \cup V(Q_1) \cup \cdots \cup V(Q_d)$. Since we have

$$|N^+(w) \cap (V(Q_1) \cup \dots \cup V(Q_d))| \le |V(Q_1) \cup \dots \cup V(Q_d)| - 1$$
$$|V(Q_1)| + \dots + |V(Q_d)| - (d-1) - 1 \le d \cdot 2b - 1 - d = d(2b-1) - 1,$$

it follows that $|N^+(w) \cap X| \ge d^+(w) - d(2b-1) + 1 \ge hd(2b-2) + h + 2$.

For each vertex $x \in X$, let y(x) denote the lowest common ancestor of u and x in the tree T. Let X' be the set of all vertices $x \in X$ such that (at least) one of the vertices u, x is at distance at most 2b - 2 from y(x) in T. We now show that $|X'| \leq hd(2b-2) + h + 1$. Let P be the unique dipath (in T) from v to u. For each $0 \leq j \leq i - 1$, let y_j be the unique element of $V(P) \cap L_j$. Observe that if $x \in X'$, then either $x \in V(P)$, or there is $0 \leq j \leq i - 1$ such that x is an internal vertex of one of the d-1 subdivision dipaths which start at y_j and are not subpaths of P. (Recall that every non-leaf branching vertex of T is the first vertex of exactly d subdivision dipaths. It is evident that for every $0 \leq j \leq i - 1$, exactly one of the d subdivision dipaths starting at y_j is a subpath of P, while the other d-1 only intersect P at y_j .) It follows that $|X'| = |V(P)| + i \cdot (d-1) \cdot (2b-2) = i \cdot (2b-1) + 1 + i \cdot (d-1) \cdot (2b-2) = i \cdot d \cdot (2b-2) + i + 1 \leq hd(2b-2) + h + 1$, as claimed.

As $|N^+(w) \cap X| \ge hd(2b-2) + h+2 > |X'|$, there exists $x \in N^+(w) \cap (X \setminus X')$. Setting y = y(x), let P'_1 (resp. P'_2) be the unique dipath (in T) from y to x (resp. u). Since $x \notin X'$, we have $|P'_1|, |P'_2| \ge 2b-1$. As $u \in L_i$, we have $|P'_2| \le i(2b-1) \le h(2b-1)$, and as $x \in X$ we have $|P'_1| \le (h+1)(2b-1)$. Let z be the second vertex of P'_2 . Now set $P_0 := P[v, y], P_1 := P'_1$ and $P_2 := P'_2[z, u] \circ Q_1 \circ (w, x)$. Observe that P_0, P_1, P_2 are internally vertex-disjoint (as y is the lowest common ancestor of x and u), and that P_1 and P_2 have length at least 2b-1 each (indeed, we have $|P_1| = |P'_1| \ge 2b-1$ and $|P_2| \ge |P'_2| - 1 + |Q_1| + 1 \ge |P'_2| \ge 2b-1$). So we see that P_1, P_2 form a type III gadget G with p(G) = y and q(G) = z, and that this gadget satisfies $V(P_0) \cap V(G) = \{y\} = \{p(G)\}$. Finally, observe that $|P_0| \le (i-1)(2b-1) \le h(2b-1), |P_1| = |P'_1| \le (h+1)(2b-1)$ and $|P_2| = (|P'_2| - 1) + |Q_1| + 1 \le h(2b-1) + 2b-2 = (h+1)(2b-1) - 1$. It follows that $v(G) = |P_1| + |P_2| + 1 \le (2h+2)(2b-1)$. This completes the proof.

2.2.4 Putting It All Together

Proof of Theorem 2.10. Let $a \ge 2$ and $b \ge 1$. Recall that we set $g := 4b^2$. We also fix an integer $k = O(ab^7)$, to be chosen later. Suppose, for the sake of contradiction, that the theorem is false, and let D be a counterexample to the theorem which minimizes v(D) + a(D). Namely, we assume that $\delta^+(D) \ge k$, $\vec{g}(D) \ge g$ and D does not contain a subdivision of $C_{a,b}$, but every digraph D' with v(D') + a(D') < v(D) + a(D), $\delta^+(D') \ge k$ and $\vec{g}(D') \ge g$ does contain a subdivision of $C_{a,b}$.

Claim 1. $d^+(v) = k$ for every $v \in V(D)$.

Proof. Suppose, by contradiction, that $d^+(v) \ge k + 1$ for some $v \in V(D)$. Let D' be the digraph obtained from D by deleting an (arbitrary) arc whose tail is v. Then $\delta^+(D') \ge k$ and $\vec{g}(D') \ge g$, but D' does not contain a subdivision of $C_{a,b}$ (as D' is a subgraph of D). This contradicts the minimality of D.

Claim 2. For every $(x, y) \in A(D)$, either D contains a directed cycle of length exactly g through (x, y), or there is $z \in V(D) \setminus \{x, y\}$ such that $(z, x), (z, y) \in A(D)$.

Proof. Let $(x, y) \in A(D)$. Suppose by contradiction that the assertion of the claim is false. Let D' be the digraph obtained from D by deleting x and adding the arc (z, y) for every $z \in N_D^-(x)$. Evidently, v(D') + a(D') < v(D) + a(D). We claim that $\delta^+(D') \ge k$ and $\vec{g}(D') \ge g$. First, note that $d_{D'}^+(y) = d_D^+(y) = k$ because $(y, x) \notin A(D)$ (as $\vec{g}(D) \ge g > 2$). Next, observe that for every $z \in V(D') \setminus \{y\} = V(D) \setminus \{x, y\}$ we also have $d_{D'}^+(z) = d_D^+(z) = k$, because z does not have both x and y as out-neighbors (by our assumption). It follows that $\delta^+(D') \ge k$. Now suppose, for the sake of contradiction, that D' contains a directed cycle C' of length at most g - 1. If there is no $z \in N_D^-(x)$ such that $(z, y) \in A(C')$, then C' is also contained in D, which is impossible as $\vec{g}(D) \ge g$. So let $z \in N_D^-(x)$ be such that $(z, y) \in A(C')$, and let C be the directed cycle obtained from C' by deleting the arc (z, y) and adding the arcs (z, x), (x, y). Then C is contained in D and has length $|C'| + 1 \le g$, implying that |C| = g. But this is impossible as we assumed that D contains no directed cycle of length g through the arc (x, y). We conclude that $\vec{g}(D') \ge g$, as claimed.

The minimality of D implies that D' contains a subdivision S' of $C_{a,b}$. If there is no $z \in N_D^-(x)$ such that $(z, y) \in A(S')$, then S' is also contained in D, contradicting our assumption that D contains no subdivision of $C_{a,b}$. We may therefore assume that the set $Z := \{z \in N_D^-(x) : (z, y) \in A(S')\}$ is non-empty. Since the maximum in-degree of $C_{a,b}$ is 2, we have $|Z| \leq 2$. Assume first that |Z| = 1, and write $Z = \{z\}$. By replacing the edge (z, y) of S' with the path (z, x), (x, y) (which is present in D), we obtain a subdivision of $C_{a,b}$ contained in D, a contradiction. Suppose now that |Z| = 2, and write $Z = \{z_1, z_2\}$. Then y must be a branch vertex in S', and we must have $d_{S'}^+(y) = 0$ (since every branch vertex of $C_{a,b}$ is either a source or a sink). Let S be the subgraph of D obtained from S' by deleting the edges $(z_1, y), (z_2, y)$ and adding the edges $(z_1, x), (z_2, x)$. Then S is a subdivision of $C_{a,b}$ in which x plays the branch-vertex role played in S' by y. Again, we arrive at a contradiction to our assumption that D contains no subdivision of $C_{a,b}$.

Let \mathcal{C} be a chain with spine $P = v_0, \ldots, v_m$ and partition $A(P) = A_1 \cup A_2$ (as in Definition 2.1). We say that \mathcal{C} is good if the following conditions are satisfied:

- (a) Every gadget in C has at most (8g+6)(2b-1) vertices;
- (b) $(v_{m-1}, v_m) \in A_2;$
- (c) Among any (4g+3)(2b-1) consecutive arcs of P, there is an arc belonging to A_2 .

Let C be a good chain contained in D of maximal length, and let $P = v_0, \ldots, v_m, A_1, A_2$ and $(G_e)_{e \in A_2}$ be as in Definition 2.1. Define $i_0 := \max \{0, m - (4g+3)(2b-1)(a+3)(b+1)\}$ and $C' := C[v_{i_0}, v_m]$. Item (a) implies that

$$|V(\mathcal{C}')| \le (8g+6)(2b-1) \cdot (m-i_0) \le 2(4g+3)^2(2b-1)^2(a+3)(b+1) \le 8b^2(4g+3)^2(a+3)(b+1).$$
(2.1)

Our choice of i_0 implies that $i_0 = 0$ and $\mathcal{C}' = \mathcal{C}$, or $i_0 = m - (4g+3)(2b-1)(a+3)(b+1)$, in which case we have by Item (c) that

$$|A_2(\mathcal{C}')| \ge \left\lfloor \frac{m - i_0}{(4g + 3)(2b - 1)} \right\rfloor = (a + 3)(b + 1).$$

Let D' be the digraph obtained from D by deleting the vertex-set $V(\mathcal{C}) \setminus \{v_m\}$.

Claim 3. Every $u \in V(D')$ which is reachable from v_m in D' satisfies $d^+_{D'}(u) \ge k - |V(\mathcal{C}')|$.

Proof. If $i_0 = 0$, then $\mathcal{C}' = \mathcal{C}$ and the claim follows directly by definition of D'. So assume now that $i_0 > 0$ and hence $|A_2(\mathcal{C}')| \ge (a+3)(b+1) \ge (a+3)(b+1) - 2$. Let Q be a dipath from v_m to u in D'. Suppose by contradiction that $d_{D'}^+(u) < k - |V(\mathcal{C}')|$. Since $d_D^+(u) \ge k$ and $V(D) \setminus V(D') \subseteq V(\mathcal{C})$, we must have $|N_D^+(u) \cap V(\mathcal{C})| > |V(\mathcal{C}')|$. Hence, there must be some $x \in V(\mathcal{C}) \setminus V(\mathcal{C}')$ such that $(u, x) \in A(D)$. Let $0 \le j \le m$ be such that $x \in G_{(v_j, v_{j+1})}$, and note that $j < i_0$ because $x \notin V(\mathcal{C}')$. Now let \mathcal{C}'' be the chain formed by concatenating $\mathcal{C}[v_j, v_m]$ with the dipath Q (in this chain, all arcs of Q belong to $A_1(\mathcal{C}'')$). This is indeed a chain because $V(Q) \cap V(\mathcal{C}) = \{v_m\}$. As \mathcal{C}' is contained in \mathcal{C}'' , we have $|A_2(\mathcal{C}'')| \ge |A_2(\mathcal{C}')| \ge (a+3)(b+1) - 2$. Observe that we are precisely in the setting of Item 1 of Lemma 2.16 with respect to the chain \mathcal{C}'' . Indeed, the last vertex of the spine of \mathcal{C}'' . So we may apply Item 1 of Lemma 2.16 to deduce that D contains a subdivision of $C_{a,b}$, a contradiction.

Claim 4. Let $u \in V(D')$ be a vertex which has distance at most (4g+3)(2b-1) from v_m in D'. Then D contains a dipath of length at most 2g from $V(\mathcal{C}')$ to u.

Proof. Let $Q = (w_0 = v_m, w_1, \ldots, w_{t-1}, w_t = u)$ be a shortest dipath from v_m to u in D'. Then $t = |Q| \le (4g+3)(2b-1)$. Claim 2 states that D satisfies the condition of Lemma 2.17. By applying Lemma 2.17 to the arc (w_{t-1}, u) , we infer that D contains a gadget G^* which is either of type I or extended type II, such that $p(G^*) = w_{t-1}$, $q(G^*) = u$ and $v(G^*) \le 2g$.

We now show that if $V(G^*) \cap V(\mathcal{C}') \neq \emptyset$, then the assertion of the claim holds. So suppose that $V(G^*) \cap V(\mathcal{C}') \neq \emptyset$, and let $x \in V(G^*) \cap V(\mathcal{C}')$. By Item 2 of Lemma 2.12, G^* contains a dipath from x to $\{w_{t-1}, u\}$, and hence also to u, as $(w_{t-1}, u) \in A(G^*)$. Evidently, this dipath has length at most $v(G^*) \leq 2g$. So we see that D contains a dipath of length at most 2g from $V(\mathcal{C}')$ to u, as required. To complete the proof, it hence suffices to show that $V(G^*) \cap V(\mathcal{C}') \neq \emptyset$. For the rest of the proof we assume, for the sake of contradiction, that $V(G^*) \cap V(\mathcal{C}') = \emptyset$. We proceed by a case analysis over the type of G^* .

Case 1. G^* is an extended gadget of type II. Let G_0^* be the basic part of G^* . We claim that $V(G_0^*) \cap V(Q) = \{w_{t-1}, u\}$. Suppose otherwise, and let $0 \leq j \leq t-2$ be such that $w_j \in V(G_0^*)$. By the definition of a basic type II gadget, every vertex in $V(G_0^*) \setminus \{u\}$ has an arc to $q(G_0^*) = u$. In particular, $(w_j, u) \in A(D)$, and hence also $(w_j, u) \in A(D')$

(as $w_j, u \in V(D')$). It follows that $w_0, \ldots, w_{j-1}, w_j, u$ is a dipath from $w_0 = v_m$ to u in D' which is shorter than Q, in contradiction to our choice of Q. So indeed we have $V(G_0^*) \cap V(Q) = \{w_{t-1}, u\}.$

We claim that $V(G^*) \cap V(\mathcal{C}) \neq \emptyset$. So suppose by contradiction that $V(G^*) \cap V(\mathcal{C}) = \emptyset$. Then one can extend the chain \mathcal{C} into a longer good chain \mathcal{C}_1 by adding the dipath Q and the gadget G_0^* ; the definition of \mathcal{C}_1 includes setting $(w_{t-1}, u) \in A_2(\mathcal{C}_1), G_{(w_{t-1}, u)}(\mathcal{C}_1) = G_0^*$, and $(w_j, w_{j+1}) \in A_1(\mathcal{C}_1)$ for every $0 \leq j \leq t-2$. Then \mathcal{C}_1 is indeed a chain because $V(G_0^*) \cap V(Q) = \{w_{t-1}, u\}$ and due to our assumption that $V(G^*) \cap V(\mathcal{C}) = \emptyset$. The goodness of \mathcal{C}_1 (i.e. that \mathcal{C}_1 satisfies Items (a)-(c) above) follows from the goodness of \mathcal{C} and the fact that $|Q| \leq (4g+3)(2b-1)$ and $v(G_0^*) \leq 2g \leq (8g+6)(2b-1)$. As the existence of \mathcal{C}_1 stands in contradiction to the maximality of \mathcal{C} , our assumption $V(G^*) \cap V(\mathcal{C}) = \emptyset$ must have been wrong, as required.

We have thus shown that $V(G^*) \cap V(\mathcal{C}) \neq \emptyset$. Since $V(G^*) \cap V(\mathcal{C}') = \emptyset$ by assumption, we must have $V(G^*) \cap (V(\mathcal{C}) \setminus V(\mathcal{C}')) \neq \emptyset$. This means that $V(G^*) \cap V(G_{(v_i,v_{i+1})}) \neq \emptyset$ for some $0 \leq i < i_0$ (as $V(\mathcal{C}) \setminus V(\mathcal{C}')$ is contained in the union of $V(G_{(v_i,v_{i+1})})$ over all $0 \leq i < i_0$). Let i_1 be the largest such $0 \leq i < i_0$, and set $G := G_{(v_{i_1},v_{i_1+1})}$. Now let \mathcal{C}_1 be the chain obtained by attaching to $\mathcal{C}[v_{i_1}, v_m]$ the dipath $Q - u = (w_0 = v_m, w_1, \dots, w_{t-1})$. This is indeed a chain because $V(Q) \cap V(\mathcal{C}) = \{v_m\}$ (as $V(Q) \subseteq V(D')$ and $(V(\mathcal{C}) \setminus \{v_m\}) \cap V(D') = \emptyset$). Then $|A_2(\mathcal{C}_1)| \geq |A_2(\mathcal{C}')| \geq (a+3)(b+1)-2$ because \mathcal{C}_1 contains \mathcal{C}' and $i_0 > 0$. Observe that Condition 2 in Lemma 2.16 holds for the chain \mathcal{C}_1 with respect to the vertex $z^* := u$ (and with $z_\ell = w_{t-1}, z_0 = v_{i_1}$ and $z_1 = v_{i_1+1}$). Indeed, there is an arc from the last vertex of the spine of \mathcal{C}_1 , namely w_{t-1} , to $u \notin V(\mathcal{C}) \cap V(\mathcal{C}^*) \neq \emptyset$ and $V(\mathcal{C}_1) \cap V(\mathcal{C}^*) \subseteq V(\mathcal{G}) \cup \{w_{t-1}\}$ (here we use our choice of i_1). By Lemma 2.16, D contains a subdivision of $C_{a,b}$, a contradiction.

Case 2. G^* is of type I, i.e., a directed cycle of length at least g through (w_{t-1}, u) . Let j be the smallest integer in $\{0, \ldots, t-1\}$ satisfying $w_j \in V(G^*)$; note that j is well-defined because $w_{t-1} \in V(G^*)$. Let w' be the vertex of the directed cycle G^* immediately following w_j , and consider the dipath $Q' := (w_0 = v_m, w_1, \ldots, w_j, w')$. Our choice of j implies that $w' \notin \{w_0, \ldots, w_{j-1}\}$ (so Q' is indeed a path) and that $V(G^*) \cap V(Q') = \{w_j, w'\}$. Note also that $j \ge 1$ because $w_0 = v_m \in V(\mathcal{C}')$ and $V(G^*) \cap V(\mathcal{C}') = \emptyset$ by assumption. Hence, we have $w_j \notin V(\mathcal{C})$.

Similarly to the previous case, if $V(G^*) \cap V(\mathcal{C}) = \emptyset$, then one can extend \mathcal{C} into a longer good chain \mathcal{C}_1 by adding the dipath Q' and the gadget G^* ; the definition of \mathcal{C}_1 includes setting $(w_j, w') \in A_2(\mathcal{C}_1)$, $G_{(w_j, w')}(\mathcal{C}_1) = G^*$, and $(w_i, w_{i+1}) \in A_1(\mathcal{C}_1)$ for every $0 \le i < j$. Then \mathcal{C}_1 is indeed a chain because $V(G^*) \cap V(Q') = \{w_j, w'\}$ and $V(G^*) \cap V(\mathcal{C}) = \emptyset$, and the goodness of \mathcal{C}_1 follows from the goodness of \mathcal{C} and the fact that $|Q| \le (4g+3)(2b-1)$ and $v(G^*) \le 2g \le (8g+6)(2b-1)$. So we see that having $V(G^*) \cap V(\mathcal{C}) = \emptyset$ contradicts the maximality of \mathcal{C} , and hence $V(G^*) \cap V(\mathcal{C}) \ne \emptyset$.

Walk along the directed cycle G^* , starting from w_j , until the first time that a vertex of $V(\mathcal{C})$ is met. Denote this vertex by x, and the preceding vertex on G^* by y. Consider the dipath $Q'' := (w_0 = v_m, w_1, \ldots, w_j) \circ G^*[w_j, y]$, and observe that $V(Q'') \cap V(\mathcal{C}) = \{v_m\}$ because $V(Q) \cap V(\mathcal{C}) = \{v_m\}$ and by our choice of x. Since $V(G^*) \cap V(\mathcal{C}') = \emptyset$, we must have $x \in V(G^*) \cap (V(\mathcal{C}) \setminus V(\mathcal{C}')) \neq \emptyset$. This means that $x \in V(G_{(v_i, v_{i+1})})$ for some $0 \leq i < i_0$. Now let \mathcal{C}_1 be the chain obtained by concatenating $\mathcal{C}[v_i, v_m]$ with the dipath Q''. This is indeed a chain because $V(Q'') \cap V(\mathcal{C}) = \{v_m\}$. Then $|A_2(\mathcal{C}_1)| \geq |A_2(\mathcal{C}')| \geq (a+3)(b+1)-2$ because \mathcal{C}_1 contains \mathcal{C}' and $i_0 > 0$. Observe that Condition 1 in Lemma 2.16 holds for the chain \mathcal{C}_1 (with y playing the role of z_ℓ). Indeed, there is an arc from the last vertex of

the spine of C_1 , namely y, to $x \in V(G_{(i,i+1)})$, and (v_i, v_{i+1}) is the first arc of the spine of C_1 . By Lemma 2.16, D contains a subdivision of $C_{a,b}$, a contradiction. This completes the proof of Claim 4.

With Claims 3-4 at hand, we can complete the proof of the theorem. To this end, we will apply Lemma 2.18. By combining Claim 4 with the fact that $\Delta^+(D) = k$, we conclude that the number of vertices of D' at distance at most (4g+3)(2b-1) from v_m (in D') is at most $|V(\mathcal{C}')| \cdot k^{2g}$. We will apply Lemma 2.18 with parameters h := 4g + 2and d := 2b(4g+3)(a+3)(b+1). To this end, we will need to verify that

$$|V(\mathcal{C}')| \ge (4g+3) \cdot (d(2b-2)+1) + d \text{ and } |V(\mathcal{C}')| \cdot k^{2g} < d^{4g+2}.$$
 (2.2)

This is the point where we choose the value of k; set $k := 12b^2(4g+3)^2(a+3)(b+1)$, noting that $k = O(ab^7)$ because $g = 4b^2$. Both inequalities in (2.2) follow from (2.1) and our choice of k and d. Indeed, we have:

$$\begin{aligned} |V(\mathcal{C}')| + (4g+3) \cdot (d(2b-2)+1) + d &\leq |V(\mathcal{C}')| + (4g+3) \cdot 2db \\ &\leq 8b^2(4g+3)^2(a+3)(b+1) + (4g+3) \cdot 2db \\ &\leq 12b^2(4g+3)^2(a+3)(b+1) = k, \end{aligned}$$

and

$$\begin{split} V(\mathcal{C}')|\cdot k^{2g} &\leq 8b^2(4g+3)^2(a+3)(b+1)\cdot k^{2g} \\ &= 8\cdot 12^{2g}\cdot b^{4g+2}(4g+3)^{4g+2}(a+3)^{2g+1}(b+1)^{2g+1} \\ &= 8\cdot 12^{2g}\cdot 2^{-4g-2}\cdot (a+3)^{-2g-1}(b+1)^{-2g-1}\cdot d^{4g+2} \\ &= 2\cdot 3^{2g}\cdot (a+3)^{-2g-1}(b+1)^{-2g-1}\cdot d^{4g+2} < d^{4g+2}. \end{split}$$

Claims 3 and 4 imply that D' satisfies Conditions 1 and 2 in Lemma 2.18, respectively, with the role of v played by v_m , and with the parameters h and d chosen above. By Lemma 2.18, D' contains a type III gadget G and a dipath P_0 from v_m to p(G) such that $V(P_0) \cap V(G) = \{p(G)\}, v(G) \leq (2h+2)(2b-1) = (8g+6)(2b-1)$ and $v(P_0) \leq$ $h(2b-1) = (4g+2)(2b-1) \leq (4g+3)(2b-1) - 1$. Now, let C_1 be the chain formed by appending to C the dipath P_0 and the gadget G; so the spine of C_1 is $P \circ P_0 \circ (p(G), q(G))$, $A_1(C_1) = A_1(C) \cup A(P_0)$ and $A_2(C_1) = A_2(C) \cup \{(p(G), q(G))\}$. It is easy to see that C_1 is indeed a chain and that it satisfies Conditions (a)-(c) above. But this contradicts the maximality of C. This final contradiction means that our initial assumption, that D is a counterexample to Theorem 2.10, was false. This completes the proof of the theorem. \Box

2.3 Disjoint Oriented Cycles

In this section, we prove Theorem 2.4. The proof relies on the notion of *directed tree-width* introduced in 1, and the *Directed Grid Theorem* by Kawarabayashi and Kreutzer [KK15]. The crucial ingredient of the proof is Proposition 2.19, a digraph-splitting result which gives a positive answer to Problem 2.1 for digraphs of bounded directed tree-width.

Proposition 2.19. Let $k, \ell, d \in \mathbb{N}$, and let D be a digraph such that $dtw(D) \leq d$ and $\delta^+(D) \geq (d+1)(\ell-1)+k$. Then there exist disjoint non-empty subsets $X_1, \ldots, X_\ell \subseteq V(D)$ such that $\delta^+(D[X_i]) \geq k$ for every $i \in [\ell]$.

Let us note at this point that in contrast to undirected graphs, digraphs of bounded directed tree-width may have arbitrarily large minimum out-degree, and hence, Proposition 2.19 addresses a non-trivial class of digraphs. An illustrating example is constructed in the following remark.

Remark 2.20. For any $k \in \mathbb{N}$ there exists a digraph F_k such that $\delta^+(F_k) = k$ and $\operatorname{dtw}(F_k) = 1$.

Proof. Let $k \in \mathbb{N}$ be fixed. Let us denote by T_k the unique k-ary out-arborescence of depth k (that is, every non-leaf vertex has k children, and every leaf has distance k from the root). Let r denote the root of T_k . Let F_k be the digraph obtained from T_k by adding the arc (u, v) for every leaf $u \in V(T_k)$ and every vertex $v \in V(T_k)$ such that u > v in T_k . Since every leaf of T_k has k ancestors, D_k has minimum out-degree k. To see that $dtw(F_k) = 1$, let us consider the directed tree-decomposition (T, β, γ) of F_k , where $T := T_k, \beta(t) := \{t\}$ for every $t \in V(T)$ and $\gamma(e) := \{tail(e)\}$ for every arc $e \in A(T)$. To see that this forms a directed tree-decomposition, let $e = (u, v) \in A(T)$ be arbitrary. Then $S := \bigcup \{\beta(t) | t \in V(T), e < t\}$ is the set of vertices in F_k contained in the subtree of T_k rooted at v. We need to show that S is $\gamma(e) = \{u\}$ -normal in D. However, in $F_k - u$, there exists no arc starting in a vertex outside S and ending in S, which directly shows that a directed walk in $F_k - u$ ending in S must be contained in β and γ we have $\Gamma(t) = \beta(t) \cup \bigcup \{\gamma(e) | e \in A(T), e \sim t\} = \{t, \operatorname{parent}(t)\}$ for every vertex $t \in V(T) \setminus \{r\}$, and $\Gamma(r) = \{r\}$. It follows that the width of (T, β, γ) is 2-1 = 1, proving that $dtw(F_k) = 1$.

Proposition 2.19 will be derived as a consequence of Lemma 2.22 below, which shows that classes of *strongly connected* digraphs possess the so-called *Erdős-Pósa-property* within digraphs of bounded directed tree-width. Here a class of (di)graphs \mathcal{H} is said to have the Erdős-Pósa-property within another class \mathcal{G} of (di)graphs, if there exists a function $f: \mathbb{N} \to \mathbb{N}$ such that for every $\ell \in \mathbb{N}$ and every digraph $D \in \mathcal{G}$ one of the following holds:

- D contains ℓ vertex-disjoint subdigraphs, all members of \mathcal{H} , or
- there exists a subset $X \subseteq V(D)$ such that $|X| \leq f(\ell)$ and D-X contains no member of \mathcal{H} as a subdigraph.

For the proof of Lemma 2.22 we need a result from [AKKW16]. In the following, given a tree-decomposition (T, β, γ) of a digraph, we use the notation

$$\beta(\geq t) := \bigcup_{t' \in V(T), t' \geq t} \beta(t')$$

for the vertices of D contained in the bags of the sub-arborescence of T rooted at t.

Lemma 2.21 (cf. Lemma 3.6 in [AKKW16]). Let (T, β, γ) be a directed tree-decomposition of a digraph D and let H be a strongly connected subdigraph of D. Let r be the root of T, and let $t^* \in V(T)$ be a node in T of maximal distance from r in T such that $\beta(\geq t^*) \supseteq V(H)$. Then for every $t \in V(T)$ with $t \geq t^*$ and $\beta(\geq t) \cap V(H) \neq \emptyset$, we have $\Gamma(t) \cap V(H) \neq \emptyset$.

Proof. Suppose towards a contradiction that $\Gamma(t) \cap V(H) = \emptyset$. This means that in particular $\beta(t) \cap V(H) = \emptyset$, and since $\beta(\geq t) \cap V(H) \neq \emptyset$, there must be a child t' of t in T such that $\beta(\geq t') \cap V(H) \neq \emptyset$. Since $t' > t \geq t^*$, the definition of t^* implies that $V(H) \setminus \beta(\geq t') \neq \emptyset$. Let us denote $e = (t, t') \in A(T)$ and note that $\beta(\geq t') = \bigcup \{\beta(s) | s \in V(T), e < s\}$ by definition of a tree-decomposition, is a $\gamma(e)$ -normal set in D.

Pick some $x \in V(H) \cap \beta(\geq t')$, $y \in V(H) \setminus \beta(\geq t')$ arbitrarily. Since H is strongly connected, there exist directed paths P_1 , P_2 in H from x to y (resp. y to x). The concatenation $W = P_1 \circ P_2$ of these two paths yields a directed walk in $H \subseteq D$ starting and ending in $\beta(\geq t')$ and intersecting $V(D) \setminus \beta(\geq t')$. Since $\beta(\geq t')$ is $\gamma(e)$ -normal, it follows that $\gamma(e) \cap V(H) \supseteq \gamma(e) \cap V(W) \neq \emptyset$. However, this contradicts the fact $\gamma(e) \subseteq \Gamma(t)$ and our assumption that $\Gamma(t) \cap V(H) = \emptyset$. This shows that this assumption was wrong, concluding the proof.

Lemma 2.22. Let \mathcal{H} be a set of strongly connected digraphs, and let $k, d \in \mathbb{N}$. The following holds for every digraph D with $dtw(D) \leq d$:

For every $\ell \in \mathbb{N}$, D contains ℓ vertex-disjoint subdigraphs contained in \mathcal{H} , or there is a subset $X \subseteq V(D)$ with $|X| \leq (d+1)(\ell-1)$ such that D - X contains no member of \mathcal{H} as a subdigraph.

Proof. Let us proof the claim by induction on ℓ . If $\ell = 1$ the claim holds trivially. Now suppose that $\ell \geq 2$ and the claim has been proved for $\ell - 1$. Let (T, β, γ) be a directed tree-decomposition of D of width at most d. We may assume w.l.o.g. that D contains at least one member of \mathcal{H} as a subdigraph, for otherwise the claim holds trivially. Since the sets $\beta(t), t \in V(T)$ partition V(D), there exists at least one $t \in V(T)$ such that $D[\beta(\geq t)]$ contains a member of \mathcal{H} . Let $t_0 \in V(T)$ be a vertex with this property maximizing the distance to the root in T. Let $D' := D - \beta \geq t_0$. Since dtw(·) is monotone under taking subgraphs, we have $dtw(D') \leq dtw(D) \leq d$. By the induction hypothesis, there exist $\ell - 1$ vertex-disjoint subdigraphs of D' which are members of \mathcal{H} , or we can hit all members of \mathcal{H} contained in D' by a set $X' \subseteq V(D')$ with $|X'| \leq (d+1)(\ell-2)$. In the first case, we can simply join a collection of $\ell - 1$ vertex-disjoint members of \mathcal{H} contained in D' by an (arbitrarily chosen) digraph in \mathcal{H} contained in $D[\beta(\geq t_0)]$ to obtain a collection of ℓ vertex-disjoint subdigraphs of D contained in \mathcal{H} , which in this case concludes the proof. So we may assume that we are in the second case and that the set X' with the stated property exists. We now define $X := X' \cup \Gamma(t_0) \subseteq V(D)$, noting that $|X| \leq |X'| + |\Gamma(t_0)| \leq (d+1)(\ell-2) + (d+1) = (d+1)(\ell-1)$, since (T, β, γ) has width at most d. We claim that D-X contains no subdigraph contained in \mathcal{H} , which if true, proves the inductive claim. Suppose towards a contradiction that there exists a digraph $H \in \mathcal{H}$ such that $H \subseteq D$ and $V(H) \cap X = \emptyset$. Since D' - X' contains no subdigraph contained in \mathcal{H} , we must have $V(H) \cap \beta \geq t_0 \neq \emptyset$. Let $t^* \in V(T)$ be such that $\beta \geq t^* \geq V(H)$ and such that the distance of the root to t^* in T is maximized. There are three possible cases concerning the relationship of t_0 and t^* in T: (1) t_0 and t^* are incomparable, (2) $t_0 < t^*$ and (3) $t_0 \ge t^*$. In case (1), we have that $\beta(\ge t_0)$ and $\beta(\ge t^*)$ are disjoint, yielding a contradiction since $V(H) \cap \beta \geq t_0 \neq \emptyset$ and $V(H) \subseteq \beta \geq t^*$. In case (2), t^* is a vertex in V(T) which is further from the root than t_0 , but still $\beta \geq t^*$ contains a member of \mathcal{H} (namely H). This yields a contradiction to the definition of t_0 . In case (3), noting that H is a strongly connected subdigraph of D, we can apply Lemma 2.21 to find that $\Gamma(t_0) \cap V(H) \neq \emptyset$. This contradicts the facts $V(H) \cap X = \emptyset$ and $\Gamma(t_0) \subseteq X$. Since all three cases lead to a contradiction, our assumption was wrong, X indeed hits all the subdigraphs of D contained in \mathcal{H} . This concludes the inductive proof of the claim.

We can now give the proof of Proposition 2.19.

Proof of Proposition 2.19. Let $k, \ell, d \in \mathbb{N}$, and let D be a digraph such that $\delta^+(D) \geq (d+1)(\ell-1) + k$ and $\operatorname{dtw}(D) \leq d$. Let \mathcal{H} be the set of all strongly connected digraphs H satisfying $\delta^+(H) \geq k$. By applying Lemma 2.22 to D and \mathcal{H} we find that either

(1) D contains ℓ vertex-disjoint subdigraphs in \mathcal{H} , or (2) there exists $X \subseteq V(D)$ such that

 $|X| \leq (d+1)(\ell-1)$ such that D-X contains no member of \mathcal{H} as a subdigraph.

If (1) occurs, then the vertex sets X_1, \ldots, X_ℓ of the ℓ vertex-disjoint subdigraphs in \mathcal{H} confirm the claim of the Proposition. Hence, to conclude the proof it suffices to rule out (2). Suppose towards a contradiction that a set $X \subseteq V(D)$ as in (2) exists. Then we have $\delta^+(D-X) \ge \delta^+(D) - |X| \ge (d+1)(\ell-1) + k - (d+1)(\ell-1) = k$. Let $U \subseteq V(D-X)$ be the vertex-set of a strong component of D-X which has no arcs leaving it (such a component always exists, as can be seen by considering a topological sorting of the strong components). Then clearly, $\delta^+(D[U]) \ge \delta^+(D-X) \ge k$ and D[U] is strongly connected. It follows that $D[U] \in \mathcal{H}$ and $D[U] \subseteq D-X$, which contradicts the assumptions on X in (2). This concludes the proof of the Proposition.

Finally, we are prepared for the proof of Theorem 2.4.

Proof of Theorem 2.4. Let F_1, \ldots, F_r be δ^+ -maderian topological grid minors. Let us denote by $d: \mathbb{N} \to \mathbb{N}$ be the function from the Directed Grid Theorem 1.3. For every $i \in [r]$ there exists $K_i \in \mathbb{N}$ such that every digraph D with $\delta^+(D_i) \geq K_i$ contains a subdivision of F_i . Further, since D_i is a topological grid minor, there is $k_i \in \mathbb{N}$ such that the elementary wall W_{k_i} contains a subdivision of F_i . Let $K := \max\{K_1, \ldots, K_r\}$, $k := k_1 + \cdots + k_r$ and let D be any given digraph such that $\delta^+(D) \ge d(2k)(r-1) + K$. We claim that then D must contain a subdivision of $F := F_1 \cup \cdots \cup F_r$, which will then show that F is δ^+ -maderian, as claimed by the Theorem. For this we apply the Directed Grid Theorem to D and find that either $dtw(D) \le d(2k) - 1$ or that $dtw(D) \ge d(2k)$, and hence, D contains a subdivision of W_k . In the first case, we have $\delta^+(D) \ge (d+1)(r-1) + K$ with d := d(2k) - 1, and hence Proposition 2.19 implies that there exist disjoint subsets X_1, \ldots, X_r such that $\delta^+(D[X_i]) \ge K \ge K_i$ for $i = 1, \ldots, r$, implying that $D[X_i]$ contains a subdivision of F_i for every *i*. The union of these subdivisions then clearly defines a subdivision of F in D and concludes the proof in this case. In the second case D contains a subdivision of W_k . Note that W_k contains r vertex-disjoint subdigraphs isomorphic to W_{k_1}, \ldots, W_{k_r} , and by definition of k_i the copy of W_{k_i} contains a subdivision of F_i for every $1 \leq i \leq k$. Hence, D contains a subdivision of W_k , which in turn contains a subdivision of F. This shows that D contains an F-subdivision in any case and concludes the proof.

2.4 Oriented Cycles with Two Blocks

In this section, we prove Theorem 2.6. We will repeatedly use the following observation:

Lemma 2.23. Let $\ell_1, \ell_2 \in \mathbb{N}$, and D a digraph with $\delta^+(D) \ge \ell_1 + \ell_2$. Then for every $v \in V(D)$, there are dipaths P_1 and P_2 in D of length ℓ_1 and ℓ_2 , respectively, which start in v and satisfy $V(P_1) \cap V(P_2) = \{v\}$.

Proof. Greedily build two disjoint dipaths starting at v by attaching out-neighbors at their ends until they have lengths ℓ_1 and ℓ_2 , respectively.

Proof of Theorem 2.6. Let D be an arbitrary digraph such that $\delta^+(D) \ge k_1 + 3k_2 - 5$. We have to show that there exist two internally vertex-disjoint dipaths in D which start and end in the same vertices, one of length at least k_1 , the other of length at least k_2 . Throughout the proof, we will say that a dipath P in D with terminal vertex x is k_2 -good if there exist dipaths P_1 and P_2 of length $k_2 - 1$ starting at x such that $V(P_1) \cap V(P_2) = \{x\}$ and $V(P_i) \cap V(P) = \{x\}$ for $i \in \{1, 2\}$. Note that D contains a k_2 -good dipath of positive length. Indeed, choose some arbitrary vertex $u \in V(D)$ and some out-neighbor v of u. Since $\delta^+(D-u) \ge \delta^+(D) - 1 \ge k_1 + 3k_2 - 6 \ge (k_2 - 1) + (k_2 - 1)$, we can apply Lemma 2.23 with $\ell_1 := \ell_2 := k_2 - 1$ to the vertex v in the digraph D - u to infer that P := (u, v) is a k_2 -good dipath. Let P_0 be a longest k_2 -good dipath in D. We have just shown that $|P_0| > 0$. Denote by x the end-vertex of P_0 and by P_1, P_2 two dipaths of length $k_2 - 1$ starting in x such that $V(P_i) \cap V(P_j) = \{x\}$ for $i \neq j \in \{0, 1, 2\}$. Let a be the terminus of P_1 and b the terminus of P_2 .

Claim 1. There exist dipaths P_a, P_b which start in a and b respectively and end in vertices $a', b' \in V(P_0) \setminus \{x\}$ such that both P_a and P_b are internally vertex-disjoint from $V(P_0) \cup V(P_1) \cup V(P_2)$.

Proof. We prove the existence of P_a and a'; the proof for the existence of P_b and b' is completely analogous. Let $D' := D - ((V(P_1) \cup V(P_2)) \setminus \{a\})$. As $|(V(P_1) \cup V(P_2)) \setminus \{a\}| =$ $2k_2-2$, we have that $\delta^+(D') \ge \delta^+(D) - (2k_2-2) \ge k_1+k_2-3 \ge 2k_2-3$. Let $R \subseteq V(D')$ be the set of vertices reachable from a by a dipath in D'. We claim that $R \cap (V(P_0) \setminus \{x\}) \neq \emptyset$. Suppose towards a contradiction that $R \cap (V(P_0) \setminus \{x\}) = \emptyset$. Since for every vertex $r \in R$ we have $N_{D'}^+(r) \subseteq R$, we see that $\delta^+(D'[R]) \geq \delta^+(D') \geq 2k_2 - 3$. We can now apply Lemma 2.23 to the vertex a of D'[R] with $\ell_1 := k_2 - 1, \ell_2 := k_2 - 2$ and find that D'[R] contains dipaths P'_1 and P'_2 of lengths ℓ_1 and ℓ_2 , respectively, which start at a and satisfy $V(P'_1) \cap V(P'_2) = \{a\}$. Let w be the end-vertex of the path P'_2 . We have $N_D^+(w) \cap (V(P_0) \setminus \{x\}) \subseteq R \cap (V(P_0) \setminus \{x\}) = \emptyset$. Since $d_D^+(w) \ge k_1 + 3k_2 - 5 > 3k_2 - 4 = 0$ $|V(P_1) \cup V(P'_1) \cup (V(P'_2) \setminus \{w\})|$, there must exist $w' \in N_D^+(w) \setminus (V(P_1) \cup V(P'_1) \cup V(P'_2))$. Now the dipaths P'_1 and $P'_2 \circ (w, w')$ are of length $k_2 - 1$, have only the starting vertex a in common and are disjoint from the set $(V(P_0) \cup V(P_1)) \setminus \{a\}$. Hence, $P_0 \circ P_1$ is a k_2 -good dipath in D which is longer than P_0 , a contradiction. This shows that indeed $R \cap (V(P_0) \setminus \{x\}) \neq \emptyset$. Hence, by the definition of R, there is a shortest dipath P_a from a to $R \cap (V(P_0) \setminus \{x\})$ in D'[R]. Write $V(P_a) \cap (V(P_0) \setminus \{x\}) =: \{a'\}$. Now P_a and a'satisfy the claimed properties.

Let $A, B \subseteq V(P_0) \setminus \{x\}$ be the sets of vertices on $P_0 - x$ reachable from a, b, respectively, by a dipath which is internally vertex-disjoint from $V(P_0) \cup V(P_1) \cup V(P_2)$. By the previous claim we have $A, B \neq \emptyset$. Let a^* respectively b^* denote the vertex in A respectively Bwhose distance from x on P_0 is maximum. By symmetry, we may assume without loss of generality that $\operatorname{dist}_{P_0}(a^*, x) \geq \operatorname{dist}_{P_0}(b^*, x)$. Hence, $B \subseteq V(P_0[a^*, x])$. Fix some dipath P_{a^*} from a to a^* in D which is internally disjoint from $V(P_0) \cup V(P_1) \cup V(P_2)$. Set $Q := P_1 \circ P_{a^*}$, and note that $|Q| = |P_1| + |P_{a^*}| = k_2 - 1 + |P_{a^*}| \geq k_2$. Let $Q' \subseteq Q$ be defined as follows: if the length of Q is at most k_1 then Q' := Q, and otherwise Q' is the unique subpath of Q which starts at x and has length exactly k_1 . In the following, let rdenote the length of Q'. Observe that $r = |Q'| = \min\{|Q|, k_1\}$, and hence $k_2 \leq r \leq k_1$. Moreover, $P_1 \subseteq Q'$ because P_1 consists of the first k_2 vertices of Q. Let $y \in V(Q)$ be the terminus of Q', and let us define B^* as the subset of B consisting of those vertices in $B \subseteq V(P_0 - x)$ which are reachable from b by a dipath which is internally vertex-disjoint from $V(P_0) \cup V(Q) \cup V(P_2)$.

Claim 2. We either have $|B^*| \ge k_1 - r + 1$, or there exists a dipath starting at b and ending in V(Q) which is vertex-disjoint from $(V(P_0) \cup V(Q') \cup V(P_2)) \setminus \{b, y\}$.

Proof. Suppose towards a contradiction that $|B^*| \leq k_1 - r$ but there exists no dipath starting at b and ending in V(Q) which is disjoint from $(V(P_0) \cup V(Q') \cup V(P_2)) \setminus \{b, y\}$. Let us consider the digraph $D'' := D - ((V(P_0) \cup V(Q') \cup V(P_2)) \setminus \{b, y\})$. Let $R \subseteq V(D'')$ denote
the set of vertices reachable from b in D". By our assumption, we have $R \cap V(Q) = \emptyset$, and hence $R \cap (V(P_0) \cup V(Q) \cup V(P_2)) = \{b\}$ (since $R \subseteq V(D'')$ and by the definition of D"). We claim that $N_D^+(u) \cap (V(P_0) \setminus \{x\}) \subseteq B^*$ for all $u \in R$. Indeed, let $u \in R$ and $v \in N_D^+(u) \cap (V(P_0) \setminus \{x\})$. By definition, there exists a b-u-dipath P_u in D", and $V(P_u) \subseteq R$. Then the dipath $P_u \circ (u, v)$ starts at b, ends in $V(P_0) \setminus \{x\}$ and is internally vertex-disjoint from $V(P_0) \cup V(Q) \cup V(P_2)$, certifying that $v \in B^*$.

Since
$$|(V(Q') \cup V(P_2)) \setminus \{y, b\}| = r + k_2 - 2$$
, for every $u \in R$ we have:
 $d_{D''}^+(u) \ge d_D^+(u) - |N_D^+(u) \cap (V(P_0) \setminus \{x\})| - |(V(Q') \cup V(P_2)) \setminus \{y, b\}|$
 $\ge k_1 + 3k_2 - 5 - |B^*| - (r + k_2 - 2) \ge 2k_2 - 3$,

where in the last inequality we used our assumption that $|B^*| \leq k_1 - r$. As $N_{D''}^+(u) \subseteq R$ for every $u \in R$ (by the definition of R), we get that $\delta^+(D''[R]) \geq 2k_2 - 3$. Applying Lemma 2.23 to the vertex b in D''[R] with $\ell_1 := k_2 - 1, \ell_2 := k_2 - 2$, we find dipaths P_1'' and P_2'' in D''[R] starting at b of lengths ℓ_1 and ℓ_2 , respectively, such that $V(P_1'') \cap V(P_2'') = \{b\}$. By the definition of D'' we have $V(P_i'') \cap (V(P_0) \cup V(P_2)) = \{b\}$ for every i = 1, 2. Let z denote the end-vertex of P_2'' . We have $|V(P_2) \cup V(P_1'') \cup (V(P_2'') \setminus \{z\})| = 3k_2 - 4$ and $|N_D^+(z) \cap (V(P_0) \setminus \{x\})| \leq |B^*| \leq k_1 - r \leq k_1 - k_2$. Here we used the fact that $z \in R$ and hence $N_D^+(z) \cap (V(P_0) \setminus \{x\}) \subseteq B^*$. So we see that

$$|N_D^+(z) \setminus (V(P_0) \cup V(P_2) \cup V(P_1'') \cup V(P_2''))| \ge k_1 + 3k_2 - 5 - (k_1 - k_2) - (3k_2 - 4) = k_2 - 1 > 0.$$

Let $z' \notin V(P_0) \cup V(P_2) \cup V(P_1'') \cup V(P_2'')$ be an out-neighbor of z. The two dipaths P_1'' and $P_2'' \circ (z, z')$ start at b and have length $k_2 - 1$ each. Moreover, the three dipaths P_1'' , $P_2'' \circ (z, z')$ and $P_0 \circ P_2$ intersect each other only in the vertex b. Hence, $P_0 \circ P_2$ is a k_2 -good dipath in D which is strictly longer than P_0 , a contradiction. This contradiction shows that our initial assumption was wrong, concluding the proof of Claim 2.

We will now show how to find a subdivision of $C(k_1, k_2)$ in D using Claim 2. Consider the two alternatives in the conclusion of this claim. The first case is that $|B^*| \ge k_1 - r + 1$. Since $B^* \subseteq B \subseteq V(P_0[a^*, x])$, this clearly implies that there exists a vertex $b^* \in B^*$ whose distance from a^* on the dipath P_0 is at least $k_1 - r$. By definition of B^* , there exists a dipath P_{b^*} in D starting in b and ending at b^* which is internally disjoint from $V(P_0) \cup V(Q) \cup V(P_2)$. Now the two dipaths $Q \circ P_0[a^*, b^*]$ and $P_2 \circ P_{b^*}$ in D both start at x and end at b^* , are internally vertex-disjoint, and have lengths $|Q| + |P_0[a^*, b^*]| \ge$ $r + k_1 - r = k_1$ and $|P_2| + |P_{b^*}| \ge k_2 - 1 + 1 = k_2$, respectively. Hence, they form a subdivision of $C(k_1, k_2)$.

The second case is that there exists a dipath in D starting at b and ending in V(Q), which is vertex-disjoint from $(V(P_0) \cup V(Q') \cup V(P_2)) \setminus \{b, y\}$. Let P^* be a shortest such dipath, and let $q \in V(Q)$ denote its end-vertex. Then clearly $V(P^*) \cap V(Q) = \{q\}$, as well as $q \notin V(Q') \setminus \{y\}$ and $q \neq a^*$ (as $a^* \in V(P_0)$). This readily implies that $Q' \neq Q$, and hence by definition of Q' we conclude that Q' has length exactly k_1 . Let us consider the two dipaths Q[x,q] and $P_2 \circ P^*$ in D, which both start in x and end in q. They are internally vertex-disjoint, and have lengths $|Q[x,q]| \geq |Q'| = k_1$ and $|P_2| + |P^*| \geq k_2 - 1 + 1 = k_2$, respectively. Hence, they form a subdivision of $C(k_1, k_2)$ in D.

Summarizing, we have shown that D contains a subdivision of $C(k_1, k_2)$ in all the cases, which concludes the proof of the theorem.

2.5 Subdivisions of $K_3 - e$

In this section we give a proof of Theorem 2.2. As it turns out, it is convenient to prove the following slightly stronger result, which clearly implies that mader_{δ^+}($\vec{K}_3 - e$) = 2. **Proposition 2.24.** Let D be a digraph and $v_0 \in V(D)$ such that $d^+(v_0) \ge 1$ and $d^+(v) \ge 2$ for every $v \in V(D) \setminus \{v_0\}$. Then D contains a subdivision of $\vec{K}_3 - e$.

Proof. Suppose towards a contradiction that the claim is false, and let D be a counterexample which minimizes v(D) with first priority and a(D) with second priority. Let $v_0 \in V(D)$ be a vertex such that $d^+(v_0) \ge 1$ and $d^+(v) \ge 2$ for all $v \in V(D) \setminus \{v_0\}$.

Claim 1. We have $d^+(v_0) = 1$ and $d^+(v) = 2$ for all $v \in V(D) \setminus \{v_0\}$.

Proof. If $d^+(v_0) > 1$ or $d^+(v) > 2$ for some $v \in V(D) \setminus \{v_0\}$, then we may delete an arc of D and be left with a digraph D' which still satisfies $d_{D'}^+(v_0) \ge 1$ and $d_{D'}^+(v) \ge 2$ for every $v \in V(D) \setminus \{v_0\}$. This contradicts the assumed minimality of D (as D' evidently contains no subdivision of $\vec{K}_3 - e$ either).

Claim 2. *D* is strongly connected.

Proof. If not, then there is $\emptyset \neq X \subsetneq V(D)$ such that no arc of D leaves X. Then clearly $d_{D[X]}^+(x) = d^+(x)$ for all $x \in X$, and hence D[X] meets the conditions of the Lemma. But as D[X] contains no subdivision of $\overset{\leftrightarrow}{K}_3 - e$ and is smaller than D, we get a contradiction to the minimality of D.

Claim 3. There exists no partition (W, K, Z) of V(D) such that $W, Z \neq \emptyset$, $v_0 \in K \cup Z$, $|K| \leq 1$ and there is no arc in D with tail in W and head in Z.

Proof. Suppose towards a contradiction that a partition (W, K, Z) with the described properties exists. Since *D* is strong, we must have |K| = 1; say $K = \{s_0\}$ for some vertex $s_0 \in V(D)$. Since $v_0 \notin W$ and since no arc of *D* goes from *W* to *Z*, every vertex in *W* has out-degree 2 in $D[W \cup \{s_0\}]$. Since *D* is strongly connected, there must be an s_0 -*W*-dipath *P* in *D*. Denoting the last vertex of *P* by $w \in W$, we note that $V(P) \setminus \{s_0, w\} \subseteq Z$. Let *D'* be the digraph obtained from $D[W \cup \{s_0\}]$ by adding the arc (s_0, w) . We clearly have $d_{D'}^+(s_0) \geq 1$, as well as $d_{D'}^+(v) = 2$ for every $v \in W$ by the above. Since v(D') < v(D), the minimality of *D* implies that *D'* contains a subdivision *S'* of $\vec{K}_3 - e$. If *S'* does not use the arc (s_0, w) then $S' \subseteq D$. And otherwise, the subdigraph $S \subseteq D$ of *D* defined by $V(S) := V(S') \cup V(P), A(S) := (A(S') \setminus \{(s_0, w)\}) \cup A(P)$ forms a subdivision of $\vec{K}_3 - e$ in *D*. In both cases we obtain a contradiction to our assumption that *D* does not contain a subdivision of $\vec{K}_3 - e$. This concludes the proof of the claim.

In the following, let $v_1 \in V(D)$ denote the unique out-neighbor of v_0 . The rest of the proof is divided into two cases depending on whether v_0 and v_1 have common in-neighbors.

Case 1. $N^{-}(v_0) \cap N^{-}(v_1) = \emptyset$. Since $d^+(v_1) = 2$, there exists $v_2 \in N^+(v_1) \setminus \{v_0\}$. Let D' be the digraph obtained from $D - v_1$ by adding the arc (v_0, v_2) and the arcs (x, v_0) for all $x \in N_D^-(v_1) \setminus \{v_0\}$. We clearly have $d_{D'}^+(v_0) = 1$ and $d_{D'}^+(v) = 2$ for all $v \in V(D') \setminus \{v_0\}$, since no vertex in D has arcs to both v_0 and v_1 . Since v(D') < v(D), there must be a subdivision S' of $\vec{K}_3 - e$ contained in D'. If $v_0 \notin V(S')$, then S' is a subdigraph of D, which contradicts our assumption that D contains no $(\vec{K}_3 - e)$ -subdivision. Hence we must have $v_0 \in S'$. Since v_2 is the only out-neighbor of v_0 in D', we must have $d_{S'}^+(v_0) = 1$ and $(v_0, v_2) \in A(S')$. We now distinguish between two subcases depending on the indegree of v_0 in S'. Note that every vertex of $\vec{K}_3 - e$ has in-degree either 1 or 2. Hence, $d_{S'}^-(v_0) \in \{1, 2\}$.

Case 1(a). $d_{S'}^{-}(v_0) = 1$. Let $x_0 \in N_{D'}^{-}(v_0)$ be the unique in-neighbor of v_0 in S'. By definition of D', we must have either $x_0 \in N_D^{-}(v_0) \setminus \{v_1\}$ or $x_0 \in N_D^{-}(v_1) \setminus \{v_0\}$. Define a subdigraph $S \subseteq D$ of D as follows: If $x_0 \in N_D^{-}(v_0) \setminus \{v_1\}$, then we put $V(S) := V(S') \cup \{v_1\}$ and $A(S) := (A(S') \setminus \{(v_0, v_2)\}) \cup \{(v_0, v_1), (v_1, v_2)\}$, and if $x_0 \in N_D^{-}(v_1) \setminus \{v_0\}$, then we put $V(S) := (V(S') \setminus \{v_0\}) \cup \{v_1\}$, $A(S) := (A(S') \setminus \{(x_0, v_0), (v_0, v_2)\}) \cup \{(x_0, v_1), (v_1, v_2)\}$. It is easy to see that in each case S is isomorphic to a subdivision of S', and hence forms a subdivision of $\vec{K}_3 - e$ contained in D, a contradiction to our assumption on D.

Case 1(b). $d_{S'}^-(v_0) = 2$. Let $x_1, x_2 \in N_{D'}^-(v_0)$ be the two in-neighbors of v_0 in S'. By definition of D', we have $x_i \in N_D^-(v_0) \setminus \{v_1\}$ or $x_i \in N_D^-(v_1) \setminus \{v_0\}$ for each i = 1, 2. Let us define a subdigraph $S \subseteq D$ of D as follows. Firstly, if $x_1, x_2 \in N_D^-(v_0) \setminus \{v_1\}$, then we set $V(S) := V(S') \cup \{v_1\}$ and $A(S) := (A(S') \setminus \{(v_0, v_2)\}) \cup \{(v_0, v_1), (v_1, v_2)\}$. Secondly, if $x_i \in N_D^-(v_0) \setminus \{v_1\}$ and $x_{3-i} \in N_D^-(v_1) \setminus \{v_0\}$ for some $i \in \{1, 2\}$, then we set $V(S) := V(S') \cup \{v_1\}$ and $A(S) := (A(S') \setminus \{(v_0, v_2), (x_{3-i}, v_0)\}) \cup \{(v_0, v_1), (v_1, v_2), (x_{3-i}, v_1)\}$. Lastly, if $x_1, x_2 \in N_D^-(v_1) \setminus \{v_0\}$ then we set $V(S) := (V(S') \setminus \{v_0\}) \cup \{v_1\}$ and $A(S) := (A(S') \setminus \{(v_0, v_2), (x_1, v_1), (x_2, v_1)\}$. It is easy to check that in each of the three cases, S is isomorphic to a subdivision of S', and hence forms a subdivision of $K_3 - e$ which is contained in D. This contradiction to our initial assumption on D rules out Case 1.

Case 2. There exists a vertex $z_0 \in N^-(v_0) \cap N^-(v_1)$. Let now $A := \{v_0, z_0\}$ and apply Theorem 1.2 to the vertex v_1 versus the set A in D. We conclude that either there are two v_1 -A-dipaths intersecting only at v_1 , or there is a set $K \subseteq V(D) \setminus \{v_1\}$ such that $|K| \leq 1$ and there is no dipath in D - K starting in v_1 and ending in A.

In the first case, let P_1 and P_2 be dipaths such that $V(P_1) \cap V(P_2) = \{v_1\}$ and such that P_1 ends in v_0 , while P_2 ends in z_0 . Now the subdigraph $S \subseteq D$ with vertex set $V(S) := V(P_1) \cup V(P_2)$ and arc-set $A(S) := A(P_1) \cup A(P_2) \cup \{(v_0, v_1), (z_0, v_0), (z_0, v_1)\}$ forms a subdivision of $K_3 - e$ with branch vertices v_0, v_1, z_0 . This is a contradiction to our initial assumption on D.

In the second case, let $W \subseteq V(D) - K$ be the subset of vertices reachable from v_1 by a dipath in D - K and let $Z := V(D) \setminus (W \cup K)$. Since there is no v_1 -A-dipath in D - K, we must have $v_0 \in A \subseteq K \cup Z$. We further have $v_1 \in W$ and $A \setminus K \subseteq Z$, hence $W, Z \neq \emptyset$. Moreover, by definition of W, no arc in D starts in W and ends in Z. All in all, this shows that the partition (W, K, Z) of V(D) yields a contradiction to Claim 3.

Since we arrived at contradictions in all possible cases, we conclude that our initial assumption about the existence of D was wrong. This concludes the proof.

2.6 Subdivisions and Arc-Connectivity

In this section we give the proofs of Propositions 2.7 and 2.8, which show that \vec{K}_4 and \vec{S}_4 are not κ' -maderian.

Proof of Proposition 2.7. A construction of Thomassen [Tho85b] shows that for every integer $k \geq 1$, there exists a digraph D_k such that $\delta^+(D_k) = k$ and D_k contains no directed cycle of even length. For every $k \geq 1$ let \overleftarrow{D}_k denote the digraph obtained from D_k by reversing all its arcs. Then clearly we have $\delta^-(\overleftarrow{D}_k) = k$. Let G'_k be the digraph obtained from the vertex-disjoint union of a copy of D_k with vertex-set A and a copy of \overleftarrow{D}_k with vertex-set B by adding all the arcs in $B \times A$ (i.e., all arcs from B to A). Note that since $|A| = |B| = v(D_k) > k$, we have $\delta^+(G'_k) = \delta^-(G'_k) = k$. Finally, we define G_k as the digraph obtained from G'_k by adding a vertex $v \notin V(G'_k)$ as well as all arcs (v, x), (x, v) for $x \in V(G'_k)$. We claim that G_k is strongly k-arc-connected. Indeed, let $E \subseteq A(G_k)$ be a

set of arcs such that |E| < k. We claim that in $G_k - E$, every vertex $x \in V(G'_k)$ can reach and is reachable from v via a dipath. This will show that $G_k - E$ is strongly connected, as required. Let $x \in V(G'_k)$ be given arbitrarily, and let $x_1, \ldots, x_k \in V(G'_k)$ be k pairwise distinct out-neighbors of x in G'_k . Consider the k arc-disjoint dipaths $P_i := (x, x_i) \circ (x_i, v)$, $i = 1, \ldots, k$. At least one of these dipaths must be disjoint from E and hence constitute an x-v-dipath in D - E. With a symmetric argument considering k distinct in-neighbors of x, we also obtain that there is a v-x-dipath in $G_k - E$, as required. We further claim that G_k contains no subdivision of K_4 . Indeed, suppose this was the case, then clearly there would be $S \subseteq G_k - v = G'_k$ such that S is a subdivision of K_3 . As is easy to see, Smust contain an even directed cycle. Since there is no arc in G'_k from A to B, we find that this cycle must be entirely contained in either $G'_k[A] \simeq D_k$ or $G'_k[B] \simeq \tilde{D}_k$. This however means that D_k contains an even directed cycle, a contradiction. This contradiction shows that G_k contains no subdivision of \tilde{K}_4 , and this concludes the proof. \Box

Proof of Proposition 2.8. A construction of Thomassen [Tho85b] shows that for every integer $k \ge 1$, there exists a digraph R_k such that $\delta^+(R_k) = k$ and R_k contains no subdivision of the bioriented 3-star \tilde{S}_3 . For $k \geq 1$, let us denote by \tilde{R}_k the digraph obtained from R_k by reversing all its arcs. Let H'_k be the digraph obtained from the disjoint union of a copy of R_k with vertex-set A and a copy of R_k with vertex-set B by adding all the arcs in $B \times A$. Since R_k and R_k have at least k vertices, we obtain that $\delta^+(H'_k) = \delta^-(H'_k) = k$. We now define H_k to be the digraph obtained from two disjoint copies of H'_k with vertex-sets X and Y by adding two distinct new vertices u and v as well as the following arcs: (u, x)and (x, v) for every $x \in X$, and (y, u) and (v, y) for every $y \in Y$. We claim that H_k is strongly k-arc-connected. Indeed, let $E \subseteq A(H_k)$ be an arbitrarily given set of arcs such that |E| < k. We must prove that $H_k - E$ is strongly connected. For this, it clearly suffices to show that in $H_k - E$, every vertex in X can reach v and is reachable from u, and every vertex in Y can reach u and is reachable from v. Let $x \in X$ be any given vertex, and let $x_1^-, \ldots, x_k^- \in X$ denote k distinct in-neighbors of x in $H_k[X] \simeq H'_k$. Among the k arc-disjoint u-x-dipaths $(u, x_i^-) \circ (x_i^-, x), i = 1, \dots, k$ in H_k , at least one must also exist in $H_k - E$, and hence x is reachable from u in $H_k - E$. Similarly, considering k distinct out-neighbors $x_1^+, \ldots, x_k^+ \in X$ of x in $H_k[X]$, and considering the arc-disjoint x-v-dipaths $(x, x_i^+), (x_i^+, v), i = 1, \dots, k$, we find that there is an x-v-dipath in $H_k - E$. With a symmetric argument for the vertices in Y, we can verify the above claim, showing that $H_k - E$ is strongly connected. This shows that indeed $\kappa'(H_k) \geq k$.

Next we claim that H_k does not contain a subdivision of \ddot{S}_4 . Suppose otherwise. Then there exists a vertex $w \in V(H_k)$ and directed cycles C_1, C_2, C_3, C_4 in H_k such that $w \in V(C_i)$ for $i = 1, \ldots, 4$, and such that the sets $V(C_i) \setminus \{w\}, 1 \leq i \leq 4$, are pairwise disjoint. Suppose first that $w \in \{u, v\}$. Without loss of generality, we may assume that w = u (the case w = v is symmetric). Then for each $1 \leq i \leq 4$, $C_i - w$ is a dipath which starts in X and ends in Y (since the vertex of C_i preceding w = u must be in Y, and the vertex of C_i succeeding w = u must be in X). It follows that $C_i - w, 1 \leq i \leq 4$, are pairwise vertex-disjoint dipaths from X to Y, contradicting the fact that X and Y can be disconnected in H_k by deleting only two vertices, namely u and v. Suppose now that $w \in X \cup Y$. Note that every directed cycle in H_k is either contained in $H_k[X]$, or contained in $H_k[Y]$, or contains both u and v. It follows that if $w \in X$, then at least three of the cycles $C_i, 1 \leq i \leq 4$, are contained in $H_k[Y] \simeq H'_k$. So we see that in each case, H'_k must contain a subdivision of \ddot{S}_3 . Since every subdivision of \ddot{S}_3 is a strongly connected digraph, and since there are no arcs from A to B in H'_k , we find that this subdivision must be entirely contained in either $H'_k[A] \simeq R_k$ or $H'_k[B] \simeq \overline{R}_k$. Since S_3 is invariant under the reversal of all arcs, we obtain that in each case R_k must contain a subdivision of S_3 . This contradicts our initial assumptions on the sequence $(R_k)_{k\geq 1}$. This contradiction proves the claim of the proposition; namely, H_k is indeed a k-strongly arc connected digraph not containing S_4 as a subdivision.

2.7 Conclusion

In this concluding section, we mention further open problems related to subdivisions in digraphs of large minimum out-degree, which were discovered during the work presented in this chapter.

Theorem 2.3 shows that orientations of cycles are δ^+ -maderian, and that for an orientation C of a cycle, mader_{δ^+}(C) grows polynomially in |C|. Aboulker et al. actually conjectured the very explicit bound of mader_{δ^+}(C) $\leq 2|C| - 1$ (cf. [ACH⁺19], Conjecture 27). However, it is unclear to us whether mader_{δ^+}(C) should be linear in |C| at all.

Problem 2.4. Does it hold that $mader_{\delta^+}(C) = O(|C|)$ for every orientation C of a cycle?

We remark that Theorem 2.10 gives a positive answer to this question when the size of a longest block in C is bounded by a constant.

Digraph subdivision is a natural graph operation under which it is plausible to expect that the δ^+ -maderian property is preserved.

Conjecture 2.5. If a digraph F is δ^+ -maderian, then all the subdivisions of F are δ^+ -maderian as well.

Conjecture 2.5 would follow if we could show that every digraph of sufficiently large outdegree contains a subdivision of some digraph of out-degree k in which every subdivision path is long.

Conjecture 2.6. There is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathbb{N}$ and for every digraph D with $\delta^+(D) \ge f(k)$, there exists a digraph D' such that $\delta^+(D') \ge k$ and D contains a subdivision of D' in which every subdivision-path has length at least two.

An important step towards Conjecture 2.2 would be to show that attaching an out-leaf to any vertex of a δ^+ -maderian digraph yields still a δ^+ -maderian digraph.

Conjecture 2.7. If F is a δ^+ -maderian digraph, $v_0 \in V(F)$ and F^* is the digraph obtained from F by adding a new vertex v_1 and the arc (v_0, v_1) , then F^* is δ^+ -maderian as well.

Conjecture 2.7 would follow directly from the following natural statement. We call a set of vertices X in a digraph D an *in-dominating set* if every $y \in V(D) \setminus X$ has at least one out-neighbor in X.

Conjecture 2.8. There exists a function $f : \mathbb{N} \to \mathbb{N}$ such that the following holds for every $k \geq 1$. If D is a digraph with $\delta^+(D) \geq f(k)$, then there exists an in-dominating set $X \subsetneq V(D)$ such that $\delta^+(D-X) \geq k$.

Another interesting direction is to characterize the undirected graphs F for which the biorientation \vec{F} of F is δ^+ -maderian. If \vec{F} is δ^+ -maderian, then F must be a forest, since every bioriented cycle has arc-connectivity two and hence is not δ^+ -maderian (see the necessary properties of δ^+ -maderian digraphs mentioned in the introduction). Furthermore,

it is known that \vec{S}_3 is not δ^+ -maderian [Tho85b]. Thus, if \vec{F} is δ^+ -maderian then F must be a path-forest. Thomassen's result [Tho83] shows that a biorientation of any matching is δ^+ -maderian. By Theorem 2.2, $\vec{S}_2 = \vec{P}_3$ is δ^+ -maderian (where P_ℓ denotes the path on ℓ vertices). The first open case is that of \vec{P}_4 .

Problem 2.5. Is $\overrightarrow{P}_4 \delta^+$ -maderian?

Finally, several open problems arise from the questions considered in Section 2.6. Given that \vec{K}_4 and \vec{S}_4 are not κ' -maderian (see Propositions 2.7-2.8), it is natural to ask whether \vec{K}_3 and \vec{S}_3 are.

Problem 2.6. Is $\vec{K}_3 \kappa'$ -maderian? Is $\vec{S}_3 \kappa'$ -maderian?

As mentioned in the introduction, every subdivision of \vec{K}_3 contains an even dicycle, and one cannot force an even dicycle by means of minimum out-degree [Tho85b]. Thus, even dicycles can be thought of as an obstacle to forcing subdivisions of \vec{K}_3 . Interestingly, this obstacle disappears when considering arc-connectivity (rather than out-degree), as a theorem of Thomassen [Tho92] shows that every digraph D with $\kappa'(D) \geq 3$ contains an even dicycle. This can be thought of as a hint that \vec{K}_3 could in fact be κ' -maderian.

A critical first step towards the resolution of Problem 2.3 for the strong vertexconnectivity would be to answer the following.

Problem 2.7. Is there a constant $K \in \mathbb{N}$ such that every K-strongly-vertex connected digraph contains two vertices $x \neq y$ and four pairwise internally vertex-disjoint dipaths, two from x to y and two from y to x?

Chapter 3

Disjoint Directed Cycles with Length Constraints

3.1 Introduction

All graphs and digraphs considered in this chapter are simple.

Cycles are amongst the most fundamental graph objects, and a large body of work has dealt with understanding which conditions force a graph to contain cycles of certain types and lengths. In many cases, the conditions require the graphs to be sufficiently dense, which usually means that they should have either large minimum degrees or should be highly connected.

The problem of packing vertex-disjoint cycles in undirected graphs of large minimum degree has received a lot of attention, see [CY18] for a survey on this topic. For example, for every $k \in \mathbb{N}$ an undirected graph whose minimum degree is sufficiently large in terms of k contains

- k vertex-disjoint cycles of equal lengths [Häg85, Alo96, Ega96],
- k vertex-disjoint cycles of pairwise distinct lengths [BHL⁺17],
- k vertex-disjoint cycles of even lengths [CFKS14].

In each case, the precise degree bounds required to obtain k disjoint cycles with these properties are known.

In contrast, the analogous problems for directed graphs in large amounts remain open. In 1981, Bermond and Thomassen [BT81] stated the following conjecture.

Conjecture 3.1. For every $k \in \mathbb{N}$, every digraph D with $\delta^+(D) \ge 2k - 1$ contains k vertex-disjoint directed cycles.

As mentioned in the previous chapter, Thomassen [Tho83] was the first to prove that for every k there exists a finite number f(k) such that all digraphs of minimum outdegree f(k) contain k disjoint directed cycles, and he showed that $f(k) \leq (k+1)!$. This estimate was improved by Alon [Alo96] to $f(k) \leq 64k$, and then further by Bucić [Buc18] to $f(k) \leq 18k$, which remains the state of the art.

Thomassen [Tho83] also conjectured that his result may be strengthened in the sense that digraphs of sufficiently large minimum out-degree as a function of k always contain kdisjoint directed cycles of equal lengths. This conjecture was disproved by Alon [Alo96], who gave a construction of digraphs with arbitrarily large minimum out-degree containing no two arc-disjoint directed cycles of equal lengths. On the other hand, Henning and Yeo [HY12] initiated the study of the question whether every digraph of sufficiently large minimum out-degree contains two disjoint directed cycles of *distinct* lengths. This was resolved by Lichiardopol [Lic14] in the positive, proving that every digraph D with $\delta^+(D) \ge 4$ contains two vertex-disjoint directed cycles of distinct lengths. Lichiardopol conjectured the following far-reaching qualitative generalization of his result:

Conjecture 3.2 (cf. [Lic14]). For every $k \in \mathbb{N}$ there exists an integer $g(k) \in \mathbb{N}$ such that every digraph D with $\delta^+(D) \ge g(k)$ contains k vertex-disjoint directed cycles of pairwise distinct lengths.

Remarkably, Lichiardopol's conjecture remains open even for k = 3. In addition to Lichiardopol's proof for the existence of q(2), quite some work has been done on forcing two vertex-disjoint directed cycles of distinct lengths in more special classes of digraphs, see [GM13, Tan14, Tan15, Tan17, Tan20] for some results on this topic. Bensmail, Harutyunyan, Le, Li, and Lichiardopol [BHL⁺17] took up Lichiardopol's conjecture and proved it in some special cases, namely for tournaments (orientations of complete graphs), for regular digraphs (all vertices have out- and in-degree exactly r for some number $r \in \mathbb{N}$) and for digraphs of bounded order. In the latter two cases, the rough idea of their probabilistic proof is as follows: We randomly split the vertex-set of the digraph into k parts. Then every vertex is expected to have roughly a $\frac{1}{k}$ -fraction of its out-neighbors in its part, and using Chernoff's bound a large deviation from the expected value is very unlikely. Finally, it can be proved using the Lovász Local Lemma (or a union bound in the bounded order case) that with positive probability, the random partition gives a splitting into vertexdisjoint subdigraphs with still large minimum out-degree, and in each of those parts we can find directed cycles of many distinct lengths. Grouping the cycles from different parts together, it is then possible to obtain many disjoint directed cycles with distinct lengths.

It however seems difficult to extend the proof methods of Bensmail et al. to general digraphs. For this recall from Chapter 2, Problem 2.1 that whether or not digraphs with large minimum out-degree contain vertex-disjoint subdigraphs with given minimum out-degree is a fundamental open problem posed by several researchers. Concrete reasons why the method of Bensmail et al. does not work for general digraphs are that it relies on (1) special properties of tournaments (such as the fact that strong tournaments contain directed cycles of every length) or (2) the assumption that the digraph is *balanced*, i.e. that its maximum in-degree is upper-bounded by a function of its minimum out-degree, a requirement for the Lovász Local Lemma in their argument for regular digraphs.

Our results. A natural strengthening of the condition of having large minimum degree is to require large connectivity. The main result of this chapter verifies Lichiardopol's conjecture for digraphs of large (strong) connectivity. This is the first result guaranteeing arbitrarily many vertex-disjoint directed cycles of distinct lengths which is applicable to general digraphs and allows for digraphs with very unbalanced out- and in-degrees.

Theorem 3.1. For every $k \ge 1$ there exists an integer s(k) such that every strongly s(k)connected digraph contains k vertex-disjoint directed cycles of pairwise distinct lengths.

The relationship between minimum degree, strong connectivity and the existence of directed cycles of particular lengths in digraphs has received attention previously. For instance, a long-standing open problem posed by Lovász [Lov75, Tho92] asked whether (1) every digraph of sufficiently large minimum out-degree contains an even directed cycle, and (2) whether every digraph of sufficiently large strong connectivity contains an even directed cycle. Both questions were resolved by Thomassen in [Tho83] and [Tho92], answering (1) in the negative and (2) in the positive.

3.1. INTRODUCTION

With Theorem 3.1 at hand, it is natural to ask whether a sufficiently high degree of strong connectivity also guarantees the existence of many vertex-disjoint directed cycles of equal length. We can show that the answer to this question is, maybe surprisingly, negative¹.

Proposition 3.2. For every $k \in \mathbb{N}$ there exists a strongly k-connected digraph D_k which contains no two arc-disjoint directed cycles (and hence no two vertex-disjoint cycles) of equal length.

Our next result makes progress on the question concerning the existence of g(3), guaranteeing three vertex-disjoint directed cycles of distinct lengths in digraphs of large outand in-degree.

Theorem 3.3. There exists $K \in \mathbb{N}$ such that every digraph D with $\delta^+(D), \delta^-(D) \ge K$ contains three vertex-disjoint directed cycles of pairwise distinct lengths.

Our proofs for Theorem 3.1 and Theorem 3.3 make use of the *Directed Flat Wall Theorem*, a tool from structural digraph theory established recently by Giannopoulou, Kawarabayashi, Kreutzer, and Kwon [GKKK20]. We believe it might be fruitful to investigate applications of this structural result to other packing problems in digraphs.

Theorem 3.1 shows that a counterexample to Lichiardopol's conjecture, if it exists, cannot contain any large well-connected parts. Intuitively, this means that the digraph does not have a "rich" directed cycle structure, and other methods might apply to fully resolve the conjecture in this case. To illustrate this intuition, we prove Lichiardopol's conjecture for digraphs whose *directed tree-width* is bounded. Recall that directed tree-width is supposed to be small for digraphs which have a "sparse" directed cycle structure (we refer to Chapter 1 for details on this parameter).

Proposition 3.4. Let $k, d \in \mathbb{N}$. Every digraph D with $\delta^+(D) > (d+2)(k-1)$ and directed tree-width at most d contains k vertex-disjoint directed cycles of pairwise distinct lengths.

Recall from Chapter 2, Remark 2.20 that digraphs of large minimum out-degree need not have large directed tree-width. Hence, Proposition 3.4 verifies Lichiardopol's conjecture for a non-trivial class of digraphs.

Structure of the chapter. In Section 11.2 we introduce important definitions and make some preliminary observations required later in the chapter. Most importantly, we explain the statement of the *Directed Flat Wall Theorem* from [GKKK20]. In Section 3.3 we give the proofs of Theorem 3.1 and Theorem 3.3. The directed flat wall theorem yields a natural division of our proofs into three different cases, which are prepared separately in the Subsections 3.3.1, 3.3.2 and 3.3.3. In Subsection 3.3.1 we study digraphs of bounded directed tree-width and give the proof of Proposition 3.4. In Subsection 3.3.2 we show that digraphs containing a large complete butterfly-minor have many disjoint directed cycles of distinct lengths. In Subsection 3.3.3 we study digraphs containing large flat walls and prove sufficient conditions for the existence of disjoint directed cycles of different lengths. In Subsection 3.3.4 we put the insights from the previous subsections together to conclude the proofs of Theorems 3.1 and 3.3. Finally, in Section 3.4, we present the construction for Proposition 3.2. We conclude with a conjecture in Section 3.5.

¹Note that Proposition 3.2 is not implied by Alon's construction of digraphs having large minimum out-degree without two disjoint cycles of equal length, since the digraphs constructed in [Alo96] contain vertices of in-degree 1, and hence are not strongly 2-connected.

3.2 Preliminaries

We start this chapter by introducing some important definitions and auxiliary results related to directed tree-width.

Similar to the Directed Grid Theorem, which relates high directed tree-width to the existence of large grid minors, there exists yet another concept dual to directed tree-width, called *havens*, which was introduced in [JRST01].

Definition 3.1. Let D be a digraph and let $k \in \mathbb{N}$. A haven of order k for D is a function h, assigning to every set $X \subseteq V(D)$ of size |X| < k a set $h(X) \subseteq V(D) \setminus X$, which forms a strong component of D - X, and which has the property that $h(X) \supseteq h(Y)$ for every $X \subseteq Y \subseteq V(D)$ with |Y| < k.

Theorem 3.5 (cf. [JRST01], Theorem 3.1). Let D be a digraph. If D admits a haven of order k, then $dtw(D) \ge k - 1$.

Let us observe the following simple consequence of Theorem 3.5 for later use.

Corollary 3.6. Let $k \in \mathbb{N}$, and let D be a strongly k-connected digraph. Then $dtw(D) \ge k$.

Proof. Let us define $h : \binom{V(D)}{\leq k} \to 2^{V(D)}$ as follows: For every $X \subseteq V(D)$ with $|X| \leq k$, we let $h(X) := V(D) \setminus X$, if |X| < k, while h(X) is defined as the vertex-set of an arbitrarily chosen strong component of D - X if |X| = k. We claim that h is a haven of order k + 1 for D. Indeed, for every $X \subseteq V(D)$ such that |X| < k, D - X is strongly connected and hence $h(X) = V(D) \setminus X$ is the unique strong component of D - X. If |X| = k, then by definition h(X) is a strong component of D - X. Let now $X \subseteq Y \subseteq V(D)$ be arbitrary such that $|Y| \leq k$. If |Y| < k, then we clearly have $h(X) = V(D) \setminus X \supseteq V(D) \setminus Y$. If |Y| = k, then either X = Y and h(X) = h(Y) or |X| < k and $h(X) = V(D) \setminus X \supseteq h(Y)$, since h(Y) is a strong component of $D - Y \subseteq D - X$. This shows that indeed, h is a haven of order k + 1 in D and the claim follows by applying Theorem 3.5.

Let us now turn to a strengthening of the Directed Grid Theorem, the so-called *Directed Flat Wall Theorem* established recently by Giannopoulou, Kawarabayashi, Kreutzer, and Kwon [GKKK20]. To do so, we need a definition of *flatness* of a wall contained in a digraph. Giannopoulou et al. (cf. [GKKK20], Definition 2.14) give a complex definition including 3 different items. For our purposes, only the properties of a flat wall guaranteed by the second item of their definition are required. To not complicate matters unnecessarily, we use a weakened definition of flatness as stated below. Intuitively, the definition asserts that directed paths in the digraph do not jump far across the wall. Recall from Chapter 1 that a *cylindrical wall* is any subdivision of an elementary cylindrical wall of arbitrary order. We start with some additional terminology.

Definition 3.2 (Perimeter, Bricks cf. [GKKK20]). Let W be a cylindrical wall of order k. The perimeter per(W) of W is defined as the union $V(Q_1) \cup V(Q_k)$, where Q_1 and Q_k are the first resp. last vertical directed cycles in the canonical embedding of W. We further define the interior of W by $int(W) := V(W) \setminus per(W)$.

A brick of W is a cycle in the canonical embedding of W induced by the boundary of an inner face of W (i.e., distinct from the two faces bounded by the cycles Q_1 and Q_k). Every brick contains exactly 6 branch vertices of W.

Definition 3.3. Let D be a digraph, and let $W \subseteq D$ be a cylindrical wall. We say that W is weakly flat in D, if for every directed path in D which is internally vertex-disjoint from int(W) with both endpoints x, y contained in int(W), there exists a brick B of W such that $x, y \in V(B)$.

We can now finally state a weakened version of the Directed Flat Wall Theorem from [GKKK20] as follows.

Theorem 3.7 (cf. [GKKK20], Theorem 2.3). For every $k, t \in \mathbb{N}$ there are integers d(k, t) and a(t) such that for every digraph D at least one of the following is true:

- (i) $\operatorname{dtw}(D) < d(k, t),$
- (ii) D contains \vec{K}_t as a butterfly-minor,
- (iii) there is a set $X \subseteq V(D)$ of order $|X| \leq a(t)$ and a cylindrical wall $W \subseteq D X$ of order k which is weakly flat in D X.

3.3 Finding Disjoint Dicycles of Distinct Lengths

In this section we give the proofs of Theorem 3.1 and Theorem 3.3.

The structure of the section follows the three possible cases given by Theorem 3.7:

- (i) digraphs of bounded directed tree-width,
- (ii) digraphs containing a large complete butterfly-minor,
- (iii) digraphs containing a weakly flat wall.

In the last subsection we then prove Theorem 3.1 and Theorem 3.3 by applying Theorem 3.7 and using the insights from the three previous subsections.

At this point it is worth pointing out why our proof strategy needs to make use of the Direced Flat Wall Theorem 3.7 to make progress on Lichiardopol's conjecture, and why its weakening, the Directed Grid Theorem 1.3, would not be as helpful. The main reason for this is that while walls of large order contain many vertex-disjoint directed cycles, they might not even contain two directed cycles whose lengths are different, as shown by the Remark 3.8 below. Hence, it seems unlikely that Theorem 1.3 will be helpful in obtaining disjoint directed cycles of distinct lengths in digraphs, since it gives no information concerning the relation between the wall contained in the digraph and the rest of the digraph. This disadvantage is improved by the Directed Flat Wall Theorem, which restricts the ways in which directed paths can intersect a weakly flat wall.

Remark 3.8. For every $k \in \mathbb{N}$ there exists a cylindrical wall of order k containing no two directed cycles of different lengths.

Proof. A subdivision of W_k is determined by the lengths of the subdivision-paths replacing its arcs. Hence, we may as well give an assignment $w : A(W_k) \to \mathbb{N}$ of positive integers to the arcs of W_k such that the sum of the labels on any directed cycle is the same. To do so, let R denote the set of arcs in W_k starting in $V(P_k^2)$ and ending in $V(P_1^1)$, and observe that every directed cycle in W_k contains *exactly* one arc in R. As a consequence, the digraph $W_k - R$ is acyclic. Let v_1, \ldots, v_{4k^2} be a topological ordering of this digraph, i.e., such that $(v_i, v_j) \notin A(W_k - R)$ whenever i > j. Let us define w(e) := j - i for every $e = (v_i, v_j) \in A(W_k) \setminus R$. Let $L := 4k^2$. By definition of the arc-weighting w, for every directed path $P = v_{i_1}, v_{i_2}, \ldots, v_{i_\ell}$ in $W_k - R$, it now holds that $\sum_{e \in A(P)} w(e) =$ $i_\ell - i_{\ell-1} + i_{\ell-1} - i_{\ell-2} \pm \ldots + i_2 - i_1 = i_\ell - i_1 < L$. In particular, the total weight of any directed path in $A(W_k)$ depends only on its endvertices and is smaller than L. We conclude that for every arc $e = (u, v) \in R$, there exists a number $L_{uv} \in \mathbb{N}, L_{uv} < L$ such that every directed path in $W_k - R$ starting in v and ending in u has total arc-weight L_{uv} . Let us now put $w(e) := L - L_{uv} \in \mathbb{N}$ for every $e = (u, v) \in R$.

Every directed cycle C in W_k intersects R in exactly one arc $e = (u, v) \in R$. Then C-e is a directed v-u-path in $W_k - R$, and we conclude $\sum_{e \in A(C)} w(e) = L - L_{uv} + L_{uv} = L$. This shows that w is an arc-weighting of W_k with positive integers in which all directed cycles have total arc-weight equal to a number $L \in \mathbb{N}$. This concludes the proof. \Box

3.3.1 Digraphs of Bounded Directed Tree-Width

Before giving the proof of Proposition 3.4, let us introduce the crucial definition of k-trains. These are digraphs containing many distinct cycle lengths, which are useful for embedding disjoint directed cycles of distinct lengths.

Definition 3.4. Let $k \in \mathbb{N}$. A k-train is a digraph consisting of a directed path P with vertex-trace u_0, u_1, \ldots, u_ℓ directed from u_0 to u_ℓ together with k arcs of the form $(u_\ell, u_{\ell_j}), j = 1, \ldots, k$, where $0 = \ell_1 < \ell_2 < \cdots < \ell_k < \ell$.

Note that every k-train contains k directed cycles of pairwise distinct lengths, namely the directed cycles through the arcs $(u_{\ell}, u_{\ell_j}), j = 1, \ldots, k$. This fact has the following useful consequences.

Observation 3.9. If a digraph D contains k vertex-disjoint k-trains as subdigraphs, then D contains k vertex-disjoint directed cycles of pairwise distinct lengths.

Proof. Let $D_1, \ldots, D_k \subseteq D$ be vertex-disjoint k-trains. For $i = 1, 2, 3, \ldots, k$ we successively pick a directed cycle C_i in D_i whose length is distinct from the lengths of the already chosen directed cycles C_1, \ldots, C_{i-1} , which is possible since D_i contains directed cycles of k different lengths. Eventually the process returns k vertex-disjoint directed cycles of pairwise distinct lengths. \Box

The next observation shows that k-trains exist in digraphs of large out-degree.

Observation 3.10. Let D be a digraph with $\delta^+(D) \ge k$. Then D contains a k-train.

Proof. Let P be the longest directed path in D. Then P together with the arcs leaving its end-vertex u contains a k-train, as all the out-neighbors of u are contained in V(P). \Box

Observation 3.9 motivates studying packings of disjoint k-trains in digraphs. Combined with Observation 3.10, one way to obtain many disjoint directed cycles with distinct lengths in digraphs is to find k vertex-disjoint subdigraphs with minimum out-degree at least k each. Recalling Proposition 2.19 from Chapter 2, we know that it is possible to find such subdigraphs if the minimum out-degree is large and the directed tree-width is bounded. Therefore we can easily obtain Proposition 3.4 as follows.

Proof of Proposition 3.4. Let D be any given digraph such that $\delta^+(D) > (d+2)(k-1)$ and $dtw(D) \leq d$. This means $\delta^+(D) \geq (d+2)(k-1) + 1 = (d+1)(k-1) + k$, and hence we may apply Proposition 2.19 to D. We find disjoint vertex-sets X_1, \ldots, X_k satisfying $\delta^+(D[X_i]) \geq k, i = 1, \ldots, k$. By Observation 3.10, every $D[X_i]$ contains a k-train. The claim now follows from Observation 3.9.

3.3.2 Digraphs Containing a Large Complete Minor

In this subsection, we show that every digraph containing as a butterfly-minor a complete digraph of sufficiently large order must also contain k disjoint directed cycles of distinct lengths. To achieve this goal, it is more convenient to prove an arc-weighted generalization of this statement as follows.

Lemma 3.11. Let $k \in \mathbb{N}$, $t = \frac{k^2 + 3k}{2}$. Let D be a digraph containing \vec{K}_t as a butterflyminor, and let $w : A(D) \to (0, \infty)$ be an arc-weighting of D. Then D contains kpairwise vertex-disjoint directed cycles C_1, \ldots, C_k such that the total weights $w(C_i) = \sum_{e \in A(C_i)} w(e), i = 1, \ldots, k$ of the cycles are pairwise distinct.

Proof. By induction on v(D)+a(D). We clearly have $v(D) \ge t$ and $a(D) \ge a(\vec{K}_t) = t(t-1)$ for every D as in the lemma, so in the base case we have $v(D) + a(D) = t^2$ and $D = \vec{K}_t$. Since $t = 2 + 3 + \cdots + k + (k + 1)$, we can partition V(D) into subsets V_1, \ldots, V_k , such that $|V_i| = i + 1$ for $1 \le i \le k$. We claim that $D[V_i]$ contains at least i directed cycles with pairwise distinct total weights, for every $i \in \{1, \ldots, k\}$. Indeed, for $1 \le i \le k$ let us pick some vertex $v_i \in V_i$ and order the vertices in $V_i \setminus \{v_i\}$ as $v_{i,1}, \ldots, v_{i,i}$ in such a way that

$$w((v_{i,1}, v_i)) \le w((v_{i,2}, v_i)) \le w((v_{i,3}, v_i)) \le \dots \le w((v_{i,i}, v_i)).$$

For $1 \leq j \leq i$, let us denote by $C_{i,j}$ the directed cycle in $D[V_i]$ induced by the arcs $(v_i, v_{i,1})$, $(v_{i,1}, v_{i,2}), \ldots, (v_{i,j-1}, v_{i,j}), (v_{i,j}, v_i)$. We then have

$$w(C_{i,j}) - w(C_{i,j-1}) = \underbrace{w((v_{i,j-1}, v_{i,j}))}_{>0} + \underbrace{(w((v_{i,j}, v_i)) - w((v_{i,j-1}, v_i)))}_{\ge 0} > 0$$

for $2 \leq j \leq i$, showing that $C_{i,1}, \ldots, C_{i,i}$ are *i* directed cycles in $D[V_i]$ of pairwise distinct total weights. For $i = 1, 2, \ldots, k$ we can now successively pick a cycle $C_i \in \{C_{i,j} | 1 \leq j \leq i\}$ whose weight is distinct from the weights of the already chosen cycles C_1, \ldots, C_{i-1} . This proves the claim in the base case of the induction.

Suppose now that $v(D) + a(D) > t^2$ and that the claim holds for all digraphs D' such that v(D') + a(D') < v(D) + a(D). Then either there is a proper subdigraph $D_1 \subsetneq D$ such that D_1 still contains \vec{K}_t as a butterfly-minor, or D contains a contractible arc $e = (u, v) \in A(D)$ such that $D_2 := D/e$ contains \vec{K}_t as a butterfly-minor. In the first case, we can apply the induction hypothesis to D_1 (with arc-weights inherited from D) to find that there exists k vertex-disjoint directed cycles in D_1 (and hence also in D) of pairwise distinct weights. In the second case, we know that either $d_D^+(u) = 1$ or $d_D^-(v) = 1$. Let us assume w.l.o.g. that e is the only arc leaving u, as the proof in the other case works symmetrically. By identifying the contraction vertex of e in D_2 with v, we may represent D_2 as follows:

$$V(D_2) = V(D) \setminus \{u\},$$

$$A(D_2) = A(D-u) \cup \{(x,v) | (x,u) \in A(D), x \neq v\}.$$

Let us define an arc-weighting $w_2 : A(D_2) \to (0, \infty)$ by $w_2(e) := w(e)$ for all $e \in A(D-u)$ and $w_2((x, v)) := w((x, u)) + w((u, v))$ for all $x \in N_D^-(u) \setminus (\{v\} \cup N_D^-(v))$. Since we have $v(D_2) + a(D_2) < v(D) + a(D)$, the induction hypothesis yields that D_2 contains k vertexdisjoint directed cycles C_1, \ldots, C_k with pairwise distinct total weights. If none of these cycles uses an arc of the form (x, v) with $x \in N_D^-(u) \setminus (\{v\} \cup N_D^-(v))$, then the cycles also exist in D with the same weights, and hence the inductive claim holds. Otherwise, exactly one of the cycles, say C_1 , uses an arc of the from (x, v) with $x \in N_D^-(u) \setminus (\{v\} \cup N_D^-(v))$. But then replacing the arc (x, v) on C_1 with the path (x, u), (u, v) in D, we find a directed cycle C'_1 in D such that $V(C'_1) = V(C_1) \cup \{u\}$, and such that the weight of C_1 in (D_2, w_2) equals the weight of C'_1 in (D, w). Hence, C'_1, C_2, \ldots, C_k are k vertex-disjoint directed cycles in D with pairwise distinct weights also in D, again proving the inductive claim. This concludes the proof by induction.

By putting w(e) := 1 for all $e \in A(D)$ in Lemma 3.11, we obtain the following.

Corollary 3.12. Let D be a digraph containing \vec{K}_t as a butterfly-minor, where $t \geq \frac{k^2+3k}{2}$. Then D contains k vertex-disjoint directed cycles of pairwise distinct lengths.

3.3.3 Digraphs Containing a Flat Wall

In this subsection we give conditions which guarantee many vertex-disjoint directed cycles of distinct lengths in digraphs of large minimum degree containing a weakly flat wall of large order. Let us start by introducing a useful notation: Let D be a digraph, and let $W \subseteq D$ be a cylindrical wall. For a vertex $w \in int(W)$, in the rest of this section we denote by $R_W^+[w]$ and $R_W^-[w]$ the sets of vertices in $V(D) \setminus int(W)$ which are reachable from w in $D - (int(W) \setminus \{w\})$, respectively which can reach w in $D - (int(W) \setminus \{w\})$.

Lemma 3.13. Let D be a digraph containing a cylindrical wall $W \subseteq D$. Let $a, b \in \mathbb{N}$ such that $\delta^+(D) \ge a + b$, and let $w \in int(W)$. Then at least one of the following holds.

- There exists $x \in V(D)$, a w-x-dipath P in D such that $V(P) \cap int(W) = \{w\}$ and distinct vertices $v_1, \ldots, v_a \in int(W) \setminus \{w\}$ such that $(x, v_i) \in A(D), i = 1, \ldots, a, or$
- $D[R_W^+[w]]$ has minimum out-degree at least b.

Proof. If $R_W^+[w] = \emptyset$, then putting x = w verifies the first item. Now suppose $R_W^+[w] \neq \emptyset$. By definition of $R_W^+[w]$, for every vertex $x \in R_W^+[w]$, we have $N_D^+(x) \setminus \operatorname{int}(W) \subseteq R_W^+[w]$. Now suppose that $D[R_W^+[w]]$ has minimum out-degree less than b. Then there exists $x \in R_W^+[w]$ such that

$$b > d^+_{D[R^+_W[w]]}(x) = |N^+_D(x) \setminus \operatorname{int}(W)|$$

= $d^+_D(x) - |N^+_D(x) \cap \operatorname{int}(W)| \ge a + b - |N^+_D(x) \cap \operatorname{int}(W)|.$

Consequently, $|N_D^+(x) \cap (\operatorname{int}(W) \setminus \{w\})| \ge |N_D^+(x) \cap \operatorname{int}(W)| - 1 \ge a$, and we find that x has a distinct out-neighbors $v_1, \ldots, v_a \in \operatorname{int}(W) \setminus \{w\}$. Since $x \in R_W^+[w]$, there exists a dipath P from w to x such that $V(P) \cap \operatorname{int}(W) = \{w\}$, and the claim follows. \Box

To state our next lemma we need the following definition.

Definition 3.5. Let W be a cylindrical wall. Let G_W denote the auxiliary undirected graph on the vertex-set int(W) in which two vertices $x \neq y \in int(W)$ are adjacent if there exists a brick B of W such that $x, y \in V(B)$. Then we define the brick-distance $dist_W(w_1, w_2)$ of two vertices $w_1, w_2 \in int(W)$ as their distance in the graph G_W .

Lemma 3.14. Let D be a digraph, and let $W \subseteq D$ be a cylindrical wall which is weakly flat in D. Let $w_1, w_2 \in int(W)$. Then

- (i) $R_W^+[w_1] \cap R_W^-[w_2] = \emptyset$, if $dist_W(w_1, w_2) \ge 2$, and
- (ii) $R_W^+[w_1] \cap R_W^+[w_2] = \emptyset$, if D is strongly connected and $dist_W(w_1, w_2) \ge 3$.

Proof.

- (i) For a contradiction suppose that $\operatorname{dist}_W(w_1, w_2) \geq 2$ and $R_W^+[w_1] \cap R_W^-[w_2] \neq \emptyset$. Pick some $x \in R_W^+[w_1] \cap R_W^-[w_2]$. Then by definition, there exists a w_1 -x-dipath P_1 and a x- w_2 -dipath P_2 with $V(P_1) \cap \operatorname{int}(W) = \{w_1\}, V(P_2) \cap \operatorname{int}(W) = \{w_2\}$. Let z be the first vertex of $V(P_2)$ we meet when traversing P_1 starting at w_1 . Then $P := P_1[w_1, z] \circ P_2[z, w_2]$ is a w_1 - w_2 -dipath in D satisfying $V(P) \cap \operatorname{int}(W) = \{w_1, w_2\}$. Since W is weakly flat in D, this means that there is a brick B of W such that $w_1, w_2 \in V(B)$, and hence $\operatorname{dist}_W(w_1, w_2) \leq 1$, a contradiction to our initial assumptions. This proves the claim.
- (ii) Suppose towards a contradiction that $R_W^+[w_1] \cap R_W^+[w_2] \neq \emptyset$, that D is strongly connected and that $\operatorname{dist}_W(w_1, w_2) \geq 3$. Pick a vertex $x \in R_W^+[w_1] \cap R_W^+[w_2]$.

By definition, there exists a w_1 -x-dipath P_1 and a w_2 -x-dipath P_2 in D such that $V(P_1) \cap \operatorname{int}(W) = \{w_1\}, V(P_2) \cap \operatorname{int}(W) = \{w_2\}$. Since D is strongly connected, there exists a directed path in D starting in x and ending in a vertex of $\operatorname{int}(W)$. Let P_3 be a shortest directed path with this property. Then $V(P_3) \cap \operatorname{int}(W) = \{y\}$, where y denotes the last vertex of P_3 . The triangle inequality yields that $3 \leq \operatorname{dist}_W(w_1, w_2) \leq \operatorname{dist}_W(w_1, y) + \operatorname{dist}_W(y, w_2)$, and hence we have $\operatorname{dist}_W(w_1, y) \geq 2$ or $\operatorname{dist}_W(w_2, y) \geq 2$. W.l.o.g. suppose that the former is true. Let z be the first vertex of $V(P_3)$ we meet when traversing P_1 starting at w_1 . Then $P := P_1[w_1, z] \circ P_3[z, y]$ is a w_1 -y-dipath in D satisfying $V(P) \cap \operatorname{int}(W) = \{w_1, y\}$. Since W is weakly flat in D, we conclude that there exists a brick B of W such that $w_1, y \in V(B)$, contradicting the fact that $\operatorname{dist}_W(w_1, y) \geq 2$. Hence our initial assumption was wrong, and the proof concludes.

The next lemma is the main technical result of this subsection, yielding conditions which guarantee a k-train intersecting a weakly flat wall only in a "thin strip".

Lemma 3.15. Let $k \in \mathbb{N}$, let D be a digraph with $\delta^+(D) \ge 7k - 5$, and let $W \subseteq D$ be a cylindrical wall which is weakly flat in D. Let w be a branch-vertex of D, whose coordinates in W are (c_1, c_2) for $c_1, c_2 \in \mathbb{N}$ such that $c_1 = 2\ell - 1$ is odd and c_2 is even.

Let $S \subseteq V(W)$ be defined as the set consisting of all branch vertices of W whose first coordinate is in $\{c_1 - 2, c_1 - 1, c_1, c_1 + 1, c_1 + 2, c_1 + 3\}$, together with all vertices contained in the interior of a subdivision-path of W spanned between two such branch vertices. If $S \cap per(W) = \emptyset$, then $D[S \cup R_W^+[w]]$ contains a k-train.

Proof. By definition of S and since $S \cap \mathbf{per}(W) = \emptyset$, we have $w \in S \subseteq \mathbf{int}(W)$. Since the definition of S depends only on the first coordinate of w, and by the cylindrical symmetry of W, we may assume w.l.o.g. $c_2 = 2$. Note that w is contained in the directed cycle Q_ℓ of W, and that $V(Q_{\ell-1}) \cup V(Q_\ell) \cup V(Q_{\ell+1}) \subseteq S$.

By Lemma 3.13 applied with a = 6k - 5 and b = k, either (1) there exists a vertex $x \in V(D) \setminus \operatorname{int}(W)$, a *w*-*x*-dipath *P* in *D* such that $V(P) \cap \operatorname{int}(W) = \{w\}$, and distinct out-neighbors $v_1, \ldots, v_{6k-5} \in \operatorname{int}(W) \setminus \{w\}$ of *x*, or (2) $\delta^+(D[R^+_W[w]]) \ge k$.

If (2) occurs, then we can apply Observation 3.10 to conclude that $D[R_W^+[w]]$ and therefore also $D[S \cup R_W^+[w]]$ contains a k-train as a subdigraph, which proves the assertion in this case. Hence, assume for the rest of the proof that (1) occurs.

Note that $V(P) \cap \operatorname{int}(W) = \{w\}$ implies that $x \in V(P) \subseteq \{w\} \cup R_W^+[w] \subseteq S \cup R_W^+[w]$. Further note that for every $i \in \{1, \ldots, 6k - 5\}$, the dipath $P \circ (x, v_i)$ intersects $\operatorname{int}(W)$

87



Figure 3.1: The local situation around the vertex w in the proof of Lemma 3.15, showing the bricks incident to w and relevant branch vertices within S. The non-filled central vertex represents w, while the green fat path represents part of the cycle Q_{ℓ} .

only in its first and last vertex, and hence the weak flatness of W in D implies that v_i is contained in one of the three bricks of W meeting at w.

For further reasoning, we fix the following notation (compare Figure 3.1): We denote by u_1, \ldots, u_{16} 16 distinct branch vertices of W contained in S, given by the following coordinates:

$$(c_1 - 1, 1), (c_1, 1), (c_1 + 1, 1), (c_1 - 2, 2), (c_1 - 1, 2), (c_1 + 1, 2), (c_1 + 2, 2),$$

 $(c_1-2,3), (c_1-1,3), (c_1,3), (c_1+1,3), (c_1+2,3), (c_1+3,3), (c_1+1,4), (c_1+2,4), (c_1+3,4).$

Let us label the bricks incident to w as B_1, B_2, B_3 , where B_1 contains the branch vertices $u_1, u_2, u_3, u_5, w, u_6, B_2$ contains $u_4, u_5, w, u_8, u_9, u_{10}$, and B_3 contains $w, u_6, u_7, u_{10}, u_{11}, u_{12}$.

In the following, for two distinct branch-vertices s, t of W which are connected by a directed subdivision-path, let us denote this path by W[s, t] = W[t, s].

Let $P_1, P_2, P_3, P_4, P_5, P_6$ be 6 directed paths in D[S] defined by

$$\begin{split} P_1 &= W[u_1, u_2] \circ W[u_2, u_3] \circ W[u_3, u_6] \circ W[u_6, w], P_2 = W[u_1, u_5], \\ P_3 &= W[w, u_5] \circ W[u_5, u_4] \circ W[u_4, u_8] \circ W[u_8, u_9] \circ W[u_9, u_{10}], \\ P_4 &= W[w, u_{10}] \circ W[u_{10}, u_{11}] \circ W[u_{11}, u_{12}], \\ P_5 &= W[u_7, u_6], P_6 = W[u_7, u_{12}]. \end{split}$$

Then we have $V(B_1) \cup V(B_2) \cup V(B_3) = V(P_1) \cup V(P_2) \cup V(P_3) \cup V(P_4) \cup V(P_5) \cup V(P_6)$. By the above, we have $\{v_1, \ldots, v_{6k-5}\} \subseteq (V(B_1) \cup V(B_2) \cup V(B_3)) \setminus \{w\}$, and hence there is $j \in \{1, \ldots, 6\}$ such that P_j contains at least k of the vertices v_1, \ldots, v_{6k} . Let $y_1, \ldots, y_k \in V(P_j) \setminus \{w\}$ be the ordering of these k vertices on the path P_j when traversing it starting from its first vertex. The rest of the proof shows how to find a k-train in $D[S \cup R_W^+[w]]$ for every value of j.

- Case j = 1. Let us consider the directed path $Q := P_1[y_1, w] \circ P$, which starts at y_1 and ends at x. Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q, showing that $D[S \cup R_W^+[w]]$ contains a k-train.
- Case j = 2. Consider the directed path $Q := P_2[y_1, u_5] \circ P_3[u_5, u_{10}] \circ Q_\ell[u_{10}, w] \circ P$, which starts at y_1 and ends at x (recall that $Q_\ell[u_{10}, w]$ denotes the directed subpath of the directed cycle Q_ℓ starting at u_{10} and ending at w). Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q. Hence, $D[S \cup R_W^+[w]]$ contains a k-train.

- Case j = 3. Consider the directed path $Q := P_3[y_1, u_{10}] \circ Q_\ell[u_{10}, w] \circ P$, which starts at y_1 and ends at x. Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q, showing that $D[S \cup R_W^+[w]]$ contains a k-train.
- Case j = 4. Consider the directed path

$$Q := P_4[y_1, u_{12}] \circ W[u_{12}, u_{13}] \circ W[u_{13}, u_{16}] \circ W[u_{16}, u_{15}] \circ W[u_{15}, u_{14}] \circ Q_\ell[u_{14}, w] \circ P,$$

which starts at y_1 and ends at x. Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q, showing that $D[S \cup R_W^+[w]]$ contains a k-train.

- Case j = 5. Consider the directed path $Q := P_5[y_1, u_6] \circ W[u_6, w] \circ P$, which starts at y_1 and ends at x. Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q, showing that $D[S \cup R_W^+[w]]$ contains a k-train.
- Case j = 6. Consider the directed path

$$Q := P_6[y_1, u_{12}] \circ W[u_{12}, u_{13}] \circ W[u_{13}, u_{16}] \circ W[u_{16}, u_{15}] \circ W[u_{15}, u_{14}] \circ Q_\ell[u_{14}, w] \circ P,$$

which starts at y_1 and ends at x. Then $Q \subseteq D[S \cup R_W^+[w]]$, and the last vertex x of Q has the k distinct out-neighbors y_1, \ldots, y_k on Q, showing that $D[S \cup R_W^+[w]]$ contains a k-train.

In every case, $D[S \cup R_W^+[w]]$ contains a k-train. This concludes the proof.

Before moving on, let us note the following symmetrical version of Lemma 3.15. In the following we refer to a digraph obtained from a k-train by reversing the orientations of all its arcs as a *reverse-k-train*.

Corollary 3.16. Let $k \in \mathbb{N}$, let D be a digraph with $\delta^{-}(D) \geq 7k - 5$, and let $W \subseteq D$ be a cylindrical wall which is weakly flat in D. Let w be a branch-vertex of D, whose coordinates in W are (c_1, c_2) for $c_1, c_2 \in \mathbb{N}$ such that $c_1 = 2\ell - 1$ and c_2 are odd.

Let $S \subseteq V(W)$ be defined as the set consisting of all branch vertices of W whose first coordinate is in $\{c_1 - 2, c_1 - 1, c_1, c_1 + 1, c_1 + 2, c_1 + 3\}$, together with all vertices contained in the interior of a subdivision-path of W spanned between two such branch vertices.

If $S \cap per(W) = \emptyset$, then $D[S \cup R_W^-[w]]$ contains a reverse-k-train.

Proof. Let m denote the order of the wall W. Let \overline{D} denote the digraph obtained from D by reversing its arcs. Then $\delta^+(\overline{D}) \geq 7k-5$ and \overline{D} contains \overline{W} , the digraph obtained from W by reversing the orientations, as a subdigraph. Then \overline{W} itself forms a wall with the following associated coordinates: Any branch-vertex of W at coordinates (t_1, t_2) receives the coordinates $(t_1, 2m + 1 - t_2)$ in \overline{W} . In particular w has cordinates $(c_1, 2m + 1 - c_2)$ in \overline{W} , where c_1 is odd and $2m + 1 - c_2$ is even. Since $\mathbf{per}(W) = \mathbf{per}(\overline{W})$, the weak flatness of W in D implies that also \overline{W} is weakly flat in \overline{D} . Also note that the set S remains invariant when changing between W and \overline{W} . We may thus apply Lemma 3.15 to \overline{D} and the vertex w in $\overline{W} \subseteq \overline{D}$ to find that $\overline{D}[S \cup R_{\overline{W}}^+[w]] = \overline{D}[S \cup R_W^-[w]]$ contains a k-train. It directly follows that $D[S \cup R_W^-[w]]$ contains a reverse-k-train, and the claim follows. \Box

The following consequences of Lemma 3.15 are used in the proofs of Theorem 3.1 and Theorem 3.3.

Lemma 3.17. Let D be a strongly connected digraph, and let $W \subseteq D$ be a cylindrical wall of order 3k + 2 which is weakly flat in D. If $\delta^+(D) \ge 7k - 5$, then D contains k vertex-disjoint directed cycles of distinct lengths.

Proof. Let us pick distinct branch vertices $w_1, \ldots, w_k \in \operatorname{int}(W)$ of W given by the following coordinates (compare Figure 3.2): For every $1 \leq i \leq k$, w_i is at position (6i - 1, 2) in W. We note that w_i is contained in the cycle Q_{3i} of W for every $1 \leq i \leq k$, and that $\operatorname{dist}_W(w_i, w_j) \geq 3$ for every $1 \leq i < j \leq k$. Using Lemma 3.14, (ii) we conclude that the sets $R_W^+[w_i]$, $i = 1, \ldots, k$ are pairwise disjoint. For every $1 \leq i \leq k$, let us define the *i*-th strip $S_i \subseteq \operatorname{int}(W)$ as the set of vertices $v \in V(W)$, such that either v is a branch vertex of W whose first coordinate is in $\{6i - 3, 6i - 2, \ldots, 6i + 2\}$, or such that v is contained in a subdivision-path of W between two branch vertices in S_i . Note that the sets S_1, \ldots, S_k are pairwise vertex-disjoint. For every $i = 1, \ldots, k$ we can now apply Lemma 3.15 to D and the vertex w_i in W, and we find that $D[S_i \cup R_W^+[w_i]]$ contains a k-train. Since the sets $S_i \cup R_W^+[w_i], i = 1, \ldots, k$ are pairwise disjoint, it follows that D contains k vertex-disjoint k-trains. The assertion now follows by applying Observation 3.9 to D.



Figure 3.2: Illustration of the placement of the vertices w_1, \ldots, w_4 in the wall W_{14} in the proof of Lemma 3.17. The vertices w_i are marked by big non-filled vertices, and the three incident bricks of every w_i are shaded. The directed cycles $Q_{3i}, i = 1, \ldots, 4$ containing the vertices w_i are marked fat and green.

Lemma 3.18. Let D be a digraph, and let $W \subseteq D$ be a cylindrical wall of order 8 which is weakly flat in D. Suppose further that $\delta^+(D), \delta^-(D) \ge 16$. Then D contains 3 vertexdisjoint directed cycles of distinct lengths.

Proof. Let us place two distinct vertices w_1, w_2 in the interior of W as follows: w_1 is the branch-vertex of W with coordinates (5, 2), and w_2 is the branch vertex of W with coordinates (11, 3). Then we have $\operatorname{dist}_W(w_1, w_2) = 3 \ge 2$, and by Lemma 2.21, (i) we have $R_W^+[w_1] \cap R_W^-[w_2] = \emptyset$. Let the strips $S_1, S_2 \subseteq \operatorname{int}(W)$ be defined as the sets consisting of the branch-vertices of W with first coordinates in $\{3, 4, \ldots, 8\}$ respectively $\{9, 10, \ldots, 14\}$ as well as the vertices on the subdivision-paths of W spanned between those vertices. It now follows from Lemma 3.15 applied to w_1 with k = 3, respectively from Corollary 3.16 applied to w_2 with k = 3, that $D[S_1 \cup R_W^+[w_1]]$ contains a 3-train and that $D[S_2 \cup R_W^-[w_2]]$ contains a reverse-3-train. Noting that $V(Q_1), S_1 \cup R_W^+[w_1]$ and $S_2 \cup R_W^-[w_2]$ are pairwise disjoint, we conclude that D contains the directed cycle Q_1 , a 3-train, and a reverse 3-train contains 3 directed cycles of distinct lengths, it follows that D contains 3 vertex-disjoint directed cycles of distinct lengths. □

3.3.4 Proofs of Theorem 3.1 and Theorem 3.3

We are now ready to give the proofs of Theorem 3.1 and Theorem 3.3.

Proof of Theorem 3.1. Let $k \in \mathbb{N}$ and put $s(k) := \max\{d(3k+2,t), a(t)+7k-5\}$, where $t := \frac{k^2+3k}{2}$ and $d(\cdot, \cdot), a(\cdot)$ are the functions from Theorem 3.7.

Let D be any given strongly s(k)-connected digraph. By Theorem 3.7, applied to D with parameters 3k + 2 and t, we find that one of the following must be true.

- (i) dtw(D) < d(3k+2,t), or
- (ii) D contains \vec{K}_t as a butterfly-minor, or
- (iii) there exists $X \subseteq V(D)$ with $|X| \leq a(t)$ and a cylindrical wall $W \subseteq D X$ of order 3k + 2 which is weakly flat in D X.

Case (i) is impossible, since we have $dtw(D) \ge s(k) \ge d(3k+2,t)$ by Corollary 3.6. If Case (ii) occurs, then we can apply Corollary 3.12 to D and we find that D contains kvertex-disjoint directed cycles of pairwise distinct length, as required. Finally, if Case (iii) occurs, we know from $s(k) \ge a(t) + 7k - 5$ that D - X is strongly (7k - 5)-connected. This in particular implies that D - X is strongly connected and $\delta^+(D) \ge 7k - 5$. Since there exists a cylindrical wall $W \subseteq D - X$ of order 3k + 2 which is weakly flat in D - X, we can apply Lemma 3.17 to D - X and find that there exist k vertex-disjoint directed cycles of pairwise distinct lengths in $D - X \subseteq D$, yielding the assertion also in this case. This concludes the proof.

Proof of Theorem 3.3. Let $K := \max\{2d(8,9) + 3, a(9) + 16\}$, where $d(\cdot, \cdot), a(\cdot)$ are the functions from Theorem 3.7.

Let D be any given digraph such that $\delta^+(D), \delta^-(D) \ge K$. By applying Theorem 3.7 to D with parameters 8 and 9, we find that one of the following must hold:

- (i) dtw(D) < d(8,9), or
- (ii) D contains \vec{K}_9 as a butterfly-minor, or
- (iii) There exists $X \subseteq V(D)$ with $|X| \le a(9)$ such that D X contains a wall W of order 8 which is weakly flat in D X.

If Case (i) occurs, then we have $dtw(D) \leq d(8,9) - 1$. Since $\delta^+(D) \geq K \geq 2d(8,9) + 3 = (d(8,9) + 1)(3 - 1) + 1 > (d(8,9) + 1)(3 - 1)$, we can apply Proposition 3.4 to D with parameters d := d(8,9) - 1 and k = 3 and find that indeed, D contains 3 vertex-disjoint directed cycles of distinct lengths in this case, proving the assertion in this case. If Case (ii) occurs, then we can apply Corollary 3.12 to D with t = 9, k = 3 and find that D contains 3 vertex-disjoint directed cycles of distinct lengths. Finally, if (iii) occurs, then we have $\delta^+(D-X) \geq \delta^+(D) - |X| \geq a(9) + 16 - a(9) = 16$ and $\delta^-(D-X) \geq \delta^-(D) - |X| \geq a(9) + 16 - a(9) = 16$. Since D - X contains a weakly flat wall W of order 8, it follows from Lemma 3.18 that $D - X \subseteq D$ contains 3 vertex-disjoint directed cycles the proof.

3.4 Equicardinal Disjoint Directed Cycles

In this section we give the proof of Proposition 3.2, by constructing for every $k \in \mathbb{N}$ a strongly k-connected digraph containing no two arc-disjoint directed cycles of equal length.

Proof of Proposition 3.2. Let $k \in \mathbb{N}$. Let $N := 4^{k^2}$ and let $a(1), \ldots, a(k^2), b(1), \ldots, b(k^2)$ be defined as $a(\ell) := N + 2^{\ell-1}, b(\ell) := N + 2^{k^2+\ell-1}, 1 \le \ell \le k^2$. Note that all possible subset-sums of $\{a(1), a(2), \ldots, a(k^2), b(1), b(2), \ldots, b(k^2)\}$ are pairwise distinct.

Let D_k denote a digraph on 2Nk vertices whose vertex-set is partitioned into 2N disjoint parts V_1, \ldots, V_{2N} , each of size exactly k, and having the following arcs:

For every $2 \leq \ell \leq N$, D_k contains all possible arcs of the form $(u, v), u \in V_\ell, v \in V_{\ell-1}$. Label the vertices in V_1 and V_N as $V_1 = \{u_1, \ldots, u_k\}, V_N = \{w_1, \ldots, w_k\}$. For every $i \in \{1, \ldots, k\}, u_i$ has k distinct out-arcs $e_{i,j}, j = 1, \ldots, k$, where $e_{i,j}$ connects u_i to a vertex chosen arbitrarily from $V_{a(k(i-1)+j)}$. Similarly, v_i has k distinct in-arcs $f_{i,j}, j = 1, \ldots, k$, where $f_{i,j}$ connects a vertex chosen arbitrarily from $V_{2N-b(k(i-1)+j)+1}$ to v_i .

This concludes the description of D_k . Please note that all out-neighbors of the vertices in V_1 are contained in $V_{N+1} \cup \cdots \cup V_{2N}$ while all the in-neighbors of the vertices in V_{2N} are contained in $V_1 \cup \cdots \cup V_N$. We next show that D_k is strongly k-vertex-connected. To this end, let $S \subseteq V(D_k)$ with |S| < k be arbitrary and let us prove that $D_k - S$ is strongly connected. Since |S| < k, we have $V_\ell \setminus S \neq \emptyset, 1 \leq \ell \leq 2N$, and hence in $D_k - S$ every vertex in $V_\ell \setminus S$ can reach every vertex in $V_{\ell'} \setminus S$ provided $1 \leq \ell' < \ell \leq 2N$. For strong connectivity of $D_k - S$ it now is sufficient to check that for every $i, j \in \{1, \ldots, k\}$ such that $u_i \in V_1 \setminus S, w_j \in V_{2N} \setminus S$, there exists a directed path P in $D_k - S$ starting at u_i and ending in w_j . Since u_i and w_j have out-degree respectively in-degree k in D_k , there exist vertices $x, y \in V(D_k) \setminus S$ such that $(u_i, x), (y, w_j) \in A(D_k)$. Note that by construction, $x \in V_\ell \setminus S, y \in V_{\ell'} \setminus S$ for some $\ell \in \{N + 1, \ldots, 2N\}, \ell' \in \{1, \ldots, N\}$. Hence, from the above it follows that x can reach y in $D_k - S$, and hence also u_i can reach w_j in $D_k - S$. This shows that $D_k - S$ is indeed strongly connected for every $S \subseteq V(D_k), |S| < k$, and hence that D_k is strongly k-vertex-connected.

Let us now show that D does not contain two arc-disjoint directed cycles of equal lengths. To this end, let us call $E := \{e_{i,j}, f_{i,j} | 1 \leq i, j \leq k\}$ the set of forward arcs of D_k and define the lengths of the forward arcs as $L(e_{i,j}) := a(k(i-1)+j), L(f_{i,j}) :=$ $b(k(i-1)+j), 1 \leq i, j \leq k$. Let now C be a directed cycle in D_k using the forward arcs $e_1, \ldots, e_r \in E$ in this cyclical order, and for $1 \leq s \leq r$ let $1 \leq t(s) < h(s) \leq 2N$ be such that e_s starts in $V_{t(s)}$ and ends in $V_{h(s)}$. Since all non-forward arcs on C connect a vertex of V_ℓ to $V_{\ell-1}$ for some $2 \leq \ell \leq 2N$, $C - \{e_1, \ldots, e_r\}$ decomposes into directed paths of lengths $h(1) - t(2), h(2) - t(3), h(3) - t(4), \ldots, h(r-1) - t(r), h(r) - t(1)$. It follows that

$$|C| = r + (h(1) - t(2)) + \ldots + (h(r-1) - t(r)) + (h(r) - t(1)) = \sum_{s=1}^{r} (h(s) - t(s) + 1).$$

Noting that $h(s) - t(s) + 1 = L(e_s)$ for every $1 \le s \le k$, we obtain that the length of any directed cycle in D_k equals the sum of the lengths of the forward arcs it is using. The multi-set of the edge lengths of forward-arcs equals $\{a(1), \ldots, a(k^2), b(1), \ldots, b(k^2)\}$. Since the subset-sums of this set are pairwise distinct, two directed cycles in D_k have the same length iff they use the same forward-arcs. This concludes the proof.

3.5 Conclusion

In this chapter, we have studied the existence of vertex-disjoint directed cycles with length constraints in digraphs of large connectivity. We have found that while we are guaranteed to find many disjoint cycles of different lengths in such digraphs, we cannot even be sure to find two arc-disjoint directed cycles of equal lengths. It would be interesting to complete this picture by understanding whether further length constraints, such as parities, can be enforced on the disjoint cycles. Clearly, bipartite digraphs of large connectivity show that we cannot expect the containment of odd dicycles. As mentioned in the introduction, a result by Thomassen [Tho85b] states that the existence of an even directed cycle cannot be forced by large minimum out-degree and in-degree. In contrast, he showed that every strongly 3-connected digraph contains an even directed cycle [Tho92]. We propose the following strengthening of Thomassen's result.

Conjecture 3.3. For every $k \in \mathbb{N}$, there exists an integer $s(k) \in \mathbb{N}$ such that every strongly s(k)-connected digraph contains k vertex-disjoint directed cycles of even lengths.

94 CHAPTER 3. DISJOINT DIRECTED CYCLES WITH LENGTH CONSTRAINTS

Chapter 4

Cycle-Reversion Distance between Graph Orientations

4.1 Introduction

Graphs and digraphs in this chapter are loopless but may have multiple parallel and anti-parallel arcs.

The term *flip* is commonly used in combinatorics to refer to an elementary, local, reversible operation that transforms one combinatorial object into another. Such flip operations naturally yield a *flip graph*, whose vertices are the considered combinatorial objects, and two of them are adjacent if they differ by a single flip. A classical example is the flip graph of triangulations of a convex polygon [STT88, Pou14]; see Figure 4.1. The vertex set of this graph are all triangulations of the polygon, and two triangulations are adjacent if one can be obtained from the other by replacing the diagonal of a quadrilateral formed by two triangles by the other diagonal. Similar flip graphs have also been investigated for triangulations of general point sets in the plane [DLRS10], triangulations of topological surfaces [Neg94], and planar graphs [BH09, BV11]. The flip distance between two combinatorial objects is the minimum number of flips needed to transform one into the other. It is known that computing the flip distance between two triangulations of a simple polygon [AMP15] or of a point set [LP15] is NP-hard. The latter is known to be fixed-parameter tractable [KSX17]. On the other hand, the NP-hardness of computing the flip distance between two triangulations of a convex polygon is a well-known open question [LZ98, Rog99, DP02, CSJ09, CSJ10, LEP10]. Flip graphs involving other geometric configurations have also been studied, such as flip graphs of non-crossing perfect matchings of a point set in the plane, where flips are with respect to alternating 4-cycles [HY02], or alternating cycles of arbitrary length [HHNRC05]. Other flip graphs include the flip graph on plane spanning trees [AAHV07], the flip graph of non-crossing partitions of a point set or dissections of a polygon [HHNOP09], the mutation graph of simple pseudoline arrangements [Rin57], the Eulerian tour graph of an Eulerian graph [ZG87], and many others. There is also a vast collection of interesting flip graphs for non-geometric objects, such as bitstrings, permutations, combinations, and partitions [FKMS18].

In essence, a flip graph provides the considered family of combinatorial objects with an underlying structure that reveals interesting properties about the objects. It can also be a useful tool for proving that a property holds for all objects, by proving that one particularly nice object has the property, and that the property is preserved under flips. Flip graphs are also an essential tool for solving fundamental algorithmic tasks such as random and exhaustive generation, see e.g. [AF96], [PW98] and [BKV19].



Figure 4.1: The flip graph of triangulations of a convex polygon.

The focus of the present chapter is on flip graphs for orientations of graphs satisfying some constraints. First, we consider so-called α -orientations, in which the outdegree of every vertex is specified by a function α , and the flip operation consists of reversing the orientation of all arcs in a directed cycle. We study the complexity of computing the flip distance between two such orientations. An interesting special case of α -orientations corresponds to perfect matchings in bipartite graphs, where flips involve alternating cycles. We also consider the dual notion of *c*-orientations, in which the number of forward arcs along each cycle is specified by a function *c*. Here a flip consists of reversing all arcs in a directed cut. We also analyze the computational complexity of the flip distance problem in *c*-orientations.

There are several deep connections between flip graphs and polytopes. Specifically, many interesting flip graphs arise as the (1-)skeleton of a polytope. For instance, flip graphs of triangulations of a convex polygon are skeletons of associahedra [CSZ15], and flip graphs of regular triangulations of a point set in the plane are skeletons of secondary polytopes (see [DLRS10, Chapter 5]). Associahedra are generalized by *quotientopes* [PS19], whose skeletons yield flip graphs on rectangulations [CSS18], bitstrings, permutations, and other combinatorial objects. Moreover, flip graphs of acyclic orientations or strongly connected orientations of a graph are skeletons of graphical and co-graphical zonotopes, respectively (see [Pos09, Section 2]). Similarly, as we show below, flip graphs on α orientations are skeletons of matroid intersection polytopes. We also consider vertex flips in *c*-orientations, inducing flip graphs that are distributive lattices and in particular subgraphs of skeletons of certain distributive polytopes. These polytopes specialize to flip polytopes of planar α -orientations, are generalized by the polytope of tensions of a digraph, and form part of the family of *alcoved polytopes* (see [FK11]).

In the next section, we give the precise statements of the computational problems we consider, connections with previous work, and the statements of our main results. In Section 4.3, we give the proof of our first main result, showing that computing the flip distance between α -orientations and between perfect matchings is NP-hard even for planar graphs. Section 4.4 presents the proof of our second main result, where we give a polynomial time algorithm to compute the vertex flip distance between *c*-orientations. Finally, in Section 4.5 we show that computing the distance between *c*-orientations, when double vertex flips are also allowed, is NP-hard.



Figure 4.2: Two α -orientations of a graph and a flip between them, where the values of α are depicted on the vertices.

4.2 Problems and Main Results

Flip distance between α -orientations. Given a graph G and some $\alpha : V(G) \to \mathbb{N}_0$, an α -orientation of G is an orientation of the edges of G in which every vertex v has outdegree $\alpha(v)$. An example for a graph and two α -orientations for this graph is given in Figure 4.2. A flip of a directed cycle C in some α -orientation X consists of the reversal of the orientation of all arcs of C, as shown in the figure. Arcs with distinct orientations in two given α -orientations X and Y induce an Eulerian subdigraph of both X and Y. They can therefore be represented by an edge-disjoint union of cycles in G which are directed in both X and Y. Hence the reversal of each such cycle in X gives rise to a flip sequence transforming X into Y and vice versa. We may thus define the *flip distance* between two α orientations X and Y to be the minimum number of cycles in a flip sequence transforming X into Y. We are interested in the computational complexity of determining the flip distance between two given α -orientations.

Problem 4.1. Given a graph G, some $\alpha : V(G) \to \mathbb{N}_0$, a pair X, Y of α -orientations of G and an integer $k \ge 0$, decide whether the flip distance between X and Y is at most k.

The crucial difficulty of this problem is that a shortest flip sequence transforming X into Y may flip arcs that are oriented the same in X and Y an even number of times, to reach Y with fewer flips compared to only flipping arcs that are oriented differently in X and Y; see the example in Figure 4.3. This motivates the following variant of Problem 4.1:

Problem 4.2. Given G, α, X, Y, k as in Problem 4.1, decide whether the flip distance between X and Y is at most k, where we may only flip arcs that are oriented differently in X and Y.

From α -orientations to perfect matchings. The flexibility in choosing a function α for a set of α -orientations on a graph allows us to capture numerous relevant combinatorial structures, some of which are listed below:

- domino and lozenge tilings of a plane region [Rém04, Thu90],
- planar spanning trees [GL86],
- (planar) bipartite perfect matchings [LZ03],
- (planar) bipartite *d*-factors [Pro02, Fel04],
- Schnyder woods of a planar triangulation [Bre00],
- Eulerian orientations of a (planar) graph [Fel04],
- k-fractional orientations of a planar graph with specified outdegrees [BF12],
- contact representations of planar graphs with homothetic triangles, rectangles, and k-gons [Fel13, GLP12, FSS18b, FSS18a].



Figure 4.3: An α -orientation X of a graph. The α -orientation Y obtained by flipping the four directed facial cycles C_1, \ldots, C_4 can be reached with fewer flips by flipping only the three directed facial cycles D_1, D_2, D_3 in this order.



Figure 4.4: An α -orientation of a bipartite graph and the corresponding perfect matching.

In the following, we focus on perfect matchings of bipartite graphs. Consider any bipartite graph G with bipartition (V_1, V_2) such that $|V_1| = |V_2|$ equipped with

$$\alpha: V(G) \to \mathbb{N}_0, \quad \alpha(x) := \begin{cases} 1 & \text{if } x \in V_1, \\ d_G(x) - 1 & \text{if } x \in V_2. \end{cases}$$

With this definition, in each α -orientation of G, the arcs directed from V_1 to V_2 form a perfect matching. This is illustrated in Figure 4.4. Conversely, given a perfect matching M of G, orienting all edges of M from V_1 to V_2 and all the other edges from V_2 to V_1 yields an α -orientation of the above type. Furthermore, the directed cycles in any α -orientation of G correspond to the alternating cycles in the associated perfect matching. Flipping an alternating cycle in a perfect matching corresponds to exchanging matching and nonmatching edges. An example of the flip graph of perfect matchings of a graph is given in Figure 4.5. In this special case, Problem 4.1 boils down to:

Problem 4.3. Given a bipartite graph G, a pair X, Y of perfect matchings in G and an integer $k \ge 0$, decide whether the flip distance between X and Y is at most k.

The example from Figure 4.3 can be easily modified to show that when transforming X into Y using the fewest number of flips, we may have to flip alternating cycles that are not in the symmetric difference of X and Y; see the example in Figure 4.6. If we restrict the flips to only use cycles in the symmetric difference of X and Y, then the problem of finding the flip distance becomes trivial, as the symmetric difference is a collection of vertex-disjoint cycles, and each of them has to be flipped, so Problem 4.2 is trivial for perfect matchings.



Figure 4.5: The flip graph of perfect matchings of a graph. The solid edges indicate flips along facial cycles, and the dashed edges indicate flips along non-facial cycles.



Figure 4.6: A perfect matching X in a graph. The perfect matching Y obtained by flipping the four alternating facial cycles C_1, \ldots, C_4 can be reached with fewer flips by flipping only the three alternating facial cycles D_1, D_2, D_3 in this order.

Flip graphs and matroid intersection polytopes. We proceed to give a geometric interpretation of the flip distance between α -orientations as the distance in the skeleton of a 0/1-polytope.

Recall that a matroid M can be described as (E, \mathcal{B}) , where E contains the elements of M and $\mathcal{B} \subseteq 2^E$ its bases. A common base of two matroids $M_1 = (E_1, \mathcal{B}_1), M_2 = (E_2, \mathcal{B}_2)$ is simply an element of $\mathcal{B}_1 \cap \mathcal{B}_2$.

It is well-known that perfect matchings in a bipartite graph $G = (V_1 \cup V_2, E)$ are the common bases of two partition matroids¹ (E, \mathcal{I}_1) and (E, \mathcal{I}_2) , in which a set of edges is independent if no two share an endpoint in V_1 , or, respectively, in V_2 .

Similarly, α -orientations can be defined as common bases of two partition matroids. In this case, every edge of the graph G is replaced by an anti-parallel pair of arcs, one arc for each possible orientation of the edge. One matroid ensures that in a basis, for every edge exactly one orientation is chosen. The second matroid encodes the constraint that in a basis, each vertex v has exactly $\alpha(v)$ outgoing arcs.

The common base polytope of two matroids is a 0/1-polytope obtained as the convex hull of the characteristic vectors of the common bases. Adjacency of two vertices of this polytope has been characterized by Frank and Tardos [FT88]. A shorter proof was given by Iwata [Iwa02]. We briefly recall their result in the next theorem. To state the theorem, consider a matroid $M = (E, \mathcal{I})$, a base $B \in \mathcal{I}$, and a subset $F \subseteq E$. The exchangeability graph G(B, F) of M is a bipartite graph with $B \setminus F$ and $F \setminus B$ as vertex bipartition, and edge set $\{ij \mid i \in B \setminus F, j \in F \setminus B, B \setminus \{i\} \cup \{j\}$ is a basis}. This definition and the theorem are illustrated in Figure 4.7 for the two partition matroids whose common bases are perfect matchings of a graph.

Theorem 4.1 ([FT88, Iwa02]). For two matroids $M^+ = (E, \mathcal{I}^+)$ and $M^- = (E, \mathcal{I}^-)$, two common bases $A, B \in \mathcal{I}^+ \cap \mathcal{I}^-$ are adjacent on the common base polytope if and only if all the following conditions hold:

- (i) the exchangeability graph G(A, B) of M^+ has a unique perfect matching P^+ ,
- (ii) the exchangeability graph G(B, A) of M^- has a unique perfect matching P^- ,
- (iii) $P^+ \cup P^-$ is a single cycle.

From this theorem we can conclude that the flip graphs we consider on perfect matchings and α -orientations are precisely the skeletons of corresponding polytopes of common bases of the two associated partition matroids.

It is interesting to compare Problems 4.1 and 4.3 with the analogous problems for other families of matroid polytopes. For instance, it is known that for two bases A, B of a matroid, the exchangeability graph G(A, B) has a perfect matching [Bru69]. Hence Acan be transformed into B by performing $|A\Delta B|/2$ exchanges of elements (where $A\Delta B$ is the symmetric difference of A and B), which is also the distance in the skeleton of the base polytope of the matroid. On the other hand, the problem of computing the flip distance between two triangulations of a convex polygon amounts to computing distances in skeletons of associahedra, which are known to be polymatroids (see [AA17] and references therein). This problem is neither known to be polynomial-time solvable nor known to be NP-hard. Also note that for other families of combinatorial polytopes, testing adjacency is already intractable. This is the case for instance for the polytope of the Traveling Salesman Problem (TSP) [Pap78], whose skeleton is known to have diameter at most 4 [RC98]. On the other hand, the corresponding polytope is known to be the common base polytope

¹In a partition matroid, we are given a partition S_1, \ldots, S_k of the ground-set E and numbers $c_1, \ldots, c_k \in \mathbb{N}_0$. A subset $I \subseteq E$ is independent in the corresponding partition matroid if $|I \cap S_i| \leq c_i$ for all $1 \leq i \leq k$.



Figure 4.7: Two common bases A and B (left and middle) of the matroids M^+ and M^- , where M^+ and M^- have as independent sets all subsets of edges of the graph where no two share an endpoint in the set of circled vertices, or the set of squared vertices, respectively. The right hand side shows the exchangeability graphs G(A, B) of M^+ (solid edges) and G(B, A) of M^- (dashed edges). As the conditions of Theorem 4.1 are met, the two bases are adjacent in the common base polytope, and adjacent in the flip graph shown in Figure 4.5.

of three matroids. Another important class of combinatorial polytopes are *alcoved poly*topes, see [LP07]. It is known that the flip graphs of planar α -orientations are skeletons of alcoved polytopes, see [FK11]. Thus, by our results below, flip distances in this class are also NP-hard to compute.

Hardness of flip distance between perfect matchings and α -orientations. We prove that Problem 4.3 is NP-complete, even for 2-connected bipartite subcubic planar graphs and k = 2. This clearly implies that Problem 4.1 is NP-complete as well.

Theorem 4.2. Given a 2-connected bipartite subcubic planar graph G and a pair X, Y of perfect matchings in G, deciding whether the flip distance between X and Y is at most two is NP-complete.

As direct consequences of the proof of Theorem 4.2 we get:

Corollary 4.3. Unless P = NP, deciding whether the flip distance between two perfect matchings is at most k is not fixed-parameter tractable with respect to parameter k.

Corollary 4.4. Unless P = NP, the flip distance between two perfect matchings is not approximable within a multiplicative factor $3/2 - \epsilon$ in polynomial time, for any $\epsilon > 0$.

We also prove that Problem 4.2 is NP-complete, even for 4-regular graphs and k = 2.

Theorem 4.5. Given a 4-regular graph G and a pair X, Y of α -orientations of G, deciding whether the flip distance between X and Y is at most two is NP-complete. Moreover, the problem remains NP-complete if we only allow flipping arcs that are oriented differently in X and Y.

The proofs of Theorem 4.2 and 4.5 are presented in Section 4.3.

From α -orientations in planar graphs to *c*-orientations. In what follows, we generalize the problem, via planar duality, to flip distances in so-called *c*-orientations.

Consider an arbitrary 2-connected plane graph G and its planar dual G^* . Then for any orientation D of the edges of G, the directed dual D^* of D is obtained by orienting any



Figure 4.8: Duality between flips in α -orientations (solid) and in *c*-orientations (dashed).

dual edge forward if it crosses a left-to-right arc in D in a simultaneous plane embedding of G and G^* , and backward otherwise; see Figure 4.8. Arc sets of directed cycles in D correspond to arc sets of *directed bonds* in D^* and vice-versa. Hence D is acyclic (respectively, strongly connected) if and only if D^* is strongly connected (respectively, acyclic). A *directed vertex cut* is a cut consisting of all arcs incident to a sink or a source vertex. Directed facial cycles in D are in bijection with the directed vertex cuts in D^* , and vice versa. The unbounded face in the plane embedding of D can be chosen such that it corresponds to a fixed vertex \top in D^* .

Let D be an α -orientation of G. For every vertex-set $U \subseteq V(D)$ we have

$$|\partial^+(U)| = \sum_{v \in U} d_D^+(v) - e(G[U]) = \sum_{v \in U} \alpha(v) - e(G[U]),$$

which only depends on α and G. Consequently, the set of orientations of G^* which are directed duals of α -orientations of G can be characterized by the property that for every cycle C in G^* , the number of arcs in clockwise direction is fixed by a certain value c(C)independent of the orientation. The flip operation between α -orientations of D consists of the reversal of a directed cycle. In the corresponding set of dual orientations of D^* , this translates to the reversal of the orientations of the arcs in a minimal directed cut, as shown on Figure 4.8.

The same notion has been investigated more generally without planarity conditions under the name of c-orientations by Propp [Pro02] and Knauer [Kna07]. Given a graph G, we can fix an arbitrary direction of traversal for each cycle C. Given a graph and an assignment $c(C) \in \mathbb{N}_0$ to each cycle in G, one may define a *c-orientation* of G to be an orientation having exactly c(C) arcs in forward direction for every cycle C in G. Note that it is sufficient to define the function c on a cycle basis of G, which consists of no more than e(G) cycles². The *flip* operation on the set \mathcal{R}_c of such *c*-orientations of a graph is defined as the reversal of all arcs in a minimal directed cut. It is not difficult to see that flips make the set of *c*-orientations of a graph connected (this will be noted in Section 4.4).

From the duality between planar α -orientations and planar *c*-orientations, determining flip distances between α -orientations of 2-connected planar graphs reduces to determining flip distances between the dual *c*-orientations. Note that planar duals of bipartite graphs are exactly the Eulerian planar graphs. Theorem 4.2 therefore directly yields:

Corollary 4.6. Given an Eulerian planar graph G and a pair X, Y of c-orientations of G, deciding whether the flip distance between X and Y is at most two is NP-complete.

²The existence of such a cycle basis can be easily seen by considering an ear-decomposition of G.

c-orientations and distributive lattices. A more local operation consists of flipping only directed vertex cuts, induced by sources and sinks, excluding a fixed vertex \top . We will refer to this special case as a *vertex flip*. Specifically, given a pair of *c*-orientations X, Y of a graph G with a fixed vertex \top , we aim to transform X into Y using only vertex flips at vertices distinct from \top .

A c-orientation X of G might contain a cycle C in G which is directed in X. According to the definition of a c-orientation, this means that C keeps the same orientation in every c-orientation of G. Consequently, any (minimal) directed cut in a c-orientation of G is disjoint from A(C). Contracting the cycle C in G, we end up with a smaller graph G' containing the same (minimal) directed cuts, such that the c-orientations of G are determined by their corresponding orientations on G'. We can therefore assume without loss of generality that the c-orientations that we consider are all acyclic. Similarly, G will be assumed to be connected.

Problem 4.4. Given a connected graph G, the fixed vertex \top and a pair X, Y of acyclic c-orientations, what is the length of a shortest vertex flip sequence from X to Y?

We should convince ourselves that under the assumptions made above, every pair of c-orientations is reachable from each other by vertex flips. This property is provided in a much stronger way by a distributive lattice structure on the set \mathcal{R}_c ; see Figure 4.9. The next theorem is a special case of Theorem 1 in Propp [Pro02] for acyclic c-orientations.

Theorem 4.7 ([Pro02, Kna07]). Let G be a graph with fixed vertex \top and \mathcal{R}_c a set of acyclic c-orientations of G. Then the partial order \leq_c on \mathcal{R}_c in which a c-orientation Y covers another c-orientation X if and only if Y can be obtained from X by flipping a source distinct from \top , defines a distributive lattice on \mathcal{R}_c .

Hence Problem 4.4 consists of finding shortest paths in the cover graph of a distributive lattice, where the size of the lattice can be exponential in the size of the input G.

Every distributive lattice is a lattice of *c***-orientations.** We next point out that every distributive lattice is isomorphic to the distributive lattice induced by a set of *c*-orientations of a graph. This relationship was described by Knauer [Kna08].

In order to represent a given distributive lattice L by an isomorphic lattice of corientations, we need to construct a corresponding digraph D(L). For this purpose, we
shortly recall a classical result from lattice theory, *Birkhoff's Theorem* (see [DP02]).

For any distributive lattice L, $\mathcal{J}(L)$ is the subposet of L induced by the set of *join-irreducible* elements, these are the elements of L covering exactly one element. On the other hand, given any poset P we may look at the distributive lattice $\mathcal{O}(P)$ formed by the downsets of P ordered by inclusion. Birkhoff's Theorem in our setting asserts that those two operations are inverse in the sense that $P \cong \mathcal{J}(\mathcal{O}(P))$ for any finite poset P and $\mathcal{O}(\mathcal{J}(L)) \cong L$ for any finite distributive lattice.

The idea is to define a digraph D(L) whose vertex set consists of the elements of $\mathcal{J}(L)$ with an additional vertex \top . The digraph is obtained from the natural upwardorientation of $\mathcal{J}(L)$ plus additional arcs from all the sinks and sources to \top . Let G(L) be the underlying graph of D(L), and for every cycle C of G(L), fix c(C) to be the number of forward-arcs on c in the orientation D(L). Let \mathcal{D}_c be the set of c-orientations of G(L). Fix \top as the unique non-flippable vertex.

We now define $L_c(D(L))$ as the distributive lattice induced on \mathcal{D} according to Theorem 4.7. An example of this construction is provided in Figure 4.10.

The next theorem (Theorem 4.8 below) is an easy consequence of Birkhoff's Theorem.



Figure 4.9: The distributive lattice induced by vertex flips in *c*-orientations. The reference orientation at the bottom is the directed dual D^* of the orientation D of the graph G used in Figures 4.4 and 4.5, where some parallel arcs incident with \top are grouped together for simplicity. The numbers depicted at the vertices are the number of times that each vertex is flipped with respect to the reference orientation.



Figure 4.10: A distributive lattice L represented by its Hasse diagram (left), the corresponding subposet of join-irreducible elements $\mathcal{J}(L)$ (middle), and the digraph D(L) associated with the lattice (right).

Theorem 4.8 ([Kna08]). Let L be any distributive lattice and D(L) be the corresponding digraph as defined above. Then $L \cong L_c(D(L))$.

Theorem 4.13 below gives a natural geometric embedding of the lattice L depending on the digraph D(L). This embedding is such that all values $z_X(x)$ are 0 or 1, and the vectors $z_X(.)$ are exactly the characteristic vectors of the downsets of $\mathcal{J}(L)$. The convex hull of those vectors is known as the order polytope of $\mathcal{J}(L)$ [Sta86], which is a particular case of the above-mentioned alcoved polytopes. The problem of computing vertex flip distances between elements of L encoded by c-orientations of G(L) therefore boils down to computing the distance between two downsets of $\mathcal{J}(L)$ in their inclusion lattice, which is a simple special case of Problem 4.4.

Facial flips in planar graphs. When we consider Problem 4.4 on planar graphs, restricting to vertex flips and considering the dual plane graph amounts to considering only flips of directed facial cycles, excluding the outer face whose dual vertex is \top . We refer to these as *facial flips*. Felsner [Fel04] considered distributive lattices induced by facial flips. The following computational problem is a special case of Problem 4.4.

Problem 4.5. Given a 2-connected plane graph G and a pair X, Y of strongly connected α -orientations, what is the length of a shortest facial flip sequence from X to Y?

Zhang, Qian, and Zhang [ZQZ19] recently provided a closed formula for this flip distance, which probably could be turned into a polynomial-time algorithm for computing a shortest face flip sequence. We prove the analogous stronger statement for Problem 4.4.

Theorem 4.9. There is an algorithm that, given a graph G with a fixed vertex \top and a pair X, Y of c-orientations of G, outputs a shortest vertex flip sequence between X and Y, and runs in time $O(m^3)$ where m is the number of edges of G.

In the planar case, this directly translates to a polynomial-time algorithm for Problem 4.5. The proof of Theorem 4.9 is presented in Section 4.4. In [FK09], the distributive lattice structure on c-orientations is generalized to so-called Δ -bonds, also known as tensions. We believe that our proof of Theorem 4.9 can be generalized to these objects.

Flip distance with larger cut sets. While computing the cut flip distance between *c*-orientations is an NP-hard problem in general (Theorem 4.2), there is a polynomial-time-algorithm for computing the distance when only using vertex flips (Theorem 4.9). It is natural to ask for a threshold between the hard and easy cases of flip distance problems. Our hardness reduction in Section 4.3 involves very long directed cycles, which correspond to flips of directed cuts in the dual *c*-orientations with cut sets of large size. Consequently, one may hope that the problem gets easier when restricting the sizes of the cut sets involved in a flip sequence. Our last result destroys this hope:

Theorem 4.10. Let X, Y be c-orientations of a connected graph G with fixed vertex \top . It is NP-hard to determine the length of a shortest cut flip sequence transforming X into Y, which consists only of minimal directed cuts with interiors of order at most two.

We will present the proof of Theorem 4.10 in Section 4.5.

4.3 Flip Distance between Perfect Matchings and between α -Orientations

The proof of Theorem 4.2 is by reduction from the following NP-complete problem.

Theorem 4.11 (Plesńik [Ple79]). Deciding directed Hamiltonicity of orientations of cubic planar graphs is NP-complete.

The above problem remains NP-complete if we additionally assume 2-connectivity of the cubic graph and that the orientation does not have sinks or sources (otherwise, there is no directed Hamiltonian cycle).

Proof of Theorem 4.2. As each flip sequence of length at most two can be used as a polynomially verifiable certificate, the problem is clearly in NP.

We now provide a reduction of the decision problem in Theorem 4.11 to Problem 4.3. So suppose we are given an orientation D of a 2-connected cubic planar graph without sinks and sources, and assume without loss of generality that $v(D) \ge 3$. Given D, we define an undirected graph G = G(D) as follows; see Figure 4.11: For each vertex $v \in V(D)$ we create a vertex x_v in G, and for each arc $e \in A(D)$ we create a pair of vertices x_e^+, x_e^- in G. The edges of G are defined as follows: For each arc $e \in A(D)$, we connect x_e^+ and x_e^- with an edge in G. Furthermore, we denote by V_1 and V_2 the vertices of D with outdegree 1 or 2, respectively. For each $v \in V_1$, if $e, f \in A(D)$ are the two incoming arcs at v and gis the outgoing arc, then we add the edges $x_e^+ x_v, x_f^+ x_v, x_e^+ x_g^-$, and $x_f^+ x_g^-$ to G. Similarly, for each $v \in V_2$, if $e, f \in A(D)$ are the two outgoing arcs at v and g is the incoming arc, then we add the edges $x_e^- x_v, x_f^- x_v, x_e^- x_g^+$, and $x_f^- x_g^+$ to G. We refer to the 4-cycles in Gformed by these edges as C_4 -gadgets. Note that G is subcubic, planar, 2-connected, and bipartite. Specifically, the bipartition is given by $\{x_v \mid v \in V_1\} \cup \{x_e^- \mid e \in A(D)\}$ and $\{x_v \mid v \in V_2\} \cup \{x_e^+ \mid e \in A(D)\}$.

We construct a pair of perfect matchings X, Y on G as follows. The first matching X is defined by fixing a particular perfect matching on each C_4 -gadget, and the second matching Y is obtained from X by flipping all cycles formed by the C_4 -gadgets; see Figure 4.11. We claim that X and Y have flip distance at most two in G if and only if D has a directed Hamiltonian cycle. From this the theorem follows.



Figure 4.11: C_4 -gadgets to construct the undirected graph G = G(D) (right) from the digraph D (left) in the proof of Theorem 4.2. The edges of the matchings X and Y in G are indicated by bold solid and dashed lines, respectively.

First, assume there is a directed Hamiltonian cycle H in D. We define a pair of cycles C_1, C_2 in G, where C_1 and C_2 both contain all the edges $\{x_e^+ x_e^- \mid e \in A(H)\}$, plus additional edges defined as follows. For each vertex $v \in V(D)$, consider the corresponding C_4 -gadget in G with the two incident edges corresponding to the arcs incident with v on H. The endpoints of those edges on the gadget divide it into two alternating paths in X, one with matching edges at both ends and one with non-matching edges at both ends. We add the edges on those two types of paths to C_1 or C_2 , respectively. Note that C_1 is an alternating cycle in X. Moreover, after flipping C_1 , the cycle C_2 is alternating, and flipping C_2 yields Y, as each edge in a C_4 -gadget gets flipped once while the remaining edges are flipped an even number of times and thus remain unchanged.

For the reverse implication, assume that X and Y are transformable into each other by flipping at most two alternating cycles. As the symmetric difference $X\Delta Y$ contains at least three disjoint cycles (recall the assumption $v(D) \geq 3$), exactly two cycles C_1, C_2 are flipped to transform X into Y, and neither C_1 nor C_2 is one of the 4-cycles formed by the C_4 -gadgets. As the edges outside the gadgets remain unchanged, they are covered by both C_1, C_2 or by neither of them. We claim that $H := \{e \in A(D) \mid x_e^+ x_e^- \in E(C_1)\}$ is the arc set of a directed Hamiltonian cycle in D. Since up to isomorphism, H is obtained from $E(C_1)$ by contraction of the C_4 -gadgets, H forms a cycle in D (here we need that D is cubic). If the cycle H is not a directed cycle in D, there would be some $v \in V_1$ with two incoming incident arcs from H. However, in this case, the path in the corresponding C_4 -gadget contained in C_1 consists of two edges, one of which is not in X, contradicting that C_1 is an alternating cycle in X. Finally, the directed cycle H has to be spanning. Indeed, if there is a C_4 -gadget not traversed by C_1 , then C_2 would be equal to the 4-cycle in the corresponding gadget, a contradiction. Proof of Theorem 4.5. We use the following hardness result of Peroche [Pér84]. Given a digraph D, where each vertex has indegree and outdegree equal to 2, it is NP-complete to decide if A(D) is the union of two directed Hamiltonian cycles. Given such a digraph D, let D be the digraph obtained from D by reversing the direction of every arc. We regard D and D as α -orientations X and Y of the same underlying graph G, where $\alpha(v) = 2$ for all $v \in V(G)$. The theorem follows by observing that the flip distance between X and Y is at most 2 if and only if A(D) is the disjoint union of two directed Hamiltonian cycles. Moreover, the same statement holds when we only allow flipping arcs that are oriented differently in X and Y.

4.4 Vertex Flip Distance between *c*-Orientations

In this section we prove Theorem 4.9.

Recall that, given a graph G with a fixed vertex \top , we only allow vertex flips at vertices distinct from \top . In the case that G is connected, we distinguish between two types of dicuts as follows: we say that a dicut S in an orientation of G is *positive* with respect to \top if and only if the uniquely determined cut set U of S does not contain \top . Otherwise the dicut is called *negative*. We also define the *interior* of S, denoted Int(S), as the cut set U of S if S is positive and as its complement \overline{U} if S is negative. That is, Int(S) is the set of vertices on the side of the cut opposite to \top .

The following lemma is needed to decompose the arc-sets of certain digraphs into dicuts with nested cut sets. Formally, for a digraph D and dicuts $S_1 = \delta(U_1)$ and $S_2 = \delta(U_2)$, the pair S_1, S_2 is called *laminar* if either $U_1 \subseteq U_2, U_2 \subseteq U_1, U_1 \cap U_2 = \emptyset$, or $U_1 \cup U_2 = V(D)$; see Figure 4.12. A family of dicuts in D is called *laminar* if all of its pairs are laminar. A *balanced digraph* is a digraph in which every oriented cycle has the same number of forward and backward arcs.

Lemma 4.12. Let D be a balanced digraph. Then A(D) can be decomposed into a laminar family of disjoint dibonds.

Proof. We prove the statement by induction on a(D). It is clearly true if $A(D) = \emptyset$. Assume for the induction step that $a(D) = k \ge 1$ and the statement holds for all digraphs with less than k arcs.

As D is balanced, it is obviously acyclic and therefore contains a source $s \in V(D)$. Each cycle in the underlying graph of D that contains s has exactly one forward and one backward arc incident to s. Therefore, the digraph D' obtained from D by contracting the set $E_D^+(s)$ of arcs incident to s is still balanced and has less than k arcs. By the induction hypothesis, there exists a laminar decomposition of $A(D') = A(D) \setminus E_D(s)$ into disjoint dibonds in D' and thus in D. Note that the dicuts in D' are exactly those of D disjoint from $E_D(s)$, and laminarity is preserved. Hence adding a decomposition of the directed vertex cut $E_D(s)$ into dibonds to the collection gives rise to a decomposition of A(D) into disjoint dibonds. This resulting decomposition is also laminar. To see this, let V_1, \ldots, V_l denote the vertex sets of the weak components of D - s (l = 1 is possible). The new minimal cuts added to the decomposition are induced by the cut sets $U_i := V(D) \setminus V_i$ for $i \in [l]$. We clearly have $U_i \cup U_j = V(D)$ for $i \neq j \in [l]$. Moreover, for any dibond S in the laminar decomposition of D', the cut set corresponding to S in D fully contains U_i or is disjoint from U_i for some $i \in [l]$ (otherwise S would not be a minimal cut). This finally implies that each pair of cuts in the new decomposition is laminar, as desired.

For every pair X, Y of c-orientations on a graph G, the difference $X \setminus Y$ denotes the set of arcs in X whose orientation is reversed in Y. Because X and Y are c-orientations, every
cycle in G has the same number of forward- and backward-arcs in $X \setminus Y$. The digraph trans(X, Y) obtained from X by contracting $\overline{X \setminus Y}$ therefore forms a balanced digraph. Consequently, Lemma 4.12 provides another proof that c-orientations can be reached from one another by flipping dibonds.

We now consider the partial order defined on acyclic *c*-orientations of an *n*-vertex graph such that the cover relation corresponds to flipping a source vertex. Recall that by Theorem 4.7, this partial order is a distributive lattice. We now reuse a result from Propp [Pro02] and Felsner and Knauer [FK09] that gives an embedding of this distributive lattice into \mathbb{N}_0^{n-1} , which led to the introduction of distributive polytopes by Felsner and Knauer [FK11]. This theorem is illustrated in Figure 4.9, where the values of the functions z_X are depicted in the vertices of the graph G.

Theorem 4.13 ([Pro02, FK09]). Let G be a graph on n vertices with a fixed vertex \top , X an acyclic c-orientation of G, and denote by X_{\min} the minimal element of the associated distributive lattice. Then the number of times $z_X(x)$ a vertex $x \in V(G) \setminus \{\top\}$ is flipped in an upward lattice path from X_{\min} to X is independent of the sequence. The resulting function $z_X : \mathcal{R}_c \to \mathbb{N}_0^{n-1}$ is a lattice embedding. That is, for every $x, y \in \mathbb{N}^{n-1}$ corresponding to c-orientations of G, the join and meet correspond to $\min(x, y)$ and $\max(x, y)$, respectively.

In other words, the distributive lattice on \mathcal{R}_c is isomorphic to an induced sublattice of the componentwise dominance order on \mathbb{N}_0^{n-1} . We call a vertex flip sequence *monotone* if every flipped vertex is either only flipped as a source or only as a sink. With this definition, Theorem 4.13 yields the following:

Corollary 4.14. Let G be a graph with fixed vertex \top and X,Y a pair of acyclic corientations on G. Then every monotone vertex flip sequence transforming X into Y has minimal length.

Consider two c-orientations X, Y of G. Our goal is to construct a monotone flip sequence from X to Y. By Lemma 4.12, there is a laminar decomposition of the arcs in trans(X, Y) into dibonds. The latter also form a laminar collection S = S(X, Y) of dibonds in X which partition $X \setminus Y$. Therefore reversing these dicuts yields Y.

We construct a poset P on S by the inclusion order of the interiors of the dibonds. That is, for $S, T \in S$, S is ordered before T in P if and only if $\operatorname{Int}(S) \subseteq \operatorname{Int}(T)$; see Figure 4.12. Since S is laminar, the cover graph of P is a forest, with the additional property that every non-maximal element S is covered by a unique other element, which we denote by $\operatorname{cov}(S)$. Moreover, for each vertex $x \in V(G)$ in the interior of at least one of the cuts in S, we let S_x be the (unique) minimal element of the poset P such that $x \in \operatorname{Int}(S_x)$. Also, for each $S \in S$ we denote by $\operatorname{Int}(S) := \{x \in V(G) \mid S_x = S\} \subseteq \operatorname{Int}(S)$ the set of vertices in the interior of S but in none of the interiors of the cuts which are covered by S in the poset.

For each dicut $S \in S$ we define an integer weight w(S) and a sign $sgn(S) \in \{+, 0, -\}$ as follows; see Figure 4.12. If $S \in S$ is a maximal element in P then we define w(S) := 1, and sgn(S) := + if S is positive and sgn(S) := - otherwise. For every sign $s \in \{+, 0, -\}$ and dicut $S \in S$, we say that S agrees with s, if either s = 0, or if s = + and S is positive, or s = - and S is negative. For every non-maximal $S \in S$, we inductively define

$$w(S) := \begin{cases} w(\operatorname{cov}(S)) + 1 & \text{if } S \text{ agrees with } \operatorname{sgn}(\operatorname{cov}(S)), \\ w(\operatorname{cov}(S)) - 1 & \text{otherwise,} \end{cases}$$



Figure 4.12: A laminar collection S of disjoint dibonds of $X \setminus Y$ (left), where the positive ones are dashed, and the negative ones are dotted, and the corresponding poset P of dicuts in S ordered by inclusion (right) with its associated signed weights $sgn(S) \cdot w(S), S \in S$.

and

$$\operatorname{sgn}(S) := \begin{cases} \operatorname{sgn}(\operatorname{cov}(S)) & \text{if } \operatorname{sgn}(\operatorname{cov}(S)) \neq 0 \text{ and } w(S) \neq 0, \\ + & \operatorname{if } \operatorname{sgn}(\operatorname{cov}(S)) = 0, w(S) \neq 0 \text{ and } S \text{ is positive,} \\ - & \operatorname{if } \operatorname{sgn}(\operatorname{cov}(S)) = 0, w(S) \neq 0 \text{ and } S \text{ is negative,} \\ 0 & \operatorname{if } w(S) = 0. \end{cases}$$

It follows from this definition that the weights are non-negative and that sgn(S) = 0if and only if w(S) = 0 for every $S \in S$. We will see that given a dibond S in S, the weight w(S) describes the number of times each vertex which lies in <u>Int</u>(S) will be flipped, whereas sgn(S) captures the direction in which (all) these vertices are flipped. That is, a positive sign means that vertices are flipped from sources to sinks, while a negative sign means that vertices are flipped from sinks to sources.

We will need the following auxiliary statement.

Lemma 4.15. Let X, Y be acyclic c-orientations of a connected graph G with fixed vertex \top . If $S = X \setminus Y$ is a positive dicut, then there is a vertex flip sequence transforming X into Y such that only vertices in Int(S) are flipped, each exactly once from source to sink.

The analogous statement for negative dicuts holds with sources and sinks exchanged.

Proof. We prove the statement by induction on |Int(S)|. If Int(S) is a single vertex, then S corresponds to the arcs incident to a source, and the statement holds.

For the induction step assume $|\operatorname{Int}(S)| = k \ge 2$ and that the claim holds for all positive cuts in c-orientations of G whose interiors have order less than k. Since X is acyclic, the induced subdigraph $X[\operatorname{Int}(S)]$ is also acyclic and thus contains a source $x \in \operatorname{Int}(S)$. Since $\operatorname{Int}(S)$ is the cut set of S, x is a source in X as well. Let Z be the c-orientation obtained from X by a vertex flip at x. It follows that the cut $\delta(\operatorname{Int}(S) \setminus \{x\})$ in Z is positive with interior of order k - 1. By induction, there is a vertex flip sequence from Z to Y such that only vertices in $\operatorname{Int}(S) \setminus \{x\}$ are flipped, each exactly once from source to sink. Starting with a vertex flip at x and continuing with this flip sequence yields a flip sequence from X via Z to Y with the desired properties. \Box

We are now in position to prove the main result of this section.

Theorem 4.16. Let X, Y be acyclic c-orientations of a connected graph G. There is a monotone vertex flip sequence transforming X into Y which can be computed in cubic time in terms of e(G).

Proof. Consider the following strengthening of the theorem:

Claim. Let X, Y be acyclic *c*-orientations of a connected graph G and S = S(X, Y) a laminar decomposition of $X \setminus Y$ into disjoint dibonds. Then there is a monotone vertex flip sequence from X to Y, such that every flipped vertex x is contained in the interior of the dicut $S_x \in S$, and x is flipped $w(S_x)$ times from source to sink if $sgn(S_x) = +$, and $w(S_x)$ times from sink to source otherwise.

We prove this claim by induction on the size of S. The statement is clearly true if |S| = 0 (which means that X = Y), settling the base case of the induction. Assume for the induction step that we are given a pair $X \neq Y$ of *c*-orientations and a laminar decomposition S of $X \setminus Y$ of size $k \geq 1$. Assume that the claim holds for all pairs of *c*-orientations with a laminar decomposition of size less than k.

In the poset P on S we consider a minimal element corresponding to a cut $S \in S$, i.e., we have $\underline{Int}(S) = Int(S)$ and all vertices $x \in Int(S)$ satisfy $S_x = S$. Lemma 4.15 gives a vertex flip sequence F_1 that flips only vertices in Int(S), each exactly once from source to sink if S is positive and from sink to source if S is negative. Applying this flip sequence to X, we obtain an intermediate c-orientation Z that differs from X only by the reversal of all arcs in S. Consequently, $S \setminus \{S\}$ is a laminar decomposition of $Z \setminus Y$ into dibonds in Z of size k - 1. By induction, we also have a vertex flip sequence F_2 transforming Zinto Y with the aforementioned properties.

Note that the weights and signs of all dicuts $T \in S \setminus \{S\}$ defined with respect to S or $S \setminus \{S\}$ are the same, so we may simply write w(T) and $\operatorname{sgn}(T)$. Furthermore, the set $\operatorname{Int}(T)$ defined with respect to S is a subset of the same set defined with respect to $S \setminus \{S\}$. To complete the induction step, we distinguish two cases.

The first case is that S is a maximal element in P, or that S agrees with $\operatorname{sgn}(\operatorname{cov}(S))$. In this case, we claim that the concatenation F of F_1 and F_2 is a flip sequence transforming X via Z into Y with the desired properties. It suffices to check this for the vertices in $\operatorname{Int}(S) = \operatorname{Int}(S)$, since for all other vertices, the claimed properties follow inductively (they are never flipped in F_1 , so their behavior in F will be the same as in F_2). If S is a maximal element in P, then w(S) = 1 and every vertex $x \in \operatorname{Int}(S)$ will be flipped exactly once. Moreover, according to Lemma 4.15, if S is positive, i.e., $\operatorname{sgn}(S) = +$, then x is flipped from source to sink, and if S is negative and $\operatorname{sgn}(S) = -$, then x is flipped sink to source. It remains to consider the subcase that S is not maximal, i.e., $\operatorname{cov}(S)$ exists. Consider any vertex $x \in \operatorname{Int}(S)$. During the flip sequence F, the vertex x is flipped once in F_1 and $w(\operatorname{cov}(S))$ times in F_2 , so $w(S) = w(\operatorname{cov}(S)) + 1$ times in total, as required. Moreover, the assumption that S agrees with $\operatorname{sgn}(\operatorname{cov}(S))$ means that either $\operatorname{sgn}(\operatorname{cov}(S)) = +$ and S is positive or $\operatorname{sgn}(\operatorname{cov}(S)) = -$ and S is negative, or $w(\operatorname{cov}(S)) = \operatorname{sgn}(\operatorname{cov}(S)) = 0$. We conclude from the inductive assumption that in those three cases, x is only flipped from source to sink in both F_1 and F_2 , or only sink to source in both, or only once in F_1 but not in F_2 , respectively. Consequently, x satisfies the inductive claim in all cases.

The second case is that S is not maximal, i.e., cov(S) exists, and that S does not agree with $\operatorname{sgn}(\operatorname{cov}(S))$. This means that $w(S) = w(\operatorname{cov}(S)) - 1$. Without loss of generality, assume that S is positive and consequently sgn(cov(S)) = - (the other case is symmetric). Consider again the vertex flip sequence F obtained by concatenating F_1 and F_2 . This flip sequence would transform X via Z into Y, however, we will not actually apply F, but modify the sequence as follows. By Lemma 4.15, F_1 flips each vertex in Int(S) exactly once from source to sink. By induction, in F_2 , each vertex in Int(S) contained in the interior of cov(S) (defined with respect to $S \setminus \{S\}$) is flipped from sink to source. Let x be the last element of F_1 and consider the subsequence x, x_1, \ldots, x_k, x of F starting with x and ending with the first occurrence of x in F_2 . None of the vertices x_1, \ldots, x_k is adjacent to x in G, because after the first vertex flip at x (from source to sink) all arcs incident with x are incoming, and in F_2 we only flip sinks to sources. This shows that deleting the first two occurrences of x from F preserves the number of and direction of all flips at vertices distinct from x, and still transforms X into Y. Repeated application of this argument produces a reduced vertex flip sequence F' transforming X into Y such that each vertex $x \in V(G) \setminus \text{Int}(S)$ is flipped the same number of times and in the same direction as in F_2 . By the inductive assumption, this means that x is flipped $w(S_x)$ times from source to sink if $sgn(S_x) = +$, and $w(S_x)$ times from sink to source if $sgn(S_x) = -$. On the other hand, every $x \in \text{Int}(S)$ is missing its first occurrence in F_2 but is flipped in the same way from sink to source for all remaining occurrences. This implies that x is flipped w(S) = w(cov(S)) - 1 times from sink to source, as it should. This proves that F' is a vertex flip sequence from X to Y satisfying the conditions in our claim. This completes the proof of our claim.

It remains to verify that the recursive algorithm obtained from this inductive argument runs in cubic time in m := e(G). First of all, the number of dicuts in any laminar decomposition, which corresponds to the number of induction steps, is bounded by the number of edges m. Consequently, it suffices to bound the number of operations needed in one induction step in terms of m. Specifically, we need to compute the cover relations of P, the weights and signs of the dicuts, find a minimal element of the poset P, test its properties for the case distinction and construct the resulting flip sequence by concatenation and possibly deletion of double occurrences, all of which can be done in time $O(m^2)$. This proves an upper bound of $O(m^3)$ for the total number of steps performed for the construction of the monotone flip sequence to transform X into Y. Finally, a laminar decomposition S of $X \setminus Y$ as guaranteed by Lemma 4.12 can be computed in time $O(m^3)$ as well, by following the recursive strategy explained in the proof of the lemma. This completes the proof.

Combining Corollary 4.14 and Theorem 4.16 yields Theorem 4.9.

4.5 Flip Distance with Larger Cut Sets

In this section we prove Theorem 4.10 by reduction from the following NP-hard problem.

Given a (finite) poset (P, \prec) , its *height* is the maximum size k of a so-called chain $x_1 \prec x_2 \prec \cdots \prec x_k$ of elements in P. A linear extension of P is a linear ordering (x_1, \ldots, x_n) of all elements of P such that $x_i \prec x_j$ implies that i < j. Given a linear extension $L = (x_1, \ldots, x_n)$ of P, a jump is a pair x_i, x_{i+1} in L for which $x_i \not\prec x_{i+1}$ in P. Conversely, a bump is a pair x_i, x_{i+1} such that $x_i \prec x_{i+1}$. The jump number s(P) of P is the minimum number of jumps among all linear extensions of P. The Jump Number Problem is the algorithmic problem of computing the jump number of a poset encoded by its comparabilities.

Theorem 4.17 ([Pul75, Mül90]). Determining the jump number of a poset of height two is an NP-hard problem.

Proof of Theorem 4.10. We provide a Turing-reduction of the Jump Number Problem for posets of height two to the problem stated in the theorem. For this purpose, assume we are given a poset (P, \prec) of height two with bipartite Hasse diagram $G = (P_1 \cup P_2, E)$ as an instance for the Jump Number Problem. We may assume that P has no isolated elements and that P_1 contains all minimal elements and P_2 all maximal elements of the poset. We construct an auxiliary graph G' from G by adding an additional unique maximal element \top , and connecting it with edges to all vertices of G. We construct two orientations X, Yof G' as follows: In both orientations all edges are oriented from P_1 to P_2 . Moreover, in Xall edges incident with \top are oriented towards \top , while in Y all these edges are oriented away from \top . As X, Y are obtained from each other by flipping all arcs incident with \top (this flip is not allowed, though, as \top is the fixed vertex), they are c-orientations with respect to the same c.

Let d denote the minimal flip distance between these c-orientations according to the conditions of the theorem. We will complete the proof by showing that s(P) = d - 1.

We first show that $s(P) \ge d-1$. For this argument, let $L = (x_1, \ldots, x_n)$ be an arbitrary linear extension of P. As P has height two, the elements $\{x_1, \ldots, x_n\}$ of P are partitioned into subsets B_1, \ldots, B_m of size one or two, such that for all B_i, B_j with i < j, the elements from B_i appear before the elements from B_j in L, and such that the two-element sets B_i contain exactly all bump pairs. We define a flip sequence that starts with the orientation X and consecutively flips the cuts induced by B_1, \ldots, B_m . Since for all $1 \le i \le m$, each of B_i and $\overline{B_i} := (P \cup \{\top\}) \setminus B_i$ induces a connected subgraph of G', these are indeed minimal cuts. Moreover, each of these cuts is flippable. This is obviously true for B_1 , as B_1 induces a dicut in X. Now assume inductively that the cuts induced by B_1, \ldots, B_{k-1} for some $k \ge 2$ have been flipped. As L is a linear extension of P, all elements in the downset of B_k in P but not in B_k are in one of the B_i with i < k. This implies that every arc between some $x \in B_k$ and $y \notin B_k$ is oriented from x to y in the current orientation, and thus B_i is indeed flippable.

In this flip sequence, every arc in X not incident to \top will be flipped zero or two times and thus maintains its original orientation, while all the arcs incident to \top get reversed, as they are incident to exactly one set B_i . Consequently, the flip sequence transforms X into Y, proving that $d \leq m$. As m equals the number of jumps in L plus 1 (every non-jump is a bump within one of the B_i), this yields $d - 1 \leq s(P)$.

We now show that $s(P) \leq d-1$. Assume that $B_1, \ldots, B_d \subseteq P$ are the cut sets of size one or two appearing (in this order) in a shortest flip sequence transforming X into Y. We may assume that among all shortest flip sequences, this sequence also minimizes $|B_1| + |B_2| + \ldots + |B_d|$. Since each vertex $x \in P$ has an outgoing arc to \top in X which

must be reversed during the flip sequence, x must be contained in at least one of the B_i . We claim that x is contained in at most one of the B_i . That is, the B_i are pairwise disjoint. Assume to the contrary that $x \in B_i \cap B_j$ for some i < j and that B_i, B_j is the only intersecting pair among $B_i, B_{i+1}, \ldots, B_j$ (by minimizing j - i). In particular, none of the cut sets B_{i+1}, \ldots, B_{j-1} contains x, and x is the only vertex flipped multiple times in this subsequence. We are then in one of the four cases $B_i = B_j = \{x\}$, or $B_i = \{x, y\}$ and $B_j = \{x\}$, or $B_i = \{x\}$ and $B_j = \{x, z\}$, or $B_i = \{x, y\}$ and $B_j = \{x, z\}$ for some elements $y, z \in P$ distinct from x. Since no vertex adjacent to x in G' can be flipped by B_{i+1}, \ldots, B_{j-1} , it follows that in each of these cases, the sequence

$$B_1, \dots, B_{i-1}, B_i \setminus \{x\}, B_{i+1}, \dots, B_{j-1}, B_j \setminus \{x\}, B_{j+1}, \dots, B_d$$

is a valid flip sequence from X to Y of length at most d and with decreased sum $|B_1| + |B_2| + \ldots + |B_d|$, a contradiction. This proves that the cut sets B_i are pairwise disjoint.

The B_i are flipped one after the other and by definition of X, the dicut induced by B_i is flippable if and only if all the elements in the downset of B_i with respect to P but not in B_i were flipped before. Therefore, by listing the elements in the sets B_1, \ldots, B_d in this relative order, and ordering the elements within each B_i according to their order in P, we obtain a linear extension L of P whose jumps are exactly those pairs having elements in two consecutive sets B_i . It follows that there are d-1 jumps in L, proving that $s(P) \leq d-1$.

Combining these arguments shows that s(P) = d - 1, and using Theorem 4.17 we obtain the claimed hardness result.

4.6 Conclusion

Recall that Problem 4.2 asks for a shortest flip sequence of directed cycles transforming one α -orientation X into another one Y, where we only allow flipping arcs that are oriented differently in X and Y. Since the set of arcs that are oriented differently in X and Y induce an Eulerian subdigraph in both X and Y, we have the following natural question:

Question 4.1. What is the smallest number of directed cycles into which an Eulerian digraph can be decomposed?

We have seen in Theorem 4.5 that from a computational point of view, this problem is hard for general digraphs, but we wonder what happens when adding planarity constraints. The aforementioned question can also be studied in terms of upper bounds as a function of the number of vertices, which is related to the famous Hajós conjecture on undirected Eulerian graphs, see [Lov68]. Another interesting variant of Question 4.1 is the following:

Question 4.2. Given a graph G with an Eulerian subgraph H, what is the smallest number of cycles of G such that their symmetric difference is H?

Concerning our proof of Theorem 4.10, we believe that for any bound on the size of the cuts, the corresponding flip distance will be NP-hard to compute. On the other hand, we use very particular graphs as gadgets, and we do not know the complexity of the corresponding problem for planar α -orientations. We think the following is an interesting special case:

Question 4.3. Let X, Y be perfect matchings of a planar bipartite 3-connected graph G. What is the complexity of determining the distance of X and Y with respect to alternating cycles that are either a face or the symmetric difference of two touching faces?

4.6. CONCLUSION

The feeling that this problem might be tractable is supported by the following observation. It is not difficult to show that every height two poset with bipartite *planar* Hasse diagram has dimension at most two. It then follows from [SS87] that the restriction of the Jump Number Problem to such posets is solvable in polynomial time, and thus, the hardness reduction presented in the previous section fails.

Chapter 5

Even Circuits in Regular Oriented Matroids

5.1 Introduction

Graphs and digraphs in this chapter are allowed to have multiple parallel arcs, anti-parallel arcs and loops. The reader is assumed to have basic familiarity with matroid theory.

Deciding whether a given digraph contains a directed cycle, briefly *dicycle*, of even length is a fundamental algorithmic problem for digraphs and often referred to as the *even dicycle problem*. The computational complexity of this problem was unknown for a long time and several interesting polynomial time equivalent problems have been found [KLM84, MS86, Tho86, McC04]. The question about the computational complexity was resolved by Robertson, Seymour, and Thomas [RST99] and independently by McCuaig [McC04] who stated polynomial time algorithms for one of the polynomially equivalent problems, and hence also for the even dicycle problem.

One of these polynomially equivalent problems makes use of the following definition.

Definition 5.1 ([ST87]). Let D be a digraph. We call D non-even, if there exists a set J of arcs in D such that every directed cycle C in D intersects J in an odd number of arcs. If such a set does not exist, we call D even.

The arc-set J can be seen as a special kind of feedback arc-set in D.

Seymour and Thomassen proved that the decision problem whether a given digraph is non-even, is polynomially equivalent to the even dicycle problem.

Theorem 5.1 ([ST87]). The problem of deciding whether a given digraph contains an even directed cycle, and the problem of deciding whether a given digraph is non-even, are polynomially equivalent.

Furthermore, the main result of Seymour and Thomassen [ST87] characterized being non-even in terms of minimal forbidden subdigraphs. Their result can however be stated much more compactly by formulating it in terms of forbidden butterfly minors instead of forbidden subgraphs. We can state the result of Seymour and Thomassen as follows.

Theorem 5.2 ([ST87]). A digraph D is non-even if and only if no butterfly minor of D is a \vec{C}_k for some odd k.

The main purpose of the results presented in this chapter is to lift the even dicycle problem to oriented matroids, and to extend Theorem 5.1 and partially Theorem 5.2 to oriented matroids.

5.1.1 The Even Directed Circuit Problem in Oriented Matroids

In what follows we introduce a generalisation of the graph theoretic notion of being noneven to oriented matroids and state the main results of this chapter. For our purposes, the most important examples of matroids are *graphical matroids* and *bond matroids*.

Digraphs can be seen as a special case of oriented matroids in the sense that every digraph D has an associated oriented graphic matroid M(D) whose signed circuits resemble the oriented cycles in the digraph D. In this spirit, it is natural to lift questions concerning cycles in directed graphs to more general problems on circuits in oriented matroids. The following algorithmic problem is the straightforward generalisation of the even dicycle problem to oriented matroids, and the main motivation of the research in this chapter.

Problem 5.1. Given an oriented matroid \vec{M} , decide whether there exists a directed circuit of even size in \vec{M} .

Our first contribution is to generalize the definition of non-even digraphs to oriented regular matroids in the following sense.

Definition 5.2. Let \vec{M} be an oriented matroid. We call \vec{M} non-even if its underlying matroid is regular and there exists a set $J \subseteq E(\vec{M})$ of elements such that every directed circuit in \vec{M} intersects J in an odd number of elements. If such a set does not exist, we call \vec{M} even.

The reader might wonder why the preceding definition concerns only regular matroids. This has several reasons. The main reason is a classical result by Bland and Las Vergnas [BLV78] which states that a binary matroid is orientable if and only if it is regular. Hence, if we were to extend the analysis of non-even oriented matroids beyond the regular case, we would have to deal with orientations of matroids which are not representable over \mathbb{F}_2 . This has several disadvantages, most importantly that cycles bases, which constitute an important tool in all of our results, are not guaranteed to exist any more. Furthermore, some of our proofs make use of the strong orthogonality property of oriented regular matroids (Equation ** from Chapter 1), which fails for non-binary oriented matroids. Lastly, since Problem 5.1 is an algorithmic question, oriented regular matroids have the additional advantage that they allow for a compact encoding in terms of totally unimodular matrices, which is not a given for general oriented matroids.

The first result of this chapter generalizes Theorem 5.1 to oriented matroids as follows:

Theorem 5.3. The problems of deciding whether an oriented regular matroid represented by a totally unimodular matrix contains an even directed circuit, and the problem of recognising whether an oriented regular matroid given by a totally unimodular matrix is noneven, are polynomially equivalent.

Theorem 5.3 motivates a structural study of the class of non-even oriented matroids, as in many cases the design of a recognition algorithm for a class of objects is based on a good structural understanding of the class. In order to state our main result, which is a generalisation of Theorem 5.2 to graphic and cographic oriented matroids, we have to introduce a new minor concept. We naturally generalize the concept of butterfly minors to regular oriented matroids, in the form of so-called *generalized butterfly minors*.

Definition 5.3. Let \vec{M} be an orientation of a regular matroid M. An element $e \in E(\vec{M})$ is called butterfly-contractible if there exists a cocircuit S in M such that $(S \setminus \{e\}, \{e\})$ forms a signed cocircuit of \vec{M} .¹ A generalized butterfly minor (GB-minor for short) of

¹For a definition of a signed (co)circuit see Chapter 1.

 \vec{M} is any oriented matroid obtained from \vec{M} by a finite sequence of element deletions and contractions of butterfly-contractible elements (in arbitrary order).

Note that the generalized butterfly-contraction captures the same fundamental idea as the initial one for digraphs while being more general: Given a butterfly-contractible element e of a regular oriented matroid \vec{M} , we cannot have a directed circuit C of \vec{M}/e such that $(C, \{e\})$ is a signed circuit of \vec{M}^2 , and hence either C or $C \cup \{e\}$ must form a directed circuit of \vec{M} .

Replacing the notion of butterfly minors by GB-minors allows us to translate Theorem 5.2 to the setting of oriented matroids in the following way:

Theorem 5.4. An oriented graphic matroid \vec{M} is non-even if and only if none of its GB-minors is isomorphic to $M(\vec{C}_k)$ for some odd $k \geq 3$.

As the main result of this chapter, we complement Theorem 5.4 by determining the list of forbidden GB-minors for cographic non-even oriented matroids. Recall the following notation from Chapter 1: For integers $m, n \ge 1$ we denote by $\vec{K}_{m,n}$ the digraph obtained from the complete bipartite graph $K_{m,n}$ by orienting all edges from the partition set of size m towards the partition set of size n.

Theorem 5.5. An oriented bond matroid \vec{M} is non-even if and only if none of its GBminors is isomorphic to $M^*(\vec{K}_{m,n})$ for any $m, n \ge 2$ such that m + n is odd.

To prove Theorem 5.5 we study those digraphs whose oriented bond matroids are noneven. Equivalently, these are the digraphs admitting an *odd dijoin*, which is an arc set hitting every directed bond an odd number of times. After translating GB-minors into a corresponding minor concept on directed graphs, which we call *cut minors*³, we show that the class of digraphs with an odd dijoin is described by two infinite families of minimal forbidden cut minors (Theorem 5.31). Finally, we translate this result to oriented bond matroids in order to obtain a proof of Theorem 5.5.

The structure of this chapter is as follows. In Section 5.2, we prove that non-even oriented matroids are closed under GB-minors (Proposition 5.6), which is then used to derive Theorem 5.4 in the same section. We start Section 5.3 by showing that the even directed circuit problem for general oriented matroids cannot be solved using only polynomially many calls to a signed circuit oracle (Proposition 5.7). The remainder of the section is devoted to the proof of Theorem 5.3. We also note that *odd* directed circuits in orientations of regular oriented matroids can be detected in polynomial time (Proposition 5.17). In Section 5.4 we characterize those digraphs that admit an odd dijoin (Theorem 5.31) and use this result to deduce our main result, Theorem 5.5. We conclude with a discussion of open problems and a conjecture in Section 5.5.

5.2 Non-Evenness and GB-minors

Our main result, Theorem 5.5, builds on the important fact that the non-even oriented matroids are closed under the GB-minor relation. In this subsection we present a proof of this fact and use it to derive Theorem 5.4 from Theorem 5.2.

Lemma 5.6. Every GB-minor of a non-even oriented matroid is non-even.

²In this case, $(C, \{e\})$ together with a signed cocircuit $(S \setminus \{e\}, \{e\})$ would contradict the orthogonality property (see Chapter 1, (*)) for oriented matroids.

 $^{^{3}}$ See the beginning of Section 5.4 for a precise definition.

Proof. It suffices to show the following two statements: For every non-even oriented matroid \vec{M} and every element $e \in E(\vec{M})$, the oriented matroid $\vec{M} - e$ is non-even as well, and for every element $e \in E(\vec{M})$ which is butterfly-contractible, the oriented matroid \vec{M}/e is non-even as well. The claim then follows by repeatedly applying these two statements.

For the first claim, note that since the underlying matroid M of \dot{M} is regular, so is the underlying matroid of $\vec{M} - e$. Let $J \subseteq E(\vec{M})$ be a set of elements intersecting every directed circuit in \vec{M} an odd number of times. Then clearly the set $J \setminus \{e\}$ intersects every directed circuit in $\vec{M} - e$ an odd number of times, proving that $\vec{M} - e$ is non-even.

For the second claim, let $e \in E(\vec{M})$ be butterfly-contractible. Let S be a cocircuit of M such that $(S \setminus \{e\}, \{e\})$ forms a signed cocircuit of \vec{M} . Then the underlying matroid of \vec{M}/e is a matroid minor of the regular matroid M and is hence regular. Let us define $J' \subseteq E(\vec{M}) \setminus \{e\}$ via

$$J' := \begin{cases} J & \text{if } e \notin J \\ J+S & \text{if } e \in J. \end{cases}$$

We claim that for every directed circuit C in \vec{M}/e , the intersection $C \cap J'$ is odd. Indeed, by definition either C is a directed circuit also in \vec{M} not containing e, or $C \cup \{e\}$ is a directed circuit in \vec{M} , or $(C, \{e\})$ is a signed circuit of \vec{M} . The last case however is impossible, as then the signed circuit $(C, \{e\})$ and the signed cocircuit $(S \setminus \{e\}, \{e\})$ in \vec{M} would yield a contradiction to the orthogonality property (*) of oriented matroids.

In the first case, since $e \notin C$, we must have $S \cap C = \emptyset$ as otherwise again C and the signed cocircuit $(S \setminus \{e\}, \{e\})$ form a contradiction to the orthogonality property (*). This then shows that indeed $|C \cap J'| = |C \cap (J' \setminus S)| = |C \cap (J \setminus S)| = |C \cap J|$ is odd, as required.

In the second case, the orthogonality property (**) of regular oriented matroids applied with the directed circuit $C \cup \{e\}$ and the signed cocircuit $(\{e\}, S \setminus \{e\})$ within \vec{M} yield that the equation $|(C \cup \{e\}) \cap (S \setminus \{e\})| = |(C \cup \{e\}) \cap \{e\}| = 1$ holds. So let $C \cap S =$ $\{f\}$ for some element $f \in E(\vec{M}) \setminus \{e\}$. By definition of J', if $e \notin J$, then we have $|C \cap J'| = |C \cap J| = |(C \cup \{e\}) \cap J|$, which is odd. If $e \in J$, then we have (modulo 2)

$$|C \cap J'| = |C \cap (J+S)| = |(C \cap J) + (C \cap S)| \equiv |C \cap J| + |\{f\}| = |(C \cup \{e\}) \cap J|,$$

which is odd. Hence, we have shown that $|C \cap J'|$ is odd in every case, which yields that \vec{M}/e is a non-even oriented matroid. This concludes the proof.

Proposition 5.6 allows us to immediately prove the correctness of Theorem 5.4.

Proof of Theorem 5.4. We prove both directions of the equivalence. Suppose first that \overline{M} is non-even. Then by Theorem 5.6 every oriented matroid isomorphic to a GB-minor of \overline{M} is non-even as well. Hence it suffices to observe that none of the matroids $M(\overrightarrow{C}_k)$ for odd $k \geq 3$ is non-even. However, this follows directly since any element set J in $M(\overrightarrow{C}_k)$ intersecting every directed circuit an odd number of times corresponds to an arc set in \overrightarrow{C}_k intersecting every directed cycle an odd number of times, which cannot exist since by Theorem 5.2 none of the digraphs \overrightarrow{C}_k is non-even for an odd $k \geq 3$.

Vice versa, suppose that no GB-minor of \vec{M} is isomorphic to $M(\vec{C}_k)$ for any odd $k \geq 3$. Let D be a digraph such that $\vec{M} \simeq M(D)$. We claim that D must be non-even. Suppose not, then by Theorem 5.2 D admits a butterfly minor isomorphic to \vec{C}_k for some odd $k \geq 3$. We now claim that M(D) has a GB-minor isomorphic to $M(\vec{C}_k)$. For this, it evidently suffices to verify the following general statement:

If an arc e of a digraph F is butterfly-contractible in F, then within M(F) the corresponding element e of M(F) is butterfly-contractible.

Indeed, let e = (u, v) for distinct vertices $u, v \in V(D)$. Then by definition either u has out-degree 1 or v has in-degree 1 in D. In the first case, e is the unique arc leaving u in the cut $\partial(\{u\})$, while in the second case e is the only arc entering v in the cut $\partial(\{v\})$. Since every cut is a disjoint union of bonds, we can find in both cases a bond containing e where e is the only arc directed away (towards, respectively) the side of the bond that contains u (v, respectively).

Since the oriented bonds in D yield the signed cocircuits of M(D), this shows that there is a cocircuit S in M(D) such that $(S \setminus \{e\}, \{e\})$ is a signed cocircuit. Hence, eis a butterfly-contractible element of M(D). This shows that $M(\vec{C}_k)$ is isomorphic to a GB-minor of $M(D) \simeq \vec{M}$ which contradicts our initial assumption that no GB-minor of \vec{M} is isomorphic to $M(\vec{C}_k)$. Hence, D is non-even, and there exists $J \subseteq A(D)$ such that every directed cycle in D contains an odd number of arcs from J. The same set J also certifies that $\vec{M} \simeq M(D)$ is non-even, and this concludes the proof of the equivalence. \Box

5.3 The Complexity of the Even Directed Circuit Problem

The formulation of Problem 5.1 is rather vague, as it is not clear by which means the oriented matroid \vec{M} is given as an input to an algorithm designed for solving the problem, and in which way we will measure its efficiency. For the latter, it is natural to aim for an algorithm which performs a polynomial number of elementary steps in terms of the number of elements of \vec{M} . This also resembles the even dicycle problem in digraphs, where we aim to find an algorithm running in polynomial time in a(D) (or, equivalently, v(D)).

For the former, it is not immediately clear how to encode the (oriented) matroid, and hence how to make information contained in the (oriented) matroid available to the algorithm. For instance the list of all circuits of a matroid, if given as input to an algorithm, will usually have exponential size in the number of elements, and therefore disgualify as a good reference value for efficiency of the algorithm. For that reason, different computational models (and efficiency measures) for algorithmic problems in matroids (see [HK81]) and oriented matroids (see [BR89]) have been proposed in the literature. These models are based on the concept of *oracles*. For a family $\mathcal{F} \subseteq 2^{E(M)}$ of objects characterising the matroid M, an oracle is a function $f: 2^{E(M)} \to \{\texttt{true}, \texttt{false}\}$ assigning to every subset a truth value indicating whether or not the set is contained in \mathcal{F} . If \mathcal{F} for instance corresponds to the collection of circuits, cocircuits, independent sets, or bases of a matroid, we speak of a *circuit*-, *cocircuit*-, *independence*-, or *basis-oracle*. Similarly, for oriented matroids we can define several oracles [BR89]. Maybe the most natural choice for an oriented matroid-oracle for Problem 5.1 is the *circuit oracle*, which given any subset of the element set together with a $\{+, -\}$ -signing of its elements, reveals whether or not this signed subset forms a signed circuit of the oriented matroid. This computational model applied to Problem 5.1 yields the following question.

Question 5.1. Does there exist an algorithm which, given an oriented matroid \vec{M} , decides whether there exists a directed circuit in \vec{M} of even size, by calling the circuit-oracle of \vec{M} only $O(|E(\vec{M})|^c)$ times for some $c \in \mathbb{N}$?

However, as it turns out, the answer to the above problem is quite trivially negative, even when the input oriented matroid \vec{M} is graphic.

Proposition 5.7. Any algorithm deciding whether a given oriented (graphic) matroid on n elements, for some $n \in \mathbb{N}$, contains an even directed circuit must use at least $2^{n-1} - 1$ calls to the circuit-oracle for some instances.

Proof. Suppose towards a contradiction there was an algorithm which decides whether a given oriented graphic matroid contains an even directed circuit and uses at most $2^{n-1}-2$ oracle calls for any input oriented graphic matroid on elements $E := \{1, \ldots, n\}$. Now, playing the role of the oracle, we will answer all of the (at most $2^{n-1}-2$) calls of the algorithm by **false**. Since there are exactly $2^{n-1}-1$ non-empty sets $Y \in 2^E$ of even size, there must be an even non-empty subset Y of E such that the algorithm did not call the oracle with any input signed set whose support is Y. But this means the algorithm cannot distinguish between the oriented graphic matroids (E, C_0) and (E, C_Y) , where $C_0 := \emptyset$ and $C_Y := \{(Y, \emptyset), (\emptyset, Y)\}$, which result in the same oracle-answers to the calls by the algorithm, while (E, C_0) contains no even directed circuit, but (E, C_Y) does. This shows that the algorithm does not work correctly, and this contradiction proves the assertion. \Box

The above result and its proof give a hint that maybe in general the use of oriented matroid-oracles to measure the efficiency of algorithms solving Problem 5.1 is doomed to fail. One should therefore look for a different encoding of the input oriented matroids in order to obtain a sensible algorithmic problem. Here we solve this issue by restricting the class of possible input oriented matroids to *oriented regular matroids*, which allow for a much simpler and compact encoding via their representation by totally unimodular matrices (cf. Theorem 1.4, Theorem 1.5). The following finally is the actual algorithmic problem we are going to discuss in this chapter.

Problem 5.2. Is there an algorithm which decides, given as input a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$, whether $\vec{M}[A]$ contains an even directed circuit, and runs in time polynomial in mn?

The alert reader might be wondering what happens if in the above problem we aim to detect odd instead of even directed circuits. The reason why this problem is not a center of study in our paper is that it admits a simple polynomial time solution, which is given in the form of Proposition 5.17 at the end of this section.

The next statement translates the main results from [RST99] and [McC04] to our setting to show that Problem 5.2 has a positive answer if we restrict to graphic oriented matroids as inputs.

Lemma 5.8. There exists an algorithm which, given as input any totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$ such that $\vec{M}[A]$ is a graphic oriented matroid, decides whether $\vec{M}[A]$ contains a directed circuit of even size, and which runs in time polynomial in mn.

Proof. The main results of Robertson et al. [RST99] and McCuaig [McC04] yield polynomial time algorithms which, given as input a digraph D (by its vertex- and arc-list) returns whether or not D contains an even directed cycle. Therefore, given a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ such that $\vec{M}[A]$ is graphic, if we can construct in polynomial time in mn a digraph D such that $\vec{M}[A] \simeq M(D)$, then we can decide whether $\vec{M}[A]$ contains a directed circuit of even size by testing whether D contains an even directed cycle using the algorithms from [McC04, RST99]. Such a digraph can be found as follows:

First, we consider the unoriented matroid M[A] defined by the matrix A, which is graphic. It follows from a result of Seymour [Sey81] that using a polynomial number in |E(M[A])| = n of calls to an independence-oracle for M[A], we can compute a connected graph G with n edges such that $M(G) \simeq M[A]$. For every given subset of columns of A we can test linear independence in polynomial time in mn, and hence we can execute the steps of Seymour's algorithm in polynomial time. Since $M(G) \simeq M[A]$, there must exist an orientation of M(G) isomorphic to $\tilde{M}[A]$, and this orientation in turn can be realized as M(D) where D is an orientation of G^4 . To find the desired orientation D of G in polynomial time, we first compute a decomposition of G into its blocks G_1, \ldots, G_k (maximal connected subgraphs without cutvertices).

Next we (arbitrarily) select for every $i \in \{1, \ldots, k\}$ a special "reference"-edge $e_i \in$ $E(G_i)$. Note that two different orientations of G obtained from each other by reversing all edges in one block result in the same oriented matroid, as cycles in G are always entirely contained in one block. Hence for every $i \in \{1, \ldots, k\}$ we can orient e_i arbitrarily and assume w.l.o.g. that this orientation coincides with the orientation in D. Note that every block of G which is not 2-connected must be a K_2 forming a bridge in G. In this case, the only edge of the block is our chosen reference-edge and already correctly oriented. Now, for every $i \in \{1, \ldots, k\}$ such that G_i is 2-connected and every edge $e \in E(G_i) \setminus \{e_i\}$ there is a cycle C in G_i containing both e_i and e. This cycle can be computed in polynomial time using a disjoint-paths algorithm between the endpoints of e and e_i . Now we consider the minimally linearly dependent set of columns in A corresponding to C, and compute the coefficients of a non-trivial linear combination resulting in 0. As we already know the orientation of $e_i \in E(C)$, this yields us the orientations of all edges on the cycle C in D and hence of the edge e. In this way, we can compute all orientations of edges in D in polynomial time in mn and find the digraph D such that $M[A] \simeq M(D)$. As discussed above, this concludes the proof.

5.3.1 Proof of Theorem 5.3

We prepare the proof by a set of useful definitions and lemmas dealing with circuit bases of regular matroids.

Definition 5.4. Let M be a binary matroid. The circuit space of M is the \mathbb{F}_2 -linear vector space generated by the incidence vectors $\mathbf{1}_C \in \mathbb{F}_2^{E(M)}$ defined by $\mathbf{1}_C(e) := 1$ for $e \in C$ and $\mathbf{1}_C(e) := 0$ for $e \notin C$ and all circuits C of M. A circuit basis of M is a set of circuits of M whose incidence vectors form a basis of the circuit space. Equivalently, we can consider the circuit space as a \mathbb{F}_2 -linear subspace of the vector space whose elements are all the subsets of E and where the sum X + Y of two sets $X, Y \subseteq E(M)$ is defined as their symmetric difference.

Definition 5.5. Let \vec{M} be a regular oriented matroid and M be its underlying regular matroid. We call a circuit basis \mathcal{B} of M directed if all its elements are directed circuits of \vec{M} .

The next proposition is a well-known fact about the circuit space of a binary matroid.

Proposition 5.9. Let M be a binary matroid. Then the dimension of the circuit space of M equals |E(M)| - r(M).

The following lemma is crucial for the proof of Theorem 5.3 as well as for our work on digraphs in Section 5.4.

Lemma 5.10. Let \vec{M} be an oriented regular matroid. If \vec{M} is totally cyclic, then the underlying matroid M admits a directed circuit basis. Furthermore, for every coindependent

⁴The fact that every orientation of M(G) can be realized as M(D) for an orientation D of G follows from a classical result by Bland and LasVergnas [BLV78], who show that regular matroids (and particularly graphic ones) have a unique reorientation class.

set A in M such that $\vec{M} - A$ is totally cyclic, there exists a directed circuit basis of M such that every $a \in A$ is contained in exactly one circuit of the basis.

Proof. We start by proving the first assertion concerning the existence of a directed circuit basis of M. We use induction on |E(M)|. If M consists of a single element, the claim holds trivially, since every circuit is a loop and thus directed. Moving on to the induction step, assume that $|E(M)| = k \ge 2$ and that the statement of the lemma holds for all oriented regular matroids on at most k - 1 elements. Choose some $e \in E(M)$ arbitrarily. Since \vec{M} is totally cyclic, there exists a directed circuit C_e containing e. Let us now consider the oriented regular matroid $\vec{M} - e$. If $\vec{M} - e$ is totally cyclic, then we can apply the induction hypothesis to $\vec{M} - e$ and find a directed circuit basis \mathcal{B}^- of M - e. Now consider the collection $\mathcal{B} = \mathcal{B}^- \cup \{C_e\}$ of directed circuits in \vec{M} . The incidence vectors of these circuits are linearly independent over \mathbb{F}_2 , as C_e is the only circuit yielding a non-zero entry at element e. Furthermore, we get by induction that $|\mathcal{B}| = |E(M)| - 1 - r(M - e) + 1 =$ |E(M)| - r(M - e) = |E(M)| - r(M). The last equality holds since e is contained in the circuit C_e and hence is not a coloop. As this matches the dimension of the circuit space of M, we have found a directed circuit basis of M, proving the inductive claim.

It remains to prove the case where $\overline{M} - e$ is not totally cyclic, i.e., there is an element not contained in a directed circuit. By Farkas' Lemma (Lemma 1.6) applied to $\overline{M} - e$ and this element there exists a directed cocircuit S in $\overline{M} - e$. Then either $(S, \emptyset), (S \cup \{e\}, \emptyset)$ or $(S, \{e\})$ form a signed cocircuit of \overline{M} . Since \overline{M} is totally cyclic, it contains no directed cocircuits, and hence only the latter case is possible, $(S, \{e\})$ must form a signed cocircuit.

Let us now consider the oriented regular matroid \vec{M}/e . Since \vec{M} is totally cyclic, so is \vec{M}/e . By the induction hypothesis there exists a directed circuit basis \mathcal{B}^- of M/e. By definition, for every directed circuit $C \in \mathcal{B}^-$, either C is a directed circuit in \vec{M} not containing e, or $C \cup \{e\}$ is a directed circuit in \vec{M} , or $(C, \{e\})$ forms a signed circuit of \vec{M} . The latter is however impossible, as in this case we can consider the signed cocircuit $X = (S, \{e\})$ and the signed circuit $Y = (C, \{e\})$ of \vec{M} , which satisfy $e \in X^- \cap Y^- \neq \emptyset$ but furthermore $(X^+ \cap Y^-) \cup (X^- \cap Y^+) = \emptyset$, violating the orthogonality property (*) of oriented matroids.

Hence, the set $\mathcal{B} := \{C \mid C \in \mathcal{B}^- \text{ circuit in } M\} \cup \{C \cup \{e\} \mid C \in \mathcal{B}^-, C \cup \{e\} \text{ circuit in } M\}$ consists of $|\mathcal{B}| = |\mathcal{B}^-| = |E(M)| - 1 - r(M/e) = |E(M)| - r(M)$ many circuits of M which are all directed ones in \vec{M} . Note that for the last equality we used that e is not a loop, as it is contained in the cocircuit $S \cup \{e\}$ of M. Finally, we claim that the binary incidence vectors of the elements of \mathcal{B} in $\mathbb{F}_2^{E(M)}$ are linearly independent. This follows since the restriction of these vectors to the coordinates $E(M) \setminus \{e\}$ equals the characteristic vectors of the elements of \mathcal{B}^- , which form a circuit basis of M/e. This shows that we have found a directed circuit basis of M, proving the inductive claim.

For the second assertion, let a coindependent set A in M be given and suppose that $\vec{M} - A$ is totally cyclic. We claim that for every $a \in A$ there exists a directed circuit C_a in \vec{M} such that $C_a \cap A = \{a\}$. Equivalently, we may show that the oriented matroid $\vec{M} - (A \setminus \{a\})$ has a directed circuit containing a. Towards a contradiction, suppose not, then by Farkas' Lemma (Lemma 1.6) there exists a directed cocircuit S in $\vec{M} - (A \setminus \{a\})$ containing a. Since A is coindependent, a is not a coloop of $M - (A \setminus \{a\})$ and hence $S \setminus \{a\} \neq \emptyset$. Every directed circuit in $\vec{M} - (A \setminus \{a\})$ must be disjoint from S, and hence no $f \in S \setminus \{a\}$ is contained in a directed circuit of $\vec{M} - A$, contradicting our assumption that $\vec{M} - A$ is totally cyclic. It follows that for each $a \in A$ a directed circuit C_a with $C_a \cap A = \{a\}$ exists.

Next we apply the first assertion of this lemma to the totally cyclic oriented matroid $\vec{M} - A$. We find that there exists a directed circuit basis \mathcal{B}_A of M - A. We claim that

 $\mathcal{B} := \mathcal{B}_A \cup \{C_a \mid a \in A\}$ forms a directed circuit basis of M satisfying the properties claimed in this lemma. Indeed, every circuit in \mathcal{B} is a directed circuit of \vec{M} , and for every $a \in A$ the circuit C_a is the only circuit in \mathcal{B} containing a. Since the characteristic vectors of the elements of \mathcal{B}_A are linearly independent as \mathcal{B}_A is a circuit basis of M - A, we already get that the characteristic vectors of elements of \mathcal{B} are linearly independent using that the characteristic vector of C_a is the only basis-vector having a non-zero entry at the position corresponding to element a. To show that \mathcal{B} indeed is a circuit basis of M, it remains to verify that it has the required size. We have $|\mathcal{B}| = |A| + |\mathcal{B}_A| =$ |A| + |E(M - A)| - r(M - A) = |E(M)| - r(M), where for the latter equality we used that r(M - A) = r(M) since A is coindependent. This concludes the proof of the second assertion. \Box

In order to prove our next lemma, we need the following result, which was already used by Seymour and Thomassen.

Lemma 5.11 ([ST87], Prop. 3.2). Let E be a finite set and \mathcal{F} a family of subsets of E. Then precisely one of the following statements holds:

- (i) There is a subset $J \subseteq E$ such that $|F \cap J|$ is odd for every $F \in \mathcal{F}$.
- (ii) There are sets $F_1, \ldots, F_k \in \mathcal{F}$, where $k \in \mathbb{N}$ is odd, such that $\sum_{i=1}^k F_i = \emptyset$.

Please note that (i) and (ii) cannot hold simultaneously because if k is odd and F_1, \ldots, F_k all have odd intersection with A, then the symmetric difference $\sum_{i=1}^k F_i$ has odd intersection with A.

We now derive the following corollary for totally cyclic oriented regular matroids by using Lemma 5.10 and applying Lemma 5.11 to a directed circuit basis.

Corollary 5.12. Let \vec{M} be a totally cyclic oriented regular matroid, and let \mathcal{B} be a directed circuit basis of M. Then there exists $J \subseteq E(\vec{M})$ such that $|C \cap J|$ is odd for every $C \in \mathcal{B}$.

Proof. The claim is that (i) in Lemma 5.11 with $E = E(\overline{M})$ and $\mathcal{F} := \mathcal{B}$ holds true, so it suffices to rule out (ii). However, the latter would contradict the linear independence of the basis \mathcal{B} .

Building on this corollary we derive equivalent properties for an oriented matroid to be non-even.

Proposition 5.13. Let \vec{M} be a totally cyclic oriented regular matroid and let \mathcal{B} be a directed circuit basis of M. Furthermore, let $J \subseteq E(M)$ be such that $|C \cap J|$ is odd for all $C \in \mathcal{B}$. Then the following statements are equivalent:

- (i) \vec{M} is non-even.
- (ii) If C_1, \ldots, C_k are directed circuits of \vec{M} where $k \in \mathbb{N}$ is odd, then $\sum_{i=1}^k C_i \neq \emptyset$.
- (iii) Every directed circuit of \vec{M} is a sum of an odd number of elements of \mathcal{B} .
- (iv) $|C \cap J|$ is odd for all directed circuits C of \vec{M} .

Proof.

"(i) \Rightarrow (ii)" This follows from Lemma 5.11 applied to the set of directed circuits of \vec{M} .

- "(ii) \Rightarrow (iii)" Let *C* be a directed circuit of \overline{M} . Since \mathcal{B} is a circuit basis of *M*, we can write $C = \sum_{i=1}^{k} C_i$ for some $k \in \mathbb{N}$ and $C_1, \ldots, C_k \in \mathcal{B}$. If k were even, then the sum $C + \sum_{i=1}^{k} C_i = \emptyset$ would yield a contradiction to (ii).
- "(iii) \Rightarrow (iv)" Let *C* be a directed circuit of \vec{M} . By assumption, $C = \sum_{i=1}^{k} C_i$ with $C_1, \ldots, C_k \in \mathcal{B}$ and $k \in \mathbb{N}$ being odd. Since *J* has odd intersection with all C_i , the set *J* has also odd intersection with *C*.
- "(iv) \Rightarrow (i)" This implication follows directly from the definition of non-even.

Before we turn towards the proof of Theorem 5.3 we need the following result, yielding a computational version of Farkas' Lemma for oriented regular matroids. Although we suspect the statement is well-known among experts, in the following we give a proof for the sake of completeness.

Lemma 5.14. There exists an algorithm that, given as input a regular oriented matroid \vec{M} represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ such that $\vec{M} \simeq \vec{M}[A]$, and an element $e \in E(\vec{M})$, outputs either a directed circuit of \vec{M} containing e or a directed cocircuit of \vec{M} containing e, and which runs in polynomial time in mn.

Proof. We first observe that we can decide in polynomial time in mn whether e is contained in a directed circuit or in a directed cocircuit of \vec{M} (by Farkas' Lemma, we know that exactly one of these two options must be satisfied). Let us denote for every element $f \in E(\vec{M})$ by $x_f \in \{-1, 0, 1\}^m$ the corresponding column-vector of A. Let us first prove the following auxiliary claim:

The element e is contained in a directed circuit of \vec{M} if and only if there exist nonnegative scalars $\alpha_f \geq 0$ for $f \in E(\vec{M}) \setminus \{e\}$ such that $-x_e = \sum_{f \in E(\vec{M} \setminus \{e\})} \alpha_f x_f$.

The necessity of this condition follows directly by definition of $\overline{M}[A]$: If e is contained in a directed circuit with elements e, f_1, \ldots, f_k , then there are coefficients $\beta_e > 0$ and $\beta_i > 0$ for $1 \leq i \leq k$ such that $\beta_e x_e + \sum_{i=1}^k \beta_i x_{f_i} = 0$, i.e., $-x_e = \sum_{i=1}^k \frac{\beta_i}{\beta_e} x_{f_i}$. On the other hand, if $-x_e$ is contained in the conical hull of $\{x_f | f \in E(\vec{M}) \setminus \{e\}\}$, then we can select an inclusion-wise minimal subset $F \subseteq E(\vec{M} \setminus \{e\})$ such that $-x_e$ is contained in the conical hull of $\{x_f | f \in F\}$. We claim that $\{e\} \cup F$ forms a directed circuit of \vec{M} . By definition of F, it suffices to verify that the vectors x_e and x_f for $f \in F$ are minimally linearly dependent. However, this follows directly by Carathéodory's Theorem: The dimension of the subspace spanned by $\{x_f | f \in F\}$ equals |F|, for otherwise we could select a subset of at most |F|-1 elements from $\{x_f | f \in F\}$ whose conical hull also contains $-x_e$, contradicting the minimality of F. This shows the equivalence claimed above.

We can now use the well-known linear programming algorithm for linear programs with integral constraints by Khachiyan [Kha79, GL81] to decide in strongly polynomial time⁵ (and hence in polynomial time in mn) the feasibility of the linear inequality system

$$\sum_{f \in E(\vec{M} \setminus \{e\})} \alpha_f x_f = -x_e, \text{ with } \alpha_f \ge 0.$$

Therefore, we have shown that we can decide in polynomial time in mn whether or not e is contained in a directed circuit of \vec{M} . Next we give an algorithm which, given that e is contained in a directed circuit of \vec{M} , finds such a circuit in polynomial time:

⁵Here we use the fact that all coefficients appearing in the linear system are -1, 0 or 1.

During the procedure, we update a subset $Z \subseteq E(\vec{M})$, which maintains the property that it contains a directed circuit including e. At the end of the procedure Z will form such a directed circuit of \vec{M} . We initialize $Z := E(\vec{M})$. During each step of the procedure, we go through the elements $f \in Z \setminus \{e\}$ one by one and apply the above algorithm to test whether $\vec{M}[Z] - f$ contains a directed circuit including e. At the first moment such an element is found, we put $Z := Z \setminus \{f\}$ and repeat. If no such element is found, we stop and output Z.

Since we reduce the size of the set Z at each round of the procedure, the above algorithm runs in at most n rounds and calls the above decision algorithm for the existence of a directed circuit including e at most n-1 times in every round. All in all, the algorithm runs in time polynomial in mn. It is obvious that the procedure maintains the property that Z contains a directed circuit including e and that at the end of the procedure all elements of Z must be contained in this circuit, i.e., Z forms a directed circuit with the desired properties.

To complete the proof we now give an algorithm which finds either a directed circuit or a directed cocircuit through a given element e of \vec{M} as follows: First we apply the first (decision) algorithm, which either tells us that e is contained in a directed circuit of \vec{M} , in which case we apply the second (detection) algorithm to find such a circuit. Otherwise we know that e is contained in a directed cocircuit of \vec{M} , in which case we compute in polynomial time a totally unimodular representing matrix A^* with at most n rows and n columns⁶ of the dual regular oriented matroid \vec{M}^* . As we know that e is included in a directed circuit of \vec{M}^* , we can apply the second (detection) algorithm to A^* and \vec{M}^* instead of A and \vec{M} to find a directed cocircuit in \vec{M} containing e in polynomial time. \Box

Given a regular oriented matroid \vec{M} we shall denote by $TC(\vec{M})$ the largest totally cyclic deletion minor of \vec{M} , i.e. the deletion minor of \vec{M} whose ground set is

$$E(TC(\vec{M})) := \bigcup \{ C \mid C \text{ is a directed circuit of } \vec{M} \}.$$

From Lemma 5.14 we directly have the following.

Corollary 5.15. Let \vec{M} be a regular oriented matroid represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ for some $m \in \mathbb{N}$ and n = |E(M)|. Then we can compute a representative matrix of $TC(\vec{M})$ in time polynomial in mn.

The last ingredient we shall need for the proof of Theorem 5.3 is a computational version of the first statement of Lemma 5.10 combined with Corollary 5.12.

Lemma 5.16. Let \vec{M} be a totally cyclic regular oriented matroid represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ for some $m \in \mathbb{N}$ and n = |E(M)|. Then we can compute a directed circuit basis \mathcal{B} of \vec{M} together with a set $J \subseteq E(\vec{M})$ such that $|J \cap B| \equiv 1 \pmod{2}$ for every $B \in \mathcal{B}$ in time polynomial in mn.

Proof. We shall follow the inductive proof of Lemma 5.10 to obtain a recursive algorithm for finding a desired directed circuit basis together with the desired set J. If n = 1, the unique element e of $E(\vec{M})$ is a directed loop, since \vec{M} is totally cyclic, and forms our desired directed circuit basis of \vec{M} . Furthermore, we may simply set $J := \{e\}$.

⁶To find such a representing matrix, one can use Gaussian elimination to compute a basis \mathcal{B} of ker(A). Since A is totally unimodular, the vectors in \mathcal{B} can be taken to be $\{-1, 0, 1\}$ -vectors such that the matrix A^* consisting of the elements of \mathcal{B} written as row-vectors is totally unimodular as well. It then follows from the orthogonality property of regular oriented matroids that A^* indeed forms a representation of \vec{M}^* , using the fact that the row spaces of A and A^* are orthogonal complements.

In the case $n \geq 2$, let us fix an arbitrary element e of $E(\vec{M})$ and compute a directed circuit C_e of \vec{M} containing e by applying Lemma 5.14. Also using Lemma 5.14, we can test in time polynomial in mn whether $\vec{M} - e$ is totally cyclic. If so, we fix C_e as an element of our desired directed circuit base \mathcal{B} of \vec{M} and proceed as before with $\vec{M} - e$ instead of \vec{M} . The set J is updated as follows: Suppose we have already computed a directed circuit base \mathcal{B}^- and a set J^- as in the statement of this lemma, but with respect to $\vec{M} - e$. Then we set $\mathcal{B} := \mathcal{B}^- \cup \{C_e\}$. Now we check the parity of $|J^- \cap C_e|$ and set

$$J := \begin{cases} J^- & \text{if } |J^- \cap C_e| \equiv 1 \pmod{2} \\ J^- \cup \{e\} & \text{if } |J^- \cap C_e| \equiv 0 \pmod{2}. \end{cases}$$

As C_e is the only element of \mathcal{B} that contains e, the set J has odd intersection with every element of \mathcal{B} , as desired.

If $\vec{M} - e$ is not totally cyclic, we compute a totally unimodular representative matrix $A' \in \{-1, 0, 1\}^{m \times (n-1)}$ of \vec{M}/e . This task can be executed in time polynomial in mn^{-7} . Now \vec{M}/e is totally cyclic as \vec{M} is totally cyclic and we proceed as before with \vec{M}/e instead of \vec{M} . However, when our recursive algorithm already yields a directed circuit basis \mathcal{B}^- of \vec{M}/e as well as a set J^- for \vec{M}/e as in the statement of this lemma, we know as argued in the proof of Lemma 5.10 that each element C of \mathcal{B}^- either is a directed circuit of \vec{M} or $C \cup \{e\}$ is a directed circuit of \vec{M} . Depending on this distinction we define our desired circuit basis \mathcal{B} of \vec{M} as in the proof of Lemma 5.10 via

$$\mathcal{B} := \{ C \mid C \in \mathcal{B}^- \text{ circuit in } M \} \cup \{ C \cup \{e\} \mid C \in \mathcal{B}^-, C \cup \{e\} \text{ circuit in } M \}.$$

To decide for each element $C \in \mathcal{B}^-$ whether C or $C \cup \{e\}$ is a directed circuit of \vec{M} we calculate $A\mathbf{1}_C$ where $\mathbf{1}_C$ denotes the incidence vector of C with respect to A. Then C forms a directed circuit of \vec{M} if and only if $A\mathbf{1}_C = 0$. As $|\mathcal{B}^-| = |\mathcal{B}| = |\mathcal{E}(\vec{M})| - r(\vec{M})$ as argued in the proof of Lemma 5.10 and by Proposition 5.9, we have to do at most n of these computations to compute \mathcal{B} from \mathcal{B}^- .

Regarding the set J we can simply set $J := J^-$.

We are now ready for the proof of Theorem 5.3.

Proof of Theorem 5.3. Assume first we have access to an oracle deciding whether an oriented regular matroid given by a representing totally unimodular matrix is non-even. Suppose we are given a regular oriented matroid \vec{M} represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ for some $m, n \in \mathbb{N}$ and we want to decide whether it contains a directed circuit of even size.

First we compute $TC(\vec{M})$, which can be done in time polynomial in mn by Corollary 5.15. Now we use Lemma 5.16 to compute a directed circuit basis of $TC(\vec{M})$ in time polynomial in mn. Then we go through the $|E(TC(\vec{M}))| - r(TC(\vec{M}))$ many elements of the basis and check whether one of these directed circuits has even size. If so, the algorithm terminates. Otherwise, every member of the basis has odd size. By Proposition 5.13 with $J := E(TC(\vec{M}))$, we know that $TC(\vec{M})$ contains no directed circuit of even size if and only if $TC(\vec{M})$ is non-even. Since $TC(\vec{M})$ is the largest deletion minor of \vec{M} , which

⁷To compute A', select a non-zero entry in the column of A belonging to the element e. Pivoting on this element and exchanging rows transforms A in polynomial time in mn into a totally unimodular matrix $A'' \in \{-1, 0, 1\}^{m \times n}$ of \vec{M} in which the column corresponding to the element e of \vec{M} is $(1, 0, \ldots, 0)^{\top}$. Then $\vec{M}[A] = \vec{M}[A'']$, and the matrix A' obtained from A'' by deleting the first row as well as the column corresponding to e is a totally unimodular representation of \vec{M}/e .

has the same directed circuits as \vec{M} , we know that $TC(\vec{M})$ is non-even if and only if \vec{M} is non-even. So we can decide the question using the oracle.

Conversely, assume we have access to an oracle which decides whether a given oriented regular matroid contains a directed circuit of even size. Again, our first step is to compute $TC(\vec{M})$ using Corollary 5.15. By Lemma 5.16 we then compute a directed circuit basis of $TC(\vec{M})$ and a set $J \subseteq E(TC(\vec{M}))$ such that every circuit in the basis has odd intersection with J. Let $\vec{M'}$ be the oriented matroid obtained from $TC(\vec{M})$ by duplicating every element $e \in E(TC(\vec{M})) \setminus J$ into two copies e_1 and e_2^8 . This way, every directed circuit in $\vec{M'}$ intersects $E(\vec{M'}) \setminus J$ in an even number of elements. Thus, for every directed circuit C in $TC(\vec{M})$, the size of the corresponding directed circuit in $\vec{M'}$ is odd if and only if $|C \cap J|$ is odd. Hence, J intersects every directed circuit. By Proposition 5.13 this shows that $TC(\vec{M})$ is non-even if and only if $\vec{M'}$ has no directed circuit of even size. Since $TC(\vec{M})$ is non-even if and only if $\vec{M'}$ is non-even, we can decide the non-evenness of \vec{M} by negating the output of the oracle with instance $\vec{M'}$.

With the tools developed in this section we can give the proof of Proposition 5.17.

Proposition 5.17. There is an algorithm which given as input a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$, either returns an odd directed circuit of $\vec{M}[A]$ or concludes that no such circuit exists, and runs in time polynomial in mn.

Proof. Let $A \in \mathbb{R}^{m \times n}$ be a totally unimdoular matrix given as input and let $\vec{M} := \vec{M}[A]$. To decide whether \vec{M} contains a directed circuit of odd size, we first use Corollary 5.15 to compute a totally unimdoular representation of $TC(\vec{M})$ in polynomial time in mn. We now apply Lemma 5.16 to compute in polynomial time a directed circuit basis \mathcal{B} of $TC(\vec{M})$. Going through the elements of \mathcal{B} one by one, we test whether one of the basiscircuits is odd, in which case the algorithm stops an returns this circuit. Otherwise, all circuits in \mathcal{B} are even. Since every circuit in the underlying matroid of $TC(\vec{M})$ can be written as a symmetric difference of elements of \mathcal{B} , every circuit in this matroid must be even. In particular, $TC(\vec{M})$ and hence \vec{M} do not contain any odd directed circuits, and the algorithm terminates with this conclusion.

5.4 Digraphs Admitting an Odd Dijoin

This section is dedicated to the proof of our main result, Theorem 5.5. The overall strategy to achieve this goal is to work on digraphs and their families of bonds directly. The object that certifies that the bond matroid of a digraph is non-even is called an *odd dijoin*.

Definition 5.6. Let D be a digraph. A subset $J \subseteq A(D)$ is called an odd dijoin if $|J \cap S|$ is odd for every directed bond S in D.

Let D be a digraph. The contraction D/A of an arc-set $A \subseteq A(D)$ in D is understood as the digraph arising from D by deleting all arcs of A and identifying each weak connected component of D[A] into a corresponding vertex. Note that this might produce new loops arising from arcs spanned between vertices incident with A but not included in A. Note that contracting a loop is equivalent to deleting the loop.

⁸In particular, we transform every signed circuit of $TC(\vec{M})$ into a signed circuit of \vec{M}' by replacing every occurrence of an element $e \in E(TC(\vec{M})) \setminus J$ in a signed partition by the two elements e_1, e_2 in the same set of the signed partition. It is not hard to see that this indeed defines an oriented matroid, which is still regular.

An arc e = (x, y) of a digraph D, which is not a loop, is said to be *deletable* (or *transitively reducible*) if there is a directed path in D starting in x and ending in y which does not use e. Note that an arc $e \in A(D)$ is deletable if and only if e is a butterfly-contractible element of $M^*(D)$.

For two digraphs D_1, D_2 , we say that D_1 is a *cut minor* of D_2 if it can be obtained from D_2 by a finite series of arc contractions, deletions of deletable arcs, and deletions of isolated vertices.

Our next lemma guarantees that the property of admitting an odd dijoin is closed under the cut minor relation.

Lemma 5.18. Let D_1, D_2 be digraphs such that D_1 is a cut minor of D_2 . If D_2 admits and odd dijoin, then so does D_1 .

Proof. The statement follows by applying Proposition 5.6 to $M^*(D_1)$ and $M^*(D_2)$, noting that deleting isolated vertices from a digraph does not change the induced oriented bond matroid, and that by definition, an arc in a digraph D is deletable if and only if it is a butterfly-contractible element of $M^*(D)$.

Our goal will be to characterize the digraphs admitting an odd dijoin in terms of forbidden cut minors. In the following, we prepare this characterization by providing a set of helpful statements. For an undirected graph G, we define the *cutspace* of G as the \mathbb{F}_2 linear vector space generated by the bonds in G, whose addition operation is the symmetric difference and whose neutral element is the empty set. The following statements are all obtained in a straightforward way by applying the oriented matroid results Lemma 5.10, Corollary 5.12 respectively Proposition 5.13 to the oriented bond matroid $M^*(D)$ of D.

Corollary 5.19. Let D be a weakly connected and acyclic digraph with underlying multigraph G. Then the cut space of G admits a basis \mathcal{B} whose elements are directed bonds in D. Moreover, if $A \subseteq A(D)$ is a set of arcs such that D/A is acyclic and G[A] is a forest, then one can choose \mathcal{B} such that every arc $e \in A$ appears in exactly one bond of the basis.

Corollary 5.20. Let D be a digraph and let \mathcal{B} be a basis of the cut space consisting of directed bonds. Then there is an arc set $J' \subseteq A(D)$ such that $|J' \cap B|$ is odd for all $S \in \mathcal{B}$.

Proposition 5.21. Let D be a digraph, \mathcal{B} be a basis of the cut space consisting of directed bonds, and let $J' \subseteq A(D)$ be such that $|B \cap J'|$ is odd for all $B \in \mathcal{B}$. Then the following statements are equivalent:

- (i) D has an odd dijoin.
- (ii) If B_1, \ldots, B_k are directed bonds of D with k odd, then $\sum_{i=1}^k B_i \neq \emptyset$.
- (iii) Every directed bond of D can be written as $\sum_{i=1}^{k} B_i$ with k odd.
- (iv) J' is an odd dijoin of D.

5.4.1 Forbidden Cut Minors for Digraphs with an Odd Dijoin

Next we characterize the digraphs admitting an odd dijoin in terms of minimal forbidden cut minors. For this purpose, we identify the digraphs without an odd dijoin for which

every proper cut minor has an odd dijoin. We call such a digraph a minimal obstruction. Recall that a digraph D = (V, E) is said to be oriented if it has no loops, no parallel, and no anti-parallel arcs. Furthermore, D is called *transitively reduced* if for every arc $e = (v, w) \in E$ the only directed path in D starting at v and ending in w consists of e itself, or equivalently, if no arc in D is deletable.

We start with the following crucial lemma, which will be used multiple times to successively find the structure minimal obstructions must have.

Lemma 5.22. Let D be a minimal obstruction. Then the underlying multi-graph G of D is 2-vertex-connected. Furthermore, D is oriented, acyclic, and transitively reduced.

Proof. Assume that D has no odd dijoin, but every proper cut minor of D has one. Then it is easy to check that $v(D) \ge 4$.

To prove that G must be 2-vertex-connected, suppose towards a contradiction that G can be written as the union of two proper subgraphs G_1, G_2 with the property that $|V(G_1) \cap V(G_2)| \leq 1$. Then the orientations D_1, D_2 induced on G_1, G_2 by D are proper cut minors of D: Indeed, for $i \in \{1, 2\}$ we can obtain D_i from D by contracting all arcs in D_{3-i} and then deleting all the resulting isolated vertices outside $V(D_i)$. Since D_1, D_2 are proper cut minors of D, they must admit odd dijoins J_1, J_2 , respectively. However, since D_1 and D_2 share at most a single vertex, the directed bonds of D are either directed bonds of D_1 or of D_2 . Hence, the disjoint union $J_1 \cup J_2$ defines an odd dijoin of D and yields the desired contradiction.

To prove acyclicity, assume towards a contradiction that there is a directed cycle C in D. Let us consider the digraph D/A(C). This is a proper cut minor of D and therefore must have an odd dijoin J. However, the directed bonds in D/A(C) equal the directed bonds in D which are arc-disjoint from C, and since C is directed, these are already all the directed bonds of D. Hence J is an odd dijoin also for D, which is a contradiction.

To prove that D is transitively reduced, assume towards a contradiction that there was an arc $e = (x, y) \in A(D)$ and a directed path P from x to y not containing e. Then e is a deletable arc and D - e is a cut minor of D, which therefore must have an odd dijoin $J \subseteq A(D) \setminus \{e\}$. Note that a directed cut in D either does not intersect $\{e\} \cup A(P)$ at all or contains e and exactly one arc from P. It follows from this that for every directed bond B in D, we get that B - e is a directed bond of D - e. This directly yields that Jis also an odd dijoin of D, contradiction.

Clearly, the fact that D is oriented follows from D being simultaneously acyclic and transitively reduced. This concludes the proof of the lemma.

From this, we directly have the following useful observations.

Corollary 5.23. Let D be a minimal obstruction. Then for every arc $e \in A(D)$, the digraph D/e is acyclic. Similarly, for every vertex $v \in V(D)$ which is either a source or a sink, the digraph $D/E_D(v)$ is acyclic.

Proof. Let e be an arc of D. Since D is a minimal obstruction, we know by Lemma 5.22 that e is no loop. Now assume towards a contradiction that there was a directed cycle in D/e. As D itself is acyclic according to Lemma 5.22, this implies that there is a directed path P in D connecting the end vertices of e, which does not contain e itself. This path together with e now either contradicts the fact that D is acyclic or the fact that D is transitively reduced, both of which hold due to Lemma 5.22.

For the second part assume w.l.o.g. (using the symmetry given by reversing all arcs) that v is a source. Suppose for a contradiction there was a directed cycle in $D/E_D(v)$.

This implies the existence of a directed path P in D - v which connects two different vertices in the neighbourhood of v, say it starts in $w_1 \in N(v)$ and ends in $w_2 \in N(v)$.

Now the directed path $(v, w_1) + P$ witnesses that the arc (v, w_2) is deletable contradicting that D is transitively reduced. This concludes the proof of the second statement. \Box

Lemma 5.24. Let D be a minimal obstruction. If $A \subseteq A(D)$ is such that D/A is acyclic and such that D[A] is an oriented forest, then there is a directed bond in D which fully contains A.

Proof. By Corollary 5.19 there is a basis \mathcal{B} of the cut space consisting of directed bonds such that each $e \in A$ is contained in exactly one of the bonds in the basis. Moreover, by Corollary 5.20 there is $J' \subseteq A(D)$ such that each $B \in \mathcal{B}$ has odd intersection with J'. Since D has no odd dijoin, there has to be a directed bond B_0 in D such that $|B_0 \cap J'|$ is even. Let $B_0 = B_1 + \ldots + B_m$ be the unique linear combination with pairwise distinct $B_1, \ldots, B_m \in \mathcal{B}$. Clearly, m must be even. Let D' be the cut minor obtained from D by contracting the arcs in $A(D) \setminus \bigcup_{i=1}^m B_i$. The bonds B_0, B_1, \ldots, B_m are still directed bonds in D' and satisfy $B_0 + \ldots + B_m = \emptyset$, while m + 1 is odd.

The equivalence of (i) and (ii) in Proposition 5.21 now yields that D' has no odd dijoin. By the minimality of D we thus must have D = D' and $\bigcup_{i=1}^{m} B_i = A(D)$. It follows that every $e \in A$ is contained in exactly one of the bonds B_i and thus also in B_0 . We obtain that A is fully included in the directed bond B_0 , as required.

Corollary 5.25. Let D = (V, A) be a minimal obstruction. For $i \in \{1, 2\}$ let $\emptyset \neq A_i \subseteq E$ be such that $D[A_i]$ is a forest and D/A_i is acyclic. Suppose there is a directed cut $\partial(X)$ in D separating A_1 from A_2 , i.e., such that $A_1 \subseteq A(D[X])$ and $A_2 \subseteq A(D[V \setminus X])$. Then there exists a directed bond in D containing $A_1 \cup A_2$.

Proof. Let $A := A_1 \cup A_2$. As A_1 and A_2 induce vertex-disjoint forests, D[A] is a forest as well. Since no arc is directed from a vertex in $V \setminus X$ to a vertex in X, no directed circuit in D/A can contain a contracted vertex from A_1 and a contracted vertex from A_2 , so every directed circuit must already exist in D/A_1 or in D/A_2 . Because these two digraphs are acyclic, D/A is acyclic. Hence, by Lemma 5.24, A is fully included in a directed bond of D. This proves the assertion.

With the next proposition we shall make the structure of minimal obstructions much more precise. To state the result, we shall make use of the following definition.

Definition 5.7. Let $n_0, n_1, n_2 \in \mathbb{N}$. Then we denote by $\mathcal{D}(n_0, n_1, n_2)$ the digraph (V, E), where $V = V_0 \cup V_1 \cup V_2$ with $V_i = [n_i]$ for $i \in \{1, 2, 3\}$, and $E = (V_0 \times V_1) \cup (V_1 \times V_2)$.

Proposition 5.26. Let D = (V, E) be a minimal obstruction. Then D is isomorphic to $\mathcal{D}(n_1, n_2, n_3)$ for some integers $n_1, n_2, n_3 \ge 0$.

Proof. We shall split the proof into several claims, starting with the following one.

Claim I. D contains no directed path of length 3.

Suppose towards a contradiction that $v_0, e_1, v_1, e_2, v_2, e_3, v_3$ is a directed path of length 3 in D with $e_1 = (v_0, v_1), e_2 = (v_1, v_2), e_3 = (v_2, v_3)$. By Corollary 5.23, D/e_1 and D/e_3 are acyclic. Moreover, because D is acyclic by Lemma 5.22, the arc e_2 is contained in a directed cut $\partial(X)$ in D, separating $\{e_1\}$ and $\{e_3\}$. By Corollary 5.25 this means that there is a directed bond $\partial(Y)$ in D containing both e_1 and e_3 . This however means that $v_0, v_2 \in Y$ and $v_1, v_3 \notin Y$. Hence, e_2 is an arc in D starting in $V(D) \setminus Y$ and ending in Y, a contradiction since $\partial(Y)$ is a directed bond. This completes the proof of Claim I.

For $i \in \{0, 1, 2\}$ let V_i denote the set of vertices $v \in V$ such that the longest directed path ending in v has length i. By definition of the V_i and since D is acyclic, there is no arc from a vertex in V_i to a vertex in V_j for $i \ge j$, as otherwise this would give rise to a directed path of length i + 1 ending in a vertex of V_j .

By Claim I we know that $V = V_0 \cup V_1 \cup V_2$. We move on by proving the following claim.

Claim II. Every vertex $v \in V_1$ is adjacent to every vertex $u \in V_0$.

Let $v \in V_1$ and $u \in V_0$. Assume for a contradiction that u is not adjacent to v. By definition of V_1 there is an arc f = (u', v) with $u' \in V_0$. By Corollary 5.23, D/f and $D/E_D(u)$ are acyclic because u is a source. Let $X \supseteq \{u', v\}$ be the set of all vertices from which v can be reached via a directed path. Clearly $\partial(X)$ is a directed cut in D. As $u \in V_0 \setminus X$ is a source which is not adjacent to v, we conclude that $\{u\} \cup N(u) \subseteq V \setminus X$. This however means that the directed cut $\partial(X)$ separates f from the arcs in $E_D(u)$. By Corollary 5.25, this means that there is a directed bond $\partial(Y)$ in D containing $E_D(u) \cup \{f\}$. Since $E_D(u) = \partial(\{u\})$ itself is a directed cut in D, this contradicts the fact that $\partial(Y)$ is an inclusion-wise minimal directed cut in D, and proves Claim II.

We proceed with another claim.

Claim III. D does not contain any arc from V_0 to V_2 .

Let $u \in V_0$ and $w \in V_2$. By definition of V_2 there is some $v \in V_1$ such that $(v, w) \in A$. By Claim II, $(u, v) \in E$. Because D is transitively reduced by Lemma 5.22, we obtain $(u, w) \notin A$. So the proof of Claim III is complete.

Now we come to the last claim we need for the proof of this proposition.

Claim IV. Every vertex $v \in V_1$ sees every vertex $w \in V_2$.

Let $v \in V_1, w \in V_2$ and suppose for a contradiction that w is not adjacent to v. Let f = (u, v) be an arc with $u \in V_0$. By Lemma 5.24, D/f and $D/E_D(w)$ are acyclic because w is a sink. Let $X \supseteq \{u, v\}$ be the set of all vertices from which v can be reached via a directed path. Again, $\partial(X)$ forms a directed cut in D. Claim III implies that $N(w) \subseteq V_1 \setminus \{v\} \subseteq V \setminus X$. This means $\partial(X)$ separates f from the arcs in $E_D(w)$, contradicting Corollary 5.25 again.

By combining all four claims we obtain $A = (V_0 \times V_1) \dot{\cup} (V_1 \times V_2)$, and the proof of this proposition is complete.

Now Proposition 5.26 puts us in the comfortable situation that the only possible minimal obstructions to having an odd dijoin are part of a 3-parameter class of simply structured digraphs. The rest of this section is devoted to determine the conditions on n_1, n_2, n_3 that need to be imposed such that $\mathcal{D}(n_1, n_2, n_3)$ is a minimal obstruction. It will be helpful to use the well-known concept of so-called *T*-joins.

Definition 5.8. Let G be an undirected graph and $T \subseteq V(G)$ be some vertex set. A subset $J \subseteq E(G)$ of edges is called a T-join, if in the spanning subgraph H := G[J] of G, every vertex in T has odd, and every vertex in $V(G) \setminus T$ has even degree.

The following result is folklore.

Lemma 5.27. A graph G with some vertex set $T \subseteq V(G)$ admits a T-join if and only if T has an even number of vertices in each connected component of G.

We continue with an observation about odd dijoins in digraphs of the form $\mathcal{D}(n_1, n_2, 0)$.

Observation 5.28. Let $n_1, n_2 \ge 1$. Then the digraph $\mathcal{D}(n_1, n_2, 0) \simeq \mathcal{D}(0, n_1, n_2)$ has an odd dijoin if and only if $\min(n_1, n_2) \le 1$ or $n_1, n_2 \ge 2$ and $n_1 \equiv n_2 \pmod{2}$.

Proof. If $\min(n_1, n_2) \leq 1$, then all directed bonds in $\mathcal{D}(n_1, n_2, 0)$ consist of single arcs and thus, $J := A(\mathcal{D}(n_1, n_2, 0))$ defines an odd dijoin. If $n_1, n_2 \geq 2$, the directed bonds in $\mathcal{D}(n_1, n_2, 0)$ are exactly those cuts with one vertex on one side of the cut and all other vertices on the other side. Hence, there is an odd dijoin if and only if the complete bipartite graph with partition classes of size n_1, n_2 has a *T*-join, where *T* contains all $n_1 + n_2$ vertices. The statement is now implied by Lemma 5.27.

Next we characterize when the digraphs $\mathcal{D}(n_1, n_2, n_3)$ admit an odd dijoin.

Proposition 5.29. Let $n_1, n_2, n_3 \ge 1$ be integers. Then $\mathcal{D}(n_1, n_2, n_3)$ has an odd dijoin if and only if one of the following holds:

- (*i*) $n_2 = 1$.
- (*ii*) $n_2 = 2$ and $n_1 \equiv n_3 \pmod{2}$.
- (*iii*) $n_2 \ge 3$, and $n_1 \equiv n_3 \equiv 1 \pmod{2}$.

Proof. If $n_2 = 1$, then $\mathcal{D}(n_1, n_2, n_3)$ is an oriented star. Clearly, here, the directed bonds consist of single arcs, and therefore, $J := A(\mathcal{D}(n_1, 1, n_3))$ defines an odd dijoin.

If $n_2 = 2$, it is easily seen that $\mathcal{D}(n_1, 2, n_3)$ is a planar digraph, which admits a directed planar dual isomorphic to a bicycle $\overrightarrow{C}_{n_1+n_3}$ of length $n_1 + n_3$.

By planar duality, we know that $\mathcal{D}(n_1, 2, n_3)$ has an odd dijoin if and only if there is a subset of arcs of $\vec{C}_{n_1+n_3}$ which intersects every directed cycle an odd number of times. By Theorem 5.2 we know that such an arc set exists if and only if $n_1 + n_3$ is even, that is, $n_1 \equiv n_3 \pmod{2}$.

Therefore, we assume that $n_2 \geq 3$ for the rest of the proof. We now first show the necessity of (iii). So assume that $D := \mathcal{D}(n_1, n_2, n_3)$ has an odd dijoin J. We observe that the underlying graph of D is 2-connected. Hence, for every vertex $x \in V_1 \cup V_2 \cup V_3$, the cut $E_D(x)$ of all arcs incident with x is a minimal cut of the underlying graph, and it is directed in D whenever $x \in V_1 \cup V_3$. Therefore, U(D[J]) must have odd degree at every vertex in $V_1 \cup V_3$. Moreover, we observe that for any proper non-empty subset $X \subsetneq V_2$, the cut in D induced by the partition $(V_1 \cup X, (V_2 \setminus X) \cup V_3)$ is minimal and directed. In the following, we denote this cut by F(X). Now for every vertex $x \in V_2$, choose some $x' \in V_2 \setminus \{x\}$ and consider the minimal directed cuts $F(\{x'\}), F(\{x, x'\})$. Both are minimal directed cuts (here, we use that $n_2 \geq 3$) and thus must have odd intersection with J. Moreover, the symmetric difference $F(\{x'\}) + F(\{x, x'\})$ contains exactly the set $E_D(x)$ of arcs incident with x in D. We conclude the following:

$$|E_D(x) \cap J| = |(F(\{x'\}) + F(\{x, x'\})) \cap J| \equiv |F(\{x'\}) \cap J| + |F(\{x, x'\}) \cap J|$$
$$\equiv 1 + 1 \equiv 0 \pmod{2}$$

As $x \in V_2$ was chosen arbitrarily, we conclude that J must be a T-join of the underlying multi-graph of $\mathcal{D}(n_1, n_2, n_3)$ where $T = V_1 \cup V_3$. Now Lemma 5.27 implies that $|T| = n_1 + n_3$ must be even and hence $n_1 \equiv n_3 \pmod{2}$.

We claim that (iii) must be satisfied, i.e., n_1 and n_3 are odd. Assume towards a contradiction that this is not the case. Hence, by our observation above both n_1 and n_3 are even. Let $x \in V_2$ be some vertex, and consider the directed bond $F(\{x\})$. We can rewrite this bond as the symmetric difference of the directed cut $\partial(V_1) = \{(v_1, v_2) \mid v_1 \in V_1, v_2 \in V_2\}$ and the cut $E_D(x)$ of all arcs incident with x. Because $|E_D(u) \cap J|$ is odd for every $u \in V_1$, we obtain that $|\partial(V_1) \cap J| = \sum_{u \in V_1} |E_D(u) \cap J|$ must be even. However, since also $|E_D(x) \cap J|$ is even, this means that $|F(\{x\}) \cap J| \equiv |\partial(V_1) \cap J| + |E_D(x) \cap J| \equiv 0 \pmod{2}$, which is the desired contradiction, as J is an odd dijoin. So (iii) must be satisfied.

To prove the reverse direction, assume that (iii) is fulfilled, i.e., $n_1 \equiv n_3 \equiv 1 \pmod{2}$. We shall construct an odd dijoin of $\mathcal{D}(n_1, n_2, n_3)$. For this purpose, we choose J to be a T-join of the underlying graph where $T = V_1 \cup V_3$. We claim that this defines an odd dijoin of $\mathcal{D}(n_1, n_2, n_3)$. It is not hard to check that the directed bonds of $\mathcal{D}(n_1, n_2, n_3)$ are the cuts $E_D(v)$ for vertices $v \in V_1 \cup V_3$ and the cuts F(X) as described above, where $\emptyset \neq X \subsetneq V_2$. By the definition of a T-join, all of the directed bonds of the first type have an odd intersection with J, so it suffices to consider the bonds of the second type. Consider again the directed cut $\partial(V_1)$ in $\mathcal{D}(n_1, n_2, n_3)$. For any $\emptyset \neq X \subsetneq V_2$, we can write F(X) as the symmetric difference $F(X) = \partial(V_1) + \sum_{x \in X} E(x)$. We therefore conclude that

$$|F(X) \cap J| \equiv |\partial(V_1) \cap J| + \sum_{x \in X} \underbrace{|E(x) \cap J|}_{\text{even}} \pmod{2}$$
$$\equiv |\partial(V_1) \cap J| = \sum_{x \in V_1} \underbrace{|E(x) \cap J|}_{\text{odd}} \equiv n_1 \equiv 1 \pmod{2}.$$

This verifies that J is an odd dijoin, and completes the proof of the proposition.

We shall now use these insights to characterize minimal obstructions. For this let us first introduce new notation.

Let D be a digraph consisting of a pair h_1, h_2 of "hub vertices" and other vertices x_1, \ldots, x_n , where $n \geq 3$, such that for every $i \in [n]$, the vertex x_i has either precisely two outgoing or precisely two incoming arcs to both h_1, h_2 , and these are all arcs of D. In this case, we refer to D as a *diamond*. Alternatively, we may define an odd diamond as any digraph isomorphic to $\mathcal{D}(n_1, 2, n_3)$ for some $n_1, n_3 \in \mathbb{N}_0$ such that $n_1 + n_2 \geq 3$. Furthermore, we call any digraph isomorphic to \vec{K}_{n_1,n_2} for some $n_1, n_2 \geq 2$, a *one-direction*. We shall call both, diamonds and one-directions, *odd* if the total number of vertices of these digraphs is odd.

Lemma 5.30. All odd diamonds and all odd one-directions are minimal obstructions.

Proof. It is directly seen from Observation 5.28 and Proposition 5.29 that indeed, odd diamonds and odd one-directions do not posses an odd dijoin. Therefore it remains to show that all proper cut minors of these digraphs have odd dijoins. Because both odd diamonds and odd one-directions are transitively reduced and acyclic, the only cut minor operation applicable to them in the first step is the contraction of a single arc. By Lemma 5.18 it therefore suffices to show that for both types of digraphs, the contraction of any arc results in a digraph admitting an odd dijoin.

We first consider odd diamonds. Let $D = \mathcal{D}(n_1, 2, n_3)$ with $n_1, n_2 \ge 1$ and $n_1 + n_2$ odd, and let $e \in A(D)$ be arbitrary. In the planar directed dual graph of D, an odd bicycle with $n_1 + n_2$ vertices, there is a directed dual arc corresponding to e. It is easily seen by duality that D/e has an odd dijoin if and only if the odd bicycle of order $n_1 + n_2 \ge 3$ with a single deleted arc has an arc set intersecting every directed cycle an odd number of times. However, this is the case, because such a digraph is non-even by Theorem 5.2.

Now we consider odd one-directions. Let $D = \mathcal{D}(n_1, n_2, 0)$ with $n_1, n_2 \ge 2$ and $n_1 + n_2$ odd, and let $e = (x, y) \in A(D)$ be arbitrary. Then in the digraph D/e, define J to be the set of all arcs incident with the contraction vertex. It is easily observed that J intersects

every minimal directed cut exactly once and thus indeed, every proper cut minor has an odd dijoin. This completes the proof. $\hfill \Box$

Now we are able to prove a dual version of Theorem 5.2 and characterize the existence of odd dijoins in terms of forbidden cut minors.

Theorem 5.31. A digraph admits an odd dijoin if and only if it does neither have an odd diamond nor an odd one-direction as a cut minor.

Proof. By Lemma 5.18, a digraph has an odd dijoin if and only if it does not contain a minimal obstruction as a cut minor. Hence it suffices to show that a digraph D is a minimal obstruction if and only if it is isomorphic to an odd diamond or an odd onedirection. The fact that these digraphs indeed are minimal obstructions was proved in Lemma 5.30. So it remains to show that these are the only minimal obstructions.

Let D be an arbitrary minimal obstruction. By Proposition 5.26 there are integers $n_1, n_2, n_3 \ge 0$ such that $D \simeq \mathcal{D}(n_1, n_2, n_3)$. By the definition of a minimal obstruction, we know that D has no odd dijoin, while for every arc $e \in A(D)$, the digraph D/e is a cut minor of D and therefore has one.

We know due to Lemma 5.22 that D is weakly 2-connected. Hence, we either have $\min(n_1, n_3) = 0$, so (by symmetry) w.l.o.g. $n_3 = 0$, or $n_1, n_3 \ge 1$ and therefore $n_2 \ge 2$.

In the first case, we know by Observation 5.28 and using that D has no odd dijoin, that $n_1, n_2 \ge 2$ and $n_1 \not\equiv n_2 \pmod{2}$. So D is an odd one-direction, which verifies the claim in the case of $\min(n_1, n_3) = 0$.

Next assume that $n_1, n_3 \ge 1$ and $n_2 \ge 2$. Let $e = (x_1, x_2) \in A(D)$ with $x_i \in V_i$ for i = 1, 2 be an arbitrary arc going from the first layer V_1 to the second layer V_2 . Denote by c the vertex of D/e corresponding to the contracted arc e. Then in the digraph D/e, all arcs $\{(c, v_3) \mid v_3 \in V_3\}$ as well as all the arcs in $\{(v_1, v_2) \mid v_1 \in V_1 \setminus \{x_1\}, v_2 \in V_2 \setminus \{x_2\}\}$ admit parallel paths since $n_2 \ge 2$ and, therefore, are deletable. Successive deletion yields a cut minor D' of D/e, and thus of D, with vertex set

$$V(D') = (V_1 \setminus \{x_1\}) \cup \{c\} \cup (V_2 \setminus \{x_2\}) \cup V_3$$

and arc set

$$\{(v_1,c) \mid v_1 \in V_1 \setminus \{x_1\}\} \cup \{(c,v_2) \mid v_2 \in V_2 \setminus \{x_2\}\} \cup \{(v_2,v_3) \mid v_2 \in V_2 \setminus \{x_2\}, v_3 \in V_3\}.$$

Now after contracting all arcs of D' in the set $\{(v_1, c) | v_1 \in V_1 \setminus \{x_1\}\}$ we find that D', and hence D, has a proper cut minor isomorphic to $\mathcal{D}(1, n_2 - 1, n_3)$ with corresponding layers $\{c\}, V_2 \setminus \{x_2\}$ and V_3 .

Applying a symmetric argument (starting by contracting an arc going from V_2 to V_3), we find that D also has a proper cut minor isomorphic to $\mathcal{D}(n_1, n_2 - 1, 1)$.

Using these insights, we now show that $n_2 = 2$ holds. Suppose for a contradiction that $n_2 \geq 3$ holds. Assume first that $n_2 \geq 4$, and therefore $n_2 - 1 \geq 3$. Using statement (iii) of Proposition 5.29 and that $\mathcal{D}(1, n_2 - 1, n_3)$ and $\mathcal{D}(n_1, n_2 - 1, 1)$ both have odd dijoins, we must have $n_1 \equiv n_3 \equiv 1 \pmod{2}$. In the case that $n_2 = 3$, we similarly observe from statement (ii) of Proposition 5.29 with the digraphs $\mathcal{D}(1, 2, n_3)$ and $\mathcal{D}(n_1, 2, 1)$ that both n_1 and n_3 must be odd.

Now using statement (iii) of Proposition 5.29 with the digraph $D \simeq \mathcal{D}(n_1, n_2, n_3)$ we can conclude that D must admit an odd dijoin as well, a contradiction.

Hence, we must have $n_2 = 2$. Using again statement (ii) of Proposition 5.29 with $D \simeq \mathcal{D}(n_1, 2, n_3)$, we get that $n_1 + n_3$ must be odd. Therefore D is isomorphic to an odd diamond with $2 + n_1 + n_3$ many vertices. This concludes the proof of the theorem. \Box

We are now ready to give the proof of Theorem 5.5.

Proof of Theorem 5.5. Let \overline{M} be an oriented bond matroid, and let D be a digraph such that $\overline{M} \simeq M^*(D)$. Let us first note that by definition, \overline{M} is non-even if and only if D has an odd dijoin. Hence, for the equivalence claimed in this theorem it suffices to show that D has an odd dijoin if and only if $M^*(D)$ does not have a GB-minor isomorphic to $\overline{K}_{m,n}$ for $m, n \ge 2$ such that m + n is odd. Suppose first that D has an odd dijoin and $M^*(D)$ is non-even. Then by Proposition 5.6, every GB-minor of $M^*(D)$ is non-even as well, and hence, no such minor can equal $M^*(\overline{K}_{m,n})$ for any $m, n \ge 2$ with m + n is odd, since $\overline{K}_{m,n}$ does not have an odd dijoin for any such m and n by Lemma 5.30. This proves the first implication of the equivalence.

Conversely, let us suppose that $M^*(D)$ does not have a GB-minor isomorphic to $M^*(\vec{K}_{m,n})$ for any $m, n \geq 2$ such that m + n is odd. We shall show that D admits an odd dijoin. For this we use Theorem 5.31 and verify that D has neither an odd diamond nor an odd one-direction as a cut minor. This however follows directly from the fact that the bond-matroid induced by any odd diamond of order n is isomorphic to $M^*(\vec{K}_{2,n-2})$ as well as the easy observation that if D' is a cut minor of D, then $M^*(D')$ is a GB-minor of $M^*(D)$. This finishes the proof of the claimed equivalence.

5.5 Conclusion

For every odd $k \geq 3$ it holds that $M(\vec{C}_k) \simeq M^*(\vec{K}_{k,2}) \simeq M^*(\vec{K}_{2,k})$, and hence, the list of smallest excluded GB-minors characterising non-evenness for cographic oriented matroids strictly extends the list for graphic ones. This is quite surprising and was not expected when we initiated our research on the subject.

Seymour [Sey80] has proved a theorem about generating the class of regular matroids, showing that every regular matroid can be built up from graphic matroids, bond matroids and a certain 10-element matroid R_{10} by certain sum operations. The matroid R_{10} is regular, but neither graphic nor cographic. It is given by the following totally unimodular representing matrix:

$$R_{10} = M \begin{bmatrix} 1 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & -1 \end{bmatrix}.$$

Seymour introduced three different kinds of sum operation which join two regular matroids M_1 and M_2 whose element sets are either disjoint (1-sum), intersect in a single non-loop element (2-sum) or in a common 3-circuit (3-sum) into a bigger regular matroid $M_1\Delta M_2$ (for a precise definition of these operations we refer to the introduction of [Sey80]).

Theorem 5.32 ([Sey80]). Every regular matroid can be built up from graphic matroids, bond matroids and R_{10} by repeatedly applying 1-sums, 2-sums and 3-sums.

This theorem shows that graphic matroids, bond matroids and R_{10} constitute the most important building blocks of regular matroids. Using a brute force implementation, we checked by computer that every orientation of R_{10} containing no $M^*(\vec{K}_{m,n})$ as a GB-minor for any $m, n \geq 2$ such that m + n is odd, is already non-even. We therefore expect the total list of forbidden minors for all non-even oriented matroids to not be larger than the union of the forbidden minors for graphic (Theorem 5.4) and cographic (Theorem 5.5) non-even oriented matroids. In other words, we conjecture the following. **Conjecture 5.1.** A regular oriented matroid M is non-even if and only if none of its *GB*-minors is isomorphic to $M^*(\vec{K}_{m,n})$ for some $m, n \ge 2$ such that m + n is odd.

A natural way of working on this conjecture would be to try and show that a smallest counterexample is not decomposable as the 1-, 2- or 3-sum of two smaller regular matroids.

Apart from the obvious open problem of resolving the computational complexity of the even circuit problem (Problem 5.2) for regular oriented matroids in general, already resolving the case of bond matroids would be interesting.

Problem 5.3. Is there a polynomially bounded algorithm that, given as input a digraph D, decides whether or not D contains a directed bond of even size? Equivalently, is there a polynomially bounded recognition algorithm for digraphs admitting an odd dijoin?

Conclusively, given our characterization of digraphs admitting an odd dijoin in terms of forbidden cut minors, the following question naturally comes up.

Problem 5.4. Let F be a fixed digraph. Is there a polynomially bounded algorithm that, given as input a digraph D, decides whether D contains a cut minor isomorphic to F?

Part II

Dichromatic Number

General Comments

Given a digraph D, an *acyclic k-coloring* of D is a mapping $c: V(D) \to [k]$ such that for every color $i \in [k]$, the *color class* $c^{-1}(i) \subseteq V(D)$ induces an acyclic subdigraph of D. The *dichromatic number* $\vec{\chi}(D)$ of D is defined as the smallest $k \in \mathbb{N}$ for which an acyclic kcoloring of D exists. Introduced in 1980 by Erdős and Neumann-Lara [Erd80, NL82], this parameter was rediscovered and popularized by Mohar and Bokal et al. [Moh03, BFJ⁺04], and since then has received further attention, see [ACH⁺19, AH15, BHL18, HLTW19, HM17, LM17, Moh16] for a selection (of only a small fraction) of recent results.

The investigation of acyclic colorings of digraphs has constituted a major part of the research I have conducted during my PhD. In this second part of the thesis, I collect several new results I obtained on this topic during the last years. I hope that they contribute to a better understanding of the *digraph chromatic number* $\vec{\chi}$.

Many of the previous results on the dichromatic number have shown that it shares qualitative properties with the chromatic number of undirected graphs, and have thereby established it as a natural directed analogue of the chromatic number.

In the same spirit, my research was often motivated by the goal of extending classical results about the chromatic number to the directed setting. Since much research in graph theory has concentrated on substructures contained in graphs with large chromatic number, a focus of the research presented in this chapter has been to relate the dichromatic number to the existence of natural substructures in digraphs such as minors, subdivisions and induced subgraphs.

Chapter 6

Excluding Strong Minors

6.1 Introduction

All graphs and digraphs considered in the chapter are simple.

In this first chapter we are concerned with the existence of *strong minors* in digraphs with a given dichromatic number.

Hadwiger's conjecture (Conjecture 0.1) for undirected graphs seeks to force minors of complete graphs through the chromatic number, claiming that every t-chromatic graph G contains K_t as a minor. However, this conjecture remains widely open for any value $t \geq 7$, and hence it is natural to aim for asymptotic versions of the conjecture first. For $t \in \mathbb{N}$ let us denote by $m_{\chi}(t)$ the least integer for which it is true that every graph with chromatic number at least $m_{\chi}(t)$ contains a K_t -minor. A commonly expressed relaxation of Hadwiger's conjecture, called *Linear Hadwiger Conjecture*, states that $m_{\chi}(t) = O(t)$ grows linearly. Eventhough this relaxed version of the conjecture remains open, through the history of research on the problem one has come quite close to a linear bound. For many years, the best general upper bound on $m_{\chi}(t)$ was due to Kostochka [Kos84] and Thomason [Tho84], who independently proved that every graph of average degree at least $O(t\sqrt{\log t})$ contains a K_t -minor, implying via degeneracy coloring that $m_{\chi}(t) = O(t\sqrt{\log t})$. Recently, however, there has been progress. First Norine, Postle, and Song [NPS19] showed that $m_{\chi}(t) = O\left(t(\log t)^{\beta}\right)$ (for any $\beta > \frac{1}{4}$), and then this was further improved by Postle [Pos20] to give the following state of the art-bound.

Theorem 6.1 (cf. [Pos20]). There exists an absolute constant C > 0 such that for $t \ge 3$:

$$m_{\chi}(t) \leq Ct(\log \log t)^6.$$

Hadwiger's famous conjecture has influenced many researchers and different variations of it have been studied in various frameworks, one of which is directed graphs. Axenovich, Girão, Snyder, and Weber [AGSW20] recently considered the analogue of Hadwiger's problem for digraphs, where the chromatic number is replaced by the dichromatic number and undirected minors by strong minors. Concretely, they raised the following problem.

Problem 6.1. For a given integer $t \ge 1$, what is the smallest integer $sm_{\vec{\chi}}(t) \ge 1$ such that every digraph D with $\vec{\chi}(D) \ge sm_{\vec{\chi}}(t)$ contains \vec{K}_t as a strong minor?

In the first place, it is not clear why such a number $sm_{\vec{\chi}}(t)$ should even exist for every $t \in \mathbb{N}$. Axenovich, Girão, Snyder, and Weber [AGSW20] showed that $sm_{\vec{\chi}}(t)$ indeed exists for every $t \geq 1$ and proved the bounds

$$t+1 \le sm_{\vec{\chi}}(t) \le t4^t$$

for every $t \in \mathbb{N}$. They expressed that they expect the exponential upper bound $t4^t$ to be quite far from the true value of $sm_{\vec{\chi}}(t)$ and that it would be very interesting to find an improved upper bound.

As the main result of this chapter, we indeed find a substantially improved upper bound by reducing the problem to the undirected setting as follows.

Theorem 6.2. For every $t \ge 1$ we have

$$m_{\chi}(t) \le sm_{\vec{\chi}}(t) \le 2m_{\chi}(t) - 1.$$

By combining Theorem 6.2 with Theorem 6.1 we get that

$$sm_{\vec{\chi}}(t) = \Theta(m_{\chi}(t)) = O\left(t(\log\log t)^6\right).$$

The lower bound in Theorem 6.2 can be quite easily observed as follows: For every graph G with $\chi(G) \geq sm_{\vec{\chi}}(t)$, we have¹ $\vec{\chi}(\vec{G}) = \chi(G) \geq sm_{\vec{\chi}}(t)$, and hence \vec{G} contains a strong \vec{K}_t -minor. Since every strongly connected branch set of this strong minor in \vec{G} corresponds to a connected subgraph of G, we can see by taking the same branch sets that G contains a K_t -minor. Hence, $m_{\chi}(t) \leq sm_{\vec{\chi}}(t)$.

Theorem 6.2 also has direct consequences for the other two minor notions in digraphs considered in this thesis, namely topological minors and butterfly-minors. These consequences can be found in Chapters 7 and 8, respectively.

6.2 Proof of Theorem 6.2

The proof of Theorem 6.2 is based on the following result.

Theorem 6.3. For every digraph D there is an undirected graph G such that

- 1. D is a strong \vec{G} -minor-model, and
- 2. $\vec{\chi}(D) \leq 2\chi(G)$.

Proof. To start with, let us first fix a partition X_1, X_2, \ldots, X_m of V(D) for some $m \in \mathbb{N}$ such that for every $i \in \{1, 2, \ldots, m\}$ the set X_i is an inclusion-wise maximal subset of V(D) such that $D[X_i]$ is strongly connected, $\vec{\chi}(D[X_i]) \leq 2$ and $X_i \cap (X_1 \cup \cdots \cup X_{i-1}) = \emptyset$.

Note that the X_i 's are well-defined since the one vertex-digraph is strongly connected and 2-colorable. Now we define G to be the undirected simple graph with vertex set $\{X_1, \ldots, X_m\}$ and $X_i X_j \in E(G)$ if and only if there is an arc in D starting in X_i and ending in X_j , as well as an arc starting in X_j and ending in X_i . Then, by definition (cf. Definition 1.2 from Chapter 1), D is a strong $\overset{\leftrightarrow}{G}$ -minor-model, as one can simply take X_1, X_2, \ldots, X_m as the branch sets.

Therefore, what remains to prove is property (2). For this let us assume that $\chi(G) = k$ and fix a proper coloring $f_G : V(G) \to \{c_1, c_2, ..., c_k\}$ of G. Now, for every i take an arbitrary acyclic two-coloring of $D[X_i]$ (which exists by assumption) with colors $\{c'_i, c''_i\}$. The rest of the proof is about showing that by putting these colorings together we obtain an acyclic coloring f_D of D with the 2k colors $\{c'_1, c''_1, c'_2, c''_2, ..., c'_k, c''_k\}$.

Assume for contradiction that this is not the case, and there is a directed cycle C in D which is monochromatic. We may, without loss of generality, assume that C is a shortest

¹For every undirected graph G we have $\vec{\chi}(\vec{G}) = \chi(G)$, since every adjacent pair of vertices in G induces a digon in \vec{G} and hence adjacent vertices must be colored distinctly in every acyclic coloring.
such cycle, in particular, it is an induced cycle². Let i_0 be the smallest index for which C contains a vertex from X_{i_0} . Note that, in particular, $V(C) \subseteq V(D) \setminus (X_1 \cup \cdots \cup X_{i_0-1})$ and, as f_D is a proper coloring on $D[X_{i_0}]$, the cycle C cannot be fully contained in X_{i_0} . It follows that C contains a subsequence $u, w_1, \ldots, w_\ell, v$ of consecutive vertices on C such that $(u, w_1), (w_1, w_2), \ldots, (w_\ell, v) \in A(C), u, v \in X_{i_0}$ (possibly u = v), as well as $w_1, \ldots, w_\ell \in X_{i_0+1} \cup \cdots \cup X_m$, and $\ell > 0$.

Let $s \in \{1, \ldots, \ell\}$ be the smallest index such that w_s has an out-neighbour in X_{i_0} , and denote this out-neighbor by $x \in X_{i_0}$.

We claim that w_s has no in-neighbor in D that is contained in X_{i_0} . Suppose towards a contradiction that there exists $y \in X_{i_0}$ such that $(y, w_s) \in A(D)$. Let $j > i_0$ be such that $w_s \in X_j$. Then, because of the arcs $(y, w_s), (w_s, x) \in A(D)$, we have $X_{i_0}X_j \in E(G)$ and hence $f_G(X_{i_0}) \neq f_G(X_j)$. This in turn implies that $f_D(u) \neq f_D(w_s)$ and $f_D(v) \neq f_D(w_s)$ which contradicts the monochromaticity of C. Hence, we may assume that w_s has no in-neighbor contained in X_{i_0} . In particular, this implies $s \geq 2$. Let us now consider

$$X := X_{i_0} \cup \{w_1, \dots, w_s\} \subseteq V(D) \setminus (X_1 \cup \dots \cup X_{i_0-1})$$

The digraph D[X] is strongly connected, as $D[X_{i_0}]$ is so and u, w_1, \ldots, w_s, x induce a directed path (or cycle in case of u = x) starting and ending in X_{i_0} . Moreover, any extension of an acyclic $\{1, 2\}$ -coloring of $D[X_{i_0}]$ to a $\{1, 2\}$ -coloring of D[X] where w_1, \ldots, w_{s-1} receive color 1 and w_s receives color 2 is acyclic. Indeed, by the definition of s, there are no arcs starting in $\{w_1, \ldots, w_{s-1}\}$ and ending in X_{i_0} , and by the inducedness of C the subdigraph of D induced by $\{w_1, \ldots, w_{s-1}\}$ is a directed path and therefore acyclic. Since also w_s has no in-neighbors in X_{i_0} , any directed cycle in D[X] is either fully contained in $D[X_{i_0}]$, or contains both w_s and at least one vertex in $\{w_1, \ldots, w_{s-1}\}$. In any case, it is not monochromatic. However, the existence of the set X then contradicts with the maximality of X_{i_0} , which completes the proof.

Now we can easily deduce Theorem 6.2 from Theorem 6.3.

Proof of Theorem 6.2. Let D be a digraph with $\vec{\chi}(D) \geq 2m_{\chi}(t) - 1$. By Theorem 6.3 there exists an undirected graph G such that $\vec{\chi}(D) \leq 2\chi(G)$ and $D \succeq \vec{G}$. This implies that $\chi(G) \geq m_{\chi}(t)$, and hence G contains a K_t -minor. Taking the same branch sets as for the K_t -minor in G also in \vec{G} shows that $\vec{G} \succeq_s \vec{K}_t$ (since the biorientation of every connected graph is a strongly connected digraph). By transitivity we obtain $D \succeq_s \vec{K}_t$.

Since D was arbitrarily chosen such that $\vec{\chi}(D) \ge 2m_{\chi}(t) - 1$, this proves that we have $sm_{\vec{\chi}}(t) \le 2m_{\chi}(t) - 1$, as required.

6.3 Conclusion

In this chapter we showed that $m_{\chi}(t) \leq sm_{\chi}(t) \leq 2m_{\chi}(t) - 1$ for any $t \geq 1$.

Therefore, $sm_{\tilde{\chi}}(t) = \Theta(m_{\chi}(t))$, and the question about the asymptotics of $sm_{\tilde{\chi}}(t)$ raised by Axenovich, Girão, Snyder and Weber, reduces to the well-studied undirected version of the problem, namely Hadwiger's conjecture. Also, as Hadwiger's conjecture is known to be true for small values, for $2 \le t \le 6$ we have

$$t+1 \le sm_{\vec{x}}(t) \le 2t-1$$

We believe that the upper bound should not be tight. To support this intuition, let us mention that a more careful analysis of our proof of Theorem 6.2 yields the stronger

²A cycle in a digraph is called induced if it forms an induced subdigraph.

statement that any digraph D with $\vec{\chi}(D) \geq 2m_{\chi}(t)-1$ contains a strong \vec{K}_t -minor model in which between any two branch sets, there are at least two arcs spanned in both directions. Under the assumption that Hadwiger's conjecture is true, the bound 2t-1 for this stronger property would be sharp, as shown by \vec{K}_{2t-2} . This indicates that our proof should not be expected to give a tight bound for the problem of forcing a strong \vec{K}_t -minor. Instead it seems plausible that $sm_{\vec{\chi}}(t) = t+1$ for any $t \geq 3$.

Conjecture 6.1. Every digraph D with $\vec{\chi}(D) \ge t + 1$ contains \vec{K}_t as a strong minor.

Already resolving the first open case t = 3 of this problem would be quite interesting.

A notion of *list colorings* of digraphs was introduced and studied by Bensmail et al. [BHL18]. A digraph is said to be k-choosable if for every assignment of color-lists of size k to its vertices there is a choice of colors from the lists resulting in an acyclic coloring. Interestingly, we do not know the answer to the following intriguing problem.

Problem 6.2. Does there exist $f : \mathbb{N} \to \mathbb{N}$ such that every digraph D excluding \vec{K}_t as a strong minor is f(t)-choosable?

Chapter 7

Excluding Topological Minors

7.1 Introduction

All graphs and digraphs considered in this chapter are simple.

The chromatic number is one of the fundamental graph parameters, and is well-known to be intractable. Therefore, meaningful sufficient and necessary conditions for it to be large are of high interest. In the previous chapter, we have discussed the famous conjecture of Hadwiger, relating chromatic number and the containment of graph minors. Hadwiger's conjecture claims that any t-chromatic graph contains a K_t -minor. An even stronger conclusion than that was suggested in a conjecture attributed possibly falsely¹ to Hajós, who conjectured that every t-chromatic graph contains a subdivision of K_t , that is, a graph which can be obtained from K_t by replacing its edges with pairwise internally vertexdisjoint paths connecting their original endpoints. Hajós' conjecture is easily verified for $t \leq 3$, and Dirac [Dir52] proved the case t = 4.

Theorem 7.1 ([Dir52]). Every graph G with $\chi(G) \ge 4$ contains a K₄-subdivision.

While the cases t = 5, 6 of Hajós' conjecture remain open (and the case t = 5 would represent a strengthening of the 4CT), it was disproved for all values $t \ge 7$ by Catlin [Cat79], who constructed explicit counterexamples, i.e., graphs with chromatic number k which contain no K_t -subdivision (see also [Tho05] for many more explicit constructions of counterexamples). An even more devastating blow to the conjecture was delivered by Erdős and Fajtlowicz [EF81], who showed that *almost all* graphs on $\Theta(t^2)$ vertices do *not* contain a K_t -subdivision, even though their chromatic number is $\Omega(t^2/\log t)$.

On the positive side, it turned out that large enough chromatic number does in fact necessitate the existence of a K_t -subdivision. As a matter of fact, the following classical result established that even large density is sufficient.

Theorem 7.2 (Bollobás and Thomason [BT98], Komlós and Szemerédi [KS96]). There exists an absolute constant C > 0 such that for every $t \in \mathbb{N}$, every graph G with minimum degree at least Ct^2 contains a subdivision of K_t .

Since (via degeneracy coloring) every graph G contains a subgraph of minimum degree at least $\chi(G) - 1$, one can deduce from Theorem 7.2 that having chromatic number larger than Ct^2 (for some absolute constant C) is sufficient to guarantee a K_t -subdivision. For $t \in \mathbb{N}$, let f(t) be the smallest integer such that every graph with chromatic number at least f(t) contains a K_t -subdivision. Theorem 7.2 then implies a quadratic upper bound

¹See the discussion in the introduction of [Tho05].

 $f(t) = O(t^2)$, while the result of Erdős and Fajtlowicz [EF81] establishes a lower bound of $f(t) = \Omega(t^2/\log t)$. These remain the best known bounds on f(t); Fox et al. [FLS13] conjectured that the truth lies with the lower bound.

The upshot of the above discussion is that a subdivision of any given graph is contained in any graph of sufficiently large chromatic number. In this chapter we investigate this phenomenon in the realm of directed graphs; we ask in what form, and to what extent, it holds. This follows up previous work of Aboulker et al. [ACH⁺19] on the same topic. As explained in Chapter 1, the notion of subdivision extends naturally to directed graphs. It is less obvious, however, what would be the most suitable digraph coloring concept, which would provide a rich family of forcible digraph subdivisions.

The chromatic number $\chi(D)$ of a digraph D is defined as the chromatic number of U(D). The fact that any graph, however high its chromatic number is, can be oriented acyclically and hence avoid containing any directed cycle, already hints that $\chi(D)$ being large might only have limited impact on digraph subdivision containment. In fact, as was noted by Aboulker, Cohen, Havet, Lochet, Moura, and Thomassé [ACH⁺19], the family of digraphs F which can be forced as a subdivision by high chromatic number is very limited: it consists of the orientations of forests, see [Bur80] and [CHLN18], respectively, for the positive and negative directions of this result. As a consequence, we see that high chromatic number of the underlying graph is not even strong enough to force the subdivision of any particular orientation of a cycle.

In that sense it is natural to investigate other digraph coloring parameters, which, in contrast to the chromatic number, take into account the direction of edges. Here we investigate the *dichromatic number* as such a parameter.

Previously, Aboulker et al. [ACH⁺19] initiated the study of the existence of various subdivisions in digraphs of large dichromatic number.

In one of their main results, they show that a subdivision of any given digraph is contained in digraphs of sufficiently large dichromatic number. This is quite remarkable, as (to the best of our knowledge) no other natural digraph parameter is known which is capable of forcing subdivisions of arbitrary digraphs.

Theorem 7.3 ([ACH⁺19], Theorem 32). Let F be a digraph with n vertices and m arcs. Then every digraph D with $\vec{\chi}(D) \ge 4^m(n-1) + 1$ contains a subdivision of F.

Following the terminology introduced in Chapter 2, for a digraph F we denote by $\operatorname{mader}_{\vec{\chi}}(F)$ the smallest integer $k \geq 1$ such that every digraph D with $\vec{\chi}(D) \geq k$ contains a subdivision of F. We call $\operatorname{mader}_{\vec{\chi}}(F)$ the *(dichromatic) Mader number of* F.

The problem of obtaining a polynomial bound (in terms of the number of vertices and arcs of F) on the Mader number remains open, and seems quite challenging. One reason for the increased difficulty compared to the undirected case is that there is no analogue of Theorem 7.2 for directed graphs. In fact, as explained in Chapter 2, it follows from a result of Thomassen [Tho85b] that there exist digraphs of arbitrarily high minimum outand in-degree, which do not even contain a subdivision of \vec{K}_3 , the bioriented triangle. Consequently, entirely new methods have to be developed to force clique-subdivisions in digraphs of large dichromatic number, since any methods for addressing this problem must differ substantially from the established density-based ideas used in the undirected theory.

In light of the difficulty of improving the bound in Theorem 7.3 in general, it is natural to vie for obtaining better upper bounds for special classes of digraphs. One appealing conjecture in this vein was raised by Aboulker et al. [ACH⁺19].

Conjecture 7.1 ([ACH⁺19], Conjecture 39). For every orientation C of C_{ℓ} , it holds that $mader_{\vec{\chi}}(C) = \ell$.

7.1. INTRODUCTION

In [ACH⁺19], Conjecture 7.1 is proved for directed cycles and an upper bound of $2\ell - 1$ is established for arbitrary orientations.

7.1.1 Our Results

Note that the Mader number of every digraph F is at least the number of its vertices. Indeed, the complete digraph of order v(F) - 1 has dichromatic number v(F) - 1, but does not have enough vertices to host a subdivision of F. Hence Conjecture 7.1 states that, in a sense, the Mader number of orientations of cycles is as small as it could be.

In this chapter we resolve Conjecture 7.1 and go on to study the more general question: for which digraphs F does it hold that $\operatorname{mader}_{\vec{\chi}}(F) = v(F)$? In the first main result of the chapter we prove that this equality holds for a large class of digraphs, which includes orientations of cactus graphs² (and hence all orientations of cycles), as well as all bioriented forests. This class of digraphs, a member of which we refer to as $octus^3$, is defined inductively as follows.

Definition 7.1. The class of octi digraphs is defined as follows.

- K_1 is an octus.
- Let F be an octus, and let $v_0 \in V(F)$. Let $P = v_1, \ldots, v_k$, $k \ge 1$, be an orientation of a path which is disjoint from V(F). Let F^* be obtained from F by adding the path P, both arcs $(v_0, v_1), (v_1, v_0)$, and exactly one of the arcs $(v_0, v_k), (v_k, v_0)$. Then F^* is also an octus.
- If F is an octus then every subdigraph of F is also an octus.



Figure 7.1: An example of an octus.

²Cactus graphs are usually defined as the graphs which do not contain a pair of cycles sharing at least two vertices (or, equivalently, as the graphs which do not contain $K_4 - e$ as a minor).

 $^{^{3}}$ The name alludes to the fact that every orientation of a cactus graph is an octus.

We should warn, however, that the class of octi is strictly larger than the class of orientations of cacti, as is explained following Theorem 7.4.

We note that the path P in the second item of Definition 7.1 is allowed to consist of a single vertex, which corresponds to attaching a digon to F at v_0 . The operation described in Item 2 of Definition 7.1 will be called *ear addition*. Our first main result is as follows:

Theorem 7.4. For every octus F, we have $mader_{\vec{\chi}}(F) = v(F)$.

This theorem has a couple of immediate consequences, each of which extends results of [ACH⁺19]. It is not difficult to see that orientations of cacti are precisely the octi which have no digons. Therefore, we have the following:

Corollary 7.5. For every orientation F of a cactus, we have $mader_{\vec{x}}(F) = v(F)$.

Since every cycle is a cactus, Corollary 7.5 immediately implies Conjecture 7.1.

Another immediate corollary of Theorem 7.4 is concerned with biorientations of forests. Every bioriented tree can be obtained from K_1 by a sequence of ear additions where, at each step, we add a new vertex and connect it by a digon to one of the vertices of the existing digraph. Hence, every bioriented of forest is an octus, and we have the following:

Corollary 7.6. If T is an (undirected) forest, then $mader_{\vec{x}}(\vec{T}) = v(\vec{T})$.

Corollary 7.6 strengthens another result from [ACH⁺19], where the conclusion was shown to hold for every orientation of a forest.

Next we discuss digraphs on a small number of vertices. The smallest digraph not covered by Theorem 7.4 is the bioriented triangle minus an edge. It turns out that this digraph, too, has the property that its Mader number equals its number of vertices.

Proposition 7.7. $mader_{\vec{x}}(\vec{K}_3 - e) = 3.$

In the second main result of this chapter, we show that the Mader number of every 4-vertex tournament is 4.

Theorem 7.8. For every orientation K of K_4 , we have that $mader_{\vec{x}}(K) = 4$.

Theorem 7.8 is a strict extension to the directed setting of Dirac's theorem on K_4 -subdivisions (namely, Theorem 7.1). In fact, Theorem 7.1 can be easily derived from Theorem 7.8 as follows. First, observe that $\vec{\chi}(\vec{G}) = \chi(G)$ for every graph G. Now, if G is an undirected graph with $\chi(G) \geq 4$, then by Theorem 7.8, \vec{G} contains a subdivision of any orientation of K_4 , which translates to a K_4 -subdivision in G.

In the third and last main result of this chapter, we prove a linear upper bound on the Mader number for another class of relatively sparse digraphs which we call $subcubic^4$. A digraph F is called *subcubic* if $\Delta(F) \leq 3$ and $\Delta^+(F), \Delta^-(F) \leq 2$.

Theorem 7.9. If F is a subcubic digraph, then $mader_{\vec{x}}(F) \leq 22 \cdot v(F)$.

The rest of the chapter is organized as follows. After establishing some preliminary results in Section 7.2, we prove Theorem 7.4 in Section 7.3. Section 7.4 is devoted to proving Theorem 7.8. In Section 7.5 we derive Theorem 7.9 from the results of Chapter 6. Finally, in Section 7.6 we conclude with a discussion of Mader numbers of biorientations of complete digraphs and cycles, give the proof of Proposition 7.7, and pose some open problems. A main focus of Section 7.6 is on digraphs which we call *Mader-perfect*; these are digraphs F with the property that every subdigraph F' of F satisfies mader $\vec{\chi}(F') = v(F')$. We propose the further study of these digraphs and establish some preliminary results.

⁴This name is a slight abuse of notation, since we do not allow arbitrary orientations of subcubic graphs. In [HMM18] subcubic digraphs are instead called *digraphs without big vertices*.

7.2 Preliminaries

In this section we gather a number of definitions, observations and auxiliary results about the dichromatic number and about subdivisions in digraphs which will be used in the course of the chapter. We start by observing that $mader_{\vec{\chi}}$ is subadditive with respect to taking disjoint unions.

Observation 7.10. Let F be the disjoint union of two digraphs F_1, F_2 . Then

 $mader_{\vec{x}}(F) \leq mader_{\vec{x}}(F_1) + mader_{\vec{x}}(F_2).$

Proof. For convenience, put $k_i := \text{mader}_{\vec{\chi}}(F_i)$, i = 1, 2. Let D be a digraph with dichromatic number at least $k_1 + k_2$. Let $A_1 \subseteq V(D)$ be such that $\vec{\chi}(D[A_1]) = k_1$ (such a set A_1 can be obtained by repeatedly deleting vertices as long as the dichromatic number of the current digraph is strictly larger than k_1). Put $A_2 := V(D) \setminus A_1$. Then $\vec{\chi}(D[A_2]) \ge k_2$, for otherwise one could color D with less than $k_1 + k_2$ colors. By our choice of k_i , we get that $D[A_i]$ contains a subdivision of F_i for each i = 1, 2. It follows that D contains a subdivision of F, as required.

Let $k \in \mathbb{N}$. A digraph D is called *k*-dicritical, if $\vec{\chi}(D) = k$, but $\vec{\chi}(D') < k$ for all proper subdigraphs $D' \subsetneq D$.

Lemma 7.11. Let D be k-dicritical. Then $\delta^+(D), \delta^-(D) \ge k-1$.

Proof. Since the reversal of all arcs preserves the k-dicriticality of D, it suffices to show that $\delta^+(D) \ge k - 1$. Suppose towards a contradiction that there exists some $v \in V(D)$ such that $d^+(v) < k - 1$. By assumption, D - v admits an acyclic coloring with color-set $\{1, \ldots, k - 1\}$. We can extend this to a (k - 1)-coloring of D by assigning to v a color in $\{1, \ldots, k - 1\}$ that does not appear on $N^+(v)$. Then the resulting coloring is an acylic (k - 1)-coloring of D (since no monochromatic directed cycle can pass through v), in contradiction to our assumption that $\vec{\chi}(D) = k$.

Lemma 7.12. Let D be k-dicritical. Then D is strongly connected.

Proof. Assume, for the sake of contradiction, that D is not strongly connected. Then there is a partition $V(D) = A \cup B$ such that A and B are non-empty and there are no arcs going from B to A. Since D is k-dicritical, both D[A] and D[B] have an acyclic (k-1)-coloring. But putting these colorings together is an acyclic (k-1)-coloring of D, since D contains no directed cycles which intersect both A and B. Thus, we have arrived at a contradiction to $\vec{\chi}(D) = k$.

We will further need the following two deep results by Mader on so-called non-critical vertices and on subdivisions in digraphs of sufficiently large out-degree.

Theorem 7.13 ([Mad91], see also Section 7.11 in [BJG08]). Let $k \in \mathbb{N}$, and let D be a strongly k-vertex-connected digraph with $\delta^+(D), \delta^-(D) \ge 2k$. Then there is $v \in V(D)$ such that D - v is (also) strongly k-vertex-connected.

Theorem 7.14 ([Mad96]). Let D be a digraph such that $\delta^+(D) \ge 3$. Then D contains a subdivision of \vec{K}_4 , the transitive tournament of order 4.

7.3 Oriented Cacti and Bioriented Forests

In this section we prove Theorem 7.4. The main step in the proof consists of showing that if F^* is a digraph obtained from a digraph F via ear addition (i.e., the operation described in the second item of Definition 7.1), then $\operatorname{mader}_{\vec{\chi}}(F^*) \leq \operatorname{mader}_{\vec{\chi}}(F) + k$ where k is the number of newly added vertices. This is done in the following theorem.

Theorem 7.15. Let F be a digraph and let $v_0 \in V(F)$. Let $P = v_1, \ldots, v_k$ be an orientation of a path disjoint from V(F). Let F^* be a digraph obtained from F by adding the path P, both arcs $(v_0, v_1), (v_1, v_0)$, and exactly one of the arcs $(v_0, v_k), (v_k, v_0)$. Then $mader_{\vec{x}}(F^*) \leq mader_{\vec{x}}(F) + k$.

To prove Theorem 7.15 we will need the following useful lemma, which describes a generalization of the idea of Kempe-switches to directed graphs.

Lemma 7.16. Let D be a digraph, $k \in \mathbb{N}$, and let $c : V(D) \to \{1, \ldots, k\}$ be an acyclic coloring of D. Let $i \neq j \in \{1, \ldots, k\}$, $D_{i,j} := D[c^{-1}(\{i, j\})]$, and let $X \subseteq c^{-1}(\{i, j\})$ be the vertex set of a strong component of $D_{i,j}$. Then the coloring $c' : V(D) \to \{1, \ldots, k\}$, defined by

$$c'(x) := \begin{cases} c(x) & \text{if } x \in V(D) \setminus X, \\ j & \text{if } x \in X \cap c^{-1}(i), \\ i & \text{if } x \in X \cap c^{-1}(j) \end{cases}$$

is an acyclic coloring of D as well.

Proof. Suppose towards a contradiction that there is a directed cycle C in D which is monochromatic under c'. If $V(C) \cap X = \emptyset$, then c and c' agree on V(C), contradicting our assumption that c is an acyclic coloring of D. Therefore $V(C) \cap X \neq \emptyset$. Since c'has only colors i or j on X, we find that C is monochromatic under c' either in color ior j. This means that $V(C) \subseteq (c')^{-1}(\{i, j\}) = c^{-1}(\{i, j\})$ according to the definition of c'. Hence, C is a directed cycle in $D_{i,j}$, and since X is a strong component of $D_{i,j}$, we conclude $V(C) \subseteq X$. By the definition of c' the colors i and j are switched in X, so Cmust have been monochromatic under c in color j or i. This contradicts to the fact that the coloring c of D is acyclic and concludes the proof.

Proof of Theorem 7.15. First, we argue that by symmetry, it is enough to handle the case that $(v_0, v_k) \in A(F^*)$. For a digraph D, denote by \overleftarrow{D} the digraph on the same vertex set obtained from it by reversing the orientations of all arcs, that is, with arc-set $A(\overleftarrow{D}) := \{(x, y) \mid (y, x) \in A(D)\}$. For all digraphs D and F, we have $\overrightarrow{\chi}(D) = \overrightarrow{\chi}(\overleftarrow{D})$ and D contains a subdivision of F if and only if \overleftarrow{D} contains a subdivision of \overleftarrow{F} . As a consequence, we have mader $\overrightarrow{\chi}(F) = \text{mader}_{\overrightarrow{\chi}}(\overleftarrow{F})$ for every digraph F. Therefore, the case $(v_k, v_0) \in A(F^*)$ follows from the case $(v_0, v_k) \in A(F^*)$ via this symmetry. So for the rest of the proof let us assume that $(v_0, v_k) \in A(F^*)$.

For brevity, in the following we put $M := \text{mader}_{\vec{\chi}}(F)$. Consider any given digraph D such that $\vec{\chi}(D) = M + k$. We have to show that D contains a subdivision of F^* .

Let us start by fixing an acyclic coloring $c_0 : V(D) \to \{1, 2, \ldots, M + k\}$ of D that maximizes $|c_0^{-1}(\{1, \ldots, k\})|$. Set $Y_1 := c_0^{-1}(\{1, \ldots, k\})$ and $Y_2 := c_0^{-1}(\{k+1, \ldots, M+k\})$. Note that $V(D) = Y_1 \cup Y_2$ is a partition of V(D). Since c_0 is an acyclic coloring of D with $\vec{\chi}(D)$ colors, we have $\vec{\chi}(D[Y_1]) = |\{1, \ldots, k\}| = k$ and $\vec{\chi}(D[Y_2]) = |\{k+1, \ldots, M+k\}| = M$.

From the definition of M we conclude that there exists a subgraph $S \subseteq D[Y_2]$ which is a subdivision of F. In the following, let us denote by $x_0 \in V(S) \subseteq Y_2$ the branch-vertex in this subdivision corresponding to $v_0 \in V(F)$. For each acyclic k-coloring $c: Y_1 \to \{1, \ldots, k\}$ of $D[Y_1]$, let $\mathbf{v}(c) \in \mathbb{Z}^k$ denote the vector defined by $\mathbf{v}(c)_i = |N^+(x_0) \cap c^{-1}(i)|$, for $i = 1, \ldots, k$. Let us consider the pre-order \prec on the set of acyclic $\{1, \ldots, k\}$ -colorings of $D[Y_1]$, where $c_1 \prec c_2$ iff $\mathbf{v}(c_1) <_{\text{lex}} \mathbf{v}(c_2)$. Here $<_{\text{lex}}$ denotes the lexicographical order on \mathbb{Z}^k . In the following, let $c: Y_1 \to \{1, \ldots, k\}$ denote an acyclic coloring of $D[Y_1]$ that is minimal with respect to \prec . For $i < j \in \{1, \ldots, k\}$, let $D_{i,j} := D[c^{-1}(\{i, j\})]$.

Claim 1. For every $1 \le i < j \le k$ and every vertex $x \in N^+(x_0) \cap c^{-1}(i)$, there is a vertex $y \in N^+(x_0) \cap c^{-1}(j)$ such that x and y lie in the same strong component of $D_{i,j}$.

Proof. Denote by $X \subseteq c^{-1}(\{i, j\})$ the unique strong component of $D_{i,j}$ containing x. Suppose towards a contradiction that $X \cap (N^+(x_0) \cap c^{-1}(j)) = \emptyset$. Let c' be the coloring of $D[Y_1]$ obtained from c by switching colors i and j within X. According to Lemma 7.16, c' is an acyclic coloring of $D[Y_1]$. By definition, we furthermore have $\mathbf{v}(c')_{\ell} = \mathbf{v}(c)_{\ell}$ for all $\ell \in \{1, \ldots, k\} \setminus \{i, j\}$, and since no vertex in $N^+(x_0)$ is switched from color j to color iwhile x is switched from color i to color j, we have $\mathbf{v}(c')_i < \mathbf{v}(c)_i$. However, since i < j, this means that $c' \prec c$, contradicting our minimality assumption on c. This shows that our assumption was wrong, namely there does exist a vertex $y \in X \cap (N^+(x_0) \cap c^{-1}(j))$. This yields the claim.

Claim 2. There are vertices $x_1, x_2, \ldots, x_k \in N^+(x_0) \cap Y_1$ such that

- $c(x_i) = i$, for i = 1, ..., k.
- There is a directed cycle C in D containing x_0 and x_1 such that $V(C) \setminus \{x_0\} \subseteq c^{-1}(1)$.
- For every $2 \leq i \leq k$, there exists a directed path $P_{i-1,i}$ in $D[Y_1]$ with endpoints x_{i-1}, x_i such that $V(P_{i-1,i}) \subseteq c^{-1}(\{i-1,i\})$. In addition, $P_{i-1,i}$ is directed from x_{i-1} to x_i if $(v_{i-1}, v_i) \in A(P)$, and directed from x_i to x_{i-1} if $(v_i, v_{i-1}) \in A(P)$.

Proof. We start by showing that there is a directed cycle C in D through x_0 such that $V(C) \setminus \{x_0\} \subseteq c^{-1}(1)$. Assume, towards a contradiction, that no such cycle exists, and consider the coloring $c'_0: V(D) \to \{1, \ldots, M+k\}$ defined by

$$c_0'(x) := \begin{cases} c(x), & \text{if } x \in Y_1 \\ 1, & \text{if } x = x_0 \\ c_0(x), & \text{if } x \in Y_2 \setminus \{x_0\}. \end{cases}$$

Our assumption implies that c'_0 is an acyclic coloring of D, because there is no directed cycle containing x_0 which is monochromatic under c'_0 . However, the coloring c'_0 has one more vertex in colors $\{1, ..., k\}$ than c_0 , contradicting our maximality assumption on c_0 . Therefore, a cycle C with the claimed properties exists.

Now define $x_1 \in N^+(x_0) \cap V(C)$ to be the unique out-neighbor of x_0 on C. We have $c(x_1) = 1$ since $x_1 \in V(C) \setminus \{x_0\}$. We now successively define vertices x_2, \ldots, x_k as follows: for $i = 2, 3, \ldots, k$, define the vertex x_i to be a vertex in $N^+(x_0) \cap c^{-1}(i)$ chosen such that x_{i-1} and x_i lie in the same strong component of $D_{i-1,i}$. Notice that such a choice is possible by Claim 1.

The first and second items of the claim follow directly from our choice of the vertices x_1, \ldots, x_k . For the last item, for each $2 \leq i \leq k$, we choose a directed path $P_{i-1,i}$ in $D_{i-1,i}$, such that $P_{i-1,i}$ is directed from x_{i-1} to x_i if $(v_{i-1}, v_i) \in A(P)$ and from x_i to x_{i-1} if $(v_i, v_{i-1}) \in A(P)$. The existence of such a path follows in each case since x_{i-1}, x_i are in the same strong component of $D_{i-1,i}$. Clearly, $V(P_{i-1,i}) \subseteq V(D_{i-1,i}) = c^{-1}(\{i-1,i\})$. This proves the last item.

Claim 3. There are vertices $z_1, z_2, \ldots, z_k \in Y_1$ such that

- $c(z_i) = i$, for i = 1, ..., k.
- $z_1 \in V(C)$ and $z_k \in N^+(x_0)$.
- For every $2 \leq i \leq k$, there exists a directed path $Q_{i-1,i}$ in $D[Y_1]$ with endpoints z_{i-1}, z_i such that $Q_{i-1,i}$ is directed from z_{i-1} to z_i if $(v_{i-1}, v_i) \in A(P)$, and directed from z_i to z_{i-1} if $(v_i, v_{i-1}) \in A(P)$.
- The paths $Q_{i-1,i}$, i = 2, ..., k are pairwise internally vertex-disjoint.
- $V(C) \cap V(Q_{1,2}) = \{z_1\}$ and $V(C) \cap V(Q_{i-1,i}) = \emptyset$ for i = 3, ..., k.

Proof. We define the vertices z_i as follows: We define $z_1 \in V(C)$ to be the unique last vertex in V(C) we meet when traversing the trace of the path $P_{1,2}$ starting from $x_1 \in V(C)$. Since $P_{1,2}$ uses only colors 1 and 2, we must have $z_1 \in V(C) \setminus \{x_0\}$ and thus $c(z_1) = 1$. For $i = 2, \ldots, k - 1$, we successively define z_i to be the first vertex of $P_{i,i+1}$ we meet when traversing the trace of the path $P_{i-1,i}[z_{i-1}, x_i]$ starting from z_{i-1} (such a vertex exists, since $x_i \in V(P_{i-1,i}) \cap V(P_{i,i+1})$ by Claim 2). Since $V(P_{i-1,i}) \subseteq c^{-1}(\{i-1,i\}),$ $V(P_{i,i+1}) \subseteq c^{-1}(\{i, i+1\})$, it follows that $c(z_i) = i$. Finally, we put $z_k := x_k \in N^+(x_0)$. For each $i \in \{2, 3, \ldots, k\}$, we define $Q_{i-1,i} := P_{i-1,i}[z_{i-1}, z_i]$.

Let us now verify the correctness of the claim. The first three items follow directly from Claim 2 and the definition of the vertices z_i and the paths $Q_{i-1,i}$.

For the fourth item, let $i < j \in \{2, \ldots, k\}$ be given. We need to show that $Q_{i-1,i}$ and $Q_{j-1,j}$ can only intersect in their endpoints. If $j - i \ge 2$, then we directly conclude that $V(Q_{i-1,i}) \cap V(Q_{j-1,j}) \subseteq V(P_{i-1,i}) \cap V(P_{j-1,j}) \subseteq c^{-1}(\{i-1,i\}) \cap c^{-1}(\{j-1,j\}) = \emptyset$. If on the other hand j = i + 1, then by definition of z_i , no vertex on $Q_{i-1,i} = P_{i-1,i}[z_{i-1}, z_i]$ except for z_i lies on $P_{i,i+1}$, and therefore also not on $Q_{j-1,j} = P_{i,i+1}[z_i, z_{i+1}]$. Hence, $V(Q_{i-1,i}) \cap V(Q_{j-1,j}) = \{z_i\} = \{z_{j-1}\}$. This concludes the proof of the fourth item. The claim that $V(C) \cap V(Q_{1,2}) = \{z_1\}$ in the fifth item directly follows from our choice of $Q_{1,2} = P_{1,2}[z_1, z_2]$ and the definition of z_1 as being the last vertex on C we meet when traversing $P_{1,2}$ starting at x_1 . For $i \in \{3, \ldots, k\}$, we can conclude the second part of the last item from the inclusion $V(C) \cap V(Q_{i-1,i}) \subseteq (c^{-1}(\{1\}) \cup \{x_0\}) \cap c^{-1}(\{i-1,i\}) = \emptyset$.

Let S^* be the subdigraph of D formed by joining $S \subseteq D[Y_2]$, the pairwise distinct vertices z_1, \ldots, z_k and the connecting dipaths $Q_{i-1,i}, i = 2, \ldots, k$, the two anti-parallel directed paths $C[x_0, z_1], C[z_1, x_0]$ between x_0 and z_1 as well as the arc (x_0, z_k) . From Claim 3 and since $\left(\bigcup_{i=2}^k V(Q_{i-1,i}) \cup (V(C) \setminus \{x_0\})\right) \cap V(S) \subseteq Y_1 \cap Y_2 = \emptyset$, it follows that S^* is isomorphic to a subdivision of F^* , with $x_0, z_1, z_2, \ldots, z_k$ playing the roles of the vertices $v_0, v_1, v_2, \ldots, v_k$ of F^* .

We have thus shown that every digraph D with $\vec{\chi}(D) = \text{mader}_{\vec{\chi}}(F) + k$ contains a subdivision of F^* , and this concludes the proof of the theorem. \Box

By definition, every octus is obtained from K_1 via a sequence of operations of two types: ear addition and taking a subdigraph. For an octus F, let s(F) be the (minimal) number of operations needed to obtain F. Let us say that F is a maximal octus if it can be obtained from K_1 by a sequence of ear additions only. By repeatedly applying Theorem 7.15, we see that mader_{\vec{X}}(F) = v(F) for every maximal octus F. To complete the proof of Theorem 7.4, we also need to address non-maximal octi. This will be done using the following two lemmas.

Lemma 7.17. Every octus is a subdigraph of a maximal octus.

Proof. The proof is by induction on s(F). If s(F) = 0 then $F = K_1$ and the assertion is trivial. Suppose then that $s(F) \ge 1$. By the definition of an octus (see Definition 7.1), either F is a subdigraph of some octus F' with s(F') < s(F), or F is obtained by ear addition from some octus F' with s(F') < s(F). In the former case, the induction hypothesis implies that F' — and hence also F — is a subdigraph of a maximal octus, as required. Suppose then that F is obtained by ear addition from some occus F' with s(F') < s(F). By the induction hypothesis, F' is a subdigraph of some maximal occus F''. By performing on F'' the same ear addition which turns F' into F, we obtain a maximal occus which contains F. This completes the proof.

Lemma 7.18. For every connected subdigraph F' of a maximal octus F, there is a maximal octus \overline{F} such that F' is a spanning subdigraph of \overline{F} .

Proof. The proof is by induction on s(F). If s(F) = 0 then $F = K_1$ and the assertion is trivial. Let then F be a maximal octus with $s(F) \geq 1$, and let F' be a connected subdigraph of F. By the definition of maximal octi, there is some maximal octus F_0 with $s(F_0) < s(F)$ and $v_0 \in V(F_0)$ such that F is obtained from F_0 by ear addition, namely, by adding an oriented path $P = v_1, \ldots, v_k$ with $\{v_1, \ldots, v_k\} \cap V(F_0) = \emptyset$, as well as the arcs $(v_0, v_1), (v_1, v_0)$ and (w.l.o.g.) (v_0, v_k) . Consider the subdigraph $F'_0 := F'[V(F') \cap V(F_0)]$ of F_0 . If $V(F') \cap V(F_0) = \emptyset$, namely if $V(F') \subseteq \{v_1, \ldots, v_k\}$, then F' is an oriented path, and is hence a spanning subgraph of a bioriented path, which is a maximal octus. Suppose then that $V(F') \cap V(F_0) \neq \emptyset$. The way F is constructed from F_0 and the assumption that F' is connected imply that F'_0 is connected as well. By the induction hypothesis (applied to F'_0 , there is a maximal octus \overline{F}_0 such that F'_0 is a spanning subdigraph of \overline{F}_0 . If $F' = F'_0$ then we are done, and otherwise we must have $V(F') \cap \{v_1, \ldots, v_k\} \neq \emptyset$, which in turn implies that $v_0 \in V(F'_0) = V(\overline{F_0})$ because F' is connected. Now, if $\{v_1, \ldots, v_k\} \subseteq V(F')$ then F' is a spanning subdigraph of the maximal octus obtained from $V(\overline{F}_0)$ by adding the path P and connecting its endpoints to $v_0 \in V(F_0)$ using the arcs $(v_0, v_1), (v_1, v_0), (v_0, v_k)$. Otherwise, i.e. if $\{v_1, \ldots, v_k\} \not\subseteq V(F')$, then there must be some $1 \leq i < j \leq k$ such that $V(F') = V(\overline{F_0}) \cup \{v_1, \ldots, v_i\} \cup \{v_j, \ldots, v_k\}$ (as F' is connected). Now, let \overline{F} be the maximal octus obtained from \overline{F}_0 by a sequence of two ear additions: we first add the path v_1, \ldots, v_i and the arcs $(v_0, v_1), (v_1, v_0), (v_0, v_i)$ and then the path v_1, \ldots, v_k and the arcs $(v_0, v_i), (v_i, v_0), (v_0, v_k)$. Then F' is a spanning subdigraph of F, as required.

Proof of Theorem 7.4. Our goal is to show that $\operatorname{mader}_{\vec{\chi}}(F) = v(F)$ for every octus F. First, observe that it suffices to prove this statement for connected F, since the general statement would then follow by invoking Observation 7.10. So let F be a connected octus. By combining Lemmas 7.17 and 7.18, we see that F is a spanning subdigraph of some maximal octus \overline{F} . As mentioned before, Theorem 7.15 implies that $\operatorname{mader}_{\vec{\chi}}(\overline{F}) = v(\overline{F}) = v(F)$. As F is a subdigraph of \overline{F} , we have $\operatorname{mader}_{\vec{\chi}}(F) \leq \operatorname{mader}_{\vec{\chi}}(\overline{F})$ and hence $\operatorname{mader}_{\vec{\chi}}(F) = v(F)$, as required.

7.4 Tournaments of Order 4

In this section we prove Theorem 7.8. We give a separate proof for each of the 4-vertex tournaments. There are exactly four non-isomorphic tournaments on 4 vertices: \vec{K}_4 , the transitive tournament of order 4; \vec{K}_4^s , the unique strongly connected tournament of order 4; and the tournaments W_3^+, W_3^- obtained from the directed triangle \vec{C}_3 by adding a dominating source or sink, respectively. See Figure 7.2 for an illustration.



Figure 7.2: The four non-isomorphic tournaments of order 4.

Since W_3^+ and W_3^- are obtained from each other by reversing the orientations of all arcs, it suffices to prove Theorem 7.8 for \vec{K}_4 , \vec{K}_4^s and W_3^+ . While we can derive the result for the transitive tournament \vec{K}_4 directly from Theorem 7.14, the proofs for \vec{K}_4^s and W_3^+ are more involved and require some preparation.

Proof of $mader_{\vec{\chi}}(\vec{K}_4) = 4$. Let D be a given digraph with $\vec{\chi}(D) \ge 4$. Then D contains a 4-dicritical subdigraph $D' \subseteq D$. By Lemma 7.11, we have $\delta^+(D') \ge 3$. We now apply Theorem 7.14 to conclude that D' and thus also D contains a subdivision of \vec{K}_4 . This completes the proof.

We prepare the proofs of $\operatorname{mader}_{\vec{\chi}}(\vec{K}_4^s) = 4$ and $\operatorname{mader}_{\vec{\chi}}(W_3^+) = 4$ with a set of useful lemmas, providing reduction operations which preserve the property that the digraph does not contain a certain subdivision.

Lemma 7.19. Let D be a digraph, let $(u, w) \in A(D)$, and let D' be the digraph obtained from D by deleting u and adding the arc (x, w) for each $x \in N_D^-(u) \setminus (\{w\} \cup N_D^-(w))$.

Let F be a sink-free orientation of a cubic graph. If D' contains a subdivision of F, then so does D.

Proof. Let $S' \subseteq D'$ be a subdivision of F contained in D'. If $(x, w) \notin A(S')$ for all $x \in N_D^-(u)$, then S' is also a subdigraph of D and hence we have found a subdivision of F in D. So suppose that $(x, w) \in A(S')$ for some $x \in N_D^-(u)$.

We now distinguish between two cases.

Case 1: There exists $x' \in N_D^-(u) \setminus \{x\}$ such that $(x', w) \in A(S')$. Then w must be a branch vertex of the subdivision S', and since F is a sink-free orientation of a 3-regular graph, there exists a third neighbor x'' of w in S' satisfying $(w, x'') \in A(S') \subseteq A(D')$. By definition of D', we have $(w, x'') \in A(D)$ as well. We now see that the subdigraph S of D defined by $V(S) := V(S') \cup \{u\}, A(S) := (A(S') \setminus \{(x, w), (x', w)\}) \cup \{(x, u), (x', u), (u, w)\}$ forms a subdivision of F in D, where the branch vertex w of S' is moved to the new branch vertex u of S (and w becomes a subdivision vertex).

Case 2: x is the unique vertex in $N_D^-(u)$ such that $(x, w) \in A(S')$. Then the subdigraph S of D defined by the vertex-set $V(S) := V(S') \cup \{u\}$ and arc-set $A(S) := (A(S') \setminus \{(x, w)\}) \cup \{(x, u), (u, w)\}$ forms a subdivision of F contained in D. \Box

Lemma 7.20. Let D be a strongly connected digraph, let $v \in V(D)$, and let (X, Y) be a non-trivial partition of $V(D) \setminus \{v\}$ such that $(x, y) \notin A(D)$ for all $x \in X, y \in Y$. Suppose further that D[X] is strongly connected. Let $D_1 := D[X \cup \{v\}]$ and let D_2 be defined by $V(D_2) := Y \cup \{v\}$ and $A(D_2) := A(D[Y \cup \{v\}]) \cup \{(y, v) \mid y \in Y, x \in X, (y, x) \in A(D)\}$. Let further F be a sink-free orientation of a cubic graph. Then

1. If D_1 or D_2 contains a subdivision of F, then so does D.

2. $\vec{\chi}(D) \le \max\{\vec{\chi}(D_1), \vec{\chi}(D_2)\}.$

Proof.

- 1. The claim is trivial for D_1 , since $D_1 \subseteq D$. Now suppose that D_2 contains a subdivision of F. The vertex v in D must have an in-neighbor in the set X, for otherwise $(X, Y \cup \{v\})$ would form a directed separation of D, contradicting the assumed strong connectivity. Since D[X] is strongly connected, it follows that there exists an inarborescence $T \subseteq D[X \cup \{v\}]$ rooted at $x_0 := v$ which spans $X \cup \{v\}$. Let n := |X|, and fix an ordering x_0, x_1, \ldots, x_n of the vertices of T such that each vertex of T appears before its children in the ordering (i.e., if $(x_i, x_j) \in A(T)$ then i > j). For every $i = n, n-1, n-2, \ldots, 1, 0$, let H_i be the digraph obtained from D by removing all arcs in $A(D[X \cup \{v\}]) \setminus A(T)$, deleting the vertices $\{x_{i+1}, \ldots, x_n\}$ and adding the arc (y, x_j) for every $y \in Y$ and $j \in \{0, 1, \dots, i\}$ such that y has an out-neighbor $x \in \{x_{i+1}, \ldots, x_n\}$ in D and the first intersection of the unique x-x₀-path in T with $\{x_0, x_1, \ldots, x_i\}$ is x_i . Note that H_n is a subdigraph of D and that $H_0 = D_2$. Further we can observe that for every $1 \leq i \leq n$, the digraph H_{i-1} is obtained from H_i by deleting x_i , and adding an arc from every $x \in N_{H_i}^-(x_i)$ to the parent of x_i in T. Hence, repeated application of Lemma 7.19 yields that if $D_2 = H_0$ contains a subdivision of F, then the same is true for all $H_i, 0 \le i \le n$. Hence $H_n \subseteq D$ contains a subdivision of F, and this proves the claim.
- 2. Let $k := \max\{\vec{\chi}(D_1), \vec{\chi}(D_2)\}$, let $c_1 : X \cup \{v\} \to \{1, \ldots, k\}$ be an acyclic coloring of D_1 and $c_2: Y \cup \{v\} \to \{1, \ldots, k\}$ an acyclic coloring of D_2 , respectively. Without loss of generality we may assume $c_1(v) = 1 = c_2(v)$. We now define a k-coloring of V(D)by putting $c(x) := c_1(x)$ for every $x \in X$, c(v) := 1, and $c(y) := c_2(y)$ for every $y \in Y$. We claim that this defines an acyclic k-coloring of D. Indeed, if not, then there exists a directed cycle C in D which is monochromatic under c. If $v \notin V(C)$, then since there is no arc from X to Y in D, we must have either $V(C) \subseteq X$ or $V(C) \subseteq Y$, which in both cases yields a contradiction to our choice of c_1 and c_2 as acyclic colorings. Hence, $v \in V(C)$ and $V(C) \subseteq c^{-1}(1)$. If $V(C) \cap Y = \emptyset$, then C is a monochromatic cycle in the coloring c_1 of $D_1 = D[X \cup \{v\}]$, a contradiction. We therefore have $v \in V(C), V(C) \cap Y \neq \emptyset$. Since there are no edges from X to Y, there must be $w \in Y$ such that $(v, w) \in A(C)$. Let C[v, w'] be a maximal directed subpath of C starting at v such that $V(C[v, w']) \setminus \{v\} \subseteq Y$. (In other words, C[v, w'] is obtained by traversing C starting from the arc (v, w) and stopping just before the cycle leaves Y.) Then either $(w', v) \in A(D)$, or $(w', x) \in A(D)$ for some $x \in X$ and hence $(w',v) \in A(D_2)$ by definition of D_2 . Therefore, C[v,w'] + (w',v) forms a directed cycle in D_2 , all of whose vertices have color 1 under c_2 , contradicting our assumption on c_2 . This contradiction shows that our initial assumption was wrong, namely that cis indeed an acyclic coloring of D, proving that $\vec{\chi}(D) \leq k = \max\{\vec{\chi}(D_1), \vec{\chi}(D_2)\}$. \Box

Lemma 7.21. Let D be a digraph, and let $u, v, w \in V(D)$ be pairwise distinct such that $v, w \in N^+(u) \cap N^-(u)$ (i.e., $\{u, v\}$ and $\{u, w\}$ induce digons). Let D^* be obtained from D by deleting v and w and adding the arcs

$$\{(u,x) \mid x \in V(D) \setminus \{u,v,w\}, N^{-}(x) \cap \{v,w\} \neq \emptyset\}$$

and

$$\{(x,u) \mid x \in V(D) \setminus \{u,v,w\}, N^+(x) \cap \{v,w\} \neq \emptyset\}.$$

Let F be an oriented cubic graph. If D^* contains a subdivision of F, then so does D.

Proof. Let S^* be a subdigraph of D^* isomorphic to a subdivision of F. If $u \notin V(S^*)$, then S^* is also a subdigraph of D and we are done. Hence, suppose in the following that $u \in V(S^*)$. If u is a subdivision vertex in S^* , then let $u^- \in N_{D^*}^-(u)$ and $u^+ \in N_{D^*}^+(u)$ denote the in- and the out-neighbor of u in S^* , respectively. By definition of D^* there exist $x^{-}, x^{+} \in \{u, v, w\}$ such that $(u^{-}, x^{-}), (x^{+}, u^{+}) \in A(D)$. Let P denote the bioriented path with vertex-trace v, u, w. Then clearly P contains a directed x^-, x^+ -path P_{x^-, x^+} . Now $A(S) := (A(S^*) \setminus \{(u^-, u), (u, u^+)\}) \cup \{(u^-, x^-), (x^+, u^+)\} \cup A(P_{x^-, x^+}) \text{ forms the arc-set}$ of a subdigraph $S \subseteq D$ isomorphic to a subdivision of F. For the next case suppose that u is a branch vertex of the subdivision S^* . For every in-neighbor $y \in N^-_{S^*}(u)$ in S^* , let $x(y) \in \{u, v, w\}$ be a vertex such that $(y, x(y)) \in A(D)$, and for every out-neighbor $y \in N^+_{S^*}(u)$, let $x(y) \in \{u, v, w\}$ be a vertex such that $(x(y), y) \in A(D)$. Let y_1, y_2, y_3 be the three distinct neighbors of u in S^* , ordered in such a way that $x(y_2)$ lies on the unique bioriented subpath $P_{x(y_1),x(y_3)}$ of P connecting the vertices $x(y_1)$ and $x(y_3)$. It is now evident that the subdigraph of D obtained from S^* by deleting u and adding $P_{x(y_1),x(y_3)}$ and the arcs $(y_i,x(y_i))$ for $y_i \in N^-_{S^*}(u)$ and $(x(y_i),y_i)$ for $y_i \in N^+_{S^*}(u)$, contains a subdivision of F with $x(y_2)$ as a branch vertex. This verifies the claim in the second case as well and concludes the proof.

Proof of $mader_{\vec{\chi}}(\vec{K}_4^s) = 4$. Suppose towards a contradiction that the claim is wrong, and let D be a counterexample minimizing lexicographically the pair (v(D), a(D)); namely, the number of vertices is minimized with first priority and the number of arcs with second priority. Clearly, $v(D) \ge 5$, D is 4-dicritical, and it contains no subdivision of \vec{K}_4^s . By Lemma 7.12, D is strongly connected.

Claim 1. *D* is strongly 2-vertex-connected.

Proof. Suppose towards a contradiction that there exists a vertex $v \in V(D)$ such that D-v is not strongly connected. This means that D-v has more than one strong component. Let $X \subseteq V(D-v)$ be the vertex set of a strong component of D-v which is a "sink" in D-v, that is, there is no arc leaving X. Let $Y := V(D) \setminus (X \cup \{v\})$. Then (X, Y) forms a partition of $V(D) \setminus \{v\}$, D[X] is strongly connected and $(x, y) \notin A(D)$ for all $x \in X, y \in Y$. We can therefore apply Lemma 7.20 with $F = \vec{K}_4^s$ to obtain a pair D_1, D_2 of digraphs with vertex-sets $X \cup \{v\}, Y \cup \{v\}$, respectively, such that neither D_1 nor D_2 contains a subdivision of \vec{K}_4^s and $4 = \vec{\chi}(D) \leq \max\{\vec{\chi}(D_1), \vec{\chi}(D_2)\}$. However, this means that there is some $i \in \{1, 2\}$ such that $\vec{\chi}(D_i) \geq 4$, D_i contains no \vec{K}_4^s -subdivision and clearly $v(D_i) < v(D)$. This contradicts the assumed minimality of D, thus showing that the assumption was wrong, namely that D is indeed strongly 2-vertex-connected.

Claim 2. $\delta^+(D), \delta^-(D) \ge 4.$

Proof. Note that \vec{K}_4^s is isomorphic to the tournament obtained from it by reversing all arcs. It follows that since D is a counterexample to the claim, so is \overleftarrow{D} , which is the digraph obtained from D by reversing all its arcs. Evidently, we have $(v(\overleftarrow{D}), a(\overleftarrow{D})) = (v(D), a(D))$, meaning that \overleftarrow{D} is also a minimal counterexample (in the sense defined above). Since $\delta^-(D) = \delta^+(\overleftarrow{D})$, it suffices to prove $\delta^+(D) \ge 4$.

Suppose towards a contradiction that there exists a vertex $u \in V(D)$ with $d^+(u) \leq 3$. Since D is 4-dicritical, Lemma 7.11 implies that $d^+(u) = 3$. We now distinguish between two cases depending on the structure of the out-neighborhood of u.

Case 1: There exists some $w \in N^+(u)$ such that $(w, u) \notin A(D)$. In this case, let D' be the digraph defined as in Lemma 7.19. Namely, D' is obtained from D by deleting u and adding the arcs (x, w) for all $x \in N^{-}(u)$. By Lemma 7.19, D' contains no subdivision of K_4^s . Since v(D') = v(D) - 1, the minimality assumption on D implies that $\vec{\chi}(D') \leq 3$. So let $c': V(D) \setminus \{u\} \to \{1, 2, 3\}$ be an acyclic 3-coloring of D'. Write $N_D^+(u) = \{w, w_1, w_2\}$. Fix a color $c_u \in \{1, 2, 3\} \setminus \{c'(w_1), c'(w_2)\}$ (which clearly exists). Let $c: V(D) \to \{1, 2, 3\}$ be the coloring of D defined by c(x) := c'(x) for all $x \in V(D) \setminus \{u\}$ and $c(u) := c_u$. Since $\vec{\chi}(D) = 4$, there has to be a directed cycle C in D which is monochromatic under c. Clearly, C has to pass through u, for otherwise it would have been a monochromatic dicycle already in the coloring c' of D'. Since none of the out-arcs $(u, w_1), (u, w_2)$ is monochromatic, we must have $(u, w) \in A(C)$. Let $u' \in N_D^-(u)$ be the unique predecessor of u on C. Then $u' \neq w$ because $(w, u) \notin A(D)$ by assumption. It follows from the definition of D' that replacing the directed subpath u', (u', u), u, (u, w), w of C with the ("direct") arc (u', w)in D' defines a directed cycle C' in D' such that $V(C') = V(C) \setminus \{u\}$. Hence, C' is a monochromatic dicycle in the acyclic coloring c' of D'. This contradiction shows that our initial assumption $d^+(u) \leq 3$ was wrong.

Case 2: $(w, u) \in A(D)$ for all $w \in N^+(u)$. We claim that in this case, we can find a pair $w_1, w_2 \in N^+(u)$ of distinct neighbors of u such that $(w_1, w_2) \notin A(D)$. Indeed, suppose this were not the case. Then the vertices $\{u\} \cup N^+(u)$ induce a \breve{K}_4 in D. However, this clearly means that D contains $\vec{K}_4^s \subseteq \vec{K}_4$ as a subdigraph, contradicting our initial assumption on D. So let us fix, in the following, a pair of distinct $w_1, w_2 \in N^+(u) \subseteq N^-(u)$ such that $(w_1, w_2) \notin A(D)$. Let D^{*} be the digraph obtained from D by applying the operation of Lemma 7.21 to $\{u, w_1, w_2\}$; that is, we delete w_1 and w_2 and add the arc (u, x) for every $x \in V(D) \setminus \{u, w_1, w_2\}$ which has an in-neighbor in $\{w_1, w_2\}$ and the arc (x, u) for every $x \in V(D) \setminus \{u, w_1, w_2\}$ which has an out-neighbor in $\{w_1, w_2\}$. By Lemma 7.21, D^* does not contain a subdivision of \vec{K}_4^s . We clearly have $v(D^*) < v(D)$ and so the minimality assumption on D yields that there is an acyclic 3-coloring $c^*: V(D^*) \to \{1,2,3\}$ of D^* . Write $N^+(u) = \{w_1, w_2, w_3\}$, and let $c_u \in \{1, 2, 3\}$ be a color distinct from both $c^*(u)$ and $c^*(w_3)$. We now define a 3-coloring c of D by putting $c(x) := c^*(x)$ for all $x \in V(D) \setminus \{u, w_1, w_2\}, c(u) := c_u, \text{ and } c(w_1) := c(w_2) := c^*(u).$ Since $\vec{\chi}(D) = 4$, there must be a dicycle C in D which is monochromatic under c. Then C cannot contain u, for otherwise it would have to leave u through one of the out-arcs $(u, w_1), (u, w_2), (u, w_3), (u, w_3$ but by the definition of the coloring c, none of these arcs is monochromatic. On the other hand, we must have $V(C) \cap \{w_1, w_2\} \neq \emptyset$, for otherwise C would be a monochromatic dicycle in (D^*, c^*) , which is impossible. Observe also that $V(C) \setminus \{w_1, w_2\} \neq \emptyset$ because $(w_1, w_2) \notin A(D)$. Let $x_0, x_1, \ldots, x_\ell = x_0$ be the vertex-trace of C in D. Now consider the closed sequence $y_0, y_1, \ldots, y_\ell = y_0$ of vertices in D^* , where $y_i := x_i$ if $x_i \notin \{w_1, w_2\}$ and $y_i := u$ if $y_i \in \{w_1, w_2\}$. The definitions of D^* and c and the fact that $u \notin V(C)$ imply that $c^*(y_i) = c(x_i)$ for every $i = 1, \ldots, \ell$, and that either $(y_{i-1}, y_i) \in A(D^*)$ or $y_{i-1} = y_i = u$ for every $i = 1, ..., \ell$. This means that in D^* there is a monochromatic closed directed walk which contains at least two vertices: it contains u because $V(C) \cap \{w_1, w_2\} \neq \emptyset$ and at least one other vertex because $V(C) \setminus \{w_1, w_2\} \neq \emptyset$ and $u \notin V(C)$. Therefore, D^* contains a monochromatic dicycle. All in all, this contradicts the fact that c^* was chosen as an acyclic coloring of D^* , implying that our initial assumption $d^+(u) \leq 3$ was wrong.

To sum up, we have arrived at a contradiction in both cases, which means that we indeed must have $\delta^+(D) \ge 4$. As argued above, we can derive $\delta^-(D) = \delta^+(D) \ge 4$ with the same arguments applied to the minimal counterexample D. This proves the claim.

With Claims 1 and 2 at hand, we can now apply Theorem 7.13 to D with k = 2, and thus obtain a vertex $v \in V(D)$ such that D - v is strongly 2-vertex-connected. We now complete the proof of the Theorem by explicitly constructing a subdivision of \vec{K}_4^s in D. We start with the following observation.

Claim 3. There are 3 directed cycles C_1 , C_2 , C_3 in D such that $V(C_i) \cap V(C_j) = \{v\}$ for any two distinct $i, j \in \{1, 2, 3\}$.

Proof. Since D is 4-dicritical, D - v admits an acyclic coloring with colors $\{1, 2, 3\}$. For every $i \in \{1, 2, 3\}$, if we try and extend this coloring to D by assigning color i to v, we have to find a monochromatic directed cycle C_i in D, which has to pass through v. Note that $V(C_i) \cap V(C_j) = \{v\}$ for all $1 \le i < j \le 3$, because all vertices in $V(C_i) \setminus \{v\}$ receive color i $(1 \le i \le 3)$.

The rest of the proof is divided into two cases depending on the lengths of the cycles C_i .

Case 1. All the three cycles C_1, C_2, C_3 have length two, i.e., are digons. Let v_1, v_2, v_3 be vertices such that $V(C_i) = \{v, v_i\}$ (i = 1, 2, 3). Since D - v is strongly connected, there has to be a directed path in D - v starting in v_1 and ending in $\{v_2, v_3\}$. Let P be a shortest such directed path, and without loss of generality assume that it ends in v_2 . By the minimality assumption on P we know that $v_3 \notin V(P)$. Now put A := V(P). We clearly have $|A| \ge |\{v_1, v_2\}| = 2$, and since D - v is strongly 2-vertex-connected we may apply Theorem 1.2 to obtain that there are two vertex-disjoint A- v_3 -dipaths P_1 and P_2 in D - v. For i = 1, 2, let us define $V(P) \cap V(P_i) =: \{w_i\}$. Then P_1 and P_2 only intersect at v_3 , and P only intersects P_i at w_i (for i = 1, 2). Without loss of generality (by relabeling if necessary), we may assume that when traversing P from v_1 towards v_2 , we first meet w_1 before we meet w_2 . Now let S be the subdigraph of D defined by the union of the following dipaths in D: P, P_1 , P_2 , $(v, (v, v_1), v_1)$, $(v_3, (v_3, v), v)$ and $(v_2, (v_2, v), v)$. It is now easy to observe that S constitutes a subdivision of \vec{K}_4^s whose branch vertices are v, w_1, w_2, v_3 .

Case 2. There is some $i \in \{1, 2, 3\}$ such that $|C_i| \geq 3$. Without loss of generality we may assume that i = 1. Let v_2 be the unique out-neighbor of v on C_2 . Define $B := V(C_1) \setminus \{v\} \subseteq V(D - v)$. Clearly, $|B| \geq 3 - 1 = 2$, and hence we may apply Theorem 1.2 to the strongly 2-connected digraph D - v to conclude that there are two vertex-disjoint v_2 -B-dipaths P_1 , P_2 in D - v. For i = 1, 2, let $x_i \in A$ and $y_i \in B$ be the endpoints of P_i . It is now clear that the union of C_1 and the internally vertex-disjoint dipaths $Q := (v, (v, v_2), v_2)$, P_1 and P_2 is a subdivision of \vec{K}_4^s in D with branch vertices v, v_2, y_1, y_2 . This contradicts our initial assumption that D contains no subdivision of \vec{K}_4^s .

Since we arrived at a contradiction in both cases, it follows that our initial assumption that there exists a (smallest) digraph D with $\vec{\chi}(D) \geq 4$ not containing a subdivision of \vec{K}_4^s was wrong. This finishes the proof.

We now move on to show that $\operatorname{mader}_{\vec{\chi}}(W_3^+) = 4$. This proof is partly inspired by a method used in [HMM18].

Proof of $mader_{\vec{\chi}}(W_3^+) = 4$. Suppose towards a contradiction that there exists a digraph D such that $\vec{\chi}(D) \ge 4$, but D contains no subdivision of W_3^+ . Assume additionally that D lexicographically minimizes the pair (v(D), a(D)) (i.e., the number of vertices is minimized with first priority, and the number of arcs is minimized with second priority). Clearly, $v(D) \ge 5$ and D is 4-dicritical. Hence, $\delta^+(D), \delta^-(D) \ge 3$ by Lemma 7.11, and D is strongly connected by Lemma 7.12.

Claim 1. D is strongly 2-vertex-connected.

Proof. Suppose towards a contradiction that there exists a vertex $v \in V(D)$ such that D-v is not strongly connected. This means that D-v has more than one strong component. Let $X \subseteq V(D-v)$ be the vertex set of a strong component of D-v which is a "sink" in D-v, that is, there is no arc leaving X. Let $Y := V(D) \setminus (X \cup \{v\})$. Then (X, Y) forms a partition of $V(D) \setminus \{v\}$, D[X] is strongly connected and $(x, y) \notin A(D)$ for all $x \in X, y \in Y$. We can therefore apply Lemma 7.20 with $F = W_3^+$ to obtain a pair D_1, D_2 of digraphs with vertex-sets $X \cup \{v\}, Y \cup \{v\}$, respectively, such that neither D_1 nor D_2 contains a subdivision of W_3^+ and $4 = \vec{\chi}(D) \leq \max\{\vec{\chi}(D_1), \vec{\chi}(D_2)\}$. However, this means that there is some $i \in \{1, 2\}$ such that $\vec{\chi}(D_i) \geq 4$, D_i contains no W_3^+ -subdivision and clearly $v(D_i) < v(D)$. This contradicts the assumed minimality of D. This contradiction shows that the assumption was wrong, namely that D is indeed strongly 2-vertex-connected.

Claim 2. The underlying graph of D is 3-vertex-connected.

Proof. Suppose towards a contradiction that there is a set $K \subseteq V(D)$ such that $|K| \leq 2$ and D-K is not weakly connected. Let (X_1, X_2) be a partition of $V(D) \setminus K$ into two non-empty sets such that there is no arc between X_1 and X_2 in D - K. Since D is strongly 2-vertex-connected, we must have |K| = 2, say $K = \{s_1, s_2\}$ for some distinct $s_1, s_2 \in V(D)$. For i = 1, 2, let D_i be the digraph with vertex-set $V(D_i) := X_i \cup K$ and arc-set $A(D_i) := A(D[X_i \cup K]) \cup \{(s_1, s_2), (s_2, s_1)\}$. We claim that none of D_1, D_2 contains a subdivision of W_3^+ . Indeed, suppose towards a contradiction that for some $i \in \{1, 2\}$, there exists a subdigraph $S \subseteq D_i$ which is isomorphic to a subdivision of W_3^+ . If $A(S) \cap \{(s_1, s_2), (s_2, s_1)\} = \emptyset$, then S would also be a subgraph of D, contradicting our assumptions on D. Hence, S has to contain an arc between s_1 and s_2 , and without loss of generality we may assume that $(s_1, s_2) \in A(S)$. Since W_3^+ contains no digons, the same is true for S and hence $(s_2, s_1) \notin A(S)$. We now claim that there exists an s_1 - s_2 -dipath in $D - X_i$. To this end, choose an arbitrary vertex $w \in X_{3-i}$. Since both $D - s_1$ and $D-s_2$ are strongly connected (by Claim 1), there are dipaths P_1 and P_2 in $D-s_1$ resp. $D - s_2$ such that P_1 starts at w and ends at s_2 , while P_2 starts at s_1 and ends at w. Since D contains no arcs between X_1 and X_2 , we must have $V(P_1) \subseteq X_{3-i} \cup \{s_2\}$ and $V(P_2) \subseteq X_{3-i} \cup \{s_1\}$. Finally, we see that the concatenation of P_1 and P_2 is a directed walk from s_1 to s_2 , implying that $P_1 \cup P_2 \subseteq D[X_{3-i} \cup \{s_1, s_2\}] = D - X_i$ contains an s_1 - s_2 -dipath P. Clearly, P is internally vertex-disjoint from all dipaths in $S - (s_1, s_2) \subseteq D$, and hence the subdigraph $S' := (S - (s_1, s_2)) \cup P$ of D is isomorphic to a subdivision of S, which in turn is a subdivision of W_3^+ . This contradicts our initial assumption that D is W_3^+ -subdivision-free. We conclude that neither D_1 nor D_2 contains a subdivision of W_3^+ , as claimed. Since clearly $v(D_1), v(D_2) < v(D)$, the assumed minimality of D yields that D_1 and D_2 admit acyclic 3-colorings c_1 and c_2 , respectively. Since the pair s_1, s_2 induces a digon in both D_1 and D_2 , we must have $c_1(s_1) \neq c_1(s_2), c_2(s_1) \neq c_2(s_2)$. Hence, possibly after permuting the color set, we may assume that $c_1(s_1) = c_2(s_1) = 1, c_1(s_2) = c_2(s_2) = 2$. We now claim that the common extension $c : V(D) \rightarrow \{1, 2, 3\}$ of c_1 and c_2 to D defines an acyclic coloring of D. Indeed, a monochromatic directed cycle C in (D, c) would have to contain vertices of both X_1 and X_2 , for otherwise it would also be a monochromatic dicycle in (D_1, c_1) or (D_2, c_2) , contradicting the assumption that these are acyclic colorings. However, since $K = \{s_1, s_2\}$ separates X_1 and X_2 , this is only possible if $\{s_1, s_2\} \subseteq V(C)$. But then C is not monochromatic because $c(s_1) = 1$ and $c(s_2) = 2$. This shows that c is indeed an acyclic coloring of D, which in turn contradicts $\vec{\chi}(D) = 4$. So we see that our initial assumption that the underlying graph of D admits a 2-separator was wrong. This concludes the proof of Claim 2.

Claim 3. For every $x \in V(D)$ there is a directed cycle C in D - x such that $|C| \ge 3$.

Proof. Let $x \in V(D)$ be given arbitrarily. Suppose towards a contradiction that every directed cycle in the digraph D - x has length two, i.e., is a digon. Recall that in a strongly connected digraph, every arc lies on a directed cycle. Since D - x is strongly connected (Claim 1), every arc of D - x is contained in a digon, and hence D - x is a symmetric digraph. Since D - x contains no directed cycle of length at least 3, this is only possible if D - x is a biorientation of a forest. But then the bipartition of this forest defines an acyclic 2-coloring of D - x. By assigning to x a distinct third color, we obtain an acyclic 3-coloring of D, a contradiction to $\vec{\chi}(D) = 4$. This proves the claim.

Claim 4. For every pair (x, C) of a vertex in D and a directed cycle C in D-x of length at least 3, there exists a partition (W, K, Z) of V(D) with the following properties:

- $x \in W$ and $V(C) \subseteq K \cup Z$
- There is no arc in D with tail in W and head in Z.
- |K| = 2.

A partition (W, K, Z) with these properties will be called a *good separation* for (x, C).

Proof. We claim that there are no three x-V(C)-dipaths in D which pairwise intersect only at x. Indeed, three such dipaths joined with C would form a subdivision of W_3^+ , which does not exist in D by assumption. By Theorem 1.2, there is a set $K \subseteq V(D) \setminus \{x\}$ of size at most 2 such that there are no x-V(C)-dipaths in D - K. Let $W \subseteq V(D) \setminus \{x\}$ be the set of vertices reachable in D - K via a dipath starting at x, and let $Z := V(D) \setminus \{W \cup K\}$. It follows now directly by definition that $x \in W$ and $V(C) \subseteq K \cup Z$, and there is no arc with tail in W and head in Z. We have $|K| \leq 2$, and since D is strongly 2-vertex-connected (by Claim 1), it follows that |K| = 2. Therefore (W, K, Z) is a good separation of (x, C).

In the following, for every pair (x, C) of a vertex $x \in V(D)$ and a directed cycle C in D - x of length at least 3, we denote by $\omega(x, C)$ the minimum of |W| over all good separations (W, K, Z) of (x, C). Let

$$\omega_0 := \min\{\omega(x, C) \mid x \in V(D), C \text{ dicycle in } D - x, |C| \ge 3\}$$

Claim 5. Let $x \in V(D)$, let C be a directed cycle in D - x of length at least 3, and let (W, K, Z) be a good separation for (x, C) such that $|W| = \omega(x, C)$. Then every vertex of W is reachable from x by a dipath in D[W].

Proof. Let W' be the set of all vertices $w \in W$ which are reachable from x in D[W]. Evidently, $x \in W' \subseteq W$. Observe that $(W', K, (W \setminus W') \cup Z)$ forms a good separation for (x, C), because D has no edge with tail in W' and head in $W \setminus W'$. It follows that $|W'| \geq \omega(x, C) = |W|$. This implies that W' = W, as required.

Claim 6. There exists a pair (x^*, C^*) of a vertex $x^* \in V(D)$, a dicycle C^* of length at least 3 in $D - x^*$, as well as a good separation (W, K, Z) of (x^*, C^*) , such that:

- $|W| = \omega_0$,
- there exist $z^* \in Z, w^* \in W \setminus \{x^*\}$ such that $(z^*, w^*) \in A(D)$.

Proof. Let (x_0, C_0) be a pair of a vertex and a disjoint dicycle in D, such that $|C_0| \ge 3$, and such that (x_0, C_0) attains the minimum $\omega_0 = \omega(x_0, C_0)$. Let (W, K, Z) be a good separation for (x_0, C_0) such that $|W| = \omega(x_0, C_0) = \omega_0$. Note that (W, K, Z) is also a good separation for every pair (x, C_0) where $x \in W \setminus \{x_0\}$.

Observe that there has to exist an arc between Z and W, for if not, then D-K is not weakly connected, contradicting the facts that |K| = 2 and that the underlying graph of D is 3-vertex-connected (Claim 2). As there are no arcs from W to Z, there has to exist an arc from Z to W. Let (z_0, w_0) be such an arc. If $w_0 \neq x_0$, then we directly obtain that (x_0, C_0) together with (W, K, Z) and the arc (z_0, w_0) satisfy all the required properties in the statement of the claim. If $x_0 = w_0$, then choose some $x_1 \in W \setminus \{x_0\}$. Such a selection is possible, since $N^+(x_0) \setminus K \subseteq W$ and $N^+(x_0) \setminus K \neq \emptyset$ as $d^+(x_0) \geq \delta^+(D) \geq 3$. Since (W, K, Z) is a good separation also for (x_1, C_0) , and since $x_1 \neq w_0 = x_0$, it follows now that (x_1, C_0) together with (W, K, Z) and the arc (z_0, w_0) have all the claimed properties. This concludes the proof of Claim 6.

In the following, let us consider a pair (x^*, C^*) together with the good separation (W, K, Z)and the arc (z^*, w^*) as given by Claim 6. Since $D - x^*$ is strongly connected (Claim 1), and since $(z^*, w^*) \in A(D - x^*)$, there exists a directed cycle C' in $D - x^*$ passing through (z^*, w^*) . As D has no arc from W to Z, the dicycle C' must use at least one vertex from K, which means that $|C'| \geq 3$. Let us write $K = \{s_1, s_2\}$ and assume (without loss of generality) that $s_1 \in V(C')$.

By Claim 4, there exists a good separation (W', K', Z') for the pair (x^*, C') ; choose it such that $|W'| = \omega(x^*, C')$, and such that it minimizes $|K' \cap Z|$ among all such good separations. We claim that $K' \cap W \neq \emptyset$. Indeed, if we had $K' \cap W = \emptyset$ then we would have $D[W] \subseteq D - K'$, which would imply that $w^* \in W$ is reachable from x^* by a dipath in D - K' (as every vertex of W is reachable from x^* by a dipath in D[W] by Claim 5). However, this would contradict the facts that $x^* \in W'$, $w^* \in V(C') \subseteq K' \cup Z'$, and that there is no path from W' to Z' in D - K' (by the definition of a good separation). Let us write $K' = \{s'_1, s'_2\}$, where $s'_1 \in W$.

Claim 7. $K' \cap Z = \emptyset$.

Proof. Suppose towards a contradiction that $K' \cap Z \neq \emptyset$, which means that $s'_2 \in Z$ (because $s'_1 \in W$). Let R be the set of vertices reachable from x^* via a dipath in the digraph $D - \{s'_1, s_2\}$. We claim that $R \subseteq W'$. Suppose towards a contradiction that

 $R \setminus W' \neq \emptyset$. Then there is an $(x^*, R \setminus W')$ -dipath P in $D - \{s'_1, s_2\}$. Note that $V(P) \subseteq R$ by the definition of R. Let $y \in R \setminus W'$ be the end-vertex of P and y' its predecessor. Then $y' \in W'$ because only the last vertex y of P is in $R \setminus W'$. Since (y', y) cannot have its tail in W' and head in Z', we must have $y \in (V(D) \setminus \{s'_1, s_2\}) \setminus (W' \cup Z') \subseteq \{s'_2\}$; hence $y = s'_2 \in Z$. Since $x^* \in W$, and since K separates W from Z, there must be a vertex on P - y which belongs to $K = \{s_1, s_2\}$. However, this vertex can be neither s_1 nor s_2 ; indeed, $s_1 \notin V(P) \setminus \{y\}$ because $V(P) \setminus \{y\} \subseteq W'$ and $s_1 \in V(C') \subseteq K' \cup Z'$, and $s_2 \notin V(P)$ because P is contained in $D - \{s'_1, s_2\}$.

This contradiction shows that $R \subseteq W'$, as claimed.

There is no arc from R to $V(D) \setminus (R \cup \{s'_1, s_2\})$ by the definition of R, which means that $|\{s'_1, s_2\}| = 2$ since D is strongly 2-vertex-connected (by Claim 1). We furthermore have $x^* \in R$ and $V(C') \subseteq V(D) \setminus W' \subseteq V(D) \setminus R$. Hence, $(R, \{s'_1, s_2\}, V(D) \setminus (R \cup \{s'_1, s_2\}))$ defines a good separation for the pair (x^*, C') . By the choice of (W', K', Z'), this means that R = W' and $|R| = \omega(x^*, C')$. We further have $|\{s'_1, s_2\} \cap Z| = 0 < 1 = |K' \cap Z|$ (because $s'_1 \in W$ and $s_2 \in K$), contradicting our choice of (W', K', Z'). This contradiction shows that the assumption $K' \cap Z \neq \emptyset$ was wrong, concluding the proof of the claim.

Claim 8. $W' \cap Z \neq \emptyset$ and $Z' \cap Z \neq \emptyset$.

Proof. We start by showing that $W' \cap Z \neq \emptyset$. Suppose towards a contradiction that $W' \cap Z = \emptyset$. Since $\{w^*, s_1\} \subseteq V(C') \subseteq K' \cup Z'$, we have $W' \cap \{w^*, s_1\} = \emptyset$ and hence $W' \subseteq V(D) \setminus (Z \cup \{w^*, s_1\}) = (W \cup K) \setminus \{w^*, s_1\} = (W \setminus \{w^*\}) \cup \{s_2\}$. On the other hand, we have $|W'| = \omega(x^*, C') \geq \omega_0 = |W|$, and hence $W' = (W \setminus \{w^*\}) \cup \{s_2\}$. In particular, $s_2 \in W'$. Since D is strongly 2-vertex-connected, s_2 must have an outneighbor $y \in Z$, for otherwise there would be no dipath from x^* to Z in $D - s_1$. Using Claim 7 and our assumption that $W' \cap Z = \emptyset$, we have $(W' \cup K') \cap Z = \emptyset$, and hence $y \in V(D) \setminus (W' \cup K') = Z'$. However, this means that (s_2, y) is an arc from a vertex in W' to a vertex in Z', a contradiction. This contradiction shows that the initial assumption $W' \cap Z = \emptyset$ was wrong, proving the first part of claim.

For the second part, recall that $z^* \in Z$ and $z^* \in V(C') \subseteq K' \cup Z'$. Since $K' \cap Z = \emptyset$ (by Claim 7), we conclude that $z^* \in Z' \cap Z$ and hence $Z' \cap Z \neq \emptyset$, as required.

Claim 9. Every dipath in D starting in $W' \cap Z \neq \emptyset$ and ending in $Z' \cap Z \neq \emptyset$ must contain s_1 .

Proof. We first establish that $s_2 \in W'$. To see this, pick some vertex $v \in W' \cap Z$. By Claim 5 and as $|W'| = \omega(x^*, C')$, there exists an x^* -v-dipath in D[W']. Since $x^* \in W$ and $v \in Z$, this dipath must contain a vertex from K. However, since $s_1 \notin W'$, this vertex must be s_2 , implying that $s_2 \in W'$.

Now to prove the claim, let P be a dipath starting in a vertex $a \in W' \cap Z$ and ending in a vertex $b \in Z' \cap Z$. Let $y \in V(P)$ be the last vertex on P contained in $W' \cup K'$ when traversing P starting from a. Let y' be the successor of y on P; then $(y, y') \in A(D)$ and $y' \in Z'$. Hence, we must have $y \in K'$ (since D has no arcs from W' to Z'). It now follows from Claim 7 that $y \in W \cup K$. Now let us consider the subpath P[y, b] starting at y and ending at b. By definition of y, no vertex on P[y, b] is contained in W', and hence s_2 does not lie on this path. However, P[y, b] starts in a vertex of $W \cup K$ and ends in a vertex of Z, which means that it must contain a vertex from $K = \{s_1, s_2\}$. Hence, $s_1 \in V(P[y, b]) \subseteq V(P)$. This proves the claim. Since s_1 is contained in none of the two non-empty sets $W' \cap Z$ and $Z' \cap Z$, Claim 9 shows that $D - s_1$ is not strongly connected, contradicting Claim 1. This shows that our very first assumption, namely that a digraph D with $\vec{\chi}(D) \ge 4$ which does not contain a subdivision of W_3^+ , exists, was wrong. This completes the proof of mader $_{\vec{\chi}}(W_3^+) = 4$. \Box

7.5 Subcubic Digraphs

In this section we give a proof of Theorem 7.9 using Theorem 6.3 from Chapter 6.

Proof of Theorem 7.9. As a first step let us note that given $n \in \mathbb{N}$, every undirected graph G with $\delta(G) \geq 10.437n$ contains every *n*-vertex subcubic graph as a minor. This follows directly from a result of Reed and Wood [RW16], who proved that every graph with average degree at least n + 6.291m contains every graph with *n* vertices and *m* edges as a minor and since $m \leq \frac{3}{2}n$ for subcubic graphs.

Let now D be any digraph with $\vec{\chi}(D) \geq 22n$, F a subcubic digraph on $n \geq 2$ vertices and H its underlying undirected subcubic graph. By Theorem 6.3 there exists an undirected graph G such that D is a strong \vec{G} -minor-model and $\chi(G) \geq 11n$. In particular, G contains a subgraph of minimum degree at least 11n - 1 > 10.437n and hence, by our earlier remark, an H-minor. This implies that \vec{G} contains a strong \vec{H} -minor-model and hence D does so. However, as $F \subseteq \vec{H}$, it also follows that D contains a strong F-minormodel. Let $\{X_f : f \in V(F)\}$ be a branch set partition of V(D) witnessing this. Recall that, by definition, for every arc $e = (u_1, u_2) \in A(F)$ there exist vertices $v(e, u_1) \in X_{u_1}$ and $v(e, u_2) \in X_{u_2}$ such that $(v(e, u_1), v(e, u_2)) \in A(D)$.

Let next $u \in V(F)$ be an arbitrary vertex with total degree $d = d_F(u) \in \{0, 1, 2, 3\}$ and let us denote the arcs incident to u by e_1, \ldots, e_d . Furthermore, for $i = 1, \ldots, d$ we put $v_i := v(e_i, u)$. We claim that there exists a vertex $b(u) \in X_u$ and for every $i = 1, \ldots, d$ a directed path P_i^u in $D[X_u]$ such that

- P_1^u, \ldots, P_d^u are internally vertex-disjoint;
- if u is the tail of e_i , then P_i^u is a directed path from b(u) to v_i ;
- if u is the head of e_i , then P_i^u is a directed path from v_i to b(u).

This claim holds trivially if d = 0, and if d = 1 then we can simply put $b(u) = v_1$ and let P_1^u be the trivial one-vertex path consisting of v_1 .

If d = 2 then, without loss of generality, by the symmetry of reversing all arcs in Dand F, we may assume that u is the head of e_1 . We then can put $b(u) := v_2$, let P_1^u be any directed path in $D[X_u]$ from v_1 to v_2 (given by strong connectivity), and take P_2^u to be the trivial one-vertex path consisting only of v_2 .

Finally suppose d = 3. Since F is subcubic, u either has in-degree one and out-degree two, or vice versa. As before, without loss of generality, by symmetry we may assume that the first case occurs, and it is e_1 that enters u and e_2 and e_3 that emanate from it. Take now P_{12} and P_{13} to be directed paths in $D[X_f]$ starting at v_1 and ending at v_2 and v_3 , respectively. We define now b(u) as the first vertex in $V(P_{12})$ that we meet when traversing P_{13} backwards (starting at v_3); P_1^u as the subpath of P_{12} directed from v_1 to b(u); P_2^u as the subpath of P_{12} directed from b(u) to v_2 ; and P_3^u as the subpath of P_{13} directed from b(u) to v_3 . It follows by definition that P_1, P_2, P_3 are internally vertex-disjoint, and hence the claim follows. To finish the proof, let $S \subseteq D$ be a subdigraph with vertex set

$$V(S) := \bigcup_{u \in V(F)} \left(\bigcup_{i=1}^{d(u)} V(P_i^u) \right),$$

and arcs

 \geq

$$A(S) := \left\{ \left(v(e, u_1), v(e, u_2) \right) \mid e = (u_1, u_2) \in A(F) \right\} \cup \left(\bigcup_{u \in V(F)} \left(\bigcup_{i=1}^{d(u)} A(P_i^u) \right) \right).$$

Then S is a digraph isomorphic to a subdivision of F in which a vertex $u \in V(F)$ is represented by the branch-vertex b(u). This concludes the proof.

7.6 Mader-Perfect Digraphs and Conclusion

In this chapter, and in particular Sections 7.3 and 7.4, we investigated under which circumstances the simple inequality mader $_{\vec{\chi}}(F) \geq v(F)$ is tight. Observe that tightness is trivially preserved under taking spanning subdigraphs. It turns out however, that the optimality of the bound does not necessarily carry over to arbitrary subdigraphs. In fact, in Proposition 7.22 below we show that for any digraph F there exists a constant k_F such that adding k_F isolated vertices to F produces a digraph whose Mader number equals its number of vertices. This suggests that the class of digraphs F which satisfy mader $_{\vec{\chi}}(F) = v(F)$ may not have a meaningful characterization. This motivates the following definition. We call a digraph F Mader-perfect if for every (induced) subdigraph F'of F, the Mader number of F' equals its order.

Proposition 7.22. For every digraph F there is $k_F \in \mathbb{N}$ such that for every $k \geq k_F$, the digraph F' obtained from F by adding k new isolated vertices satisfies $mader_{\vec{\chi}}(F') = v(F')$. In fact, it suffices to take $k_F = 2 \cdot mader_{\vec{\chi}}(F) - v(F) - 1$.

Proof. Let $k \ge k_F := 2 \cdot \text{mader}_{\vec{\chi}}(F) - v(F) - 1$ be arbitrary. Consider any given digraph D such that $\vec{\chi}(D) \ge k + v(F)$. We need to show that D contains a subdivision of F which misses at least k of the vertices of D.

Let $X \subseteq V(D)$ be a vertex set such that $\vec{\chi}(D[X]) = \text{mader}_{\vec{\chi}}(F)$. Then D[X] contains a subdivision of F, and we have $m := \vec{\chi}(D - X) \ge k + v(F) - \text{mader}_{\vec{\chi}}(F)$, for otherwise we could color D with less than k + v(F) colors. Let Y_1, \ldots, Y_m be a partition of $V(D) \setminus X$ into m acyclic sets. Let us first consider the case that at most v(F) - 1 of these sets are singletons. Then

$$v(D) - |X| = |Y_1| + \ldots + |Y_m| \ge 2m - (v(F) - 1)$$
$$2k + 2v(F) - 2 \cdot \operatorname{mader}_{\vec{\chi}}(F) - (v(F) - 1) = k + (k - k_F) \ge k.$$

Evidently, the subdivision of F contained in D[X] does not use any of the $\geq k$ vertices in $V(D) \setminus X$, concluding the proof in this case.

In the other case, at least v(F) of the sets Y_i are singletons; without loss of generality, say $Y_i = \{y_i\}$ for $1 \leq i \leq v(F)$. Since $Y_1 \ldots, Y_m$ form an optimal acyclic coloring of D - X, we cannot merge any two color classes to obtain an acyclic coloring with m - 1colors. It follows that $(y_i, y_j), (y_j, y_i) \in A(D)$ for every $1 \leq i < j \leq v(F)$. This implies that D contains a copy of F on the vertices $y_1, \ldots, y_{v(F)}$. By deleting the v(F) vertices $y_1, \ldots, y_{v(F)}$, we obtain a digraph of dichromatic number at least $\vec{\chi}(D) - v(F) \geq k$, and hence, the remaining digraph consists of at least k vertices. So we see that D indeed contains a subdivision of F missing at least k vertices in this second case as well. This concludes the proof.

The first two main results of this chapter — namely Theorems 7.4 and 7.8 — can be restated as saying that all octi digraphs and all tournaments of order 4 are Mader-perfect. For octi digraphs this follows immediately from Theorem 7.4, since octi are closed under taking subdigraphs; and for 4-vertex tournaments this follows from the fact that every nonspanning subdigraph F' of a 4-vertex tournament is a subgraph of an oriented triangle, and for those F' the equality mader $_{\vec{X}}(F') = v(F')$ follows from Theorem 7.4. In a similar vein, Proposition 7.7 implies that $\vec{K}_3 - e$ is Mader-perfect. Although Proposition 7.7 can be obtained as a consequence of Theorem 2.2, let us give an independent shorter proof.

Proof of Proposition 7.7. It is sufficient to show that every 3-dicritical digraph D contains a subdivision of $\vec{K}_3 - e$. So let D be a 3-dicritical digraph. Then $\delta^+(D), \delta^-(D) \ge 2$ and Dis strongly-connected, as guaranteed by Lemmas 7.11 and 7.12, respectively. By Theorem 7.13 applied for k = 1, there is $v \in V(D)$ such that D - v is strongly connected. Since D is 3-dicritical, there exists an acyclic 2-coloring $c: V(D) \setminus \{v\} \to \{1,2\}$. Evidently, c cannot be extended to an acyclic 2-coloring of D. This means that for each i = 1, 2, D contains a directed cycle C_i which contains v, such that all vertices in $V(C_i) \setminus \{v\}$ are colored with color i (under c). Note that $V(C_1) \cap V(C_2) = \{v\}$. Since D - v is strongly-connected, there is a path in D from $V(C_1)$ to $V(C_2)$ which avoids v. Let P be a shortest path from $V(C_1)$ to $V(C_2)$ avoiding v, and let us denote the endpoints of P by x_1, x_2 (where $x_i \in V(C_i)$). The minimality of P implies that $V(P) \cap V(C_i) = \{x_i\}$ (for each i = 1, 2), since otherwise P could be replaced by a shorter path. Now it is easy to see that the vertices v, x_1, x_2 and the (internally vertex-disjoint) dipaths $C_1[v, x_1], C_1[x_1, v], C_2[v, x_2], C_2[x_2, v], P$ form a subdivision of $\vec{K}_3 - e$, as required.

Altogether, we see that the class of Mader-perfect digraphs is quite rich. We believe it would be interesting to obtain a precise characterization of this class.

Problem 7.1. Characterize Mader-perfect digraphs.

On the negative side, \vec{K}_3 is the smallest digraph F satisfying $\operatorname{mader}_{\vec{\chi}}(F) > v(F)$, hence no bioriented clique of order at least 3 is Mader-perfect. In fact for any $t \geq 3$, the digraph obtained from \vec{K}_{t+2} by removing a bioriented \vec{C}_5 has dichromatic number t but contains no subdivision of \vec{K}_t . This shows that the Mader number of \vec{K}_t is at least t+1.

The analogous problem for undirected graphs seems to be interesting as well: Call an undirected graph F Mader-perfect if for every subgraph $F' \subseteq F$, every graph G of chromatic number at least v(F') contains an F'-subdivision.

Problem 7.2. Characterize Mader-perfect graphs.

Since we have $\vec{\chi}(\vec{G}) = \chi(G)$ for every undirected graph, it follows that if F is an undirected graph such that at least one of the orientations of F is Mader-perfect, then F is Mader-perfect. In particular, it follows that every forest, every cactus graph, and K_4 are Mader-perfect graphs. Using our terminology, Catlin's counterexamples to Hajós' conjecture say that K_k is not Mader-perfect for every $k \geq 7$.

Already determining mader_{$\vec{\chi}$}(\vec{K}_3) exactly seems to be a challenging problem. From above we can only show that mader_{$\vec{\chi}$}(\vec{K}_3) \leq 9, where the upper bound follows from a combination of Proposition 7.7 and Theorem 7.23 below. We believe that the truth lies with the lower bound, provided by the above construction. **Conjecture 7.2.** We have $mader_{\vec{\chi}}(\vec{K}_3) = 4$, *i.e.*, every digraph D with no \vec{K}_3 -subdivision admits an acyclic 3-coloring.

It is natural to ask how dense Mader-perfect digraphs can be. For $k \in \mathbb{N}$, let m(k) denote the maximum possible number of arcs of a Mader-perfect digraph of order k. Using a variant of the classical probabilistic argument of Erdős and Fajtlowicz [EF81], we can show that $m(k) = O(k^{3/2}\sqrt{\log k})$, which means that Mader-perfect digraphs have to be (at least somewhat) sparse. In fact, let us show the slightly more general claim that

$$\mathrm{mader}_{\vec{\chi}}(F) \ge \frac{cm^2}{k^2 \log m} \tag{7.1}$$

for every digraph F on k vertices and $m \ge c_1 k \log k$ arcs, where c_1 is a suitably large absolute constant. The bound (7.1) shows that if $\operatorname{mader}_{\vec{\chi}}(F) = v(F) = k$ (which has to be the case if F is Mader-perfect), then $m = O(k^{3/2} \sqrt{\log k})$, as claimed.

In order to prove (7.1), consider any fixed digraph F consisting of $k \ge 2$ vertices and $m \ge c_1 k \log k$ arcs. Let D(n, p) be the random digraph⁵ with parameters $n = \lfloor \frac{m}{2} \rfloor$ and $p = \frac{m}{4k^2}$. We claim that with positive probability, D(n, p) contains neither a set of k vertices spanning at least m/2 arcs nor an acyclic set of more than $\frac{c_2k^2\log m}{m}$ vertices for some suitable absolute constant c_2 . To see this, note that the expected number of arcs spanned by some fixed set of k vertices in D(n, p) is $pk(k-1) < \frac{m}{4}$, and hence the Chernoff-bound yields that the probability that some k fixed vertices span at least $\frac{m}{2}$ arcs is bounded by $\exp(-\frac{m}{12})$. Therefore the probability that there are k vertices spanning at least $\frac{m}{2}$ arcs is at most $\binom{n}{k} \exp(-\frac{m}{12}) \le (m/2)^k \exp(-\frac{m}{12}) = \exp(k \log(m/2) - \frac{m}{12}) < \frac{1}{2}$, provided c_1 is chosen large enough.

Similarly, the probability that any fixed set of $\alpha > \frac{c_2 k^2 \log m}{m}$ vertices is acyclic in D(n,p) is at most $\alpha!(1-p)^{\binom{\alpha}{2}}$. Hence, the probability that D(n,p) contains an acyclic set of size at least α is at most $\binom{n}{\alpha} \alpha!(1-p)^{\binom{\alpha}{2}} \leq \exp(\alpha \log n - p\binom{\alpha}{2}) < 1/2$, provided c_2 is chosen large enough (where in the last inequality we plugged in our choice of n, p and α).

We conclude that there exists a digraph D on $n = \lfloor \frac{m}{2} \rfloor$ vertices containing no k vertices spanning at least $\frac{m}{2}$ arcs and whose dichromatic number is at least

$$\vec{\chi}(D) \ge \frac{n}{\left(\frac{c_2k^2\log m}{m}\right)} = \Omega\left(\frac{m^2}{k^2\log m}\right).$$

Observe that D contains no subdivision of F; indeed, if D contained a subdivision of F, then since D has at most $\frac{m}{2}$ vertices, at least $\frac{m}{2}$ of the subdivision paths would have to be of length 1, i.e. be "direct" arcs between the k branch vertices of the subdivision. But this is impossible as D contains no set of k vertices spanning at least $\frac{m}{2}$ arcs, a contradiction. This proves (7.1).

We note that if F is symmetric, i.e. if it is a biorientation of an undirected graph, then we can improve the bound (7.1) to $\operatorname{mader}_{\vec{\chi}}(F) \geq \Omega(m/\log m)$. To see this, let D be a tournament of order m/2 and dichromatic number $\Omega(m/\log m)$ (it is well-known that such tournaments exist, see [Erd80]). Then D contains no subdivision of F. Indeed, since D contains no digons, any subdivision of F in D must contain at least m/2 subdivision vertices, one per every digon in F. But as v(D) = m/2 < m/2 + v(F), there are not enough vertices in D to fit a subdivision of F. As a corollary, we see that if F is Mader-perfect and symmetric, then $a(F) \leq O(k \log k)$, where k = v(F).

⁵Recall that D(n,p) is the random digraph on the vertex-set $\{1,\ldots,n\}$, where for each $1 \le i \ne j \le n$ we put the arc (i,j) independently with probability p.

So far we have shown that $m(k) = O(k^{3/2}\sqrt{\log k})$. As for a lower bound, consider the digraph obtained from $\vec{K}_3 - e$ by performing k - 3 ear additions, where at each step we attach a digon to the existing digraph. Then the resulting digraph F_k has k vertices and 2k - 1 arcs. By combining Proposition 7.7 with Theorem 7.15, we see that F_k is Mader-perfect. Hence, $m(k) \ge 2k - 1$. It would be interesting to close the gap between the upper and lower bounds. We conjecture that the truth lies with the latter.

Conjecture 7.3. m(k) = O(k).

Aboulker et al. [ACH⁺19] studied the behaviour of the Mader number with respect to the insertion of arcs, and proved the following bound using a beautiful argument based on breadth-first-search trees.

Theorem 7.23 ([ACH⁺19], Lemma 31). If F is a digraph and $e \in A(F)$, then

$$mader_{\vec{\chi}}(F) \leq 4 \cdot mader_{\vec{\chi}}(F-e) - 3.$$

Using this upper bound and a lower-bound-construction for tournaments, they obtained the following bounds on mader $\vec{k}(\vec{K}_t)$.

Theorem 7.24. For every integer $t \geq 2$,

$$\Omega\left(\frac{t^2}{\log t}\right) \le mader_{\vec{\chi}}(\vec{K}_t) \le 4^{t^2 - 2t + 1}(t - 1) + 1.$$

Consider the digraph F_t on t vertices and 2t - 1 as constructed above. Then we have $\operatorname{mader}_{\vec{\chi}}(F_t) = t$, and hence by starting from F_t and repeatedly applying Theorem 7.23, we get the (slightly) improved bound $\operatorname{mader}_{\vec{\chi}}(\vec{K}_t) \leq 4^{t^2 - 3t + 1}(t - 1) + 1$.

Still, the gap between the lower and upper bounds on $\operatorname{mader}_{\vec{\chi}}(\vec{K}_t)$ remains huge. Unfortunately, the techniques used in this chapter to tackle sparse digraphs do not seem to allow for substantial improvements of the upper bound. Our attempts to improve the lower bound to a super-quadratic growth have also been unsuccessful. It is tempting to conjecture the following:

Conjecture 7.4. There exists an absolute constant c > 0 such that $mader_{\vec{\chi}}(\vec{K}_t) \leq ct^2$ for every positive integer t.

It is worth noting that by a result of Girão et al. [GPS21], every tournament T with minimum out-degree at least ct^2 (for some absolute constant c) contains a subdivision of \vec{K}_t . This implies that *in tournaments*, having dichromatic number larger than ct^2 forces a subdivision of \vec{K}_t . As mentioned in the introduction, however, extending the result of [GPS21] to general digraphs is impossible, since having large minimum out- and in-degree does not force \vec{K}_t -subdivisions in general digraphs for any $t \geq 3$.

Another intriguing question is to determine the Mader number of bioriented cycles. Specifically, is it the case that $\operatorname{mader}_{\vec{X}}(\vec{C}_{\ell}) = \ell$ for all $\ell \geq 4$?

Problem 7.3. What is mader_{\vec{x}}(\vec{C}_{ℓ})?

A related problem is to determine the maximum possible chromatic number of a digraph which does not contain a subdivision of *any* bioriented cycle. We conjecture that the answer is 2.

Conjecture 7.5. Let D be a digraph with $\vec{\chi}(D) \geq 3$. Then there is $\ell \geq 3$ such that D contains a subdivision of \vec{C}_{ℓ} .

Chapter 8

Excluding Butterfly Minors

8.1 Introduction

Graphs and digraphs considered in this chapter are simple.

One of the arguably most influential problems in graph theory was the Four-Color-Conjecture, answered positively by Appel and Haken in 1976. A directed version of this famous problem, the Two-Color-Conjecture posed by Erdős and Neumann-Lara and independently by Skrekovski (see [BFJ⁺04, Erd80]), still stands open. Recall that a digraph D is called *oriented* if its underlying undirected graph is simple.

Conjecture 8.1. Every oriented planar digraph D is acyclically 2-colorable.

Although this conjecture appears quite innocent, there seems to be a lack of methods for attacking it. The strongest partial result proved so far is due to Mohar and Li [LM17], who showed the following:

Theorem 8.1. Every oriented planar digraph without directed triangles is acyclically 2colorable.

While the class of oriented planar digraphs is a very natural class of digraphs to consider, structurally it does not behave quite as nicely as undirected planar graphs: While the planar graphs can be described as the minor-closed class of graphs excluding K_5 and $K_{3,3}$ and are therefore linked to Hadwiger's conjecture, it is not possible to describe oriented planar digraphs by forbidden digraph minors. Indeed, for any of the three minor concepts considered in this thesis, an oriented planar digraph may have a \vec{K}_2 as a minor, which is not an oriented graph.

This difference between the undirected and directed setting leads us to the natural question what actually *are* the classes of directed graphs closed under minors that can be acyclically colored with a given number of colors. In the undirected case, Hadwiger's conjecture states that for every $t \ge 1$, the largest minor-closed class of t-colorable graphs equals the K_{t+1} -minor-free graphs. In this chapter we raise the analogous problem for directed graphs, where undirected minors are replaced by butterfly-minors.

Problem 8.1. Given $t \in \mathbb{N}$, what is the inclusion-wise largest butterfly-minor closed class \mathcal{D}_t of acyclically t-colorable digraphs?

It follows directly from the definition that \mathcal{D}_1 equals the class of acyclic digraphs (since every butterfly-minor of an acyclic digraph is still acyclic). Already characterizing the next class, \mathcal{D}_2 , is non-trivial. As the first main contribution of this chapter we solve this problem by proving that \mathcal{D}_2 is exactly the class of the *non-even digraphs*, which we have previously encountered in Chapters 2 and 5. Recall that a digraph D is non-even if there exists an arc-weighting $w: E(D) \to \{0, 1\}$ such that every directed cycle has odd total arc-weight. Let us restate Theorem 5.2 here for convenience.

Theorem 8.2 (Seymour and Thomassen [ST87]). A directed graph is non-even if and only if it does not contain an odd bicycle as a butterfly minor.

Since $\vec{\chi}(\vec{C}_k) = 3$ for every odd $k \ge 3$, it is clear that \mathcal{D}_2 must exclude every odd bicycle as a butterfly-minor, and hence that every digraph in \mathcal{D}_2 is non-even. The non-trivial part of our characterization of \mathcal{D}_2 therefore amounts to the following result.

Theorem 8.3. Let D be a non-even digraph. Then $\vec{\chi}(D) \leq 2$.

As illustrated by Figure 8.1, there are both non-planar non-even digraphs and planar even digraphs. In that sense, Theorem 8.3 does not contribute to a resolution of Conjecture 8.1. On the other hand, \mathcal{D}_2 contains a relatively large and interesting class of planar digraphs, which are called *strongly planar* (compare [Gue01], [GT11]). A digraph D is said to be strongly planar if it admits a crossing-free embedding into the plane such that around every vertex $v \in V(D)$, the incident out- and in-arcs form intervals in the rotational order around v. As we will see further below, every strongly planar digraph is non-even, and in that sense, Theorem 8.3 verifies Conjecture 8.1 for an interesting subclass of the planar digraphs, which are allowed to contain many directed triangles and in fact, even digons. The strongly planar digraphs will be further discussed in Chapter 14.





Problem 8.1 is closely related to the question of forcing complete butterfly-minors in digraphs. Analogously to the definition of $sm_{\vec{\chi}}(t)$ in Chapter 6, for a given integer $t \geq 1$ let us denote by $bm_{\vec{\chi}}(t)$ the smallest integer such that every digraph D with $\vec{\chi}(D) \geq bm_{\vec{\chi}}(t)$ contains the complete digraph \vec{K}_t as a butterfly-minor. Since every topological minor is also a butterfly-minor, we can see from the remarks following Theorem 7.24 that $bm_{\vec{\chi}}(t)$ indeed exists for every $t \geq 1$ and that $bm_{\vec{\chi}}(t) \leq 4^{t^2-3t+1}(t-1)+1$. Define the integer inverse function of $bm_{\vec{\chi}}(\cdot)$ by

$$b(x) := \max\{t \ge 1 \mid bm_{\vec{x}}(t) \le x\}$$

Let us further denote, for any $t \ge 1$, by \mathcal{K}_t the class of all digraphs with no $\overset{\leftrightarrow}{K}_t$ as a butterfly-minor. Then, on the one hand, every digraph excluding $\overset{\leftrightarrow}{K}_{b(t+1)}$ as a butterfly-minor is colorable with $bm_{\vec{\chi}}(b(t+1)) - 1 \le t$ colors. On the other hand, every digraph

in \mathcal{D}_t must exclude \vec{K}_{t+1} as a butterfly-minor, since its dichromatic number exceeds t. Therefore, for every t we have the inclusions

$$\mathcal{K}_{b(t+1)} \subseteq \mathcal{D}_t \subseteq \mathcal{K}_{t+1}.$$

These inclusions show that the classes \mathcal{D}_t and \mathcal{K}_t are sandwiched between one another and hence can be expected to share similar structural properties. To see how tight the above inclusions are one needs to obtain good lower bounds on $b(\cdot)$, or equivalently good upper bounds on $bm_{\vec{\chi}}(\cdot)$. In this direction, we have the following second main result of this chapter, which is an application of Theorem 6.2 from Chapter 6.

Theorem 8.4. For every $t \ge 1$, we have $bm_{\vec{\chi}}(t) \le 2m_{\chi}(2t) - 1 = O(t(\log \log t)^6)$.

A non-even digraph D may be characterised by the property that there exists a labelling of its arcs with elements from the cyclic group $(\mathbb{Z}_2, +)$ such that the sum of the arc-labels of any directed cycle equals 1 (is non-zero). Replacing the group \mathbb{Z}_2 by an arbitrary finite Abelian group in this definition, we obtain a family of digraph classes which naturally generalize non-even digraphs as follows.

Definition 8.1. Let (A, +) be any finite Abelian group. We say that an arc-labelling $w: A(D) \to A$ of D using elements from A is zero-sum free if

$$\sum_{e \in A(C)} w(a) \neq 0$$

(summation as in (A, +)) holds for every directed cycle C in D.

We say that a digraph D is A-non-zero if it admits a zero-sum free A-arc-labelling, and otherwise we say that D is A-zero.

With this terminology, a digraph is even (non-even) iff it is \mathbb{Z}_2 -zero (non-zero).

As a qualitative generalization of Theorem 8.3 to these more general classes, we prove the following result as a consequence of Theorem 8.4 and the main result from [MS21b].

Corollary 8.5. There exists an absolute constant C > 0 such that the following holds for every finite Abelian group (A, +) with $|A| \ge 3$: If a digraph D is A-non-zero, then

$$\vec{\chi}(D) \le C \cdot |A| (\log \log |A|)^6.$$

So-called *list colorings* naturally generalize several types of colorings of graphs and have been widely investigated. The basic idea of list coloring is that every vertex of a graph receives its own respective palette of colors, from which it is supposed to pick its own color for a proper coloring of the graph. The main difference to the chromatic number is that different vertices may have completely different lists of colors. The *choice number* $\chi_{\ell}(G)$ of an undirected graph G is the smallest integer k such that for every assignment of color lists of length at least k to the vertices, a proper graph coloring with colors chosen from the lists exists. Clearly, we have $\chi(G) \leq \chi_{\ell}(G)$. Maybe more surprisingly there are graphs with $\chi(G) = 2$, in particular the complete bipartite graphs, for which $\chi_{\ell}(G)$ can grow arbitrarily large.

It is natural to apply the concept of list coloring also to acyclic colorings of digraphs. Indeed, such a notion was investigated in [BHL18]. Therein, for a given digraph D equipped with an assignment of finite color lists $\mathcal{L} = \{L(v) | v \in V(D)\}$ to the vertices, an \mathcal{L} -listcoloring of D is defined to be a choice function $c : V(D) \to \bigcup \mathcal{L}$ such that for any vertex $v \in V(D)$, we have $c(v) \in L(v)$, and moreover, c defines an acyclic digraph coloring, that is, $D[c^{-1}(i)]$ is acyclic for all $i \in \bigcup \mathcal{L}$. Putting $L(v) := \{1, \ldots, k\}$ for each vertex simply yields the definition of a usual digraph k-coloring. In [BHL18], a digraph D is called k-list colorable (also k-choosable) if for any list assignment \mathcal{L} , where $|L(v)| \ge k$ for every $v \in V(D)$, there is an \mathcal{L} -list coloring of D. The smallest integer $k \ge 1$ for which a digraph D is k-choosable now is defined to be the list dichromatic number (also choice number) $\vec{\chi}_{\ell}(D)$. Again, we have $\vec{\chi}(D) \le \vec{\chi}_{\ell}(D)$ for every digraph. However, as pointed out in [BHL18], this estimate can be arbitrarily bad for bipartite oriented graphs.

It is therefore desirable to identify classes of digraphs with bounded choice number. In the context of Conjecture 8.1, the authors of [BHL18] observed that oriented planar digraphs are 3-choosable and posed the question whether they are even 2-choosable.

From Theorem 8.3 we know that all non-even digraphs are 2-colorable, and so it is natural to ask whether they are also 2-choosable. This question can rather easily be answered in the negative, see Figure 8.2 for an example of a strongly planar digraph with choice number 3. In the last main result of this chapter that concerns directed graphs, we show that 3 is the (best possible) upper bound for the choice number of non-even digraphs.

Theorem 8.6. Let D be a non-even digraph. Then $\vec{\chi}_{\ell}(D) \leq 3$.



Figure 8.2: A non-2-choosable strongly planar digraph.

Applications to Matching theory. We have seen in Chapter 5 that non-even digraphs and their recognition problem are equivalent to the even cycle problem, and that there are many polynomial-time equivalent versions of this problem. One of these variations amounts to a famous problem from structural matching theory which we describe in the following. An undirected graph G is called matching covered if G is connected and for every edge $e \in E(G)$ there is some $M \in \mathcal{M}(G)$ with $e \in M$, where $\mathcal{M}(G)$ denotes the set of all perfect matchings of G. A set $S \subseteq V(G)$ of vertices is called conformal if G - S has a perfect matching. A subgraph $H \subseteq G$ is conformal if V(H) is a conformal set and H has a perfect matching. A cycle C in G is called M-alternating if it alternately uses edges from M and $E(G) \setminus M$. Clearly, the conformal cycles of G are exactly the cycles occurring as an alternating cycle in at least one perfect matching.

Counting the number of perfect matchings in a given graph (also known as the *dimer* problem) is an important and well-known task which is known to be #P-hard on general graphs [Val79]. However, there is a rather rich class of graphs for which the number of perfect matchings can be expressed in terms of the determinant of a well-known matrix and can thus be computed in polynomial time [Kas67, Lit75, Tho06]. These graphs are known as the *Pfaffian* graphs:

A graph G is called Pfaffian if there exists an orientation \vec{G} of G such that every conformal cycle of G contains an odd number of directed edges going in one and an odd

8.1. INTRODUCTION

number of directed edges going in the other direction in \overline{G} . Such an orientation is also called *Pfaffian*. It is well-known that any planar graph is Pfaffian (see [Kas67]). Since edges that are not contained in a perfect matching do not contribute to a Pfaffian orientation in any way, one usually just considers matching covered graphs in this context. Similar to non-even digraphs, bipartite matching covered Pfaffian graphs can be described by excluded minors. To state the complete theorem, we need a connection between directed graphs and bipartite graphs with perfect matchings, as well as the definition of minors in the context of matching covered graphs.

Let G be a matching covered graph and let v_0 be a vertex of G of degree two incident to the edges $e_1 = v_0v_1$ and $e_2 = v_0v_2$. Let H be obtained from G by contracting both e_1 and e_2 and deleting all resulting parallel edges. We say that H is obtained from G by *bicontraction* or *bicontracting the vertex* v_0 . We say that H is a matching minor of G if H can be obtained from a conformal subgraph of G by repeatedly bicontracting vertices of degree two. Similar to how topological minors specialize graph minors, there is the following specialisation of matching minors: A *bisubdivision* of an edge is a subdivision using an even number of subdivision-vertices (possibly 0). We call H_2 a *bisubdivision of* H_1 if H_1 is a matching covered graph and H_2 can be obtained by bisubdividing the edges of H_1 . If a matching covered graph G contains a conformal bisubdivision of a matching covered graph H, then H is a matching minor of G, but the converse is not true. If G contains no conformal bisudivision of H, it is called H-free.

The following operation, called M-direction, describes a fundamental bijection between bipartite graphs G equipped with perfect matchings and directed graphs which reveals intimate relationships between these objects and opens the doors for applications of our digraph-theoretic results in this setting. Vice versa, this bijection makes tools from matching theory available to the proof of our first main result, Theorem 8.3.

Definition 8.2. Let $G = (A \cup B, E)$ be a bipartite graph and let $M = \{a_1b_1, \ldots, a_{|M|}b_{|M|}\}$ with $a_i \in A, b_i \in B$ for $1 \le i \le |M|$ be a perfect matching of G. The M-direction $\mathcal{D}(G, M)$ of G is a digraph defined as follows.

- 1. $V(\mathcal{D}(G, M)) := \{v_1, \dots, v_{|M|}\}$ and
- 2. $A(\mathcal{D}(G, M)) := \{(v_i, v_j) \mid a_i b_j \in E(G), i \neq j\}.$

Note that the above operation is reversible and that every digraph D is the M-direction of its bipartite splitting-graph equipped with the canonical perfect matching.

The *M*-directions of a bipartite matching covered graph *G* inherit important properties from *G*. Most importantly, the directed cycles in an *M*-direction are in bijection with the *M*-alternating cycles of *G*. Another relation is about connectivity. A graph *G* is called *k*-extendable if it is connected, has at least 2k + 2 vertices and every matching of size *k* is contained in a perfect matching of *G*. The following statement is folklore, it is mentioned in [RST99] and a proof can be derived from [AHLS03].

Theorem 8.7 (Robertson et al. [RST99], Aldred et al. [AHLS03]). Let G be a bipartite matching covered graph and M a perfect matching of G. Then G is k-extendable if and only if $\mathcal{D}(G, M)$ is strongly k-(vertex-)connected.

The next result describes an intimate connection between matching minors of bipartite graphs and butterfly-minors of directed graphs.

Lemma 8.8 (McCuaig [McC00]). Let G and H be bipartite matching covered graphs. Then H is a matching minor of G if and only if there exist perfect matchings $M \in \mathcal{M}(G)$ and $M' \in \mathcal{M}(H)$ such that $\mathcal{D}(H, M')$ is a butterfly minor of $\mathcal{D}(G, M)$. Finally, non-even digraphs relate to Pfaffian bipartite graphs via the following important theorem of Little [Lit75], which characterizes bipartite Pfaffian graphs by excluding the single graph $K_{3,3}$ as a matching minor.

Theorem 8.9 (Little [Lit75], Seymour and Thomassen [ST87]). Let G be a bipartite graph with a perfect matching M. The following statements are equivalent.

- 1. G is Pfaffian.
- 2. G does not contain $K_{3,3}$ as a matching minor.
- 3. $\mathcal{D}(G, M)$ is non-even.
- 4. $\mathcal{D}(G, M)$ does not contain an odd bicycle as a butterfly minor.

It is not hard to see that a further correspondence induced by the bijection from Definition 8.2 is between strongly planar digraphs and planar bipartite graphs equipped with perfect matchings. In other words, given a bipartite graph G and a perfect matching M of G, the digraph $\mathcal{D}(G, M)$ is strongly planar if and only if G is a planar graph. Instead of proving this formally, we have illustrated this relation in Figure 8.3.



Figure 8.3: Left: An oriented grid where all arcs point from the empty to the filled vertices, equipped with a perfect matching. Right: The arising M-direction, a strongly planar digraph.

Given the correspondence between directed cycles in non-even digraphs and alternating cycles in bipartite Pfaffian graphs equipped with perfect matchings, it is natural to also translate Theorem 8.3 to bipartite graphs. For that purpose, the dichromatic number can be translated to a coloring concept for perfect matchings in graphs as follows.

Given a graph G and a perfect matching $M \in \mathcal{M}(G)$, an *M*-coloring of G with k colors is a function $c: M \to [k]$. An *M*-coloring is called *proper* if there is no *M*-alternating cycle whose matching edges are all of the same color, i.e., $c^{-1}(i)$ is the unique perfect matching of the subgraph of G induced by the endpoints of the edges in $c^{-1}(i)$ for all *i*. The *M*-chromatic number $\chi(G, M)$ of G is the smallest integer k such that G has a proper *M*coloring with k colors. By the correspondence of *M*-alternating cycles in G and directed cycles in $\mathcal{D}(G, M)$, we have $\chi(G, M) = \chi(\mathcal{D}(G, M))$ for any bipartite graph G with a perfect matching *M*.

We can now translate Theorem 8.3 using Theorem 8.9 to derive the following corollary.

Corollary 8.10. Let G be a bipartite graph with a perfect matching M. If $\chi(G, M) \ge 3$, then G contains $K_{3,3}$ as a matching minor.

In the context of *M*-colorings of graphs one can identify certain subsets of perfect matchings, namely the *forcing sets*. Given a perfect matching *M* of a graph, a subset $S \subseteq M$ of edges is called *forcing* if *M* is the unique perfect matching extending *S*. *S* is called *rigid* if *S* is the only matching in *G* on its vertex-set (in other words, there is

no *M*-alternating cycle with vertices in V(S)). Pause to note that given $S \subseteq M$, the matching S is forcing if and only if $M \setminus S$ is rigid. With these definitions, an *M*-coloring with k colors corresponds to a partition $M = S_1 \cup \cdots \cup S_k$ such that for any i, S_i is rigid.

The forcing number f(G, M) of a perfect matching M denotes the size of a smallest forcing set for M. This notion arises from resonance theory in chemistry. Bounds on the forcing numbers of perfect matchings have attracted quite some interest in the last three decades. We refer to [CC11] for a comprehensive survey on this topic.

With the above terminology we can reformulate Corollary 8.10 as follows:

Corollary 8.11. Every perfect matching M of a Pfaffian bipartite graph G can be partitioned into two disjoint partial matchings, which are simultaneously rigid and forcing.

This directly yields the following corollary.

Corollary 8.12. For any Pfaffian bipartite graph G and every perfect matching M of G, we have $f(G, M) \leq \frac{|M|}{2} = \frac{v(G)}{4}$.

This generalizes Theorem 2.9 in [CC11] from bipartite graphs without a(n ordinary) $K_{3,3}$ -minor to Pfaffian bipartite graphs, which is a more general class of graphs.

The definition of the matching chromatic number is not limited to bipartite graphs only. In Section 8.5 we therefore consider a generalisation of the above results to non-bipartite matching covered graphs. As these graphs bare a much more complicated structure than their bipartite cousins, we are not able to extend our coloring results in their full strength to the non-bipartite world. Even in the planar case there are graphs with perfect matchings that are not 2-colorable. A smallest example of such a graph is found in the *triangular prism*, which is the complement of C_6 . However, we are able to bring down the planar case to exactly this graph in the sense that planar graphs excluding $\overline{C_6}$ as a conformal bisubdivision are matching-2-colorable.

Theorem 8.13. Let G be a planar graph, and M a perfect matching of G. If G contains no conformal bisubdivision of $\overline{C_6}$, then $\chi(G, M) \leq 2$.

Structure of the chapter. In Section 8.2 we prepare and present the proof of Theorem 8.3. In Section 8.3 we derive Theorem 8.4 from Theorem 6.2 and prove Corollary 8.5. In Section 8.4 we prove Theorem 8.6. Finally, in Section 8.5 we study matching-colorings of non-bipartite graphs and prove Theorem 8.13. Section 8.6 contains concluding remarks.

8.2 2-Colorings of Non-Even Digraphs

This section is dedicated to the proof of Theorem 8.3. The key idea of our proof is to consider a minimal (with respect to the number of vertices) non-2-colorable non-even digraph. We introduce a number of local reductions of digraphs transporting 2-colorability while ensuring that the reduced digraph is still non-even and prove that for any non-even digraph with at least 3 vertices one of our reductions is applicable.

Each of our reductions can be applied in polynomial time and thus this technique implies a polynomial time algorithm for 2-coloring a non-even digraph.

We start with two splitting operations, reducing the 2-coloring problem to the strongly 2-connected non-even digraphs.

Definition 8.3. Let D, D_1 and D_2 be digraphs. Then D is called a 0-sum of D_1 and D_2 if there is a partition of V(D) into non-empty sets X and Y such that no arc of D has its head in X and its tail in Y, and $D_1 = D[X]$, $D_2 = D[Y]$.

We call a strongly connected digraph D the 1-sum of D_1 and D_2 at a vertex $v \in V(D)$ if there is a partition of $V(D) \setminus \{v\}$ into non-empty sets X and Y such that no arc in D has its head in X and its tail in Y, and such that D_1 arises from D by identifying $Y \cup \{v\}$ into a single vertex and D_2 arises by identifying $X \cup \{v\}$ into a single vertex. In both cases, we unify possible multiple occurences of parallel arcs into single arcs.

In the context of perfect matchings in bipartite graphs, the described reduction of Dto D_1 and D_2 corresponds to a so-called *tight cut contraction*. Let G be an undirected graph and $X \subseteq V(G)$. Recall that the *cut* around X, denoted by $\partial(X)$, is the set of all edges in G with exactly one endpoint in X. If G is matching covered and $|\partial(X) \cap M| = 1$ for every perfect matching $M \in \mathcal{M}(G)$, we call $\partial(X)$ a *tight cut*. If $\partial(X)$ is a tight cut and $|X|, |\overline{X}| \geq 2$, it is *non-trivial*. Identifying the cut set X of a non-trivial tight cut $\partial(X)$ into a single vertex (ignoring arising parallel edges) is called a *tight cut contraction* and the resulting graph G' can easily be seen to be matching covered again. Among many other things, tight cut contractions can be used to produce reductions of Pfaffian graphs as shown by Vazirani and Yannakakis.

Theorem 8.14 ([VY89], Theorem 4.2). Let G be a matching covered graph, $X \subseteq V(G)$ such that $\partial(X)$ is a non-trivial tight cut and G_1 , G_2 the two graphs obtained by the tight cut contractions of X and \overline{X} in G respectively. Then G is Pfaffian if and only if G_1 and G_2 are Pfaffian.

To combine the theory of tight cuts and digraphs we need to be able to translate between the two more smoothly. Given a bipartite graph G with bipartition $V(G) = A \cup B$ and a set $X \subseteq V(G)$ such that $|X \cap A| < |X \cap B|$, we call A the *minority* and B the *majority* of X, and analogously if the roles of A and B are reversed. Consider the following characterization of tight cuts in bipartite graphs.

Lemma 8.15 ([LDCM15], Proposition 5). Let $G = (A \cup B, E)$ be a bipartite matching covered graph and $X \subseteq V(G)$ of odd size. Then $\partial(X)$ is a tight cut if and only if $||X \cap A| - |X \cap B|| = 1$ and no vertex of the minority of X has a neighbour in \overline{X} .

Given a digraph D, in the following let us call a pair (X, Y) of vertex subsets a *directed* separation if $X \cup Y = V(D)$ and there is no edge with tail in $Y \setminus X$ and head in $X \setminus Y$. The *order* of the separation is defined as $|X \cap Y|$. While the following result is folklore, we provide a proof for completeness.

Lemma 8.16. Let $G = (A \cup B, E)$ be a bipartite matching covered graph, M a perfect matching in G and let $X \subseteq V(G)$. Moreover let $M_S := (E(G[S]) \cup \partial(X)) \cap M$ for $S \in \{X, \overline{X}\}$ and let v_e for $e \in M$ denote the vertex of the M-direction of G corresponding to the edge e. Then $\partial(X)$ is tight if and only if $(\{v_e \mid e \in M_X\}, \{v_e \mid e \in M_{\overline{X}}\})$ or $(\{v_e \mid e \in M_{\overline{X}}\}, \{v_e \mid e \in M_X\})$ is a directed separation of order 1 in $\mathcal{D}(G, M)$.

Proof. First suppose $\partial(X)$ is tight. By Lemma 8.15 no vertex of the minority of X has a neighbour in \overline{X} . By symmetry, we may assume that $B \cap X$ is the minority of X. By Theorem 8.7, the *M*-direction of *G* must be strongly connected, however there cannot exist an arc in $\mathcal{D}(G, M)$ with head v_e and tail $v_{e'}$ where $e \subseteq X$ and $e' \subseteq \overline{X}$ since the corresponding edge in *G* would link a vertex of $X \cap B$ to a vertex of $\overline{X} \cap A$. Hence every directed path from v'_e to v_e must contain the vertex v_f where *f* is the unique edge of *M* in $\partial(X)$. Thus $(\{v_e \mid e \in M_X\}, \{v_e \mid e \in M_{\overline{X}}\})$ is a directed separation and v_f is the unique vertex in the intersection of the two sets.

For the other direction let $(\{v_e \mid e \in M_X\}, \{v_e \mid e \in M_{\overline{X}}\})$ be a directed separation of order 1 in $\mathcal{D}(G, M)$. The other case follows analogously. Let f be the unique matching edge

corresponding to the cut vertex. Then every directed cycle in $\mathcal{D}(G, M)$ must contain v_f and has exactly one edge with endpoints in $\{v_e \mid e \in M_X\} \setminus \{v_f\}$ and $\{v_e \mid e \in M_{\overline{X}}\} \setminus \{v_f\}$. This means that every *M*-alternating cycle in *G* contains exactly two edges of $\partial(X)$, namely f and one non-matching edge. We have established that $|\partial(X) \cap M| = |\{f\}| = 1$, and so to prove that $\partial(X)$ is tight, we must show that any other perfect matching M' of G has the same number of edges on $\partial(X)$ as M. For this, observe that the symmetric difference $M\Delta M'$ decomposes into a vertex-disjoint union of cycles C_1, \ldots, C_t which are simultaneously M- and M'-alternating. Consequently, exchanging matching with nonmatching edges for each C_i one after the other ("flipping") transforms M into M'. Clearly, this operation can change the number of matching edges on $\partial(X)$ only if a cycle containing vertices of both X and \overline{X} is flipped, but according to the above, each such cycle must contain f, and so at most one C_j can intersect $\partial(X)$, and $E(C_j) \cap \partial(X) = \{f, f'\}$ for a non-matching edge f'. Flipping C_j now makes f' into a matching and f into a nonmatching edge. In any case, after having performed the sequence of flips, we thus obtain that $M' \cap \partial(X)$ consists of a single edge, and, hence, $\partial(X)$ must be tight.

From Theorem 8.14 and Lemma 8.16 we obtain the following corollary.

Corollary 8.17. Let D be a digraph and $i \in \{0, 1\}$ such that D is the *i*-sum of the digraphs D_1 and D_2 . Then D is non-even if and only if D_1 and D_2 are non-even.

Proof. For i = 0, this can be seen directly from the definition of an even digraph: D is non-even if and only if there is a subset $A \subseteq A(D)$ of arcs intersected an odd number of times by each directed cycle. However, the set of directed cycles in D consists of the directed cycles in $D[X] = D_1$ and $D[Y] = D_2$ for a partition (X, Y) as in Definition 8.3, because no directed cycle can pass trough X and Y at the same time. Thus, the above is the same as saying that there are arc sets $A_i \subseteq E(D_i)$, i = 1, 2, intersecting each directed cycle in D_i an odd number of times, which is the same as saying that D_1, D_2 are non-even.

For i = 1, the claim is a direct consequence of Lemma 8.16 and Theorem 8.14.

So 0- and 1-sums preserve non-evenness. Next, we need to make sure we can obtain an acyclic 2-coloring of D from acyclic 2-colorings of its sumands D_1 and D_2 .

Lemma 8.18. Let D be a non-even digraph and D_1 , D_2 digraphs such that D is the *i*-sum of D_1 and D_2 for $i \in \{0,1\}$. If D_1 and D_2 are 2-colorable, so is D.

Proof. Assume first that D is the 0-sum of $D_1 = D[X]$, $D_2 = D[Y]$ for a partition X, Y of V(D). Then the directed cycles in D are exactly the directed cycles in D_1 together with the directed cycles in D_2 , and thus any acyclic 2-coloring of D_1 joined with an acyclic 2-coloring of D_2 yields an acyclic 2-coloring of D.

Now assume D is the 1-sum of D_1 and D_2 at a vertex v, and let (X, Y) be the corresponding directed separation such that $X \cap Y = \{v\}$. Let v_1 be the vertex of D_1 obtained from identifying $Y \cup \{v\}$, and let v_2 be the vertex in D_2 obtained by identifying $X \cup \{v\}$. For $i \in \{1, 2\}$ let $c_i \colon V(D_i) \to \{0, 1\}$ be an acyclic 2-coloring of D_i . By possibly exchanging 0 and 1 in c_2 , we may assume that $c_1(v_1) = c_2(v_2)$. We define a coloring c for D as follows.

$$c(u) := \begin{cases} c_1(u), & u \in X \\ c_1(v_1) = c_2(v_2), & u = v \\ c_2(u), & u \in Y \end{cases}$$

To see that this defines an acyclic 2-coloring of D, assume towards a contradiction that C is a monochromatic directed cycle in D. If C stays within $X \cup \{v\}$ or $Y \cup \{v\}$, then it also

appears as a directed cycle in D_1 , or D_2 respectively, contradicting the feasibility of the 2-colorings c_1 and c_2 . Otherwise, C traverses vertices of both X and Y and thus, as there are no arcs starting in X and ending in Y, C also contains v. Moreover, C - v can be decomposed into exactly two directed paths P_1 and P_2 , one contained in X and the other in Y. Hence C corresponds to the directed cycles $C_i = P_i + v_i$ in D_i for each $i \in \{1, 2\}$ and both C_i must be monochromatic under their respective colorings c_i . This again violates the feasibility of the c_i . Consequently, c defines a coloring of D as desired.

Robertson et al. [RST99] defined in total five different sum operations which they used to prove a generation theorem for non-even digraphs. One of the consequences of this generation theorem is the following.

Theorem 8.19 ([Tho06], Corollary 5.4). Let D be a strongly 2-connected non-even digraph. Then $a(D) \leq 3v(D) - 4$.

Corollary 8.20. Any strongly 2-connected, non-even digraph D contains at least two vertices of out-degree 2.

Proof. Let n := v(D). By Theorem 8.19 we have a(D) < 3(n-1). If at most one vertex in D had out-degree less than 3 we would have $a(D) = \sum_{v \in V(D)} d^+(v) \ge 0 + 3(n-1)$, a contradiction, and so there are at least two vertices of out-degree at most, and thus, because D is strongly 2-connected, exactly two.

Besides arc deletions, butterfly contractions and 0- and 1-sums, we will use another special operation in order to further reduce our digraphs. If we encounter an out-degree 2 vertex v in a digraph D such that v is contained in at most one digon, we will need to delete some arcs incident with v in order to create a butterfly contractible arc. However, if v is contained in two different digons, we will directly contract the three digon vertices, namely v and the two vertices with which v forms a digon each, into a single vertex. While this is not a standard butterfly contraction, it is natural in the context of our proof and it preserves the property of being non-even, which we show below by using matching theory.

Note that bicontractions in matching covered graphs are a special case of tight cut contractions. To see this, consider X as the set of size 3 containing a degree 2 vertex v together with its two neighbours. Then $\partial(X)$ is tight since every perfect matching must match v to one of its neighbours and thus exactly one matching edge can and must leave X. Thus one can derive the following corollary from Theorem 8.14 or Theorem 8.9.

Corollary 8.21. Let G be a Pfaffian matching covered graph. Then every matching minor of G is Pfaffian.

Lemma 8.22. Let D be a non-even digraph with a vertex $v \in V(D)$ with $N^+(v) = \{v_1, v_2\}$ such that v induces a digon together with v_i for both $i \in \{1, 2\}$. Then the digraph D^* , obtained by first deleting all arcs of the form (u, v) with $u \notin \{v_1, v_2\}$, and then identifying v_1 , v and v_2 into a single vertex (while deleting occurring loops and identifying occurring parallel arcs into single arcs afterwards), is non-even as well.

Proof. Let D be the digraph together with the vertices v, v_1 , and v_2 as in the assertion. By Theorem 8.2, when deleting all incoming arcs of v with tails other than v_1 or v_2 we obtain a subdigraph D' which is non-even as well. Moreover, by Lemma 8.17, D' is non-even if and only if every strongly connected component of D' is non-even. Since v, v_1 and v_2 are contained in two digons sharing a vertex, they all must appear in the same strong component of D', say, D'_0 . It suffices to show that the identification of the three vertices $\{v, v_1, v_2\}$ in D'_0 preserves non-eveness.
8.2. 2-COLORINGS OF NON-EVEN DIGRAPHS

With D'_0 being strongly connected, there exists a bipartite matching covered graph G together with a perfect matching $M \in \mathcal{M}(G)$ such that $D'_0 = \mathcal{D}(G, M)$. We identify the vertices v, v_1 and v_2 of D'_0 with edges e_v, e_{v_1} and e_{v_2} , respectively, in M. Additionally let A and B be the two color classes of G. Then we denote by a_x the vertex of e_x in A and by b_x the vertex in B for all $x \in \{v, v_1, v_2\}$. Since v and v_1 form a digon in D'_0 , the edges $a_v b_{v_1}$ and $a_{v_1} b_v$ exist in G and, thus, together with e_v and e_{v_1} they form a conformal cycle of length 4. Therefore we can obtain a new perfect matching from M as follows.

$$M' := (M \setminus \{e_v, e_{v_1}\}) \cup \{a_v b_{v_1}, a_{v_1} b_v\}$$

Now consider $G - e_v$ and note that it still has M' as a perfect matching and that it is a matching minor of G (see Figure 8.4 for an illustration). By our assumptions, v has exactly two out- and two in-neighbours in D'_0 and therefore the two vertices a_v and b_v must be of degree 2 in $G - e_v$. Hence we can bicontract these two vertices and identify b_{v_1} , a_v , and b_{v_2} into $b_{v_1vv_2}$ and the other three vertices into $a_{v_1vv_2}$ respectively. Let us call the resulting bipartite graph G^* and denote the edge $a_{v_1vv_2}b_{v_1vv_2}$ by $e_{v_1vv_2}$. One can easily check that G^* still is matching covered and since it is a matching minor of G it must be Pfaffian by Corollary 8.21. Moreover, the strongly connected digraph $D^*_0 := \mathcal{D}(G^*, M^*)$ must be non-even. Since $M^* \setminus \{e_{v_1vv_2}\} = M' \setminus \{a_v b_{v_1}, a_{v_1} b_v, e_{v_2}\} = M \setminus \{e_{v_1}, e_v, e_{v_2}\}$ and the two edges e_{v_1} and $e_{v_1vv_2}$ can be identified (again see Figure 8.4) D^*_0 is isomorphic to the digraph obtained from D'_0 identifying the three vertices v, v_1 , and v_2 into one, and so the latter has to be non-even as well. From this we deduce that all strong components of D^* are non-even, proving the assertion.



Figure 8.4: The four steps of the contraction of v, v_1 , and v_2 in Lemma 8.22. The matching M' is given by dashed edges while the edges of M are thicker.

We are now ready to prove our main theorem, concluding this section.

Proof of Theorem 8.3. Assume towards a contradiction that there is a non-even digraph D that is not 2-colorable. Furthermore, let us assume D to be minimal (with respect to v(D)) with this property. Clearly $v(D) \geq 3$.

First observe that, due to Lemma 8.18, D is neither a 0-sum nor a 1-sum of some other non-even digraphs D_1 and D_2 . Hence, D does not have a directed cut or a cut vertex, and must therefore be strongly 2-connected. By Corollary 8.20 there exists a vertex $v \in V(D)$ with $d^+(v) = 2$. Let $e_1 = (v, v_1)$ and $e_2 = (v, v_2)$ be the two outgoing arcs of v. We now distinguish two cases:

Case 1: Both arcs e_1 and e_2 are contained in digons.

If e_1 and e_2 are contained in digons, we can construct a non-even digraph D^* from D by applying the operation from Lemma 8.22 on v and its two out-neighbours. First, we delete all incoming arcs of v except (v_1, v) and (v_2, v) from the digraph and then contract v_1 , v, and v_2 into a single vertex. Since $|V(D^*)| = |V(D)| - 2$ and D^* is non-even, by the minimality of D, D^* admits an acyclic 2-coloring $c^* \colon V(D^*) \to \{0,1\}$. Denote by $u_{v_1v_2}$

the vertex of D^* into which v_1 , v and v_2 were identified. We now define a 2-coloring for the vertices $x \in V(D)$ as follows.

$$c(x) := \begin{cases} c^*(u_{v_1vv_2}), & x \in \{v_1, v_2\} \\ 1 - c^*(u_{v_1vv_2}), & x = v \\ c^*(x), & \text{otherwise} \end{cases}$$

By assumption, D is not acyclically 2-colorable and thus there must be a directed cycle C whose vertices receive the same color from c. Moreover, C must avoid v, since any directed cycle in D containing v must either contain v_1 or v_2 and thus, by the definition of c, cannot be monochromatic. Consequently, C must be contained in D - v. By identifying possible occurrences of v_1 or v_2 with $u_{v_1vv_2}$, the existence of a closed directed monochromatic walk C^* in D^* follows. Note that v_1 and v_2 do not form a digon, as otherwise v, v_1 and v_2 would be an odd bicycle in D, contradicting the assumption that D is non-even. Hence, the walk C^* is non-trivial (consists of at least two vertices) and therefore contains a directed cycle which, in turn, must also be monochromatic with respect to c^* . However, the existence of such a cycle contradicts the choice of c^* .

Case 2: At least one of the arcs e_1 or e_2 is not contained in a digon.

Without loss of generality assume e_1 to not be part of a digon in D. We now delete all arcs with endpoints v and v_2 , thereby obtaining a non-even digraph in which v has a single out-going arc, which is e_1 . With this, e_1 is now butterfly contractible. Let $D' = D/e_1$ be the digraph obtained by contracting e_1 and let x_{e_1} denote the contraction vertex. Since the non-even digraphs are closed under butterfly-minors, D' is again non-even. Moreover, as v(D') = v(D) - 1, D' must admit an acyclic 2-coloring $c' \colon V(D') \to \{0,1\}$ by the minimality of D. Similar to the first case we use c' to define a 2-coloring c of D as follows:

$$c(x) := \begin{cases} c'(x_{e_1}), & x = v_1 \\ 1 - c'(v_2), & x = v \\ c'(x), & \text{otherwise.} \end{cases}$$

Again, we assumed D to not be 2-colorable and thus there must be a monochromatic (with respect to c) directed cycle C in D. If C contains v, it cannot contain v_2 as $c(v_2) \neq c(v)$. Therefore, it must contain the arc e_1 . Since e_1 is not contained in a digon we have $|C| \geq 3$ and thus there exists a directed cycle C' in D' through x_{e_1} with $V(C') \setminus \{x_{e_1}\} = V(C) \setminus \{v, v_1\}$. By definition of c, C' must be monochromatic with respect to c', which yields the desired contradiction in this case. Otherwise, C does not contain v. Then, possibly after replacing v_1 with x_e , C again yields a directed cycle in D' which, again, has to be monochromatic with respect to c', contradicting our choice of c'.

The proof of Theorem 8.3 yields a polynomial time algorithm to find an acyclic 2coloring of a non-even digraph. One first reduces a digraph D into its strong components, then finds the cut vertices and decomposes D into strongly 2-connected digraphs using 1-sums. Each of these strongly 2-connected digraphs can then be colored as follows:

Repeatedly, one searches for an out-degree 2 vertex. If such a vertex is found, we can reduce the digraph as described in Case 1 and Case 2 of the proof.

After reiterating these reduction steps we have reduced the strongly 2-connected digraph to a digraph on one or two vertices, which we can trivially 2-color.

Then, by reversing the reductions step by step, we can successively extend the acyclic 2-coloring to finally obtain acyclic 2-colorings of each of the strongly 2-connected digraphs we have decomposed D into. Finally, we can join these acyclic 2-colorings together to obtain an acyclic 2-coloring of D as described in Lemma 8.18.

Additionally, the works of Robertson et. al. and McCuaig [RST99, McC04] imply polynomial time algorithms to recognize non-even digraphs. Hence, given a digraph D we can decide whether it is non-even and then find an acyclic 2-coloring in polynomial time.

8.3 Complete Butterfly Minors and Zero Sum Cycles via Dichromatic Number

In this section we give the proofs of Theorem 8.4 and Corollary 8.5.

Let us first consider Theorem 8.4. We prove the following Lemma, which states that every sufficiently large strong complete minor also contains a complete butterfly-minor of given order. Note that once having established the Lemma, the assertion of Theorem 8.4 is then a direct consequence of Theorem 6.2 from Chapter 6.

Lemma 8.23. Let $t \in \mathbb{N}$, and let D be a strong \vec{K}_{2t} -minor-model. Then D contains \vec{K}_t as a butterfly-minor.

Proof. Let D be a strong \vec{K}_{2t} -minor model and let $\{X_1^+, X_1^-, \ldots, X_t^+, X_t^-\}$ be a corresponding partition of V(D) into 2t branch sets. Then for every $i \in \{1, \ldots, t\}$ there exist $r_i^+ \in X_i^+$ and $r_i^- \in X_i^-$ such that $(r_i^-, r_i^+) \in A(D)$. Since $D[X_i^-]$ and $D[X_i^+]$ are strongly connected digraphs, there exist¹ oriented spanning trees $T_i^- \subseteq D[X_i^-]$ and $T_i^+ \subseteq D[X_i^+]$ such that T_i^- is an in-arborescence rooted at r_i^- and T_i^+ is an out-arborescence rooted at r_i^+ . Let us consider the spanning subdigraph D' of D consisting of the arcs contained in

$$T := \bigcup_{i=1}^{t} \left(\{ (r_i^-, r_i^+) \} \cup A(T_i^+) \cup A(T_i^-) \right),$$

as well as all arcs of D starting in X_i^+ and ending in X_j^- for $i \neq j$. Then every arc of D' contained in T is either the unique arc in D' emanating from its tail or the unique arc in D' entering its head. It follows that all arcs in T are butterfly-contractible. Note that the contraction of an arc does not affect the butterfly-contractibility of other arcs, hence the digraph D'/T, obtained from D' by successively contracting all arcs in T, is a butterfly-minor of D. The vertices of D'/T can be labelled v_1, \ldots, v_t , where v_i denotes the vertex corresponding to the contraction of the (weakly) connected component of D' inside $X_i^+ \cup X_i^-$. As D is a strong \vec{K}_{2t} -minor model, by definition of D' for every $(i, j) \in \{1, \ldots, k\}^2$ with $i \neq j$, there exists an arc in D' starting in X_i^+ and ending in X_j^- . Therefore, D'/T is a butterfly-minor of D isomorphic to \vec{K}_t , concluding the proof.

Proof of Theorem 8.4. From the preceding lemma and Theorem 6.2 it follows that

$$bm_{\vec{\chi}}(t) \le sm_{\vec{\chi}}(2t) \le 2m_{\chi}(2t) - 1 = O(2t(\log\log(2t))^6) = O(t(\log\log t)^6).$$

Let us now move on to colorings of A-non-zero digraphs. In order to establish Corollary 8.5, we use of the following result proved recently by Tamás Mészáros and myself.

Theorem 8.24 (cf. [MS21b]). Let (A, +) be a finite Abelian group. Then $K_{8|A|}$ is A-zero. In other words, for every A-arc-labelling of $\vec{K}_{8|A|}$ there exists a directed cycle whose arc-labels sum up to zero.

¹Such trees can easily be obtained by considering a breadth-first in-search (resp. out-search) starting from r_i^- (resp. r_i^+).

The next Lemma shows that in the same way as non-even digraphs form a butterflyminor closed class of digraphs, for every fixed Abelian group (A, +) also the A-non-zero digraphs are closed under taking butterfly minors.

Lemma 8.25. Let (A, +) be a finite Abelian group, and let D, D' be digraphs such that $D \succeq_b D'$. If D is A-non-zero, then so is D'.

Proof. It is clear that any zero-sum free A-arc-labelling of a digraph, restricted to the arcset of any subdigraph, forms a zero-sum free A-arc-labelling of this subdigraph as well. It is therefore sufficient to prove the assertion of the lemma in the case that D' is obtained from D by the contraction of one contractible arc.

So suppose that $e = (u, v) \in A(D)$ is a butterfly-contractible arc in D such that D' = D/e. W.l.o.g. we may assume $d_D^+(u) = 1$. Then D' may be represented as follows.

$$V(D') = V(D) \setminus \{u\}, A(D') = A(D-u) \cup \{(x,v) | x \in N_D^-(u) \setminus \{v\}\}.$$

Now suppose that D is A-non-zero and let $w : A(D) \to A$ be a zero-sum free A-arclabelling of D. Define an arc-labelling w' of D' by putting w'(e) := w(e) for $e \in A(D-u)$ and w'(x,v) := w(x,u) + w(u,v) (summation in (A, +)) for every $x \in N_D^-(u) \setminus \{v\}$. We claim that w' is a zero-sum free arc-labelling of D'. Indeed, let C by any directed cycle in D'. If C does not contain an arc of the form (x,v) with $x \in N_D^-(u) \setminus \{v\}$, then C is also a directed cycle in D whose arc-labels in w and w' are equivalent, and hence the sum of its arc-labels must be non-zero. Next suppose that $(x,v) \in A(C)$ for some $x \in N_D^-(u) \setminus \{v\}$. By replacing the arc (x,v) of C with the directed path x, (x,u), u, (u,v), v in D, we obtain a directed cycle C^* in D. It now follows by definition of w' that the total sum of arc-labels of C with respect to w' equals the total sum of arc-labels of C^* with respect to w, and hence must be non-zero. This shows that w' is indeed zero-sum free, and as required, it follows that D' is A-non-zero as well. This completes the proof of the lemma.

We now conclude this section by deriving Corollary 8.5 as a direct consequence of Theorem 8.4, Theorem 8.24, and Lemma 8.25.

Proof of Corollary 8.5. Let (A, +) be a finite Abelian group such that $|A| \geq 3$, and let D be any A-non-zero digraph. By Lemma 8.25, every butterfly-minor of D is A-non-zero. By Theorem 8.24, the complete digraph $\vec{K}_{8|A|}$ is A-zero, and hence it follows that $D \not\geq_b \vec{K}_{8|A|}$. We may thus apply Theorem 8.4 to D and obtain:

$$\vec{\chi}(D) < bm_{\vec{\chi}}(8|A|) = O(8|A|(\log\log(8|A|))^6) = O(|A|(\log\log|A|)^6),$$

proving the assertion of the corollary.

8.4 List Colorings of Non-Even Digraphs

In this section, we prove Theorem 8.6. It will be convenient to prove the following (slightly) strengthened version.

Theorem 8.26. Let D be a non-even digraph. Then for any choice of a designated vertex $v_0 \in V(D)$, D is \mathcal{L} -list colorable for every list assignment $\mathcal{L} = \{L(v) | v \in V(D)\}$ fulfilling $|L(v_0)| = 1$ and $|L(v)| \ge 3$ for all $v \in V(D) \setminus \{v_0\}$.

Proof. Assume towards a contradiction that there is a non-even digraph D which does not satisfy the assertion, and assume D to be chosen minimal with respect to the number

of vertices. Let in the following \mathcal{L} be a fixed list assignment for D, where $|L(v_0)| = 1$ for some designated $v_0 \in V(D)$, $|L(v)| \geq 3$ for all $v \in V(D) \setminus \{v_0\}$, and such that D is not \mathcal{L} -choosable. Clearly, we have $v(D) \geq 3$.

We first show that D must be strongly 2-connected: Assume for a contradiction that there is a directed separation of order $i \in \{0, 1\}$ in D. By Lemma 8.17, we find that there are non-even digraphs D_1 and D_2 with fewer vertices than D such that D is the *i*-sum of D_1 and D_2 . By the assumed minimality of D, we know that D_1 and D_2 both satisfy the assertion of the Theorem.

If i = 0, consider a partition (X, Y) of V(D) such that $D_1 = D[X], D_2 = D[Y]$ and no arc in D starts in Y and ends in X.. Restricting \mathcal{L} to X resp. Y defines list assignments for D_1 and D_2 (each with at most one list of size less than 3), and we find that D_j admits a choice function c_j for j = 1, 2 that defines an acyclic digraph coloring and satisfies $c_j(x) \in L(x)$ for all $x \in V(D_j)$. Let c be the common extension of c_1, c_2 to D, i.e.,

$$c(x) := \begin{cases} c_1(x), & x \in X \\ c_2(x), & x \in Y \end{cases}$$

now defines a valid choice of colors for D without a monochromatic directed cycle, proving that D is \mathcal{L} -choosable. This is a contradiction to our initial assumption.

If i = 1, let $w \in V(D)$ be such that D is the 1-sum of D_1 and D_2 along w. Consider a partition (X, Y) of $V(D) \setminus \{w\}$ such that no arc in D has its head in X and its tail in Y, and such that D_1 arises from D by identification of $Y \cup \{w\}$ into a single vertex v_1 , and D_2 by identification of $X \cup \{w\}$ into a vertex v_2 .

We have that $v_0 \in X \cup \{w\}$ or $v_0 \in Y \cup \{w\}$. Assume for the following that $v_0 \in X \cup \{w\}$, the other case works symmetrically. Define an assignment \mathcal{L}_1 of lists to the vertices of D_1 according to $L_1(x) := L(x)$ for all $x \in X$ and $L_1(v_1) := L(w)$. Because D_1 satisfies the assertion, we find a choice function c_1 which defines an acyclic digraph coloring of D_1 while satisfying $c_1(x) \in L(x), x \in X$, and $\tilde{c} := c_1(v_1) \in L(w)$. Now define a list assignment \mathcal{L}_2 for D_2 according to $L_2(x) := L(x)$ for $x \in Y$ and $L_2(v_2) := \{\tilde{c}\}$. Because we have $|L_2(x)| = |L(x)| \ge 3$ for all $x \in Y = V(D_2) \setminus \{v_2\}$, we can apply the assertion to D_2 and thus find a choice function c_2 on $V(D_2)$ satisfying $c_2(x) \in L(x)$ for all $x \in Y$ and $c_2(v_2) = \tilde{c} = c_1(v_1)$. Now define a choice function c on V(D) by

$$c(x) := \begin{cases} c_1(x), & x \in X \\ \tilde{c}, & x = w \\ c_2(x), & x \in Y \end{cases}$$

By the above it is clear that we have $c(x) \in L(x)$ for all $x \in V(D)$. Because D is not \mathcal{L} -choosable, this implies that there is a directed cycle C in D which is monochromatic under c. Because c_1 and c_2 are a digraph colorings of D_1 and D_2 , C must contain vertices of both X and Y and therefore must visit w as well as exactly one arc with tail in X and head in Y. Therefore, identifying all vertices in $Y \cup \{v\}$ on C into a single vertex results in a directed cycle in D_1 , which has to be monochromatic as well. This finally is a contradiction to the definition of c_1 .

As both cases led to a contradiction, for the rest of the proof we may assume that $v(D) \geq 3$ and D is strongly 2-connected. Applying Corollary 8.20 we find that there is a vertex $u \in V(D) \setminus \{v_0\}$ of out-degree two. Clearly, D - u is non-even as well and has less vertices, so the minimality of D implies that, for the induced list assignment $\mathcal{L}' := \{L(x) | x \in V(D) \setminus \{u\}\}$, there is a choice function c' which defines an acyclic digraph coloring of D-u. Let u_1, u_2 be the two out-neighbours of u. As $|L(u) \setminus \{c'(u_1), c'(u_2)\}| \geq 1$, we can extend c' to a choice function c on V(D) such that $c(x) = c'(x) \in L(x)$ for all

 $x \in V(D) \setminus \{u\}$ and $c(u) \in L(u) \setminus \{c(u_1), c(u_2)\}$. Because D is by initial assumption not \mathcal{L} -choosable, this implies that there is a directed cycle in D which is monochromatic with respect to c. Since c' defined a valid digraph coloring, this is only possible if the cycle traverses u and thus one of the arcs (u, u_1) or (u, u_2) . However, this gives a contradiction to the fact that both of these arcs are bi-colored.

This final contradiction shows that our initial assumption was false and concludes the proof of the Theorem. $\hfill \Box$

8.5 Non-Bipartite Graphs

In the previous sections we were concerned with digraphs which, via the bijection described in Definition 8.2 correspond exactly to the bipartite graphs with perfect matchings. However, a matching covered graph does not need to be bipartite. In fact, most parts of (bipartite) matching theory directly translate into the world of general matching covered graphs. This includes, especially, tight cuts, their contractions, and Pfaffian orientations.

In particular, the *M*-chromatic number is defined on all graphs. By Corollary 8.10 every bipartite Pfaffian graph has *M*-chromatic number at most 2 for every perfect matching. A natural question to ask would be whether this generalizes to all (also non-bipartite) Pfaffian graphs. To this question there exists a rather easy negative answer:

The triangular prism is the complement $\overline{C_6}$ of the 6-cycle. It is planar and therefore



Figure 8.5: The triangular prism $\overline{C_6}$ together with a perfect matching M.

Pfaffian, but when considering the perfect matching M from Figure 8.5, one can see that any two of the three edges in M lie together on a 4-cycle. Hence, in any M-coloring, no two of the three edges may receive the same color and therefore $\chi(\overline{C_6}, M) = 3$.

In Corollary 8.10 we studied the matching chromatic number for a class closed under matching minors, so a natural next step would be to consider a subclass of the $\overline{C_6}$ -matching minor-free graphs. The triangular prism is one of two graphs appearing in a fundamental theorem by Lovász on non-bipartite matching covered graphs.

Theorem 8.27 (Lovász [Lov87]). Every non-bipartite matching covered graph contains a conformal bisubdivision of K_4 or $\overline{C_6}$.

A matching covered graph without a non-trivial tight cut is called a *brace* if it is bipartite and a *brick* otherwise. In his seminal paper [Lov87], Lovász introduced a decomposition procedure, known under the name *tight cut decomposition*, which, given a matching covered graph, searches for non-trivial tight cuts, computes both tight cut contractions, and iterates this for both reduced matching covered graphs, until a list of bricks and braces, which are not reducible any more, is obtained. Among many other things, Lovász proved that the list of bricks and braces does not depend on the chosen order in which the tight cuts are contracted. As the following theorem shows, braces correspond exactly to the strongly 2-connected digraphs.

Theorem 8.28 (Lovász and Plummer [LP68]). A bipartite graph G is a brace if and only if it is 2-extendable.

Bricks have a more complicated structure and although every 2-extendable graph is either a brick or a brace as seen in Theorem 8.29 below, there are bricks that are not 2-extendable. For an example of such a brick consider the triangular prism.

Theorem 8.29 (Plummer [Plu80]). Let G be a 2-extendable graph. Then, G is either a brace or a brick.

There exists a generalisation of tight cuts that is useful for the study of bricks. Given a matching covered graph G and a set $X \subseteq V(G)$ we call the graph G_X obtained from Gby identifying X into a single vertex the X-contraction of G. Now a cut $\partial(X)$ is called separating if both G_X and $G_{\overline{X}}$ are matching covered.

Theorem 8.30 (de Carvalho, Lucchesi, and Murty [dCLM02b]). Let G be a matching covered graph and $X \subseteq V(G)$. The cut $\partial(X)$ is separating if and only if for every edge $e \in E(G)$ there is a perfect matching M_e of G containing e such that $|\partial(X) \cap M_e| = 1$.

In what follows we call a matching covered graph *solid* if every non-trivial separating cut already forms a tight cut.

One can easily check the following lemma on bipartite graphs, showing that any bipartite matching covered graph is solid.

Lemma 8.31 (de Carvalho, Lucchesi, Kothari, and Murty [LDCKM18]). Let G be a bipartite matching covered graph. Then $\partial(X)$ is separating if and only if it is tight.

Moreover, being solid is preserved by tight cut contractions (cf. [dCLM02b]) and thus a matching covered graph is solid of and only if all of its bricks are solid.

Please note that even bricks may contain non-trivial separating cuts. Again consider the triangular prism from Figure 8.5 and take a cut around one of the two triangles. Such a cut is separating. In fact, the existence of a prism as a conformal bisubdivision immediately implies the existence of a non-trivial and non-tight separating cut.

Lemma 8.32 (cf. [LDCKM18]). Every solid graph is $\overline{C_6}$ -free.

The goal of this section is to establish an extension of Corollary 8.10 to non-bipartite matching covered graphs in the form of a conjecture.

Conjecture 8.2. Let G be a solid and Pfaffian graph and M a perfect matching of G. Then $\chi(G, M) \leq 2$.

To provide some evidence towards Conjecture 8.2, the remainder of this section is dedicated to proving Theorem 8.13, which settle the planar case of this conjecture.

For this we first establish a more general version of Lemma 8.18 by proving it directly for tight cut contractions. We will need a bit of notation here. If G is matching covered, M a perfect matching, and G_X is a tight cut contraction of $\partial(X)$ with contraction vertex v_X , we denote by M_X the perfect matching $\{e \in M \mid e \subseteq V(G_X)\} \cup \{uv_X\}$ where u is the unique vertex of X covered by the edge of M in $\partial(X)$.

Lemma 8.33. Let G be a matching covered graph, $\partial(X)$ a non-trivial tight cut in G and M a perfect matching. If $\chi(G_X, M_X) \leq 2$ and $\chi(G_{\overline{X}}, M_{\overline{X}}) \leq 2$, then $\chi(G, M) \leq 2$.

Proof. For $S \in \{X, \overline{X}\}$ let c_S be a matching 2-coloring of M_S in G_S . Let $e_S \in M_S$ be the edge covering the contraction vertex. Then we can rename the colors for c_X and $c_{\overline{X}}$ such that $c_X(e_X) = c_{\overline{X}}(e_{\overline{X}})$ and we define a coloring for M as follows.

$$c(e) := \begin{cases} c_X(e), & e \in M_X \\ c_X(e_X) = c_{\overline{X}}(e_{\overline{X}}), & e \in \partial(X) \cap M \\ c_{\overline{X}}(e), & e \in M_{\overline{X}} \end{cases}$$

Suppose G contains an M-alternating cycle C that is monochromatic with respect to c. If V(C) is a subset of either X or \overline{X} , by definition of c, C must be a monochromatic cycle in either G_X or $G_{\overline{X}}$ and, thus, C must cross $\partial(X)$. Since $\partial(X)$ is tight, $C - (\partial(X) \cap E(C))$ contains exactly 2 components. Each of them is a path of even length and M covers all vertices but exactly one endpoint. Moreover, each of these paths forms, together with the corresponding edges in $\partial(X)$, an M_S -alternating cycle in their respective contraction G_S , for both $S \in \{X, \overline{X}\}$. By definition of c, these two cycles must also be monochromatic which ultimately contradicts the choice of the colorings $c_X, c_{\overline{X}}$ and completes the proof. \Box

Using the tight cut decomposition and the above Theorem, Theorem 8.13 reduces to the task of showing that every perfect matching of a solid planar brick or planar brace is 2-colorable. The brace case is of course taken care of by Corollary 8.10 and thus our only concern are the solid planar bricks². By Lemma 8.32 we only have to consider $\overline{C_6}$ -free planar bricks. Kothari and Murty (cf. [KM16]) gave a precise description of these bricks. To state this description, we need some additional terminology.

A wheel graph W_k with $k \ge 3$ is called odd if k is odd. It is not hard to see that every odd wheel is a brick.



Figure 8.6: The tricorn with one perfect matching of type I (solid) and one of type II (dashed).

Let (u_1, u_2, \ldots, u_k) and (v_1, v_2, \ldots, v_k) be two disjoint paths with $k \ge 2$. The graph SC_{2k+2} obtained from the union of these paths by adding the edges $u_i v_i$ for all $i \in \{1, \ldots, k\}$, two new vertices x and y joined by an edge and the edges xu_1, xv_1, yu_k, yv_k , is called a *staircase of order* 2k + 2. The graph SC_{2k+2} is a brick and SC_6 is isomorphic to the triangular prism.

The following is the main result of Kothari and Murty from [KM16].

Theorem 8.34 (Kothari and Murty [KM16]).

- 1. A matching-covered graph is $\overline{C_6}$ -free if and only if all the bricks and braces in its tight cut decomposition are $\overline{C_6}$ -free.
- 2. The only planar $\overline{C_6}$ -free bricks are the odd wheels, the staircases of order divisible by 4 and the tricorn (see Figure 8.6).

If we have a planar and matching covered graph G that does not contain a conformal bisubdivision of $\overline{C_6}$, by Theorem 8.34 the only bricks G can have are odd wheels, staircases of orders divisible by 4 and tricorns. Along with these bricks, G can have any set of planar braces. Planar braces are Pfaffian and thus by Corollary 8.10 matching-2-colorable. In the following we prove Theorem 8.13 by showing that all three families of bricks mentioned above are matching 2-colorable as well.

²Here we use the fact that given a planar matching-covered graph G, arbitrary contractions, and thus also all bricks and braces in the tight cut decomposition of G, are planar as well.

Proof (of Theorem 8.13). As we have seen, it suffices to show that the perfect matchings of the odd wheels, staircases of order 4k and the tricorn are 2-colorable.

Odd Wheels. For $K_4 = W_3$ we have exactly two edges in every perfect matching and thus are done. Let $k \ge 4$ be any odd number. For the odd wheel W_k on k + 1 vertices, let x be the unique vertex of degree k. Clearly every perfect matching M has to cover x with an edge, say, e_x^M , and every other matching edge lies on the cycle induced by the neighbourhood of x. Consider the graph induced by $\bigcup M \setminus \{e_x^M\}$. Since N(x) induces a cycle, this graph is a path and thus every M-alternating cycle in W_k must contain e_x^M . Hence, by coloring e_x^M with 0 and every other edge of M with 1 we have found a matching 2-coloring for M in W_k .

Staircases of Order 4k. For the staircases SC_{4k} we give a 2-coloring $c: E(SC_{4k}) \rightarrow$ $\{0,1\}$ of the edges that induces a matching-2-coloring on every perfect matching. Let xy be the unique edge with endpoints in two disjoint triangles. Let (u_1, \ldots, u_{2k-1}) be the path from the construction of S_{4k} not on the outer face and assume xu_1 to be an edge of SC_{4k} . We color xy with 0. Then, going counter-clockwise around the outer face, we assign 0 as the color of the edges xv_1 and v_1v_2 , the next two edges receive the color 1, then two times color 0 and so forth until the edge $v_{2k-1}y$ is colored. Since SC_{4k} is of order 4k we color 2k-2 edges this way and the last two edges receive color 1. With this the path $(x, v_1, \ldots, v_{2k-1}, y)$ on the outer face is colored. We set $c(u_i u_{i+1}) := 1 - c(v_i v_{i+1})$ for $i \in \{1, ..., 2k - 2\}$ and $c(xu_1) := 1$ while $c(yu_{2k-1}) := 0$. At last we need to color the spokes. Let $c(v_i u_i) := i \mod 2$ for $i \in \{1, \ldots, 2k-1\}$. For an illustration consider Figure 8.7. To show that c induces a matching-2-coloring for every perfect matching, we must show that there is no conformal cycle C such that every second edge has the same color. Assume for a contradiction that $SC_{4k}[c^{-1}(0)]$ contains a conformal cycle C. However, this graph contains a single even length cycle and this cycle contains exactly the vertices incident with at most one edge of color 1 in G. Therefore $V(G) \setminus V(C)$ is a stable set and thus C is not conformal, a contradiction. Thus C must contain an edge of color 1 and therefore, by construction, also two consecutive such edges. Consequently, we must have i = 1, and every second edge must be of color 1. There does not exist a path of length 5 in SC_{4k} such that the first, third, and fifth edge are colored with 1, hence C must have length 4. Clearly none of the 4-cycles contains two disjoint edges of the same color, which yields the desired contradiction, a cycle C with the stated properties cannot exist.



Figure 8.7: The staircase of order 16 together with a 2-coloring of the edges inducing a matching-2-coloring for every perfect matching. The solid edges are considered to be of color 0, while the dashed ones are of color 1.

Tricorn. For the tricorn we first observe that we can classify its perfect matchings into two types. Any perfect matching either contains exactly one edge on the outer face (compare Figure 8.6) that belongs to a triangle or none. If we fix such an edge e on the outer face belonging to a triangle for our perfect matching M_1 , the remaining edges of M_1 are uniquely determined. This can be seen as follows: Taking an edge from one of the triangles forces us to match the remaining vertex of said triangle to the middle vertex. Then the remaining neighbors of the middle vertex have to be matched within their respective triangles in such a way that the remaining two vertices are adjacent. There is only one way to do this after e has been chosen and thus $\{e\}$ is in fact a forcing set for M_1 . Hence coloring e with 0 and all other edges of M_1 with 1 yields the desired coloring. We call such a matching type I.

A matching of type II is a matching not containing any edge on the outer face belonging to a triangle. Note that any perfect matching must contain two edges of the outer face. So let e_1 and e_2 be these two edges. One of the three triangles contains an endpoint from both e_1 and e_2 , and its third vertex has to be matched to the middle one. This is already enough to determine the last two edges and we obtain M_2 . Hence $\{e_1, e_2\}$ is a forcing set of M_2 and, since the tricorn contains no 4-cycle, by coloring e_1 and e_2 with 0 and the rest of M_2 with 1 we are done.

By the above discussion it is clear that any perfect matching of the tricorn is either of type I or II and this concludes the proof. $\hfill \Box$

It is easy to see that any cut around a triangle in the tricorn or a staircase is separating. Moreover, one can check that the odd wheels are solid. Hence we have the following.

Corollary 8.35 (de Carvalho, Lucchesi, and Murty [dCLM02a]). The only planar solid bricks are the odd wheels.

Corollary 8.36. It holds that $\chi(G, M) \leq 2$ for every planar solid graph G and a perfect matching M of G.

8.6 Conclusion

In Section 8.3 we have shown that $bm_{\vec{\chi}}(t) \leq 2m_{\chi}(2t) - 1 = O(t(\log \log t)^6)$. Similar as in the case of strong minors, for every undirected graph G with $\chi(G) \geq bm_{\vec{\chi}}(t)$ we have $\vec{\chi}(\vec{G}) = \chi(G)$ and hence $\vec{G} \succeq_b \vec{K}_t$, implying that $G \succeq K_t$. This shows that $m_{\chi}(t) \leq bm_{\vec{\chi}}(t)$ for any $t \geq 1$, and it follows that the asymptotic behavior of $bm_{\vec{\chi}}(t)$ is the same as that of $m_{\chi}(t)$. Concerning explicit lower bounds, the best we can show is that $bm_{\vec{\chi}}(t) \geq t + 1$ for any $t \geq 3$, which follows by considering the biorientation of the undirected graph $G = K_{t+2} - C_5$ obtained from the complete graph on t+2 vertices by removing the edges of a 5-cycle. In fact, $\vec{\chi}(\vec{G}) = \chi(G) = t$ but \vec{G} does not contain \vec{K}_t as a butterfly-minor. As Hadwiger's conjecture is known to be true for $t \leq 6$, we have the bounds

$$t+1 \le bm_{\vec{\mathbf{v}}}(t) \le 4t-1$$

for $3 \le t \le 6$. The upper bound should certainly not be tight, and it seems plausible that the truth lies with the lower bound.

Conjecture 8.3. For every $t \ge 3$, we have $bm_{\vec{x}}(t) = t + 1$.

Note that since topological minors specialise butterfly-minors, the case t = 3 of Conjecture 8.3 would follow from Conjecture 7.2 stated in Chapter 7.

In Section 8.4 we have shown that the list dichromatic number of non-even digraphs is at most 3, that is, digraphs in the class \mathcal{D}_2 have bounded choice number. It is natural to ask whether this result at least qualitatively generalizes to the minor-closed classes $\mathcal{D}_k, k \geq 3$. Since the classes \mathcal{D}_k are sandwiched between the classes \mathcal{K}_t of digraphs excluding K_t as a butterfly-minor (compare the discussion in the introduction of this chapter), the following question comes up naturally.

Problem 8.2. Is there a function $f : \mathbb{N} \to \mathbb{N}$ such that every digraph D with $\vec{\chi}_{\ell}(D) \ge f(t)$ contains \vec{K}_t as a butterfly minor?

In the undirected case, it is known that K_t -minor free graphs have choice number at most $O(t(\log \log t)^6)$ (cf. [Pos20]), so here the same asymptotic bound as for the chromatic number applies. Yet, problem 8.2 seems to be extremely tricky, since none of the standard approaches for bounding the choice number applies. As mentioned earlier, already noneven digraphs may have unbounded minimum out- and in-degree, and so a degeneracycoloring approach seems infeasible. Also the methods for bounding $sm_{\vec{\chi}}(t)$ and $bm_{\vec{\chi}}(t)$ we have seen previously do not apply for list coloring, since they are eventually based on the fact that $\vec{\chi}(D) \leq \vec{\chi}(D[X]) + \vec{\chi}(D[Y])$ holds for any digraph D and a partition $X \cup Y = V(D)$ of its vertex-set. The same inequality however is false for the list dichromatic number: Bipartite digraphs can be partitioned into two independent sets but still may have arbitrarily large choice number [BHL18].

Due to the existence of infinite antichains (such as the odd bicycles) in the butterflyminor order of digraphs, we believe that for larger values of k, possibly no very simple description of the forbidden butterfly minors for \mathcal{D}_k can be obtained. Looking at the case k = 2, this drastically changed when moving from digraphs to the corresponding bipartite graphs, where we only needed to exclude $K_{3,3}$ as a matching minor. While by now the $K_{3,3}$ -matching minor-free bipartite graphs (that is, the Pfaffian bipartite graphs) have many equivalent characterizations and can be recognized in polynomial time, not much is known about the classes of $K_{t,t}$ -matching minor-free graphs with $t \geq 4$. Clearly, the complete bipartite graph $K_{t,t}$ has M-chromatic number t for any perfect matching. Inspired by Corollary 8.10, we think that the following analogue of Hadwiger's Conjecture for M-colorings of bipartite graphs could be true.

Conjecture 8.4. Let $t \in \mathbb{N}$, G be a bipartite graph and M an arbitrary perfect matching of G, such that $\chi(G, M) \geq t$. Then G contains $K_{t,t}$ as a matching minor.

While for t = 1, 2, the statement is trivial, the case t = 3 amounts to Corollary 8.10. At the current state, we do not have a good approach for proving this conjecture even in the first open case of t = 4. One of the reasons for this is that our proof for t = 3 relied on a certain sparsity of Pfaffian bipartite graphs, in the sense that every bipartite graph of sufficiently large extendability and minimum degree is not Pfaffian. Without this fact we could not have been sure to be able to apply our reduction operations to smallest counterexamples. Hence, it would be very helpful to prove similar results also for classes excluding larger complete bipartite graphs as matching minors.

Question 8.1. Is there a function $f: \mathbb{N} \to \mathbb{N}$ such that every f(t)-extendable bipartite graph G contains $K_{t,t}$ as a matching minor?

The following observation, which is a direct consequence of Theorem 8.4, provides some evidence towards Conjecture 8.4.

Corollary 8.37. There is an absolute constant C > 0 such that for any sufficiently large $t \in \mathbb{N}$ every bipartite graph G with a perfect matching M satisfying $\chi(G, M) \geq Ct(\log \log t)^6$ contains $K_{t,t}$ as a matching minor.

Proof. Set $f(t) := bm_{\vec{\chi}}(t)$ and let G be a bipartite graph with a perfect matching M such that $\chi(G, M) \ge f(t)$. We deduce from Theorem 8.4 that $\vec{\chi}(\mathcal{D}(G, M)) = \chi(G, M) \ge f(t)$ implies that $\mathcal{D}(G, M)$ contains \vec{K}_t , which is the unique perfect matching-direction of $K_{t,t}$, as a butterfly-minor. The claim now follows from Lemma 8.8.



Figure 8.8: A planar bipartite graph such that for any 2-coloring of its edges, there is a perfect matching with a monochromatic alternating cycle.

Considering the notion of M-colorings, it is natural to ask whether it is necessary to have different colorings of the matching edges for every perfect matching, or whether one might strengthen Corollary 8.10 by finding a single 2-coloring of all edges in a bipartite Pfaffian graph, such that for any perfect matching M the induced 2-coloring on the matching edges yields a proper M-coloring. For an example consider the 2-coloring of the staircase in Figure 8.7. Although it seems to be possible to find such a "super"-coloring for many interesting bipartite Pfaffian graphs such as the Heawood graph or grid graphs, there are small examples of (even planar) Pfaffian bipartite graphs without such a coloring, one of which is depicted in Figure 8.8.

Chapter 9

Excluding Induced Subdigraphs

9.1 Introduction

All graphs and digraphs considered in this chapter are simple.

In the previous chapters we have investigated bounds on the dichromatic number of a digraph provided that it excludes certain substructures such as minor models or subdivisions of other digraphs. Each of these notions of substructure generalizes subdigraphs in the sense that if D' is a subdigraph of another digraph D, then D' is also a strong (topological, butterfly) minor of D. The alert reader might therefore wonder about the following natural question: Given a set of digraphs \mathcal{F} , under which circumstances do the digraphs not containing any $F \in \mathcal{F}$ as a subdigraph have an interesting structure which yields a bound on the dichromatic number?

The reason why the above question has not been considered so far is that it is hard to provide a satisfying answer: If we would allow \mathcal{F} to be of infinite size, then the question is too general, as every class of digraphs closed under taking subdigraphs can be described as those digraphs excluding the minimal non-members of the class as subdigraphs. For instance, for any $k \in \mathbb{N}$ we might simply define \mathcal{F} to contain all (k+1)-dicritical digraphs F. Then a digraph D excludes all members of \mathcal{F} as subdigraphs if and only if it is acyclically k-colorable. This class of digraphs has bounded dichromatic number, but we cannot hope for a nice characterization of such digraphs, as the problem of deciding whether a digraph has dichromatic number at most k is NP-hard for any $k \geq 2$ (we will further elaborate on this topic in the next Chapter 10).

In contrast, if \mathcal{F} is finite, then the answer to our question is unsatisfyingly simple:

Observation 9.1. Let \mathcal{F} be a finite set of digraphs. Then the class of the digraphs excluding all digraphs isomorphic to members of \mathcal{F} as subdigraphs has bounded dichromatic number if and only if at least one $F \in \mathcal{F}$ is an orientation of a forest.

Proof. Suppose first that \mathcal{F} does not contain an oriented forest. Then every member of F contains an oriented cycle, let g be the maximum length of any oriented cycle appearing in any member of F (here we use the finiteness of \mathcal{F}). A result of Harutyunyan and Mohar [HM12b] now shows that there exist digraphs of arbitrarily large dichromatic number not containing any oriented cycle of length at most g, and hence not containing any member of \mathcal{F} as a subdigraph.

Therefore the digraphs excluding \mathcal{F} have unbounded dichromatic number.

Next suppose that $F \in \mathcal{F}$ is an oriented forest. Then any digraph D excluding all members of \mathcal{F} as subdigraphs must also exclude F. Then $\min\{\delta^+(D), \delta^-(D)\} < v(F) - 1$. For if $\delta^+(D), \delta^-(D) \ge v(F) - 1$, we could easily construct a copy of F in D by picking some starting vertex, identifying it as the root of F, and successively expanding a subtree of F by attaching out- and in-leafs (the degree conditions ensure that we can pick a leaf we want to attach disjoint from the already built subtree).

This implies that every digraph excluding the members of \mathcal{F} has dichromatic number at most v(F)-1, for otherwise it would contain a v(F)-dicritical subdigraph, which according to Lemma 7.11 from Chapter 7 must have minimum out- and in-degree at least v(F)-1, a contradiction. This shows that the dichromatic number of the digraphs excluding \mathcal{F} is bounded, concluding the proof.

Hence, the exclusion of subdigraphs is maybe not the right track towards obtaining interesting classes of digraphs with bounded dichromatic number. A more restricted concept of substructure is that of an *induced subdigraph*. Indeed, the classes of digraphs obtained by excluding induced subdigraphs are much richer.

For instance, while the digraphs excluding \vec{K}_2 (the single arc) as a subdigraph consist only of isolated vertices, the digraphs excluding \vec{K}_2 as an induced subdigraph constitute all biorientations of undirected graphs.

Aboulker, Charbit, and Naserasr [ACN20] recently initiated a systematic study of the relation between excluded induced subdigraphs and the dichromatic number and asked the following intriguing question.

Problem 9.1. For which finite sets \mathcal{F} of digraphs does there exist a constant C such that digraphs D without an induced subdigraph isomorphic to a member of \mathcal{F} satisfy $\vec{\chi}(D) \leq C$?

Following the terminology introduced by Aboulker et al. [ACN20], we denote by $\operatorname{Forb}_{\operatorname{ind}}(\mathcal{F})$ the set of digraphs containing no induced subdigraph isomorphic to a member of \mathcal{F} . We say that \mathcal{F} is *heroic* if the digraphs in $\operatorname{Forb}_{\operatorname{ind}}(\mathcal{F})$ have bounded dichromatic number and in this case we denote $\vec{\chi}(\operatorname{Forb}_{\operatorname{ind}}(\mathcal{F})) := \max{\{\vec{\chi}(D)|D \in \operatorname{Forb}_{\operatorname{ind}}(\mathcal{F})\}}$.

The following analogue of Problem 9.1 for undirected graphs and the chromatic number is quite famous and was the main inspiration of the work by Aboulker et al.

Problem 9.2. For which finite sets \mathcal{F} of graphs does there exist a constant C such that graphs G without an induced subgraph isomorphic to a member of \mathcal{F} satisfy $\chi(G) \leq C$?

Let us observe some necessary conditions on a finite set of graphs \mathcal{F} whose exclusion bounds the chromatic number. First of all, the complete graphs $(K_n)_{n \in \mathbb{N}}$ have unbounded chromatic number and since the only induced subgraphs of complete graphs are complete again, a necessary condition on \mathcal{F} is that it contains at least one clique graph. Secondly, a famous result by Erdős [Erd59] states the existence of graphs which simultaneously have arbitrarily large chromatic number and large girth, hence making it necessary for \mathcal{F} to contain at least one forest. After having made these simple observations, it is natural to ask whether any further graphs must be contained in \mathcal{F} to guarantee bounded chromatic number. Quite interestingly, Gyárfás [Gyá75] and Sumner [Sum81] independently conjectured that no further conditions are required.

Conjecture 9.1 (Gyárfás 1975, Summer 1980). If F is a forest and $k \in \mathbb{N}$, then the graphs excluding F and K_k as induced subgraphs have bounded chromatic number.

Over the years, despite its innocent appearance and attacks by several renowned researchers, the $Gy\acute{a}rf\acute{a}s$ -Sumner-Conjecture still stands unresolved. Known special cases include when F is a path ([Gy\acute{a}75]), a subdivided star ([Sco97]), a tree of radius two ([KP94]) and special kinds of caterpillars ([CSS19b]).

Just as in the undirected case, Aboulker et al. [ACN20] observed several necessary conditions for a finite set \mathcal{F} of digraphs to be heroic, which we summarize in the following.

Proposition 9.2 (cf. [ACN20]). Let \mathcal{F} be a heroic set of digraphs. Then \mathcal{F} must contain

- a bioriented clique \vec{K}_k for some $k \in \mathbb{N}$,
- a biorientation of a forest,
- an orientation of a forest,
- a tournament, i.e., an orientation of a complete graph.

These items can be respectively verified by considering the following families of digraphs with unbounded dichromatic number:

- The complete digraphs $(\vec{K}_n)_{n \in \mathbb{N}}$,
- biorientations of graphs with large chromatic number and large girth (cf. [Erd59]),
- oriented digraphs with large dichromatic number and large girth (cf. [HM12b]),
- tournaments with large dichromatic number.

Inspired by yet another important conjecture in graph theory, the Erdős-Hajnal-Conjecture, in [BCC⁺13] Berger, Choromanski, Chudnovsky, Fox, Loebl, Scott, Seymour, and Thomassé studied the dichromatic number of tournaments which exclude a single fixed tournament H as a(n induced) subdigraph. In this paper, the authors defined a hero to be a tournament H such that the tournaments excluding isomorphic copies of H have bounded dichromatic number. In other words, a digraph H is a hero if the set $\{\vec{K}_2, \vec{K}_2, H\}$ is heroic, where \vec{K}_{α} for $\alpha \geq 1$ denotes the arcless digraph on α vertices. The main result of Berger et al. in [BCC⁺13] was a recursive characterization of heroes as follows.

Theorem 9.3. A tournament H is a hero if and only if either H is the single-vertex digraph, or one of the following holds:

- There are heroes H₁, H₂ such that H is obtained from the disjoint union of H₁ and H₂ by adding all arcs in V(H₁) × V(H₂).
- There exist heroes H_1, H_2 , at least one of which is a transitive tournament, and a vertex $v \notin V(H_1) \cup V(H_2)$ such that H is obtained from the disjoint union of H_1 , H_2 and v by adding all the arcs in $V(H_1) \times V(H_2)$, $\{v\} \times V(H_1)$ and $V(H_2) \times \{v\}$.

It follows directly from this characterization that every transitive tournament and all tournaments on at most four vertices are heroes.

It is a natural aim to characterize the finite heroic sets \mathcal{F} of digraphs similar to what is claimed by the Gyárfás-Sumner-Conjecture for undirected graphs. In contrast to undirected graphs, only heroic sets of size at least 3 are interesting to consider, as the necessary conditions from Proposition 9.2 directly imply that $\{\vec{K}_2, \vec{K}_2\}$ is the only heroic set of size two (and is so trivially). Aboulker et al. in [ACN20] proved that every heroic set of size three must be of one of the following three types:

- $\{\vec{K}_2, \vec{F}, \vec{K}_k\}$ for a forest F and a number $k \in \mathbb{N}$,
- $\{ \overrightarrow{K}_k, \overline{K}_\alpha, H \}$ for $k, \alpha \in \mathbb{N}$ and a hero H such that k = 2 or H is transitive, or
- $\{ \stackrel{\leftrightarrow}{K_2}, F, H \}$ for some oriented star forest¹ F and a hero H, or

¹An oriented star forest is a disjoint union of orientations of stars.

• $\{\vec{K}_2, F, \vec{K}_k\}$ for some oriented forest F and some $k \in \mathbb{N}$.

They then ventured to propose the conjecture that every one of the above triples is indeed heroic, thus claiming a complete description of the heroic sets of size 3.

Note that since \overline{K}_2 -free digraphs amount exactly to the biorientations of undirected graphs, and since dichromatic number and chromatic number coincide on these, the conjecture of Aboulker et al. corresponds exactly to the undirected Gyárfás-Sumner-Conjecture when restricting to the triples of the first type above. Triples of the second type as above were shown to be indeed heroic by Aboulker et al. (cf. [ACN20], Theorem 4.1). Finally, for the third and fourth types of triples we deal with oriented graphs. Let us explicitly state the conjecture for these cases.

Conjecture 9.2. For every oriented star forest F and every hero H the oriented graphs excluding F and H as induced subdigraphs have bounded dichromatic number.

Conjecture 9.3. For every oriented forest and every $k \in \mathbb{N}$ the oriented graphs excluding F and \vec{K}_k as induced subdigraphs have bounded dichromatic number.

In the main results of this chapter we will solve several special cases of Conjecture 9.2 and Conjecture 9.3.

Aboulker et al. [ACN20] noted that in case that H is a transitive tournament, Conjecture 9.2 follows from a result of Chudnovsky, Scott, and Seymour [CSS19a]. They observed that Conjecture 9.2 is true in the case that F has at most two vertices. Finally they focused on the case when $H = \vec{C}_3$ is the smallest non-trivial hero and F has 3 vertices. Then F must be one of the following:

- \overline{K}_3 , the forest consisting of three isolated vertices,
- \vec{P}_3 , the directed path on three vertices,
- $\vec{K}_2 + K_1$, the oriented star forest consisting of an arc plus an isolated vertex,
- S_2^+ , the 2-out-star, or
- S_2^- , the 2-in-star.

They proved that $\{\vec{K}_2, F, \vec{C}_3\}$ is indeed heroic if F is one of the first three star forests. Already in the case $F \in \{S_2^+, S_2^-\}$ however, the could not to prove heroicness. Nevertheless, they made the following explicit conjecture.

Conjecture 9.4 (cf. [ACN20], Conjecture 6.2).

$$\vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^+, \vec{C}_3)) = \vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^-, \vec{C}_3)) = 2.$$

Note that by symmetry of reversing all arcs, it suffices to prove Conjecture 9.4 for the out-star S_2^+ . The digraphs in Forb_{ind}($\vec{K}_2, S_2^+, \vec{C}_3$) are exactly the directed triangle-free oriented graphs such that the out-neighborhood of every vertex induces a tournament. As the first main result of this chapter, we prove Conjecture 9.4.

Theorem 9.4.

$$\vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^+, \vec{C}_3)) = \vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^-, \vec{C}_3)) = 2$$

In fact, we deduce Theorem 9.4 as an immediate Corollary of the following stronger result involving the hero W_3^+ obtained from the directed triangle \vec{C}_3 by attaching a dominating source (cf. Figure 1.1).

9.1. INTRODUCTION

Theorem 9.5. $\vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)) = 2.$

In order to prove Theorem 9.5, we need to establish several auxiliary results which deal with the structure of digraphs in the class $\operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$, which is surprisingly complicated (these are exactly the oriented graphs in which the out-neighborhood of every vertex induces a transitive tournament).

Considering Theorem 9.5 it is natural to try and verify Conjecture 9.2 for more triples of the form $\{\vec{K}_2, S_2^+, H\}$, where H is some hero. In this direction, we can prove the following result involving the hero W_3^- obtained from \vec{C}_3 by attaching a dominating sink.

Theorem 9.6. $\vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^+, W_3^-)) \le 4.$

Our last result concerning Conjecture 9.2 generalizes Theorem 9.6 qualitatively and proves that for every $k \in \mathbb{N}$, the triple $\{\vec{K}_2, S_2^+, H_k\}$ is heroic, where H_k is the hero on k vertices obtained from the disjoint union of W_3^+ and \vec{K}_{k-4} by adding all possible arcs from W_3^+ towards \vec{K}_{k-4} .

Theorem 9.7. Let H be a hero and let H^- be the hero obtained from H by adding a dominating sink. If $\{\vec{K}_2, S_2^+, H\}$ is heroic, then so is $\{\vec{K}_2, S_2^+, H^-\}$.

Our last new result in this Chapter concerns Conjecture 9.3. As mentioned above, Conjecture 9.3 holds true whenever F is an oriented star forest, and therefore particularly for forests on at most 3 vertices. The first open cases therefore appear when F is an orientation of the P_4 . Aboulker et al. considered the directed path \vec{P}_4 and showed in one of their main results that the set $\{\vec{K}_2, \vec{P}_4, \vec{K}_3\}$ is heroic. It remains an open problem whether $\{\vec{K}_2, \vec{P}_4, \vec{K}_k\}$ is heroic for $k \ge 4$. There are three other oriented paths on four vertices. Two of them, which are called $P^+(2, 1)$ and $P^-(2, 1)$ in [ACN20], consist of two oppositely



oriented dipaths of length two and one, respectively. Chudnovsky, Scott, and Seymour proved in [CSS19a] that for every $k \in \mathbb{N}$, digraphs in the set $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, P, \vec{K}_k)$ have underlying graphs with bounded chromatic number (and thus bounded dichromatic number) for $P \in \{P^+(2,1), P^-(2,1)\}$. Hence, Conjecture 9.3 holds for these two orientations of P_4 . The same result however is wrong for the remaining orientation of P_4 , which we denote by $P^+(1,1,1)$, as it consists of 3 alternatingly oriented arcs. Here we complement the result of Aboulker et al. [ACN20] concerning the directed path \vec{P}_4 and k = 3 by showing that also the set $\{\vec{K}_2, P^+(1,1,1), \vec{K}_3\}$ is heroic.

Theorem 9.8. $\vec{\chi}(Forb_{ind}(\vec{K}_2, P^+(1, 1, 1), \vec{K}_3)) = 2.$

We remark that the class $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, P^+(1, 1, 1), \vec{K}_3)$ is quite rich, as it (among others) contains all oriented *line digraphs*.

Structure of the chapter. In Section 9.2 we investigate the structure of digraphs in the class $\operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$ and use these insights to prove Theorem 9.5. In Section 9.3 we give the proof of Theorem 9.6. In Section 9.4 we prove Theorem 9.7. Finally, in Section 8.5 we conclude this chapter by proving Theorem 9.8.

9.2 $\{S_2^+, W_3^+\}$ -Free Oriented Graphs

In this section, we will prove Theorem 9.5 and thereby show that $\{\vec{K}_2, S_2^+, W_3^+\}$ is a heroic set. Note that $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^+)$ is the class of oriented graphs D with the property that the out-neighbourhood of every vertex in D induces a transitive tournament. Given $D \in \operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^+)$, we define F = F(D) to be the spanning subdigraph of D consisting of the arcs $(x, y) \in A(D)$ such that y is the source in the transitive tournament induced by the out-neighbourhood of x in D. Observe that for every $x \in V(D)$, if $d_D^+(x) \geq 1$ then $d_F^+(x) = 1$, and otherwise $d_F^+(x) = 0$. From the definition of F(D) we immediately obtain the following property:

Fact 9.1. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$ and $(x, y) \in A(F(D))$. Then we have

$$N_D^+(x) \subseteq N_D^+(y) \cup \{y\}$$

The next fact follows immediately from Fact 9.1 via induction.

Fact 9.2. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$ and let x_1, \ldots, x_k be the vertex-trace of a dipath in F(D). Then

$$N_D^+(x_1) \setminus \{x_2, \dots, x_k\} \subseteq N_D^+(x_k).$$

From Fact 9.2 we can derive that the vertex-sets of directed cycles in F(D) form so-called *out-modules* in D. Modules in digraphs will further be discussed in Chapter 10.

Definition 9.1. Let D be a digraph, and $\emptyset \neq M \subseteq V(D)$. We say that M is an outmodule in D if $(x, z) \in A(D)$ implies that $(y, z) \in V(D) \setminus M$ for every $x, y \in M$ and $z \in V(D) \setminus M$.

Phrased differently, the definition says that a non-empty vertex-set M is an out-module if $N_D^+(x) \setminus M = N_D^+(y) \setminus M$ holds for all $x, y \in M$.

Fact 9.3. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$, and let C be a directed cycle in F(D). Then V(C) is an out-module in D.

Proof. Let $x_1, x_2, \ldots, x_k, x_1$ be the vertex-trace of C. Let $y \in V(D) \setminus \{x_1, \ldots, x_k\}$ and $1 \leq i \leq k$ be arbitrary such that $(x_i, y) \in A(D)$. Let $j \in [k] \setminus \{i\}$. By Fact 9.2, applied to the directed subpath of C starting in x_i and ending in x_j , we know that $N_D^+(x_i) \setminus \{x_1, \ldots, x_k\} \subseteq N_D^+(x_j) \setminus \{x_1, \ldots, x_k\}$. Hence $(x_j, y) \in A(D)$. This shows that $\{x_1, \ldots, x_k\}$ is indeed an out-module.

For a non-empty vertex-set U in a digraph D, we denote by D/U the digraph obtained by *identifying* U, that is, the digraph with vertex set $(V(D) \setminus U) \cup \{x_U\}$ where $x_U \notin V(D)$ is some newly added vertex representing U, and the following arcs: the arcs of D inside $V(D) \setminus U$, the arc (x_U, v) for every $v \in N_D^+(U)$, and the arc (v, x_U) for every $v \in N_D^-(U)$.

In the following we prepare the proof of Theorem 9.5 with a set of useful Lemmas. We start with two lemmas yielding modifications of digraphs which preserve the containment in the class $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^+)$.

Lemma 9.9. For every $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$ and for every out-module $U \subseteq V(D)$ it holds that $D/U \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$.

Proof. We need to show that D/U is induced $\{\vec{K}_2, S_2^+, W_3^+\}$ -free. We argue by contradiction. Suppose first that D/U contains a \vec{K}_2 , namely a pair of vertices x, y with $(x, y), (y, x) \in A(D/U)$. If $x, y \neq x_U$ then x, y also span a copy of \vec{K}_2 in D, a contradiction. Next assume that $x = x_U$ or $y = x_U$; say $x = x_U$. By the definition of D/U, there are (not necessarily distinct) $u_1, u_2 \in U$ such that $(u_1, y), (y, u_2) \in A(D)$. Since U is an out-module, $(u_2, y) \in A(D)$. Hence, u_2, y span a copy of \vec{K}_2 in D, a contradiction.

Suppose next that D/U contains an induced copy of S_2^+ , namely, distinct vertices $x, y, z \in V(D/U)$ with $(x, y), (x, z) \in A(D/U)$ and with no arc in D/U between y and z. If $x, y, z \neq x_U$ then x, y, z also span an induced S_2^+ in D, a contradiction. Now consider the case that $y = x_U$, and let $u \in U$ be such that $(x, u) \in A(D)$. We have $(u, z), (z, u) \notin A(D)$ because $(x_U, z), (z, x_U) \notin A(D/U)$. Hence, x, u, z span an induced S_2^+ in D, a contradiction. The case $z = x_U$ is analogous. Suppose now that $x = x_U$. Since $(x_U, y), (x_U, z) \in A(D/U)$ and U is an out-module, we must have $(u, y), (u, z) \in A(D)$ for every $u \in U$, implying that u, y, z span an induced S_2^+ in D for every such u, again yielding the desired contradiction.

Now let us consider the last case, i.e., that D/U contains a copy of W_3^+ with vertices x, y, z, w and arcs (x, y), (x, z), (x, w), (y, z), (z, w), (w, y). Again, if $x, y, z, w \neq x_U$ then D also has a copy of W_3^+ , a contradiction. Suppose now that $x = x_U$, and fix any $u \in U$. Since U is an out-module, we have $(u, y), (u, z), (u, w) \in A(D)$, implying that u, y, z, w span a copy of W_3^+ in D, a contradiction. Suppose finally that one of y, z, w equals x_U , say $y = x_U$ (without loss of generality). Since $(x, x_U), (x_U, z), (w, x_U) \in A(D/U)$, there are $u_1, u_2, u_3 \in U$ (not necessarily distinct) such that $(x, u_1), (u_2, z), (w, u_3) \in A(D)$. Since U is an out-module, we have $(u_1, z), (u_3, z) \in A(D)$. Since $(x, w), (x, u_1) \in A(D)$ and D is induced S_2^+ -free, we must have either $(w, u_1) \in A(D)$ or $(u_1, w) \in A(D)$. If $(u_1, w) \in A(D)$ then also $(u_3, w) \in A(D)$ because U is an out-module, but this is impossible as then u_3, w would induce a digon in D. Finally, if $(w, u_1) \in A(D)$ then x, u_1, z, w span a copy of W_3 in D, again yielding the desired contradiction.

As the assumption $D/U \notin \text{Forb}_{\text{ind}}(K_2, S_2^+, W_3^+)$ led to a contradiction in every case, we may conclude the proof of the lemma.

Lemma 9.10. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$, and let $(x, y) \in A(F(D))$. Let $z \in N_D^+(y)$ such that $(x, z), (z, x) \notin A(D)$. Then the digraph D + (x, z) obtained from D by adding the arc (x, z) is contained in $Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$.

Proof. We need to show that D + (x, z) is induced $\{\vec{K}_2, S_2^+, W_3^+\}$ -free. Again, we argue by contradiction. Clearly D + (x, z) does not contain a \vec{K}_2 , since $(z, x) \notin A(D)$ by assumption. Suppose next that D + (x, z) contains an induced copy of S_2^+ , i.e. distinct vertices a, b, c such that $(a, b), (a, c) \in A(D + (x, z))$, and $(b, c), (c, b) \notin A(D + (x, z))$. If $(x, z) \notin \{(a, b), (a, c)\}$, then a, b, c induce a copy of S_2^+ also in D, a contradiction. We may therefore assume w.l.o.g. that (x, z) = (a, b). Then we have $c \neq y$, since $(c, b) \notin A(D)$, but $(y, b) = (y, z) \in A(D)$ by assumption. Since $(x, c) = (a, c) \in A(D)$, we have $c \in N_D^+(x)$. But $(x, y) \in A(F(D))$, and hence Fact 9.1 implies that $c \in N_D^+(x) \subseteq \{y\} \cup N_D^+(y)$. It follows that $(y, c) \in A(D)$. We further have $(y, b) = (y, z) \in A(D)$ and $(b, c), (c, b) \notin A(D)$. Hence, y, b, c induce an S_2^+ in D, a contradiction.

Moving on, suppose that D + (x, z) contains an induced copy of W_3^+ , i.e., distinct vertices a, b, c, d such that $(a, b), (a, c), (a, d), (b, c), (c, d), (d, b) \in A(D) \cup \{(x, z)\}$.

Suppose first that $(x, z) \notin \{(a, b), (a, c), (a, d)\}$. Then $(a, b), (a, c), (a, d) \in A(D)$ and since D does not contain an induced copy of S_2^+ , the vertices b, c, d are pairwise adjacent. Since x and z are non-adjacent in D, it follows that $\{x, z\} \not\subseteq \{b, c, d\}$. Hence we have $(b, c), (c, d), (d, b) \in A(D)$, yielding that a, b, c, d induce a copy of W_3^+ in D, a contradiction. Hence we may suppose that $(x, z) \in \{(a, b), (a, c), (a, d)\}$. By symmetry we may assume that (x, z) = (a, b) w.l.o.g. Again using Fact 9.1 we then have $c, d \in N_D^+(x) \subseteq N_D^+(y) \cup \{y\}$.

Let us first consider the case that $c, d \neq y$. Then $(y, c), (y, d) \in A(D)$, and since $(y, b) = (y, z) \in A(D)$ by assumption, it follows that the vertices y, b, c, d induce a copy of W_3^+ in D, a contradiction. For the next case suppose that $y \in \{c, d\}$. The first option, namely that y = c, is impossible, since then we would have $(y, z) \in A(D)$ (by assumption) and $(z, y) = (b, c) \in A(D)$, a contradiction since D is \vec{K}_2 -free. Therefore, we must have y = d. Then $c \in N_D^+(y)$ as well as $(c, y) = (c, d) \in A(D)$. It follows that y, c induce a \vec{K}_2 in D, so again, we conclude with a contradiction.

Having reached a contradiction in all possible cases, it follows that our initial assumption, was wrong, indeed, $D + (x, z) \in \text{Forb}_{\text{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$. This concludes the proof. \Box

The next lemma shows the existence of out-modules with special properties.

Lemma 9.11. Let $D \in Forb_{ind}(\overset{\circ}{K}_2, S_2^+, W_3^+)$, and let $v \in V(D)$. If $N_D^-(v) \neq \emptyset$, then there exists an out-module M in D such that $M \subseteq N_D^-(v)$ and $N_D^+(M) \subseteq N_D^+(v) \cup \{v\}$.

Proof. We prove by induction on $n \ge 1$ the statement of the lemma for all digraphs $D \in \operatorname{Forb}_{\operatorname{ind}}(K_2, S_2^+, W_3^+)$ and vertices $v \in V(D)$ such that $d_D^-(v) \le n$.

If n = 1, then $N_D^-(v) = \{w\}$ for a vertex $w \in V(D)$. Then $M := \{w\}$ is an outmodule of D. Hence, it suffices to verify that $N_D^+(w) \subseteq N_D^+(v) \cup \{v\}$. Suppose towards a contradiction that $(w, v') \in A(D)$ for $v' \in N_D^+(w) \setminus (N_D^+(v) \cup \{v\})$. Since $N_D^-(v) = \{w\}$, we have $v' \notin N_D^-(v)$, and hence v, v' are non-adjacent in D, while $(w, v), (w, v') \in A(D)$. Hence, w, v, v' induce an S_2^+ in D, a contradiction.

Now let $D \in \text{Forb}_{\text{ind}}(\mathring{K}_2, S_2^+, W_3^+)$ and $v \in V(D)$ such that $d_D^-(v) = n \ge 2$, and assume that the claim holds for all pairs of digraphs in $\text{Forb}_{\text{ind}}(\mathring{K}_2, S_2^+, W_3^+)$ and vertices whose in-degree is less than n.

First let us assume that there exists a vertex $w \in N_D^-(v)$ such that $(w, v) \in A(F(D))$. Then $M := \{w\}$ is an out-module of D, and by Fact 9.1 we have $N_D^+(w) \subseteq N_D^+(v) \cup \{v\}$. This proves the assertion in this case.

Hence, for the rest of this proof we may suppose that $(w, v) \notin A(F(D))$ for every $w \in N_D^-(v)$. For any $w \in N_D^-(v)$ we clearly have $d_D^+(w) \ge 1$ and hence it follows that $d_{F(D)}^+(w) = 1$. Furthermore, for every arc $(w, w') \in A(F(D))$ such that $w \in N_D^-(v)$ by Fact 9.1 we must have $v \in N_D^+(w) \subseteq N_D^+(w') \cup \{w'\}$. Since $(w, v) \notin A(F(D))$, we have $v \neq w'$ and hence $(w', v) \in A(D)$. This shows that the out-neighbor in F(D) of any vertex in $N_D^-(v)$ is again contained in $N_D^-(v)$. It follows that F(D) restricted to $N_D^-(v)$ has minimum out-degree 1 and therefore contains a directed cycle C such that $V(C) \subseteq N_D^-(v)$.

By Fact 9.3, N := V(C) is an out-module in D. Consider the digraph D' := D/N, which by Lemma 9.9 is a member of $\operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$. Then by definition of D/Nand since $N \subseteq N_D^-(v)$, we have $v \in V(D')$ and $N_{D'}^-(v) = (N_D^-(v) \setminus N) \cup \{x_N\} \neq \emptyset$. Since $|N| = |V(C)| \ge 3$, this implies that $d_{D'}^-(v) = 1 + d_D^-(v) - |N| \le d_D^-(v) - 2 < n$. We may therefore apply the induction hypothesis to the digraph $D' \in \operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$ and the vertex $v \in V(D')$. We thus find an out-module M' in D' with the properties $M' \subseteq N_{D'}^-(v) = (N_D^-(v) \setminus N) \cup \{x_N\}$ and $N_{D'}^+(M') \subseteq N_{D'}^+(v) \cup \{v\} = N_D^+(v) \cup \{v\}$. Let us define the set $M \subseteq N_D^-(v)$ as M := M', if $x_N \notin M'$, and $M := (M' \setminus \{x_N\}) \cup N$ if $x_N \in M'$. We claim that M satisfies the assertions of the Lemma with respect to D and v. In the following, we verify both parts of the inductive claim separately.

Claim 1. $N_D^+(M) \subseteq N_D^+(v) \cup \{v\}.$

Proof. In the proof we will use the fact that

$$N_{D'}^+(M') \subseteq N_{D'}^+(v) \cup \{v\} = N_D^+(v) \cup \{v\},\$$

which holds by the induction hypothesis.

Let $x \in N_D^+(M)$ be given arbitrarily. Let $m \in M$ such that $(m, x) \in A(D)$. Our goal is to show that $x \in N_D^+(v) \cup \{v\}$.

Let us first consider the case that $x_N \notin M'$ and hence M = M'. By definition of D' = D/N we either have $x \notin N$ and $(m, x) \in A(D')$, or $x \in N$ and $(m, x_N) \in A(D')$. Then since $m \in M = M'$ and $x, x_N \notin M = M'$, we obtain that either $x \notin N$ and $x \in N_{D'}^+(M')$, or $x \in N$ and $x_N \in N_{D'}^+(M')$. As $N_{D'}^+(M') \subseteq N_D^+(v) \cup \{v\}$, in the first case we have $x \in N_D^+(v) \cup \{v\}$, as desired. The second case does not occur, since it yields $x_N \in N_D^+(v) \cup \{v\}$, which is impossible as x_N is not a vertex of D. Hence, we have shown the claim that $x \in N_D^+(v) \cup \{v\}$.

For the second case, suppose that $x_N \in M'$ and hence $M = (M' \setminus \{x_N\}) \cup N$. Note that $x \notin N$, since $x \notin M \supseteq N$. Hence, the existence of the arc $(m, x) \in A(D)$ yields that either $m \in N$ and $(x_N, x) \in A(D')$, or $m \notin N$ and $(m, x) \in A(D')$. In both cases, this implies that $x \in N^+_{D'}(M') \subseteq N^+_D(v) \cup \{v\}$, proving the assertion.

Claim 2. M is an out-module in D.

Proof. Let $x \neq y \in M$ and $z \in V(D) \setminus M$ arbitrary, and assume that $(x, z) \in A(D)$. We need to show that also $(y, z) \in A(D)$. Note that by Claim 1 we have $z \notin N_D^-(v)$, as otherwise $z \in N_D^+(M) \cap N_D^-(v) = \emptyset$. In particular, $z \notin N$.

Observe that $z \in N_{D'}^+(M')$. Indeed, if $x \notin N$ then $x \in M'$ and $(x, z) \in A(D')$, and if $x \in N$ then $x_N \in M'$ and $(x_N, z) \in A(D')$; in any case, $z \in N_{D'}^+(M')$.

Since $y \in M$, we have either $y \in M'$ or $y \in N$ and $x_N \in M'$. Suppose first that $y \in M'$. Then $(y, z) \in A(D')$ because $z \in N_{D'}^+(M')$ and M' is an out-module. Hence, in this case $(y, z) \in A(D)$, as required. Now suppose that $y \in N$ and $x_N \in M'$. Since $z \in N_{D'}^+(M')$ and M' is an out-module, we have $(x_N, z) \in A(D')$. This means that there is $w \in N$ such that $(w, z) \in A(D)$. Now, as N is itself is an out-module in D and $y \in N, z \notin N$, we have $(y, z) \in A(D)$, as required.

By Claim 1 and 2 the out-module M certifies that the pair (D, v) satisfies the inductive claim. This concludes the proof of the Lemma by induction.

Lemma 9.12. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$, let $M \subseteq V(D)$ be an out-module in D and let $v \in V(D) \setminus M$. Let T be the set of vertices defined by

$$T := \{t \in M | \exists u \in V(D) \setminus M : (v, u), (u, t) \in A(D)\}.$$

If $T \neq \emptyset$, then D[T] is a transitive tournament.

Proof. Suppose that $T \neq \emptyset$. The assertion will follow directly from the following two claims and the fact that every directed triangle-free tournament is transitive.

Claim 1. If $t_1 \neq t_2 \in T$, then t_1 and t_2 are adjacent in D.

Proof. By definition of T there exist vertices $u_1, u_2 \in V(D) \setminus M$ (not necessarily distinct) such that $(v, u_i), (u_i, t_i) \in A(T)$ for i = 1, 2. If $u_1 = u_2$, then t_1, t_2 must be adjacent, for otherwise the vertices u_1, t_1, t_2 would induce an S_2^+ in D, a contradiction. Suppose now that $u_1 \neq u_2$. Since $u_1, u_2 \in N^+(v)$, they must be adjacent, w.l.o.g. let $(u_1, u_2) \in A(D)$.

Then $t_1 \in M$ and $u_2 \in V(D) \setminus M$ are distinct out-neighbors of u_1 , and hence they must be adjacent in D. If $(t_1, u_2) \in A(D)$, then M being an out-module implies that also $(t_2, u_2) \in A(D)$, yielding a K_2 in D induced by t_2 and u_2 , a contradiction. Therefore we have $(u_2, t_1) \in A(D)$. Then t_1 and t_2 are distinct out-neighbors of u_2 in D, which implies that they must be adjacent. This concludes the proof.

Claim 2. D[T] contains no directed triangle.

Proof. Suppose towards a contradiction that there are three distinct $t_1, t_2, t_3 \in T$ inducing a directed triangle in D. Let $u_1, u_2, u_3 \in V(D) \setminus M$ be (not necessarily distinct) such that $(v, u_i), (u_i, t_i) \in A(D), i = 1, 2, 3$. We distinguish three different cases depending on the size of the set $\{u_1, u_2, u_3\}$.

For the first case, suppose that $u_1 = u_2 = u_3$. Then t_1, t_2, t_3 are three distinct outneighbors of u_1 spanning a directed triangle. Hence, u_1, t_1, t_2, t_3 induce a W_3^+ in D, a contradiction to our assumption that $D \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+)$.

For the second case, suppose that $\{u_1, u_2, u_3\}$ contains exactly two distinct vertices, w.l.o.g. $u_1 \neq u_2 = u_3$. Since u_1 and u_2 are two distinct out-neighbors of v in D, they must be adjacent. Suppose first that $(u_1, u_2) \in A(D)$. Then t_1 and u_2 are two distinct out-neighbors of u_1 in D, and hence they must be adjacent. If $(t_1, u_2) \in A(D)$, then by the module-property of M, also $(t_2, u_2) \in A(D)$, and hence t_2, u_2 induce a K_2 in D, a contradiction. If $(u_2, t_1) \in A(D)$, then $t_1, t_2, t_3 \in N_D^+(u_2)$ and hence u_2, t_1, t_2, t_3 induce a W_3^+ in D, a contradiction. Next suppose that $(u_2, u_1) \in A(D)$. Then u_1, t_2, t_3 are three distinct out-neighbors of u_2 in D, and hence u_1 must be adjacent to both t_2 and t_3 . If $(t_2, u_1) \in A(D)$ or $(t_3, u_1) \in A(D)$, then $(t_1, u_1) \in A(D)$ since M is an out-module, and hence u_1, t_1 induce a K_2 in D, a contradiction. Otherwise, we have $(u_1, t_2), (u_1, t_3) \in A(D)$ and hence u_1, t_1, t_2, t_3 induce a W_3^+ in D, again yielding the desired contradiction.

For the third case, suppose that u_1, u_2, u_3 are pairwise distinct. Since u_1, u_2, u_3 are three distinct vertices in the transitive tournament $D[N^+(v)]$, they form a transitive triangle, and we may assume W.l.o.g. that $(u_1, u_2), (u_1, u_3), (u_2, u_3) \in A(D)$. Then t_1 and u_3 are distinct out-neighbors of u_1 in D, while t_2 and u_3 are distinct out-neighbors of u_2 in D. Hence, u_3 must be adjacent to both t_1 and t_2 . If $(t_i, u_3) \in A(D)$ for some i = 1, 2,then we also have $(t_3, u_3) \in A(D)$ since M is an out-module, and hence u_3, t_3 induce a K_2 in D, contradiction. Finally, if $(u_3, t_1), (u_3, t_2) \in A(D)$, then u_3, t_1, t_2, t_3 induce a W_3^+ in D, yielding again a contradiction to the containment of D in Forb_{ind} (K_2, S_2^+, W_3^+) .

Since we arrived at a contradiction in each case, we conclude that the initial assumption concerning the existence of t_1, t_2, t_3 was wrong. This concludes the proof of Claim 2.

We are now sufficiently prepared to give the proof of Theorem 9.5. In fact, we will prove the following slightly stronger version of the result, which allows to enforce a monochromatic coloring on the closed out-neighborhood of an arbitrarily chosen vertex.

Theorem 9.13. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^+)$, and $v \in V(D)$. Then there exists an acyclic coloring $c: V(D) \to \{1, 2\}$ of D such that c(u) = c(v) for every $u \in N_D^+(v)$.

Proof. Suppose towards a contradiction that the claim is wrong, and let D be a counterexample to the claim minimizing v(D). Let $v \in V(D)$ be a vertex such that D does not admit an acyclic 2-coloring c with the property that c(u) = c(v) for every $u \in N_D^+(v)$.

Claim 1. $N_D^-(v) \neq \emptyset$.

Proof. Suppose towards a contradiction that $N_D^-(v) = \emptyset$. If also $N_D^+(v) = \emptyset$, then v is an isolated vertex of D. Then any acyclic 2-coloring of D - v could be extended to an acyclic 2-coloring of D by coloring v with color 1, and the statement that v has the same color as its out-neighbors would hold vacuously. Since this is impossible, we must have $\vec{\chi}(D-v) \geq 3$, which however contradicts the minimality of D as a counterexample. This shows that $N_D^+(v) \neq \emptyset$. Let $u \in N_D^+(v)$ be the unique out-neighbor of v in F(D). Then $N_D^+(v) \subseteq N_D^+(u) \cup \{u\}$ by Fact 9.1. The minimality of D as a counterexample now implies that the digraph $D - v \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+)$ admits an acyclic 2-coloring $c^-: V(D) \to \{1,2\}$ satisfying $c^-(x) = c^-(u)$ for every $x \in N^+_{D-v}(u) = N^+_D(u)$. Let $c: V(D) \to \{1, 2\}$ be defined as $c(x) := c^{-}(x)$ for every $x \in V(D) \setminus \{v\}$ and $c(v) := c^{-}(u)$. Since c restricted to $V(D) \setminus \{v\}$ is an acyclic coloring, and no directed cycle in D contains v (recall $N_D^-(v) = \emptyset$), it follows that c is an acyclic coloring of D. Moreover, for every $x \in N_D^+(v) \subseteq N_D^+(u) \cup \{u\}$ we have $c(x) = c^-(x) = c^-(u) = c(v)$. This is a contradiction to our initial assumption that D does not admit an acyclic 2-coloring with this property. This shows that our assumption $N_D^-(v) = \emptyset$ was wrong, concluding the proof.

By Lemma 9.11 applied to the vertex v of D, there exists an out-module M in D such that $M \subseteq N_D^-(v)$ and $N_D^+(M) \subseteq N_D^+(v) \cup \{v\}$. Let $T \subseteq M$ be the set of vertices $t \in M$ for which there exists $u \in N_D^+(v)$ such that $(u,t) \in A(D)$. Since $N_D^+(v) \cap M = \emptyset$, the definition of T here coincides with the one in Lemma 9.12. Now, Lemma 9.12 implies that either $T = \emptyset$ or D[T] is a transitive tournament.

Claim 2. The digraph D[M] admits an acyclic 2-coloring $c_M : M \to \{1, 2\}$ satisfying $c_M(t) = 2$ for all $t \in T$.

Proof. Since $v(D[M]) \leq v(D-v) < V(D)$, the minimality of the counterexample D implies that $D[M] \in \operatorname{Forb}_{\operatorname{ind}}(K_2, S_2^+, W_3^+)$ satisfies the assertion of Theorem 9.13. If $T = \emptyset$, Claim 2 is satisfied by an arbitrary choice of an acyclic 2-coloring for D[M]. If $T \neq \emptyset$, let $t_0 \in T$ be the source of the transitive tournament D[T]. Applying the assertion of the theorem to D[M] and the vertex t_0 , we find that there exists an acyclic 2-coloring of D[M] in which t_0 has the same color as all its out-neighbors. W.l.o.g. we may choose this color to be 2, and since $N_{D[M]}^+(t_0) \supseteq T$, the claim follows.

Claim 3. D[M] contains a directed cycle.

Proof. Suppose towards a contradiction that D[M] is acyclic. Let D' := D - M. Clearly, $D' \in \operatorname{Forb}_{\operatorname{ind}}(\tilde{K}_2, S_2^+, W_3^+)$. Since $v(D') \leq v(D) - 1$ and by the minimality of D as a counterexample, we know that D' admits an acyclic 2-coloring $c' : V(D) \setminus M \to \{1,2\}$ in which c'(v) = c'(u) = 1 for every $u \in N_D^+(v)$. Let $c : V(D) \to \{1,2\}$ be defined by c(x) := c'(x) for all $x \in V(D) \setminus M$ and c(x) := 2 for all $x \in M$. We claim that c is an acyclic coloring of D. Suppose towards a contradiction that C is a directed cycle in D which is monochromatic in the coloring c. We must have $V(C) \cap M \neq \emptyset$, for otherwise C would form a monochromatic directed cycle in the coloring c' of D'. Since D[M] is acyclic, we must also have $V(C) \setminus M \neq \emptyset$. It follows that there exists an arc $(x, y) \in A(C)$ such that $x \in M$ and $y \in V(D) \setminus M$. Then $y \in N_D^+(M) \subseteq N_D^+(v) \cup \{v\}$, and therefore c(y) = c'(y) = 1, while c(x) = 2 by definition. This contradicts the fact that C is monochromatic, and hence we have shown that indeed c is an acyclic coloring of D. Moreover, c(v) = c(u) = 1 for every $u \in N_D^+(v)$. This contradicts our initial assumptions on D that such a coloring does not exist. Hence, D[M] cannot be acyclic, proving the claim.

Claim 3 in particular implies that $|M| \ge 3$ and $M \setminus T \ne \emptyset$.

Let us further note that since M forms an out-module in $D, M \setminus T \neq \emptyset$ is an out-module in the digraph $D-T \in \operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$, and hence by Lemma 9.9 we also have $D_0 := (D-T)/(M \setminus T) \in \operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, W_3^+)$. Also note that since $T \subseteq M \subseteq N_D^-(v)$, we still have $N_D^+(v) \cup \{v\} \subseteq \{x_{M \setminus T}\} \cup (V(D) \setminus M) = V(D_0)$, where we denote by $x_{M \setminus T}$ the vertex in D_0 obtained by identifying $M \setminus T$.

Claim 4. We have $N_{D_0}^+(v) = N_D^+(v)$, $(x_{M\setminus T}, v) \in A(F(D_0))$, and for every $u \in N_{D_0}^+(v)$, we have $(u, x_{M\setminus T}) \notin A(D_0)$.

Proof. The very first claim follows directly from the definition of D_0 .

We have $M \setminus T \subseteq N_D^-(v)$ and $N_D^+(M) \subseteq \{v\} \cup N_D^+(v)$. This directly implies that $(x_{M\setminus T}, v) \in A(D_0)$ and that $N_{D_0}^+(x_{M\setminus T}) \subseteq N_D^+(M) \subseteq N^+(v) \cup \{v\} = N_{D_0}^+(v) \cup \{v\}$. Hence, $v \in N_{D_0}^+(x_{M\setminus T})$ has an out-arc to every other out-neighbor of $x_{M\setminus T}$ in D_0 , and this shows (by definition) that $(x_{M\setminus T}, v) \in A(F(D_0))$.

For the second claim, suppose towards a contradiction that there exists $u \in N_{D_0}^+(v)$ such that $(u, x_{M\setminus T}) \in A(D_0)$. By definition of D_0 , this means that $u \in N_D^+(v)$ and that there exists a vertex $m \in M \setminus T$ such that $(u, m) \in A(D)$. By definition of T, this however shows that $m \in T$, a contradiction.

In the following, let D^* be the digraph defined by

$$V(D^*) := V(D_0), A(D^*) := A(D_0) \cup \{(x_{M \setminus T}, u) | u \in N_{D_0}^+(v)\}.$$

Claim 5. $D^* \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+).$

Proof. Let $e_i = (x_{M \setminus T}, u_i), i = 1, ..., k$ be a list of the arcs contained in $A(D^*) \setminus A(D_0)$ for some $k \ge 0$. For $0 \le i \le k$ let D_i denote the digraph defined by $V(D_i) := V(D_0)$ and $A(D_i) := A(D_0) \cup \{e_1, \ldots, e_i\}$. Note that $D_k = D^*$.

We now claim that $D_i \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+)$ and $(x_{M\setminus T}, v) \in A(F(D_i))$ for every $i \in \{0, 1, \ldots, k\}$ and prove this claim by induction on i.

For i = 0 the claim holds true by the previous discussions and Claim 4. Now let $1 \le i \le k$ and suppose we know that the claim holds for D_{i-1} .

Note that D_i is the digraph obtained from D_{i-1} by adding the arc $e_i = (x_{M\setminus T}, u_i)$, where $u_i \in N_{D_0}^+(v) = N_{D_{i-1}}^+(v)$, $(x_{M\setminus T}, v) \in A(F(D_{i-1}))$. Further note that $e_i \notin A(D_{i-1})$, as well as $(u_i, x_{M\setminus T}) \notin A(D_{i-1})$ by Claim 4. Hence, we may apply Lemma 9.10 to the digraph D_{i-1} with $x = x_{M\setminus T}, y = v, z = u_i$ to find that indeed $D_i \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+)$. It remains to show that $(x_{M\setminus T}, v) \in A(F(D_i))$. However, the only new out-neighbor of $x_{M\setminus T}$ in D_i compared to D_{i-1} is the vertex u_i , which is still dominated by the vertex $v \in N_{D_i}^+(x_{M\setminus T})$ via the arc $(v, u_i) \in A(D_i)$, and hence v still dominates all other outneighbors of $x_{M\setminus T}$ in D_i . This shows that D_i satisfies the induction claim.

We have proved that $D^* = D_k \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^+)$ and $(x_{M\setminus T}, v) \in A(F(D_i))$. This concludes the proof of Claim 5.

The number of vertices of D^* satisfies

$$v(D^*) = v(D_0) = v(D) - |T| - (|M \setminus T| - 1) \le v(D) - (|M| - 1) \le v(D) - 2 < v(D)$$

since $|M| \geq 3$ by Claim 3. Hence, the minimality of D implies that the assertion of the theorem holds for D^* . Applying this assertion to the vertex $x_{M\setminus T}$ in D^* , we find that there exists an acyclic 2-coloring $c^* : V(D^*) \to \{1,2\}$ of D^* such that $c^*(x_{M\setminus T}) = 1 = c^*(u)$ for every $u \in N_{D^*}^+(x_{M\setminus T})$. Using the facts $N_{D_0}^+(x_{M\setminus T}) \subseteq N_D^+(v) \cup \{v\}, N_{D_0}^+(v) = N_D^+(v)$ and $(x_{M\setminus T}, v) \in A(D_0)$, the definition of D^* yields that $N_{D^*}^+(x_{M\setminus T}) = N_D^+(v) \cup \{v\}$. Hence, we have $c^*(x_{M\setminus T}) = c^*(v) = c^*(u) = 1$ for every $u \in N_D^+(v)$.

Let $c: V(D) \to \{1,2\}$ be the coloring of D defined by $c(x) := c_M(x)$ for every $x \in M$, and $c(x) := c^*(x)$ for every $x \in V(D) \setminus M$. We note that c(v) = c(u) for all $u \in N_D^+(v)$. Hence, by the initial assumption on D, the coloring c cannot be acyclic, i.e., there is a directed cycle C in D which is monochromatic in the coloring c. Then we must have $V(C) \setminus M \neq \emptyset$, for otherwise C would be a monochromatic directed cycle in the acyclic coloring c_M of D_M . Analogously, if $V(C) \cap M = \emptyset$, then C would be a directed cycle in $D - M \subseteq (D - T)/(M \setminus T) = D_0 \subseteq D^*$, a contradiction. Therefore we also have $V(C) \cap M \neq \emptyset$, and hence there must be an arc $(x, y) \in A(C)$ such that $x \in M$ and $y \notin M$. However, this means that $y \in N_D^+(M) \subseteq \{v\} \cup N_D^+(v)$, and hence c(y) = 1 by the above. Since C is monochromatic, it follows that $V(C) \subseteq c^{-1}(1)$. In particular, since $c(t) = c_M(t) = 2$ for every $t \in T$, it follows that C is a directed cycle in D-T. Let z be the first vertex of M we meet when traversing the directed cycle C in forward-direction, starting at y. Then $z \in M \setminus T$, and P := C[x, y] is a monochromatic directed x, z-path in D-T of length at least two, such that $V(P) \cap M = \{x, z\}$ and $V(P) \subseteq c^{-1}(1)$. Now $(V(P) \setminus \{x, z\}) \cup \{x_{M \setminus T}\}$ forms the vertex-set of a directed cycle C^* in $(D-T)/(M \setminus T) = D_0 \subseteq D^*$ containing $x_{M \setminus T}$, and we have $c^*(x) = c(x) = 1$ for every vertex $x \in V(C^*) \setminus \{x_{M \setminus T}\} = V(P) \setminus \{x, z\} \subseteq V(D) \setminus M$. We have $c^*(x_{M \setminus T}) = 1$ by definition of c^* , and hence C^* forms a monochromatic directed cycle of color 1 in the acyclic coloring c^* of D^* . This contradiction finally shows that our very first assumption, namely that a (smallest) counterexample D to the claim of the theorem exists, was wrong. This concludes the proof of the theorem.

9.3 $\{S_2^+, W_3^-\}$ -Free Oriented Graphs

In this section we prove Theorem 9.6, showing that every digraph in $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^-)$ is acyclically 4-colorable. Note that $\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^-)$ is the class of oriented graphs D such that the out-neighborhood of any vertex in D spans a tournament, and the inneighborhood of any vertex spans a directed triangle-free graph. In fact, we show the following strengthened statement.

Theorem 9.14. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+, W_3^-)$ and let $(u, v) \in A(D)$. Then D admits an acyclic coloring $c : V(D) \to \{1, 2, 3, 4\}$ satisfying the additional conditions c(u) = 1, c(x) = 1 for all $x \in N_D^+(u) \setminus N_D^+(v)$ (so c(v) = 1) and $c(x) \in \{1, 2\}$ for all $x \in N_D^+(v)$.

Proof. Suppose towards a contradiction the claim of the theorem was wrong, and let D be a counterexample minimizing v(D). Then there exists an arc $(u, v) \in A(D)$ such that D does not admit an acyclic coloring $c : V(D) \to \{1, 2, 3, 4\}$ satisfying the additional conditions c(u) = 1, c(x) = 1 for all $x \in N_D^+(u) \setminus N_D^+(v)$ and $c(x) \in \{1, 2\}$ for all $x \in N_D^+(v)$. Let us define $A := N_D^+(u) \setminus (N_D^+(v) \cup \{v\})$ and $B := N_D^-(u) \cap N_D^+(v)$. We start with some useful observations concerning these sets.

Claim 1. $A \subseteq N_D^-(v)$, and D[A] and D[B] are (possibly empty) transitive tournaments.

Proof. To show $A \subseteq N_D^-(v)$, let $x \in A = N_D^+(u) \setminus (N_D^+(v) \cup \{v\})$ be arbitrary. Since $(u, x), (u, v) \in A(D)$ and x, u, v cannot induce an S_2^+ in D, the vertices x and v must be equal or adjacent in D. Since $x \notin N_D^+(v) \cup \{v\}$, it follows that $x \in N_D^-(v)$, as claimed.

Since $D[N_D^+(u)]$ is a tournament and $A \subseteq N_D^+(u)$, also D[A] is a tournament. Furthermore $D[N_D^-(v)]$ is directed triangle-free, and with $A \subseteq N_D^-(v)$ also D[A] is directed triangle-free, i.e., a transitive tournament, as claimed.

Similarly, since $D[N_D^-(u)]$ is directed triangle-free, and since $D[N_D^+(v)]$ is a tournament, $B = N_D^-(u) \cap N_D^+(v)$ implies that D[B] must be both directed triangle-free and a tournament, i.e., a transitive tournament.

In the following, let us denote by $D' := D - (N_D^-(u) \cup \{u\})$ the induced subdigraph of D obtained by deleting the closed in-neighborhood of u. We clearly have $D' \in \operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, W_3^-)$ and v(D') < v(D), and hence, by minimality of D, the theorem statement holds for D'.

Claim 2. There exists an acyclic coloring $c' : V(D') \to \{1, 2, 3, 4\}$ of D' such that c'(v) = 1, c'(x) = 1 for all $x \in A$ and $c'(x) \in \{1, 2\}$ for all $x \in N_{D'}^+(v)$.

Proof. We distinguish the two cases $A = \emptyset$ and $A \neq \emptyset$.

Suppose first that $A = \emptyset$. If $N_{D'}^+(v) = \emptyset$, then applying the theorem statement to D'(for an arbitrarily chosen arc) yields that $\vec{\chi}(D') \leq 4$, and hence there exists an acyclic coloring $c': V(D') \to \{1, 2, 3, 4\}$. Since v is a sink in D', no directed cycle in D' contains v. Consequently, we may assume w.l.o.g. (possibly by recoloring) that c'(v) = 1. In particular, since $A = \emptyset, N_{D'}^+(v) = \emptyset$, the remaining two statements of Claim 2 are satisfied vacuously for c', concluding the proof in this case.

On the other hand, if $N_{D'}^+(v) \neq \emptyset$, then there exists an arc in D' leaving v. Fix an arbitrary such arc (v, y). Applying the Theorem statement to this arc in D', we find that there is an acyclic coloring $c': V(D') \rightarrow \{1, 2, 3, 4\}$ such that c'(v) = 1, c'(x) = 1 for all $x \in N_{D'}^+(v) \setminus N_{D'}^+(y)$ and $c'(x) \in \{1, 2\}$ for all $x \in N_{D'}^+(y)$; in particular, and $c'(x) \in \{1, 2\}$ for all $x \in N_{D'}^+(v)$. Again, this shows that the claim holds true.

Next suppose that $A \neq \emptyset$. By Claim 1, D[A] is a transitive tournament. Let a be the unique source-vertex of this tournament. Since $a \in A \subseteq N_D^-(v)$, it follows that $(a, v) \in A(D')$. Hence, we may apply the theorem statement to the arc (a, v) in D' and find that there exists an acyclic coloring $c' : V(D') \rightarrow \{1, 2, 3, 4\}$ such that c'(a) = 1, c'(x) = 1 for all $x \in N_{D'}^+(a) \setminus N_{D'}^+(v)$ (in particular c'(v) = 1), and $c'(x) \in \{1, 2\}$ for all $x \in N_{D'}^+(v)$. Since a is the source of D[A] = D'[A] and since $A \cap N_D^+(v) = \emptyset$, we have $\{a\} \cup (N_{D'}^+(a) \setminus N_{D'}^+(v)) \supseteq A$ and thus c'(x) = 1 for all $x \in A$, as required. This shows the assertion of Claim 2 and concludes the proof.

Claim 3. There exists an acyclic coloring $c^-: N_D^-(u) \to \{2, 3, 4\}$ of $D[N_D^-(u)]$ such that $c^-(x) = 2$ for all $x \in B$ and $c^-(x) \in \{3, 4\}$ for all $x \in N_D^-(u) \setminus B$.

Proof. Since $D \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, W_3^-)$, we have $D[N_D^-(u)] \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, \vec{C}_3)$. By Theorem 9.4 there exists an acyclic coloring of $D[N_D^-(u)] - B$ using only colors 3 and 4. Clearly, D[B] as a transitive tournament (see Claim 1) admits an acyclic coloring only with color 2. Putting these colorings together yields an acyclic coloring c' of $D[N_D^-(u)]$ with the required properties.

Let $c: V(D) \to \{1, 2, 3, 4\}$ be the coloring defined by c(u) := 1, $c(x) := c^-(x)$ for all $x \in N_D^-(u)$ and c(x) := c'(x) for all $x \in V(D) \setminus (N_D^-(u) \cup \{u\})$.

Note that from the properties of c' given by Claim 2 we have c(u) = c(v) = 1 and c(x) = 1 for all $x \in A = N_D^+(u) \setminus (N_D^+(v) \cup \{v\})$. Furthermore, since $N_D^+(v) = N_{D'}^+(v) \cup B$, the properties of c' and c^- imply that $c(x) \in \{1, 2\}$ for all $x \in N_D^+(v)$.

Given these properties, our initial assumption concerning D implies that c cannot be an acyclic coloring of D, that is, there is a directed cycle C in D which is monochromatic under c. Since $c' = c|_{V(D')}$ and $c^- = c|_{N_D^-(u)}$, we must have $V(C) \cap (N_D^-(u) \cup \{u\}) \neq \emptyset$ and $V(C) \setminus N_D^-(u) \neq \emptyset$, for otherwise c' resp. c⁻ would not be acyclic. Further note that $u \notin V(C)$, for every arc in D entering u has its tail colored with either 2,3 or 4, while its head, u, receives color 1 under c (so a directed cycle containing u cannot be monochromatic). Hence, there must exist an arc $(x, y) \in A(C)$ such that $x \in N_D^-(u)$ and $y \in V(D) \setminus (N_D^-(u) \cup \{u\})$. Since $(x, u), (x, y) \in A(D)$ and D is induced S_2^+ -free, u and y must be equal or adjacent, and since $y \notin N_D^-(u) \cup \{u\}$, we have $y \in N_D^+(u)$. By the properties of c' and c^- , we have $N_D^+(u) \setminus N_D^+(v) = A \cup \{v\} \subseteq c^{-1}(\{1\}), B \subseteq c^{-1}(\{2\})$ and $N_D^-(u) \setminus B \subseteq c^{-1}(\{3,4\})$. The cycle C is monochromatic, therefore c(x) = c(y). From this we conclude that $y \in N_D^+(v)$, and hence $c(y) \in \{1, 2\}$. This is only possible if c(x) = c(y) = 2, and hence $x \in B$. It follows that $x, y, u, v \in V(D)$ are distinct vertices satisfying $(x, y), (u, y), (v, y) \in A(D)$, as well as $(x, u), (u, v), (v, x) \in A(D)$ (here we used that $x \in B = N_D^-(u) \cap N_D^+(v)$. This however means that x, y, u, v induce a copy of $W_3^$ in D, which is absurd considering that $D \in \operatorname{Forb}_{\operatorname{ind}}(K_2, S_2^+, W_3^-)$. This shows that our very first assumption concerning the existence of a smallest counterexample D was wrong. This concludes the proof of the Theorem.

9.4 Adding a Dominating Sink to a Hero

In this section our goal is to prove Theorem 9.7. Let us first prove the following lemma.

Lemma 9.15. Let $D \in Forb_{ind}(\vec{K}_2, S_2^+)$ and let $C \in \mathbb{N}$ be such that $\vec{\chi}(D[N_D^-(x)]) \leq C$ for every $x \in V(D)$. Let $u, v \in V(D)$ and let P be a shortest u-v-dipath in D. Let $X := V(P) \cup N_D^-(V(P))$. Then $\vec{\chi}(D[X]) \leq 3C + 2$.

Proof. Let $u = x_0, x_1, \ldots, x_{\ell-1}, x_\ell = v$ be the vertex-trace of P and consider the partition $(A_i)_{i=1}^{\ell}$ of $N_D^-(V(P))$ where $A_i := N_D^-(x_i) \setminus (V(P) \cup \bigcup_{1 \le j \le i} A_j), i = 0, \ldots, \ell$.

Claim. Let $0 \le i < j \le \ell$ with $j - i \ge 3$. Then there exists no arc in D starting in A_i and ending in A_j .

Proof. Suppose towards a contradiction that there are vertices $x \in A_i$, $y \in A_j$ with $(x, y) \in A(D)$. Since $(x, x_i) \in A(D), (x, y) \in A(D)$ and $x_i \neq y$ (since $x_i \in V(P)$ and $y \notin V(P)$), x_i and y must be adjacent in D. By definition of A_j we have $A_j \cap N_D^-(x_i) = \emptyset$ and hence $(x_i, y) \in A(D)$. However, now the directed path described by the vertices $u = x_0, x_1, \ldots, x_i, y, x_j, \ldots, x_\ell = v$ is a u-v-dipath in D shorter than P, a contradiction. This proves the claim.

For every $0 \leq i \leq \ell$ we have $\vec{\chi}(D[A_i]) \leq \vec{\chi}(D[N_D^-(x_i)]) \leq C$. Let us define the set $B_r := \bigcup \{A_i | i \equiv r \pmod{3}\}$ for every $r \in \{0, 1, 2\}$. From the above claim it follows that no directed cycle in $D[B_r]$ intersects two different sets A_i, A_j . Hence, we have

$$\vec{\chi}(D[B_r]) \le \max\{\vec{\chi}(D[A_i]) | i \equiv r \pmod{3}\} \le C$$

for r = 0, 1, 2. Further note that the two sets

$$V_0 := \{x_i | i \in \{0, \dots, \ell\} \text{ even}\}, V_1 := \{x_i | i \in \{0, \dots, \ell\} \text{ odd}\}$$

both induce acyclic subdigraphs of D, for otherwise D would not be a shortest *u-v*-dipath in D. Since X is the disjoint union of B_0, B_1, B_2, V_0, V_1 , we conclude

$$\vec{\chi}(D[X]) \le \vec{\chi}(D[B_0]) + \vec{\chi}(D[B_1]) + \vec{\chi}(D[B_2]) + \vec{\chi}(D[V_0]) + \vec{\chi}(D[V_1]) \le 3C + 2,$$

as required.

Proof of Theorem 9.7. Let $\{ \vec{K}_2, S_2^+, H \}$ be heroic and $C := \vec{\chi}(\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, H)).$

We claim that every digraph $D \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, H^-)$ admits an acyclic coloring with $C^- := v(H)(C+1) + 3C + 2$ colors.

Suppose towards a contradiction that there exists some $D \in \operatorname{Forb}_{\operatorname{ind}}(\overline{K}_2, S_2^+, H^-)$ with $\vec{\chi}(D') > C'$, and choose such a D minimizing v(D). Then we have $\vec{\chi}(D) > C' \geq C$ and hence there is $Y \subseteq V(D)$ such that D[Y] is isomorphic to H. Furthermore, since the dichromatic number of D is the maximum of the dichromatic numbers of its strong components, the minimality of v(D) implies that D is strongly connected.

Let $S \supseteq Y$ denote a set of vertices in D defined as follows:

If D[Y] (resp. H) is strongly connected, put S := Y. Otherwise, let Y_1, \ldots, Y_t be a partition of Y into the $t \ge 2$ strong components of D[Y] such that all arcs between Y_i and Y_j start in Y_i and end in Y_j , for any $1 \le i < j \le t$ (note that since D[Y] is a tournament all elements of $Y_i \times Y_j$ are arcs of D[Y] for $1 \le i < j \le t$). Now pick $u \in Y_t, v \in Y_1$ arbitrarily, let P be a shortest u-v-dipath in D and put $S := V(P) \cup Y$. Let us note that in any case, D[S] is strongly connected.

Let $Z := S \cup N_D^-(S)$. Then we have $Z = X \cup Y \cup N_D^-(Y)$, where $X := V(P) \cup N_D^-(V(P))$. For every $x \in V(D)$ we know that since D is H^- -free, the digraph $D[N_D^-(x)]$ is contained in Forb_{ind} (\vec{K}_2, S_2^+, H) , and hence $\vec{\chi}(D[N_D^-(x)]) \leq C$. Using Lemma 9.15 we obtain that $\vec{\chi}(D[X]) \leq 3C + 2$. Putting it all together, we find that

$$\vec{\chi}(D[Z]) \le \sum_{y \in Y} \underbrace{\vec{\chi}(D[\{y\} \cup N_D^-(y)])}_{\le C+1} + \vec{\chi}(D[X]) \le v(H)(C+1) + 3C + 2 = C'.$$

Claim. No arc in D leaves Z.

Proof. We first show that there do not exist $z \in S$, $w \in V(D) \setminus Z$ such that $(z, w) \in A(D)$. Suppose towards a contradiction that such an arc (z, w) exists. We claim that then $(s, w) \in A(D)$ for every $s \in S$. Consider $s \in S$ arbitrarily. Since D[S] is strongly connected, there exist vertices $z = s_0, s_1, \ldots, s_k = s$ in S such that $(s_{i-1}, s_i) \in A(D)$, $1 \leq i \leq k$. We show $(s_i, w) \in A(D)$ for all $i = 0, \ldots, k$ by induction on i. Clearly it is true for i = 0, so suppose that $1 \leq i \leq k$ and we have established that $(s_{i-1}, w) \in A(D)$. Since $w \notin Z$, $s_i \in Z$, we have $w \neq s_i$ and $(s_{i-1}, w), (s_{i-1}, s_i) \in A(D)$. Since D is S_2^+ -free, it follows that s_i and w are adjacent. However, since $w \notin Z = S \cup N_D^-(S) \supseteq N_D^-(s_i)$, we must have $(s_i, w) \in A(D)$, as claimed.

This shows that indeed $(s, w) \in A(D)$ for all $s \in S$. However, since $S \supseteq Y$ and since D[Y] is isomorphic to H, it follows that $D[Y \cup \{w\}]$ is an induced subdigraph of Disomorphic to H^- , a contradiction to $D \in \operatorname{Forb}_{\operatorname{ind}}(\overset{\leftrightarrow}{K}_2, S_2^+, H^-)$. This shows that there are no arcs from S to $V(D) \setminus Z$.

To complete the proof, let us show that there are no arcs starting in $Z \setminus S = N_D^-(S)$ that end in $V(D) \setminus Z$. Suppose towards a contradiction that there exist $z \in N_D^-(S)$ and $w \in V(D) \setminus Z$ with $(z, w) \in A(D)$. Then there is a vertex $s \in S$ such that $(z, s) \in A(D)$. Since $s \neq w$, $(z, s), (z, w) \in A(D)$ and D is S_2^+ -free we find that s and w are adjacent. Since $w \notin Z \supseteq N_D^-(s)$, it follows that $(s, w) \in A(D)$. However, this yields a contradiction, since we showed above that no arc in D starts in S and ends in $V(D) \setminus Z$.

All in all, the claim follows.

Since D is strongly connected and $Z \neq \emptyset$ (since $Z \supseteq Y$), it follows that Z = V(D), and hence that $\vec{\chi}(D) = \vec{\chi}(D[Z]) \leq C'$, a contradiction to our initial assumption. This concludes the proof of the theorem.

9.5 Oriented 4-Vertex-Paths

In this section we establish that $\{\vec{K}_2, \vec{K}_3, P^+(1, 1, 1)\}$ is heroic, proving Theorem 9.8.

Proof of Theorem 9.8. We prove by induction on n that every directed graph on n vertices $D \in \text{Forb}_{\text{ind}}(\vec{K}_2, \vec{K}_3, P^+(1, 1, 1))$ admits an acyclic 2-coloring. The claim trivially holds for n = 1, so let $n \ge 2$ and suppose that every digraph in $\text{Forb}_{\text{ind}}(\vec{K}_2, \vec{K}_3, P^+(1, 1, 1))$ having less than n vertices is 2-colorable. Pick some $x \in V(D)$ arbitrarily. Let us define a sequence X_0, X_1, X_2, \ldots of subsets of V(D) as follows:

$$X_{i} := \begin{cases} \{x\}, & \text{if } i = 0, \\ N^{+}(X_{i-1}) \setminus \bigcup_{j=0}^{i-1} X_{j}, & \text{if } i \text{ odd}, \\ N^{-}(X_{i-1}) \setminus \bigcup_{j=0}^{i-1} X_{j}, & \text{if } i \ge 2 \text{ even} \end{cases}$$

The sets $(X_i)_{i\geq 0}$ are pairwise disjoint by definition, and so there exists $k \geq 1$ such that $X_1, \ldots, X_k \neq \emptyset$ and $X_i = \emptyset$ for all i > k.

Claim. X_i is an independent set of D for every $i \ge 0$.

Proof. We prove the claim by induction on i. The claim trivially holds for i = 0 since $X_0 = \{x\}$, and since D does not contain a transitive triangle, also $X_1 = N^+(x)$ must be an independent set in D. Now let $i \geq 2$ and suppose that we already established that X_0, \ldots, X_{i-1} are independent. To show that X_i is independent, let us suppose towards a contradiction that there are $x, y \in X_i$ such that $(x, y) \in A(D)$. By definition of the sets X_i there are vertices $x_1, y_1 \in X_{i-1}$ and $x_2, y_2 \in X_{i-2}$ such that the following holds: $(x_1, x_2), (x_1, x), (y_1, y_2), (y_1, y) \in A(D)$ if i is odd, respectively $(x_2, x_1), (x, x_1), (y_2, y_1), (y_1, y_2), (y_1, y_2), (y_2, y_1), (y_2, y_2), (y_1, y_2), (y_2, y_2), (y_2, y_2), (y_1, y_2), (y_2, y_2),$ $(y, y_1) \in A(D)$ if i is even. We must have $x_1 \neq y_1$ in any case, since otherwise the vertices $x_1 = y_1, x, y$ would induce a \vec{K}_3 in D. Let us now consider the oriented 4-vertexpath P in D defined as $P = x, (x, y), y, (y_1, y), y_1, (y_1, y_2), y_2$ if i is odd, respectively as $P = x_2, (x_2, x_1), x_1, (x, x_1), x, (x, y), y$ if i is even. In order for this path not to be an induced copy of $P^+(1,1,1)$, two non-consecutive vertices of the path must be adjacent. However, since D does not contain transitive triangles, this is only possible if x and y_2 (i odd) respectively x_2 and y (*i* even) are adjacent. Since $x \notin X_{i-1}$, we have $x \notin N^-(X_{i-2})$ if *i* is odd and $y \notin N^+(X_{i-2})$ if *i* is even. Since $x_2, y_2 \in X_{i-2}$ we find that $(y_2, x) \in A(D)$ if *i* is odd and $(y, x_2) \in A(D)$ if i is even. In both cases we conclude that $x_2 \neq y_2$, since otherwise the vertices $x_2 = y_2, x_1, x$ respectively $x_2 = y_2, y_1, y$ would induce a transitive triangle in D. Now consider the oriented path Q in D defined as $Q = y_2, (y_2, x), x, (x_1, x), x_1, (x_1, x_2), x_2$ if i is odd and as $Q = y_2, (y_2, y_1), y_1, (y, y_1), y, (y, x_2), x_2$ if i is even. In order for Q not to be an induced copy of $P^+(1,1,1)$ in D and since D does not contain transitive triangles, this implies in both cases that the endpoints x_2 and y_2 of Q must be adjacent. This contradicts the induction hypothesis that X_{i-2} is an independent set. Hence, our assumption was wrong, X_i is indeed independent. This concludes the proof of the claim.

Let $X := X_0 \cup \cdots \cup X_k$ and D' := D - X. By the induction hypothesis D' admits an acyclic 2-coloring $c': V(D') \to \{1, 2\}$. Let us now define $c: V(D) \to \{1, 2\}$ by c(x) := c'(x)for every $x \in V(D) \setminus X$, c(x) := 1 for every $x \in X_i$ such that i is even, and c(x) := 2for every $x \in X_i$ such that i is odd. We claim that D defines an acyclic coloring of D. Suppose towards a contradiction that there exists a monochromatic directed cycle C in (D, c). Since c' is an acyclic coloring of D', we must have $V(C) \cap (X_0 \cup \cdots \cup X_k) \neq \emptyset$. Note that by definition of the sets $(X_i)_{i\geq 0}$ we have $N^+(\bigcup_{i \text{ even }} X_i), N^-(\bigcup_{i \text{ odd }} X_i) \subseteq X$. Hence, no arc of D starts in $c^{-1}(\{1\}) \cap X$ and ends in $V(D) \setminus X$, and no arc of D starts in $c^{-1}(\{2\}) \cap X$ and ends in $V(D) \setminus X$. Since $V(C) \subseteq c^{-1}(t)$ for some $t \in \{1, 2\}$, the strong connectivity of C shows that in fact $V(C) \subseteq c^{-1}(t) \cap X$ for some $t \in \{1, 2\}$. Let $i_0 \geq 0$ be smallest such that $X_{i_0} \cap V(C) \neq \emptyset$. Let $u \in X_{i_0} \cap V(C) \neq \emptyset$, and let $u^{-}, u^{+} \in V(C)$ be such that $(u^{-}, u), (u, u^{+}) \in A(C)$. Since X_{i_0} is an independent set, and by definition of the coloring c, we must have $u^- \in X_{i_0+2s^-}, u^+ \in X_{i_0+2s^+}$ for integers $s^-, s^+ \geq 1$. On the other hand, we have $u^+ \in N^+(X_{i_0}) \setminus \bigcup_{j=0}^{i_0-1} X_j = X_{i_0+1}$ if i_0 is even and $u^- \in N^-(X_{i_0}) \setminus \bigcup_{j=0}^{i_0-1} X_j = X_{i_0+1}$ if i_0 is odd, in both cases yielding a contradiction since $X_{i_0+2s^+}, X_{i_0+2s^-}$ are disjoint from X_{i_0+1} . This shows that our assumption was wrong, indeed, c is an acyclic coloring of D. Hence, $\vec{\chi}(D) \leq 2$, concluding the proof.

9.6 Conclusion

In the first three sections of this chapter we have proved that set $\{\vec{K}_2, S_2^+, H\}$ is heroic for several small heroes H and, in particular, we resolved Conjecture 9.4. It would be interesting to prove that in fact, for any hero H, $\{\vec{K}_2, S_2^+, H\}$ is heroic, as this would be a generalization of the main results of Berger et al. [BCC⁺13] from tournaments to *locally out-complete* oriented graphs, i.e., oriented graphs in which the out-neighborhood of every vertex induces a tournament. The smallest open case of this problem, would be to show that $\{\vec{K}_2, S_2^+, \vec{K}_4^s\}$ is heroic, where \vec{K}_4^s denotes the unique strong tournament on four vertices. It seems that already for this small case a new method is required. We do however believe that the following is true.

Conjecture 9.5. $\vec{\chi}(Forb_{ind}(\vec{K}_2, S_2^+, \vec{K}_4^s)) = 3.$

Here, a tight lower bound would be provided by the following construction: Take a 3-fold blow-up of a directed 4-cycle (every arc being replaced by a $\vec{K}_{3,3}$) and connect each of the three blow-up triples by a directed triangle. This oriented graph is contained in Forb_{ind}($\vec{K}_2, S_2^+, \vec{K}_4^s$) and has dichromatic number 3.

Let us further remark at this point that there exists a very simple proof that if we exclude both S_2^+ and S_2^- , i.e., we consider *locally complete* oriented graphs (where the inand out-neigborhood of every vertex induces a tournament), then we can show that the exclusion of any hero indeed bounds the dichromatic number as follows.

Remark 9.16. For any hero H, we have

$$\vec{\chi}(\mathit{Forb}_{\mathit{ind}}(\overset{\leftrightarrow}{K}_2,S_2^+,S_2^-,H)) \leq 2\vec{\chi}(\mathit{Forb}_{\mathit{ind}}(\overset{\leftrightarrow}{K}_2,\overline{K}_2,H)) < \infty$$

Proof. By the result of Berger et al. [BCC⁺13] we have $C := \vec{\chi}(\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, \overline{K}_2, H)) < \infty$. Let us now prove that $\vec{\chi}(\operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, S_2^-, H)) \leq 2C$. Towards a contradiction suppose that $\vec{\chi}(D) > 2C$ for some $D \in \operatorname{Forb}_{\operatorname{ind}}(\vec{K}_2, S_2^+, S_2^-, H)$, and pick D such that v(D) is minimum. Pick $v \in V(D)$ arbitrarily and define $D' := D - (\{v\} \cup N_D(v))$.

Since v(D') < v(D), there exists an acyclic 2*C*-coloring $c': V(D') \to \{1, \ldots, 2C\}$ of D'. Since D is induced S_2^+, S_2^- -free, we further know that $D^+ := D[\{v\} \cup N_D^+(v)]$ and

 $D^{-} := D[N_{D}^{-}(v)] \text{ are tournaments excluding } H. \text{ It follows from the definition of } C \text{ that there exists an acyclic } C\text{-coloring } c^{+} : V(D^{+}) \to \{1, \ldots, C\} \text{ of } D^{+} \text{ as well as an acyclic coloring } c^{-} : V(D^{-}) \to \{C + 1, \ldots, 2C\} \text{ of } D^{-}. \text{ Let } c \text{ be the } 2C\text{-coloring of } D \text{ defined as the common extension of } c', c^{+}, c^{-} \text{ to } V(D). \text{ Since } \vec{\chi}(D) > 2C \text{ there exists a directed cycle } C \text{ which is monochromatic under } c. \text{ Since } c', c^{+}, c^{-} \text{ are acyclic colorings and since the color sets used by } c^{+} \text{ and } c^{-} \text{ are disjoint, we must have } V(C) \cap (\{v\} \cup N_D(v)) \neq \emptyset, V(C) \setminus (\{v\} \cup N_D(v)) \neq \emptyset. \text{ Since all in-neighbors of } v \text{ have a distinct color from } v, we further have } v \notin V(C). We conclude that there are vertices <math>x_1, x_2 \in V(C) \cap N_D(v), y_1, y_2 \in V(C) \setminus (\{v\} \cup N_D(v)) \text{ such that } (x_1, y_1), (y_2, x_2) \in A(C). \text{ We claim that } x_1 \in N_D^+(v) \text{ and } x_2 \in N_D^-(v). \text{ Indeed, otherwise we would have } (x_1, v) \in A(D) \text{ or } (v, x_2) \in A(D), \text{ and then either the vertices } x_1, v, y_1 \text{ induce an } S_2^+ \text{ in } D, \text{ or } x_2, v, y_2 \text{ induce an } S_2^- \text{ in } D, \text{ in each case yielding a contradiction to } D \in \text{Forb}_{\text{ind}}(\vec{K}_2, S_2^+, S_2^-, H). \text{ Finally, we conclude that } c(x_1) = c^+(x_1) \leq C < c^-(x_2) = c(x_2), \text{ a contradiction to the facts that } C \text{ is monochromatic and } x_1, x_2 \in V(C). \text{ This shows that our initial assumption concerning the existence of } D \text{ was wrong, concluding the proof of the remark.}$

In the last section of this chapter we investigated oriented graphs excluding the antioriented 4-vertex-path $P^+(1,1,1)$. It would certainly be interesting and insightful to generalize both Theorem 9.8 as well as the result of Aboulker et al. concerning \vec{P}_4 by proving that $\{\vec{K}_2, \vec{P}_4, \vec{K}_k\}$ and $\{\vec{K}_2, P^+(1,1,1), \vec{K}_k\}$ are heroic for any $k \ge 4$.

Chapter 10

Complexity Results

10.1 Introduction

The graphs and digraphs considered in this chapter are simple.

In the previous chapters, we have investigated theoretical upper and lower bounds on the dichromatic number. However, given a concrete digraph, these upper bounds might not be tight enough to help us determine its dichromatic number, or an optimal acyclic coloring of the digraph, in a reasonable amount of time. While determining the dichromatic number of a fixed digraph is certainly a finite problem (at worst we could enumerate all possible colorings of the digraph with a given number of colors and check for each such coloring whether every color class induces an acyclic digraph), the running time of a brute force search is at least exponential and therefore explodes rather quickly. For that reason, in this chapter, we study the dichromatic number from an algorithmic point of view. Concretely, for a fixed number $k \in \mathbb{N}$, we aim at finding algorithms solving the following decision problem (possibly provided that the input digraphs D have some additional structure).

Problem 10.1 (Digraph k-Coloring (k-DCP)). Given as input a digraph D, does there exist a proper k-coloring for D?

It is a well-known NP-hardness result [GJ90] that testing whether a given undirected graph G is properly k-colorable is NP-complete for any fixed value of $k \ge 3$.

Along similar lines, the problem of deciding whether a given digraph D has dichromatic number at most k has been shown to be NP-complete for all $k \ge 2$, see [BFJ⁺04, FHM03] for the first proofs of this fact.

Although NP-complete problems do not admit polynomial-time solution algorithms for general inputs provided that $P \neq NP$, if the input objects are guaranteed to have additional structure, polynomial-time algorithms for these restricted input classes may still be found. This idea has given rise to a whole branch in the field of algorithm design, which is still quite actively pursued at present. In particular, much research has focused on the existence of *parametrized algorithms* for NP-hard problems, whose most prominent representatives are the so-called XP-and FPT-algorithms. The following is the formal definition of these concepts which we will use in this chapter.

Definition 10.1 (XP- and FPT-algorithms). Let \mathcal{P} be an algorithmic problem whose inputs come from a set \mathcal{X} of objects, and let $\ell : \mathcal{X} \to \mathbb{N}_0$ be a function measuring the coding length of the objects in \mathcal{X} . Let further $p : \mathcal{X} \to \mathbb{N}$ be a parameter for problem \mathcal{P} .

An algorithm solving problem \mathcal{P} correctly on the input set \mathcal{X} is called an XP-algorithm with respect to parameter p if there exists a computable function $f : \mathbb{N} \to \mathbb{N}$ and an absolute constant C > 0 such that for every $x \in \mathcal{X}$ the running time¹ r(x) of the algorithm satisfies

$$r(x) \le C\ell(x)^{f(p(x))}$$

Further, if there exist $k \in \mathbb{N}$ and a computable function $f : \mathbb{N} \to \mathbb{N}$ such that

$$r(x) \le f(p(x))\ell(x)^{\ell}$$

for all $x \in \mathcal{X}$, then the algorithm is called an FPT-algorithm with respect to parameter p.

It follows directly from the definition that every FPT-algorithm is also an XP-algorithm, however with a qualitatively better running-time guarantee. The main benefit of having an XP-algorithm for a particular problem is that it becomes possible to solve the problem in polynomial time on any class $\{x \in \mathcal{X} | p(x) \leq k\}$ of inputs for which the *parameter* p(x) is bounded by a *fixed* number $k \in \mathbb{N}$ (this also explains the term fixed-parametertractability). For further details and more extensive background on this topic we refer the interested reader to [DF99].

A famous result in the theory of parametrized algorithms, *Courcelle's Theorem* [Cou90] implies that a wide range of problems², and in particular the computation of the chromatic number of a given graph, admit FPT-algorithms with respect to *undirected* tree-width as a parameter. Unfortunately, a similarly powerful *meta theorem* is not known for the directed relatives of tree-width such as directed tree-width, and in fact, such a meta-theorem is unlikely to exist as justified in [GHK⁺10].

Still it is natural to explore interesting algorithmic problems for directed graphs whose directed tree-width is bounded. Intuitively, one might guess that the dichromatic number problem 10.1 should be particularly suited for a parametrization by directed tree-width: Directed tree-width is small for digraphs which are structurally similar to acyclic digraphs, and for acyclic digraphs, computing the dichromatic number is utterly trivial.

However, quite surprisingly the first main result of this chapter refutes this intuition by showing that the k-DCP remains NP-complete, even if the input digraphs are of bounded directed tree-width. In fact, we can show a stronger result³: The problem is NP-complete for input digraphs admitting a feedback vertex set of bounded size. Given $d \in \mathbb{N}$, a digraph D is called *d-out-degenerate* if it can be reduced to the empty digraph by successively deleting vertices of out-degree at most d, and the *out-degeneracy* dgn⁺(D) is the smallest integer d such that D is d-out-degenerate.

Theorem 10.1. For every $k \ge 2$, DIGRAPH k-COLORING is NP-hard even if restricted to input digraphs D satisfying $\tau(D) \le k + 4$ and $dgn^+(D) \le k + 1$.

Theorem 10.1 is relatively tight with respect to $\tau(D)$ and $dgn^+(D)$: If $\tau(D) \leq k - 1$, we can find a feedback vertex set S of size at most k - 1 in time $f(k)n^{O(1)}$ [CLL⁺08], assign each vertex of S a different color in [k-1] and the remaining vertices the remaining color k. Further, one can easily find an acyclic $(dgn^+(D) + 1)$ -coloring of a digraph by successively reinserting vertices of out-degree at most $dgn^+(D)$ and greedily assigning the vertex a color which does not appear on its out-neighborhood. Hence, if $dgn^+(D) \leq k - 1$ or $\tau(D) \leq k - 1$, finding an acyclic k-coloring of D can be done in $f(k)n^{O(1)}$ time. In contrast, our hardness result excludes the existence of an $n^{f(k)}$ -time algorithm if we only

¹i.e., number of elementary operations required to execute the algorithm with input x, for an appropriate definition of *elementary*.

²so-called MSO_2 -definable problems

³It is a simple exercise to show that $dtw(D) \leq \tau(D)$ for every digraph D, as the addition of a vertex to a digraph increases the directed tree-width by at most one.

215

assume $\tau(D) \leq k+4$ and $dgn^+(D) \leq k+1$ instead, leaving only the cases $k \leq \tau(D) \leq k+3$ and $dgn^+(D) = k$ open. These cases were recently resolved in a paper by Harutyunyan, Lampis, and Melissinos [HLM20]. Let us mention their improvement of Theorem 10.1 here for the sake of completeness.

Theorem 10.2 (cf. [HLM20]). For every $k \ge 2$, DIGRAPH k-COLORING is NP-hard even if $\tau(D) \le k$, where D is the input digraph.

Note that $dgn^+(D) \leq \tau(D)$ for every digraph D, and hence Theorem 10.2 also closes the $\{k, k+1\}$ -gap for the degeneracy. Among other results, Harutyunyan et al. also showed that the k-DCP does admit an FPT-algorithm with respect to undirected tree-width, which reveals yet another interesting example where undirected and directed structural parameters show a significant difference in behavior.

Given a boolean formula in conjunctive normal form in which each clause contains at most k literals, k-SAT is the problem of deciding whether or not this formula admits a truthful variable assignment. Famously, k-SAT is NP-complete for every fixed value $k \geq 3$, and our proof of Theorem 10.1 actually works by reducing from the satisfiability problem. While $P \neq NP$ would imply the non-existence of polynomial-time algorithms for k-SAT, in fact, not even a *subexponential-time* algorithm, i.e., an algorithm with running time $2^{o(n)}n^{O(1)}$ is known⁴. Impagliazzo and Paturi [IP01] provided evidence that no such algorithm for k-SAT exists, and formulated the following famous hypothesis.

Hypothesis (Exponential Time Hypothesis (ETH) [IP01]). For each $k \ge 3$ there is some $s_k > 0$ such that no $2^{s_k n} n^{O(1)}$ -time algorithm for k-SAT exists.

Assuming the ETH, we can use our reduction in the proof of Theorem 10.1 to also show lower bounds on the running time of algorithms solving the k-DCP.

Theorem 10.3. For each $k \ge 2$ there is some $\epsilon > 0$ such that no $2^{\epsilon n} n^{f(\tau)}$ algorithm for DIGRAPH k-COLORING exists, where D is the input digraph, $\tau = \tau(D)$ and f is some function, unless the ETH is false.

In the second part of this chapter, we present a positive algorithmic result for DIGRAPH *k*-COLORING based on another parameter called *directed modular width*. This parameter was introduced in [SW19, SW20] inspired by module-decompositions of directed graphs as studied in [MdM05] as well as the corresponding notion of *modular width* for undirected graphs from [GLO13]. See also [MR84, Möh85] for results on modular decompositions of more general combinatorial structures.

Modules in graphs are sets of vertices which have the same relations to vertices outside the set. For digraphs, we have the following similar definition.

Definition 10.2. Let D be a digraph. A subset $\emptyset \neq M \subseteq V(D)$ of vertices is called a module, if all the vertices in M have the same sets of out-neighbours and the same sets of in-neighbours outside the module. Formally, we have $N_D^+(u_1) \setminus M = N_D^+(u_2) \setminus M$ and $N_D^-(u_1) \setminus M = N_D^-(u_2) \setminus M$ for all $u_1, u_2 \in M$. Consider Figure 10.1 for an illustration.

The following gives the precise recursive definition of directed modular width. For an illustration see Figure 10.2.

Definition 10.3 (Directed Modular Width). Let $k \in \mathbb{N}_0$, and let D be a digraph. We say that D has directed modular width at most k, if one of the following holds:

⁴Here n is the number of variables in the input formula.



Figure 10.1: A digraph D together with a module M.

- $v(D) \leq k$, or
- There exists a partition of V(D) into $\ell \in \{2, \ldots, k\}$ modules M_1, \ldots, M_ℓ such that for every $i, D[M_i]$ has directed modular width at most k.

The least $k \ge 1$ for which a digraph D has directed modular width at most k is defined to be the directed modular width, denoted by dmw(D), of D.



Figure 10.2: A digraph D together with a decomposition into modules (left) and the corresponding module-digraph (right) together with the modules M_i represented by vertices v_i .

Intuitively, the directed modular width is small on dense but structured networks, and is thus quite different from tree-width based parameters which are small for digraphs with a sparse directed cycle structure. In [SW19, SW20] it was shown that many NP-hard problems which are intractable by means of most other known structural digraph parameters, admit FPT-algorithms with respect to the parameter *dmw*.

Here we present only the corresponding result for digraph coloring, which opposes the negative result from Theorem 10.1.

Theorem 10.4. The dichromatic number of a given digraph D can be computed in time

$$O(n^3 + f(\omega)n\log^2 n),$$

where n := v(D), $\omega := dmw(D)$ and $f(\omega) = 2^{O(\omega 2^{\omega})}$.

Again, there has been follow-up work inspired by Theorem 10.4. Recently, Gurski, Komander, and Rehs [GKR20] studied the parameter *directed clique-width* which is a lower bound for directed modular width. They extended the positive result from Theorem 10.4 to this parameter, however, with a worse running time bound, as they only constructed an XP-algorithm. In fact, they argued that no FPT-algorithm for the k-DCP with respect to directed clique-width should exist under reasonable complexity theoretic assumptions.

The rest of this chapter is organized as follows. In Section 10.2 we give the proofs of Theorem 10.1 and Theorem 10.3. In Section 10.3 we then explain the necessary background on directed modular width and give the proof of Theorem 10.4.
10.2 Parametrization by Feedback Vertex Set

We start this section by proving Theorem 10.1, which will be a simple consequence of the following Lemma. It establishes the case k = 2 of Theorem 10.1 by reducing from the satisfiability problem.

Lemma 10.5. DIGRAPH 2-COLORING is NP-hard even if $\tau(D) \leq 6$ and $dgn^+(D) \leq 3$, where D is the input digraph.

Proof. We provide a reduction from SAT to DIGRAPH 2-COLORING. Let C_1, C_2, \ldots, C_m denote the clauses and X_1, X_2, \ldots, X_n the variables in the SAT instance. We construct a digraph D which is 2-colorable if and only if there is a satisfying assignment for the SAT instance. For each clause C_i we add the vertex c_i to D, and for each literal $L_j \in C_i$ we add the vertex $l_{j,i}$. That is, we add the vertex $x_{j,i}$ if $X_j \in C_i$ and the vertex $\overline{x}_{j,i}$ if $\overline{X}_j \in C_i$. To simplify our notation, we assume that a literal L_j is associated with the variable X_j , that is $L_j = X_j$ or $L_j = \overline{X}_j$, and that l_j corresponds to the lower-case variant of L_j , that is $l_j = x_j$ if $L_j = X_j$ and $l_j = \overline{x}_j$ if $L_j = \overline{X}_j$. We want the color of a vertex $x_{j,i}$ to correspond to an assignment of the variable X_i . To this end, we add a set $S = \{t_1, t_2, t_3, f_1, f_2, f_3\}$ of vertices which will correspond to a feedback vertex set in D. Furthermore, for each literal L_j we add a vertex l_j . We now add directed cycles to D in such a way that any acyclic coloring $c: V(D) \to \{0, 1\}$ of D must have the following properties.

(i) $c(l_{j,i}) = c(l_{j,h})$ for all $j \in [n]$ and $i, h \in [m]$, and

(ii)
$$c(\overline{x}_{j,h}) \neq c(x_{j,i})$$
 for all $j \in [n]$ and $i, h \in [m]$.

These properties will allow us to obtain a variable assignment from any acyclic 2-coloring.

To ensure (i), we construct a *literal gadget* (illustrated in Figure 10.3(a)). First, we add the digon t_1, f_1 . Then, for each literal L_j and each clause C_i with $L_j \in C_i$ we add the directed triangles $l_j, l_{j,i}, t_1$ and $l_j, l_{j,i}, f_1$. In any acyclic 2-coloring of the literal gadget, if there are $i, h \in [n]$ such that $l_{j,i}$ and $l_{j,h}$ have different colors, one of them, say, $l_{j,i}$, must have the same color as l_j . Since t_1, f_1 are connected by a digon, they must have different colors in any acyclic 2-coloring. Hence, the cycle $l_j, l_{j,i}, t_1$ or the cycle $l_j, l_{j,i}, f_1$ is monochromatic if $l_{i,j}$ and $l_{i,h}$ have different colors. This proves (i).

For (ii), we construct a variable gadget (illustrated in Figure 10.3(b)). First, we add the digon t_2, f_2 . Then, we add the directed triangles x_j, \overline{x}_j, t_2 and x_j, \overline{x}_j, f_2 for each $j \in [n]$ where both X_j and \overline{X}_j appear in the formula. In any acyclic 2-coloring of the gadget, if x_j and \overline{x}_j receive the same color, then one of the added cycles is monochromatic as t_2 and f_2 must receive different colors. Because of the literal gadgets, we know that l_j and $l_{j,i}$ have different colors for all $j \in [n]$ and $i \in [m]$. As x_j and \overline{x}_j have different colors, it follows from (i) that $x_{j,i}$ and $\overline{x}_{j,h}$ have different colors for all $j \in [n]$ and all $h, i \in [m]$. This implies (ii).

We now construct a *clause gadget* (illustrated in Figure 10.3(c)) that ensures that each clause is satisfied by at least one of its literals. We first add the digon t_3, f_3 . Then, for each clause C_i we add the digon c_i, t_3 . Finally, we add the directed cycle $c_i, l_{j_1,i}, l_{j_2,i}, \ldots, l_{j_h,i}, f_3$, where $l_{j_1,i}, l_{j_2,i}, \ldots, l_{j_h,i}$ are the literals of C_i . We sort the literals in such a way that $j_1 < j_2 < \cdots < j_h$ and such that positive literals are ordered before negative literals.

This concludes the construction of the digraph D.

We first show that $\tau(D) \leq 6$. We claim that the set $S = \{t_1, t_2, t_3, f_1, f_2, f_3\}$ is a feedback vertex set of D. We prove that D - S is acyclic by providing a topological ordering of its vertices. We first take the positive literal vertices x_j and the clause vertices c_i into the ordering, as these are sources in D - S. Removing these vertices, all negative



Figure 10.3: Variable, literal and clause gadgets of the proof of Lemma 10.5 for the variable X_1 and the clause $(X_1 \vee \overline{X}_2)$ in the SAT formula $(X_1 \vee X_2) \wedge (X_1 \vee \overline{X}_2) \wedge (\overline{X}_1 \vee X_2)$.

literal vertices \overline{x}_j become sources, which we then add to the end of the current topological ordering. The only remaining vertices are the variable vertices $l_{j,i}$. It follows from the construction of the clause gadget that ordering the $l_{j,i}$ monotonically in j, with positive literals preceding corresponding negative literals, yields a topological ordering of D-S.

To show that the degeneracy of D is 3, we construct a linear ordering of the vertices as follows. The first vertices of the ordering are t_1, f_1, t_2, f_2, t_3 and f_3 . These have at most one outgoing arc to vertices which are smaller. Afterwards come all positive literal vertices x_j , then all negative literal vertices \overline{x}_j , followed by the variable vertices $l_{j,i}$. The vertices \overline{x}_j have arcs to t_2 and f_2 , and x_j has no arc to smaller vertices. Hence, they have at most two arcs to smaller vertices. The vertices $l_{j,i}$ have arcs to t_1 , f_1 and potentially to some other $l_{h,i}$ or to f_3 , but never both. Hence, they have at most 3 arcs to smaller vertices. The last vertices in the ordering are the clause vertices c_i . These have an arc to t_3 and another to some $l_{j,i}$. Hence, the out-degeneracy of D is at most 3.

We now prove that D is acyclically 2-colorable if there is a truth assignment of the variables satisfying all clauses.

Let $\beta : \{X_j \mid j \in [n]\} \to \{0, 1\}$ be a satisfying truth assignment of the variables. We construct a coloring $c \colon V(D) \to \{0, 1\}$ as follows.

- 1. $c(f_i) := 0$ and $c(t_i) := 1$ for $i \in [3]$.
- 2. $c(c_i) := 0$ for $i \in [m]$.
- 3. $c(x_{j,i}) := \beta(X_j)$ for all $j \in [n]$ and $i \in [m]$ with $X_j \in C_i$.
- 4. $c(\overline{x}_{j,i}) := 1 \beta(X_j)$ for all $j \in [n]$ and $i \in [m]$ with $\overline{X}_j \in C_i$.
- 5. $c(x_i) := 1 \beta(X_i)$ and $c(\overline{x}_i) := \beta(X_i)$ for all $j \in [n]$.

This concludes the construction of c.

We now argue that each color class induces an acyclic digraph in D.

Let $d \in \{0, 1\}$ be some color. Since S is a feedback vertex set of D, it suffices to show that there are no directed cycles using vertices of $S_d := c^{-1}(d) \cap S$ in $D[c^{-1}(d)]$.

Assume, without loss of generality, that $t_1, t_2 \in c^{-1}(d)$, i.e., d = 1. The case $f_1, f_2 \in c^{-1}(d)$, i.e., d = 0, follows analogously. We prove that no cycle contains t_1 or t_2 by progressively identifying and removing sinks from $D[c^{-1}(d)]$. As for all $j \in [n]$ and $i \in [m]$ we have $c(x_j) \neq c(\overline{x}_j) = c(x_{j,i})$, it follows that all x_j are sinks in $D[c^{-1}(d)]$. Removing all x_j , we can see that t_2 is now a sink. Hence, no directed cycle in $D[c^{-1}(d)]$ contains t_2 . As $c(\overline{x}_j) \neq c(\overline{x}_{j,i})$, it follows that \overline{x}_j is now a sink and we can remove it. Without literal vertices, t_1 becomes a sink, implying no cycle goes through t_1 in $D[c^{-1}(d)]$, as desired. Consequently, for any $d \in \{0, 1\}$, no directed cycle in $D[c^{-1}(d)]$ can possibly use one of the vertices t_1, t_2, f_1, f_2 and therefore must either contain t_3 or f_3 .

If $t_3 \in c^{-1}(d)$, then $c_i \notin c^{-1}(d)$ for all $i \in [m]$, as $c(t_3) = 1$ and $c(c_i) = 0$. Hence, t_3 has no neighbours in $D[c^{-1}(d)]$ and cannot be in any directed cycle. If $f_3 \in c^{-1}(d)$, assume towards a contradiction that there is a directed cycle C in $D[c^{-1}(d)]$ containing f_3 . Note that this cycle must also contain c_i for some $i \in [m]$, as these are the only out-neighbours of f_3 in $D[c^{-1}(d)]$. Furthermore, the out-neighbour of c_i in C is some $l_{j,i}$, and the only out-neighbours of $l_{j,i}$ are t_1 and potentially some $l_{h,i}$ or f_3 , as these were the arcs added in the clause gadgets. The vertices $l_{j,i}$ in C correspond to the literals in c_i . In order to form a monochromatic directed cycle, all literals in c_i must be in C. However, this means that $c(x_{j,i}) = 0$ for all X_j in clause C_i and $c(\overline{x}_{j,i}) = 0$ for all \overline{X}_j in clause C_i . By construction of c, this implies that all literals in C_i are set to false, which means that the clause is not satisfied, a contradiction to our initial assumption. Hence, the digraph $D[c^{-1}(d)]$ is acyclic, and D is acyclically 2-colorable.

We now show that the formula is satisfiable if $\vec{\chi}(D) \leq 2$ by constructing a satisfying variable assignment β from an acyclic 2-coloring of D. Let $c: V(D) \to \{0, 1\}$ be an acyclic coloring of D. Without loss of generality, we assume that $c(t_3) = 1$, which implies $c(f_3) = 0$ and $c(c_i) = 0$ for $i \in [m]$. We set $\beta(X_j)$ to true if $c(x_j) = 0$ and to false if $c(x_j) = 1$.

Assume towards a contradiction that there is some clause C_i which is not satisfied by β . By simply renaming the variables, we can assume without loss of generality that the literals of C_i are L_1, L_2, \ldots, L_a . As C_i is not satisfied, it follows that all L_j evaluate to false with β . By construction of the literal gadget, $c(l_j) \neq c(l_{j,i})$ for all $i \in [m]$ with $L_j \in C_i$. From (i) and (ii), for all $j \in [n]$ it follows that $c(l_{j,i}) = 1$ if the literal L_j is true, and that $c(l_{j,i}) = 0$ if the literal L_j is false. As C_i is not satisfied, $c(c_i) = c(f_3) = c(l_{j,i}) = 0$ for all $j \in [a]$. Hence, the directed cycle $C = c_i, l_1, l_2, \ldots, l_a, f_3$ is monochromatic, contradicting our assumption that c is an acyclic coloring. This implies that β is a satisfying variable assignment, concluding our proof.

With a simple self-reduction, we can extend the previous result to all $k \ge 2$.

Proof of Theorem 10.1. We prove the statement by induction on k. The case k = 2 follows from Lemma 10.5. We provide a reduction from DIGRAPH (k-1)-COLORING to DIGRAPH k-COLORING such that $\tau(D') \leq \tau(D) + 1$ and $dgn^+(D') \leq dgn^+(D) + 1$, where D is the input instance and D' is the reduced instance. We obtain D' by adding a new vertex x to D, together with the edges $\{(x, v), (v, x) \mid v \in V(D)\}$. If D is (k - 1)-colorable, then setting the color of x to k gives an acyclic k-coloring for D'. If D' is k-colorable, then no vertex in D has the same color as x. Hence, D is (k - 1)-colorable. Furthermore, all new cycles created by adding x go through x. If D - S is acyclic for some vertex set S, then $D' - (S \cup \{x\}) = D - S$ is also acyclic. Hence, $\tau(D') \leq \tau(D) + 1$. To show that the degeneracy of D' increased by at most one, we consider some ordering of V(D) certifying that $dgn^+(D) \leq k$. By placing v as the smallest vertex with respect to the ordering, we increase the outdegree of the vertices in D by one. Hence, the degeneracy of D' is at most $dgn^+(D) + 1 = k + 1$, as desired.

As an immediate consequence of the above theorem arises the following corollary.

Corollary 10.6. There is no $n^{f(k,\tau)}$ -time algorithm deciding DIGRAPH k-COLORING where $\tau = \tau(D)$ and f is some function, unless P=NP.

A finer analysis of the reduction provided in Lemma 10.5 gives us the stronger hardness result 10.3 under the assumption of the Exponential Time Hypothesis.

Note that the ETH only considers the running time with respect to the number of variables in the input formula, not the number of clauses. In several reductions, however,

it is difficult to ensure that the size of the reduced instance depends only on the number of variables. For example, the reduction in Lemma 10.5 contains one vertex for each clause. This would prevent us from directly applying the ETH. Fortunately, [IPZ01] showed that it is possible to assume that $m \in O(n)$, where m is the number of clauses, by proving the following lemma.

Lemma 10.7 (Sparsification Lemma, Impagliazzo, Paturi, and Zane [IPZ01]). For all $\epsilon > 0$ and k > 0 there is a constant C so that any k-SAT formula Φ with n variables can be expressed as $\Phi = \bigvee_{i=1}^{t} \Psi_i$, where $t \leq 2^{\epsilon n}$ and each Ψ_i is a k-SAT formula with at most Cn clauses such that each variable appears in constantly many clauses. Moreover, this disjunction can be computed by an algorithm running in time $2^{\epsilon n} n^{O(1)}$.

By first applying the sparsification lemma to the input formula and then the reduction from Theorem 10.1, we can show the following.

Proof of Theorem 10.3. First note that the reduction from DIGRAPH (k-1)-COLORING to DIGRAPH k-COLORING from Theorem 10.1 increases the input instance by one vertex. Hence, it suffices to show the statement for k = 2, as the remaining cases follow by induction. We first use the sparsification lemma to obtain at most $2^{\epsilon n}$ many 3-SAT instances where each variable appears in constantly many clauses. Applying the reduction from Lemma 10.5 to each instance, we obtain at most $2^{\epsilon n}$ many digraphs where for each variable we have constantly many vertices and for each clause we have one vertex. This means that the number of vertices on the reduced instances is linear in the number of variables of the formula. Hence, a subexponential-time algorithm for DIGRAPH 2-COLORING implies a subexponential-time algorithm for 3-SAT, which would contradict the ETH.

Note that an algorithm solving the k-DCP with running time $O(k^n \cdot (n+m))$ (n the number of vertices, m the number of arcs) is simple: We can test all k^n colorings of the vertices of D, and then check if each color class is an acyclic digraph in linear time O(m+n) by computing a topological ordering.

10.3 Parametrization by Modular Width

In this section, we present a polynomial time-algorithm to compute an optimal acyclic coloring of a given directed graph, provided that its *directed modular width* is bounded by a constant. In the following we prepare the proof of Theorem 10.4 by introducing some notation and mentioning simple but important properties of directed modular width.

While the directed modular width can increase when taking subdigraphs, it is monotone with respect to taking induced subdigraphs. More precisely, we have the following.

Fact 10.1. Let D be a digraph, and let D' be an induced subdigraph of D. Then

$$dmw(D') \le dmw(D).$$

Proof. We prove the statement by induction on the number of vertices of D. The statement clearly holds true when v(D) = 1, so suppose for the inductive step that $v(D) = n \ge 2$, and that the statement is true for all digraphs on less than n vertices (and for all their induced subdigraphs).

Let D' = D[X] with $X \subseteq V(D)$ be an induced subdigraph of D, and let $\omega := \operatorname{dmw}(D)$, $\omega' := \operatorname{dmw}(D')$. If $v(D) \leq \omega$, then clearly, we also have $\omega' \leq v(D') \leq v(D) \leq \omega$, proving the claim. Otherwise, let M_1, \ldots, M_ℓ denote a partition of V(D) into modules such that $2 \leq \ell \leq \omega$ and such that $\dim(D[M_i]) \leq \omega$ for all $i \in [\ell]$. For every $i \in [\ell]$, $D'[X \cap M_i] = D[X \cap M_i]$ is an induced subdigraph of $D[M_i]$, and since $v(D[M_i]) < v(D)$, the induction hypothesis tells us that $\dim(D'[X \cap M_i]) \leq \dim(D[M_i]) \leq \omega$ for all $i \in [\ell]$. Clearly, $X \cap M_i$ defines a module in D' for each $i \in [\ell]$. Because $\ell \leq \omega$, the definition of directed modular width now implies

$$\omega' = \operatorname{dmw}(D') \le \max\left\{\ell, \operatorname{dmw}(D'[X \cap M_1]), \dots, \operatorname{dmw}(D'[X \cap M_\ell])\right\} \le \omega$$

which yields the claim also in this case.

Given a digraph D and a partition M_1, \ldots, M_ℓ of V(D) into modules, we will use D_M to denote the *module-digraph* of D corresponding to the module-decomposition $\{M_1, \ldots, M_\ell\}$: D_M is obtained from D by identifying $M_i, i \in [\ell]$ each into a single vertex $v_i \in V(D_M)$ and deleting parallel directed edges afterwards. Equivalently, an edge (v_i, v_j) lies in $E(D_M)$ if and only if in D, there is at least one directed edge starting in M_i and ending in M_j . Due to the modular property, this is equivalent to the fact that $(u, w) \in A(D)$ for all $u \in M_i, w \in M_j$. For an example of a module-digraph see Figure 10.2.

Throughout this section, given a module-decomposition M_1, \ldots, M_ℓ of a directed graph D, we will denote by $\eta : V(D) \to V(D_M)$ the mapping defined by $\eta(z) := v_k$ for all $z \in M_k$.

The most important tool involved in the algorithm of Theorem 10.4 is a subroutine to find a non-trivial decomposition of the vertex set of a given digraph into modules, in polynomial time. In fact, this task can be executed in a much stronger form. In [MdM05], it was shown that a so-called *canonical module-decomposition* of a given digraph can be obtained in linear time. For us, the following weaker form of their result will be sufficient.

Theorem 10.8 ([MdM05]). There is an algorithm that, given a digraph D on at least two vertices as input, returns a decomposition of V(D) into $\ell \in \{2, ..., dmw(D)\}$ modules. This algorithm runs in time O(n + m), where n := v(D) and m := a(D).

To do the runtime-analysis of our algorithm, we will use a rooted model-tree T which resembles the structure of recursive calls. Every vertex $q \in V(T)$ has either no children or at least two. It furthermore admits a labelling of its vertices of the following kind:

The root of the tree is labelled with a finite ground set Ω (in our case the vertex set of the considered digraph). Every other vertex $q \in V(T)$ is labelled with a subset $\emptyset \neq \Omega(q) \subseteq \Omega$, and for every branching vertex, the associated subset is the disjoint union of the subsets associated to its children. Finally, the leafs of the tree are labelled with the singletons $\{v\}, v \in \Omega$. A tree which admits a labelling of this type will be called a *decomposition tree*.

Fact 10.2. If T is a decomposition tree with ground-set Ω , then $v(T) \leq 2|\Omega| - 1$.

Proof. First note that we can reduce to the case where T is a rooted binary tree: If there is a branching vertex $q \in V(T)$ with $b \geq 3$ children q_1, \ldots, q_b , we can locally replace this branching by a binary tree with b leafs, where instead of directly splitting $\Omega(q)$ into $\Omega(q_1), \ldots, \Omega(q_b)$, we first split-off $\Omega(q_1)$, then $\Omega(q_2)$, and so on. Clearly, successive application of this operation to every branching with more than two children yields a binary decomposition-tree T' with ground set Ω and $v(T') \geq v(T)$.

Now if T is a binary-tree, because the leafs of T are labelled by the singletons of Ω , T has $|\Omega|$ leafs and therefore $2|\Omega| - 1$ vertices.

In our algorithm, we reduce the computation of the dichromatic number to the same problem on an input digraph with a bounded number of vertices, but equipped with additional information (such as weightings) of polynomial-size in the original input. We will then make use of an integer program reformulation of the problem, in which we have a bounded number of constraints and variables, but possibly entries in the input matrices and vectors of polynomial size. To solve this ILP efficiently we make use of the following powerful result from the theory of Integer Programming, which shows that the feasibility of a given ILP can be decided in polynomial time for a fixed number of variables.

Theorem 10.9 (cf. [FLM⁺08], Theorem 1). There exists an algorithm that, given as input a matrix $A \in \mathbb{Z}^{n \times p}$ and a vector $b \in \mathbb{Z}^n$, decides whether there is a feasible solution to

$$Ax \ge b, x \in \mathbb{Z}^p$$

(and returns a solution if applicable) in time $O(p^{2.5p+o(p)}L)$ where L denotes the coding length of the input (A, b).

Solving an ILP can be easily reduced to checking the feasibility of several ILP-s using binary search.

Corollary 10.10 (cf. [FLM⁺08], Theorem 12). There exists an algorithm that, given as input a matrix $A \in \mathbb{Z}^{n \times p}$, vectors $c \in \mathbb{Z}^p, b \in \mathbb{Z}^n$, and some $U_1, U_2 \in \mathbb{Z}_+$, tests feasibility and if applicable outputs an optimal solution of the ILP

$$\min c^T x \tag{10.1}$$

subj. to
$$Ax \ge b, x \in \mathbb{Z}^p$$
 (10.2)

in time $O(p^{2.5p+o(p)}L\log(U_1U_2))$ where L denotes the coding length of the input (A, b, c). Here we assume that the optimal value of the program lies within $[-U_1, U_1]$ and that U_2 upper-bounds the largest absolute value any entry in an optimal solution vector can take.

Our algorithm for computing the dichromatic number parametrized by modular width will work recursively. In order to enable the kind of recursion we aim for, it will be useful to solve a more general problem, in which we want to find an acyclic coloring of a given digraph where vertices have to be assigned *lists* of colors⁵.

Definition 10.4. Let D be a digraph equipped with an assignment $\mathcal{N} : V(D) \to \mathbb{N}$ of positive integers to the vertices. An \mathcal{N} -coloring with $k \in \mathbb{N}$ colors of D is an assignment of lists $c(v) \subseteq [k]$ of colors to every vertex $v \in V(D)$ such that $|c(v)| = \mathcal{N}(v)$ for all $v \in V(D)$, and moreover for every $i \in [k]$ there is no directed cycle C in D such that $i \in c(v)$ for every $v \in V(C)$.

For a fixed number-assignment \mathcal{N} , we define the \mathcal{N} -dichromatic number $\vec{\chi}_{\mathcal{N}}(D)$ of a digraph D to be the minimum k such that an \mathcal{N} -coloring with k colors of D exists.

As an additional input for our generalized coloring problem, we also have a threshold $\tau \in \mathbb{N}$, which bounds the total size of the color lists which have to be assigned. For bounded directed modular width, the algorithm runs in polynomial time in τ and n := v(D).

Problem 10.2 (Weighted Digraph Coloring (wDCP)). Given as input digraph D, a natural number $\tau \in \mathbb{N}$, and an assignment $\mathcal{N} : V(D) \to \mathbb{N}$ such that $\sum_{v \in V(D)} \mathcal{N}(v) \leq \tau$, determine the value of $\vec{\chi}_{\mathcal{N}}(D)$ and an optimal \mathcal{N} -coloring of D.

⁵Note however that this generalization of the dichromatic number is fundamentally different from the list dichromatic number, in which we need to pick only one color per vertex from a given list of colors.

In the remainder of this section, let us prove the following result. Theorem 10.4 will then follow directly from this result by putting $\mathcal{N}(v) := 1$ for all $v \in V(D)$ and $\tau := n$.

Theorem 10.11. There is an algorithm that, given a digraph D on n vertices and an assignment $\mathcal{N} : V(D) \to \mathbb{N}$ of numbers to the vertices such that $\sum_{v \in V(D)} \mathcal{N}(v) \leq \tau$, outputs the value of $\vec{\chi}_{\mathcal{N}}(D)$ together with a certifying assignment of color lists to the vertices. The running time of the algorithm is $O(n^3 + f(\omega)n\log^2 \tau + n\tau)$, where n := v(D), $\omega := dmw(D)$, and $f(\omega) = 2^{O(\omega 2^{\omega})}$.

We prepare the proof of Theorem 10.11 with some auxiliary statements.

Lemma 10.12. Let D be a digraph equipped with a module-decomposition $\{M_1, \ldots, M_\ell\}$ of the vertex set. If C is an induced directed cycle in D, then either there is some $i \in [\ell]$ such that C is contained in $D[M_i]$, or D uses at most one vertex from each module.

Proof. Assume towards a contradiction that there was a directed cycle C, such that for some $i \in [\ell]$ we have $|V(C) \cap M_i| \geq 2$, and $V(C) \setminus M_i \neq \emptyset$. Let $x \in V(C) \setminus M_i$ be some vertex, and let y_1 be the closest vertex after x in the cyclic directed order along Cwhich is contained in M_i , and let y_2 be the closest vertex before x contained in M_i in the cyclic order. Because of $|V(C) \cap M_i| \geq 2$, we know that $y_1 \neq y_2$. Let $x_1 \in V(C) \setminus M_i$ be the predecessor of y_1 on C, and let $x_2 \in V(C) \setminus M_i$ be the successor of y_2 on C. This means that $(x_1, y_1), (y_2, x_2) \in E(C)$. By the modular property, this implies that also $(x_1, y_2), (y_1, x_2) \in A(D)$. Because C was assumed to be induced, this implies that $A(C) \supseteq \{(x_1, y_1), (y_1, x_2), (y_2, x_2), (x_1, y_2)\}$, contradicting that C is a directed cycle. \Box

Lemma 10.13. Let D be a digraph, let M_1, \ldots, M_ℓ be a partition of V(D) into modules, and let D_M denote the corresponding module-digraph.

Let $\mathcal{N}: V(D) \to \mathbb{N}$ be an assignment of numbers to the vertices.

Denote by $v_i \in V(D_M)$ for every $i \in [\ell]$ the vertex of the module-digraph representing M_i , and define an assignment $\mathcal{N}_M : V(D_M) \to \mathbb{N}$ according to $\mathcal{N}_M(v_i) := \vec{\chi}_{\mathcal{N}|_{M_i}}(D[M_i])$ for each $i \in [\ell]$. Then we have

$$\vec{\chi}_{\mathcal{N}}(D) = \vec{\chi}_{\mathcal{N}_M}(D_M).$$

Moreover, given an optimal \mathcal{N}_M -coloring of D_M (i.e., with a minimal total number of colors), we can construct an optimal \mathcal{N} -coloring of D in time $O(\ell v(D))$.

Proof. Let $k := \vec{\chi}_{\mathcal{N}}(D), k_M := \vec{\chi}_{\mathcal{N}_M}(D_M)$. We prove $k \leq k_M$ and $k_M \leq k$ separately.

To prove the first inequality, consider an assignment $c_M : V(D_M) \to 2^{[k_M]}$ of color lists that define an optimal \mathcal{N}_M -coloring of D_M . We therefore have $|c_M(v_i)| = \vec{\chi}_{\mathcal{N}|M_i}(D[M_i])$ for all $i \in [\ell]$. Hence, for every $i \in [\ell]$, the digraph $D[M_i]$ admits an $\mathcal{N}|_{M_i}$ -coloring c_i using $|c_M(v_i)|$ colors in total. By relabeling, we may assume that the set of colors used by c_i is exactly $c_M(v_i)$. Consider now $c : V(D) \to 2^{[k_M]}$ defined by $c(v) := c_i(v)$ whenever $v \in M_i$. This assignment has the property that $|c(v)| = \mathcal{N}|_{M_i}(v) = \mathcal{N}(v)$ for all $v \in M_i \subseteq V(D), i \in [\ell]$. We claim that it defines a valid \mathcal{N} -coloring of D. Assume towards a contradiction that there exists a directed cycle C in D such that $\bigcap_{v \in V(C)} c(v) \neq \emptyset$. W.l.o.g. we may assume that C is chosen with V(C) inclusion-wise minimal, i.e., C is induced (otherwise we could find a directed cycle using a proper subset of the vertices). By Lemma 10.12 C either is contained in $D[M_i]$ for some $i \in [\ell]$, or it uses at most one vertex from each module. In the first case, we obtain that C is a directed cycle in $D[M_i]$ with $\bigcap_{v \in V(C)} c_i(v) \neq \emptyset$, which contradicts the choice of c_i as a proper $\mathcal{N}|_{M_i}$ -coloring of $D[M_i]$. In the second case, C in a natural way yields a directed cycle C_M in the moduledigraph D_M such that for every vertex $v_i \in V(C_M)$, there is a unique corresponding vertex $w_i \in V(C) \cap M_i$ from the module. By the choice of the colorings c_i , we find that $c(w_i) = c_i(w_i) \subseteq \bigcup_{z \in M_i} c_i(z) = c_M(v_i)$ for all $i \in [\ell]$, and therefore we have

$$\bigcap_{v_i \in V(C_M)} c_M(v_i) \supseteq \bigcap_{w_i \in V(C)} c(w_i) \neq \emptyset,$$

contradicting the choice of c_M as a valid \mathcal{N}_M -coloring of D_M . As all cases led to a contradiction, we conclude that indeed c is a proper \mathcal{N} -coloring of D whose color sets are contained in $[k_M]$, and therefore $k \leq k_M$ as claimed.

To prove the second inequality, consider an optimal \mathcal{N} -coloring $c: V(D) \to 2^{[k]}$ of D. Clearly, for any i, the restriction $c|_{M_i}$ defines a valid $\mathcal{N}|_{M_i}$ -coloring of $D[M_i]$, and therefore has to fulfil

$$\left| \bigcup_{z \in M_i} c(z) \right| \ge \vec{\chi}_{\mathcal{N}|_{M_i}}(D[M_i]) = \mathcal{N}_M(v_i).$$

Now for each $i \in [\ell]$, choose some subset $L_i \subseteq \bigcup_{z \in M_i} c(z)$ with $|L_i| = \mathcal{N}_M(v_i)$ and define an assignment $c_M : V(D_M) \to 2^{[k]}$ of color lists to the vertices of D_M according to $c_M(v_i) := L_i$. We claim that this defines a proper \mathcal{N}_M -coloring of D_M . Assume towards a contradiction that there exists a directed cycle C in D_M and a color $\tilde{c} \in [k]$ such that $\tilde{c} \in L_i$ for every $v_i \in V(C)$. By the definition of L_i , this implies that for every i such that $v_i \in V(C)$, we can find a vertex $w_i \in M_i$ such that $\tilde{c} \in c(w_i)$. By the properties of the modules, we now immediately conclude that $\{w_i \mid i \in [\ell], v_i \in V(C)\}$ forms the vertex set of a directed cycle in D such that \tilde{c} is contained in all color sets of its vertices. This contradicts the fact that c was chosen as a valid \mathcal{N} -coloring of D. Finally, since c_M uses at most k colors in total, this shows $k_M \leq k$ and concludes the proof of the claimed equality.

Concerning the (algorithmic) construction of an optimal \mathcal{N} -coloring of D given an optimal \mathcal{N}_M -coloring of D_M (using k_M colors in total), we can compute the \mathcal{N} -coloring of D as defined in the proof of $k \leq k_M$. This requires at most O(v(D)) operations for each module, and so at most $O(\ell v(D))$ in total.

Lemma 10.14. Given a digraph D on at most ω vertices, an assignment $\mathcal{N} : V(D) \to \mathbb{N}$ and some $\tau \in \mathbb{N}$ such that $\sum_{v \in V(D)} \mathcal{N}(v) \leq \tau$, we can compute $\vec{\chi}_{\mathcal{N}}(D)$ and a corresponding optimal \mathcal{N} -coloring in time $O(f(\omega) \log^2 \tau)$, where $f(\omega) = 2^{O(\omega 2^{\omega})}$.

Proof. We reformulate the problem of determining the \mathcal{N} -dichromatic number as a linear integer program to enable an application of 10.10. For this purpose, note that we can alternatively represent a color-list assignment $c: V(D) \to 2^{[k]}$ by the collection of "color classes" $A_i := \{v \in V(D) | i \in c(v)\}$ for all colors $i \in [k]$, each of which (by the definition of an \mathcal{N} -coloring) induces an acyclic subdigraph of D.

We can therefore associate an \mathcal{N} -coloring with a set of variables $x_A, A \in \mathcal{A}(D)$, where $\mathcal{A}(D)$ is the collection of acyclic vertex sets in D, and x_A counts the number of $i \in [k]$ such that $A = A_i$. The condition that the assigned color lists are of the sizes required by \mathcal{N} can be formulated as a linear equation for each vertex. This shows that we can compute $\vec{\chi}_{\mathcal{N}}(D)$ as the optimal value of the following ILP:

$$\min \sum_{A \in \mathcal{A}(D)} x_A$$
(10.3)
subj. to
$$\sum_{A \ni v} x_A = \mathcal{N}(v) \text{ for all } v \in V(D)$$
$$x_A \ge 0 \text{ for all } x_A \in \mathbb{Z}$$

This ILP in canonical form has $p = O(|\mathcal{A}(D)|) = O(2^{\omega})$ variables. The coding length L of the matrix and the vectors describing this ILP is clearly bounded by $(\omega + 1)2^{\omega} \log \tau$. Setting up the ILP requires enumerating the subsets $A \subseteq V(D)$ for which D[A] is acyclic. As we can test whether D[A] is acyclic in time $O(a(D)) \leq O(\omega^2)$, the ILP can be set up in time $O(\omega^2 2^{\omega} + \omega 2^{\omega} \log \tau) \leq O(\omega^2 2^{\omega} \log \tau)$.

It is readily verified that the optimal value $\vec{\chi}_{\mathcal{N}}(D)$ of the program is bounded from above by $U_1 := \tau$, more generally, it holds that $\vec{\chi}_{\mathcal{N}}(D) \leq \sum_{v \in V(D)} \mathcal{N}(v)$ (assign disjoint color sets to the different vertices). In an optimal solution to program 10.3, we certainly have $x_A \leq \tau$ for all $A \in \mathcal{A}(D)$. Therefore we can put $U_2 := \tau$. Application of 10.10 now yields that there is an algorithm for determining $\vec{\chi}_{\mathcal{N}}(D)$ in time $O(f(\omega) \log \tau \log \tau^2) =$ $O(f(\omega) \log^2 \tau)$ for some function f. In fact, we may take $f(\omega) = p^{O(p)} + \omega^2 2^{\omega} \leq 2^{O(\omega 2^{\omega})}$. This proves the claim.

Proof of Theorem 10.11. We follow a recursive approach which makes use of Theorem 10.8, Lemma 10.13 and Lemma 10.14.

Assume we are given a digraph D, the assignment $\mathcal{N} : V(D) \to \mathbb{N}$ and a natural number $\tau \in \mathbb{N}$ such that $\sum_{v \in V(D)} \mathcal{N}(v) \leq \tau$ as input. Let $\omega := \operatorname{dmw}(D)$ be the directed modular width of D.

If v(D) = 1, say $V(D) = \{v\}$, we return $\vec{\chi}_{\mathcal{N}}(D) = \mathcal{N}(v)$ and a color-list of $\mathcal{N}(v)$ different colors.

If $v(D) \geq 2$, we first apply the algorithm from [MdM05] to D, in order to obtain a partition of V(D) into modules M_1, \ldots, M_ℓ , where $2 \leq \ell \leq \omega$. We compute the digraphs $D[M_1], \ldots, D[M_\ell]$ as well as the module-digraph D_M .

We now recursively apply the algorithm to $(D[M_1], \mathcal{N}|_{M_1}, \tau), \ldots, (D[M_\ell], \mathcal{N}|_{M_\ell}, \tau)$. Each of the digraphs $D[M_i]$ (according to Fact 10.1) has directed modular width at most ω and less vertices than D.

Now given the outputs $\vec{\chi}_{\mathcal{N}|M_i}(D)$, $i = 1, \ldots, \ell$ of these recursive calls (and the corresponding optimal color-list assignments), as in Lemma 10.13, we define $\mathcal{N}_M : V(D_M) \to \mathbb{N}$ according to $\mathcal{N}_M(v_i) := \vec{\chi}_{\mathcal{N}|M_i}(D)$ for all $i \in [\ell]$. Because of $v(D_M) = \ell \leq \omega$ we can now apply the algorithm from Lemma 10.14 to the instance $(D_M, \mathcal{N}_M, \tau)$ in order to obtain the value of $\vec{\chi}_{\mathcal{N}_M}(D_M)$ and a corresponding optimal \mathcal{N}_M -coloring of D_M . Note that the instance $(D_M, \mathcal{N}_M, \tau)$ is feasible, as we have the estimate

$$\sum_{v \in V(D_M)} \mathcal{N}_M(v) = \sum_{i=1}^{\ell} \vec{\chi}_{\mathcal{N}|_{M_i}}(D[M_i]) \le \sum_{i=1}^{\ell} \sum_{v \in M_i} \mathcal{N}(v) = \sum_{v \in V(D)} \mathcal{N}(v) \le \tau.$$

By the procedure explained in (the proof of) Lemma 10.13, from an optimal color-list assignment for D_M with respect to \mathcal{N}_M we obtain an optimal color-list assignment for D, and furthermore can calculate the \mathcal{N} -dichromatic number of D via $\vec{\chi}_{\mathcal{N}}(D) = \vec{\chi}_{\mathcal{N}_M}(D_M)$.

It remains to argue for the correctness of the algorithm. First of all, the algorithm returns an optimal \mathcal{N} -coloring of D in finite time: In each of the recursive calls, the number of vertices of the digraphs in the instances is strictly smaller than the number of vertices in the current digraph. At some point, we therefore have reduced the task to such instances where the digraphs consist of a single vertex. In this case, the algorithm outputs a solution without further recursion.

For the runtime-analysis, we again consider a rooted tree T corresponding to the execution of the algorithm, where the root vertex is identified with the digraph D in the instance, and the remaining vertices are each identified with a different induced subdigraph D' which appears in a recursive call during the execution. The children of a vertex corresponding to such a digraph D' are associated with the induced subdigraphs $D'[M'_1], \ldots, D'[M'_s]$ for the corresponding module-decomposition $\{M'_1, \ldots, M'_s\}$. The leafs of this tree correspond to single-vertex digraphs.

The runtime consumed by a vertex in T corresponding to a call with instance (D', \mathcal{N}', τ) (disregarding the time needed to execute the recursive calls corresponding to its successors) involves

- Computing a directed modular decomposition $\{M'_1, \ldots, M'_s\}$ of D' with the algorithm from [MdM05], which takes time $O(v(D') + a(D')) \leq O(|V(D')|^2)$.
- Computing the digraphs $D'[M'_1], \ldots, D'[M'_s]$ and the module-digraph D'_M . This certainly can be executed in time $\leq O(v(D')^2)$.
- Applying the algorithm from Lemma 10.14 to the instance $(D'_M, \mathcal{N}'_M, \tau)$, which can be executed $(D'_M \text{ has } s \leq \omega \text{ vertices})$ in time $O(f(\omega) \log^2 \tau)$.
- Constructing an optimal \mathcal{N}' -coloring of D' given an optimal \mathcal{N}'_M -coloring of D'_M and optimal $\mathcal{N}'|_{M'_i}$ -colorings of $D'[M'_i]$ for all $i \in [s]$. By Lemma 10.13, we get $O(\omega v(D'))$ as an upper bound for the required runtime.

This yields an upper bound of $O(v(D')^2 + f(\omega)\log^2 \tau) \leq O(n^2 + f(\omega)\log^2 \tau)$ for the runtime needed for the computations corresponding to the non-leaf vertex D' of the tree. For a leaf vertex corresponding to a digraph with unique vertex v, we only need to output $\mathcal{N}(v) \leq \tau$ different colors, which requires linear time in τ . The number of leafs of T clearly is the number of vertices of D, which is n. Therefore, summing over all vertices of T, we conclude that the runtime of our algorithm is at most

$$O(\underbrace{v(T)(n^2 + f(\omega)\log^2 \tau)}_{\text{branchings}} + \underbrace{n\tau}_{\text{leafs}}).$$

By Fact 10.2, we have $v(T) \leq 2n - 1$. This finally yields the desired upper bound of $O(n^3 + f(\omega)n\log^2 \tau + n\tau)$ for the total run-time.

As in Lemma 10.14, we can upper-bound f by $f(\omega) = 2^{O(\omega 2^{\omega})}$.

Part III

Other Coloring Concepts for Digraphs

Chapter 11

Complete Colorings

Graphs and digraphs in this chapter are allowed to have multiple parallel edges as well as (anti-)parallel arcs, but are loopless. A parallel pair of edges in an undirected multi-graph will be called a *bigon* and treated as an undirected cycle of length 2. The character of our problems often depends on the existence of digons in digraphs and bigons in undirected multi-graphs, respectively. We will emphasize this distinction at the respective points.

11.1 Introduction

A complete coloring of a graph is a proper vertex coloring such that the identification of any two colors produces a monochromatic edge. The *achromatic number* $\Psi(G)$ is the maximum number of colors in a complete coloring. There has been a substantial amount of research on the achromatic number since its introduction in [HHP67], we refer to [HM97] and [Edw97] for survey articles on this topic.

In the same spirit, for most coloring parameters an associated notion of *complete* coloring and an "a-parameter" may be defined. The dichromatic number studied in Part II is the smallest size of a partition of the vertex set of a directed graph into acyclic subsets. Similarly, for an undirected graph G, the vertex arboricity va(G) is defined to be the minimal number of induced forests which cover all the vertices. This is another wellstudied graph parameter, see e.g. [GKW68, GK69, Škr02, RW08, HM12a, RW12]. Note that the dichromatic number of a digraph is at most the vertex-arboricity of its underlying undirected graph, while the dichromatic number of a bioriented digraph coincides with the chromatic number of the underlying undirected graph.

In this chapter, we investigate complete colorings corresponding to the above two coloring parameters, resulting in the adichromatic number of directed graphs and the a-vertex arboricity of undirected graphs. More precisely, the *adichromatic number* $\operatorname{adi}(D)$ of a directed graph D is the largest number of colors its vertices can be colored with such that every color induces an acyclic subdigraph but in the merge of any two color classes there is a directed cycle. We refer to such a coloring as a *complete (acyclic) coloring* of D. Similarly, the *a-vertex arboricity* $\operatorname{ava}(G)$ of an undirected graph G is the largest number of colors that can be used such that every color class induces a forest but in the merge of any two color classes there is a cycle. Such a coloring will be referred to as a *complete (arboreal) coloring* of G.

Similar parameters have been introduced in [Edw13, Sop14]. In particular, the *diachromatic number* of digraphs was introduced in [APMBORM18] and sparked the investigations presented in this chapter. While it is closest to our parameter, the adichromatic

number, it still behaves quite differently, as we will note in Proposition 11.9 that transitive tournaments have arbitrarily large diachromatic number but adichromatic number one.

In this chapter, we initiate the study of adi and ava, their relation to other graph parameters, and their behavior with respect to graph and digraph operations.

For other notions of complete colorings, so-called *interpolation theorems* have been shown. See [APMBORM18, Theorem 22] and [HHP67] for the diachromatic and achromatic versions, respectively. As our first main result of this chapter we extend these results to the adichromatic number and a-vertex arboricity as follows.

Theorem 11.1.

- 1. Let D be a digraph and let $\ell \in \mathbb{N}$. Then there exists a complete acyclic coloring of D using exactly ℓ colors if and only if $\vec{\chi}(D) \leq \ell \leq \operatorname{adi}(D)$.
- 2. Let G be a graph and let $\ell \in \mathbb{N}$. Then there exists a complete arboreal coloring of G using exactly ℓ colors if and only if $\operatorname{va}(G) \leq \ell \leq \operatorname{ava}(G)$.

The main focus of this chapter is on the relation of adi and ava to important graph and digraph parameters such as the degeneracy and more importantly the size τ of a smallest feedback vertex set. It is not hard to see that both adi and ava are bounded from above by $\tau + 1$, as we will show in Proposition 11.5. An interesting and natural question to ask then is whether also an inverse relationship between adi, ava and the respective feedback-vertex-set parameters exists.

In Proposition 11.18 we answer this question negatively by providing a construction of digraphs showing that there is no function f such that $\tau \leq f(\text{adi})$ in the directed setting, not even for digraphs of large digirth.

In contrast to this, our second main result in this chapter gives a positive answer to this question for *simple* undirected graphs, showing that the parameters τ and ava are qualitatively tied to each other.

Theorem 11.2. There is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every simple graph G

$$\tau(G) \le f(\operatorname{ava}(G)).$$

While the best asymptotic lower bound on the function f we know of is $f(k) = \Omega(k^2)$, the upper bound on f we could obtain from our proof of Theorem 11.2 is far from polynomial. It is therefore desirable to find improved (polynomial) upper bounds on $\tau(G)$ in terms of ava(G). Interestingly, if we restrict our attention to graphs contained in any fixed non-trivial minor-closed class such as the planar graphs, we can actually give an asymptotically almost tight upper bound on τ in terms of ava. Moreover, we can show that in contrast to general digraphs, the parameters adi and τ are tied to one another for digraphs whose underlying graphs are contained in such a minor-closed class of graphs.

These results are summarized in the following Theorem, which is our third and last main result of this chapter.

Theorem 11.3. Let \mathcal{G} be a minor-closed class of simple graphs which is non-trivial (that is, it does not contain all graphs).

 There is a function f : N → N (depending on G), such that for any digraph D whose simple underlying graph (obtained from ignoring parallel edges) is contained in G, we have

$$\tau(D) \le f(\operatorname{adi}(D)).$$

2. There is a constant C > 0 (depending on \mathcal{G}) such that for every graph G whose simplification (identifying parallel edges) lies in \mathcal{G} , we have

$$\tau(G) \le C \cdot \operatorname{ava}(G)^2 \log(\operatorname{ava}(G)).$$

An interesting consequence of Theorem 11.2 is that there exists a function f such that $adi(D) \leq f(ava(G))$ for every orientation D of a simple graph G (Corollary 11.23). Conversely, we show that any graph with sufficiently high a-vertex arboricity has an orientation with large adichromatic number (Proposition 11.24).

We further investigate the relationship between the parameters adi, ava and the degeneracy of the corresponding (di)graphs. Again, we exhibit a difference between directed and undirected graphs: While there are arbitrarily dense digraphs with bounded adichromatic number (Proposition 11.18), the degeneracy of a simple graph is bounded from above by a function of ava (Theorem 11.20).

From a more general point of view, we discuss the relations of our findings to the Erdős-Pósa property and introduce a strengthening of the latter that we call τ -boundedness.

Structure of the chapter. Section 11.2 contains first observations and results. It consists of Subsection 11.2.1, which studies some relations between the parameters and contains the proof of Theorem 11.1, and of Subsection 11.2.2, which contains results on the behavior of our parameters with respect to graph and digraph operations. Section 11.3 contains the proofs of Theorem 11.2 and Theorem 11.3, and the above mentioned relations of ava and adi with τ and the degeneracy are given. Finally, Section 11.4 contains the discussion of τ -boundedness and the Erdős-Pósa property.

11.2 First Observations and Results

In this section, we present some basic properties of the adichromatic number and the avertex arboricity, discuss the relation of the adichromatic number and some other notions of complete coloring and make observations which will be used in the rest of the chapter.

11.2.1 Relations between Parameters

A fundamental property of an optimal acyclic coloring of a digraph (i.e., without monochromatic directed cycles) is that we cannot improve the coloring by any merge of two colors. Therefore, any acyclic coloring of a digraph D with $\vec{\chi}(D)$ colors must be complete. With the analogous argument for the a-vertex arboricity, we get the following basic relations.

Observation 11.4.

- $\vec{\chi}(D) \leq \operatorname{adi}(D)$ for any digraph D.
- $\operatorname{va}(G) \leq \operatorname{ava}(G)$ for any graph G.

Clearly, the only digraphs with adichromatic number 1 are acyclic digraphs, and similarly the only graphs with a-vertex arboricity 1 are forests. The same is true for the dichromatic number and the vertex arboricity, and so in these very special cases, the concepts of coloring and complete coloring coincide. The examples from Proposition 11.7 below show that in general, no such relation exists, as there are digraphs with dichromatic number 2 and arbitrarily large adichromatic number and graphs with vertex arboricity 2 but arbitrarily large a-vertex arboricity. Before turning to these examples, let us note the following fundamental relation of the treated coloring parameters to feedback vertex sets.

Proposition 11.5.

- For any digraph D, we have $\operatorname{adi}(D) \leq \tau(D) + 1$.
- For any graph G, we have $\operatorname{ava}(G) \leq \tau(G) + 1$.

Proof. We give the proof for the directed case, the undirected case is completely analogous. Denote by $F \subseteq V(D)$ a directed feedback vertex set of D with minimal size. Assume towards a contradiction that $\operatorname{adi}(D) \ge |F| + 2$. Then there is a complete acyclic partition (V_1, \ldots, V_k) of D with $k \ge |F| + 2$ colors. Consequently, there are at least two colors $i, j \in \{1, \ldots, k\}$ which do not appear on any vertex of F. However, this implies that $V_i \cup V_j \subseteq V(D) \setminus F$ induces an acyclic subdigraph of D, contradicting the definition of a complete acyclic coloring.

The given upper bounds are easily seen to be tight for acyclic digraphs, directed cycles of arbitrary length, and complete digraphs in the directed case; respectively for forests, cycles and multi-graphs obtained from complete graphs by replacing each simple edge by a bigon in the undirected case. As mentioned in the introduction, it is a natural question whether there is also an inverse relationship between the parameters $\operatorname{adi}(D)$ and $\tau(D)$ (respectively $\operatorname{ava}(G)$ and $\tau(G)$). This question is the center of the investigations in this chapter and will be dealt with in Section 11.3.

Since digons and bigons are counted as (directed) cycles of length two, it is easily seen that both the adichromatic number and the a-vertex arboricity form a proper generalization of the achromatic numbers of graphs when allowing multiple edges.

Observation 11.6. For a simple graph G, let 2G be the multi-graph obtained from G by replacing every edge with a bigon connecting the same endpoints. Then we have the following equalities involving the achromatic number of G:

$$\Psi(G) = \operatorname{adi}(\vec{G}) = \operatorname{ava}(2G).$$

Proof. The three parameters are defined as the maximal size of a partition into independent (or acyclic) vertex sets of maximal size such that the union of any two partition classes is not independent (or acyclic, respectively) any more. The claim therefore follows from observing that a vertex set in $\overset{\leftrightarrow}{G}$ induces an acyclic subdigraph if and only if it induces a forest in 2G and if and only if it is independent in G.

As independent sets form a special case of acyclic vertex sets in graphs, which again define acyclic vertex sets in any orientation of that graph, it is easily seen that for any digraph D with underlying graph G, we have $\vec{\chi}(D) \leq \operatorname{va}(G) \leq \chi(G)$. It is therefore natural to ask whether similar relationships between the adichromatic number, a-vertex arboricity and achromatic number exist. The following presents a set of canonical graphs and digraphs with their a-coloring parameters, which show that in general, both $\operatorname{adi}(D)$ and $\operatorname{ava}(G)$ are not bounded from above in terms of $\Psi(G)$, where G is the underlying graph of D. Note that vice-versa, $\Psi(G)$ cannot be bounded from above in terms of $\operatorname{adi}(D)$ or $\operatorname{ava}(G)$, as it is easily seen to be unbounded already for matchings.

Proposition 11.7. For any $m, n \in \mathbb{N}$, $m, n \ge 1$ we have

- 1. $\operatorname{ava}(K_{m,n}) = \min\{m, n\}, while \Psi(K_{m,n}) = 2.$
- 2. $\operatorname{ava}(K_n) = \left\lceil \frac{n}{2} \right\rceil$.

- 3. Let D_n be the orientation of $K_{n,n}$ in which a perfect matching is directed from the first to the second class of the bipartition, while all non-matching edges emanate from the second class. Then $\operatorname{adi}(D_n) = n$.
- 4. Let G (respectively D) be the vertex-disjoint union of $\binom{n}{2}$ cycles (respectively directed cycles), $n \ge 2$. Then $\operatorname{ava}(G) = \operatorname{adi}(D) = n$.

Proof.

- 1. Assume that $m \leq n$. We first observe that $\operatorname{ava}(K_{m,m}) \geq m$: Consider a perfect matching and assign m different colors to the pairs of matched vertices. Now the union of any two color classes produces a cycle of length 4. Because $K_{m,m}$ is an induced subgraph of $K_{m,n}$, we have $\operatorname{ava}(K_{m,n}) \geq \operatorname{ava}(K_{m,m}) \geq m = \min\{m,n\}$ (this will be noted later in Corollary 11.12). On the other hand, removing all but 1 vertex from the smaller partite class of $K_{m,n}$ shows that $\tau(K_{m,n}) \leq \min\{m,n\} - 1$. The claim follows from Proposition 11.5. The fact that $\Psi(K_{m,n}) = 2$ is well-known and follows by observing that no two distinct colors can appear on the same partition class in any complete proper coloring.
- 2. In any complete arboreal coloring of K_n , each color class has size at most 2 and there is at most one singleton-color class (otherwise two singleton-colors could be merged). We therefore have $\operatorname{ava}(K_n) \leq \lceil \frac{n}{2} \rceil$.

If n is even, K_n contains $K_{n/2,n/2}$ as an induced subgraph and we therefore have ava $(K_n) \ge ava(K_{n/2,n/2}) = \frac{n}{2}$. If n is odd, we can use $\frac{n-1}{2}$ colors to color paired vertices of K_{n-1} and then add an extra color for the remaining vertex. This yields a complete arboreal coloring with $\frac{n}{2} + 1$ colors. This verifies the claim in both the even and the odd case.

- 3. To verify $\operatorname{adi}(D_n) \ge n$, we observe that by taking the two vertices of each matching edge as a color class we define a partition into n acyclic subdigraphs such that the union of any two creates a directed cycle of length 4. Deleting all vertices of one partite class except one shows that $\tau(D_n) \le n-1$. Again, the claim follows using Proposition 11.5.
- 4. In any complete arboreal (respectively acyclic) coloring of G (respectively D), we must have a cycle (respectively directed cycle) in the union of any two color classes, and so if we use k colors, we need at least $\binom{k}{2}$ distinct cycles. This proves the inequalities $\operatorname{ava}(G), \operatorname{adi}(D) \leq n$. On the other hand, assigning to each pair $\{i, j\} \in \binom{[n]}{2}$ a different cycle and coloring this cycle using only i and j defines a complete arboreal (respectively acyclic) coloring of G (respectively D) and proves that $\operatorname{ava}(G), \operatorname{adi}(D) \geq n$.

Recently, the concept of the diachromatic number was introduced in [APMBORM18]. Given a digraph D, the diachromatic number dac(D) of D is defined as the maximum number k of colors that can be used in an acyclic coloring of D with color classes V_1, \ldots, V_k , such that for every ordered pair $(i, j) \in [k]^2$, $i \neq j$, there exists at least one arc of D with head in V_i and tail in V_j .

In a complete acyclic coloring of a digraph, the union of any two color classes contains the vertex set of a directed cycle, and therefore arcs in both directions between the two acyclic color classes must exist. Hence, any complete acyclic coloring in our sense also defines a complete coloring as defined in [APMBORM18]. We therefore have **Observation 11.8.** For any digraph D, it holds that $\operatorname{adi}(D) \leq \operatorname{dac}(D)$.

In general however, the above estimate is far from being tight. For example, the diachromatic number is not bounded for directed acyclic digraphs, which always have adichromatic number equal to 1. Together with [APMBORM18, Corollary 12] we get:

Proposition 11.9. For the transitive tournament \vec{K}_n on n vertices it holds that

$$\operatorname{adi}(\vec{K}_n) = 1, \operatorname{dac}(\vec{K}_n) = \left\lceil \frac{n}{2} \right\rceil.$$

Our next goal is to prove the interpolation Theorem 11.1.

In order to do so, we consider generalizations of the dichromatic number and the vertex arboricity, where we want to minimize the number of colors in an acyclic coloring with the additional restriction that certain vertices must be colored the same.

Definition 11.1.

- 1. Let D be a digraph, and let $\mathcal{P} = \{P_1, \ldots, P_t\}$ be a partition of the vertex set such that $D[P_i]$ is acyclic for all $i \in [t]$. Then we define the \mathcal{P} -dichromatic number of D, denoted by $\vec{\chi}_{\mathcal{P}}(D)$, to be the least number of colors required in an acyclic digraph coloring of D such that for any i, the vertices in P_i receive the same color.
- 2. Let G be a graph, and let $\mathcal{P} = \{P_1, \ldots, P_t\}$ be a partition of the vertex set such that $G[P_i]$ is a forest for all $i \in [t]$. Then we define the \mathcal{P} -vertex arboricity of G, denoted by $\operatorname{va}_{\mathcal{P}}(G)$, to be the least number of colors required in an arboreal coloring of G such that for any i, the vertices in P_i receive the same color.

We prepare the proof of Theorem 11.1 with the following simple lemma.

Lemma 11.10.

1. Let D be a digraph with a partition $\mathcal{P} = \{P_1, \ldots, P_t\}$ into acyclic vertex sets. Assume that also $D[P_1 \cup P_2]$ is acyclic, and let $\mathcal{Q} := \{P_1 \cup P_2, P_3, \ldots, P_t\}$. Then we have

$$\vec{\chi}_{\mathcal{P}}(D) \le \vec{\chi}_{\mathcal{Q}}(D) \le \vec{\chi}_{\mathcal{P}}(D) + 1.$$

2. Let G be a graph with a partition $\mathcal{P} = \{P_1, \ldots, P_t\}$ into vertex sets inducing forests. Assume that also $G[P_1 \cup P_2]$ is a forest, and let $\mathcal{Q} := \{P_1 \cup P_2, P_3, \ldots, P_t\}$. Then

$$\operatorname{va}_{\mathcal{P}}(G) \le \operatorname{va}_{\mathcal{Q}}(G) \le \operatorname{va}_{\mathcal{P}}(G) + 1.$$

Proof. We give a proof for the directed case, the undirected case is analogous. Clearly, any acyclic coloring of D which is compatible with Q is also compatible with \mathcal{P} . This directly yields $\vec{\chi}_{\mathcal{P}}(D) \leq \vec{\chi}_{\mathcal{Q}}(D)$.

On the other hand, let $c: V(D) \to [\ell]$ be an acyclic coloring of D using $\ell = \chi_{\mathcal{P}}(D)$ colors such that the vertices in P_i are colored the same, for all $i \in [t]$. It is now easily seen that coloring all vertices in $P_3 \cup P_4 \cup \cdots \cup P_t$ as in c and giving color $\ell + 1$ to all vertices in $P_1 \cup P_2$ defines a proper digraph coloring of D which is compatible with \mathcal{Q} and uses at most $\ell + 1$ colors. This proves the inequality $\chi_{\mathcal{Q}}(D) \leq \chi_{\mathcal{P}}(D) + 1$.

Proof of Theorem 11.1. We prove the first part of the Theorem, the proof of the second part is completely analogous.

First of all, the conditions on ℓ are necessary: Any complete acyclic coloring also is an acyclic digraph coloring, so it uses at least $\vec{\chi}(D)$ colors, and by definition, it cannot use more than $\operatorname{adi}(D)$ colors.

So let now $\ell \in {\vec{\chi}(D), \ldots, \operatorname{adi}(D)}$ be given. Define $\mathcal{P}_0 := {\{v\} | v \in V(D)}$ to be the partition of V(D) into singletons, and let \mathcal{P} denote the partition of V(D) into the $\operatorname{adi}(D)$ color classes corresponding to a complete acyclic coloring of D with the maximum number of colors. Looking at Definition 11.1, it is readily verified that $\vec{\chi}_{\mathcal{P}_0}(D) = \vec{\chi}(D)$ and $\vec{\chi}_{\mathcal{P}}(D) = \operatorname{adi}(D)$ (for the latter note that different partition classes must be colored differently, as their union is cyclic). Now consider a sequence $\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_r = \mathcal{P}$ consisting of partitions of V(D) into acyclic vertex sets such that for any $i = 0, 1, \ldots, r-1, \mathcal{P}_{i+1}$ is obtained from \mathcal{P}_i by merging a pair of partition classes. The existence is easily seen by successively splitting partition classes, starting from the back with \mathcal{P} . Applying Lemma 11.10 we get that $\vec{\chi}_{\mathcal{P}_i}(D) \leq \vec{\chi}_{\mathcal{P}_{i+1}}(D) \leq \vec{\chi}_{\mathcal{P}_i}(D) + 1$ for all $0 \leq i < r$. It follows from this and from $\vec{\chi}_{\mathcal{P}_0}(D) \leq \ell \leq \vec{\chi}_{\mathcal{P}_r}(D)$ that there exists some $i \in [r]$ such that $\ell = \vec{\chi}_{\mathcal{P}_i}(D)$. Let now $c: V(D) \to [\ell]$ denote an optimal digraph coloring compatible with \mathcal{P}_i which uses the fewest number ℓ of colors. This coloring must be a complete acyclic coloring: If the union of two color classes was acyclic, we could improve the number of colors used by merging these color classes, still keeping it compatible with \mathcal{P}_i . As we have found a complete acyclic coloring using exactly ℓ colors, this verifies the claim.

11.2.2 Behavior with Respect to Graph Operations

Most standard coloring parameters such as the chromatic number are monotone under subgraphs. While this is not the case for the adichromatic number and a-vertex arboricity in general (consider a bidirected C_4 in the directed case and the multi-graph obtained from C_4 by replacing all edges with bigons for small examples, or alternatively the digraphs described in Proposition 11.18 in the next section), we can establish a monotonicity under taking induced subgraphs.

Lemma 11.11. Let D be a digraph and G a graph.

- For any $v \in V(D)$, we have $\operatorname{adi}(D) 1 \leq \operatorname{adi}(D v) \leq \operatorname{adi}(D)$.
- For any $v \in V(G)$, we have $\operatorname{ava}(G) 1 \leq \operatorname{ava}(G v) \leq \operatorname{ava}(G)$.

Proof. We prove the claim for the directed case, the undirected case is analogous. To prove that $\operatorname{adi}(D-v) \leq \operatorname{adi}(D)$, consider an optimal complete acyclic coloring of D-v with color classes $V_1, \ldots, V_l, \ell := \operatorname{adi}(D-v)$. If there is a color $i \in [\ell]$ such that $V_i \cup \{v\}$ induces an acyclic subdigraph, we can join v to this color class in order to obtain a complete acyclic coloring of D with ℓ colors. If on the other hand $V_i \cup \{v\}$ contains a directed cycle for every $i \in [\ell]$, we can give v the new color $\ell + 1$ and see that this defines a complete acyclic coloring of D using $\ell + 1$ colors. This implies that $\operatorname{adi}(D) \geq \ell = \operatorname{adi}(D-v)$ in every case.

For the second inequality, we have to prove that $\operatorname{adi}(D) - 1 \leq \operatorname{adi}(D-v)$. To see this, consider a complete acyclic coloring of D using $r := \operatorname{adi}(D)$ colors. There are at least r-1 color classes in this coloring which were not affected by the deletion of v from D, and therefore, the union of any two of these color classes still contains a directed cycle. Clearly, each of the r-1 color classes still induces an acyclic subdigraph. Now simply give a unique new color to each vertex in V(D-v) which is in none of the r-1 color classes. Perform a greedy merging process in which, as long as there exists a pair of color classes whose union is acyclic, we merge them. In the end, we obtain a complete coloring of D-v, in which no two of the r-1 color classes considered above was merged. Therefore, we have found a complete acyclic coloring of D using at least r - 1 = adi(D) - 1 colors. This concludes the proof.

Corollary 11.12. Let D_1, D_2 be digraphs and G_1, G_2 graphs.

- If D_1 is an induced subdigraph of D_2 , then $\operatorname{adi}(D_1) \leq \operatorname{adi}(D_2)$.
- If G_1 is and induced subgraph of G_2 , then $\operatorname{ava}(G_1) \leq \operatorname{ava}(G_2)$.

A special role in the theory of acyclic digraph colorings is played by directed separations. If $S = \partial^+(X)$ forms a directed cut for some partition $X, \overline{X} := V(D) \setminus X$ of the vertexset, no directed cycle in D can use an edge from S and therefore either stays in D[X] or in $D[\overline{X}]$, which implies that $\overline{\chi}(D) = \max\{\overline{\chi}(D[X]), \overline{\chi}(D[\overline{X}])\}$. Iterating this argument, one can see that the dichromatic number of a digraph can be computed as the maximum over the dichromatic numbers of its strongly connected components. For the adichromatic number, such a simple relation does not hold true, as even for the disjoint union of two digraphs, there is no explicit way of computing the adichromatic number in terms of the adichromatic numbers of the two components (cf. Proposition 11.7, 4.). However, we can bring it down to exactly this case, by noting that the adichromatic number of a digraph equals the adichromatic number of the disjoint union of its strong components.

Observation 11.13. Let D be a digraph, and let $S \subseteq A(D)$ be a directed cut. Then adi(D) = adi(D - S).

Proof. This follows directly from the fact that the complete acyclic colorings of D are the same as those of D - S, because a vertex subset $X \subseteq V(D)$ is acyclic in D if and only if it is acyclic in D - S.

We can go one step further and consider cuts which are almost directed, i.e., they have only a single arc in forward-direction. In this case, we can contract this forward-arc without increasing the adichromatic number.

Lemma 11.14. Let D be a digraph with a non-trivial partition (X, \overline{X}) of the vertex set such that $\partial^+(X) = \{e\}$ for some arc $e \in A(D)$ which does not form part of a digon in D. Then we have $\operatorname{adi}(D/e) \leq \operatorname{adi}(D)$.

Proof. Let $c: V(D/e) \to [\ell]$ be a complete acyclic coloring of D/e using $\ell = \operatorname{adi}(D/e)$ colors. Let c' be the vertex-coloring of D in which every vertex not incident to e is colored as by c, and where the endpoints of e both receive the color which is given to the contraction vertex of e under c. It is clear that this still defines an acyclic coloring, as any directed cycle in D, after contracting e, still yields a directed cycle in D/e using exactly the same colors. On the other hand, if $i \neq j \in [\ell]$ is a pair of colors, then there exists a directed cycle in D/e which uses only colors i and j. Possibly after re-inserting the arc ein case the cycle uses the contraction vertex of e now defines a directed cycle in D which also only uses colors i and j (the fact that the cycle remains directed after reinserting efollows since there is no directed path in D - e starting in tail(e) and ending in head(e), since $\partial^+(X) = \{e\}$). We therefore have found a complete acyclic coloring of D which uses ℓ colors. This yields that $\operatorname{adi}(D) \geq \ell = \operatorname{adi}(D/e)$, as required. \Box

This operation from Lemma 11.14 resembles the generalized butterfly-contractions from Chapter 5. In particular, if $e \in A(D)$ is a butterfly-contractible arc with tail uand head v, then we have $\partial^+(\{u\}) = \{e\}$ if e is the only arc emanating from u, and $\partial^+(V(D) \setminus \{v\}) = \{e\}$ if e is the only arc entering v. Therefore, the following is a direct consequence of Corollary 11.12 and Lemma 11.14: **Corollary 11.15.** Let D_1 be an induced butterfly-minor of D_2 , i.e., D_1 is obtained from an induced subdigraph of D_2 by repeatedly contracting butterfly-contractible arcs not contained in digons. Then

$$\operatorname{adi}(D_1) \leq \operatorname{adi}(D_2).$$

The following statements, which yield lower bounds on the a-vertex arboricity of a graph in terms of the a-vertex arboricities of induced minors, will form a central tool in the proof of our second main result, Theorem 11.2.

Lemma 11.16. Let G be a (multi-)graph and let T = G[X] be an induced subtree of G. Let G/T denote the (multi-)graph obtained from G by deleting all edges of T from G and identifying X into a single vertex v_X . Then

$$\operatorname{ava}(G/T) \leq \operatorname{ava}(G).$$

Proof. Let $\ell := \operatorname{ava}(G/T)$ and let $c : V(G/T) \to [\ell]$ be a complete arboreal coloring of G/T using all ℓ colors. We claim that the coloring $c' : V(G) \to [\ell]$,

$$c'(v) := \begin{cases} c(v), & \text{if } v \notin X, \\ c(v_X), & \text{if } v \in X, \end{cases}$$

is also a complete arboreal coloring of G that uses ℓ colors. For this purpose, we must verify that there are no monochromatic cycles in G with respect to c and that in the union of any two color classes, there is a cycle. For the first part, suppose there was a cycle Cin G all whose vertices are colored i in c'. Because T is an induced tree, C must contain a vertex outside X. Consequently, after the contraction of X, the cycle C yields a closed walk of positive length in G/T, which by definition of c' must still be monochromatic, a contradiction. On the other hand, given any pair $i \neq j \in [\ell]$ of colors, there is a cycle in G/T which uses only colors i and j according to c. If the cycle does not use v_X , this yields also a cycle in G which only uses colors i and j according to c', as desired. In the case that the cycle traverses v_X , let e, f be the two incident edges of v_X on the cycle. By connecting the endpoints of e and f in T by the unique monochromatic connection path in T if necessary, we find that also in this case there is a cycle in G which uses only colors i, j according to c'. Hence, c' defines a complete arboreal coloring of G with ℓ colors, and ava $(G) \geq \ell = ava(G/T)$.

Corollary 11.17. Let G and H be simple graphs such that G contains a subdivision of H as an induced subgraph. Then $\operatorname{ava}(G) \ge \operatorname{ava}(H)$.

Proof. Repeated application of Lemma 11.16 to contractions of subdivision edges yields that the a-vertex arboricity of any subdivision of a graph is lower bounded by the a-vertex arboricity of the graph itself. The claim now follows from Corollary 11.12. \Box

11.3 Upper Bounds for Minimum Feedback Vertex Sets

The main goal of this section is to complement the lower bounds on τ via ava and adi from Proposition 11.5 with upper bounds, and in particular, to give the proofs of Theorem 11.2 and Theorem 11.3.

By Observation 11.6, the adichromatic number of the biorientation $\overset{\leftrightarrow}{K}_{n,n}$ of the complete bipartite graph is given by the achromatic number of $K_{n,n}$, which is 2. However, the size of a smallest feedback vertex set equals n. Similarly, the multi-graph $2K_{n,n}$ has a-vertex arboricity 2 but $\tau(2K_{n,n}) = n$ for any $n \ge 1$. These observations show that there is no qualitative upper bound on the size of a smallest feedback vertex set in terms of ava or adi for simple digraphs or multi-graphs.

In the rest of this section, we therefore focus on oriented digraphs and simple graphs and demonstrate that while there are simple digraphs D with bounded adichromatic number and unbounded $\tau(D)$ (Proposition 11.18), $\tau(G)$ is bounded above in terms of ava(G)for simple graphs (by proving Theorem 11.2). On the way to this result we achieve an upper bound for the degeneracy by a function of ava in the form of Corollary 11.20. We also give the proof of Theorem 11.3 and thereby establish a relationship of adi and τ within minor-closed classes. As consequences of the results mentioned above we show non-trivial relations between ava and adi in Subsection 11.3.1.

The following construction gives a family of oriented digraphs with an unbounded size of the feedback vertex set but bounded adichromatic number. Additionally, these digraphs can have arbitrarily large directed girth.

Let D(n,k) with $n \ge 1, k \ge 3$ denote the *n*-fold blow-up of \vec{C}_k , that is, the *k*-partite digraph whose vertex set consists of *k* disjoint partition classes V_1, \ldots, V_k of size *n* each and where $A(D) = \bigcup_{i=1}^k (V_i \times V_{i+1}) \ (k+1 := 1)$. Hence D(n,k) is obtained from \vec{C}_k by replacing each vertex by *n* independent copies.

Proposition 11.18. For $n \ge 1, k \ge 3$, we have $\operatorname{adi}(D(n,k)) \le k$ while $\tau(D(n,k)) = n$.

Proof. Clearly, we can find a packing of n vertex-disjoint directed cycles in D(n, k), and so $\tau(D(n,k)) \ge n$. On the other hand, V_1 forms a feedback vertex set, and we conclude that $\tau(D(n,k)) = n$. To see that $\operatorname{adi}(D(n,k)) \le k$, let $c: V(D(n,k)) \to [\ell]$ be a complete acyclic coloring using ℓ colors, and assume towards a contradiction that $\ell \ge k + 1$. For each $i \in [\ell]$, there is at least one partition class in which i does not appear, otherwise there would be a directed cycle colored i. By the pigeon-hole principle, we therefore find a pair $i \neq j \in [\ell]$ of colors such that both do not appear in a certain partition class. However, there must be a directed cycle in D(n,k) using only vertices with color i or j. This contradiction shows $\operatorname{adi}(D(n,k)) \le k$.

Our next goal will be to prove Theorem 11.2. The proof needs to be prepared with some auxiliary statements. For our next result we will combine Corollary 11.17 with the following strong result from the literature obtained by Kühn and Osthus [KO04]:

Theorem 11.19 ([KO04]). For every $s \in \mathbb{N}$ and any simple graph K there is some $d = d(s, K) \in \mathbb{N}$ such that every simple graph G with minimum degree greater than d contains $K_{s,s}$ as a subgraph or a subdivision of K as an induced subgraph.

For a simple graph G, the *degeneracy* of G is defined as $dgn(G) := \max_{H \subseteq G} \delta(H)$, where the maximum is taken over all subgraphs (or, equivalently, all induced subgraphs) of G. We call G d-degenerate for $d \in \mathbb{N}$, if $dgn(G) \leq d$. The following shows that simple graphs of bounded a-vertex-arboricity have bounded degeneracy.

Corollary 11.20. There exists a function $g : \mathbb{N} \to \mathbb{N}$ such that $dgn(G) \leq g(ava(G))$ for every simple graph G.

Proof. For any $k \in \mathbb{N}$, define g(k) as the integer d(s, K) from Theorem 11.19 where s = k + 1 and $K = K_{k+1,k+1}$.

Now let G be an arbitrary simple graph and let $k := \operatorname{ava}(G)$. We have to prove that $\operatorname{dgn}(G) \leq g(k)$. Assume towards a contradiction that $\operatorname{dgn}(G) > g(k)$, i.e., there exists an induced subgraph H of G such that the minimum degree of H is greater than $g(k) = d(k + 1, K_{k+1,k+1})$. By Corollary 11.12, we have $k = \operatorname{ava}(G) \ge \operatorname{ava}(H)$. By Theorem 11.19, H either contains $K_{k+1,k+1}$ as a subgraph or as an induced subdivision. In the first case, let $X \subseteq V(H)$ be the set of vertices of the subgraph. On the one hand, we know that $k \ge \operatorname{ava}(H) \ge \operatorname{ava}(H[X])$. On the other hand, H[X] is a simple graph which contains $K_{k+1,k+1}$ as a spanning subgraph. Consider a perfect matching of $K_{k+1,k+1}$ and color the vertices in X with k + 1 different colors, such that end vertices of the same matching edge have the same color, and all matching edges are colored differently. It is now easily seen that this defines a complete arboreal coloring of H[X] with more than kcolors, which yields the desired contradiction in this case. In the second case we directly apply Corollary 11.17 to obtain the contradiction $k \ge \operatorname{ava}(H) \ge \operatorname{ava}(K_{k+1,k+1}) = k+1$. \Box

The last ingredient of our proof is the following well-known theorem of Erdős and Pósa, which relates the maximum size of a vertex-disjoint cycle packing in a graph to the minimum size of a feedback vertex set. For a graph G let $\nu(G)$ denote the maximal size of a collection of pairwise vertex-disjoint cycles in G.

Theorem 11.21 ([EP65]). There is an absolute constant c > 0 such that for every $k \in \mathbb{N}$, every graph G with $\tau(G) > ck \log(k)$ fulfills $\nu(G) \ge k$.

We are now prepared for the proof of Theorem 11.2.

Proof of Theorem 11.2. For a clearer presentation, we prove the theorem by contradiction. From a finer analysis of the proof, one could derive an explicit expression for f(k), the bound would be rather bad however. So assume for the rest of the proof that such a function f as claimed does not exist. This means that there is a fixed $A \in \mathbb{N}$ and an infinite sequence $(G_s)_{s=1}^{\infty}$ of simple graphs such that $\operatorname{ava}(G_s) < A$ for all $s \in \mathbb{N}$ but $\tau(G_s) \to \infty$. From Theorem 11.21 we directly conclude that also $\nu(G_s) \to \infty$.

From Corollary 11.20 we get that there exists a constant $d := \max_{l=1,\ldots,A-1} g(l) > 0$ such that all the graphs G_s are d-degenerate.

For each $s \ge 1$, we fix a packing C_s of induced (that is, chordless) and pairwise vertexdisjoint cycles in G_s of size $\nu(G_s)$.

For each $s \ge 1$, we associate with C_s a model graph M_s which has $|C_s|$ vertices, one for each cycle in C_s , and an edge between two vertices for every edge spanned between the corresponding cycles in G_s (so this might be a multi-graph). Because the cycles were assumed to be induced, we know that M_s is loopless.

Claim 1. $\alpha(M_s) < \binom{A}{2}$ for all $s \ge 1$.

Proof. Assume towards a contradiction the statement was false. Consequently, we can find some $s \ge 1$ such that M_s contains an independent set I of size $\binom{A}{2}$. Let H be the subgraph of G_s induced by the union of the vertex sets of cycles in \mathcal{C}_s corresponding to the vertices in I. Consider some bijection which maps each pair $\{i, j\} \in \binom{[A]}{2}$ to one of the $\binom{A}{2}$ cycles corresponding to I.

Since *H* is the disjoint union of $\binom{A}{2}$ cycles, it follows directly from item 4 of Proposition 11.7 that $\operatorname{ava}(H) \ge A$. Applying Lemma 11.11 to the induced subgraph *H* of *G* now yields $A > \operatorname{ava}(G_s) \ge \operatorname{ava}(H) \ge A$, which is the desired contradiction.

Applying Ramsey's Theorem to each of the graphs $M_s, s \ge 1$, we find that

$$R\left(\omega(M_s)+1, \begin{pmatrix} A\\2 \end{pmatrix}\right) > v(M_s) = \nu(G_s) \to \infty,$$

where for any $r, b \in \mathbb{N}, r, b \geq 1$, R(r, b) denotes the well-known Ramsey number. Therefore, the size $\omega(M_s)$ of a maximum clique in M_s tends to infinity for $s \to \infty$. For each $s \geq 1$, consider a clique W_s in M_s of maximum size and let G'_s be the subgraph of G_s induced by the vertices contained in the cycles corresponding to the vertices in W_s . Clearly, the subcollection \mathcal{C}'_s of \mathcal{C}_s corresponding to W_s defines a decomposition of G'_s in induced vertexdisjoint cycles of size $|\mathcal{C}'_s| = \omega(W_s) \to \infty$ in G'_s . Moreover, we have $\operatorname{ava}(G'_s) \leq \operatorname{ava}(G_s) < A$ and $\operatorname{dgn}(G'_s) \leq \operatorname{dgn}(G_s) \leq d$ for all $s \in \mathbb{N}$. In the following, we will continue working with the sequence $(G'_s)_{s=1}^{\infty}$ of simple graphs.

For a fixed $s \ge 1$ consider the graph G'_s with the cycle-decomposition $\mathcal{C}'_s = \{C_1, \ldots, C_k\}$. By the definition of \mathcal{C}'_s , for every pair C_j , C_l of cycles, there is an edge in G'_s with endpoints in $V(C_j)$ and $V(C_l)$. Let us select and fix one such edge $e_{jl} \in E(G'_s)$ for every $\{j, l\} \in \binom{[k]}{2}$.

Claim 2. There are less than R := 2A cycles $C \in \mathcal{C}'_s$ with $|C| \ge R$.

Proof. Assume towards a contradiction that there were at least R cycles in C'_s with at least R vertices each, say C_1, \ldots, C_R . For each $i \in [R]$, we can find a vertex $v_i \in V(C_i)$ which is not incident to any of the edges $\{e_{jl}|j, l \in [R]\}$. Let $X := \bigcup_{i=1}^R (V(C_i) \setminus \{v_i\})$ and consider the induced subgraph $G'_s[X]$. For every i, let $P_i := C_i - v_i$. P_1, \ldots, P_R defines a vertex-partition of $G'_s[X]$ into induced paths. Let M_X be the model (multi-)graph on R vertices obtained from $G'_s[X]$ by identifying each of P_1, \ldots, P_R into a single vertex. By Corollary 11.12 and Lemma 11.16, we know that $\operatorname{ava}(M_X) \leq \operatorname{ava}(G'_s[X]) \leq \operatorname{ava}(G'_s) < A$. Because all the edges $e_{jl}, 1 \leq j < l \leq R$ still exist in $G'_s[X]$, we know that the vertices of M_X are mutually adjacent. It now follows directly from the monotonicity of the vertex-arboricity va under taking subgraphs that $\operatorname{ava}(M_X) \geq \operatorname{va}(M_X) \geq \operatorname{va}(K_R) = \lceil \frac{R}{2} \rceil = A$. This contradiction shows that our assumption was wrong, concluding the proof.

For each $s \geq 1$, consider the subset $\mathcal{C}_s^R \subseteq \mathcal{C}_s'$ of cycles of length less than R, and consider the induced subgraph H_s of G_s' with vertex set $\bigcup_{C \in \mathcal{C}_s^R} V(C)$. For each $s \geq 1$, define $k_s := |\mathcal{C}_s^R|$, then $k_s > |\mathcal{C}_s'| - R$. By the above, we have $k_s \to \infty$ for $s \to \infty$. Because G_s' is d-degenerate, so is H_s , and therefore we have $e(H_s) \leq dn_s$, where n_s is the number of vertices of H_s . By definition, we have $n_s = \sum_{C \in \mathcal{C}_s^R} |C| \leq Rk_s$. On the other hand, all the distinct edges e_{jl} with $C_j, C_l \in \mathcal{C}_s^R$ are contained in $E(H_s)$, and so we get the estimate

$$\binom{k_s}{2} \le e(H_s) \le dRk_s$$

for all $s \ge 1$. This clearly contradicts the fact that k_s can grow arbitrarily large. This concludes the proof of the theorem.

The given examples for digraphs and multi-graphs with small complete coloring parameters but without small feedback vertex sets are based on very dense (di)graphs. However, for many investigations, *minor-closed* classes of graphs such as planar graphs, which are rather sparse, are also important. In the following we give the proof of Theorem 11.3 and thereby show that for orientations of graphs in a fixed non-trivial minor-closed class, also for digraphs it is possible to establish an upper bound on the feedback vertex set in terms of the adichromatic number. Moreover, we give more explicit bounds for undirected graphs within such a class. To prove the first part of Theorem 11.3, we need a directed version of Theorem 11.21. This result is not trivial at all. Before its resolution in [RRST96], it was known as Younger's Conjecture. No good (polynomial) upper bounds on the function g are known as of today. In this case, $\nu(D)$ denotes the maximal size of a collection of pairwise vertex-disjoint *directed* cycles in D. **Theorem 11.22** ([RRST96]). There exists a function $g : \mathbb{N} \to \mathbb{N}$ such that for any digraph D, we have

$$\tau(D) \le g(\nu(D)).$$

Proof of Theorem 11.3. We start by noting that the graphs in \mathcal{G} have bounded chromatic number: Since \mathcal{G} is non-trivial, there is a graph $H \notin \mathcal{G}$ which is a forbidden minor for all members of \mathcal{G} . Therefore, all graphs in \mathcal{G} are $K_{v(H)}$ -minor free. By a classical result of Mader ([Mad67]), these graphs have bounded degeneracy and therefore bounded chromatic number. In the following, let d > 0 denote a constant such that $\chi(G) \leq d$ for all $G \in \mathcal{G}$. It follows from the estimate $\frac{v(G)}{\alpha(G)} \leq \chi(G)$ that $\alpha(G) \geq \frac{1}{d}n$ for all $G \in \mathcal{G}$ on n vertices.

1. Let $g: \mathbb{N} \to \mathbb{N}$ be the function from Theorem 11.22, and let D be a given digraph whose simplified underlying graph is in \mathcal{G} . Let $\nu(D) = k$ and let $\{C_1, \ldots, C_k\}$ be a largest collection of vertex-disjoint and w.l.o.g induced directed cycles. Consider the simple model-graph M which has k vertices, one for each cycle C_i , and an edge between two vertices if the corresponding cycles are connected by an edge. Because the cycles C_i are all induced, M is obtained from the simplified underlying graph of D by first deleting all the vertices not on any of the cycles and then contracting the cycles into vertices. These are graph minor operations, and therefore we have $M \in \mathcal{G}$. We conclude that $\alpha(M) \geq \frac{1}{d}k$. Let $I \subseteq V(M)$ be an independent set of size $|I| \geq \frac{1}{d}k$ in M and consider the subdigraph D' of D induced by the union of the vertex sets of cycles $C_i \in \mathcal{C}$ corresponding to vertices in I. We know that D' is the disjoint union of at least $\frac{1}{d}k$ directed cycles. Therefore, if $\frac{1}{d}k \geq \binom{\operatorname{adi}(D')+1}{2}$, then D' contains an induced subdigraph which is the disjoint union of $\binom{\operatorname{adi}(D')+1}{2}$ directed cycles, and then part 4 of Proposition 11.7 and Corollary 11.12 would imply that $\operatorname{adi}(D') \geq \operatorname{adi}(D') + 1$, a contradiction. Hence, $\frac{1}{d}k < \binom{\operatorname{adi}(D')+1}{2} \leq \binom{\operatorname{adi}(D)+1}{2}$. We finally conclude that (w.l.o.g. assuming g to be monotone)

$$\tau(D) \le g(\nu(D)) \le g\left(d\binom{\operatorname{adi}(D)+1}{2}\right) =: f(\operatorname{adi}(D)),$$

which proves the claim.

2. The proof works completely analogous to the directed case, and we obtain the estimate

$$\nu(G) < d\left(\frac{\operatorname{ava}(G) + 1}{2}\right) \le d \cdot \operatorname{ava}(G)^2$$

Finally, this implies using Theorem 11.21 that

$$\tau(G) \le c \cdot (\nu(G) + 1) \log(\nu(G) + 1) \le C \cdot \operatorname{ava}(G)^2 \log(\operatorname{ava}(G)),$$

where C > 0 is a constant which only depends on \mathcal{G} .

11.3.1 Interplay of ava and adi

A consequence of Theorem 11.2 is the following one-sided relationship between the adichromatic number of an oriented digraph and the a-vertex arboricity of its underlying graph.

Corollary 11.23. There exists a function $h_1 : \mathbb{N} \to \mathbb{N}$ such that for any oriented digraph D with underlying graph G, we have

$$\operatorname{adi}(D) \le h_1(\operatorname{ava}(G)).$$

Proof. This follows directly from Proposition 11.5 (applied to D) and Theorem 11.2, as we have

$$adi(D) \le \tau(D) + 1 \le \tau(G) + 1 \le f(ava(G)) + 1.$$

The above estimate cannot be reversed when looking at a fixed digraph (consider transitive tournaments). However, if a graph has large a-vertex arboricity, then it is possible to find an orientation of G with large adichromatic number.

Proposition 11.24. There is a function $h_2 : \mathbb{N} \to \mathbb{N}$ such that

$$\operatorname{ava}(G) \le h_2\left(\max_{D \in \mathcal{O}(G)} \operatorname{adi}(D)\right)$$

for every graph G, where $\mathcal{O}(G)$ denotes the set of digraphs whose underlying graph is G.

Proof. Because $\operatorname{ava}(G)$ is bounded from above in terms of $\tau(G)$, which again (by Theorem 11.21) is bounded from above by a function of $\nu(G)$, it suffices to prove that graphs with a sufficiently large cycle packing have an orientation with large adichromatic number. So let $k \in \mathbb{N}$ be arbitrary and let G be a graph with $\nu(G) \geq \binom{k}{2}$. Let us consider a cycle packing $\mathcal{C} = C_1, \ldots, C_{\binom{k}{2}}$ of induced cycles. Let D be an orientation of G such that:

- all cycles in \mathcal{C} are made directed,
- all edges with endpoints in $V(C_i)$ and $V(C_j)$ are oriented towards C_j if i < j,
- all edges in G with one endpoint on a cycle C_i and the other outside the cycle-packing are directed outwards, and
- $D \bigcup_{C \in \mathcal{C}} V(C)$ is oriented acyclically.

Then every edge of G not contained in one of the cycles C_i is contained in a directed cut in D, and we may conclude from Observation 11.13 that $\operatorname{adi}(D) = \operatorname{adi}(D')$, where D' is obtained from D by deleting all arcs not contained in any of the cycles in \mathcal{C} . As D' is the disjoint union of $\binom{k}{2}$ directed cycles and possibly some isolated vertices (which do not affect the adichromatic number), we conclude from Proposition 11.7 that $\operatorname{adi}(D') \geq k$. As k was arbitrary, this proves the claim. \Box

11.4 Conclusion

In this final section we generalize some of the results established in the previous sections for ava and adi to a more general class of complete coloring parameters defined by forbidden monochromatic patterns. Moreover, we take a more general point of view by relating our investigations to the Erdős-Pósa property and introducing a novel notion which we call τ -boundedness.

Many of the results obtained in this chapter fit into the following more general setting. Let \mathcal{H} be a class of guest (di)graphs, then for a host (di)graph define $\tau_{\mathcal{H}}(G)$ to be the size of a minimum $F \subseteq V$ such that G - F is \mathcal{H} -free, that is, it contains no element of \mathcal{H} as an induced sub(di)graph. Moreover, denote by $\nu_{\mathcal{H}}(G)$ the size of a largest packing of elements of \mathcal{H} as induced sub(di)graphs in G. While clearly $\nu_{\mathcal{H}}(G) \leq \tau_{\mathcal{H}}(G)$ for all (di)graphs G, recall from Chapter 2 that a guest class \mathcal{H} has the Erdős-Pósa property within a class of host (di)graphs \mathcal{G} if there is a function f such that $\tau_{\mathcal{H}}(G) \leq f(\nu_{\mathcal{H}}(G))$ for all $G \in \mathcal{G}$. Define an \mathcal{H} -coloring of G to be a partition of V(G) into sets that induce \mathcal{H} -free (di)graphs. Call an \mathcal{H} -coloring complete if the union of any two color classes contains a member of \mathcal{H} as induced sub(di)graph. The \mathcal{H} -chromatic number $\chi_{\mathcal{H}}(G)$ and the \mathcal{H} -achromatic number $\Psi_{\mathcal{H}}(G)$ are the smallest (the largest, respectively) number of colors that can be used in a complete \mathcal{H} -coloring of G. With a completely analogous proof to Theorem 11.1 we get the following more general interpolation theorem.

Theorem 11.25. Let G be a (di)graph, \mathcal{H} a class of (di)graphs and let $\ell \in \mathbb{N}$. Then there exists a complete \mathcal{H} -coloring of G using exactly ℓ colors if and only if $\chi_{\mathcal{H}}(G) \leq \ell \leq \Psi_{\mathcal{H}}(G)$.

More importantly, the arguments of Proposition 11.5 go through to show:

Proposition 11.26. For any (di)graph G and any class \mathcal{H} of (di)graphs, we have

$$\Psi_{\mathcal{H}}(G) \le \tau_{\mathcal{H}}(G) + 1.$$

Conversely, we say that a host class \mathcal{G} is $\tau_{\mathcal{H}}$ -bounded if there is a function f such that $\tau_{\mathcal{H}}(G) \leq f(\Psi_{\mathcal{H}}(G))$ for all $G \in \mathcal{G}$. By pairing the parts of a maximum complete \mathcal{H} -coloring one obtains a set of $\lfloor \frac{\Psi_{\mathcal{H}}(G)}{2} \rfloor$ disjoint members of \mathcal{H} in G. Thus, $\Psi_{\mathcal{H}}(G) \leq 2\nu_{\mathcal{H}}(G) + 1$ and we get that $\tau_{\mathcal{H}}$ -boundedness is a strengthening of the Erdős-Pósa property:

Proposition 11.27. Let \mathcal{G} and \mathcal{H} be classes of (di)graphs. If \mathcal{G} is $\tau_{\mathcal{H}}$ -bounded, then \mathcal{H} has the Erdős-Pósa property within \mathcal{G} .

A classical result for the achromatic number states that the size of a minimum vertex cover is bounded in terms of the achromatic number of a graph ([FHHM86]), i.e, the class \mathcal{G} of all graphs is τ_{K_2} -bounded. Theorem 11.2 shows that the class of simple graphs is also $\tau_{\mathcal{C}}$ -bounded with respect to the class \mathcal{C} of cycles, thus in a sense strengthening the classical Erdős-Pósa result [EP65]. Furthermore, Theorem 11.3 can be generalized in a straight-forward way to yield:

Theorem 11.28. Let \mathcal{G} be (the orientations of) a non-trivial minor-closed class of simple undirected graphs, and let \mathcal{H} be a class of weakly connected (di)graphs. If \mathcal{H} has the Erdős-Pósa property within \mathcal{G} , then \mathcal{G} is $\tau_{\mathcal{H}}$ -bounded.

On the other hand, while the class \mathcal{C} of all directed cycles has the Erdős-Pósa property within the class $\vec{\mathcal{G}}$ of all digraphs, see [RRST96], our construction in Proposition 11.18 shows that \mathcal{C}_g , the class of digraphs with directed girth at least g, is not $\tau_{\vec{\mathcal{C}}}$ -bounded, for any fixed $g \geq 2$. Hence, the strengthening of the Erdős-Pósa property claimed in Proposition 11.27 is strict and leaves open finding good lower bounds for the adichromatic number. In general, we believe that τ -boundedness deserves further investigation in particular with respect to the numerous Erdős-Pósa properties that have been studied, see e.g. [Ray] for a dynamic listing of Erdős-Pósa-type results from the literature.

Chapter 12

Majority Colorings

12.1 Introduction

The graphs and digraphs considered in this chapter are simple.

A majority coloring of a digraph D with k colors is an assignment $c : V(D) \rightarrow \{1, \ldots, k\}$ such that for every $v \in V(D)$, we have c(w) = c(v) for at most half of all the out-neighbors $w \in N^+(v)$. This notion of coloring was first introduced and investigated by Kreutzer, Oum, Seymour, van der Zypen, and Wood [KOS⁺17] and was inspired by a theoretical problem related to neural networks raised in 2016 by van der Zypen [vdZ19]. Similar questions concerning splittings of digraphs with degree restrictions have been widely studied in the literature, see for instance [Alo96, Alo06, NBJB20].

The main result obtained by Kreutzer et al. shows that every digraph has a majority 4-coloring. Their elegant argument is based on the observation that every acyclic digraph can be majority 2-colored. The relevant property of an acyclic digraph is that there is an ordering of its vertices, in which every vertex is preceded by its complete out-neighborhood. Then coloring vertices along this ordering with two colors such that each vertex is assigned the color that appears least frequently in the (already colored) out-neighborhood will produce a majority 2-coloring.

It is easy to construct digraphs which require three colors for a majority coloring. The canonical examples are the odd directed cycles $\vec{C}_{2k+1}, k \geq 1$, which are not majority 2-colorable since for digraphs with maximum out-degree 1 majority-coloring and proper graph coloring of the underlying graph are equivalent. As of today, however, no digraphs are known that require the use of four colors. Kreutzer et al. claimed that there are none.

Conjecture 12.1 ([KOS⁺17]). Every digraph is majority 3-colorable.

Kreutzer et al. [KOS⁺17] also provide ample evidence for their conjecture by establishing that it holds for "most" digraphs. They show, using the Lovász Local Lemma, that the uniform random 3-coloring is a majority 3-coloring with non-zero probability if certain local density conditions hold, namely if

- $\delta^+(D) > 72 \log(3v(D))$, or
- $\delta^+(D) \ge 1200$ and $\Delta^-(D) \le \frac{\exp(\delta^+(D)/72)}{12\delta^+(D)}$.

In [KOS⁺17] it is also mentioned at the end that a more careful analysis of the Local Lemma approach works for r-regular digraphs provided $r \geq 144$. Subsequently, Girão, Kittipassorn, and Popielarz [GKP17] studied tournaments in particular, and showed, also

using the probabilistic method, that every tournament with minimum out-degree at least 55 is majority 3-colorable.

These are all the results I am aware of about Conjecture 12.1. All the proofs use the Local Lemma for a random coloring and hence require some upper bound on the maximum in-degree in terms of the minimum out-degree (in order to control the number of "bad" events that are adjacent to any fixed bad event in some dependency graph of the events). As it is the case in many related open problems on splitting/coloring digraphs with large minimum out-degree [Alo06, YBWW18, AL89, BHL⁺17], large maximum in-degrees seem to be outside the realm of any such probabilistic approach and it looks like it constitutes the main difficulty of the problem. This is also illustrated by the fact that it was not even known whether planar digraphs are majority 3-colorable.

In this chapter our main motivation is to complement the existing results on digraphs with balanced in- and out-degrees, and provide approaches for natural, broad families of digraphs, without any restrictions of balancedness or on the maximum in-degree.

12.1.1 Our Results

Majority 3-Colorability

Since a proper coloring of the undirected graph underlying a digraph is also a majority coloring, Conjecture 12.1 is immediately true for digraphs with chromatic number at most three. For 4-chromatic digraphs this is already not obvious. Our first result resolves the conjecture for digraphs with bounded chromatic number, including planar digraphs (in the following we denote $\chi(D) := \chi(U(D))$ for every digraph D).

Theorem 12.1. Let D be a digraph such that $\chi(D) \leq 6$. Then D is majority 3-colorable.

The most commonly used digraph coloring concept capturing the orientations of arcs is the dichromatic number, which we have investigated thoroughly in Part II of this thesis.

In the introduction above we mentioned how to give a majority 2-coloring of acyclic digraphs, i.e., digraphs with dichromatic number 1. In our second main result we prove Conjecture 12.1 for digraphs with dichromatic number at most three.

Theorem 12.2. Let D be a digraph such that $\vec{\chi}(D) \leq 3$. Then D is majority 3-colorable.

The Proofs and Majority List Coloring

For our proofs it will be crucial to work in a more general framework, involving the list coloring version of majority coloring. This allows us to formulate appropriately loaded inductive statements from which our theorems follow.

The notion of majority choosability of digraphs was first proposed in [KOS⁺17]. For an assignment $L(\cdot)$ of finite sets L(v) of colors to each vertex $v \in V(D)$, we call a mapping $f : V(D) \to \bigcup_{v \in v(D)} L(v)$ an *L*-coloring if $f(v) \in L(v)$ for every $v \in V(D)$. When L(v) = [k] for every $v \in V(D)$, then *L*-coloring and *k*-coloring coincide. We call a digraph majority *k*-choosable if for every *k*-list assignment (i.e., assignment *L* with |L(v)| = kfor every $v \in V(D)$) there is a majority *L*-coloring. Hence, majority *k*-choosability is a stronger property than majority *k*-colorability.

It was noted in [KOS⁺17] that all the results about dense digraphs using the Local Lemma remain valid for majority 3-choosability (instead of majority 3-colorability). Moreover, Anholcer, Bosek, and Grytczuk [ABG17] gave a beautiful proof to show that every digraph is majority 4-choosable (not only majority 4-colorable).

The following theorem is at the heart of our proofs and is interesting in its own right.

Theorem 12.3. If there is a 2-list assignment $L(\cdot)$ of D such that there is no odd directed cycle in D all whose vertices are assigned the same list, then D has a majority L-coloring.

This statement has several nice consequences, some immediate, some less so. We collect these in the next subsection.

Consequences for majority 3- and 2-colorings

We start by stating choosability analogues of our first two theorems. The analogue of Theorem 12.1 connects the choosability of the underlying graph to majority choosability.

Theorem 12.4. Let D be a digraph whose underlying undirected graph is 6-choosable. Then D is majority 3-choosable. In particular any digraph with a 5-degenerate underlying graph is majority 3-choosable.

Recall from Chapter 8 that the *list dichromatic number* $\vec{\chi}_{\ell}(D)$ of a digraph D is defined as the minimum integer $k \geq 1$ such that for any k-list assignment, we can choose colors from the respective lists without producing monochromatic directed cycles. We have the following analogue of Theorem 12.2 involving this parameter.

Theorem 12.5. Let D be a digraph with $\vec{\chi}_{\ell}(D) \leq 3$. Then D is majority 3-choosable.

The results of [KOS⁺17] and [GKP17] cited in the introduction indicate that the case of *r*-regular digraphs for constant *r* constitute an important benchmark in the study of Conjecture 12.1. Recall in particular that the Local Lemma approach works for *r*-regular digraphs provided $r \ge 144$. Next we obtain conditions at the other end of the local density spectrum, which imply that *r*-regular digraphs are majority 3-colorable for $r \le 4$.

Note first the crucial non-monotonicity in the problem: even though we do not know whether Conjecture 12.1 is true for r = 143, it does hold (quite easily) for r = 1 and 2. Indeed, a 1-regular digraph is the disjoint union of directed cycles, and hence we can 3-color it properly to obtain a majority-coloring. Then Conjecture 12.1 also follows for 2regular digraphs. Even more generally, the validity of the conjecture for any odd regularity r - 1 implies it for the next even regularity r. This is the consequence of the fact¹ that for even r any r-regular digraph D contains a 1-regular spanning subdigraph F and any majority 3-coloring of the (r - 1)-regular digraph D - A(F) is also a majority coloring of D. Most generally, if a digraph D' is obtained from a digraph D by adding an arc (u, v)whose tail has odd out-degree $d_D^+(u)$, then any majority coloring of D is also a majority coloring of D'.

From our next result it follows that 3- and 4-regular digraphs are majority 3-choosable and hence Conjecture 12.1 holds for them as well. We use the following notation: Given a digraph D, we denote by $\Delta_s(D)$ the maximum degree of the simple graph obtained from U(D) by ignoring parallel edges.

Theorem 12.6. Each of the following conditions implies that D is majority 3-choosable:

- $\Delta^+(D) \le 4$,
- $\Delta_s(D) \le 6$,
- $\Delta(D) \leq 7$.

¹To see this, consider the undirected bipartite graph obtained from D by splitting every vertex v into two vertices v^+, v^- and adding an edge u^+v^- for every arc $(u, v) \in A(D)$. Then this bipartite graph is r-regular and hence has a perfect matching, which yields a 1-regular spanning subdigraph of D.

An open question posed in [KOS⁺17] asked whether there is a characterization of digraphs that have a majority 2-coloring (or a polynomial time algorithm to recognize such digraphs). This was answered (most likely) in the negative by Bang-Jensen, Bessy, Havet, and Yeo [BJBHY18] who showed that deciding whether a 3-out-regular digraph is majority 2-colorable is NP-complete. With no hope for an efficient characterization of majority 2-colorability, any simple sufficient condition comes in handy.

For a condition, it is natural to exclude odd directed cycles, as they are canonical examples of graphs with no majority 2-coloring. It turns out that excluding them already implies majority 2-choosability.

Theorem 12.7. Any digraph without odd directed cycles is majority 2-choosable.

Fractional Majority Colorings

The concept of fractional majority coloring emerges as the natural LP-relaxation of the problem of majority coloring, much in the same way as the usual fractional colorings of graphs. This notion was first introduced in $[KOS^+17]$. The definition is somewhat technical and we postpone it to Section 12.4. To appreciate our results here, it is sufficient to keep in mind that the minimum total weight of a fractional majority coloring is at most the majority chromatic number.

Kreutzer et al. $[KOS^+17]$ ask what is the smallest constant K such that every digraph admits a fractional majority coloring with total weight at most K. This is yet another direction to approach Conjecture 12.1 from. Proving that there is a fractional majority coloring with total weight 3 for every digraph would certainly be an easier task. Here we take the first step in this direction and show that the upper bound of 4, which follows from the fact that every digraph is majority 4-colorable, can be slightly improved.

Theorem 12.8. Every digraph D admits a fractional majority coloring with total weight at most 3.9602.

Our proof is the combination of an intricate probabilistic coloring with some deterministic alteration.

In the second theorem of the section we show that digraphs with sufficiently large minimum out-degree have fractional majority colorings with total weight arbitrarily close to 2. This improves the corresponding result in [KOS⁺17] obtained using the Local Lemma, as the upper bound on the maximum in-degree is not necessary here.

Theorem 12.9. There exists a constant C > 0 such that for every $0 < \varepsilon < 1$ and every digraph D with $\delta^+(D) \ge C(1/\varepsilon)^2 \ln(2/\varepsilon)$, there exists a fractional majority coloring of D with total weight at most $2 + \varepsilon$.

Structure of the chapter. In Section 12.2 we obtain Theorem 12.7 as a consequence of a more general result (Theorem 12.3). This result is crucial for the proofs of Theorems 12.1, 12.2, 12.6, 12.4, 12.5, which are presented in Section 12.3. In Section 12.4 we treat fractional majority colorings and prove Theorems 12.8 and 12.9. We conclude with final remarks and some open problems in Section 12.5.

12.2 Digraphs without Odd Directed Cycles

We have seen that acyclic digraphs as well as bipartite digraphs are majority 2-colorable. We have also seen that odd directed cycles are canonical examples of digraphs having no majority 2-coloring. It is therefore natural to try unifying these results and ask whether every digraph without an odd directed cycle is majority 2-colorable. In this section, we answer this question positively. We start with a simple observation:

Lemma 12.10. A digraph D contains no odd directed cycles if and only if all its strong components are bipartite.

Proof. The sufficiency of this condition is obvious, as a directed cycle is always contained in a single strong component. For the reverse direction, it suffices to observe that if Dis strongly connected and all directed cycles have even length, then D is bipartite. By Lemma 5.10 from Chapter 5 applied to the totally cyclic² oriented matroid $\vec{M}(D)$, the arc-set of every oriented cycle in D can be written as the symmetric difference of arc-sets of directed cycles in D, all of which must have even size by assumption. Hence, every cycle in D has even length and therefore D must indeed be bipartite, concluding the proof. \Box

Proposition 12.11. Let D be a digraph which contains no odd directed cycles. Then D is majority 2-colorable. Moreover, any given pre-coloring of the sinks of D can be extended to a majority 2-coloring of D.

Proof. We prove the statement by induction on the number $s \ge 1$ of strong components of D. Suppose first that s = 1, i.e., D is strongly connected. Then by Lemma 12.10 D is bipartite and therefore majority 2-colorable. Since D is either a single vertex or contains no sinks, the claim follows.

Now let $s \ge 2$ and suppose that the statement holds true for all digraphs with at most s-1 strong components. We now distinguish two cases: Either, D is an independent set of s vertices, and therefore, the claim holds trivially true. If there exists at least one arc in D, there has to be a strong component of D containing no sinks such that there are no arcs entering the component. Let X be the vertex set of this component.

Now let a pre-coloring of the sinks of D with 1, 2 be given. By the choice of X, D - X has the same set of sinks as D and s-1 strong components. By the inductive assumption, there exists a majority 2-coloring $c : V(D) \setminus X \to \{1,2\}$ of D - X which extends the pre-coloring of the sinks. By Lemma 12.10, there exists a bipartition $\{A, B\}$ of D[X].

For any subset $W \subseteq X$ equipped with a vertex-coloring $c_W : V(D) \setminus W \to \{1, 2\}$ of D - W, any vertex $x \in W$, and any $i \in \{1, 2\}$, denote by $d(c_W, i, x)$ the number of out-neighbors of x which lie in $V(D) \setminus W$ and have color i under c_W .

We now claim that there exists a subset $U \subseteq X$ and a $\{1,2\}$ -coloring c_U of D - U which extends c, such that

- Every vertex $x \in V(D) \setminus U$ has at least $\frac{1}{2}d^+(x)$ out-neighbors in $V(D) \setminus U$ with a color different from $c_U(x)$.
- Every vertex $x \in U$ fulfills $\max\{d(c_U, 1, x), d(c_U, 2, x)\} < \frac{1}{2}d^+(x)$.

In order to find such a set, we apply the following procedure:

We keep track of a pair (W, c_W) , consisting of a subset $W \subseteq X$ and a vertex-coloring $c_W : V(D) \setminus W \to \{1, 2\}$ extending c. As an invariant, we will keep the first of the two

²Since D is strongly connected, every arc of D is contained in a directed cycle. Hence, $\vec{M}(D)$ is totally cyclic.

above properties, i.e. we ensure that every vertex $x \in V(D) \setminus W$ has at least $\frac{1}{2}d^+(x)$ out-neighbors with a different color according to c_W .

We initialize W := X, $c_W := c$. It is clear that this assignment satisfies the invariant (since c is a majority coloring of D - X and since there are no arcs entering X).

As long as a vertex $x_0 \in W$ with $\max\{d(c_W, 1, x_0), d(c_W, 2, x_0)\} \geq \frac{1}{2}d^+(x_0)$ exists, we choose such a vertex. We put $W' := W \setminus \{x_0\}$, and define a coloring $c_{W'}$ of D - W' according to

$$c_{W'}(x) := \begin{cases} c_W(x), & \text{if } x \neq x_0, \\ 1, & \text{if } x = x_0, \ d(c_W, 1, x_0) < d(c_W, 2, x_0), \\ 2, & \text{if } x = x_0, \ d(c_W, 1, x_0) \ge d(c_W, 2, x_0). \end{cases}$$

It is easily verified that the coloring $c_{W'}$ also fulfils the invariant, since by definition x_0 has at least $\max\{d(c_W, 1, x_0), d(c_W, 2, x_0)\} \ge \frac{1}{2}d^+(x_0)$ out-neighbors in D-W' of different color. Furthermore, for every vertex $x \in W'$, the number of out-neighbors of different color does not decrease by coloring x_0 .

Finally we update according to $(W, c_W) := (W', c_{W'})$.

In the moment the procedure terminates, we have found a subset $U := W \subseteq X$ and a $\{1,2\}$ -coloring c_U of D-U extending c with the property that every vertex $x \in V(D) \setminus U$ has at least $\frac{d^+(x)}{2}$ out-neighbors with different color according to c_U . Since the procedure terminated, we furthermore have $\max\{d(c_U, 1, x), d(c_U, 2, x)\} < \frac{1}{2}d^+(x)$ for every vertex $x \in U$. This shows that U satisfies both of the conditions stated above.

We now finally extend the coloring c_U of $V(D) \setminus U$ to a $\{1,2\}$ -coloring of D by giving color 1 to each vertex in $A \cap U$ and color 2 to every vertex in $B \cap U$. This coloring extends c and therefore the initial pre-coloring of the sinks, and is a majority coloring: By the first of the two conditions, every vertex $x \in V(D) \setminus U$ has at least $\frac{d^+(x)}{2}$ out-neighbors with a different color. For each vertex $x \in U$, since $\{A, B\}$ is a bipartition of D[X], all out-neighbors in U have a different color, and among the out-neighbors in D-X, at most $\max\{d(c_U, 1, x), d(c_U, 2, x)\} < \frac{1}{2}d^+(x)$ can share its color. Therefore every vertex satifies the condition for a majority-coloring, and this concludes the proof of the claim. \Box

We are now ready for the proof of Theorem 12.3.

Proof of Theorem 12.3. We may assume w.l.o.g. that color lists of adjacent vertices always intersect: Otherwise, we remove all edges between vertices with disjoint color lists to obtain a digraph D'. Any majority-coloring of D' with colors chosen from the lists will also be a majority-coloring of D.

For a pair $\{a, b\}$ of colors let us denote $X_{\{a,b\}} := \{x \in V(D) | L(x) = \{a,b\}\}$. By assumption $D[X_{\{a,b\}}]$ contains no odd directed cycles. Let $D'_{\{a,b\}}$ be the digraph obtained from $D[X_{\{a,b\}}]$ by adding all arcs $(x, y) \in A(D)$ with $x \in X_{\{a,b\}}$ and $y \notin X_{\{a,b\}}$ and their endpoints. Since we only add sinks to $D[X_{\{a,b\}}]$, also $D'_{\{a,b\}}$ contains no odd directed cycles. For each vertex $y \in N^+(X_{\{a,b\}})$, there is a unique color $p_{\{a,b\}}(y)$ in $L(y) \cap \{a,b\}$. Pre-color the sinks of $D'_{\{a,b\}}$ in such a way that every vertex $y \in N^+(X_{\{a,b\}})$ receives color $p_{\{a,b\}}(y)$. By Proposition 12.11 we can now find a majority-coloring $c_{\{a,b\}}$ of $D'_{\{a,b\}}$ extending this pre-coloring with colors a and b.

Now define a coloring c of all vertices in D by setting $c(x) := c_{\{a,b\}}(x)$ if $L(x) = \{a,b\}$. Clearly, we have $c(x) \in L(x)$ for all $x \in V(D)$. We claim that c is a majority-coloring of D. Indeed, for any vertex $x \in V(D)$, if $L(x) = \{a,b\}$, then we have $N^+(x) = N^+_{D'_{\{a,b\}}}(x)$, and $\{y \in N^+(x) \mid c(y) = c(x)\} \subseteq \{y \in N^+_{D'_{\{a,b\}}}(x) \mid c_{\{a,b\}}(x) = c_{\{a,b\}}(y)\}$. Hence, at most half of the out-neighbors of x share its color, and the claim follows. Theorem 12.7 is now obtained from Theorem 12.3 as a direct consequence.

12.3 Majority 3-Colorings of Sparse Digraphs

As a consequence of Theorem 12.3, we obtain our main result:

Theorem 12.12. Let D be a digraph. Suppose there is a partition $\{X_1, X_2, X_3\}$ of the vertex set such that for every $i \in \{1, 2, 3\}$, $D[X_i]$ contains no odd directed cycles. Then D is majority 3-colorable.

Proof. We assign lists of size two to the vertices of D, namely, we assign the list $\{2, 3\}$ to all vertices in X_1 , the list $\{1, 3\}$ to all vertices in X_2 , and the list $\{1, 2\}$ to all vertices in X_3 . Because $D[X_i]$, i = 1, 2, 3 contains no odd directed cycle, we can apply Theorem 12.3 to conclude that there exists a majority-coloring of D which only uses the colors 1, 2 and 3. This proves the claim.

From this we now directly derive Theorems 12.1 and 12.2.

Proof of Theorem 12.1. If $\chi(D) \leq 6$, then D admits a partition Y_1, \ldots, Y_6 into independent sets. Using the partition $\{Y_1 \cup Y_2, Y_3 \cup Y_4, Y_5 \cup Y_6\}$ of the vertex set to apply Theorem 12.12 now shows that D is indeed majority 3-colorable.

Proof of Theorem 12.2. If $\vec{\chi}(D) \leq 3$, then there exists a partition $\{X_1, X_2, X_3\}$ of the vertex set such that $D[X_i]$ contains no directed cycles, for i = 1, 2, 3. The claim now follows by Theorem 12.12.

The fact that Theorem 12.3 deals with an assignment of lists can be further exploited to show analogues of Theorem 12.12, Theorems 12.1 and 12.2 for list colorings.

For this purpose we need the following notion: Call a digraph D OD-3-choosable if for any assignment of color lists L(x), $x \in V(D)$ of size 3 to the vertices, there exists a choice function c (i.e. $c(x) \in L(x)$ for all $x \in V(D)$) such that no odd directed cycle in (D, c) is monochromatic.

Theorem 12.13. If a digraph is OD-3-choosable, then it is also majority 3-choosable.

Proof. Let L(v) for all $v \in V(D)$ be a given color list of size three. We have to show that there is a majority-coloring c of D such that $c(v) \in L(v)$ for all $v \in V(D)$. For every $v \in V(D)$, we let $L^*(v) := \{\{C_1, C_2\} | C_1 \neq C_2 \in L(v)\}$ contain all three unordered color-pairs in L(v). Since D is OD-3-choosable, there exists a choice function c^* on V(D)such that $c^*(v) \in L^*(v)$ for each vertex $v \in V(D)$ is a subset of L(v) of size two and such that there exists no odd directed cycle in D which is monochromatic with respect to c^* . If we now consider $c^*(v), v \in V(D)$ as a 2-list assignment of D, we can apply Theorem 12.3 to conclude that there is a majority-coloring c of D such that $c(v) \in c^*(v) \subseteq L(v)$ for every vertex $v \in V(D)$. As $L(\cdot)$ was arbitrary, we conclude that D is majority 3-choosable. \Box

We are now ready to prove Theorems 12.4 and 12.5.

Proof of Theorem 12.4. We show that D is OD-3-choosable, the claim then follows by Theorem 12.13. Let L(v) for each vertex $v \in V(D)$ be an assigned list of three colors. For each color C used in one of the lists, let C' be a distinct copy of this color. We now consider the assignment $L_6(\cdot)$ of lists of size 6 to the vertices of D, where for each vertex $v \in V(D), L_6(v) := \{C_1, C'_1, C_2, C'_2, C_3, C'_3\}$ if C_1, C_2, C_3 denote the colors contained in L(v). Because the underlying graph of D is 6-choosable, there is a proper coloring c_6 of D such that $c_6(v) \in L_6(v)$ for all $v \in V(D)$. Now consider the coloring c of D obtained from c_6 by identifying each copy C' of an original color C with C again. We then have $c(v) \in L(v)$ for every $v \in V(D)$. Since c_6 was a proper coloring of the undirected underlying graph of D, each color class with respect to c induces a bipartite subdigraph of D, and hence there are no monochromatic odd directed cycles in (D, c). Hence, D is OD-3-choosable.

Proof of Theorem 12.5. This follows directly since any digraph with $\vec{\chi}_{\ell}(D) \leq 3$ is clearly OD-3-choosable.

The rest of this section is devoted to proving Theorem 12.6. The proof uses the following Lemma, which in turn uses Theorems 12.4 and 12.5.

Lemma 12.14. Let D be a digraph such that $\min\{d^+(x), d^-(x)+1\} \le 3$ for all $x \in V(D)$. Then D is OD-3-choosable.

Proof. Suppose the claim was false. Let us consider a counterexample D minimizing v(D)+a(D). We have $v(D) \ge 4$, D is connected and we know that every proper subdigraph of D must be OD-3-choosable.

We first consider the case that there is a vertex v with $d^-(v) \leq 2$. Since D - v is OD-3-choosable, given any assignment $L(v), v \in V(D)$ of lists of size at least 3 to the vertices, we can choose colors $c(w) \in L(w)$ from the lists for every $w \in V(D) \setminus \{v\}$ such that in D - v, there exists no monochromatic odd directed cycle. Now assign to v a color $c(v) \in L(v) \setminus \{c(w) \mid w \in N^-(v)\}$. We claim that c is a coloring of D without monochromatic odd directed cycles. In fact, such a cycle would have to pass v, however no edge entering v is monochromatic. Therefore D is OD-3-choosable, a contradiction.

Hence we know for every $x \in V(D)$ that $d^{-}(x) \ge 3$. Since $\min\{d^{+}(x), d^{-}(x) + 1\} \le 3$, we also must have $d^{+}(x) \le 3$. We conclude

$$3v(D) \le \sum_{v \in V(D)} d^{-}(x) = a(D) = \sum_{v \in V(D)} d^{+}(x) \le 3v(D)$$

and thus we have $d^+(x) = d^-(x) = 3$ for all $x \in V(D)$. Consequently, the underlying simple graph U(D) has maximum degree $\Delta(U(D)) \leq 6$. If U(D) is 6-choosable, then it follows as in the proof of Theorem 12.4 that D is OD-3-choosable, a contradiction.

Therefore, by the list coloring version of Brooks' Theorem [Viz76], we must have $U(D) = K_7$. Since D is 3-out- and 3-in-regular, it follows that D is a tournament on 7 vertices. However, every tournament on 7 vertices has list dichromatic number at most 3 and is therefore OD-3-choosable according to Theorem 12.5. This can be seen using two results from [BHL18]. Clearly, we have $\vec{\chi}(D) \leq 3$. If $\vec{\chi}(D) = 3$, then we have $v(D) = 7 \leq 2\vec{\chi}(D) + 1$, and by Theorem 2.2 in [BHL18] we conclude that $\vec{\chi}_{\ell}(D) = \vec{\chi}(D) = 3$. Otherwise, we have $\vec{\chi}(D) \leq 2$. In this case, we can apply Theorem 3.3 in [BHL18] to conclude $\vec{\chi}_{\ell}(D) \leq 2\log(7) < 4$. Therefore we have $\vec{\chi}_{\ell}(D) \leq 3$ in each case.

Finally, since we obtained that D is OD-3-choosable in each case, the initial assumption was wrong, which concludes the proof by contradiction.

Corollary 12.15. Let D be a digraph with $\min\{d^+(x), d^-(x)+2\} \le 4$ for every $x \in V(D)$. Then D is majority 3-choosable.

Proof. For a proof by contradiction, suppose the claim was false and consider a counterexample D minimizing the number of edges.

Consider first the case that there is a $v \in V(D)$ with $d^+(v) = 4$. Let e be an arc leaving v and put D' := D - e. By the minimality of D, D' is majority 3-choosable. We now claim
that any majority-coloring of D' also defines a majority-coloring of D. Clearly, such a coloring satisfies the condition for a majority-coloring at any vertex distinct from v. Since v has out-degree 3 in D', it has at most one out-neighbor in D' of the same color. Thus there are at most two out-neighbors of v in D which share its color, and so the majority condition is fulfilled at v. We conclude that also D must be majority 3-choosable, which gives the desired contradiction.

Now for the second case, assume that no vertex has out-degree 4. This means that for every $x \in V(D)$, we either have $d^+(x) \leq 3$ or $d^+(x) \geq 5$ and therefore $d^-(x) \leq 2$. We can therefore apply Lemma 12.14 to D, which shows that D is OD-3-choosable. From Theorem 12.13 we get that D is majority 3-choosable. This again is a contradiction to Dbeing a counterexample to the claim.

Therefore the initial assumption was wrong, and this concludes the proof.

Proof of Theorem 12.6. If $\Delta^+(D) \leq 4$ or $\Delta(D) \leq 7$, then the claim follows by applying Corollary 12.15. If $\Delta(U(D)) \leq 6$, then by the list coloring version of Brook's Theorem either U(D) is 6-choosable, and then the claim follows from Theorem 12.4, or $U(D) = K_7$. Now let $L(v_1), \ldots, L(v_7)$ be lists of size three assigned to the vertices $\{v_1, \ldots, v_7\}$ of D.

We first consider the case that all lists are equal, i.e., show that D is majority 3-colorable.

If there exists a vertex $v \in V(D)$ which is contained in at most 3 digons, then there are vertices $u_1 \neq u_2 \in V(D) \setminus \{v\}$ such that u_1, u_2, v do not form a directed triangle. Therefore, any partition $\{X_1, X_2, X_3\}$ of V(D) where $X_1 = \{v, u_1, u_2\}$ and $|X_2| = |X_3|$ shows, by Theorem 12.12, that D is majority 3-colorable. Otherwise, every vertex in D is contained in at least 4 digons and thus has out-degree at least 4. Now any 3-coloring of D with color classes of sizes 2, 2, 3 defines a majority-coloring of D.

Now suppose that not all lists are equal. In this case we can choose for each vertex v_i a sublist $L_2(v_i) \subseteq L(v_i)$ of size two such that no three vertices are assigned the same sublist (minimize the number of edges whose ends are assigned the same sublist). By Theorem 12.3 we obtain a majority-coloring c of D where $c(v_i) \in L_2(v_i) \subseteq L(v_i)$ for every $i \in [7]$. Hence, D is majority 3-choosable in each case, which concludes the proof.

12.4 Fractional Majority Colorings

Another concept introduced in [KOS⁺17] is that of a fractional majority coloring. Fractional majority colorings represent a linear relaxation of majority colorings, in the same way that the well-known fractional chromatic number (which we well discuss in detail in the next chapter) forms a linear relaxation of the chromatic number.

Given a subset $S \subseteq V(D)$, a vertex v is popular in S if $v \in S$ and more than half of its out-neighbors are in S. A subset $S \subseteq V(D)$ is stable if it contains no popular vertices. Let S(D) be the set of all stable sets of D, and S(D, v) the set of all stable sets containing v.

A fractional majority coloring is a function that assigns a weight $w_T \ge 0$ to every set $T \in S(D)$, satisfying $\sum_{T \in S(D,v)} w_T \ge 1$ for every $v \in V(D)$. The total weight of a fractional majority coloring is simply $\sum_{T \in S(D)} w_T$.

Kreutzer et al. asked for the minimum constant K such that every digraph admits a fractional majority coloring with total weight at most K.

We will show two results related to this question, namely Theorem 12.8 and Theorem 12.9. The proof of these two theorems will be based on the dual of the linear program defined by the restrictions on a fractional majority coloring: **Observation 12.16.** For a digraph D, the minimum possible total weight of a fractional majority coloring equals the maximum total weight $\sum_{v \in V(D)} w_v$ in a non-negative weight assignment of V(D) in which every stable set T satisfies $\sum_{v \in T} w_v \leq 1$.

The main idea of the proof of both theorems is that, given any choice of weights on V(D), we can construct a stable set in which the weight is at least a given fraction of the total weight, using the probabilistic method.

Lemma 12.17. Let D be a digraph and let $0 . Suppose that one can take a random subset <math>X \subseteq V(D)$ with the property that, for every $v \in V(D)$, the probability that v is in X but not popular in X is at least p. Then D admits a fractional majority coloring with total weight at most $\frac{1}{p}$.

Proof. Suppose that D is a counterexample to our statement, and we will reach a contradiction. By Observation 12.16, we can assign weights to V(D) so that the total weight is $w > \frac{1}{p}$, and every stable set in D has a sum of weights at most one. Let Y be the set of popular vertices in X. By linearity of expectation, the expected total weight of $X \setminus Y$ is at least pw > 1.

Take an instance of $X \setminus Y$ with weight greater than 1. Every vertex in $X \setminus Y$ has at least half of its out-neighbors outside of X, which implies that it is not popular in $X \setminus Y$. Hence $X \setminus Y$ is stable in D and has total weight greater than 1, producing a contradiction. \Box

The proof of Theorem 12.9 is a straightforward application of this lemma:

Proof of Theorem 12.9. Let N be a large enough positive integer. Let D be a digraph with $\delta^+(D) \ge N$. Set $p = \frac{1}{2} - \sqrt{\frac{\log N}{N}}$. Let X be a random subset of V(D) in which every element is included independently with probability p. By Hoeffding's inequality, for any vertex v the probability that at least half of its out-neighbors are in X is at most

$$\Pr\left(|X \cap N^+(v)| \ge \frac{1}{2}d^+(v)\right) \le e^{-2\left(\frac{1}{2}-p\right)^2 d^+(v)} \le e^{-2\log N} = N^{-2}$$

Setting $q = N^{-2}$, from Lemma 12.17 we find a fractional majority coloring of total weight at most $\frac{1}{p-q} = 2 + O\left(\sqrt{\frac{\log N}{N}}\right)$. For a large enough absolute constant C > 0, this implies that $\frac{1}{p-q} < 2 + \varepsilon$ whenever $N > C\left(\frac{1}{\varepsilon}\right)^2 \ln\left(\frac{2}{\varepsilon}\right)$.

For Theorem 12.8, we need to be more careful. Consider again the set X containing each vertex independently with probability p, where p is slightly lower than $\frac{1}{2}$. If the out-degree of v is not 1, one can show that the probability that v is popular in X is upper-bounded by a constant, strictly smaller than $p - \frac{1}{4}$. However, if v has out-degree 1, the probability that v is popular in X is $p^2 > p - \frac{1}{4}$. For this reason, the vertices with out-degree 1 deserve extra consideration.

Observe that, in the graph induced by the vertices of out-degree 1, all cycles are directed, pairwise disjoint and act as sink components. Consequently, removing one vertex from each directed cycle produces an acyclic digraph, where the vertices can be given a topological ordering in which every arc goes from a larger vertex to a smaller one.

Proof of Theorem 12.8. Set $p_1 = 0.4594$ and $p_2 = 0.4503$. Assign independently to each vertex v a random indicating variable X_v , which takes the value 1 with probability p_1 if $d^+(v) = 1$ and with probability p_2 otherwise. Now construct a random set X as follows:

• Add to X all vertices v with $d^+(v) \neq 1$ and $X_v = 1$.

- For every cycle C formed by vertices with $d^+(v) = 1$ and $X_v = 1$, select a vertex $v \in C$ uniformly at random and set $X_v = 0$.
- Take an ordering of the vertices v with $d^+(v) = 1$ and $X_v = 1$, in which if we have an arc (v, w) then v comes after w (this is possible because these vertices form an acyclic digraph). Following this order, add v to X if its out-neighbor is not in X.

We will show that, for every vertex v, the probability that v is in X but not popular in X is at least $\frac{1}{4} + \varepsilon$, for a fixed value of $\varepsilon > 0$. Suppose first that $d^+(v) = 1$. Note that if $v \in X$, then the unique out-neighbor of v is not in X and hence v is not popular in X. If we draw the vertices with out-degree 1 as red boxes and those with other out-degrees as blue circles, then the successive out-neighborhoods of v must have one of these forms:



Figure 12.1: The six possible out-neighborhoods of a red vertex. The black vertex at the end of the path in Case 4 can be either red or blue.

We denote $v = v_0$, and v_{i+1} as the out-neighbor of v_i , if it is unique. We go through each case and give a lower bound on the probability that $v \in X$ and v is not popular.

- If v is in Case 1, then whenever $X_v = 1$ and $X_{v_1} = 0$ we have $v \in X$. This happens with probability $p_1(1-p_2)$.
- If v is in Case 2, then whenever $X_v = 1$ and $X_{v_1} = 0$, or whenever X_v, X_{v_1} and X_{v_2} all equal 1, we have $v \in X$. This happens with probability $p_1(1-p_1) + p_1^2 p_2$.
- If v is in Case 3, then whenever $X_v = 1$ and $X_{v_1} = 0$, or if X_v and X_{v_1} initially equal 1 and X_{v_1} is selected to be modified (with respect to the cycle v, v_1), then we have $v \in X$. This happens with probability at least $p_1(1-p_1) + \frac{1}{2}p_1^2$.
- If v is in Case 4, then whenever $X_v = 1$ and $X_{v_1} = 0$, or whenever $X_v = 1$, $X_{v_1} = 1$, $X_{v_2} = 1$ and $X_{v_3} = 0$ we have $v \in X$. This happens with probability at least $p_1(1-p_1) + p_1^3(1-p_1)$.
- If v is in Case 5, if $X_v = 1$ and $X_{v_1} = 0$, or if X_v, X_{v_1} and X_{v_2} all initially equal 1 and X_{v_1} is selected to be modified (with respect to the cycle v_1, v_2), then we have $v \in X$. This happens with probability at least $p_1(1-p_1) + \frac{1}{2}p_1^3$.
- If v is in Case 6, if $X_v = 1$ and $X_{v_1} = 0$, or if X_v , X_{v_1} and X_{v_2} all initially equal 1 and X_{v_1} is selected to be modified (with respect to the cycle v, v_1, v_2), then we have $v \in X$. This happens with probability $p_1(1-p_1) + \frac{1}{3}p_1^3$.

Suppose now that $d^+(v) \neq 1$. The probability that $v \in X$ is p_2 . If v is popular in X, then over half of its out-neighbors w have $X_w = 1$ (this is necessary for $w \in X$). Since the X_w are independent, and each of them takes the value 1 with probability at most p_1 , the probability that v is popular on X, conditioned on $v \in X$, is at most $\Pr\left(B(d^+(v), p_1) > \frac{d^+(v)}{2}\right)$. For $d^+(v) = 3$, this probability is $3p_1^2 - 2p_1^3$. We claim that this is the worst case:

Proposition 12.18. For every $k \neq 1$,

$$\Pr\left(B(k, p_1) > \frac{k}{2}\right) \le \Pr(B(3, p_1) \ge 2).$$

Proof. Consider an infinite sequence X_1, X_2, \ldots of indicating random variables, each taking value 1 independently with probability p_1 . Let I_i be the event "among the first *i* variables more than half take value 1". Then $\Pr(I_k) = \Pr\left(B(k, p_1) > \frac{k}{2}\right)$. Clearly $\Pr(I_0) = 0$. Moreover, if k is even then I_k implies I_{k+1} , so we can restrict ourselves to odd k.

We will prove our statement by induction, by showing that $Pr(I_{2k+1}) < Pr(I_{2k-1})$ for $k \ge 2$. Indeed, the event $I_{2k-1} \setminus I_{2k+1}$ is precisely the case in which exactly k of the first 2k - 1 variables take value 1, and $X_{2k} = X_{2k+1} = 0$. Thus

$$\Pr(I_{2k-1} \setminus I_{2k+1}) = \binom{2k-1}{k} p_1^k (1-p_1)^{k+1}.$$

Similarly, the event $I_{2k+1} \setminus I_{2k-1}$ is precisely the case in which exactly k-1 of the first 2k-1 variables take value 1, and $X_{2k} = X_{2k+1} = 1$. Thus

$$\Pr(I_{2k+1} \setminus I_{2k-1}) = \binom{2k-1}{k-1} p_1^{k+1} (1-p_1)^k$$

Using that the equality $\mathbb{P}(A \setminus B) - \mathbb{P}(B \setminus A) = \mathbb{P}(A) - \mathbb{P}(B)$ holds for any two probability events A, B, we finally conclude that

$$P(I_{2k-1}) - P(I_{2k+1}) = {\binom{2k-1}{k}} p_1^k (1-p_1)^k (1-2p_1) > 0,$$

as required.

With this, we know that for every vertex v the probability that v is in X and not popular in X is at least

$$\min\{p_1(1-p_2), p_1(1-p_1) + p_1^2 p_2, (p_1+p_1^3)(1-p_1), p_1(1-p_1) + \frac{1}{3}p_1^3, p_2(1-3p_1^2+2p_1^3)\} = p_2(1-3p_1^2+2p_1^3) = 0.252513 =: p.$$

Applying Lemma 12.17, there is a fractional majority coloring of D with total weight at most $\frac{1}{p} < 3.9602$.

12.5 Conclusion

Girão et al. [GKP17] and independently Knox and Šámal [KS18] investigated a natural generalization of majority colorings: For any $\alpha \in [0, 1]$, define an α -majority coloring of a digraph D to be a vertex-coloring in which for every vertex v, at most $\alpha \cdot d^+(v)$ vertices in

 $N^+(v)$ have the same color as v. If such a coloring can be found for any ℓ -list-assignment, we call the digraph α -majority ℓ -choosable.

Generalizing the result by Anholcer et al. [ABG17] it was proved both in [GKP17] and [KS18] that for every integer $k \ge 1$, every digraph is $\frac{1}{k}$ -majority 2k-choosable. Girão et al. proposed the following generalization of Conjecture 12.1:

Conjecture 12.2. For every integer $k \ge 1$, every digraph D has a $\frac{1}{k}$ -majority (2k-1)coloring. In fact, every digraph is $\frac{1}{k}$ -majority (2k-1)-choosable.

It is natural to try and generalize the results presented in this chapter for majority colorings with $\alpha = \frac{1}{2}$ to arbitrary values $\alpha \in [0, 1]$. Among our results, we can only generalize a special case of Theorem 12.2, namely for digraphs of dichromatic number 2, we verify the first part of Conjecture 12.2 for all $k \geq 1$.

Proposition 12.19. Let D be a digraph with $\vec{\chi}(D) \leq 2$. Then for every $k \in \mathbb{N}$, $k \geq 2$, D admits a $\frac{1}{k}$ -majority coloring using 2k - 1 colors.

Proof. We use a similar idea as in the proof of Theorem 12.3. Consider first an acyclic digraph F with a pre-coloring of its sinks using colors from $\{1, \ldots, k\}$. We claim that such a coloring can always be extended to a $\frac{1}{k}$ -majority coloring of F also using colors from $\{1, \ldots, k\}$. To find such a coloring, we take a topological ordering x_1, \ldots, x_n of the vertices (i.e. $(x_i, x_j) \notin A(D)$ for all $i \leq j$) such that $\{x_1, \ldots, x_t\}$ are the pre-colored sinks. Now we color the vertices one by one, starting with x_{t+1} , then x_{t+2} etc. When coloring the vertex x_i with i > t, we assign to it a color from $\{1, \ldots, k\}$ appearing least frequently among its (already colored) out-neighbors. This procedure eventually yields a k-coloring of F where any vertex has at most a $\frac{1}{k}$ -fraction of its out-neighbors with the same color.

Now let $\{X_1, X_2\}$ be a partition of V(D) such that $D[X_1]$, $D[X_2]$ are acyclic. For i = 1, 2 let D'_i be the digraph obtained from $D[X_i]$ by adding all arcs in D leaving X_i together with their endpoints. Clearly, also D'_1 and D'_2 are acyclic. By the above observation, D'_i for i = 1, 2 has a majority $\frac{1}{k}$ -coloring c_i with k colors in which all sinks receive color 1. After renaming we may suppose that c_1 uses colors from $\{1, 2, \ldots, k\}$, while c_2 uses colors from $\{1, k + 1, k + 2, \ldots, 2k - 1\}$. We now define a (2k - 1)-coloring of all vertices in D by putting $c(x) := c_i(x)$ for $x \in X_i$. For any vertex $x \in X_i$, we have that $N^+(x) = N^+_{D'_i}(x)$, and, since all vertices in $V(D'_i) \setminus X_i$ received color 1 under c_i , it follows that $\{y \in N^+(x) | c(y) = c(x)\} \subseteq \{y \in N^+_{D'_i}(x) | c_i(y) = c_i(x)\}$. Therefore, and since c_i is a majority $\frac{1}{k}$ -coloring of D'_i , at most a $\frac{1}{k}$ -fraction of vertices in $N^+(x)$ have the same color as x. This shows that c is a coloring as requested and concludes the proof.

It is worth noting that the above bound is tight. Consider for example the circulant digraph³ $\vec{C}(2k-1,k)$ which has as vertex set \mathbb{Z}_{2k-1} , and where we have an edge (i, j) if and only if $j-i \in \{1, 2, 3, ..., k-1\}$. It is easy to see that in any majority $\frac{1}{k}$ -coloring of D, the 2k-1 vertices must receive pairwise distinct colors, however, partitioning the vertex-set into $X_1 = \{0, 1, ..., k-1\}$ and $X_2 = \{k, k+1, ..., 2k-2\}$ shows that $\vec{\chi}(\vec{C}(2k-1,k)) = 2$.

The methods used in this chapter are unlikely to resolve Conjecture 12.1 for the next open cases of 5- and 6-regular digraphs. One possible approach could be via an extension to hypergraphs: Given a 5-regular digraph D, consider the hypergraph $\mathcal{H}(D)$ with vertex set V(D) and whose edges are $\{v\} \cup N^+(v), v \in V(D)$. This hypergraph is 6-regular and 6-uniform⁴. If we could now find a vertex-3-coloring of $\mathcal{H}(D)$ such that no hyperedge

³We will revisit the circulant digraphs in the next Chapter.

⁴That is, every vertex is included in exactly 6 hyperedges and every hyperedge is of size 6.

contains four vertices of the same color, this coloring would certainly be a majority coloring of D. We are therefore interested in deciding the following question.

Problem 12.1. Let H be a 6-regular 6-uniform hypergraph. Is there a 3-coloring of V(H) such that no hyperedge contains four vertices of the same color?

The setting of k-regular k-uniform hypergraphs could be fruitful, as it is known that these hypergraphs have property B for all $k \ge 4$ (as noted in [Vis03])⁵. We conclude with a small selection of open questions.

- Is every 5-regular digraph $\frac{1}{3}$ -majority 5-colorable? There exists an argument showing that it is possible to color with 5 colors such that in each connected component, at most one vertex violates the majority condition.
- Does every digraph with $\chi(D) \leq 6$ have a $\frac{1}{3}$ -majority 5-coloring?
- Does every digraph D with $\vec{\chi}(D) \leq 3$ have a $\frac{1}{k}$ -majority (2k-1)-coloring?

 $^{{}^{5}}A$ hypergraph is said to have property B if its vertex-set can be 2-colored such that every hyperedge includes at least one vertex of each color.

Chapter 13

Complexity of Fractional and Circular Colorings

13.1 Introduction

Graphs and digraphs considered in this chapter are loopless, but may have parallel edges and (anti-)parallel arcs, respectively.

Throughout this chapter, whenever we write a + b or a - b for elements $a, b \in \mathbb{Z}_k$, this is meant as in the group $(\mathbb{Z}_k, +)$. For elements x of \mathbb{Z}_k or \mathbb{Z} we will use $x \mod k$ to denote the unique element within $\{0, \ldots, k-1\} \subseteq \mathbb{Z}$ equivalent to $x \mod k$.

In addition, for every $k \in \mathbb{N}$ and elements $x, y \in \{0, \ldots, k-1\} \simeq \mathbb{Z}_k$, let us denote by $\operatorname{dist}_k(x, y) := |(x - y) \mod k|_k$, where $|a|_k := \min\{|a|, |k - a|\}$ for all $a = 0, \ldots, k-1$, the circular k-distance between x and y.

In this chapter, we are concerned with several notions of acyclic colorings of directed and undirected graphs which, in contrast to the integer coloring parameters treated in previous chapters of this thesis, may take on fractional values. We will be mainly concerned with two different (but related) such coloring concepts, called *circular* and *fractional* colorings. Before we state the main results of this chapter, which establish hardness results for the computation of these parameters, let us give a short introduction of the coloring concepts for (di)graphs considered in this chapter.

Circular colorings of undirected graphs were introduced by Vince [Vin88], where the concept of the *star chromatic number*, nowadays also known as the *circular chromatic number* of a graph, made its first appearance. The original definition of the star chromatic number by Vince is based on so-called (k, d)-colorings, where colors at adjacent vertices are not only required to be distinct as usual but moreover "far apart" in the following sense:

Definition 13.1 (Vince, [Vin88]). Let G be a graph and $(k, d) \in \mathbb{N}^2, k \geq d$. A (k, d)coloring of G is an assignment $c : V(G) \to \{0, \ldots, k-1\} \simeq \mathbb{Z}_k$ of colors to the vertices so
that $dist_k(c(u), c(w)) \geq d$ whenever u, w are adjacent.

The circular chromatic number $\chi_c(G) \geq 1$ of the graph G is defined as the infimum over all values of $\frac{k}{d}$ for which (k, d)-colorings exist.

The most important properties of $\chi_c(G)$ proved in [Vin88] are that $\chi_c(G)$ is always a rational number and that $\lceil \chi_c(G) \rceil = \chi(G)$ for every graph G. Intuitively, one may think that the closer the value of $\chi_c(G)$ is to its lower bound $\chi(G) - 1$, the "closer" the graph G is to being colorable with $\chi(G) - 1$ colors. A good illustrating example are the odd cycles: For every $k \ge 1$, we have $\chi(C_{2k+1}) = 3$, but C_{2k+1} can be 3-colored such that only one

vertex out of the 2k + 1 receives the third color. In that sense one might feel that C_{2k+1} is "almost" 2-colorable for large k. This is reflected in the value $\chi_c(C_{2k+1}) = 2 + \frac{1}{k}$ as well.

The following canonical construction related to the circular chromatic number will be used in Section 13.4.

Definition 13.2. For any given natural numbers $(k, d) \in \mathbb{N}^2$ with $k \geq 2d$, we denote by C(k, d) the circulant graph with vertex set \mathbb{Z}_k where vertices $i \neq j \in \mathbb{Z}_k$ are adjacent if and only if $dist_k(i, j) \geq d$.

As was shown in Theorem 6 of [Vin88], for any k, d it holds that $\chi_c(C(k, d)) = \frac{k}{d}$ and therefore $\chi(C(k, d)) = \lceil \frac{k}{d} \rceil$.

In this chapter, we focus on fractional colorings related to (directed) cycles in graphs and digraphs. For further background on circular colorings of undirected graphs we refer to the survey article of Zhu [Zhu01].

Circular Chromatic Number of Digraphs. The *circular chromatic number* of digraphs was introduced by Bokal, Fijavz, Juvan, Kayll, and Mohar $[BFJ^+04]$ as a refinement of the dichromatic number, which is capable of distinguishing between digraphs with the same dichromatic number by taking on (arbitrary) rational values. Instead of integer pairs as in the definition of the circular chromatic number by Vince, Bokal et al. use real numbers as colors in their definition:

Given a real number $p \ge 1$, consider a plane-circle S_p of perimeter p and define a weak circular p-coloring as a color-map $c : V(D) \to S_p$, such that equal colors at both ends of an arc, i.e., c(u) = c(w) where $e = (u, w) \in A(D)$, are allowed, but at the same time, the clockwise distance¹ from c(u) to c(w) on S_p is at least 1 whenever $c(u) \neq c(w)$. Additionally, each so-called color class, i.e., a set $c^{-1}(t), t \in S_p$, has to be an acyclic set.

The circular dichromatic number $\vec{\chi}_c(D)$ now is defined as the infimum over all values $p \geq 1$ providing weak circular *p*-colorings of *D*. This infimum was shown in [BFJ⁺04] to be always attained as a minimum.

For any natural numbers $k \ge d \ge 1$, denote by $\vec{C}(k,d)$ the digraph with vertex set \mathbb{Z}_k in which there is an arc (i,j) for $i,j \in \mathbb{Z}_k$ if and only if $j-i \in \{d, d+1, \ldots, k-1\}$. As these digraphs will play an important role in this chapter, in the following let us sum up properties of the circular dichromatic number and of these special circulant digraphs.

Theorem 13.1 (Bokal et al., $[BFJ^+04]$). Let D be a digraph. Then the following holds:

- (i) $\vec{\chi}_c(D) \geq 1$ is a rational number with numerator at most v(D).
- (*ii*) $[\vec{\chi}_c(D)] = \vec{\chi}(D), i.e., \vec{\chi}_c(D) \in (\vec{\chi}(D) 1, \vec{\chi}(D)].$
- (iii) $\vec{C}(k,d)$ has circular dichromatic number exactly $\frac{k}{d}$ for any $k \ge d \in \mathbb{N}$.
- (iv) D is weakly circularly p-colorable for $p \ge 1$ if and only if for every $(k, d) \in \mathbb{N}^2$ with $\frac{k}{d} \ge p$, D admits a coloring $c_{k,d} : V(D) \to \mathbb{Z}_k$ with the following properties:

For any arc $(u, w) \in A(D)$, either $c_{k,d}(u) = c_{k,d}(w)$ or $(c_{k,d}(w) - c_{k,d}(u)) \mod k \ge d$, and $c_{k,d}^{-1}(i)$ is acyclic for every $i \in \mathbb{Z}_k$.

¹By the clockwise distance of an element $x \in S_p$ to an element $y \in S_p$ we mean the length of the unique subarc of S_p connecting x to y in clockwise direction.

Graph Homomorphisms and Acyclic Homomorphisms. Given a pair of undirected graphs G, H, a graph homomorphism from G to H is a mapping $\phi : V(G) \to V(H)$ which preserves adjacency, i.e., for every pair of adjacent vertices u, v of G, the vertices $\phi(u)$ and $\phi(G)$ are adjacent in H. It is well-known that graph homomorphisms generalize graph colorings in the following way: Given a fixed graph H, for any graph G, an H-coloring is defined to be a graph homomorphism $\phi : V(G) \to V(H)$. The H-coloring problem then asks for a given graph G whether it is H-colorable. If we take H to be the complete graph on k vertices, this is just the k-coloring problem for graphs which is known to be polynomially solvable for k = 2 and NP-complete for $k \ge 3$. It was a long-standing open problem to determine the complexity of H-colorability for arbitrary graphs H. This was finally resolved by Hell and Nešetřil [HN90] who proved the following:

Theorem 13.2 (Hell and Nešetřil, [HN90]). The H-coloring problem is polynomially solvable if H is bipartite, and it is NP-complete if H is non-bipartite.

It is natural to ask for a definition of homomorphisms acting on digraphs which resembles acyclic colorings of digraphs in a similar way. One such notion, which has received quite some attention in past years, are *acyclic homomorphisms*. Given a pair D_1, D_2 of digraphs, an acyclic homomorphism from D_1 to D_2 is defined to be a mapping $\phi: V(D_1) \to V(D_2)$ with the property that for any arc (u, w) in D, either $\phi(u) = \phi(w)$ or $(\phi(u), \phi(w))$ is an arc in D_2 , and additionally, for every vertex $v \in V(D_2)$, the vertex set $\phi^{-1}(v)$ is an acyclic set of D_1 . The following statement describes the relation of (circular) digraph colorings and acyclic homomorphisms and shows that for digraph colorings, the circulant digraphs $\vec{C}(k, d)$ as defined above take the role of the complete graphs for usual graph colorings.

Proposition 13.3 (Bokal et al., [BFJ⁺04]). Let $p = \frac{k}{d} \ge 1$ with $k, d \in \mathbb{N}$ be a rational number. Then for any digraph D, we have $\vec{\chi}_c(D) \le p$ if and only if there is an acyclic homomorphism mapping D to the digraph $\vec{C}(k, d)$.

The question of determining the complexity of the decision problem whether or not $\vec{\chi}_c(D) \leq p$ for a fixed rational number $p \geq 1$ was raised by Bokal et al. [BFJ⁺04] and answered by Feder, Hell and Mohar [FHM03] in form of a much more general statement which can be seen as a variant of Theorem 13.2 for acyclic homomorphisms:

Theorem 13.4 (Feder, Hell and Mohar, [FHM03]). Let F be a digraph. Then the acyclic F-coloring problem, i.e., deciding whether a given digraph D admits an acyclic homomorphism to F, is polynomially solvable if F is acyclic and NP-complete otherwise.

Applying Theorem 13.4 with $F = \vec{C}(k, d)$ directly yields the following hardness result.

Corollary 13.5 (Feder, Hell and Mohar, [FHM03]). Given a rational number p > 1, deciding whether a digraph satisfies $\vec{\chi}_c(D) \leq p$ is NP-complete.

Star Dichromatic Number. Another related concept of circular colorings of digraphs was proposed by Hochstättler and the author [HS19] under the name of the *star dichromatic number* $\chi^*(D)$ of a digraph. Again, for a coloring, real numbers associated with a plane circle are used, but instead of looking at circular distances between adjacent vertices, an *acyclic p-coloring* of a digraph D for any $p \ge 1$ requires pre-images of cyclic open subintervals of length 1 to be acyclic. Alternatively, one may use pairs of integers to define the star dichromatic number: **Definition 13.3** ([HS19]). Let D be a digraph, $(k,d) \in \mathbb{N}^2, k \geq d$. An acyclic (k,d)coloring of D is an assignment $c: V(D) \to \mathbb{Z}_k$ of colors to the vertices such that for every $i \in \mathbb{Z}_k$, the pre-image of the cyclic interval $A_i := \{i, i+1, \ldots, i+d-1\} \subseteq \mathbb{Z}_k$ of colors, $c^{-1}(A_i) \subseteq V(D)$, is an acyclic set in D. The infimum of the values $\frac{k}{d}$ for which an acylic (k,d)-coloring exists is defined to be the star dichromatic number $\chi^{*}(D) \geq 1$ of D.

Similar to the circular dichromatic number, the star dichromatic number fulfils the following series of interesting properties.

Theorem 13.6 ([HS19]). Let D be a digraph. Then the following holds:

- (i) $\vec{\chi}^*(D) \ge 1$ is a rational number with numerator at most v(D).
- (*ii*) $[\vec{\chi}^*(D)] = \vec{\chi}(D), i.e., \vec{\chi}^*(D) \in (\vec{\chi}(D) 1, \vec{\chi}(D)].$
- (iii) $\vec{C}(k,d)$ admits star dichromatic number exactly $\frac{k}{d}$ for any $k \ge d \in \mathbb{N}$.
- (iv) For all $k \ge d \in \mathbb{N}$, D admits an acyclic (k, d)-coloring if and only if $\vec{\chi}^*(D) \le \frac{k}{d}$.

Although the star dichromatic number and the circular dichromatic number share many similar properties, in some cases they may behave very differently. While the star dichromatic number is immune to the addition of sinks and sources (as directed cycles may never pass them), this may have a significant effect on the circular dichromatic number, see Figure 13.1 for an illustration.

As a by-product of Section 13.2 we introduce a notion of homomorphisms for digraphs, so-called *circular homomorphisms*, which appropriately generalize the star dichromatic number in the same way that acyclic homomorphisms generalize circular digraph colorings.

Fractional Dichromatic Number. The last fractional coloring notion for digraphs we want to discuss in this (as well as the next) chapter is the *fractional dichromatic number* of a digraph D, denoted by $\vec{\chi}_f(D) \ge 1$. As its analogue for graphs, the well-known *fractional chromatic number* $\chi_f(G)$ of a graph G, it may be defined as the optimal value of a linear program. Here, acyclic vertex sets in digraphs play the role of independent vertex sets in undirected graphs.



Figure 13.1: Left: The directed cycle \vec{C}_4 , which has fractional, star and circular dichromatic number $\frac{4}{3}$. While the addition of a dominating source does not change the fractional and the star dichromatic number, the circular dichromatic number jumps to 2 (Right).

Definition 13.4 (cf. Severino, [Sev]). Let D be a digraph. Denote by $\mathcal{A}(D)$ the collection of vertex subsets of D inducing an acyclic subdigraph, and for each $v \in V(D)$, let $\mathcal{A}(D, v) \subseteq \mathcal{A}(D)$ be the subset containing only those sets including v. The fractional

dichromatic number $\vec{\chi}_f(D)$ of D is now defined as the value of

$$\min \sum_{A \in \mathcal{A}(D)} x_A$$
subj. to
$$\sum_{A \in \mathcal{A}(D,v)} x_A \ge 1, \text{ for all } v \in V(D)$$

$$x \ge 0.$$
(13.1)

The fractional dichromatic number has received some attention in recent years, it has e.g. proved to be useful for proving a fractional version of the *Erdős-Neumann-Lara-conjecture* mentioned in the introduction, see Mohar and Wu [Moh16], and has been related to acyclic homomorphisms by Severino [Sev].

By applying duality to the linear program defining $\vec{\chi}_f(D)$, we obtain the following alternative definition:

Proposition 13.7. The fractional dichromatic number of a digraph D can be computed as the optimal value of

$$\max \sum_{v \in V} y_v$$
subj. to $\sum_{v \in A} y_v \le 1$, for all $A \in \mathcal{A}(D)$
 $y \ge 0$.
$$(13.2)$$

Since the linear programs are bounded, they attain their optimal value $\vec{\chi}_f(D)$, and in fact there always exists a rational optimal solution (since the same holds true for every vertex of the underlying polyhedron using Cramer's rule). Therefore $\vec{\chi}_f(D)$ is always a rational number. It also follows directly by definition that $\vec{\chi}_f(\cdot)$ is monotonous with respect to taking subdigraphs. Further properties of the fractional dichromatic number, in particular of planar digraphs, will be investigated in Chapter 14. The following puts the three presented fractional digraph coloring parameters in relation and establishes direct relations to corresponding notions for graphs.

Proposition 13.8 (cf. [HS19]).

- (i) Let D be a digraph. Then $\vec{\chi}_f(D) \leq \vec{\chi}^*(D) \leq \vec{\chi}_c(D)$.
- (ii) For any graph G, we have $\vec{\chi}^*(\vec{G}) = \vec{\chi}_c(\vec{G}) = \chi_c(G)$ and $\vec{\chi}_f(\vec{G}) = \chi_f(G)$.

Circular Vertex Arboricity. The counterpart of digraph colorings for undirected graphs is known as the *vertex arboricity*, which we alrady encountered in Chapter 11. Given some $k \in \mathbb{N}$, let us define an *arboreal k-coloring* of a (multi-)graph G as a coloring of the vertices of G using colors $\{0, \ldots, k-1\}$ such that there are no monochromatic cycles, i.e., $G[c^{-1}(i)]$ is a forest for any $i \in \{0, \ldots, k-1\}$. The vertex arboricity va(G) is then the minimal number k for which G admits an arboreal k-coloring. Similar to the notions of circular colorings of graphs and digraphs, it is also possible to investigate a circular version of the vertex arboricity, which was introduced by Wang, Zhou, Liu and Wu [WZLW11] under the name *circular vertex arboricity*. For this purpose, the notion of an *arboreal* (k, d)-coloring of a graph G is defined. Similarly to acyclic (k, d)-colorings of digraphs, this is a mapping $c : V(G) \to \mathbb{Z}_k$ with the property that for any cyclic subinterval $A_i := \{i, i + 1, \ldots, i + d - 1\}$ of \mathbb{Z}_k , the subgraph of G induced by $c^{-1}(A_i)$ is a forest.

The circular vertex arboricity $va_c(G)$ of the graph G is now defined as the infimum of the values $\frac{k}{d}$ for which an arboreal (k, d)-coloring of G exists. The circular vertex arboricity has the following basic properties.

Theorem 13.9 (Wang et al. [WZLW11]). Let G be a (multi-)graph. Then

- (i) $va_c(G) \ge 1$ is a rational number with numerator at most v(G).
- (*ii*) $\lceil va_c(G) \rceil = va(G), i.e., va_c(G) \in (va(G) 1, va(G)].$
- (iii) For all $k \ge d \in \mathbb{N}$, D admits an arboreal (k, d)-coloring if and only if $va_c(G) \le \frac{k}{d}$.

This chapter is divided into three sections, studying the complexity of decision problems for the *Star Dichromatic Number*, the *Fractional Dichromatic Number* respectively the *Circular Vertex Arboricity* as defined above.

The following main results are proved:

Theorem 13.10.

- For any fixed rational number p > 1, deciding whether a given (multi-)digraph D fulfills $\vec{\chi}^*(D) \leq p$ is NP-complete.
- For any fixed real number $p > 1, p \neq 2$, deciding whether a given (multi-)digraph D fulfills $\vec{\chi}_f(D) \leq p$ is NP-complete.
- For any fixed rational number p > 1, deciding whether a given (multi-)graph G fulfills $va_c(G) \le p$ is NP-complete.

This theorem answers an open question from [HS19] as well as questions in the context of the work by Wang et al. [WZLW11]. It further resembles the hardness result achieved by Feder, Hell, and Mohar [FHM03] for the circular dichromatic number (cf. Corollary 13.5).

The proof of the NP-hardness requires different techniques for each case. As a tool to prove Theorem 13.10, in Section 13.2 the notion of *circular homomorphisms* acting between digraphs is introduced and might be of independent interest.

13.2 Complexity of the Star Dichromatic Number

In this section, we deal with decision problems for the star dichromatic number analogous to those considered by Feder, Hell and Mohar [FHM03] for the circular dichromatic number. The problem of determining the computational complexity of the following decision problem was posed by Hochstättler and the author in [HS19]:

Problem 13.1. Let $p \ge 1$ be a fixed rational number.

Instance: A (multi-)digraph D. Decide whether $\vec{\chi^*}(D) < p$.

For p = 1, the problem is to decide whether $\vec{\chi}^*(D) = 1$, which is equivalent to D being acyclic, and hence this can be solved in time linear in v(D) + a(D).

We now introduce *circular homomorphisms* as an extension of the well-known *acyclic homomorphisms* defined in the previous section.

Definition 13.5. Let D_1, D_2 be digraphs. A mapping $\phi : V(D_1) \to V(D_2)$ is called a circular homomorphism, if for all $A \subseteq V(D_2)$ such that $D_2[A]$ is acyclic, $\phi^{-1}(A)$ is acyclic in D_1 . Equivalently, for any directed cycle C in $D_1, D_2[\phi(V(C))]$ contains a directed cycle. It is obvious that D_1 admits the injection $\mathrm{id}|_{V(D_1)}$ as a circular homomorphism to D_2 whenever D_1 is a subdigraph of D_2 and that the composition of two circular homomorphisms remains a circular homomorphism.

Note that this definition is one natural way to generalize graph homomorphisms to digraphs, as the former may be characterized by the property that pre-images of *independent* sets remain independent. This similarity is made precise by the following.

Proposition 13.11. Let G_1, G_2 be graphs. Then a mapping $\phi : V(G_1) \to V(G_2)$ is a graph homomorphism if and only if it is a circular homomorphism from \vec{G}_1 to \vec{G}_2 .

Proof. This follows immediately from the characterizations of graph resp. circular homomorphisms in terms of independent resp. acyclic vertex sets and the fact that for any graph G, the acyclic vertex sets of \tilde{G} are exactly the independent vertex sets of G.

The following, which is similar to Proposition 13.3, reformulates acyclic (k, d)-colorings in terms of circular homomorphisms.

Proposition 13.12. Let $p = \frac{k}{d} \ge 1, k, d \in \mathbb{N}$. Then $\vec{\chi}^*(D) \le p$, i.e., there is an acyclic (k, d)-coloring of D, if and only if there is a circular homomorphism from D to $\vec{C}(k, d)$.

Proof. Recall that $\vec{C}(k,d)$ was defined to be the digraph with vertex set $\mathbb{Z}_k \simeq \{0,\ldots,k-1\}$ where there is an arc (i,j) between two elements if and only if $j-i \in \{d,\ldots,k-1\} \subseteq \mathbb{Z}_k$. To prove the claim, we need the following property: A vertex set $A \subseteq \vec{C}(k,d)$ is acyclic if and only if it is contained in a set of d consecutive vertices, i.e., $A \subseteq \{i, i+1, \ldots, i+d-1\} \subseteq \mathbb{Z}_k$ with some $i \in \mathbb{Z}_k$.

For the first implication, assume that A is acyclic. Then $\vec{C}(k,d)[A]$ must contain a sink $i \in A$ (i.e., *i* has no out-neighbors in A), which means that none of the vertices $i + d, \ldots, i + k - 1 \in \mathbb{Z}_k$ can be contained in A, and so $A \subseteq \{i, \ldots, i + d - 1\}$. For the reverse, since $\vec{C}(k,d)$ is circulant, it is enough to show that $\{0,\ldots,d-1\} \subseteq \mathbb{Z}_k$ is acyclic. However, by definition, this interval can only contain backward arcs, and so the subdigraph of $\vec{C}(k,d)$ induced by $\{0,\ldots,d-1\}$ admits a topological ordering.

Consequently, the circular homomorphisms $\phi : V(D) \to \mathbb{Z}_k = V(\vec{C}(k,d))$ from any digraph D to $\vec{C}(k,d)$ are exactly those mappings for which $\phi^{-1}(\{i,\ldots,i+d-1\})$ is acyclic for all $i \in \mathbb{Z}_k$, and this is just the same as an acyclic (k,d)-coloring of D.

This proves the claim of the proposition.

Furthermore, the well-studied acyclic homomorphisms between digraphs appear as a special case of circular homomorphisms:

Proposition 13.13. Let D_1, D_2 be two digraphs. Then every acyclic homomorphism $\phi: V(D_1) \to V(D_2)$ is a circular homomorphism.

Proof. Let C be any directed cycle in D_1 . We need to show that $\phi(V(C))$ contains the vertex set of a directed cycle. Since pre-images of single vertices under ϕ are acyclic in D_1 , $\phi(V(C))$ needs to contain at least two vertices. Any arc (x, y) on C is either mapped to a single vertex $\phi(x) = \phi(y)$ or to an arc $(\phi(x), \phi(y))$ of D_2 , which implies that $D_2[\phi(V(C))]$ contains a closed directed walk visiting at least two vertices and thus also a directed cycle. This proves the claim.

However, the converse of this statement is not true. This follows from the fact that there are digraphs with $\vec{\chi}_c(D) > \vec{\chi}^*(D)$, Proposition 13.3, and Proposition 13.12. Examples of such digraphs are e.g. directed cycles with an additional dominating source (cf. [HS19]), see also Figure 13.1.

We conclude the discussion of circular homomorphisms with the following observation, which identifies them as interlacing structures between digraphs in terms of their star and fractional dichromatic numbers:

Proposition 13.14. Let D_1, D_2 be digraphs such that there is a circular homomorphism $\phi: V(D_1) \to V(D_2)$. Then $\vec{\chi}^*(D_1) \leq \vec{\chi}^*(D_2)$ and $\vec{\chi}_f(D_1) \leq \vec{\chi}_f(D_2)$.

Proof. The inequality for the star dichromatic number follows from Proposition 13.12 and the fact that the composition of two circular homomorphisms remains a circular homomorphism. The inequality for the fractional dichromatic number can be seen from the definition in terms of the linear program (13.1) as follows: Given any optimal solution $x' \geq 0$ of the program with respect to D_2 , define a corresponding instance $x \geq 0$ of the program for D_1 by assigning the value²

$$x_A := \sum_{\substack{A' \in \mathcal{A}(D_2):\\ \phi^{-1}(A') = A}} x'_A$$

for every acyclic vertex set $A \in \mathcal{A}(D_1)$. It is now easily verified using the fact that $\phi^{-1}(A') \in \mathcal{A}(D_1)$ for any $A' \in \mathcal{A}(D_2)$, that x is a legal instance of (13.1) for D_1 with

$$\vec{\chi}_f(D_1) \le \sum_{A \in \mathcal{A}(D_1)} x_A = \sum_{A' \in \mathcal{A}(D_2)} x'_{A'} = \vec{\chi}_f(D_2).$$

Given a fixed digraph F, any other digraph D will be called *circularly* F-colorable if there exists a circular homomorphism mapping D to F. The following decision problem, which can be seen as a directed analogue of the H-coloring problem for graphs then generalizes Problem 13.1.

Problem 13.2. Let F be a fixed (multi-)digraph. Instance: A (multi-)digraph D. Decide whether D is circularly F-colorable.

As in the graph coloring problem, there is a trivial case: Only acyclic digraphs map circularly to acyclic digraphs:

Observation 13.15. The circular F-coloring problem is polynomially solvable for any acyclic digraph F.

We conjecture that this simple observation covers already all polynomially solvable cases under the assumption $P \neq NP$. In other words,

Conjecture 13.1. Let F be a digraph which contains a directed cycle. Then the circular F-coloring problem is NP-complete.

Our main result of this section is the following theorem, which shows that this conjecture holds true in almost all the cases. Given a (multi-)digraph, the symmetric part of D is defined to be the simple graph on the same vertex set as D which contains an edge xy if and only if there is an arc from x to y and from y to x in D.

Theorem 13.16. Let F be a digraph containing a directed cycle such that at least one of the following holds:

²Throughout this thesis we define the value of an empty sum to be 0.

- (i) The symmetric part of F is edgeless, i.e., F is digon-free, or
- (ii) The symmetric part of F contains an odd cycle, or

(*iii*) $\vec{\chi}(F) = 2$.

Then the circular F-coloring problem is NP-complete.

Proof. We start by observing that the problem is contained in the complexity class NP: Given a digraph D, any circular homomorphism from D to F can be used as a certificate that D is circularly F-colorable. Note that the digraph F itself defines the coloring problem and is not considered as an instance, and so for checking whether a given mapping $\phi : V(D) \rightarrow V(F)$ defines a circular homomorphism, it suffices to compute at most constantly many inverse images under ϕ and verify that the corresponding subdigraphs of D are indeed acyclic, which can be done in polynomial time in v(D).

(i), (ii) Assume that the symmetric part of F is either edgeless or non-bipartite. Let us denote $k := \vec{g}(F)$. We will define a graph H_F with vertex set V(F) and in which two vertices $u \neq v$ are adjacent if and only if there is a directed cycle of length k containing both u and v in F. We now give a polynomial reduction of the H_F coloring problem to the circular F-coloring problem. For a given instance G of the H_F -coloring problem, we will construct a polynomial-sized instance D_G for the circular F-coloring problem and prove that G is H_F -colorable if and only if the digraph D_G is circularly F-colorable.

If the symmetric part of F is empty, then there is a directed cycle of length $k \geq 3$ in F which forms a clique of size k in H_F , and therefore H_F contains a triangle. Otherwise, the symmetric part of F is non-empty but non-bipartite. In this case, we have k = 2 and the vertices of any odd cycle in the symmetric part of F will form an odd cycle in H_F . In any case, H_F is non-bipartite, and thus, the decision problem of H_F -colorability is NP-hard according to Theorem 13.2.

Let now G be an instance of the H_F -coloring problem. We construct the digraph D_G by first choosing some acyclic orientation \vec{G} of G and then attaching to every arc (x, y) of \vec{G} a directed path of length k - 1 in reverse direction whose only common vertices with \vec{G} are x and y, so that each arc e = (x, y) in \vec{G} is contained in a directed cycle C(e) of length k in D_G . The set of the k - 2 extra vertices that are added is pairwise disjoint for distinct arcs. Clearly, this construction is polynomial in a(G).

We now show that there is a graph homomorphism from G to H_F if and only if there is a circular homomorphism from D_G to F. Since the H_F -coloring problem is NP-hard, this will prove NP-hardness of the circular F-coloring problem, as desired.

For the first implication let $\phi: V(G) \to V(H_F)$ be a graph homomorphism mapping G to H_F . Then for any arc e = (x, y) in \vec{G} , $\phi(x)\phi(y)$ is an edge in H_F , i.e., there is a directed cycle C'(e) of length k containing $\{\phi(x), \phi(y)\}$ in F. Moreover, since $|V(C(e)) \setminus \{x, y\}| = |V(C'(e)) \setminus \{\phi(x), \phi(y)\}| = k - 2$, we find that there are bijections $f_e: V(C(e)) \setminus \{x, y\} \to V(C'(e)) \setminus \{\phi(x), \phi(y)\}$ for every edge e of G.

Let now $\phi' : V(D_G) \to V(F)$ be the mapping defined by $\phi'(u) := \phi(u)$ for any $u \in V(G) \subseteq V(D_G)$ and $\phi'(u) := f_e(u)$ for any $u \in V(C(e)) \setminus \{x, y\}$ and any edge $e = xy \in E(G)$. We claim that this defines a circular homomorphism from D_G to F: If C is any directed cycle in D_G , since \vec{G} is an acyclic orientation, C needs to contain a full attachment path and thus the vertex set of C(e) for at least one edge

 $e \in E(G)$. This implies $\phi'(V(C)) \supseteq \phi'(V(C(e))) = V(C'(e))$ by definition of ϕ' , and thus C'(e) is a directed cycle contained in the image of V(C), as required.

Conversely, assume there is a circular homomorphism ϕ' mapping D_G to F. We claim that the restriction $\phi := \phi'|_{V(G)}$ is a graph homomorphism from G to H_F . For this purpose, let e = xy be any edge of G. Then, since ϕ' is a circular homomorphism, $\phi'(V(C(e)))$ contains the vertex set of a directed cycle in F, which must have length at least k. However, $|\phi'(V(C(e)))| \leq |V(C(e))| = k$ by definition of C(e), so this directed cycle has exactly $\phi'(V(C(e)))$ as vertex set, which contains $\phi'(x)$ and $\phi'(y)$. Hence, ϕ' restricted to V(C(e)) must be injective, and so $\phi'(x) \neq \phi'(y)$. According to the definition of H_F , this finally implies that $\phi(x)\phi(y) = \phi'(x)\phi'(y)$ is an edge of H_F , and so ϕ is indeed a graph homomorphism as required. This settles the proof in the case where the symmetric part of F is empty or non-bipartite.

(iii) Now, let F be acyclically 2-colorable. Referring to (i) and (ii), we may assume that the symmetric part of F is non-empty and bipartite. Since F is 2-colorable, by Proposition 13.12 there is a circular homomorphism from F to \vec{C}_2 . On the other hand, since the symmetric part of F is non-empty, F contains a digon, and thus, there also is a circular homomorphism from \vec{C}_2 to F. Hence, in this case the circular F-coloring and the circular \vec{C}_2 -coloring problem are equivalent. However, as we have discussed in Chapter 10, deciding 2-colorability of a digraph is NP-hard.

Applying this result to the star dichromatic number, we finally obtain the desired hardness result:

Theorem 13.17. Let p > 1 be a rational number. Deciding whether D admits an acyclic p-coloring, i.e., whether $\vec{\chi}^*(D) \leq p$, is NP-complete.

Proof. Let $p = \frac{k}{d}$ with $k, d \in \mathbb{N}$. By Proposition 13.12, the decision problem is equivalent to the circular $\vec{C}(k, d)$ -coloring problem.

To prove NP-completeness, we distinguish between $p \leq 2$ and p > 2. In the first case, $\vec{C}(k, d)$ is not acyclic and 2-colorable, and thus the claim follows from Theorem 13.16 (iii).

In the case p > 2, the symmetric part of $\vec{C}(k, d)$ is given by the circulant graph C(k, d) (cf. Definition 13.2). C(k, d) has chromatic number $\lceil p \rceil \geq 3$ and thus is not bipartite. Consequently, another application of Theorem 13.16 yields the claimed result. \Box

13.3 Complexity of the Fractional Dichromatic Number

Continuing in the spirit of the previous section, we now want to deal with decision problems for the fractional dichromatic number as follows.

Problem 13.3. Let $p \ge 1$ be a fixed real number. Instance: A (multi-)digraph D. Decide whether $\vec{\chi}_f(D) \le p$.

Again, it is clear that for p = 1, $\vec{\chi}_f(D) \leq p$ if and only if the digraph D is acyclic, and this can be decided in linear time in v(D) + a(D). Conversely, we want to show in the following that for all real numbers $p > 1, p \neq 2$, this problem is NP-complete. It is indeed always contained in NP:

Observation 13.18. For any $p \ge 1$, the Problem 13.3 is in NP.

Proof. Let D be a digraph given as an instance of the problem. Let again $\mathcal{A}(D)$ denote the set of acyclic vertex sets in D. We have to prove the existence of a certificate polynomiallysized in v(D) which is verifiable in polynomial time. For this purpose, we repeat some standard arguments from linear programming for (13.1). Clearly, any optimal solution of (13.1) satisfies³ $x \leq 1$. Thus, adding the constraints $x_A \leq 2$ for all $A \in \mathcal{A}(D)$ yields an equivalent bounded feasible program. As the optimal solution is attained by a vertex xof the corresponding polytope there is a subset of $|\mathcal{A}(D)|$ inequality-constraints which are satisfied by x with equality, and the corresponding linear system uniquely determines x. Since any optimal solution x satisfies $x \leq 1$, none of the additional constraints is tight, implying that the size of the support of x (the set of $A \in \mathcal{A}(D)$ such that $x_A > 0$) is $m := |\operatorname{supp}(x)| \leq v(D)$. Denote by $x' \in \mathbb{R}^m$ the subvector of x restricted to the support. As x' is the unique solution of a regular linear system, according to Cramer's rule, there are matrices $B_1, \ldots, B_m, B \in \{0, 1\}^{m \times m}$ such that $x_i = \frac{\det(B_i)}{\det(B)}, i = 1, \ldots, m$. According to Hadamard's inequality, we have $|\det(B_i)|, |\det(B)| \le m^{m/2} \le v(D)^{v(D)/2}, i = 1, \dots, m.$ This finally implies that there exists an optimal solution to the linear program (13.1)whose support is of size at most v(D) and where the non-zero values in the solution are rational numbers, each of which can be stored using at most $v(D) \log v(D)$ bits. Such a solution can thus be described using $\mathcal{O}(v(D)^2 \log v(D))$ bits. As we can verify all the constraints and the inequality $\sum_{i=1}^{m} x_i \leq p$, certifying that $\vec{\chi}_f(D) \leq p$, in polynomial time in $m \leq v(D)$, this finally proves that we can use optimal solutions of this form as polynomial-time verifiable NP-certificates. This concludes the proof.

We start our proof of the hardness with the following simple observation derived from the relation of the fractional chromatic and the fractional dichromatic number:

Observation 13.19. Let $p \in \mathbb{R}$, p > 2. Then Problem 13.3 is NP-complete.

Proof. It is well-known (see e.g. [SU97], Theorem 3.9.2) that the problem of deciding whether $\chi_f(G) \leq p$ for a given graph G is NP-hard for any real number p > 2. However, this problem admits a polynomial reduction to Problem 13.3 for p: For any graph G, the biorientation \dot{G} fulfils $\chi_f(G) \leq p \Leftrightarrow \chi_f(\dot{G}) = \chi_f(G) \leq p$. This proves the claim. \Box

It thus suffices to prove the hardness in the case $p \in (1, 2)$. For any given p, we will reduce one of the decision problems proved to be hard in Observation 13.19 to Problem 13.3 with p. For this purpose, we introduce a certain operation on digraphs reducing its fractional dichromatic number:

Definition 13.6. Let D be a digraph. For every $\ell \ge 1$, we denote by D_{ℓ} a digraph called ℓ -split of D obtained from D by replacing each vertex by a directed path of length $\ell - 1$ as follows: Each vertex $x \in V(D)$ is assigned a directed path $P(x) = x_1, \ldots, x_{\ell}$ in D_{ℓ} . The remaining adjacencies within D_{ℓ} are given as follows: For each arc e = (u, w) in D, we have a corresponding arc (u_{ℓ}, w_1) in D_{ℓ} . Thus, in a path $P(x), x_1, \ldots, x_{\ell-1}$ have out-degree 1 while x_2, \ldots, x_{ℓ} have each exactly one incoming arc.

It follows by definition that each directed cycle in D_{ℓ} contains the whole path P(x) or none of its vertices, for all $x \in V(D)$. This means that there is a bijection between the directed cycles in D and those in D_{ℓ} by replacing each vertex $x \in V(D)$ contained in a directed cycle by P(x) in D_{ℓ} and vice versa. The following makes the relation between the fractional dichromatic numbers of D and D_{ℓ} precise.

³By $\mathbf{1}$ we denote vectors all whose entries equal 1.

Proposition 13.20. For each digraph D and every integer $\ell \geq 1$, the following holds:

$$\vec{\chi}_f(D_\ell) = \frac{\ell \vec{\chi}_f(D)}{(\ell - 1)\vec{\chi}_f(D) + 1}.$$

Proof. Assume $\ell \geq 2$ (for $\ell = 1$ we have $D_1 = D$ and the claim holds trivially). For the proof we use the alternative representation of $\vec{\chi}_f(D)$ as the maximal value of the dual program (13.2) in Proposition 13.7.

Throughout the rest of the proof, the following relation between acyclic sets of Dand its ℓ -split will be crucial: Define a mapping $f : \mathcal{A}(D_{\ell}) \to \mathcal{A}(D)$ such that for all $B \in \mathcal{A}(D_{\ell}), f(B) := \{x \in V(D) | V(P(x)) \subseteq B\}$. Furthermore, define $g : \mathcal{A}(D) \to \mathcal{A}(D_{\ell})$ by $g(A) := \bigcup_{x \in A} V(P(x))$. These mappings are well-defined due to the bijection between directed cycles in D resp. D_{ℓ} described above. We clearly have $f \circ g = \mathrm{id}_{\mathcal{A}(D)}$ and thus, gis injective while f is surjective.

We start by showing that $\vec{\chi}_f(D_\ell) \geq \frac{\ell \vec{\chi}_f(D)}{(\ell-1)\vec{\chi}_f(D)+1}$. For this purpose, let $y_v, v \in V(D)$ be an optimal instance for the dual program (13.2) for D, i.e., $\sum_{v \in V(D)} y_v = \vec{\chi}_f(D)$. We define an instance of the dual problem for D_ℓ as follows: For every $w \in V(D_\ell), v \in V(D)$ with $w \in V(P(v))$, let $y'_w := \frac{y_v}{(\ell-1)\vec{\chi}_f(D)+1} \geq 0$. Obviously,

$$\sum_{w \in V(D_{\ell})} y'_w = \frac{\sum_{y \in V(D)} \ell y_v}{(\ell - 1)\vec{\chi}_f(D) + 1} = \frac{\ell \vec{\chi}_f(D)}{(\ell - 1)\vec{\chi}_f(D) + 1}$$

Furthermore, for each $B \in \mathcal{A}(D_{\ell})$, we have

$$\sum_{w \in B} y'_w \le \sum_{v \in f(B)} \ell \frac{y_v}{(\ell - 1)\vec{\chi}_f(D) + 1} + \sum_{v \in V(D) \setminus f(B)} (\ell - 1) \frac{y_v}{(\ell - 1)\vec{\chi}_f(D) + 1}$$
$$= \frac{1}{(\ell - 1)\vec{\chi}_f(D) + 1} \left((\ell - 1) \sum_{\substack{v \in V(D) \\ = \vec{\chi}_f(D)}} y_v + \sum_{\substack{v \in f(B) \\ \le 1}} y_v \right) \le 1.$$

Thus, the y'_w are admissible for the program, which proves the first inequality.

For the reverse inequality, we want to show

$$\vec{\chi}_f(D_\ell) \le \frac{\ell \vec{\chi}_f(D)}{(\ell-1)\vec{\chi}_f(D)+1} \quad \text{or equivalently}$$
$$\vec{\chi}_f(D) \ge \frac{\vec{\chi}_f(D_\ell)}{\ell - (\ell-1)\vec{\chi}_f(D_\ell)}.$$

Notice that always $\vec{\chi}_f(D_\ell) < \frac{\ell}{\ell-1}$ and thus $\ell - (\ell-1)\vec{\chi}_f(D_\ell) > 0$. To see this, note that for every vertex $v \in V(D)$ and any selection of vertices $r(w) \in V(P(w))$ for $w \in V(D) \setminus \{v\}$, the set $V(P(v)) \cup \bigcup_{w \in V(D) \setminus \{v\}} (V(P(w)) \setminus \{r(w)\})$ is acyclic in D. There are $v(D)\ell^{v(D)-1}$ such acyclic sets, and every vertex in $V(D_\ell)$ is contained in all but $(v(D) - 1)\ell^{v(D)-2}$ of these sets. Hence assigning weight $\frac{1}{v(D)\ell^{v(D)-1}-(v(D)-1)\ell^{v(D)-2}}$ to every such acyclic set (and 0 to any other acyclic set) gives a legal instance of the program 13.1, showing that indeed

$$\vec{\chi}_f(D) \le \frac{v(D)\ell^{v(D)-1}}{v(D)\ell^{v(D)-1} - (v(D)-1)\ell^{v(D)-2}} < \frac{\ell}{\ell-1}.$$

Assume now that $y'_w, w \in V(D_\ell)$ is an optimal solution of the dual program (13.2) for D_ℓ , which means $\sum_{w \in V(D_\ell)} y'_w = \vec{\chi}_f(D_\ell)$. We define an instance of the dual program for D according to

$$y_v := \frac{\sum_{w \in V(P(v))} y'_w}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)} \ge 0$$

for each $v \in V(D)$. First of all, with this definition, we have

$$\sum_{v \in V(D)} y_v = \frac{\sum_{w \in V(D_\ell)} y'_w}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)} = \frac{\vec{\chi}_f(D_\ell)}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)}.$$

For the above inequality, it thus suffices to verify that the y_v define a legal instance for the dual program: Let $A \in \mathcal{A}(D)$ be arbitrary. Then

$$\sum_{v \in A} y_v = \frac{\sum_{v \in A} \sum_{w \in V(P(v))} y'_w}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)} = \frac{\sum_{w \in g(A)} y'_w}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)}$$

For each $v \in V(D) \setminus A$, we choose exactly one vertex $w_v \in V(P(v))$ with minimal value within P(v) and consider the acyclic vertex subset $X := V(D_\ell) \setminus \bigcup_{v \in V(D) \setminus A} \{w_v\}$ which contains g(A). According to our choice of the w_v , we know that

$$1 \ge \sum_{w \in X} y'_w = \sum_{w \in g(A)} y'_w + \sum_{v \in V(D) \setminus A} \left(\sum_{w \in V(P(v)), w \neq w_v} y'_w \right)$$
$$\ge \sum_{w \in g(A)} y'_w + \sum_{v \in V(D) \setminus A} \frac{\ell - 1}{\ell} \sum_{w \in V(P(v))} y'_w = \left(\frac{\ell - 1}{\ell} + \frac{1}{\ell}\right) \sum_{w \in g(A)} y'_w + \frac{\ell - 1}{\ell} \sum_{w \in V(D_\ell) \setminus g(A)} y'_w$$
$$= \frac{\ell - 1}{\ell} \vec{\chi}_f(D_\ell) + \frac{1}{\ell} \sum_{w \in g(A)} y'_w.$$

Multiplying the inequality with ℓ and subtracting $(\ell - 1)\vec{\chi}_f(D_\ell)$ now yields that indeed

$$\sum_{v \in A} y_v = \frac{\sum_{w \in g(A)} y'_w}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)} \le 1,$$

and thus $\vec{\chi}_f(D) \ge \frac{\vec{\chi}_f(D_\ell)}{\ell - (\ell - 1)\vec{\chi}_f(D_\ell)}$ as claimed. Finally, this proves $\vec{\chi}_f(D_\ell) = \frac{\ell \vec{\chi}_f(D)}{(\ell - 1)\vec{\chi}_f(D) + 1}$.

The following is now an immediate consequence of the above:

Theorem 13.21. Problem 13.3 is NP-complete for every real number $p > 1, p \neq 2$.

Proof. The case p > 2 was proved in Observation 13.19, so let now $p \in (1,2)$ be arbitrary. Then there is an $\ell \in \mathbb{N}, \ell \geq 2$ only dependent on p such that $p \in (\frac{2\ell}{2\ell-1}, \frac{\ell}{\ell-1})$. Choose such an ℓ and define $p' := \frac{p}{\ell-(\ell-1)p}$. Then p' > 2 and thus, Problem 13.3 is NP-hard for p'. However, since the function $x \to \frac{\ell x}{(\ell-1)x+1}$ is strongly increasing for positive values of x, we have for any digraph D that $\vec{\chi}_f(D) \leq p'$ if and only if $\vec{\chi}_f(D_\ell) = \frac{\ell \vec{\chi}_f(D)}{(\ell-1)\vec{\chi}_f(D)+1} \leq \frac{\ell p'}{(\ell-1)p'+1} = p$. This thus provides a polynomial reduction of Problem 13.3 with p' to the one with p, proving the NP-hardness (and thus -completeness) of the latter. \Box

13.4 Complexity of the Circular Vertex Arboricity

In this section, we treat the analogue of the decision Problems 13.1 and 13.3 for the circular vertex arboricity $va_c(G)$ of graphs.

Problem 13.4. Let $p \ge 1$ be a fixed rational number. Instance: A (multi-)graph G. Decide whether $va_c(G) \le p$.

As acyclicity of a graph can be tested in polynomial time, it is easy to see that in the case that $p = \frac{k}{d} \ge 1$ is a rational number, any arboreal (k, d)-coloring of a graph can be used as a polynomially verifiable certificate for $\operatorname{va}_c(G) \le p$, and so the above decision problem for p is contained in NP.

In order to prove complexity results, as in the case of the star dichromatic number, we could introduce a notion of circular homomorphisms between graphs analogous to Definition 13.5 and consider corresponding homomorphism-coloring problems. However, unlike in the case of digraphs, no easy interpretation of the arboreal (k, d)-coloring problem of a graph as such a homomorphism problem seems possible in general. This goes along with the fact that no simple canonical constructions of graphs with circular vertex arboricity $\frac{k}{d}$ for every pair (k, d) similar to the circulant (di)graphs $\vec{C}(k, d), C(k, d)$ are known so far.

It is again easily observed that $va_c(G) = 1$ for any graph G if and only if it is a forest, so the above decision problem is polynomially solvable for p = 1.

In the following we prove that similar to the cases of the circular and star dichromatic numbers, Problem 13.4 is NP-complete for all rational numbers p > 1. We prepare the proof with the following observation.

Lemma 13.22. Let $(k,d) \in \mathbb{N}^2$, k > d. Let I(k,d) denote the minimal size of a subset of \mathbb{Z}_k which is not contained in a cyclic subinterval of size d. Then

$$I(k,d) := \min\left\{|A| \middle| A \subseteq \mathbb{Z}_k, \forall i \in \mathbb{Z}_k : A \not\subseteq \{i, i+1, \dots, i+d-1\}\right\} = \left\lceil \frac{k}{k-d} \right\rceil.$$

Proof. The complements of the cyclic subintervals of \mathbb{Z}_k of size d are the cyclic subintervals of size k - d. Thus, a set $A \subseteq \mathbb{Z}_k$ is not contained in a cyclic subinterval of size d if and only if any two consecutive points in A according to the cyclic ordering of \mathbb{Z}_k have cyclic distance at most k - d. Consequently, $(k - d)|A| \ge k$ implying $|A| \ge \left\lceil \frac{k}{k-d} \right\rceil$, and thus $I(k,d) \ge \left\lceil \frac{k}{k-d} \right\rceil$. On the other hand, we may define $\left\lceil \frac{k}{k-d} \right\rceil$ points in \mathbb{Z}_k according to $a_i := ((k - d)i) \mod k$, for each $i \in \{0, \ldots, \left\lceil \frac{k}{k-d} \right\rceil - 1\} \subseteq \mathbb{Z}_k$, and it is easily seen that $\left\{a_0, \ldots, a_{\left\lceil \frac{k}{k-d} \right\rceil - 1}\right\}$ defines a set as required, proving $I(k,d) \le \left\lceil \frac{k}{k-d} \right\rceil$.

For any pair $(k, d) \in \mathbb{N}^2$, k > d we now define a simple auxiliary graph H(k, d) which has vertex set \mathbb{Z}_k and in which a pair $i \neq j \in \mathbb{Z}_k$ of vertices is adjacent if and only if there is a subset $A \subseteq \mathbb{Z}_k$ not contained in any cyclic subinterval of size d such that $\{i, j\} \subseteq A$ and |A| = I(k, d). It is easy to see that whenever $\frac{k}{d} \geq 2$, H(k, d) is just the circulant graph C(k, d) introduced in Definition 13.2. More generally, it follows from the definition that adjacency in H(k, d) only depends on the circular distance of the respective vertices. Hence, H(k, d) is always a circulant graph. For instance, H(5, 3) is the complete graph K_5 , H(6, 4) is the disjoint union of two triangles, and H(8, 5) admits an edge between vertices $i, j \in \mathbb{Z}_8$ if and only if $|i - j|_8 \in \{2, 3\}$.

We are now prepared to prove the following NP-hardness result.

Theorem 13.23. For any rational number p > 1, Problem 13.4 is NP-complete.

Proof. The NP-membership of the problem was verified above. So let now 1 be arbitrary but fixed. We distinguish between the cases <math>p = 2 and $p \neq 2$.

Assume first that $p \neq 2$. We prove the claimed NP-hardness by describing a polynomial reduction of the H(k, d)-coloring problem (in terms of graph homomorphisms) to Problem 13.4 with $p = \frac{k}{d}$. To do so, given any graph G as an instance of the H(k, d)-coloring problem, we construct (in polynomial time) a graph $G_{k,d}$ of size polynomial in v(G) and prove that G maps to H(k, d) if and only if $va_c(G_{k,d}) \leq \frac{k}{d}$, which is equivalent to $G_{k,d}$ admitting an arboreal (k, d)-coloring.

The graph $G_{k,d}$ is obtained from G by replacing any edge $e \in E(G)$ by a bunch of 2^k parallel paths of length I(k,d) - 1 each connecting the end vertices of e. The vertex sets of different replacement-paths are disjoint except for common end vertices.

To prove the first direction of the claimed equivalence, assume there is a graph homomorphism $\phi: V(G) \to \mathbb{Z}_k = V(H(k,d))$. This means that for any edge $e = xy \in E(G)$, $\phi(x) \neq \phi(y)$ are contained in a subset $A(e) \subseteq \mathbb{Z}_k$ of size I(k,d) which is not contained in a cyclic subinterval of \mathbb{Z}_k of size d. We now define a coloring $c: V(G_{k,d}) \to \mathbb{Z}_k$ of $G_{k,d}$ as follows: Any vertex $v \in V(G_{k,d})$ originally contained in G gets color $c(v) := \phi(v)$. For any replacement-path P of an edge $e \in E(G)$, we assign all the I(k,d) - 2 elements of $A(e) \setminus \{\phi(x), \phi(y)\}$ to the I(k, d) - 2 internal vertices of P (in arbitrary order). We claim that this defines an arboreal (k, d)-coloring of $G_{k,d}$: For any cycle C in $G_{k,d}, V(C)$ contains the vertex set of a whole replacement-path of an edge $e \in E(G)$, and thus $c(V(C)) \supseteq A(e)$. As A(e) is not contained in any cyclic subinterval of \mathbb{Z}_k of size d, the same is true for c(V(C)). This proves the validity of c as an arboreal (k, d)-coloring and we conclude $\operatorname{va}_c(G_{k,d}) \leq \frac{k}{d}$.

To prove the reverse implication, assume $\operatorname{va}_c(G_{k,d}) \leq p = \frac{k}{d}$, i.e., $G_{k,d}$ admits an arboreal (k, d)-coloring c. We define $\phi : V(G) \to \mathbb{Z}_k = V(H(k, d))$ by restriction of c to the vertices originally contained in G. We claim that this defines a graph homomorphism. To prove this, let $e = xy \in E(G)$ be an edge. Any of the 2^k replacementpaths of e in $G_{k,d}$ receives a non-empty subset of \mathbb{Z}_k of colors according to c. Applying the pigeon-hole principle we find a pair $P_1(e) \neq P_2(e)$ of replacement-paths of e such that $c(V(P_1(e))) = c(V(P_2(e)))$. As the union of $P_1(e)$ and $P_2(e)$ forms a cycle in $G_{k,d}$, according to the definition of an arboreal (k, d)-coloring, it follows that $c(V(P_1(e))) \cup c(V(P_2(e))) = c(V(P_1(e)))$ is not contained in a cyclic subinterval of \mathbb{Z}_k of size d and is of size at most $|c(V(P_1(e)))| \leq |V(P_1(e))| = I(k, d)$. By definition of I(k, d), this implies $|c(V(P_1(e)))| = I(k, d)$. Consequently, all the I(k, d) colors assigned to the vertices of $P_1(e)$ are pairwise distinct, and thus, $\phi(x), \phi(y) \in c(V(P_1(e)))$ are distinct. According to the definition of H(k, d), this implies that $\phi(x)\phi(y)$ forms an edge in H(k, d), i.e., ϕ indeed is a graph homomorphism mapping G to H(k, d).

We finally conclude the correctness of the reduction. For the NP-hardness, it remains to verify that the H(k, d)-coloring problem is NP-hard. According to Theorem 13.2, it suffices to prove that H(k, d) is non-bipartite. As $p = \frac{k}{d} \neq 2$, we either have $\frac{k}{d} < 2$, which implies $I(k, d) = \left\lceil \frac{k}{k-d} \right\rceil \geq 3$, and thus, H(k, d) contains a clique of size at least 3 and is thus not bipartite. Otherwise, we have $\frac{k}{d} > 2$ and thus, $\chi(H(k, d)) = \chi(C(k, d)) = \lceil p \rceil \geq 3$. This finally yields the claimed NP-hardness (and thus -completeness) in the case $p \neq 2$.

In the remaining case of p = 2, deciding Problem 13.4 is the same as deciding whether a given graph G fulfills va $(G) \leq 2$. However, it is not hard to see that a planar cubic 3-connected graph G admits a Hamiltonian cycle if and only if its planar dual graph G^* admits vertex arboricity at most 2 (cf. [HS89]). Consequently, the NP-hardness of the decision problem 13.4 in this case follows from the NP-hardness of the Hamiltonicity problem restricted to planar cubic 3-connected graphs [GJ90]. \Box

13.5 Conclusion

The complexity results achieved in this chapter together with the results by Feder, Hell, and Mohar [FHM03], clarify our view on fractional and circular coloring parameters related to acyclic vertex sets in digraphs and graphs in terms of computational complexity. Maybe surprisingly, for all of the considered coloring notions, deciding *p*-colorability remains NPcomplete even for values of *p* arbitrarily close to 1. Looking at related notions such as the *fractional and circular arboricity* of graphs, which can be computed in polynomial time using the Matroid Partitioning Algorithm (cf. [SU97], Chapter 5), those results show that circular and fractional arboricities behave differently with respect to complexity depending on whether they are based on vertices or edges.

Theorem 13.21 unfortunately does not treat the case p = 2. Thus, we cannot rule out the possibility that there is some clever way to algorithmically decide whether a given digraph has fractional dichromatic number at most 2. Still, the author strongly believes that it is possible to prove NP-completeness for this case as well.

A natural question left open in this chapter concerns restrictions of the treated decision problems to specialized inputs. An interesting special case consists of planar (di)graphs. It is clear that deciding the problems 13.1, 13.3 and 13.4 will now be trivially polynomial-time solvable for large values of p, as for instance, the 2-Color-Conjecture (Conjecture 8.1) states that oriented planar digraphs are 2-dichromatic, while an upper bound of 2.5 for each of the three notions studied in this chapter is known for oriented planar (di)graphs (cf. [WZLW11],[HS19]).

It appears to be hard to use the reductions provided in this chapter to achieve hardness results for planar inputs. This is mostly due to the fact that the complexity of H-coloring planar graphs is very poorly understood. While the K_4 -coloring problem is trivially contained in P (output true), only for few graphs H such as odd cycles (cf. [MS09]) hardness results are known, while for many non-trivial graphs such as the *Clebsch graph*, the H-coloring problem becomes solvable in polynomial time. Another problem is that the l-split-operation used in Section 13.3 for fractional colorings does not preserve planarity.

Nevertheless, we may deduce the following special cases:

Theorem 13.24.

- Deciding whether a given oriented planar digraph D fulfils $\vec{\chi}^*(D) \leq \frac{3}{2}$ is NP-complete.
- Deciding whether a given simple planar graph G fulfils $va_c(G) \leq \frac{3}{2}$ is NP-complete.

Proof. Both problems are clearly in NP(using acyclic/arboreal (3, 2)-colorings as certificates). Note that $H_{\vec{C}(3,2)} = K_3$ for the graph defined in the proof of Theorem 13.16 and $H(3,2) = K_3$ for the auxiliary graph defined in Section 13.4. It is easy to see that the digraph D_G defined in the proof of Theorem 13.16 as well as the graph $G_{3,2}$ as defined in the proof of Theorem 13.23 are both planar and simple for any planar and simple graph G. Deciding K_3 -colorability of planar graphs can thus be polynomially reduced to each of the above decision problems. The NP-hardness of 3-colorability of planar graphs (cf. [GJ90]) now yields the claim.

It would be furthermore interesting to study the notion of circular homomorphisms in more detail. Natural questions consider for instance descriptions of the *cores* of such homomorphisms, which enable the study of a corresponding *homomorphism order*. For graphs, this is a wide and active field of research, we refer to the book by [HN04] for a comprehensive survey of the topic.

In this context, a graph G is called a *core* if it does not admit a graph homomorphism to a proper subgraph. Equivalently, one may define a core to be a graph G such that every homomorphism $\phi: V(G) \to V(G)$ is a bijection. The interest in cores comes from their role as minimal representatives of homomorphic equivalence classes of graphs. Harutyunyan, Kayll, Mohar, and Rafferty [HKMR12] considered the following corresponding definition of digraph cores: A digraph D is called an *acyclic core*, if every acyclic homomorphism $\phi: V(D) \to V(D)$ of D to itself is a bijection. Similarly, we here define a digraph D to be a *circular core* if any circular homomorphism $\phi: V(D) \to V(D)$ is bijective. We want to conclude with some first observations concerning this notion.

Proposition 13.25.

- A graph G is a core if and only if \vec{G} is a circular core.
- If D is a circular core, then D is an acyclic core.
- For any integers $k \ge d \ge 1$, the circulant digraph $\vec{C}(k,d)$ is a circular core if and only if k and d are coprime.

Proof.

- This is a direct consequence of Proposition 13.11.
- This follows from Proposition 13.13.
- Assume for the first direction that gcd(k, d) = l > 1, let k' := k/l, d' := d/l and consider the mapping φ : Z_k → Z_{k'} defined by φ(i) := ⌊i mod k/l for all i ∈ Z_k. We claim that this defines a circular homomorphism from C(k, d) to C(k', d'). However, φ is easily seen to be an acylic (k', d')-coloring of C(k, d), and according to Proposition 13.12, this already means that φ is a circular homomorphism. As C(k', d') is isomorphic to the proper subdigraph of C(k, d) induced by the vertices il, i = 0, ..., k' − 1, this proves that C(k, d) is circularly homomorphic to a proper induced subdigraph and thus no circular core.

To prove the converse, let gcd(k,d) = 1 and assume that, contrary to the assertion, $\vec{C}(k,d)$ admits a circular homomorphism $\phi : \vec{C}(k,d) \to \vec{C}(k,d)$ which is not bijective. Let D be the subdigraph of $\vec{C}(k,d)$ induced by $Im(\phi)$. Then according to Proposition 13.14, we have $\frac{k}{d} = \vec{\chi}^*(\vec{C}(k,d)) \leq \vec{\chi}^*(D) \leq \vec{\chi}^*(\vec{C}(k,d)) = \frac{k}{d}$. However, referring to Theorem 13.6 (i), we also know that $\frac{k}{d} = \vec{\chi}^*(D)$ can be represented as a fraction with numerator at most $v(D) < v(\vec{C}(k,d)) = k$. This finally contradicts the assumption that k and d are coprime, and we deduce the claimed equivalence.

Chapter 14

Fractional Colorings of Planar Digraphs

14.1 Introduction

Graphs and digraphs considered in this chapter are simple.

In this last (and relatively short) chapter, we continue the topic of fractional colorings of digraphs which we introduced in Chapter 13. In this chapter, however, we are particularly concerned with bounds on the fractional dichromatic number of *planar digraphs*.

For undirected planar graphs, the 4-Color-Theorem yields a tight upper bound of 4 on the (circular, fractional) chromatic number. A classical example achieving tightness is the K_4 . One reason why the K_4 achieves the maximum (circular, fractional) chromatic number among all planar graphs might be that it is very compact in the sense that many edges and triangles are concentrated on a small set of vertices. In contrast, sparse graphs such as forests, which do not contain cycles at all, can be easily 2-colored. In general, when trying to color small graphs by hand, one might get the impression that in most cases the only real problem restraining us from properly coloring that graph with few colors is the existence of many short cycles and local forcing patterns which eventually yield a conflict. Conversely, one may think that the chromatic number of a graph should be small given its girth (i.e., the length of its shortest cycle) is sufficiently large. Maybe surprisingly, this intuition is well-known to be completely misleading for general graphs: Erdős [Erdő9] proved in 1959 that for every pair k, g of natural numbers there exists a graph G of chromatic number at least k and girth at least g. However, the intuition might still apply to the setting of planar graphs. Namely, a well-known result of Grötzsch [Grö59] asserts that every planar graph of girth at least four (i.e., without triangles) is properly 3-colorable. Even with larger girth requirements, one clearly cannot hope to improve this bound on the chromatic number, as shown by long odd cycles. However, we have seen in Chapter 13 that the fractional and circular chromatic number of very long odd cycles approaches 2. It is therefore reasonable to expect that there is a (smallest) function $g: \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathbb{N}$, a planar graph of girth at least g(k) satisfies $\chi_c(G) \leq 2 + \frac{1}{k}$. In that sense, Grötzsch's Theorem asserts that q(1) = 4. A challenging and quite famous open problem raised in 1984 by Jaeger states a generalization of Grötzsch's result as follows.

Conjecture 14.1 (cf. Jaeger [Jae84]). Let $k \in \mathbb{N}$. Then $g(k) \leq 4k$, i.e., every planar graph G of girth at least 4k satisfies $\chi_c(G) \leq 2 + \frac{1}{k}$.

The existence of g(k) can be quite easily observed. It is not difficult to show that every 2-connected planar graph G of sufficiently large girth contains a long subpath all whose

vertices are of degree 2, and repeatedly removing such paths one can easily construct a circular coloring inductively. The best known upper bound on g(k) for general k currently is $g(k) \leq 6k$ and follows from the paper of Lóvasz, Thomassen, Wu and Zhang [LTWZ13] on modulo k-orientations.

Also for the dichromatic number, the only "conflicts" that can occur in a vertexcoloring are directed cycles all in one color. Following the arguments from the above discussion, it is natural to expect that acyclic colorings would require fewer colors provided we exclude the existence of very short directed cycles. An extension of the negative result by Erdős [Erd59] to directed graphs was proved first by Harutyunyan and Mohar [HM12b] and later strengthened by Harutyunyan, Kayll, Mohar and Rafferty [HKMR12]: For every $k, g \in \mathbb{N}$ there exists a digraph D whose underlying graph U(D) has girth at least g but still its dichromatic number exceeds k.

Again, the situation for planar digraphs is different. That the exclusion of short directed cycles may help to bound the dichromatic number was demonstrated by Li and Mohar [LM17], who proved that every planar digraph with digirth at least 4 is acyclically 2-colorable (cf. Theorem 8.1 in Chapter 8). As the star dichromatic number of the directed k-cycle \vec{C}_k is $\frac{k}{k-1}$, it again seems reasonable to expect that planar digraphs of sufficiently large digirth should admit acyclic $(1 + \varepsilon)$ -colorings for any $\varepsilon > 0$. The following very explicit bound was conjectured in [HS19].

Conjecture 14.2. Every planar digraph D of directed girth at least $g \ge 3$ satisfies

$$\vec{\chi}^*(D) \le \frac{g-1}{g-2} = 1 + \frac{1}{g-2}.$$

Note that the statement of this conjecture means that D admits an acyclic (g-1, g-2)coloring, which in turn is equivalent to the following statement.

Conjecture 14.3. The vertices of every planar digraph D of directed girth at least $g \ge 3$ can be (g-1)-colored such that every directed cycle in D contains a vertex of every color.

The case q = 3 of the conjecture amounts exactly to the 2-Color-Conjecture.

Conjecture 14.3 should be expected to be quite difficult in general, even qualitatively. The reason for that is the following planar version of a well-known open problem in digraph theory known as *Woodall's conjecture*.

Conjecture 14.4 (cf. [Egr17]). There exists a function $f : \mathbb{N} \to \mathbb{N}$ such that the arcs of every planar digraph D of directed girth at least $f(k) \geq 3$ can be k-colored such that every directed cycle in D contains an arc of every color.

While it is easy to partition the arc-set of any digraph into two directed cycle-free subsets¹ (and hence, f(2) = 2), Conjecture 14.4, to the best of the author's knowledge, remains open already when k = 3, and is therefore expected to be quite difficult.

On the other hand, if Conjecture 14.3 holds true, then for $k \ge 2$ every planar digraph of digirth at least k + 1 admits a vertex-coloring with k colors such that every directed cycle contains vertices of all k colors. Hence, coloring every arc of the digraph with the color of its tail would give rise to a k-coloring of the arcs of D such that every directed cycle contains all colors. Hence, Conjecture 14.3 would imply Conjecture 14.4 and show that $f(k) \le k + 1$ for every $k \ge 3$.

Unfortunately, in this chapter we do not provide new results towards the resolution of Conjecture 14.4. However, we make progress towards a weakening of Conjecture 14.2, where we replace the star dichromatic number by the fractional dichromatic number.

¹Consider a linear order of the vertices of the digraph. Then the sets of forward-arcs and backward-arcs are both directed cycle-free.

Conjecture 14.5. Every planar digraph D of directed girth at least $g \ge 3$ satisfies

$$\vec{\chi}_f(D) \le \frac{g-1}{g-2} = 1 + \frac{1}{g-2}.$$

If true, then the bound claimed in Conjecture 14.2 would be optimal, as observed in [HS19]: For every $\varepsilon > 0, g \ge 3$ there exists a planar digraph D of directed girth g such that $\vec{\chi}_f(D) > \frac{g-1}{g-2} - \varepsilon$. As the main contributions of this chapter, we (1) show a qualitative version of Con-

As the main contributions of this chapter, we (1) show a qualitative version of Conjecture 14.5 and (2) verify it in a stronger form for strongly planar digraphs. Recall from Chapter 8 that a digraph D is called strongly planar if it can be embedded into the plane without crossings such that the incident out- and in-arcs of any vertex in D form intervals in the rotational order around this vertex.

Theorem 14.1. Let $g \ge 6$ and let D be a planar digraph of directed girth at least g. Then

$$\vec{\chi}_f(D) \le \frac{g}{g-5}.$$

Theorem 14.2. Let D be a strongly planar digraph of directed girth exactly $g \ge 2$. Then

$$\vec{\chi}_f(D) = \frac{g}{g-1}.$$

Interestingly, Theorem 14.2 implies that the fractional dichromatic number can be computed in polynomial time for strongly planar digraphs², which stands in contrast to the negative complexity result (Theorem 13.10) from Chapter 13.

For the proofs of Theorems 14.1 and 14.2, which are given in Section 14.2 below, we combine results from the theory of so-called *clutters* with a classical min-max-result obtained by Lucchesi and Younger.

14.2 Proofs

Let us prepare the proofs of Theorems 14.1 and 14.2 by describing some necessary background from Clutter Theory.

A *clutter* is defined to be a collection \mathcal{C} of subsets of a finite ground set S such that $C_1 \notin C_2$ for any $C_1 \neq C_2 \in \mathcal{C}$. We refer to the first chapter of [Cor01] for a short and comprehensible introduction to the topic.

Associated with any clutter C over the ground set S we have a *clutter matrix* M_C whose columns are indexed by the elements of S and whose rows correspond to the characteristic vectors of the members of C with respect to S. The following primal-dual pair (14.1), (14.2) of linear optimization programs resembles natural covering and packing problems related to clutters. Here, $w \geq 0$ denotes a fixed row vector whose entries are non-negative real numbers or possibly $+\infty$, and **1** denotes the vector with all entries equal to 1.

$$\min\left\{wx \mid x \ge 0, \ M_{\mathcal{C}}x \ge \mathbf{1}\right\} \tag{14.1}$$

$$= \max\left\{y\mathbf{1} \mid y \ge 0, \ yM_{\mathcal{C}} \le w\right\}.$$
(14.2)

=

²This is because the directed girth of a given digraph D can be computed in polynomial time as follows: For every arc $e = (u, v) \in A(D)$, apply one of the standard polynomial-time algorithms to determine the length of a shortest directed *v*-*u*-dipath in D. This number plus 1 is then the length of a shortest directed cycle through e. Taking the minimum of these values over all arcs $e \in A(D)$, we obtain the digirth of D.

In the following, we introduce a number of important notions for clutters related to integral solutions of the linear programs (14.1) and (14.2).

Given a clutter C, we will say that it admits the *Max-Flow-Min-Cut-Property* (MFMC for short) if, for any non-negative w with integral entries, there exists a primal-dual pair of integral optimal solutions to the linear programs (14.1) and (14.2).

We say that C packs if the same holds true at least for w = 1, and we call C packing if an integral primal-dual solution exists for all vectors w with entries 0, 1 or $+\infty$.

It is not hard to see that if a clutter has the MFMC-property, it is packing, and, clearly, any packing clutter also packs. While there are examples of clutters that pack but do not have the packing property, it is a famous open problem due to Conforti and Cornuejols to show that, in fact, the packing property and the MFMC-property are equivalent.

Conjecture 14.6 (Conforti and Cornuejols, cf. [Cor01]). A clutter has the packing property if and only if it has the MFMC property.

For the following, we will furthermore need the notion of *idealness* for clutters. A clutter is said to be *ideal* if, for any real-valued vector $w \ge 0$, the primal linear program (13.1) has an integral optimal solution vector x. It is not hard to show that the MFMC-property implies idealness of a clutter.

A famous example of a clutter related to digraphs is the clutter of all directed bonds of a fixed directed graph D. This clutter is actually MFMC for every digraph D, as stated by the following classical min-max-relation proved by Lucchesi and Younger.

Theorem 14.3 (Lucchesi and Younger [LY78]). Let D be a digraph and $w: A(D) \to \mathbb{N}_0$ a non-negative integral arc-weighting. Then the minimum weight of a dijoin in D equals the maximum size of a collection³ of directed bonds in D so that any arc $e \in A(D)$ is contained in at most w(e) of them.

Using planar duality of directed graphs as described e.g. in Chapter 4, Theorem 14.3 restricted to planar digraphs reformulates as follows.

Corollary 14.4. Let D be a planar digraph and $w: E(D) \to \mathbb{N}_0$ a non-negative integral arc-weighting. Then the minimum weight of a feedback arc-set in D equals the maximal size of a collection of directed cycles containing any arc $e \in A(D)$ at most w(e) times.

Corollary 14.4 in other words states that for any planar digraph D, the clutter of all arc-sets of directed cycles in D is MFMC (and thus ideal).

Corollary 14.4 can also be used to prove the following packing result.

Proposition 14.5. Let D be strongly planar. Then for any non-negative integral vertexweighting $w: V(D) \to \mathbb{N}_0$, the minimal weight of a feedback vertex set in D equals the maximal size of a collection of (induced) directed cycles in D together containing any vertex $x \in V(D)$ at most w(x) times.

Proof. We construct an auxiliary splitting-digraph⁴ D' by replacing each vertex $x \in V(D)$ by a directed arc $e_x \in A(D')$ in such a way that all the incoming arcs originally incident to x in D are now incident to $tail(e_x)$ while all the outgoing arcs of x in D are now emanating from head (e_x) . By contracting the (butterfly-contractible) arcs e_x for each $x \in V(D)$, it is clear that the directed cycles in D' are in one-to-one correspondence with the directed cycles of D. Moreover, the vertex-intersection of a pair of directed cycles in D yields a

³A collection is a multi-set, i.e., we may take the same elements multiple times into the collection.

⁴This digraph is the same as the 2-split of D as introduced in Definition 13.6

subset of the arc-intersection of the corresponding directed cycles in D'. It is furthermore easy to see from the fact that the outgoing and incoming arcs incident to any vertex in Dare separated in the cyclic ordering in any strongly planar embedding, that D' indeed is a planar digraph. We now define a corresponding weighting of the arcs of D' by setting $w'(e_x) := w(x)$ for any $x \in V(D)$ and w'(e) := M for a large natural number $M \in \mathbb{N}$ for any other arc of D'. If we choose M large enough, we find that the minimal arc-weight of a feedback arc set in D' is exactly the minimal vertex-weight of a feedback vertex set in D. Corollary 14.4 now tells us that the latter is the same as the maximal size of a collection of directed cycles in D' in which any e_x is contained at most w(x) times while any other arc is contained at most M times. As the latter condition becomes redundant for M large enough, this again is the same as the maximal size of a collection of directed cycles in Din which any vertex $x \in V(D)$ is contained at most w(x) times. As we may assume all the directed cycles in an optimal collection to be induced, this implies the claim.

As a consequence, the clutter of vertex sets of induced directed cycles of a strongly planar digraph D is also MFMC and ideal.

Given any clutter (S, \mathcal{C}) , we may define a corresponding dual clutter (called *blocking* clutter and denoted by (S, \mathcal{C}^*)) which contains all the inclusion-wise minimal subsets $X \subseteq S$ with the property that $X \cap C \neq \emptyset$ for all $C \in \mathcal{C}$. It is clear that the blocking clutter of the clutter of arc sets of directed cycles (vertex sets of induced directed cycles) of a digraph is just the clutter of inclusion-wise minimal feedback arc (vertex) sets. To proceed, we will need the following theorem of Lehman.

Theorem 14.6 (Lehman, [Leh79], and [Cor01], Theorem 1.17). A clutter is ideal if and only if its blocking clutter is.

Applying Lehman's theorem to the clutter of arc-sets of directed cycles in a planar digraph and to the clutter of vertex-sets of induced directed cycles in a strongly planar digraph, this implies that the corresponding linear optimization problems (14.1) admit integer optimal solutions $x \ge 0$ for any real-valued vector $w \ge 0$. By setting $w := \mathbf{1}^T$ in both cases, we obtain the following two results.

Corollary 14.7. Let D be a planar digraph and let g be the digitth of D. Then there is a collection F_1, \ldots, F_m of feedback arc sets of D and a weighting $y_1, \ldots, y_m \in \mathbb{R}_{\geq 0}$ such that $y_1 + \cdots + y_m = g$ and for any arc $e \in A(D)$, we have $\sum_{\{j|e \in F_i\}} y_j \leq 1$.

Corollary 14.8. Let D be a strongly planar digraph and let g be the digirth of D. Then there is a collection F_1, \ldots, F_m of feedback vertex sets of D and a weighting $y_1, \ldots, y_m \in \mathbb{R}_{\geq 0}$ such that $y_1 + \cdots + y_m = g$ and for any vertex $v \in V(D)$, we have $\sum_{\{j|v \in F_i\}} y_j \leq 1$.

In the following, let us only show how to derive Corollary 14.8, as the proof of Corollary 14.7 is completely analogous.

Proof of Corollary 14.8. Let $x \ge 0$ be an integer-valued optimal solution of the linear program (14.1) corresponding to the clutter of inclusion-wise minimal feedback vertex sets of D and $w = \mathbf{1}^T$. It is easy to see from the definition of the linear program (14.1) that, in any optimal solution, we have $x \le \mathbf{1}$ (component-wise), as otherwise one could replace x with min $\{x, \mathbf{1}\}$ and obtain a better solution to the linear program, contradicting the optimality. Consequently, we know that x has only 0 and 1 as entries and is thus determined by its support $X := \operatorname{supp}(x) \subseteq V(D)$.

From the conditions in the program (14.1) we derive that X has a common intersection with any feedback vertex set of D and thus must contain a directed cycle (as $V(D) \setminus X$ cannot be a feedback vertex set). Hence $wx = |X| \ge g$. On the other hand, the $\{0, 1\}$ -vector whose support is given by the vertex set of some directed cycle of length g clearly has value g and also satisfies the conditions of the program, which makes it an optimal solution. Consequently, also the optimal value of the dual program (14.2) is g and thus there is an optimal solution vector $y \ge 0$ with $y\mathbf{1} = g$. This implies the claim. \Box

We now conclude this section by giving the proofs of the Theorems 14.1 and 14.2.

Proof of Theorem 14.1. As $g \geq 3$, D is an oriented digraph. Let $F_1, \ldots, F_m, y_1, \ldots, y_m$ be as given by Corollary 14.7. Use the 5-degeneracy of the underlying simple planar graph U(D) to derive an ordering $v_1, \ldots, v_n, n := v(D)$ of the vertices so that for each $i \in \{1, \ldots, n\}, v_i$ has degree at most 5 in $D_i := D[v_1, \ldots, v_i]$. For each v_i , let $c(v_i)$ denote the set of $j \in \{1, \ldots, m\}$ so that v_i has an incident arc in $F_j \cap A(D_i)$. Then clearly

$$\sum_{j \in c(v_i)} y_j \le \sum_{e \in E_{D_i}(v)} \sum_{j: e \in F_j} y_j \le d_{D_i}(v_i) \le 5$$

for each v_i . Furthermore, the vertex set $X_j := \{x \in V(D) | j \notin c(x)\}$ is acyclic in D for all $j = 1, \ldots, m$: In any directed cycle C in D we find an arc contained in F_j , and thus, j is contained in at least one of the c-sets of its end-vertices.

We now define an instance of the linear optimization program (13.1) defining $\vec{\chi}_f(D)$ according to $x_A = \frac{i_A}{g-5}$, where $i_A = \sum_{j \in \{1,...,m\}: A = X_j} y_j$ for each $A \in \mathcal{A}(D)$. Then those variables are non-negative and for each vertex v, we have

$$\sum_{A \in \mathcal{A}(D,v)} i_A = \sum_{j \in \{1,\dots,m\}: v \in X_j} y_j = \sum_{j \notin c(v)} y_j = \sum_{j=1}^m y_j - \sum_{j \in c(v)} y_j \ge g - 5.$$

Hence, this is a legal instance proving $\vec{\chi}_f(D) \leq \sum_{A \in \mathcal{A}} \frac{i_A}{g-5} = \frac{\sum_{j=1}^m y_j}{g-5} = \frac{g}{g-5}$.

Proof of Theorem 14.2. Let D be strongly planar and let $g \ge 2$ denote the directed girth of D. We show that $\vec{\chi}_f(D) = \frac{g}{g-1}$. First of all, the fractional dichromatic number clearly cannot increase by taking subdigraphs, and so we have $\vec{\chi}_f(D) \ge \vec{\chi}_f(\vec{C}_g) = \frac{g}{g-1}$. It remains to prove $\vec{\chi}_f(D) \le \frac{g}{g-1}$. For this purpose we construct a feasible instance of the linear optimization program (13.1) with value at most $\frac{g}{g-1}$. To do so, let F_1, \ldots, F_m be a collection of feedback vertex sets as given by Corollary 14.8 with a corresponding weighting $y_1, \ldots, y_m \ge 0$. The complements $V(D) \setminus F_j$ are clearly acyclic for any $j \in \{1, \ldots, m\}$. For any acyclic set $A \in \mathcal{A}(D)$, we now define the value of the corresponding variable to be

$$x_A := \frac{1}{g-1} \sum_{\{j \mid A = V(D) \setminus F_j\}} y_j \ge 0.$$

We then have for any vertex $v \in V(D)$:

$$\sum_{A \in \mathcal{A}(D,v)} x_A = \frac{1}{g-1} \sum_{\{j | v \notin F_j\}} y_j = \frac{1}{g-1} \left(\underbrace{\sum_{j=1}^m y_j}_{=g} - \underbrace{\sum_{j | v \in F_j\}}_{\leq 1} y_j}_{\leq 1} \right) \ge \frac{g-1}{g-1} = 1$$

So this is indeed a feasible instance of the program (13.1) and we obtain the desired bound

$$\vec{\chi}_f(D) \le \sum_{A \in \mathcal{A}(D)} x_A = \frac{1}{g-1} \sum_{j=1}^m y_j = \frac{g}{g-1}.$$

Appendix A

Open Problems Collection

"In mathematics the art of proposing a question must be held of higher value than solving it"

David Hilbert

In this conclusive part of the thesis I would like to list some of the conjectures and open problems raised in this thesis, which I personally consider specially intriguing and interesting, and to explicitly restate them here once again.

With this problem list I hope that it will be easier for the readers to get an impression of the current state of the research on the problems discussed in this thesis without having to go through many details in the respective chapters. I would be very delighted to see progress being made on the resolution of any of the problems.

Open Problems from Chapter 2: A *subdivision* of a digraph is any digraph obtained from it by replacing its arcs with internally vertex-disjoint directed paths connecting the endpoints in the same directions.

Problem A.1. Does there exist $\alpha > 0$ such that for every integer $\ell \geq 3$ every digraph with minimum out-degree at least $\alpha \ell$ contains a subdivision of any orientation of C_{ℓ} ?

Conjecture A.1. There exists a function $f : \mathbb{N} \to \mathbb{N}$ such that the following holds for every $k \in \mathbb{N}$. If D is a digraph with minimum out-degree f(k), then there exists a digraph D' of minimum out-degree k such that D contains a subdivision of D' in which every subdivision-path has length at least two.

Conjecture A.2. There exists a function $f : \mathbb{N} \to \mathbb{N}$ such that the following holds for every integer $k \in \mathbb{N}$. If a digraph has minimum out-degree f(k), then we can partition its vertex set into two non-empty parts A and B such that every vertex in A has at least k out-neighbors in A and at least one out-neighbor in B.

Problem A.2. Does there exist $K \in \mathbb{N}$ such that every digraph with minimum out-degree K contains 3 distinct directed cycles C_1, C_2, C_3 such that C_1 and C_3 are disjoint, C_1 and C_2 share exactly one vertex u, C_2 and C_3 share exactly one vertex v, and $u \neq v$?

Problem A.3. Does there exist $K \in \mathbb{N}$ such that every strongly K-arc-connected digraph contains 3 distinct directed cycles such that they pairwise share exactly one vertex and such that the shared vertices are pairwise distinct?

Problem A.4. Does there exist $K \in \mathbb{N}$ such that every strongly K-vertex connected digraph contains two vertices $x \neq y$ and four pairwise internally vertex-disjoint dipaths, two from x to y and two from y to x?

Open Problems from Chapter 3:

Conjecture A.3. For every $k \in \mathbb{N}$ there exists $s(k) \in \mathbb{N}$ such that every strongly s(k)-vertex-connected digraph contains k vertex-disjoint directed cycles of even lengths.

Open Problems from Chapter 4:

Problem A.5. What is the computational complexity of partitioning the arc-set of a planar Eulerian digraph into the smallest number of directed cycles?

Problem A.6. What is the computational complexity of the following problem: Given a planar graph G with an Eulerian subgraph H, determine the smallest number of cycles of G such that their symmetric difference is H.

Open Problems from Chapter 5: An oriented regular matroid \vec{M} is called *non-even* if it admits a set of elements intersecting every directed circuit an odd number of times. A GB-minor of an oriented matroid is any oriented matroid obtained from it by repeatedly deleting arbitrary elements or contracting elements e which are contained in signed cocircuits of the form $(S \setminus \{e\}, \{e\})$, i.e., in which e is the only element with a negative sign. $M^*(\vec{K}_{m,n})$ denotes the oriented bond matroid of the digraph $\vec{K}_{m,n}$ obtained from the complete bipartite graph $K_{m,n}$ by directing all arcs from one to the other partite set.

Conjecture A.4. A regular oriented matroid M is non-even if and only if none of its *GB*-minors is isomorphic to $M^*(\vec{K}_{m,n})$ for some $m, n \geq 2$ such that m + n is odd.

Problem A.7. Is there a polynomial-time algorithm that, given as input a digraph D, decides whether or not D contains an inclusion-wise minimal directed cut of even size?

The dichromatic number $\vec{\chi}(D)$ of a digraph D is the smallest number of colors required in an acyclic coloring, i.e., a vertex-coloring avoiding monochromatic directed cycles.

Open Problems from Chapter 6: We say that a digraph D contains another digraph H as a *strong minor* if H can be obtained from G by repeatedly deleting vertices, arcs or contracting directed cycles into single vertices.

Conjecture A.5. Every digraph D excluding $\overset{\leftrightarrow}{K}_t$ as a strong minor satisfies $\vec{\chi}(D) \leq t$.

A digraph is said to be k-choosable if for every assignment of color-lists of size k to its vertices their is a choice of colors from the lists resulting in an acyclic coloring.

Problem A.8. Does there exist $f : \mathbb{N} \to \mathbb{N}$ such that every digraph D excluding K_t as a strong minor is f(t)-choosable?

Open Problems from Chapter 7: Given a digraph F, mader $_{\vec{\chi}}(F)$ denotes the smallest integer k such that every digraph D satisfying $\vec{\chi}(D) \ge k$ contains a subdivision of F.

Conjecture A.6. There exists an absolute constant c > 0 such that $mader_{\vec{\chi}}(\vec{K}_n) \leq cn^2$ for every positive integer n.

We say that a digraph F is *Mader-perfect* if mader_{$\vec{\chi}$}(F') equals the number of vertices of F' for every subdigraph $F' \subseteq F$.

Problem A.9. Give a structural characterization of Mader-perfect digraphs.

Conjecture A.7. The maximum number m(k) of arcs in a Mader-perfect digraph on k vertices satisfies $m(k) = \Theta(k)$.

Conjecture A.8. Every digraph D with $\vec{\chi}(D) \ge 4$ contains three distinct directed cycles sharing pairwise exactly one vertex such that the shared vertices are pairwise distinct.

Conjecture A.9. Let D be a digraph with $\vec{\chi}(D) \geq 3$. Then there is $\ell \geq 3$ such that D contains a cyclic chain of ℓ directed cycles as follows: There are directed cycles C_1, \ldots, C_ℓ and pairwise distinct vertices v_1, \ldots, v_ℓ such that $V(C_i) \cap V(C_{i+1}) = \{v_i\}, i = 1, \ldots, \ell$ (where $\ell + 1$ is identified with 1) and such that $V(C_i) \cap V(C_j) = \emptyset$ whenever i and j are non-consecutive in the cyclical ordering $1, 2, \ldots, \ell, 1$.

Open Problems from Chapter 8: An arc in a digraph is called *contractible* if it is the only arc leaving its tail or the only arc entering its head. We say that a digraph D contains another digraph H as a *butterfly-minor* if H can be obtained from D by a sequence of vertex-deletions, arc-deletions and contractions of contractible arcs.

Conjecture A.10. Every digraph D excluding \vec{K}_t as a butterfly-minor satisfies $\vec{\chi}(D) \leq t$.

Problem A.10. Is there a function $f : \mathbb{N} \to \mathbb{N}$ such that every digraph D excluding \vec{K}_t as a butterfly-minor is f(t)-choosable?

Let G be a graph and $H \subseteq G$. We say that H is a conformal subgraph of G if both G - V(H) and H have perfect matchings. A matching minor of a graph G is any graph obtainable from a conformal subgraph of G by repeatedly picking a vertex of degree two and simultaneously contracting its two incident edges.

Given a graph G equipped with a perfect matching M, a proper M-coloring of G is a coloring of the edges of M such that no M-alternating cycle in G uses only matching edges of the same color. By $\chi(G, M)$ we denote the minimum number of colors that can be used for a proper M-coloring.

Conjecture A.11. Let $t \in \mathbb{N}$, let G be a bipartite graph and M an arbitrary perfect matching of G such that $\chi(G, M) \geq t$. Then G contains $K_{t,t}$ as a matching minor.

Let $k \in \mathbb{N}$. A graph G is called k-extendable if for every matching S in G with $|S| \leq k$ there exists a perfect matching M of G such that $S \subseteq M$.

Problem A.11. Is there a function $f : \mathbb{N} \to \mathbb{N}$ such that every f(t)-extendable bipartite graph G contains $K_{t,t}$ as a matching minor?

Open Problems from Chapter 9: A tournament H is called a *hero* if every tournament with sufficiently large dichromatic number contains H as a subdigraph. Heroes were characterized in [BCC⁺13]. A digraph is called *locally out-complete* if the out-neighbors of any vertex are pairwise adjacent.

Conjecture A.12. If H is a hero, then every locally-out-complete oriented graph of sufficiently large dichromatic number contains H as a subdigraph.

Conjecture A.13. Every locally-out-complete oriented graph of dichromatic number 4 contains \vec{K}_4^s , the unique strongly connected tournament on four vertices, as a subdigraph.

Open Problems from Chapter 11: A complete arboreal coloring of an undirected graph G is a vertex-coloring without monochromatic cycles such that for every pair of colors i, j used in the coloring, there is a cycle in G using only colors i and j. The *a*-vertex arboricity $\operatorname{ava}(G)$ is the largest number of colors that can be used in a complete arboreal coloring of G. In Chapter 11 we prove that for every $k \in \mathbb{N}$ there exists a smallest integer f(k) such that every simple graph G with $\operatorname{ava}(G) \leq k$ admits a set of at most f(k) vertices whose deletion turns G into a forest. Our best lower bound on f(k) is quadratic and the best upper-bound non-polynomial.

Problem A.12. Determine the asymptotic growth of f(k).

Open Problems from Chapter 12:

Problem A.13. Let H be a 6-regular 6-uniform hypergraph (that is, every hyperedge contains 6 vertices and every vertex is contained in 6 hyperedges). Is there a 3-coloring of the vertices of H such that no hyperedge contains four vertices of the same color?

Problem A.14. Let D be a 5-regular digraph, that is, every vertex has in- and out-degree exactly 5. Is there a 5-coloring of the vertices of D such that every vertex in D sees at most one vertex of his own color?

Open Problems from Chapter 14:

Conjecture A.14. Let D be a planar digraph in which every directed cycle has length at least $g \in \mathbb{N}$. Then D admits a (g-1)-vertex-coloring such that every directed cycle in D contains a vertex of every color.

Bibliography

[AA17]	M. Aguiar and F. Ardila. Hopf monoids and generalized permutahedra. arXiv preprint, 2017. arXiv:1709.07504.
[AAHV07]	O. Aichholzer, F. Aurenhammer, C. Huemer, and B. Vogtenhuber. Gray code enumeration of plane straight-line graphs. <i>Graphs and Combinatorics</i> , 23(5):467–479, 2007.
[ABG17]	M. Anholcer, B. Bosek, and J. Grytczuk. Majority choosability of di- graphs. <i>The Electronic Journal of Combinatorics</i> , 24(3), 2017. Article number P3.57.
[ACH ⁺ 19]	P. Aboulker, N. Cohen, F. Havet, W. Lochet, P. Moura, and S. Thomassé. Subdivisions in digraphs of large out-degree or large dichromatic number. <i>The Electronic Journal of Combinatorics</i> , 26(3), 2019. Article Number P3.19.
[ACH ⁺ 21]	O. Aichholzer, J. Cardinal, T. Huynh, K. Knauer, R. Steiner, and B. Vogtenhuber. Flip distances between graph orientations. <i>Algorithmica</i> , 83(1):116–143, 2021.
[ACN20]	P. Aboulker, P. Charbit, and R. Naserasr. Extension of Gyárfás-Sumner conjecture to digraphs. <i>arXiv preprint</i> , 2020. arXiv:2009.13319.
[AF96]	D. Avis and K. Fukuda. Reverse search for enumeration. Discrete Applied Mathematics, $65(1-3)$:21–46, 1996.
[AGSW20]	M. Axenovich, A. Girão, R. Snyder, and L. Weber. Strong complete minors in digraphs. <i>arXiv preprint</i> , 2020. arXiv:2010.05643.
[AH76]	K. Appel and W. Haken. Every planar map is four colorable. I. Discharging. <i>Illinois Journal of Mathematics</i> , 21(3):429–490, 1976.
[AH15]	S. D. Andres and W. Hochstättler. Perfect digraphs. <i>Journal of Graph Theory</i> , 79:21–29, 2015.
[AHK77]	K. Appel, W. Haken, and J. Koch. Every planar map is four colorable. II. Reducibility. <i>Illinois Journal of Mathematics</i> , 21(3):491–567, 1977.
[AHLS03]	R. E. L. Aldred, D. A. Holton, D. Lou, and A. Saito. <i>m</i> -alternating paths in <i>n</i> -extendable bipartite graphs. <i>Discrete Mathematics</i> , 269(1–3):1–11, 2003.
[AKKW16]	S. Amiri, K. Kawarabayashi, S. Kreutzer, and P. Wollan. The Erdős-Pósa-property for directed graphs. <i>arXiv preprint</i> , 2016. arXiv:1603.02504.

[AL89]	N. Alon and N. Linial. Cyles of length zero modulo k in directed graphs. Journal of Combinatorial Theory, Series B, 47(1):114–119, 1989.
[Alo96]	N. Alon. Disjoint directed cycles. Journal of Combinatorial Theory, Series B, 68(2):167–178, 1996.
[Alo06]	N. Alon. Splitting digraphs. Combinatorics, Probability and Computing, 15(6):933–937, 2006.
[ALSS19]	M. Anastos, A. Lamaison, R. Steiner, and T. Szabó. Majority colorings of sparse digraphs. <i>arXiv preprint</i> , 2019. arXiv:1911.01954, submitted.
[AMP15]	O. Aichholzer, W. Mulzer, and A. Pilz. Flip distance between triangu- lations of a simple polygon is NP-complete. <i>Discrete & Computational Geometry</i> , 54(2):368–389, 2015.
[APMBORM18]	G. Araujo-Pardo, J. J. Montellano-Ballesteros, M. Olsen, and C. Rubio- Montiel. The diachromatic number of digraphs. <i>The Electronic Journal</i> of <i>Combinatorics</i> , 25(3), 2018. Article number P3.51.
[BCC ⁺ 13]	E. Berger, K. Choromanski, M. Chudnovsky, J. Fox, M. Loebl, A. D. Scott, P. Seymour, and S. Thomassé. Tournaments and colouring. <i>Journal of Combinatorial Theory, Series B</i> , 103(1):1–20, 2013.
[BF12]	O. Bernardi and É. Fusy. A bijection for triangulations, quadrangulations, pentagulations, etc. Journal of Combinatorial Theory, Series A, 119(1):218–244, 2012.
$[BFJ^+04]$	D. Bokal, G. Fijavz, M. Juvan, P. Kayll, and B. Mohar. The circular chromatic number of a digraph. <i>Journal of Graph Theory</i> , 46:227–240, 2004.
[BH09]	P. Bose and F. Hurtado. Flips in planar graphs. Computational Geometry, $42(1):60-80$, 2009.
[BHL ⁺ 17]	J. Bensmail, A. Harutyunyan, N. K. Le, B. Li, and N. Lichiardopol. Disjoint cycles of different lengths in graphs and digraphs. <i>The Electronic Journal of Combinatorics</i> , 24(4), 2017. Article Number P4.37.
[BHL18]	J. Bensmail, A. Harutyunyan, and N. Le. List coloring digraphs. <i>Journal of Graph Theory</i> , 87:492–508, 2018.
[BJBHY18]	J. Bang-Jensen, S. Bessy, F. Havet, and A. Yeo. Out-degree reducing partitions of digraphs. <i>Theoretical Computer Science</i> , 719:64–72, 2018.
[BJG08]	J. Bang-Jensen and G. Z. Gutin. <i>Digraphs: Theory, Algorithms and Applications</i> . Springer, 2008.
[BKV19]	S. Blind, K. Knauer, and P. Valicov. Enumerating <i>k</i> -arc-connected orientations. <i>arXiv preprint</i> , 2019. arXiv:1908.02050.
[BLV78]	R. G. Bland and M. Las Vergnas. Orientability of matroids. <i>Journal of Combinatorial Theory, Series B</i> , 24(1):94–123, 1978.
[BLVS ⁺ 99]	A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. Oriented Matroids, volume 46 of Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2nd edition, 1999.
------------------------	---
[BR89]	A. Bachem and A. Reinhold. On the complexity of the Farkas-property of oriented matroids. <i>Technical Preprint</i> , 1989. http://e-archive.informatik.uni-koeln.de/65.
[Bre00]	 E. Brehm. 3-Orientations and Schnyder-3-tree-decompositions, 2000. Diploma Thesis, Freie Universität Berlin, https://page.math. tu-berlin.de\T1\guilsinglrightbrehm.ps.gz.
[Bru69]	R. A. Brualdi. Comments on bases in dependence structures. <i>Bulletin of the Australian Mathematical Society</i> , 1:161–167, 1969.
[BT81]	J. C. Bermond and C. Thomassen. Cycles in digraphs–a survey. <i>Journal of Graph Theory</i> , 5:1–43, 1981.
[BT96]	B. Bollobás and A. Thomason. Highly linked graphs. <i>Combinatorica</i> , 16(3):313–320, 1996.
[BT98]	B. Bollobás and A. Thomason. Proof of a conjecture by Mader, Erdős and Hajnal on topological complete subgraphs. <i>European Journal of Combinatorics</i> , 19(8):883–887, 1998.
[Buc18]	M. Bucić. An improved bound for disjoint directed cycles. <i>Discrete Mathematics</i> , 341(8):2231–2236, 2018.
[Bur80]	S. Burr. Subtrees of directed graphs and hypergraphs. <i>Congressus Numerantium</i> , 28:227–239, 1980.
[Bur82]	S. Burr. Antidirected subtrees of directed graphs. Canadian Mathemat- ical Bulletin, 25(1):119–120, 1982.
[BV11]	P. Bose and S. Verdonschot. A history of flips in combinatorial triangula- tions. In <i>Proceedings Computational Geometry, EGC 2011</i> , pages 29–44. Springer, LNCS, 2011.
[Cat79]	P. A. Catlin. Hajós' graph coloring conjecture: variations and coun- terexamples. <i>Journal of Combinatorial Theory, Series B</i> , 26(2):268–274, 1979.
[CC11]	Z. Che and Z. Chen. Forcing on perfect matchings–a survey. <i>Match</i> , 66(1):93–136, 2011.
[CFKS14]	S. Chiba, S. Fujita, K. Kawarabayashi, and T. Sakuma. Minimum degree conditions for vertex-disjoint even cycles in large graphs. <i>Advances in Applied Mathematics</i> , 54:105–120, 2014.
[CHLN18]	N. Cohen, F. Havet, W. Lochet, and N. Nisse. Subdivisions of oriented cycles in digraphs with large chromatic number. <i>Journal of Graph Theory</i> , 89(4):439–456, 2018.
[CLL+08]	J. Chen, Y. Liu, S. Lu, B. O'Sullivan, and I. Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. <i>Journal of the</i> ACM , 55(5), 2008. Article number 21.

[Cor01]	G. Cornuéjols. Combinatorial Optimization: Packing and Covering, vol- ume 74 of CBMS-NSF Regional Conference Series in Applied Mathemat- ics. SIAM, 2001.
[Cou90]	B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. <i>Information and Computation</i> , 85(1):12–75, 1990.
[CSJ09]	S. Cleary and K. St. John. Rotation distance is fixed-parameter tractable. Information Processing Letters, 109(16):918–922, 2009.
[CSJ10]	S. Cleary and K. St. John. A linear-time approximation for rotation distance. <i>Journal of Graph Algorithms and Applications</i> , 14(2):385–390, 2010.
[CSS18]	J. Cardinal, V. Sacristán, and R. I. Silveira. A note on flips in diag- onal rectangulations. <i>Discrete Mathematics and Theoretical Computer</i> <i>Science</i> , 20(2), 2018.
[CSS19a]	M. Chudnovsky, A. D. Scott, and P. D. Seymour. Induced subgraphs of graphs with large chromatic number. XI. Orientations. <i>European Journal of Combinatorics</i> , 76:53–61, 2019.
[CSS19b]	M. Chudnovsky, A. D. Scott, and P. D. Seymour. Induced subgraphs of graphs with large chromatic number. XII. Distant stars. <i>Journal of Graph Theory</i> , 92(3):237–254, 2019.
[CSZ15]	C. Ceballos, F. Santos, and G. M. Ziegler. Many non-equivalent realizations of the associahedron. <i>Combinatorica</i> , 35:513–551, 2015.
[CY18]	S. Chiba and T. Yamashita. Degree conditions for the existence of vertex- disjoint cycles and paths: A survey. <i>Graphs and Combinatorics</i> , 34:1–83, 2018.
[dCLM02a]	M. H. de Carvalho, C. L. Lucchesi, and U. S. R. Murty. How to build a brick. <i>Discrete Mathematics</i> , 306(19–20):2383–2410, 2002.
[dCLM02b]	M. H. de Carvalho, C. L. Lucchesi, and U. S. R. Murty. On a conjecture of Lovász concerning bricks, I. The characteristic of a matching covered graph. <i>Journal of Combinatorial Theory, Series B</i> , 85(1):94–136, 2002.
[DF99]	R. G. Downey and M. R. Fellows. <i>Parameterized Complexity</i> . Monographs in Computer Science. Springer, 1st edition, 1999.
[dFOdM01]	H. de Fraysseix and P. Ossona de Mendez. On topological aspects of orientations. <i>Discrete Mathematics</i> , 229(1):57–72, 2001.
[Die17]	R. Diestel. <i>Graph Theory</i> , volume 173 of <i>Graduate Texts in Mathematics</i> . Springer, 5th edition, 2017.
[Dir52]	G. A. Dirac. A property of 4-chromatic graphs and some remarks on critical graphs. <i>Journal of the London Mathematical Society</i> , 27:85–92, 1952.
[DKMR11]	D. Dellamonica, V. Koubek, D. M. Martin, and V. Rödl. On a conjecture of Thomassen concerning subgraphs of large girth. <i>Journal of Graph Theory</i> , 67(4):316–331, 2011.

BIBLIOGRAPHY

[DLRS10]	J. A. De Loera, J. Rambau, and F. Santos. <i>Triangulations: Structures for Algorithms and Applications</i> . Number 25 in Algorithms and Computation in Mathematics. Springer, 2010.
[DMMS12]	M. DeVos, J. McDonald, B. Mohar, and D. Scheide. Immersing complete digraphs. <i>European Journal of Combinatorics</i> , 33(6):1294–1302, 2012.
[DP02]	B. A. Davey and H. A. Priestley. <i>Introduction to Lattices and Order</i> . Cambridge University Press, New York, 2nd edition, 2002.
[Edw97]	K. J. Edwards. The harmonious chromatic number and the achromatic number. <i>Surveys in Combinatorics 1997</i> , pages 13–47, 1997.
[Edw13]	K. J. Edwards. Harmonious chromatic number of directed graphs. <i>Discrete Applied Mathematics</i> , 161(3):369–376, 2013.
[EF81]	P. Erdős and S. Fajtlowicz. On the conjecture of Hajós. <i>Combinatorica</i> , 1(2):141–143, 1981.
[Ega96]	Y. Egawa. Vertex-disjoint cycles of the same length. Journal of Combi- natorial Theory, Series B, 66:168–200, 1996.
[Egr17]	Egres Open. Woodall's conjecture, 2017. Retrieved from http://lemon.cs.elte.hu/egres/open/Woodall%27s_conjecture.
[EP65]	P. Erdős and L. Pósa. On independent circuits contained in a graph. Canadian Journal of Mathematics, 17:347–352, 1965.
[Erd59]	P. Erdős. Graph theory and probability. Canadian Journal of Mathe- matics, 11:34–38, 1959.
[Erd80]	P. Erdős. Problems and results in number theory and graph theory. Pro- ceedings of the Ninth Manitoba Conference on Numerical Mathematics and Computing, pages 3-21, 1980. https://www.renyi.hu/~p_erdos/ 1980-30.pdf.
[Fel04]	S. Felsner. Lattice structures from planar graphs. <i>The Electronic Journal of Combinatorics</i> , 11(1), 2004. Research Article R15.
[Fel13]	S. Felsner. Rectangle and square representations of planar graphs. In <i>Thirty Essays on Geometric Graph Theory</i> , pages 214–248. Springer, New York, 2013.
[FHHM86]	M. Farber, G. Hahn, P. Hell, and D. Miller. Concerning the achromatic number of graphs. <i>Journal of Combinatorial Theory, Series B</i> , 40:21–39, 1986.
[FHKS20]	S. Felsner, W. Hochstättler, K. Knauer, and R. Steiner. Complete acyclic colorings. <i>The Electronic Journal of Combinatorics</i> , 27(2), 2020. Article Number P2.40.
[FHM03]	T. Feder, P. Hell, and B. Mohar. Acyclic homomorphisms and circular colorings of digraphs. <i>SIAM Journal on Discrete Mathematics</i> , 17(1):161–169, 2003.

292	BIBLIOGRAPHY
[FK09]	S. Felsner and K. Knauer. ULD-lattices and Δ -bonds. Combinatorics, Probability and Computing, 18(5):707–724, 2009.
[FK11]	S. Felsner and K. Knauer. Distributive lattices, polyhedra, and general- ized flows. <i>European Journal of Combinatorics</i> , 32(1):45–59, 2011.
[FKMS18]	S. Felsner, L. Kleist, T. Mütze, and L. Sering. Rainbow cycles in flip graphs. In 34th International Symposium on Computational Geometry, SoCG 2018, pages 38:1–38:14, 2018.
[FLM ⁺ 08]	 M. R. Fellows, D. Lokshtanov, N. Misra, F. A. Rosamond, and S. Saurabh. Graph layout problems parameterized by vertex cover. In Proceedings of the 19th International Symposium on Algorithms and Computation, ISAAC 2008, pages 294–305. Springer, 2008.
[FLS13]	J. Fox, C. Lee, and B. Sudakov. Chromatic number, clique subdivisions, and the conjectures of Hajós and Erdős-Fajtlowicz. <i>Combinatorica</i> , 33:181–197, 2013.
[FSS18a]	S. Felsner, H. Schrezenmaier, and R. Steiner. Equiangular contact repre- sentations. In <i>Graph-Theoretic Concepts in Computer Science - 44th In-</i> <i>ternational Workshop, WG 2018, Proceedings</i> , pages 203–215. Springer, 2018.
[FSS18b]	S. Felsner, H. Schrezenmaier, and R. Steiner. Pentagon contact representations. <i>The Electronic Journal of Combinatorics</i> , 25(3), 2018. Article Number P3.39.
[FT88]	A. Frank and É. Tardos. Generalized polymatroids and submodular flows. <i>Mathematical Programming</i> , 42(3, Ser. B):489–563, 1988.
[GHK ⁺ 10]	R. Ganian, P. Hliněný, J. Kneis, D. Meister, J. Obdržálek, P. Ross- manith, and S. Sikdar. Are there any good digraph width measures? In <i>Proceedings of the International Symposium on Parameterized and Exact</i> <i>Computation</i> , pages 135–146. Springer, 2010.
[GJ90]	M. R. Garey and D. S. Johnson. <i>Computers and Intractability: A Guide to the Theory of NP-Completeness.</i> W. H. Freeman & Co., New York, USA, 1990.
[GK69]	C. G. and H. V. Kronk. The point-arboricity of planar graphs. <i>Journal of the London Mathematical Society</i> , s1-44(1):612–616, 1969.
[GKKK20]	A. C. Giannopoulou, K. Kawarabayashi, S. Kreutzer, and Oj. Kwon. The directed flat wall theorem. In <i>Proceedings of the Thirty-First Annual</i> <i>ACM-SIAM Symposium on Discrete Algorithms SODA 2020</i> , pages 239– 258. Association for Computing Machinery, 2020.
[GKP17]	A. Girão, T. Kittipassorn, and K. Popielarz. Generalized majority colour- ings of digraphs. <i>Combinatorics, Probability and Computing</i> , 26(6):850– 855, 2017.
[GKR20]	F. Gurski, D. Komander, and C. Rehs. Acyclic coloring of special di- graphs. <i>arXiv preprint</i> , 2020. arXiv:2006.13911.

[GKW68]	C. G., H. V. Kronk, and C. E. Wall. The point-arboricity of a graph. <i>Israel Journal of Mathematics</i> , 6(2):169–175, 1968.
[GL81]	P. Gács and L. Lovász. Khachiyan's algorithm for linear programming. Mathematical Programming Studies, 14:61–68, 1981.
[GL86]	P. M. Gilmer and R. A. Litherland. The duality conjecture in formal knot theory. Osaka Journal of Mathematics, 23(1):229–247, 1986.
[GLO13]	J. Gajarský, M. Lampis, and S. Ordyniak. Parameterized algorithms for modular-width. In <i>Proceedings of Parameterized and Exact Computation</i> , <i>IPEC 2013</i> , volume 8246 of <i>Lecture Notes in Computer Science</i> , pages 163–176. Springer, 2013.
[GLP12]	D. Gonçalves, B. Lévêque, and A. Pinlou. Triangle contact representa- tions and duality. <i>Discrete & Computational Geometry</i> , 48(1):239–254, 2012.
[GM13]	Y. Gao and D. Ma. Disjoint cycles with different length in 4-arc- dominated digraphs. <i>Operations Research Letters</i> , 41(6):650–653, 2013.
[Gör00]	F. Göring. Short proof of Menger's theorem. <i>Discrete Mathematics</i> , 219:295–296, 2000.
[GPS21]	A. Girão, K. Popielarz, and R. Snyder. Subdivisions of digraphs in tour- naments. <i>Journal of Combinatorial Theory, Series B</i> , 146:266–285, 2021.
[Grö59]	H. Grötzsch. Ein Dreifarbensatz für dreikreisfreie Netze auf der Kugel. Wissenschaftliche Zeitschrift, Martin-Luther-Universität Halle-Wittenberg, Mathematisch-naturwissenschaftliche Reihe, 8:109– 120, 1959.
[GSS20a]	L. Gishboliner, R. Steiner, and T. Szabó. Dichromatic number and forced subdivisions. <i>arXiv preprint</i> , 2020. arXiv:2008.09888, submitted.
[GSS20b]	L. Gishboliner, R. Steiner, and T. Szabó. Oriented cycles in digraphs of large out-degree. <i>arXiv preprint</i> , 2020. arXiv:2008.13224, submitted.
[GT11]	B. Guenin and R. Thomas. Packing directed circuits exactly. <i>Combinatorica</i> , 31(4):397–421, 2011.
[Gue01]	B. Guenin. Circuit Mengerian directed graphs. In <i>Integer Programming</i> and Combinatorial Optimization, pages 185–195. Springer, 2001.
[Gyá75]	A. Gyárfás. On Ramsey covering-numbers. Infinite and finite sets, Colloquia Mathematica Societatis János Bolyai, 10:801–806, 1975.
[Had43]	H. Hadwiger. Über eine Klassifikation der Streckenkomplexe. Viertel- jahrsschrift der Naturforschenden Gesellschaft in Zürich, 88:133–143, 1943.
[Häg85]	R. Häggkvist. Equicardinal disjoint cycles in sparse graphs. Annals of Discrete Mathematics, 27:269–274, 1985.

[HHNOP09]	C. Huemer, F. Hurtado, M. Noy, and E. Omaña-Pulido. Gray codes for non-crossing partitions and dissections of a convex polygon. <i>Discrete</i> <i>Applied Mathematics</i> , 157(7):1509–1520, 2009.
[HHNRC05]	M. E. Houle, F. Hurtado, M. Noy, and E. Rivera-Campo. Graphs of triangulations and perfect matchings. <i>Graphs and Combinatorics</i> , 21(3):325–331, 2005.
[HHP67]	F. Harary, S. Hedetniemi, and G. Prins. An interpolation theorem for graphical homomorphisms. <i>Portugaliae Mathematica</i> , 26:453–462, 1967.
[HK81]	D. Hausmann and B. Korte. Algorithmic versus axiomatic definitions of matroids. <i>Mathematical Programming Studies</i> , 14:98–111, 1981.
[HKL20]	J. Haslegrave, J. Kim, and H. Liu. Extremal density for sparse minors and subdivisions. <i>arXiv preprint</i> , 2020. arXiv:2012.02159.
[HKMR12]	A. Harutyunyan, P. Kayll, B. Mohar, and L. Rafferty. Uniquely <i>d</i> -colourable digraphs with large girth. <i>Canadian Journal of Mathematics</i> , 64(6):1310–1328, 2012.
[HLM20]	A. Harutyunyan, M. Lampis, and N. Melissinos. Digraph coloring and distance to acyclicity. <i>arXiv preprint</i> , 2020. arXiv:2010.06317.
[HLTW19]	A. Harutyunyan, TN. Le, S. Thomassé, and H. Wu. Coloring tourna- ments: From local to global. <i>Journal of Combinatorial Theory, Series</i> <i>B</i> , 138:166–171, 2019.
[HM97]	J. Hughes and G. MacGillivray. The achromatic number of graphs: A survey and some new results. <i>Bulletin of the Institute of Combinatorics and its Applications</i> , 19:27–56, 1997.
[HM12a]	A. Harutyunyan and B. Mohar. Planar graphs have exponentially many 3-arboricities. <i>SIAM Journal on DIscrete Mathematics</i> , 26(3):1269–1280, 2012.
[HM12b]	A. Harutyunyan and B. Mohar. Two results on the digraph chromatic number. <i>Discrete Mathematics</i> , 312(10):1823–1826, 2012.
[HM17]	A. Harutyunyan and B. Mohar. Planar digraphs of digirth five are 2-colorable. <i>Journal of Graph Theory</i> , 84:408–427, 2017.
[HMM18]	F. Havet, A. K. Maia, and B. Mohar. Finding a subdivision of a pre- scribed digraph of order 4. <i>Journal of Graph Theory</i> , 87(4):536–560, 2018.
[HN90]	P. Hell and J. Nešetřil. On the complexity of H -colouring. Journal of Combinatorial Theory, Series B, $48(1)$:92–110, 1990.
[HN04]	P. Hell and J. Nešetřil. <i>Graphs and Homomorphisms</i> , volume 28 of <i>Oxford Lecture Series in Mathematics and its Applications</i> . Oxford University Press, 2004.
[HS89]	S. L. Hakimi and E. F. Schmeichel. A note on the vertex arboricity of a graph. <i>SIAM Journal on Discrete Mathematics</i> , 2(1):64–67, 1989.

[HS19]	W. Hochstättler and R. Steiner. The star dichromatic number. <i>Discussiones Mathematicae Graph Theory</i> , 2019.
[HSS20]	W. Hochstättler, F. Schröder, and R. Steiner. On the complexity of di- graph colourings and vertex arboricity. <i>Discrete Mathematics and The-</i> <i>oretical Computer Science</i> , 22(1), 2020.
[HSW20]	K. Heuer, R. Steiner, and S. Wiederrecht. Even circuits in oriented matroids. <i>arXiv preprint</i> , 2020. arXiv:2010.08988, submitted.
[HY02]	M. Henning and A. Yeo. Graphs of non-crossing perfect matchings. <i>Graphs and Combinatorics</i> , 18(3):517–532, 2002.
[HY12]	M. Henning and A. Yeo. Vertex disjoint cycles of different length in digraphs. <i>SIAM Journal on Discrete Mathematics</i> , 26(2):687–694, 2012.
[IP01]	R. Impagliazzo and R. Paturi. On the complexity of k-SAT. Journal of Computer and System Sciences, 62(2):367–375, 2001.
[IPZ01]	R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? <i>Journal of Computer and System Sciences</i> , 63(4):512–530, 2001.
[Iwa02]	S. Iwata. On matroid intersection adjacency. <i>Discrete Mathematics</i> , 242(1–3):277–281, 2002.
[Jae84]	F. Jaeger. On circular flows in graphs. <i>Finite and Infinite Sets, Colloquia Mathematica Societatis János Bolyai</i> , 37:391–402, 1984.
[Jag96]	C. Jagger. An extremal function for digraph subcontraction. <i>Journal of Graph Theory</i> , 21(3):343–350, 1996.
[JRST01]	T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. Directed tree-width. <i>Journal of Combinatorial Theory, Series B</i> , 82(1):138–154, 2001.
[Kas67]	P. Kasteleyn. Graph theory and crystal physics. Discrete & Computa- tional Geometry, pages 43–110, 1967.
[Kha79]	L. G. Khačijan. A polynomial algorithm in linear programming. <i>Doklady Akademii Nauk</i> , 244(5):1093–1096, 1979.
[KK15]	K. Kawarabayashi and S. Kreutzer. The directed grid theorem. In Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, pages 655–664. Association for Computing Machinery, 2015.
[KLM84]	V. Klee, R. Ladner, and R. Manber. Signsolvability revisited. <i>Linear Algebra and its Applications</i> , 59:131–157, 1984.
[KM16]	N. Kothari and U. S. R. Murty. K_4 -free and $\overline{C_6}$ -free planar matching covered graphs. Journal of Graph Theory, 82(1):5–32, 2016.
[Kna07]	K. Knauer. Partial orders on orientations via cycle flips. Master's thesis, Faculty of Mathematics, TU Berlin, 2007.

[Kna08]	K. Knauer. Distributive lattices on graph orientations. In <i>Semigroups, acts and categories with applications to graphs</i> , volume 3, pages 79–91. Estonian Mathematical Society, 2008.
[KO04]	D. Kühn and D. Osthus. Induced subdivisions in $K_{s,s}$ -free graphs of large average degree. <i>Combinatorica</i> , 24(2):287–304, 2004.
[Kos84]	A. V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. Combinatorica, $4(4)$:307–316, 1984.
[KOS ⁺ 17]	S. Kreutzer, Si. Oum, P. Seymour, D. van der Zypen, and D. Wood. Majority colourings of digraphs. <i>The Electronic Journal of Combinatorics</i> , 24(2), 2017. Article Number P2.25.
[KP94]	H. A. Kierstead and S. G. Penrice. Radius two trees specify χ -bounded classes. Journal of Graph Theory, 18(2):119–129, 1994.
[KS96]	J. Komlós and E. Szemerédi. Topological cliques in graphs II. Combina- torics, Probability and Computing, 5(1):79–90, 1996.
[KS15]	I. Kim and P. D. Seymour. Tournament minors. <i>Journal of Combinato-</i> rial Theory, Series B, 112:138–153, 2015.
[KS18]	F. Knox and R. Sámal. Linear bound for majority colourings of digraphs. <i>The Electronic Journal of Combinatorics</i> , 25(3), 2018. Article number P3.29.
[KSX17]	I. Kanj, E. Sedgwick, and G. Xia. Computing the flip distance be- tween triangulations. <i>Discrete & Computational Geometry</i> , 58(2):313– 344, 2017.
[Kur30]	K. Kuratowski. Sur le problème des courbes gauches en topologie. <i>Fundamenta Mathematicae</i> , 15:271–283, 1930.
[LDCKM18]	C. L. Lucchesi, M. H. De Carvalho, N. Kothari, and U. S. R. Murty. On two unsolved problems concerning matching covered graphs. <i>SIAM Journal on Discrete Mathematics</i> , 32(2):1478–1504, 2018.
[LDCM15]	C. L. Lucchesi, M. H. De Carvalho, and U. S. R. Murty. Thin edges in braces. <i>The Electronic Journal of Combinatorics</i> , 22(4), 2015. Article number P4.14.
[Leh79]	A. Lehman. On the width-length inequality. <i>Mathematical Programming</i> , 16(1):245–259, 1979.
[LEP10]	F. Luccio, A. M. Enriquez, and L. Pagli. Lower bounds on the rotation distance of binary trees. <i>Information Processing Letters</i> , 110(21):934–938, 2010.
[Lic14]	N. Lichiardopol. Proof of a conjecture of Henning and Yeo on vertex- disjoint directed cycles. <i>SIAM Journal on Discrete Mathematics</i> , 28(3):1618–1627, 2014.
[Lit75]	C. H. C. Little. A characterization of convertible (0, 1)-matrices. Journal of Combinatorial Theory, Series B, 18(3):187–208, 1975.

296

[LM17]	Z. Li and B. Mohar. Planar digraphs of digirth four are 2-colorable. SIAM Journal on Discrete Mathematics, 31(3):2201–2205, 2017.
[Lov68]	L. Lovász. On covering of graphs. In <i>Proceedings of the Colloquium on the Theory of Graphs 1966</i> , pages 231–236, 1968.
[Lov75]	L. Lovász. Problem 2. In Symposium on Recent Advances in Graph Theory, Prague, pages 655–664. Academia, Praha, 1975.
[Lov87]	L. Lovász. Matching structure and the matching lattice. Journal of Combinatorial Theory, Series B, 43(2):187–222, 1987.
[LP68]	L. Lovász and M. D. Plummer. <i>Matching Theory</i> , volume 29 in Annals of Discrete Mathematics. North-Holland, 1968.
[LP07]	T. Lam and A. Postnikov. Alcoved polytopes, I. Discrete & Computa- tional Geometry, 38(3):453–478, 2007.
[LP15]	A. Lubiw and V. Pathak. Flip distance between two triangulations of a point set is NP-complete. <i>Computational Geometry</i> , 49:17–23, 2015.
[LTWZ13]	L. M. Lovász, C. Thomassen, Y. Wu, and CQ. Zhang. Nowhere-zero 3-flows and modulo k-orientations. <i>Journal of Combinatorial Theory,</i> Series B, 103(5):587–598, 2013.
[LY78]	C. L. Lucchesi and D. H. Younger. A minimax theorem for directed graphs. <i>Journal of the London Mathematical Society</i> , 2(3):369–374, 1978.
[LZ98]	M. Li and L. Zhang. Better approximation of diagonal-flip transformation and rotation transformation. In <i>Proceedings COCOON 1998</i> , pages 85–94, 1998.
[LZ03]	P. C. B. Lam and H. Zhang. A distributive lattice on the set of perfect matchings of a plane bipartite graph. <i>Order</i> , $20(1)$:13–29, 2003.
[Mad67]	W. Mader. Homomorphie eigenschaften und mittlere Kantendichte von Graphen. Mathematische Annalen, 174:265–268, 1967.
[Mad85]	W. Mader. Degree and local connectivity in digraphs. <i>Combinatorica</i> , 5:161–165, 1985.
[Mad91]	W. Mader. Ecken von kleinem Grad in kritisch <i>n</i> -fach zusammenhängenden Digraphen. Journal of Combinatorial Theory, Series B, $43(2)$:260–272, 1991.
[Mad95]	W. Mader. Existence of vertices of local connectivity k in digraphs of large outdegree. Combinatorica, 15:533–539, 1995.
[Mad96]	W. Mader. On topological tournaments of order 4 in digraphs of outde- gree 3. Journal of Graph Theory, 21:371–376, 1996.
[McC00]	W. McCuaig. Even dicycles. Journal of Graph Theory, 35(1):46–68, 2000.
[McC04]	W. McCuaig. Pólya's permanent problem. <i>The Electronic Journal of Combinatorics</i> , 11(1), 2004. Research Article 79.

298	BIBLIOGRAPHY
[MdM05]	R. M. McConnella and F. de Montgolfier. Linear time modular decompo- sition of directed graphs. <i>Discrete Applied Mathematics</i> , 145(2):198–209, 2005.
[Men27]	K. Menger. Zur allgemeinen Kurventheorie. <i>Fundamenta Mathematicae</i> , 10(1):96–115, 1927.
[Möh85]	R. H. Möhring. Algorithmic aspects of the substitution decomposition in optimization over relations, set systems and boolean functions. <i>Annals of Operations Research</i> , 4:195–225, 1985.
[Moh03]	B. Mohar. Circular colorings of edge-weighted graphs. <i>Journal of Graph Theory</i> , 43:107–116, 2003.
[Moh16]	B. Mohar. Dichromatic number and fractional chromatic number. <i>Forum of Mathematics, Sigma</i> , 4, 2016. Article number E32.
[MR84]	R. H. Möhring and F. J. Radermacher. Substitution decomposition for discrete structures and connections with combinatorial optimization. <i>North-Holland Mathematics Studies, Algebraic and Combinatorial Meth-</i> <i>ods in Operations research</i> , 95:257–355, 1984.
[MS86]	R. Manber and J. Y. Shao. On digraphs with the odd cycle property. Journal of Graph Theory, $10(2)$:155–165, 1986.
[MS09]	G. MacGillivray and M. Siggers. On the complexity of <i>H</i> -colouring planar graphs. <i>Discrete Mathematics</i> , 309(18):5729–5738, 2009.
[MS21a]	T. Mészáros and R. Steiner. Complete minors in digraphs with given dichromatic number. <i>arXiv preprint</i> , 2021. arXiv:2101.04590, submitted.
[MS21b]	T. Mészáros and R. Steiner. Zero sum cycles in complete digraphs. <i>arXiv</i> preprint, 2021. arXiv:2103.04359, submitted.
[MSW19]	M. G. Millani, R. Steiner, and S. Wiederrecht. Colouring non-even di- graphs. <i>arXiv preprint</i> , 2019. arXiv:1903.02872, submitted.
[Mül90]	H. Müller. Alternating cycle-free matchings. Order, 7(1):11–21, 1990.
[NBJB20]	A. N., J. Bang-Jensen, and S. Bessy. Out-colorings of digraphs. <i>Journal of Graph Theory</i> , 93:88–112, 2020.
[Neg94]	S. Negami. Diagonal flips in triangulations of surfaces. Discrete Mathematics, $135(1-3)$:225–232, 1994.
[NL82]	V. Neumann-Lara. The dichromatic number of a digraph. Journal of Combinatorial Theory, Series B, 33(3):265–270, 1982.
[NPS19]	S. Norine, L. Postle, and Z. X. Song. Breaking the degeneracy barrier for coloring graphs with no K_t -minor. arXiv preprint, 2019. arXiv:1910.09378.
[Oxl11]	J. Oxley. <i>Matroid Theory</i> , volume 21 of <i>Oxford Graduate Texts in Mathematics</i> . Oxford University Press, Oxford, 2nd edition, 2011.

[Pap78]	C. H. Papadimitriou. The adjacency relation on the traveling salesman polytope is NP-complete. <i>Mathematical Programming</i> , 14(3):312–324, 1978.
[Pér84]	B. Péroche. NP-completeness of some problems of partitioning and covering in graphs. <i>Discrete Applied Mathematics</i> , 8(2):195–208, 1984.
[Ple79]	J. Plesník. The NP-completeness of the Hamiltonian cycle problem in planar digraphs with degree bound two. <i>Information Processing Letters</i> , 8(4):199–201, 1979.
[Plu80]	M. D. Plummer. On <i>n</i> -extendable graphs. <i>Discrete Mathematics</i> , 31(2):201–210, 1980.
[Pos09]	A. Postnikov. Permutohedra, associahedra, and beyond. International Mathematics Research Notices, (6):1026–1106, 2009.
[Pos20]	L. Postle. An even better density increment theorem and its application to Hadwiger's conjecture. <i>arXiv preprint</i> , 2020. arXiv:2006.14945.
[Pou14]	L. Pournin. The diameter of associahedra. <i>Advances in Mathematics</i> , 259:13–42, 2014.
[Pro02]	J. Propp. Lattice structure for orientations of graphs. <i>arXiv preprint</i> , 2002. arXiv:math/0209005.
[PS19]	V. Pilaud and F. Santos. Quotientopes. Bulletin of the London Mathematical Society, 51:406–420, 2019.
[Pul75]	W. R. Pulleyblank. On minimizing setups in precedence constrained scheduling. 1975. Technical Report 81105-OR.
[PW98]	J. Propp and D. Wilson. Coupling from the past: a user's guide. In <i>Microsurveys in discrete probability (Princeton, NJ, 1997)</i> , volume 41 of <i>DIMACS Series on Discrete Mathematics and Theoretical Computer Science</i> , pages 181–192. American Mathematical Society, 1998.
[Ray]	JF. Raymond. Dynamic Erdős-Pósa listing. Website: https://perso. limos.fr/~jfraymon/Erd%C5%91s-P%C3%B3sa.
[RC98]	F. J. Rispoli and S. Cosares. A bound of 4 for the diameter of the symmetric traveling salesman polytope. <i>SIAM Journal on Discrete Mathematics</i> , 11(3):373–380, 1998.
[Rém04]	E. Rémila. The lattice structure of the set of domino tilings of a polygon. <i>Theoretical Computer Science</i> , 322(2):409–422, 2004.
[Rin57]	Ringel, G. Über Geraden in allgemeiner Lage. <i>Elemente der Mathematik</i> , 12:75–82, 1957.
[Rog99]	R. O. Rogers. On finding shortest paths in the rotation graph of binary trees. In <i>Proceedings of the Thirtieth Southeastern International Con-</i> <i>ference on Combinatorics, Graph Theory, and Computing</i> , volume 137, pages 77–95, 1999.

[RRST96]	B. Reed, N. Robertson, P. D. Seymour, and R. Thomas. Packing directed circuits. <i>Combinatorica</i> , 16(4):535–554, 1996.
[RS86]	N. Robertson and Seymour. Graph minors. V. Excluding a planar graph. <i>Journal of Combinatorial Theory, Series B</i> , 41:92–114, 1986.
[RS04]	N. Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. Journal of Combinatorial Theory, Series B, 92(2):325–357, 2004.
[RST93]	N. Robertson, P. D. Seymour, and R. Thomas. Hadwiger's conjecture for K_6 -free graphs. <i>Combinatorica</i> , 13:279–361, 1993.
[RST99]	N. Robertson, P. D. Seymour, and R. Thomas. Permanents, Pfaffian ori- entations, and even directed circuits. <i>Annals of Mathematics</i> , 150(3):929– 975, 1999.
[RW08]	A. Raspaud and W. Wang. On the vertex-arboricity of planar graphs. <i>European Journal of Combinatorics</i> , 29(4):1064–1075, 2008.
[RW12]	A. Raspaud and W. Wang. Vertex-arboricity of planar graphs without intersecting triangles. <i>European Journal of Combinatorics</i> , 33(5):905–923, 2012.
[RW16]	B. Reed and D. R. Wood. Forcing a sparse minor. <i>Combinatorics, Probability and Computing</i> , 25(2):300–322, 2016.
[Sco97]	A. D. Scott. Induced trees in graphs of large chromatic number. <i>Journal of Graph Theory</i> , 24(4):297–311, 1997.
[Sev]	M. Severino. <i>Digraphs and Homomomorphisms: Cores. Colorings, and Constructions, 2014.</i> PhD thesis, University of Montana, Missoula, Faculty of Mathematics.
[Sey80]	P. D. Seymour. Decomposition of regular matroids. <i>Journal of Combi-</i> natorial Theory, Series B, 28(3):305–359, 1980.
[Sey81]	P. D. Seymour. Recognizing graphic matroids. <i>Combinatorica</i> , 1(1):75–78, 1981.
[Škr02]	R. Škrekovski. On the critical point-arboricity graphs. Journal of Graph Theory, 39(1):50–61, 2002.
[Sop14]	É. Sopena. Complete oriented colourings and the oriented achromatic number. <i>Discrete Applied Mathematics</i> , 173:102–112, 2014.
[SS87]	G. Steiner and L. K. Stewart. A linear time algorithm to find the jump number of 2-dimensional bipartite partial orders. <i>Order</i> , 3(4):359–367, 1987.
[ST87]	P. D. Seymour and C. Thomassen. Characterization of even directed graphs. <i>Journal of Combinatorial Theory, Series B</i> , 42(1):36–45, 1987.
[Sta86]	R. P. Stanley. Two poset polytopes. <i>Discrete & Computational Geometry</i> , 1(1):9–23, 1986.

[Ste20]	R. Steiner. Disjoint cycles with length constraints in digraphs of large connectivity or minimum degree. <i>arXiv preprint</i> , 2020. arXiv:2011.11605, submitted.
[Ste21]	R. Steiner. On coloring digraphs with forbidden induced subgraphs. <i>arXiv preprint</i> , 2021. arXiv:2103.04191, submitted.
[STT88]	D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. <i>Journal of the American Mathematical Society</i> , 1(3):647–681, 1988.
[SU97]	R. Scheinerman, E and D. H. Ullman. Fractional Graph Theory: A Rational Approach to the Theory of Graphs. Wiley, 1997.
[Sum81]	D. P. Sumner. Subtrees of a graph and the chromatic number. In <i>The theory and applications of graphs</i> , pages 557–576. Wiley, 1981.
[SW19]	R. Steiner and S. Wiederrecht. Parametrised algorithms for directed modular width. <i>arXiv preprint</i> , 2019. arXiv:1905.13203.
[SW20]	R. Steiner and S. Wiederrecht. Parameterized algorithms for directed modular width. In <i>Algorithms and Discrete Applied Mathematics. CAL-DAM 2020</i> , pages 415–426. Springer, 2020.
[Tan14]	N. D. Tan. Vertex disjoint cycles of different lengths in <i>d</i> -arc-dominated digraphs. <i>Operations Research Letters</i> , 42(5):351–354, 2014.
[Tan15]	N. D. Tan. On vertex-disjoint cycles of different lengths in 3-regular digraphs. <i>Discrete Mathematics</i> , 338(12):2485–2491, 2015.
[Tan17]	N. D. Tan. On 3-regular digraphs without vertex disjoint cycles of different lengths. <i>Discrete Mathematics</i> , 340(8):1933–1943, 2017.
[Tan20]	N. D. Tan. On 3-regular digraphs of girth 4. <i>Discrete Mathematics</i> , 343(1):111632, 2020.
[Tho83]	C. Thomassen. Disjoint cycles in digraphs. <i>Combinatorica</i> , 3:393–396, 1983.
[Tho84]	C. Thomassen. An extremal function for contractions of graphs. <i>Mathematical Proceedings of the Cambridge Philosophical Society</i> , 95(2):261–265, 1984.
[Tho85a]	C. Thomassen. The 2-linkage problem for acyclic digraphs. <i>Discrete Mathematics</i> , 55(1):73–87, 1985.
[Tho 85b]	C. Thomassen. Even cycles in directed graphs. European Journal of Combinatorics, 6(1):85–89, 1985.
[Tho86]	C. Thomassen. Sign-nonsingular matrices and even cycles in directed graphs. <i>Linear Algebra and its Applications</i> , 75:27–41, 1986.
[Tho91]	C. Thomassen. Highly connected non-2-linked digraphs. <i>Combinatorica</i> , 11(4):393–395, 1991.

[Tho92]	C. Thomassen. The even cycle problem for directed graphs. Journal of the American Mathematical Society, 5(2):217–229, 1992.
[Tho 05]	C. Thomassen. Some remarks on Hajós' conjecture. Journal of Combinatorial Theory, Series B, $93(1)$:95–105, 2005.
[Tho06]	R. Thomas. A survey of Pfaffian orientations of graphs. <i>Proceedings of the International Congress of Mathematicians</i> , 3:963–984, 2006.
[Thu90]	W. P. Thurston. Conway's tiling groups. American Mathematical Monthly, 97(8):757–773, 1990.
[Tut58]	W. T. Tutte. A homotopy theorem for matroids. I, II. Transactions of the American Mathematical Society, 88:144–174, 1958.
[Val79]	L. G. Valiant. The complexity of computing the permanent. <i>Theoretical Computer Science</i> , 8(2):189–201, 1979.
[vdZ19]	D. van der Zypen. Personal communication. 2019.
[Vin88]	A. Vince. Star chromatic number. Journal of Graph Theory, 12(4):551–559, 1988.
[Vis03]	S. Vishwanathan. On 2-coloring certain k -uniform hypergraphs. Journal of Combinatorial Theory, Series A, 101(1):168–172, 2003.
[Viz76]	V. G. Vizing. Vertex colorings with given colors. <i>Diskretnyi Analiz i Issledovanie Operatsii</i> , 29:3–10, 1976.
[VY89]	V. V. Vazirani and M. Yannakakis. Pfaffian orientations, 0–1 permanents, and even cycles in directed graphs. <i>Discrete Applied Mathematics</i> , 25(1–2):179–190, 1989.
[Wag37]	K. Wagner. Über eine Eigenschaft der ebenen Komplexe. Mathematische Annalen, 114(1):570–590, 1937.
[Wel76]	D. J. A. Welsh. <i>Matroid Theory</i> , volume 21 of <i>L. M. S. Monographs</i> . Academic Press, London-New York, 2nd edition, 1976.
[Woo78]	D. R. Woodall. Menger and Kőnig systems. In <i>Theory and applications of graphs</i> , pages 620–635. Springer, 1978.
[WZLW11]	G. Wang, S. Zhou, G. Liu, and J. Wu. Circular vertex arboricity. <i>Discrete Applied Mathematics</i> , 159(12):1231–1238, 2011.
[YBWW18]	D. Yang, Y. Bai, G. Wang, and J. Wu. On splitting digraphs. <i>European Journal of Combinatorics</i> , 71:174 – 179, 2018.
[ZG87]	F. J. Zhang and X. F. Guo. Hamilton cycles in directed Euler tour graphs. <i>Discrete Mathematics</i> , 64(2-3):289–298, 1987.
[Zhu01]	X. Zhu. Circular chromatic number: a survey. <i>Discrete Mathematics</i> , 229(1–3):371–410, 2001.
[ZQZ19]	W. J. Zhang, J. G. Qian, and F. J. Zhang. Distance between α -orientations of plane graphs by facial cycle reversals. <i>Acta Mathematica Sinica, English Series</i> , 2019.

BIBLIOGRAPHY

302