Foundations of computing Volume 17

Meike Hatzel

Dualities in graphs and digraphs

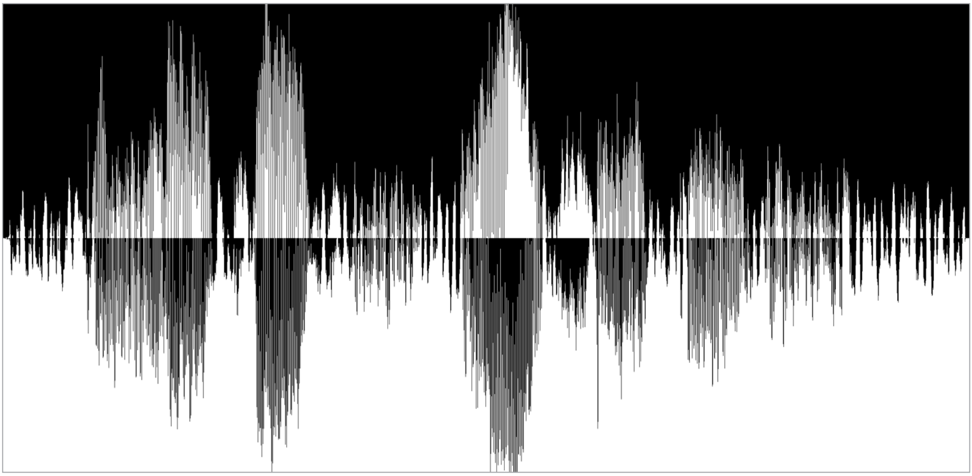

Universitätsverlag der TU Berlin

Technische Universität Berlin

Meike Hatzel

Dualities in graphs and digraphs

The scientific series *Foundations of computing* of the Technische Universität Berlin is edited by:
Prof. Dr. Stephan Kreutzer
Prof. Dr. Uwe Nestmann
Prof. Dr. Rolf Niedermeier

Foundations of computing | 17

Meike Hatzel

Dualities in graphs and digraphs

Universitätsverlag der TU Berlin

Bibliographic information published by Deutschen Nationalbibliothek
The German National Library lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at http://dnb.dnb.de.

Universitätsverlag der TU Berlin, 2023
https://verlag.tu-berlin.de/

Fasanenstr. 88, 10623 Berlin
Tel.: +49 (0)30 314 76131
Email: publikationen@ub.tu-berlin.de

Zugl.: Berlin, Techn. Univ., Diss., 2022
Gutachter: Prof. Dr. Martin Grohe (RWTH Aachen)
Gutachter: Prof. Dr. Stephan Kreutzer
Gutachter: Prof. Dr. Dimitrios Thilikos (LIRMM, France)
Die Arbeit wurde am 11. Mai 2022 an der Fakultät IV unter
Vorsitz von Prof. Dr. Markus Brill erfolgreich verteidigt.

This work – except where otherwise noted – is licensed under the
Creative Commons License CC BY 4.0.
http://creativecommons.org/licenses/by/4.0/

Cover image: Meike Hatzel

Print: docupoint GmbH
Layout/Typesetting: Meike Hatzel

ORCID iD Meike Hatzel: 0000-0003-3249-1169
https://orcid.org/0000-0003-3249-1169

ISBN 978-3-7983-3291-1 (print)
ISBN 978-3-7983-3292-8 (online)

ISSN 2199-5249 (print)
ISSN 2199-5257 (online)

Published online on the institutional repository of the
Technische Universität Berlin:
DOI 10.14279/depositonce-16147
http://dx.doi.org/10.14279/depositonce-16147

Abstract

Dualitity results are a central tool of bidimensionality theory and structure theory. They generally state that every graph has one of two properties, making it either algorithmically approachable or structurally rich. Due to the flexibility in the choice of the two properties, duality theorems occur in most areas of research in graph theory. We present a number of duality results in both graphs and digraphs, focusing on two kinds of properties: width measures and the exclusion of certain minors or subgraphs.

We introduce new width measures based on directed separations for directed graphs. In this context, we define directed versions of tangles, which are obstructions to our width measures being small. We consider the properties of these new concepts and prove some parallels to DAG-width. The duality theorem we prove is a very general one, it holds for all width measures that satisfy certain criteria.

Another directed width measure we introduce is cyclewidth. Cyclewidth is a branch decomposition for digraphs that uses families of disjoint directed cycles. We prove its parametric equivalence to directed treewidth and make use of its close relation to matching width measures in order to characterise the digraphs of small cyclewidth.

A current central goal in digraph structure theory is to describe the structure of digraphs excluding a fixed digraph as a butterfly minor. Much progress has been made over the past years, transferring comparable undirected results step by step into the directed setting. A recent step is a directed flat wall theorem. However, there are problems that arise when trying to push further towards a structure theorem based on that theorem. We suggest an alternative directed flat wall theorem and give some background and intuition as to why we think it provides a better base for a fully fledged structure theorem.

On undirected graphs we change the scope of our results a little. Here, we consider smaller, more specific graph classes and exclude induced subgraphs rather than minors. Two different generalisation of simplicial vertices are considered: moplexes and avoidable vertices. We identify the position of graphs with at most two moplexes in the hierarchy of hereditary graph classes and also prove that such graphs always admit a Hamiltonian path. The concept of avoidable vertices can be further generalised to avoidable paths. We prove that every graph with an induced path on k vertices contains such an avoidable path on k vertices.

Perfect graphs are a graph class on which solving some hard problems becomes tractable. We characterise the graphs with perfect linegraph squares by describing a family of forbidden induced subgraphs. Every graph inherits some properties from its linegraph and the square of its linegraph. In particular, some algorithmic problems can be reduced to solving a different problem on the (square of the) linegraph. The graphs with perfect linegraph squares thus form a class on which some problems that are hard in general, like Strong Edge Colouring, become feasible.

Kurzfassung

Dualitäten sind ein zentrales Werkzeug in der Bidimensionality-Theorie und Strukturtheorie. Dualitäten beschreiben in der Regel, dass eine Klasse von Graphen eine von zwei Eigenschaften besitzt: Man kann algorithmische Probleme auf dieser Klasse effizient lösen, oder die Klasse zeigt bestimmte Struktur. Aufgrund der Flexibilität in der Wahl dieser beiden Parameter, kommen Dualitäten in fast allen Bereichen der Graphentheorie vor. Wir präsentieren im Rahmen dieser Dissertation verschiedene Dualitäten in sowohl Graphen als auch gerichteten Graphen, und konzentrieren uns dabei auf zwei Arten von Eigenschaften: Weiteparameter und das Ausschließen bestimmter Minoren oder Teilgraphen.

Wir führen neue Weiteparameter für gerichtete Graphen ein, welche gerichtete Separationen benutzen. In diesem Kontext definieren wir auch eine gerichtete Version von Tangles, deren Vorkommen in einem Graph impliziert, dass die Weite des Graphen bezüglich der neuen Parameter nicht klein sein kann. Wir untersuchen die Eigenschaften dieser neuen Konzepte und ihre Verwandschaft zum bekannten Weiteparameter DAG-Weite. Dann beweisen wir eine Dualität, und zwar nicht bloß direkt für die eingeführten Konzepte, sondern basierend auf allgemeinen Eigenschaften, die diese haben sollen. Das führt zu einem sehr allgemeinen Dualitätsergebnis, das potentiell auch für andere Definitionen als unsere genutzt werden kann.

Ein weiterer gerichteter Weiteparameter, den wir einführen, ist die Cyclewidth. Cyclewidth zerlegt den Graphen mit Hilfe einer Branch-Zerlegung für gerichtete Graphen und nutzt Familien von Kreisen, um für diese Zerlegungen eine Weite zu bestimmen. Wir zeigen, dass Cyclewidth parametrisch äquivalent zum bekannten Parameter der gerichteten Baumweite ist. Außerdem nutzen wir die enge Verwandschaft zwischen Cyclewidth und Weiteparametern, die auf Matchings im Graphen basieren, um die Klasse von gerichteten Graphen mit kleiner Cyclewidth zu charakterisieren.

Ein großer Teil der Forschung im Rahmen der gerichteten Strukturtheorie konzentriert sich zur Zeit auf das Finden einer Beschreibung der Struktur von gerichteten Graphen, die einen bestimmten Graphen als Butterflyminor verbieten. Über die vergangenen Jahre hinweg gab es einigen Fortschritt darin die bestehenden Ergebnisse auf ungerichteten Graphen Schritt für Schritt auch für gerichtete Graphen zu beweisen. Zuletzt wurde in diesem Kontext ein gerichtetes Flat-Wall-Theorem bewiesen. Allerdings stößt

man auf einige Komplikationen, wenn man basierend auf diesem Flat-Wall-Theorem versucht ein Strukturtheorem zu beweisen. Wir schlagen hier ein alternatives Flat-Wall-Theorem für gerichtete Graphen vor und begründen wieso sich dieses besser als Grundlage für ein Strukturtheorem eignet.

Bei der Betrachtung von Dualitäten für ungerichteten Graphen betrachten wir dann eher spezifische Graphklassen und nutzen verbotene induzierte Teilgraphen statt Minoren. Wir betrachten zwei verschiedene Verallgemeinerungen für das Konzept von simplizialen Knoten in chordalen Graphen auf allgemeine Graphen: Moplexe und vermeidbare Knoten. Wir reihen die Klasse der Graphen mit höchstens zwei Moplexen in die Hierarchie der bekannten hereditären Graphklassen ein und beweisen außerdem, dass jeder solche Graph einen Hamilton-Pfad besitzt. Das Konzept von vermeidbaren Knoten kann weiter verallgemeinert werden zu vermeidbaren Pfaden. Wir zeigen, dass jeder Graph, der einen induzierten Pfad auf k Knoten enthält, auch einen vermeidbaren Pfad auf k Knoten enthält.

Perfekte Graphen sind eine bekannte Graphklasse, auf der manche schweren Probleme in polynomieller Zeit lösbar sind. Wir characterisierene die Graphen mit perfekten Quadratgraphen ihres Kantengraphen durch verbotene induzierte Teilgraphen. Jeder Graph erbt gewisse Eigenschaften von seinem Kantengraphen oder dessen Quadratgraphen. Insbesondere können manche algorithmische Probleme auf das Lösen verwandter Probleme auf dem (Quadratgraphen des) Kantengraphen reduziert werden. Damit sind die Graphen mit perfektem Quadratgraphen ihres Kantengraphens eine Klasse, auf der das Problem Strong Edge Colouring in Polynomialzeit lösbar ist.

Acknowledgements

First, I want to thank my supervisor Stephan Kreutzer for introducing me to the fascinating topic of digraph structure theory in my first semester of master studies and giving me the chance to further extend my interest in it in form of a PhD. Whenever I got stuck or needed advise, he was there for me with encouragement and guidance. I am especially thankful for all the freedom he gave me in choosing what to work on. He gave me the opportunity to attend the conferences and workshops of my choice, enabling me to make many valuable contacts and meet many interesting people.

Very special thanks go also to Sebastian Wiederrecht, who is not only a highly valued co-author and collaborator, but also supported me throughout my whole PhD with kind words, ideas, fruitful discussions, and (not to be underestimated) fun. Another person I want to give special thanks to is Irene Muzi. I am very sorry none of our joint research ended up in this thesis, still we frequently collaborated during the last few years and I greatly enjoyed it. In Irene I found a valuable colleague and friend, and she always knows how to lift me up when I seem in need of support.

Of course I thank all the authors involved in the work that went into this thesis. Marthe Bonamy especially for inviting me to an amazing workshop in Pessac and for further research visits afterwards.

I also want to thank the "Logic and Semantics" group at TU Berlin in all the constellations it had while I was part of it, that includes Jakub Gajarský, Archontia Gianopoulou, Maximilian Gorsky, Karl Heuer, O-joung Kwon, Marcelo Garlet Milani, Jana Pilz, Roman Rabinovich and Jean-Florent Raymond. I thank Jakub Gajarský in particular, with whom I shared an office for many years and who was always willing to help me getting a hang of the research and teaching world.

I thank my co-authors at the university of Primorska, Clément Dallard, Matjaž Krnc and Martin Milanič, for many interesting and productive research visits in Slovenia.

My thanks go also to the untiring and thorough proofreaders, who undoubtedly made this thesis more enjoyable to read: Aljoscha Meyer, Marcelo Garlet Milani, Sebastian Wiederrecht, Henning Seidler, and Matjaž Krnc.

Last but certainly not least, I want to thank my family for their patience and support.

Contents

1 Introduction

Dualities are omnipresent tools in graph structure theory. Put very simply, a duality theorem is of the form “Every graph has property A or it has property B.” There are many different kinds of dualities depending on the tools at hand. This is because there is a lot of freedom in how to formalise the properties A and B and also whether or not to restrict the class of graphs considered. Often the property A describes the graph to have a structure that allows us to solve certain problems more efficiently, so it says “this graph is somehow nice”. The property B on the other side often describes a concrete structural object the graph contains that witnesses its complexity, so it says “this graph is quite complicated”. The field concerned with finding and using dualities produces strong results that can be used in the algorithmic context as well as in the structural context. In particular, the algorithmic field of *bidimensionality theory* makes frequent use of dualities.

Many duality results make use of width parameters or tree-like decompositions of graphs. The most famous width parameter is probably treewidth, introduced in its current form by Robertson and Seymour [RS84]. Bounding the treewidth of a graph allows to use techniques like dynamic programming and therefore helps solving many problems that are hard in general [Bod96, Bod97, Bod05, DF13]. Robertson and Seymour also prove a number of dualities that are used frequently in research involving structures and algorithms for graphs. A famous example is the grid theorem [RS86], which states that every undirected graph has small treewidth or contains a large grid as a minor. Here the property helping to solve problems efficiently is having low treewidth, which essentially means that the graph is structurally similar to a tree. The structure preventing this is the grid graph.

Theorems describing the structure of graph classes excluding a certain graph as a minor are called *excluded minor theorems* and there are enough results like this that this area of duality results is often called *excluded minor theory*. The best known excluded minor result is Kuratowski’s theorem stating that a graph is planar if and only if it does not contain a complete graph on 5 vertices or a complete bipartite graph with 3 vertices on each side as a minor. This is a duality in the sense described above as we either get a nice graph property, i.e. being planar, which allows for many problems to be solved efficiently, or we find one of two minors that witness the non-planarity of the graph. Robertson and Seymour proved that every minor-closed graph class has a characterisation by forbidden minors [RS04b], which is equivalent to the graph minor relation forming a well-quasi-ordering. This is a very impactful result with numerous applications, see for example the survey by Kawarabayashi and Mohar [KM07] for an overview. One of the most imminent and known applications is the first polynomial time algorithm for the k-Disjoint Paths problem [RS04b].

Containing another graph as a minor is only one possibility of "containing a complicated structure", there are many other options, like subgraphs, induced subgraphs, or topological minors, just to name a few. Characterisation of graph classes by forbidden (induced) subgraphs for example is a very popular and active field. For example, forests are the class of graphs excluding cycles as subgraphs and bipartite graphs are characterised by forbidding odd cycles as subgraphs. For induced subgraphs a popular example is the class of chordal graphs, which are by definition the graphs with no induced cycles of length four or more. There are many problems that are related to such characterisations such as the Erdős-Hajnal-conjecture [EH89, Chu14] stating that the graphs of every graph class that has a characterisation by forbidden induced subgraphs either contain large cliques or large independent sets. Also the research area around χ-boundedness, which is a generalisation of perfect graphs, often makes use of characterisations through forbidden induced subgraphs, see [SR19] for a survey, in particular cycles and their complements (see, e.g. [SS16, RS84, CSSS20]). The Gyárfás-Sumner-conjecture, a very central question of this field, states that the class of graphs forbidding a tree as induced subgraphs is χ-bounded.

As width parameters often play a central role for dualities, they also play a very central role in this thesis. We mention and introduce a number of width parameters for undirected, directed and bipartite graphs. Having so many different width measures at hand, it is important to have a way to compare them. To this end, we call two width measures parametrically equivalent, if they are bound in a function of each other.

With duality results having such a large impact on graph theory, there was also soon a keen interest in generalisations to directed graphs. A directed analogue of treewidth was introduced by Reed [Ree99] and Johnson, Robertson, Seymour and Thomas [JRST01]. They conjectured that for this directed width measure and an adequate choice of a grid structure and a directed minor relation a similar duality as between treewidth and grids would hold. It took about 15 years to obtain this directed grid theorem [KK15] and the functions relating the directed treewidth and the size of the obtained grid structure is much faster growing than in the undirected case.

This thesis presents duality results from different areas of graph theory, digraph structure theory and even matching theory. There are several directed width measures that we use or even introduce. On undirected graphs, we consider dualities involving induced subgraphs and forbidden induced subgraphs. We provide a short overview over the results here before giving global definitions and a short background overview.

The results presented in this thesis origin from several research projects conducted with different groups of co-authors. All results obtained by the groups mentioned here were obtained through cooperative research. When stated within the corresponding chapters themselves these results are not further attributed. Although it is always difficult in retrospect to determine who came up with which idea in a research meeting, this thesis only contains results to which the author of this thesis made major contributions. Any results not obtained in the scope of any of these projects is attributed to the original authors at the point where the results are mentioned, stated or used.

1.1 Ganglions

In undirected graphs there is the concept of tangles, which yield an obstruction to small treewidth. This gives a way to decompose graphs into dense, structurally rich areas. Though there is much understanding on how such decompositions can be found and deployed in undirected graphs, the picture is more complicated for directed graphs. Erde introduces directed blockages [Erd20], which are an obstruction to directed pathwidth. Moreover, Giannopoulou, Kawarabayashi, Kreutzer and Kwon introduce a directed definition for tangles [GKK+20], which yield an obstruction to small directed treewidth.

We are interested in directed width measures in between directed pathwidth and directed treewidth. One example for such a width measure is DAG-width. There are digraph classes of constant directed treewidth and unbounded DAG-width as well as digraph classes of constant DAG-width and unbounded directed pathwidth. In Chapter 3 we find a generalisation of tangles to directed graphs, called *ganglions*, using DAGs as the base of decompositions via separations. We try to make use of the similarity to DAG-width in order to obtain a duality between the existence of a decomposition of small width and the existence of a ganglion of high order, while maintaining a strong correspondence to strategies in the cops and robber reachability game.

To this end we define two width measures and consider their properties and behaviour. While the first allows for a duality theorem, it is too weak to provide monotone strategies, the second is strong enough to ensure monotonicity but we pay for it by not being able to prove the duality any more.

The work presented in Chapter 3 is joint work with Roman Rabinovich and Sebastian Wiederrecht and based on unpublished work.

1.2 Cyclewidth

In Chapter 4 we introduce a new directed width measure, the cyclewidth. We prove that it is parametrically equivalent to directed treewidth and at the same time closely related to width measures for undirected graphs with perfect matchings. Currently, the main application for cyclewidth is a grid theorem for bipartite graphs [HRW19].

Cyclewidth is closed under taking butterfly minors, a property that directed treewidth only has when choosing a very specific and restrictive alternative definition. We use its relation to matching theory in order to characterise the digraphs of cyclewidth one. This characterisation is a rather unusual one as it is by allowed strongly 2-connected butterfly minors. Therefore, it also yields a rather unusual duality property, which is that a digraph either has large cyclewidth or only allows for specific strongly 2-connected butterfly minors.

The work presented in Chapter 4 is joint work with Archontia Giannopoulou, Roman Rabinovich and Sebastian Wiederrecht and based on [HRW19, GHW19]. The results come from two different projects which are both mainly concerned with matching theory. The first [HRW19] introducing cyclewidth and providing the matching grid theorem, the second [GHW19] giving the characterisation of bipartite graphs of small perfect matching width. Here we extract the directed consequences and statements of both.

1.3 Towards a directed structure theorem

In Chapter 5 we deal with the duality between non-planar butterfly minors and embedding oriented decompositions of digraphs.

The graph minor project contains the most influential results in recent undirected graph theory research. Naturally, the desire to achieve insight in the structure of directed graphs excluding a fixed minor is strong. There has been progress over recent years by achieving a directed grid theorem [KK15] and a directed flat wall theorem [GKKK20].

We discuss the two different versions of the existing directed flat wall theorem and their drawbacks. Then, we present an alternative directed flat wall theorem, excluding a different digraph as a butterfly minor. This new theorem lies “in between” the two existing ones and as such does have neither of these drawbacks. Motivated by early results by Robertson and Seymour on single crossing minor free graphs and Wagner’s theorem for graphs excluding K_5, we also present a flat wall theorem excluding the unique strongly 2-connected orientation of K_5.

These results given in Chapter 5 consist of preliminary thoughts in direction of a structure theorem for directed graphs, and indicate where the main difficulties lie in finding one. They are also meant to lay a foundation of definitions and concepts.

The work presented in Chapter 5 is joint work with Maximilian Gorski, Ken-ichi Kawarabayashi, Stephan Kreutzer and Sebastian Wiederrecht and based on unpublished work. Results presented here were found within different subgroups and different meetings. The flat wall proof is based on an original one by Giannopoulou et al. [GKKK20] and adapted by Giannopoulou and Wiederrecht [GW21] and then again adapted by the author of this thesis to our setting. The remaining proofs in this section are written and developed by the author of this thesis based on discussion among the co-authors.

1.4 Moplexes

A moplex is a maximal clique module in a graph. Moplexes yield a generalisation of Dirac's classical theorem that every chordal graph contains a simplicial vertex [Dir61] from chordal graphs to general graphs. In Section 6.1 we investigate k-moplex graphs, which are defined as graphs containing at most k moplexes. In particular we study the smallest nontrivial case $k = 2$, which forms a counterpart to the class of interval graphs. As the main structural result in that section, we show that the class of connected 2-moplex graphs is sandwiched between the classes of proper interval graphs and cocomparability graphs; moreover, both inclusions are tight for hereditary classes.

This leads to the natural question of whether having at most two moplexes guarantees a sufficiently strong structure to efficiently solve problems that are known to be intractable on cocomparability graphs, but not on proper interval graphs. We provide reductions that answer this question negatively for two prominent problems fitting this profile, namely Graph Isomorphism and Max-Cut. On the other hand, we prove that every connected 2-moplex graph contains a Hamiltonian path, generalising the same property of connected proper interval graphs. Furthermore, for graphs with a higher number of moplexes, we lift the previously known result that graphs without asteroidal triples have at most two moplexes to the more general setting of larger asteroidal sets.

The work presented in Section 6.1 is joint work with Clément Dallard, Robert Ganian, Matjaž Krnc and Martin Milanič and based on [DGH+21]. The results of this section have been developed in many online meetings involving all authors and during a number of research visits of the author of the thesis in Koper working with Clément Dallard, Matjaž Krnc and Martin Milanič in person.

1.5 Avoidable paths

An avoidable vertex is a different generalisation of Dirac's classical theorem that every chordal graph contains a simplicial vertex. Ohtsuki, Cheung and Fujisawa [OCF76] show that every graph contains at least two avoidable vertices. Beisegel et al. [BCG+19] generalise this result to edges, showing that every graph that contains at least one edge also contains an avoidable edge. They conjecture this to be true for any length of paths. In Section 6.2 we present a proof for this conjectured duality, showing that every graph either does not contain an induced path on k vertices or contains an avoidable path on k vertices.

The work presented in Section 6.2 is joint work with Marthe Bonamy, Oscar Defrain and Jocelyn Thiebaut and based on [BDHT20]. This result was achieved during a one-week-workshop in Pessac in 2019. Everything about it was done in very close teamwork within this week, from the first ideas to the write-up.

1.6 Linegraph squares

A strong edge colouring is a proper colouring of the edges of a graph such that no two edges that are incident to a common edge receive the same colour. The square of a graph G is obtained from G by adding edges between vertices of distance exactly 2. Therefore the strong edge colouring problem can be transformed to the problem of finding a proper vertex colouring of the squared linegraph. This shows that the linegraph square can be useful for proving structural properties of the graph itself and it is of interest to characterise the graphs with a squared linegraph in a specific class of graphs.

Proper vertex colouring is not the only problem that becomes solvable in polynomial time in the class of perfect graphs, so do Clique and Independent Set [PW10] for example. Section 6.3 obtains a characterisation of graphs with perfect linegraph squares by forbidden induced subgraphs. In addition we are able to observe that this class, among some other graph classes with specific linegraph squares, is χ-bounded. The concept of χ-boundedness is a generalisation of perfect graphs.

The work presented in Section 6.3 is joint work with Sebastian Wiederrecht and based on [HW18]. The proofs and extensive case distinctions were done on joint effort within many hours spent in front of a white board. The credit for most of the write-up goes to Sebastian Wiederrecht and it was only slightly adapted for presentation here.

2 Background and preliminaries

Many definitions are specific to certain results or concepts discussed in the single chapters. Therefore they are placed within the relevant chapter. Here we put in place the definitions and concepts used throughout the whole thesis and take a short look into the background of the results presented throughout this thesis.

For a set X we denote the set of all subsets of X that have size exactly k by $\binom{X}{k}$. Similarly, we denote the set of all subsets of X that have size at most k by $\binom{X}{\leq k}$. Two sets X and Y are called *comparable* if $X \subseteq Y$ or $Y \subseteq X$. For two sets X and Y we write $X \uplus Y$ for the disjoint union of X and Y. Additionally, we define $[k] \coloneqq \{1, \ldots, k\}$ for all $k \in \mathbb{N}$. For two sets X and Y and a function $\alpha : X \to Y$, we often use the shortcut $\alpha(X') \coloneqq \{\alpha(x) \mid x \in X'\}$ for a subset $X' \subseteq X$.

$\binom{X}{k}$

$X \uplus Y$

$[k]$

The natural numbers include 0 and are denoted $\mathbb{N}$, and $\mathbb{N}^+ \coloneqq \mathbb{N} \setminus \{0\}$.

$\mathbb{N}$

$\mathbb{N}^+$

All graphs and digraphs in this thesis are finite and simple. We assume familiarity with basic concepts in graph theory as used, e.g. by West or Diestel [Wes96, Die17] and for directed graphs in [BG02]. This thesis contains duality results on undirected, directed and bipartite graphs. Directed graphs can be considered a generalisation of undirected graphs in a natural way. In fact, directed graphs can be further generalised by bipartite graphs, which is less straightforward and explained in Chapter 4. This generalisation provides a lot of insight in the structure of behaviour of directed graphs. In order to make it easier to see what type of graph is used, we consequently use G to refer to undirected graphs, B to refer to bipartite graphs and D to refer to refer to directed graphs.

2.1 Undirected graphs

A graph G has a vertex set $V(G)$ and an edge set $E(G) \subseteq \binom{V(G)}{2}$, note that this implies the graph to be simple, we write $G = (V(G), E(G))$. We sometimes write $G \coloneqq (V, E)$ to fix $V = V(G)$ and $E = E(G)$ if the context does not allow for confusion with other graphs. The union of two graphs G_1 and G_2 is defined by $G_1 \cup G_2 \coloneqq (V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$.

$G_1 \cup G_2$

$N_G(v)$ $N_G[v]$

Let G be a graph. Let $e = \{u, v\} \in E(G)$ be an edge of G, then e is *incident* to u and v. We call u and v the *endpoints* of e and say u is *adjacent* to v, and we sometimes write uv instead of $\{u, v\}$. If $\{u, v\} \notin E(G)$, then u and v are *non-adjacent*. For a vertex $v \in V(G)$ we define its *neighbourhood* $N_G(v) := \{u \mid \{u, v\} \in E(G)\}$. The set $N_G[v] := \{v\} \cup N_G(v)$ is the *closed neighbourhood* of v. We generalise that definition to sets $S \subseteq V(G)$ by $N_G[S] := \bigcup_{s \in S} N_G[s]$ and $N_G(S) := N_G[S] \setminus S$.

$\overline{G}$ $\overline{S}$ $\partial_G(S)$

We define the *complement* of G as $\overline{G} := \left(V(G), \binom{V(G)}{2} \setminus E(G)\right)$. For a set $S \subseteq V(G)$, define by $\overline{S} := V(G) \setminus S$ the *complement set* of S. The *edge-cut* of a set $S \subseteq V(G)$ is defined as $\partial_G(S) := \{\{u, v\} \in E(G) \mid u \in S, v \in V(G) \setminus S\}$, we call S and $\overline{S}$ the *shores* of this edge cut. A *cut-edge* of $\partial_G(S)$ is an edge $uv \in E(G)$ with $u \in S$ and $v \in \overline{S}$. The *size* if the edge-cut $\partial_G(S)$ is the number of cut-edges it has. In general we say a set of edges $E' \subseteq E(G)$ is an edge-cut in G if there exists a set $S \subseteq V(G)$ such that $\partial_G(S) = E'$. Two edge-cuts $\partial(S)$ and $\partial(S')$ are called *laminar* if S and S' are comparable or S and $\overline{S'}$ are comparable.

$G' \subseteq G$ $G[X]$

Every subset $V' \subseteq V(G)$ together with a subset $E' \subseteq \binom{V'}{2} \cap E$ constitutes a *subgraph* G' of G, also notated as $G' \subseteq G$ with vertex set $V(G') = V'$ and edge set $E(G') = E'$. For $X \subseteq V(G)$ we define $G[X] = (X, \{e \in E(G) \mid e \subseteq X\})$ to be the subgraph of G induced by X. Any subgraph of G induced by a subset of $V(G)$ is an *induced subgraph* of G. A graph H is said to be *contained* in G if H is isomorphic to an induced subgraph of G. If H is not contained in G, we say G is *H-free*, respectively for graphs $H_1, \ldots, H_t$ that are not contained in G we say that G is *$\{H_1, \ldots, H_t\}$-free*.

$G/(e \to x_{uv})$ G/e $H \preccurlyeq G$

Let G be a graph and $e = \{u, v\} \in E(G)$ one of its edges. We say that the graph G' with vertex set $V(G') = (V(G) \setminus \{u, v\}) \cup \{x_{uv}\}$ and edge set $E(G') = \left(E(G) \cap \binom{V(G')}{2}\right) \cup \{\{x_{uv}, y\} \mid \{u, y\} \in E(G) \text{ or } \{v, y\} \in E(G)\}$ is obtained from G by *contracting* the edge e into the vertex x_{uv}. If the name of the new vertex is not relevant, we also write $G' = G/(e \to x_{uv})$, or only $G' = G/e$. A graph H that is obtained from G by deleting vertices and contracting edges is called a *minor* of G, written $H \preccurlyeq G$.

A *path* P of *length* k in G is a sequence of distinct vertices $v_1, \ldots, v_{k+1}$ such that $\{v_i, v_{i+1}\} \in E(G)$ for $1 \leq i \leq k$. The vertices v_1 and v_{k+1} are the *end-vertices* of P and we call P a v_1-v_{k+1}-path. In favour of readability as well as consistency with related literature, we use the notation P_k for the path on k vertices. Note that this is a path of length $k+1$. For two vertices $x, y \in V(G)$ we define the *distance* $\mathsf{dist}_G(x, y)$ between x and y to be the length of a shortest path in G with endpoints x and y. A *cycle* C in G of length k is a path of length $k-1$ and an edge between its two end-vertices. For $k \geq 3$, we denote by C_k the cycle of length k. We often identify a path P or a cycle C with the subgraph $G[P]$ or $G[C]$ it induces.

By adding edges between vertices of distance at most k, we obtain the *k-th distance power* G^k, or simply *k-th power* of G. Formally G^k is defined by $V(G^k) := V$ and $E(G^k) := \{xy \mid \mathsf{dist}(x,y) \leq k\}$. The *linegraph* of a graph G is the graph $\mathsf{LG}(G)$ defined by G^k $\mathsf{LG}(G)$

$$V(\mathsf{LG}(G)) := E(G) \text{ and}$$
$$E(\mathsf{LG}(G)) := \{e_1e_2 \mid e_1 \cap e_2 \neq \emptyset,\ e_1, e_2 \in E(G)\}.$$

A set $X \subseteq V(G)$ is an independent set in G if $\{x,y\} \notin E(G)$ for all $x, y \in X$. The graph G is called *complete* if all its vertices in it are pairwise adjacent. The complete graph on n vertices is denoted by K_n. Consistent with this, we denote by $3K_1$ the edgeless graph with three vertices. We call G *bipartite* if its vertex set can be partitioned into two independent sets A and B and *complete bipartite* if G contains all possible edges between A and B. If G is complete bipartite, then we denote G by $K_{a,b}$, where $a := |A|$ and $b := |B|$. The *independence number* of G is the maximum size of an independent set in G. A *clique* in G is a subgraph induced by a set of pairwise adjacent vertices. A *(connected) component* of a graph is a maximal connected subgraph. We sometimes identify components of a graph with their vertex sets. Given two vertex sets A and B in G, we say that A *dominates* B if every vertex in B has a neighbour in A or is an element of A. K_n $K_{a,b}$

A *vertex cover* of a graph G is a set of vertices $S \subseteq V(G)$ such that for every edge $e \in E(G)$ we have $e \cap S \neq \emptyset$. The size of the minimum vertex cover is called the *vertex cover number* of G, denoted $\tau(G)$. A *matching* in G is a set of edges $M \subseteq E(G)$ such that no two edges in M intersect. The maximum size of a matching in G is called the *maximum matching number* of G, denoted $\nu(G)$. $\tau(G)$ $\nu(G)$

A *cut-vertex* in a graph is a vertex whose deletion increases the number of connected components. A graph without any cut-vertices is called *2-connected* and the maximal 2-connected subgraphs of a graph are called *blocks* For a graph G and $S \subseteq V(G)$, let $G - S$ be the subgraph of G induced by $V \setminus S$; if $S = \{v\}$ then we use $G - v$ as shorthand for $G - \{v\}$. For two vertices $u, v \in V(G)$, a set $S \subseteq V(G) \setminus \{u, v\}$ is a *u,v-separator* if u and v belong to different components of $G - S$, and it is a *minimal u,v-separator* if additionally no proper subset of S is a u,v-separator. A u,v-separator S is minimal if and only if the two components of $G - S$ containing u and v are *S-full*, that is, these components dominate S. A *minimal separator* is a minimal u,v-separator for two non-adjacent vertices u and v.

A *star* is a complete bipartite graph $K_{1,k}$ for $k \in \mathbb{N}$, and the *claw* is the star on 4 vertices. A graph is a *tree* if for every two vertices u, v in it there is exactly one u-v-path. A tree is *rooted* if it has a distinguished vertex called the *root*.

A *(proper) vertex colouring* of a graph assigns a colour to every vertex of the graph such that no two adjacent vertices receive the same colour. Similarly, an *edge colouring* assigns colours to the edges of a graph, such that two edges incident to a common vertex receive different colours. By $\omega(G)$ we denote the size of a maximum clique in a graph G, the *clique number* of G and by $\chi(G)$ the minimum number of colours required for a vertex colouring of G, the *chromatic number* of G. A graph G is *perfect* if $\omega(H) = \chi(H)$ holds for all induced subgraphs $H \subseteq G$.

$\omega(G)$

$\chi(G)$

2.1.1 Graph classes and their properties

A graph G is *cobipartite* if its complement is bipartite. Furthermore, G is said to be a *cochain graph* if there are two disjoint cliques X and Y with $V(G) = X \cup Y$ and $X = \{x_1, \ldots, x_k\}$ such that $N[x_i] \subseteq N[x_j]$ for all $1 \leq i < j \leq k$. A graph G is an *interval graph* if it has an *interval representation*, that is, if its vertices can be put in a one-to-one correspondence with a family of closed intervals on the real line such that two distinct vertices are adjacent if and only if the corresponding intervals intersect. If G has an interval representation in which no interval contains another interval, then G is said to be a *proper interval graph*. A vertex set A in a graph G is an *asteroidal set* if for every $a \in A$, the vertices in $A \setminus \{a\}$ are contained in a single connected component of $G - N(a)$ [Wal78]. The *asteroidal number* of G is the maximum size of an asteroidal set in G (see, e.g. [LMW98, KKM01, Alc14]). Of particular interest are the graphs of asteroidal number at most 2, these are also called *AT-free* graphs, as they do not contain any asteroidal set of size 3, which are called *asteroidal triples*. A prominent subclass of AT-free graphs is the class of *cocomparability graphs*, which are graphs whose complements allow for a transitive orientation of their edges. For further background on graph classes, we refer to [BLS99].

Chordal graphs and simplicial vertices

A graph is *chordal* if it does not contain any induced cycle of length greater than three. The class of chordal graphs is a subclass of perfect graphs on which many problems that are usually hard to solve are solvable in polynomial time, e.g. maximal clique [EHPS05] and graph colouring [Gol80]. Chordal graphs are therefore well studied and in particular of interest in the context of algorithmic applications.

The *chordality* of a graph G is defined as the minimum $k \in \mathbb{N}$ such that $E(G) = E_1 \cap \ldots \cap E_k$ and $(V(G), E_i)$ is a chordal graph for every $i \in [k]$.

One of the main reasons for chordal graphs having such algorithmically useful behaviour is the existence of simplicial vertices in them. A vertex is *simplicial* if its

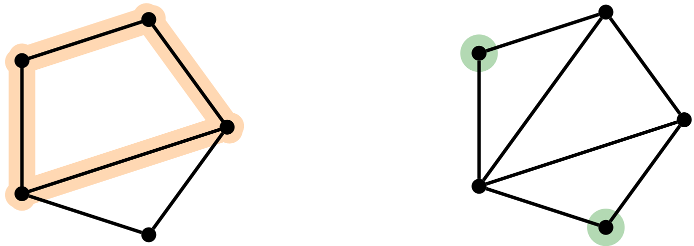

Figure 2.1: The left graph is not chordal as it contains an induced cycle of length four. The right graph is chordal and the two marked vertices are its simplicial vertices.

neighbourhood forms a clique, see, for example, Figure 2.1. In 1961, Dirac proved the following result for chordal graphs.

Theorem 2.1.1 (Dirac [Dir61])**.** Every non-complete chordal graph contains at least two non-adjacent simplicial vertices.

This leads to the characterisation of chordal graphs by the existence of a perfect elimination ordering [Dir61, Ros70]. A perfect elimination ordering is an ordering of the vertices of the graph such that every vertex is simplicial in the graph induced by this vertex and all the vertices to its left in the ordering. Such an ordering can be computed in linear time using algorithms like lexicographic breath first search [RTL76], which are considered in more detail in Subsection 6.1.4.

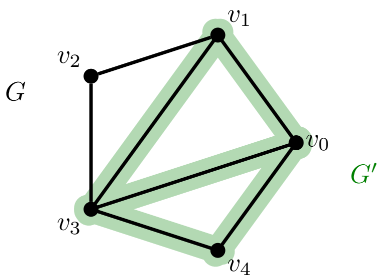

Figure 2.2: The graph $G' := G[\{v_0, v_1, v_3, v_4\}]$ is the subgraph induced by the first four vertices of the perfect elimination ordering $(v_4, v_0, v_3, v_1, v_2)$. The vertex v_1 is the last vertex of G' with respect to the ordering and thus is simplicial in G'.

Two adjacent simplicial vertices have exactly the same closed neighbourhood, which is a clique. Thus, they behave similarly and for many applications the existence of two independent simplicial vertices is more useful.

A vertex set M is a *module* of a graph G if each vertex $v \in V(G) \setminus M$ is either adjacent to every vertex in M or not adjacent to any vertex in M. A *clique module* is a module that induces a clique. A *simplicial module* is an inclusion-maximal clique module containing a simplicial vertex. Note that all vertices in a simplicial module are simplicial. Dirac's theorem is equivalent to the statement that every non-complete chordal graph contains at least two simplicial modules [Cao16, Shi88].

Theorem 2.1.2 (Cao [Cao16], Shibata [Shi88])**.** Every non-complete chordal graph contains at least two simplicial modules.

Another characterisation of chordal graphs is based on subtree intersections. A graph G is chordal if and only if there is a tree T such that G is the intersection graph of the subtrees of T [Gav74, Bun74]. This leads to a definition introduced by Lin, McKee and West [LMW98]. The *leafage* $\mathsf{leaf}(G)$ of a chordal graph G is the minimum number of leafs in a tree T such that G is representable as an intersection graph of subtrees of T. The leafage divides the class of chordal graphs naturally into subclasses of bounded leafage. For example interval graphs are the chordal graphs of leafage at most two.

Similarly one can now consider subclasses of chordal graphs of bounded number of simplicial modules. As the number of simplicial modules yields an upper bound on the leafage of a graph [LMW98], we obtain that every chordal graph with at most two simplicial modules is an interval graph.

2.1.2 Treewidth

Many problems that are hard in general are comparably easy to solve on trees and forests. The width measure *treewidth* describes the similarity of a given graph to a forest, which gives rise to many algorithmic applications using dynamic programming approaches. Therefore, treewidth is a popular parameter in the field of parametrised complexity. In its current form treewidth was introduced by Robertson and Seymour in 1984 [RS84]. But the concept was known long before this. Originally it was introduced by Bertelè and Brioschi in 1972 [BB72].

Definition 2.1.3 (Treewidth)**.** Let G be a graph. A *tree decomposition* of G is a tuple (T, β) , where T is a tree and $\beta : V(T) \to 2^{V(G)}$ a function mapping every vertex of T to a set of vertices in G such that the following properties hold.

(Tw1) Every vertex of G is contained in $\beta(t)$, which is called the *bag* of t, for some $t \in V(T)$.

(Tw2) For every edge $uv \in E(G)$ there exists a $t \in V(T)$ such that $u, v \in \beta(t)$.

(Tw3) For every pair of vertices $t_i, t_j \in V(T)$ and every $t \in V(T)$ lying on the unique path from t_i to t_j in T holds $\beta(t_i) \cap \beta(t_j) \subseteq \beta(t)$.

The *width* of a tree decomposition (T, β) is given by $\max\{|\beta(t)| \mid t \in T\} - 1$. The *treewidth* of G, denoted $\mathsf{tw}(G)$, is then defined as the minimum width of any tree decomposition of G. $\mathsf{tw}(G)$

The graphs of treewidth one are exactly the trees and forests while the graphs of treewidth two are exactly the series-parallel graphs [Bod98]. More general, the graphs of treewidth at most k are the partial k-trees, which have an inductive definition.

Classes of small treewidth also have characterisations by forbidden minors. For treewidth one it is K_3, for treewidth two it is K_4 [Bod88]. This does not continue as nicely, for treewidth three it is already four forbidden minors [APC90, ST90].

2.1.3 Obstructions to small treewidth

Definition 2.1.4 (Bramble)**.** A *bramble* in a graph G is a collection $\mathfrak{B} := (B_1, \ldots, B_k)$ of connected subgraphs of G such that for every pair B_i and B_j of these subgraphs $B_i \cap B_j \neq \emptyset$ or there is an edge $e \in E(G)$ with one endpoint in B_i and the other in B_j. The *order* of $\mathfrak{B}$ is the size of a minimum hitting set of all subgraphs.

Definition 2.1.5 (Haven)**.** A *haven* of *order* k in a graph G is a function $\eta : \binom{V(G)}{<k} \to \mathcal{C}$, where $\mathcal{C}$ is the collection of all connected subgraphs of G, such that $\eta(X)$ is a connected component of $G - X$ for all $X \in \binom{V(G)}{<k}$. Additionally, if $X \subseteq Y$, then $\eta(Y) \subseteq \eta(X)$ for every two sets $X, Y \in \binom{V(G)}{<k}$.

Theorem 2.1.6 (Seymour and Thomas [ST93])**.** A graph G contains a bramble of order k if and only if it has treewidth at least $k - 1$ if and only if G contains a haven of order k.

A *grid of order k* or *$k \times k$ grid* is the graph G with the vertex set $V(G) = \{v_{i,j} \mid 1 \leq i, j \leq k\}$ and the edge set $E(G) = \{\{v_{i,j}v_{i',j}\}, \{v_{i,j}v_{i,j'}\} \mid j' = j + 1, i' = i + 1\}$, see Figure 2.3 for an illustration.

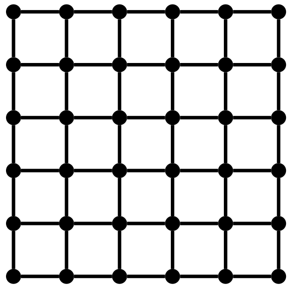

Figure 2.3: The 6×6 grid.

Theorem 2.1.7 (Robertson and Seymour [RS86])**.** There is a function $f : \mathbb{N} \to \mathbb{N}$ such that every graph of treewidth at least $f(k)$ contains a $k \times k$-grid as a minor.

Originally, Robertson and Seymour proved an exponential bound for the function f. This was improved, first to be polynomial [CC16] and then again several times optimising the polynomial [Chu15, CT21] and the current bound is $\mathcal{O}(k^9 \text{polylog} k)$, for planar graphs it is even linear: every planar graph of treewidth at least $6k$ contains a grid of order k as a minor [RST94]. There is also a significantly simplified version of the grid theorem proof due to Diestel, Jensen, Gorbunov and Thomassen [DJGT99] that should not remain unmentioned in this context.

Another obstruction to small treewidth that was introduced by Robertson and Seymour is called *tangle* and uses the concept of separations. Let G be a graph. Two sets $X, Y \subseteq V(G)$ are a separation, written $\{X, Y\}$ if there is no edge between $X \setminus Y$ and $Y \setminus X$. Let $\mathcal{S}_k$ be the family of separations of order less than k. An *orientation* of $\mathcal{S}_k$ contains (X, Y) or (Y, X) for every $\{X, Y\} \in \mathcal{S}_k$.

Definition 2.1.8 (Tangle)**.** Let G be a graph. A *tangle* of order k is an orientation $\mathcal{O}$ of $\mathcal{S}_k$ such that for every three separations $(X_1, Y_1), (X_2, Y_2), (X_3, Y_3) \in \mathcal{O}$ we have $X_1 \cup X_2 \cup X_3 \neq V(G)$. $\dashv$

Robertson and Seymour proved the following duality between tangles and treewidth.

Theorem 2.1.9. Let G be a graph, and $k \in \mathbb{N}$. Then, there is a tangle of order k in G if and only if $\text{tw}(G) \geq k$.

The cops and robber game

Another possibility to characterise treewidth is in terms of graph searching games. These games typically involve two parties the robber and some cops that are positioned in the graph, on vertices or edges or sets of such and the cops try to catch the robber while the robber tries to escape the cops.

The graph searching game that characterises treewidth is given by the following definition.

Definition 2.1.10 (Cops and robber game)**.** Given a graph G, the *k-cops and robber game* on G is played between two players, the *cop* and the *robber* player. Positions of the game are pairs (C, R) where $C \in \binom{V(G)}{\leq k}$ is the position of the cops and R is a connected component of $G - C$, the component in which the robber is hiding. The component R can be empty, that is, we consider $\emptyset$ to be a connected component of every graph for this purpose. The game is played as follows:

1. The cop player starts by choosing a location $C_0 \in \binom{V(G)}{\leq k}$ of the cops, and the robber player chooses a connected component R_0 of $G - C_0$.
2. From position (C_i, R_i) the cop player chooses $C_{i+1} \in \binom{V(G)}{\leq k}$, and the robber player chooses a component R_{i+1} of $G - C_{i+1}$ such that $(V(R_i) \cap V(R_{i+1})) \setminus (C_i \cap C_{i+1}) \neq \emptyset$.
3. The cop player wins the game if, after a finite number of turns, the robber player selects $\emptyset$ as the new component, otherwise the robber player wins.

A *play* π is a maximal sequence $\pi := ((C_0, R_0), (C_1, R_1), \ldots)$ of positions given by the rules above. A *(k-cop) strategy* for the cop player is a function $f : \binom{V(G)}{\leq k} \times \mathcal{G} \to \binom{V(G)}{\leq k}$ where $\mathcal{G}$ denotes the set of all connected subgraphs of G. A play π is *consistent* with a strategy f if $f((C_i, R_i)) = C_{i+1}$ for all i. The strategy f is called *winning* for the cop player if every play consistent with f is won by the cop player.

A play is *robber monotone* if $R_{i+1} \subseteq R_i$ for all i. A strategy f is *robber monotone* if every play consistent with f is robber monotone.

A play is called *cop monotone* if for all i, j with $i < j$ and for every vertex $v \in C_i \cap C_j$ we have $v \in C_h$ for all $i \leq h \leq j$. A strategy f is *cop monotone* if every play consistent with f is cop monotone.

A strategy f is *monotone* if it is both robber and cop monotone.

The (robber monotone, cop monotone) cop number of a graph G is the smallest number k such that there exists a (robber monotone, cop monotone) winning strategy for the k-cops and robber game.

For this game there are no monotonicity costs, that is, if k cops have a winning strategy, then they also have a monotone winning strategy.

Theorem 2.1.11 (Seymour and Thomas [ST93])**.** A graph has treewidth at most k if and only if $k+1$ cops can capture the robber.

2.2 Directed graphs

$V(D)$ $E(D)$ $D_1 \cup D_2$

A *directed graph*, or *digraph*, D has a vertex set $V(D)$ and an edge set $E(D) \subseteq (V(D) \times V(D)) \setminus \{(v,v) \mid v \in V(D)\}$. We sometimes write $D := (V,E)$ to fix $V = V(D)$ and $E = E(D)$ if the context does not allow for confusion with other graphs. The union of two digraphs D_1 and D_2 is defined by $D_1 \cup D_2 := (V(D_1) \cup V(D_2), E(D_1) \cup E(D_2))$.

$e \sim u$ $N_D^{\mathsf{out}}(v)$ $N_D^{\mathsf{in}}(v)$ $N_D^{\mathsf{out}}[v]$ $N_D^{\mathsf{in}}[v]$

Let D be a digraph. Let $e = (u,v) \in E(D)$ be an edge of D, then we call u the *tail* of e and v the *head* of e, we say e is *incident* to u and v, written $e \sim u$ and $e \sim v$. For a vertex $v \in V(D)$ we define its *out-neighbourhood* $N_D^{\mathsf{out}}(v) := \{u \mid (v,u) \in E(D)\}$ and its *in-neighbourhood* $N_D^{\mathsf{in}}(v) := \{u \mid (u,v) \in E(D)\}$. The set $N_D^{\mathsf{out}}[v] := \{v\} \cup N_D^{\mathsf{out}}(v)$ is the *closed out-neighbourhood* of v and the set $N_D^{\mathsf{in}}[v] := \{v\} \cup N_D^{\mathsf{in}}(v)$ is the *closed in-neighbourhood* of v. We generalise these definitions to sets $S \subseteq V(D)$ by $N_D^{\mathsf{out}}[S] := \bigcup_{s \in S} N_D^{\mathsf{out}}[s]$, $N_D^{\mathsf{out}}(S) := N_D^{\mathsf{out}}[S] \setminus S$, $N_D^{\mathsf{in}}[S] := \bigcup_{s \in S} N_D^{\mathsf{in}}[s]$ and $N_D^{\mathsf{in}}(S) := N_D^{\mathsf{in}}[S] \setminus S$. If $N_D^{\mathsf{out}}(v) = \emptyset$ for a vertex $v \in V(D)$, then it is called a *sink* of D, and it is called a *source* if $N_D^{\mathsf{in}}(v) = \emptyset$. If the digraph D is apparent from the context, we omit the index.

$\partial_D^{\mathsf{out}}(v)$ $\deg_D^{\mathsf{out}}(v)$ $\partial_D^{\mathsf{in}}(v)$ $\deg_D^{\mathsf{in}}(v)$ $\partial_D(v)$

Let v be a vertex in D, by $\partial_D^{\mathsf{out}}(v)$ we denote the set of *out-edges* of v, that is $\partial_D^{\mathsf{out}}(v) := \{(v,u) \in E(D)\}$. The *out-degree* of v is defined by $\deg_D^{\mathsf{out}}(v) := |\partial_D^{\mathsf{out}}(v)|$. Similarly, by $\partial_D^{\mathsf{in}}(v)$ we denote the set of *in-edges* of v, that is $\partial_D^{\mathsf{in}}(v) := \{(u,v) \in E(D)\}$ and define the *in-degree* of v by $\deg_D^{\mathsf{in}}(v) := |\partial_D^{\mathsf{in}}(v)|$. We also define $\partial_D(v) := \partial_D^{\mathsf{out}}(v) \cup \partial_D^{\mathsf{in}}(v)$. Again, we omit the index, if the context clearly provides the digraph D.

$\mathsf{start}(P)$ $\mathsf{end}(P)$

A *(directed) path* P of *length* k in a directed graph D is a sequence of distinct vertices $v_1, \dots, v_{k+1}$ such that $(v_i, v_{i+1}) \in E(D)$ for all $1 \leq i \leq k$. The vertex v_1 is called the *start-vertex* of P, denoted $\mathsf{start}(P)$, while the vertex v_{k+1} is the *end-vertex* of P, denoted $\mathsf{end}(P)$. We call P a $\mathsf{start}(P)$-$\mathsf{end}(P)$-path. We often identify P

with the subgraph $(\{v_1, \ldots, v_{k+1}\}, \{(v_i, v_{i+1}) \mid 1 \leq i \leq k\})$. Let $X \subseteq V(D)$ be a set of vertices in D. A directed path P is *disjoint* from X if $V(P) \cap X = \emptyset$. It is *internally disjoint* from X if $V(P) \cap X \subseteq \{\mathsf{start}(P), \mathsf{end}(P)\}$. An X*-path* is a directed path P of length at least one that is internally disjoint from X and $\mathsf{start}(P), \mathsf{end}(P) \in X$. For a subgraph $D' \subseteq D$ we also write D'-path instead of $V(D')$-path. Two paths P and P' are *disjoint* if $V(P) \cap V(P') = \emptyset$ and they are *internally disjoint* if their only intersections are start- and end-vertices, that is, $V(P) \cap V(P') \subseteq \{\mathsf{start}(P), \mathsf{start}(P'), \mathsf{end}(P), \mathsf{end}(P')\}$. A *(directed) cycle* of length k is a directed path of length $k-1$ such that $(\mathsf{end}(P), \mathsf{start}(P)) \in E(D)$. A collection $\mathcal{L}$ of pairwise disjoint paths is called a *linkage*. We say $\mathcal{L}$ is an A-B-linkage if every $L \in \mathcal{L}$ is an a-b-path for some $a \in A$ and $b \in B$. We call a collection $\mathcal{L}$ of paths a *half-integral linkage* if every vertex of the graph occurs in at most two paths of $\mathcal{L}$.

A digraph D that is obtained from an undirected graph G by adding exactly one direction of every edge, that is, $V(D) = V(G)$ and for all $\{u, v\} \in E(G)$ we have either $(u, v) \in E(D)$ or $(v, u) \in E(D)$, is called an *orientation* of G. A digraph D is a *(directed) tree* or *out-branching* rooted at a vertex r if it is the orientation of an undirected tree T rooted at r such that $\deg_D^{\mathsf{in}}(v) = 1$ for every $v \in V(D) \setminus \{r\}$.

The *underlying undirected graph* of a digraph D is the undirected graph $\mathfrak{G}(D)$ with the vertex set $V(\mathfrak{G}(D)) \coloneqq V(D)$ and the edge set $E(\mathfrak{G}(D)) \coloneqq \{\{u, v\} \mid (u, v) \in E(D) \text{ or } (v, u) \in E(D)\}$. The digraph D is *strongly connected* if for every two vertices $x, y \in V(D)$ there is an x-y-path and a y-x-path in D. The digraph D is *weakly connected* if $\mathfrak{G}(D)$ is connected. A *strongly connected component* or *strong component* of D is a maximal strongly connected subgraph of D. $\mathfrak{G}(D)$

In directed graphs there are several ways to define containment relations that generalise minors in undirected graphs. The most prominent one in structure theory is the one of butterfly minors. The idea behind the definition is that the contraction of an edge should not create new paths, that is, if there is no path between two vertices before the contraction, then there is no such path after the contraction either. To this end, we call an edge e *butterfly contractible* if e is the only out-edge of its tail u or the only in-edge of its head v. For a butterfly contractible edge $e = (u, v)$ in D we define the digraph obtained by contracting it into the vertex x_{uv}, which we call the *contraction vertex*, as follows.

$$\begin{aligned} D/e \to x_{uv} \coloneqq \Big(& V(D) \setminus \{u, v\} \cup \{x_{uv}\}, \\ & E(D) \setminus \{(x, y) \in E(D) \mid \{x, y\} \cap \{u, v\} \neq \emptyset\} \\ & \quad \cup \{(x, x_{uv}) \mid (x, u) \in E(D) \text{ or } (x, v) \in E(D)\} \\ & \quad \cup \{(x_{uv}, x) \mid (u, x) \in E(D) \text{ or } (v, x) \in E(D)\} \Big) \end{aligned}$$

If we do not require a specific name for the contraction vertex, then we also write D/e D/e. A digraph D' is a butterfly minor of D if it can be obtained from a subgraph of D by butterfly contractions.

We can make an edge butterfly contractible by deleting some edges. We say a graph D' is obtained by *out-contracting* an edge e if it is obtained by first deleting all other out-edges of its tail and then butterfly contracting it. Similarly, we can *in-contract* an edge e by first deleting all other in-edges of its head and then butterfly contracting e.

An alternative way to describe this concept is via minor models. An *in-branching* is the *reverse graph* (obtained by reversing the direction of all edges) of an out-branching, and an *in-out-branching* is obtained by identifying the root of an in-branching and the root of an out-branching. A *butterfly (minor) model* of a digraph D' in a digraph D is a function μ mapping every vertex of D' to a subgraph of D and every edge of D' to an edge of D such that

1. $\mu(v)$ is an in-out-branching for every $v \in V(D')$,
2. $\mu(v)$ is disjoint from $\mu(u)$ for distinct $u, v \in V(D')$,
3. $\mu(e) \in E(D)$ for all $e \in E(D')$,
4. $\mu(e) \neq \mu(e')$ for distinct $e, e' \in E(D')$, and
5. if $e = (u, v) \in E(D')$, then $\mathsf{start}(\mu(e))$ lies in the out-branching of $\mu(u)$ and $\mathsf{end}(\mu(e))$ lies in the in-branching of $\mu(v)$.

2.2.1 Directed separations

Let D be a directed graph and $A, B \subseteq V(D)$. The tuple (A, B) is a *directed separation* if $A \cup B = V(D)$ and there are no edges with tail in $A \setminus B$ and head in $B \setminus A$ or no edges with tail in $B \setminus A$ and head in $A \setminus B$. In case there are no edges with tail in $A \setminus B$ and head in $B \setminus A$, we write $\overleftarrow{(A, B)}$ indicating that edges are allowed to $\overrightarrow{(A, B)}$ go from $B \setminus A$ to $A \setminus B$. Similarly, we write $\overrightarrow{(A, B)}$ if no edge in D has its tail in $B \setminus A$ and its head in $B \setminus A$. We call $\overrightarrow{(A, B)}$ or $\overleftarrow{(A, B)}$ the *direction* of (A, B), see Figure 2.4 for an illustration. It is possible that for a directed separation (A, B) there are no edges at all between $A \setminus B$ and $B \setminus A$, in that case we call (A, B) a *weak separation*. Note that this naming is due to $A \setminus B$ and $B \setminus A$ not being weakly connected in $D - (A \cap B)$, though in fact this notion is more restrictive, so in a sense stronger than directed separations.

The set $S := A \cap B$ is called the *separator* of the separation, and we sometimes write (A, S, B) to emphasise the separator. The *order* of a directed separation (A, B),

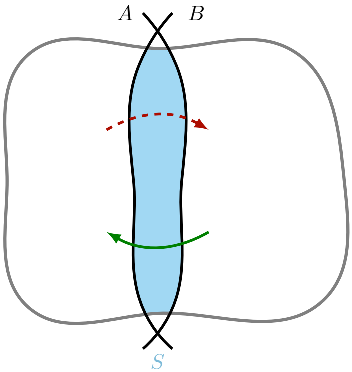

Figure 2.4: Depicted is the separation $\overleftarrow{(A,B)}$. The green solid arrow symbolises the direction of the allowed edges from $B \setminus A$ to $A \setminus B$, and the red dashed arrow symbolises the forbidden edges from $A \setminus B$ to $B \setminus A$. The separator $S = A \cap B$ is marked in black.

denoted by $|(A,B)|$, is the size of S. The set of all directed separations of order less than k in D is denoted by $\vec{\mathcal{S}}_k$, here for weak separations (A,B) we consider two distinct copies $\overrightarrow{(A,B)}$ and $\overleftarrow{(A,B)}$, both of which are elements in $\vec{\mathcal{S}}_k$. This way we obtain that every directed separation in $\vec{\mathcal{S}}_k$ has a uniquely defined direction. We also define the set of all separations $\vec{\mathcal{S}} := \bigcup_{k \in \mathbb{N}} \vec{\mathcal{S}}_k$. $\vec{\mathcal{S}}_k$ $\vec{\mathcal{S}}$

Let D be a digraph. A directed separation (A,B) is *trivial*, if $A \setminus B = \emptyset$, or $B \setminus A = \emptyset$. We call a set $S \subseteq V(G)$ a *separator*, if there is a non-trivial directed separation (A,S,B) in D.

If D is strongly connected and $D - S$ is not strongly connected, then the set S is called a *strong separator*. Let S be a strong separator in D and $C_1, \dots, C_\ell$ the strongly connected components of $D - S$. We call an ordering $\sigma = (C_1, \dots, C_\ell)$ a *topological ordering* of the strong components of $D - S$, if for all $(u,v) \in E(D)$ we have $u \in V(C_i)$ and $v \in V(C_j)$ implies that $i \leq j$. Note that every topological ordering can be associated with a number of directed separations as follows. For every $j \in \{1, \dots, \ell - 1\}$, there is the non-trivial directed separation $\left(S \cup \bigcup_{i=1}^{j} V(C_i), S \cup \bigcup_{i=j+1}^{\ell} V(C_i)\right)$. This implies that every strong separator also is a separator.

2.2.2 Directed treewidth

As a directed analogue to treewidth Reed [Ree99] and Johnson, Robertson, Seymour and Thomas [JRST01] introduced the concept of directed treewidth and conjectured that a directed version of the grid theorem holds for this directed width measure and a directed version of a grid.

$x \preccurlyeq_D y$

Let D be a digraph. A set Z *guards* a set X if there is no walk in $D - Z$ that starts and ends in X and visits vertices from $D - X$. For two sets $X, Y \subseteq V(D)$ we say that Y *strongly guards* X if $u \in X$ implies $v \in X \cup Y$ for all edges $(u, v) \in E(D)$. A vertex y is *reachable* from a vertex x in D, written as $x \preccurlyeq_D y$, if there exists a directed path from x to y. We say that y is reachable from a set of vertices X in D if there is a vertex $x \in X$ such that y is reachable from x. If neither y is reachable from x, nor x is reachable from y, we say that x and y are *incomparable*. The set of vertices reachable from X in D is denoted by $\mathsf{Below}_D(X)$. Similarly the set of vertices v such that $X \subseteq \mathsf{Below}_D(v)$ is denoted by $\mathsf{Above}_D(X)$.

$T_l(e)$ $T_r(e)$ T_d

For an arborescence, or directed tree, T, the removal of an edge $e = (t_1, t_2)$ splits T into two sub-arborescence: $T_l(e)$ containing t_2 and $T_r(e)$ containing t_1 as well as the root of T. We also write T_d for the sub-arborescence rooted at the vertex d.

Definition 2.2.1 (Directed treewidth)**.** A *directed tree decomposition* of a digraph D is a triple (T, β, γ) where T is a directed tree, $\beta : V(T) \to 2^{V(D)}$ maps every vertex t of T to a set $\beta(t) \subseteq V(D)$ called the *bag at* t and $\gamma : E(T) \to 2^{V(D)}$ maps every edge e of T to a set $\gamma(e) \subseteq V(D)$ called the *guard at* e such that the following hold

1. $\{\beta(t) \mid t \in V(T)\}$ is a partition of $V(D)$ (with possibly empty classes), and
2. for all $e \in E(T)$ the guard $\gamma(e)$ guards $\beta(T_l(e))$.

For every vertex $t \in V(T)$ we define $\Gamma(t) \coloneqq \beta(t) \cup \bigcup_{e \sim t} \gamma(e)$. The *width* of a directed tree decomposition is defined by $\max\{|\Gamma(t)| \mid t \in V(T)\}$. The directed treewidth of D, denoted $\mathsf{dtw}(D)$, is defined as the smallest $k \in \mathbb{N}$ such that D has a directed tree decomposition of width k.

A *cylindrical grid* $\mathsf{D}_k^{\circlearrowright}$ of order k consists of k concentric directed cycles and $2k$ paths connecting the cycles in alternating directions, see Figure 2.5 for an example, as follows. The k directed disjoint cycles $C_1, \ldots, C_k$ have length $2k$ each. The cycle C_i has the vertex set $\{v_1^i, \ldots, v_{2k}^i\}$ with the natural cyclic ordering. The paths are of two different kinds, we have the *in-paths* $P_1^{\mathsf{i}}, \ldots, P_k^{\mathsf{i}}$ and the *out-paths* $P_1^{\mathsf{o}}, \ldots, P_k^{\mathsf{o}}$ as follows.

$$P_j^{\mathsf{i}} = v_{2j}^k, v_{2j}^{k-1}, \ldots, v_{2j}^2, v_{2j}^1 \text{ for all } j \in \{1, \ldots, k\}$$
$$P_j^{\mathsf{o}} = v_j^1, v_j^2, \ldots, v_j^{k-1}, v_j^k \text{ for all } j \in \{1, \ldots, k\}.$$

The paths P_i^{i} and P_i^{o} together build the i-th *row* of W, which we also denote R_i.

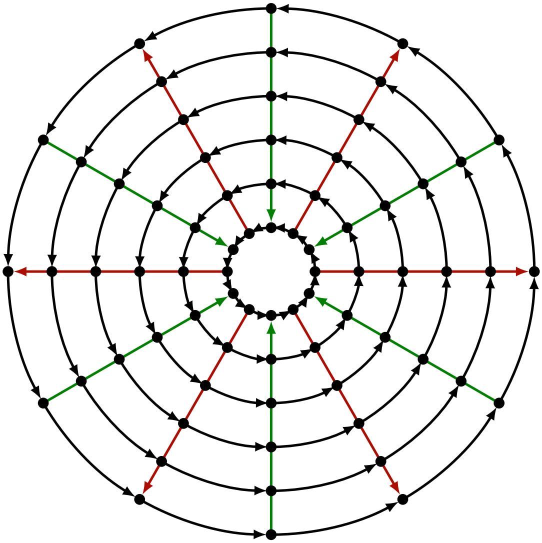

Figure 2.5: A cylindrical grid of order 6 with in-paths and out-paths.

After being open for over 15 years this conjecture was finally confirmed to be true by Kawarabayashi and Kreutzer [KK15].

Theorem 2.2.2 (Kawarabayashi and Kreutzer [KK15])**.** There is a function $f_{\mathsf{grid}} : \mathbb{N} \to \mathbb{N}$ such that every digraph D either satisfies $\mathsf{dtw}(D) \leq k$, or contains the cylindrical grid of order $f_{\mathsf{grid}}(k)$ as a butterfly minor.

While the function they provide is exponential, there is a polynomial bound for planar directed graphs [HKK19]. Additionally, Campos, Lopes, Maia, and Sau [CLMS22] adapted the result into an FPT algorithm.

2.2.3 Obstructions to small directed treewidth

There are a number of other obstructions and corresponding duality theorems that can be transferred from treewidth to directed treewidth.

Definition 2.2.3 ((Directed) bramble)**.** A *(directed) bramble* in a digraph D is a collection $\mathfrak{B} := (B_1, \ldots, B_k)$ of strongly connected subgraphs of D such that for every pair B_i and B_j of these subgraphs $B_i \cap B_j \neq \emptyset$ or there are edges $e_1, e_2 \in E(D)$ with $\mathsf{head}(e_1), \mathsf{tail}(e_2) \in V(B_i)$ and $\mathsf{tail}(e_1), \mathsf{head}(e_2) \in V(B_j)$. The *order* of $\mathfrak{B}$ is the size of a minimum hitting set of all subgraphs.

Definition 2.2.4 ((Directed) haven)**.** A *(directed) haven* of *order* k in a digraph D is a function $\eta : \binom{V(D)}{<k} \to \mathcal{C}$, where $\mathcal{C}$ is the collection of all strongly connected subgraphs of D, such that $\eta(X)$ is a strongly connected component of $D - X$ for all $X \in \binom{V(D)}{<k}$. Additionally, if $X \subseteq Y$, then $\eta(Y) \subseteq \eta(X)$ for every two sets $X, Y \in \binom{V(D)}{<k}$.

Both of these concepts have been shown to be dual to the directed treewidth.

Theorem 2.2.5 ([KO11,Ree99,JRST01])**.** Let D be a digraph. The directed treewidth of D, the maximum order of a bramble in D and the highest order of a haven in D all lie within a constant factor of each other. More precisely:

- If D has directed treewidth less than k, then D contains no haven of order $k+1$.
- If D contains no haven of order k, then D has directed treewidth at most $3k-2$.
- If D contains a bramble of order k, then it contains a haven of order $k+1$.
- If D contains a haven of order $k+1$, then D contains a bramble of order $k/2$.

See also [KO14] for details.

The directed cops and robber game

For directed graphs there exists a corresponding definition for a version of the cops and robber games, introduced by Johnson et al. [JRST01], which is closely related to directed treewidth.

Definition 2.2.6 (Directed cops and robber game)**.** The *directed k-cops and robber game* on a digraph D is played between two players, the *cop* and the *robber* player. Positions of the game are pairs (X, R) where $X \in \binom{V(D)}{\leq k}$ is the position of the cops and R is a strongly connected component of $D - X$, in which the robber is hiding. The component R can be empty, that is, we consider $\emptyset$ to be a strongly connected component of every digraph for this purpose. The game is played as follows:

1. The cop player starts by choosing a location $X_0 \in \binom{V(D)}{\leq k}$ of the cops, and the robber player chooses a strongly connected component R_0 of $D - X_0$.
2. From position (X_i, R_i) the cop player chooses $X_{i+1} \in \binom{V(D)}{\leq k}$, and the robber player chooses a component R_{i+1} of $D - X_{i+1}$ such that $V(R_i) \cup V(R_{i+1}) \subseteq V(R)$ for some strongly connected component R of $D - (X_i \cap X_{i+1})$.
3. The cop player wins the game if, after a finite number of turns, the robber player selects $\emptyset$ as the new component, otherwise the robber player wins.

A *play* π is a maximal sequence $\pi := ((X_0, R_0), (X_1, R_1), \ldots)$ of positions given by the rules above. A *(k-cop) strategy* for the cop player is a function $f : \binom{V(D)}{\leq k} \times \mathcal{D} \to \binom{V(D)}{\leq k}$ where $\mathcal{D}$ denotes the set of all strongly connected subgraphs of D. A play π is *consistent* with a strategy f if $f((X_i, R_i)) = X_{i+1}$ for all i. The strategy f is called *winning* for the cop player if every play consistent with f is won by the cop player.

We say that a vertex $v \in V(D)$ is *reachable* from the position (X, R) if there is a directed path from a vertex of R to v in $D - X$. We denote the set of all vertices reachable from (X, R) by $\mathsf{Below}(X, R)$. A play is *robber monotone* if $\mathsf{Below}(X_{i+1}, R_{i+1}) \subseteq \mathsf{Below}(X_i, R_i)$ for all i. A strategy f is *robber monotone* if every play consistent with f is robber monotone. A play is called *cop monotone* if for all i, j with $i < j$ and for every vertex $v \in X_i \cap X_j$ we have $v \in X_h$ for all $i \leq h \leq j$. A strategy f is *cop monotone* if every play consistent with f is cop monotone. A strategy f is *monotone* if it is both robber and cop monotone.

The (robber monotone, cop monotone) cop number of a digraph D is the smallest number k such that there exists a (robber monotone, cop monotone) winning strategy for the directed k-cops and robber game. ⊣

The duality between these games and directed treewidth is not as tight as for the undirected version.

Theorem 2.2.7 (Johnson et al. [JRST01])**.** Let D be a digraph and $k \in \mathbb{N}$. If $\mathsf{dtw}(D) \leq k$, then $k + 1$ cops have a winning strategy in the directed cops and robber game, and if k cops have a winning strategy, then $\mathsf{dtw}(D) \leq 3k - 2$.

2.3 Bipartite graphs and perfect matchings

Let G be a graph. A *perfect matching* in G is a matching M in G such that for every vertex $v \in V(G)$ there is an edge $e \in M$ with $v \in e$. We say that M *covers* the vertex

$\mathcal{M}(G)$ v (with the edge e). By $\mathcal{M}(G)$ we denote the set of all perfect matchings in G. If G is connected and every edge of G is contained in some perfect matching of G we say that G is *matching covered*.

Let $M \in \mathcal{M}(G)$. A cycle C in G is *M-alternating* if every second edge of it belongs to M, it additionally is *M-M'-alternating* if $E(C) \setminus M \subseteq M'$ for some $M' \in \mathcal{M}(G)$, see Figure 2.6 for an example. If C is an M-alternating cycle, then we say the matching $M'' := (M \setminus E(C)) \cup (E(C) \setminus M)$ is obtained by *switching M along C*.

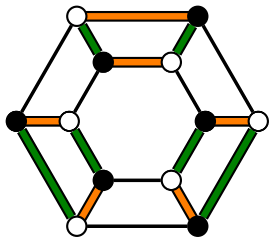

Figure 2.6: A bipartite graph with two perfect matchings M and M' building three M-M'-alternating cycles.

In the context of perfect matchings we cannot use the same definition for minors as in undirected graphs. This is because contracting an edge changes the parity of the vertex set and thus a graph with a perfect matching does not have any perfect matchings after a contraction. Thus, matching theory works with its own version of contractions and minors. A *bicontraction* is the operation of contracting both edges incident to a vertex of degree two at the same time.

For minors we need a little bit more. Let G be a graph with a perfect matching M. A set $X \subseteq V(G)$ is *conformal* if $G - X$ has a perfect matching, it is said to be *M-conformal* if M contains a perfect matching of $G - X$. Similarly, we call a subgraph $H \subseteq G$ *conformal* if $V(H)$ is, and we call it *M-conformal* if M contains perfect matchings of $G - V(H)$ and of H. A *matching minor* of G is a graph H that can be obtained by a series of bicontractions from a conformal subgraph of G.

The significance of matching minors, especially for the study of matching theoretic properties of bipartite graphs, was first observed by Little [Lit75]. But he does not call it a "matching minor" yet, the name is only introduced later by Norine and Thomas [Nor05, NT07].

Bipartite graphs play a central role in matching theory. As mentioned above, a graph B is *bipartite* if its vertex set can be partitioned into two independent sets V_0 and V_1, which we call the *colour classes* of B. We write $B = (V_0, V_1, E)$ to fix $V_0 \uplus V_1 = V(B)$ and $E = E(B)$ if the context does not allow for confusion with other graphs. In figures we depict the vertices of V_0 as white and the vertices of V_1 as black. V_0 V_1

3 Ganglions

A *tangle* in an undirected graph is an orientation of all separations of small order. Tangles are a tool introduced by Robertson and Seymour [RS91] as part of their graph minor series. A tangle yields an obstruction to a graph having small treewidth, that is, a graph has high treewidth if and only if it admits a tangle of large order. We can imagine this as the tangle pointing us in direction of the large gridminor the graph contains. Additionally, treewidth and tangles admit a game-theoretic characterisation by cops and robber games.

For many years now, researchers have been trying to find directed concepts with comparable properties. There are many possibilities to translate the undirected definitions to directed graphs and it is far from obvious which is the correct choice. In structure theory, (undirected) treewidth is often generalised to directed treewidth, which shares many properties with treewidth. Giannopoulou et. al introduced a generalisation of tangles to directed graphs by defining *directed tangles* [GKK⁺20]. They obtain a canonical directed tree decomposition, which generalises a result by Robertson and Seymour, distinguishing any two tangles in a given digraph.

Bienstock et. al [BRST91] introduce a variant of tangles, called *blockages*, that corresponds to pathwidth instead of treewidth. This was transferred to directed graphs by Erde [Erd20].

This chapter is about lifting the result due to Erde from directed pathwidth to more general digraph measures. We provide a definition of tangle-like structures in digraphs that yield a duality to a digraph width measure that uses directed separations. We call the structure we introduce a ganglion, which is closely related to the concept of tangles in undirected graphs. We introduce two width measures: the ν-DAG-width and the μ-DAG-width. Both of them use a DAG as their main structure and map its edges to directed separations in the digraph.

While there is a correspondence between directed treewidth and strategies in the directed cops and robber game, the canonical strategies provided by the decompositions are not necessarily monotone. This is different for DAG-width and the cops and robber reachability game; here a DAG-decomposition yields a monotone cop strategy analogously to undirected tree decompositions. We obtain the desired duality between ν-DAG-width and ganglions, but unfortunately the tie to the cops and robber reachability game we consider is weaker than for DAG-width, as the obtained strategies are not necessarily monotone. In an attempt to fix this, we introduce μ-DAG-width, demanding an additional property for the decomposition DAGs, which then yields monotone strategies but is too strong to still yield the duality with ganglions. Yet, the non-monotonicity that occurs in the strategies obtained for ν-DAG-width is quite restricted, so the gap in which a width measures with both properties can lie is rather small. The end of this chapter contains a discussion of the insights we obtain from the presented results.

3.1 Duality of DAG-width and the cops and robber reachability game

As mentioned before, the cops and robber game is originally a game-theoretic equivalent to treewidth on undirected graphs. It yields a useful tool for designing algorithms exploiting bounded treewidth as well as for finding obstructions to the treewidth of a class of graphs being bounded.

Obdržálek and Berwanger et. al. [Obd06, BDHK06, BDH$^+$12] introduced a directed width measure, the DAG-width, which corresponds to a directed variant of the cops and robber game, the cops and robber reachability game.

3.1.1 DAG-Width

The width measure DAG-width [Obd06, BDHK06, BDH$^+$12] is a natural way to transfer the undirected definitions for treewidth to directed graphs measuring the distance to an acyclic graph rather than to a directed tree as directed treewidth does. DAG-width lies between the two width measures directed treewidth and directed pathwidth. It uses a decomposition that assigns bags to the nodes of a DAG. One crucial reason to consider DAG-width is that it yields a polynomial time algorithm for solving parity games [BDH$^+$12] on classes of bounded width. The width measures for separations that we introduce have many properties in common with DAG-width.

$\mathsf{sources}_T$ sinks_T

For any DAG T let $\mathsf{sources}_T$ denote the set of sources in T and sinks_T denote the set of sinks in T.

Definition 3.1.1 (DAG-width)**.** Let D be a digraph. A *DAG-decomposition* of D is a pair (T, β) where T is a DAG and $\beta\colon V(T) \to 2^{V(D)}$ is a function mapping every node of T to a set of vertices, called a *bag*, such that

(Dag1) $\bigcup_{t \in V(T)} \beta(t) = V(D)$,

(Dag2) for all nodes $t_1, t_2, t_3 \in V(T)$ with $t_1 \preccurlyeq_T t_2$ and $t_2 \preccurlyeq_T t_3$ we have $\beta(t_1) \cap \beta(t_3) \subseteq \beta(t_2)$, and

(Dag3) for all edges $(t, t') \in E(T)$, $\beta(t) \cap \beta(t')$ strongly guards $\beta(\mathsf{Below}_T(t')) \setminus \beta(t)$, where $\beta(X) = \bigcup_{x \in X} \beta(x)$ for all $X \subseteq V(T)$.

The *width* of (T, β) is defined as $\max\{|\beta(t)| \mid t \in V(T)\}$. The *DAG-width* of D, denoted by $\mathsf{DAGw}(D)$, is defined as the minimum width of any DAG-decomposition of D. ⊣

As we are in search of an obstruction using directed separations, we are particularly interested in how DAG-width interacts with the directed separations of a digraph. The following lemma shows that there is a natural way to obtain directed separations from a DAG-decomposition. This behaves similarly to the intersection of two bags in an undirected tree decomposition and makes DAG-decompositions interesting in the context of decompositions using directed separations.

Lemma 3.1.2. Let (T, β) be a DAG-decomposition of a digraph D. For every edge $(t, t') \in E(T)$, the following are directed separations in D:

$$\overrightarrow{((V(D) \setminus \beta(\mathsf{Below}_T(t'))) \cup (\beta(t) \cap \beta(t')), \beta(\mathsf{Below}_T(t')))}.$$

Proof. Since $\beta(t) \cap \beta(t')$ strongly guards $\beta(\mathsf{Below}_T(t')) \setminus \beta(t)$, there is no edge from $\beta(\mathsf{Below}_T(t')) \setminus \beta(t)$ to $V(D) \setminus \beta(\mathsf{Below}_T(t'))$. Therefore, every path starting in $\beta(\mathsf{Below}_T(t'))$ and ending in $(V(D) \setminus \beta(\mathsf{Below}_T(t'))) \cup \beta(t)$ contains a vertex of $\beta(t) \cap \beta(t')$ and so the claim follows with $\beta(t) \cap \beta(t') = \beta(t) \cap \beta(\mathsf{Below}_T(t'))$. □

This shows that every edge in a DAG-decomposition naturally corresponds to a directed separation of order at most the width of the decomposition. However, the number of incoming and outgoing edges of a node in the decomposition DAG T can be unbounded. In order to relate DAG-width to our later definition of ganglions, we must restrict the number of such edges.

We can restrict the number of outgoing edges at each node of T by using a special type of DAG-decomposition. A DAG-decomposition (T, β) of a directed graph D is called *nice*, due to [BDH+12], if the following requirements are met:

(Nice1) T has a unique source,

(Nice2) every $t \in V(T)$ has at most two out-neighbours,

(Nice3) if t_1 and t_2 are distinct out-neighbours of $t \in V(T)$, then $\beta(t) = \beta(t_1) = \beta(t_2)$, and

(Nice4) if t_1 is the unique out-neighbour of t, then $\big||\beta(t)| - |\beta(t_1)|\big| \leq 1$.

Theorem 3.1.3 (Berwanger et al. [BDH+12]). Let D be a digraph. If D has a DAG-decomposition of width k, then it has a nice DAG-decomposition of width k.

Note that the number of incoming edges at the nodes of T remains unbounded.

To the best of our knowledge there is no tangle-like concept so far that relates to DAG-width in a similar fashion as undirected tangles relate to treewidth.

3.1.2 The cops and robber reachability game

There are several natural adaptions of cops and robber games to directed graphs. Depending on the definition, the strategies describe different properties of the digraph they are played on. The following variant, in which the robber moves along directed paths was introduced by Berwanger et al. [BDHK06] and corresponds to *DAG-width*.

Definition 3.1.4. Given a digraph D, the *directed k-cops and robber reachability game* on D, refer to it as the *cops and robber reachability game*, is played between two players, the *cop* and the *robber* player. Positions of the game are pairs (X, R) where $X \in \binom{V(D)}{\leq k}$ is the position of the cops and R is a strongly connected component of $D - X$, the component in which the robber is hiding. The component R can be empty, that is, we consider $\emptyset$ to be a strongly connected component of every digraph for this purpose. The game is played as follows:

1. The cop player starts by choosing a location $X_0 \in \binom{V(D)}{\leq k}$ of the cops, and the robber player chooses a strongly connected component R_0 of $D - X_0$.
2. From position (X_i, R_i) the cop player chooses $X_{i+1} \in \binom{V(D)}{\leq k}$, and the robber player chooses a component R_{i+1} of $D - X_{i+1}$ such that there is a directed path from a vertex of R_i to a vertex of R_{i+1} in $D - (X_i \cap X_{i+1})$.
3. The cop player wins the game if, after a finite number of turns, the robber player selects $\emptyset$ as the new component, otherwise the robber player wins.

A *play* π is a maximal sequence $\pi := ((X_0, R_0), (X_1, R_1), \ldots)$ of positions given by the rules above. A *(k-cop) strategy* for the cop player is a function $f : \binom{V(D)}{\leq k} \times \mathcal{D} \to \binom{V(D)}{\leq k}$ where $\mathcal{D}$ denotes the set of all strongly connected subgraphs of D. A play π is *consistent* with a strategy f if $f((X_i, R_i)) = X_{i+1}$ for all i. The strategy f is called *winning* for the cop player if every play consistent with f is won by the cop player.

We say that a vertex $v \in V(D)$ is *reachable* from the position (X, R) if there is a directed path from a vertex of R to v in $D - X$. We denote the set of all vertices reachable from (X, R) by $\mathsf{Below}(X, R)$. A play is *robber monotone* if $\mathsf{Below}(X_{i+1}, R_{i+1}) \subseteq \mathsf{Below}(X_i, R_i)$ for all i. A strategy f is *robber monotone* if every play consistent with f is robber monotone.

A play is called *cop monotone* if for all i, j with $i < j$ and for every vertex $v \in X_i \cap X_j$ we have $v \in X_h$ for all $i \leq h \leq j$. A strategy f is *cop monotone* if every play consistent with f is cop monotone.

A strategy f is *monotone* if it is both robber and cop monotone.

The (robber monotone, cop monotone) cop number of a digraph D is the smallest number k such that there exists a (robber monotone, cop monotone) winning strategy for the k-cops and robber reachability game. $\dashv$

A particularly elegant property of DAG-width is that its decompositions directly correspond to winning strategies for the cop player in the cops and robber reachability game. This is not the case for directed treewidth and the directed cops and robber game, where the best known upper bounds one the necessary number of cops and the directed treewidth use functions of each other as seen in Theorem 2.2.7.

Theorem 3.1.5 (Berwanger et al. [BDH+12])**.** For any digraph D there is a DAG-decomposition of width at most k if and only if the cop player has a monotone winning strategy in the k-cops and robber reachability game.

It is a standing open problem whether the costs for monotonicity of this game is bounded, that is, whether there is a function bounding how many extra cops are needed for monotone strategies. While this is still unknown, there is a conjecture for such a bound to even be linear.

Conjecture 3.1.6 (Berwanger et. al [BDH+12])**.** There exists a linear function f such that if k cops have a winning strategy for the cops and robber reachability game, then $f(k)$ cops have a monotone winning strategy in the cops and robber reachability game.

3.2 Systems of directed separations

In order to define decomposition that make use of directed separations, we need to know more about how directed separations relate to each other. Similar to separations in undirected graphs there are concepts of crossing and laminar separations. However, as we can have edges from one part of the separation to the other, working with these concepts becomes more complicated in directed graphs.

3.2.1 Uncrossing of directed separations

An important notion to describe the relation between separations is the property of being laminar. It describes that two separations are aligned with each other; they agree in the direction of the crossing edges as well as fulfil certain subset properties.

We introduce a partial order on the separations $\vec{\mathcal{S}}$ as follows. Let $(A, B), (J, K) \in \vec{\mathcal{S}}$, $(A,B) \leq (J,K)$ then we say $(A, B) \leq (J, K)$ if and only if

1. $\overrightarrow{(A, B)}, \overrightarrow{(J, K)}$, $A \subseteq J$ and $K \subseteq B$, or
2. $\overleftarrow{(A, B)}, \overrightarrow{(J, K)}$, $B \subseteq J$ and $K \subseteq A$, or
3. $\overrightarrow{(A, B)}, \overleftarrow{(J, K)}$, $A \subseteq K$ and $J \subseteq B$, or
4. $\overleftarrow{(A, B)}, \overleftarrow{(J, K)}$, $B \subseteq K$ and $J \subseteq A$.

Two directed separations (A, B) and (J, K) *cross* if neither $(A, B) \leq (J, K)$, nor $(J, K) \leq (A, B)$, that is, the separations are incomparable with respect to the partial order $\leq$. Otherwise, they are *laminar* (see Figure 3.1 for an illustration).

For two crossing directed separations (A, B) and (J, K) there is a directed separation of order not greater than the order of the initial two separations that is laminar to both of them. These yield local suprema and infima with respect to the partial order $\leq$, thus, we use the conventional notation ($\vee$,$\wedge$). For $(A, B), (J, K) \in \vec{\mathcal{S}}_k$ we define the $\wedge$ operator as follows:

$$
\begin{aligned}
\overrightarrow{(A, B)} \wedge \overrightarrow{(J, K)} &:= \overrightarrow{(A \cap J, B \cup K)}, \\
\overrightarrow{(A, B)} \wedge \overleftarrow{(J, K)} &:= \overrightarrow{(A \cap K, B \cup J)}, \\
\overleftarrow{(A, B)} \wedge \overrightarrow{(J, K)} &:= \overrightarrow{(B \cap J, A \cup K)}, \text{ and} \\
\overleftarrow{(A, B)} \wedge \overleftarrow{(J, K)} &:= \overrightarrow{(B \cap K, A \cup J)}.
\end{aligned}
$$

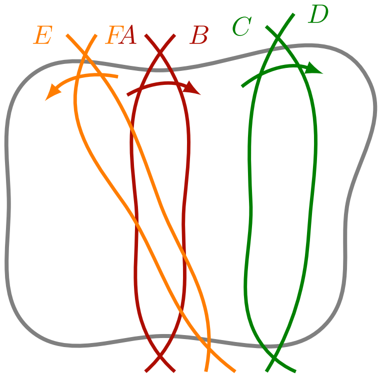

Figure 3.1: The figure shows three directed separations $\overrightarrow{(A,B)}$, $\overrightarrow{(C,D)}$ and $\overleftarrow{(E,F)}$. As $A \subseteq C$ and $D \subseteq B$, the first two are laminar: $\overrightarrow{(A,B)} \leq \overrightarrow{(C,D)}$. On the other hand, (E,F) crosses both (A,B) and (C,D).

The $\vee$ operator we define in the corresponding fashion:

$$\begin{aligned}
\overrightarrow{(A,B)} \vee \overrightarrow{(J,K)} &:= \overrightarrow{(A \cup J, B \cap K)}, \\
\overrightarrow{(A,B)} \vee \overleftarrow{(J,K)} &:= \overrightarrow{(A \cup K, B \cap J)}, \\
\overleftarrow{(A,B)} \vee \overrightarrow{(J,K)} &:= \overrightarrow{(B \cup J, A \cap K)}, \text{ and} \\
\overleftarrow{(A,B)} \vee \overleftarrow{(J,K)} &:= \overrightarrow{(B \cup K, A \cap J)}.
\end{aligned}$$

The results obtained by using the operators $\wedge$ and $\vee$ are called the *uncrossings* of the two original directed separations. See Figure 3.2 for an illustration of the four cases. The uncrossings of two given directed separations are again directed separations.

Lemma 3.2.1 (Erde [Erd20]**).** Let D be a digraph and (A,B), (J,K) two directed separations in D. Then, $(A,B) \wedge (J,K)$ and $(A,B) \vee (J,K)$ are also directed separations in D.

Moreover, uncrossing two directed separations of order at most k yields at least one directed separation that again is of order at most k. This fact is well-known, however we provide a proof for one of the four cases for the sake of completeness.

Lemma 3.2.2. Let D be a digraph and (A,B), (J,K) two directed separations in D, then

$$|(A,B)| + |(J,K)| = |(A,B) \wedge (J,K)| + |(A,B) \vee (J,K)|. \tag{3.1}$$

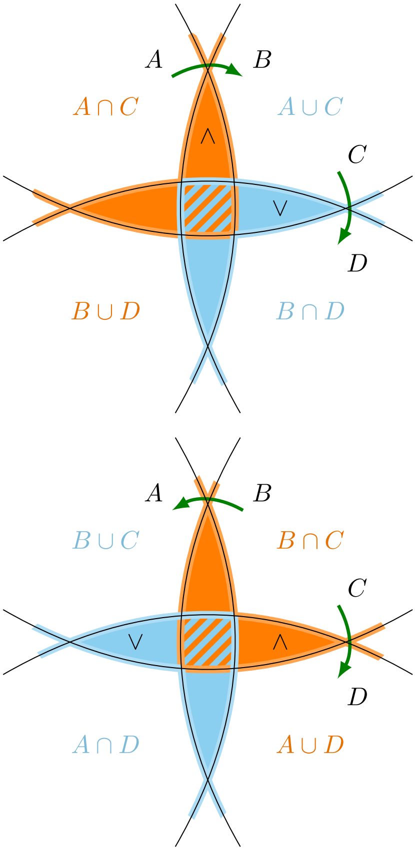

Figure 3.2: The four possible uncrossings of the directed separations (A, B) and (C, D), the separators for the $\wedge$-uncrossing is depicted in orange, and the separator for the $\vee$-uncrossing is depicted in blue.

Proof. For each of the four possibilities for the directions of (A, B) and (J, K) we consider one case. Let us investigate the case where $(A, B) = (A, \overleftarrow{B})$ and $(J, K) = (J, \overleftarrow{K})$. The proofs for the other three cases work similarly. Then,

$$(A, \overleftarrow{B}) \wedge (J, \overleftarrow{K}) = (A \cup J, B \cap \overleftarrow{K}) \text{ and}$$
$$(A, \overleftarrow{B}) \vee (J, \overleftarrow{K}) = (A \cap J, B \cup \overleftarrow{K}).$$

Now, $|\,(A,B)\,| = |A \cap B|$ and $|\,(J,K)\,| = |J \cap K|$. Moreover,

$$\left|\overleftarrow{(A,B)} \wedge \overleftarrow{(J,K)}\right| = |(A \cup J) \cap B \cap K| \text{ and}$$
$$\left|\overleftarrow{(A,B)} \vee \overleftarrow{(J,K)}\right| = |A \cap J \cap (B \cup K)|\,.$$

So a vertex v lies in the separator of both uncrossings if and only if it lies in $A \cap B \cap J \cap K$. Hence, we count a vertex twice on the left side of (3.1) if and only if it is counted twice on the right side of (3.1). Finally, observe that every vertex of $A \cap B \setminus (J \cap K)$ is contained in exactly one of the separators and the same holds for every vertex of $J \cap K \setminus (A \cap B)$. Thus, the equality (3.1) follows. □

Let us make some observations that describe how uncrossings affect the relation between separations. The following lemma investigates what happens if two laminar separations are uncrossed with the same separation.

Lemma 3.2.3. Let D be a digraph and $\overrightarrow{(J,K)}\,, \overrightarrow{(A,B)}\,, \overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$. Then, $\overrightarrow{(A,B)} \leq \overrightarrow{(J,K)}$ implies $\overrightarrow{(A,B)} \wedge \overrightarrow{(X,Y)} \leq \overrightarrow{(J,K)} \wedge \overrightarrow{(X,Y)}$ and $\overrightarrow{(A,B)} \vee \overrightarrow{(X,Y)} \leq \overrightarrow{(J,K)} \vee \overrightarrow{(X,Y)}$.

Proof. Both claims follow immediately from the fact that $A \subseteq J$ and $K \subseteq B$. □

The next lemma yields the distributivity of the two operations $\vee$ and $\wedge$ for uncrossings of separations.

Lemma 3.2.4. Let D be a digraph and $\overrightarrow{(J_1,K_1)}\,, \overrightarrow{(J_2,K_2)} \in \vec{\mathcal{S}}$ two crossing separations. Then, for all $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$ that cross both $\overrightarrow{(J_1,K_1)}$ and $\overrightarrow{(J_2,K_2)}$ we have

- $$\left(\overrightarrow{(J_1,K_1)} \wedge \overrightarrow{(J_2,K_2)}\right) \vee \overrightarrow{(X,Y)} = \left(\overrightarrow{(J_1,K_1)} \vee \overrightarrow{(X,Y)}\right) \wedge \left(\overrightarrow{(J_2,K_2)} \vee \overrightarrow{(X,Y)}\right),$$
- $$\left(\overrightarrow{(J_1,K_1)} \vee \overrightarrow{(J_2,K_2)}\right) \wedge \overrightarrow{(X,Y)} = \left(\overrightarrow{(J_1,K_1)} \wedge \overrightarrow{(X,Y)}\right) \vee \left(\overrightarrow{(J_2,K_2)} \wedge \overrightarrow{(X,Y)}\right),$$
- $$\left(\overrightarrow{(J_1,K_1)} \vee \overrightarrow{(J_2,K_2)}\right) \vee \overrightarrow{(X,Y)} = \left(\overrightarrow{(J_1,K_1)} \vee \overrightarrow{(X,Y)}\right) \vee \left(\overrightarrow{(J_2,K_2)} \vee \overrightarrow{(X,Y)}\right), \text{ and}$$

- $\left(\overrightarrow{(J_1, K_1)} \wedge \overrightarrow{(J_2, K_2)}\right) \wedge \overrightarrow{(X, Y)} = \left(\overrightarrow{(J_1, K_1)} \wedge \overrightarrow{(X, Y)}\right) \wedge \left(\overrightarrow{(J_2, K_2)} \wedge \overrightarrow{(X, Y)}\right).$

Proof. We only prove the first statement, the others can be proven in a similar way.

$$\begin{aligned}
&\left(\overrightarrow{(J_1, K_1)} \wedge \overrightarrow{(J_2, K_2)}\right) \vee \overrightarrow{(X, Y)} \\
=\ &\overrightarrow{(J_1 \cap J_2, K_1 \cup K_2)} \vee \overrightarrow{(X, Y)} \\
=\ &\overrightarrow{((J_1 \cap J_2) \cup X, (K_1 \cup K_2) \cap Y)} \\
=\ &\overrightarrow{((J_1 \cup X) \cap (J_2 \cup X), (K_1 \cap Y) \cup (K_2 \cap Y))} \\
=\ &\overrightarrow{(J_1 \cup X, K_1 \cap Y)} \wedge \overrightarrow{(J_2 \cup X, K_2 \cap Y)} \\
=\ &\left(\overrightarrow{(J_1, K_1)} \vee \overrightarrow{(X, Y)}\right) \wedge \left(\overrightarrow{(J_2, K_2)} \vee \overrightarrow{(X, Y)}\right)
\end{aligned}$$

□

3.2.2 Orientations of directed separations

As defined before, we denote directed separations as ordered tuples (A, B). Note that so far the order of A and B in the tuple does not make a difference to us. In particular, it is independent of the direction of the considered separation. Similarly as on undirected graphs, we introduce a meaning to the order of A and B now, by using it to denote the *orientation* of a separation in $\vec{\mathcal{S}}_k$. This does not have to coincide with the direction. So, (A, B) means that the separation is oriented *from A to B*, but its direction could be $\overleftarrow{(A, B)}$.

The concepts presented here closely follow the ideas from [DO21,Die18] together with the work of Erde [Erd20], who made first steps in adapting the notion of separation systems to the setting of directed graphs. We adapt some of the original definitions to our notation, which might seem more technical on occasion, but has its advantages when embedding the results stated into the greater context of directed graphs and orientations of directed separations.

(Partial) orientations Let D be a digraph. A *partial orientation* of $\vec{\mathcal{S}}_k$ is a set $\mathcal{O} \subseteq \vec{\mathcal{S}}_k$ such that for every $(A, B) \in \vec{\mathcal{S}}_k$, $\mathcal{O}$ contains at most one of (A, B) and (B, A).

We say that $\mathcal{O}$ is an *orientation* of $\vec{\mathcal{S}}_k$ if additionally $|\mathcal{O}| = \frac{|\vec{\mathcal{S}}_k|}{2}$, i.e. for every $(A, B) \in \vec{\mathcal{S}}_k$, $\mathcal{O}$ contains exactly one of (A, B) and (B, A). An orientation $\mathcal{O}$ *extends*

a partial orientation $\mathcal{P}$ if $\mathcal{P} \subseteq \mathcal{O}$. If $\mathcal{O}$ is a partial orientation of $\vec{\mathcal{S}}_k$ and $(A, B) \in \mathcal{O}$, we say that B is the *big side* of (A, B) and A is the *small side*. Thus, an oriented directed separation always *points away* from the small side and *towards* the big side. Using this terminology, an orientation of $\vec{\mathcal{S}}_k$ assigns to every directed separation of order less than k a unique big side towards which the separation points. If $\mathcal{O}$ is a partial orientation of $\vec{\mathcal{S}}_k$ and $\ell \leq k$, then we denote by $\mathcal{O}|_\ell$ the restriction of $\mathcal{O}$ to $\vec{\mathcal{S}}_\ell$. $\mathcal{O}|_\ell$
Formally, $\mathcal{O}|_\ell := \mathcal{O} \cap \vec{\mathcal{S}}_\ell$.

We say that a partial orientation $\mathcal{O}$ is *consistent* if it satisfies the following conditions:

1. if $\overset{\curvearrowright}{(A, B)} \in \mathcal{O}$ and $(J, K) \leq (A, B)$ for some $(J, K) \in \vec{\mathcal{S}}_k$, then
   - $(J, K) = \overset{\curvearrowright}{(J, K)}$ implies $(J, K) \in \mathcal{O}$, and
   - $(J, K) = \overset{\curvearrowleft}{(J, K)}$ implies $(K, J) \in \mathcal{O}$, and
2. if $\overset{\curvearrowleft}{(B, A)} \in \mathcal{O}$ and $(A, B) \leq (J, K)$ for some $(J, K) \in \vec{\mathcal{S}}_k$, then
   - $(J, K) = \overset{\curvearrowright}{(J, K)}$ implies $(K, J) \in \mathcal{O}$, and
   - $(J, K) = \overset{\curvearrowleft}{(J, K)}$ implies $(J, K) \in \mathcal{O}$.

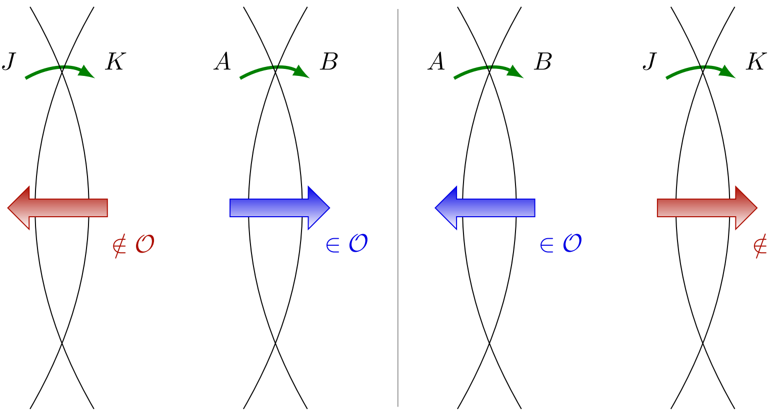

Figure 3.3: If the orientation $\mathcal{O}$ is fixed for (A, B) and consistent, then $\mathcal{O}$ cannot orient (J, K) in the indicated way.

Generally, $\vec{\mathcal{S}}_k$ admits multiple orientations. As for tangles in undirected graphs, we want the orientation to point towards the part of the graph containing the more complicated structure. So, the big side of an orientation shall contain most of the

structure of the graph. Thus, deciding on an orientation is straightforward for trivial separations: The separation should be oriented away from the side that is completely contained in the separator. This can be generalised to separations that are not trivial, but in which one of the sides contains only few vertices, where "few" simply means under a given threshold. These should also be oriented away from the side containing only few vertices. For $\omega \geq k$ we define the *initial ω-orientation* of $\vec{\mathcal{S}}_k$, denoted by $\mathcal{I}_k^\omega$
$\mathcal{I}_k^\omega$ as follows.

$$\mathcal{I}_k^\omega := \left\{ \overrightarrow{(A,B)} \in \vec{\mathcal{S}}_k \;\middle|\; |A| < \omega \right\} \cup \left\{ \overleftarrow{(B,A)} \in \vec{\mathcal{S}}_k \;\middle|\; |B| < \omega \right\}.$$

Using ω-orientations one can define an ω-diblockage as follows (also see Erde [Erd20]). An ω-diblockage in a digraph D for fixed k is an orientation $\mathcal{O}$ of $\vec{\mathcal{S}}_k$ such that

1. $\mathcal{O}$ extends $\mathcal{I}_k^\omega$,
2. $\mathcal{O}$ is consistent, and
3. if $\overrightarrow{(A,B)} \in \mathcal{O}$ and $\overrightarrow{(A,B)} \leq \overleftarrow{(K,J)} \in \mathcal{O}$ then $|B \cap J| \geq \omega$.

Intuitively, the definition expresses for any two laminar separations where the orientation of the smaller one respects its direction and the orientation of the larger one does not that the two big sides meet in a large number of vertices. The notion of diblockages corresponds to the digraph width measure *directed pathwidth* in the following way.

Theorem 3.2.5 (Erde [Erd20])**.** A digraph has directed pathwidth at least $\omega - 1$ if and only if it has an ω-diblockage of $\vec{\mathcal{S}}_k$.

As mentioned before, this result by Erde [Erd20] is the generalisation of the correspondence between blockages and pathwidth in undirected graphs. We refer the reader to his paper for the definition of directed pathwidth. We seek to further generalise this to cover weaker directed width measures than directed pathwidth.

3.2.3 Introducing $\vec{\mathcal{S}}$-DAGs

We now introduce the structure that provides the general backbone to our decompositions.

Definition 3.2.6 (consistency)**.** Let D be a digraph and $\vec{\mathcal{S}}' \subseteq \vec{\mathcal{S}}$ a set of directed separations in D. An *$\vec{\mathcal{S}}'$-DAG* for D is a tuple (T, σ), where T is a DAG of maximum out-degree at most two and $\sigma\colon E(T) \to \vec{\mathcal{S}}'$ is a function such that all for $t_1, t_2, t_3 \in V(T)$ with $(t_1, t_2), (t_2, t_3) \in E(T)$ we have $\sigma((t_1, t_2)) \leq \sigma((t_2, t_3))$. We call this property the *consistency* of $\vec{\mathcal{S}}'$-DAGs. $\dashv$

Given an $\vec{\mathcal{S}}$-DAG (T,σ) for a digraph D and a node $d \in V(T)$, we define two DAGs which are subgraphs of T as follows. By $T^{\uparrow}(d)$ we denote $T - (\mathsf{Above}_T(d) \setminus \{d\})$, and by $T^{\downarrow}(d)$ we denote $T - (\mathsf{Below}_T(d) \setminus \{d\})$. If $T' \subseteq T$, then we relax our notation a little and write (T',σ) for $\left(T', \sigma|_{E(T')}\right)$. $T^{\uparrow}(d)$ $T^{\downarrow}(d)$

Let (T,σ) be an $\vec{\mathcal{S}}'$-DAG for D and $t \in V(T)$. We can find a canonical directed separation $\top^{(T,\sigma)}(t)$ for every node of T by using the directed separations given by σ. If $t \in \mathsf{sources}_T$, we set $\top^{(T,\sigma)}(t) = \overrightarrow{(\emptyset, V(D))}$. For non-source nodes $t \in V(T)$, we define $\top^{(T,\sigma)}(t)$

$$\top^{(T,\sigma)}(t) := \overrightarrow{\left(\top_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)} := \bigvee_{d \in N_T^{\mathsf{in}}(t)} \sigma((d,t))\,.$$

$\top_A^{(T,\sigma)}(t)$ $\top_B^{(T,\sigma)}(t)$

If the $\vec{\mathcal{S}}$-DAG is provided by the context, we also write $\top(t)$, and $\overrightarrow{(\top_A(t), \top_B(t))}$. We also write $\overrightarrow{(\sigma_C(e), \sigma_D(e))} := \sigma(e)$ for the separation mapped to $e \in E(T)$ by the function σ. $\overrightarrow{(\sigma_C(e), \sigma_D(e))}$

Lemma 3.2.7. Let D be a digraph and (T,σ) be an $\vec{\mathcal{S}}$-DAG for D. Also, let $d \in V(T)$ be a node of T with an in-neighbour t_1 and an out-neighbour t_2. Then, $\sigma((t_1,d)) \le \top(d) \le \sigma((d,t_2))$.

Proof. The first inequality follows immediately from the definition of $\top(d)$. We have $\overrightarrow{(\sigma_C((t_1,d)), \sigma_D((t_1,d)))} \le \overrightarrow{(\sigma_C((d,t_2)), \sigma_D((d,t_2)))}$ for all $(t_1,d) \in E(T)$ due to the consistency of $\vec{\mathcal{S}}$-DAGs. Thus,

$$\top_A(d) = \bigcup_{(t_1,d)\in E(T)} \sigma_C((t_1,d)) \subseteq \sigma_C((d,t_2)) \text{ and}$$
$$\sigma_D((d,t_2)) \subseteq \bigcap_{(t_1,d)\in E(T)} \sigma_D((t_1,d)) = \top_B(d)\,,$$

and therefore we obtain

$$\begin{aligned}\top(d) &= \overrightarrow{(\top_A(d)\,, \top_B(d))}\\ &\le (\sigma_C((d,t_2))\,, \sigma_D((d,t_2))) = \overrightarrow{(\sigma_C((d,t_2))\,, \sigma_D((d,t_2)))}\,.\end{aligned}$$ □

Let D be a digraph and (T,σ) an $\vec{\mathcal{S}}$-DAG for D. With $\top^{(T,\sigma)}(t)$ we have derived a canonical directed separation for every node of T from the directed separations given by σ. Let $d \in V(T)$ be a node with out-neighbours t_1 and t_2, then $\top^{(T,\sigma)}(t) \le \top^{(T,\sigma)}(t_i)$ for $i \in \{1,2\}$, but $\top^{(T,\sigma)}(t_1)$ and $\top^{(T,\sigma)}(t_2)$ are not necessarily laminar.

With the operators $\wedge$ and $\vee$ we are given two possible ways to uncross the $\top(t_i)$. One way is given by $\perp^{(T,\sigma)}(t)$, defined as

$\perp^{(T,\sigma)}(t)$

$$\perp^{(T,\sigma)}(t) := \bigwedge_{d \in N_T^{\mathsf{out}}(t)} \top^{(T,\sigma)}(d) =$$

$$\left(\bigcap_{d \in N_T^{\mathsf{out}}(t)} \bigcup_{\ell \in N_T^{\mathsf{in}}(d)} \sigma_C((\ell, d)), \bigcup_{d \in N_T^{\mathsf{out}}(t)} \bigcap_{\ell \in N_T^{\mathsf{in}}(d)} \sigma_D((\ell, d)) \right)$$

for all non-sink nodes t in T. We choose this way of uncrossing as the obtained separation lies nearer to $\top(t)$. If t is a sink, we set $\perp^{(T,\sigma)}(t) := \overrightarrow{(V(D), \emptyset)}$. Again, we write $\perp(t)$ and $\overrightarrow{(\perp_A(t), \perp_B(t))}$ if the $\vec{\mathcal{S}}$-DAG is clear from the context.

Lemma 3.2.8. Let D be a digraph and (T, σ) be an $\vec{\mathcal{S}}$-DAG for D. Moreover, let $(d, t) \in E(T)$. Then, $\overrightarrow{(\top_A(d), \top_B(d))} \le \overrightarrow{(\perp_A(d), \perp_B(d))} \le \overrightarrow{(\top_A(t), \top_B(t))}$.

Proof. The second inequality holds by definition, therefore we only need to validate the first one.

Consider a node $d \in V(T)$. By definition, we have $\top_A(d) = \bigcup_{d' \in N_T^{\mathsf{in}}(d)} \sigma_C((d', d))$. As Lemma 3.2.7 implies that for all $d' \in N_T^{\mathsf{in}}(d)$ and for all $t \in N^{\mathsf{out}}(d)$ holds $\sigma_C((d', d)) \subseteq \sigma_C((d, t))$, we obtain $\bigcup_{d' \in N_T^{\mathsf{in}}(d)} \sigma_C((d', d)) \subseteq \bigcap_{t \in N_T^{\mathsf{out}}(d)} \sigma_C((d, t))$. Using that $\bigcap_{t \in N_T^{\mathsf{out}}(d)} \sigma_C((d, t)) \subseteq \bigcap_{t \in N_T^{\mathsf{out}}(d)} \bigcup_{t' \in N_T^{\mathsf{in}}(t)} \sigma_C((t', t)) = \perp_A(d)$, we obtain $\top_A(d) \subseteq \perp_A(d)$.

We know $\perp_B(d) = \bigcup_{t \in N_T^{\mathsf{out}}(d)} \bigcap_{t' \in N_T^{\mathsf{in}}(t)} \sigma_D((t', t)) \subseteq \bigcup_{t \in N_T^{\mathsf{out}}(d)} \sigma_D((d, t))$. As Lemma 3.2.7 implies that $\sigma_D((d, t)) \subseteq \sigma_D((d'd))$ for all $d' \in N_T^{\mathsf{in}}(d)$ and for all $t \in N_T^{\mathsf{out}}(d)$, we obtain that $\bigcup_{t \in N_T^{\mathsf{out}}(d)} \sigma_D((d, t)) \subseteq \bigcap_{d' \in N_T^{\mathsf{in}}(d)} \sigma_D((d', d)) = \top_B(d)$. Thus, $\perp_B(d) \subseteq \top_B(d)$.

Together $\top_A(d) \subseteq \perp_A(d)$ and $\perp_B(d) \subseteq \top_B(d)$ imply $\top(d) \le \perp(d)$. $\square$

We associate two special separations with each $\vec{\mathcal{S}}$-DAG (T, σ), namely the *initial separation* $\mathsf{Initial}(T, \sigma)$ and the *terminal separation* $\mathsf{Terminal}(T, \sigma)$.

$$\mathsf{Initial}(T, \sigma) := \bigwedge_{r \in \mathsf{sources}_T} \overrightarrow{\left(\perp_A^{(T,\sigma)}(r), \perp_B^{(T,\sigma)}(r)\right)}, \text{ and}$$

$$\mathsf{Terminal}(T, \sigma) := \bigvee_{s \in \mathsf{sinks}_T} \overrightarrow{\left(\top_A^{(T,\sigma)}(s), \top_B^{(T,\sigma)}(s)\right)}.$$

These separations yield an overall maximum and an overall minimum for all separations in the $\vec{\mathcal{S}}_k$-DAG, which we utilise later in order to combine $\vec{\mathcal{S}}_k$-DAGs.

For two laminar directed separations $\overrightarrow{(A,B)}$ and (X,Y), the directed separation of smallest order that lies between $\overrightarrow{(A,B)}$ and (X,Y) with respect to the partial order $\leq$ has useful and interesting properties. For $\overrightarrow{(A,B)} \leq \overrightarrow{(X,Y)}$ we define a way to obtain the order of such a separation by

$$\lambda\Big(\overrightarrow{(A,B)}, \overrightarrow{(X,Y)}\Big) := \min\Big\{|(J,K)| \;\Big|\; \overrightarrow{(A,B)} \leq \overrightarrow{(J,K)} \leq \overrightarrow{(X,Y)}\Big\}.$$

We say that $\overrightarrow{(X,Y)}$ is *down-linked* to $\overrightarrow{(A,B)}$ if $\overrightarrow{(A,B)} \leq \overrightarrow{(X,Y)}$ and additionally $|(X,Y)| = \lambda((A,B),(X,Y))$. Analogously, we say that $\overrightarrow{(X,Y)}$ is *up-linked* to $\overrightarrow{(A,B)}$ if $\overrightarrow{(X,Y)} \leq \overrightarrow{(A,B)}$ and $|(X,Y)| = \lambda((X,Y),(A,B))$.

The following lemma shows that uncrossing a separation $\overrightarrow{(J,K)}$ that is comparable to a given separation $\overrightarrow{(A,B)}$ and larger with respect to $\leq$ with a separation $\overrightarrow{(X,Y)}$ that is down-linked to $\overrightarrow{(A,B)}$ cannot result in a separation of higher order.

Lemma 3.2.9. Let D be a digraph and $\overrightarrow{(A,B)}, \overrightarrow{(J,K)} \in \vec{\mathcal{S}}$ with $\overrightarrow{(A,B)} \leq \overrightarrow{(J,K)}$. Then, for every $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$ that is down-linked to $\overrightarrow{(A,B)}$ we have

$$\left|\overrightarrow{(J,K)} \vee \overrightarrow{(X,Y)}\right| \leq \left|\overrightarrow{(J,K)}\right|.$$

Proof. As $\overrightarrow{(X,Y)}$ is down-linked to $\overrightarrow{(A,B)}$, we know that $\overrightarrow{(A,B)} \leq \overrightarrow{(X,Y)}$. Together with $\overrightarrow{(A,B)} \leq \overrightarrow{(J,K)}$, this implies that

$$\overrightarrow{(A,B)} \leq \overrightarrow{(J,K)} \wedge \overrightarrow{(X,Y)} \leq \overrightarrow{(X,Y)}.$$

Due to $|(X,Y)| = \lambda((A,B),(X,Y))$, we obtain

$$\left|\overrightarrow{(J,K)} \wedge \overrightarrow{(X,Y)}\right| \geq \left|\overrightarrow{(X,Y)}\right|.$$

Then, using Lemma 3.2.2, we conclude

$$\left|\overrightarrow{(J,K)} \vee \overrightarrow{(X,Y)}\right| \leq \left|\overrightarrow{(J,K)}\right|.$$ □

A symmetric bound can also be obtained for the up-linked case.

Lemma 3.2.10. Let D be a digraph and $\overrightarrow{(J,K)}, \overrightarrow{(A,B)} \in \vec{\mathcal{S}}$ with $\overrightarrow{(J,K)} \leq \overrightarrow{(A,B)}$. Then, for every $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$ that is up-linked to $\overrightarrow{(A,B)}$, we have

$$\left|\overrightarrow{(J,K)} \wedge \overrightarrow{(X,Y)}\right| \leq \left|\overrightarrow{(J,K)}\right|.$$

Proof. As $\overrightarrow{(X,Y)}$ is up-linked to $\overrightarrow{(A,B)}$, we have $\overrightarrow{(X,Y)} \leq \overrightarrow{(A,B)}$. Thus, $\overrightarrow{(J,K)} \leq \overrightarrow{(A,B)}$ implies

$$\overrightarrow{(X,Y)} \leq \overrightarrow{(J,K)} \vee \overrightarrow{(X,Y)} \leq \overrightarrow{(A,B)}.$$

We also have $|\,(X,Y)\,| = \lambda((X,Y),(A,B))$, so

$$\left|\overrightarrow{(X,Y)}\right| \leq \left|\overrightarrow{(J,K)} \vee \overrightarrow{(X,Y)}\right|$$

Then, using Lemma 3.2.2, we conclude

$$\left|\overrightarrow{(J,K)} \wedge \overrightarrow{(X,Y)}\right| \leq \left|\overrightarrow{(J,K)}\right|. \qquad \square$$

We can manipulate a $\vec{\mathcal{S}}$-DAG (T,σ) by uncrossing all separations lying on edges above or below a certain node in T with a given separation, which then becomes the initial or terminal separation of the new $\vec{\mathcal{S}}$-DAG. In order to formalise the concept of uncrossing these separations, we introduce the following two definitions.

Definition 3.2.11 (down-shift)**.** Let (T,σ) be an $\vec{\mathcal{S}}$-DAG for a digraph D, $d \in V(T)$, and $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$ with $\mathsf{Initial}(T^{\uparrow}(d),\sigma) \leq \overrightarrow{(X,Y)}$. The *down-shift* of (T,σ) onto $\overrightarrow{(X,Y)}$ at d is the $\vec{\mathcal{S}}$-DAG $\left(T^{\uparrow}(d), \sigma^{\uparrow}(d)\right)$ where for all $(t,t') \in E(T^{\uparrow}(d))$ we define

$$\sigma^{\uparrow}(d)((t,t')) := \sigma((t,t')) \vee \overrightarrow{(X,Y)} = \overrightarrow{(\sigma_C((t,t')) \cup X, \sigma_D((t,t')) \cap Y)}. \qquad \dashv$$

Note that after this down-shift $\overrightarrow{(X,Y)}$, is the initial separation.

Observation 3.2.12. Let (T,σ) be an $\vec{\mathcal{S}}$-DAG for a digraph D, $d \in V(T)$, and $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}$ such that $\mathsf{Initial}(T^{\uparrow}(d),\sigma) \leq \overrightarrow{(X,Y)}$. If $\left(T^{\uparrow}(d), \sigma^{\uparrow}(d)\right)$ is a down-shift of (T,σ) onto $\overrightarrow{(X,Y)}$ at d, then $\overrightarrow{(X,Y)}$ is its initial separation, that is,

$$\mathsf{Initial}\left(T^{\uparrow}(d), \sigma^{\uparrow}(d)\right) = \overrightarrow{(X,Y)}.$$

Similarly, we define the up-shift, which changes the terminal separation.

Definition 3.2.13 (up-shift)**.** Let (T, σ) be an $\vec{\mathcal{S}}$-DAG for a digraph D, $d \in V(T)$, and $\overrightarrow{(X, Y)} \in \vec{\mathcal{S}}$ such that $\overrightarrow{(X, Y)} \leq \mathsf{Terminal}(T^{\downarrow}(d), \sigma)$. Then, the *up-shift* of (T, σ) onto $\overrightarrow{(X, Y)}$ at d is the $\vec{\mathcal{S}}$-DAG $\big(T^{\downarrow}(d), \sigma^{\downarrow}(d)\big)$ where for all $(t, t') \in E(T^{\downarrow}(d))$ we define

$$\sigma^{\downarrow}(d)((t, t')) := \sigma((t, t')) \wedge \overrightarrow{(X, Y)} = \overrightarrow{(\sigma_C((t, t')) \cap X, \sigma_D((t, t')) \cup Y)}. \qquad \dashv$$

Observation 3.2.14. Let (T, σ) be an $\vec{\mathcal{S}}$-DAG for a digraph D, $d \in V(T)$, and $\overrightarrow{(X, Y)} \in \vec{\mathcal{S}}$ a separation such that $\overrightarrow{(X, Y)} \leq \mathsf{Terminal}(T^{\downarrow}(d), \sigma)$. If $\big(T^{\downarrow}(d), \sigma^{\downarrow}(d)\big)$ is an up-shift of (T, σ) onto the separation $\overrightarrow{(X, Y)}$ at d, then $\overrightarrow{(X, Y)}$ is the terminal separation of the up-shift, that is,

$$\mathsf{Terminal}\big(T^{\downarrow}(d)\,, \sigma^{\downarrow}(d)\big) = \overrightarrow{(X, Y)}.$$

Given a directed separation $\overrightarrow{(X, Y)}$ and a node $d \in V(T)$, we say that (X, Y) is *down-linked to d* if it is down-linked to $\mathsf{Initial}(T^{\uparrow}(d), \sigma)$. Similarly, (X, Y) is *up-linked to d* if it is up-linked to the separation $\mathsf{Terminal}(T^{\downarrow}(d), \sigma)$. Sometimes, we do not wish to specify whether we are dealing with an up-shift or down-shift, and we refer to a *shift*[1] of an $\vec{\mathcal{S}}$-DAG instead. In a similar way, we call a separation $\overrightarrow{(X, Y)}$ *linked to d* if it is up-linked or down-linked to d. A down-shift of (T, σ) is $\overrightarrow{(X, Y)}$*-d-admissible* if $\overrightarrow{(X, Y)}$ is down-linked to d. An up-shift of (T, σ) is $\overrightarrow{(X, Y)}$-d-admissible if $\overrightarrow{(X, Y)}$ is up-linked to d.

Shifts are defined such that they preserve certain properties of the $\vec{\mathcal{S}}$-DAG that was shifted. A rather straightforward observation to make is the following.

Lemma 3.2.15. Let (T, σ) be an $\vec{\mathcal{S}}$-DAG of digraph D. Let (T', σ') be the $\vec{\mathcal{S}}$-DAG obtained by a down-shift of (T, σ) onto $\overrightarrow{(X_1, Y_1)}$ at a node $d \in V(T)$ and let (T'', σ'') be the $\vec{\mathcal{S}}$-DAG obtained by an up-shift of (T, σ) onto $\overrightarrow{(X_2, Y_2)}$ at a node $d \in V(T)$. Then

1. $\perp^{(T', \sigma')}(t) = \perp^{(T, \sigma)}(t) \vee \overrightarrow{(X_1, Y_1)}$ for all $t \in V(T')$,
2. $\top^{(T', \sigma')}(t) = \top^{(T, \sigma)}(t) \vee \overrightarrow{(X_1, Y_1)}$ for all $t \in V(T')$,
3. $\perp^{(T'', \sigma'')}(t) = \perp^{(T, \sigma)}(t) \wedge \overrightarrow{(X_2, Y_2)}$ for all $t \in V(T'')$, and
4. $\top^{(T'', \sigma'')}(t) = \top^{(T, \sigma)}(t) \wedge \overrightarrow{(X_2, Y_2)}$ for all $t \in V(T'')$.

[1] Note that up- and down-shifts are named the other way round by Erde [Erd20], but as we consider DAGs to have their sources at the top and their sinks at the bottom this naming would be very unintuitive here.

Proof. If $t \in \mathsf{sinks}_{T'}$, then the first statement holds. So we may assume that $t \in V(T') \setminus \mathsf{sinks}_{T'}$. Then,

$$\begin{aligned}
&\perp^{(T',\sigma')}(t) \\
&= \left(\bigcap_{d \in N_T^{\mathsf{out}}(t)} \bigcup_{\ell \in N_T^{\mathsf{in}}(d)} \sigma'_C((\ell,d)),\ \bigcup_{d \in N_T^{\mathsf{out}}(t)} \bigcap_{\ell \in N_T^{\mathsf{in}}(d)} \sigma'_D((\ell,d)) \right) \\
&= \left(\bigcap_{d \in N_T^{\mathsf{out}}(t)} \bigcup_{\ell \in N_T^{\mathsf{in}}(d)} (\sigma_C((\ell,d)) \cup X_1),\ \bigcup_{d \in N_T^{\mathsf{out}}(t)} \bigcap_{\ell \in N_T^{\mathsf{in}}(d)} (\sigma_D((\ell,d)) \cap Y_1) \right) \\
&= \left(\bigcap_{d \in N_T^{\mathsf{out}}(t)} \bigcup_{\ell \in N_T^{\mathsf{in}}(d)} \sigma_C((\ell,d)) \cup X_1,\ \left(\bigcup_{d \in N_T^{\mathsf{out}}(t)} \bigcap_{\ell \in N_T^{\mathsf{in}}(d)} \sigma_D((\ell,d)) \right) \cap Y_1 \right) \\
&= \perp^{(T,\sigma)}(t) \vee \widehat{(X_1, Y_1)}.
\end{aligned}$$

The remaining three statements can be proven similarly. □

3.3 A general duality theorem

In order to use $\vec{\mathcal{S}}$-DAGs to define width measures, we need an analogue to a bag-function that is obtained from the separations in the $\vec{\mathcal{S}}$-DAG, and possibly restrictions on which $\vec{\mathcal{S}}$-DAGs even qualify as base of a decomposition. In this section, we determine the definitions that yield a width measure. Then, we prove a quite general duality theorem, in Theorem 3.3.13, between a type of ganglion and a type of width measure based on the same set of properties, provided that the width measure also fulfils certain robustness criteria.

To this end we define a prop-$\vec{\mathcal{S}}$*-DAG* (or only prop-*DAG*) to be an $\vec{\mathcal{S}}$-DAG satisfying some additional properties specified by prop. If these properties also yield a bound k on the order of the separations involved, we call these prop-k*-DAGs*. In the remainder of this section, we consider prop to be a specific set of properties.

During this section we collect the definitions needed for a width measure. For our purposes we need prop to provide the following requirements for any given prop-DAG (T,σ) :

$\beta_{\mathsf{prop}}^{(T,\sigma)}$ **(P1)** a *canonical bag function* $\beta_{\mathsf{prop}}^{(T,\sigma)} : V(T) \to 2^{V(D)}$,

$\mathsf{eval}_{\mathsf{prop}}$ **(P2)** an *evaluation function* $\mathsf{eval}_{\mathsf{prop}} : 2^{V(D)} \to \mathbb{N}$ for the bags.

This list is still incomplete and we complete it throughout this section.

The simplest evaluation function is just taking the bag size, which we also do throughout this chapter, but we leave the possibility to choose something more elaborate as well. The prop-*width* of a prop-k-DAG (T, σ) is the maximum $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d))$ for any $d \in V(T)$. Further, the prop-*DAG-width* of a digraph D is the minimum prop-width of any prop-k-DAG of D.

We want the evaluation of a bag to have the following quite reasonable properties. It should not grow when adding another incoming edge to its node as this only moves the top and the bottom separation "closer" to each other. Similarly, if we uncross all separations on the in-edges with a separation that is already "pushed towards" the terminal separation of the DAG above the node as far as possible, then the bag size also should not change.

Formally, let (T, σ) be an $\vec{\mathcal{S}}_k$-DAG of width at most ω and let $v \in V(T)$ be a node in a prop-k-DAG with $\overrightarrow{(J_1, K_1)}$ and $\overrightarrow{(J_2, K_2)}$ being the two separations on the out-edges of v.

(P3) If we add an edge to T with head v, then $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}(v)) < \omega$.

(P4) If we replace all separations on the in-edges of v by their uncrossing with a separation $\overrightarrow{(X, Y)} \in \vec{\mathcal{S}}_k$ that is up-linked to a separation $\overrightarrow{(J_2', K_2')}$ with $\overrightarrow{(J_2, K_2)} \leq \overrightarrow{(J_2', K_2')}$, then $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}(v)) < \omega$.

Above we describe shifts on $\vec{\mathcal{S}}$-DAGs. We would like prop-DAGs to remain prop-DAGs and also to retain additional features throughout a shifting process. This is formalised by the following definition.

Definition 3.3.1. We say that prop-DAG-width is *shifting closed* if for every digraph D, every prop-DAG (T, σ) for D, every $d \in V(T)$, and every directed separation $\overrightarrow{(X, Y)}$, every $\overrightarrow{(X, Y)}$-d-admissible shift (T', σ') of (T, σ) satisfies the following requirements.

(S1) (T', σ') is a prop-DAG,

(S2) if (T, σ) is a prop-k-DAG, then so is (T', σ'), and

(S3) $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T',\sigma')}(d)) \leq \mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d))$ for all $d \in V(T')$. $\dashv$

This is a rather strong demand on the prop-DAG-width and it turns out that a slightly weaker notion suffices to prove the general duality theorem. We do not need to consider all prop-DAGs, but we can restrict the shifts we need to perform to specific ones. A

DAG is called *closed* if it has exactly one source and exactly one sink. Similarly, we call a prop-DAG (T, σ) *closed* if T is closed. This definition allows us to introduce a slightly weaker version of being shifting closed.

Definition 3.3.2. We say that prop-DAG-width is *weakly shifting closed* if for every digraph D, every closed prop-DAG (T, σ) for D, every $d \in \mathsf{sources}_T \cup \mathsf{sinks}_T$, and every directed separation $\overrightarrow{(X, Y)}$, every $\overrightarrow{(X, Y)}$-d-admissible shift (T', σ') of (T, σ) meets the following requirements.

(W1) (T', σ') is a prop-DAG,

(W2) if (T, σ) is a prop-k-DAG, then so is (T', σ'), and

(W3) $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T',\sigma')}(d)) \leq \mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d))$ for all $d \in V(T')$. ⊣

Note that every shifting closed prop-DAG-width is also weakly shifting closed, but the reverse is not true in general. Being weakly shifting closed is a property that is only defined for closed prop-DAGs. So, we need a way of obtaining closed prop-DAGs from general ones. For a given DAG it is rather straightforward to find a closed DAG containing it as a subgraph. We explain how this is done and then also generalise the procedure to prop-DAGs.

Let T be a DAG. A DAG T' is a *closure* of T if T' is a closed DAG and $T \subseteq T'$. Similarly, a *closure* of (T, σ) is a closed prop-DAG (T', σ') such that $T \subseteq T'$ and $\sigma'|_T = \sigma$. The closure of a closed DAG is the DAG itself. For any other DAG, a closure can be easily obtained as follows. Let T be a DAG that is not closed. The *up-closing* of T at $r_1, r_2 \in \mathsf{sources}_T$ is obtained by adding a new node r^+ together with the two edges (r^+, r_1) and (r^+, r_2) to the DAG T. Clearly, the resulting digraph is again a DAG. Similarly, the *down-closing* of T at the sinks $s_1, s_2 \in \mathsf{sinks}_T$ is obtained by adding a new node s^- together with the edges (s_1, s^-) and (s_2, s^-) to T. Repeatedly applying these two operations yields a closure of T.

In order to define define a similar operation for prop-DAGs, every prop-DAG-width must provide for a fixed pair of sources or sinks a new directed separation to which the newly added edges from the up- or down-closing operation of the DAG can be mapped. So, we enlarge the list of requirements the width measure has to provide. If (T, σ) is a prop-DAG with two sources $r_1, r_2 \in \mathsf{sources}_T$ and two sinks $s_1, s_2 \in \mathsf{sinks}_T$, and T' is the up-closing of T at r_1 and r_2 and T'' is the down-closing of T at s_1 and s_2, then prop has to provide

(P5) a prop-*up-closing* of (T, σ) at r_1 and r_2 which is a prop-DAG (T', σ') where $\sigma = \sigma'|_T$, and

(P6) a prop-*down-closing* of (T, σ) at s_1 and s_2 which is a prop-DAG (T'', σ'') where $\sigma = \sigma''|_T$.

Let (T, σ) be a prop-DAG. A prop-*up-closure* of (T, σ) is a prop-DAG (T', σ') that has a unique source and $T \subseteq T'$ as well as $\sigma = \sigma'|_T$. Similarly, a prop-*down-closure* of (T, σ) is a prop-DAG (T', σ') that has a unique sink and $T \subseteq T'$ as well as $\sigma = \sigma'|_T$. Note that a prop-DAG (T', σ') is a closure of (T, σ) if and only if it is both a prop-up-closure and a prop-down-closure of (T, σ) . Similarly as for DAGs, we can now obtain a closure of a prop-DAG by repeatedly applying prop-up-closings and prop-down-closings.

This finishes the list of requirements prop has to provide in order for prop-DAG-width to be eligible for the general duality theorem and directly leads us to additional robustness criteria these definitions have to fulfil in order to prove the main theorem of this chapter, the general duality theorem.

Definition 3.3.3. The width measure prop-DAG-width is called *closed* if it meets the following requirements:

(C1) For every digraph D and every prop-DAG (T, σ) there exists a closure of (T, σ) of the same prop-width,

(C2) for every digraph D and every closed prop-DAG (T, σ) of prop-width less than k there exists a closed prop-k-DAG (T', σ') of the same prop-width,

(C3) for every digraph D, every closed prop-k-DAG (T, σ) with unique source r, and every $\overrightarrow{(X, Y)} \in \vec{\mathcal{S}}_k$ with $\overrightarrow{(X, Y)} \leq \mathsf{Initial}(T, \sigma)$, the $\vec{\mathcal{S}}_k$-DAG (T', σ') obtained from (T, σ) by introducing a new node r^+ together with the edge (r^+, r) and by extending σ to σ' by setting $\sigma'((r^+, r)) = \overrightarrow{(X, Y)}$, is a prop-k-DAG,

(C4) prop-DAG-width is weakly shifting closed. ⊣

We also want some connection between the evaluation of the bags and the order of the separations used in the decomposition. In order to describe this demand formally we obtain a function $\beta_{\mathsf{prop}} : \vec{\mathcal{S}} \times \vec{\mathcal{S}} \times \vec{\mathcal{S}} \to 2^{V(D)}$, which allows us to use the evaluation function in order to evaluate triples of separations as well. This we do by constructing a special prop-DAG (T', σ') and using the canonical bag function of prop as follows. We define $T = (\{t_1, t_2, t_3, t_4\}, \{e_1 = (t_1, t_2), e_2 = (t_2, t_3), e_3 = (t_2, t_4)\})$, which allows us to define $\beta_{\mathsf{prop}}(S_1, S_2, S_3) := \beta_{\mathsf{prop}}^{(T, \sigma = \{e_i \mapsto S_i \mid i \in \{1,2,3\}\})}(t_2)$ for $S_1, S_2, S_3 \in \vec{\mathcal{S}}$. If $S_1 = S_2 = S_3$, then we simply write $\beta_{\mathsf{prop}}(S_1)$ instead of $\beta_{\mathsf{prop}}(S_1, S_2, S_3)$ in favour of readability. β_{prop}

For separations of small order, we also want the resulting bag to be small in order to obtain the wanted relation between the bags given by the canonical bag function and the order of the separations of the decomposition.

Definition 3.3.4. The width measure prop-DAG-width is called *contained* if

(Cnt) $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}(\overrightarrow{(X,Y)})) < k$ for all $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}_k$. ⊣

If the measure prop-DAG-width is closed and contained, we say that prop-DAG-width is *complete*. By now, we have collected a number of properties and robustness criteria for the width measure. We now also call prop *complete* if it fulfils (P1) to (P6) and gives rise to a complete width measure prop-DAG-width.

We now define the objects that yield obstructions to small prop-DAG-width.

Definition 3.3.5. Let D be a digraph and $k, \omega \in \mathbb{N}$ with $\omega \geq k$. We say that an orientation $\mathcal{O}$ of $\vec{\mathcal{S}}_k$ is an *ω-big* prop-*ganglion* of order k if $\mathcal{O}$ meets the following requirements.

- **($\mathcal{O}$1)** $\mathcal{O}$ extends $\mathcal{I}_k^\omega$,
- **($\mathcal{O}$2)** $\mathcal{O}$ is consistent, and
- **($\mathcal{O}$3)** for every triple of directed separations $\overrightarrow{(A,B)}, \overleftarrow{(K_1,J_1)}, \overleftarrow{(K_2,J_2)} \in \mathcal{O}$ with $\overrightarrow{(A,B)} \leq \overleftarrow{(K_1,J_1)}$ and $\overrightarrow{(A,B)} \leq \overleftarrow{(K_2,J_2)}$ we have
$$\mathsf{eval}_{\mathsf{prop}}\Big(\beta_{\mathsf{prop}}\Big(\overrightarrow{(A,B)}, \overleftarrow{(K_1,J_1)}, \overleftarrow{(K_2,J_2)}\Big)\Big) \geq \omega. \qquad ⊣$$

The overall goal is to obtain a duality theorem between prop-DAG-width and ω-big prop-ganglions. The next lemma shows that containing an ω-big prop-ganglion prevents the existence of a prop-DAG of small width and thus proves that ω-big prop-ganglions are indeed an obstruction.

Lemma 3.3.6. Let $\mathcal{D}$ be a ω-big prop-ganglion of order k. Then, for every closed prop-k-DAG (T,σ) of D, there exists a node $d \in V(T)$ such that $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d)) \geq \omega$.

Proof. Let (T,σ) be a closed prop-k-DAG for D with the unique source r and the unique sink s. We know that $\bot(r), \top(s) \in \vec{\mathcal{S}}_k$. If for some $v \in \{r,s\}$ we have $\mathsf{eval}_{\mathsf{prop}}(\beta_\nu^{(T,\sigma)}(v)) \geq \omega$, then we are done. So assume $\mathsf{eval}_{\mathsf{prop}}(\beta_\nu^{(T,\sigma)}(v)) < \omega$ for every $v \in \{r,s\}$. Due to $\mathcal{D}$ extending $\mathcal{I}_\omega^k$, that is, ($\mathcal{O}$1), this implies

$$\overrightarrow{(\bot_A(r), \bot_B(r))}, \overleftarrow{(\top_B(s), \top_A(s))} \in \mathcal{D}.$$

Thus, because $\mathcal{D}$ is consistent, that is, ($\mathcal{O}2$), there exists a node $d \in V(T)$ such that $\overrightarrow{(\top_A(d), \top_B(d))} \in \mathcal{D}$ and $\overleftarrow{(\top_B(t), \top_A(t))} \in \mathcal{D}$ for all $t \in N_T^{\mathsf{in}}(t)$. Due to ($\mathcal{O}3$), this implies $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d)) \geq \omega$. □

In order to obtain the duality we introduce the following definition which intuitively describes prop-DAGs for which the bags of the internal nodes are of small size, while the initial and the terminal separation both are either oriented "inwards" with respect to the given prop-DAG, or they are "trivial".

Definition 3.3.7. Let D be a digraph and $\mathcal{P}$ a partial orientation of $\vec{\mathcal{S}}_k$. A prop-k-DAG (T, σ) is $(\mathsf{prop}, \omega, \mathcal{P})$*-admissible* if it meets the following requirements.

(A.1) For all $d \in V(T) \setminus (\mathsf{sources}_T \cup \mathsf{sinks}_T)$ we have $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d)) < \omega$,

(A.2) $\mathsf{Initial}(T, \sigma) = \overrightarrow{(A, B)} \in \mathcal{P} \cup \mathcal{I}_k^{\omega}$,

(A.3) $\mathsf{Terminal}(T, \sigma) = \overleftarrow{(Z, Y)} \in \mathcal{P} \cup \mathcal{I}_k^{\omega}$,

(A.4) $|\mathsf{sources}_T| = 1$, and

(A.5) for the source r and every sink s we have

$$\sigma((r, d)) = \sigma((r, d')) \text{ for all } d, d' \in N_T^{\mathsf{out}}(r) \text{ and}$$
$$\sigma((d, s)) = \sigma((d', s)) \text{ for all } d, d' \in N_T^{\mathsf{in}}(s)\,.$$

As the separations mapped to the in-edges of every sink are the same, we abuse notation and write $\sigma(s)$ for the separation mapped to all in-edges of s for all sinks s of T. Note that if (T, σ) is closed with a single sink s, then $\mathsf{Terminal}(T, \sigma) = \sigma(s)$. ⊣

We say that a separation $\overrightarrow{(A, B)}$ is oriented *upwards* by an orientation $\mathcal{P}$ if $\overleftarrow{(B, A)} \in \mathcal{P}$.

Definition 3.3.8 (flipping)**.** Let $\mathcal{P}$ be a consistent orientation of $\vec{\mathcal{S}}_k$ of a digraph D and let $\overrightarrow{(A, B)} \in \vec{\mathcal{S}}_k$.

If (A, B) is oriented upwards by $\mathcal{P}$, that is $\overleftarrow{(B, A)} \in \mathcal{P}$, then we say the orientation $\mathcal{P}'$ is *obtained from* $\mathcal{P}$ *by flipping* $\overleftarrow{(B, A)}$ if we replace $\overleftarrow{(B', A')}$ with $\overrightarrow{(A', B')}$ for all $\overleftarrow{(B', A')} \in \mathcal{P}$ with $\overleftarrow{(B', A')} \leq \overleftarrow{(B, A)}$ and refer to $\mathcal{P}'$ as a *flip* of $\mathcal{P}$. ⊣

Lemma 3.3.9. Let D be a digraph and $\mathcal{P}$ a consistent orientation of $\vec{\mathcal{S}}_k$. Then, every orientation $\mathcal{P}'$ obtained from $\mathcal{P}$ by flipping a separation $(X, Y) \notin \mathcal{I}_k^{\omega}$ is again consistent.

Proof. Suppose towards a contradiction that $\mathcal{P}'$ is not consistent. This means there are two separations $\overrightarrow{(A,B)}\,, \overrightarrow{(J,K)}$ such that $\overrightarrow{(A,B)} \leq \overrightarrow{(J,K)}$ and $\overleftarrow{(B,A)}\,, \overrightarrow{(J,K)} \in \mathcal{P}'$. As $\mathcal{P}$ is consistent, at least one of these two separations is oriented differently in $\mathcal{P}$. Note that indeed exactly $\overleftarrow{(K,J)}$ is affected by the flip, as all separations affected by a flip are oriented upwards by $\mathcal{P}$. Then we know that $\overleftarrow{(K,J)} \leq \overleftarrow{(X,Y)}$. This implies that $\overleftarrow{(B,A)} \leq \overleftarrow{(X,Y)}$, thus $\overleftarrow{(B,A)}$ should have been affected as well, a contradiction. ☐

Lemma 3.3.10. Let prop be complete, D be a digraph and $\mathcal{P}$ an orientation of $\vec{\mathcal{S}}_k$. For every closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG (T,σ) with source r and every separation $\overrightarrow{(X,Y)} \in \mathcal{P}$ such that $\sigma(r) \leq \overrightarrow{(X,Y)}$, the down-shift $\left(T^{\uparrow}, \sigma^{\uparrow}\right)$ of (T,σ) onto $\overrightarrow{(X,Y)}$ at r is again a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

Proof. First, as r is the source of T, we have $T^{\uparrow} = T$ and thus the new DAG is again closed and we also obtain (A.4).

Second, since prop-DAG-width is shifting closed, the DAG $\left(T^{\uparrow}, \sigma^{\uparrow}\right)$ is again a prop-k-DAG and has prop-DAG-width at most ω, which yields (A.1).

Next, by definition of down-shifts and due to Observation 3.2.12, we obtain that $\mathsf{Initial}(T^{\uparrow}, \sigma^{\uparrow}) = \overrightarrow{(X,Y)} \in \mathcal{P} \subseteq \mathcal{P} \cup \mathcal{I}_k^{\omega}$, thus (A.2) holds.

For the terminal separation we obtain $\mathsf{Terminal}(T^{\uparrow}, \sigma^{\uparrow}) = \mathsf{Terminal}(T,\sigma) \vee \overrightarrow{(X,Y)}$. Thus $\mathsf{Terminal}(T,\sigma) \leq \mathsf{Terminal}(T^{\uparrow}, \sigma^{\uparrow})$. As $\mathsf{Terminal}(T,\sigma) = \overleftarrow{(Z,Y)} \in \mathcal{P}$ and $\mathcal{P}$ is consistent, we can derive $\mathsf{Terminal}(T^{\uparrow}, \sigma^{\uparrow}) = \overleftarrow{(Z',Y')} \in \mathcal{P}$, and thus (A.3).

Finally, $(A,B) = (A',B')$ implies $(A,B) \vee \overrightarrow{(X,Y)} = (A',B') \vee \overrightarrow{(X,Y)}$ for all separations $(A,B)\,, (A',B')$. Thus, (A.5) for $\left(T^{\uparrow}, \sigma^{\uparrow}\right)$ follows directly from (A.5) for (T,σ) . ☐

Lemma 3.3.11. Let prop be complete, (T,σ) be a prop-k-DAG of prop-width less than ω, $d \in V(T)$, $\overrightarrow{(X,Y)}$ up-linked to d and $\left(T^{\downarrow}(d), \sigma^{\downarrow}(d)\right)$ the up-shift onto $\overleftarrow{(X,Y)}$ at d. Define (T',σ') by $T' := T$ and $\sigma' := \sigma^{\downarrow}(d) \cup \sigma|_{T - T^{\downarrow}(d)}$. Then, (T',σ') , which we call the *extended up-shift of* (T,σ) *onto* $\overleftarrow{(X,Y)}$ *at* d, is again a prop-k-DAG of prop-width less than ω.

Proof. There are two points we have to check. First, that (T',σ') is indeed again a prop-k-DAG and second, that the bags evaluate to less than ω with respect to $\mathsf{eval}_{\mathsf{prop}}$.

For all nodes for which all adjacent edges lie in $T^{\downarrow}(d)$ this is due to (T, σ) being a prop-k-DAG of prop-width less than ω. For the nodes with all adjacent edges in $T - T^{\downarrow}(d)$ this follows from prop-DAG-width being shifting closed.

Note that the only possible remaining case is that one of the out-edges (v, v_1) lies in $T - T^{\downarrow}(d)$ while all other adjacent edges lie in $T^{\downarrow}(d)$.

As $\sigma^{\downarrow}(d)(e) \leq \sigma(e)$ for all in-edges e of v and $\sigma(e) \leq \sigma((v, v_1))$, we obtain $\sigma^{\downarrow}(d)(e) \leq \sigma((v, v_1))$ for all in-edges. Thus, the obtained DAG is indeed again a prop-k-DAG.

By (P4) the bag of v evaluates to less than ω as well. $\square$

In favour of readability we extract the induction base of the proof for the duality theorem as a separate lemma which works with an orientation of all elements of $\vec{\mathcal{S}}_k$.

Lemma 3.3.12. Let prop be complete, D be a digraph that does not contain an ω-big prop-ganglion of order k, and $\mathcal{P}$ a consistent orientation of $\vec{\mathcal{S}}_k$. Then, there exists a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

Proof. For every consistent orientation $\mathcal{O}$ of $\vec{\mathcal{S}}_k$ we define

$$\mathsf{b}(\mathcal{O}) := \left|\{\overrightarrow{(A, B)} \in \vec{\mathcal{S}}_k \mid \overleftarrow{(B, A)} \in \mathcal{O}\}\right|.$$

We prove the statement by induction on $\mathsf{b}(\mathcal{P})$.

For a given prop-k-DAG (T, σ) and a separation $\overrightarrow{(A, B)} \in \vec{\mathcal{S}}_k$ with $\mathsf{Initial} \leq \overrightarrow{(A, B)}$ let

$$\begin{aligned}\mathsf{Att}_{T,\sigma}\left(\overrightarrow{(A, B)}\right) := \Big\{ & d \in V(T) \mid \overrightarrow{(A, B)} \leq \bot^{(T,\sigma)}(d), \\ & \text{and there is no } d' \in V(T) \text{ with } \overrightarrow{(A, B)} \leq \bot^{(T,\sigma)}(d') < \bot^{(T,\sigma)}(d) \Big\}.\end{aligned}$$

Let $\mathcal{P}$ be a consistent orientation of $\vec{\mathcal{S}}_k$. As $\mathcal{P}$ is not an ω-big prop-ganglion, by $(\mathcal{O}3)$, there are three directed separations $\overrightarrow{(A, B)}, \overleftarrow{(K_1, J_1)}, \overleftarrow{(K_2, J_2)} \in \mathcal{P}$ with $\overrightarrow{(A, B)} \leq \overleftarrow{(K_i, J_i)}$ for $i \in \{1, 2\}$ and $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}(\overrightarrow{(A, B)}, \overleftarrow{(K_1, J_1)}, \overleftarrow{(K_2, J_2)})) < \omega$. Choose such a triple with maximal $\overrightarrow{(A, B)}$.

We define a prop-k-DAG (T_0, σ_0) as follows. Let T be the DAG with four nodes t_0, t_1, t_2, t_3 and edges $\{(t_0, t_1), (t_1, t_2), (t_1, t_3)\}$. We define

$$\sigma_0((t_0, t_1)) := \overrightarrow{(A, B)},\ \sigma_0((t_1, t_2)) := \overrightarrow{(J_1, K_1)},\ \text{and } \sigma_0((t_1, t_3)) := \overrightarrow{(J_2, K_2)}.$$

Note that (T_0, σ_0) is $(\mathsf{prop}, \omega, \mathcal{P})$-admissible. Also, define $\mathcal{P}_0 := \mathcal{P}$.

For the following construction we are now given an orientation $\mathcal{P}_n$ with a given $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG (T_n, σ_n) with the following properties:

(IH.a) $\mathsf{Initial}(T_n, \sigma_n) \leq \mathsf{Initial}(T, \sigma)$,

(IH.b) $\sigma_n(\ell)$ is oriented upwards by $\mathcal{P}$ for all leaves ℓ of T_n, and

(IH.c) for every closed $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG (T^c, σ^c) such that $\mathsf{Initial}(T, \sigma) \leq \mathsf{Initial}(T^c, \sigma^c)$ the set $\mathsf{Att}_{T_n, \sigma_n}(\mathsf{Initial}(T^c, \sigma^c))$ is not empty and for every $d \in \mathsf{Att}_{T_n, \sigma_n}(\mathsf{Initial}(T^c, \sigma^c))$ there is a closed $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG $\left(T'^c, \sigma'^c\right)$ with $\mathsf{Initial}(T^c, \sigma^c) \leq \mathsf{Initial}(T'^c, \sigma'^c) = \perp^{(T_n, \sigma_n)}(d)$, $\mathsf{Terminal}(T^c, \sigma^c) \leq \mathsf{Terminal}(T'^c, \sigma'^c)$ and $T'^c \subseteq T_n$.

Claim 1. (T_0, σ_0) and $\mathcal{P}_0$ fulfils (IH.a) to (IH.c).

Proof. The properties (IH.a) and (IH.b) hold by definition and (IH.c) holds as $\overrightarrow{(A, B)}$ in σ is chosen maximally. ■

We now construct a new orientation $\mathcal{P}_{n+1}$ with a new $(\mathsf{prop}, \omega, \mathcal{P}_{n+1})$-admissible prop-$k$-DAG again having these properties, or find a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG on the way.

STEP 1: TWINS As long as possible, choose two sinks t_i, t_j with $\sigma(t_i) = \sigma(t_j)$. We add a new node t_{ij} and edges (t_i, t_{ij}) and (t_j, t_{ij}) which are mapped to the separation $\sigma(t_i)$.

Claim 2. If (T_n, σ_n) is now closed, then (T_n, σ_n) is a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

Proof. By (IH.a), $\mathsf{Initial}(T_n, \sigma_n) \leq \mathsf{Initial}(T_0, \sigma_0)$ holds and thus, the orientation of $\mathsf{Initial}(T_n, \sigma_n)$ in $\mathcal{P}$ agrees with the direction of its edges, that is, it is oriented downwards. As all leaf-separations are oriented upwards before adding extra leaves and because $\mathcal{P}$ is consistent, the new leaf is oriented upwards as well. So (T_n, σ_n) is a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG. ■

STEP 2: TRIVIAL LEAVES We claim that if we can find a leaf $\dot{t}$ with $\sigma(\dot{t}) \in \mathcal{I}_k^\omega$, then we can also find a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

Claim 3. If (T_n, σ_n) contains a leaf $\dot{t}$ with $\sigma(\dot{t}) \in \mathcal{I}_k^\omega$, then there exists a is a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

Proof. We construct such a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG from (T_n, σ_n) as follows. Let $T'_n := T_n^{\downarrow}(\dot{t})$ and let σ'_n be obtained from σ_n by reducing to T'_n. Now, we observe that (T'_n, σ'_n) is not a prop-k-DAG as T'_n possibly contains nodes X that had an out-neighbour in $V(T_n) \setminus V(T'_n)$. For every node $x \in X$ we add the edge $(x, \dot{t})$ to which we assign the separation $\perp^{((T_n, \sigma_n))}(x)$, obtaining a closed $\vec{\mathcal{S}}$-DAG (T''_n, σ''_n). It remains to show that (T''_n, σ''_n) is a $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG. Note that $\beta_{\mathsf{prop}}^{(T''_n, \sigma''_n)}(d) = \beta_{\mathsf{prop}}^{(T_n, \sigma_n)}(d)$ for all $d \in V(T''_n)$ and thus, $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T''_n, \sigma''_n)}(d)) < \omega$ for all $d \in V(T''_n)$. Also the initial separation of (T''_n, σ''_n) remains the same as that of (T_n, σ_n). The terminal separation of (T''_n, σ''_n) is $\sigma''(\dot{t}) \vee \bigvee_{\perp^{((T_n, \sigma_n))}(x)}$ and as such again trivial. ∎

So from now on we can assume that (T_n, σ_n) does not contain any trivial leaves.

STEP 3: FLIPPING Obtain $\mathcal{P}'_{n+1}$ from $\mathcal{P}_n$ by flipping the separations of all leaves. By Lemma 3.3.9, $\mathcal{P}'_{n+1}$ is again a consistent orientation .

STEP 4: CATERPILLAR Note that $\mathsf{b}(\mathcal{P}'_{n+1}) < \mathsf{b}(\mathcal{P})$ and thus we can apply the induction hypotheses.

By induction hypothesis there is at least one closed $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible prop-$k$-DAG. Let $\mathcal{T}_{n+1}$ be the set of all closed $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible prop-$k$-DAGs.

For every separation σ^* that is the initial separation of some prop-k-DAG in $\mathcal{T}_{n+1}$, let $\mathcal{T}_{n+1}(\sigma^*)$ be the subset of $\mathcal{T}_{n+1}$ with σ^* as initial separation.

Claim 4. $\sigma^* \leq \sigma_n(\ell)$ for some $\ell \in \mathsf{sinks}_{T_n}$.

Proof. Suppose $\mathcal{T}_{n+1}$ contains a $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible (T^c, σ^c). By (IH.c), we know that for every $d \in \mathsf{Att}_{T_n, \sigma_n}(\sigma^*)$ there is a closed $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG (T'^c, σ'^c) such that $\mathsf{Initial}(T^c, \sigma^c) \leq \mathsf{Initial}(T'^c, \sigma'^c) = \perp^{(T_n, \sigma_n)}(d)$, $\mathsf{Terminal}(T^c, \sigma^c) \leq \mathsf{Terminal}(T'^c, \sigma'^c)$ and $T'^c \subseteq T_n$. This contradicts that $\mathcal{P}'_{n+1}$ was obtained from $\mathcal{P}_n$ by flipping all leaves, thus all elements (T^c, σ^c) in $\mathcal{T}_{n+1}$ are $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible but not $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible. This implies that there is a separation in (T^c, σ^c) which is comparable to a leaf of T_n and thus $\sigma^* \leq \sigma_n(\ell)$ for some $\ell \in \mathsf{sinks}_{T_n}$. ■

For every $d \in \mathsf{Att}_{T_n,\sigma_n}(\sigma^*)$, let σ' be of minimal order with $\sigma^* \leq \sigma' \leq \perp^{(T_n,\sigma_n)}(d)$. This choice ensures that σ' is down-linked to σ^* and up-linked to d. Next, perform a down-shift on all elements of $\mathcal{T}_{n+1}(\sigma^*)$ onto σ' at their root, obtaining the set of DAGs $\mathcal{T}_{\sigma'}$. By Lemma 3.3.10, the elements of $\mathcal{T}_{\sigma'}$ are again closed $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible prop-$k$-DAGs.

We construct a new DAG by first splitting d into three nodes d', d_1 and d_2 with edges (d', d_1) and (d', d_2) both with separation $\perp^{(T_n,\sigma_n)}(d)$. All in-edges into d are now in-edges to d'. All out-edges from d are now out-edges from d_1.

Now, d_2 is a new leaf and σ' is up-linked to d_2. Next, we perform an extended up-shift of (T_n, σ_n) onto σ' at d_2. By Lemma 3.3.11, this yields again a $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible prop-$k$-DAG.

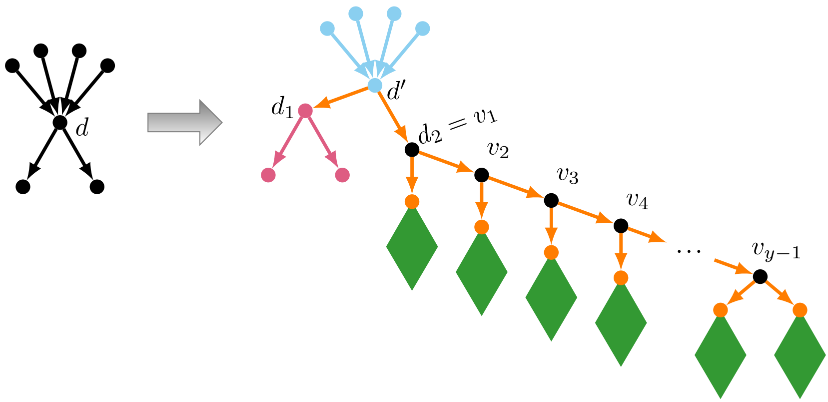

Figure 3.4: The caterpillar construction. The separation on all orange edges is the new separation σ'. The DAGs from $\mathcal{T}_{\sigma'}$ are shown in green. The part of the DAG marked in blue is shifted up, while the part marked in red is shifted down.

Now let $y := |\mathcal{T}_{\sigma'}|$. We build a path on $y-1$ nodes $v_1, \dots, v_{y-1}$, then we decide on a mapping between the nodes of the paths and the elements in $\mathcal{T}_{\sigma'}$, mapping the last node v_{y-1} to two elements and all remaining nodes to exactly one element. We add edges from each path-node to the root of the corresponding element in $\mathcal{T}_{\sigma'}$, two edges for the node v_{y-1}. To these new edges we map the separation σ'.

Finally, we identify the node d_2 with the root v_1 of this new DAG. See Figure 3.4 for an illustration of this construction.

Additionally, we add σ' to a set of separations S and continue until we handled all σ^* that are root-separation for some closed $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible prop-$k$-DAG. This in the end gives us the prop-k-DAG (T_{n+1}, σ_{n+1}).

Finally, we obtain $\mathcal{P}_{n+1}$ from $\mathcal{P}_n$ by flipping all separations in S downwards.

Claim 5. (T_{n+1}, σ_{n+1}) is a $(\mathsf{prop}, \omega, \mathcal{P}_{n+1})$-admissible prop-$k$-DAG.

Proof. We start with checking (A.1). For all nodes $d \in V(T_{n+1}) \cap V(T_n)$ the bag only changes due to shifts. As prop-DAG-width is shifting closed, we know (S3) holds. Additionally, by induction, we have $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T_n,\sigma_n)}(d)) < \omega$. This implies $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T_{n+1},\sigma_{n+1})}(d)) < \omega$.

As the elements of $\mathcal{T}_{\sigma'}$ are prop-k-DAGs, for all $d \in V(T^c)$ for some $(T^c, \sigma^c) \in \mathcal{T}_{\sigma'}$ with $\sigma' \in S$ we also obtain that $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T_{n+1},\sigma_{n+1})}(d)) < \omega$.

So we need to consider the nodes in $\{d', d_1, d_2, v_1, \ldots, v_{y-1}\} \cup \{r \in \mathsf{sources}_{T^c} \mid (T^c, \sigma^c) \in \mathcal{T}_{\sigma'}, \sigma' \in S\}$ during one of the construction steps. Let $d \in V(T_n)$ be the node that is replaced during the considered step. For these nodes we obtain $|\beta_{\mathsf{prop}}^{(T_{n+1},\sigma_{n+1})}(d)| < \omega$ by (Cnt). Therefore (A.1) holds.

By (IH.a) we have that $\mathsf{Initial}(T_{n+1}, \sigma_{n+1}) \leq \mathsf{Initial}(T_n, \sigma_n)$ and thus (A.2) holds. The property (A.3) holds because $(A, \overleftarrow{B)} \in \mathcal{P}_{n+1}$ implies $(A, \overleftarrow{B)} \in \mathcal{P}_0$ for all separations $(A, B) \in \vec{\mathcal{S}}_k$. By construction (T_{n+1}, σ_{n+1}) has only one source, thus (A.4) holds. And finally, (A.5) holds for all old sinks and the source by induction and for the new sinks because the added DAGs are $(\mathsf{prop}, \omega, \mathcal{P}'_{n+1})$-admissible. ∎

Now we prove that the new DAG (T_{n+1}, σ_{n+1}) fulfils the induction criteria.

Claim 6. (T_{n+1}, σ_{n+1}) fulfils (IH.a) to (IH.c).

Proof. **(IH.a):** $\mathsf{Initial}(T_{n+1}, \sigma_{n+1}) \leq \mathsf{Initial}(T_n, \sigma_n)$, as the only shifts are performed on the sub-DAG containing the root of T_n are up-shifts. By induction hypothesis $\mathsf{Initial}(T_n, \sigma_n) \leq \mathsf{Initial}(T, \sigma)$, which implies $\mathsf{Initial}(T_{n+1}, \sigma_{n+1}) \leq \mathsf{Initial}(T, \sigma)$.

(IH.b): By Claim 5, (T_{n+1}, σ_{n+1}) is $(\mathsf{prop}, \omega, \mathcal{P}_{n+1})$-admissible. Thus, all separations at leaves of T_{n+1} are oriented upwards by $\mathcal{P}_{n+1}$. As we only ever flip separations downwards, this implies that all separations at leaves of T_{n+1} are oriented upwards by $\mathcal{P}$ as well.

(IH.c): By Claim 4, for every closed $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG (T^c, σ^c) with $\mathsf{Initial}(T, \sigma) \leq \mathsf{Initial}(T^c, \sigma^c)$ the set $\mathsf{Att}_{T_n, \sigma_n}(\mathsf{Initial}(T^c, \sigma^c))$ is not empty. Thus the caterpillar construction added a down-shift of (T^c, σ^c) with $\mathsf{Initial}(T, \sigma) \leq \mathsf{Initial}(T^c, \sigma^c)$ as a sub-DAG. Therefore, for every $d \in \mathsf{Att}_{T_n, \sigma_n}(\mathsf{Initial}(T^c, \sigma^c))$ there is a closed $(\mathsf{prop}, \omega, \mathcal{P}_n)$-admissible prop-$k$-DAG $\left(T'^c, \sigma'^c\right)$ with $\mathsf{Initial}(T^c, \sigma^c) \leq \mathsf{Initial}(T'^c, \sigma'^c) = \perp^{(T_n, \sigma_n)}(d)$, $\mathsf{Terminal}(T^c, \sigma^c) \leq \mathsf{Terminal}(T'^c, \sigma'^c)$ and $T'^c \subseteq T_n$. ■

As $\mathsf{b}(\mathcal{P}_n)$ decreases with every walkthrough, this iteration ends in a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG eventually. □

For usage in the next proof we define the following notation.

$$\vec{\mathcal{S}}_{\mathcal{O},k} := \vec{\mathcal{S}}_k \setminus \left\{ (A, \overrightarrow{B)}, (B, \overleftarrow{A)} \;\middle|\; \{(A, \overrightarrow{B)}, (B, \overleftarrow{A)}\} \cap \mathcal{O} \neq \emptyset \right\}$$

It denotes the family of directed separations of order less than k not oriented by a partial orientation $\mathcal{O}$ of $\vec{\mathcal{S}}_k$.

Now, we have assembled all the definitions needed to finally prove the duality theorem between prop-DAG-width and ω-big prop-ganglions.

Theorem 3.3.13. Let prop be complete, $\omega \geq k$, and let D be a digraph with at least k vertices. Then, exactly one of the following holds:

(i) D admits a prop-k-DAG of prop-width less than ω, or

(ii) there is an ω-big prop-ganglion of order k in D.

Proof. We prove the slightly different statement that for every consistent partial orientation $\mathcal{P}$ of $\vec{\mathcal{S}}_k$ exactly one of the following holds:

(i') there exists a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG (T, σ) with source r and sink s such that

$$\begin{aligned} \sigma((r, d)) = \sigma((r, d')) &\text{ for all } d, d' \in N_T^{\mathsf{out}}(r) \text{ and} \\ \sigma((d, s)) = \sigma((d', s)) &\text{ for all } d, d' \in N_T^{\mathsf{in}}(s)\,, \end{aligned} \tag{3.2}$$

or

(ii) there is an ω-big prop-ganglion of order k in D.

Note, that (i') directly implies (i) as it holds in particular for $\mathcal{P} = \emptyset$.

Note that every closed $(\mathsf{prop}, \omega, \mathcal{I}_k^\omega)$-admissible prop-$k$-DAG is of prop-width less than ω. Moreover, as the existence of an ω-big prop-ganglion implies that every closed prop-k-DAG (T, σ) of D contains some $d \in V(T)$ with $\mathsf{eval}_{\mathsf{prop}}(\beta^{(T,\sigma)}_{\mathsf{prop}}(d)) \geq \omega$ (see Lemma 3.3.6) both statements (i) and (ii) cannot be true at the same time. Thus it remains to prove that if D does not contain an ω-big prop-ganglion of order k, then for every consistent partial orientation $\mathcal{P}$ of $\vec{\mathcal{S}}_k$ there exists a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG.

First, suppose that $\mathcal{I}_k^\omega \not\subseteq \mathcal{P}$ and $\mathcal{P} \cup \mathcal{I}_k^\omega$ is not a consistent partial orientation of $\vec{\mathcal{S}}_k$. Then, no consistent orientation $\mathcal{O}$ of $\vec{\mathcal{S}}_k$ with $\mathcal{P} \subseteq \mathcal{O}$ can be an ω-big prop-ganglion, since each ω-big prop-ganglion must contain $\mathcal{I}_k^\omega$. We thus have to show that there exists a $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG fulfilling (3.2). Without loss of generality there is a directed separation $\overleftarrow{(B, A)} \in \mathcal{P}$ with $\mathsf{eval}_{\mathsf{prop}}(A) < k \leq \omega$, because the two cases are symmetric. In this case, let T be the directed path consisting of exactly the two nodes t_1 and t_2 with the edge (t_1, t_2), and define $\sigma((t_1, t_2)) \coloneqq \overrightarrow{(A, B)}$. Then, (T, σ) is a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG that fulfils (3.2), thus we obtain (i'). Hence, from now on we may assume that $\mathcal{I}_k^\omega \subseteq \mathcal{P}$.

Now, we prove the claim by induction over $|\vec{\mathcal{S}}_{\mathcal{P},k}|$, that is, the number of non-oriented separations with respect to $\mathcal{P}$. The induction base, that is, $\vec{\mathcal{S}}_{\mathcal{P},k} = \emptyset$, is covered by Lemma 3.3.12.

Now, assume $|\vec{\mathcal{S}}_{\mathcal{P},k}| > 0$, thus there exists some separation $\overrightarrow{(E, F)} \in \vec{\mathcal{S}}_{\mathcal{P},k}$ that is not oriented by $\mathcal{P}$. As there is no ω-big prop-ganglion of order k that extends $\mathcal{P}$, we can choose a maximal directed separation $\overrightarrow{(A, B)} \in \vec{\mathcal{S}}_{\mathcal{P},k}$ with $\overrightarrow{(E, F)} \leq \overrightarrow{(A, B)}$. Additionally, we choose two (not necessarily distinct) minimal directed separations $\overrightarrow{(J_i, K_i)} \in \vec{\mathcal{S}}_{\mathcal{P},k}$ for $i \in \{1, 2\}$ such that $\overrightarrow{(J_1, K_1)}, \overrightarrow{(J_2, K_2)} \leq \overrightarrow{(E, F)}$ and $\overrightarrow{(J_1 \cap J_2, K_1 \cup K_2)} \in \vec{\mathcal{S}}_k$. We claim that $\mathcal{P}_0 \coloneqq \mathcal{P} \cup \{\overleftarrow{(B, A)}\}$ and $\mathcal{P}_i \coloneqq \mathcal{P} \cup \{\overrightarrow{(J_i, K_i)}\}$ for $i \in \{1, 2\}$ are all consistent partial orientations of $\vec{\mathcal{S}}_k$. Indeed, by the maximality of $\overrightarrow{(A, B)}$, every directed separation $\overrightarrow{(U, V)}$ with $\overrightarrow{(A, B)} < \overrightarrow{(U, V)}$ has to be oriented by $\mathcal{P}$. So, because $\mathcal{P}$ is consistent and $\overrightarrow{(A, B)}$ is not oriented by $\mathcal{P}$, we have $\overleftarrow{(V, U)} \in \mathcal{P}$. Similar arguments can be made for $\overrightarrow{(J_i, K_i)}$, and thus, $\mathcal{P}_i$ is a consistent partial orientation of $\vec{\mathcal{S}}_k$ for $i \in \{1, 2\}$. Now, for every $i \in \{0, 1, 2\}$ holds $|\vec{\mathcal{S}}_{\mathcal{P}_i,k}| < |\vec{\mathcal{S}}_{\mathcal{P},k}|$, and thus, we can apply the induction hypothesis to $\mathcal{P}_i$.

Every ω-big prop-ganglion of order k extending one of the $\mathcal{P}_i$ also extends $\mathcal{P}$ and thus we may assume that no such ω-big prop-ganglion exists. Hence, by induction hypothesis, for every $i \in \{0, 1, 2\}$ there exists a $(\mathsf{prop}, \omega, \mathcal{P}_i)$-admissible and prop-closed prop-k-DAG (T_i, σ_i) which fulfils (3.2). If (T_i, σ_i) is $(\mathsf{prop}, \omega, \mathcal{P})$-admissible for an $i \in \{0, 1, 2\}$, then we obtain (i). Thus, we may assume that (T_i, σ_i) is not $(\mathsf{prop}, \omega, \mathcal{P})$-admissible for all $i \in \{0, 1, 2\}$. Hence, $\overrightarrow{(A, B)} = \mathsf{Terminal}(T_0, \sigma_0)$ and $\overrightarrow{(J_i, K_i)} = \mathsf{Initial}(T_i, \sigma_i)$ for $i \in \{1, 2\}$.

Next, we combine (T_1, σ_1) and (T_2, σ_2) into a single closed prop-k-DAG (T', σ'). To this end, we consider two different cases.

If $\overrightarrow{(J_1, K_1)} = \overrightarrow{(J_2, K_2)}$, then we may assume $(T_1, \sigma_1) = (T_2, \sigma_2)$ and can choose $(T', \sigma') := (T_1, \sigma_1)$. Moreover, we choose $\mathcal{P}' := \mathcal{P}_1 = \mathcal{P}_2$ and the separation $\overrightarrow{(J', K')} := \mathsf{Initial}(T', \sigma') \in \mathcal{P}'$.

If $\overrightarrow{(J_1, K_1)} \neq \overrightarrow{(J_2, K_2)}$, then let r_i be the unique source of T_i and let (T'', σ'') be the prop-up-closing of the union of (T_1, σ_1) and (T_2, σ_2) at r_1 and r_2. The obtained (T'', σ'') is a prop-k-DAG, as prop-DAG-width is complete and thus closed. Then, let (T', σ') be a prop-down-closure of (T'', σ''), which exists by the completeness of prop-DAG-width and is again a prop-k-DAG. As $\overrightarrow{(J_i, K_i)}$ are distinct and minimal, they are incomparable and so $\mathcal{P}' := \mathcal{P}_1 \cup \mathcal{P}_2$ is a consistent partial orientation of $\vec{\mathcal{S}}_k$. In particular, we have $\overrightarrow{(J', K')} := \mathsf{Initial}(T', \sigma') \in \mathcal{P}'$.

In both cases (T', σ') is a closed $(\mathsf{prop}, \omega, \mathcal{P}')$-admissible prop-$k$-DAG.

By our construction so far, we now have two closed prop-k-DAGs: (T_0, σ_0) with the terminal prop-separation $\overrightarrow{(A, B)}$, and (T', σ') with the initial prop-separation $\overrightarrow{(J', K')}$. Additionally, both still fulfil (3.2). What is left to do is to combine these two prop-k-DAGs into a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG (T, σ) fulfilling (3.2).

As $\overrightarrow{(J_i, K_i)} \leq \overrightarrow{(A, B)}$ for both $i \in \{1, 2\}$, we know that $\overrightarrow{(J', K')} \leq \overrightarrow{(A, B)}$. We choose a directed separation $\overrightarrow{(X, Y)}$ such that

$$\overrightarrow{(J', K')} \leq \overrightarrow{(X, Y)} \leq \overrightarrow{(A, B)} \text{ and } |(X, Y)| = \lambda\Big(\overrightarrow{(J', K')}, \overrightarrow{(A, B)}\Big).$$

Then, $\overrightarrow{(X, Y)}$ is down-linked to $\overrightarrow{(J', K')}$ and up-linked to $\overrightarrow{(A, B)}$.

Let s be the unique sink of (T_0, σ_0) and r the unique source of (T', σ'). Then, let $\left(T^{\downarrow}, \sigma^{\downarrow}\right)$ be the up-shift of (T_0, σ_0) onto $\overrightarrow{(X, Y)}$ at s, which is $\overrightarrow{(X, Y)}$-s-admissible. Furthermore, let $\left(T^{\uparrow}, \sigma^{\uparrow}\right)$ be the down-shift of (T', σ') onto $\overrightarrow{(X, Y)}$ at r, which is $\overrightarrow{(X, Y)}$-r-admissible.

Since prop-DAG-width is complete and therefore weakly shifting closed, we obtain that $\left(T^{\downarrow}, \sigma^{\downarrow}\right)$ is a closed $(\mathsf{prop}, \omega, \mathcal{P}_0)$-admissible prop-$k$-DAG, and that $\left(T^{\uparrow}, \sigma^{\uparrow}\right)$

is a closed $(\mathsf{prop}, \omega, \mathcal{P}')$-admissible prop-$k$-DAG. Moreover, by observations 3.2.12 and 3.2.14, we have

$$\mathsf{Terminal}\big(T^{\downarrow}, \sigma^{\downarrow}\big) = (X, \overrightarrow{Y)} = \mathsf{Initial}\big(T^{\uparrow}, \sigma^{\uparrow}\big)\,.$$

Let $s^{\downarrow}$ be the unique sink of $\big(T^{\downarrow}, \sigma^{\downarrow}\big)$ and $r^{\uparrow}$ the unique source of $\big(T^{\uparrow}, \sigma^{\uparrow}\big)$. We create a new $\vec{\mathcal{S}}$-DAG (T, σ) by identifying $s^{\downarrow}$ and $r^{\uparrow}$ into the node $t^{\updownarrow}$ and by defining for every $e \in V(T)$

$$\sigma(e) := \begin{cases} \sigma^{\downarrow}(e), & e \in E(T^{\downarrow}) \\ \sigma^{\uparrow}(e), & e \in E(T^{\uparrow}). \end{cases}$$

Claim 1. (T, σ) is a closed $(\mathsf{prop}, \omega, \mathcal{P})$-admissible prop-$k$-DAG fulfilling (3.2).

Proof. By construction, we know that (T, σ) is a closed prop-k-DAG. So what we need to show is that it also is $(\mathsf{prop}, \omega, \mathcal{P})$-admissible.

For every $d \in V(T) \setminus \{t^{\updownarrow}\}$, we have $\beta_{\mathsf{prop}}^{(T,\sigma)}(d) = \beta_{\mathsf{prop}}^{(T^{\downarrow},\sigma^{\downarrow})}(d)$ if $d \in V(T^{\downarrow})$, and $\beta_{\mathsf{prop}}^{(T,\sigma)}(d) = \beta_{\mathsf{prop}}^{(T^{\uparrow},\sigma^{\uparrow})}(d)$ if $d \in V(T^{\uparrow})$. Thus, for all $d \in V(T) \setminus \{t^{\updownarrow}\}$, we have $\mathsf{eval}_{\mathsf{prop}}(\beta_{\mathsf{prop}}^{(T,\sigma)}(d)) < \omega$, because $\big(T^{\downarrow}, \sigma^{\downarrow}\big)$ is a closed $(\mathsf{prop}, \omega, \mathcal{P}_0)$-admissible prop-$k$-DAG and $\big(T^{\uparrow}, \sigma^{\uparrow}\big)$ is a closed $(\mathsf{prop}, \omega, \mathcal{P}')$-admissible prop-$k$-DAG. For $t^{\updownarrow}$, because prop-DAG-width is complete and therefore closed, (Cnt) and (3.2) imply

$$\mathsf{eval}_{\mathsf{prop}}\Big(\beta_{\mathsf{prop}}^{(T,\sigma)}\big(t^{\updownarrow}\big)\Big) = \mathsf{eval}_{\mathsf{prop}}\Big(\beta_{\mathsf{prop}}\big((X, \overrightarrow{Y)}\big)\Big) < k \leq \omega.$$

Thus, we obtain (A.1).

Let $(U, \overrightarrow{V)} := \mathsf{Initial}(T^{\downarrow}, \sigma^{\downarrow})$. By Lemma 3.2.4, $\mathsf{Initial}(T, \sigma) = (U \cap X, V \cup \overrightarrow{Y)} \leq (U, \overrightarrow{V)}$. As $(U, \overrightarrow{V)} \in \mathcal{P}$, we can use the consistency of $\mathcal{P}$ to obtain $\mathsf{Initial}(T, \sigma) \in \mathcal{P}$ and thus (A.2). Also (3.2) directly implies (A.5) and (A.4) holds by construction.

Similar arguments show that $\mathsf{Terminal}(T, \sigma) \in \mathcal{P}$ and thus (A.3). ■

Thus, we obtain (i'), which finishes the proof. □

3.4 The ν-property

We now introduce a width measure that allows us to apply the general duality theorem Theorem 3.3.13 we prove above: the ν-DAG-width. The width measure ν-DAG-width uses $\vec{\mathcal{S}}$-DAGs and provides all the definitions required in Section 3.3, see Subsection 3.4.1. In Subsection 3.4.2, we introduce the corresponding concept of a ganglion. The duality between these two concepts, that is, the applicability of Theorem 3.3.13, is established in Subsection 3.4.3. In Subsection 3.4.5, we consider the relation of ν-DAG-width to the cops and robber reachability game.

3.4.1 The ν-DAG-width

We start out with the definition of ν-DAG-width. In Section 3.3 we describe a set of tools which suffices to obtain the duality in Theorem 3.3.13. We make sure that ν-DAG-width does provide all of them.

From Lemma 3.2.8 we can derive the following corollary providing us with an additional separation we can associate with a node of a $\vec{\mathcal{S}}$-DAG.

Corollary 3.4.1. Let D be a digraph and (T, σ) be an $\vec{\mathcal{S}}$-DAG for D. For every node $t \in V(T)$ holds $\overrightarrow{(\bot_A(t), \top_B(t))} \in \vec{\mathcal{S}}$.

Proof. Let $t \in V(T)$. If t is a sink, then $\bot_B(t) = V(D)$ and the statement holds. If t is not a sink, then, by Lemma 3.2.8, we have $\bot_B(t) \subseteq \top_B(t)$. As, $\bot(t) \in \vec{\mathcal{S}}$, this implies the statement. □

$\beta_\nu^{(T,\sigma)}$ We use the separator of this new separation in order to define tangible instantiations of the requirements (P1) and (P2) in Section 3.3 to obtain the ν-property. For the bag function $\beta_\nu^{(T,\sigma)}$, based on $(\bot_A(t), \top_B(t))$, we define $\beta_\nu^{(T,\sigma)} : V(T) \to 2^{V(D)}$ as

$$\beta_\nu^{(T,\sigma)}(t) := \bot_A(t) \cap \top_B(t)\,.$$

eval_ν For the evaluation function eval_ν we use the size of the bag and simply write it as $|\cdot|$. As defined in Section 3.3, this allows us to obtain the *ν-width* of an $\vec{\mathcal{S}}$-DAG (T, σ) with $\max_{t \in V(T)} |\beta_\nu^{(T,\sigma)}(t)|$, and the *$\nu$-DAG-width* of a digraph D, denoted by $\nu(D)$, being the minimum ν-width of an $\vec{\mathcal{S}}$-DAG for D.

We observe that for every $\vec{\mathcal{S}}$-DAG (T, σ) the separators of the canonical $\top$- and $\bot$-separation of a node are contained in the bag of that node given by $\beta_\nu^{(T,\sigma)}$.

Observation 3.4.2. Let D be a digraph and (T, σ) be an $\vec{\mathcal{S}}$-DAG for D. Then, for every node $t \in V(T)$ holds $\top_A(t) \cap \top_B(t) \subseteq \beta_\nu^{(T,\sigma)}(t)$ and $\bot_A(t) \cap \bot_B(t) \subseteq \beta_\nu^{(T,\sigma)}(t)$.

We do not restrict the $\vec{\mathcal{S}}$-DAGs when considering the ν-DAG-width, thus every $\vec{\mathcal{S}}_k$-DAG is a *ν-DAG*. We say that an $\vec{\mathcal{S}}$-DAG (T, σ) is a *ν-k-DAG* if we have

- $\overrightarrow{(\top_A(d), \top_B(d))}, \overrightarrow{(\bot_A(d), \bot_B(d))} \in \vec{\mathcal{S}}_k$ for all $d \in V(T)$ and
- $\mathsf{Initial}(T, \sigma), \mathsf{Terminal}(T, \sigma) \in \vec{\mathcal{S}}_k$.

For every two nodes in the decomposition tree of a tree decompositions, all bags of nodes on the unique path between them contain the intersection of the bags of the two nodes, see (Tw3). We can prove a similar property for the bag function β_ν.

Lemma 3.4.3. Let D be a digraph, (T, σ) be an $\vec{\mathcal{S}}$-DAG for D, and $t_1, t_2, t_3 \in V(T)$ such that $t_1 \preccurlyeq_T t_2$ and $t_2 \preccurlyeq_T t_3$. Then, $\beta_\nu^{(T,\sigma)}(t_1) \cap \beta_\nu^{(T,\sigma)}(t_3) \subseteq \beta_\nu^{(T,\sigma)}(t_2)$.

Proof. Let $(t_1', t_2), (t_2, t_3') \in E(T)$ such that $t_1 \preccurlyeq_T t_1'$ and $t_3' \preccurlyeq_T t_3$. Now, let $x \in \beta_\nu^{(T,\sigma)}(t_1) \cap \beta_\nu^{(T,\sigma)}(t_3)$, then $x \in \top_B(t_3) \cap \bot_A(t_1)$. With Lemma 3.2.8 we have $\top_B(t_3) \subseteq \top_B(t_2)$ and $\bot_A(t_1) \subseteq \bot_A(t_2)$ and so $x \in \top_A(t_2) \cap \bot_A(t_2) = \beta_\nu(t_2)$. □

We also establish the two requirements (P3) and (P4) for ν-DAG-width.

Lemma 3.4.4. Let (T, σ) be an $\vec{\mathcal{S}}_k$-DAG of width at most ω. Let $v \in V(T)$ be a node in an $\vec{\mathcal{S}}_k$-DAG and $\overrightarrow{(J_1, K_1)}$ and $\overrightarrow{(J_2, K_2)}$ the two separations on the out-edges of v. If we add an in-edge of v, then $|\beta_\nu(v)| < \omega$.

Proof. Let $\overrightarrow{(J, K)}$ be the label of the new in-edge. Then, the bag for v becomes

$$\bot^{((T,\sigma))}(v) \cap J \cap \bot^{((T,\sigma))}(v) \subseteq \bot^{((T,\sigma))}(v) \cap \bot^{((T,\sigma))}(v)$$

which is smaller than ω by assumption. □

Lemma 3.4.5. Let (T, σ) be an $\vec{\mathcal{S}}_k$-DAG of ν-width at most ω. Let $v \in V(T)$ be a node in an $\vec{\mathcal{S}}_k$-DAG and $\overrightarrow{(J_1, K_1)}$ and $\overrightarrow{(J_2, K_2)}$ the two separations on the out-edges of v. If we replace all separations on the in-edges of v and $\overrightarrow{(J_2, K_2)}$ by their uncrossing with a separation $\overrightarrow{(X, Y)} \in \vec{\mathcal{S}}_k$ that is up-linked to a separation $\overrightarrow{(J_2', K_2')}$ with $\overrightarrow{(J_2, K_2)} \leq \overrightarrow{(J_2', K_2')}$, then $|\beta_\nu(v)| < \omega$.

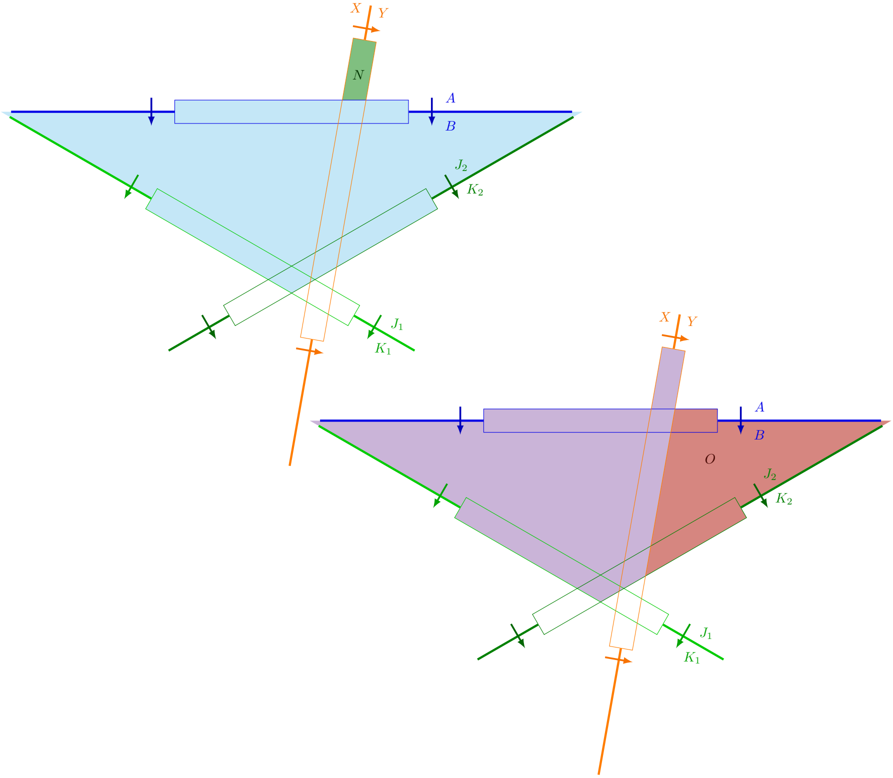

Figure 3.5: In the upper part of the figure we see the old bag and the set N that we add to it. In the lower part we see the new bag and the set O, marked in red, that is deleted.

Proof. We consider the set N of vertices that lie in the new but not in the old bag:

$$N := (X \cap Y) \cap \left(\top_A^{((T,\sigma))}(v) \setminus \top_B^{((T,\sigma))}(v)\right).$$

Note that $N \subseteq X \cap Y$ and let $Z := (X \cap Y) \setminus N$. Additionally, we consider the set O of vertices that are in the old bag, but do not lie in the new one:

$$O := \left(\top_B^{((T,\sigma))}(v) \cap J_1 \cap J_2\right) \cap (Y \setminus X).$$

See Figure 3.5 for an illustration.

Suppose $|N| > |O|$. Then, $O \cup Z$ is a directed separator of size less than k and in particular $|O \cup Z| < |X \cap Y|$. This implies that $|\overrightarrow{(X,Y)} \vee \overrightarrow{(J_2,K_2)}| < |\overrightarrow{(X,Y)}|$.

But as $\overrightarrow{(X,Y)} \leq |\overrightarrow{(X,Y)} \vee \overrightarrow{(J_2,K_2)}| \leq \overrightarrow{(J_2',K_2')}$, this contradicts $\overrightarrow{(X,Y)}$ being downlinked to $\overrightarrow{(J_2',K_2')}$. □

3.4.2 ν-ganglions

We adapt the definition of ω-diblockages to our setting by involving three separations, obtaining ω-big ν-ganglions. The term “big” emphasises that we use the big sides of the separations. We explicitly state the definition for an ω-big ν-ganglion, which is obtained from the abstract one given in Definition 3.3.5 by replacing $\beta_{\mathsf{prop}}^{(T,\sigma)}$ and $\mathsf{eval}_{\mathsf{prop}}$ with $\beta_{\nu}^{(T,\sigma)}$ and eval_{ν}.

Definition 3.4.6. Let D be a digraph and $k, \omega \in \mathbb{N}$ with $\omega \geq k$. An *ω-big ν-ganglion* of order k of D is an orientation $\mathcal{D}$ of $\vec{\mathcal{S}}_k$ such that

- **($\mathcal{O}$1-ν)** $\mathcal{D}$ extends $\mathcal{I}_k^{\omega}$,
- **($\mathcal{O}$2-ν)** $\mathcal{D}$ is consistent, and
- **($\mathcal{O}$3-ν)** if $\overrightarrow{(A,B)}, \overleftarrow{(K_1,J_1)}, \overleftarrow{(K_2,J_2)} \in \mathcal{D}$ with $\overrightarrow{(A,B)} \leq \overleftarrow{(K_i,J_i)}$ for both $i \in \{1,2\}$, then
$$\mathsf{eval}_{\nu}(B \cap J_1 \cap J_2) = |B \cap J_1 \cap J_2| \geq \omega.$$

If $\mathcal{D}$ is a k-big ν-ganglion of order k, we drop the redundant k at the beginning and call $\mathcal{D}$ a big ν-ganglion of order k. ⊣

The key property of a tangle in an undirected graph is that the union of three small sides never covers the whole graph. The definition of directed tangles introduced by Giannopoulou et. al [GKK+20] as an orientation of directed separations makes use of the small sides as well. In a similar fashion, we also introduce a small version of ν-ganglions in order to make this structure more comparable to existing ones. In contrast to directed tangles, we only demand for certain sets of separations that their small sides do not span the whole graph.

Definition 3.4.7. Let D be a digraph and $k \in \mathbb{N}$. A *small ν-ganglion* of order k of D is an orientation $\mathcal{D}$ of $\vec{\mathcal{S}}_k$ such that if $\overrightarrow{(A,B)} \in \mathcal{D}$ and $\overleftarrow{(K_1,J_1)}, \overleftarrow{(K_2,J_2)} \in \mathcal{D}$ with $\overrightarrow{(A,B)} \leq \overleftarrow{(K_i,J_i)}$ for both $i \in \{1,2\}$ then
$$A \cup K_1 \cup K_2 \neq V(D)\,.$$
⊣

Lemma 3.4.8. Let D be a digraph and $k \in \mathbb{N}$. If $\mathcal{D}$ is a small ν-ganglion of order k in D, then it is consistent and extends $\mathcal{I}_{\omega}^{k}$.

Proof. Suppose $\mathcal{D}$ is not consistent. Then, there are two separations $(A, B), (K, J) \in \vec{\mathcal{S}}_k$ with $\overrightarrow{(A, B)}, \overrightarrow{(K, J)} \in \mathcal{D}$ and $\overrightarrow{(J, K)} \leq \overrightarrow{(A, B)}$. This means that $J \subseteq A$ and $B \subseteq K$, implying $A \cup K = V(D)$, but A and K are the two small sides of our separations, a contradiction.

Now suppose $\mathcal{D}$ does not extend $\mathcal{I}_\omega^k$. Then, one of the following two cases holds. The first case is that there is a separation $\overrightarrow{(A, B)} \in \mathcal{I}_k^\omega$ with $\overleftarrow{(B, A)} \in \mathcal{D}$. Being an orientation of $\vec{\mathcal{S}}_k$, $\mathcal{D}$ also has to orient $\overrightarrow{(A, B \cup A)}$. Suppose that $\overrightarrow{(A, B \cup A)} \in \mathcal{D}$, then we have $\overrightarrow{(A, B \cup A)} \leq \overrightarrow{(A, B)}$ as well as $\overleftarrow{(B, A)} \in \mathcal{D}$, which contradicts that $\mathcal{D}$ is a ν-ganglion, as $A \cup B = V(D)$. So, we have $\overleftarrow{(B \cup A, A)} \in \mathcal{D}$, which directly contradicts the definition of a small ν-ganglion. The second case is that there is a separation $\overleftarrow{(A, B)} \in \mathcal{I}_k^\omega$ with $\overrightarrow{(B, A)} \in \mathcal{D}$. Because $\mathcal{D}$ is consistent, this implies $\overrightarrow{(B \cup A, A)} \in \mathcal{D}$, which again contradicts the definition of a small ν-ganglion. $\square$

The two notions of small and big ν-ganglions are closely related. The existence of a small ν-ganglion of high order implies the existence of a big ν-ganglion of comparable order and vice versa.

Lemma 3.4.9. Let D be a digraph and $\mathcal{D}$ a small ν-ganglion of order k in D. Then, $\mathcal{D}|_{\lfloor k/2 \rfloor}$ is a big ν-ganglion of order $\left\lfloor \frac{k}{2} \right\rfloor$ in D.

Proof. Towards a contradiction, suppose $\mathcal{D}|_{\lfloor k/2 \rfloor}$ is not a big ν-ganglion of order $\left\lfloor \frac{k}{2} \right\rfloor$. So one of the following three cases holds: $\mathcal{D}|_{\lfloor k/2 \rfloor}$ does not extend $\mathcal{I}_\omega^{\lfloor \frac{k}{2} \rfloor}$, or it is not consistent, or there exist directed separations $\overrightarrow{(A, B)}, \overleftarrow{(K_1, J_1)}, \overleftarrow{(K_2, J_2)} \in \mathcal{D}$ of order less than $\left\lfloor \frac{k}{2} \right\rfloor$ with $\overrightarrow{(A, B)} \leq \overleftarrow{(K_i, J_i)}$ for $i \in \{1, 2\}$ such that $|B \cap J_1 \cap J_2| < \left\lfloor \frac{k}{2} \right\rfloor$. Since $\mathcal{D}|_{\lfloor k/2 \rfloor}$ is a small ν-ganglion, by Lemma 3.4.8, the first two cases cannot apply and thus we only have to consider the third case.

Consider the uncrossing $\overleftarrow{(K_1, J_1)} \wedge \overleftarrow{(K_2, J_2)} = \overleftarrow{(K_1 \cup K_2, J_1 \cap J_2)}$. We know that $\overrightarrow{(A, B)} \leq \overleftarrow{(K_1, J_1)} \wedge \overleftarrow{(K_2, J_2)}$ and therefore $A \subseteq J_1 \cap J_2$ and $K_1 \cup K_2 \subseteq B$. Thus, $(K_1 \cup K_2) \cap (J_1 \cap J_2) \subseteq B \cap J_1 \cap J_2$ and so $\overleftarrow{(K_1, J_1)} \wedge \overleftarrow{(K_2, J_2)} \in \vec{\mathcal{S}}_{\lfloor k/2 \rfloor}$.

From $\overleftarrow{(K_1, J_1)} \wedge \overleftarrow{(K_2, J_2)}$ we obtain the following directed separation by adding further vertices to the separator. Define

$$\overleftarrow{(K, J)} := \overleftarrow{(K_1 \cup K_2 \cup (B \cap J_1 \cap J_2), (B \cap J_1 \cap J_2) \cup (J_1 \cap J_2))}.$$

The separator of $\overleftarrow{(K, J)}$ is exactly the set $(B \cap J_1 \cap J_2)$, thus $\overleftarrow{(K, J)} \in \vec{\mathcal{S}}_{\lfloor k/2 \rfloor}$. Note that $\overrightarrow{(A, B)} \leq \overleftarrow{(K, J)}$ holds as well as $\overrightarrow{(J, K)} \leq \overleftarrow{(K_i, J_i)}$ for both $i \in \{1, 2\}$.

As $\mathcal{D}$ is an orientation of $\vec{\mathcal{S}}_k$, we have $\overleftarrow{(K, J)} \in \mathcal{D}$ or $\overrightarrow{(J, K)} \in \mathcal{D}$. Suppose $\overleftarrow{(K, J)} \in \mathcal{D}$, then $A \cup K = V(D)$, because $K = K_1 \cup K_2 \cup (B \cap J_1 \cap J_2) = (K_1 \cup K_2 \cup B) \cap V(D) = K_1 \cup K_2 \cup B \supseteq B$. Thus, two small sides are spanning, which contradicts $\mathcal{D}$ being a small ν-ganglion. So we may assume $\overrightarrow{(J, K)} \in \mathcal{D}$. With $J \cup K_1 \cup K_2 = V(D)$ we have three spanning small sides, which again is a contradiction to $\mathcal{D}$ being a small ν-ganglion.

Therefore, $\mathcal{D}|_{\lfloor k/2 \rfloor}$ is a big ν-ganglion. $\square$

Lemma 3.4.10. Let D be a digraph and $\mathcal{D}$ a big ν-ganglion of order k in D. Then, $\mathcal{D}|_{\lfloor \frac{k}{3} \rfloor}$ is a small ν-ganglion of order $\lfloor k/3 \rfloor$ in D.

Proof. Let $\mathcal{D}$ be a big ν-ganglion of order k in D. Suppose towards a contradiction that $\mathcal{D}|_{\lfloor k/3 \rfloor}$ is not a small ν-ganglion of order $\lfloor \frac{k}{3} \rfloor$ in D. Then, there exist three directed separations $\overrightarrow{(A, B)}, \overleftarrow{(K_1, J_1)}, \overleftarrow{(K_2, J_2)} \in \mathcal{D}|_{\lfloor k/3 \rfloor}$ with $\overrightarrow{(A, B)} \leq \overleftarrow{(K_i, J_i)}$ for $i \in \{1, 2\}$ and $|A \cup K_1 \cup K_2| = V(D)$.

The intersection of the three big sides has to therefore lie within $|A \cup K_1 \cup K_2|$ and thus, completely lies in the three separators.

$$B \cap J_1 \cap J_2 \subseteq (A \cap B) \cup (J_1 \cap K_1) \cup (J_1 \cap K_2) .$$

This implies $|B \cap J_1 \cap J_2| < 3 \cdot \frac{k}{3} = k$, because all three separators have size less than $\lfloor \frac{k}{3} \rfloor$. Thus, we obtain a contradiction to $\mathcal{D}$ being a big ν-ganglion and therefore $\mathcal{D}|_{\lfloor k/3 \rfloor}$ is a small ν-ganglion. $\square$

3.4.3 Duality for ν-DAG-width

Our main goal in this section is to apply Theorem 3.3.13 in order to obtain the following duality theorem for ν-DAG-width.

Theorem 3.4.11. Let D be a digraph and $\omega, k \in \mathbb{N}$ with $\omega \geq k$. Then, exactly one of the following holds:

(i) D has a ν-k-DAG of ν-width less than ω, or

(ii) there is an ω-big ν-ganglion of order k in D.

By Definition 3.4.6 and Lemma 3.3.6, we know that ω-big ν-ganglions indeed implement Definition 3.3.5, so we have the desired ganglion definition we need in order to apply Theorem 3.3.13. We already have established that (P1) to (P4) are fulfilled. So in this subsection we prove (P5) and (P6) and that ν-DAG-width is complete, allowing us to apply Theorem 3.3.13, which then implies Theorem 3.4.11. The first step is to show that ν-DAG-width is shifting closed.

Lemma 3.4.12. The digraph width measure ν-DAG-width is shifting closed.

Proof. We prove the three requirements individually.

Because every $\vec{\mathcal{S}}$-DAG is a ν-DAG, by Lemma 3.2.3 every $\overrightarrow{(X,Y)}$-d-admissible shift of an $\vec{\mathcal{S}}$-DAG is again an $\vec{\mathcal{S}}$-DAG, so (S1) holds.

Let (T',σ') be a down-shift of a given ν-k-DAG (T,σ) onto the separation $\overrightarrow{(X,Y)}$ at a node $d \in V(T)$ to which $\overrightarrow{(X,Y)}$ is down-linked. By Lemma 3.2.8, we know

$$\overrightarrow{\left(\perp_A^{(T',\sigma')}(t)\,,\perp_B^{(T',\sigma')}(t)\right)} \leq \overrightarrow{\left(\top_A^{(T',\sigma')}(t')\,,\top_B^{(T',\sigma')}(t')\right)}$$

for all $(t,t') \in E(T')$ and thus lemmata 3.2.4 and 3.2.9 yield (S2) for down-shifts. In case that (T',σ') is an up-shift of a given ν-k-DAG (T,σ) onto the separation $\overrightarrow{(X,Y)}$ at a node $d \in V(T)$ to which $\overrightarrow{(X,Y)}$ is up-linked, we can again use Lemma 3.2.8 to obtain that $\overrightarrow{\left(\perp_A^{(T',\sigma')}(t),\perp_B^{(T',\sigma')}(t)\right)} \leq \overrightarrow{\left(\top_A^{(T',\sigma')}(t'),\top_B^{(T',\sigma')}(t')\right)}$ for all $(t,t') \in E(T')$. Then, (S2) for up-shifts follows from lemmata 3.2.4 and 3.2.10.

Let $\left(T^{\uparrow}(d),\sigma^{\uparrow}(d)\right)$ be the down-shift of some ν-k-DAG (T,σ) onto the separation $\overrightarrow{(X,Y)}$ at a node $d \in V(T)$ to which $\overrightarrow{(X,Y)}$ is down-linked. Let $t \in V(T^{\uparrow}(d))$. By Lemma 3.2.15, we know

$$\begin{aligned}\left|\beta_\nu^{(T^{\uparrow}(d),\sigma^{\uparrow}(d))}(t)\right| &= \left|\overrightarrow{\left(\perp_A^{(T^{\uparrow}(d),\sigma^{\uparrow}(d))}(t)\,,\top_B^{(T^{\uparrow}(d),\sigma^{\uparrow}(d))}(t)\right)}\right| \\ &= \left|\left(\perp_A^{(T,\sigma)}(t)\cup X\right)\cap\left(\top_B^{(T,\sigma)}(t)\cap Y\right)\right| \\ &= \left|\overrightarrow{\left(\perp_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\vee\overrightarrow{(X,Y)}\right|.\end{aligned}$$

Because $\overrightarrow{(X,Y)}$ is down-linked to $\mathsf{Initial}(T^{\uparrow}(d),\sigma)$ and

$$\mathsf{Initial}\left(T^{\uparrow}(d)\,,\sigma\right) \leq \overrightarrow{\left(\perp_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)},$$

Lemma 3.2.9 implies

$$\left|\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\vee\overrightarrow{(X,Y)}\right|\le\left|\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\right|.$$

Therefore, we obtain

$$\left|\beta_\nu^{\left(T^\uparrow(d),\sigma^\uparrow(d)\right)}(t)\right|\le\left|\beta_\nu^{(T,\sigma)}(t)\right|,$$

and thus (S3) holds for down-shifts.

Finally, with similar arguments, we show that (S3) holds for up-shifts as well. Let $\left(T^\downarrow(d),\sigma^\downarrow(d)\right)$ be the up-shift of some ν-k-DAG (T,σ) onto the separation $\overrightarrow{(X,Y)}$ at a node $d\in V(T)$ to which it is up-linked. Let $t\in V(T^\downarrow(d))$. We know that

$$\left|\beta_\nu^{\left(T^\downarrow(d),\sigma^\downarrow(d)\right)}(t)\right|=\left|\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\wedge\overrightarrow{(X,Y)}\right|.$$

Due to $\overrightarrow{(X,Y)}$ being up-linked to $\mathsf{Terminal}(T^\downarrow(d),\sigma)$ and

$$\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\le\mathsf{Terminal}\left(T^\downarrow(d)\,,\sigma\right),$$

Lemma 3.2.10 implies

$$\left|\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}\wedge\overrightarrow{(X,Y)}\right|\le\overrightarrow{\left(\bot_A^{(T,\sigma)}(t)\,,\top_B^{(T,\sigma)}(t)\right)}.$$

Thus, we obtain

$$\left|\beta_\nu^{(T',\sigma')}(t)\right|\le\left|\beta_\nu^{(T,\sigma)}(t)\right|,$$

and therefore (S3) holds for up-shifts. $\square$

Next, we instantiate ν-up-closings and ν-down-closings, providing the two missing requirements (P5) and (P6) for ν-DAG-width. Let (T,σ) be a ν-DAG with two sources $r_1,r_2\in\mathsf{sources}_T$ and two sinks $s_1,s_2\in\mathsf{sinks}_T$, and let T' be the up-closing of T at r_1 and r_2 and let T'' be the down-closing of T at s_1 and s_2. The *ν-up-closing* (T',σ') of (T,σ) at r_1 and r_2 is defined by the extension σ' of σ,

$$\sigma'\left((r^+,r_i)\right):=\overrightarrow{\left(\bot_A^{(T,\sigma)}(r_1)\,,\bot_B^{(T,\sigma)}(r_1)\right)}\wedge\overrightarrow{\left(\bot_A^{(T,\sigma)}(r_2)\,,\bot_B^{(T,\sigma)}(r_2)\right)}$$

which maps the new edges (r^+, r_i) to separations. The *ν-down-closing* (T'', σ'') of (T, σ) at s_1 and s_2 is defined by σ'' extending σ to the new edges via

$$\sigma''\big((s_i, s^-)\big) \coloneqq \overrightarrow{\Big(\top_A^{(T,\sigma)}(s_1)\,, \top_B^{(T,\sigma)}(s_1)\Big)} \vee \overrightarrow{\Big(\top_A^{(T,\sigma)}(s_2)\,, \top_B^{(T,\sigma)}(s_2)\Big)}\,.$$

That these definitions indeed satisfy (P5) and (P6) follows from the fact that the uncrossings of two directed separations are again directed separations and because every $\vec{\mathcal{S}}$-DAG is a ν-DAG.

We can additionally prove that the ν-width of a ν-up-closing or a ν-down-closing is not larger than the ν-width of the original ν-DAG.

Lemma 3.4.13. Let D be a digraph and (T, σ) be a ν-DAG for D. If (T', σ') is obtained from (T, σ) by a ν-up-closing or a ν-down-closing and (T, σ) has ν-width ω, then (T', σ') has ν-width at most ω.

Proof. We only consider the case of ν-up-closings, because the case of ν-down-closings follows along the same lines.

Let $r_1, r_2 \in \mathsf{sources}_T$ be the pair of distinct sources such that (T', σ') is the ν-up-closing at r_1 and r_2. Assume that (T, σ) has ν-width ω. We claim that (T', σ') has ν-width at most ω.

For all nodes t of $V(T) \setminus \{r_1, r_2\}$, we have $\beta_\nu^{(T',\sigma')}(t) = \beta_\nu^{(T,\sigma)}(t)$, so the size of the bag at node t does not increase.

For r_i with $i \in \{1, 2\}$, we obtain

$$\begin{aligned}
\top^{(T',\sigma')}(r_i) &= \overrightarrow{\Big(\bot_A^{(T,\sigma)}(r_1) \cap \bot_A^{(T,\sigma)}(r_2)\,, \bot_B^{(T,\sigma)}(r_1) \cup \bot_B^{(T,\sigma)}(r_2)\Big)}\,,\\
&\text{and thus}\\
\beta_\nu^{(T',\sigma')}(r_i) &= \Big(\Big(\bot_B^{(T,\sigma)}(r_1) \cup \bot_B^{(T,\sigma)}(r_2)\Big) \cap \bot_A^{(T,\sigma)}(r_i)\Big)\\
&\subseteq \bot_A^{(T,\sigma)}(r_i)\,.
\end{aligned}$$

Therefore, the bag size of r_i does not increase for $i \in \{1, 2\}$.

Now consider r^+. As r^+ is a source, its bag is given by $\bot_A^{(T',\sigma')}(r^+) = \bot_A^{(T,\sigma)}(r_1) \cap \bot_A^{(T,\sigma)}(r_2)$, so $\beta_\nu^{T',\sigma'}(r')$ is at most of size $\max\{|\beta_\nu^{(T,\sigma)}(r_1)|, |\beta_\nu^{(T,\sigma)}(r_2)|\}$.

Thus, the ν-width of (T', σ') is at most ω. ☐

Corollary 3.4.14. Any ν-up-closure, ν-down-closure or ν-closure of a ν-DAG of ν-width ω has ν-width at most ω.

Finally, it remains to show that ν-DAG-width is complete in order to have Theorem 3.3.13 apply to ν-DAG-width. We need one further lemma to do so, which states that if the ν-width of an $\vec{\mathcal{S}}$-DAG is at most k, then we can also manipulate the separations so that they are from $\vec{\mathcal{S}}_k$.

Lemma 3.4.15. Let D be a digraph and (T, σ) be an $\vec{\mathcal{S}}$-DAG for D. If (T, σ) is of ν-width less than k, then for all $t \in V(T)$ holds

- **(i)** $\perp^{(T,\sigma)}(t) \in \vec{\mathcal{S}}_k$, and
- **(ii)** $\perp^{(T,\sigma)}(t) \in \vec{\mathcal{S}}_k$.

Additionally, there exists an $\vec{\mathcal{S}}_k$-DAG (T, σ') of the same width such that for all $t \in V(T)$ holds

- **(iii)** $\top^{(T,\sigma)}(t) = \top^{(T,\sigma')}(t)$, and
- **(iv)** $\perp^{(T,\sigma)}(t) = \perp^{(T,\sigma')}(t)$ for all $t \in V(T)$.

Proof. The size of the bag $\beta_\nu^{(T,\sigma)}(t)$ equals $|\perp_A^{(T,\sigma)}(t) \cap \top_B^{(T,\sigma)}(t)|$ and we know that $\perp_B^{(T,\sigma)}(t) \subseteq \top_B^{(T,\sigma)}(t)$ due to $\top^{(T,\sigma)}(t) \leq \perp^{(T,\sigma)}(t)$. Therefore, (T, σ) having ν-width less than k implies $|\overrightarrow{\left(\perp_A^{(T,\sigma)}(t), \perp_B^{(T,\sigma)}(t)\right)}| \leq |\beta_\nu^{(T,\sigma)}(t)| < k$ for all $t \in V(T)$ and thus (i) follows.

By Observation 3.4.2, we have $\top_A^{(T,\sigma)}(t) \cap \top_B^{(T,\sigma)}(t) \subseteq \beta_\nu^{(T,\sigma)}(t)$ and thus $\top^{(T,\sigma)}(t) \in \vec{\mathcal{S}}_k$ for all $t \in V(T)$ as well, implying (ii).

Let us define σ' for every $(d, t) \in E(T)$ as $\sigma'((d, t)) := \top^{(T,\sigma)}(t)$. By Lemma 3.2.7 and the discussion above, (T, σ') is an $\vec{\mathcal{S}}_k$-DAG.

By definition,

$$\begin{aligned}
\top_A^{(T,\sigma')}(t) &= \bigcup_{(d,t) \in E(T)} \sigma'_C((d, t)) \\
&= \bigcup_{(d,t) \in E(T)} \top_A^{(T,\sigma)}(t) = \top_A^{(T,\sigma)}(t)
\end{aligned}$$

$$\text{and} \qquad \top_B^{(T,\sigma')}(t) = \bigcap_{(d,t)\in E(T)} \sigma'_D((d,t)) = \bigcap_{(d,t)\in E(T)} \top_B^{(T,\sigma)}(t) = \top_B^{(T,\sigma)}(t)\,.$$

Thus, $\top^{(T,\sigma)}(t) = \overrightarrow{\left(\top_A^{(T,\sigma')}(t), \top_B^{(T,\sigma')}(t)\right)}$ for all $t \in V(T)$, and so (iii) holds.

As the $\perp$-separations are derived from the $\top$-separations, this implies

$$\overrightarrow{\left(\perp_A^{(T,\sigma)}(t)\,, \perp_B^{(T,\sigma)}(t)\right)} = \overrightarrow{\left(\perp_A^{(T,\sigma')}(t)\,, \perp_B^{(T,\sigma')}(t)\right)}$$

for all $t \in V(T)$, i.e. (iv).

With this, $\beta_\nu^{(T,\sigma)}(t) = \beta_\nu^{(T,\sigma')}(t)$ follows for all $t \in V(T)$, and thus, (T, σ') is of the same ν-width as (T, σ) . □

Corollary 3.4.16. Let D be a digraph. Then, there exists an $\vec{\mathcal{S}}$-DAG of width less than k for D if and only if there exists an $\vec{\mathcal{S}}_k$-DAG of width less than k for D.

We are now equipped with all necessary prerequisites to show that ν-DAG-width is complete.

Lemma 3.4.17. The digraph width measure ν-DAG-width is complete.

Proof. We need to show that ν-DAG-width is closed and contained. By Lemma 3.4.12, ν-DAG-width is shifting closed and therefore in particular weakly shifting closed, so (C4) holds. Moreover, Corollary 3.4.14 guarantees ν-closures of the same width for all ν-k-DAGs, and Lemma 3.4.15 allows us to transform these closures into ν-k-DAGs of the same width, implying (C1) to (C3). Additionally, we have $|\beta_\nu(\overrightarrow{(X,Y)})| = |X \cap Y| < k$ for all $\overrightarrow{(X,Y)} \in \vec{\mathcal{S}}_k$, which yields (Cnt). □

3.4.4 Ganglions yield robber strategies

Next we show that ganglions yield a strategy for the robber in the cops and robber reachability game. We start with a lemma about separations with the same separator and what the edge-directions for these can look like.

Lemma 3.4.18. Let D be a digraph, $k \in \mathbb{N}$, $X \subseteq V(D)$ with $|X| < k$, $\mathcal{D}$ a small ν-ganglion of order k, and $L = \{C_1, \ldots, C_\ell\}$, $P = \{K_1, \ldots, K_p\}$ where the C_i and K_j are strongly connected components of $D' := D - X$. Let $H_L := \bigcup_{i=1}^{\ell} \mathsf{Below}_{D'}(C_i)$ and $H_P := \bigcup_{j=1}^{p} \mathsf{Below}_{D'}(K_j)$. If

$$
\begin{aligned}
\overleftarrow{(B_1, A_1)} &:= \overleftarrow{(H_L \cup X, V(D) \setminus H_L)} \in \mathcal{D}, \text{ and} \\
\overleftarrow{(B_2, A_2)} &:= \overleftarrow{(H_P \cup X, V(D) \setminus H_P)} \in \mathcal{D}, \text{ then} \\
\overleftarrow{(Y, X)} &:= \overleftarrow{(H_L \cup H_P \cup X, V(D) \setminus (H_L \cup H_P))} \in \mathcal{D}.
\end{aligned}
$$

Proof. With $\overrightarrow{(X, Y)}$ being a directed separation of order less than k, in fact $\overrightarrow{(X, Y)} = \overrightarrow{(A_1, B_1)} \wedge \overrightarrow{(A_2, B_2)}$, it must be oriented by $\mathcal{D}$. If $\overleftarrow{(Y, X)} \in \mathcal{D}$ we are done, so we may assume $\overrightarrow{(X, Y)} \in \mathcal{D}$. Since, $\overrightarrow{(X, Y)} \leq \overrightarrow{(A_i, B_i)}$ for both i, and with $X \cup B_1 \cup B_2 = V(D)$ we now have three small sides covering all of $V(D)$ which is a contradiction to $\mathcal{D}$ being a small ganglion of order k. □

Lemma 3.4.19. Let D be a digraph and $\mathcal{D}$ a small ganglion of order $k + 1$, then the robber player has a winning strategy in the (not necessarily monotone) k-cops and robber reachability game.

Proof. Consider a game position for the robber turn in which the cops move from position X_1 to position X_2, and R is the robber component when X_2 was announced. Now, the robber may move within $\mathsf{Below}_{D-Y}(R)$, where $Y := X_1 \cap X_2$. We show that there is always a strongly connected component R' of $D - X_2$ such that $R' \subseteq \mathsf{Below}_{D-Y}(R)$ and $\overrightarrow{(V(D) \setminus \mathsf{Below}_{D-X_2}(R'), \mathsf{Below}_{D-X_2}(R') \cup X_2)} \in \mathcal{D}$. If this is the case, the robber can escape to a non-empty component and thus the game continues, meaning the robber player wins eventually.

For the first round let X_0 be the initial cop position, then the directed separation representing this state is $\overrightarrow{(X_0, V(D))} \in \mathcal{I}_{k+1}^{k+1}$. By Lemma 3.4.8, $\mathcal{D}$ extends $\mathcal{I}_{k+1}^{k+1}$ and thus $\overrightarrow{(X_0, V(D))} \in \mathcal{D}$. Hence, every strongly connected component R of $D - X_0$ is contained in the big side of $\overrightarrow{(X_0, V(D))}$. Suppose for all such R we have

$$
\overleftarrow{(\mathsf{Below}_{D-X_0}(R) \cup X_0, V(D) \setminus \mathsf{Below}_{D-X_0}(R))} \in \mathcal{D}.
$$

Then, Lemma 3.4.18 implies that $\overleftarrow{(V(D), X_0)} \in \mathcal{D}$ which contradicts $\mathcal{D}$ extending $\mathcal{I}_{k+1}^{k+1}$. Hence, there is at least one strongly connected component R_0 of $D - X_0$ such that

$$
\overrightarrow{(V(D) \setminus \mathsf{Below}_{D-X_0}(R_0), \mathsf{Below}_{D-X_0}(R_0) \cup X_0)} \in \mathcal{D},
$$

which we choose as the first robber position.

Let us now assume we are further along in the game. Let $X_i \subseteq V(D)$ be the current cop position and $X_{i+1} \subseteq V(D)$ be the announced new cop position. Moreover, let R_i be the current robber component with

$$(V(D) \setminus \mathsf{Below}_{D-X_i}(R_i)\,, \mathsf{Below}_{D-X_i}(R_i) \cup X_i\vec{)} \in \mathcal{D}. \tag{3.3}$$

Let $\mathcal{R}$ be the set of strong components of $D - X_{i+1}$ that intersect $\mathsf{Below}_{D-Y}(R_i)$, where $Y := X_i \cap X_{i+1}$. Suppose towards a contradiction that for all $R \in \mathcal{R}$ we have

$$\big(\mathsf{Below}_{D-X_{i+1}}(R) \cup X_{i+1}, V(D) \setminus \mathsf{Below}_{D-X_{i+1}}(R)\overleftarrow{\big)} \in \mathcal{D}.$$

So for every strongly connected component $R \in \mathcal{R}$, there is a directed separation whose small side consists of $\mathsf{Below}_{D-X_{i+1}}(R) \cup X_{i+1}$. Hence, by Lemma 3.4.18, there is a directed separation with separator X_{i+1} whose small side with respect to $\mathcal{D}$ is exactly the set

$$X_{i+1} \cup \bigcup_{R \in \mathcal{R}} \mathsf{Below}_{D-X_{i+1}}(R)\,. \tag{3.4}$$

Observe that due to $\mathsf{Below}_{D-X_i}(R_i) \subseteq \mathsf{Below}_{D-Y}(R_i)$, the following holds:

$$\begin{aligned}&(V(D) \setminus \mathsf{Below}_{D-Y}(R_i)\,, X_i \cup \mathsf{Below}_{D-Y}(R_i)\vec{)} \\ \leq\ &(V(D) \setminus \mathsf{Below}_{D-X_i}(R_i)\,, X_i \cup \mathsf{Below}_{D-X_i}(R_i)\vec{)}\,.\end{aligned}$$

By (3.3) and the consistency of $\mathcal{D}$, see Lemma 3.4.8, we obtain that

$$(V(D) \setminus \mathsf{Below}_{D-Y}(R_i)\,, X_i \cup \mathsf{Below}_{D-Y}(R_i)\vec{)} \in \mathcal{D}. \tag{3.5}$$

Next suppose that $(Y \cup \mathsf{Below}_{D-Y}(R_i), V(D) \setminus \mathsf{Below}_{D-Y}(R_i)\overleftarrow{)} \in \mathcal{D}$. Then, due to (3.5), we obtain two small sides spanning the whole graph, a contradiction to Definition 3.4.7. Thus, we obtain that

$$(V(D) \setminus \mathsf{Below}_{D-Y}(R_i)\,, Y \cup \mathsf{Below}_{D-Y}(R_i)\vec{)} \in \mathcal{D}. \tag{3.6}$$

Moreover, we have

$$\begin{aligned}&\left(X_{i+1} \cup \bigcup_{R \in \mathcal{R}} \mathsf{Below}_{D-X_{i+1}}(R)\,, V(D) \setminus \bigcup_{R \in \mathcal{R}} \mathsf{Below}_{D-X_{i+1}}(R)\overleftarrow{\right)} \\ \leq\ &(X_{i+1} \cup \mathsf{Below}_{D-Y}(R_i)\,, V(D) \setminus \mathsf{Below}_{D-Y}(R_i)\overleftarrow{)}\end{aligned}$$

which, due to $\mathcal{D}$ being consistent (see Lemma 3.4.8) and (3.4) implies

$$\overleftarrow{(X_{i+1} \cup \mathsf{Below}_{D-Y}(R_i), V(D) \setminus \mathsf{Below}_{D-Y}(R_i))} \in \mathcal{D}. \tag{3.7}$$

Finally we obtain a contradiction to $\mathcal{D}$ being a small ν-ganglion, as we have

$$\begin{aligned} &\overrightarrow{(V(D) \setminus \mathsf{Below}_{D-Y}(R_i), \mathsf{Below}_{D-Y}(R_i) \cup Y)} \\ \leq\ &\overleftarrow{(X_{i+1} \cup \mathsf{Below}_{D-Y}(R_i), V(D) \setminus \mathsf{Below}_{D-Y}(R_i))} \end{aligned}$$

and (3.6) as well as (3.7).

Therefore, there exists some $R \in \mathcal{R}$ such that $\mathsf{Below}_{D-X_{i+1}}(R) \cup X_{i+1}$ is the big side with respect to $\mathcal{D}$. So we can choose $R_{i+1} := R$ ensuring again that

$$\overrightarrow{\big(V(D) \setminus \mathsf{Below}_{D-X_{i+1}}(R_{i+1}), \mathsf{Below}_{D-X_{i+1}}(R) \cup X_{i+1}\big)} \in \mathcal{D}. \qquad \square$$

3.4.5 A non-monotone cop strategy

Here, we prove that a bound on the ν-DAG-width of a digraph yields a bound on the number of cops needed to win the cops and robber reachability game. Unfortunately, as described later on, the obtained strategies are not monotone. We need a few statements to show that every ν-DAG-decomposition yields a strategy for the cop player.

Lemma 3.4.20. Let D be a digraph, (T, σ) be an $\vec{\mathcal{S}}$-DAG for D, and $x \in V(D)$. If there is some $t \in V(T)$ with $x \in \top_B(t)$, then there is $t' \in V(T)$ with $t \preccurlyeq_T t'$ such that $x \in \beta_\nu^{(T,\sigma)}(t')$.

Proof. Let $P = t_0, \ldots, t_\ell$ be a longest directed path in T such that $x \in \top_B(t_i)$ for all $0 \leq i \leq \ell$ and $t_0 = t$. If t_ℓ is a sink, then $x \in \beta_\nu^{(T,\sigma)}(t_\ell)$ by definition, so suppose t_ℓ has at least one out-neighbour. Since $x \notin \top_B(d)$ for every $d \in N_T^{\mathsf{out}}(t_\ell)$ by choice of ℓ, we have $x \in \bigcap_{(t_\ell,d) \in E(T)} \top_A(d) = \bot_A(t_\ell)$ and thus $x \in \top_B(t_\ell) \cap \bot_A(t_\ell) = \beta_\nu^{(T,\sigma)}(t_\ell)$. $\square$

Because $\top_B(r) = V(D)$ for all sources r of T, we obtain the following corollaries.

Corollary 3.4.21. Let D be a digraph, (T, σ) be an $\vec{\mathcal{S}}$-DAG for D. For every $x \in V(D)$ there is a $t \in V(T)$ with $x \in \beta_\nu^{(T,\sigma)}(t)$.

Corollary 3.4.22. Let D be a digraph, (T, σ) be an $\vec{\mathcal{S}}$-DAG for D, and $t \in V(T)$. Then, $\beta_\nu^{(T,\sigma)}(\mathsf{Below}_T(t)) = \top_B(t)$.

Proof. From Lemma 3.4.20 we directly obtain $\top_B(t) \subseteq \beta_\nu^{(T,\sigma)}(\mathsf{Below}_T(t))$. Hence, we prove $\beta_\nu^{(T,\sigma)}(\mathsf{Below}_T(t)) \subseteq \top_B(t)$. If t is a source, then this holds due to $\top_B(r) = V(D)$. If t is not a source, then $\beta_\nu^{(T,\sigma)}(t') = \top_B(t') \cap \bot_A(t') \subseteq \top_B(t')$ for all $t' \in \mathsf{Below}_T(t)$. By Lemma 3.2.8, we know that $\top_B(t') \subseteq \top_B(t)$ for all $t' \in \mathsf{Below}_T(t)$. Thus, $\beta_\nu^{(T,\sigma)}(\mathsf{Below}_T(t)) \subseteq \top_B(t)$. □

Now we can prove that bounded ν-DAG-width provides a strategies for the cop player.

Theorem 3.4.23. Let D be a digraph and (T, σ) an $\vec{\mathcal{S}}_k$-DAG of width $\omega \geq k$. Then, the cop player has a (not necessarily monotone) winning strategy for the 3ω-cops and robber reachability game.

Proof. For every $t \in V(T)$ we define

$$\alpha(t) := \beta_\nu^{(T,\sigma)}(t) \cup \bigcup_{d \in N_T^{\mathsf{out}}(d)} \beta_\nu^{(T,\sigma)}(d) .$$

Since every node of T has at most two out-neighbours and $|\beta_\nu^{(T,\sigma)}(t)| \leq \omega$ for all t, we have $|\alpha(t)| \leq 3\omega$ for all $t \in V(T)$.

As initial cop position we choose $\alpha(r)$ for some source r of T. The robber must choose a position within $\top_B(r)$, because $\top_B(r) = V(D)$.

We now describe a strategy for the following situation: The cops are placed on all vertices of $\alpha(t)$ for some $t \in V(T)$ and the robber component R_t is contained in $\top_B(t)$.

In case t is a sink we know that $\alpha(t) = \beta_\nu^{(T,\sigma)}(t) = \top_B(t)$ and the robber is captured. So, if we can force the situation above after each move of the cop player, then the cops eventually capture the robber, because T is finite.

Thus, we may assume that t is not a sink node. Let t_1 and t_2 be the out-neighbours of t. The robber chooses a strongly connected component R_t in $D - \alpha(t)$. By Lemma 3.4.20 and Corollary 3.4.22, the robber must choose a position in $\top_B(t_1)$ or $\top_B(t_2)$. Let R_t be the strongly connected component of the robber in $D - \alpha(t)$. Without loss of generality, assume that $R_t \subseteq \top_B(t_1)$. We choose $\alpha(t_1)$ as the next cop position. Moving the cops from $\alpha(t)$ to $\alpha(t')$ may remove cops from $\alpha(t)$. Still, the robber cannot leave $\top_B(t_1)$, because $\beta_\nu^{(T,\sigma)}(t_1) \subseteq \alpha(t) \cap \alpha(t_1)$ and, by Corollary 3.4.1, $\widehat{(\bot_A(t_1), \top_B(t_1))}$ is a directed separation. Thus, the new robber component R_{t_1} is contained in $\top_B(t_1)$.

This procedure can be iterated until the cop player reaches a sink. Moreover, we never use more than 3ω cops and thus the proof is complete. □

3.5 Monotonicity

Unfortunately, the strategies obtained from ν-DAG-decompositions described in Subsection 3.1.2 are not monotone. In order solve Conjecture 3.1.6 in the positive, we would need a width measure having the duality as well as yielding monotone strategies as the ganglions give a lower bound on the number of cops linear in the number of cops sufficient for a monotone strategy. In this section, we introduce a restriction on $\vec{\mathcal{S}}$-DAGs that provides a width measure that yields a monotone strategy for the cops in the cops and robber reachability game. We also prove that this width measure is parametrically equivalent to DAG-width. We finish by a discussion why we cannot apply our general duality result for this width measure.

In the preceding subsection we established that on a digraph of ν-DAG-width at most k the cop player has a winning strategy using at most $3k$ cops. For DAG-width (in contrast) this is an equality instead of a factor of 3. This is due to the fact that DAG-decompositions demand the intersection of two consecutive bags to strongly guard the vertices in the bags of the subtree defined by the smaller of the two nodes in the decomposition. This is not the case for $\vec{\mathcal{S}}$-DAGs in general. If (T, σ) is an $\vec{\mathcal{S}}$-DAG for a digraph D and $(d, t) \in E(T)$, then the tuple $\Big((V(D) \setminus \top_B(t)) \cup \beta_\nu^{(T,\sigma)}(d), \top_B(t)\Big)$ is not necessarily a directed separation. To be more precise, if t' is the other out-neighbour of d in T, the sets $(\top_B(t') \setminus \top_A(t')) \cap (\top_A(t) \cap \top_B(t))$ and $(\top_B(t) \setminus \top_A(t)) \cap (\top_A(t') \cap \top_B(t'))$ can be non-empty. Hence, as there can be edges from $\top_A(t) \cap \top_B(t)$ to $\top_B(t') \setminus (\top_A(t') \cup \top_B(t))$, there can be edges in both directions between $\big((V(D) \setminus \top_B(t)) \cup \beta_\nu^{(T,\sigma)}(d)\big) \setminus \top_B(t)$ and $\top_B(t) \setminus \big((V(D) \setminus \top_B(t)) \cup \beta_\nu^{(T,\sigma)}(d)\big)$. See the areas marked in orange in Figure 3.6.

Definition 3.5.1. Let D be a digraph and $\vec{\mathcal{S}}' \subseteq \vec{\mathcal{S}}$. An $\vec{\mathcal{S}}'$-DAG (T, σ) is called a *μ-$\vec{\mathcal{S}}$-DAG*, or μ-DAG, if for all $d \in V(T)$ we have

$$\bigcup_{(d,t) \in E(T)} \left(\top_A(t) \cap \bigcap_{(d,t') \in E(T)} \top_B(t') \right) \subseteq \bot_A(d),$$

which we call the *μ-property*, see Figure 3.6 for an illustration. The *μ-DAG-width* of D, denoted $\mu(D)$, is defined as the minimum ν-width over all μ-$\vec{\mathcal{S}}$-DAGs for D. $\dashv$

$\mu(D)$

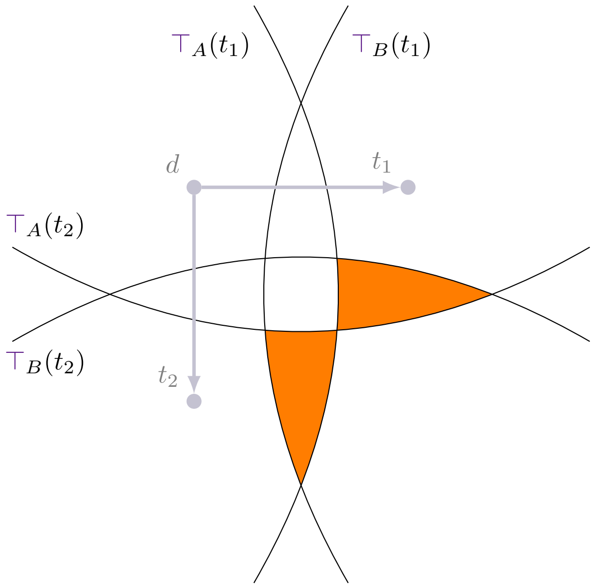

Figure 3.6: The orange areas are ensured to be empty.

Even though we restrict the $\vec{\mathcal{S}}$-DAGs by imposing a new property on them, the definitions for the requirements (P1) to (P6) stay the same as for ν-DAG-width. The two lemmata 3.4.4 and 3.4.5 still hold for ν-DAG-width. That is, $\beta_\mu = \beta_\nu$ and $\mathsf{eval}_\mu = \mathsf{eval}_\nu = |\cdot|$. Also, the μ-up-closing is defined to be the ν-up-closing at the corresponding sources and the μ-down-closing is defined to be the ν-down-closing at the corresponding sinks.

These definitions ensure that μ-$\vec{\mathcal{S}}$-DAGs contain strong guards as the following lemma shows.

Lemma 3.5.2. Let D be a digraph, (T, σ) be a μ-$\vec{\mathcal{S}}$-DAG for D and $(d, t) \in E(T)$. Then, $\top_A(t) \cap \top_B(t) \subseteq \beta_\mu^{(T,\sigma)}(d)$ and $\beta_\mu^{(T,\sigma)}(d) \cap \beta_\mu^{(T,\sigma)}(t)$ strongly guards $\top_B(t)$.

Proof. Suppose there is a vertex $x \in (\top_A(t) \cap \top_B(t)) \setminus \beta_\mu^{(T,\sigma)}(d)$. Then, by definition, $x \in (\top_A(t) \cap \top_B(t)) \setminus \beta_\nu^{(T,\sigma)}(d)$. We have $(\top_A(d), \top_B(d)) \leq \top(t)$ by

Lemma 3.2.8, and therefore $\top_B(t) \subseteq \top_B(d)$. Thus, x cannot lie in $\bot_A(d)$, as this would also mean that x lies in $\beta_\nu^{(T,\sigma)}(d)$, a contradiction. Hence, we may assume $x \in \bot_B(d) = \bigcup_{(d,t') \in E(T)} \top_B(t')$. If t is the only out-neighbour of d, then $\bot_A(d) = \bigcap_{(d,t')} \top_A(t') = \top_A(t)$ and therefore $x \in \beta_\nu^{(T,\sigma)}(d)$, a contradiction. Thus, we consider the case in which there exists an out-neighbour t' of d that is distinct from t with $x \notin \top_B(t')$. Therefore, $x \in \top_A(t) \cap \top_B(t) \cap \top_B(t')$. Due to $x \notin \bot_A(d)$, this implies $\top_A(t) \cap \top_B(t) \cap \top_B(t') \nsubseteq \bot_A(d)$, contradicting (T, σ) being a μ-$\vec{\mathcal{S}}$-DAG. Hence, the first part of the statement holds.

Then, the second part of the statement follows, because, by Observation 3.4.2, we know that $\top_A(t) \cap \top_B(t) \subseteq \beta_\nu^{(T,\sigma)}(t) = \beta_\mu^{(T,\sigma)}(t)$ and, by the above, $\top_A(t) \cap \top_B(t) \subseteq \beta_\mu^{(T,\sigma)}(d)$. Therefore, the fact that $\top_A(t) \cap \top_B(t)$ strongly guards $\top_B(t)$ implies that $\beta_\mu^{(T,\sigma)}(d) \cap \beta_\mu^{(T,\sigma)}(t)$ strongly guards $\top_B(t)$. $\square$

The μ-property allows us to translate μ-$\vec{\mathcal{S}}$-DAGs into DAG-decompositions.

Lemma 3.5.3. Let D be a digraph and (T, σ) a μ-$\vec{\mathcal{S}}$-DAG for D. Then, $\left(T, \beta_\nu^{(T,\sigma)}\right)$ is a DAG-decomposition of D.

Proof. Since T is a DAG and $\beta_\nu^{(T,\sigma)}$ a mapping of $V(T)$ to sets of vertices of D, it remains to check the three properties (Dag1) to (Dag3) required for a DAG-decomposition.

By Corollary 3.4.21, the union of all $\beta_\nu^{(T,\sigma)}(t)$ equals $V(D)$, so (Dag1) holds. Let $t_1, t_2, t_3 \in V(T)$ such that $t_1 \preccurlyeq_T t_2$ and $t_2 \preccurlyeq_T t_3$. By Lemma 3.4.3, we know $\beta_\nu^{(T,\sigma)}(t_1) \cap \beta_\nu^{(T,\sigma)}(t_3) \subseteq \beta_\nu^{(T,\sigma)}(t_2)$, thus (Dag2) holds. In order to see that (Dag3) holds as well, let $(d, t) \in E(T)$. By Lemma 3.5.2, the set $\beta_\nu^{(T,\sigma)}(d) \cap \beta_\nu^{(T,\sigma)}(t)$ strongly guards $\top_B(t)$ which equals $\beta_\nu^{(T,\sigma)}(\mathsf{Below}_T(t))$ by Corollary 3.4.22. $\square$

The previous lemma shows that μ-DAG-width yields an upper bound on DAG-width. Next, we show the other direction, that is, how to obtain a μ-DAG from a nice DAG-decomposition in order to obtain the following theorem.

Theorem 3.5.4. For all digraphs D holds $\mathsf{DAGw}(D) \leq \mu(D) \leq 3\mathsf{DAGw}(D)$.

We have already seen in Lemma 3.1.2 that a DAG-decomposition naturally provides a way to obtain separations that can be assigned to the edges of the DAG it uses.

Let D be a digraph and (T, β) a nice DAG-decomposition of D. Then, we define $\sigma_\beta \colon E(T) \to \vec{\mathcal{S}}$ by

$$\sigma_\beta((d,t)) := \overrightarrow{((V(D) \setminus \beta(\mathsf{Below}_T(t))) \cup (\beta(d) \cap \beta(t)), \beta(\mathsf{Below}_T(t)))}.$$

The following lemma shows that we can obtain an $\vec{\mathcal{S}}$-DAG from a nice DAG-decomposition. The obtained $\vec{\mathcal{S}}$-DAG has several additional properties.

Lemma 3.5.5. Let D be a digraph and (T, β) a nice DAG-decomposition of D. Then,

(i) (T, σ_β) is an $\vec{\mathcal{S}}$-DAG.

Additionally, the following properties hold for every $d, t \in V(T)$ and $x \in V(D)$.

(ii) If $x \in \top_B(d)$, but not in $\bot_B(d)$, then $x \in \beta(d) \cap \bot_A(d)$.

(iii) If $x \in \beta(d)$, $d \preccurlyeq_T t$ and $x \in \top_B(t)$, then $x \in \beta(t)$.

(iv) If $x \in \beta(d) \cap \bot_B(d)$, then $x \in \bot_A(d)$.

(v) If $x \in (\bot_A(d) \cap \bot_B(d)) \setminus \beta(d)$, then $x \in \bigcup_{(d,t') \in E(T)} \beta(t')$. In particular, $\top_A(t') \cap \top_B(t') \subseteq \beta_\nu^{T,\sigma_\beta}(d)$ for all $t' \in N_T^{\mathsf{out}}(d)$.

Proof. We prove the properties one by one.

Since (T, β) is a nice DAG-decomposition, T has maximum out-degree at most two. Hence, it suffices to show that σ_β satisfies the $\vec{\mathcal{S}}$-DAG consistency. By Lemma 3.1.2, $\sigma_\beta(e)$ is indeed a directed separation for all $e \in E(T)$. Moreover, for all (t_1, t_2), $(t_2, t_3) \in E(T)$, we have

$$\beta(\mathsf{Below}_T(t_3)) \subseteq \beta(\mathsf{Below}_T(t_2)). \tag{3.8}$$

Due to (T, β) being a DAG-decomposition, we know $\beta(t_1) \cap \beta(t_3) \subseteq \beta(t_2)$, and thus, $x \in \beta(\mathsf{Below}_T(t_3)) \cap \beta(t_1)$ implies $x \in \beta(t_3)$, because T is a DAG and therefore $t_1 \notin \mathsf{Below}_T(t_3)$. Thus, $\beta(t_1) \setminus \beta(t_3) \subseteq V(D) \setminus \beta(\mathsf{Below}_T(t_3))$. Therefore, we obtain

$$\begin{aligned}
&(V(D) \setminus \beta(\mathsf{Below}_T(t_2))) \cup (\beta(t_1) \cap \beta(t_2)) \\
\subseteq\ &(V(D) \setminus \beta(\mathsf{Below}_T(t_3))) \cup (\beta(t_1) \cap \beta(t_2) \cap \beta(t_3)) \cup \\
&((\beta(t_1) \cap \beta(t_2)) \setminus \beta(t_3)) \\
\subseteq\ &(V(D) \setminus \beta(\mathsf{Below}_T(t_3))) \cup (\beta(t_2) \cap \beta(t_3)) \cup (\beta(t_1) \setminus \beta(t_3)) \\
\subseteq\ &(V(D) \setminus \beta(\mathsf{Below}_T(t_3))) \cup (\beta(t_2) \cap \beta(t_3)).
\end{aligned} \tag{3.9}$$

Together, containments (3.8) and (3.9) imply that $\sigma_\beta((t_1, t_2)) \leq \sigma_\beta((t_2, t_3))$. Thus, all properties of an $\vec{\mathcal{S}}$-DAG hold and therefore (i) does.

Next, let $d \in V(T)$ and choose $x \in \top_B(d) \setminus \bot_B(d)$. By definition of σ_β, we have $\top_B(d) = \beta(\mathsf{Below}_T(d))$ and $\bot_B(d) = \bigcup_{(d,t)\in E(T)} \beta(\mathsf{Below}_T(t))$. Hence, $x \in \beta(d)$. Moreover, due to $\overrightarrow{(\bot_A(d), \bot_B(d))}$ being a directed separation, $x \notin \bot_B(d)$ implies $x \in \bot_A(d)$. Thus, (ii) holds.

Let $d, t \in V(T)$ be such that $d \preccurlyeq_T t$. Then, choose $x \in \beta(d) \cap \top_B(t)$. We can choose a maximal directed path P starting in t such that $x \in \top_B(t'')$ for all $t'' \in V(P)$. Let t' be the end-node of P. By the maximality of P, we obtain $x \notin \top_B(t^*)$ for all $t^* \in N_T^{\mathsf{out}}(t')$, thus $x \notin \bot_B(t')$ and therefore, by (ii), $x \in \beta(t')$. Since t lies on a directed path between d and t' and $x \in \beta(d) \cap \beta(t')$, (iii) follows from (T, β) being a DAG-decomposition.

Let $d \in V(T)$ and choose $x \in \beta(d) \cap \bot_B(d)$. Suppose $x \notin \bot_A(d)$. Then, there is an out-neighbour t of d such that $x \in \top_B(t) \setminus \top_A(t)$. By (iii), this implies $x \in \beta(t)$ and thus $x \in \beta(d) \cap \beta(t)$. From the definitions of σ_β and $\top_A(t)$ we obtain $\beta(d) \cap \beta(t) \subseteq \top_A(t)$, contradicting $x \notin \bot_A(d)$. Therefore, $x \in \bot_A(d)$ and thus (iv) holds.

Finally, let $d \in V(T)$ and $x \in (\bot_A(d) \cap \bot_B(d)) \setminus \beta(d)$. As $x \in \bot_A(d) \cap \bot_B(d)$, there is an out-neighbour t of d such that $x \in \top_A(t) \cap \top_B(t)$. We have

$$
\begin{aligned}
&\top_A(t) \cap \top_B(t) \\
&= \left(\bigcup_{(t',t)\in E(T)} (V(D) \setminus \beta(\mathsf{Below}_T(t))) \cup (\beta(t') \cap \beta(t)) \right) \cap \beta(\mathsf{Below}_T(t)) \\
&= \bigcup_{(t',t)\in E(T)} (\beta(t') \cap \beta(t)) \subseteq \beta(t)\,.
\end{aligned}
$$

So, $x \in \beta(t) \subseteq \bigcup_{(d,t')\in E(T)} \beta(t')$. For every d with exactly one out-neighbour t, by (Nice4) we have $||\beta(d)| - |\beta(t)|| \leq 1$. If $\beta(t) \subseteq \beta(d)$, then the above immediately implies $\top_A(t) \cap \top_B(t) \subseteq \beta(d) \cap \beta(t)$ and thus $\top_A(t) \cap \top_B(t) \subseteq \beta_\nu^{(T,\sigma_\beta)}(d)$. Hence, there exists an $x \in \beta(t) \setminus \beta(d)$. For this vertex holds $x \in \top_B(t) \setminus \top_A(t)$ and therefore, we still obtain $\top_A(t) \cap \top_B(t) \subseteq \beta(d) \cap \beta(t)$. For every d with two out-neighbours t and t', we have $\beta(t) = \beta(d) = \beta(t')$, by (Nice3), and thus the above again implies $\top_A(t) \cap \top_B(t) \subseteq \beta(d) \cap \beta(t)$. In every case $\top_A(t) \cap \top_B(t) \subseteq \beta_\nu^{T,\sigma_\beta}(d)$ and thus (v) holds. $\square$

These properties in fact suffice to ensure the μ-property.

Corollary 3.5.6. Let D be a digraph and (T, β) a nice DAG-decomposition of D. Then, (T, σ_β) has the μ-property.

Proof. Let $d \in V(T)$ and $t \in N_T^{\mathsf{out}}(d)$. Let $x \in \top_A(t) \cap \bigcap_{(d,t') \in E(T)} \top_B(t')$. As $\top_A(t) \cap \bigcap_{(d,t') \in E(T)} \top_B(t') \subseteq \top_A(t) \cap \top_B(t)$ and, by (v) of Lemma 3.5.5, $\top_A(t) \cap \top_B(t) \subseteq \beta_\nu^{(T,\sigma)}(d) = \bot_A(d) \cap \top_B(d)$, this implies $x \in \bot_A(d)$. Therefore, (T, σ_β) has the μ-property. $\square$

The $\vec{\mathcal{S}}_k$-DAG obtained by Lemma 3.4.15 from some $\vec{\mathcal{S}}$-DAG (T, σ) of ν-width less than k has the μ-property if and only if (T, σ) has the μ-property. Thus, we obtain the following corollary.

Corollary 3.5.7. Let D be a digraph and (T, σ) a μ-$\vec{\mathcal{S}}$-DAG of ν-width less than k. Then, there exists a μ-$\vec{\mathcal{S}}_k$-DAG of the same width.

Combining the above results yields the reverse of Lemma 3.5.3.

Corollary 3.5.8. Let D be a digraph and (T, β) be a nice DAG-decomposition of width k for D. Then, (T, σ_β) is a μ-$\vec{\mathcal{S}}$-DAG of ν-width at most $3k$.

Hence, we have established that μ-DAG-width is parametrically equivalent to DAG-width.

Theorem 3.5.4. For all digraphs D holds $\mathsf{DAGw}(D) \leq \mu(D) \leq 3\mathsf{DAGw}(D)$.

Thus, due to Theorem 3.1.5, every μ-DAG-decomposition yields a monotone strategy for the cop player.

In order to be able to apply Theorem 3.3.13 to μ-DAG-width, we would need to show that it is complete. We can instantiate ω-big μ-ganglions in the same way as ω-big ν-ganglions, and so we would only need to show that μ-DAG-width is closed. As every closed μ-k-DAG is a closed $\vec{\mathcal{S}}_k$-DAG, by Lemma 3.4.17, the existence of an ω-big ν-ganglion implies that every closed μ-k-DAG has a bag of size at least ω. We can prove that every μ-DAG has a closure as well as that μ-DAG-width is closed under down-shifts. But unfortunately, μ-DAG-width does not seem to be closed under up-shifts.

Lemma 3.5.9. The μ-DAG-width measure is closed under down-shifts.

Proof. Let (T', σ') be a down-shift of (T, σ) onto $\overset{\frown}{(X, Y)}$ at node $d \in V(T)$. As the bag function as well as the evaluation function are the same for μ-DAG-width as for ν-DAG-width, (S3) follows from Lemma 3.4.12.

Using Lemma 3.2.15 and that (T, σ) has the μ-property, we obtain that for all $t \in V(T')$ holds

$$\begin{aligned}
&\bigcup_{(t,t_i)\in E(T')} \left(\top_A^{(T',\sigma')}(t_i) \cap \bigcap_{(t,t_j)\in E(T')} \top_B^{(T',\sigma')}(t_j) \right) = \\
&\bigcup_{(t,t_i)\in E(T')} \left(\left(\top_A^{(T,\sigma)}(t_i) \cup X\right) \cap \bigcap_{(t,t_j)\in E(T')} \left(\top_B^{(T,\sigma)}(t_j) \cap Y\right) \right) \subseteq \\
&\bigcup_{(t,t_i)\in E(T')} \left(\left(\top_A^{(T,\sigma)}(t_i) \cap \bigcap_{(t,t_j)\in E(T')} \top_B^{(T,\sigma)}(t_j)\right) \cup X \right) \subseteq \\
&\bot_A^{(T,\sigma)}(d) \cup X = \bot_A^{(T',\sigma')}(t)\,.
\end{aligned} \tag{3.10}$$

Thus (S1) holds. If (T, σ) is a μ-k-DAG, then due to Lemma 3.4.12, (T', σ') is a ν-k-DAG and together with (3.10), this implies (S2). □

Lemma 3.5.10. If (T, σ) is a μ-DAG, then so is any ν-up-closing (T', σ') and any ν-down-closing (T'', σ'') of (T, σ).

Proof. First, we show the statement for ν-up-closings. Let (T', σ') be an up-closing at r_1 and r_2. Then, we have to show that the μ-property holds at the new source r^+. As $\top^{(T',\sigma')}(r_1) = \sigma'((r^+, r_1)) = \sigma'((r^+, r_2)) = \top^{(T',\sigma')}(r_2)$, we obtain that

$$\begin{aligned}
\bigcup_{(r^+,r_i)\in E(T')} \left(\top_A^{(T',\sigma')}(r_i) \cap \bigcap_{(r^+,r_j)\in E(T')} \top_B^{(T',\sigma')}(r_j) \right) &= \\
\top_A^{(T',\sigma')}(r_1) \cap \top_B^{(T',\sigma')}(r_1) &\subseteq \\
\top_A^{(T',\sigma')}(r_1) &= \\
\top_A^{(T',\sigma')}(r_1) \cap \top_A^{(T',\sigma')}(r_2) &= \\
\bot_A^{(T',\sigma')}(r^+)\,.
\end{aligned}$$

Let (T'', σ'') be a ν-down-closing at s_1 and s_2. The μ-property has to be proven for s_1 and s_2 only, because it trivially holds for the new sink s^+, for which we have $\perp_A^{(T'',\sigma'')}(s^+) = V(D)$. Since $|N^{\mathsf{out}}_{T''}(s_i)| = 1$, the μ-property also holds there. $\square$

Together, Corollary 3.4.14 and Lemma 3.5.10 imply the following.

Corollary 3.5.11. Let D be a digraph and (T, σ) a ν-DAG for D. Every (T', σ') where T' has a unique source, and (T', σ') is obtained from (T, σ) by ν-up-closings, is a ν-up-closure of (T, σ) that additionally is a μ-DAG. Moreover, every (T'', σ''), where T'' has a unique sink and (T'', σ'') is obtained from (T, σ) by ν-down-closings, is a ν-down-closure of (T, σ) that additionally is a μ-DAG. Also, if (T, σ) is of μ-width ω, then so are (T', σ') and (T'', σ'').

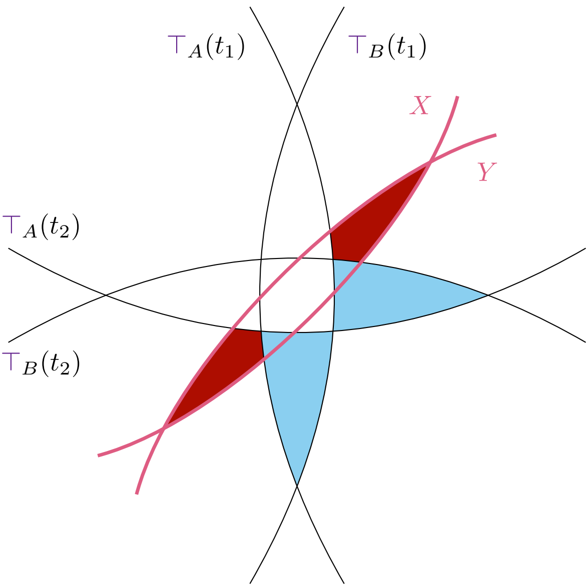

Figure 3.7: By uncrossing with (X, Y) the parts in the separators that the μ-property ensures to be empty (marked in blue) are replaced by parts of the separator $X \cap Y$, which we do not know to be empty (marked in red).

The missing piece for μ-DAG-width to admit μ-ganglions as obstructions is the closure under up-shifts. This would imply that Conjecture 3.1.6 holds. Unfortunately, we do not believe that this is true, because the uncrossings made during the up-shift affect the part of the separator used for the μ-property, see Figure 3.7 for an illustration. Thus, we cannot apply Theorem 3.3.13 to μ-ganglions and μ-DAG-width.

3.6 Conclusion and open questions

We set out to find a directed generalisation of tangles with a similar relation to a width measure as tangles have to undirected treewidth. Additionally, we would have liked the decompositions of the width measure to provide monotone cop strategies in the cops and robber reachability game, similar to how tree-decompositions do for the undirected cops and robber game.

Overall, we introduced two width measures, the ν-DAG-width and the μ-DAG-width. For ν-DAG-width we proved the duality to the defined ω-big ν-ganglions. However, we noticed that ν-DAG-width is too weak, because the cop-strategies it provides are not necessarily monotone. On the other hand, μ-DAG-width does provide monotone strategies. However, it is too strong, because we are no longer able to obtain the duality to ω-big μ-ganglions. Yet, the introduced width measures are not too far away from fulfilling the desired properties. The non-monotonicity in the strategies derived from the ν-DAGs is local. This can be seen by taking a closer look at the strategy found in the proof of Theorem 3.4.23. Considering a game played according to such a strategy, in the step corresponding to a vertex in the ν-DAG, the cops occupy the three corresponding separations. The robber then chooses one of the two subtrees and the cops remove the part of the separator for the other subtree. So the strategy places the cops in one turn, and in their next turn immediately removes them again. This is the only non-monotonic behaviour of the strategy.

However, it is not surprising that finding the desired width measure is not easy, as we discussed the close relation of these concepts to the long-standing open conjecture, Conjecture 3.1.6, by Berwanger et. al [BDH+12]. It states that the monotonicity cost for the cops and robber reachability game is linear and proving the closure of μ-DAG-width under up-shifts would imply the conjecture to be true.

4 Digraphs of cyclewidth one

In this chapter we introduce another directed width measure, the *cyclewidth*, which is parametrically equivalent to directed treewidth. The two major advantages cyclewidth has over directed treewidth are that it is closed under taking butterfly minors (which directed treewidth is not in its standard definition) and its close relation to matching theory, which is also explained within this chapter.

Cyclewidth was introduced in order to prove a matching grid theorem for bipartite graphs [HRW19]. Its rather elegant properties make it worthy of study in its own rights.

The classes having small undirected treewidth are rather well understood. Graphs of treewidth one are forest and graphs of treewidth at most two are series-parallel graphs. There are also characterisations by forbidden minors for these classes. For most directed width measures this kind of understanding is not yet established. Only recently Wiederrecht [Wie20] achieved a characterisation of the digraphs with directed treewidth one. For cyclewidth we can give the following characterisation of graphs with cyclewidth one by allowed strongly 2-connected butterfly minors. This is the main result of this chapter.

Theorem 4.0.1. Let D be a digraph. Then, D has cyclewidth exactly one if and only if every strongly 2-connected butterfly minor of D is isomorphic to the digon or the bi-directed K_3.

The proof of this theorem makes use of the interrelatedness between strongly connected digraphs and matching covered bipartite graph, which we explain closer in Section 4.2.

Branch decompositions were introduced by Robertson and Seymour with the definition of branchwidth which is parametrically equivalent to treewidth on undirected graphs [RS91]. These decompositions were also used to obtain further graph parameters, a popular one being the maximum matching width [Vat12], which is parametrically equivalent to treewidth as well and can be used to relate treewidth to other parameters and also yields algorithmic improvements [JST18].

Cyclewidth is a directed width measure based on a branch decomposition, that is, the vertices of a given digraph are mapped one-to-one to the leaves of a cubic tree and a function to measure the width at an edge in the tree is provided. In order to achieve a connection to cops and robber games (which corresponds to directed treewidth), the width at an edge should be at least the size of a hitting set for all cycles crossing the edge cut induced by this edge.

At first glance, our definition might seem unintuitive, as we count several edges for a single cycle. However, we cannot use the size of a maximum cycle packing instead. This is due to a result by Kawarabayashi et al. [KKKK13] showing that directed cycles through a specific set do not have the Erdős-Pósa-property, thus the maximum cycle packing does not yield a bound on the smallest hitting set for all cycles. So, we need something else to identify an upper bound on a hitting set of all cycles.

We introduce a few definitions for working with the decomposition trees. For every tree $\mathsf{L}(T)$ T, we define the *leaves* of T by $\mathsf{L}(T) := \{v \in V(T) \mid |N_T(v)| = 1\}$. A tree T is *cubic* if all *inner vertices*, that is, vertices in $V(T) \setminus \mathsf{L}(T)$, have degree exactly 3. It is called *subcubic* if the degree of these vertices is at most 3. If T is a subcubic tree, we can obtain a cubic tree T' from it by iteratively contracting one of the incident edges of a degree-2-vertex. We say T' is obtained from T by *trimming*. Also, by identifying the contraction vertex of an edge incident to a leaf with the original leaf, we can assume $\mathsf{L}(T) = \mathsf{L}(T')$. Let T be an undirected tree and X a set. Then, for every function $\alpha : \mathsf{L}(T) \to X$ and every subtree T' of T, we define $\alpha(T') := \{\alpha(t) \mid t \in \mathsf{L}(T) \cap V(T')\}$. Removing an edge $e = \{t_1, t_2\} \in E(T)$ from the tree T splits T into two subtrees: T_1 containing $T \ltimes e$ the vertex t_1 and T_2 containing the vertex t_2. We define $T \ltimes e := (T_1, T_2)$. If we are also given a bijection $\alpha : \mathsf{L}(T) \to V(G)$ for some graph or digraph G, we additionally $\partial(t_1t_2)$ define the edge cut $\partial(t_1t_2)$ in G that contains all edges with one endpoint in $\alpha(T_1)$ and the other in $\alpha(T_2)$, i.e. $\partial(t_1t_2) := \partial(\alpha(T_1))$. While working with cyclewidth we use this for digraphs, later on when defining matching width we make use of this for undirected graphs.

Intuitively, a cycle decomposition assigns every vertex of a given digraph to a leaf in a cubic tree. This way every edge in the tree induces an edge cut in the digraph. We define the porosity of such an edge cut by the maximum number of edges a family of disjoint cycles can have in it. Formally, let D be a digraph and $X \subseteq V(D)$. We define the *cycle porosity* of the cut $\partial_D(X)$ as follows.

$\mathsf{cp}(\partial(X))$

$$\mathsf{cp}(\partial(X)) := \max_{\substack{\mathcal{C} \text{ family of pairwise} \\ \text{disjoint directed cycles} \\ \text{in } D}} \left| \partial(X) \cap \bigcup_{C \in \mathcal{C}} E(C) \right| .$$

This then naturally yields the width measure cyclewidth.

Definition 4.0.2 (Cyclewidth)**.** Let D be a digraph. A *cycle decomposition* of D is a tuple (T, φ) , where T is a cubic tree (i.e. all inner vertices have degree three) and $\varphi : \mathsf{L}(T) \rightarrow V(D)$ is a bijection. The *width* of a cycle decomposition (T, φ) is given by

$$\max_{e \in E(T)} \frac{\mathsf{cp}(\partial(e))}{2}$$

and the *cyclewidth* of D is defined as

$$\mathsf{cyw}(D) := \min_{\substack{(T,\varphi) \text{ cycle decomposition} \\ \text{of } D}} \max_{t_1 t_2 \in E(T)} \mathsf{cp}(\partial(t_1 t_2))/2. \qquad \dashv \quad \mathsf{cyw}(D)$$

This chapter now first considers the relations between cyclewidth and directed treewidth in more detail. Then, we look at the similarities between cyclewidth and width measures in graphs with perfect matchings, identifying one width measure in particular, the M-perfect matching width, that is parametrically equivalent to cyclewidth. We then explain how a result characterising the graphs with perfect matchings that have small M-perfect matching width can be translated back into the directed setting and yield the characterisation in Theorem 4.0.1. Then, we finish by providing the proof for the characterisation of the graphs with perfect matchings that have small M-perfect matching width.

4.1 Relation to directed treewidth

In this section we describe how cyclewidth relates to directed treewidth.

First, we prove that it is closed under taking butterfly minors. This is a useful property and not the case for directed treewidth (that is, in its standard definition; there is a parametrically equivalent alternative definition that is closed under taking butterfly minors).

Theorem 4.1.1. Let D and D' be digraphs with $D' \preccurlyeq_b D$. Then, $\mathsf{cyw}(D') \leq \mathsf{cyw}(D)$.

Proof. We split the proof into two parts. First, we show that cyclewidth is closed under taking subgraphs. Second, we show that it is closed under butterfly contractions.

Towards the first point, let $D' \subseteq D$ and (T, φ) be a cycle decomposition of D. We obtain a cycle decomposition (T', φ') of D' as follows. To this end, we first obtain the subtree T'' of T by deleting all leaves t with $\varphi(t) \in V(D) \setminus V(D')$. The tree T' is then obtained by trimming T''. Note that $\mathsf{L}(T') \subseteq \mathsf{L}(T)$, so we simply define φ' to be the restriction of φ to $\mathsf{L}(T')$. Now, every family of pairwise disjoint cycles in D' is a family of pairwise disjoint cycles in D. Additionally, for every partition $(X', Y)'$ of D' induced by an edge in T' there exists an edge in T that induces a partition (X, Y) in D with $X' \subseteq X$ and $Y' \subseteq Y$. Thus, the width of (T', φ') is at most the width of (T, φ) .

To prove that cyclewidth is closed under butterfly contractions, let $D' := D/e \to x_e$ for a butterfly contractible edge $e = (u, v) \in E(D)$ and (T, φ) be a cycle decomposition of D. Since e is butterfly contractible, $\partial_D^{\mathsf{out}}(u) = \{e\}$ or $\partial_D^{\mathsf{in}}(v) = \{e\}$.

Consider the case that $\partial_D^{\mathsf{out}}(u) = \{e\}$. This implies that every cycle containing u contains v as well. We build a cycle decomposition (T', φ') from (T, φ) as follows. The tree T' is obtained by deleting the leaf ℓ of T with $\varphi(\ell) = u$ and then trimming. We set $\varphi'(\varphi^{-1}(v)) := x_{u,v}$ and $\varphi'(t) = \varphi(t)$ for all $t \in \mathsf{L}(T')$ with $t \neq \varphi^{-1}(v)$.

Every cut $\partial_{D'}(X')$ induced by an edge in T' corresponds to a cut $\partial_D(X)$ induced by an edge of T. For those cuts with $u, v \in X$, that is, $X' = (X \setminus u, v) \cup \{x_e\}$, or $u, v \in \overline{X}$, that is, $X' = X$, we have $\mathsf{cp}(\partial_{D'}(X')) = \mathsf{cp}(\partial_D(X))$.

So consider a cut $\partial(X)$ induced by some edge f with $v \in X$ and $u \in \overline{X}$, then $X' = (X \setminus \{v\}) \cup \{x_e\}$ and $V(D') \setminus X' = V(D) \setminus (X \cup \{u\})$. Let $\mathcal{C}'$ be a family of disjoint directed cycles witnessing the cycle porosity of $\partial_{D'}(X')$. If no cycle in $\mathcal{C}'$ contains x_e, then it is a family of disjoint directed cycles in D as well, implying $\mathsf{cp}(\partial_{D'}(X')) \leq \mathsf{cp}(\partial_D(X))$. So assume there is a cycle C' in $\mathcal{C}'$ that contains x_e. We can obtain the cycle C in D from it by replacing x_e with u and v, or only v. Due to $x_e \in X'$ and $v \in X$, we obtain that C has as many edges in $\partial_D(X)$ as C' has in $\partial_{D'}(X')$. Thus, again $\mathsf{cp}(\partial_{D'}(X')) \leq \mathsf{cp}(\partial_D(X))$.

Next, we observe that the cyclewidth does not change if we reverse all directions of the edges in the graph. This holds because the cuts stay the same and the cycles are still cycles with reversed direction and crossing the same cuts. Therefore the decomposition stays exactly the same with the same porosities for all cuts. Using this the case that $\partial_D^{\mathsf{in}}(v) = \{e\}$ follows directly from the case $\partial_D^{\mathsf{out}}(u) = \{e\}$.

Therefore, $\mathsf{cyw}(D') \leq \mathsf{cyw}(D)$ holds for every butterfly minor D' of D. □

Cyclewidth and directed treewidth are parametrically equivalent. To show this, we first consider how to obtain a cycle decomposition from a directed tree decomposition. There are two major differences between the trees underlying these two decompositions. One is that cyclewidth demands a cubic tree and the other is that directed tree decompositions have bags at potentially every vertex of the tree, but cycle decompositions map the vertices of the digraph to the leaves only. So we start by manipulating directed tree decompositions to look more like cycle decompositions.

Definition 4.1.2 (Leaf Directed Tree Decomposition)**.** A directed tree decomposition (T, β, γ) of a digraph D is called a *leaf directed tree decomposition* if $\beta(t) = \emptyset$ for all $t \in V(T) \setminus \mathsf{L}(T)$. $\dashv$

Lemma 4.1.3. Let D be a digraph and $k \in \mathbb{N}$. If D has a directed tree decomposition (T, β, γ) of width k, then one can construct a leaf directed tree decomposition (T', β', γ') of the same width from (T, β, γ) in linear time.

Proof. We construct a leaf directed tree decomposition (T', β', γ') of D as follows. Let $N \subseteq V(T)$ be the set of vertices in T with non-empty bags. For every $n \in N$ we introduce a new vertex n' and define T' by $V(T') \coloneqq V(T) \cup \{n' \mid n \in N\}$ and $E(T') \coloneqq E(T) \cup \{(n, n') \mid n \in N\}$. For the bags we define

$$\beta'(t) \coloneqq \begin{cases} \beta(t) = \emptyset, & \text{if } t \in V(T) \setminus N \\ \beta(n), & \text{if } t = n' \text{ for some } n \in N \\ \emptyset, & \text{otherwise, that is, if } t \in N. \end{cases}$$

And for the guards we define

$$\gamma'(e) \coloneqq \begin{cases} \gamma(e), & \text{if } e \in E(T) \\ \beta(n), & \text{if } e = (n, n') \text{ for } n \in N. \end{cases}$$

By construction, the bags of all non-leaf vertices are empty. Note that for all edges $e \in E(T) \cap E(T')$, we have $\beta'(T'_l(e)) = \beta(T_l(e))$ as well as $\gamma'(e) = \gamma(e)$, thus $\gamma'(e)$ guards $\beta'(T'_l(e))$. For all edges $e = (n, n')$ for some $n \in N$, we have $\beta'(T'_l(e)) = \beta(n) = \gamma'(e)$, thus $\gamma'(e)$ guards $\beta'(T'_l(e))$. So (T', β', γ') is indeed a leaf directed tree decomposition.

Now consider the width of (T', β', γ') . For all $t \in V(T) \subseteq V(T')$, we have $\Gamma'(t) = \Gamma(t) \leq k$. For all $t \in V(T') \setminus V(T)$, we have $t = n'$ for some $n \in N$, thus $\Gamma'(t) = \gamma'((n, n')) \cup \beta'(t) = \beta(n) \leq k$. Thus, (T', β', γ') has width at most k. $\square$

Observation 4.1.4. The construction used in the proof of Lemma 4.1.3 only changes the degree for inner vertices with non-empty bags. Additionally, in these cases the degree is increased by exactly one.

We bring the directed tree decomposition in a form with slightly stronger properties next. A directed tree decomposition (T, β, γ) of a digraph D is called *strong* if for all $t \in V(T)$ with children $t_1, \ldots t_\ell$, the bag $\beta(T_{t_i})$ is a strongly connected component of $D - \Gamma(t)$.

Lemma 4.1.5 (Bang-Jensen and Gutin [BG18]). Let D be a digraph and $k \in \mathbb{N}$. If D has a directed tree decomposition of width k, then one can construct a strong directed tree decomposition of D of width at most $3k + 2$ in linear time.

We need this stronger version of a directed tree decomposition, because we consider the subtrees with roots in the out-neighbourhood of a vertex in the decomposition in a topological ordering.

Observation 4.1.6. Let D be a digraph, (T, β, γ) a strong directed tree decomposition, and $t \in V(T)$ with children $t_1, \ldots t_\ell$. Then, as $\beta(T_{t_i})$ is a strongly connected component of $D - \Gamma(t)$ for all $i \in \{1, \ldots, \ell\}$, there is a topological ordering on these components. This yields an ordering on $t_1, \ldots t_\ell$, and without loss of generality we can assume them to already be given in that order. That is, if $i < j$, then every path from $\beta(T_{t_j})$ to $\beta(T_{t_i})$ intersects $\Gamma(t)$.

This now allows us to approach the difference in degree between directed tree decompositions and cycle decompositions.

Lemma 4.1.7. Let D be a digraph and $k \in \mathbb{N}$. If D has a strong directed tree decomposition (T, β, γ) of width k, then it also has a subcubic directed leaf tree decomposition (T', β', γ') of width k which can be computed from (T, β, γ) in linear time. Additionally, all bags of (T', β', γ') have size at most one.

Proof. We split the proof into three steps. The first step yields a subcubic directed tree decomposition (T'', β'', γ'') , which the second step turns into a subcubic directed leaf tree decomposition $(T''', \beta''', \gamma''')$ using Lemma 4.1.3. The third step then takes care of all bags having size at most one.

Step 1: cubification. We traverse the tree T top-down. Thus, whenever reaching a vertex, we can assume that the subtree rooted at it is yet unchanged. Let $t \in V(T)$ be the next encountered vertex. If t has out-degree at most two, then we leave t as it is in this step. So, assume t has out-degree $d > 2$. Let $c_1, \ldots, c_d$ be the children of t. As the subtrees $T_{c_1}, \ldots, T_{c_d}$ and their guards are yet untouched by our construction and (T, β, γ) is a strong directed tree decomposition, we can assume that the children are given in a topological order of their subtrees, as seen in Observation 4.1.6. Removing all outgoing edges of t splits T into several subtrees: T_r containing the root and t as a leaf, and T_{c_i} for every child of t. Next, define $t_0 := t$ and add new vertices $t_1, \ldots, t_{d-1}$ as follows. We add the edges (t_i, t_{i+1}) , for all $0 \leq i < d-1$, with guard $\gamma'((t_i, t_{i+1})) = \Gamma(t)$. Finally, we re-attach the subtrees by adding the edges (t_i, c_i) for all $1 \leq i < d-1$, and the edge (t_{d-1}, c_d) . The in-edge to c_i obtains the guard

$\gamma((t, c_i))$. All new vertices have empty bags and out-degree two. The vertex t, whose bag remains unchanged, has in- and out-degree of one, which we make use of next.

For all the subtrees T_{c_i} with $1 \le i \le d$, every walk starting and ending in $\beta(T_{c_i})$ intersects $\gamma((t, c_i))$ and thus the guard on the incoming edge to c_i. However, we also introduce new subtrees which contain only a subset of the child-subtrees of t. Suppose there was a subtree T_{t_i} with in-edge e and a walk W starting and ending in $\beta(T_{t_i})$, containing vertices outside of $\beta(T_{t_i})$ and not intersecting $\gamma(e)$. If W contains a vertex of $T - T_t$, then we directly obtain a contradiction to (T, β, γ) being a proper directed tree decomposition. Thus, W contains a vertex from $\beta(T_{t_j})$ for some $j < i$. But, then W contains a subwalk W' from $\beta(T_{t_j})$ to $\beta(T_{t_i})$ which intersects $\gamma'(e) = \Gamma(t)$, a contradiction to the topological ordering.

We can continue this construction within the subtrees of the children of t and in subtrees not containing t. As we traverse the tree top-down, we can in each step use the properties of the strong directed tree decomposition of the subtrees of the children although we break this property during our construction. So, at this point we have obtained a subcubic directed tree decomposition (T'', β'', γ'') of D of width k.

Step 2: leaf decomposition. By Lemma 4.1.3 we can turn (T'', β'', γ'') into a leaf directed tree decomposition $(T''', \beta''', \gamma''')$ of width k. Due to Observation 4.1.4 this only increases the degree of inner vertices with non-empty bags. Recall that we established these to have in- and out-degree of one, thus the result is a subcubic leaf directed tree decomposition of width k.

Step 3: leaves. Next, we achieve that the leaves have bags of size one. We use a similar approach to before: replacing the current leaves by paths. Let $t \in \mathsf{L}(T''')$ with bag $\beta'''(t) = \{v_1, \ldots, v_\ell\}$. For each of these vertices v_i we remove them from the bag of t, i.e. $\beta'(t) := \emptyset$, and create a new vertex c_i with $\beta'(c_i) := \{v_i\}$. We also introduce new vertices $t_1, \ldots, t_{\ell-1}$ with empty bags. Now, we add edges (t, t_1) and (t_i, t_{i+1}) for $1 \le i < \ell - 1$. For each such edge e we set $\gamma'(e) := \Gamma(t)$. Then, we attach the vertices with non-empty bags as leaves of T' by adding the edges (t_i, c_i) for $1 \le i < \ell - 1$ and the edge $(t_{\ell-1}, c_\ell)$. Let e_i be the in-edge to c_i, then we define $\gamma'(e_i) := \beta'(c_i)$. By the same arguments as in the first step, the obtained decomposition still has all the properties of a leaf directed tree decomposition.

So, after all three steps we obtain a leaf directed decomposition (T', β', γ') of width k, in which all leaves have a bag of size one. $\square$

These manipulations allow us to prove an upper bound on the cyclewidth in terms of directed treewidth.

Theorem 4.1.8. For every digraph D holds

$$\mathsf{cyw}(D) \leq 2\,(3\mathsf{dtw}(D) + 2)\,.$$

Proof. Let $k := \mathsf{dtw}(D)$. Then, there is a directed tree decomposition of width k of D. By Lemma 4.1.5, there is a strong directed tree decomposition of D of width at most $3k+2$. Using Lemma 4.1.7, we can obtain a subcubic leaf directed tree decomposition (T, β, γ) of D of width $3k+2$ from this in which every leaf has a bag of size one.

We use the underlying undirected graph of T and the function β, that is, $(\mathfrak{G}(T), \beta)$, to obtain a cycle decomposition. By construction, we already know that β yields a bijection between $\mathsf{L}(\mathfrak{G}(T))$ and $V(D)$.

First, we consider the cycle porosity of the cuts induced by the edges in $\mathfrak{G}(T)$. Let $e \in E(\mathfrak{G}(T))$, and let X_1 and X_2 be the two shores of $\partial(e)$, where X_1 corresponds to the subtree not containing the root of T. Consider a maximal family of pairwise disjoint cycles $\mathcal{C}$ and let $Y_1 \subseteq V(\mathcal{C}) \cap X_1$ be the set of vertices of the cycles in $\mathcal{C}$ incident to an edge of $\partial(e) \cap E(\mathcal{C})$. Let $\mathcal{W}$ be the set of nontrivial walks starting in some $v \in Y_1$ and going along the cycle in $\mathcal{C}$ that contains v through vertices of X_2 until reaching a vertex w in X_1 again. That is, $\mathcal{W}$ is the collection of all directed walks $W_{v,w}$ from a vertex $v \in Y_1$ to a vertex $w \in Y_1$ such that

1. $W_{v,w}$ is a subwalk of some cycle in $\mathcal{C}$, and
2. $\emptyset \neq V(W_{v,w}) \setminus \{v, w\} \subseteq X_2$.

Note that the walks in $\mathcal{W}$ are not necessarily vertex disjoint as the paths may share common endpoints in Y_1. If two walks from $\mathcal{W}$ intersect, then there is a cycle in $\mathcal{C}$ that contains both walks, thus they intersect in at most two vertices. Also, every walk in $\mathcal{W}$ contains exactly two edges from $\partial(e) \cap E(\mathcal{C})$. As $\gamma(e)$ guards X_1, it must contain a vertex of every walk in $\mathcal{W}$. By our observations above, every vertex of $\gamma(e)$ can hit at most two walks in $\mathcal{W}$. Thus, $\mathsf{cp}(\partial(e)) = 2|\mathcal{W}| \leq 4|\gamma(e)|$.

What is left to be taken care of is that $(\mathfrak{G}(T), \beta)$ is not necessarily cubic, but might still contain vertices of degree two. Note that those vertices are adjacent to two edges inducing the same cut. We obtain T' from $\mathfrak{G}(T)$ by contracting one of these edges for each such vertex. As $\mathsf{L}(T') = \mathsf{L}(\mathfrak{G}(T))$, β is still well-defined and (T', β) is a cycle decomposition. Every cut has cycle porosity at most $4(3k+2)$, thus the width of this decomposition is $2(3k+2)$. $\square$

There is also a lower bound on the cyclewidth in terms of directed treewidth. A straightforward way to obtain such a bound it is to show that the cylindrical grid has high cyclewidth. We call an edge set $E \subseteq E(D)$ in a digraph D a *balanced cut* if no strongly connected component of $D - E$ contains more than $\frac{2}{3}|V(D)|$ vertices.

Lemma 4.1.9. The cylindrical grid of order k has cyclewidth at least $k/3$.

Proof. Let $\mathsf{D}_k^{\circlearrowright}$ be the cylindrical grid of order k and let (T, φ) be an optimal cycle decomposition of $\mathsf{D}_k^{\circlearrowright}$.

First we show that T contains an edge inducing a balanced cut in $\mathsf{D}_k^{\circlearrowright}$. To this end, direct every edge e of T such that it points in direction of the subtree $T' \in T - e$ such that $|\varphi(T')| > |\overline{\varphi(T')}|$. If there is an edge that cannot be oriented this way, that is, both sides contain exactly half the vertices, we found the balanced cut. Otherwise, every edge can be directed and no two edges with a common vertex can point away from each other. Additionally, all leaf edges point away from the leaf. So there has to be a unique inner vertex v with only ingoing edges. Let e_1, e_2 and e_3 be the edges incident to v, and T_i the subtree of $T - e_i$ not containing v. Any two of these subtrees contain together at least half of the leaves of T. If there is a subtree with less than one third of the leaves of T, then the other two edges induce balanced cuts. Otherwise, all three subtrees contain exactly one third of the leaves of T and all three edges induce balanced cuts in $\mathsf{D}_k^{\circlearrowright}$.

Thus, we can choose $e \in E(T)$ such that $\partial(e)$ is a balanced cut in $\mathsf{D}_k^{\circlearrowright}$. We consider two cases: either each shore of $\partial(e)$ contains one of the concentric cycles of $\mathsf{D}_k^{\circlearrowright}$ or one of its shores does not contain any of the concentric cycles of the grid completely.

In the case where each shore of $\partial(e)$ contains one of the concentric cycles of $\mathsf{D}_k^{\circlearrowright}$, we construct a cycle C that contains $2k$ edges of $\partial(e)$. Let C_x and C_y be two concentric cycles of $\mathsf{D}_k^{\circlearrowright}$ completely contained in different shores of $\partial(e)$ and assume without loss of generality that $x < y$. Let X_i be the shore containing C_x and X_o be the shore containing C_y.

We construct C by starting on a vertex of C_y where it intersects an in-path P^{i} of $\mathsf{D}_k^{\circlearrowright}$. Then, we walk along P^{i} until meeting C_x and walk along C_x for one edge. There we meet an out-path P^{o} and walk along it until we intersect C_y again. Now, we walk along C_y for one edge where it meets the next in-path. We repeat this until closing C. This way we use all in- and out-paths of $\mathsf{D}_k^{\circlearrowright}$. Thus, there are $2k$ subpaths of C crossing $\partial(e)$ at least once. Therefore C contains at least $2k$ edges of $\partial(e)$ as desired. This implies that $\mathsf{cp}(\partial(e)) \geq 2k > \frac{2}{3}k$, and therefore, (T, φ) has width more than $k/3$.

In the case that there is a shore of $\partial(e)$ that does not contain any concentric cycle of $\mathsf{D}_k^{\circlearrowright}$, the other shore of $\partial(e)$ contains at most two third of the concentric cycles of $\mathsf{D}_k^{\circlearrowright}$, as $\partial(e)$ is balanced. Therefore, the remaining at least $\frac{k}{3}$ cycles of $\mathsf{D}_k^{\circlearrowright}$ cross $\partial(e)$. Each of them meets $\partial(e)$ in at least two edges and thus they build a family of disjoint cycles witnessing that $\mathsf{cp}(\partial(e)) \geq \frac{2}{3}k$, so (T, φ) has width at least $k/3$.

So, we obtain $\mathsf{cyw}(\mathsf{D}_k^{\circlearrowright}) \geq \frac{k}{3}$. $\square$

Together with Theorem 4.1.1 this implies the following corollary.

Corollary 4.1.10. If a digraph D has the cylindrical grid of order k as a butterfly minor, then its cyclewidth is at least $k/3$.

And thus we obtain the following.

Theorem 4.1.11. A class $\mathcal{D}$ of digraphs is a class of bounded directed treewidth if and only if it is a class of bounded cyclewidth.

Proof. Let $\mathcal{D}$ be a class of digraphs. Suppose $\mathcal{D}$ has unbounded directed treewidth, then for each $n \in \mathbb{N}$ there is a digraph $D'_n \in \mathcal{D}$ such that $\mathsf{dtw}(D') \geq n$. By Theorem 2.2.2, we can conclude that for every $n \in \mathbb{N}$ there is a digraph $D_n \in \mathcal{D}$ that contains the cylindrical grid of order n as a butterfly minor. Therefore, $\mathsf{cyw}(D_n) \geq \frac{n}{3}$ by Lemma 4.1.9 and Theorem 4.1.1. Thus, $\mathcal{D}$ has also unbounded cyclewidth. Vice versa, assume $\mathcal{D}$ is of bounded directed treewidth. Then, it also is of bounded cyclewidth due to Theorem 4.1.8. $\square$

Giannopoulou, Kreutzer and Wiederrecht obtained a polynomial bound using techniques from matching theory.

Theorem 4.1.12 (Giannopoulou, Kreutzer and Wiederrecht [GKW21])**.** There is a polynomial function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathsf{dtw}(D) \leq f(\mathsf{cyw}(D))$ for all digraphs D.

We conjecture that the relation between cyclewidth and directed treewidth is linear with an even tighter lower bound.

Conjecture 4.1.13. [1] For all digraphs D holds $\mathsf{cyw}(D) \pm 1 \leq \mathsf{dtw}(D) \leq c \cdot \mathsf{cyw}(D)$ for come constant c.

The relation between cyclewidth and directed treewidth established by Theorems 4.1.11 and 4.1.12 allows us to immediately defer a grid theorem for cyclewidth.

Corollary 4.1.14. There is a function $f : \mathbb{N} \to \mathbb{N}$ such that every digraph D either satisfies $\mathsf{cyw}(D) \leq f(k)$, or contains the cylindrical grid of order k as a butterfly minor.

[1]There is a, so far unpublished, proof for this by Nathan Bowler, Ann-Kathrin Elm, Florian Gut, Raphael Jacobs, Marcel Koloschin, and Florian Reich.

4.2 Digraphs and matchings

Before we compare cyclewidth to width measures on graphs with perfect matchings, we explain how digraphs relate to graphs with perfect matchings at all.

There is a close connection between strongly connected digraphs and bipartite graphs with perfect matchings. Every bipartite graph with a perfect matching corresponds to a class of digraphs. This is because we can obtain a digraph for every perfect matching in the bipartite graph, by contracting the matching edges and directing every remaining edge from the vertex it has in the colour class V_1 to the vertex it has in the colour class V_0. The following definition describes this formally.

Definition 4.2.1. Let B be a bipartite graph and let $M \in \mathcal{M}(B)$ be a perfect matching of B. The *M-direction* $\mathcal{D}(G, M)$ of G is defined by $\mathcal{D}(G, M)$

$$\begin{aligned} V(\mathcal{D}(G, M)) &:= M \text{ and} \\ E(\mathcal{D}(G, M)) &:= \{(e, f) \mid e \neq f, \\ &\qquad uv \in E(G) \text{ for } u \in e \cap V_1 \text{ and } v \in f \cap E(V_0)\}. \end{aligned}$$ ⊣

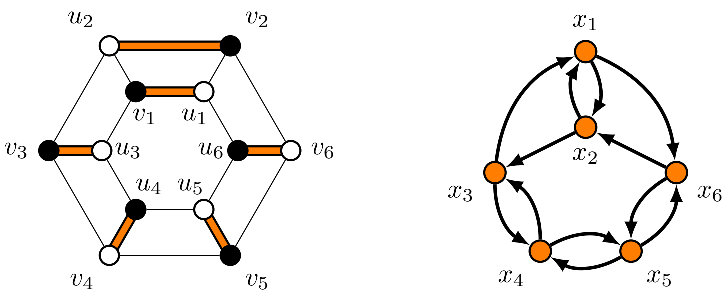

Figure 4.1: On the left we see a bipartite graph with a fixed matching M and on the right the corresponding M-direction.

The direction of the edges is obtained by directing from the vertices in V_1, which we depict in figures as black vertices, to the vertices in V_0, which we depict as white vertices. See Figure 4.1 for an illustration of this definition.

The M-directions of a bipartite matching covered graph B inherit some properties of B.

Theorem 4.2.2 (McCuaig [McC00])**.** Let B and H be bipartite matching covered graphs. Then, H is a matching minor of B if and only if there exist a perfect matchings M in B and a perfect matching M' in H such that $\mathcal{D}(H, M')$ is a butterfly minor of $\mathcal{D}(D, M)$.

4.2.1 Perfect matching width

While graphs with perfect matchings are also undirected graphs, and therefore, we can use all tools and measures for undirected graphs to work with them, there exist width measures that are specifically designed for the matching setting and describe the matchings structure of the graph.

The perfect matching width is a width measure using perfect matchings in a graph. It was originally introduced by Norin [Nor05] in the context of Pfaffian orientations. It makes use of the matching porosity of edge cuts in the graph, which is similar to the cycle porosity we use for cyclewidth. Let G be a graph and $X \subseteq V(G)$. We define the *matching porosity* of $\partial(X)$ as

$\mathsf{mp}(\partial(X))$

$$\mathsf{mp}(\partial(X)) \coloneqq \max_{M \in \mathcal{M}(G)} |M \cap \partial(X)| \,.$$

Definition 4.2.3 (Perfect Matching Width)**.** Let G be graph. A *perfect matching decomposition* of G is a tuple (T, δ) where T is a cubic tree and $\delta\colon \mathsf{L}(T) \to V(G)$ a bijection. The *width* of (T, δ) is given by $\max_{e \in E(T)} \mathsf{mp}(\partial(e))$ and the *perfect matching width* of G is then defined as

$\mathsf{pmw}(G)$

$$\mathsf{pmw}(G) \coloneqq \min_{\substack{(T,\delta) \text{ perfect matching} \\ \text{decomposition of } G}} \max_{e \in E(T)} \mathsf{mp}(\partial(e)) \,. \qquad \dashv$$

When considering the M-directions of a graph G with $M \in \mathcal{M}(G)$, every cycle decomposition (T, φ) of $\mathcal{D}(G, M)$ can be interpreted as a decomposition of G where φ is a bijection between $\mathsf{L}(T)$ and the edges in M. Then, every edge in T induces a bipartition of $V(G)$ into M-conformal sets. The next definition captures this observation in terms of specific perfect matching decompositions.

Definition 4.2.4 (M-Perfect Matching Width)**.** Let G be a graph with a perfect matching $M \in \mathcal{M}(G)$. A perfect matching decomposition (T, δ) of G is an *M-perfect matching decomposition* of G if for every inner edge e with $(T_1, T_2) = T \ltimes e$ we have $\delta(\mathsf{L}(T_1))$ and $\delta(\mathsf{L}(T_2))$ are M-conformal. Define S as the set of all M-perfect matching decompositions (T, δ) of G. The *M-perfect matching width*, pmw_M, is defined as

$$\mathsf{pmw}_M(G) := \min_{(T,\delta)\in S} \max_{e\in E(T)} \mathsf{mp}(\partial(e))\,. \qquad \dashv$$

The relation of these two width measures for graphs with perfect matchings is described by the following theorem, which we prove later within this chapter.

Theorem 4.2.5. Let G be a graph with a perfect matching M. Then

$$\mathsf{pmw}(G) \leq \mathsf{pmw}_M(G) \leq 2\mathsf{pmw}(G)\,.$$

The M-alternating cycles of a bipartite graph B with perfect matching M correspond to the directed cycles of $\mathcal{D}(B, M)$. Every M-alternating cycle becomes a directed cycle in $\mathcal{D}(B, M)$, as the matching edges are contracted into vertices and the edges directed from V_1 to V_0, an example for this is given in the upper part of Figure 4.2. This yields the following relation between M-perfect matching width and cyclewidth.

Lemma 4.2.6. Let B be a bipartite and matching covered graph and $M \in \mathcal{M}(B)$. Then, $\mathsf{pmw}_M(B) = 2\mathsf{cyw}(\mathcal{D}(B, M))$.

Proof. We prove the statement by proving both inequalities. For readability define $D := \mathcal{D}(B, M)$. We refer the reader to Figure 4.2 for an illustration of the following arguments by an example.

We first prove that $\mathsf{pmw}_M(G) \geq 2\mathsf{cyw}(D)$. To this end assume $\mathsf{pmw}_M(G) = k$ for some $k \in \mathbb{N}$. Then, there is a perfect matching decomposition (T, δ) of width k such that all shores of the cuts induced by inner edges of T are M-conformal. In particular, the leaves of T mapped by δ to two vertices matched by M share a neighbour. We construct a cycle decomposition (T', φ) of D as follows. We start by defining $T' := T - \mathsf{L}(T)$. Note that $V(T') \subseteq V(T)$. Recall that matching edges become vertices in D. For $xy \in M$ let t_{xy} be the common neighbour of $\varphi^{-1}(x)$ and $\varphi^{-1}(y)$. We define $\varphi(t_{xy}) := xy$.

Suppose towards a contradiction that the obtained decomposition (T', φ) has an edge $e \in T'$ that induces a cut $\partial_D(X)$ of cycle porosity more than $2k$ and let $\mathcal{C}$ be a family of directed cycles in D witnessing this. The cycles from $\mathcal{C}$ correspond to M-alternating cycles $\mathcal{C}'$ in G that also have more than $2k$ edges in the cut $\partial_B(X')$ induced by $e \in T$. Since X' is M-conformal, none of the cut edges is from M, that is, $M \cap (E(\mathcal{C}') \cap \partial_B(X')) = \emptyset$. Let M' be the matching we obtain by switching M along all the cycles in $\mathcal{C}'$, that is, $M' := (M \setminus E(\mathcal{C}')) \cup (E(\mathcal{C}') \setminus M)\,.$ Now, M' has

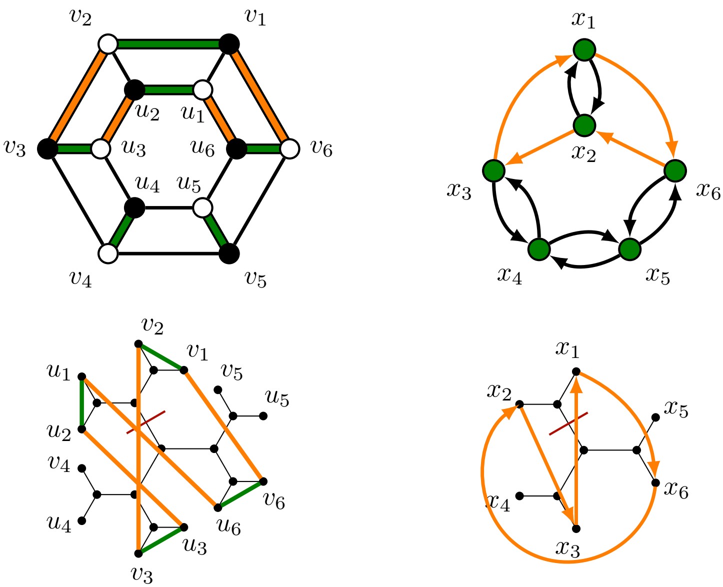

Figure 4.2: We reuse the example from Figure 4.1. The upper part of the figure shows how the alternating cycle $v_6u_6u_1u_2u_3v_3v_2v_1$ in B corresponds to the directed cycle $x_6x_2x_3x_1$ in $\mathcal{D}(B, M)$. The lower part of the figures shows these two cycles transferred to the corresponding decomposition. We consider the cut induced by the edge crossed by the red line and we can see that every orange edge in the alternating cycle in B remains a directed edge in the cut in $\mathcal{D}(B, M)$.

at least $2k + 1$ edges in $\partial_B(e)$ contradicting that (T, δ) has width k. Therefore (T', φ) is a cycle decomposition of D of width at most k.

Next, we prove that $\mathsf{pmw}_M(G) \leq 2\mathsf{cyw}(D)$. Let $\mathsf{cyw}(D) = k$ for some $k \in \mathbb{N}$. Then, there is a cycle decomposition (T, φ) of D of width k. We construct a perfect matching decomposition (T', δ) of B as follows. Every leaf t in T is mapped to a vertex corresponding to an edge in M, that is, $\varphi(t) = t_{xy}$ with $xy \in M$. We construct T' from T by introducing two new vertices t_x and t_y, making them adjacent to t_{xy}, and defining $\delta(t_x) := x$ and $\delta(t_y) := y$. Formally, $V(T') := V(T) \cup \{t_x, t_y \mid t_{xy} \in \mathsf{L}(T)\}$ and $E(T') := E(T) \cup \{t_{xy}t_x, t_{xy}t_y \mid t_{xy} \in \mathsf{L}(T)\}$. All pairs of vertices that are matched by M have a common neighbour in T', so the shores of the cuts induced by inner edges of T' are M-conformal. Therefore, (T', δ) is indeed a M-perfect matching decomposition.

Suppose towards a contradiction that there is an edge $e \in E(T')$ and a matching M' such that $|M' \cap \partial_B(e)| > 2k$. We consider the subgraph of B only containing edges from M and M'. As such, it only consists of disjoint cycles and independent edges. Because none of the edges in M lie in $\partial_B(e)$, all edges of $M' \cap \partial_B(e)$ lie on M-M'-alternating cycles. Therefore, there is a family of M-conformal cycles with more than $2k$ edges in $\partial_B(e)$. This corresponds to a family of directed cycles in D having more than $2k$ edges in the cut induced by e in D. This yields a contradiction to (T, φ) having width k. Therefore (T', δ) is a perfect matching decomposition of G of width k. $\square$

Corollary 4.2.7. Let B be a bipartite graph and M a perfect matching of B. Then

$$\mathsf{pmw}(B)/2 \leq \mathsf{cyw}(\mathcal{D}(B, M)) \leq \mathsf{pmw}(B)\,.$$

4.3 Matching decompositions

Having established these connections between the matching width measures and cyclewidth we move our focus towards matching theory. In fact, in order to prove Theorem 4.0.1, we can instead prove the following theorem. It characterises bipartite matching covered graphs of low M-perfect matching width by *braces*, which is a concept similar to blocks in undirected graphs and formally defined later on.

Theorem 4.3.1. Let B be a bipartite matching covered graph, then the following statements are equivalent.

(1) $\mathsf{pmw}_M(B) = 2$ for an $M \in \mathcal{M}(B)$,

(2) $\mathsf{pmw}_M(B) = 2$ for all $M \in \mathcal{M}(B)$, and

(3) Every brace of B is either isomorphic to C_4 or to $K_{3,3}$.

Both graphs C_4 and $K_{3,3}$ have only one single digraph as an M-direction. All matchings in C_4 are symmetric and yield the bidirected K_2, also called digon, as an M-direction, see Figure 4.3. For the $K_{3,3}$ the unique directed graph obtained for every matching is the bidirected K_3, see Figure 4.4. Therefore, Theorem 4.3.1 implies our main theorem.

Theorem 4.0.1. Let D be a digraph. Then, D has cyclewidth exactly one if and only if every strongly 2-connected butterfly minor of D is isomorphic to the digon or the bi-directed K_3.

Figure 4.3: The C_4 with the matching M on the left and $\mathcal{D}(C_4, M)$.

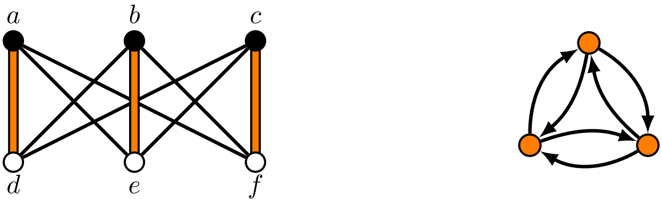

Figure 4.4: The $K_{3,3}$ with a matching M on the left and $\mathcal{D}(K_{3,3}, M)$.

In the remainder of this chapter we now purely consider the matching setting in order to prove Theorem 4.3.1. First, we need to learn a bit more about perfect matching decompositions and their structure.

We start by looking at the structure of the trees that are the base of the decompositions. What properties can they have? How do these relate to properties of the decomposed graph? Second, we consider the concept of tight cut contractions and show that the M-perfect matching width cannot grow under this operation. Finally, we also consider a generalisation of tight cuts and use it to learn about how the vertices of the two colour classes of a bipartite graph can be distributed in a decomposition.

4.3.1 Cubic trees and spines

We establish some observations on the cubic trees that appear as the base structure of perfect matching decompositions.

For a cubic tree T we define the *spine* of T by $\mathsf{spine}(T) := T - \mathsf{L}(T)$. The edges in $E(T) \setminus E(\mathsf{spine}(T))$ are called *trivial*. We say an edge $e \in E(\mathsf{spine}(T))$ is *even*, if the two trees of $T \ltimes e$ each contain an even number of leaves of T, and it is *odd* otherwise.

We observe that the trees occurring in perfect matching decompositions have an even number of leaves. This is because graphs with perfect matchings have an even number of vertices. This implies that in such a tree T a nontrivial edge e is odd if and only if it the two trees of $T \ltimes e$ each contain an odd number of leaves of T. We make a few observations on cubic trees with an even number of leaves.

Lemma 4.3.2. Let T be a cubic tree with $|\mathsf{L}(T)| = \ell$ even. Then, the following statements are true.

(C1) $|V(T)| = 2\ell - 2$,

(C2) $\mathsf{spine}(T)$ has an even number of vertices,

(C3) $\mathsf{spine}(T)$ has an even number of vertices of degree 2, and

(C4) $e \in E(\mathsf{spine}(T))$ is an odd edge of T if and only if the two trees of $T \ltimes e$ each contain an even number of vertices.

Proof. We use induction on the size of T. Split T by removing an inner edge e to obtain $(T_1, T_2) = T \ltimes e$. As T has an even number of leaves either both T_1 and T_2 have an even number of leaves or both have an odd number of leaves.

In case both T_i have an odd number of leaves, for $i \in \{1, 2\}$, we obtain T_i' from T_i by adding an extra leaf adjacent to the endpoint of e being contained in T_i. This makes T_i' again a cubic tree with an even number ℓ_i of leaves. By induction T_i' has $2\ell_i - 2$ vertices. Then, T has $|V(T_1')| - 1 + |V(T_1')| - 1 = 2\ell_1 - 2 - 1 + 2\ell_2 - 2 - 1 = 2(\ell_1 + \ell_2 - 2) - 2 = 2\ell - 2$ vertices. Thus, (C1) holds.

In case both T_i have an even number of leaves, for $i \in \{1, 2\}$, we obtain T_i' from T_i by trimming, that is, contracting one of the two edges in T_i adjacent to the endpoint of e being contained in T_i. This makes T_i' again a cubic tree with an even number ℓ_i of leaves. By induction T_i' has $2\ell_i - 2$ vertices. Then, T has $|V(T_1')| + 1 + |V(T_1')| + 1 = 2\ell_1 - 2 + 1 + 2\ell_2 - 2 + 1 = 2(\ell_1 + \ell_2) - 2 = 2\ell - 2$ vertices. Thus, (C1) holds in this case as well.

(C2) follows directly from (C1) and the number of leaves being even.

Every degree-2-vertex in $\mathsf{spine}(T)$ is adjacent to exactly one leaf of T. All other vertices are either adjacent to exactly two leaves or to no leaf at all. Thus, the other vertices are adjacent to an even number of leaves. So the total number of leaves not sharing a neighbour with another leaf is even as well. Thus, the number of degree-2-vertices is even, yielding (C3).

Let e be an odd edge of T and $(T_1, T_2) \coloneqq T \ltimes e$. For $i \in \{1, 2\}$, obtain T_i' from T_i by adding an extra leaf adjacent to the endpoint of e being contained in T_i. This makes

T_i' again a cubic tree with an even number of leaves. By (C2), this implies that it has an even number of vertices. Thus T_i, having one less vertex, has an odd number of vertices. This implies (C4). $\square$

If T is a cubic tree, then $\mathsf{spine}(T)$ is a subcubic tree. There is a close correspondence between the occurrences of odd edges in T and vertices of degree 2 in its spine.

Lemma 4.3.3. Let T be a cubic tree with an even number of leaves and $v \in V(T)$. Then, the following statements hold.

(E1) If $\deg_{\mathsf{spine}(T)}(v) = 1$, then v is not incident to an odd edge of T.

(E2) If $\deg_{\mathsf{spine}(T)}(v) = 2$, then v is incident to exactly one odd edge of T.

(E3) If $\deg_{\mathsf{spine}(T)}(v) = 3$, then v is either incident to exactly two odd edges of T or with none.

Proof. First, consider some $v \in V(T) \setminus \mathsf{L}(T)$ with $\deg_{\mathsf{spine}(T)}(v) = 1$. Then, as T is cubic, v is adjacent to two leaves of T. Let e be the unique nontrivial edge v is incident to. Then, one of the trees in $T \ltimes e$ has exactly two leaves, and the other has $|V(G)| - 2$ leaves, which implies that e is not odd. So, v is incident to two trivial edges and one even edge, which implies (E1).

Second, consider a vertex $v \in V(T) \setminus \mathsf{L}(T)$ with $\deg_{\mathsf{spine}(T)}(v) = 2$. So, there are two edges $e_1, e_2 \in E(\mathsf{spine}(T))$ that are incident to v. Thus, the third edge e_3 incident to v is trivial. So, the partitions of the leaves induced by e_1 and e_2 differ by exactly one, and therefore, e_1 is even if and only if e_2 is odd. This implies that v is adjacent to exactly one trivial, exactly one even and exactly one odd edge, which implies (E2).

Finally, consider a vertex $v \in V(T) \setminus \mathsf{L}(T)$ with $\deg_{\mathsf{spine}(T)}(v) = 3$. So, no neighbour of v is a leaf. Let e_1, e_2, e_3 be the three edges incident to v. Removing v and e_1, e_2, e_3 splits the remaining tree into three subtrees which contain all leaves of T. As there is an even number of leaves, two of these trees contain an odd number of them and one an even number. Therefore, v is incident to two odd edges and one even edge, again implying (E3). $\square$

Corollary 4.3.4. Let T be a cubic tree with an even number of leaves. Then, $\mathsf{spine}(T)$ is cubic if and only if T has no odd edges.

Additionally, the odd edges of a cubic tree T induce a subforest of $\mathsf{spine}(T)$, the leaves of which are exactly the degree-2-vertices of $\mathsf{spine}(T)$. This subforest is in fact a collection of paths, because no vertex of $\mathsf{spine}(T)$ can be incident to three odd edges, see Figure 4.5 for an example.

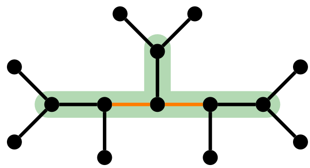

Figure 4.5: An example for a cubic tree T with its spine and its odd edges.

Corollary 4.3.5. Let T be a cubic tree with an even number of leaves and $E_O \subseteq E(T)$ the set of odd edges of T. Then, $T[E_O]$ is a collection of pairwise disjoint paths. Moreover, the set of endpoints of these paths is exactly the set of degree-2-vertices in $\mathsf{spine}(T)$.

The following lemma shows that the odd edges of the tree underlying a perfect matching decomposition influence the width of the decomposition.

Lemma 4.3.6. Let G be a graph with a perfect matching and $X \subseteq V(G)$. Then, $\mathsf{mp}(\partial(X))$ is odd if and only if $|X|$ is odd.

Proof. Let $M \in \mathcal{M}(G)$ be a perfect matching of G that maximises $\partial(X)$ and define $k := |M \cap \partial(X)|$. Then, the graph $G[X] - V(\partial(X) \cap M)$ has a perfect matching, and therefore an even number of vertices. Let $n := |V(G[X] - V(\partial(X) \cap M))|$. So in total $|X| = n + k$.

This implies that $|X| \equiv k \pmod 2$. And thus, $\mathsf{mp}(\partial(X)) = k$ is odd if and only if $|X|$ is odd. □

Lemma 4.3.6 implies that the tree of every perfect matching decomposition (T, δ) of odd width contains an odd edge. Additionally, from the proof of Lemma 4.3.6, we obtain the following statement, which informally says that changing the position of a single vertex within a decomposition barely changes its width.

Corollary 4.3.7. Let G be a graph with a perfect matching, $X \subseteq V(G)$ and $x \in V(G) \setminus X$. Then

$$\mathsf{mp}(\partial(X)) - 1 \leq \mathsf{mp}(\partial(X \cup \{x\})) \leq \mathsf{mp}(\partial(X)) + 1.$$

As parity is of high importance in the study of perfect matchings, it can be useful to know and manipulate the occurrence of odd edges in the tree of a perfect matching decomposition. Moreover, these insights now allow us to prove the relation between M-perfect matching width and perfect matching width claimed previously.

Theorem 4.2.5. Let G be a graph with a perfect matching M. Then

$$\mathsf{pmw}(G) \leq \mathsf{pmw}_M(G) \leq 2\mathsf{pmw}(G)\,.$$

Proof. Let G be a graph with a perfect matching M. The inequality $\mathsf{pmw}(G) \leq \mathsf{pmw}_M(G)$ holds by definition.

So we have to prove that $\mathsf{pmw}_M(G) \leq 2\mathsf{pmw}(G)$. Let (T, δ) be an optimal perfect matching decomposition of G. We choose a set X that contains exactly one endpoint of every edge in M. For every $x \in X$ let $M(x)$ be the vertex x is matched to by M, that is, $\{x, M(x)\} \in M$. We consider the subset $X' \subseteq X$ that contains all vertices x such that $\delta^{-1}(x)$ and $\delta^{-1}(M(x))$ do not have a common neighbour in T.

Next, we build an M-perfect matching decomposition (T', δ') of G. First, we define T'' by

$$\begin{aligned} V(T'') &\coloneqq V(T) \setminus \{\delta^{-1}(M(x)) \mid x \in X'\} \cup \{t_x, t'_x \mid x \in X'\}, \text{ and} \\ E(T'') &\coloneqq \{e \in E(T) \mid e \subseteq V(T')\} \cup \\ &\qquad \{\delta^{-1}(x)\, t_x, \delta^{-1}(x)\, t'_x \mid x \in X'\}, \end{aligned}$$

where t_x, t'_x are newly introduced vertices. Then, we obtain T' from T'' by trimming, thus $\mathsf{L}(T') = \mathsf{L}(T'')$. For δ' we define $\delta'(t_x) \coloneqq x$ and $\delta'(t'_x) \coloneqq M(x)$ for all $x \in X'$, this covers all leaves in $\mathsf{L}(T') \setminus V(T)$. For all leaves $t \in \mathsf{L}(T') \cap \mathsf{L}(T)$ we define $\delta'(t) = \delta(t)$.

Due to the trimming, T' is cubic and by definition every vertex of G is mapped to a leaf of T' by δ'. The construction ensures that for all edges $e = \{u, v\} \in M$ the vertices $\delta'^{-1}(u)$ and $\delta'^{-1}(v)$ have a common neighbour in T', thus (T', δ') is an M-perfect matching decomposition.

Next, we consider the width of (T', δ') . The only new inner edges we construct in T' induce cuts with one shore being of size two and the two vertices in it being matched by M. Thus the matching porosity of these cuts is 2 and the shores are both M-conformal. Now consider one of the remaining inner edges e' of T', it corresponds to an inner edge e in T which induces an edge cut $\partial(X_e)$. The matching M has at most $\mathsf{pmw}(G)$ many edges in this cut, so we changed the position of at most $\mathsf{pmw}(G)$ vertices with respect to this cut. Therefore, by Corollary 4.3.7, the cut $\partial(X_{e'})$ induced by e' is at most $2\mathsf{pmw}(G)$. Therefore the width of the obtained M-perfect matching decomposition is at most $2\mathsf{pmw}(G)$.

This yields the wanted inequality $\mathsf{pmw}_M(G) \leq 2\mathsf{pmw}(G)$. $\square$

4.3.2 Tight cut contractions

Edge cuts in graphs with perfect matchings are central to working with perfect matching decompositions. Here we consider specifically the cuts that contain exactly one edge from each perfect matching in the graph. We observe that these have specific and useful properties and yield a specific kind of matching minor.

Let $\partial_G(X)$ be an edge cut in a graph G that has a perfect matching. The edge cut is called *tight* if $\mathsf{mp}(\partial_G(X)) = 1$, we also say that X *induces* a tight cut in this case. It is called *trivial* if $|X| = 1$ or $|\overline{X}| = 1$. Note that if $\mathsf{mp}(\partial_G(X)) = 1$, then $|X|$ is odd, and thus, in particular, $|\partial_G(X) \cap M| = 1$ for all $M \in \mathcal{M}(G)$.

Definition 4.3.8 (Tight cut contractions)**.** If Z is the shore of a tight cut, we call the operation of identifying Z into a single vertex v_Z and deleting all resulting loops and parallel edges a *tight cut contraction*, we write $G/(Z \to v_Z)$ and call the obtained graph G_Z. Note that G having a perfect matching implies that G_Z has a perfect matching as well. If M is a perfect matching of G, then let e be the unique edge in $\partial_G(Z) \cap M$ with endpoints $u_Z \in Z$ and $u_{\overline{Z}} \in \overline{Z}$. The matching $M|_{G_Z} \coloneqq (M \cap E(G_Z)) \cup \{u_{\overline{Z}} v_Z\}$ is a perfect matching of G_Z. $\dashv$

G_Z

$M|_{G_Z}$

The matching covered graphs that do not contain any tight cuts yield the building blocks of matching covered graphs. They are the matching analogue of *blocks* (maximal 2-connected components) in general graph theory. A matching covered graph without nontrivial tight cuts is called a *brace* if it is bipartite and a *brick* otherwise. By repeatedly choosing tight cuts and performing tight cut contractions in a graph one obtains a list of bricks and braces. This is called a *tight cut decomposition procedure* and any possible choice of a laminar family of tight cuts used for it is called a *tight cut decomposition* of the graph. A well-known theorem by Lovász states that the list of bricks and braces obtained by a tight cut decomposition procedure is always the same, independent of the choice of the laminar family of tight cuts made. So, every matching covered graph G has a uniquely defined list of bricks and braces, which we also call the *bricks and braces of* G.

Theorem 4.3.9 (Lovász [Lov87])**.** Any two tight cut decomposition procedures of a matching covered graph G yield the same list of bricks and braces.

Lucchesi et al. [LdM15] proved that tight cut contractions and therefore braces are special cases of matching minors.

Lemma 4.3.10 (Lucchesi et al. [LdM15])**.** Let B be a bipartite matching covered graph and $\partial(Z)$ a nontrivial tight cut in B. Then, the tight cut contractions of Z and $\overline{Z}$ are matching minors of B.

Corollary 4.3.11 (Lucchesi et al. [LdM15])**.** If H is a brace of a bipartite matching covered graph B, then H is a matching minor of B.

We show that the M-perfect matching width of a tight cut contraction cannot be larger than the M-perfect matching width of the original graph. The main issue is where in the reduced decomposition the leaf is added that is mapped to the contraction vertex without increasing the width of the decomposition. The advantage of working with M-perfect matching decompositions is that, because the two shores are M-conformal, the cuts given by edges are even. The following lemma shows that this allows us to determine a position for the leaf mapped to the contraction vertex obtaining a decomposition of width at most the width of the original decomposition.

Lemma 4.3.12. Let G be a graph with a perfect matching and $\partial(Z)$ a nontrivial tight cut of G as well as G_Z the tight cut contraction obtained by contracting Z into the vertex v_Z. For every $X \subseteq V(G)$ of even size where $|X \cap Z|$ is odd we have $\mathsf{mp}(\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})) \leq \mathsf{mp}(\partial_G(X))$.

Proof. Let $X \subseteq V(G)$ be a set of even size with $|X \cap Z|$. Suppose towards a contradiction $\mathsf{mp}(\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})) > \mathsf{mp}(\partial_G(X))$. Consider a perfect matching $M \in \mathcal{M}(G_Z)$ that witnesses the matching porosity of $\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})$ and let $e \in M$ be the edge covering v_Z. Let M^+ be a perfect matching in G with $M^+|_{G_Z} = M$. All edges of $M \setminus \{e\}$ that lie in $\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})$ also lie in $\partial_G(X) \cap M^+$. Thus, we obtain $\mathsf{mp}(\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})) = \mathsf{mp}(\partial_G(X)) + 1$.

By Lemma 4.3.6, $\partial_G(X)$ has even matching porosity. Therefore, the porosity of $\partial_{G_Z}((X \setminus Z) \cup \{v_Z\})$ is odd. But due to $|X \cap Z|$ being odd, $|\,(X \cap Z) \cup \{v_Z\}|$ must be even, which yields a contradiction to Lemma 4.3.6. $\square$

The following concept formalises how Lemma 4.3.12 provides a position for the leaf mapped to the contraction vertex of a tight cut contraction.

Definition 4.3.13 (Z-orientations)**.** Let G be a graph with a perfect matching and $\partial(Z)$ a nontrivial tight cut of G. For every perfect matching decomposition (T, δ)
$\overrightarrow{T_Z}$ of G we define the *Z-orientation* $\overrightarrow{T_Z}$ of T as the following orientation of the edges of T. For every edge $t_1t_2 \in E(\mathsf{spine}(T))$, we define $(t_1, t_2) \in E(\overrightarrow{T_Z})$ if and only

if $|\delta(T_{t_2}) \cap Z|$ is odd. For every edge $\ell t \in E(T)$, where ℓ is a leaf of T, we define $(\ell, t) \in E(\overrightarrow{T_Z})$.

Note that $|Z|$ is odd for every tight cut $\partial(Z)$. Thus, the Z-orientation of the edge $t_1 t_2$ is well defined, see Figure 4.6 for an example. A vertex $t \in V(\overrightarrow{T_Z})$ with at least two outgoing edges in $\overrightarrow{T_Z}$ is called an *inconsistency*. ⊣

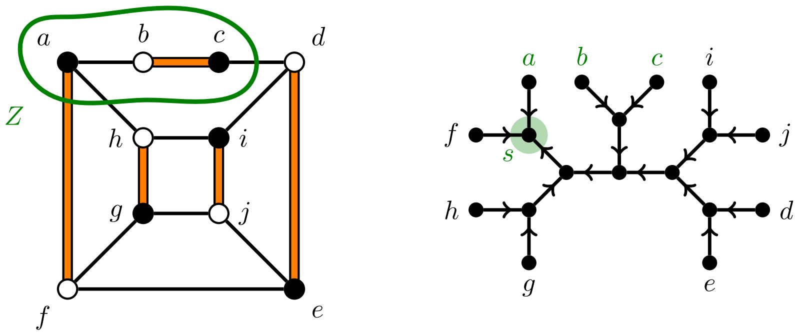

Figure 4.6: A matching covered graph G with a nontrivial tight cut $\partial(Z)$, a perfect matching $M \in \mathcal{M}(G)$, and an M-perfect matching decomposition (T, δ) of width four. The arrows in T are the edges forming the Z-orientation of T, note that it is free of inconsistencies and has a unique sink s.

The idea behind this definition is that, with a Z-orientation that does not have any inconsistencies, the orientation of the edges guides you to the position where the contraction vertex can be placed. For general perfect matching decompositions this is not necessarily the case, but we show that it works for M-perfect matching decompositions.

Lemma 4.3.14. Let G be a graph with a perfect matching, $\partial(Z)$ a nontrivial tight cut in G, $M \in \mathcal{M}(G)$, and (T, δ) an M-perfect matching decomposition of G. Then, $\overrightarrow{T_Z}$ is free of inconsistencies and has a unique sink that is adjacent to two leaves.

Proof. As $\partial(Z)$ is tight, the perfect matching M has a unique edge xy in $\partial(Z)$. Without loss of generality assume that $x \in Z$. For every edge e in T the subtree of $T \ltimes e$ containing an odd number of leaves mapped to vertices of Z contains the leaf mapped to x. Thus all inner edges of T are oriented towards the leaf mapped to x.

So consider the three edges e_1, e_2 and e_3 incident to some vertex $v \in V(T) \setminus \mathsf{L}(T)$. For exactly one of them the subtree of $T \ltimes e_i$ that does not contain v contains the leaf mapped to x. Thus, exactly one of them is oriented away from v unless this edge is itself trivial and incident to the leaf mapped to x. Therefore, $\overrightarrow{T_Z}$ is free of inconsistencies and the unique vertex of $\mathsf{spine}(T)$ that is adjacent to the leaf mapped to x is the only sink and adjacent to two leaves x and y. □

Finally, we obtain that for every tight cut contraction there is an M-perfect matching decomposition that does not have larger width than the optimal M-perfect matching decomposition of the original graph.

Theorem 4.3.15. Let G be a graph with a perfect matching, $\partial(Z)$ a nontrivial tight cut in G, $M \in \mathcal{M}(G)$, and (T, δ) an M-perfect matching decomposition of G of width k. Moreover, let G_Z be the matching covered graph obtained by the tight cut contraction of Z into the vertex v_Z. Then, there is an $M|_{G_Z}$-perfect matching decomposition of G_Z of width at most k.

Proof. Let $\overrightarrow{T_Z}$ be the Z-orientation of T. By Lemma 4.3.14, $\overrightarrow{T_Z}$ has a unique sink s, which is adjacent to two leaves ℓ_x and ℓ_y. These leaves are mapped to two vertices x and y, $\delta(\ell_x) = x$ and $\delta(\ell_y) = y$, where xy is the unique edge in $M \cap \partial(Z)$. Assume without loss of generality that $x \in Z$ and $y \in \overline{Z}$.

We define an $M|_{G_Z}$-perfect matching decomposition (T', δ') as follows. First we obtain a subtree T'' of T by deleting all leaves ℓ with $\delta(\ell) \in Z \setminus \{x\}$. Then, we obtain T' from T'' by first removing leaves that are not leaves of T as long a possible and then trimming. Note that every bipartition of $\mathsf{L}(T')$ induced by an inner edge in T' is also induced by an edge in T. To obtain δ' we define

$$\delta'(t) := \begin{cases} v_Z, & \text{if } \delta(t) = x, \text{ and} \\ \delta(t), & \text{otherwise} \end{cases}$$

for $t \in \mathsf{L}(T')$. The perfect matching $M|_{G_Z}$ of G_Z contains all edges of M with both endpoints in $\overline{Z}$ and the edge $\{y, v_Z\}$, so, by construction, (T', δ') is an $M|_{G_Z}$-perfect matching decomposition of G_Z.

Next, consider the width of (T', δ'). Let $t_1t_2 \in E(T')$ be an inner edge and X_1' the shore of $\partial_{G_Z}(t_1t_2)$ containing v_Z. Then, there is an edge $e \in E(T)$ such that X_1 is a shore of $\partial_G(e)$ and $(X_1 \setminus Z) \cup \{v_Z\} = X_1'$. As $v_Z \in X_1'$, by construction, we have $x \in X_1$. Additionally, we know that $|X_1|$ is even, because (T, δ) is an M-perfect matching decomposition. So $|X_1 \cap Z|$ is odd. By Lemma 4.3.12, we obtain $\mathsf{mp}(\partial_{G_Z}(X_1')) = \mathsf{mp}(\partial_{G_Z}((X_1 \setminus Z) \cup \{v_Z\})) \leq \mathsf{mp}(\partial_G(X_1)) \leq k$.

Thus, (T', δ') is an $M|_{G_Z}$-perfect matching decomposition of G_Z of width at most k. □

As bicontractions are a special case of tight cut contractions, we obtain the following corollary. We need this in order to prove Theorem 4.3.1 later on.

Corollary 4.3.16. Let G be a graph with a perfect matching M and H a brick, brace or a matching minor obtained by a series of bicontractions from an M-conformal subgraph of G. Then, $\mathsf{pmw}_{M|_H}(H) \leq \mathsf{pmw}_M(G)$.

4.3.3 Imbalances of sets

In bipartite graphs a set does not have to contain as many vertices of one colour class as of the other. In the context of matchings the difference and what neighbours these vertices have inside of the set and outside are interesting as they influence how a matching interacts with the set. Thus, we introduce the following definition.

Definition 4.3.17. Let B be a bipartite graph, and $X \subseteq V(G)$. If $|X \cap V_1| = |X \cap V_0|$ we say that X is *balanced*, otherwise it is *unbalanced*.

If X is unbalanced, then there are $i, j \in \{1, 2\}$, and $k \in \mathbb{N}$ such that $|X \cap V_i| = |X \cap V_j| + k$. In this case we call $\mathsf{Maj}(X) := X \cap V_i$ the *majority* of X, and we call $\mathsf{Min}(X) := X \cap V_j$ the *minority* of X. We say that k is the *imbalance* of X.

In general we define

$$\mathsf{imbalance}(X) := \begin{cases} 0, & \text{if } X \text{ is balanced, or} \\ k, & \text{if the imbalance of } X \text{ is } k. \end{cases}$$

We can make the following observation on the relation between the parity of X and its imbalance.

Observation 4.3.18. Let B be a bipartite graph, and $X \subseteq V(G)$. Then,

$$|X| \equiv \mathsf{imbalance}(X) \pmod 2.$$

We aim to prove that the vertices of the minority in a cut induced by the edge of a width-2-matching decomposition of a brace cannot have neighbours on the other side of the cut.

Lemma 4.3.19. Let B be a brace with $\mathsf{pmw}(B) = 2$ and (T, δ) be an optimal perfect matching decomposition of B. If X is a shore of $\partial_B(e)$ for an edge $e \in E(\mathsf{spine}(\mathsf{spine}(T)))$, then no vertex of the minority of X has a neighbour in $\overline{X}$.

There is an equivalent way of looking at braces via the extendability of matchings. Let G be a graph with a perfect matching and $F \subseteq E(G)$ a matching. We say that F is *extendable* if there exists $M \in \mathcal{M}(G)$ such that $F \subseteq M$. For any positive integer $k \in \mathbb{N}$, G is k-*extendable* if it is connected, has at least $2k+2$ vertices, and every matching of size k in G is extendable. The following statement by Lovász and Plummer shows that all braces except C_4 are 2-extendable.

Theorem 4.3.20 (Lovász and Plummer [LP09])**.** A bipartite graph B is a brace if and only if it is either isomorphic to C_4, or it is 2-extendable.

We consider the imbalance of sets in the context of k-extendable graphs, which we do in the following subsection and then apply it to braces in Subsection 4.3.5.

4.3.4 Imbalance and spines in k-extendable bipartite graphs

We consider the case of k-extendable bipartite graphs in this subsection. Later on we only make use of the case $k = 2$ in order to prove Lemma 4.3.19. Plummer proves that Theorem 4.3.20 generalises to k-extendable graphs.

Theorem 4.3.21 (Plummer [Plu86])**.** Let B be a bipartite graph and $k \in \mathbb{N}$ a positive integer. The following statements are equivalent.

1. B is k-extendable.
2. $|V_1| = |V_0|$, and for all non-empty $S \subseteq V_1$ with $|S| \leq |V_1| - k$ we have $|N_B(S)| \geq |S| + k$.
3. For all sets $S_1 \subseteq V_1$ and $S_2 \subseteq V_0$ with $|S_1| = |S_2| \leq k$ the graph $B - S_1 - S_2$ has a perfect matching.

He also proved the two following properties of k-extendable graphs.

Theorem 4.3.22 (Plummer [Plu80])**.** Let $k \in \mathbb{N}$ be a positive integer. Then, every k-extendable graph is also $(k-1)$-extendable.

Theorem 4.3.23 (Plummer [Plu80])**.** Let $k \in \mathbb{N}$ be a positive integer. Then, every k-extendable graph is $(k+1)$-connected.

Especially the latter is useful when working with the matching porosity of cuts. In order to see why, we make use of the following well known theorem proving that the vertex cover number and the size of a maximum matching coincide in bipartite graphs. This property was discovered by Kőnig [Kő31] and Egerváry [Ege31] independently in 1931 (see for example [LP09]).

Theorem 4.3.24 (Kőnig's Theorem)**.** If B is a bipartite graph, then $\tau(B) = \nu(B)$.

Using this and Plummer's result we can establish that in bipartite k-extendable graphs no cut of small porosity can have two large shores. To this end we define for every graph G and cut $\partial_G(X)$ in G the graph $G[\partial_G(X)]$ as the subgraph of G induced by all the edges in $\partial_G(X)$. $G[\partial_G(X)]$

Lemma 4.3.25. Let $k \in \mathbb{N}$ and B be a bipartite k-extendable graph and $X \subseteq V(B)$, then one of the following holds for every $k' \leq k$:

1. $\nu(B[\partial_B(X)]) > k'$,
2. $|X| \leq k'$, or
3. $|\overline{X}| \leq k'$.

Proof. We assume $\nu(B[\partial_B(X)]) \leq k'$. As B is bipartite, the graph $B[\partial_B(X)]$ is as well. By Theorem 4.3.24, we thus obtain $\tau(B[\partial_B(X)]) = \nu(B[\partial_B(X)]) \leq k'$. So there is a set of vertices S of size at most k' hitting all edges crossing $\partial_B(X)$. This means S is a separator of size at most k' in B separating $X \setminus S$ from $\overline{X} \setminus S$. By Theorem 4.3.23, B is $k+1$-connected, therefore $X \subseteq S$ and thus $|X| \leq k'$, or $\overline{X} \subseteq S$ and thus $|\overline{X}| \leq k'$. □

We establish a connection between the matching porosity of a cut and its imbalance in k-extendable bipartite graphs.

Lemma 4.3.26. Let $k \in \mathbb{N}$ be a positive integer, B be a k-extendable and bipartite graph, and $X \subseteq V(G)$ such that $\mathsf{mp}(\partial_B(X)) = k$ and $k+2 \leq |X| \leq |V(B)| - (k+2)$. Then, $\mathsf{imbalance}(X) = k$.

Proof. We start by observing that every perfect matching of B has at most $|\mathsf{Min}(X)|$ many edges matching two vertices of X. Thus, the matching porosity of $\partial_B(X)$ yields an upper bound on the imbalance of X, that is, $k = \mathsf{mp}(\partial_B(X)) \geq |X| - 2|\mathsf{Min}(X)| = \mathsf{imbalance}(X)$. By Lemma 4.3.6 and Observation 4.3.18, we know that $k \equiv |X| \equiv \mathsf{imbalance}(X) \pmod 2$.

Suppose towards a contradiction that $k' \coloneqq \mathsf{imbalance}(X) \leq k-2$. Additionally, we may assume without loss of generality that $\mathsf{Maj}(X) \subseteq V_0$. We split $B[\partial_B(X)]$ into the following two subgraphs.

$$\begin{aligned} B_1 &\coloneqq B[\mathsf{Min}(X) \cup (V_0 \setminus \mathsf{Maj}(X))]\,, \text{ and} \\ B_2 &\coloneqq B[(V_1 \setminus \mathsf{Min}(X)) \cup \mathsf{Maj}(X)]\,. \end{aligned}$$

We have $B_1 \cup B_2 = B[\partial_B(X)]\,.$

Suppose $\nu(B_1) \geq \frac{k-k'}{2} + 1$, then there is a matching F of size $\frac{k-k'}{2} + 1$ in B_1. As $|F| = \frac{k-k'}{2} + 1 \leq k$ and B is k-extendable, there is a perfect matching M_F of B with $F \subseteq M_F$. Due to $\mathsf{mp}(\partial_B(X)) = k$, at most $k - \left(\frac{k-k'}{2} + 1\right) = \frac{k+k'}{2} - 1$ edges of $M_F \cap \partial_B(X)$ have an endpoint in $\mathsf{Maj}(X)$. Thus we obtain

$$\begin{aligned} |\mathsf{Maj}(X) \setminus V(M_F \cap \partial_B(X))| &\geq |\mathsf{Maj}(X)| - \left(\frac{k+k'}{2} - 1\right) \\ &= |\mathsf{Maj}(X)| - \frac{k}{2} - \frac{k'}{2} + 1 \\ &= |\mathsf{Min}(X)| + k' - \frac{k}{2} - \frac{k'}{2} + 1 \\ &> |\mathsf{Min}(X)| - \frac{k-k'}{2} - 1 \\ &\geq |\mathsf{Min}(X) \setminus V(M_F \cap \partial_B(X))|\,. \end{aligned}$$

Therefore, $\mathsf{imbalance}(X \setminus V(M_F \cap \partial_B(X))) \geq 1$, contradicting M_F being a perfect matching. Thus, $\nu(B_1) \leq \frac{k-k'}{2}$. With similar arguments $\nu(B_2) \geq \frac{k+k'}{2} + 1$ yields a contradiction. Thus, $\nu(B_2) \leq \frac{k+k'}{2}$. It follows that

$$\nu(B[\partial_B(X)]) = \nu(B_1) + \nu(B_2) \leq \frac{k-k'}{2} + \frac{k+k'}{2} = k.$$

Together with $k+2 \leq |X| \leq |V(B)| - (k+2)$, this contradicts Lemma 4.3.25. Thus, we obtain that $\mathsf{imbalance}(X) = k$. □

Next, we show that decompositions with the special structure of the spine of the spine of the decomposition tree being a path have the following property. The edges of this path having matching porosity exactly k induce shores that have imbalance k and the neighbourhood of the minority of these shores is completely contained in the shore. So the shores have a kind of closure property when it comes to the minority: no vertex of the minority has neighbours outside the shore.

Theorem 4.3.27. Let $k \geq 2$ be an integer and B be a k-extendable bipartite graph with a perfect matching decomposition (T, δ) of width k such that $\mathsf{spine}(\mathsf{spine}(T))$ is a path. Then, for all $e \in \mathsf{spine}(\mathsf{spine}(T))$ with $\mathsf{mp}(\partial_B(e)) = k$, every shore X of $\partial_B(e)$ satisfies

1. $\mathsf{imbalance}(X) = k$, and
2. $N_B(\mathsf{Min}(X)) \subseteq X$.

Proof. We first consider how the colour classes can be distributed in the shores of two cuts X and Y corresponding to two adjacent edges in the path $\mathsf{spine}(\mathsf{spine}(T))$ such that $X \subseteq Y$. We claim that X and Y differ by exactly one vertex from each colour class.

Claim 1. If $|V(B)| \geq 2k+4$ and e_1, e_2 two adjacent edges of $\mathsf{spine}(\mathsf{spine}(T))$ such that the cut $\partial_B(e_1)$ has a shore X_1 and the cut $\partial_B(e_2)$ has a shore X_2 with $X_1 \subseteq X_2$ and $\mathsf{mp}(\partial_B(X_1)) = \mathsf{mp}(\partial_B(X_2)) = k$. Then, we have $|X_1 \cap V_i| + 1 = |X_2 \cap V_i|$ for both $i \in \{0, 1\}$. See Figure 4.7 for an illustration.

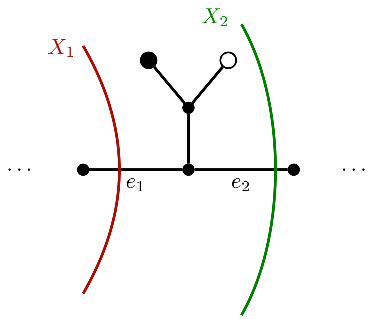

Figure 4.7: There are exactly two vertices in $X_2 \setminus X_1$ and they come from different colour classes of B.

Proof. Lemma 4.3.6 and $\mathsf{spine}(\mathsf{spine}(T))$ being a path together imply $|X_2| - |X_1| = 2$. Suppose towards a contradiction that both vertices in $X_2 \setminus X_1$ are from the same colour class, without loss of generality say V_1, that is, $|X_1 \cap V_1| + 2 = |X_2 \cap V_1|$. Due to $\mathsf{mp}(\partial_B(X_2)) = k$ we know $|\overline{X_1}| \geq k+2$ and $|X_2| \geq k+2$. So, by Lemma 4.3.26, we have $\mathsf{imbalance}(X_1) = \mathsf{mp}(\partial_B(X_1)) = k$, or $|X_1| = k$ for X_1 and we have $\mathsf{imbalance}(X_2) = \mathsf{mp}(\partial_B(X_2)) = k$, or $|\overline{X_2}| = k$.

If $|X_1| = k$ and $|\overline{X_2}| = k$, then $|V(B)| = |X_1| + 2 + |\overline{X_2}| = 2k + 2$, which contradicts $|V(B)| \geq 2k+4$. So we have $\mathsf{imbalance}(X_1) = k$ or $\mathsf{imbalance}(X_2) = k$. Suppose only one of them holds, so without loss of generality consider the case that $\mathsf{imbalance}(X_1) \leq k-2$ and $\mathsf{imbalance}(X_2) = k$. This implies that $\mathsf{Maj}(X_1) \subseteq V_1$ and $\mathsf{Maj}(X_2) \subseteq V_1$. We split $B[\partial_B(X_2)]$ into the two subgraphs

$$B_1 := B\big[\mathsf{Maj}(X_2) \cup \mathsf{Maj}\big(\overline{X_2}\big)\big] \text{ and}$$
$$B_2 := B\big[\mathsf{Min}(X_2) \cup \mathsf{Min}\big(\overline{X_2}\big)\big] .$$

We have $B_1 \cup B_2 = B[\partial_B(X_2)]$, and B_1 and B_2 are disjoint. Due to $|X_1| = k$ and $\mathsf{imbalance}(X_1) \leq k-2$, we know $|\mathsf{Min}(X_2)| = |\mathsf{Min}(X_1)| = 2$ and thus, $\nu(B_2) \leq 2$.

Suppose $\nu(B_1) \geq k$, that is, there is a matching F of size k in B_1. By Theorem 4.3.22, B has a perfect matching M_F containing F. The set $X_1 \setminus V(F)$ contains two vertices of V_0 and no vertex of V_1. So M_F maps both these vertices to vertices in $\overline{X_2}$. Thus, $|M_F \cap \partial_B(X_2)| = |X_2| = k+2$, which contradicts $\mathsf{mp}(\partial_B(X_2)) = k$. Therefore, we have $\nu(B_1) \leq k-1$.

Similarly, suppose there is a matching F of size two in B_2. Then, by Theorem 4.3.22, B has a perfect matching M_F containing F. The set $X_2 \setminus V(F)$ contains k vertices of V_1 and no vertex of V_0. So M_F maps all these k vertices to vertices in $\overline{X_2}$. Thus, $|M_F \cap \partial_B(X_2)| = |X_2| = k+2$, which contradicts $\mathsf{mp}(\partial_B(X_2)) = k$. Therefore, we have $\nu(B_2) \leq 1$.

So we obtain $\nu(B[\partial_B(X_2)]) = \nu(B_1)+\nu(B_2) \leq k$, which together with $|X_2| \geq k+2$ and $|\overline{X_2}| \geq k+2$ contradicts Lemma 4.3.25.

Thus, we know that $\mathsf{imbalance}(X_1) = \mathsf{imbalance}(X_2) = k$. But this contradicts that the two vertices in $X_2 \setminus X_1$ come from the same colour class. So, we obtain $|X_1 \cap V_i| + 1 = |X_2 \cap V_i|$ for both $i \in \{0, 1\}$. ■

We define $P := \mathsf{spine}(\mathsf{spine}(T))$ with the two endpoints $p_\triangleleft$ and $p_\triangleright$. We order the edges $(e_1, \ldots, e_\ell)$ of P by occurrence along P when traversing it from $p_\triangleleft$ to $p_\triangleright$. Next, we show that P contains two edges that induce a cut with one shore building a star with $k+1$ leaves.

Claim 2. The tree T contains edges $e_\triangleleft$ and $e_\triangleright$ such that for $\diamond \in \{\triangleleft, \triangleright\}$ the cut $\partial_B(e_\diamond)$ has a shore $X_\diamond$ of size $k+2$ satisfying the following conditions.

(i) $e_\triangleleft = e_\triangleright$ if and only if $|V(B)| = 2k+4$,

(ii) $X_\triangleleft \cap X_\triangleright = \emptyset$,

(iii) if $\mathsf{Maj}(X_\triangleleft) \subseteq V_i$, then $\mathsf{Min}(X_\triangleright) \subseteq V_{|1-i|}$ for $i \in \{0, 1\}$, and

(iv) $B[X_\diamond]$ is a star such that its central vertex has no neighbour in $\overline{X_\diamond}$ for both $\diamond \in \{\triangleleft, \triangleright\}$.

Proof. Choose j minimal with $\mathsf{mp}(\partial_B(e_j)) = k$. Let $T_\triangleleft$ be the subtree in $T \ltimes e_j$ that contains $p_\triangleleft$ and $X'_\triangleleft := \delta(T_\triangleleft)$.

Consider the size of $X'_\triangleleft$. Due to Lemma 4.3.6, we can defer $|X'_\triangleleft| \neq k+1$. Suppose towards a contradiction that $|X'_\triangleleft| \geq k+2$. By Lemma 4.3.26, this implies $\mathsf{imbalance}(X'_\triangleleft) = k$. Consider the shore $X''_\triangleleft$ of $\partial_B(e_{j-1})$ with $X''_\triangleleft \subseteq X'_\triangleleft$. By minimality of j, we know that $\mathsf{mp}(\partial_B(e_{j-1})) \leq k-1$, and thus, $\mathsf{imbalance}(X''_\triangleleft) \leq k-1$. By $|X'_\triangleleft| - |X''_\triangleleft| \leq 2$, we know that $|X''_\triangleleft| \geq k$, and by $\mathsf{mp}(\partial_B(X'_\triangleleft)) = k$, we obtain $|\overline{X''_\triangleleft}| \geq k$ as well. By Lemma 4.3.25, this implies that $\partial_B(X''_\triangleleft)$ contains a matching F of size k. As B is k-extendable there is a perfect matching M_F of B containing F. This yields a contradiction to $\mathsf{mp}(\partial_B(X''_\triangleleft)) \leq k-1$. Thus, $|X'_\triangleleft| = k$.

We define $e_\triangleleft := e_{j+1}$ and additionally $X_\triangleleft$ to be the shore of $\partial_B(e_\triangleleft)$ containing $X'_\triangleleft$. By Lemma 4.3.6, we obtain $|X_\triangleleft| = k+2$ as desired. Moreover, $\mathsf{mp}(\partial_B(X_\triangleleft)) = k = \mathsf{imbalance}(X'_\triangleleft) = \mathsf{imbalance}(X_\triangleleft)$. Thus, by Claim 1, $|\,(X_\triangleleft \setminus X'_\triangleleft) \cap V_1| = |\,(X_\triangleleft \setminus X'_\triangleleft) \cap V_0| = 1$. We know that $X_\triangleleft \subseteq V_i$ for some $i \in \{0,1\}$, because $\mathsf{imbalance}(X_\triangleleft) = |X_\triangleleft|$, so let us assume that $i = 1$ without loss of generality. Together with $\mathsf{mp}(\partial_B(X_\triangleleft)) = k$, this implies that the only vertex w of V_0 in $X_\triangleleft$ has no neighbours in $\overline{X_\triangleleft}$. As B is $(k+1)$-connected, by Theorem 4.3.23, $N_B(w) = X_\triangleleft \setminus \{w\}$. Hence, $B[X_\triangleleft]$ forms the desired star.

Now, choose j' maximal with $\mathsf{mp}(\partial_B(e_{j'})) = k$. Let $T_\triangleright$ be the subtree in $T \ltimes e_{j'}$ that contains $p_\triangleright$ and $X'_\triangleright := \delta(T_\triangleright)$. Also, we define $e_\triangleright := e_{j'-1}$ and $X_\triangleright$ to be the shore of $\partial_B(e_\triangleright)$ disjoint from $X_\triangleleft$. This ensures that $X_\triangleleft \cap X_\triangleright = \emptyset$. By symmetric arguments to the above we obtain $|X'_\triangleright| = k$ and $|X_\triangleright| = k+2$. Additionally, $\mathsf{mp}(\partial_B(X_\triangleright)) = k = \mathsf{imbalance}(X'_\triangleright) = \mathsf{imbalance}(X_\triangleright)$ and $X_\triangleleft \subseteq V_i$ for some $i \in \{0,1\}$. Due to Claim 1, the decomposition gains a vertex from each colour class with every edge of matching porosity k. By Lemma 4.3.25, the edges along P lying between $e_\triangleleft$ and $e_\triangleright$ all have matching porosity k. Thus, $X_\triangleright \subseteq V_0$, as $|V_0| = |V_1|$. So, $B[X_\triangleright]$ forms the desired star with the centre vertex b being from V_1.

Finally, $X_\triangleleft \cap X_\triangleright = \emptyset$ implies that $e_\triangleleft = e_\triangleright$ if and only if $X_\triangleleft \cup X_\triangleright = V(B)$, that is, $|V(B)| = 2k+4$, because $|X_\triangleleft| = |X_\triangleright| = k+2$. ■

Using the edges found in Claim 2 and the subpath P' of P starting with $e_\triangleleft$ and ending with $e_\triangleright$ we now conclude the proof.

If $|V(B)| = 2k+4$, then, by Claim 2, $e_\triangleleft = e_\triangleright$ and these are the only edges inducing cuts of matching porosity k, thus the statement holds.

So, assume that $|V(B)| \geq 2k+6$ and $e_\triangleleft \neq e_\triangleright$. We prove the statement for the shores of any edge e_i lying between $e_\triangleleft$ and $e_\triangleright$ assuming that the shores of the adjacent edge e_{i-1} fulfil the statement, that is, assuming $\mathsf{imbalance}(X_{e_{i-1}}) = k$ and $N_B(\mathsf{Min}(X_{e_{i-1}})) \subseteq X_{e_{i-1}}$. Let us assume without loss of generality that $\mathsf{Min}(X_{e_i}) \subseteq V_1$. By Claim 1, we know that $\mathsf{imbalance}(X_{e_i}) = \mathsf{imbalance}(X_{e_{i-1}}) = k$ and there is a unique vertex a in $\left(X_{e_i} \setminus X_{e_{i-1}}\right) \cap V_1$. Suppose a has a neighbour b in $\overline{X_{e_i}}$. Then, there is a perfect matching M of B containing ab and $M \cap \partial_B(X_{e_i})$ contains at least $k+2$ edges, because Claim 1 implies $\mathsf{Min}(X_{e_{i-1}}) \subseteq V_1$, a contradiction. Thus, $N_B(\mathsf{Min}(X_{e_i})) \subseteq X_{e_i}$ and we are done. $\square$

4.3.5 Application to braces of perfect matching width 2

We now apply Theorem 4.3.27 to the braces of perfect matching width two. In order to do so we need to first show that the decompositions of such braces indeed have the property that the spine of the spine is always a path. We start by considering the spine and proving that it is cubic. This holds for bricks as well as for braces.

Lemma 4.3.28. Let G be a brick or brace of perfect matching width two and (T, δ) be an optimal perfect matching decomposition. Then, $\mathsf{spine}(T)$ is cubic.

Proof. By Corollary 4.3.4, $\mathsf{spine}(T)$ is cubic if and only if T does not have any odd edges. Suppose there is an odd edge $e = t_1t_2 \in E(T)$ and $(T_1, T_2) := T \ltimes e$. Then, both shores $X_1 := \delta(T_1)$ and $X_2 := \delta(T_2)$ contain an odd number of vertices. So, by Lemma 4.3.6, the matching porosity of X_1 is odd as well, that is, $\mathsf{mp}(\partial_G(X_1))$ is odd. Due to (T, δ) having width 2, this implies $\mathsf{mp}(\partial_G(X_1)) = 1$. But e is an edge of the spine, so $|X_1| \geq 3$ as well as $|X_2| \geq 3$. This makes $\partial_G(X_1)$ a nontrivial tight-cut in G, contradicting G being a brick or brace. $\square$

Using this we can prove that the spine of the spine of the decomposition tree of every brace is always a path, which allows us to apply Theorem 4.3.27.

Lemma 4.3.29. Let B be a brace of perfect matching width two and (T, δ) a perfect matching decomposition of minimum width for B. Then, $\mathsf{spine}(\mathsf{spine}(T))$ is a path.

Proof. Suppose towards a contradiction $\mathsf{spine}(\mathsf{spine}(T))$ contains a vertex t with three neighbours t_1, t_2 and t_3. By Lemma 4.3.28, we know $\mathsf{spine}(T)$ is cubic. As $t_1, t_2, t_3 \in V(\mathsf{spine}(\mathsf{spine}(T)))$ and the spine of T is cubic, we obtain that the degree of t_i is again exactly 3 in $\mathsf{spine}(T)$ for all $i \in \{1, 2, 3\}$. This implies that every t_i again has two neighbours other than t in $\mathsf{spine}(T)$, which are not leaves, because

$t_i \in V(\mathsf{spine}(\mathsf{spine}(T)))$. These neighbours of t_i again, due to $\mathsf{spine}(T)$ being cubic, have two neighbours other than t_i. So, if we consider T_i to be the tree of $T \ltimes tt_i$ that contains t_i and $X_i := \delta(T)$, then $|X_i| \geq 4$. By Corollary 4.3.4, the tree T does not contain odd edges. So additionally, we obtain $\mathsf{mp}(\partial_B(X_i)) = 2$. Thus, Lemma 4.3.26 implies $\mathsf{imbalance}(X_i) = 2$.

So every X_i has a majority in one of the colour classes. If this was always the same colour class, then we would obtain that the whole brace B has an imbalance of 6, a contradiction. Thus, exactly two of them have their majority in the same colour class, let us assume without loss of generality that $\mathsf{Maj}(X_1), \mathsf{Maj}(X_2) \subseteq V_1$ and $\mathsf{Maj}(X_3) \subseteq V_0$. But then, we still obtain $|V_1| = |V_0| + 2$, so the whole brace has an imbalance of 2, which still yields a contradiction.

Thus, no such vertex can exist and $\mathsf{spine}(\mathsf{spine}(T))$ is a path. $\square$

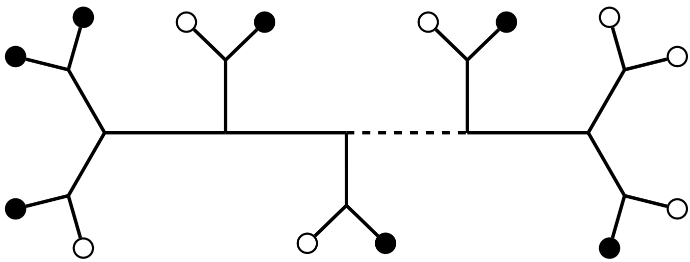

Figure 4.8: The linear structure of a perfect matching decomposition of width 2 with a claw on each side and two vertices from different colour classes added with each step. The black vertices in the figure represent the leaves mapped to a vertex from V_1 and the white vertices represent the leaves mapped to vertices from V_0.

Finally, Theorem 4.3.27 and Lemma 4.3.29 together imply our desired result and that every width-2-perfect matching decomposition of a brace looks as depicted in the example in Figure 4.8.

Lemma 4.3.19. Let B be a brace with $\mathsf{pmw}(B) = 2$ and (T, δ) be an optimal perfect matching decomposition of B. If X is a shore of $\partial_B(e)$ for an edge $e \in E(\mathsf{spine}(\mathsf{spine}(T)))$, then no vertex of the minority of X has a neighbour in $\overline{X}$.

4.4 Graphs of M-perfect matching width 2

Now, we have gathered everything we need to prove the following characterisation of braces with M-perfect matching width two. From the braces we can then obtain a characterisation for all bipartite graphs with a perfect matching of M-perfect matching width two, which, as observed in Section 4.3, implies Theorem 4.0.1.

Theorem 4.4.1. Let B be a brace, then the following statements are equivalent.

(1) $\mathsf{pmw}_M(B) = 2$ for an $M \in \mathcal{M}(B)$,

(2) $\mathsf{pmw}_M(B) = 2$ for all $M \in \mathcal{M}(B)$, and

(3) B is isomorphic to C_4 or $K_{3,3}$.

Proof. Clearly, (2) implies (1). So, we start by proving that (1) implies (3). To this end, assume (1) holds, that is, there is a perfect matching M in B such that $\mathsf{pmw}_M(B) = 2$. By Theorem 4.2.5 and because B is a brace, $\mathsf{pmw}(B) = 2$. So, there is an M-perfect matching decomposition (T, δ) of width two of B which is an optimal perfect matching decomposition as well.

Suppose towards a contradiction that $|V(B)| \geq 8$. By Claim 2, there is an edge $e \in E(\mathsf{spine}(\mathsf{spine}(T)))$ with $\partial_B(e)$ having a shore X of size four such that $B[X]$ is a claw. Using Lemma 4.3.19, we obtain that $\mathsf{imbalance}(X) = 2$ and therefore X is not M-conformal. This contradicts the definition of an M-perfect matching decomposition as e is an inner edge of T.

Thus, $|V(B)| \leq 6$. On six vertices there are exactly two braces: C_4 and $K_{3,3}$, which yields (3).

Finally we prove that (3) implies (2). So, assume (3) holds. We consider C_4 and $K_{3,3}$ as separate cases.

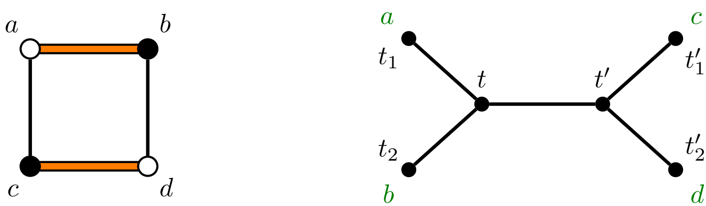

Figure 4.9: The C_4 with the matching M on the left and its M-perfect matching decomposition (T, δ) of width 2 on the right.

We start with the case that $B \cong C_4$. Let M be a perfect matching in C_4. Note that the two possible matchings in C_4 are symmetric, thus we can assume without loss of generality that $V(C_4) = \{a, b, c, d\}$ and $M = \{ab, cd\}$. We construct an M-perfect matching decomposition (T, δ) for C_4. The tree T is defined by $V(T) := \{t, t', t_1, t_2, t'_1, t'_2\}$ and $E(T) := \{tt', tt_1, tt_2, t't'_1, t't'_2\}$. For δ we define $\delta(t_1) := a$, $\delta(t_2) := b$, $\delta(t'_1) := c$ and $\delta(t'_2) := d$, also see Figure 4.9 for an illustration. Note that (T, δ) is indeed an M-perfect matching decomposition of C_4. Also its width is two as the matching porosity of every cut induced by an edge of T is either one or two.

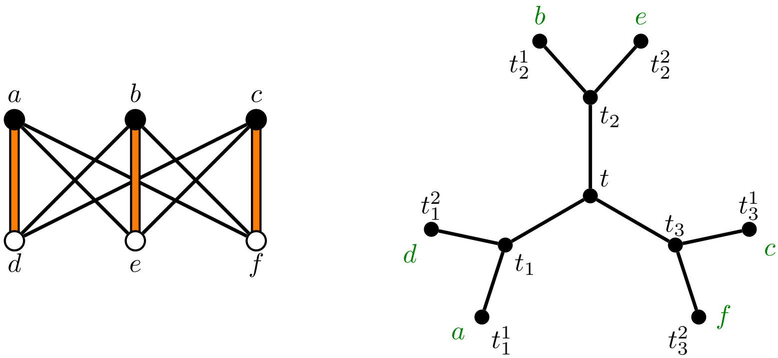

Figure 4.10: The $K_{3,3}$ with the matching M on the left and its M-perfect matching decomposition (T, δ) of width 2 on the right.

Next, we consider the case that $B \cong K_{3,3}$. We again start by choosing a perfect matching M in $K_{3,3}$ and, noticing that all matchings in $K_{3,3}$ are symmetric, we can assume without loss of generality that $V_1 = \{a, b, c\}$, $V_0 = \{d, e, f\}$, and $M = \{ad, be, cf\}$. We construct an M-perfect matching decomposition (T, δ) for $K_{3,3}$ as follows. The cubic tree T is given by

$$V(T) := \{t\} \cup \{t_i, t_i^1, t_i^2 \mid 1 \leq i \leq 3\}, \text{ and}$$
$$E(T) := \{tt_i \mid 1 \leq i \leq 3\} \cup \{t_i t_i^1, t_i t_i^2 \mid 1 \leq i \leq 3\}.$$

Then, we map the vertices of $K_{3,3}$ to the leaves of T by defining $\delta(t_1^1) := a$, $\delta(t_1^2) := d$, $\delta(t_2^1) := b$, $\delta(t_2^2) := e$, $\delta(t_3^1) := c$, and $\delta(t_3^2) := f$. See Figure 4.10 for an illustration. Note that (T, δ) is indeed an M-perfect matching decomposition of $K_{3,3}$. Also its width is two as the matching porosity of every cut induced by an edge of T is either one or two.

So, (2) holds, which finishes the proof. $\square$

We now obtain the general characterisation of bipartite graphs with a perfect matching of M-perfect matching width two.

Theorem 4.3.1. Let B be a bipartite matching covered graph, then the following statements are equivalent.

(1) $\mathsf{pmw}_M(B) = 2$ for an $M \in \mathcal{M}(B)$,

(2) $\mathsf{pmw}_M(B) = 2$ for all $M \in \mathcal{M}(B)$, and

(3) Every brace of B is either isomorphic to C_4 or to $K_{3,3}$.

Proof. By Theorem 4.3.9, all maximal families of pairwise laminar nontrivial tight cuts in B have the same size. Let $\mathcal{C}$ be a maximal family of pairwise laminar tight cuts in B. We prove the statement by induction over $|\mathcal{C}|$. In case that $|\mathcal{C}| = 0$, B is a brace and thus, the statement follows from Theorem 4.4.1.

So assume that $|\mathcal{C}| \geq 1$. Thus, there is a nontrivial tight cut $\partial_B(Z)$ in $\mathcal{C}$. Let $B_Z := B/\left(Z \rightarrow v_Z\right)$, and $B_{\overline{Z}} := B/\left(\overline{Z} \rightarrow v_{\overline{Z}}\right)$ be the two tight cut contractions.

Clearly, (2) implies (1).

We prove that (1) implies (3). Assume (1) holds, that is, $\mathsf{pmw}_M(B) = 2$ for some $M \in \mathcal{M}(B)$. By Corollary 4.3.16, we obtain $\mathsf{pmw}_{M|_{B_Z}}(B_Z) = 2 = \mathsf{pmw}_{M|_{B_{\overline{Z}}}}(B_{\overline{Z}})$. Using the induction hypothesis we obtain that all braces of B_Z and $B_{\overline{Z}}$ are isomorphic to C_4 or $K_{3,3}$. As every brace of B is a brace of B_Z or $B_{\overline{Z}}$, this implies (3).

Next, we prove that (3) implies (2). Assume (3) holds, that is, all braces of B are isomorphic to C_4 or $K_{3,3}$. Choose an $M \in \mathcal{M}(B)$ arbitrarily. Let $m_Z \in M|_{B_Z}$ be the matching edge covering v_Z in B_Z and let u_Z be its other endpoint. Similarly, let $m_{\overline{Z}} \in M|_{B_{\overline{Z}}}$ the matching edge covering $v_{\overline{Z}}$ in $B_{\overline{Z}}$ and let $u_{\overline{Z}}$ be its other endpoint, that is, $M = \left(\left(M|_{B_Z} \cup M'|_{B_{\overline{Z}}}\right) \setminus \{m_Z, m_{\overline{Z}}\}\right) \cup \{u_Z u_{\overline{Z}}\}$.

By induction hypothesis, $\mathsf{pmw}_{M|_{B_Z}}(B_Z) = 2 = \mathsf{pmw}_{M|_{B_{\overline{Z}}}}(B_{\overline{Z}})$. Thus, there is an $M|_{B_Z}$-perfect matching decomposition (T_Z, δ_Z) of B_Z and an $M|_{B_{\overline{Z}}}$-perfect matching decomposition $\left(T_{\overline{Z}}, \delta_{\overline{Z}}\right)$ of $B_{\overline{Z}}$, both of width two. Let t_Z be the vertex of T_Z, that is, the common neighbour of the two leaves $\delta_Z^{-1}(v_Z)$ and $\delta_Z^{-1}(u_Z)$. Similarly, let $t_{\overline{Z}}$ be the vertex of $T_{\overline{Z}}$ that is, a common neighbour of $\delta_{\overline{Z}}^{-1}(v_{\overline{Z}})$ and $\delta_{\overline{Z}}^{-1}(u_{\overline{Z}})$.

Now, we construct an M-perfect matching decomposition of B. We first obtain a tree T'_Z from T_Z by removing the two leaves $\delta_Z^{-1}(v_Z)$ and $\delta_Z^{-1}(u_Z)$. And from $T_{\overline{Z}}$ we obtain $T'_{\overline{Z}}$ by removing $\delta_{\overline{Z}}^{-1}(v_{\overline{Z}})$ and $\delta_{\overline{Z}}^{-1}(u_{\overline{Z}})$. Then, from T_Z and $T'_{\overline{Z}}$ we obtain a tree T' by identifying t_Z and $t_{\overline{Z}}$ into a single vertex t and adding a new vertex t' as well as the edge tt'. From T' we obtain T'' by adding two new leaves t_1 and t_2 attached by the edges $t't_1$ and $t't_2$. The obtained tree T is cubic and has as many leaves as B has vertices. Finally, we define

$$\delta(d) := \begin{cases} \delta_Z(d), & \text{if } d \in \mathsf{L}(T_Z) \setminus \{\delta_Z^{-1}(v_Z)\}, \\ \delta_{\overline{Z}}(d), & \text{if } d \in \mathsf{L}(T_{\overline{Z}}) \setminus \{\delta_{\overline{Z}}^{-1}(v_{\overline{Z}})\}, \\ u_{\overline{Z}}, & \text{if } d = t_1, \text{ and} \\ u_Z, & \text{if } d = t_2. \end{cases}$$

This way we obtain an M-perfect matching decomposition (T, δ) of B.

Let e be an inner edge of T. Then, either $e = tt'$, or e is also an inner edge of T_Z or $T_{\overline{Z}}$. If $e = tt'$, then one of the shores of $\partial_B(e)$ is $\{u_Z, u_{\overline{Z}}\}$ and thus $\mathsf{mp}(\partial_B(e)) = 1$. If e is an inner edge of T_X for $X \in \{Z, \overline{Z}\}$, then $\mathsf{mp}(\partial_{B_X}(e)) = 1$ and thus, due to $\partial_B(Z)$ being a tight cut, we also obtain $\mathsf{mp}(\partial_B(e)) = 1$. Therefore, (T, δ) is of width 2, which implies $\mathsf{pmw}_M(B) = 2$ yielding (2). Which finishes the proof. ☐

4.5 Conclusion

So in the end we proved the characterisation of bipartite matching covered graphs with M-perfect matching width two.

Theorem 4.3.1. Let B be a bipartite matching covered graph, then the following statements are equivalent.

- **(1)** $\mathsf{pmw}_M(B) = 2$ for an $M \in \mathcal{M}(B)$,
- **(2)** $\mathsf{pmw}_M(B) = 2$ for all $M \in \mathcal{M}(B)$, and
- **(3)** Every brace of B is either isomorphic to C_4 or to $K_{3,3}$.

As we discussed earlier every matching in C_4 yields the digon as M-direction and every matching in $K_{3,3}$ yields the bidirected K_3 as M-direction. Therefore, Theorem 4.3.1 implies our main theorem.

Theorem 4.0.1. Let D be a digraph. Then, D has cyclewidth exactly one if and only if every strongly 2-connected butterfly minor of D is isomorphic to the digon or the bi-directed K_3.

A natural next step is to ask for characterisations of classes of graphs of cyclewidth two or three. One could also look for other more traditional characterisations of the graphs of cyclewidth one, for example by forbidden butterfly minors, or subgraphs.

5 Towards a directed structure theorem

The graph minor structure theorem is a cornerstone in the graph minor series by Robertson and Seymour. It is concerned with the question of what the structure of graphs excluding a fixed minor looks like. If one excludes a planar graph, one obtains a class of graphs of bounded treewidth [RS86]. But the class of grids shows that the treewidth is not bounded when excluding a non-planar graph. So, what is the structure of graphs excluding non-planar graphs? For K_5, this is answered by Wagner's theorem [Wag37], which states that every four-connected graph is planar if and only if it excludes K_5 as a minor. This can equivalently be stated as: every graph excluding K_5 can be obtained from the moebius ladder on eight vertices, also called the Wagner graph, and planar graphs by an operation called *small clique sums*. See Figure 5.1 for an illustration of the Wagner graph.

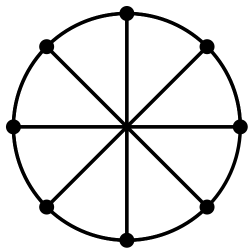

Figure 5.1: The Wagner graph.

The graph minor structure theorem [RS99] states that every graph that excludes a non-planar graph H as a minor can be decomposed by a tree decomposition with additional requirements. The bags of this tree decomposition do not have to have bounded width, but the adhesion, that is, the intersection between two bags, has bounded size. Additionally, the bags induce subgraphs that can be further decomposed into several 4-connected graphs. These 4-connected graphs, after deleting a small set of apex vertices, admit a drawing into a surface in which H does not have a drawing without crossings, such that there are only few areas, the *vortices*, that contain crossings.

The structure described by the graph minor structure theorem is parametrically necessary and sufficient. Not only does every graph that excludes a non-planar graph as a minor have the described structure, additionally graphs of the described structure do not contain a large clique as a minor.

In directed structure theory, there is an analogue to the grid theorem establishing that the cylindrical grid is an obstruction to directed treewidth [KK15]. But the property that a digraph is planar if and only if it is a minor of the cylindrical grid does not hold. The cylindrical grid does play a role similar to the grid in undirected graphs and thus, the cylindrical wall a similar one to the wall. However, every planar undirected graph is contained as a minor in an undirected wall. This is not true for digraphs. The class of cylindrical walls does not contain all planar digraphs, not even all strongly planar digraphs, as butterfly minors. In general, highly non-planar digraphs can be of directed treewidth one, as is witnessed by the class of DAGs. So, neither does excluding a planar digraph imply small directed treewidth, nor does even small directed treewidth imply planarity.

Directed structure theory evolved further in the past few years, introducing directed tangles [GKK$^+$20] and also a directed flat wall theorem [GKKK20]. However, building a directed structure theorem based on this flat wall theorem leads to difficulties as its properties differ from the undirected flat wall theorem. The main obstacle to transferring the results leading to the undirected structure theorem into the directed setting is the fact that there is no two-path theorem. A digraph can be highly connected and still not allow for two disjoint paths between two given pairs of vertices [Tho91]. In the undirected case, the two-path theorem yields that certain parts of a graph allow for a drawing without crosses into a closed disk, which makes up a crucial step in the undirected structure theorem.

Within this chapter we present an alternative directed flat wall theorem and provide some intuition why we think it yields a better base for a structure theorem. We also present a rather simple proof for a flat wall theorem for digraphs excluding a specific orientation of K_5. This yields valuable insight into the non-planar behaviour of digraphs.

5.1 The directed wall

Let us take a closer look at the existing directed flat wall theorem [GKKK20]. In order to do so, we have to introduce how a wall in the directed setting is defined alongside some additional related definitions.

An *elementary cylindrical wall* of order k is a graph, which is the union of k directed disjoint cycles $C_1, \dots, C_k$ of length $4k$ each and $2k$ disjoint paths. The cycle C_i has the vertex set $\{v_1^i, \dots, v_{4k}^i\}$ with the natural cyclic ordering. The paths are of two different kinds, we have the *in-paths* $P_1^{\mathsf{i}}, \dots, P_k^{\mathsf{i}}$ and the *out-paths* $P_1^{\mathsf{o}}, \dots, P_k^{\mathsf{o}}$ as follows, also see Figure 5.2 for an illustration.

$$P_j^{\mathsf{i}} = v_{4j-1}^k, v_{4j}^k, v_{4j-1}^{k-1}, v_{4j}^{k-1}, \dots, v_{4j-1}^1, v_{4j}^1 \text{ for all } j \in [k]$$

$$P_j^{\mathsf{o}} = v_{4(j-1)+1}^1, v_{4(j-1)+2}^1, v_{4(j-1)+1}^2, v_{4(j-1)+2}^2, \dots, v_{4(j-1)+1}^k, v_{4(j-1)+2}^k$$
$$\text{for all } j \in [k].$$

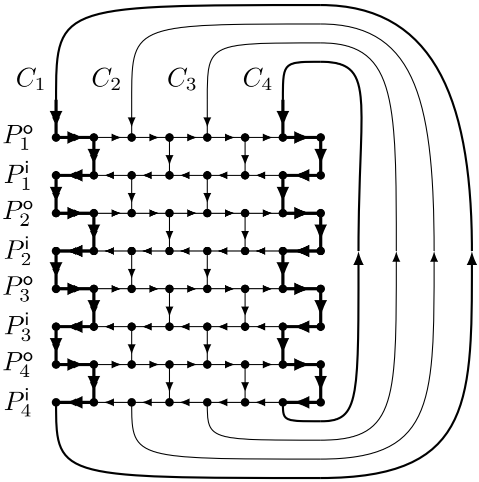

Figure 5.2: The cylindrical wall of order four. The perimeters are depicted using thick edges.

The paths P_i^{i} and P_i^{o} together make up the i-th *row* of W, which we also denote R_i. We consider rows modulo the wall size, that is, for $\ell > k$ we consider $P_\ell^{\mathsf{i}} \coloneqq P_{(\ell-1) \bmod k+1}^{\mathsf{i}}$, $P_\ell^{\mathsf{o}} \coloneqq P_{(\ell-1) \bmod k+1}^{\mathsf{o}}$ as well as $R_\ell \coloneqq R_{(\ell-1) \bmod k+1}$. The *perimeter* of W, denoted $\mathsf{per}(W)$, contains all vertices of C_1 and C_k. The vertices of C_1 we also refer to as the *inner perimeter*, $\mathsf{per}^{\mathsf{in}}(W)$, and we call the vertices of C_k the *outer perimeter*, $\mathsf{per}^{\mathsf{out}}(W)$. The graph that is obtained from W by removing the perimeter (and the edges incident to its vertices) is called the *interior* of the wall, we write $\mathsf{int}(W)$. A *subwall* $W[i,j]$ of W for $1 \le i < j \le k$ is the graph obtained by the union of the cycles $C_i, \dots, C_j$ and the subpaths $P_x^{\mathsf{i}}[C_i, C_j]$ and $P_x^{\mathsf{o}}[C_i, C_j]$ for all $1 \le x \le k$. We sometimes use $Q_1, \dots, Q_k$ to refer to the subpaths of $C_1, \dots, C_k$ starting on P_1^{o} and ending on P_k^{i}.

A *(cylindrical) wall* is a subdivision of an elementary cylindrical wall. The *branch vertices* of a wall are the vertices that have out- or in-degree two. Let W be a cylindrical wall of order k. We say that W *grasps* a butterfly minor model μ of $\overleftrightarrow{K}_t$ if for every $v \in V(\overleftrightarrow{K}_t)$ there exists a pair $i_v, j_v \in [k]$ such that $V(Q_{i_v}) \cap V(P^{\mathsf{i}}_{j_v}) \subseteq V(\mu(v))$ or $V(Q_{i_v}) \cap V(P^{\circ}_{j_v}) \subseteq V(\mu(v))$, that is, the model of every vertex contains a branch vertex of W.

Definition 5.1.1 (Strip)**.** A *strip* between column i and j of a wall W, for $1 \leq i < j \leq k$, is the subgraph of W containing all rows $R_i, \dots, R_j$ and the subpaths $Q_1[R_i, R_j], \dots, Q_k[R_i, R_j]$. The *height* of such a strip is $j - i + 1$. ⊣

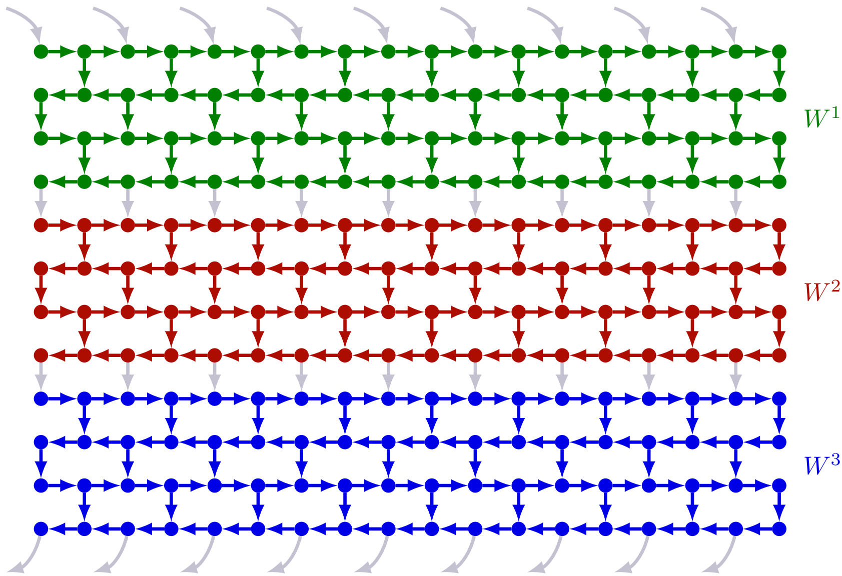

Figure 5.3: The wall divided into three strips W^1, W^2 and W^3.

The boundaries of the faces of a wall, except for the two faces containing C_1 and C_k, are called *bricks*.

Definition 5.1.2 (Slice)**.** Let $k \in \mathbb{N}$ be a positive integer and W be a cylindrical wall of order k. A *slice* W' of W is a cylindrical wall containing the vertical paths $Q_i, \dots, Q_{i+\ell}$ for all $i \in [k]$ and some $\ell \in [k-i]$, and the horizontal paths $P^1_1[Q_i, \dots, Q_{i+\ell}], \dots, P^2_k[Q_i, \dots, Q_{i+\ell}]$. We say that W' is the *slice of* W *between* Q_i *and* $Q_{i+\ell}$ and that W' is of *width* $\ell + 1$. ⊣

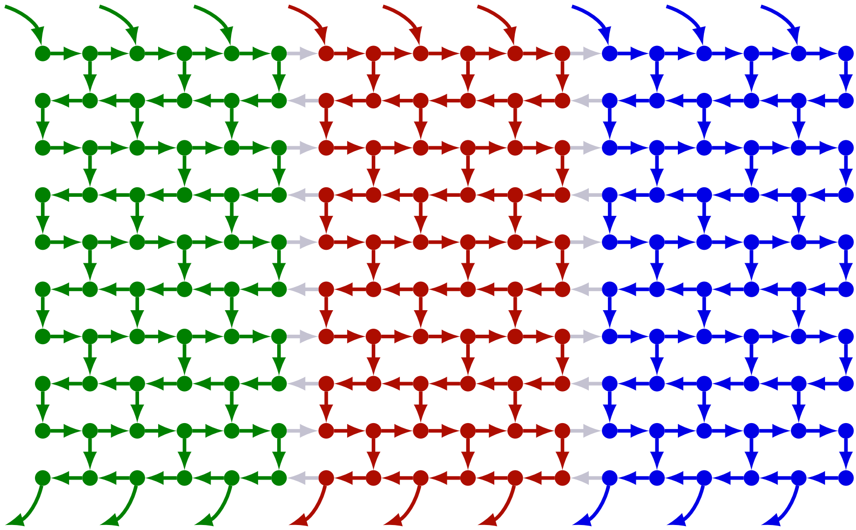

Figure 5.4: The wall divided into three slices W_1, W_2 and W_3.

We define a distance measure with respect to a cylindrical wall based on how many paths lie between two given vertices.

Definition 5.1.3 (W-Distance)**.** Let $k \in \mathbb{N}$ be a positive integer and W be a cylindrical wall of order k. Given two vertices $u, v \in V(W)$, we say that they have W*-distance* at least i if there exist i distinct vertical or i distinct horizontal paths whose removal separates u and v in W. ⊣

Most of these definitions can naturally be used for cylindrical grids as well and we do so at times. Cylindrical grids and walls are closely related. From a wall one can obtain a cylindrical grid of the same order by butterfly contraction. From a cylindrical grid one can obtain a wall by deleting subpaths of the in- and out-paths alternatingly, thereby loosing a factor two in the order. Therefore, we can obtain the following theorem using the function $f_{\mathsf{wall}} : \mathbb{N} \to \mathbb{N}$ obtained from the function f_{grid} provided by Theorem 2.2.2: f_{wall}

$$f_{\mathsf{wall}}(k) := 2f_{\mathsf{grid}}(k)\,.$$

Theorem 5.1.4 (Kawarabayashi and Kreutzer, 2015 [KK15])**.** There is a function $f_{\mathsf{wall}} : \mathbb{N} \to \mathbb{N}$ such that every digraph D either satisfies $\mathsf{dtw}(D) \leq k$, or contains the cylindrical wall of order $f_{\mathsf{wall}}(k)$ as a butterfly minor.

The flat wall theorem for directed graphs by Giannopoulou et al. [GKKK20] ensures the existence of a wall with a bounded number of cross-rows when excluding a large clique as a minor. We prove a statement for digraphs excluding a cross-row grid instead, which allows us to get rid of all forward crosses in the wall by deleting a small apex set. We use the following definitions that were introduced by Giannopoulou et al.

Definition 5.1.5 (Triadic Partitions)**.** Let $k \in \mathbb{N}$ be a positive integer and W be a cylindrical wall of order $3k$. The *triadic partition* of W is the tuple

$$\mathcal{W} = \left(W, k, W_1, W_2, W_3, W^1, W^2, W^3\right)$$

such that for each $i \in [3]$, W_i denotes the slice of W between $Q_{k(i-1)+1}$, see Figure 5.4 and Q_{ik}, and W^i denotes the strip of W between the rows $k\,(i-1)+1$ and ik, see Figure 5.3. ⊣

We often consider paths starting and ending in a brick. In order to be able to construct paths to and from their start- and end-vertices, we need parts of the wall around the brick, thus there is the concept of tiles centred at a brick.

Definition 5.1.6 (Tiles)**.** For $i, j \in [k]$ and $d \geq 1$, the *tile* $\mathsf{T}_{i,j,d}$ of W is defined as the subgraph of W obtained by the union

$$\bigcup_{i \leq \ell \leq i+2d+1} Q_\ell[R_j, R_{j+2d+1}] \cup \bigcup_{j \leq \ell \leq j+2d+1} R_\ell[Q_i, Q_{i+2d+1}]\,.$$

We say i is the *column index* of $\mathsf{T}_{i,j,d}$, and j is the *row index* of $\mathsf{T}_{i,j,d}$. Also, d is the *width* of $\mathsf{T}_{i,j,d}$. See Figure 5.5 for an illustration.

The *perimeter* of $\mathsf{T}_{i,j,d}$ is given by

$$\mathsf{T}_{i,j,d} \cap \left(Q_i \cup Q_{i+2d+1} \cup P^{\mathsf{o}}_j \cup P^{\mathsf{i}}_{j+2d+1}\right).$$

We call Q_i the *left path* of the perimeter, Q_{i+2d+1} its *right path*, P^{o}_j the *upper path* of the perimeter, and finally P^{i}_{j+2d+1} its *lower path*.

The *corners* of a tile are the vertices $a, b, c, d \in V(\mathsf{T}_{i,j,d})$ where

- a, the *upper left corner*, is the common starting point of $\mathsf{T}_{i,j,d} \cap Q_i$ and $\mathsf{T}_{i,j,d} \cap P^{\mathsf{o}}_j$,
- b, the *upper right corner*, is the end of $\mathsf{T}_{i,j,d} \cap P^{\mathsf{o}}_j$ and the starting point of $\mathsf{T}_{i,j,d} \cap Q_{i+2d+1}$,

- c, the *lower left corner*, is the common end of $\mathsf{T}_{i,j,d} \cap Q_i$ and $\mathsf{T}_{i,j,d} \cap P^{\mathsf{i}}_{j+2d+1}$, and
- d, the *lower right corner*, is the end of $\mathsf{T}_{i,j,d} \cap Q_{i+2d+1}$ and the starting point of $\mathsf{T}_{i,j,d} \cap P^{\mathsf{i}}_{j+2d+1}$.

The *centre* of $\mathsf{T}_{i,j,d}$ is the boundary of the unique brick $\mathsf{C}_{\mathsf{T}_{i,j,d}}$ of W whose boundary consists of vertices from Q_{i+d+1}, Q_{i+d+2}, P^{i}_{j+d+1}, and P^{o}_{j+d+2}. All vertices of $\mathsf{T}_{i,j,d}$ which are not in the centre and not on the perimeter of $\mathsf{T}_{i,j,d}$ are called *internal*. $\dashv$

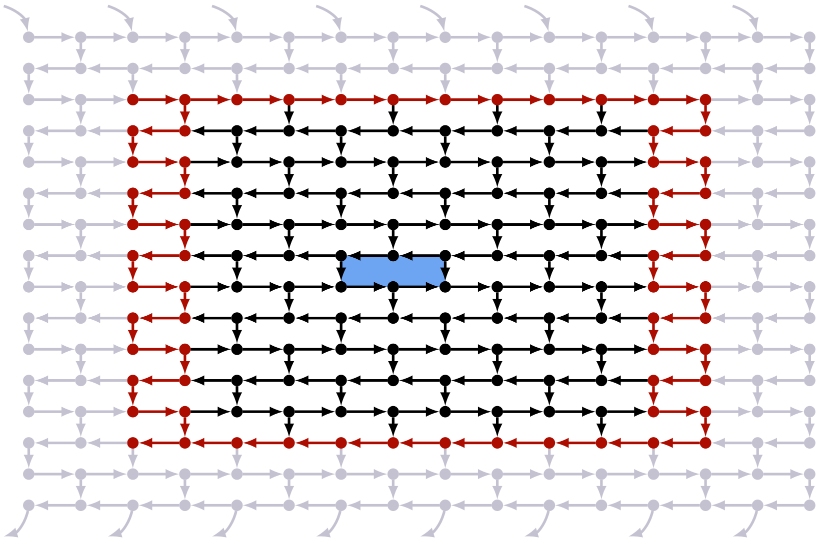

Figure 5.5: A tile of of width two in a wall of order eight. The centre brick is filled in blue and the perimeter is drawn in red.

Please note that by this definition, only bricks lying between P^{i}_i and P^{o}_{i+1} for some $i \in [k]$ can be the centre of a tile. However, if we take the mirror image of the unique embedding of the wall along Q_1, we obtain a new embedding, switching in- and out-paths. We call these two possible embeddings the two *parametrisations* of a wall, we write $\pi(W) = \mathsf{o}$, if the cycles encounter the out-path of a row before the in-path $\pi(W)$ and $\pi(W) = \mathsf{i}$ otherwise. This means that we can define for every brick F of W a tile T_F such that F is the centre of T_F.

Giannopoulou et al. use the following concept of flatness.

Definition 5.1.7 (weakly flat)**.** Let D be a digraph, W a wall in D, $a : \mathbb{N} \to \mathbb{N}$ a function and $t \in \mathbb{N}$. The wall W is $a(t)$-*weakly flat* with directed treewidth bounded by d if there is a subset $A \subseteq V(D)$ with $|A| \leq a(t)$ such that W is a wall in $D - A$ and the following hold:

(i) There is an undirected separation (X, Y) in D such that $A \cup \mathsf{per}(W) = X \cap Y$ and $\mathsf{int}(W)$ is contained in $X \setminus Y$, as well as every vertex in $X \setminus Y$ reaches a vertex in $\mathsf{int}(W)$ or is reachable from it.

(ii) For every path Q in $D - A$ that has both endpoints in $\mathsf{int}(W)$ but is internally disjoint from $\mathsf{int}(W)$ there is a brick B in W such that the boundary of B contains both endpoints of Q. Moreover, for every brick B of W let V_B be the set of vertices of $D - A$ that appear as internal vertices of a path Q with both endpoints on B but internally disjoint from B. The strongly connected components of $D[C]$, which are called the *extensions* of B, have directed treewidth at most d.

(iii) If T is a tile of width five of W and c_1 is its upper left corner, c_2 its upper right corner, d_1 its lower left corner and d_2 its lower right corner, then there are no two paths Q_1 and Q_2 in $(D - A) - (W - \mathsf{T})$ such that Q_1 is a c_1-d_2-path and Q_2 is a c_2-d_1-path.

For a function $G : \mathbb{N} \to \mathbb{N}$, the wall W is $g(t)$-*nearly*-$a(t)$-*weakly flat* with directed treewidth bounded by d if it satisfies (i) and (ii) and there are at most $g(t)$ rows whose tiles do not satisfy (iii), we also refer to such a row as *cross-row*. ⊣

Giannopoulou et al. obtain two different flat wall theorems for directed graphs. The first excludes a transitive tournament in order to get rid of cross-rows entirely. The second only excludes the clique, thus cross-rows may still exist.

Theorem 5.1.8 (Directed weakly flat wall theorem)**.** There exist functions $d \colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ and $a \colon \mathbb{N} \to \mathbb{N}$ such that for every directed graph D and all $k, t \in \mathbb{N}$ one of the following is true:

(i) $\mathsf{dtw}(D) < d(k, t)$,

(ii) D contains a tournament of order t as a butterfly minor

(iii) there is a directed cylindrical wall W of order k in D that is $a(t)$-weakly flat with directed treewidth bounded by $d(k, t)$.

Theorem 5.1.9 (Directed nearly-weakly flat wall theorem)**.** There exist functions $d \colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, $a \colon \mathbb{N} \to \mathbb{N}$ and $g \colon \mathbb{N} \to \mathbb{N}$ such that for every directed graph D and all $k, t \in \mathbb{N}$ one of the following is true:

(i) $\mathsf{dtw}(D) < d(k,t)$,

(ii) D contains $\overleftrightarrow{K}_t$ as a butterfly minor

(iii) there is a directed cylindrical wall W of order k in D that is $g(t)$-nearly-$a(t)$-weakly flat with directed treewidth bounded by $d(k,t)$.

Both of these theorems have their drawbacks. Theorem 5.1.9 only finds a wall that still contains cross-rows. So, it still contains highly non-planar behaviour. Theorem 5.1.8 ensures a wall that does not contain any cross-rows, but the excluded structure of a transitive tournament is in itself not strongly connected. So, the transitive tournament is of directed treewidth one. This means that any structure theorem using the transitive tournament as the excluded butterfly minor does not provide sufficiency, that is, every graph having the described structure also excludes the transitive tournament as a butterfly minor.

5.2 The cross-row grid

Here, we consider a structure lying in between the transitive tournament and the clique: the *cross-row grid*. It lies in between the two other structures in the sense that every transitive tournament is a butterfly minor of a cross-row grid and every cross-row grid is a butterfly minor of a clique.

Let D be cylindrical grid of order k with cycles $C_1, \dots, C_k$ and in-paths $P_1^{\mathsf{i}}, \dots, P_k^{\mathsf{i}}$ and out-paths $P_1^{\mathsf{o}}, \dots, P_k^{\mathsf{o}}$. Let $a_1, \dots, a_k$ be the vertices of P_1^{o} and $b_1, \dots, b_k$ be the vertices of P_1^{i} in order of occurrence along the path. A *cross-row grid* of order k, written $D_k^{\boxtimes}$, is obtained from D by adding the edges $\{(a_j, b_{k-j}) \mid 1 \leq j \leq k-1\} \cup \{\big(a_j, b_{k-(j-2)}\big) \mid 2 \leq j \leq k\}$. See Figure 5.6 for an example.

We start by proving that every transitive tournament is contained in a large enough cross-row grid as a butterfly minor. In order to construct such a transitive tournament minor we introduce the following tool. Let D be a digraph with $2t$ designated vertices $(i_1, \dots, i_t, o_1, \dots, o_t)$. If for every $s \in [t-1]$ there exist t disjoint paths $P_1, \dots, P_t$ in D such that P_s is an i_s-o_{s+1}-path, P_{s+1} is an i_{s+1}-o_s-path and for all $s' \notin \{s, s+1\}$ the path $P_{s'}$ is an $i_{s'}$-$o_{s'}$-path, then D is called a $(i_1, \dots, i_t)$-$(o_1, \dots, o_t)$-*switch*.

We define the function

$$f_{\mathsf{switch}}(t) := \frac{1}{2}\left(t^2 - t\right) - 2,$$ $f_{\mathsf{switch}}(t)$

which gives for every number t the number of switches necessary to a build transitive tournament minor on t vertices.

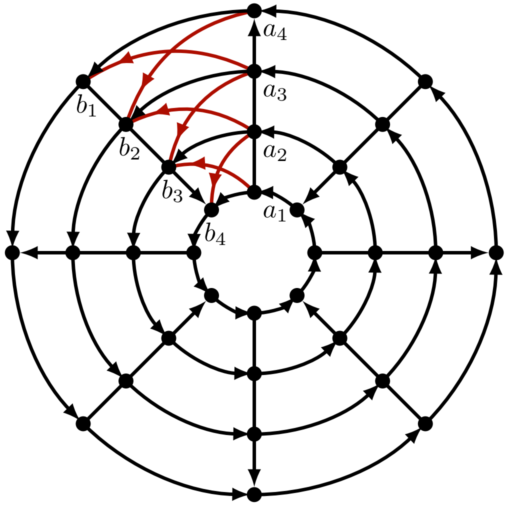

Figure 5.6: A cross-row grid of order four.

Lemma 5.2.1. Let D be a digraph that contains $f_{\mathsf{switch}}(t)$ disjoint subgraphs $D_1, \dots,$ $D_{f(t)}$ such that D_j is an $(i_1^j, \dots, i_t^j)$-$(o_1^j, \dots, o_t^j)$-switch. For every $1 \leq j < f_{\mathsf{switch}}(t)$ let $\mathcal{L}_j = (P_1^j, \dots, P_t^j)$ be an $(o_1^j, \dots, o_t^j)$-$(i_1^{j+1}, \dots, i_t^{j+1})$-linkage with P_s^j is an o_s^j-i_s^{j+1}-path and let $\mathcal{L}_0 = \left(P_1^0, \dots, P_t^0\right)$ be a linkage ending in the vertices $\left(i_1^1, \dots, i_t^1\right)$ and $\mathcal{L}_{f_{\mathsf{switch}}(t)+1} = (P_1^{f_{\mathsf{switch}}(t)+1}, \dots, P_t^{f_{\mathsf{switch}}(t)+1})$ be a linkage starting in the vertices $(i_1^{f_{\mathsf{switch}}(t)}, \dots, i_t^{f_{\mathsf{switch}}(t)})$ such that $\mathcal{L}_0, \dots, \mathcal{L}_{f_{\mathsf{switch}}(t)}$ are internally disjoint from the switches. Also, let $\mathcal{R}_i := \{R_i^{a,a+1} \mid 1 \leq a < t\}$ for all $1 \leq i \leq f_{\mathsf{switch}}(t)$ such that $R_i^{a,a+1}$ starts on the a-th path and ends on the $a+1$-th path of $\mathcal{L}_i$ and is disjoint to the switches and internally disjoint to the linkages $\mathcal{L}_0, \dots, \mathcal{L}_{f_{\mathsf{switch}}(t)}$. Additionally, we require that the linkages $\mathcal{L}_1, \dots, \mathcal{L}_{f_{\mathsf{switch}}(t)}$ are pairwise disjoint. Then, D contains a $\vec{K}_t$ as a butterfly minor.

Proof. We build butterfly contractible models for every vertex of $\vec{K}_t$. Let $v_1, \dots, v_t$ be the vertices of $\vec{K}_t$ such that all edges are of the form (v_i, v_j) with $i < j$. The model of v_k consists of a path M_k, $k-1$ edges $e_1^{\mathsf{in}(k)}, \dots, e_{k-1}^{\mathsf{in}(k)}$ and $t-k$ edges $e_{k+1}^{\mathsf{out}(k)}, \dots, e_t^{\mathsf{out}(k)}$ such that the vertices $\mathsf{head}(e_1^{\mathsf{in}(k)}), \dots, \mathsf{head}(e_{k-1}^{\mathsf{in}(k)}), \mathsf{tail}(e_{k+1}^{\mathsf{out}(k)}),$ $\dots, \mathsf{tail}(e_t^{\mathsf{out}(k)})$ occur on M_k in that order, see Figure 5.7 for an illustration.

We construct $M_1, \dots, M_t$ inductively. We also define a function f_k helping us to keep track of the path we construct. The path M_k^1 is defined as P_k^0, which ends in the vertex $i_{f_k(1)}^1$, and we define $f_k(1) := k$ for all $1 \leq k \leq t$.

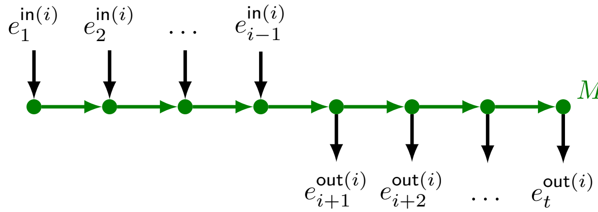

Figure 5.7: The butterfly minor model for a vertex of the transitive tournament.

Now assume we have constructed a path M_k^j ending in $i^j_{f_k(j)}$ for all $1 \leq k \leq t$. Let

$$x := \begin{cases} j \bmod (t-1) & \text{, if } j \bmod (t-1) < 0 \\ t-1 & \text{, if } j \bmod (t-1) = 0. \end{cases}$$

We use D_j to switch the paths arriving at i^j_x and i^j_{x+1}. So let $\mathcal{Q}^j$ be the $(i^j_1, \dots, i^j_t)$-$(o^j_1, \dots, o^j_t)$-linkage within D_j that contains an i^j_y-o^j_y-path for every $y \notin \{x, x+1\}$ and an i^j_x-o^j_{x+1}-path as well we an i^j_{x+1}-o^j_x-path. Obtain M_k^{j+1} from M_k^j by appending the path Q from $\mathcal{Q}^j$ starting in $\mathsf{end}(M_k^j)$ and then the path $P^j_{k'} \in \mathcal{L}_{j+1}$ starting in $\mathsf{end}(Q)$. Define $f_k(j+1) := k'$.

Finally, we define $M_k := M_k^{f_{\mathsf{switch}}(t)}$. As $M_1, \dots, M_t$ each contain exactly one path from each linkage $\mathcal{L}_1, \dots, \mathcal{L}_t, \mathcal{Q}_1, \dots, \mathcal{Q}_t$ they yield a $(i^1_1, \dots, i^1_t)$-$(o_1^{f_{\mathsf{switch}}(t)}, \dots, o_t^{f_{\mathsf{switch}}(t)})$-linkage.

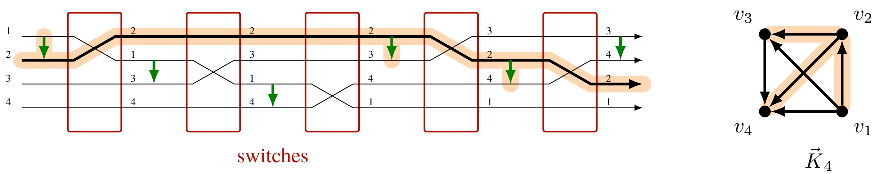

Figure 5.8: An example how the construction looks for finding a $\vec{K}_4$ butterfly minor in the given graph, the model for the vertex v_2 is highlighted in orange.

We divide the linkages $\mathcal{L}_0, \dots, \mathcal{L}_{f_{\mathsf{switch}}(t)+1}$ into $t-1$ consecutive blocks $\mathcal{B}_1, \dots, \mathcal{B}_{t-1}$ such that $\mathcal{B}_i$ contains $t-i$ linkages. Let $\mathcal{B}_i = \left(\mathcal{L}_{i_1}, \dots, \mathcal{L}_{i_{t-i}}\right)$ and $\mathcal{R}_{i_j} := \{R^{a,a+1}_{i_j} \mid 1 \leq a < t\}$. We consider the path $R^{f_i(i_j), f_i(i_j)}$, which starts in M_i and ends in M_{i+j} due to our construction earlier. We add its first edge as e^{out}_{i+j} to the model of v_i and its last edge as e^{in}_i to the model of v_{i+j}. The remaining path is the model of $(i, i+j)$.

As we collect all the in-edges for the model of v_k within the blocks $\mathcal{B}_1, \dots, \mathcal{B}_{k-1}$ and only collect the out-edges in block $\mathcal{B}_k$, all in-edges are met by M_k before the out-edges, making the model of v_i butterfly contractible. See Figure 5.8 for an illustration of the construction. □

Now, we can prove that for every transitive tournament there is a sufficiently large cross-row grid containing it as a butterfly minor.

Lemma 5.2.2. The cross-row grid of order $f_{\mathsf{switch}}(t) \cdot t + 2$ contains $\vec{K}_t$ as a butterfly minor.

Proof. We define $k := f_{\mathsf{switch}}(t) \cdot t + 2$. Let $a_1, \dots, a_k$ be the vertices of P_1^{o} and $b_1, \dots, b_k$ be the vertices of P_1^{i} in order of appearance along the paths. We define $f_{\mathsf{switch}}(t)$ many subgraphs $D_1, \dots, D_{f_{\mathsf{switch}}(t)}$, where D_i is the subgraph of D containing the vertices $a_{k-(i-1)t}, \dots, a_{k-it+1}, b_{(i-1)t}, \dots, b_{it}$ and all edges between them. Note that D_i is a $\big(a_{k-(i-1)t}, \dots, a_{k-it+1}\big)$-$\big(b_{(i-1)t}, \dots, b_{it}\big)$-switch.

We choose $\mathcal{L}_0$ to be the subpaths of $C_k, \dots, C_{k-t+1}$ starting in P_k^{i} and ending in the vertices $a_k, \dots, a_{k-t+1}$. For $\mathcal{L}_i$ with $1 \leq i \leq f_{\mathsf{switch}}(t)$ we construct paths from $(b_{(i-1)t}, \dots, b_{it})$ along $(C_{(i-1)t}, \dots, C_{it})$ until meeting $(P^{\mathsf{i}}_{k-1-(i-1)t}, \dots, P^{\mathsf{i}}_{k-it})$, then following the paths $(P^{\mathsf{i}}_{k-1-(i-1)t}, \dots, P^{\mathsf{i}}_{k-it})$ until meeting the cycles $(C_{k-(i)t}, \dots, C_{k-(i+1)t+1})$ which the paths of $\mathcal{L}_i$ then follow until finally reaching $(a_{k-(i)t}, \dots, a_{k-(i+1)t+1})$. Next, we define the linkage $\mathcal{L}_{f_{\mathsf{switch}}(t)+1}$ as the subpaths of $(C_{(f_{\mathsf{switch}}(t)-1)t}, \dots, C_{f_{\mathsf{switch}}(t)t})$ starting in $(b_{(f_{\mathsf{switch}}(t)-1)t}, \dots, b_{f_{\mathsf{switch}}(t)t})$ and ending on P_2^{i}.

Finally we construct families $\mathcal{R}_i$ for all $1 \leq i \leq f_{\mathsf{switch}}(t) + 1$. For $1 < i \leq f_{\mathsf{switch}}(t) + 1$ we obtain the subpaths between the paths in $\mathcal{L}_i$ from P_2^{i}. And for $i = 1$ we obtain them from P_k^{i}.

Then, we have everything we need to apply Lemma 5.2.1 and thus, obtain the desired $\vec{K}_t$ as a butterfly minor. □

Clearly, for every cross-row grid there is bidirected clique containing it as a minor.

Lemma 5.2.3. The bidirected clique $\overleftrightarrow{K}_{2\cdot t^2}$ contains the cross-row grid $D_t^{\boxtimes}$ as a butterfly minor.

Proof. This directly follows from the cross-row grid of order t having $2t^2$ vertices. We fix a bijection $\varphi : V(\overleftrightarrow{K}_{2\cdot t^2}) \to V(D_t^{\boxtimes})$ and then delete from $\overleftrightarrow{K}_{2\cdot t^2}$ every edge (u, v) with $(\varphi(u), \varphi(v)) \notin E(D_t^{\boxtimes})$. This yields the desired butterfly minor. □

Basically, a cross-row grid consists of many slices in a wall that each contain a local cross. We formalise this concept by the following definition.

Definition 5.2.4. A *cross* in a slice of a wall W between Q_i and $Q_{i+\ell}$ are two paths J_1 and J_2 with $\mathsf{start}(J_1), \mathsf{end}(J_2) \in V(Q_i)$ and $\mathsf{start}(J_2), \mathsf{end}(J_1) \in V(Q_{i+\ell})$ and for the rows R_a containing $\mathsf{start}(J_1)$, $R_{a'}$ containing $\mathsf{end}(J_1)$, R_b containing $\mathsf{start}(J_2)$ and $R_{b'}$ containing $\mathsf{end}(J_2)$ we have $\max\{a, a'\} - \min\{b, b'\} > 1$. We call a slice with a cross in it a *cross slice*. $\dashv$

In order to construct the cross-row grid as a minor in a given digraph it suffices to show that there is a wall in the graph that contains enough cross slices and still has some rows that remain untouched by the crosses.

Lemma 5.2.5. Let D be a digraph, W a cylindrical wall in D and $i \in \mathbb{N}$. If W contains t disjoint slices which each contain a cross, and there are t rows in W such that all t crosses lie in a strip not containing these t rows, then W contains $D_t^{\boxtimes}$ as a butterfly minor.

Proof. Let $S_i, \ldots, S_t$ be t slices such that S_i is the slice between Q_{j_i} and $Q_{j_i+\ell_i}$ and contains a cross J_1^i and J_2^i. Let W_R be the strip containing all crosses but not the rows $R'_1, \ldots, R'_t$ and let R_a and R_b such that W_R is the strip between R_a and R_b. We construct a new row containing t crosses by valid butterfly minor operations within the subwall between P_a^{o} and P_b^{i} as follows. First, we delete all in- and out-paths within W_R except for P_a^{o} and P_b^{i}. Next, for every slice S_i we can contract the following paths into single vertices:

1. the subpath of Q_{j_i} starting on P_a^{o} and ending in $\mathsf{start}(J_1^i)$, we refer to this contraction vertex as $\mathsf{start}(J_1^i)'$,
2. the subpath of Q_{j_i} starting in $\mathsf{end}(J_2^i)$ and ending on P_b^{i}, we refer to this contraction vertex as $\mathsf{end}(J_2^i)'$,
3. the subpath of $Q_{j_i+\ell_i}$ starting on P_a^{o} and ending in $\mathsf{start}(J_2^i)$, we refer to this contraction vertex as $\mathsf{start}(J_2^i)'$, and
4. the subpath of $Q_{j_i+\ell_i}$ starting in $\mathsf{end}(J_1^i)$ and ending on P_b^{i}, we refer to this contraction vertex as $\mathsf{end}(J_1^i)'$.

This leaves us with one row consisting of out-path P_a^{o} and in-path P_b^{i} containing t crosses. In order to make this a proper cross-row grid, we delete all vertical paths but $Q_{j_1}, \ldots, Q_{j_t}$ and $Q_{j_t+\ell_t}$. Now, for every $i \in [t]$, we can contract the subpath of P_a^{o} that starts in $\mathsf{start}(J_2^i)'$ and ends in $\mathsf{start}(J_1^{i+1})'$ into a single vertex. Similarly, for every $i \in [t]$, we can contract the subpath of P_b^i that starts in $\mathsf{end}(J_2^{i+1})'$ and ends in $\mathsf{end}(J_1^i)'$ into a single vertex.

This newly constructed row together with the rows $R_1', \dots, R_t'$ and the vertical paths $Q_{j_1}, \dots, Q_{j_t}$ and $Q_{j_t+\ell_t}$ now yields the claimed $D_t^{\nwarrow\!\!\!\swarrow}$ butterfly minor. $\square$

5.3 Drawings and weak renditions

Now that we have seen how the cross-row grid relates to the traditional excluded minors, we give some context for surfaces and embeddings. Formally, a *surface* is a compact 2-dimensional manifold with or without a boundary. However, the only two surfaces we explicitly use in this chapter are the sphere $\mathbb{S}^2$ and the torus.

Definition 5.3.1. A *drawing (with crossings)* of a digraph D on a surface Σ is a tuple $\Gamma = (U, V, E, \varphi)$ such that

- $\varphi : V \cup E \to V(D) \cup V(E)$ is a bijection such that $\varphi|_E$ is a bijection between E and $E(D)$ and $\varphi|_V$ is a bijection between V and $V(D)$
- $V \subseteq U \subseteq \Sigma$ and $V \cup \bigcup_{e \in E} e = U$
- for every $e \in E$, $e = h_e((0,1))$, where $h_e : [0,1] \to U$ is a homeomorphism onto its image with $h_e(0), h_e(1) \in V$
- V is disjoint from every $e \in E$
- $(u, v) \in E(D)$ if and only if $\varphi(h_e(0)) = u$ and $\varphi(h_e(1)) = v$
- if $e, e' \in E$ are distinct, then $e \cap e'$ is finite.

If $e, e' \in E$ with $e \cap e' \neq \emptyset$ and $e \neq e'$, then we say e and e' *cross*. A drawing is *cross-free* if there are no two elements in E that cross. $\dashv$

We say a drawing $\Gamma_1 = (U_1, V_1, E_1, \varphi_1)$ of a digraph D is *consistent* with a drawing $\Gamma_2 = (U_2, V_2, E_2, \varphi_2)$ of a subgraph $D' \subseteq D$ if $U_2 \subseteq U_1$, $V_2 \subseteq V_1$, $E_2 \subseteq E_1$ and $\varphi_1(v) = \varphi_2(v)$ for all $v \in V(D')$.

We often consider drawings of parts of some digraph into a closed disk, in that case we are interested in which vertices are drawn into the boundary in order to describe interaction with the remaining graph. To this end we introduce the following definition, which basically generalises the concept of *societies* by Kawarabayashi, Thomas and Wollan [KTW20].

Definition 5.3.2 (Society)**.** Let Ω be a cyclic order of the elements of some set and let $V(\Omega)$ denote this set. A *society* is a pair (D, Ω), where D is a digraph, and Ω is a cyclic order with $V(\Omega) \subseteq V(D)$. (D, Ω)

A *cylindrical society* is a tuple (D, Ω_1, Ω_2), where D is a digraph and Ω_1, Ω_2 are cyclic orders with $V(\Omega_i) \subseteq V(D)$ for both $i \in [2]$ and $V(\Omega_1) \cap V(\Omega_2) = \emptyset$. (D, Ω_1, Ω_2)

A subset $X \subseteq V(\Omega)$ is called a segment if there are no vertices $x_1, x_2 \in X$ and $y_1, y_2 \in V(\Omega)$ such that x_1, y_1, x_2, y_2 occur in Ω in that order. The ordering Ω naturally induces a linear ordering on its segments. We write $a\Omega b$ for the unique segment of Ω that has a as its first vertex and b as its last vertex. $\dashv$ $a\Omega b$

The following definition helps us to describe the structure inside a society, especially non-planar structure.

Definition 5.3.3 (Transactions)**.** Let (D, Ω) be a society. A path P is an Ω*-path* if $V(P) \cap V(\Omega) = \{\mathsf{start}(P), \mathsf{end}(P)\}$, that is, P is a $V(\Omega)$-path. A linkage $\mathcal{P}$ in D is called a *transaction* in (D, Ω) if every $P \in \mathcal{P}$ is an Ω-path and there are two disjoint segments X and Y of Ω such that $\{\mathsf{start}(P) \mid P \in \mathcal{P}\} \subseteq X$ and $\{\mathsf{end}(P) \mid P \in \mathcal{P}\} \subseteq Y$. The *depth* of (D, Ω) is defined as the maximum order of a transaction in (D, Ω).

The two endpoints $\mathsf{start}(P)$ and $\mathsf{end}(P)$ of an Ω-path P split Ω into two segments. If another Ω-path P' has its start-vertex in the one and its end-vertex in the other segment, then P and P' build a *cross* in (D, Ω). A transaction is called *planar* if no two paths in it build a cross. $\dashv$

For societies in undirected graphs having no cross and being 4-connected suffices to ensure a planar embedding of the graph into a disk. This is not sufficient in digraphs. Still, we would like to reduce the non-planarity of the graphs to smaller, more controllable regions. The next definition allows us to fix a part of a given graph that allows for a drawing without crossings in the sphere.

Definition 5.3.4 (Skeleton)**.** Let D be a digraph, $r \in \mathbb{N}$ and W be a wall of order r in D. A skeleton of D is a tuple $\mathcal{S} = (\Gamma, W, \mathsf{L})$ with

1. L is a family of linkages in D,
2. Γ is a planar drawing of the graph $D_{\mathcal{S}} \coloneqq W \cup \bigcup_{\mathcal{L} \in \mathsf{L}} \mathcal{L}$,

3. if $\mathsf{L} = \emptyset$, then the two faces bound by the wall perimeters are called the *big faces* of $\mathcal{S}$, otherwise there exists a linkage $\mathcal{L} \in \mathsf{L}$ such that

$$\mathcal{S}' = \left(\Gamma|_{W \cup (\mathsf{L} \setminus \{\mathcal{L}\})}, W, (\mathsf{L} \setminus \{\mathcal{L}\})\right)$$

is a skeleton and $\mathcal{L}$ is a planar transaction on a big face f of $\mathcal{S}'$. In the ladder case, the *big faces* of $\mathcal{S}$ are the big faces of $\mathcal{S}'$ without f but adding all faces bound by f and one of the two outer paths of $\mathcal{L}$.

$D_{\mathcal{S}}$ We call the skeleton $\mathcal{S}$ *centred at the wall* W. We call $D_{\mathcal{S}}$ the *rigid subgraph* of $\mathcal{S}$. The *order* of $\mathcal{S}$ is half the minimum order of an element in $\mathsf{L} \cup \{W\}$. ⊣

Intuitively, in order to describe that non-planar behaviour of the digraph around the rigid subgraph of a skeleton is restricted, we want to say that closed curves in the rigid part have a separating property. We cannot demand them to yield proper separations, so we describe a slightly weaker notion where we cut out the curve together with some part of the rigid graph around it.

A *noose* of a digraph D within a fixed drawing Γ into a surface Σ is a closed curve that bounds a disk and only intersects Γ in vertices of D.

Definition 5.3.5. Let $\mathcal{S} := (\Gamma, W, \mathsf{L})$ be a skeleton of a digraph D in the sphere. Let Γ^+ be a drawing of D in Σ that is consistent with Γ and C a noose of $D_{\mathcal{S}}$ within Γ^+. Then, the graph $D_{\mathcal{S}} - C$ consists of at most two weakly 2-connected components H_1' and H_2'. For $i \in [2]$, let H_i be the weakly 2-connected component not containing C obtained by removing the face of H_i' that is no face in Γ from $D_{\mathcal{S}}$. We define H_i^+ to be the subgraph of D induced by $V(H_i) \cup \bigcup_{f \text{ face in } \Gamma^+|_{H_i} \text{ and in } \Gamma^+} \{v \mid$ the vertex v is drawn into f by $\Gamma^+\}$. ⊣

H_i^+

Using this we can define our demands on a drawing of the digraph around a wall in the sphere.

Definition 5.3.6 (Σ-decomposition)**.** Let $\mathcal{S} := (\Gamma, W, \mathsf{L})$ be a skeleton of a digraph D in the surface Σ and $A \subseteq V(D)$ a set of vertices in D such that $V(D_{\mathcal{S}}) \subseteq V(D) \setminus A$. Let Γ^+ be a drawing of the strongly connected component of $D - A$ containing $D_{\mathcal{S}}$ in Σ that is consistent with Γ. The tuple $\rho = (\Gamma^+, \mathcal{S})$ is a Σ*-decomposition* of D in Σ with apex set A, if For all undirected cycles C in $D_{\mathcal{S}}$ there is no directed path from H_1^+ to H_2^+ as well as no directed path from H_2^+ to H_1^+.

Let $N(\rho)$, the set of *nodes* of ρ, be the set of all vertices in $D_{\mathcal{S}}$. If $\Gamma = (U, V, E, \varphi)$, we refer to φ^{-1} by φ_ρ.

Let Σ and Σ' be two surfaces, where Σ' has strictly higher genus. Let $\rho = (\Gamma^+, \mathcal{S})$ be a Σ-decomposition and $\rho' = (\Gamma'^+, \mathcal{S}')$ be a Σ'-decomposition such that $D_{\mathcal{S}} = D_{\mathcal{S}'}$. If Γ'^+ is consistent with Γ^+, then we say ρ' is *centred at* ρ. $\dashv$

There are certain crosses we cannot forbid by excluding a cross-row grid, so we need a notion of flatness that accommodates these crosses. In order to achieve this, we define renditions into a disk that internally are nearly planar except for small exceptions which cover the mentioned crosses.

Definition 5.3.7 (Weak rendition)**.** Let (D, Ω) be a society and Δ be a closed disk, possibly with an open hole. A weak rendition of (D, Ω) into Δ is a $\mathbb{S}^2$-decomposition $\rho = (\Gamma^+, \mathcal{S} = (\Gamma, W, \emptyset))$ such that

1. $V(\Gamma^+) \subseteq \Delta$,
2. one of the cyclic orderings of $\mathsf{bd}(\Delta)$ maps to the image of $\varphi_\rho(N(\rho) \cap \mathsf{bd}(\Delta)) = V(\Omega)$.

Now, let Δ' be obtained from a closed disk Δ'' by removing an open disk disjoint from the boundary of Δ'' and let (D, Ω_1, Ω_2) be a cylindrical society. Let B_1 and B_2 be the two closed curves in Δ' whose union equals $\mathsf{bd}(\Delta')$. Observe that for each $i \in [1, 2]$ the curve B_i bounds a closed disk Δ'_i with an open hole that contains Δ'. A *weak rendition in* Δ' is a $\mathbb{S}^2$-decomposition $(\Gamma^+, \mathcal{S} = (\Gamma, W, \emptyset, \emptyset))$ such that for both $i \in [2]$, δ is a cylindrical rendition of (D, Ω_i) in the disk Δ'_i. We say that (D, Ω_1, Ω_2) has a *weak rendition in the disk* if there exists Δ' as above such that (D, Ω_1, Ω_2) has a weak rendition in Δ'. $\dashv$

These tools enable us to describe drawings of digraphs that are consistent with the unique embedding of some wall they contain and additionally does have restricted non-planar behaviour with respect to this wall.

5.4 Flat wall theorem

Having introduced weak renditions we can define our notion of flatness, which is slightly different from the notion of Definition 5.1.7. In particular, it uses two directed separations instead of one undirected separation and it uses our concept of weak renditions, that is, removing a noose in the wall separates the two remaining parts of the wall.

Definition 5.4.1 (Bridge)**.** Let H be a subgraph of a digraph D. A weakly connected component C in $D - H$ is called an *H-bridge* if

1. there is a non-empty set $O \subseteq V(C)$ such that for every $o \in O$ there is a vertex $h \in H$ with $(o, h) \in E(D)$, h is called an *out-attachment*,
2. there is a non-empty set $I \subseteq V(C)$ such that for every $i \in I$ there is a vertex $h \in H$ with $(h, i) \in E(D)$, h is called an *in-attachment*, and
3. for every vertex $x \in V(C)$ there is a directed path P in C such that $\mathsf{start}(P) \in I$, $\mathsf{end}(P) \in O$ and $x \in V(P)$. ⊣

Definition 5.4.2 (Flat wall)**.** Let D be a digraph, $A \subseteq V(D)$ be a set of vertices and $W \subseteq D - A$ be a wall of order $k+4$. We define $W^- \subseteq W$ to be the wall of order k such that W^- is disjoint from Q_1, Q_2, Q_{k+3} and Q_{k+4}. Moreover, we define the
$\mathsf{border}(W)$ *border* of W as $\mathsf{border}(W) := W - W^-$. Let D' be the strongly connected component
$\mathsf{compass}(W)$ of D containing W. The *compass* of W, written $\mathsf{compass}(W)$ is the union of the border of W and all W^--bridges in $D' - \mathsf{border}(W)$. We say that W is a *flat wall under* A if

(F1) $V(W) \cap A = \emptyset$,

(F2) there are two directed separations $\overrightarrow{(Y_1, X_1)}$ and $\overrightarrow{(X_2, Y_2)}$ in D' such that X_1 and X_2 both contain W and $Y_1 \cap X_1 = A \cup \mathsf{border}(W) = X_2 \cap Y_2$ and additionally in $D' - (A \cup \mathsf{border}(W))$ every vertex in X_1 is reachable from W^- and every vertex in X_2 reaches W^-,

(F3) the cylindrical society $(\mathsf{compass}(W), \Omega_1, \Omega_2)$ with the sets $V(\Omega_1) = V(Q_1)$ and $V(\Omega_2) = V(Q_{k+4})$ has a weak rendition $\rho = (\Gamma^+, \mathcal{S} = (\Gamma, W, \emptyset))$ in the disk,

(F4) D' has a torus decomposition centred at ρ.

In case $A = \emptyset$, we say that W is *flat*. ⊣

The main theorem of this chapter is a directed flat wall theorem that excludes the cross-row grid as a butterfly minor.

Theorem 5.4.3 (Directed flat wall theorem)**.** There exist functions $\rho_W : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ and $\alpha_W : \mathbb{N} \to \mathbb{N}$ such that for all integers $r, t \geq 1$ and all digraphs D that do not contain $D_t^{\bowtie}$ for every $\rho_W(r,t)$-wall W in D there exist a set $A \subseteq V(D)$ with $|A| \leq \alpha_W(t)$ and an r-wall $W' \subseteq W - A$ which is flat under A.

We refer to the functions ρ_W and α_W from Theorem 5.4.3 as such globally. Please note that this yields a in some respect stronger statement than Theorems 5.1.8 and 5.1.9, because it refers to every wall in the digraph. While Theorems 5.1.8 and 5.1.9 provide the existence of their respective flat wall in the given digraph D, our result Theorem 5.4.3 provides the flat wall inside *every* wall that is large enough. This way one can start out with any large enough wall in a given digraph and can be sure that it contains a flat wall, which is a desirable property especially for algorithmic usage. Additionally Theorem 5.4.3 obtains a wall that does allow for any cross-rows by excluding a strongly connected digraphs as a butterfly minor, thus combining the features of Theorems 5.1.8 and 5.1.9. We mostly follow the arguments and techniques from [GKKK20]. Some refinements are due to Gianopoulou and Wiederrecht [GW21]. In some places the proofs need adjustments propagating through non-trivial steps of the arguments, therefore, we present the whole proof containing parts from both [GKKK20, GW21].

Definition 5.4.4 (Tiling)**.** A *tiling* is a family of pairwise disjoint tiles, and a tiling is said to *cover* a subwall W' of W if every branch vertex of W' occurs in one of the tiles of the tiling. For every function $f_w : \mathbb{N} \to \mathbb{N}$ and all $\xi, \xi' \in [f_w(t) + 1]$ we define the tiling $\mathcal{T}_{W,k,f_w(t),\xi,\xi'}$ with tiles of width $f_w(t)$. In order to do so, we define two function, the *column function*

$$\mathsf{c}_{\xi,f_w}(p) := (k + 1 - \xi) + (p - 1)\,(2 \cdot f_w(t) + 1)\,,$$

and the *row function*

$$\mathsf{r}_{\xi',f_w}(q) := (1 + \xi') + (q - 1)\,(2 \cdot f_w(t) + 1)\,.$$

For both the column and the row function we omit f_w, ξ and ξ' from the indices if they are clearly provided by the context. Then, define

$$\begin{aligned}\mathcal{T}_{W,k,f_w(t),\xi,\xi'} := \Big\{ \mathsf{T}_{\mathsf{c}(p),\mathsf{r}(q),f_w(t)} \mid & 1 \leq p \leq \left\lceil \frac{k + \xi - 1}{2 \cdot f_w(t) + 1} + 1 \right\rceil, \\ & 1 \leq q \leq \left\lceil \frac{3k - (1 + \xi')}{2 \cdot f_w(t) + 1} + 1 \right\rceil \Big\}.\end{aligned}$$

See Figure 5.9 for an illustration. ⊣

Note that every tiling $\mathcal{T}_{W,k,f_w(t),\xi,\xi'}$ covers W_2. Moreover, every brick of W_2 that lies between the two paths of R_i for some $i \in [3k]$ is the centre of some tile T' of some tiling $\mathcal{T}' \in \mathcal{T}_{W,k,f_w(t),\xi,\xi'}$. Hence, if we switch the parametrisation of W, we are able to find in total $2\,(f_w(t) + 1)^2$ many tilings that cover W_2, and every brick of W_2 is the centre of some tile in one of these tilings. This number of tilings becomes relevant in the proof of Theorem 5.4.3.

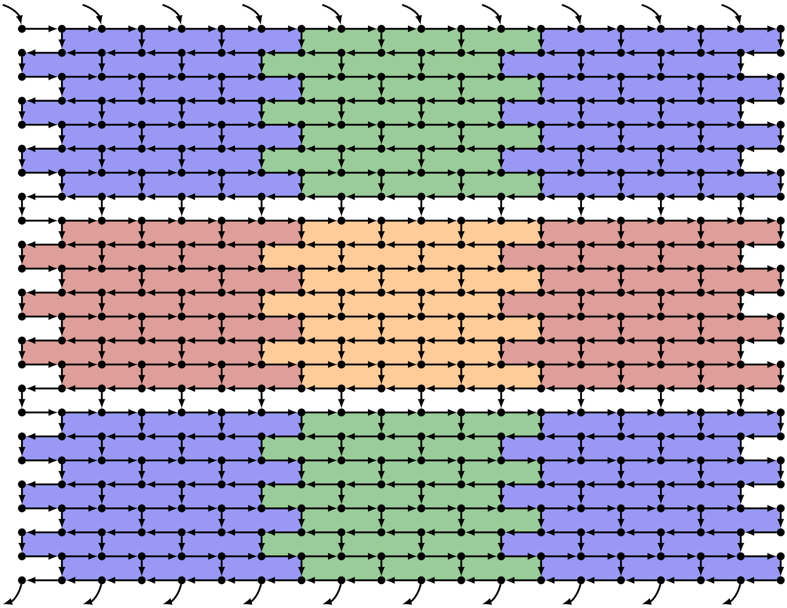

Figure 5.9: A tiling of a wall together with a four-colouring of it. We also see that the yellow tile in the middle is surrounded by eight tiles of different colours.

A *colouring* of $\mathcal{T}$ is a partition of $\mathcal{T}$ into four classes, namely $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$, and $\mathcal{C}_4$ as follows. For every $i \in \left[\left\lceil \frac{k+\xi-1}{2f_w(t)+1} \right\rceil + 1\right]$ and every $j \in \left[\left\lceil \frac{3k-\xi'-1}{2f_w(t)+1} \right\rceil + 1\right]$ we assign to $\mathsf{T}_{\mathsf{c}(i),\mathsf{r}(j),f_w(t)}$ the colour $(i \bmod 2) + 2\,(j \bmod 2) + 1$. This means that to tiles where $\mathsf{c}(i)$ and $\mathsf{r}(j)$ are even we assign the colour one, to tiles where $\mathsf{r}(j)$ is even but $\mathsf{c}(i)$ is odd we assign the colour three, and so on, see Figure 5.9 for an example. Hence, every column and every row is two-chromatic, and between each pair of tiles from the same colour that share a row or a column, there is a tile of a different colour that separates those tiles in their respective row or column. Additionally, the eight tiles surrounding a tile T are all of a different colour than T itself.

We use tilings in several different ways, and sometimes it is necessary to "zoom out" of our current wall, i.e. to abstract over some of the horizontal paths and vertical cycles in order to obtain a more streamlined version of our wall.

Definition 5.4.5 (Walls from a Tiling)**.** Let $k, d \in \mathbb{N}$ be positive integers, W a cylindrical wall of order $3k$ with the triadic partition $\mathcal{W} = (W, k, W_1, W_2, W_3, W^1, W^2, W^3)$, and $\mathcal{T}$ a tiling of width d that covers W_2. Moreover, let $\widetilde{W}$ be some slice of W_2 and let I_Q be the largest set of integers such that for every $i \in I_Q$ the vertical cycle Q_i

contains vertices of a tile from $\mathcal{T}$ which intersects $\widetilde{W}$. Let $\widetilde{W}(\mathcal{T})$ be the union of the cycles Q_i with $i \in I_Q$, and the paths $P^{\mathsf{i}}_j[\{Q_h \mid h \in I_Q\}]$ and $P^{\mathsf{o}}_j[\{Q_h \mid h \in I_Q\}]$ for every $j \in [3k]$. We call $\widetilde{W}(\mathcal{T})$ the *extension of* $\widetilde{W}$ *that covers* $\mathcal{T}$. $\widetilde{W}(\mathcal{T})$

Now, let $\{\mathcal{C}_1, \dots, \mathcal{C}_4\}$ be a four-colouring of $\mathcal{T}$ and $i \in [4]$ be a fixed colour. Then, let $J_Q \subseteq [3k]$ be the largest set of integers such that for all $j \in J_Q$ the vertical cycle Q_j of $\widetilde{W}(\mathcal{T})$ does not contain a vertex of some tile from $\mathcal{C}_i$. Similarly, let $J_P \subseteq [3k]$ be the largest set of integers such that for every $j \in J_P$, none of the two paths from R_j contains a vertex of a tile from $\mathcal{C}_i$.

By $\widetilde{W}[\mathcal{T}, i]$ we denote the subgraph of W induced by the union of the cycles Q'_i with $i' \in J_Q$, and the paths $P^{\mathsf{i}}_j[\{Q_h \mid h \in J_Q\}]$ and $P^{\mathsf{o}}_j[\{Q_h \mid h \in J_Q\}]$ for every $j \in J_P$. We say that $\widetilde{W}[\mathcal{T}, i]$ is the *i-th $\mathcal{T}$-slice* of $\widetilde{W}$. ⊣ $\widetilde{W}[\mathcal{T}, i]$

Note that the width of $\widetilde{W}(\mathcal{T})$ is at most by $2d$ greater than the width of $\widetilde{W}$. In $\widetilde{W}[\mathcal{T}, i]$, we essentially cut away the tiles of $\mathcal{C}_i$. This operation gives us a slice W' of some cylindrical wall for which the perimeter of every tile in $\mathcal{C}_i$ is the perimeter of some brick. Next, we find a tiling of W' such that every tile of $\mathcal{C}_i$ that belongs to $\widetilde{W}$ is captured by the centre of some tile in the new tiling.

Definition 5.4.6 (Tier II Tiling)**.** Let $t, k, k' \in \mathbb{N}$ be positive integers with $k \geq k'$ and $f\colon \mathbb{N} \to \mathbb{N}$ be some function. Let W be a cylindrical wall of order $3k$ with its triadic partition $\mathcal{W} = \left(W, k, W_1, W_2, W_3, W^1, W^2, W^3\right)$, and $\mathcal{T} = \mathcal{T}_{W,k,f,\xi,\xi'}$ for some $\xi, \xi' \in [f(t)+1]$, as well as $\{\mathcal{C}_1, \dots, \mathcal{C}_4\}$ be a four-colouring of $\mathcal{T}$ and $i \in [4]$ be a fixed colour. Moreover, let $\widetilde{W}$ be a slice of W_2 of width k' such that no tile of $\mathcal{C}_i$ contains a vertex of $\mathsf{per}(W)$ and $\widetilde{\mathcal{T}}$ be the collection of all tiles from $\mathcal{T}$ that contain a vertex of $\widetilde{W}$.

The *tier II tiling* of width $f(t)$ for the slice $\widetilde{W}[\mathcal{T}, i]$, denoted by $(\mathcal{T}, i, f)_{\mathrm{II}}\left[\widetilde{W}\right]$, is defined as the unique tiling of $\widetilde{W}[\mathcal{T}, i]$ such that every $\mathsf{T} \in \mathcal{C}_i \cap \widetilde{\mathcal{T}}$ is in the interior of the centre of some tile of $(\mathcal{T}, i, f(t))_{\mathrm{II}}\left[\widetilde{W}\right]$. ⊣ $(\mathcal{T}, i, f(t))_{\mathrm{II}}$

Since every tile in $\mathcal{T}$ consists of $2f(t)+2$ path pairs, the tiling $(\mathcal{T}, i, f(t))_{\mathrm{II}}\left[\widetilde{W}\right]$ is well defined and does in fact cover all of $\widetilde{W}[\mathcal{T}, i]$.

Now, we fix some terminology for paths in the graph that cause non-planar behaviour with respect to the wall. Let $k, w \in \mathbb{N}$ be positive integers, W be a wall of order k and W' be a slice of W. We call a C_1-C_k- or C_k-C_1-path P a *jump over all of* W if it is internally disjoint from W. A *perimeter jump* over a wall W is a path that starts or ends at a tile of W-distance at least 2 to $\mathsf{per}(W)$ and is otherwise disjoint of W. A $V(W')$-path P is called a *jump over* W' if $E(P) \cap E(W') = \emptyset$. We say that a directed $V(W')$-path P is a *w-long jump over* W' if for all $\xi, \xi' \in [w+1]$ the start-vertex and the end-vertex of P belong to distinct tiles T_s and T_t of the tiling $\mathcal{T}_{W,k,w,\xi,\xi'}$.

5.4.1 Selected results from [GKKK20]

In this subsection we state some results from [GKKK20] which provide the corner stones for our proof of Theorem 5.4.3.

Theorem 5.4.7 (Giannopoulou et al. [GKKK20])**.** Let D be a digraph and $X \subseteq V(D)$. For all positive $k \in \mathbb{N}$, there are k pairwise vertex disjoint X-paths in D, or there exists a set $S \subseteq V(D)$ of size at most $2k$ such that every X-path in D contains a vertex of S.

Furthermore, there is a polynomial time algorithm which, given a digraph D and a set $X \subseteq V(D)$, outputs k pairwise disjoint X-paths, or a set $S \subseteq V(D)$ of size at most $2k$ as above.

Theorem 5.4.8 (Giannopoulou et al. [GKKK20])**.** Let $k \in \mathbb{N}$ be a positive integer, D be a digraph, and $X, Y \subseteq V(D)$. If $\mathcal{P}$ is a half-integral X-Y-linkage of order $2k$ in D, then there exists an X-Y-linkage $\mathcal{J}$ of order k such that $V(\mathcal{J}) \subseteq V(\mathcal{P})$.

The following is a combination of Lemmas 4.3 to 4.8 from [GKKK20] and a proof can be found in the proof of Lemma 4.9 in [GKKK20]. The only difference between Lemma 4.9 from [GKKK20] and the statement below is that we extract the last subcase of Case 1 in its proof as a potential outcome. In fact the statement contains two lemmata that are merely two analogue cases of the same situation. We mark the changes one needs to perform to obtain the second statement in blue and in parentheses.

Lemma 5.4.9 (Giannopoulou et al. [GKKK20])**.** There exist functions $f_w \colon \mathbb{N} \to \mathbb{N}$, $f_P \colon \mathbb{N} \to \mathbb{N}$, and $f_W \colon \mathbb{N} \to \mathbb{N}$ such that for every $t \in \mathbb{N}$ the following holds. Let

- D be a digraph,
- W be a cylindrical wall of order $3k$ with $k \geq f_W(t)$ in D,
- $\mathcal{W} = \big(W, k, W_1, W_2, W_3, W^1, W^2, W^3\big)$ be the triadic partition of W, and
- $\mathcal{T} = \mathcal{T}_{W,k,f_w(t),\xi,\xi'}$ for some $\xi, \xi' \in [f_w(t)+1]$.

If there exists a subfamily $\mathcal{T}'$ of $\mathcal{T}$ and a family $\mathcal{J}$ of pairwise disjoint directed paths in D with the following properties:

1. Every member of $\mathcal{J}$ is a $V(W)$-path,
2. $|\mathcal{T}'| = |\mathcal{J}| = f_P(t)$,
3. for every $\mathsf{T}_{\mathsf{c}(p),\mathsf{r}(q),f_w(t)} \neq \mathsf{T}_{\mathsf{c}(p'),\mathsf{r}(q'),f_w(t)} \in \mathcal{T}'$ we have $\max\{|p - p'|, |q - q'|\} \geq 2$,
4. there exists a bijection $\mathsf{leave}\colon \mathcal{T}' \to \mathcal{J}$ such that the start-vertex of the path $\mathsf{leave}(\mathsf{T})$ belongs to the centre of T for every $\mathsf{T} \in \mathcal{T}'$, (there exists a bijection $\mathsf{arrive}\colon \mathcal{T}' \to \mathcal{J}$ such that the end-vertex of the path $\mathsf{arrive}(\mathsf{T})$ belongs to the centre of T for every $\mathsf{T} \in \mathcal{T}'$,)
5. for all $\mathsf{T} \in \mathcal{T}'$, where $V(\mathcal{T}') = \bigcup_{\mathsf{T}' \in \mathcal{T}'} V(\mathsf{T}')$ the intersection $V(\mathsf{leave}(\mathsf{T})) \cap V(\mathcal{T}')$ contains exactly the end-vertex of $\mathsf{leave}(\mathsf{T})$ ($V(\mathsf{arrive}(\mathsf{T})) \cap V(\mathcal{T}')$ contains exactly the start-vertex of $\mathsf{arrive}(\mathsf{T})$), and finally
6. the end-vertices (start-vertices) of the paths in $\mathcal{J}$ are of mutual W-distance at least four.

Then, at least one of the following is true.

(i) D has a $\overleftrightarrow{K}_t$-butterfly minor grasped by W,

(ii) there exists a family of tiles $\mathcal{T}'' \subseteq \mathcal{T}'$ all contained in a single strip $S \subseteq W$ of height equal to the height of the tiles in $\mathcal{T}$ such that

- we can number $\mathcal{T}'' = \{\mathsf{T}_1, \dots, \mathsf{T}_h\}$ such that $S - \mathsf{T}_i$ has one component containing exactly the tiles $\mathsf{T}_1, \dots, \mathsf{T}_{i-1}$ for each $i \in [h]$,
- $|\mathcal{T}''| \geq f_P(t)^{\frac{1}{4}}$,
- for every $i \in [h-1]$ the tiles T_i and T_{i+1} are separated in W by a slice of width equal to the width of the tiles in $\mathcal{T}$, and
- there is a family $\mathcal{J}' \subseteq \mathcal{J}$ with $|\mathcal{J}'| = |\mathcal{T}''|$ such that for each $\mathsf{T} \in \mathcal{T}''$ we have $\mathsf{leave}(\mathsf{T}) \in \mathcal{J}'$ ($\mathsf{arrive}(\mathsf{T}) \in \mathcal{J}'$), and
- for each $i \in [h]$ the end-vertex of $\mathsf{leave}(\mathsf{T}_i)$ (start-vertex of $\mathsf{arrive}(\mathsf{T}_i)$) lies in the component of $S - \mathsf{T}_i$ that contains no tiles of $\mathcal{T}''$ if $i = 1$, or in the slice of S separating T_{i-1} and T_i otherwise.

(iii) there exists a family of tiles $\mathcal{T}'' \subseteq \mathcal{T}'$ all contained in a single strip $S \subseteq W$ of height equal to the height of the tiles in $\mathcal{T}$ such that $\mathcal{T}''$ and S meet the properties of outcome (ii) after switching the parametrisation of W.

We fix the functions f_w, f_P, and f_W from Lemma 5.4.9 and, by using the bounds provided by the original proofs [GKKK20], we obtain the following rough estimates:

(i) $f_w(t) = 2^9 t^{10}$,

(ii) $f_P(t) = 2^7 t^8$, and

(iii) $f_W(t) = 2^{32+t^{30}}$.

Lemma 5.4.9 is already powerful enough to guarantee a large cross-row grid as a butterfly minor in case we find many long jumps over a sufficiently large wall. This is due to the following proof sketch, which we formalise over the remainder of the section. In case Lemma 5.4.9 yields the existence of a $\overleftrightarrow{K}_{2\cdot t^2}$ butterfly minor, this minor contains a cross-row grid of order t as a butterfly minor, by Lemma 5.2.3. Otherwise, Lemma 5.4.9 produces many pairwise disjoint slices of the wall W, mutually far apart from each other, and each of them having a neighbouring slice together with a fairly long jump into this neighbouring slice. Such a slice and its neighbour together yield a slice of twice the width that contains a cross. By Lemma 5.2.5, this yields a cross-row grid as a butterfly minor.

5.4.2 Removing long jumps

In order to provide some formal basis to the proof sketch from before, we prove two auxiliary results, each providing ways to find cross-row grid minors from long jumps over our wall.

Definition 5.4.10 (Auxiliary Digraph Type I)**.** Let $t, k, k', w \in \mathbb{N}$ be positive integers such that $k \geq k' \geq 2f_W(t) + 4f_P(t) \cdot (2w+1)$, $w \geq 2f_w(t)$, and $\xi, \xi' \in [w+1]$. Let D be a digraph containing a cylindrical wall W of order $3k$ with its triadic partition $\mathcal{W} = \left(W, k, W_1, W_2, W_3, W^1, W^2, W^3\right)$, and a tiling $\mathcal{T} = \mathcal{T}_{W,k,w,\xi,\xi'}$. Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_4\}$ be a four-colouring of $\mathcal{T}$, $i \in [4]$ and $W' \subseteq W$ be a slice of width k' of W_2. At last, let us denote by $\mathcal{T}'$ the family of tiles from $\mathcal{T}$ that share a vertex with W' and let $\mathcal{C}'_i \coloneqq \mathcal{T}' \cap \mathcal{C}_i$. Then, $D^{\mathrm{I}}_i(W')$ is the digraph obtained from D by performing the following construction steps for every $\mathsf{T} \in \mathcal{C}'_i$:

$D^{\mathrm{I}}_i(W')$

1. add new vertices $x^{\mathrm{in}}_{\mathsf{T}}$ and $x^{\mathrm{out}}_{\mathsf{T}}$,
2. for every vertex u in the centre of T introduce the edges $\left(u, x^{\mathrm{in}}_{\mathsf{T}}\right)$ and $\left(x^{\mathrm{out}}_{\mathsf{T}}, u\right)$, and then
3. delete all internal vertices of T.

Additionally, we define the sets

$$X_{\mathrm{I}}^{\mathrm{out}} := \{x_{\mathsf{T}}^{\mathrm{out}} \mid \mathsf{T} \in \mathcal{C}_i'\}, \text{ and}$$
$$X_{\mathrm{I}}^{\mathrm{in}} := \{x_{\mathsf{T}}^{\mathrm{in}} \mid \mathsf{T} \in \mathcal{C}_i'\}. \qquad \dashv$$

We establish that every outcome of Lemma 5.4.9 yields a sufficiently large cross-row grid.

Lemma 5.4.11. Every outcome of Lemma 5.4.9 yields $D^{\boxtimes}_{\sqrt{t/2}}$ as a butterfly minor.

Proof. If the outcome is (i), then, by Lemma 5.2.3, we can find a butterfly model of $D^{\boxtimes}_{\sqrt{t/2}}$ within the $\overleftrightarrow{K}_t$ butterfly minor model.

Hence, we may assume the outcome is (ii) or (iii). Since these two cases are symmetric it suffices to only consider outcome (ii). In that case, we find a family $\mathcal{T}'' \subseteq \mathcal{T}'$ of size $h \geq f_P(t)^{\frac{1}{4}} \geq \frac{1}{3}t$ contained in a single strip $S \subseteq W$ which is of the same height as the tiles in $\mathcal{T}''$. We can number the tiles $\mathcal{T}'' = \{\mathsf{T}_1, \dots, \mathsf{T}_h\}$ such that $S - \mathsf{T}_i$ has one component containing exactly the tiles $\mathsf{T}_1, \dots, \mathsf{T}_{i-1}$ for each $i \in [h]$. For each $i \in [h]$ let $S_i \subseteq W$ be the slice of width equal to the width of the tiles in $\mathcal{T}''$ whose intersection with S is exactly the tile T_i. Additionally, for each $i \in [h]$ there is a slice H_i of W containing S_i and both the start- and the end-vertex of the jump $J_i := \mathsf{leave}(\mathsf{T}_i)$ such that H_i and H_j are disjoint if $i \neq j$. By Lemma 5.4.15, J_i yields a cross slice H_i' in H_i such that the cross lies within the strip S. As there are more than t rows that do not lie in S and we have more than $\sqrt{t/2}$ such slices, by Lemma 5.2.5 we obtain a $D^{\boxtimes}_{\sqrt{t/2}}$ as a butterfly minor. $\square$

Let L and P be directed paths. We say that P is a *long jump of* L if P is a w-long jump over W and $P \subseteq L$. Additionally, we say that P is a *jump of* L if P is a directed W-path.

The following lemma deals with long jumps between one colour class and the remaining three colour classes of a tiling colouring.

Lemma 5.4.12. Let $t, k, k', w \in \mathbb{N}$ be positive integers with $k \geq k' \geq 2f_W(2 \cdot t^2) + 2^{16} f_P(2 \cdot t^2) + 2 + 2w$, $w \geq 2f_w(2 \cdot t^2) + 2^7 f_P(2 \cdot t^2)$, and $\xi, \xi' \in [w+1]$. Let D be a digraph containing a cylindrical wall W of order $3k$ with its triadic partition $\mathcal{W} = (W, k, W_1, W_2, W_3, W^1, W^2, W^3)$, and a tiling $\mathcal{T} = \mathcal{T}_{W,k,w,\xi,\xi'}$. Let $\{\mathcal{C}_1, \dots, \mathcal{C}_4\}$ be a four-colouring of $\mathcal{T}$, $i \in [4]$ and $W' \subseteq W$ be a slice of width k' of W_2.

Now let $\mathcal{T}'$ be the family of all tiles of $\mathcal{T}$ that are completely contained in W' and let $\widetilde{W}$ be the smallest slice of W that contains all tiles from $\mathcal{T}'$.

Consider the auxiliary digraph $D_i^{\mathrm{I}}(\widetilde{W})$ with $X_{\mathrm{I}}^{\mathrm{out}}$ and $X_{\mathrm{I}}^{\mathrm{in}}$. Additionally, we construct the set Y_{I} as follows. Let Q and Q' be the two cycles of $\mathsf{per}(W)$. For every $j \in \left[\frac{3k}{4}\right]$, Y_{I} contains exactly one arbitrarily chosen vertex of $Q \cap P_{4j}^{\mathsf{o}}$, $Q \cap P_{4j+2}^{\mathsf{i}}$, $Q' \cap P_{4j}^{\mathsf{o}}$, and $Q' \cap P_{4j+2}^{\mathsf{i}}$ each.

If there exists a family $\mathcal{L}$ of pairwise disjoint directed paths with $|\mathcal{L}| = 2^7 f_P(2 \cdot t^2)$ such that either

- $\mathcal{L}$ is a family of directed $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-paths, or
- $\mathcal{L}$ is a family of directed Y_{I}-$X_{\mathrm{I}}^{\mathrm{in}}$-paths,

then D contains $D_t^{\boxtimes}$ as a butterfly minor.

Proof. In the statement we insert $2 \cdot t^2$ as the arguments for all functions. This means, due to Lemma 5.4.11, any application of Lemma 5.4.9 yields a $D_t^{\boxtimes}$ as a butterfly minor. Thus, we construct a cylindrical wall $W''' \subseteq W$ of sufficient size, together with a family of $f_P(2 \cdot t^2)$ directed W'''-paths that meet the requirements of Lemma 5.4.9.

Without loss of generality, let us assume $\mathcal{L}$ is a family of directed $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-paths. The other case, Y_{I}-$X_{\mathrm{I}}^{\mathrm{out}}$-paths, follows by similar arguments.

Short overview of the proof: Towards our goal, we first show that we can use $\mathcal{L}$ to construct a half-integral $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-linkage $\mathcal{L}_1$ such that

1. $|\mathcal{L}_1| = 2^7 f_P(2 \cdot t^2)$,
2. there exists a family $\mathcal{F} \subseteq \mathcal{T}'$ with $|\mathcal{F}| \leq 2^7 f_P(2 \cdot t^2)$, and
3. for every $L \in \mathcal{L}_1$, every end-vertex u of a jump of L with $u \in V(\widetilde{W})$ belongs to a tile from $\mathcal{C}_i' \cup \mathcal{F}$.

Second, we use Theorem 5.4.8 to obtain a family $\mathcal{L}_2$ of pairwise disjoint directed $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-paths of size $2^6 f_P(2 \cdot t^2)$ from $\mathcal{L}_1$. Third, we remove the cycles and paths of $\widetilde{W}$ that meet tiles from $\mathcal{F}$ and obtain a new slice W''' of some cylindrical wall. For this slice, we construct a tiling and a tier II tiling as well as a half-integral linkage $\mathcal{L}_3$ of size $2^4 f_P(2 \cdot t^2)$ from $\mathcal{L}_2$. The linkage $\mathcal{L}_3$ connects the centres of some tiles in the tier II tiling to vertices of $\widetilde{W}'$ such that their end-vertices and start-vertices are mutually far enough apart and every path in $\mathcal{L}_3$ is internally disjoint from the new wall W'''. Another application of Theorem 5.4.8 then yields the family $\mathcal{L}_4$ of long jumps necessary for an application of Lemma 5.4.9.

The construction of $\mathcal{L}_1$ and $\mathcal{F}$. We do construct $\mathcal{L}_1$ and $\mathcal{F}$ iteratively starting with $\mathcal{L}' := \mathcal{L}$, $\mathcal{L}_1 := \emptyset$, and $\mathcal{F} := \emptyset$. As long as $\mathcal{L}'$ is non-empty, perform the following actions.

Select some path $L \in \mathcal{L}'$. In case L is internally disjoint from $\widetilde{W}$, add L to $\mathcal{L}_1$ and remove it from $\mathcal{L}'$. Otherwise, let v_L be the first vertex of L that belongs to $\widetilde{W}$, but not to a tile from $\mathcal{C}'_i$.

(i) If v_L does not belong to any tile from $\mathcal{F}$, let $\mathsf{T} \in \mathcal{T} \setminus \mathcal{C}_i$ be the tile that contains v_L and add T to $\mathcal{F}$. Let R be a shortest directed path from v_L to Y_{I} in W such that R avoids all vertices of W that are contained in two different paths of $\mathcal{L}_1$ and that is internally disjoint from Lv_L. Now add Lv_LR to $\mathcal{L}_1$ and remove L from $\mathcal{L}'$. Note that such a path R must exist because the paths in $\mathcal{L}$ are pairwise disjoint, T was not used for such a re-routing before, and w and k' are chosen sufficiently large in proportion to $2^7 f_P(2 \cdot t^2)$. Also note that the path R is exactly the subpath that might cause $\mathcal{L}_1$ to be half-integral. However, because the paths in $\mathcal{L}'$ are pairwise disjoint, we can be sure that R never meets a vertex contained in two distinct paths of $\mathcal{L}_1$.

(ii) If v_L belongs to a tile from $\mathcal{F}$, follow along v_LL until we encounter a vertex u_L for which one of the following is true:

a) u_L belongs to a tile T from $\mathcal{T} \setminus (\mathcal{C}_i \cup \mathcal{F})$, or

b) every internal vertex of u_LL belongs to $W - \widetilde{W}$ or to some tile from $\mathcal{C}_i \cup \mathcal{F}$.

If (ii)a, add T to $\mathcal{F}$ and then repeat the instructions from (i) but replace v_L with u_L. Otherwise, (ii)b holds and we remove L from $\mathcal{L}'$ and add it to $\mathcal{L}_1$.

During this construction, for every $L \in \mathcal{L}$, we added at most one tile to $\mathcal{F}$ and thus $|\mathcal{F}| \leq |\mathcal{L}|$. Note that, by construction, $\mathcal{L}_1$ is indeed a half-integral linkage from $X_{\mathrm{I}}^{\mathrm{out}}$ to Y_{I}. Moreover, we may assume that every L meets each tile in $\mathcal{F}$ in at most $2^7 f_P(2 \cdot t^2) + 1$ horizontal path pairs and vertical cycles, because otherwise we can replace it by a shorter path through W itself.

Obtaining $\mathcal{L}_2$. We apply Theorem 5.4.8 to obtain a family $\mathcal{L}_2$ of pairwise disjoint directed $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-paths with $V(\mathcal{L}_2) \subseteq V(\mathcal{L}_1)$ and $|\mathcal{L}_2| = 2^6 f_P(2 \cdot t^2)$.

Constructing $\mathcal{L}_3$. Let us consider $W'' := \widetilde{W}[\mathcal{T}, i]$ together with the tiling $\mathcal{T}'' := (\mathcal{T}, i, w)_{\mathrm{II}}\left[\widetilde{W}\right]$ and a four-colouring $\{\widetilde{\mathcal{C}}_1, \ldots, \widetilde{\mathcal{C}}_4\}$. Note that the width of $\widetilde{W}$ is at most $2w$ smaller than the width of W'. Thus, by choice of k', we obtain that W'' is a slice of width $k'' \geq f_W(2 \cdot t^2) + 2^7 f_P(2 \cdot t^2)\,(2w + 1) + 1$ of some cylindrical wall of order $3k''$ that is completely contained in W. For each $L \in \mathcal{L}_2$ let $\mathsf{T}_L^1 \in \mathcal{C}_i$ such that

$\mathsf{start}(L)$ belongs to T_L^1. Let $K_L^1 \in \mathcal{T}''$ be the tile whose centre is the perimeter of T_L^1. Choose any vertex $\mathsf{start}(L)'$ of degree three in W'' that is not contained in any path of $\mathcal{L}_2$, and let R_L be a directed path from $\mathsf{start}(L)'$ to $\mathsf{start}(L)$ within T_L^1. Let $\mathcal{L}_3$ be the, possibly again half-integral, family of directed paths resulting from concatenating the new paths and the corresponding original path from $\mathcal{L}_2$.

Finding $\mathcal{L}_4$ and W'''. There exists $j \in [4]$ such that at least $2^4 f_P(2 \cdot t^2)$ of the paths from $\mathcal{L}_3$ start at the centre of a tile from $\widetilde{\mathcal{C}}_j$. Let $\mathcal{L}_3' \subseteq \mathcal{L}_3$ be a family of exactly $2^4 f_P(2 \cdot t^2)$ such paths. Next, let us consider the family $\mathcal{F}$. Let W''' be the subgraph of W'' induced by all vertical cycles and horizontal path pairs in W'' that do not contain a vertex of some tile in $\mathcal{F}$ that belongs to a path in $\mathcal{L}_3'$. Since $|\mathcal{F}| \leq 2^7 f_P(2 \cdot t^2)$ and each tile in $\mathcal{F}$ meets a path in $\mathcal{L}_3'$ in at most $2^7 f_P(2 \cdot t^2) + 1$ such cycles and pairs of horizontal paths, it follows that W''' is a slice of width $k''' \geq f_W(2 \cdot t^2) + 2$ of some cylindrical wall $W^* \subseteq W$ of order $3k'''$. Moreover, W^* can be partitioned into three slices of width k''' as in its triadic partition, such that W''' is the slice in the middle. Let us rename the paths and cycles of W^* such that $Q_1^*, \ldots, Q_{3k'''}^*$ are the vertical paths of W^*, $P_1^{*\mathsf{i}}, \ldots, P_{3k'''}^{*\mathsf{i}}$ its in-paths and $P_1^{*\mathsf{o}}, \ldots, P_{3k'''}^{*\mathsf{o}}$ its out-paths. We construct the set Y^* as follows: For every $j \in \left[\frac{3k'''}{4}\right]$, Y^* contains exactly one vertex of $Q_1^* \cap P_{4j}^{*1}$, $Q_1^* \cap P_{4j+2}^{*2}$, $Q_{3k'''}^* \cap P_{4j}^{*1}$, and $Q_{3k'''}^* \cap P_{4j+2}^{*2}$ each.

Similarly to $\mathcal{L}_1$, we now construct $\mathcal{L}_3''$ iteratively from $\mathcal{L}_3'$. We start with $\mathcal{L}_3''$ being empty. Let $L \in \mathcal{L}_3'$ be any path and t_L be the first vertex after $\mathsf{start}(L)$ that L shares with either W''' or $W^* - W'''$. If $t_L \in V(W''')$, add Lt_L to $\mathcal{L}_3''$. Otherwise, let b_L be the last vertex of L in $W^* - W'''$. Then, we can find a path R_L in W from b_L to a vertex t_L^* of Y^* such that t_L^* is of W^*-distance at least four to every endpoint of every path already in $\mathcal{L}_3''$, R_L is internally disjoint from L, and R_L does not contain a vertex that is contained in two distinct paths from $\mathcal{L}_3''$. Add LR_L to $\mathcal{L}_3''$. Finally, $\mathcal{L}_3''$ is a half-integral linkage from the set $S^* := \mathsf{start}(\mathcal{L}_3')$ to Y^* of size $2^4 f_P(2 \cdot t^2)$, and thus by Theorem 5.4.8 we can find a family $\mathcal{L}_4$ of pairwise disjoint directed paths from S^* to Y^* with $V(\mathcal{L}_4) \subseteq V(\mathcal{L}_3'')$ that is of size $2^3 f_P(2 \cdot t^2)$. It follows that all paths in $\mathcal{L}_4$ are internally disjoint from W'''.

Let us consider the tiles of $\widetilde{\mathcal{C}}_i$ whose centres contain a vertex of S^*. Since W''' might be a proper subgraph of W'', $\mathcal{T}''$ is not necessarily a tiling of W'''. Each tile $\mathsf{T} \in \mathcal{T}''$, however, contains a tile T' of width $f_w(2 \cdot t^2)$ with the same centre. Since T is surrounded by at most 8 tiles from $\mathcal{F}$ in W', we may find, among the $2^3 f_P(2 \cdot t^2)$ many such tiles, a family $\mathcal{J}$ of $f_P(2 \cdot t^2)$ tiles that are pairwise disjoint. Thus, because they all are constructed from the family $\widetilde{\mathcal{C}}_i$, they meet the distance requirements of the tiles in Lemma 5.4.9. Hence, we can apply Lemma 5.4.9 and by Lemma 5.4.11 we obtain the desired butterfly minor $D_t^{\boxtimes}$. $\square$

The above lemma allows us to argue that enough long jumps that all start, respectively end, in the centre of tiles from a single colour class but end, respectively start, in tiles from the remaining three classes yield a cross-row grid as a butterfly minor. By utilising a second auxiliary digraph we can make use of Theorem 5.4.7 to also prove this for long jumps between tiles of the same colour.

Definition 5.4.13 (Auxiliary Digraph Type II)**.** Let $t, k, k', w \in \mathbb{N}$ be positive integers such that $k \geq k' \geq 2f_W(t)$, $w \geq 2f_w(t)$, and $\xi, \xi' \in [w+1]$. Let D be a digraph containing a cylindrical wall W of order $3k$ with its triadic partition $\mathcal{W} = \big(W, k, W_1, W_2, W_3, W^1, W^2, W^3\big)$, and a tiling $\mathcal{T} = \mathcal{T}_{W,k,w,\xi,\xi'}$. Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_4\}$ be a four-colouring of $\mathcal{T}$, $i \in [4]$ and $W' \subseteq W$ be a slice of width k' of W_2 such that no tile of $\mathcal{C}_i$ contains a vertex of the perimeter of W'. Then, $D_i^{\mathrm{II}}(W')$ is the digraph obtained from D by performing the following construction steps. $D_i^{\mathrm{II}}(W')$

For every $\mathsf{T} \in \mathcal{C}_i$, such that T contains a vertex of W', we do the following:

1. add a new vertex x_{T}, and
2. for every vertex v that belongs to the interior or the centre of T, introduce the edges (x_{T}, v) and (v, x_{T}).

Once this is done, delete all vertices of W' that do not belong to tiles of $\mathcal{C}_i$ that are contained in W'. Let X_{II}^i be the collection of all newly introduced vertices x_{T}. ⊣

Lemma 5.4.14. Let $t, k, k', w \in \mathbb{N}$ be positive integers, and $\xi, \xi' \in [w+1]$ where $w \geq 2f_w(2 \cdot t^2)$. Let D be a digraph containing a cylindrical wall W of order $3k$ with vertical paths $Q_1, \ldots, Q_{3k}$, in-paths $P_1^{\mathsf{i}}, \ldots, P_{3k}^{\mathsf{i}}$ and out-paths $P_1^{\mathsf{o}}, \ldots, P_{3k}^{\mathsf{o}}$, where $k \geq k' \geq 4f_W(2 \cdot t^2)^2$. Also, let $\mathcal{W} = (W_0, k, W_1, W_2, W_3, W^1, W^2, W^3)$ be its triadic partition, $W^* \subseteq W_2$ be a slice of width k', $\mathcal{T} = \mathcal{T}_{W,k,w,\xi,\xi'}$ be a tiling, $\{\mathcal{C}_1, \ldots, \mathcal{C}_4\}$ a four-colouring and $i \in [4]$ a fixed colour.

Then, D contains $D_t^{\boxtimes}$ as a butterfly minor, or there exists a set $\mathcal{Z}_{i,\xi,\xi'}^2 \subseteq \mathcal{T}$ with $|\mathcal{Z}_{i,\xi,\xi'}^2| \leq 8f_P(2 \cdot t^2)$ and a set $Z_{i,\xi,\xi'}^2 \subseteq V(D - W)$ with $|Z_{i,\xi,\xi'}^2| \leq 8f_P(2 \cdot t^2)$ such that every directed $V(W_0)$-path in $D - Z_{i,\xi,\xi'}^2$ whose start- and end-vertex belong to different tiles of $\mathcal{C}_i$ contains a vertex of some tile in $\mathcal{Z}_{i,\xi,\xi'}^2$.

Proof. Let W' be the largest slice of W^* such that no tile of $\mathcal{C}_i$ contains a vertex of $\mathrm{per}(W)$. Let us consider the auxiliary digraph $D_i^{\mathrm{II}}(W')$ with the set X_{II}^i of newly added vertices. By applying Theorem 5.4.7 to the set X_{II}^i in $D_i^{\mathrm{II}}(W')$, we either find a set Z of size at most $8f_P(2 \cdot t^2)$ that hits all directed X_{II}^i-paths, or there exists a family $\mathcal{J}'$ of $4f_P(2 \cdot t^2)$ pairwise disjoint directed X_{II}^i-paths in $D_i^{\mathrm{II}}(W')$.

In case we obtain the latter, by construction of $D_i^{\mathrm{II}}(W')$, no path in $\mathcal{J}'$ contains a vertex of W_2. Back in the digraph D, let us consider the tier II tiling $\mathcal{T}'' := (\mathcal{T}, i, w)_{\mathrm{II}}[W']$ of width w of $W'' := W'[\mathcal{T}, i]$. Note that the width of W' is at most $2w$ smaller than the width of W^*. Thus, by choice of k, W'' contains a cylindrical wall W''' of order $f_W(2 \cdot t^2)$ such that the perimeter of every tile $\mathsf{T} \in \mathcal{C}_i$ for which x_T is an endpoint of some path in $\mathcal{J}'$ bounds a cell of W'''. Let $\mathcal{T}'''$ be a tiling of W''' such that the perimeter of every $\mathsf{T} \in \mathcal{C}_i$ for which x_T is a start- or end-vertex of a path in $\mathcal{J}'$ is the centre of some tile in $\mathcal{T}''$. We now consider a four-colouring $\{\mathcal{C}'_1, \dots, \mathcal{C}'_4\}$ of $\mathcal{T}'''$. Then, there exists $j \in [4]$ and a family $\mathcal{J}''$ of size $f_P(2 \cdot t^2)$ such that the start-vertex of every path in $\mathcal{J}''$ belongs to a tile of $\mathcal{C}_i$ whose perimeter is the centre of a tile in $\mathcal{C}'_j$. For every $J'' \in \mathcal{J}''$ do the following: Let $\mathsf{T}_1, \mathsf{T}_2 \in \mathcal{C}_i$ be the two tiles such that J'' is a directed x_{T_1}-x_{T_2}-path. Then, find a directed W'''-path J that starts on the perimeter of T_1 and ends on the perimeter of T_2. Add the path J to a family $\mathcal{J}$. Hence, $\mathcal{J}$ is a family of pairwise disjoint directed W'''-paths whose start- and end-vertices all lie on the centres of distinct tiles of $\mathcal{T}''$ and that all start at the centres of tiles from $\mathcal{C}'_j$. Thus, we can apply Lemma 5.4.9 and by Lemma 5.4.11 we find $D_t^{\bowtie}$ as a butterfly minor.

Therefore, we may assume that we find a set Z of size at most $8 f_P(2 \cdot t^2)$ that hits all directed X_{II}^i-paths. Let $Z^2_{i,\xi,\xi'} := Z \cap V(D)$, and $\mathcal{Z}^2_{i,\xi,\xi'} := \{\mathsf{T} \in \mathcal{T} \mid x_\mathsf{T} \in Z\}$. Since $|Z| \leq 2 f_P(2 \cdot t^2)$, the demanded bounds on the sizes of the two sets are met. Moreover, because Z meets every directed X_{II}^i-path in $D_i^{\mathrm{II}}(W')$, every directed path with start- and end-vertex in distinct tiles of $\mathcal{C}_i$ which is otherwise disjoint from W contains a vertex from $Z^2_{i,\xi,\xi'}$ or meets a tile from $\mathcal{Z}^2_{i,\xi,\xi'}$. □

5.4.3 Proof of Theorem 5.4.3

As we have seen in Lemma 5.2.5, finding sufficiently many cross slices in a wall yields a cross-row grid as a minor. Thus, we introduce the following lemmata, which describe ways to find cross slices in a wall.

Lemma 5.4.15. Let W be a wall and $\mathcal{T}$ a tiling of W and $d \in \mathbb{N}^+$. Let P be a path with start-vertex in the centre of a tile $\mathsf{T}_{i,j,d} \in \mathcal{T}$ and end-vertex in the centre of a tile $\mathsf{T}_{i',j',d} \in \mathcal{T}$ with $|i - i'| \geq d + 1$, then the slice between $Q_{\min\{i,i'\}}$ and $Q_{\max\{i,i'\}+d}$ contains a cross which lies in the strip between $R_{\min\{j,j'\}-1}$ and $R_{\max\{j,j'\}+d+1}$.

Proof. Without loss of generality we fix the parametrisation of W such that the centres of the tiles have an out-path as upper and an in-path as lower perimeter. The proof for the other case can be obtained by swapping in- and out-paths in the proof. Furthermore, we assume without loss of generality that $i > i'$, the other case work analogously.

Let $\ell := \max\{j, j'\} + d$ and $u := \min\{j, j'\}$, that is, the strip between row R_ℓ and R_u contains all vertices of $\mathsf{T}_{i,j,d}$ and $\mathsf{T}_{i',j',d}$. Now we build a path P' extending P by starting in the upper right corner of $\mathsf{T}_{i,j,d}$, then within $\mathsf{T}_{i,j,d}$ reach $\mathsf{start}(P)$, follow P to $\mathsf{T}_{i',j',d}$ and then, within $\mathsf{T}_{i',j',d}$, reach the lower left corner of $\mathsf{T}_{i',j',d}$.

As $|i - i'| \geq d + 1$, there is a vertical cycle Q separating the two tiles.

Next, we build a path P'' by starting in the intersection of $Q_{i'}$ and $P^{\circ}_{\ell-1}$, following along $P^{\circ}_{\ell-1}$ until meeting Q, then following along Q until meeting P°_{u+1} and finally following P°_{u+1} until meeting Q_{i+d}.

The two paths P' and P'' yield the desired cross. $\square$

Lemma 5.4.16. Let W be a wall and $\mathcal{T}$ a tiling of W. If there is a perimeter jump P over W, then W contains a cross slice.

Proof. We consider the case where P ends on the perimeter of W (the case where P starts on the perimeter works similarly). Let $\mathsf{T} \in \mathcal{T}$ be the tile in which P starts. There are two vertical paths, Q_i and Q_{i+1}, separating T from $\mathsf{per}(W)$. Without loss of generality, assume that the column index of T is larger than $i + 1$ and $\mathsf{end}(P)$ lies on Q_1. We consider the slice between Q_{i+1} and the left path Q_ℓ of the perimeter of T. Let P^{i}_j be the lower path of the perimeter of T. Then, P°_{j+1} contains a path P_1 from Q_{i+1} to Q_ℓ that is disjoint from P. Next, we extend P to a path P_2 by starting with the shortest path within T from Q_ℓ to $\mathsf{start}(P)$, then we add P, from $\mathsf{end}(P)$ follow along Q_1 until meeting P°_{j+2}, and then follow P°_{j+2} until reaching Q_{i+1}. Together, P_1 and P_2, yield a cross over the slice between Q_{i+1} and Q_ℓ. $\square$

Cross slices are the obstruction to flatness that we use in our proofs in order to construct the cross-row grid. Thus, we prove next that walls without cross slices do not allow for certain non-planar behaviour.

Lemma 5.4.17. Let W be a wall of order $k \geq 5$ in a digraph D that is free of cross slices. Then, there is no directed path that is internally disjoint from W whose start- and end-vertex are separated by two vertical paths of W.

Proof. Suppose such a path P exists and let Q_i and Q_{i+1} be the vertical paths of the wall separating the start- and the end-vertex of P. By possibly extending P we can assume that P starts and ends in branch vertices of W. Additionally, assume without loss of generality that P starts on a vertical path Q_r with $r \geq i + 1$.

Let Q_ℓ be the vertical path containing $\mathsf{end}(Q)$ and R_u be the row containing $\mathsf{end}(Q)$. Moreover, let R_b be the row containing $\mathsf{start}(Q)$. We show that the slice between Q_ℓ

and Q_r contains a cross. To this end we construct a path P' starting in the intersection of P°_{u-1}, from there following P°_{u-1} until reaching Q_i (this subpath might be empty as $i = \ell$ is possible), then we follow Q_i until reaching P°_{b+1}, then we follow along P°_{b+1} until reaching Q_r (again this part might be empty as possibly $r = i + 1$). Together, P and P' yield a cross in the slice between Q_ℓ and Q_r, a contradiction. □

Lemma 5.4.18. Let W be a wall of order $k \geq 5$ in a digraph D that is free of cross slices. Then, there is no directed path that is internally disjoint from W whose start- and end-vertex are separated by a row of W and at least one of the endpoints does not lie on $\mathsf{per}(W)$.

Proof. Suppose there is such a path P and let R_j be the row separating its start- and end-vertex. Let Q_i be the vertical path of W containing $\mathsf{start}(P)$. By Lemma 5.4.17, we obtain that $\mathsf{end}(P)$ lies on C_{i+1}, C_i or C_{i+1}.

Consider $\mathsf{end}(P)$ lying on C_{i-1}, then P together with the subpath of P°_j yield a cross in the slice between C_{i-1} and C_i, a contradiction.

Similarly, in case that $\mathsf{end}(P)$ lies on C_{i+1} we obtain a cross in the slice between C_i and C_{i+1} built by P and the subpath of P^{i}_j starting on C_{i+1} and ending on C_i, which again yields a contradiction.

Next, consider the case that $\mathsf{end}(P)$ lies on C_i as well. By assumption, we have $1 < i < k$. Let R_m be the row containing $\mathsf{start}(P)$ and $R_{m'}$ be the row containing $\mathsf{end}(P)$, so we have $|m - m'| \geq 2$. We show that there is a cross in the slice between C_{i-1} and C_{i+1}. First, we construct a path P_1 from C_{i-1} to C_{i+1}. We start in the intersection of C_{i-1} and P°_m and follow P°_m until reaching C_i, we follow C_i until reaching $\mathsf{start}(P)$ (this subpath is empty if $\mathsf{start}(P)$ lies on P°_m), then we append P, next, from $\mathsf{end}(P)$ we follow C_i again until reaching $P^{\circ}_{m'+1}$, we then follow $P^{\circ}_{m'+1}$ until reaching C_{i+1}. Second, we choose P_2 to be the subpath of P^{i}_{m+1} that starts in C_{i+1} and ends in C_{i-1}. Together P_1 and P_2 yield a cross over the slice between C_{i-1} and C_{i+1}, a contradiction. □

Lemma 5.4.19. Let D be a digraph and W a wall in D. If W does not contain any cross slices, then W is flat.

Proof. In favour of readability we use D to refer to the strongly connected component of the digraph D containing W. Let $k + 4$ be the order of W and let $W^- \subseteq W$ be the subwall of order k contained in $W - \mathsf{border}(W)$.

Let X_1' be the minimal set of vertices containing every vertex that is reachable from W^- in $D - \mathsf{border}(W)$ and let $Y_1 := V(D) \setminus X_1'$. Then, (Y_1, X_1) where $X_1 :=$

$X_1' \cup V(\mathsf{border}(W))$ is one of the required separations for (F2). Let X_2' be the minimal set of vertices containing every vertex that reaches W^- in $D - \mathsf{border}(W)$ and let $Y_2 := V(D) \setminus X_2'$ Then, $\overrightarrow{(X_2, Y_2)}$ with $X_2 := X_2' \cup V(\mathsf{border}(W))$ is the other separation required for (F2).

Consider the cylindrical society $(\mathsf{compass}(W), \Omega_1, \Omega_2)$ with sets $V(\Omega_1) = V(Q_1)$ and $V(\Omega_2) = V(Q_{k+4})$. Let Δ' be obtained from a closed disk Δ'' by removing an open disk disjoint from the boundary of Δ''. Fix an embedding Γ of W with $V(\Omega_1)$ being exactly the vertices in $\mathsf{bd}(\Delta'')$ and $V(\Omega_2)$ being exactly the vertices in $\mathsf{bd}(\Delta')$, both consistent with the appearance of the vertices along the vertical cycles in W. Now, $\mathcal{S} := (\Gamma, W, \emptyset)$ is a skeleton. We obtain Γ^+ from Γ by embedding a bridge of a brick with in- and out-attachments in more than one perimeter into such a brick. Otherwise, embed bridges into any brick they have an attachment in.

Claim 1. The tuple $\rho := (\Gamma^+, \mathcal{S})$ is a $\mathbb{S}^2$-decomposition.

Proof. Suppose there is an undirected cycle C in $D_{\mathcal{S}}$ such that there is a directed path P between H_1^+ and H_2^+. By definition of H_1^+ and H_2^+, $\mathsf{start}(P)$ can be reached from a vertex of H_1^+ or H_2^+ and $\mathsf{end}(P)$ can reach a vertex of H_1^+ or H_2^+. Thus, we can assume that $\mathsf{start}(P)$ and $\mathsf{end}(P)$ are vertices of W. If there are at least two vertical paths separating $\mathsf{start}(P)$ and $\mathsf{end}(P)$, then by Lemma 5.4.17, we obtain a cross slice, a contradiction. So both $\mathsf{start}(P)$ and $\mathsf{end}(P)$ lie in the slice between Q_i and Q_{i+2} for some i. By possibly prolonging P at both ends, we can also assume that both $\mathsf{start}(P)$ and $\mathsf{end}(P)$ lie in a row of W and still in the slice between Q_i and Q_{i+2}. By definition of H_1^+ and H_2^+ there is a whole row R_j separating $\mathsf{start}(P)$ and $\mathsf{end}(P)$. By Lemma 5.4.18, this implies that there is a cross slice in W, a contradiction. ■

Claim 1 ensures (F3).

Claim 2. The strongly connected component D' has a torus decomposition centred at ρ.

Proof. Suppose there is an undirected cycle C in $D_{\mathcal{S}}$ such that there is a directed path P between H_1^+ and H_2^+. Assume without loss of generality that H_2^+ contains $D - \mathsf{compass}(W)$. By, Claim 1, there are no jumps between H_1^+ and $H_2^+ \cap \mathsf{compass}(W)$, so P is a path between H_1^+ and $D - \mathsf{compass}(W)$. As D is strongly connected it can be extended into a perimeter jump over W and thus, by Lemma 5.4.16, W contains a cross slice, a contradiction. ■

Claim 2 ensures (F4), so in the end we obtain that W is flat. □

We need the following local version of Menger's Theorem in order to complete the proof of Theorem 5.4.3.

Theorem 5.4.20 (Menger's Theorem [Men27])**.** Let D be a digraph and $X, Y \subseteq V(D)$ be two sets of vertices, then the maximum number of pairwise disjoint directed X-Y-paths in D equals the minimum size of a set $S \subseteq V(G)$ such that every directed X-Y-path in D contains a vertex of S.

Now we want state the proof of our main theorem. Simply put, we consider two cases: either there are enough long jumps to build a cross-row grid or we can find a large slice in the wall that does not contain any long jumps. Then, we divide this slice into smaller parts. Either we find short jumps in enough such parts to build a cross-row grid again, or we find a part free of any jumps which then yields the flat wall.

Based on this idea the proof is split into two phases. A vertex of a wall W is said to be *marked* if it belongs to a separator obtained from Theorem 5.4.20 and Lemma 5.4.12, or through Lemma 5.4.14. A tile is *marked* if it contains a marked vertex or is replaced by a marked vertex in the construction of an auxiliary graph (type I or II). That is, we "mark" every vertex or tile involved in a long jump. A slice of W is said to be *clear* if it does not contain vertices that are marked. This happens in two steps, one for jumps between tiles of different colour and one for jumps between tiles of the same colour. In each of the two steps we introduce families of marked tiles and vertices. In the end we can bound the number of columns containing marked tiles. This gives us a clear slice of W, which we then use in phase two of the proof. In this second phase we split the clear slice into smaller parts, identify among those a part which does not contain any jumps or crossings, and prove that this part is flat.

Theorem 5.4.3 (Directed flat wall theorem)**.** There exist functions $\rho_W\colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ and $\alpha_W\colon \mathbb{N} \to \mathbb{N}$ such that for all integers $r, t \geq 1$ and all digraphs D that do not contain $D_t^{\boxtimes}$ for every $\rho_W(r,t)$-wall W in D there exist a set $A \subseteq V(D)$ with $|A| \leq \alpha_W(t)$ and an r-wall $W' \subseteq W - A$ which is flat under A.

Proof. Let $r, t \in \mathbb{N}$ be positive integers, D be a digraph and W be a cylindrical wall of order $\rho_W(t,r)$, where $\rho_W(t,r)$ will be determined throughout the proof. So, we introduce constants d_1 and d_2, for which we make more and more assumptions in the form of lower bounds. Let us assume

$$\rho_W(t,r) \geq 3d_1.$$

Then, W is a cylindrical wall of order $3d_1$ with vertical paths $Q_1, \ldots, Q_{3k}$, in-paths $P_1^{\mathsf{i}}, \ldots, P_{3k}^{\mathsf{i}}$ and out-paths $P_1^{\mathsf{o}}, \ldots, P_{3k}^{\mathsf{o}}$ and the triadic partition $\mathcal{W} = (W, d_1, \widetilde{W}_1,$

$\widetilde{W}_2, \widetilde{W}_3, \widetilde{W}^1, \widetilde{W}^2, \widetilde{W}^3)$. Throughout the proof let us fix $\widetilde{W}_2$

$$w \coloneqq 2f_w\big(2 \cdot t^2\big) + 2^7 f_P\big(2 \cdot t^2\big)\,.$$

Phase I Phase I is divided into $2^{11+8f_P\left(2\cdot t^2\right)} f_P(2 \cdot t^2)$ *rounds*, each of which is divided into two steps, *Step I* and *Step II*. Let $i \in [2^{11+8f_P\left(2\cdot t^2\right)} f_P(2 \cdot t^2)]$. After round i is complete we require the following sets and graphs as its *output* which then can be used in round $i+1$.

- $F_{\mathrm{I},i} \subseteq V(D)$ such that

$$|F_{\mathrm{I},i}| \leq i\left(2^{11}\,(w+1)^2\, f_P\big(2\cdot t^2\big) + 2^5\,(w+1)^2\, f_P\big(2\cdot t^2\big)\right),$$

  which are the vertices marked in round i

- $D_i \coloneqq D - F_{\mathrm{I},i}$,

- $\mathcal{F}_{\mathrm{I},i} \subseteq \bigcup_{\xi,\xi'\in[w+1]} \mathcal{T}_{\widetilde{W}_2,d_1,w,\xi,\xi'}$ of size at most

$$i\left(2^{11}\,(w+1)^2\, f_P\big(2\cdot t^2\big) + 2^5\,(w+1)^2\, f_P\big(2\cdot t^2\big)\right),$$

  which are the tiles marked in round i, and

- a slice W_i of W_{i-1} of width

$$\left(\left(2^{12}\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\right)\cdot \left(2^{6}\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\right)\right)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)-i} \cdot d_2$$

  that is clear with respect to $F_{\mathrm{I},i}$ and $\mathcal{F}_{\mathrm{I},i}$, such that every long jump over W_i in D_i contains a vertex of some tile in $\mathcal{F}_{\mathrm{I},j} \setminus \mathcal{F}_{\mathrm{I},j-1}$ for every $j \in [i]$.

For $i = 0$ we define $D_0 \coloneqq D$, $W_0 \coloneqq \widetilde{W}_2$ as a slice of itself, and $F_{\mathrm{I},i} \coloneqq \emptyset$ as well as $\mathcal{F}_{\mathrm{I},i} \coloneqq \emptyset$. In this context, whenever we ask for a clear slice of the current slice W_i or W_i' we ask for a slice W' such that there do not exist $\xi, \xi' \in [w+1]$ whose corresponding tiling of $\widetilde{W}_2$ has a tile T that is marked or contains a vertex of any separator set found so far, which satisfies $V(\mathsf{T}) \cap V(W') \neq \emptyset$.

To be able to find a slice of width d_2 after the last round we therefore must fix

$$d_1 \geq \Big(\Big(2^{12}\,(w+1)^3\, f_P\big(2 \cdot t^2\big)+1\Big) \cdot \\ \Big(2^6\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\Big)\Big)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)} \cdot d_2,$$

and we further assume $d_2 \geq 2^{16} f_W(2 \cdot t^2)^2$ to make sure we can apply Lemma 5.4.12 and Lemma 5.4.14 in every round. Note that this is not yet the final lower bound on d_2, just an intermediate assumption.

Next we describe the steps we perform in every round. Let $i \in [2^{11+8f_P\left(2\cdot t^2\right)} f_P(2\cdot t^2)]$ and suppose we are given sets $F_{I,i-1}$, $\mathcal{F}_{I,i-1}$ and graphs D_{i-1}, W_{i-1} as input that satisfy the required invariants.

Step I: Let k_i^{I} be defined as follows.

$$k_i^{\mathrm{I}} := \Big(\Big(2^{12}\,(w+1)^3\, f_P\big(2 \cdot t^2\big)+1\Big) \cdot \\ \Big(2^6\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\Big)\Big)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)-(i-1)} \cdot d_2$$

For each of the two possible parametrisations of W, and for every possible choice of $\xi, \xi' \in [w+1]$, we consider the tiling $\mathcal{T} := \mathcal{T}_{W_{i-1},k_i^{\mathrm{I}},w,\xi,\xi'}$ together with its four-colouring $\{\mathcal{C}_1, \dots, \mathcal{C}_4\}$. For each $j \in [4]$, we consider the smallest slice W' of W that contains all vertices which belong to some tile of $\mathcal{T}$. Consider $D_j^{\mathrm{I}}(W')$ to be the auxiliary digraph of type I obtained from D_{i-1} with the sets

$$X_{\mathrm{I}}^{\mathrm{out}} := \left\{x_T^{\mathrm{out}} \,\middle|\, T \in \mathcal{C}_j\right\}, \text{ and} \\ X_{\mathrm{I}}^{\mathrm{in}} := \left\{x_T^{\mathrm{in}} \,\middle|\, T \in \mathcal{C}_j\right\}.$$

Additionally, we construct the set Y_{I} as follows: Let Q and Q' be the two cycles of $\mathsf{per}(W)$. For every $j \in \left[\frac{3d_1}{4}\right]$, Y_{I} contains exactly one vertex of $Q \cap P_{4j}^{\mathsf{o}}$, $Q \cap P_{4j+2}^{\mathsf{i}}$, $Q' \cap P_{4j}^{\mathsf{o}}$, and $Q' \cap P_{4j+2}^{\mathsf{i}}$ each. Then, remove all vertices of Y_{I} that do not belong to D_{i-1}. Note that, by choice of d_1 and the bound on $F_{\mathrm{I},i-1}$, this does not significantly decrease the size of Y_{I}.

Then, if there is a family of $2^7 f_P(2 \cdot t^2)$ pairwise disjoint directed $X_{\mathrm{I}}^{\mathrm{out}}$-$Y_{\mathrm{I}}$-paths in $D_j^{\mathrm{I}}(W')$, Lemma 5.4.12 implies the existence of a $D_t^{\boxtimes}$ butterfly minor, a contradiction. So, we may assume that there does not exist such a family and thus, by Theorem 5.4.20, we find a set $Z_1 \subseteq V(D_j^{\mathrm{I}}(W'))$ of size at most $2^7 f_P(2 \cdot t^2)$ that meets all these paths. With a similar argument, we either find a $D_t^{\boxtimes}$ butterfly minor, which would yield a contradiction, or a set $Z_2 \subseteq V(D_j^{1}(W'))$ of size at most $2^7 f_P(2 \cdot t^2)$ that meets all directed Y_{I}-$X_{\mathrm{I}}^{\mathrm{in}}$-paths in $D_j^{\mathrm{I}}(W')$. Let $\pi \in [2]$ indicate which of the two parametrisations of W we are currently considering. We define the following two sets:

$$\begin{aligned} Z_{\pi,\xi,\xi',j} &:= (Z_1 \cup Z_2) \cap V(D)\,, \text{ and} \\ \mathcal{Z}_{\pi,\xi,\xi',j} &:= \big\{ \mathsf{T} \in \mathcal{T} \bigm| \big(V(\mathsf{T}) \cup \{x_{\mathsf{T}}^{\mathrm{out}}, x_{\mathsf{T}}^{\mathrm{in}}\}\big) \cap (Z_1 \cup Z_2) \neq \emptyset \big\}. \end{aligned}$$

Note that $\max\{|Z_{\pi,\xi,\xi',j}|, |\mathcal{Z}_{\pi,\xi,\xi',j}|\} \leq 2^8 f_P(2 \cdot t^2)$.

As we do not find a $D_t^{\boxtimes}$ butterfly minor at any point, the sets $Z_{\pi,\xi,\xi',j}$ and $\mathcal{Z}_{\pi,\xi,\xi',j}$ are well defined for every possible choice of $\pi \in [2]$, $\xi, \xi' \in [w+1]$, and $j \in [4]$. Using these we define the following two sets of marked vertices and tiles:

$$\begin{aligned} F'_{\mathrm{I},i} &:= \bigcup_{\pi \in [2]} \bigcup_{\xi,\xi' \in [w+1]} \bigcup_{j \in [4]} Z_{\pi,\xi,\xi',j}, \text{ and} \\ \mathcal{F}'_{\mathrm{I},i} &:= \bigcup_{\pi \in [2]} \bigcup_{\xi,\xi' \in [w+1]} \bigcup_{j \in [4]} \mathcal{Z}_{\pi,\xi,\xi',j}. \end{aligned}$$

Consequently, we have

$$\max\{\left|F'_{\mathrm{I},i}\right|, \left|\mathcal{F}'_{\mathrm{I},i}\right|\} \leq 2^{11} \left(w+1\right)^2 f_P\big(2 \cdot t^2\big)\,.$$

Removing all vertices in $F'_{\mathrm{I},i}$ and tiles in $\mathcal{F}'_{\mathrm{I},i}$ yields a clear slice $W'_i \subseteq W_{i-1}$ of width

$$\begin{aligned} &\Big(2^{12} \left(w+1\right)^3 f_P\big(2 \cdot t^2\big) + 1\Big)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right) - i} \cdot \\ &\Big(2^{6} \left(w+1\right)^3 f_P\big(2 \cdot t^2\big) + 1\Big)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right) - i + 1} \cdot d_2 \end{aligned}$$

which does not contain a marked vertex. Note that we loose the additional factor of $2\,(w+1)$ because we remove whole tiles of width w from W_{i-1}. Let $D'_i := D_{i-1} - F'_{\mathrm{I},i}$. So now we have a preliminary version for all the structures we need to start the next round. This concludes Step I of round i.

Claim 1. Every long jump J over W'_i in D'_i whose endpoints belong to tiles of different colour contains a vertex of a tile from $\mathcal{F}_{\mathrm{I},j} \setminus \mathcal{F}_{\mathrm{I},j-1}$ for every $j \in [i-1]$, and it contains a vertex of a tile from $\mathcal{F}'_{\mathrm{I},i}$.

Proof. Suppose J is also a long jump over W_{i-1} in D_{i-1}, then, as J still exists in D'_i, the tile in whose centre J starts, or the tile in whose centre J ends for some choices of $\pi \in [2]$, $\xi, \xi' \in [w+1]$, and $j \in [4]$, are marked and therefore do not belong to W'_i.

Hence, J contains some vertex of W_{i-1} as an internal vertex. Let $\mathsf{T_s}$ be the tile of W'_i in whose centre J starts, and let T be the first tile from the same tiling of W_{i-1}, that J meets after $\mathsf{T_s}$. Let J' be the shortest subpath of J with endpoints in $\mathsf{T_s}$ and T. Then, J' is a long jump over W_{i-1} in D_{i-1}. Therefore, by our assumptions on the input of round i of Phase I, the first part of our claim is satisfied. Moreover, if T has a different colour than $\mathsf{T_s}$, then T is marked. So assume T has the same colour as $\mathsf{T_s}$. Nonetheless, because $\mathsf{T_s}$ and the tile $\mathsf{T_t}$ that contains the endpoint of J in the current tiling have different colours, J contains a directed subpath J'' which is a long jump over W_{i-1} and attaches to tiles of different colour. Hence our claim follows. ■

With this we are ready for Step II of round i.

Step II: For this step let

$$k_i^{\mathrm{II}} := \left(2^{12}\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\right)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)-i} \cdot$$
$$\left(2^{6}\,(w+1)^3\, f_P\big(2\cdot t^2\big)+1\right)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)-i+1} \cdot d_2.$$

We are mainly concerned with the digraph D'_i. In Step II it suffices to fix one parametrisation of W because the construction of the type II auxiliary digraph leaves the complete interior of tiles that are in the same colour class intact instead of only their centres. For every pair of $\xi,\xi'\in[w+1]$ we consider the tiling $\mathcal{T}:=\mathcal{T}_{W'_i,k_i^{\mathrm{II}},w,\xi,\xi'}$ together with a four-colouring $\{\mathcal{C}_1,\dots,\mathcal{C}_4\}$. Then, for every $j\in[4]$ we apply Lemma 5.4.14, which, as it cannot yield a $D_t^{\boxtimes}$, produces two sets

$$Z^2_{\xi,\xi',j}\subseteq V(D'_i) \text{ of size at most } 2^3 f_P\big(2\cdot t^2\big)\text{, and}$$
$$\mathcal{Z}^2_{\xi,\xi',j}\subseteq \mathcal{T} \text{ of size at most } 2^3 f_P\big(2\cdot t^2\big)\,,$$

such that every directed $V(W'_i)$-path whose endpoints belong to different tiles of $\mathcal{C}_j$, contains a vertex of some tile in $\mathcal{Z}^2_{\xi,\xi',j}$. This allows us to form the two following sets of marked vertices and tiles:

$$F''_{\mathrm{I},i} := \bigcup_{\xi,\xi'\in[w+1]}\ \bigcup_{j\in[4]} Z^2_{\xi,\xi',j}\text{, and}$$
$$\mathcal{F}''_{\mathrm{I},i} := \bigcup_{\xi,\xi'\in[w+1]}\ \bigcup_{j\in[4]} \mathcal{Z}^2_{\xi,\xi',j}.$$

As a result we obtain $\max\{|F''_{\mathrm{I},i}|,|\mathcal{F}''_{\mathrm{I},i}|\}\leq 2^5\,(w+1)^2\, f_P(2\cdot t^2)$, and we are able to produce the two sets $F_{\mathrm{I},i}$ of marked vertices and $\mathcal{F}_{\mathrm{I},i}$ of marked tiles, which are passed on to the next round.

$$F_{\mathrm{I},i} \coloneqq F'_{\mathrm{I},i} \cup F''_{\mathrm{I},i} \cup F_{\mathrm{I},i-1}, \text{ and}$$
$$\mathcal{F}_{\mathrm{I},i} \coloneqq \mathcal{F}'_{\mathrm{I},i} \cup \mathcal{F}''_{\mathrm{I},i} \cup \mathcal{F}_{\mathrm{I},i-1}.$$

The bounds on $F_{\mathrm{I},i}$ and $\mathcal{F}_{\mathrm{I},i}$ follow from the bounds on $F'_{\mathrm{I},i}$ and $F''_{\mathrm{I},i}$, $\mathcal{F}'_{\mathrm{I},i}$ and $\mathcal{F}''_{\mathrm{I},i}$, and the assumptions on the input of round i respectively.

The pigeon hole principle allows us to find a clear slice $W_i \subseteq W'_i$ of width

$$\left(\left(2^{12}\,(w+1)^3\, f_P\!\left(2\cdot t^2\right)+1\right)\cdot \left(2^{6}\,(w+1)^3\, f_P\!\left(2\cdot t^2\right)+1\right)\right)^{2^{11+8f_P\left(2\cdot t^2\right)} f_P\left(2\cdot t^2\right)-i}\cdot d_2,$$

which does not contain a marked vertex, that is, no vertex from $F_{\mathrm{I},i}$ or $\mathcal{F}_{\mathrm{I},i}$. Similar to Step I, we loose the additional factor of $2\,(w+1)$ because we remove whole tiles of width w from W'_i. Finally, let $D_i \coloneqq D'_i - F''_{\mathrm{I},i}$, which concludes Step II of round i.

Claim 2. Every long jump over W_i in D_i contains a vertex of some tile in $\mathcal{F}_{\mathrm{I},j} \setminus \mathcal{F}_{\mathrm{I},j-1}$ for every $j \in [i]$.

Proof. Let J be a long jump over W_i in D_i, and let $\mathcal{T}$ be a tiling of $\widetilde{W}_2$ defined by w and some $\xi, \xi' \in [w+1]$ such that J starts at the centre of some tile $\mathsf{T}_{\mathsf{s}} \in \mathcal{T}$. Suppose all tiles of $\mathcal{T}$ that contain vertices of J belong to the same colour. Then, J must have existed during the corresponding part of Step II of round i and thus either T_{s} or $\mathsf{T}_{\mathsf{t}} \in \mathcal{T}$, which is the tile that contains the endpoint of J, must have been marked, a contradiction. Therefore J must contain at least one tile of a colour different than the one of T_{s}. Moreover, we may assume T_{s} and T_{t} to be of the same colour as otherwise we would be done by Claim 1. Next, suppose J is also a long jump over W_{i-1}, then again J would have been considered during Step II as a long jump connecting two tiles of the same colour and thus T_{s} or T_{t} would have been marked, a contradiction. Therefore, J contains a vertex of some tile from W_{i-1}. Let J' be a shortest subpath from T_{s} to some tile T of W_{i-1}, then J' is a long jump over W_{i-1} and thus J contains a vertex of some tile of $\mathcal{F}_{\mathrm{I},j} \setminus \mathcal{F}_{\mathrm{I},j-1}$ for every $j \in [i-1]$ by our assumptions on the input of round i. If T has a different colour than T_{s}, then T would have been marked in Step I of round i, and if T shares the colour of T_{s}, then it must have been marked in Step II of round i. Either way our claim follows. ■

From Claim 2 it follows that we satisfy all requirements for the output of round i and thus, round i is complete. We continue until we finish round $2^{11+8f_P\left(2\cdot t^2\right)} f_P(2\cdot t^2)$ and obtain the following four objects as its output:

- a slice $W_{\mathrm{I}} := W_{2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)}$ of width d_2,
- a set $A_{\mathrm{I}} := F_{\mathrm{I},2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)}$ with

$$|A_{\mathrm{I}}| \leq (2 \cdot t^2)^{28} 2^{60+2^{10}(2\cdot t^2)^8},$$

- a digraph $D_{\mathrm{I}} := D_{2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)} = D - A_{\mathrm{I}}$, and
- a sequence $\mathcal{F}_{\mathrm{I},1} \subseteq \mathcal{F}_{\mathrm{I},2} \subseteq \cdots \subseteq \mathcal{F}_{\mathrm{I},2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)}$ such that for every $i \in [2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)]$ every long jump over W_{I} in the graph D_{I} contains a vertex of some tile in $\mathcal{F}_{\mathrm{I},i} \setminus \mathcal{F}_{\mathrm{I},i-1}$.

This brings us to the final claim of Phase I.

Claim 3. If there is a long jump over W_{I} in D_{I}, then there exists a $D_t^{\bowtie}$ as a butterfly minor in D.

Proof. Let J be a long jump over W_{I} in D_{I}. We fix a parametrisation of W, $\xi, \xi' \in [w+1]$, and $c \in [4]$ such that there exists a tile $\mathsf{T_s} \in \mathcal{T} := \mathcal{T}_{W_0,d_1,w,\xi,\xi'}$ of colour c whose centre contains the start-vertex of J. Let $\mathsf{T_t} \in \mathcal{T}$ be the tile that contains the end-vertex of J. As J is a long jump, note that $\mathsf{T_s} \neq \mathsf{T_t}$. We claim that every internal vertex of J that belongs to W belongs to some tile from $\mathcal{F}_{\mathrm{I},2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)}$. This is because otherwise we could find a directed path from the centre of $\mathsf{T_s}$ to the perimeter of $\widetilde{W}_2$, contradicting the construction in Step I of Phase I, or we would have a directed path between two tiles of the same colour, where both of them are unmarked. This second outcome contradicts the construction in Step II of Phase I.

Constructing $\mathcal{L}_0$. Now, we create a family $\mathcal{L}_0$ of $2^{9+8f_P(2\cdot t^2)} \cdot f_P(2\cdot t^2)$ pairwise disjoint subpaths of J with the following properties:

1. for every $L \in \mathcal{L}_0$, let $\mathsf{T}_{L,1}$ be the tile of $\mathcal{T}$ containing $\mathsf{start}(L)$ and $\mathsf{T}_{L,2}$ be the tiles of $\mathcal{T}$ containing $\mathsf{end}(L)$, then there exist distinct $i_{L,1}, i_{L,2} \in [2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)]$ such that $\mathsf{start}(L)$ is a vertex of a tile from $\mathcal{F}_{\mathrm{I},i_{L,1}} \setminus \mathcal{F}_{\mathrm{I},i_{L,1}-1}$, and $\mathsf{end}(L)$ is a vertex of some tile in $\mathcal{F}_{\mathrm{I},i_{L,2}} \setminus \mathcal{F}_{\mathrm{I},i_{L,2}-1}$, and
2. if $L, L' \in \mathcal{L}_0$ are distinct, then $\{i_{L,1}, i_{L,2}\} \cap \{i_{L',1}, i_{L',2}\} = \emptyset$.

We do so iteratively. We start by initialising $\mathcal{L}_0 = \emptyset$ and $\mathcal{I}_0 := [2^{11+8f_P(2\cdot t^2)}f_P(2\cdot t^2)]$ and for every subset $\mathcal{I}' \subseteq \mathcal{I}_0$, we define the family $\mathcal{F}_{\mathcal{I}'} := \bigcup_{i\in\mathcal{I}'} \mathcal{F}_{\mathrm{I},i} \setminus \mathcal{F}_{\mathrm{I},i-1}$. Next, we add new paths to $\mathcal{L}_0$ while taking smaller and smaller subsets of $\mathcal{I}_0$. Also, define $t_{L_0} := \mathsf{start}(J)$.

Let $q \in [2^{9+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)]$ and assume that the paths $L_1, \dots, L_{q-1}$ together with the tiles, indices and the set $\mathcal{I}_{q-1}$ have already been constructed. Follow along J, starting from $t_{L_{q-1}}$, until the next time we encounter the last vertex s_{L_q} of some tile from $\mathcal{F}_{\mathcal{I}_{q-1}}$ before J leaves said tile again. Let $\mathsf{T}_{L_q,1} \in \mathcal{T}$ be the tile that contains s_{L_q}, and let $i_{L_q,1} \in \mathcal{I}_{q-1}$ be the integer such that s_{L_q} belongs to a tile of $\mathcal{F}_{\mathrm{I},i_{L_q,q}} \setminus \mathcal{F}_{\mathrm{I},i_{L_q,q}-1}$. Then, let L_q be the shortest subpath of J that starts in s_{L_q} and ends in a vertex t_{L_q} which belongs to a tile from $\mathcal{F}_{\mathrm{I},i_{L_q,2}} \setminus \mathcal{F}_{\mathrm{I},i_{L_q,2}-1}$, where $i_{L_q,2} \in \mathcal{I}_{q-1} \setminus \{i_{L-q,1}\}$. We choose $\mathsf{T}_{L_q,2} \in \mathcal{T}$ to be the tile that contains t_{L_q} and set $\mathcal{I}_q := \mathcal{I}_{q-1} \setminus \{i_{L_q,1}, i_{L_q,2}\}$. Note that $t_{L_q} J$ still contains a vertex from some tile in $\mathcal{F}_{\mathrm{I},j} \setminus \mathcal{F}_{\mathrm{I},j-1}$ for every $j \in \mathcal{I}_q$. Add L_q to $\mathcal{L}_0$.

With every iteration we remove exactly two members from $\mathcal{I}_0$ and, due to $|\mathcal{I}_0| = 2^{11+8f_P(2\cdot t^2)} f_P(2\cdot t^2)$, this means that by the time we reach some q for which $\mathcal{I}_q = \emptyset$, we have indeed constructed $2^{10+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ paths as required.

Obtaining $\mathcal{L}_5$. There exist $c' \in [4]$ and a linkage $\mathcal{L}_1 \subseteq \mathcal{L}_0$ of size $2^{8+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ such that each path $L \in \mathcal{L}_1$ has at least its start- or end-vertex in $\mathcal{C}_{c'}$. Thus, we can find a linkage $\mathcal{L}_2 \subseteq \mathcal{L}_1$ of size $2^{7+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ such that every path in $\mathcal{L}_2$ starts in a tile of $\mathcal{C}_{c'}$, or every path in $\mathcal{L}_2$ ends in a tile of $\mathcal{C}_{c'}$. Without loss of generality, we may assume that every path in $\mathcal{L}_2$ starts in a tile of $\mathcal{C}_{c'}$, because the other case follows with similar arguments.

Let $\widetilde{W}'$ be the smallest slice of W such that $\widetilde{W}'$ contains all tiles from $\mathcal{C}_{c'}$, but no tile from $\mathcal{C}_{c'}$ meets the perimeter of $\widetilde{W}'$. Then, let $\widetilde{\mathcal{T}} := (\mathcal{T}, c', w)_{\mathrm{II}}\left[\widetilde{W}'\right]$ be the tier II tiling of $\widetilde{W} := \widetilde{W}'[\mathcal{T}, c', w]$. Since the paths in $\mathcal{L}_2$ are pairwise disjoint, we can extend each $L \in \mathcal{L}_2$ such that it starts on the centre of the tile of $\widetilde{\mathcal{T}}$ which encloses its endpoint in W, while making sure that the resulting family of paths is still at least half-integral. Similarly, wherever necessary, we may extend the paths through W such that each of them also ends in a tile of $\widetilde{\mathcal{T}}$. Indeed, we can even guarantee that the start- and end-vertices of the resulting paths are mutually at $\widetilde{W}$-distance at least four. Let $\mathcal{L}_3$ be the resulting half-integral linkage.

Next, consider the four-colouring $\{\widetilde{\mathcal{C}}_1, \dots, \widetilde{\mathcal{C}}_4\}$ of $\widetilde{\mathcal{T}}$. There exists $\widetilde{c} \in [4]$ and a family $\mathcal{L}_4 \subseteq \mathcal{L}_3$ of size $2^{5+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ such that every path in $\mathcal{L}_4$ starts at the centre of some tile from $\mathcal{C}_{\widetilde{c}}$. It follows from the construction of $\mathcal{L}_0$ that no two paths in $\mathcal{L}_4$ start in the same tile.

By a similar argument, there exists a family $\mathcal{L}_5 \subseteq \mathcal{L}_4$ of size $2^{4+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ such that either none, or all paths in $\mathcal{L}_5$ end in tiles of $\widetilde{\mathcal{C}}_{\widetilde{c}}$.

Obtaining $D_t^{\boxtimes}$. If none of the paths in $\mathcal{L}_5$ end in tiles of $\widetilde{\mathcal{C}}_{\widetilde{c}}$, we can extend every path in $\mathcal{L}_5$ towards the perimeter of $\widetilde{W}$ such that the resulting family $\mathcal{L}_6$ of paths remains at worst half-integral, and the start- and end-vertices of the resulting paths are mutually at $\widetilde{W}$-distance at least four. By Theorem 5.4.8, we obtain a family $\mathcal{L}_7$ of size $2^{3+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ such that $V(\mathcal{L}_7) \subseteq V(\mathcal{L}_6)$, and the paths in $\mathcal{L}_7$ are pairwise vertex disjoint. Hence, Lemma 5.4.12 yields the existence of a $D_t^{\boxtimes}$ butterfly minor and our claim follows.

If all of the paths in $\mathcal{L}_5$ end in tiles of $\widetilde{\mathcal{C}}_{\widetilde{c}}$, we consider two subcases. Let $\mathcal{X}$ be the family of all tiles of $\widetilde{\mathcal{T}} \setminus \widetilde{\mathcal{C}}_{\widetilde{c}}$ that contain an internal vertex of some path in $\mathcal{L}_5$ but no start- or end-vertex of any path in $\mathcal{L}_5$.

Recall the following two definitions:

1. If L and P are directed paths, we say that P is a *long jump of* L if P is a w-long jump over W and $P \subseteq L$.
2. P is a *jump of* L, if P is a directed $V(W)$-path.

If $|\mathcal{X}| \geq 2^8 f_P(2 \cdot t^2)$, then we can use the technique from the first part of the proof of Lemma 5.4.12 to construct a half-integral family $\mathcal{L}_6$ such that

1. $|\mathcal{L}_6| = 2^{4+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$, and
2. for every $L \in \mathcal{L}_6$, every endpoint u of a jump of L with $u \in V(\widetilde{W})$ belongs to a tile from $\widetilde{\mathcal{C}}_{\widetilde{c}} \cup \mathcal{X}$.

Then, we apply Theorem 5.4.8 to obtain a linkage $\mathcal{L}_7$ of size $2^{3+8f_P(2\cdot t^2)} f_P(2 \cdot t^2)$ with $V(\mathcal{L}_7) \subseteq V(\mathcal{L}_6)$ such that the paths in $\mathcal{L}_7$ link the same two sets of vertices as the paths in $\mathcal{L}_6$ do. Finally, Lemma 5.4.12 yields the existence of a $D_t^{\boxtimes}$ as a butterfly minor.

If $|\mathcal{X}| < 2^8 f_P(2 \cdot t^2)$, then we find a subwall $\widetilde{W}'$ of $\widetilde{W}$ of order $d_1 - 2^8 f_P(2 \cdot t^2)\,(2w+1)$ that does not contain a vertex of any tile in $\mathcal{X}$. We do this by removing, for every tile $\mathsf{T} \in \mathcal{X}$, all edges and vertices of the horizontal and vertical paths of T that are not used by other cycles or paths. For each tile we remove during this procedure, we remove a row and a column of tiles and thereby reduce the number of distinct tiles which contain start-vertices of paths in $\mathcal{L}_5$ by a factor of at most $\frac{1}{2}$. However, because $|\mathcal{X}| < 2^8 f_P(2 \cdot t^2)$, we can still find, after potentially expanding the start and end sections of the paths in $\mathcal{L}_5$ in order to reach the slightly shifted perimeters of their tiles, a half-integral family $\mathcal{L}_6$ of size $2^4 f_P(2 \cdot t^2)$ of paths that start and end in tiles of $\widetilde{\mathcal{C}}_{\widetilde{c}}$ and that are otherwise disjoint from $\widetilde{W}'$. By applying Theorem 5.4.8 we can transform this family into a linkage $\mathcal{L}_7$ of size $2^3 f_P(2 \cdot t^2)$ and thus an application of Lemma 5.4.14 yields a $D_t^{\boxtimes}$ as a butterfly minor. ■

Concluding Phase I, Claim 3 either yields $D_t^{\boxtimes}$ as a butterfly minor and therefore finishes the proof, or W_{I} is in fact clean, meaning that W_{I} has no long jump in D_{I}. Consequently we may bound the function α_W from the statement of Theorem 5.4.3 as follows:

$$\alpha_W(t) \leq \left(2 \cdot t^2\right)^{28} \cdot 2^{61+2^{10}\left(2 \cdot t^2\right)^8}.$$

Phase II With W_{I} we have found a wall of still sufficient size but without any long jumps. Therefore, me may now find t slices of W_{I}, mutually still far enough apart from each other within W_{I}, and can ask if among them there is one that is flat. If so, then we have found the desired flat wall within W. If not, each of the t slices contains a cross and an application of Lemma 5.2.5 yields the desired $D_t^{\boxtimes}$ butterfly minor. The only technical part that remains is to provide sufficient definitions for these slices and their mutual distance.

To meet the requirements from Phase I and have enough space left in W_{I}, let us make the following assumption:

$$d_2 \geq t\left(r + 4 + 2^{32+(3t)^{30}}\right).$$

We partition W_{I} further into smaller slices. First we partition W_{I} into t slices S_i of width $r + 4 + 2^{32+(3t)^{30}}$. Each S_i is then partitioned into a slice H_i of width $r + 4 + 2^{31+(3t)^{30}}$ that contains the left perimeter cycle of the slice S_i, and a slice G_i of width 2 containing the right perimeter cycle of S_i. For every $i \in [t]$ we may now further partition H_i. Let $N_{i,L} \subseteq H_i$ be the slice of width $2^{30+t^{30}}$ containing the left cycle of $\mathsf{per}(W)$, let $N_{i,R}$ be the slice of width $2^{30+t^{30}}$ containing the right cycle of $\mathsf{per}(W)$, and let $N_i' := H_i - N_{L,i} - N_{R,i}$ be the remaining slice of width $r + 4$. Finally, let N_i be the slice obtained from N_i' by removing the two leftmost and the two rightmost vertical cycles. Then, N_i is a slice of width r.

For every $i \in [t]$ we show that if there exists a directed path P_i with one endpoint in N_i, the other endpoint in $W_{\mathrm{I}} - N_i'$, and which is internally disjoint from W_{I}, then H_i contains a cross slice such that the cross lies in a strip of size at most $2^{30+t^{30}} + 2$. In this case, there is a vertical path Q in one of the two components of $N_i' - N_i$ which separates the start- and the end-vertex of P_i in W_{I}. As there are no long jumps over W_{I} in D_{I}, we further know that there exists $Y \in \{L, R\}$ such that the endpoint of P_i not in N_i lies in $N_{Y,i}$ within a strip of height at most $2^{30+t^{30}} + 2$. By Lemma 5.4.15, this yields a cross slice within H_i. Let $J \subseteq [t]$ contain every $i \in [t]$ for that such a P_i exists.

Suppose $J = [t]$, then we can use Lemma 5.2.5 and obtain a $D_t^{\boxtimes}$ butterfly minor, a contradiction. The row condition is met as the strips are all of height at most $2^{30+t^{30}}+2$ leaving a strip of height at least t.

So, there is at least one $i \in [t]$ for which such a path does not exist, let $J' := [t] \setminus J$. Suppose for every $i \in J'$ the strong component of $D_{\mathrm{I}} - \mathsf{per}(W)$ that contains N_i has a cross slice. Since there is no long jump over W_{I}, these components are pairwise disjoint and also disjoint from the cross slices found in the S_j, $j \in J$, and thus we can again apply Lemma 5.2.5 to obtain a $D_t^{\boxtimes}$ butterfly minor, a contradiction.

Thus, there must exist some $i \in J'$ for which N_i' has no cross slice. In particular this means that the component of $D_{\mathrm{I}} - \mathsf{per}(W)$ containing the remaining vertices of $N_i - \mathsf{per}(W)$ must be free of cross slices. Hence, by Lemma 5.4.19, N_i is a flat wall of order r in $D_{\mathrm{I}} = D - A$ which completes the proof.

Let us combine all assumptions on the d_1 and d_2 to obtain the following bound on $\rho_W(t,r)$:

$$\rho_W(t,r) \leq \left(2^{140}t^{72}\right)^{2^{10}t^8+2^{12}} \left(r+4+2^{32+(3t)^{30}}\right). \qquad \square$$

Note that this proof is close to the original ones [GKKK20, GW21] and thus, we often obtain clique minors, when with a probably better function, we could obtain a cross-row grid instead. Therefore, it might be possible to obtain overall better functions by choosing a procedure more specific to the structure of $D_t^{\boxtimes}$.

5.5 A flat wall theorem excluding $K_5^{\rightarrow}$

Before proving the graph structure theorem, Robertson and Seymour proved a similar statement for single-crossing minor-free graphs [RS93]. To this end, Robertson and Seymour already use similar techniques as later for the proof of the general graph structure theorem.

The first result in this direction was the proof of Wagner's theorem [Wag37] about graphs excluding an undirected clique on five vertices. This yields a class of graphs that can be obtained from planar graphs and the Wagner graph by an operation called *small clique sums*. Inspired by this, we consider the structure of digraphs which exclude the unique strongly 2-connected orientation of the K_5. We give this orientation in
$K_5^{\rightarrow}$ Figure 5.10 and refer to it as $K_5^{\rightarrow}$ throughout the remainder of the chapter. Excluding

the $K_5^{\rightarrow}$ as a butterfly minor yields a slightly stronger flat wall theorem with a much simpler proof.

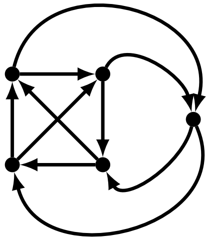

Figure 5.10: The strongly 2-connected orientation $K_5^{\rightarrow}$ of K_5.

As the following observation shows, excluding $K_5^{\rightarrow}$ directly excludes $D^{\times} := D_2^{\times}$ $D^{\times}$
and $D^{\otimes}$, which is obtained from a wall of order 2 by adding a cross over one of the $D^{\otimes}$
perimeters, see Figure 5.11 for an illustration of both digraphs.

Observation 5.5.1. We have $K_5^{\rightarrow} \preccurlyeq_b D^{\times}$ and $K_5^{\rightarrow} \preccurlyeq_b D^{\otimes}$.

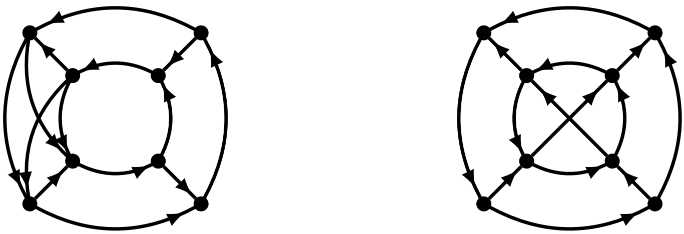

Figure 5.11: On the left the digraph $D^{\times}$ on the right the digraph $D^{\otimes}$. Both contain $K_5^{\rightarrow}$ as a butterfly minor.

Surprisingly, even for such a small case like $K_5^{\rightarrow}$, there remain certain kinds of jumps in a wall that cause the digraph containing it to be non-planar, but which do not produce a $K_5^{\rightarrow}$ butterfly minor. So, excluding the $K_5^{\rightarrow}$ always leaves the possibility of their existence. These are short backwards jumps with respect to the direction of the wall, see Figure 5.12 for an example. The start- and end-vertex of these jumps are not separated by any cycle of the wall and by at most one in- or out-path.

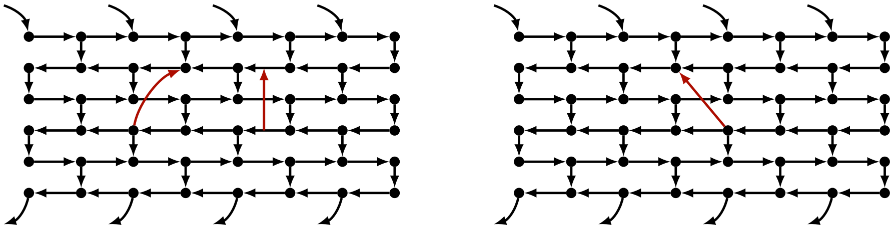

Figure 5.12: Backwards jumps in the wall that do not produce a $\overrightarrow{K_5}$ butterfly minor.

Still, most occurrences of non-planarity under the infrastructure offered by a cylindrical wall create butterfly minor models of $\overrightarrow{K_5}$. So, for digraphs excluding $\overrightarrow{K_5}$ we can prove that already for small walls there are only restricted jumps and we can even prove that there are no jumps from one side to the other side of the wall.

A wall W is *separating* if there is no *jump over all of* W, that is, there is no path with start-vertex on the one perimeter of W, end-vertex on the other and otherwise internally disjoint from W.

Lemma 5.5.2. Let W be a wall of order $k \geq 5$ in a digraph D excluding $D^{\searrow\!\!\!\!\nearrow}$. Then, W is separating.

Proof. Suppose there is a path Q starting in a vertex x on $\mathsf{per}^{\mathsf{in}}(W)$ and ending in a vertex y on $\mathsf{per}^{\mathsf{out}}(W)$ and being internally disjoint from W. The case that the path goes in the other direction works analogously. Let $P^{\mathsf{i}}(x)$ be the next in-path encountered when following C_1 backwards from x and let $P^{\mathsf{i}}(y)$ be the next in-path encountered when following C_k forward from y.

As there are at least five in-paths we have (I) another in-path starting along C_k after $P^{\mathsf{i}}(y)$ and before $P^{\mathsf{i}}(x)$, or (II) three other in-paths starting along C_k after $P^{\mathsf{i}}(x)$ and before $P^{\mathsf{i}}(y)$.

We obtain a subgraph $D' \subseteq D$ as follows, see Figure 5.13 for an illustration of this construction. We completely include the two cycles C_1 and C_4. Then, we define the path Q_a by starting with Q and then continuing along the subpath of C_5 starting in $\mathsf{end}(Q) = y$ and ending in $\mathsf{start}(P^{\mathsf{i}}(y))$ (this second part might be empty). We also add the subpath $Q_{a,a'}$ of $P^{\mathsf{i}}(y)$ starting on C_5 and ending in a vertex a' on C_4. In case (I), let b be the intersection vertex of $P^{\mathsf{i}}(x)$ and C_4 and $Q_{b'}$ be the subpath of $P^{\mathsf{i}}(y)$ starting in C_2 and ending on C_1. In case (II), let b be the intersection vertex between the in-path one index smaller than $P^{\mathsf{i}}(y)$ and C_4, moreover, let $Q_{b'}$ be the subpath of the in-path one index larger than $P^{\mathsf{i}}(x)$ that starts on C_2 and ends on C_1. We then define the path $Q_{b,b'}$ by starting at b, following the in-path b lies on until meeting C_2

and then following C_2 until $\mathsf{start}(Q_{b'})$. From a' we add a path $Q_{a',b'}$ that starts by following $P^{\mathsf{i}}(y)$ until meeting C_3 and then following along C_3 (this might be just a single vertex) until meeting the in-path $Q_{b,b'}$ is a subpath of, which we then follow until $\mathsf{start}(Q_{b,b'})$. We add the path $Q_{a,b}$ which starts in $\mathsf{end}(Q_a)$ and walks along C_5 until meeting the in-path b lies on, which it follows until reaching b. Now, note that, after the in-path containing $Q_{b'}$ and before the in-path containing b, the cycles meet another in-path P^{i}_j and before that the out-path P°_{j-1}. So finally, we add the paths $Q_{c,c'}$ and $Q_{d,d'}$ to D', where $Q_{c,c'}$ is the subpath of P°_{j-1} starting on C_1 and ending on C_4 and $Q_{d,d'}$ is the subpath of P^{i}_j starting on C_4 and ending on C_1.

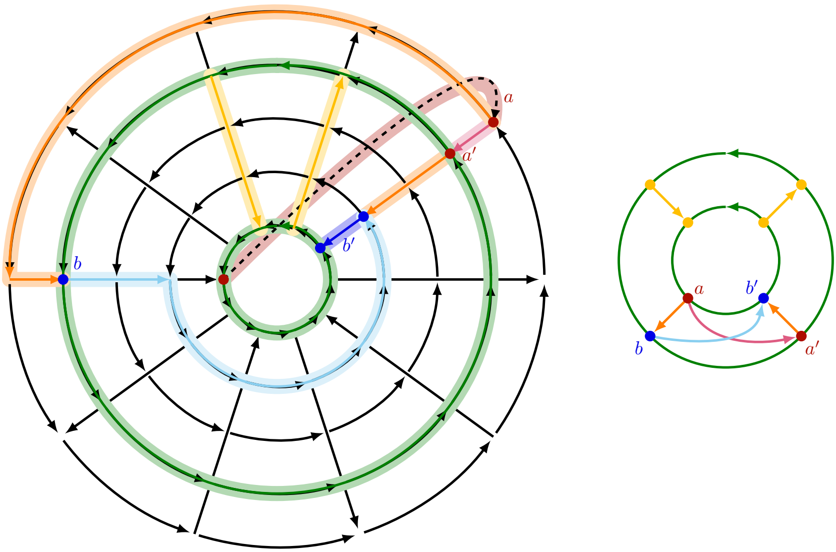

Figure 5.13: How to construct $D^{\boxtimes}$ as a butterfly minor in the wall of order five with a jump from the inner to the outer perimeter.

Note that Q_a and $Q_{b'}$ are contractible into a single vertex in D' as $\mathsf{end}(Q_a)$ has no further in-edges and $\mathsf{start}(Q_{b'})$ has no further out-edges. Contracting the remaining paths we added into single edges, we obtain $D^{\boxtimes}$, a contradiction. □

Theorem 5.5.3. Let D be a digraph excluding $D^{\boxtimes}$ and $D^{\otimes}$ as butterfly minors, then every wall of order $k \geq 5$ in D is flat and separating.

Proof. Let W be a wall of order $k \geq 5$ in D. By Lemma 5.5.2, W is separating.

We claim that every cross slice contains $D^{\boxtimes}$ as a butterfly minor and thus $K_5^{\leftrightarrow}$. To see this, let J_1 and J_2 build a cross in the slice between Q_i and $Q_{i+\ell}$. Let R_a be the row containing $\mathsf{start}(J_1)$, $R_{a'}$ be the row containing $\mathsf{end}(J_1)$, R_b be the row containing $\mathsf{start}(J_2)$ and $R_{b'}$ be the row containing $\mathsf{end}(J_2)$. Then, we have $\min\{a, a'\} - \max\{b, b'\} > 1$. Thus, there is a row R_j such that $b, b' \leq j \leq a, a'$ separating the start-vertices of J_1 and J_2 from their end-vertices. We take the subgraph D' of D containing Q_i,$Q_{i+\ell}$,J_1,J_2, the subpath of P_j^{o} lying between Q_i and $Q_{i+\ell}$, the subpath of P_j^{i} lying between Q_i and $Q_{i+\ell}$, the subpath of $P_{\min\{a,a'\}}^{\mathsf{o}}$ lying between Q_i and $Q_{i+\ell}$, and the subpath of $P_{\max\{b,b'\}}^{\mathsf{i}}$ lying between Q_i and $Q_{i+\ell}$. By butterfly contractions we can obtain $D^{\boxtimes}$ and finally $K_5^{\leftrightarrow}$, a contradiction.

Therefore, W does not contain cross slices and thus is flat by Lemma 5.4.19. $\square$

Finally, we exclude another type of jumps. Let W be a wall of order k in a digraph D. Then, a directed path that is internally disjoint from W in D with start-vertex in a row R_j, with $1 \leq j < k$, and end-vertex in the row R_{j+1}, such that both endpoints lie in $W - \mathsf{per}(W)$ and are separated by an in- or out-path of W, is called a *short forward jump* with respect to W. We can even exclude the existence of short forward jumps when excluding $K_5^{\leftrightarrow}$ as a butterfly minor.

Lemma 5.5.4. Let W be a wall of order $k \geq 5$ in a digraph D excluding $D^{\boxtimes}$ and $D^{\otimes}$. Then, there is no short forward jumps with respect to W.

Proof. Suppose such a path Q exists and let $\mathsf{start}(Q)$ lie on the cycle C_i and in the row R_j, thus $\mathsf{end}(Q)$ lies in the row R_{j+1}. By Lemma 5.4.17, we additionally know that $\mathsf{end}(Q)$ lies on C_{i-1}, C_i or C_{i+1}.

If $\mathsf{start}(Q)$ lies on P_j^{i} and $\mathsf{end}(Q)$ lies on P_{j+1}^{o}, then there is a face containing both endpoints of Q, contradicting it being separated by an in- or out-path. If $\mathsf{start}(Q)$ lies on P_j^{o} and $\mathsf{end}(Q)$ lies on P_{j+1}^{i}, then we can switch the parametrisation of the wall such that there lies a whole row between the start- and end-vertex of Q. This allows us to use Lemma 5.4.18, which yields the existence of $D^{\boxtimes}$ as a butterfly minor in D, a contradiction.

Thus, we can assume that both endpoints lie in the in-path of their respective row or both endpoints lie on the out-path of their respective row. Assume without loss of generality that both lie on the out-path. We continue with a case distinction on where the end-vertex of Q lies.

In case $\mathsf{end}(Q)$ **lies on** C_{i-1} **or** C_i**, for** $2 < i < k$, we construct two paths P_a and P_b as follows. The path P_a starts in the intersection of C_{i+1} and P_{j-1}^{i}, from there it follows P_{j-1}^{i} until meeting C_i, which it follows until meeting $\mathsf{start}(Q)$, then it follows

Q until its end, which lies on P^{o}_{j+1}, then it follows P^{o}_{j+1} until meeting C_{i+1} again. The path P_b starts on the intersection of C_{i+1} and P^{i}_j, from there it follows P^{i}_j until C_{i-2}, which it follows until reaching P^{o}_{j+2}, which it follows until it reaches C_{i+1} again.

The cycles C_i and C_{i+1} together with the paths P_a and P_b yield $D^{\otimes}$ as a butterfly minor, a contradiction. See Figure 5.14 for an example with $\mathsf{end}(Q)$ on C_{i-1}.

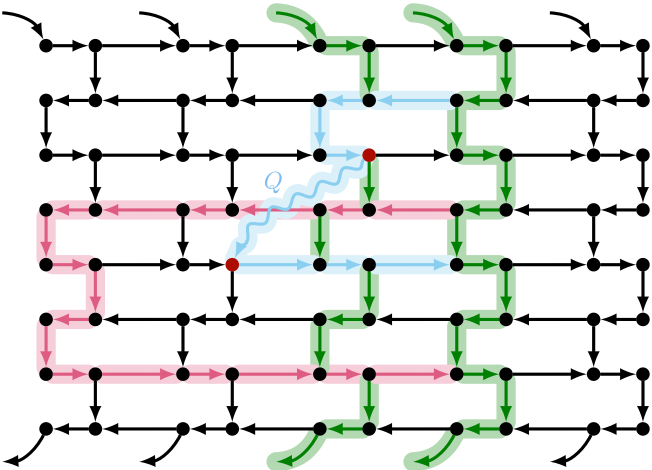

Figure 5.14: How $D^{\otimes}$ is obtained as a butterfly minor in case $\mathsf{end}(Q)$ lies on C_{i-1}.

In case $\mathsf{end}(Q)$ **lies on** C_{i+1}, we prove that the graph contains $D^{\boxtimes}$ as a butterfly minor. To this end we identify a butterfly model μ for the vertices and edges in $D^{\boxtimes}$ named as in Figure 5.15. For the vertex a we define $\mu_a := \mathsf{start}(Q)$ to $\mu(D^{\boxtimes})$ and for the edge (a, a') we define $\mu_{a,a'} := Q$. Next, we choose the subpath $\mu_{a,b}$ of P^{o}_j that starts in μ_a and ends on C_{i+1}. We define μ_b to be the subpath of C_{i+1} starting in $\mathsf{end}(\mu_{a,b})$ and ending on P^{i}_j. Choose $\mu_{b,b'}$ to be the subpath of P^{i}_j starting in $\mathsf{end}(\mu_b)$ and ending on C_i. Next, we define $\mu_{b'}$ to be the subpath of C_i that starts in $\mathsf{end}(\mu_{b,b'})$ and ends on P^{i}_{j+1}. Additionally, we choose $\mu_{a'}$ to be the subpath of C_{i+1} starting in $\mathsf{end}(\mu_{a,a'})$ and ending on P^{i}_{j+1}. We choose $\mu_{a',b'}$ to be the subpath of P^{i}_{j+1} starting in $\mathsf{end}(\mu_{a'})$ and ending in $\mathsf{end}(\mu_{b'})$. Now choose $\mu_{c,c'}$ to be the subpath of P^{o}_{j+2} starting on C_i and ending on C_{i+1} and finally, choose $\mu_{d,d'}$ to be the subpath of P^{i}_{j+2} starting on C_{i+1} and ending on C_i.

The subgraph D' of D containing C_i, C_{i+1} and all the above defined paths: μ_a, $\mu_{a'}$, μ_b, $\mu_{b'}$, $\mu_{a,a'}$, $\mu_{b,b'}$, $\mu_{a,b}$, $\mu_{a',b'}$, $\mu_{c,c'}$, and $\mu_{d,d'}$, now yields the desired model μ of $D^{\boxtimes}$ as illustrated in Figure 5.15. □

Thus, we obtain that in a digraph excluding $K_5^{\rightarrow}$ not only every wall of order at least five is flat and separating, Theorem 5.5.3, but it additionally excludes short forward jumps.

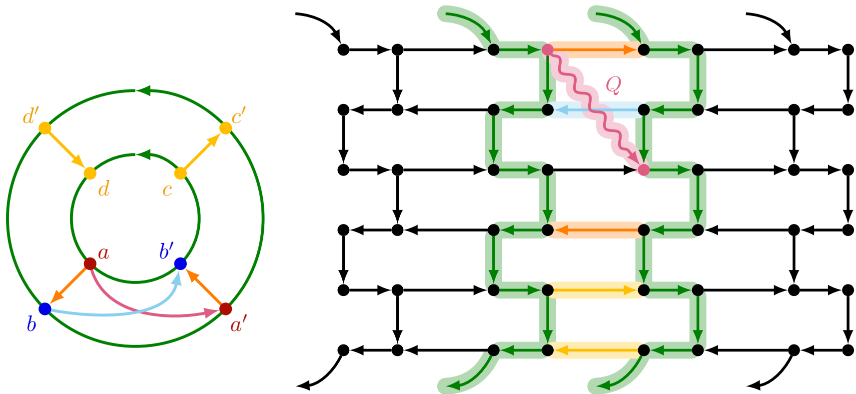

Figure 5.15: The digraph $D^{\nwarrow\!\!\nearrow}$ on the left and its model in case $\mathsf{end}(Q)$ lies on C_{i+1} on the right.

5.6 Where to go from here?

This chapter only considers some first steps in direction of the much larger problem of giving a good description of the structure of digraphs excluding a fixed butterfly minor. But it also lays some foundations one can build on when proceeding to a structure theorem.

The definition for skeletons stated here already has the potential to be extended to contain more than just the wall in its rigid subgraph. Skeletons give the possibility to also add planar transactions to the rigid subgraph, thus fixing a planar drawing for larger parts of the digraph. In undirected graphs, such a transaction splits an existing face into two parts that we can consider separately and which are each again surrounded by a family of cycles, which is called a *nest*, see the left side of Figure 5.16 for an illustration. This is not the case with the faces defined by the perimeters of the directed wall. Here a transaction splits a face into two parts that are of different structure, as we can see on the right side of Figure 5.16, where only one side builds cycles again. The boundary of the other resulting face, marked in the figure in orange, is bound by undirected cycles that split into two directed paths. Adding further transactions within such faces further complicates the boundaries of the faces. Another complication is caused if the transaction has further intersections with the already fixed planar structure. In that case one obtains several new faces instead of only two, which can be separated by much or only very little of the collected structure so far. Thus, a directed concept of nests has to grasp a lot more structure than the undirected notion does. Once identified, the definition of flatness needs to be extended to cover the structures that can appear as rigid subgraphs in skeletons. This step seems to be rather straightforward with our definition of flatness, as the condition on the separating property of nooses can be used for more general rigid graphs as well. For the traditional definition of flatness one has to consider the "direction" of the rigid graph, which might be more tricky to obtain when the rigid graphs more complicated than a wall.

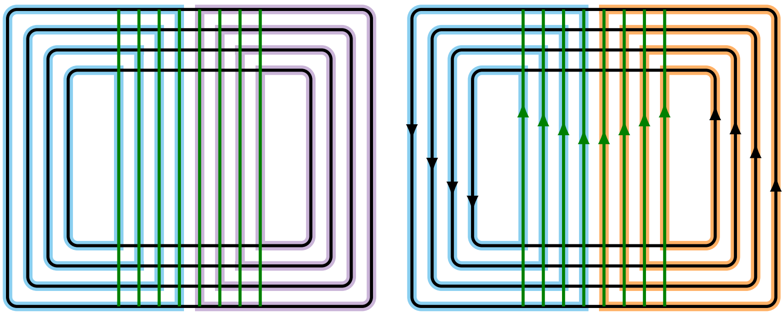

Figure 5.16: On the left we depict an undirected cycle family with a planar transaction. Note that the transaction builds new cycle families, marked in blue and violet, around both new faces. On the right we depict a directed cycle family with a planar transaction. While the family of undirected cycles marked in blue still yields a family of directed cycles, the undirected cycle family marked in orange does not yield a family of directed cycles around the new face.

The central idea of such an approach would be to add transactions to the skeleton as often as possible and then consider the parts that remain drawn into the faces of the rigid graph. Here, big and small faces have to be considered and the structure obtained from the rigid subgraph has to be used to restrict the number of faces containing extensive non-planar behaviour. This leads to a natural analogue of vortices and nests around them as they are used by Kawarabayashi, Thomas and Wollan [KTW20].

Dealing with vortices, areas of non-planarity, turns out to be more difficult in digraphs. This is mainly due to the lack of a 2-paths-theorem, that is, not containing a cross over a society does not necessarily imply that there exists a planar embedding. It is surprising that even when excluding $K_5^{\overrightarrow{}}$ there can still be non-planar behaviour as seen in figures 5.12 and 5.17a. The question that arises is how we can get rid of all non-planar behaviour.

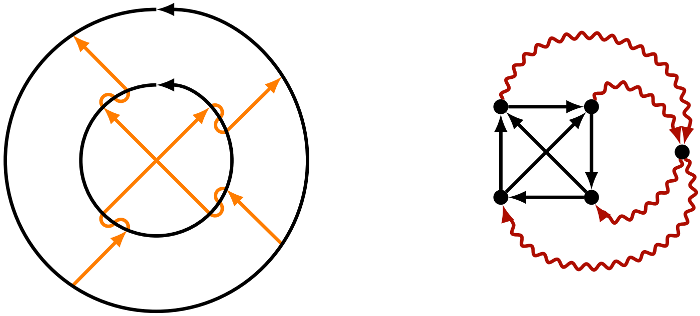

(a) A cross over a society that does not yield a $K_5^{\overrightarrow{}}$.

(b) The weak $K_5^{\overrightarrow{}}$. The wiggly red edges visualise paths.

Figure 5.17: Crosses and $K_5^{\overrightarrow{}}$.

Giannopoulou and Wiederrecht [GW21] consider excluding not only a single digraph as a butterfly minor, but a whole family related to it, called the corresponding *antichain*. This corresponds to forbidding the bipartite graph obtained by splitting the vertices of the digraph (as seen in Chapter 4) as a matching minor. They prove a directed flat wall theorem for excluding this possibly infinitely large class of digraphs. Unfortunately, these digraphs may not only be infinitely many but they are also structurally hard to grasp and describe from a digraph theoretic point of view.

Another option to handle the problematic crosses which are demonstrated in figures 5.12 and 5.17a would be to exclude something we call a *weak* $K_5^{\overrightarrow{}}$, which, as illustrated in Figure 5.17b uses a butterfly minor of an orientation of K_4 as its base and then models the remaining four edges as four disjoint paths that are allowed specific intersections with the butterfly model of the K_4 orientation.

Yet another way of dealing with problematic non-planar behaviour would be to slightly adjust the minor relation we consider. However, butterfly minors have the advantage of being related to matching minors, as seen in Theorem 4.2.2, which yields a number of valuable structural insights that could prove useful in proving a directed structure theorem.

6 Induced subgraphs

In contrast to the other chapters, this one is mostly concerned with undirected graphs and induced subgraphs. Induced subgraphs are a stronger way of describing that one graph is contained in another than general subgraphs or minors. In graph structure theory they are often used to describe classes of graphs. There are two ways to do so, one is by demanding the containment of certain induced subgraphs, e.g. a graph is distance hereditary if and only if every path in it is induced. The other is by forbidding certain induced subgraphs. The most famous example for that is probably the strong perfect graph theorem by Chudnovsky, Robertson, Seymour and Thomas, who proved the following statement after it was open for about 40 years.

Theorem 6.0.1 (Strong Perfect Graph Theorem [CRST06])**.** A graph is perfect if and only if it contains neither C_k nor $\overline{C}_k$ as induced subgraph for odd $k \geq 5$.

Many interesting problems in graph theory involve induced subgraphs. An example is a more restrictive version of DISJOINT PATH, the INDUCED DISJOINT PATHS problem, which asks to find k mutually induced paths between k given pairs of vertices, that is, any two paths are not allowed to have common vertices or adjacent vertices. This problem was introduced by Kawarabayashi and Kobayashi [KK08] in 2008 and is NP-hard even for $k = 2$ on general graphs. If we consider the class of AT-free graphs, see the next paragraph for a formal definition, then the problem becomes solvable in polynomial time [GPv12]. Köhler [Kö99] proved that the class of AT-free graphs is characterisable by a family of forbidden induced subgraphs.

6.1 Graphs with at most two moplexes

A moplex is a natural graph structure that arises when lifting Dirac's theorem, Theorem 2.1.1, from chordal graphs to general graphs. The notion is known to be closely related to lexicographic searches in graphs as well as to asteroidal triples, and has been applied in several algorithms related to graph classes such as interval graphs, claw-free, and diamond-free graphs. However, though every non-complete graph has at least two moplexes, little is known about structural properties of graphs with a bounded number of moplexes. The study of these graphs is, in part, motivated by the parallelism between moplexes in general graphs and simplicial modules in chordal graphs: unlike in the moplex setting, properties of chordal graphs with a bounded number of simplicial modules are well understood. For instance, chordal graphs having at most two simplicial modules are interval graphs.

Definition 6.1.1 (Moplex and simplicial moplex)**.** A *moplex* in a graph G is an inclusion-maximal clique module $X \subseteq V(G)$ such that $N_G(X)$ is either empty or a minimal separator in G. A moplex X is *simplicial* if $N_G(X)$ is a clique and a vertex is *moplicial* if it belongs to a moplex. ⊣

We use the definition from Meister [Mei05], which differs from the one by Berry and Bordat [BB98] by considering the vertex set of any complete graph to be a moplex (see also [BP11, XLL13]). This difference is purely technical and allows for less special cases in the proofs. An illustration of the notions of moplexes and avoidable vertices that are not moplicial is provided in Figure 6.1.

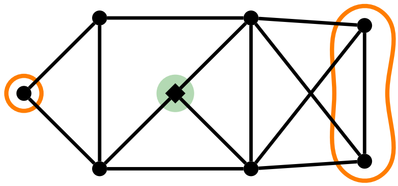

Figure 6.1: A graph with exactly two moplexes (circled in orange) containing an avoidable vertex (the diamond marked in green) that is neither moplicial nor simplicial.

Dirac proved, Theorem 2.1.1, that every non-complete chordal graph contains at least two simplicial vertices. This result was generalised by Berry and Bordat to moplexes as follows.

Theorem 6.1.2 (Berry and Bordat [BB01a, Mei05])**.** Every non-complete graphs contain at least two moplexes.

Moplexes have a strong algorithmic connection to lexicographic searches which we analyse more carefully in Subsection 6.1.4. Moreover, there are various algorithms that make use of moplexes, e.g. for computing a minimal completion to an interval graph [RST08], for computing minimal triangulations of claw-free graphs [BW12], and for recognising diamond-free graphs without induced cycles of length at least five [BBGM15].

Despite these fundamental connections and useful applications, we know little about the connection between structural properties of graphs and their moplexes. This section approaches this problem by investigating the effect of using moplexes as a complexity measure on the class of general graphs. Do graphs with a bounded number of moplexes have useful structural or algorithmic properties? And if chordal graphs with at most two simplicial modules form a natural subclass of the fundamental class of interval graphs, what can we say about graphs with at most two moplexes?

For a positive integer k, a *k-moplex graph* is a graph that contains at most k moplexes. Moreover, the *moplex number* of a graph is the number of moplexes it contains. The first result in this section provides a link between the moplex number and the *asteroidal number* (see Subsection 2.1.1) of a graph [KKM01, COS97], generalising an earlier result of Berry and Bordat for graphs with at most two moplexes [BB01a].

Theorem 6.1.3. The asteroidal number of a graph is a lower bound on its moplex number.

Theorem 6.1.3 immediately implies that the graphs with a bounded number of moplexes inherit the nice algorithmic properties of graphs with bounded asteroidal number. This includes polynomial-time algorithms for various algorithmic problems [FKM04, BKKM99, KM12, KMT08], existence of a spanning tree approximating vertex distances up to a constant additive term [KKM01], a constant factor approximation algorithm for treewidth [BKMT01], and an upper bound on the treewidth in terms of the maximum degree [BT97]. We remark that though computing the asteroidal number of a graph is NP-hard [KKM97], the moplex number of a graph is polynomial-time computable [BB01a].

A graph class is *hereditary* if it is closed under vertex deletion. The class of 1-moplex graphs is hereditary, but not of particular interest, as it is precisely the class of complete graphs. Unfortunately, as one can verify using the family of paths, the class of k-moplex graphs is not hereditary for any $k \geq 2$. The graph in Figure 6.1 shows that even the connected components of a graph obtained by deleting a vertex from a k-moplex graph do not need to be k-moplex graphs. This makes understanding the structure of k-moplex graphs significantly more challenging.

Even the structure of 2-moplex graphs is not yet fully understood. Berry and Bordat [BB01a] showed that 2-moplex graphs are AT-free and that all connected induced subgraphs of a graph G are 2-moplex if and only if G is a proper interval graph. This already gives some clue to how this class of graphs relates to the known hereditary graph classes. We strengthen the former result and complement the latter by proving further results relating the class of 2-moplex graphs to the hierarchy of hereditary graph classes. More precisely, any graph class $\mathcal{G}$ can be naturally mapped to the following two hereditary graph classes, one contained in $\mathcal{G}$ and one containing $\mathcal{G}$:

(i) the class of graphs all of whose induced subgraphs belong to $\mathcal{G}$, or, equivalently, the largest hereditary graph class contained in $\mathcal{G}$, and

(ii) the class of all induced subgraphs of graphs in $\mathcal{G}$, or, equivalently, the smallest hereditary graph class containing $\mathcal{G}$.

Furthermore, if the graph class $\mathcal{G}$ is not closed under disjoint union (as is the case for the class of 2-moplex graphs), it is also natural to consider the class $\mathcal{G}_c$ of all graphs G such that every connected component of G belongs to $\mathcal{G}$ and the corresponding two hereditary graph classes, one contained in $\mathcal{G}_c$ and one containing $\mathcal{G}_c$:

(iii) the class of graphs all of whose induced subgraphs belong to $\mathcal{G}_c$, and

(iv) the class of all induced subgraphs of graphs in $\mathcal{G}_c$.

As previously mentioned, when $\mathcal{G}$ is the class of 2-moplex graphs, a result of Berry and Bordat shows that the corresponding class from (iii) is the class of proper interval graphs. We determine the remaining three hereditary classes related to 2-moplex graphs. We show that the corresponding class from (i) is the class of cochain graphs by proving the following theorem.

Theorem 6.1.4. Every induced subgraph of a graph G is a 2-moplex graph if and only if G is a cochain graph.

Additionally, we establish that the classes from (ii) and (iv) both coincide with the class of cocomparability graphs. To this end we prove the following theorem.

Theorem 6.1.5. The smallest hereditary graph class containing the class of 2-moplex graphs is the class of cocomparability graphs.

After establishing that all 2-moplex graphs are cocomparability graphs and the class contains the proper interval graphs, it is natural to ask for the complexity of problems that are intractable for general cocomparability graphs but yet unknown or polynomial time solvable for the class of proper interval graphs. Along this line we develop reductions showing that two prominent examples of such problems, namely MAX-CUT and GRAPH ISOMORPHISM, both remain as hard on 2-moplex graphs as they are on cocomparability graphs. For proper interval graphs, the complexity of MAX-CUT is still open, and GRAPH ISOMORPHISM is solvable in linear time [LB79].

Theorem 6.1.6. MAX-CUT is NP-complete on cobipartite 2-moplex graphs.

Theorem 6.1.7. Graph Isomorphism is GI-complete on cobipartite 2-moplex graphs.

Theorems 6.1.6 and 6.1.7 provide some indication that the class of 2-moplex graphs is a significant generalisation of the class of connected proper interval graphs. Nevertheless, as the final result on moplexes we show that 2-moplex graphs share the well-known structural property of proper interval graphs that connectedness is a sufficient condition for the existence of a Hamiltonian path [Ber83].

Theorem 6.1.8. Every connected 2-moplex graph has a Hamiltonian path.

The proof of this theorem is an interplay between properties of the class of cocomparability graphs, the Lexicographic Depth First Search algorithm, and the concept of avoidable vertices (also known as OCF-vertices) [BCG+19, BDHT20, OCF76].

6.1.1 Bounding the moplex number

In this subsection we establish a few structural properties for general k-moplex graphs for a fixed $k \in \mathbb{N}$. We begin by recalling a result on minimal separators and moplexes.

Theorem 6.1.9 (Berry and Bordat [BB01a])**.** For every minimal separator S in a graph G, each component of $G - S$ contains a moplex in G.

This result strengthens an earlier result by the same authors, which used the additional assumption that the graph G is chordal.

Theorem 6.1.10 (Berry and Bordat [BB98])**.** Let S be a minimal separator in a chordal graph H. Then, each connected component of $H - S$ contains at least one moplex in H.

As Theorem 6.1.9 is stated in [BB01a] without a proof, we give a short proof here for the sake of completeness. In order to do so, we make use of two more results from the literature. We call a chordal graph G' a *minimal triangulation* of a graph G if $V(G) = V(G')$, $E(G) \subseteq E(G')$, and for all $F \subsetneq E(G') \setminus E(G)$, the graph $(V(G), E(G) \setminus F)$ is not chordal.

Lemma 6.1.11 (Berry and Bordat [BB98])**.** Let H be a minimal triangulation of a graph G and U be a moplex of H. Then, U is a moplex of G.

The following theorem is an immediate consequence of [PS97, Theorem 4.6].

Theorem 6.1.12. Let S be a minimal separator in a graph G. Then, there exists a minimal triangulation H of G such that the vertex sets of the connected components of $H - S$ are the same as the vertex sets of the connected components of $G - S$.

Using these results the proof of Theorem 6.1.9 is easily obtained.

Proof of Theorem 6.1.9. Let S be a minimal separator in a graph G. Using Theorem 6.1.12, we obtain that there exists a minimal triangulation H of G such that for each connected component C_i in $G - S$ there is a connected component C_i' in H with $V(C_i) = V(C_i')$. This implies that S is a minimal separator in H. Then, using Theorem 6.1.10, we know there is a moplex of H in each connected component of $H - S$ and by Lemma 6.1.11 each of these moplexes is a moplex in G as well. □

Using Theorem 6.1.9, we obtain some preliminary results on graphs with at most two moplexes which become useful in Subsection 6.1.5.

Lemma 6.1.13. Let G be a non-complete 2-moplex graph and denote by U and W its two moplexes. Then, the following two properties hold:

(Prop1) U and W are disjoint simplicial moplexes, and

(Prop2) for every minimal separator S in G, the graph $G - S$ contains exactly two connected components, one of which contains U and the other one W.

Proof. First, we show that (Prop1) holds. The fact that U and W are disjoint follows directly from the definition of a moplex, and in particular because every moplex is a maximal set of vertices with the same closed neighbourhood. Suppose without loss of generality that U is not a simplicial moplex, that is, there exist two non-adjacent vertices $a, b \in N(U)$. Let S be a minimal a,b-separator. Then, S contains U, as every vertex in U is adjacent to both a and b. However, following Theorem 6.1.9, the graph $G - S$ contains at least two moplexes in G, say M_1 and M_2. But then M_1, M_2, and U are three distinct moplexes in G, a contradiction to G being a 2-moplex graph.

Second, we show (Prop2). Let S be a minimal separator in G. If there exists a connected component of $G - S$ that does not contain U or W, then Theorem 6.1.9 implies that G contains at least three moplexes, a contradiction. Thus, $G - S$ must contain exactly two connected components, one of which contains U and the other one W. □

While there is a lot of work on how to utilise moplexes algorithmically [BB01b, ST09], little is known about the structural properties of graphs with few moplexes.

As we have seen in Subsection 2.1.1, chordal graphs with a bounded number of simplicial modules have been studied. Thus the natural question arises what the structure of graphs with bounded moplex number might look like.

For $k = 1$, the answer is not that interesting, as 1-moplex graphs are simply the class of all complete graphs. For $k = 2$ it becomes more complicated to answer the question. By Theorem 6.1.9, there is only one way for a 2-moplex graph to not be connected, which is that it is the disjoint union of two cliques. Therefore, we mainly concentrate on the class of connected 2-moplex graphs.

First, we make some observations for general k.

The concept of asteroidal sets can be generalised to moplexes. An *asteroidal set of moplexes* [BB01a] in a graph G is a set $\{X_1, \dots, X_k\}$ of pairwise disjoint moplexes in G such that for each $i \in \{1, \dots, k\}$, all moplexes X_j, $j \neq i$, lie in the same connected component of the graph $G - N(X_i)$.

Berry and Bordat [BB01a] proved that a graph has an asteroidal triple of vertices if and only if it has an asteroidal triple of moplexes. This corresponds to the case $k = 3$ of the following more general statement.

Theorem 6.1.14. A graph has an asteroidal set of vertices of size k if and only if it has an asteroidal set of moplexes of size k.

Proof. Let G be a graph and let k be a positive integer. If G has an asteroidal set of moplexes of size k, then G has an asteroidal set of vertices of size k.

Let $A = \{a_1, \dots, a_k\}$ be an asteroidal set of size k in G such that A contains as many moplicial vertices as possible. First, note that because A is an independent set and every moplex is a clique, no two vertices in A belong to the same moplex. Thus, to complete the proof it suffices to show that every vertex $a_i \in A$ is moplicial. Indeed, denoting by M_i the moplex of G containing a_i, for all $i \in \{1, \dots, k\}$, we would obtain that $\{M_1, \dots, M_k\}$ is an asteroidal set of moplexes of size k.

Suppose that some $a_i \in A$ is not part of a moplex. Let C be the component of $G - N[a_i]$ containing $A \setminus \{a_i\}$ and let D be the component of $G - N[V(C)]$ containing a_i. Let $S = N(V(C))$ and observe that $S \subseteq N(a_i)$. This implies that both C and D are S-full components of $G - S$, and thus that S is a minimal a_i,a_j-separator for any $a_j \in A \setminus \{a_i\}$. Hence, we can use Theorem 6.1.9 and obtain that D contains a moplex M_D in G.

Let x be a vertex in M_D and $A' = (A \setminus \{a_i\}) \cup \{x\}$. We show next that A' is an asteroidal set in G. Fix two vertices $v, w \in A' \setminus \{x\}$. Note that $N_G[x] \subseteq D \cup S$, and because both v and w belong to C, there exists a v-w-path that does not contain any vertex from $N_G[x]$. Furthermore, observe that there exists a path P between a_i and x in D, which does not contain any vertex in $N[w]$. Also, because A is an asteroidal set, there exists a path P' between a_i and v which does not contain any vertex in $N[w]$. Hence, the subgraph $P \cup P'$ contains a walk between x and v in $G - N[w]$. Therefore, A' is an asteroidal set in G of size k. However, the number of vertices in A' that belong to a moplex is strictly larger than the number of vertices in A that belong to a moplex, contradicting the choice of A. This shows that every vertex in A is moplicial. $\square$

That the moplex number of a graph is at least its asteroidal number is an immediate consequence of Theorem 6.1.14.

Theorem 6.1.3. The asteroidal number of a graph is a lower bound on its moplex number.

On the other hand, the gap between the moplex and asteroidal number can be arbitrarily large. This is easily seen if we, for example, consider the class of stars. An asteroidal set of a star has size 2, as we can only contain two of the leaves. Removing the closed neighbourhood of the leaf disconnects the graph into components of size 1, thus no other two vertices lie in the same component. But the moplex number of stars is equal to the number of leaves it has, as every leaf is its own moplex.

Theorem 6.1.3 implies that the asteroidal number is computable in polynomial time in every class of graphs with bounded moplex number as one can verify in polynomial time whether a given set of vertices is an asteroidal set. Combining this with the results by Fomin, Kratsch and Müller [FKM04], Theorem 6.1.3 also implies that Dominating Set and Total Dominating Set can be solved in polynomial time in classes of graphs of bounded moplex number. The same holds for Independent Set, Independent Dominating Set, and Efficient Dominating Set, along with their weighted variants [BKKM99], for k-Colouring (for any fixed k [KM12]) and for Weighted Feedback Vertex Set [KMT08]. Furthermore, for graphs of bounded moplex number, bounded degree implies bounded treewidth, as a consequence of the fact that graphs with asteroidal number at most k have chordality at most $2k+1$ and of a result by Bodlaender and Thilikos [BT97].

For later use, we explicitly state the previously mentioned result of Berry and Bordat on 2-moplex graphs, which is now an immediate consequence of Theorem 6.1.3.

Corollary 6.1.15 (Berry and Bordat [BB01a])**.** Every 2-moplex graph is AT-free.

6.1.2 Avoidable vertices in moplexes

Another way to generalise simplicial vertices to general graphs is the concept of *avoidable vertices*. Avoidable vertices go back to the work of Ohtsuki, Cheung, and Fujisawa [OCF76], who proved that every graph has an avoidable vertex, that is, avoidable vertices also yield a generalisation of Theorem 2.1.1. Later, these vertices were called *OCF-vertices* [BBHP04, Heg06, BBH+06, YHF14, BBBS10], and very recently the notion has reappeared under the name of avoidable vertices [BCG+19, BDHT20].

Definition 6.1.16 (Avoidable vertices)**.** Let v be a vertex in a graph G. An *extension* of v in G is an induced P_3 in G having v as midpoint. Such an extension is *failing* if it is not contained in an induced cycle. A vertex v is *avoidable* in G if none of its extensions is failing. ⊣

Note that every simplicial vertex is avoidable, and a vertex in a chordal graph is avoidable if and only if it is simplicial.

Theorem 6.1.17 (Ohtsuki, Cheung, and Fujisawa [OCF76])**.** Every graph contains an avoidable vertex.

For simplicial vertices in chordal graphs we know there are always at least two. Similarly, one can infer from Theorem 6.1.17 the following statement.

Corollary 6.1.18 (Ohtsuki, Cheung, and Fujisawa [OCF76])**.** Every graph on at least two vertices contains at least two avoidable vertices.

Equivalently, every non-complete graph contains at least two *avoidable modules*, that is, maximal clique modules containing an avoidable vertex.

The two concepts of moplexes and avoidable vertices are closely related. Note that if X is a moplex in a graph G, then the graph $G - N_G[X]$ contains an $N_G(X)$-full component. Using such a component, it can be shown that every extension of a vertex $v \in X$ is contained in an induced cycle in G. This leads to the following observation (as noted already in [BBB+05], and perhaps earlier).

Observation 6.1.19 (Berry et. al. [BBB+05])**.** Every moplicial vertex in a graph is avoidable.

We use this observation frequently throughout the proofs of this section without explicit reference. It also makes the following an immediate corollary of Theorem 6.1.3.

Corollary 6.1.20. The asteroidal number of a graph G is a lower bound on the number of avoidable vertices in G.

The converse however is not true, as Figure 6.2 shows not every avoidable vertex is moplicial. We call vertices that are avoidable but not moplicial *purely avoidable*.

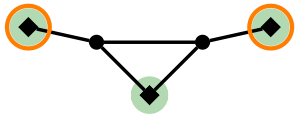

Figure 6.2: The bull graph has exactly two moplexes (circled in orange) but three avoidable vertices (the diamonds marked in green).

Since the class of k-moplex graphs is not hereditary for any $k \geq 2$, it is a natural question whether for all k, every k-moplex graph contains a vertex whose removal results in a k-moplex graph. Unfortunately, this is not the case. It is not difficult to show that for every $k \geq 2$ there exists a k-moplex graph G_k such that deleting any vertex results in a graph that is not a k-moplex graph. See Figure 6.3 for examples of such graphs for $k \in \{2, 3, 4\}$; the construction can be easily generalised to larger values of k.

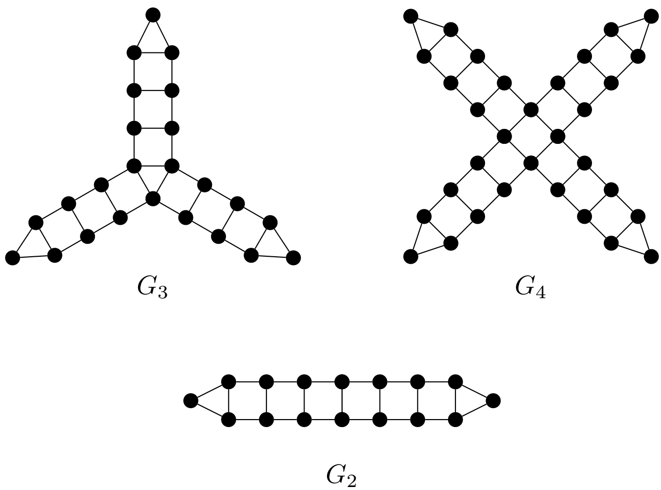

Figure 6.3: Examples of graphs in which deleting any vertex increases the moplex number.

Nevertheless, we show that there are still certain vertices that, if existent, can be removed without leaving the class of k-moplex graphs. To this end, we first prove the following stronger statement.

Theorem 6.1.21. Let G be a graph, $v \in V(G)$ a purely avoidable vertex of G, and $M \subseteq V(G)$. Then, M is a moplex in G if and only if it is a moplex in $G - v$.

Proof. Observe that if there is a vertex $v' \in V(G) \setminus \{v\}$ such that $N_G[v'] = N_G[v]$, then there is a natural correspondence between the clique modules of G and those of $G - v$, which preserves, in both directions, the property of being a moplex. Thus, from now on, we assume that for any vertex $v' \neq v$ in G, we have $N_G[v'] \neq N_G[v]$.

Let M be a moplex in G. Suppose that M is not a moplex in $G - v$, and let $S = N_G(M)$. By assumption, v is not moplicial in G, and thus $v \notin M$. As M is a clique and a module in $G - v$, the reason for M no being a moplex in $G - v$ is that there is no S-full component of $G - (S \cup \{v\})$ other than M itself. Because such a component exists in $G - S$, call it H, we infer that v belongs to this component and that there is a vertex $x \in S$ such that v is the only neighbour of x in H, that is, $N_G(x) \cap V(H) = \{v\}$. Let H' be the connected component of $G - N_G[v]$ containing M. We claim that H' is $N_G(v)$-full. We know that $N_G[v] \subseteq S \cup V(H)$. Since M lies in H', every vertex in $N_G(v) \cap S$ has a neighbour in H'. Consider now a vertex $w \in N_G(v) \setminus S$. Then, $w \in V(H)$ and, because v is the only vertex in H that is adjacent to x, we infer that w is not adjacent to x. Thus, we obtain an extension wvx and, because v is avoidable in G, there is a w-x-path P in $G - (N_G[v] \setminus \{w, x\})$. Let y be the first vertex of P not in H. Then, y belongs to S and consequently to H'. Furthermore, because v is the only neighbour of x in H, we have $y \neq x$. Hence, the path P' obtained from P by removing x and w lies in H', and so w has a neighbour in H'. We obtain that every neighbour of v has a neighbour in H', and thus that H' is $N_G(v)$-full. It follows that $\{v\}$ is a moplex in G, a contradiction to v being purely avoidable. Thus, M remains a moplex in $G - v$.

For the backward direction let M be a moplex in G_v. Suppose towards a contradiction that M is not a moplex in G. Let S be the neighbourhood of M in $G - v$. Since S is a minimal separator in $G - v$, there exists an S-full component H in $(G - v) - N_G[M]$. Assume first that v has no neighbours in M. Then, $N_G(M) = S$, and hence the component of $G - S$ containing the vertices of H contains no vertices of M and is S-full. This implies that M is a moplex in G, a contradiction. Hence, there exists a vertex $u \in N_G(v) \cap M$.

We claim that $M \subseteq N_G(v)$. Towards a contradiction, suppose there exists a vertex $u' \in M \setminus N_G(v)$. Let H' be the component of $G - N_G[v]$ containing u'. We show that H' is $N_G(v)$-full in G. Since $u' \in M$, all the vertices in $(M \cup S) \setminus \{u'\}$ are adjacent to u'. Thus, as u' is in H', every vertex in $N_G(v) \cap (M \cup S)$ has a neighbour in H'. So, let $z \in N_G(v) \setminus (M \cup S)$. Note that $u \in M$ and $z \notin S \cup \{v\} = N_G(M)$, and hence the vertices u and z are non-adjacent in G, which implies that uvz is an extension of v. Since v is avoidable, there is a u-z-path in $G - (N_G[v] \setminus \{u, z\})$. Any such path intersects $S \setminus N_G(v)$, and thus all its internal vertices are in the same component of $G - N_G[v]$ as u', namely H'. In particular, this is the case for the neighbour of z on the path. This shows that every vertex in $N_G(v) \setminus (M \cup S)$ has a neighbour in H'. We conclude that H' is $N_G(v)$-full in G, as claimed. Thus, $\{v\}$ is a moplex in G, a contradiction.

From the previous observation that $M \subseteq N_G(v)$, we readily get that v does not have any neighbours in H. Indeed, if v has a neighbour in H, then M would be a clique module in G with neighbourhood $S \cup \{v\}$ and H would be an $(S \cup \{v\})$-full component, and hence M would be a moplex in G, a contradiction.

We show next that $N_G(v) \setminus (M \cup S)$ is non-empty. Suppose that $N_G(v) \subseteq M \cup S$. Recall that $N_G[v] \neq N_G[u]$. Thus, because $N_G[u] = M \cup S \cup \{v\}$, there exists a vertex $x \in N_G[u] \setminus N_G[v]$. Then, $x \in M \cup S$; however, as $M \subseteq N_G(v)$, we must have $x \in S$. Following the fact that $S = N_G(V(H))$ and $M \subseteq N_G(x)$, we obtain that $(S \cup M) \setminus \{x\} \subseteq N_G(V(H) \cup \{x\})$, and hence $N_G(v) \subseteq N_G(V(H) \cup \{x\})$. As previously shown, v has no neighbours in H. Hence, the vertices in $V(H) \cup \{x\}$ all belong to the same connected component of $G - N_G(v)$. However, this implies that $N_G(v)$ is a minimal separator in G, and thus $\{v\}$ is a moplex in G, a contradiction. Thus, $N_G(v) \setminus (M \cup S)$ is non-empty.

Let $z \in N_G(v) \setminus (M \cup S)$. Then, u and z are non-adjacent in G and uvz is an extension of v. Since v is avoidable, there exists an induced u-z-path in G having all internal vertices in $G - N_G[v]$. In particular, this path must contain a vertex in $S \setminus N_G(v)$, and thus $S \setminus N_G(v) \neq \emptyset$. Following the fact that $N_G(v) \cap V(H) = \emptyset$, there exists a component H' of $G - N_G[v]$ that contains all vertices of H. We show that H' is an $N_G(v)$-full component in G. As every vertex of S has a neighbour in H, we infer that every vertex in $S \setminus N_G(v)$ belongs to H', and every vertex in $N_G(v) \cap S$ has a neighbour in H'. Furthermore, because $\emptyset \neq S \setminus N_G(v) \subseteq V(H')$, every vertex of $N_G(v) \cap M = M$ has a neighbour in H'. Finally, as v is avoidable, for every vertex $y \in N_G(v) \setminus (M \cup S)$ there is a u-y-path in $G - (N_G[v] \setminus \{u, y\})$. Any such path intersects $S \setminus N_G(v)$, and thus all its internal vertices are contained in H'. Hence, $\{v\}$ is a moplex in G, a contradiction. $\square$

We can now establish the announced claim, which is later applied in Subsection 6.1.7.

Corollary 6.1.22. For every positive integer k, the class of k-moplex graphs is closed under deletion of an avoidable vertex that is not moplicial.

Proof. Let G be a k-moplex graph, and let $v \in V(G)$ be an avoidable vertex that is not moplicial. By Theorem 6.1.21, every moplex in $G - v$ is a moplex in G. Thus, because G has at most k moplexes, so does $G - v$. □

6.1.3 Connections with proper interval graphs and cochain graphs

Berry and Bordat characterised graphs in which every connected induced subgraph has at most two moplexes by the well-known class of proper interval graphs. We complement this result by characterising the graphs in which every induced subgraph has most two moplexes by the class of cochain graphs, that is, proving Theorem 6.1.4.

Theorem 6.1.23 (Berry and Bordat [BB01a])**.** Let G be a graph. Then, each connected induced subgraph of G has at most two moplexes if and only if G is a proper interval graph.

Since Theorem 6.1.23 was stated in [BB01a] without a detailed proof, we summarise here the main ideas leading to this result. Roberts proved that proper interval graphs are exactly the claw-free interval graphs [Rob69] (see also [Gar07, BW99]). This result, together with the characterisation of interval graphs due to Lekkerkerker and Boland stating that interval graphs are exactly the AT-free chordal graphs [LB62], implies the following.

Corollary 6.1.24. A graph G is a proper interval graph if and only if G is a claw-free AT-free chordal graph.

Corollary 6.1.24 implies the forward implication in the equivalence given by Theorem 6.1.23. The backward implication states that every connected proper interval graph has at most two moplexes, which can also be derived using known results from the literature:

- a result of Roberts [Rob69] stating that every connected proper interval graph in which no two distinct vertices have the same closed neighbourhoods has at most two extreme vertices, where an extreme vertex is a simplicial vertex s such that every pair of neighbours of s have a common neighbour outside $N[s]$, and
- the fact that for every minimal separator S in a chordal graph G, every S-full component of $G - S$ has a vertex dominating S (see, e.g. [KM98]).

For completeness, we also offer a short proof of Theorem 6.1.23.

Proof of Theorem 6.1.23. Suppose first that each connected induced subgraph of G has at most two moplexes. Note that for every $k \geq 4$, the cycle C_k is a connected graph in which every vertex forms a moplex. Furthermore, the claw is a connected graph in which every vertex of degree one forms a moplex. Thus, G must be claw-free and chordal. By Corollary 6.1.15 this implies that G is a claw-free AT-free chordal graph. Using Corollary 6.1.24 we conclude that G is a proper interval graph.

For the converse direction, let G be a proper interval graph and let H be a connected induced subgraph of G. Then, H is a connected proper interval graph. We need to show that H has at most two moplexes. Since both the connectedness and the moplex number are preserved upon deleting a vertex from a pair of vertices with the same closed neighbourhoods, we may assume that no two vertices in H have the same closed neighbourhood. Under this assumption, every moplex in H has size one. To complete the proof, fix a proper interval model of H and let $I_1, \ldots, I_n$ be the ordering of the intervals according to their left endpoints. For $j \in \{1, \ldots, n\}$, let v_j be the vertex represented by I_j. It now suffices to show that for all $j \in \{2, \ldots, n-1\}$, the set $\{v_j\}$ is not a moplex in H. Suppose towards a contradiction that $\{v_j\}$ is a moplex. Then, the graph $H - N[v_j]$ contains an $N(v_j)$-full component C. Since C contains no vertex from the closed neighbourhood of v_j, no interval representing a vertex in C intersects I_j. Furthermore, because C is a connected graph, we may assume that all its vertices are represented by intervals whose left endpoints are strictly larger than the right endpoint of I_j. The connectedness of H and the ordering of the intervals imply that I_{j-1} intersects I_j, that is, v_{j-1} is adjacent to v_j. However, because all intervals representing vertices in C are disjoint from I_j and lie entirely to the right of I_j, the fact that I_{j-1} ends before I_j ends implies that no vertex in C can be adjacent to $v_{j-1} \in N(v_j)$. This contradicts the assumption that C is an $N(v_j)$-full component of $H - N[v_j]$. □

As mentioned before, there are four nested hereditary graph classes naturally associated with the graph class $\mathscr{M}$ of 2-moplex graphs:

class $\mathscr{M}^-$**:** the largest hereditary graph class contained in $\mathscr{M}$, that is, the class of graphs all of whose induced subgraphs are 2-moplex graphs,

class $\mathscr{M}^+$**:** the smallest hereditary graph class containing $\mathscr{M}$, that is, the class of all induced subgraphs of 2-moplex graphs,

class $\mathscr{M}_c^-$**:** the largest hereditary graph class contained in $\mathscr{M}_c$, that is, the class of graphs all of whose induced subgraphs only have 2-moplex graphs as connected components, and

class $\mathscr{M}_c^+$**:** the smallest hereditary graph class containing $\mathscr{M}_c$, that is, the class of all induced subgraphs of graphs in $\mathscr{M}_c$.

The following inclusion relations follow directly from the definitions:

$$\mathscr{M}^- \subseteq \mathscr{M} \subseteq \mathscr{M}^+, \qquad \mathscr{M}_c^- \subseteq \mathscr{M}_c \subseteq \mathscr{M}_c^+,$$
$$\mathscr{M}^- \subseteq \mathscr{M}_c^-, \quad \mathscr{M} \subseteq \mathscr{M}_c, \quad \text{and} \quad \mathscr{M}^+ \subseteq \mathscr{M}_c^+.$$

Theorem 6.1.23 implies that $\mathscr{M}_c^-$ is the class of proper interval graphs. We determine the remaining three hereditary classes, $\mathscr{M}^-$, $\mathscr{M}_c^+$, and $\mathscr{M}^+$. Theorem 6.1.4 shows that $\mathscr{M}^-$ is the class of cochain graphs, while the results in Subsection 6.1.5 show that $\mathscr{M}_c^+$ and $\mathscr{M}^+$ both coincide with the class of cocomparability graphs, yielding

$$\underbrace{\mathscr{M}^-}_{\text{cochain}} \subset \overbrace{\mathscr{M}}^{\text{2-moplex}} \subset \underbrace{\mathscr{M}^+}_{\text{cocomparability}} \quad \text{and} \quad \underbrace{\mathscr{M}_c^-}_{\text{proper interval}} \subset \overbrace{\mathscr{M}_c}^{\text{each component is 2-moplex}} \subset \underbrace{\mathscr{M}_c^+}_{\text{cocomparability}}. \tag{6.1}$$

Let us also comment on the remaining three inclusions $\mathscr{M}^- \subseteq \mathscr{M}_c^-$, $\mathscr{M} \subseteq \mathscr{M}_c$, and $\mathscr{M}^+ \subseteq \mathscr{M}_c^+$. Since $3K_1$ is a proper interval graph but not a 2-moplex graph, the inclusions $\mathscr{M}^- \subseteq \mathscr{M}_c^-$ and $\mathscr{M} \subseteq \mathscr{M}_c$ are both proper. On the other hand, the inclusion $\mathscr{M}^+ \subseteq \mathscr{M}_c^+$ holds with equality.

Theorem 6.1.4. Every induced subgraph of a graph G is a 2-moplex graph if and only if G is a cochain graph.

Proof. Suppose first that each induced subgraph of G has at most two moplexes. Since for every $k \geq 4$ the cycle C_k is a connected graph in which every vertex forms a moplex, G is a chordal graph. Furthermore, because the graph $3K_1$ has three moplexes, G has independence number at most two. As G is a chordal graph, it is also perfect, and hence the vertex set of G can be covered with two disjoint cliques X and Y (see [Lov72]), that is, G is cobipartite. Furthermore, for any two vertices x_1, x_2 in X we must have $N(x_1) \cap Y \subseteq N(x_2) \cap Y$ or $N(x_2) \cap Y \subseteq N(x_1) \cap Y$, since otherwise G would contain an induced 4-cycle. Using the fact that X is a clique, we infer that $N[x_1] \subseteq N[x_2]$ or $N[x_2] \subseteq N[x_1]$; since this holds for any two vertices in X, we conclude that G is a cochain graph.

For the converse direction it suffices to show that every cochain graph has at most two moplexes, because the class of cochain graphs is hereditary, that is, every induced subgraph of a cochain graph is again a cochain graph. Let G be a cochain graph. If G is disconnected, then G is isomorphic to the disjoint union of two complete graphs, and hence has moplex number two. So we may assume that G is connected. By definition of cochain graphs, G is a chordal graph. Furthermore, because G is a cobipartite graph, G is also claw-free and AT-free. Thus, G is a proper interval graph by Corollary 6.1.24. Finally, by Theorem 6.1.23, G has at most two moplexes. □

6.1.4 Lexicographic searches and vertex orderings

Vertex orderings are a central tool when working with graph classes.

A *(vertex) ordering* of a graph G is a total order on its vertex set. Given an ordering σ of a graph G and two distinct vertices x and y, we write $x <_\sigma y$ if x precedes y in σ. If $x <_\sigma y$ and there is no vertex $z \in V(G) \setminus \{x, y\}$ such that $x <_\sigma z$ and $z <_\sigma y$, then we say y is the *direct successor* of x. A *bump* of σ is a pair of vertices x and y such that y is the direct succesor of x in σ but $xy \notin E(G)$. Three vertices $x <_\sigma y <_\sigma z$ of G with $xy \notin E(G)$, $yz \notin E(G)$, and $xz \in E(G)$ are said to form an *umbrella* in σ. An ordering is *umbrella-free* if no three vertices form an umbrella in it. As shown by Kratsch and Stewart [KS93], the existence of an umbrella-free ordering characterises cocomparability graphs.

Theorem 6.1.25 (Kratsch and Stewart [KS93])**.** A graph G is a cocomparability graph if and only if it has an umbrella-free ordering.

Especially on graph classes that are characterisable via vertex orderings, these can be very helpful to prove further structural properties.

Theorem 6.1.26 (Deogun and Steiner [DS94])**.** It can be decided in polynomial time whether a cocomparability graph has a Hamiltonian path.

A concept very closely related to vertex orderings are graph search algorithms as every graph search algorithm yields a vertex ordering of the input graph. A *DFS ordering* of a connected graph G is any ordering of $V(G)$ obtained by a depth first search. We are particularly interested in one particular graph search algorithm, the lexicographic depth first search (LDFS). Informally speaking, LDFS can be seen as a special version of Depth First Search (DFS) with a tie-breaking rule favouring vertices with recently visited neighbours. These are determined using labels that are assigned to the vertices during runtime and the lexicographic relation between them.

When introduced in 2008 by Corneil and Krueger [CK08] the LDFS algorithm did not have any application. It was simply introduced as the natural analogue to LBFS, which is used for recognising chordal graphs. Only after that its strong relation to Hamiltonian properties was discovered [CDH13a]. Both LDFS and LBFS are mostly studied on graphs closely related to chordal graphs.

Of much greater interest to us is the characterisation of LDFS via a constraint on the ordering it produces (see [CK08, Theorem 2.7]). Thus we use this characterisation as definition of an LDFS.

Algorithm 1 Lexicographic Depth First Search.
Input: graph $G = (V, E)$, start vertex $s \in V(G)$
Out: an ordering σ of V

$U \leftarrow V$
for all $v \in V(G)$ **do**
    $label(v) \leftarrow \epsilon$
$label(s) \leftarrow (0)$
for all $1 \leq i \leq n$ **do**
    choose vertex v from U with lexicographic largest label
    $\sigma(i) \leftarrow v$
    $U \leftarrow U \setminus \{v\}$
    for all $u \in U$ **do**
        if $uv \in E$ **then**
            $label(u) \leftarrow i \cdot label(v)$ ▷ considering i to be a letter

Definition 6.1.27 (Corneil and Krueger [CK08, Theorem 2.7])**.** An ordering σ of G is an *LDFS ordering* if the following holds: if $a <_\sigma b <_\sigma c$ and $ac \in E(G)$ and $ab \notin E(G)$, then there exists a vertex d such that $a <_\sigma d <_\sigma b$ and $db \in E$ and $dc \notin E$. See Figure 6.4a for an illustration. ⊣

(a) LDFS orderings

(b) LDFS umbrella-free ordering

Figure 6.4: If an ordering is both LDFS and umbrella free, then the umbrella-freeness implies the existence of the orange edges, thus implying Lemma 6.1.29.

Note that every LDFS ordering is a DFS ordering. We recall some well-known properties of all DFS orderings. Essentially, considering the vertices ordered from left to right, if the graph is connected, every vertex except the first has a neighbour to its left and the left endpoint of every bump does not have any neighbours to its right.

Observation 6.1.28. Let G be a connected graph and let $\sigma = (v_1, v_2, \ldots, v_n)$ be a DFS ordering of G. Then, for every $i \in \{2, \ldots, n\}$ there exists $j < i$ such that $v_i v_j \in E(G)$. Also, if $v_i v_{i+1}$ is a bump of σ for some $i \in \{1, 2, \ldots, n-1\}$, then $i \geq 2$ and $v_i v_j \notin E(G)$ for every $j > i$. See Figure 6.5 for an illustration.

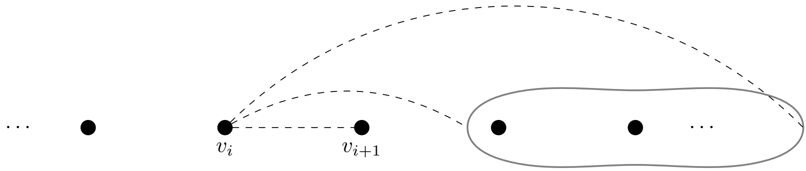

Figure 6.5: The left vertex of a bump has no neighbours to the right.

A result of Köhler and Mouatadid [KM14] states that every cocomparability graph admits an umbrella-free LDFS ordering. In Proposition 6.1.35 we prove that all 2-moplex graphs are cocomparability graphs. Thus we consider the properties of these orderings that are both LDFS and umbrella-free.

Lemma 6.1.29 (Corneil et al. [CDH13b])**.** Let σ be an umbrella-free LDFS ordering of a cocomparability graph G and let a, b, c be vertices such that $a <_\sigma b <_\sigma c$, and $ac \in E(G)$, and $ab \notin E(G)$. Then, there exists a vertex d such that $a <_\sigma d <_\sigma b$, and $\{a, b, c, d\}$ induces a C_4 in G. (See Figure 6.4b for an illustration.)

Next, we look at the local structure around bumps in an umbrella-free LDFS ordering.

Lemma 6.1.30. Let G be a connected cocomparability graph, let $\sigma = (v_1, \ldots, v_n)$ be an umbrella-free LDFS ordering of G, and let $v_i v_{i+1}$ be a bump in σ. Then, v_i is avoidable in G.

Proof. Let v_j and v_k be two non-adjacent neighbours of v_i. By Observation 6.1.28, we know that $v_j <_\sigma v_i$ and $v_k <_\sigma v_i$. We may assume without loss of generality that $v_j <_\sigma v_k <_\sigma v_i$. By Lemma 6.1.29, there exists a vertex v_ℓ such that $\{v_i, v_j, v_k, v_\ell\}$ induces a C_4 in G. Thus v_i is avoidable. □

In applications for graph searches we often find that an algorithm is applied more than once such that the latter applications use the results of the former ones as input. For example this leads to a simple recognition algorithm for unit interval graphs [Cor04a]. This technique is often called a *multisweep* and was introduced by Simon to recognise interval graphs [Sim91].

One specific tie breaking strategy due to Simon [Sim91], sometimes referred to as a "+ sweep", breaks ties in a DFS (or any other vertex ordering) procedure by prioritising the greatest eligible element with respect to some already given ordering. This can be combined with multisweep techniques in a very natural way by using an earlier sweep as the input ordering.

Definition 6.1.31. Let σ be an arbitrary ordering of a graph G. Then, (L)DFS$^+(\sigma)$ is the (L)DFS ordering of G obtained by starting an (L)DFS in the last vertex of σ and breaking ties by choosing the largest vertex with respect to $<_\sigma$, that is, if in line 6 of Algorithm 1 there is more than one eligible vertex, then take the largest with respect to $<_\sigma$. In particular, (L)DFS$^+(\sigma)$ always starts with the last vertex of σ. We provide an example in Figure 6.6.

Note that (L)DFS$^+(\sigma)$ is the (L)DFS ordering of G which is lexicographically maximal with respect to $<_\sigma$ (where the first vertex has highest significance) among all (L)DFS orderings of G. ⊣

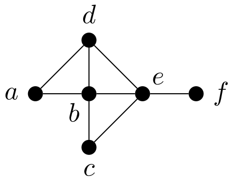

$$\sigma = (d, b, a, e, f, c)$$

$$\text{DFS}^+(\sigma) = (c, e, f, b, a, d)$$

$$\text{LDFS}^+(\sigma') = (c, e, b, d, a, f)$$

Figure 6.6: A graph and a given ordering σ together with the ordering DFS$^+(\sigma)$ and LDFS$^+(\sigma')$.
The first vertex of DFS$^+(\sigma)$ is c, as it is the last of σ. Next, there is the choice between b and e. Since $b <_\sigma e$, we choose e. The largest neighbour of e is f, so that one is chosen next. As f does not have further neighbours we backtrack to e and choose the next largest neighbour b. Finally, because $d <_\sigma a$, we visit a and then d.

We make use of the following three key results on LDFS vertex orderings in the class of cocomparability graphs.

(LDFS1) As mentioned before, Köhler and Mouatadid [KM14] showed that every cocomparability graph admits an LDFS ordering that is umbrella-free.

(LDFS2) Corneil, Dalton and Habib [CDH13b] showed that if σ is an umbrella-free ordering of a graph, then so is LDFS$^+(\sigma)$.

(LDFS3) Xu, Li and Liang [XLL13] showed that every LDFS ordering ends in a moplicial vertex.

These three results imply the following corollary.

Corollary 6.1.32. Every cocomparability graph G has an umbrella-free LDFS ordering σ such that both its first and last vertex are moplicial.

Proof. We first apply the approach of [KM14] to obtain an umbrella-free LDFS ordering σ' of G and then we compute $\sigma = \mathrm{LDFS}^+(\sigma')$. By [CDH13b], σ is an umbrella-free LDFS ordering of G. Since σ starts with the last vertex of σ', it follows from [XLL13] that both first and last vertices of σ are moplicial. $\square$

This implies that we can find the two moplexes of a 2-moplex graph in linear time. There is a lot of research on the connection between moplexes and lexicographic searches, see for example [BB01b, Cor04b, XLL13, Tro16].

6.1.5 2-moplex graphs are cocomparabilty

In this subsection, we show that the smallest hereditary graph class containing the class of 2-moplex graphs is the class of cocomparability graphs. In order to prove that every 2-moplex graph G is a cocomparability graph, we identify a property common to all minimal x,y-separators, for any two non-adjacent vertices x and y of G. We then exploit this property to orient the edges of the complement of G in a transitive way.

Let G be a non-complete 2-moplex graph with the two moplexes $U, W \subseteq V(G)$. For two non-adjacent vertices x and y we denote by $\mathcal{S}_G(x, y)$ (or simply $\mathcal{S}(x, y)$ if the graph is clear from the context) the set of all minimal x,y-separators in G. Given $M \in \{U, W\}$ and $S \in \mathcal{S}(x, y)$, we say that M *prefers* x *to* y *with respect to* S if M and x lie in the same connected component of $G - S$. By (Prop2) of Lemma 6.1.13, either M prefers x to y with respect to S or M prefers y to x with respect to S. As the following key lemma shows, which of these two cases occurs is actually independent of the choice of S.

Lemma 6.1.33. Let U and W be the two moplexes of a non-complete 2-moplex graph G. Then, for each $M \in \{U, W\}$ and for every two non-adjacent vertices x and y exactly one of the following conditions holds:

- M prefers x to y with respect to all $S \in \mathcal{S}(x, y)$, or
- M prefers y to x with respect to all $S \in \mathcal{S}(x, y)$.

Proof. Without loss of generality, assume that $M = U$. Suppose towards a contradiction that there exist two minimal x,y-separators S and S' such that U and x lie in the same component of $G - S$, and U and y lie in the same component of $G - S'$. Then, due to Lemma 6.1.13, W and y lie in the same component of $G - S$, and W and x lie in the same component of $G - S'$. We have $x \notin U$, because x and U lie in different components of $G - S'$. In a similar way we conclude that neither x nor y can belong to $U \cup W$. Now fix any $u \in U$ and $w \in W$. Observe that $\{x, u, w\}$ is an independent set in G. We conclude the proof by deriving a contradiction with Corollary 6.1.15. To this end, it is enough to show that $\{x, u, w\}$ is an asteroidal triple in G.

The removal of $N_G[u]$ does not affect the component of $G - S'$ containing both x and w. Similarly, the removal of $N_G[w]$ does not affect the component of $G - S$ containing both x and u. Finally, consider the graph $G - N_G[x]$ and first observe that it contains a u-y-path, because the removal of $N_G[x]$ does not affect the component of $G - S'$ containing y and u. Similarly, $G - N_G[x]$ contains a y-w-path, because the removal of $N_G[x]$ does not affect the component of $G - S$ containing y and w. Thus, the vertices u and w are in the same component of $G - N_G[x]$ and $\{x, u, w\}$ is an asteroidal triple, as claimed. □

It follows for every non-complete 2-moplex graph G with the moplexes U and W: if for *some* minimal x,y-separator the vertex x belongs to the component containing the moplex U and y belongs to the component containing the moplex W, then this is the case for *every* minimal x,y-separator. Let $M \in \{U, W\}$. If M prefers x to y with respect to all $S \in \mathcal{S}(x, y)$, we say that *M prefers x to y*. Using this we define a binary relation R_M over $V(G)$ as follows:

$$xR_My \quad \text{if and only if} \quad M \text{ prefers } x \text{ to } y.$$

By Lemma 6.1.33 the two relations R_U and R_W are well-defined. Furthermore, either U prefers x to y or U prefers y to x (in which case W prefers x to y). We thus have xR_Uy if and only if yR_Wx, that is, $R_W = R_U^{-1}$. Note also that by definition, xR_Uy or xR_Wy implies that x and y are distinct and non-adjacent. R_U R_W

Lemma 6.1.34. Let U and W be the two moplexes of a non-complete 2-moplex graph G. Then, the relations R_U and R_W are transitive.

Proof. Let x, y, z be vertices such that U prefers x to y as well as y to z, that is, xR_Uy and yR_Uz. Furthermore, let $u \in U$ and $w \in W$ and observe that $\{u, w, y\}$ is an independent set. We first show that if $xz \in E(G)$, then $\{u, w, y\}$ must be an asteroidal triple.

Notice that yR_Uz implies that the removal of $N_G[w]$ preserves a u-y-path. Similarly, xR_Uy implies that yR_Wx, and thus the removal of $N_G[u]$ preserves a y-w-path. The graph $G - N_G[y]$ contains a u, x-path, because xR_Uy, and, similarly, $G - N_G[y]$

contains a z, w-path, because zR_Wy. Suppose $xz \in E(G)$, then it connects both paths to a u-w-path in $G - N_G[y]$, yielding that $\{u, w, y\}$ is an asteroidal triple. However, according to Corollary 6.1.15, 2-moplex graphs are AT-free, a contradiction. Thus, $\{x, y, z\}$ is an independent set.

To prove the transitivity of R_U, we need to show that xR_Uz. Note first that vertices x and z are distinct, because otherwise conditions xR_Uy and yR_Ux would hold simultaneously. Suppose towards a contradiction that U does not prefer x to z. Then, because x and z are non-adjacent, U prefers z to x. Furthermore, because $\{x, y, z\}$ is an independent set, there exists a minimal x,z-separator S such that $y \notin S$. Recall that U and W lie in the same component of $G - S$ as z and x, respectively. Now notice that because S separates y from either x or z, it follows that some $S' \subseteq S$ is either a minimal x,y-separator, or a minimal y,z-separator.

Suppose first that S' is a minimal x,y-separator. As zR_Ux and $S' \subseteq S$, the vertex x remains in the same component together with W in $G - S'$, and thus xR_Wy. But then xR_Uy violates the fact that $R_W = R_U^{-1}$, a contradiction. The case when S' is a minimal y,z-separator is similar. This concludes the proof that xR_Uy and yR_Uz indeed imply xR_Uz. Therefore, the relation R_U is transitive. By symmetry, so is R_W. □

We remark that R_U is a strict partial order on the vertices of G. Furthermore, because R_U is an orientation of the edges of the complement of G, Lemma 6.1.34 implies the main result of this section.

Proposition 6.1.35. Every 2-moplex graph is a cocomparability graph.

Proof. Let G be a 2-moplex graph. If G is a complete graph, then G is cocomparability. Otherwise, G has exactly two moplexes U and W. Consider the orientation of the edges of the complement of G obtained by orienting each edge $\{x, y\} \in E(\overline{G})$ from x to y if and only if xR_Uy. By Lemma 6.1.33, this orientation is well-defined. Furthermore, by Lemma 6.1.34, it is a transitive orientation. Thus, G is a cocomparability graph. □

Proposition 6.1.35 is a strengthening of Corollary 6.1.15 with algorithmic consequences for the class of 2-moplex graphs. First, Weighted Independent Set is solvable in linear time in the class of cocomparability graphs [KM16], and thus also in the class of 2-moplex graphs. In the more general class of AT-free graphs, this problem is only known to be solvable in time $\mathcal{O}(|V(G)|^3)$ [Kö99]. Furthermore, Proposition 6.1.35 implies that the class of 2-moplex graphs is a subclass of the class of perfect graphs, and thus Clique (and its weighted generalisation), Clique Cover, and Colouring are all solvable in polynomial time in the class of

2-moplex graphs [GLS88]. Note that this conclusion cannot be derived from Corollary 6.1.15: because Independent Set and Colouring are NP-hard in the class of C_3-free graphs (see [Ale03, KKTW01]), Clique and Clique Cover are NP-hard in the class of $3K_1$-free graphs (and thus in the more general class of AT-free graphs), whereas the complexity of Colouring is still open in the class of AT-free graphs (see [BKKM99, KM12]).

We show next that the result of Proposition 6.1.35 is best possible in the sense that one cannot find a smaller hereditary graph class containing all 2-moplex graphs.

Proposition 6.1.36. Every cocomparability graph is an induced subgraph of some connected 2-moplex graph.

Proof. Let G be a cocomparability graph. By Theorem 6.1.25, G has an umbrella-free ordering σ. We write $\sigma = (v_1, \ldots, v_n)$ where $i < j$ if and only if $v_i <_\sigma v_j$. Consider the graph G' obtained from G as follows (see Figure 6.7 for an illustration of the construction):

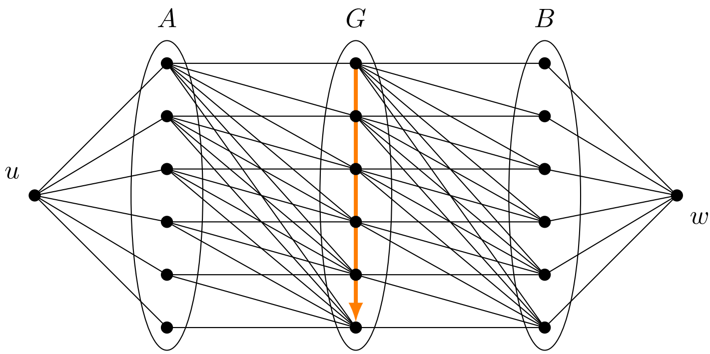

Figure 6.7: The graph G' constructed from the graph G whose vertices are arranged according to a cocomparability ordering.

- add a set $A = \{a_1, \ldots, a_n\}$ of vertices and a vertex u such that $A \cup \{u\}$ is a clique;
- add a set $B = \{b_1, \ldots, b_n\}$ of vertices and a vertex w such that $B \cup \{w\}$ is a clique;
- for all $i, j \in \{1, \ldots, n\}$ such that $i \leq j$, add an edge from a_i to v_j and an edge from v_i to b_j.

The constructed graph G' contains G as an induced subgraph and G' is connected. Because $N_{G'}(u) = A$ and $N_{G'}(w) = B$ are cliques, $U = \{u\}$ and $W = \{w\}$ are moplexes in G'. To complete the proof, it suffices to show that these are the only moplexes of G'. Since no two distinct vertices of G' have the same closed neighbourhood, every moplex in G' consists of a single vertex.

We show that for all $v \in V(G') \setminus \{u, w\}$, the set $\{v\}$ is not a moplex in G'. By symmetry, we may assume that $v \in \{a_i, v_i \mid 1 \leq i \leq n\}$. First, consider the case that $v = a_i$. Since u is a neighbour of a_i not adjacent to any vertex in $G' - N_{G'}[a_i]$, no component of $G' - N_{G'}[a_i]$ can dominate $N_{G'}(a_i)$. Thus, $\{a_i\}$ is not a moplex. Next, assume that $v = v_i$. To show that $\{v_i\}$ is not a moplex, we need to verify that no component of $G' - N_{G'}[v_i]$ dominates $N_{G'}(v_i)$. Consider the vertex sets

$$X = \{u\} \cup \{a_j \mid i < j \leq n\} \cup (\{v_j \mid i < j \leq n\} \setminus N_{G'}(v_i)) \text{ and}$$
$$Y = (\{v_j \mid 1 \leq j < i\} \setminus N_{G'}(v_i)) \cup \{b_j \mid 1 \leq j < i\} \cup \{w\},$$

and let C and D denote the subgraphs of $G' - N_{G'}[v_i]$ induced by X and Y, respectively. By construction, C and D are connected. Note that G' contains no edges from a vertex in $A \cup \{u\}$ to a vertex in $B \cup \{w\}$ and also no edge from a vertex in $\{v_j \mid 1 \leq j < i\} \setminus N_{G'}(v_i)$ to a vertex in $\{v_j \mid i < j \leq n\} \setminus N_{G'}(v_i)$, because $(v_1, \ldots, v_n)$ is an umbrella-free ordering of G. It follows that the graph $G' - N_{G'}[v_i]$ has exactly two components, namely C and D. By construction, C contains no vertex adjacent to b_i and D contains no vertex adjacent to a_i. Since a_i and b_i are adjacent to v_i, we infer that $\{v_i\}$ is not a moplex. Therefore, G' is connected 2-moplex, as claimed. □

Let us explain how the above results establish that both $\mathscr{M}^+$ and $\mathscr{M}_c^+$ coincide with the class of cocomparability graphs, as stated in (6.1).

Corollary 6.1.37. $\mathscr{M}^+ = \mathscr{M}_c^+ = \mathcal{C}$, where $\mathcal{C}$ is the class of cocomparability graphs.

Proof. Since $\mathscr{M}^+ \subseteq \mathscr{M}_c^+$, it suffices to show that $\mathcal{C} \subseteq \mathscr{M}^+$ and $\mathscr{M}_c^+ \subseteq \mathcal{C}$.

The inclusion $\mathcal{C} \subseteq \mathscr{M}^+$ is an immediate consequence of Proposition 6.1.36. For the inclusion $\mathscr{M}_c^+ \subseteq \mathcal{C}$, consider an arbitrary graph G in $\mathscr{M}_c^+$. Then, there exists a graph G' such that every connected component of G' is a 2-moplex graph and G is an induced subgraph of G'. By Proposition 6.1.35, every connected component of G' is a cocomparability graph. Theorem 6.1.25 implies that the class of cocomparability graphs is closed under disjoint union. It follows that G' is a cocomparability graph, and, since the class of cocomparability graphs is hereditary, so is G. We conclude that $\mathscr{M}_c^+ \subseteq \mathcal{C}$, as desired. □

In particular, we obtain the following result as announced above.

Theorem 6.1.5. The smallest hereditary graph class containing the class of 2-moplex graphs is the class of cocomparability graphs.

This is very intuitive as it immediately implies our observation from earlier that chordal graphs with two simplicial modules are interval graphs, as interval graphs are the intersection of chordal graphs and cocomparability graphs.

6.1.6 Hardness results

As seen in Subsection 6.1.5, the class of 2-moplex graphs is a proper subclass of the class of cocomparability graphs. In this subsection we show that two classical problems, namely Max-Cut and Graph Isomorphism, remain as hard on the class of 2-moplex graphs as they are on cocomparability graphs.

Hardness of Max-Cut on 2-moplex graphs

The Max-Cut problem is defined as follows.

Max-Cut	
Input:	An undirected graph G and $k \in \mathbb{N}$.
Question:	Does G contain an edge-cut of size at least k?

Recall that we have established that the class of connected 2-moplex graphs is sandwiched between the classes of connected proper interval and cocomparability graphs.

It is known that Max-Cut is NP-complete on cobipartite graphs [BJ00], and thus on cocomparability graphs. Interestingly, the complexity of Max-Cut is still open on proper interval graphs [BDFG$^+$04, KMN20, ABMR21, dFdMdSOS21].

We show that the problem remains NP-complete on cobipartite graphs with only two moplexes. The hardness reduction is based on the following construction (see also Figure 6.8).

Construction 6.1.38. Let $G = (A \cup B, E)$ be a cobipartite graph such that A and B are disjoint cliques. We define the graph G' obtained from G as follows:

- add a set A' containing $|A|$ vertices and a vertex u such that $\{u\} \cup A \cup A'$ is a clique;
- add a set B' containing $|B|$ vertices and a vertex w such that $\{w\} \cup B \cup B'$ is a clique;
- fix a vertex $a^* \in A'$ and connect it to every vertex in $B \cup B'$;
- fix a vertex $b^* \in B'$ and connect it to every vertex in $A \cup A'$.

Lemma 6.1.39. Let $G = (A \cup B, E)$ be a cobipartite graph such that A and B are disjoint cliques and G' be the graph obtained from G by Construction 6.1.38. Then, G' is a cobipartite graph with exactly two moplexes.

Proof. First, notice that the sets $A \cup A' \cup \{u\}$ and $B \cup B' \cup \{w\}$ are cliques and that hey form a partition of the vertex set of G'. This implies that G' is cobipartite. Furthermore, by construction $\{u\}$ and $\{w\}$ are simplicial moplexes. Next, fix a vertex $a \in A \cup A'$. We observe that $N_{G'}[u] \neq N_{G'}[a]$, because a is adjacent to b^* but u is not. Additionally, as u is a neighbour of a and $N_{G'}[u] \subseteq N_{G'}[a]$, no component of $G' - N_{G'}[a]$ dominates $N_{G'}(a)$. Hence, a does not belong to any moplex. A similar argument can be used to show that every vertex $b \in B \cup B'$ does not belong to any moplex. Thus, G contains exactly two moplexes. $\square$

Theorem 6.1.6. Max-Cut is NP-complete on cobipartite 2-moplex graphs.

Proof. As one can verify in polynomial time whether a given edge set is a maximum cut of the graph, the problem is in NP. To prove NP-hardness, we reduce from Max-Cut on cobipartite graphs, which is known to be NP-hard [BJ00]. Let $G = (A \cup B, E)$ be a cobipartite graph with disjoint cliques A and B. Then, let G' be the graph obtained from G using Construction 6.1.38. Note that G' can be obtained in polynomial time. By Lemma 6.1.39, we have that G' is a cobipartite graph with exactly two moplexes. We complete the proof by showing that there exists an edge-cut of size at least k in G if and only if there exists an edge-cut of size at least $(|A| + 1)^2 + (|B| + 1)^2 + k$ in G'.

First, assume we have an edge-cut $\partial_G(Z)$ in G of size at least $k \geq 1$. Without loss of generality, we assume that $|A \cap Z| \geq 1$ and $|B \cap \overline{Z}| \geq 1$. Indeed, if, say, $A \cap Z = \emptyset$, then $Z \subseteq B$ and $A \subseteq \overline{Z}$, in which case we can achieve the desired inequalities by swapping the roles of Z and $\overline{Z}$. We define a set $Z' \subseteq V(G')$ as follows. The first step is to choose a subset $A'_1 \subseteq A' \setminus \{a^*\}$ such that $|\,(Z \cap A) \cup \{a^*\} \cup A'_1| = |A| + 1$ and a subset $B'_2 \subseteq B' \setminus \{b^*\}$ such that $|\,\big(\overline{Z} \cap B\big) \cup \{b^*\} \cup B'_2| = |B| + 1$. Then, we define $Z' :=$

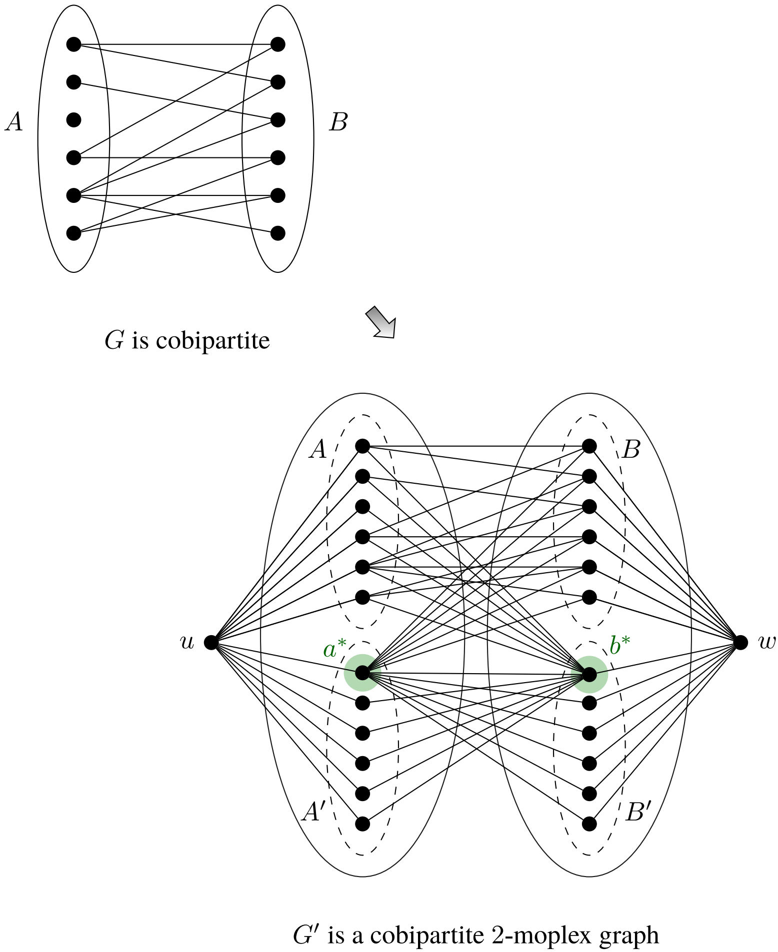

Figure 6.8: An example of Construction 6.1.38 used to prove Theorem 6.1.6 with the vertices a^* and b^* marked in green. The ellipses represent cliques.

$Z \cup \{a^*, w\} \cup A'_1 \cup (B' \setminus B'_2)$, which implies $\overline{Z'} = \overline{Z} \cup \{b^*, u\} \cup B'_2 \cup (A' \setminus A'_1)$. This gives us that $|\,(A \cup A') \cap Z'| = |A| + 1$ (and consequently $|\,(A \cup A') \cap \overline{Z'}| = |A| - 1$) as well as $|\,(B \cup B') \cap Z'| = |B| - 1$ (and consequently $|\,(B \cup B') \cap \overline{Z'}| = |B| + 1$).

Now, we reason about the size of the edge-cut $\partial_{G'}(Z')$. Consider the set of cut-edges with both endpoints in $A \cup A' \cup \{u, b^*\}$. It contains $2\,(|A| + 1)$ edges between $A \cup A'$ and $\{u, b^*\}$ and, because $A \cup A'$ is a clique, $(|A| + 1)\,(|A| - 1)$ edges within $A \cup A'$.

Thus, there are exactly $2(|A|+1)+(|A|+1)(|A|-1) = (|A|+1)^2$ cut-edges whose endpoints both lie in $A \cup A' \cup \{u, b^*\}$. Similarly, there are $(|B|+1)^2$ cut-edges whose endpoints both lie in $B \cup B' \cup \{w, a^*\}$. Hence, adding the (at least k) cut-edges between A and B, we obtain that $\partial_{G'}(Z')$ is an edge-cut of size at least $(|A|+1)^2 + (|B|+1)^2 + k$ in G'.

Second, assume that we have an edge-cut $\partial_{G'}(Z')$ in G' of size at least $(|A|+1)^2 + (|B|+1)^2 + k$. Consider the cut $\partial_G(Z)$ in G defined by $Z := Z' \cap (A \cup B)$ which implies $\overline{Z} = \overline{Z'} \cap (A \cup B)$. There are three possible cases:

(Case1) $u, b^* \in Z'$;

(Case2) $u \in Z'$ and $b^* \in \overline{Z'}$, or $u \in \overline{Z'}$ and $b^* \in Z'$;

(Case3) $u, b^* \in \overline{Z'}$.

Consider the subgraph H'_A of G' induced by $A \cup A' \cup \{u, b^*\}$ and let $p = |(A \cup A') \cap Z'| - |A|$. Note that $|(A \cup A') \cap Z'| = |A| + p$ and $|(A \cup A') \cap \overline{Z}| = |A| - p$. In (Case1), the number of cut-edges in H'_A is exactly $(|A|+p+2)(|A|-p) = (|A|+1)^2 - (p+1)^2$. In (Case2), the number of cut-edges lying in H'_A is exactly $(|A|+p)(|A|-p) + (|A|+p) + (|A|-p) = (|A|+1)^2 - \left(p^2+1\right)$. And finally, in (Case3), the number of cut-edges lying in H'_A is exactly $(|A|+p)(|A|-p+2) = (|A|+1)^2 - (p-1)^2$. Thus, at most $(|A|+1)^2$ cut-edges can lie in H'_A. Similar arguments show that the number of cut-edges lying in the subgraph induced by the vertices in $B \cup B' \cup \{a^*, w\}$ is bounded by $(|B|+1)^2$. By assumption, the size of $\partial_{G'}(Z')$ is at least $(|A|+1)^2 + (|B|+1)^2 + k$, which implies that the remaining k cut-edges must lie between the sets A and B, and thus $\partial_G(Z)$ is a cut of size at least k in G. $\square$

Hardness of Graph Isomorphism on 2-moplex graphs

The Graph Isomorphism problem is defined as follows.

Graph Isomorphism	
Input:	Two graphs G_1 and G_2.
Question:	Are G_1 and G_2 isomorphic to each other?

The Graph Isomorphism problem is solvable in linear time in the class of interval graphs [LB79]. Since the problem is GI-complete on bipartite graphs [UTN05], it is also GI-complete on cobipartite graphs, and thus on cocomparability graphs. Using a reduction from the Graph Isomorphism problem in the class of bipartite graphs, we show that the problem remains hard on cobipartite graphs with at most two moplexes.

Theorem 6.1.7. Graph Isomorphism is GI-complete on cobipartite 2-moplex graphs.

Proof. We reduce from the GI-complete isomorphism problem on connected bipartite graphs [UTN05]; note that the authors claim the GI-completeness only for bipartite graphs but the construction ensures that the obtained graph is connected.

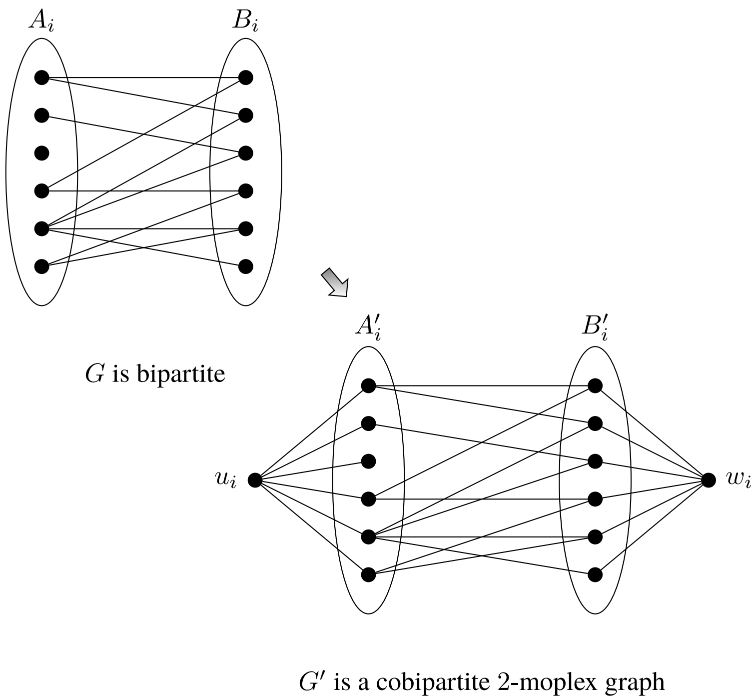

Figure 6.9: An example for the construction used to prove Theorem 6.1.7. The ellipses represent independent sets in G and cliques in G'.

Let G_1, G_2 be connected bipartite graphs with colour classes A_1 and B_1, A_2 and B_2 respectively. We construct two graphs G'_1 and G'_2 as follows. The vertex set of G'_i for $i \in \{1,2\}$ consists of the vertex set of G_i together with two extra vertices u_i and w_i. Each of the sets $A_i \cup \{u_i\}$ and $B_i \cup \{w_i\}$ forms a clique in G'_i, for $i \in \{1,2\}$.

Furthermore, we add an edge to G'_i between $x \in A_i$ and $y \in B_i$ if and only if x and y are adjacent in G_i. See Figure 6.9 for an illustration of this construction.

The obtained graphs are cobipartite. We prove that they are also 2-moplex graphs. In each of the obtained graphs the two new vertices u_i and w_i form simplicial moplexes (of size one). So what remains to show is that there are no additional moplexes in the constructed graphs. Let $a \in A_i$ and suppose it belongs to a moplex M in G'_i. Since G_i is connected, a has a neighbour in B_i, which implies $N_{G'_i}[a] \neq N_{G'_i}[u_i]$, thus we know that $u_i \notin M$. Also, $N_{G'_i}(M)$ is a minimal x,y-separator for two non-adjacent vertices $x, y \in V(G'_i)$. Since $a \in A_i \cap M$, we have $A_i \cup \{u_i\} \subseteq N_{G'_i}[M]$, thus $\{x, y\} \subseteq B_i \cup \{w_i\}$. However, $B_i \cup \{w_i\}$ is a clique in G'_i, contradicting the fact that x and y are non-adjacent in G'_i. By similar arguments no vertex of B_i is moplicial.

To complete the proof, we show that there is an isomorphism $f : G_1 \to G_2$ if and only if there is an isomorphism $f' : G'_1 \to G'_2$. Assume there is an isomorphism $f : G_1 \to G_2$. Since G_1 and G_2 are connected, we must have either $f(A_1) = A_2$ (and then $f(B_1) = B_2$) or $f(A_1) = B_2$ (and then $f(B_1) = A_2$). We extend f to an isomorphism $f' : G'_1 \to G'_2$ by setting

$$(f'(u_1)\,, f'(w_1)) = \begin{cases} (u_2, w_2)\,, & \text{if } f(A_1) = A_2; \\ (w_2, u_2)\,, & \text{if } f(A_1) = B_2. \end{cases}$$

Now, assume there is an isomorphism $f' : G'_1 \to G'_2$. First note that u_i and w_i are the only simplicial vertices in G'_1, and thus we have either $f'(u_1) = u_2$ and $f'(w_1) = w_2$, or $f'(u_1) = w_2$ and $f'(w_1) = u_2$. This immediately implies that their neighbourhoods are also mapped to each other: $f'(A_1) = A_2$ and $f'(B_1) = B_2$, or $f'(A_1) = B_2$ and $f'(B_1) = A_2$. Since f' is an isomorphism, for every $x, y \in A_1 \cup B_1$, we have $\{x, y\} \in E(G'_1)$ if and only if $\{f'(x), f'(y)\} \in E(G'_2)$. Additionally using that A_i and B_i are cliques in G'_i and independent sets in G_i, we infer that $f := f'|_{V(G_1)}$ is an isomorphism between G_1 and G_2. □

We note that the proofs of Theorems 6.1.6 and 6.1.7 also imply stronger statements, namely that Max-Cut is NP-complete and Graph Isomorphism is GI-complete even for cobipartite graphs with at most 2 avoidable vertices (recall Observation 6.1.19).

6.1.7 Hamiltonian properties of 2-moplex graphs

A graph is *traceable* if it contains a Hamiltonian path. It is well-known that every connected proper interval graph is traceable [Ber83]. In this subsection, we generalise this result by proving Theorem 6.1.8, i.e. that every connected 2-moplex graph is traceable. The presented approach relies on the fact that every 2-moplex graph is a cocomparability graph (as proved in Proposition 6.1.35).

Using the results developed so far, we can establish traceability of the following special case of 2-moplex graphs.

Proposition 6.1.40. Every connected graph containing at most two avoidable modules has a Hamiltonian path.

Proof. Let G' be a connected graph containing at most two avoidable modules. We can remove all but one vertex from every avoidable module in G' to obtain a graph G with exactly two avoidable vertices. There is a Hamiltonian path in G' if and only if there is a Hamiltonian path in G.

Applying Observation 6.1.19 and Proposition 6.1.35, we infer that G is a cocomparability graph. By Corollary 6.1.32, the graph G has an umbrella-free LDFS ordering $\sigma = (v_1, \ldots, v_n)$ such that the vertices v_1 and v_n are moplicial. We claim that $(v_1, \ldots, v_n)$ is a Hamiltonian path in G. Suppose that $v_i v_{i+1}$ is a bump in σ. Then, by Lemma 6.1.30, $i \geq 2$ and the vertex v_i is avoidable in G. Since v_1 and v_n are moplicial vertices, they are also avoidable, and thus G contains three distinct avoidable vertices, a contradiction. □

An important ingredient in extending the statement of Proposition 6.1.40 to the whole family of 2-moplex graphs is the following theorem.

Theorem 6.1.41 (Corneil et al. [CDH13b]). Let G be a cocomparability graph with an umbrella-free LDFS vertex ordering σ. If G admits a Hamiltonian path, then one such path corresponds to $\mathrm{DFS}^+(\sigma)$.

We proceed with the proof of traceability for general 2-moplex graphs.

Theorem 6.1.8. Every connected 2-moplex graph has a Hamiltonian path.

Proof. Aiming towards a contradiction, fix a connected 2-moplex graph G of minimum order that does not admit a Hamiltonian path. By Corollary 6.1.32 and Proposition 6.1.35, the graph G has an umbrella-free LDFS ordering $\sigma = (v_1, \ldots, v_n)$ such that v_1 and v_n are moplicial. For $i \in \{1, \ldots, n\}$, we denote by σ_i the ordering of $G - v_i$ obtained from σ by removing v_i. Note that $n \geq 2$ because otherwise G would have a Hamiltonian path. σ_i

The minimality of G implies that no two distinct vertices of G have the same closed neighbourhood. In particular, every moplex in G is of size one. Hence, because G is a 2-moplex graph, v_1 and v_n are the only two moplicial vertices of G. Because $N_G[v_1] \neq N_G[v_n]$, these two vertices belong to distinct moplexes.

Note that, as G does not admit a Hamiltonian path, there is at least one bump in σ. As we show next, no two bumps are consecutive in the ordering and deleting the left endpoint of a bump retains the desirable properties of the ordering σ.

Claim 1. Suppose that $v_i v_{i+1}$ is a bump in σ. Then, σ_i is an umbrella-free LDFS ordering of $G - v_i$, and still starts and ends in moplicial vertices. Furthermore, we have that $v_{i-1}v_i \in E(G)$.

Proof. By Observation 6.1.28, we have $i \geq 2$. Every subsequence of an umbrella-free vertex ordering remains umbrella-free. To prove that σ_i remains an LDFS ordering observe that, by Observation 6.1.28, there is no neighbour $v \in N_G(v_i)$ such that $v_i <_\sigma v$, so all triplets satisfying Definition 6.1.27 remain unaffected. As $v_n \neq v_i \neq v_1$, the vertex v_i is not moplicial. By Lemma 6.1.30, v_i is avoidable in G, and thus, by Theorem 6.1.21, v_1 and v_n are moplicial in $G - v_i$.

We prove that $v_{i-1}v_i \in E(G)$ by contradiction, so suppose $v_{i-1}v_i \notin E(G)$. Now, observe that $N_G(v_i) \subseteq N_G(v_{i-1})$. Indeed, Suppose there exists a vertex $v_j <_\sigma v_i$ with $v_j v_i \in E(G)$ and $v_j v_{i-1} \notin E(G)$, then the vertices v_j, v_{i-1}, and v_i form an umbrella, a contradiction. So, $N_G(v_i) \subseteq N_G(v_{i-1})$ holds. Then, $N_G(v_i)$ is a minimal v_{i-1},v_i-separator in G, and thus $\{v_i\}$ is a moplex. Recall that $i \geq 2$, and thus vertices v_1, v_i, and v_n belong to three distinct moplexes, a contradiction. ■

$\prec$ Let $\prec$ be the predecessor relation given by $\mathrm{DFS}^+(\sigma)$, i.e. for all $u, v \in V(G)$ we have $u \prec v$ if and only if u is the immediate predecessor of v in $\mathrm{DFS}^+(\sigma)$. The next claim looks at how the left vertex of a bump behaves under this relation.

Claim 2. Let $v_i v_{i+1}$ be a bump in σ, and let $x, y \in V(G)$ such that $x \prec v_i \prec y$ and $xv_i, yv_i \in E(G)$. Then, $xy \notin E(G)$.

Proof. Note that, by Observation 6.1.28, the vertices x and y appear before v_i in the ordering σ. Suppose that $xy \in E(G)$ (see Figure 6.10 for an illustration). By Claim 1 we have that σ_i is an umbrella-free LDFS ordering of $G - v_i$. Following Lemma 6.1.30, the vertex v_i is avoidable. Hence v_i is not a cut-vertex, that is, $G - v_i$ is connected. Furthermore, because v_i is not moplicial, Corollary 6.1.22 implies that $G - v_i$ is a

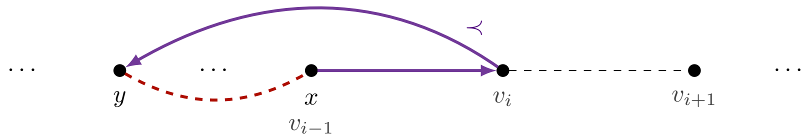

Figure 6.10: The vertices are drawn in order of σ. Edges are solid and non-edges are dashed. The purple, directed edges stand for the $\prec$-relation; they are solid as we know them to be edges. Claim 2 proves that xy is not an edge, thus it is depicted in red.

2-moplex graph. Now, observe that by the minimality of G, the graph $G - v_i$ contains a Hamiltonian path, and thus by Theorem 6.1.41,

$$\mathrm{DFS}^+(\sigma_i) = (v_n, \ldots, x, y, \ldots)$$

corresponds to a Hamiltonian path of $G - v_i$ which extends to a Hamiltonian path $(v_n, \ldots, x, v_i, y, \ldots)$ in G, contradicting the choice of G. ■

Now fix k to be the maximal integer in $\{2, \ldots, n-1\}$ such that $v_k v_{k+1}$ is a bump of σ, and let $G' = G - v_k$. Applying Claim 1 yields that $\sigma_k = (v_1, v_2, \ldots, v_{k-1}, v_{k+1}, \ldots, v_n)$ is an umbrella-free LDFS ordering of G' starting and ending in moplicial vertices.

The graph G' is connected, because v_k is an avoidable vertex in G, and thus not a cut-vertex. Furthermore, because v_k is not moplicial, Corollary 6.1.22 implies that G' is a 2-moplex graph. Thus, the minimality of G and Theorem 6.1.41 imply that $\mathrm{DFS}^+(\sigma_k)$ yields a Hamiltonian path of $G - v_k$.

We consider what this implies about the structure of $\mathrm{DFS}^+(\sigma)$. At this point, we know that $\mathrm{DFS}^+(\sigma)$ and $\mathrm{DFS}^+(\sigma_k)$ share the prefix $(v_n, \ldots, v_{k+1})$, and thus $v_n \prec \ldots \prec v_{k+1}$. Since $v_k v_{k+1}$ is a bump, Observation 6.1.28 implies that v_{k+1} has at least one neighbour in $\{v_1, \ldots, v_{k-1}\}$, and hence v_{k+1} must be adjacent in G to its successor in $\mathrm{DFS}^+(\sigma)$. We show that this successor is indeed v_{k-1}.

Claim 3. We have $v_{k+1} \prec v_{k-1}$.

Proof. Let j be the largest integer such that $j \leq k$, and $v_j v_{k+1} \in E(G)$. The choice of j implies that $v_{k+1} \prec v_j$ and $j \neq k$.

Suppose towards a contradiction that $j \neq k - 1$, that is $v_{k+1} v_{k-1} \notin E(G)$ (see Figure 6.11 for an illustration). This implies that v_j is adjacent to both v_{k-1} and v_k, as otherwise there is an umbrella in σ containing the vertices v_j and v_{k+1}. By Claim 1,

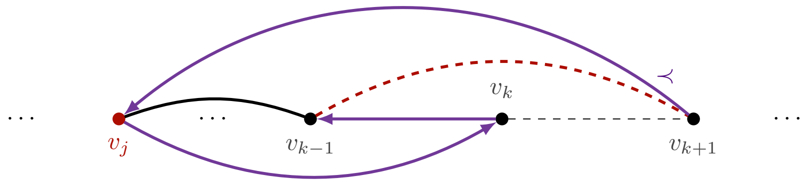

Figure 6.11: The vertices are drawn in order of σ. Edges are solid and non-edges are dashed. The purple directed edges stand for the $\prec$-relation; they are solid as we know them to be edges. Claim 3 proves that $\mathrm{DFS}^+(\sigma)$ visits v_{k-1} after v_{k+1}, as a different vertex v_j would yield a contradiction to Claim 2.

we know $v_k v_{k-1}$ is not a bump, i.e. $v_k v_{k-1} \in E(G)$, which implies $v_{k+1} \prec v_j \prec v_k \prec v_{k-1}$. As $v_j v_{k-1} \in E(G)$ this contradicts Claim 2. ■

From $v_k v_{k-1} \in E(G)$ and the definition of DFS^+ it follows that v_k is visited immediately after v_{k-1}, i.e. $v_{k-1} \prec v_k$. So by now we have established $v_n \prec \ldots \prec v_{k+1} \prec v_{k-1} \prec v_k$. This corresponds to a path of length $n-k+1$ in G, so in the following we assume $k > 2$. We next prove that left of v_{k-1} follows another bump.

Claim 4. We have that $v_{k-2}v_{k-1}$ is a bump in σ.

Proof. Suppose $v_{k-2}v_{k-1}$ is an edge (see Figure 6.12 for an illustration). Then, $v_{k-2}v_k \notin E(G)$, as otherwise $v_{k+1} \prec v_{k-1} \prec v_k \prec v_{k-2}$, which contradicts Claim 2 (with $i = k$). However, if $v_{k-2}v_k \notin E(G)$, then observe that $N_G(v_k) \subseteq N_G(v_{k-2})$, as otherwise v_{k-2} and v_k would form an umbrella in σ with any $x \in N_G(v_k) \setminus N_G(v_{k-2})$, because, by Observation 6.1.28 (with $i = k$) and the fact that $x \neq v_{k-1}$, we have $x <_\sigma v_{k-2}$. The inclusion $N_G(v_k) \subseteq N_G(v_{k-2})$ implies that $\{v_k\}$ is a moplex. However, because v_1 and v_n belong to distinct moplexes in G other than $\{v_k\}$, this contradicts that G is a 2-moplex graph. ■

Since $v_{k-2}v_{k-1}$ is a bump in σ, we can use Observation 6.1.28 and Claim 1 to obtain $k \geq 4$ and $v_{k-2}v_{k-3} \in E(G)$, respectively. Now, we consider which vertex $\mathrm{DFS}^+(\sigma)$ visits after v_k. As v_k is not moplicial, it has degree at least two. Fix v_j to be its neighbour such that the value of $j < k-1$ is maximised, i.e. $v_k \prec v_j$. We show that no value of j is realisable.

As $v_{k-2}v_{k-1} \notin E(G)$ is a bump in σ, $v_{k-2}v_k \notin E(G)$ (see Observation 6.1.28). Thus, we have $j \neq k-2$.

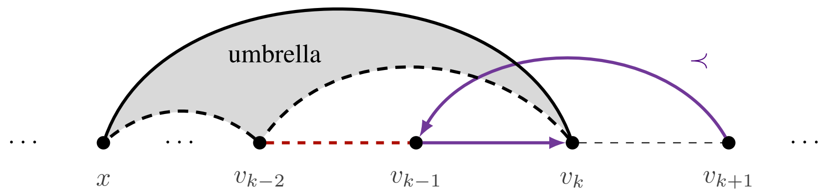

Figure 6.12: The vertices are drawn in order of σ. Edges are solid and non-edges are dashed. The purple, directed edges stand for the $\prec$-relation; they are solid as we know them to be edges. Claim 4 shows that $v_{k-2}v_{k-1}$ is a non-edge as otherwise G contains another moplex.

Suppose next that $j = k-3$, or equivalently $v_{k-3}v_k \in E(G)$. Then, $v_{k-1}v_{k-3} \in E(G)$, as otherwise the vertices $v_{k-3} <_\sigma v_{k-1} <_\sigma v_k$ violate Definition 6.1.27. Since $v_{k-1} \prec v_k \prec v_{k-3}$, Claim 2 implies $v_{k-1}v_{k-3} \notin E(G)$, a contradiction.

Finally, suppose that $j < k-3$ and observe that $v_{k-2}v_j$ and $v_{k-3}v_j$ are edges, as otherwise the corresponding two vertices would form an umbrella in σ with v_k. In particular, we have

$$v_{k+1} \prec v_{k-1} \prec v_k \prec v_j \prec v_{k-2} \prec v_{k-3}\,.$$

By Claim 2 (with $i = k-2$) we get $v_jv_{k-3} \notin E(G)$, a contradiction. This concludes the proof of Theorem 6.1.8. □

Due to [KM14], computing a minimum path cover can be done in linear time for cocomparability graphs. This result, along with Theorem 6.1.8 and Proposition 6.1.35, implies the following.

Corollary 6.1.42. Given a connected 2-moplex graph G, a Hamiltonian path in G can be computed in linear time.

Layering

The presented proof for the existence of a Hamiltonian path in every 2-moplex graph is essentially an inductive argument using Proposition 6.1.40 as its induction basis. There is an alternative way to prove this base staement, that is, every graph with at most two avoidable modules is traceable. Although it more complicated than the straightforward proof provided above, it reveals additional information about the structure of 2-moplex graphs.

Before we proceed to the result itself, we establish a technical lemma that we use in the proof.

Lemma 6.1.43. Let G be a graph with exactly two avoidable modules U and W. Then, for every non-avoidable vertex $x \in V(G)$ and every failing extension x_1xx_2 there exists an induced U-W-path containing $\{x_1, x, x_2\}$.

Proof. First, note that $N_G[x] \setminus \{x_1, x_2\}$ is an x_1,x_2-separator. As every component of $G - N_G[x]$ contains an avoidable vertex of G [BDHT20], the graph $G' := G - (N_G[x] \setminus \{x_1, x_2\})$ contains exactly two connected components, one containing U and the other containing W. Assume without loss of generality that x_1 is in the connected component containing U and x_2 is in the connected component containing W. Fix an induced path Q_1 between some $u \in U$ and x_1, and an induced path Q_2 between x_2 and some $w \in W$. We claim that $P = Q_1 - x - Q_2$ is an induced path in G. As observed earlier, $Q_1 \cup Q_2$ is disconnected in G. Hence, the only obstruction to P being an induced path in G is the edge x_1x_2, which does not exist because x_1xx_2 is an extension of x. Thus, P is an induced path in G. □

Proposition 6.1.40. Every connected graph containing at most two avoidable modules has a Hamiltonian path.

Proof. Let G be a connected graph with at most two avoidable modules. Because the statement is trivial if G is a complete graph, we assume that G is not complete. Thus, G has exactly two avoidable modules. It suffices to consider the case in which the two avoidable modules are of size 1, because they are cliques and a Hamiltonian path containing one of their vertices can easily be extended to contain all their vertices. So let G be a connected graph with exactly two avoidable vertices u and w which are not adjacent.

Since every vertex in a moplex is avoidable, $U = \{u\}$ and $W = \{w\}$ are the only two moplexes of G. Furthermore, as G is a 2-moplex graph, Proposition 6.1.35 and its proof provide a transitive orientation R_U of the complement of G. This yields a partial order $\leq$ on the vertices of G in which without loss of generality u is a minimal and w is a maximal element.

We define a partition of the partial order into pairwise disjoint and non-empty *layers* $L_1, \ldots, L_k$ of pairwise incomparable elements with respect to $\leq$ so that every vertex of G lies in the lowest possible layer, that is, if $x \in L_i$ then for all $j \in \{1, \ldots, i-1\}$ there exists an $x' \in L_j$ with $x' \leq x$. Any two vertices that are not connected by an edge in G are comparable with respect to $\leq$, thus the layers are cliques in G.

By Lemma 6.1.13, vertex u is simplicial in G. Therefore, the first layer L_1 consists precisely of vertex u and its neighbourhood. Next, we show that vertex w lies in the last layer L_k. Suppose otherwise, let $x \in L_k$ and $w \in L_j$ with $j < k$. If $xw \notin E(G)$, then U prefers x to w, that is, $xR_U w$, which contradicts the layer construction. So, x is connected to w by an edge. The layer L_j is a clique and because $N(w)$ is a clique as well x is also connected to every other vertex in L_j. Thus, there is no $x' \in L_j$ with $x' \leq x$, a contradiction to the layer definition.

We call an edge of G an *L_i-jump* if it has one endpoint, called the *lower endpoint*, in a layer L_j with $j < i$, and the other endpoint, called the *upper endpoint*, in a layer $L_{j'}$ with $j' > i$.

Claim 1. For all $x \in L_i$, $1 \leq i \leq k$ there is no failing extension axb of x with $a \in L_j, b \in L_{j'}$ and $j, j' < i$.

Proof. If $i = 1$ this is obvious, so suppose $i > 1$. Suppose there is a vertex $x \in L_i$ with a failing extension axb such that $a \in L_j$ and $b \in L_{j'}$ with $j, j' < i$. Since axb is failing, we have $ab \notin E(G)$, and thus $j \neq j'$. Let us assume without loss of generality that $j' < j$. By the definition of the layers, there is a vertex $x' \in L_j$ with $x' \leq x$ and therefore $xx' \notin E(G)$. We have $x'b \in E(G)$, because otherwise $b \leq x'$ and $x' \leq x$ but x and b are incomparable, contradicting the transitivity. Because L_j is a clique, $ax' \in E(G)$ which closes the failing extension axb, a contradiction. ■

Observe that if $x \in L_i$ has a neighbour $x' \in L_j$ with $j > i$, then it has a neighbour in L_{i+1} due to the definitions of the layers and the transitivity of $\leq$.

Claim 2. For all $1 \leq i \leq k$ a vertex $b \in \bigcup_{j<i} L_j$ can never be the lower endpoint of an L_i-jump and part of a failing extension axb of a vertex $x \in L_i$ with $N_G(x) \subseteq \bigcup_{j \leq i} L_j$.

Proof. Suppose there is a vertex $b \in \bigcup_{j \leq i} L_j$ and a vertex $x \in L_i$ with $N_G(x) \subseteq \bigcup_{j<i} L_j$ such that b is part of a failing extension axb of x and b is also the lower endpoint of an L_i-jump $e = bc$. By definition of L_i-jumps $c \in \bigcup_{j>i} L_j$ and by Claim 1, we have $a \in L_i$. By the assumptions on x, we have that $xc \notin E(G)$. In case $ac \in E(G)$ we obtain a cycle closing the extension axb, a contradiction. And in case $ac \notin E(G)$ the non-edges ba and ac together with the edge e contradict the transitivity of $\leq$. See Figure 6.13 for an illustration. ■

Claim 3. For all L_i such that $1 \leq i < k$, there exists an $x \in L_i$ such that there is some L_j with $j > i$, $x' \in L_j$ and $xx' \in E(G)$.

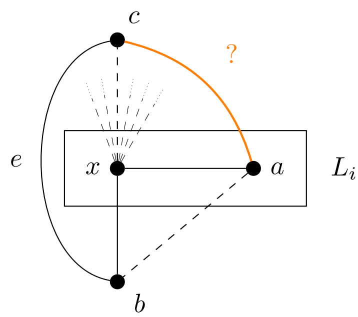

Figure 6.13: An illustration for the proof of Claim 2. The layer L_i with a jump e. Between a and c (marked in orange) we can neither have an edge nor a non-edge, yielding the contradiction.

Proof. Assume $N(L_i) \subseteq \bigcup_{j \leq i} L_j$ for some $i \in \{1, \ldots, k-1\}$. That is, vertices of L_i only have neighbours in the same and lower layers. Since G is connected, it contains an L_i-w-path and thus, because $w \in L_k$ and $k > i$, there is at least one L_i-jump.

Consider some $x \in L_i$ and a failing extension axb of x. By Claim 1, the fact that L_i is a clique in G, and the assumption on the neighbourhood of L_i, we may assume that $a \in L_i$ and $b \notin L_i$. By Lemma 6.1.43 there is an induced u-w-path P containing either axb or bxa as a subpath. Since L_i has no neighbours in higher layers, P has to contain an L_i-jump $e = cd$ with lower vertex d. We know $xd \in E(G)$, because otherwise the non-edges dx and xc together with edge e contradict the transitivity of $\leq$. But neither a nor b, by Claim 2, can be an endpoint of e. Thus x has three neighbours on P contradicting P to be induced. ■

Note that for $1 < i < k$ we have that $|L_i| > 1$, because otherwise Claim 3 and the earlier observation would imply that the unique vertex in L_i has a neighbour in L_{i+1} that is adjacent to all vertices in L_i, contradicting the layer definition.

For $1 \leq i < k$ let $U(L_i)$ be the set of neighbours of L_i in the next layer, that is, $U(L_i) = N(L_i) \cap L_{i+1}$. Clearly, if $|U(L_i)| \geq 2$ for all $1 \leq i < k$ there is a Hamiltonian path in G visiting the layers in increasing order. Also $|U(L_{i-1})| = 1$ still allows for such a path, if $U(L_i)$ has a neighbour in $L_i \setminus U(L_{i-1})$ (the Hamiltonian path reaches L_i in a different vertex).

Thus, we suppose there is a layer L_i such that $x \in L_i$ is the only element in $U(L_{i-1})$ and x' is one of its neighbours in L_{i+1}, $|L_i| > 1$, and $N(U(L_i)) \cap L_i = \{x\}$. This implies that no vertex in $L_i \setminus \{x\}$ has any neighbours in $\bigcup_{j>i} L_j$.

Since $|L_i| > 1$, there exists $y \in L_i \setminus \{x\}$. Let ayb be a failing extension of y. As y does not have neighbours in higher layers and by Claim 1 no failing extension of y can have both vertices in lower layers, we infer that a or b lies in L_i. Without loss of generality let $a \in L_i$ and b lies in a layer L_j with $j < i$. By Lemma 6.1.43 there is an induced u-w-path P containing either ayb or bya as a subpath. If $a \neq x$ then P cannot contain x, because x is adjacent to y, thus P contains an L_i-jump $e = cd$ with lower vertex d. By Claim 2, $d \neq b$, and thus $dy \notin E(G)$, because P is induced. But then the non-edges dy and yc together with edge dc contradict the transitivity of $\leq$. Thus we can assume that $a = x$. Additionally, because $\{x\} = U(L_{i-1}))$, $b \notin L_{i-1}$ with $xb \notin E(G)$.

Let $x'' \in L_{i-1} \cap N_G(x)$. Since y does not have any neighbours in L_{i-1}, we have $yx'' \notin E(G)$. If $x''b \notin E(G)$, then the non-edges bx'' and $x''y$ together with edge by contradict the transitivity of $\leq$. Thus, $x''b \in E(G)$ closing the extension xyb, a contradiction to it being failing. □

Algorithm 2 Hamiltonian path.
Input: a graph $G = (V, E)$ with at most two avoidable modules U and W, layers $L_1, \ldots, L_k$
Out: a Hamiltonian path P in G

---

if $k = 1$ **then**
    return a random Hamiltonian path on the clique L_1
choose $w \in W$
$P = (w)$
$i = k$
while $i \neq 1$ **do**
    let $x \in U(L_{i-1}) \setminus V(P)$
    let Q be a path from x to the first vertex of P visiting every vertex in the clique L_i
    $P = Q \cdot P$
    if $L_{i-1} \setminus U(L_{i-2}) \neq \emptyset$ **then**
        choose $y \in L_{i-1} \setminus U(L_{i-2})$
    else
        choose $y \in N(x) \cap L_{i-1}$
    $i = i - 1$
choose $u \in U$
let Q be a path from u to the first vertex of P visiting every vertex in the clique L_1
$P = Q \cdot P$
return P

The proof of Theorem 6.1.8 not only shows that there always is a Hamiltonian path from one avoidable module U to the other W but also that this path is layer-monotone with respect to the layers given by the transitive ordering of the non-edges in G. Algorithm 2 provides the Hamiltonian path given that the input graph has at most two avoidable modules.

6.1.8 Outlook on moplexes

Graphs with a bounded number of moplexes form interesting graph classes, because moplexes provide a tool that has the potential to lift the beneficial structural properties of simplicial modules in chordal graphs to the setting of all graphs, see, e.g. [BBBS10, BB01a, BB01b]. We introduce the moplex number of a graph, focusing our study on properties of graphs with moplex number 2, the smallest nontrivial class in the moplex-number hierarchy. Some of the questions we answered for the case $k = 2$ can also be asked for $k > 2$. For instance, how do the classes of k-moplex graphs relate with the hierarchy of hereditary graph classes? Also, what is the complexity of Clique, Clique Cover, and Colouring for k-moplex graphs?

In Theorem 6.1.21, we identify a graph operation, the removal of an avoidable non-moplicial vertex, which preserves the moplexes. One can easily determine other such operations (e.g. removing a universal vertex or a true twin). It would be interesting to characterise the class of k-moplex graphs by identifying a set of operations and a base class that can be used to generate every member of the class. Even for $k = 2$ the existence of such a characterisation is still open.

In Theorem 6.1.8 we show that every connected 2-moplex graph is traceable. We conjecture the following strengthening of Theorem 6.1.8.

Conjecture 6.1.44. Every 2-connected 2-moplex graph has a Hamiltonian cycle.

Even for the stronger property of pancyclicity there is no known counterexample.

6.2 Avoidable paths

In [BCG$^+$19], the authors consider a generalisation of the concept of avoidable vertices to edges. An edge uv in a graph G is an *avoidable edge* if every induced path on four vertices with middle edge uv is contained in an induced cycle in G. They generalised Theorem 6.1.17 to that notion.

Theorem 6.2.1 (Beisegel et al. [BCG+19])**.** Every graph that has an edge has an avoidable edge.

This notion naturally generalises to paths. We again define extensions as we did for avoidable vertices. Given an induced path P in a graph G, an *extension* of P is an induced path xPy in G for some vertices x, y. The extension xPy is *failing* if there is no induced cycle of G containing xPy.

Definition 6.2.2 (Avoidable path)**.** A path P in a graph G is an *avoidable path* if it is induced and has no failing extension. Given a subgraph G' of G, we say that P is an *avoidable path of G in G'* if it is avoidable in G and $V(P) \subseteq V(G')$. ⊣

This section presents a result proving that every graph containing an induced P_k also contains an avoidable induced P_k. This was conjectured by Beisegel et al. and established for $k \in \{1, 2\}$ (Ohtsuki et al. 1976 [OCF76], and Beisegel et al. 2019 [BCG+19]) motivated by the following result of Chvátal et al. [CRS02], which generalises Dirac's theorem, Theorem 2.1.1.

Theorem 6.2.3 (Chvátal et al. [CRS02])**.** For every positive integer k, every $C_{\geq k+3}$-free graph either is P_k-free or contains an avoidable path on k vertices.

Originally, Theorem 6.2.3 states the existence of a simplicial path in the class of $C_{\geq k+3}$-free graphs. A *simplicial* path is an induced path with no extension, we also say it is avoidable *by vacuity*. Note that these two definitions coincide on the considered graph class, as no cycle on at most $k + 2$ vertices can contain the extension of an induced path on k vertices.

In this section we mainly prove the following result.

Theorem 6.2.4. For every positive integer k, every graph either is P_k-free or contains an avoidable P_k.

6.2.1 A stronger induction hypothesis

We prove Theorem 6.2.4 using a stronger induction hypothesis, in the exact same flavour as [CRS02], see Theorem 6.2.9 in Subsection 6.2.1. To this end we first define two useful properties.

Definition 6.2.5 (Basic property H_B)**.** Given a positive integer k and a graph G, the property $H_B(G, k)$ holds if either G is P_k-free or there is an avoidable P_k in G. ⊣

Definition 6.2.6 (Refined property H_R)**.** Given a positive integer k, a graph G and a vertex $u \in V(G)$, the property $H_R(G, k, u)$ holds if either $G - N[u]$ is P_k-free or there is an avoidable P_k of G in $G - N[u]$.

Given a positive integer k and a graph G, the property $H_R(G, k)$ holds if $H_R(G, k, u)$ holds for every $u \in V(G)$. $\dashv$

Note that property $H_R(G, k)$ does not directly imply property $H_B(G, k)$. We also emphasise the fact that an avoidable path of a subgraph is not necessarily an avoidable path of the whole graph.

The next lemma states a form of heredity in H_R.

Lemma 6.2.7. Let k be a positive integer, G a graph and u_1u_2 an edge of G and $G' = G/\left(u_1u_2 \to u\right)$. If $G' - N[u]$ contains a P_k, then $H_R(G', k, u)$ implies $H_R(G, k, u_1)$.

Proof. Suppose towards a contradiction that $G' - N[u]$ contains a P_k, and that $H_R(G', k, u)$ holds but $H_R(G, k, u_1)$ does not. Since $G' - N[u]$ is not P_k-free, there is an avoidable P_k of G' in $G' - N[u]$. Call it Q. The path Q is contained in $G' - N[u] = G - N[\{u_1, u_2\}]$, so in particular in $G - N[u_1]$. Since $H_R(G, k, u_1)$ does not hold, Q is not an avoidable P_k of G. Thus, there is a failing extension xQy of Q in G. Note that x, y, u_1, and u_2 are all pairwise distinct.

Hence, xQy is an extension of Q in G', and, because Q is avoidable in G', there is an induced cycle C in G' containing the path xQy. If $u \notin V(C)$, then the cycle C is also an induced cycle in G containing xQy, a contradiction. Therefore, $u \in V(C)$. By replacing u with either u_1, u_2 or the edge u_1u_2 as appropriate, we obtain an induced cycle in G containing xQy, a contradiction. $\square$

This can be strengthened to a connected subset of vertices in a graph instead of a single edge.

Lemma 6.2.8. Let G be a graph, X a subset of its vertices such that $G[X]$ is connected and $G' = G/\left(X \to x\right)$. Assume that $G - N[X]$ contains a P_k. Then $H_R(G', k, x)$ implies that there is an avoidable P_k of G in $G - N[X]$.

Proof. Suppose $H_R(G', k, x)$ holds but there is no avoidable P_k of G in $G - N_G[X]$. Since $G' - N_{G'}[x]$ is not P_k-free, there is an avoidable P_k of G' in $G' - N_{G'}[x]$, call it Q, that is not an avoidable P_k of G. Thus, there is a failing extension aQb of Q in G. Note that a, b, and x are all pairwise distinct.

Hence, aQb is an extension of Q in G', and, because Q is avoidable in G', there is an induced cycle C in G' containing the path aQb. If $x \notin V(C)$, then the cycle C is also an induced cycle in G containing aQb, a contradiction. Therefore, $x \in V(C)$. Now let a', b' be the two neighbours of x on C. By replacing $a'xb'$ in C by a shortest a'-b'-path within X (which exists because $G[X]$ is connected) we obtain an induced cycle C' in G containing aQb, a contradiction. □

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

Theorem 6.2.9. For every $k \in \mathbb{N}^+$ and every graph G, both properties $H_B(G,k)$ and $H_R(G,k)$ hold.

Proof. Suppose the statement is false and consider a counterexample G which is minimal with respect to the number of vertices.

Lemma 6.2.10. The property $H_R(G,k)$ holds for every k.

Proof. We proceed by contradiction. Suppose that $H_R(G,k,u)$ does not hold for some k and some vertex $u \in V(G)$, that is, there exists a P_k in $G - N[u]$, and every P_k in $G - N[u]$ has a failing extension in G. We prove the following.

Claim 1. Every P_k in $G - N[u]$ dominates $N(u)$.

Proof. Assume towards a contradiction that there is a P_k in $G - N[u]$, call it Q, which is not adjacent to some vertex $v \in N(u)$. Then $G - N[\{u,v\}]$ contains a P_k. Let G' be the graph obtained from G by merging u and v into a vertex u'. Since G' has fewer vertices than G, property $H_R(G',k,u')$ holds by minimality of G. Then, by Lemma 6.2.7, also $H_R(G,k,u)$ holds, a contradiction. ■

Let $G' := G - N[u]$. Then G' contains a P_k. As G' has fewer vertices than G, the property $H_B(G',k)$ holds. Let Q be an avoidable P_k of G'. By assumption, Q is not an avoidable P_k of G. So there is a failing extension xQy of Q in G. Since Q has no failing extension in G', we can assume without loss of generality that $y \in N(u)$. It follows that $x \notin N(u)$: otherwise the cycle $xQyu$ contradicts the fact that xQy is failing. By definition of an extension, xQy is an induced path. Let z be the only neighbour of y in Q, and let us now consider the path $xQ - z$ (which is the path obtained from Q by first removing z from one end and then adding x to the other end). It is a P_k, and it does not intersect $N[u]$. However, no vertex in it is adjacent to y which lies in $N(u)$, contradicting Claim 1. ⌟

Lemma 6.2.11. The property $H_B(G, k)$ holds for every k.

Proof. Assume towards a contradiction that for some k, property $H_B(G, k)$ does not hold. By Lemma 6.2.10, the property $H_R(G, k, u)$ holds for every vertex $u \in V(G)$. In other words, the graph G contains a P_k but no avoidable P_k, and for every vertex $u \in V(G)$, either $G - N[u]$ is P_k-free or there is an avoidable P_k of G in $G - N[u]$.

We derive the following claim.

Claim 1. Every P_k in G dominates $V(G)$.

Proof. Suppose there is a P_k, call it Q, that does not dominate some vertex u of G. Since $H_R(G, k)$ holds, either $G - N[u]$ is P_k-free or there is an avoidable P_k of G in $G - N[u]$. The first case contradicts the existence of Q, and the second contradicts the fact that $H_B(G, k)$ does not hold. ■

Since $H_B(G, k)$ does not hold, G contains a P_k, call it Q, that is not avoidable. So it has a failing extension xQy. Let z be the only neighbour of y in Q, and consider the path $xQ - z$. It is an induced P_k and none of its vertices are adjacent to y. This contradicts Claim 1. ⌟

Finally, lemmata 6.2.10 and 6.2.11 together contradict G being a counterexample. □

Theorem 6.2.4 directly follows from Theorem 6.2.9. We point out that the proof of Theorem 6.2.4 is self-sufficient, thus it implies Theorems 6.2.1, 6.2.3 and 6.1.17.

6.2.2 Consequences

Note that Dirac's theorem, Theorem 2.1.1, implies the existence of two simplicial vertices and the results of Beisegel et al. [BCG+19] as well as Ohtsuki, Cheung and Fujisawa [OCF76] also state the existence of two avoidable vertices or edges. By using ingredients of Theorem 6.2.9 (namely Lemma 6.2.7), we also obtain a way to build more than one avoidable P_k. The following corollary follows from Lemma 6.2.8 and Theorem 6.2.9.

Corollary 6.2.12. Let $k \in \mathbb{N}^+$, G be a graph and $X \subseteq V(G)$ such that $G[X]$ is connected, then either $G - N[X]$ is P_k-free or there is an avoidable P_k of G in $G - N[X]$.

We call two paths Q_1 and Q_2 *adjacent* if there is an edge between a vertex of Q_1 and a vertex of Q_2. This allows us to also infer the following statement.

Corollary 6.2.13. Let $k \in \mathbb{N}^+$ and G be a graph. Either G does not contain two non-adjacent P_k, or it contains two non-adjacent avoidable P_k.

Proof. Let Q_1 and Q_2 be two non-adjacent P_k. By Corollary 6.2.12, either $G - N[Q_1]$ is P_k-free or there is an avoidable P_k of G in $G - N[Q_1]$. The first outcome is ruled out by the existence of Q_2. Let Q_2' be an avoidable P_k of G in $G - N[Q_1]$. We repeat the argument with Q_2' instead of Q_1, and obtain an avoidable P_k of G in $G - N[Q_2']$, call it Q_1'. The two paths Q_1' and Q_2' are two non-adjacent avoidable P_k, as desired. $\square$

This raises the question whether every graph G either does not contain two disjoint P_k, or contains two disjoint avoidable P_k.

We know the answer to be positive in case $k \in \{1, 2\}$, due to [BCG+19, Theorems 3.3 and 6.4]. The following counterexample shows that the answer is negative for all $k \geq 3$. Let G be a graph which consists of a cycle on $2k - 1$ vertices with an added vertex adjacent to two consecutive vertices on the cycle (see Figure 6.14 for the case $k = 3$). The graph G contains two disjoint P_k, and it has $2k$ vertices, so the vertex sets of any two disjoint P_k are in fact complementary in the graph. Suppose that the graph contains two disjoint avoidable P_k, and note that each intersects the triangle (otherwise the complement would not be a path). Since there are three vertices in the triangle, there is an avoidable P_k containing a single vertex in the triangle. This P_k has a failing extension containing one of the non-cycle edges from the triangle, a contradiction.

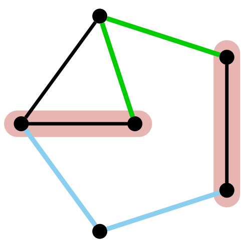

Figure 6.14: A graph that contains two disjoint P_3 (in blue and in green) but no two disjoint avoidable P_3 (there is a unique partition into two disjoint P_3, up to symmetry). In red, a failing extension of the blue path.

6.2.3 An algorithm for Theorem 6.2.9

Here, we consider the algorithmic point of view on finding avoidable paths in graphs. Note that there is a straightforward naive algorithm checking for every subset of size k if it corresponds to an avoidable path.

Our proof of Theorem 6.2.9 is constructive and yields an elementary algorithm that has comparable complexity to the very naive one, but we believe that it provides an outline of the proof, which might be helpful to the reader.

By going through the proof and extracting the key ingredients, we obtain a straightforward algorithm verifying both properties $H_R(G, k)$ and $H_B(G, k)$ (see Algorithm 3). The algorithm has the complexity $\mathcal{O}(n^{k+2})$ which, though naive, is has the right order of magnitude under the ETH.

Algorithm 3 Avoidable path.
Input: graph $G = (V, E)$, integer $k \in \mathbb{N}$
Out: an avoidable P_k in G or `null` if there is none

---

procedure FindAvoidablePathRefined(G, k, u)
    for all $v \in N(u)$ **do**
        if InducedPath($G - N[\{u, v\}], k$) $\neq$ `null` **then**
            $G' \leftarrow G$ with u and v merged into u'
            return FindAvoidablePathRefined(G', k, u')
    return FindAvoidablePath($G - N[u], k$)

procedure FindAvoidablePath(G, k)
    for all $u \in V(G)$ **do**
        if InducedPath($G - N[u], k$) $\neq$ `null` **then**
            return FindAvoidablePathRefined(G, k, u)
    return InducedPath(G, k)

It suffices to consider connected graphs here, because if the graph is not connected, then its components can be computed in linear time and the algorithm can be run on the components separately. The algorithm uses the subprocedure InducedPath that, given a graph G and a positive integer k, decides whether G contains a P_k. If G contains a P_k, the procedure returns a P_k, otherwise it returns `null`. The naive algorithm for that (testing all subsets of size k) has complexity $\mathcal{O}(n^k)$. However, this is nearly optimal. Indeed, the problem of finding a P_k in a given graph is $W[1]$-hard[1] when parametrised by k (see [CFK+15, Ex. 13.16, p. 460]). In fact, the reduction

[1] see, e.g. [CFK+15] for definitions around complexity

there has a linear blow-up, so it follows that there is no $f(k) \cdot n^{o(k)}$ algorithm under ETH.

Fix a positive integer k, then for $n \in \mathbb{N}^+$ we define $B_k(n)$ (resp. $R_k(n)$) to be the worst case complexity of FindAvoidablePath (resp. FindAvoidablePathRefined) on an n-vertex graph with parameter k. We have $B_k(n) \leq n \cdot n^k + \max(R_k(n), n^k)$, and $R_k(n) \leq n \cdot n^k + \max\{R_k(n-1), B_k(n-1)\}$ (here the recursive instances are smaller by one when merging two vertices and smaller by at least 2, when removing the closed neighbourhood because in a connected graph every vertex has at least one neighbour). We obtain $R_k(n) \in \mathcal{O}(n^{k+2})$ and $B_k(n) \in \mathcal{O}(n^{k+2} + n^{k+1})$. While this may well be improved, the known limitations for finding an induced path on k vertices also apply for an induced avoidable path on k vertices (by Theorem 6.2.4, if the first exists, then so does the second). Therefore, this naive algorithm has the same order of magnitude as an optimal solution.

6.2.4 Conclusion

Based on the techniques used in the proof presented above, Gurvich, Krnc, Milanič and Vyalyi [GKMV21] recently strengthened the statement further by showing that one can "shift" any given P_k onto an avoidable one. A shift of a path here, means to iteratively delete an edge at one side and add the incident edge of an extension on the other side.

Theorem 6.2.14 ([GKMV21])**.** Every induced path in a graph can be shifted to an avoidable one.

They also consider other structures such as walks, trails, paths that are not induced, and isometric paths. For walks and paths they obtain the same result for fitting definitions of extensions and shifting, however for isometric paths and trails even the statement that an avoidable such structure exists fails.

It is natural to wonder whether we can find avoidable structures that are further removed from paths. However, for these structures, already the definition of an extension is a challenge. What, for example, would the extension of a clique or a claw look like? Quite likely we would loose the relation to simplicial structures in chordal graphs at this point. What happens when allowing for a family of graphs instead of fixing a single graph? This motivates us to formulate the following question. Does there exist a family $\mathcal{H}$ of connected graphs, not containing any path (or path-like structure), such that any graph is either $\mathcal{H}$-free or contains an avoidable element of $\mathcal{H}$?

Another question arising from the discussions in Subsection 6.2.2 is: When does a graph admit three (or more) disjoint (resp. pairwise non-adjacent) avoidable paths? Note that although Corollary 6.2.12 states sufficient conditions for there to be more than two avoidable P_k, we do not believe that the corresponding sufficient conditions for the case of three or more paths are necessary. However, for this question even the answer for chordal graphs is unclear.

6.3 Perfect linegraph squares

The colouring of graphs is a well known and highly active area of research that has many applications. The concept of colourings for both vertices and edges can be generalised by adding a distance constraint. A *strong vertex colouring* is a proper vertex colouring such that vertices within distance at most 2 of each other receive different colours and a *strong edge colouring* is a proper edge colouring where additionally every two edges sharing end-vertices with the same edge are coloured differently. In other words, the colour classes of a strong edge colouring form an *induced matching*.

The strong edge colouring problem and the related MAXIMUM INDUCED MATCHING problem have received a lot of attention from the network community as this problem appears in the context of interference-free channel assignments (see [RL93], [Ram97], [NKGB00], [JN01], and [AM17]).

It is known that STRONG VERTEX COLOURING [LR92], STRONG EDGE COLOURING [Mah02] and MAXIMUM INDUCED MATCHING [Cam89, SV82] are all NP-complete. This is especially surprising for MAXIMUM INDUCED MATCHING, as MAXIMUM MATCHING is long known to be in P. Moreover, MAXIMUM INDUCED MATCHING stays NP-hard even on bipartite graphs [Cam89] and on planar graphs of maximum degree four [KS94].

While many problems become solvable in more reasonable running time if parametrised, STRONG VERTEX COLOURING remains $W[1]$-hard, even when it is parametrised by treewidth [FGK11].

Finding a strong vertex colouring of G is equivalent to finding a vertex colouring of G^2. That is, if we are able to ensure G^2 to be in a class of graphs where VERTEX COLOURING is known to be in P, we can use this for the construction of new algorithms for STRONG VERTEX COLOURING.

A classical example of graphs on which many NP-complete problems, such as VERTEX COLOURING, become easy to solve are chordal graphs. While squares of chordal graphs are not necessarily chordal again, Duchet proved the following result.

Theorem 6.3.1 (Duchet [Duc84]). If G^k is chordal, then so is G^{k+2}.

An edge colouring of a graph corresponds to a vertex colouring of its linegraph. In the same fashion a strong edge colouring corresponds to a vertex colouring of the squared linegraph. So, in order to find graphs on which the strong edge colouring problem becomes accessible we are interested in $\mathsf{LG}(G)^2$ belonging to a class of graphs that we can colour in polynomial time. Again chordal graphs provide a nice example.

Theorem 6.3.2 (Cameron [Cam89]). If G is a chordal graph, then $\mathsf{LG}(G)^2$ is chordal.

In this section we investigate the structure of graphs whose squared linegraphs have certain properties with respect to the vertex colouring problem. Namely, we find graph families $\mathcal{G}$ and $\mathcal{S}$ such that for any graph G that is $\mathcal{G}$-free, the squared linegraph of G is $\mathcal{S}$-free. In particular, this yields a characterisation of graphs with perfect linegraph squares. Therefore, we obtain a class of graphs on which the strong edge colouring problem can be solved in polynomial time. We extend the results from [SW18] by using similar techniques which are also used in [CST03].

We introduce a few more definitions. For $k \geq 5$ the complement of a cycle C_k, namely $\overline{C_k}$, is called an *antihole*. If we are given a set S of vertices and a vertex $x \in S$, we refer to x as an *S-vertex*. Similarly, if S is a set of edges and $e \in S$, then we refer to e as *S-edge*.

The Strong Perfect Graph Theorem by Chudnovsky et al. allows the description of perfect graphs in terms of forbidden induced subgraphs.

Theorem 6.0.1 (Strong Perfect Graph Theorem [CRST06]). A graph is perfect if and only if it contains neither C_k nor $\overline{C}_k$ as induced subgraph for odd $k \geq 5$.

We are aiming to forbid induced cycles and antiholes on an odd number that is at least five of vertices in the squared linegraph. A possible generalisation of perfect graphs is the concept of χ-boundedness. A class of graphs $\mathcal{C}$ is χ-bounded if there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such all $G \in \mathcal{C}$ satisfy $\chi(G) \leq f(\omega(G))$. A famous result by Gyárfás [Gya87] states that the class of P_t-free graphs is χ-bounded for every $t \geq 1$.

A first step towards this was already taken by Scheidweiler and Wiederrecht, who introduced the structure of *flowers*, the graphs generating induced cycles in graph squares.

F_n **Definition 6.3.3** (Flower)**.** A *flower* of order n is a graph $\mathsf{F}_n = (U \cup W, E)$ with $U = \{u_0, \dots, u_{n-1}\}$ and $W = \{w_0, \dots, w_{q-1}\}$, $\lfloor \frac{n}{2} \rfloor \leq q \leq n$, satisfying the following conditions considering the indices for W-vertices modulo q:

(F1) There is a cycle C containing all vertices of W in the order $w_0, \dots, w_{q-1}$.

(F2) The elements of the set $U = \{u_0, \dots, u_{n-1}\}$ are sorted such that the vertices in $V(C) \cap U$ appear along C in that order. Moreover, we have $u_0 w_{q-1}, u_1 w_0 \in E$ and $u_i u_j \notin E$ for all $j \neq i \pm 1 \pmod{n}$.

(F3) For all $0 \leq i < q$, we have if $w_i w_{i+1} \in E(C)$, then there exists exactly one $u \in U \setminus V(C)$ such that $N_{\mathsf{F}_n}(u) = \{w_i, w_{i+1}\}$. Such a vertex u is called *pending*.

(F4) If $w_i w_{i+1} \notin E(C)$, then there either is exactly one $u \in U \cap V(C)$ adjacent to w_i and w_{i+1}, or there are exactly two vertices $u, t \in U \cap V(C)$, such that $w_i u t w_{i+1}$ is a subpath of C.

(F5) The pending vertices are pairwise non-adjacent and all vertices $u \in U$ that are not pending are contained in C.

If F_n is contained in a graph G and there exists an additional vertex v with $vu_i, vu_j \in E(G)$ for some $j \neq i \pm 1 \pmod{n}$, then F_n is called a *withered flower* or just *withered*. ⊣

The correspondence between a flower and the induced cycle in the graph square provided by the following theorem is used in the proof of Theorem 6.3.13.

Theorem 6.3.4 (Scheidweiler, Wiederrecht [SW18])**.** Let G be a graph, then G^2 contains an induced cycle on $k \geq 4$ vertices if and only if it contains an unwithered flower of order k.

Inspired by this, we start our discussion on induced graphs in G^2 by investigating graphs with P_t-free squares.

6.3.1 Induced paths in G^2

This subsection provides some insight into the structure of the subgraphs that are responsible for the existence of induced paths in the square of a graph G. Our main method is quite technical and uses a lot of case distinctions. The proof of Lemma 6.3.6, which characterises the structures that give rise to induced paths in G^2, is presented in order to illustrate this on a comparatively easy structure.

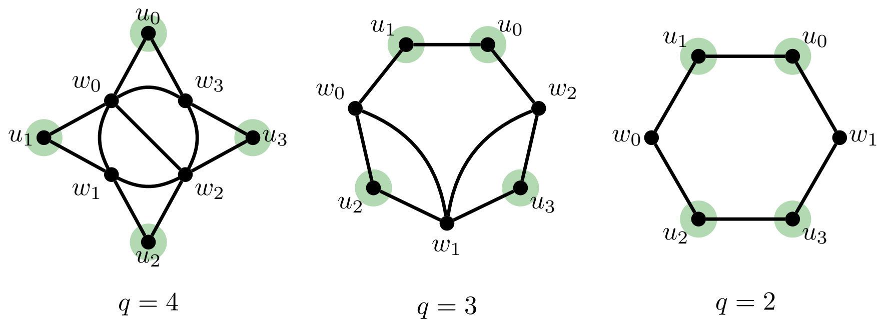

Figure 6.15: Some examples of flowers of order four with different values for q. The marked vertices induce a cycle of length four in the squared graphs.

Definition 6.3.5 (Spire)**.** A *spire* of order n is a graph $\mathsf{S}_n = (U \cup W, E)$ with $U = \{u_0, \ldots, u_{n-1}\}$ and $W = \{w_0, \ldots, w_{q-1}\}$ such that the following conditions hold considering the indices for W-vertices modulo q and for U-vertices modulo n: S_n

(S1) S_n contains a path P with $W \subseteq V(P)$ and $U \subseteq V(P)$. Moreover, the graph $G[U]$ only contains edges that lie on P.

(S2) The path P has the endpoints u_0 and u_{n-1}, and all U- and W-vertices are ordered by their appearance along P.

(S3) No three U-vertices form a subpath of length 2 on P.

(S4) No two W-vertices are consecutive on P.

(S5) For all $0 \leq i < n$, we either have $u_i u_{i+1} \in E(\mathsf{S}_n)$, or $u_i w_j u_{i+1}$ is an induced subpath of P for some j.

(S6) Every $w_j \in W$ is adjacent to exactly two U-vertices and those are consecutive on P.

If S_n is contained in some graph G and there exists an additional vertex v with $vu_i, vu_j \in E(G)$ for some $j \neq i \pm 1$, then S_n is called a *withered spire* or just *withered*.

We refer to U as the *base* of S_n and if $w_j u_i w_{j+1} \subseteq P$ and $w_j w_{j+1} \in E(\mathsf{S}_n)$, u_i is called *pending*. $\dashv$

We prove that unwithered spires indeed ensure the existence of induced paths in the square of a graph and that we can always find one if the square contains an induced path, see Figure 6.16 for some examples.

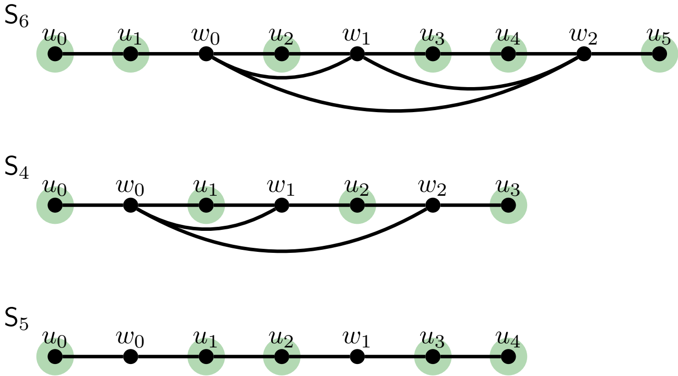

Figure 6.16: Some examples for spires of order four, five and six. The marked vertices induce paths in the respective graph squares.

Lemma 6.3.6. Let G be a graph and $t \geq 1$. There is a set of vertices $U \subseteq V(G)$ such that $G^2[U] \cong P_t$ if and only if G contains an unwithered spire of order t with base U.

Proof. Let $U = \{u_0, \dots, u_{t-1}\} \subseteq V(G)$ such that $G^2[U] \cong P_t$. Assume that the vertices of U are ordered by their appearance along P_t. If $\mathsf{dist}_G(u_i, u_{i+1}) = 2$ for some $0 \leq i < t-1$, then there is a vertex $w \in V(G)$ with $u_i, u_{i+1} \in N(w)$. Suppose there is some $j \in \{0, \dots, t-1\} \setminus \{i, i+1\}$ with $u_j \in N(w)$. Then, $\mathsf{dist}_G(u_i, u_j) \leq 2$ and $\mathsf{dist}_G(u_{i+1}, u_j) \leq 2$ contradicting that P_t is an induced path. Hence, we can collect a set W of such vertices by choosing exactly one such w for every pair of vertices u_i, u_{i+1} of distance 2 in G. This set W immediately satisfies (S6). In addition, we now have a path P on the vertices of U and W in which two consecutive vertices u_i and u_{i+1} are not adjacent if and only if there is some $w \in W$ adjacent to both of them. Thus, (S2) to (S4) hold as well.

As we have seen, all W-vertices connect U-vertices of distance exactly 2 in G. So, if $e \in E(G[U]) \setminus E(P)$, then e joins two U-vertices that do not have a common neighbour in W and are not consecutive on P. This contradicts P_n being induced, thus such edges do not exists and, as P is a path on the vertices of U and W, (S1) holds. (S5) holds by construction and therefore we are done.

For the reverse direction let $\mathsf{S}_t = (U \cup W, E)$ be an unwithered spire in G and P the corresponding path. By (S5), we have $u_i u_{i+1} \in E(G^2)$ for all $0 \leq i < t-1$. So, $G^2[U]$ contains a path of length $t-1$ on the vertices of U. Suppose towards a contradiction that this path is not induced. Hence, there are two U-vertices u_i and u_j with $j \neq i \pm 1$ and $\mathsf{dist}_G(u_i, u_j) \leq 2$. By (S1) and (S2), u_i and u_j cannot be adjacent and because S_t is not withered, there is no $v \in V(G) \setminus V(\mathsf{S}_t)$ adjacent to both. Therefore, there is some $x \in V(\mathsf{S}_t)$ with $u_i, u_j \in N(x)$. By (S1) and (S3), $x \notin U$ and thus $x \in W$. But by (S6), u_i and u_j now have to be consecutive, a contradiction. □

Consider the class $\mathcal{S}_t$ of graphs G such that G^2 excludes P_t as induced subgraph. There has to be an additional vertex in the graph that does not lie within the spire, in order for the spire to be withered. So, our result does not allow us to conclude the χ-boundedness of $\mathcal{S}_t$. However, consider the class $\mathcal{S}_t^2$ of graphs G for which some graph H exists with $H^2 = G$ and all S_t in H are withered. By Gyárfás' Theorem, the class $\mathcal{S}_t^2$ is χ-bounded for every $t \geq 1$.

6.3.2 Induced odd antiholes and perfect squares

We are looking for a way to characterise graphs with perfect graph squares. A first step is Theorem 6.3.4, as it characterises graphs that exclude induced cycles in their square. Here we give a similar way of excluding induced odd antiholes in the graph square.

As antiholes are more complex than paths, we need some more insight regarding their structure. In order to give a compact description of the graphs generating odd antiholes in their squares, we show that a lot of the required structure can be captured by generating cliques of a specific size. The following is a folklore result, but we provide the proof for sake of completeness.

Lemma 6.3.7. Let $t = 2k+1$ for $k \geq 2$. Then, $\overline{C_t}$ has t maximum cliques of size k and for each vertex $v \in V(\overline{C_t})$ there are exactly two cliques K, K' amongst them, such that $V(K) \cap V(K') = \{v\}$.

Proof. Let $V(\overline{C_t}) = \{u_0, \ldots, u_{t-1}\}$ and $C_t = (u_0, \ldots, u_{t-1})$ that is we consider the vertices ordered by appearance along C_t, and we consider the indices module t. Let $0 \leq i < t$, then $C_t - u_i$ is a path of even length, so in particular a bipartite graph with colour classes A containing u_{i-1} and B containing u_{i+1}.

Both A and B induce cliques of size $k = \left\lfloor \frac{t}{2} \right\rfloor$ in $\overline{C_t}$ because they are independent sets in $C_t - u_i$. In addition those cliques are maximum cliques in $\overline{C_t} - u_i$ and neither of them contains u_i. They are both also maximum cliques in $\overline{C_t}$, because $\overline{C_t}$ does not contain edges between neighbours on C_t, neither $u_i u_{i+1}$ nor $u_i u_{i-1}$ is an edge in $\overline{C_t}$. For every $i \in \{1, \ldots, t\}$ let K_i be the clique created by the vertices of the colour class in $C_t - u_i$ containing u_{i+1}. Note that $K_1, \ldots, K_t$ are t distinct cliques. K_i

Next, we show that for every $0 \leq j < t-1$ there are two of these cliques that are disjoint except for containing u_j. We have seen that the K_{j-1} contains u_j and

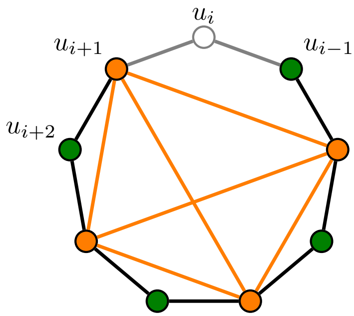

Figure 6.17: The path $C_9 - u_i$ together with a 2-colouring of its vertices and a clique on one of its colour classes.

$k-1 = \frac{1}{2}|N_{\overline{C_t}}(u_j)|$ of its neighbours. We claim that the K_{j+2} also contains u_j together with exactly the other half of $N_{\overline{C_t}}(u_j)$. In order to see this consider the 2-colouring of $C_t - u_{j+2}$. By construction, none of the two colour classes contains u_{j+2}. In addition, the colour class containing u_{j+1}, which is one of the two endpoints of the path $C_t - u_{j+2}$, cannot contain u_j. The path $C_t - u_{j+2}$ is of odd length and thus u_{j+3} is coloured differently from u_{j+1}, which means that it has the same colour as u_j and therefore u_j is contained in K_{j+2}. So, K_{j-1} and K_{j+2} both contain u_j. At last consider K_{j-1}, which contains both u_j and u_{j+2}, so in particular not u_{j+3}. In both cases vertices are collected into the cliques in steps of two along C_t and thus the two cliques are vertex disjoint except for u_j. □

As an immediate consequence we can describe the neighbourhood of every vertex $v \in V(\overline{C_t})$ as two disjoint cliques that can be completed to a maximum clique of our antihole by adding v. For every antihole $\overline{C_t}$, we keep referring to the cliques constructed in the proof of Lemma 6.3.7 by $K_1(\overline{C_t}), \ldots, K_t(\overline{C_t})$.

$K_i(\overline{C_t})$

Corollary 6.3.8. Let $t = 2k+1$ for $k \geq 2$ and $v_1, \ldots, v_t$ the vertices of $V(\overline{C_t})$ sorted by appearance along C_t. Then, for every $0 \leq i < t$ we can cover $N[v_i]$ with exactly two cliques of size k such that v_i is the only vertex appearing in both cliques. These are exactly $K_{i-1}(\overline{C_t})$ and $K_{i+2}(\overline{C_t})$.

For the construction of antiholes, we again introduce sets U and W, but this time the construction focusses on a family of sets $\{U_0, \ldots, U_{t-1}\}$ where $t = 2k+1$ as in Lemma 6.3.7. Each of these sets U_i is going to represent exactly the clique $K_i(\overline{C_t})$.

Definition 6.3.9 (Thornbush)**.** A *thornbush* of order $n = 2k+1$, $k \geq 2$, is a graph $\mathsf{T}_n = (U \cup W, E)$ with $U = \{u_0, \ldots, u_{n-1}\}$ such that the following conditions hold

T_n

considering the indices of U-vertices modulo n:

(T1) The vertices of U are given in a fixed order which we describe by a cycle $C_U = (U, \{u_i u_{i+1} \mid 0 \leq i < n\})$.

(T2) There exists a family $\{U_0, \dots, U_{n-1}\} \subseteq \binom{U}{k}$ such that U_i is the colour class of the bipartite graph $C_U - u_i$ containing u_{i+1}.

(T3) For all $u \in U$ and all $v, v' \in N[u] \cap U$ there is some j such that $v, v' \in U_j$.

(T4) There are sets $W_0, \dots, W_{n-1}$, which are possibly empty and not necessarily disjoint, with $\bigcup_{i=0}^{n-1} W_i = W$ and for all $w \in W_i$ there are $u, v \in U_i$ such that $uw, vw \in E$ and $\mathsf{dist}_{\mathsf{T}_n - w}(u, v) > 2$.

(T5) For all $w \in W$ and $u, v \in U$ with $u, v \in N(w)$ there is some j such that $u, v \in U_j$ and $w \in W_j$.

(T6) For all $u, u' \in U_i$ with $uu' \notin E$ there is a $w \in W_i$ with $uw, u'w \in E$.

We call U the *base* of T_n. Let T_n be contained in some graph G. If u_i is not adjacent to any other u_j and $W \cap N(u_i)$ forms a clique in G, u_i is called *pending*. If G contains a vertex v with $u_i, u_{i+1} \in N(v)$ for some i, we call T_n a *withered thornbush* or just *withered*. ⊣

Lemma 6.3.10. Let G be a graph and $n = 2k+1 \geq 5$, then there is a set of vertices $U \subset V(G)$ such that $G^2[U] \cong \overline{C_n}$ if and only if G contains an unwithered thornbush T_n with base U.

Proof. Assume there is an unwithered thornbush T_n with base $U = \{u_0, \dots, u_{n-1}\}$ in G, and let the u_i be ordered according to the cyclic ordering provided by (T1). Furthermore, let $\{U_0, \dots, U_{n-1}\}$ be the family provided by (T2).

By (T6), we obtain $\mathsf{dist}_G(u', u) \leq 2$ for all $u, u' \in U_i$ and all i. Hence, the U_i form cliques of size k in G^2. Additionally, by the construction of the U_i in (T2), $u_j \in U_i$ implies $u_{j-1}, u_{j+1} \notin U_i$ for all i and j. Hence, due to (T3), $\mathsf{dist}_G(u_j, u_{j-1}) \geq 2$ and $\mathsf{dist}_G(u_j, u_{j+1}) \geq 2$. Also, for every $i \notin \{j-1, j+1\}$ there is some h such that $u_i, u_j \in U_h$ which follows from (T2) and the proof of Lemma 6.3.7.

Suppose there is some w and some i such that $w \in N(u_i) \cap N(u_{i+1})$. If $w \in V(G) \setminus V(\mathsf{T}_n)$, then T_n is withered in G contradicting our assumption, so $w \in U \cup W$. If $w \in U$, by (T3), there is some h such that $u_i, u_{i+1} \in U_h$, which yields a contradiction to the construction of U_h, that is (T2). If $w \in W$, then w belongs to some W_j and, due to (T5), $u_i, u_{i+1} \in U_j$, which again yields a contradiction. Therefore, no such w exists and $N(u_i) \cap N(u_{i+1}) = \emptyset$ for all i.

Hence, for all i and all $j \neq i \pm 1$ we have $u_i u_j \notin E(G^2)$ and therefore $G^2[U] \cong \overline{C_n}$.

Next, let G^2 contain a $\overline{C_n}$ with $U := V(\overline{C_n})$. We order the vertices of $U = \{u_0, \ldots, u_{n-1}\}$ by occurrence along the cycle C_n. This gives us the required ordering for (T1).

By Corollary 6.3.8, we can cover all vertices in U by a cyclic arrangement of exactly n cliques of size k, let $U_0, \ldots, U_{n-1}$ be those cliques. By construction, these cliques fulfil exactly the requirements of (T2).

Since the vertices of every U_i form a clique in G^2 we obtain $\mathsf{dist}_G(u, v) \leq 2$ for all $u, v \in U_i$ for every i. Hence, for every such pair with $uv \notin E(G)$ the set $N(u) \cap N(v)$ is not empty. Now choose W to be a set containing exactly one vertex from $N(u) \cap N(v)$ for every pair of vertices $u, v \in U_i$ with $uv \notin E(G)$ for some i. Consider a vertex $w \in W$ and the set $U_w := N(w) \cap U$. Note that, by construction, U_w forms a clique in G^2. For each pair of vertices $u', v' \in U_w$ we add w to all W_i with $u', v' \in W_i$. Hence, (T4) and (T6) are satisfied.

If there is some $w \in W_i$ and $u \in N(w)$ with $u \in U \setminus U_i$, then, by construction, there is some U_j with $u, v \in U_j$ and $v \in N(w)$. So, (T5) is satisfied as well.

At last, let $u \in U$ be a vertex with neighbours $v, v' \in U$. By Corollary 6.3.8, all neighbourhoods of the vertices of $\overline{C_n}$ are covered by the U_j, hence there is some h with $v, v' \in U_h$ and therefore (T3) holds. $\square$

So, thornbushes have the claimed property to be sufficient and necessary for antiholes in the square of the graph.

6.3.3 Linegraphs

The problem of structures like spires and thornbushes that are responsible for the existence of prescribed induced subgraphs in G^2 is that an additional vertex in the graph G is required to render them withered.

Flotow [Flo97] gives the following construction to show that there is no finite family of forbidden induced subgraphs describing a class $\mathcal{C}'$ such that $G^2 \in \mathcal{C}$ if and only if $G \in \mathcal{C}'$ if we put no bound on the clique number. Suppose there was a finite family $\mathcal{F}$ characterising $\mathcal{C}$ in this way. We construct the graph G by taking a copy of every graph in $\mathcal{F}$ together with one additional vertex v adjacent to everything else. Thus, G^2 is complete, that is, $G^2 \in \mathcal{C}$, contradicting the $\mathcal{F}$-free graphs to be exactly those with squares in $\mathcal{C}$.

This construction exploits the problem of additional vertices responsible for the existence of paths of length 2. This problem does not occur if we consider linegraphs.

In order for a vertex to be responsible for a structure being withered in $\mathsf{LG}(G)$, an edge in G has to contain endpoints of two other edges that are not supposed to be connected in $\mathsf{LG}(G)^2$. For unwithered spires and induced paths this is illustrated in Figure 6.18. Hence, the edge is part of the subgraph induced by the vertex set of the edges producing the structure in $\mathsf{LG}(G)$.

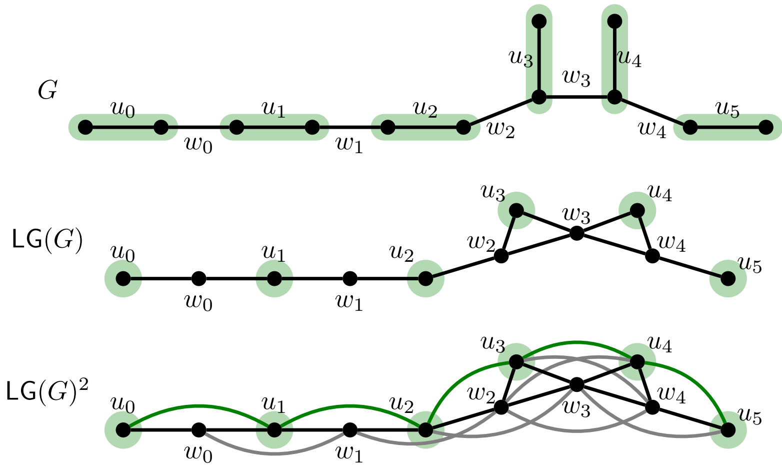

Figure 6.18: A graph G containing an unwithered spire in its linegraph together with the induced P_6 in $\mathsf{LG}(G)^2$.

When considering induced cycles in the squared linegraph of some graph G, then we are interested in structures in G that yield unwithered flowers in $\mathsf{LG}(G)$. That is, we need a structure having the same connection to cycles as spires have to paths.

Definition 6.3.11 (Sprout [SW18]). A *sprout* of order n is defined as a graph $\mathsf{ST}_n = (V, U \cup W \cup E)$ with $|U| = n$ and $|W| = q$, U, W and E having a pairwise empty intersection and $\lceil \frac{n}{2} \rceil \leq q \leq n$ satisfying the following conditions where we consider the indices of U-vertices modulo n and indices of W-vertices modulo q: ST_n

(ST1) There is a cycle C with $E(C) \supseteq W$ containing the edges of W in the order $w_0, \ldots, w_{q+1}$.

(ST2) The elements of $U = \{u_0, \ldots, u_{n-1}\}$ are sorted by appearance along C with $u_0 \cap w_{q-1} \neq \emptyset$ and $u_1 \cap w_0 \neq \emptyset$. In addition, $u_i \cap u_j = \emptyset$ for all $j \neq i \pm 1 \pmod{n}$.

(ST3) If $w_i \cap w_{i+1} \neq \emptyset$, then there is exactly one $u \in U$ with $(w_i \cap w_{i+1}) \cap u \neq \emptyset$. These edges are called *pending*.

(ST4) If $w_i \cap w_{i+1} = \emptyset$, then there either is one $u \in U$ connecting w_i and w_{i+1} in C, or there are exactly two edges $t, t' \in U$, such that the graph induced by t and t' is a path starting on w_i, ending on w_{i+1} and being part of C.

(ST5) The pending U-edges are pairwise non-adjacent and every U-edge that is not pending is an edge in $E(C)$.

If a sprout $\mathsf{ST}_n = (V, U \cup W \cup E)$ contains an edge $e \in E$ connecting two non-consecutive u-edges, we say ST_n is *infertile*, otherwise ST_n is called *fertile*. ⊣

Lemma 6.3.12 (Scheidweiler, Wiederrecht [SW18])**.** Let G be a graph and $n \in \mathbb{N}$. Then, $\mathsf{LG}(G)$ contains an unwithered flower of order n if and only if G contains a fertile sprout of order n.

Theorem 6.3.13 (Scheidweiler, Wiederrecht [SW18])**.** Let G be a graph and $n \in \mathbb{N}$. Then, $\mathsf{LG}(G)^2$ contains a cycle of length n if and only if G contains a fertile sprout of order n.

By extensive case analysis it is possible to make use of the phenomenon described above in order to characterise the graphs with a certain pre-described structure, such as chordal graphs, more succinctly. This is illustrated by the following theorem.

Theorem 6.3.14 (Scheidweiler, Wiederrecht [SW18])**.** Let G be a graph. Then, $\mathsf{LG}(G)^2$ is chordal if and only if G does not contain a C_n with $n \geq 6$, or one of the graphs in Figure 6.19.

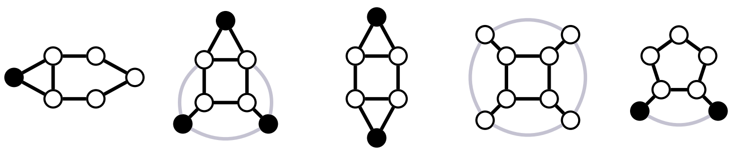

Figure 6.19: The 5 types of forbidden induced subgraphs (the grey edges may or may not exist) from Theorem 6.3.14.

6.3.4 Induced paths in linegraph squares

In order to forbid P_t in $\mathsf{LG}(G)^2$ we translate spires into the world of linegraphs. We call the structure a spire corresponds to in its linegraph a *plantlet*.

Similar to sprouts, the definition is very technical but for $t = 4$ we are able to make use of it in order find a small family of four types of graphs to forbid in G.

Theorem 6.3.15. Let G be a graph. Then, $\mathsf{LG}(G)^2$ is P_4-free if and only if G does not contain a graph of type A), B), C), or D).

We come back to the identification of these graphs after defining plantlets and proving their correspondence to spires.

Definition 6.3.16 (Plantlet)**.** A *plantlet* of order n, $\mathsf{PL}_n = (V, U \cup W \cup R)$, is a graph with $U = \{u_0, \dots, u_{n-1}\}$, $W = \{w_0, \dots, w_{q-1}\}$ and U, W and R pairwise disjoint such that the following conditions hold considering the indices of U-vertices modulo n and the indices of W-vertices module q: PL_n

- **(P1)** $V = \bigcup_{u \in U} u$.
- **(P2)** If $u \in U$ is pending, then none of its endpoints belongs to another U-edge.
- **(P3)** There is a path P with $E(P) \subseteq U \cup W$ and $W \subseteq E(P)$ such that for every $u_i \in U \setminus E(P)$ there is a unique vertex $\{v\} = u_i \cap V(P)$. U-edges with a vertex that is not on P are called *pending*. Every pending U-edge is adjacent to two W-edges. Both, U- and W-edges, are ordered by their appearance along P.
- **(P4)** No U-edge contains vertices of two other U-edges.
- **(P5)** Either $u_i \cap u_{i+1} \neq \emptyset$, or there exists a unique $w_j \in W$ such that $u_i \cap w_j \neq \emptyset$ and $u_{i+1} \cap w_j \neq \emptyset$.

If there is an edge $e \in R$ and $i, j \in \{0, \dots, n-1\}$ with $|i - j| \geq 2$ such that $e \cap u_i \neq \emptyset$ and $e \cap u_j \neq \emptyset$, then PL_n is called *infertile*, otherwise it is called *fertile*. ⊣

We show that fertile plantlets correspond to unwithered spires in the linegraph, see Figure 6.18 again for an illustration.

Lemma 6.3.17. A graph G contains a fertile plantlet $\mathsf{PL}_n = (V, U \cup W \cup R)$ if and only if $\mathsf{LG}(G)$ contains an unwithered spire $\mathsf{S}_n = (U \cup W, E)$ for some $E \subseteq E(\mathsf{LG}(G))$.

Proof. Let $\mathsf{PL}_n = (V, U \cup W \cup R)$ be a fertile plantlet of G. By (P3) there is a path P in PL_n consisting only of U- and W-edges. Clearly, u_0 is the first edge of P. Suppose u_{n-1} is not the last edge of P, then the last edge is w_{q-1} and u_{n-1} is pending at the vertex in which w_{q-2} and w_{q-1} meet. But then the edge w_{q-1} has an endpoint in P that is not in any U-edge, which contradicts (P1). Hence, u_{n-1} is in fact the last edge of P. Now, consider $\mathsf{LG}(G)$. Let P'' be the path in $\mathsf{LG}(G)$ starting on u_0, ending on u_{n-1} and visiting all W-vertices and all non-pending U-vertices as they are visited as edges along P. Let u be a pending U-vertex, i.e. a pending U-edge in G. By definition

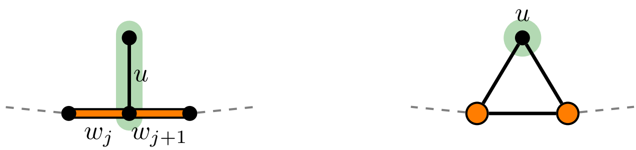

Figure 6.20: A pending U-edge in G as a pending U-vertex in $\mathsf{LG}(G)$.

of pending U-edges, u shares a vertex with two consecutive W-edges w_j and w_{j+1} on P. Let v be the vertex contained in u, w_j and w_{j+1}. By construction, the path P'' contains both w_j and w_{j+1} and because v is contained in all three edges, u, w_j and w_{j+1} form a clique in $\mathsf{LG}(G)$. Moreover, w_j and w_{j+1} are adjacent on P''. Now, obtain a path P' from P'' by replacing the edge w_jw_{j+1} by the edges w_ju and uw_{j+1}. The constructed path P' contains all U- and W-vertices and no other vertices. Also, no two W-vertices are adjacent on P' and, by (P4), no three consecutive U-vertices form a subpath of P'. Furthermore, if two U-edges in G share a vertex, they are adjacent, as vertices, on P'. So (S1) to (S4) are satisfied. The property (P5) implies (S5) and (S6) and thus the subgraph of $\mathsf{LG}(G)$ induced by $U \cup W$ is a spire of order n. Suppose it is withered. Then, there is a vertex x in $\mathsf{LG}(G)$ adjacent to two non-consecutive U-vertices, that is, there is an edge x in G containing vertices of two non-consecutive U-edges. Thus, our PL_n is infertile. This contradicts our assumption and thus $\mathsf{LG}(G)$ contains an unwithered spire of order n.

Now, assume there is an unwithered spire $\mathsf{S}_n = (U \cup W, E)$ in $\mathsf{LG}(G)$. By (S2), S_n contains a path P' with endpoints u_0 and u_{n-1} containing, by (S1), all U- and W-vertices in their respective order. Due to (S3), no three consecutive U-vertices form a subpath of length 2 of P'. Also, by (S4), no two W-vertices are adjacent on P'. Furthermore, all edges of $\mathsf{LG}(G)[U]$ are on P'.

Hence, in G we can construct a path P containing u_0, u_{n-1} and all W-edges as follows. Let $P_0 = u_0$ be the path with exactly one edge. If $u_0 \cap u_1 \neq \emptyset$, then obtain P_1 from P_0 by appending u_1, otherwise append w_0 instead (now the end-vertex of P_1 is the vertex contained in u_1 and w_0). Now, assume that for some $j \geq 1$ the path P_j is already constructed to be starting with the edge u_0 and ending in a vertex of u_j. In order to construct P_{j+1} we have to consider three cases (illustrated in Figure 6.21).

(Case1) If the last two edges of P_j are u_{j-1} and u_j, then the next edge should be w_h for some h, followed by either w_{h+1} or u_{j+1}.

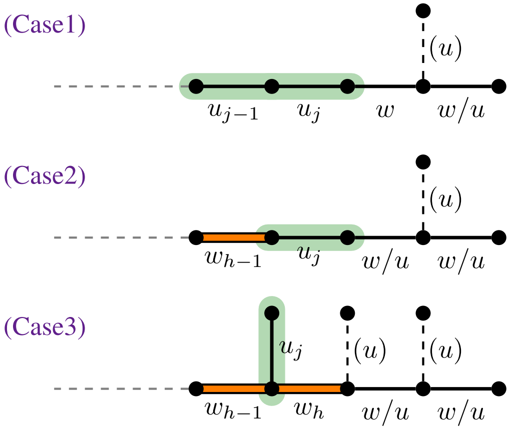

Figure 6.21: The three cases in the construction of the path P_j.

(Case2) If the last two edges of P_j are w_{h-1} for some h and u_j, then the next edge may either be w_h, or u_{j+1}, followed by either w_h, if the first was u_{j+1}, or w_{h+1} or u_{j+1}, if the first was w_h.

(Case3) If the last edge of P_j is w_{h-1} for some h. Then, by our construction, u_j was not added to P because w_h shares a vertex with w_{h-1} and thus u_j is pending. So the next edge is definitely w_h followed by either w_{h+1}, or u_{j+1}.

Considering P_j ending on a vertex of u_j, we either have u_j being an edge of P_j leading to (Case1) and (Case2), or this is the only vertex of u_j on P_j and thus u_j is a pending edge and we have (Case3).

First, consider (Case1), that is, both u_{j-1} and u_j are edges of P_j and consecutive. Then, u_{j-1} and u_j are also neighbours on P' and so there exists some w_h that is next on P' and thus shares a vertex with u_j in G. By (S6), w_h is not adjacent to u_{j-1} and thus $w_h \cap u_{j-1} = \emptyset$. The vertex of w_h that is not shared with u_j is contained in u_{j+1}. If it is also contained in w_{h+1}, then u_{j+1} is pending and we obtain P_{j+1} from P_j by appending w_h only. Otherwise, we obtain P_{j+1} from P_j by appending first w_h and then u_{j+1}.

Next, consider (Case2) and the last two edges of P_j are w_{h-1} and u_j. If u_j and u_{j+1} share a vertex we obtain P_{j+1} from P_j by appending u_{j+1}. Otherwise, the endpoint of P_j belongs to w_h. Again we distinguish the cases from before, either u_{j+1} is pending and we obtain P_{j+1} from P_j by appending w_h only, or it is not pending and we obtain P_{j+1} from P_j by appending first w_h and then u_{j+1}.

At last, consider (Case3), so u_j is not an edge of P_j. Then, u_j was not added to P_j during the construction, meaning that it is pending. So, we know that there is some w_h sharing a vertex with both u_j and the last edge of P_j, namely w_{h-1}. Again we consider the two possible cases: Either u_{j+1} is pending, then we obtain P_{j+1} from P_j by adding w_h, or, otherwise, we obtain P_{j+1} from P_j by adding first w_h and then u_{j+1}. This construction we iterate until obtaining P_{n-1} and then we define $P \coloneqq P_{n-1}$.

At no point in this construction we add a vertex to P that does not belong to U and the only vertices not belonging to $V(P)$ but contained in some edge of $U \cup W$ are the endpoints of the pending U edges, so (P1) is satisfied. Also, we satisfy (P2) and (P3) for the same reasons. If (P4) was not satisfied we would have an edge joining two non-consecutive U-vertices in $\mathsf{LG}(G)$ contradicting S_n to be unwithered. Finally, (P5) follows again from the construction of P and (S6).

Note that any edge rendering the resulting plantlet of order n in G infertile would result in a vertex in $\mathsf{LG}(G)$ responsible for S_n being withered. Since this is not possible, the constructed plantlet of order n in G is fertile. □

To reach a characterisation of a graph class excluding an induced path of fixed length in $\mathsf{LG}(G)$ in terms of a succinct list of forbidden induced subgraphs, one usually needs a large case distinction. Excluding the P_4 results in a class of perfect graphs [Sei74]. Furthermore, excluding any single induced subgraph not contained in the P_4 results in a class of graphs for which no linear χ-bounding function can exist [RS04a]. So, considering the class of graphs whose squared linegraphs exclude P_4 seems natural in the context of investigating perfect linegraph squares in general. We also obtain the following general observations on graph classes whose linegraph squares exclude induced paths of a certain length.

Lemma 6.3.18. Let $n \geq 2$ and G be a graph such that $\mathsf{LG}(G)^2$ is P_n-free. Then, G is $P_{\lceil 3n/2 \rceil}$-free.

Proof. Assume G contains the path $P \coloneqq P_{\lceil 3n/2 \rceil}$ as induced subgraph.

First, consider the case that $n = 2k$ for some k. Then, $\frac{3}{2}n = 3k$. So P has length $3k-1$. Let e_1, e_2 be the first two edges of P and add them to a set U, the remaining path P' has length $3k-3$. Now, divide P' into paths of length 3 and add the leading edge to a set W, while adding the two other edges to U as well. There are $k-1$ such paths, so in total $|U| = 2 + 2k - 2 = 2k = n$. Hence, P is a plantlet of order n.

Second, consider the case that $n = 2k+1$ for some k. Then, $\left\lceil \frac{3}{2}n \right\rceil = 3k+2$, so P is of length $3k+1$. Again, we add the first two edges to a set U, but now we also add the last edge to U and the second to last one to W. What remains is again a path of length

$3k - 3$ for which we use the same decomposition as in the first case. We end up with $|U| = 2 + 2k + 1 - 2 = 2k + 1$, so in this case P is a plantlet of order n as well.

Since P is an induced subgraph of G, this plantlet is fertile and thus, by lemmata 6.3.6 and 6.3.17, $\mathsf{LG}(G)^2$ contains P_n. □

There are two main characteristics by which we can distinguish different plantlets of the same order: The length of the path P of (P3), and the number and position of pending edges along P. By Lemma 6.3.18, we have a lower bound of $\left\lceil \frac{3}{2}n \right\rceil - 1$ on the length of P in a plantlet of order n that does not have any pending U-edges. Next, we show that $2n - 1$ is an upper bound. As long as there are no pending U-edges, these bounds are strict. By allowing pending edges we obtain a lower bound of $n + 2$. To do so, we need two additional lemmata.

Lemma 6.3.19. Let $n \in \mathbb{N}$. Then, for every $i \in \{\left\lceil \frac{3}{2}n \right\rceil, \ldots, 2n\}$ there is a fertile plantlet of order n that does not have any pending U-edges, such that its path P, provided by (P3), is isomorphic to P_i.

Proof. We prove this by induction on i. Lemma 6.3.18 yields the basis for the induction.

So, for the induction step, assume that the claim holds for $i \in \{\left\lceil \frac{3}{2}n \right\rceil, \ldots, 2n - 1\}$. By induction hypothesis, there is a plantlet PL_n of order n whose base path P is isomorphic to P_i and that does not have any pending edges. If PL_n has any edges besides U- and W-edges, we can simply delete them and obtain a smaller plantlet of them same order, hence we can assume $P = \mathsf{PL}_n$.

Since PL_n does not have any pending U-edges, all U edges are part of P and thus there are exactly $i - 1 - n \leq n - 2$ many W-edges in PL_n. So, P contains two consecutive U-edges, call them u and u'. Let v'' be the vertex shared by u and u'. We replace the vertex v'' by the two vertices v and v' such that $v \in u$ and $v' \in u'$ and add the edge vv'. The result is a path $P' \cong P_{i+1}$. If we add vv' to W, P' satisfies all properties of a fertile plantlet of order n, which concludes the proof. □

Lemma 6.3.20. Let $n \in \mathbb{N}$. Then, there is a fertile plantlet PL_n of order n such that its (P3)-path is isomorphic to P_i and it has exactly $2n - i$ pending U-edges if and only if $i \in \{n + 2, \ldots, 2n\}$.

Proof. Assume PL_n is a fertile plantlet of order n such that its (P3)-path P is isomorphic to P_i and it has exactly $2n - i$ pending U-edges for some i. Since there cannot be a negative number of pending edges, $i \leq 2n$ holds. Assume that $i < n + 2$. Then, there are at least $n - 1$ pending U-edges on P which has length at most n. But because the first and last edge of the path are U-edges, there cannot be any pending edges incident to the first two and last two vertices of the path. So, P only has $n-2$ vertices that can be incident to a pending edge. But due to (P2) there can be at most one edge pending on any one of these, therefore there are not enough vertices for the pending edges.

The other direction we again prove by induction on i. We prove a slightly stronger statement, namely that there even is a plantlet PL_n that has next to the properties described in the theorem no additional edges besides U- and W-edges.

For $i = n+2$ consider $P := P_{n+2}$. We add the first and the last edge of P to a set U. The remaining $n-1$ edges are added to a set W. By deleting the endpoints of P, we obtain a path P' consisting entirely of W-edges. This path has exactly $n-2$ internal vertices, let them be numbered $v_1, \dots, v_{n-2}$ as they appear along P'. For each of these vertices v_j, $j \in \{1, \dots, n-2\}$, we introduce a vertex x_j and the edge v_jx_j which is added to U. The resulting graph fulfils all requirements of a fertile plantlet of order n. Note that it also only contains edges that are either in U or in W, which proves the induction basis.

So, for the induction step, assume that the statement holds for $i \in \{n+2, \dots, 2n-1\}$. We construct a plantlet of order n with a (P3)-path of length $i-1$. Let P the (P3)-path of length $i-1$ of the plantlet provided by the induction hypothesis. Also, let uu' be the first pending edge along P with u being a vertex on P and u' the one not on P. The vertex u is incident to exactly two W edges on P, namely w and w'. Let x be the other endpoint of w and y the other endpoint of w'. Then, delete the edge uy and add the edge $u'y$. If we take the path P' that results from P in this construction by adding the edges uu' and $u'y$ together with the remaining pending edges, the result is the desired plantlet. □

Corollary 6.3.21. Let $n \in \mathbb{N}$, then there is no plantlet of order n such that its (P3)-path P is isomorphic to P_i for some $i \in \mathbb{N} \setminus \{n+2, \dots, 2n\}$.

Let us return to P_4 and the identification of the four types of graphs needed for Theorem 6.3.15. In the case of P_4, the (P3)-paths of a plantlet has between six and eight vertices. We categorise these plantlets into seven different types as depicted in Figure 6.22.

We claim that, up to isomorphism, these are all possible plantlets of order four. To prove this, we partition them into three families based on the length of their (P3)-path. By Corollary 6.3.21, there are no other plantlets of order four, so this case distinction suffices.

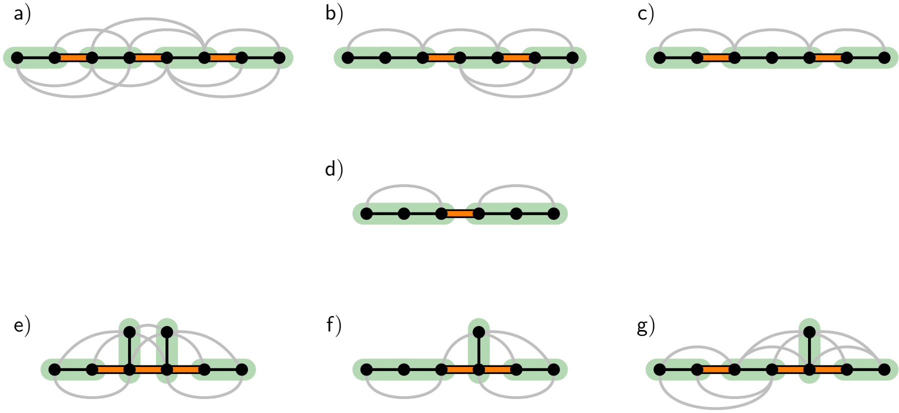

Figure 6.22: The seven types of plantlets of order four. The edges marked in green are the U-edges, the thicker orange edges belong to W and the grey edges are those that may exist in R, because they do not stop the graph from being a fertile plantlet.

When describing the (P3)-path P of a plantlet with $P \cong P_i$ for some i we, from now on, assume $V(P) = \{v_0, v_2, \ldots, v_{i-1}\}$ to be the vertex set of P ordered by appearance along (P3). We refer to the end v_0 of P as the *left* end and then consider the ordering to go *from left to right*.

Lemma 6.3.22. Let H be a fertile plantlet of order four. If its (P3)-path P is isomorphic to P_6, then H is of type d), f), or e) as seen in Figure 6.22.

Proof. First, consider the case that H does not have any pending edges. By the definition of plantlets, v_0v_1 and v_4v_5 are U-edges and, furthermore, v_0v_1 is the first, and v_4v_5 is the last one with respect to their ordering. This leaves three more edges, two of which are U-edges and the last one is a W-edge. Let $P' = (P - v_0) - v_5$ and suppose the two remaining U-edges appear consecutively on P', then either v_1, or v_4 is a vertex of their subpath of length 2 of P'. Say this is true for v_1, then $\{v_0, v_1, v_2, v_3\}$ form a subpath of length 3 on P, contradicting (P4). Therefore, the middle edge v_2v_3 is the only W-edge in P. Because H is fertile, the only additional edges allowed are v_0v_2 and v_3v_5, indeed every other possible edge would connect vertices of two non-consecutive U-edges. Hence, in this case H is of type d).

Second, consider the case that there is exactly one pending edge uv_i with endpoint $u \in V(H) \setminus V(P)$. By (P2), the vertex v_i has to be contained in two consecutive

W-edges in P. Hence, $i \in \{3,4\}$ and the three remaining U-edges provide three edges in P that are not in W. Without loss of generality, say that v_0v_1, v_1v_2 and v_4v_5 are U-edges. Then, $i = 4$ and, for because H is fertile, u can only be adjacent to v_2, v_4 and v_5. In addition the edges v_0v_2 and v_3v_5 may exist as well. All other edges would connect two non-consecutive U-edges contradicting that H is fertile. Therefore H is of type f).

At last, consider the case that there are two pending edges with endpoints $u, u' \in V(H) \setminus V(P)$. By the observations made in the previous case, the two pending edges are incident to v_2 and v_3, so let uv_2 and $u'v_3$ be the pending U-edges. Then, v_0v_1 and v_4v_5 are the two remaining U-edges. Further case distinction reveals that the only additional edges that may exist without rendering H infertile are those allowed in a type e) plantlet. This concludes the proof, because there cannot be more than 2 pending edges in total. ☐

Lemma 6.3.23. Let H be a fertile plantlet of order four. If its (P3)-path P is isomorphic to P_7, then H is of type b), c) or g) as seen in Figure 6.22.

Proof. This proof is, in its structure and the kind of case distinctions, analogous to the proof of Lemma 6.3.22. We consider cases of no, one or two pending edges, determine their possible positions – if there are any – consider all possible distributions of the remaining U-edges in P, using the knowledge that v_0v_1 and v_5v_6 have to be U-edges, and consider possible additional edges.

There are two types of fertile plantlets without any pending edge. In total we have six edges, so four U- and 2 W-edges and the 2 W-edges cannot share a vertex on P. Thus, there is one pair of consecutive U-edges sharing a vertex and this pair either appears at one of the ends, or in the middle of P. In this case we have a plantlet of type b) or c).

There is only one type with exactly one pending edge, which is due to the two W-edges sharing a vertex that are necessary for the edge to be pending. The remaining U-edges on P cannot share vertices and thus these plantlets are of type g).

Finally, we prove that H cannot have two pending edges. Suppose there are two pending edges, then no inner edge of P can be a U-edge, because there only are four in total. Also, for every vertex of P contained in two W-edges we need a pending U-edge. Thus, in order to have a plantlet of order four with 2 pending edges, the path from (P3) can have at most 5 edges, which contradicts the fact that the path in H has length six. ☐

Lemma 6.3.24. Let H be a fertile plantlet of order four and P its (P3)-path such that $P \cong P_8$, then H is of type a) as seen in Figure 6.22.

Proof. Let $P \cong P_8$, then H has at least three W-edges.

Suppose there are at least four W-edges on P. Since the end edges of P belong to U, the only possible positions for the W-edges are the five inner edges of P. With four out of five edges, we would obtain a P_4 subpath of P consisting only of edges from W. By (P5), this guarantees the existence of at least two pending U-edges. If all the five inner edges of P belong to W, we have six U-edges in total, which yields a contradiction to H being of order four. So, there are three U-edges that are part of P and additionally two U-edges pending making a total of 5, again contradicting the order of H.

Therefore, H has exactly three W-edges and its path P is of the form as seen in type a). Further case distinction reveals that the allowed additional edges are exactly those depicted in Figure 6.22, because any other edge would connect two non-consecutive U-edges. Thus, H is of type a). □

Still, it is clear that some of these fertile plantlets of order four are subgraphs of plantlets of the other types, so this family is not minimal with respect to excluding all fertile plantlets of order four. The next step is to reduce the number of forbidden subgraphs to the desired four types A), B), C) and D). Figure 6.23 depicts this smaller family of forbidden subgraphs for excluding plantlets of order four.

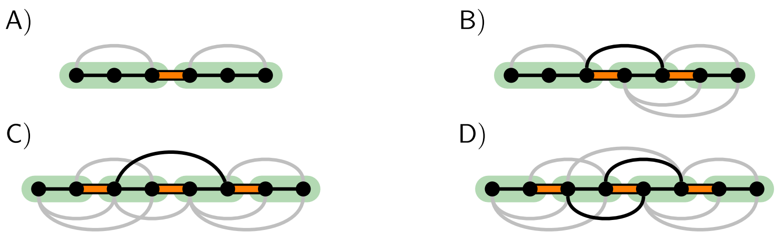

Figure 6.23: The four minimal types of plantlets of order four. Unmarked black edges must necessarily exist, all other colours/patterns are chosen as in Figure 6.22.

Type A) plantlets correspond exactly to the type d) plantlets of Figure 6.22. The type B) plantlets are a specific version of type b) plantlets where the v_2v_4-edge has to exist in order to distinguish them from plantlets of type A). Type C) is obtained from type a) by requiring the additional edge v_2v_5 and allowing all possible combinations of all other allowed edges except for the existence of v_2v_4 and v_3v_5 at the same time. The type D) is obtained by requiring v_2v_4 and v_3v_5 to exist at the same time, distinguishing them from plantlets of type B).

Lemma 6.3.25. Let H be a fertile plantlet of order four, then it contains a plantlet of type $\mathsf{A})$, $\mathsf{B})$, $\mathsf{C})$, or $\mathsf{D})$.

Proof. Let H be a type-$\mathsf{A})$-free fertile plantlet of order four, then in particular H does not contain P_6.

The proof is structured as follows: we discuss the plantlets of order four in order of the length of their (P3)-path. In all cases we only consider those additional edges that have both endpoints on the (P3)-path P.

We start with the plantlets with (P3)-paths of length 5. So we have to consider the plantlets of types $\mathsf{d})$, $\mathsf{f})$, and $\mathsf{e})$. Since none of these plantlets can contain a type $\mathsf{B})$, $\mathsf{C})$ or $\mathsf{D})$ subgraph, we show the existence of a plantlet of type $\mathsf{A})$. Type $\mathsf{d})$ already is exactly type $\mathsf{A})$. So, suppose H is of type $\mathsf{f})$. In this case let u be the unique endpoint of the single pending edge that is not on P and consider $H - u = H[V(P)]$. The only two additional edges allowed are v_0v_2 and v_3v_5, thus H contains a type $\mathsf{A})$ subgraph in any case. Thus, we can assume H to be of type $\mathsf{e})$ with u and u' being the endpoints of the two pending edges that are not on P. Again we consider $H - u - u' = H[V(P)]$ and again the only additional edges allowed are v_0v_2 and v_3v_5, so in this case H contains a type $\mathsf{A})$ plantlet as well.

Next, consider the case where P is of length six and thus H is of type $\mathsf{b})$, $\mathsf{c})$, or $\mathsf{g})$. We show that either an $\mathsf{A})$-, or a $\mathsf{B})$-type induced subgraph exists. Suppose H is of type $\mathsf{b})$. If it contains the edge v_2v_4, we have a type $\mathsf{B})$ subgraph, so we can assume that this edge does not exist. Consider $H' := H[\{v_0, v_1, \ldots, v_5\}]$, disregarding the edges v_3v_6 and v_4v_6. In H' the only edges we can have in addition to the necessary ones are v_0v_2 and v_3v_5 and thus H' contains a type $\mathsf{A})$ subgraph. Suppose H is of type $\mathsf{c})$. If the edge v_2v_4 does not exist, we obtain a subclass of graphs of type $\mathsf{b})$ with a type $\mathsf{A})$ subgraph. If v_2v_4 is an edge of H we obtain exactly the graphs of type $\mathsf{B})$ not containing the edge v_3v_6. Suppose H to be of type $\mathsf{g})$ and let u be its vertex not on P, then $H - u$ is a plantlet of type $\mathsf{b})$. Thus, by the case above, H contains a type $\mathsf{A})$ or a type $\mathsf{B})$ plantlet.

Finally, let H be of type $\mathsf{a})$. If the edge v_2v_5 exists, H is of type $\mathsf{C})$, or, if also v_2v_4 and v_3v_5 exist, we have a plantlet of type $\mathsf{D})$. So, suppose v_2v_5 does not exist. If we have both v_2v_4 and v_3v_5, we have a subgraph of type $\mathsf{D})$ again, so assume one of them, without loss of generality v_3v_5, to be definitely excluded. If also v_2v_4 does not exist, we consider $H[\{v_1, v_2, \ldots, v_6\}]$, which is exactly a plantlet of type $\mathsf{A})$. If $v_2v_4 \in E(H)$, then consider $H - v_7$ is a type $\mathsf{B})$ graph. □

To summarise, we obtain the following theorem for graphs without an induced P_4 in their squared linegraph.

Theorem 6.3.15. Let G be a graph. Then, $\mathsf{LG}(G)^2$ is P_4-free if and only if G does not contain a graph of type $\mathsf{A})$, $\mathsf{B})$, $\mathsf{C})$, or $\mathsf{D})$.

Proof. As the graphs of type $\mathsf{A})$, $\mathsf{B})$, $\mathsf{C})$ and $\mathsf{D})$ are themselves fertile plantlets, Lemma 6.3.25 implies that G contains a graph of type $\mathsf{A})$, $\mathsf{B})$, $\mathsf{C})$, or $\mathsf{D})$, if and only if it contains a fertile plantlet of order four.

Thus, we know that G does not contain a graph of type $\mathsf{A})$, $\mathsf{B})$, $\mathsf{C})$, or $\mathsf{D})$ if and only if G does not contain any fertile plantlet of order four. By Lemma 6.3.17, this is the case if and only if $\mathsf{LG}(G)$ does not contain an unwithered spire of order four. And finally, by Lemma 6.3.6, this is equivalent to $\mathsf{LG}(G)^2$ not containing any induced P_4. $\square$

6.3.5 Induced odd antiholes and perfect linegraph squares

In order to describe the class of graphs with perfect linegraph squares, we need to find a structure similar to sprouts. The edges that will become the vertices of an induced antihole in $\mathsf{LG}(G)^2$ are ordered in a cyclic fashion and we use a cycle to represent this ordering. When deleting a vertex from this cycle we obtain a path, hence a bipartite graph. Whenever we talk about the colour classes of a path obtained this way, we refer to the two classes of the unique 2-colouring of this path.

Definition 6.3.26 (Meristem)**.** A *meristem* of order $n = 2k+1$, $k \geq 2$, is a graph $\mathsf{M}_n = (V, U \cup W \cup R)$ with $U = \{u_0, \ldots, u_{n-1}\}$ such that the following conditions are satisfied: M_n

(M1) $V = \bigcup_{u \in U} u$.

(M2) The edges of U are given in a fixed order which we describe by a cycle $C_U = (U, \{u_i u_{i+1} \mid 0 \leq i < n\})$.

(M3) There exists a family $\{U_0, \ldots, U_{n-1}\} \subseteq \binom{U}{k}$ such that the edges in U_i are identified with the vertices of $C_U - u_i$ from the same colour class as u_{i+1}.

(M4) For all $u \in U$ and $v, v' \in U$ with $v \cap u \neq \emptyset$ and $v' \cap u \neq \emptyset$ there is some $j \in \{1, \ldots, n\}$ with $v, v' \in U_j$.

(M5) There are sets $W_0, \ldots, W_{n-1}$, possibly empty and not necessarily disjoint, with $\bigcup_{i=0}^{n-1} W_i = W$. Additionally, for all $w = xy \in W_i$ there are $u, v \in U_i$ such that $x \in u$, $y \in v$, $u \cap v = \emptyset$ and $e \not\subseteq u \cup v$ for all $e \in (U \cup W) \setminus \{u, v, w\}$.

(M6) For all $w \in W$ and $u, v \in U$ with $u \cap w \neq \emptyset$ and $v \cap w \neq \emptyset$ there is some j such that $u, v \in U_j$ and $w \in W_j$.

(M7) For all $u, u' \in U_i$ with $u \cap u' = \emptyset$ there is a $w \in W_i$ with $u \cap w \neq \emptyset$ and $u' \cap w \neq \emptyset$.

We call U the *base* of M_n. If $u_i \cap u = \emptyset$ for all $u \in U \setminus \{u_i\}$, u_i is called *pending*. If there is an edge $e \in R$ with $e \cap u_i \neq \emptyset$ and $e \cap u_{i+1} \neq \emptyset$ for some $i \in \{1, \ldots, n\}$, M_n is called *infertile*. ⊣

In order to show that fertile meristems are sufficient and necessary for the existence of antiholes in the square of the linegraph, we need two things. First, we need a family of graphs whose squares contain antiholes, this we have already done in Subsection 6.3.2. Such a graph is called a thornbush of order n, where n is the size of the induced antihole it generates when being squared. Second, we need that a fertile meristem of order n in G corresponds to an unwithered thornbush in $\mathsf{LG}(G)$, which we proof in the following lemma.

Proof. Let $\mathsf{M}_n = (V, U \cup W \cup R)$ be a fertile meristem in G. (T1) and (T2) follow directly from the corresponding (M2) and (M3). Let $u, v, v' \in U$. If $v \cap u \neq \emptyset$ and $v' \cap u \neq \emptyset$, then $v, v' \in N_{\mathsf{LG}(G)}(u)$ and by (M4) there is some j with $v, v' \in U_j$, hence (T3) holds. (M5) gives us the existence of sets $W_1, \ldots, W_n$, possibly empty and not necessarily disjoint, with $\bigcup_{i=1}^{n} W_i = W$. This immediately translates into the required sets of (T4). For all $w = xy \in W_i$ there are $u, v \in U_i$ such that $x \in u$, $y \in v$ and $u \cap v = \emptyset$. Hence, u and v are not adjacent in $\mathsf{LG}(G)$, but have distance 2 connected via the vertex w. Since $e \nsubseteq u \cup v$ for all $(U \cup W) \setminus \{u, v, w\}$, there is no other vertex in $\mathsf{LG}(G)$ adjacent to both u and v and thus (T4) holds. Now, let $w \in W$ and $u, v \in N_{\mathsf{LG}(G)}(w)$. Then, $u \cap w \neq \emptyset$ and $v \cap w \neq \emptyset$ and thus (M6) implies the existence of some j such that $u, v \in U_j$ and $w \in W_j$, therefore (T5) holds. At last, let $u, u' \in U_i$ with $uu' \notin E(\mathsf{LG}(G))$, then $u \cap u' = \emptyset$ and therefore (M7) provides some $w \in W_i$ intersecting both u and u'. Hence, $u, u' \in N_{\mathsf{LG}(G)}(w)$ and so (T6) is satisfied.

Thus, we obtain a thornbush $\mathsf{T}_n = (U \cup W, E)$ in $\mathsf{LG}(G)$. Suppose this thornbush is withered, then there is a vertex $v \in V(\mathsf{LG}(G)) = E(G)$ with $u_i, u_{i+1} \in N_{\mathsf{LG}(G)}(v)$. So, v intersects both u_i and u_{i+1} in G contradicting that M_n is fertile.

For the reverse direction let $\mathsf{T}_n = (U \cap W, E)$ be an unwithered thornbush in $\mathsf{LG}(G)$. As before, we use the sets $U, W, U_0, \ldots, U_{n-1}$ and the cyclic ordering of the elements in U provided by (T1) and (T2) in order to satisfy the corresponding (M2) and (M3). Let $u \in U$. If there are any $v, v' \in U$ both intersecting u as edges in G, then they are adjacent in the linegraph and therefore, due to (T3), there exists some j such that $v, v' \in U_j$. Hence, (M4) is satisfied. (T4) implies the existence of sets $W_0, \ldots, W_{n-1}$ with $\bigcup_{i=0}^{n-1} W_i = W$ and for every $w \in W_i$ there are two U-vertices u and v adjacent to w with $\mathsf{dist}_{\mathsf{T}_n - w}(u, v) > 2$, hence u and v are not adjacent and thus $u \cap v = \emptyset$ in

G. So $w = xy$ for vertices $x, y \in V(G)$ and $x \in u$ as well as $y \in v$. Moreover, there is no other path of length 2 in $\mathsf{T}_n \subseteq \mathsf{LG}(G)$, so no edge $e \in (E \cup W) \setminus \{u, v, w\}$ is contained in $u \cup v$ and thus (M5) is satisfied. As before (M6) and (M7) follow directly from (T5) and (T6). Now, we prove (M1). We claim that $w \in \bigcup U$ for all $w \in W$. By (M7), for every pair $u, u' \in U_i$ with $u \cap u' = \emptyset$ there is a $w \in W_i \subseteq W$ that intersects both u and u'. By (M5) and (M6), the reverse is also true, so if there is a $w \in W_i$, then there are two U-edges intersecting w in different endpoints. Therefore, for every $w \in W$ there are $u, v \in U$ such that $w \subseteq u \cup v$ and thus our claim holds.

So let $V := \bigcup U$, then $\mathsf{M}_n = (V, U \cup W \cup R)$ is a meristem, where $R := E(G[V]) \setminus (U \cup W)$.

Suppose M_n is infertile, then there is an edge $e \in R$ intersecting u_i and u_{i+1} for some $i \in \{0, \ldots, n-1\}$. In the linegraph e is adjacent to both u_i and u_{i+1} and so T_n is withered, which contradicts our assumption. □

Theorem 6.3.27. A graph G contains a fertile meristem of order n if and only if $\mathsf{LG}(G)^2$ contains an antihole of size n.

Proof. This follows from lemmata 6.3.10 and 6.3.28. □

Lemma 6.3.28. A graph G contains a fertile meristem $\mathsf{M}_n = (V, U \cup W \cup R)$ if and only if $\mathsf{LG}(G)$ contains an unwithered thornbush $\mathsf{T}_n = (U \cup W, E)$ for some $E \subseteq E(\mathsf{LG}(G))$.

Theorem 6.3.27 allows us to state a first, straightforward characterisation of graphs with perfect linegraph squares by combining it with Theorem 6.3.13.

Corollary 6.3.29. Let G be a graph, then $\mathsf{LG}(G)^2$ is perfect if and only if G does not contain a fertile sprout or fertile meristem of order $n = 2k + 1$ for any $k \geq 2$.

We now further refine this result by taking a closer look at the structure of sprouts and meristems. We make use of the following three lemmata by Scheidweiler and Wiederrecht.

Lemma 6.3.30 (Scheidweiler, Wiederrecht [SW18])**.** Let G be a graph and $k \geq 2$ an integer. If C is an induced cycle in G^k, then G^r cannot contain two consecutive edges of C for all $r \leq \left\lfloor \frac{k}{2} \right\rfloor$.

Lemma 6.3.31 (Scheidweiler, Wiederrecht [SW18])**.** A graph G contains a cycle of length at least four as a, not necessarily induced, subgraph if and only if $\mathsf{LG}(G)$ contains an induced cycle of the same length.

Lemma 6.3.32 (Scheidweiler, Wiederrecht [SW18])**.** Let $k \geq 4$ be an integer. If a graph G does not contain induced cycles of length $\ell \geq k$, then $\mathsf{LG}(G)^2$ contains no induced cycles of length $\ell \geq k$.

First, we show that cycles are fertile sprouts.

Lemma 6.3.33. For all $n \geq 4$ the cycle C_j with $n + \left\lceil \frac{n}{2} \right\rceil \leq j \leq 2n$ is a fertile sprout of order n.

Proof. By Lemma 6.3.30 there are at most $\left\lfloor \frac{n}{2} \right\rfloor$ pairs of adjacent U-edges in a fertile sprout of size n.

First, consider $j = n + \left\lceil \frac{n}{2} \right\rceil$. If there are exactly $\left\lfloor \frac{n}{2} \right\rfloor$ pairs of adjacent U-edges, then there are $\left\lceil \frac{n}{2} \right\rceil - \left\lfloor \frac{n}{2} \right\rfloor \in \{0, 1\}$ remaining U-edges. These cannot be adjacent to any other U-edge. Hence, we need exactly $\left\lceil \frac{n}{2} \right\rceil$ many W-edges to complete the sprout. By alternating between U-edge pairs, W-edges and possibly one single U-edge (not adjacent to other U-edges) we obtain a cycle of length $n + \left\lceil \frac{n}{2} \right\rceil$. The sprout definition allows some additional chords, but no such edge is necessary, so the C_j is a fertile sprout of order n.

Second, we note that each of the $\left\lfloor \frac{n}{2} \right\rfloor$ pairs of U-edges may be split by an additional W-edge, hence with $k \leq \left\lfloor \frac{n}{2} \right\rfloor$ such splits we can produce a cycle of length $n + \left\lceil \frac{n}{2} \right\rceil + k \leq 2n$ which again is a fertile sprout of order n. $\square$

Corollary 6.3.34. A graph G with an induced cycle of length $\ell \geq 8$ contains a fertile sprout of odd order.

Proof. By Lemma 6.3.33, cycles of length $8, 9$ or 10 contain a fertile sprout of order 5. So consider $\ell \geq 11$. We observe that for $n \geq 7$ we have $n + \left\lceil \frac{n}{2} \right\rceil \leq 2\ (n-2) + 1$ and $7 + \left\lceil \frac{7}{2} \right\rceil = 11$. By this observation and Lemma 6.3.33, if ℓ is odd, then the cycle is a fertile sprout of order $\frac{\ell - 1}{2} + 2$. If ℓ is even, then the cycle is a fertile sprout of order $\frac{\ell}{2}$. $\square$

As the squared linegraph of C_7 yields an antihole of size 7 (see Figure 6.24), the following lemma reduces the length of allowed induced cycles even further.

Lemma 6.3.35. If a graph G contains a C_7, $\mathsf{LG}(G)^2$ contains an antihole of size 7.

Proof. If G contains a C_7 Lemma 6.3.31 yields the existence of an induced cycle C of the same length in $\mathsf{LG}(G)$. For each pair of non-adjacent vertices $u, w \in V(\mathsf{LG}(G))$ of C that do not have a common neighbour on C, $\mathsf{dist}_{\mathsf{LG}(G)}(u, w) \geq 3$ holds, because a path of length 2 between two such vertices would correspond to a chord in the C_7, which does not exist. In $\mathsf{LG}(G)^2$ each vertex of C is adjacent to its four nearest vertices on C and not adjacent to the two opposite vertices of C. By reordering the vertices C^2 is an antihole. $\square$

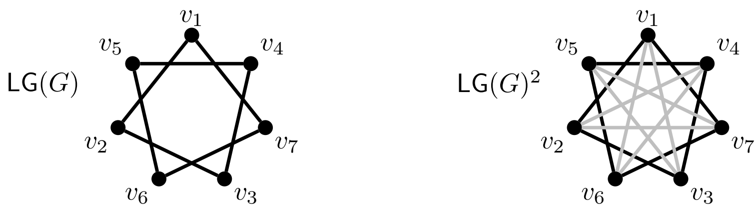

Figure 6.24: The cycle C in the linegraph and its squared linegraph, which is isomorphic to an antihole.

Lemma 6.3.32 implies that a graph G without induced cycles of length at least seven does not contain a fertile sprout of order at least seven, hence $\mathsf{LG}(G)^2$ only contains holes of size at most five. In order to forbid the holes of size five, we exclude all fertile sprouts of order five in G. Since we can exclude the existence of induced cycles of length at least seven, it suffices to consider sprouts of order five with a longest induced cycle, or base cycle, of length five and six. Figure 6.25 depicts the three possible types of sprouts of order five with a base cycle of length five. We proceed by discussing the case of a base cycle of length six.

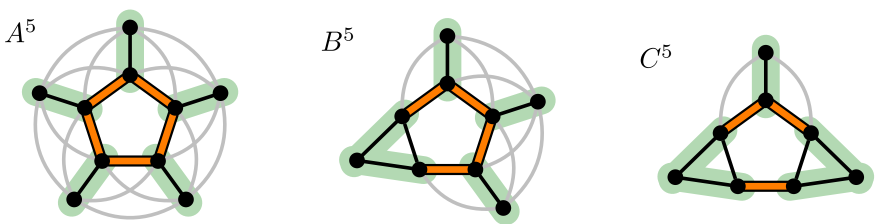

Figure 6.25: The sprouts of order five with a base cycle of length five. Colours are chosen as in figures 6.22 and 6.23.

Lemma 6.3.36. Let $\mathsf{ST}_5 = (V, U \cup W \cup R)$ be a fertile sprout of order five with a base cycle C of length six and $E(C) \cap R = \emptyset$. Then, ST_5 has either three or four pending edges, which are incident to consecutive vertices of C.

Proof. Because ST_5 is a fertile sprout of order five, the number of pending edges is at most five. Assume that ST_5 has five pending edges. Then, the cycle consists of exactly one U-edge and five W-edges. Since pending U-edges can only be incident to vertices on C that are not contained in any other U-edge, the number of pending edges in ST_5 is at most four contradicting the assumption.

Suppose ST_5 has no pending edge, then C contains five U-edges. But, due to Lemma 6.3.30, there are at least $\lceil \frac{5}{2} \rceil = 3$ many W-edges, which must be contained in C as well, because C is the base cycle of ST_5. This contradicts the fact that $|C| = 6$. For the same reason, ST_5 does not have exactly one pending edge. There remain four U-edges and at least three W-edges on C, which contradicts $|C| = 6$ again.

Next, suppose there are exactly two pending edges. There are at least three W-edges necessary for the pending edges and because three U-edges may not form a path on C we need at least one additional W-edge, thus C must now consist of three U-edges and at least four W-edges, again exceeding the length of C.

So there only are two possibilities: Either ST_5 has three or four pending U-edges. These are incident to consecutive vertices of C. Suppose they were not, then the base cycle must contain at least six W-edges and an additional U-edge and if there are just three pending edges not being adjacent to consecutive vertices of C, five W-edges and two additional U-edges are required. □

Notice that every fertile sprout of type B^6 contains a type A^6-sprout. Hence the fertile ST_5 not containing a type A^6-sprout but still having a base cycle of length six are the types C^6, D^6 and E^6 from Figure 6.26.

Corollary 6.3.37. Let $\mathsf{ST}_5 = (V, U \cup W \cup R)$ be a fertile sprout of order five with a longest induced cycle C of length six, then ST_5 contains a fertile sprout ST'_5 of type A^6, C^6, D^6, or E^6 (see Figure 6.26).

Finally, we observe that sprouts of type C^6 are also of type B^6, and sprouts of type D^6 or E^6 certainly contain type A^6 sprouts. The complement of a C_5 is again a C_5, hence the family of meristems of order five is exactly the family of sprouts of order five. With these last observations we can further reduce the number of obstructions to perfect linegraph squares.

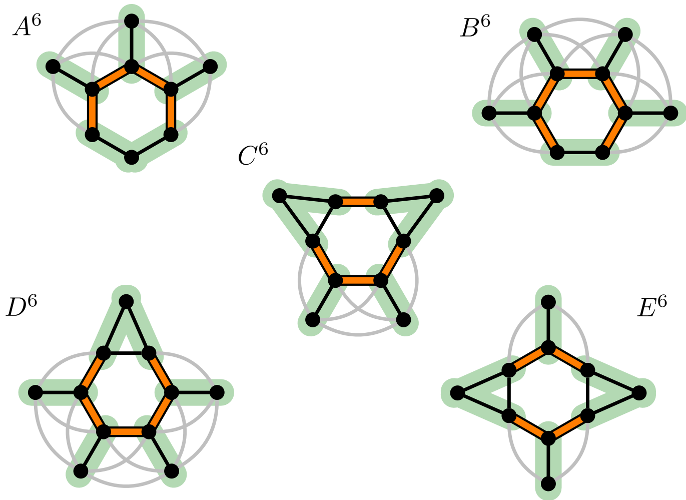

Figure 6.26: The sprouts of order five with a base cycle of length six. Colours are chosen as in figures 6.22 and 6.23.

Theorem 6.3.38. Let G be a graph. Then, $\mathsf{LG}(G)^2$ is perfect if and only if G does not contain a cycle of length $\ell \geq 7$, a fertile sprout of type A^5, B^5, or C^5 (see Figure 6.25), a fertile sprout of type A^6 (see Figure 6.26), or a fertile meristem of order $n = 2k + 1$ with $k \geq 3$.

This implies the following succinct sufficiency condition for perfect linegraph squares.

Corollary 6.3.39. Let G be a graph. If G does not contain a cycle of length $\ell \geq 4$, then $\mathsf{LG}(G)^2$ is perfect.

6.3.6 Concluding χ-boundedness

Theorem 6.3.38 states that the class of graphs with perfect linegraph squares excludes induced cycles of length $\ell \geq 7$. Similarly, if we exclude induced cycles of a certain length in the squared linegraph of a graph G the graph G itself also excludes cycles of some length, this is provided by Lemma 6.3.33. Formally, let $\mathcal{C}_n$ is the class of graphs G such that $\mathsf{LG}(G)^2$ does not contain an induced cycle of length $\ell \geq n$. Then, the graphs in $\mathcal{C}_n$ exclude cycles of length $\ell \geq n + \left\lceil \frac{n}{2} \right\rceil$.

A similar statement holds for graph classes that exclude an induced paths of a certain length in their squared linegraphs, see Lemma 6.3.18. With Gyárfás' Theorem on classes excluding induced paths and the following theorem by Chudnovsky et al. we are able to deduce the χ-boundedness of graph classes with certain excluded induced subgraphs in their linegraph squares.

Theorem 6.3.40 (Chudnovsky et al. 2016 [CSS17])**.** If $\mathcal{C}_n$ is a class excluding induced cycles of length $\ell \geq n$, $\mathcal{C}_n$ is χ-bounded.

Theorem 6.3.41. The following classes are χ-bounded:

- **(i)** the class $\mathcal{C}_{\text{chordal}}$ of graphs G with $\mathsf{LG}(G)^2$ chordal,
- **(ii)** the class $\mathcal{C}_{\text{perfect}}$ of graphs G with $\mathsf{LG}(G)^2$ perfect,
- **(iii)** the classes $\mathcal{C}_{P_t}$ of graphs G with $\mathsf{LG}(G)^2$ being P_t-free, for $t \geq 1$, and
- **(iv)** the class $\mathcal{C}_n$ of graphs G with $\mathsf{LG}(G)^2$ excluding induced cycles of length $\ell \geq n$.

The families of forbidden induced subgraphs for all four of these classes contain many more graphs than just cycles and paths of a certain length. In fact, we forbid plantlets, sprouts or meristems in all of them, which have a far more complicated structure.

Therefore, it might be possible to derive much better χ-bounding functions for some of those classes than those provided by Theorem 6.3.40 and Gyárfás' Theorem. This seems particularly likely for classes like $\mathcal{C}_{P_4}$, which can be described by a finite family of forbidden induced subgraphs as we saw in Theorem 6.3.15.

7 Conclusion

As the different chapters of this thesis are mostly independent, they each contain a conclusion. Still we want to give a short overview on the results and some natural follow-up questions of this thesis.

We introduced a new form of directed tangles, the ganglions. The duality theorem of Chapter 3 proves a very general duality between ganglions directed width measures using $\vec{\mathcal{S}}_k$-DAGs having certain properties. The width measure ν-DAG-width fulfils all of these, and therefore has the described duality to ganglions. However, its decompositions do not yield monotone strategies in the cops and robber reachability game. The other width measure we introduce, μ-DAG-width, yields monotone strategies but does not fulfil all the properties required for the duality to ganglions. The question remains of whether there is a width measure that has both decompositions that yield monotone strategies and a duality to a ganglion.

In Chapter 4 we defined cyclewidth, another digraph width measure, and proved that it is parametrically equivalent to directed treewidth, as well as that it has a close connection to width measures from matching theory. By using this connection we characterised digraphs of small cyclewidth.

We gave the proof for a new directed flat wall theorem and some intuition on how this theorem can be used as the base for a directed structure theorem. The next steps towards such a structure theorem are to formalise these ideas and to identify the exact statement for such a structure theorem. However, we also saw that there is still certain non-planar behaviour we could not bound with these tools. As suggested in Section 5.6, there are a number of options one could consider utilising in order to bound such behaviour as well.

We proved that the class of graphs with at most two moplexes lies between the class of proper interval graphs and the class of cocomparability graphs. Additionally, we proved that Max Cut and Graph Isomorphism, both problems that are hard on cocomparability graphs and tractable or unknown on proper interval graphs, remain hard for the class of graphs with at most two moplexes. However, every graph with at most two moplexes admits a Hamiltonian path, which is not the case for cocomparability graphs in general. Here, there remain the specific question of whether 2-connected graphs with at most two moplexes admit a Hamiltonian cycle, and the very general question about which properties can be determined for graphs with higher moplex number.

Beisegel et al. [BCG+19] generalised the concept of avoidable vertices to edges and proved that every graph that contains edges contains an avoidable edge. They also conjectured that this holds for paths of any length. We proved this conjecture, that is, that every graph with an induced P_k contains an avoidable P_k. We also discussed a few further recent results that consider objects other than paths. However, these objects are still very similar to paths, a natural question is whether we can also find a concept of avoidability for very different structural objects? Another possible direction to consider would be a strengthening of the avoidability definition. For example finding a path for which all extensions close in the same connected component after removing the path and its neighbourhood.

The results in Section 6.3 provide a characterisation via forbidden induced subgraphs of graphs that have a perfect linegraph square. This has not only algorithmic consequences, as for example for Strong Edge Colouring, but also yields some χ-boundedness results. Here, it is natural to ask which other graph classes can be characterised this way. We think that such a characterisation can always be found when the desired class of linegraph squares can be characterised by forbidden induced subgraphs. However, our way of proving such a characterisation by extensive case distinctions does not seem very feasible to repeat for every such class. So the question is whether there is a more general proof characterising a graphs class that excludes certain induced subgraphs in the linegraph square.

Bibliography

[ABMR21] Ranendu Adhikary, Kaustav Bose, Satwik Mukherjee, and Bodhayan Roy. Complexity of Maximum Cut on Interval Graphs. In Kevin Buchin and Éric Colin de Verdière, editors, *37th International Symposium on Computational Geometry (SoCG 2021)*, volume 189 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:11, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[Alc14] Liliana Alcón. On asteroidal sets in chordal graphs. *Discrete Applied Mathematics*, 164:482–491, 2014.

[Ale03] Vladimir E. Alekseev. On easy and hard hereditary classes of graphs with respect to the independent set problem. *Discrete Applied Mathematics*, 132(1-3):17–26, 2003.

[AM17] I. Annammal Arputhamary and M. Helda Mercy. An analytical discourse on strong edge coloring for interference-free channel assignment in interconnection networks. *Wireless Personal Communications*, 94(4):2081–2094, 2017.

[APC90] Stefan Arnborg, Andrzej Proskurowski, and Derek G. Corneil. Forbidden minors characterization of partial 3-trees. *Discrete Math.*, 80(1):1–19, 1990.

[BB72] Umberto Bertele and Francesco Brioschi. *Nonserial dynamic programming*, volume 91. Elsevier, Amsterdam, 1972.

[BB98] Anne Berry and Jean-Paul Bordat. Separability generalizes Dirac's theorem. *Discrete Applied Mathematics*, 84(1-3):43–53, 1998.

[BB01a] Anne Berry and Jean-Paul Bordat. Asteroidal triples of moplexes. *Discrete Applied Mathematics*, 111(3):219–229, 2001.

[BB01b] Anne Berry and Jean-Paul Bordat. Moplex elimination orderings. *Electronic Notes in Discrete Mathematics*, 8:6–9, 2001.

[BBB$^+$05] Anne Berry, Jean Blair, Jean-Paul Bordat, Richard Krueger, and Geneviève Simonet. Extremities and orderings defined by generalized graph search algorithms. *Electronic Notes in Discrete Mathematics*, 22:413–420, 2005.

[BBBS10] Anne Berry, Jean R. S. Blair, Jean Paul Bordat, and Geneviève Simonet. Graph extremities defined by search algorithms. *Algorithms*, 3(2):100–124, 2010.

[BBGM15] Anne Berry, Andreas Brandstädt, Vassilis Giakoumakis, and Frédéric Maffray. Efficiently decomposing, recognizing and triangulating hole-free graphs without diamonds. *Discrete Applied Mathematics*, 184:50–61, 2015.

[BBH$^+$06] Anne Berry, Jean-Paul Bordat, Pinar Heggernes, Geneviève Simonet, and Yngve Villanger. A wide-range algorithm for minimal triangulation from an arbitrary ordering. *Journal of Algorithms*, 58(1):33–66, 2006.

[BBHP04] Anne Berry, Jean R. S. Blair, Pinar Heggernes, and Barry W. Peyton. Maximum cardinality search for computing minimal triangulations of graphs. *Algorithmica*, 39(4):287–298, 2004.

[BCG$^+$19] Jesse Beisegel, Maria Chudnovsky, Vladimir Gurvich, Martin Milanič, and Mary Servatius. Avoidable vertices and edges in graphs. In *Algorithms and data structures. 16th international symposium, WADS 2019, Edmonton, AB, Canada, August 5–7, 2019. Proceedings*, pages 126–139. Cham: Springer, 2019.

[BDFG$^+$04] Hans L. Bodlaender, Celina M. H. De Figueiredo, Marisa Gutierrez, Ton Kloks, and Rolf Niedermeier. Simple Max-Cut for split-indifference graphs and graphs with few P_4's. In *International Workshop on Experimental and Efficient Algorithms*, pages 87–99. Springer, 2004.

[BDH$^+$12] Dietmar Berwanger, Anuj Dawar, Paul Hunter, Stephan Kreutzer, and Jan Obdržálek. The dag-width of directed graphs. *J. Comb. Theory, Ser. B*, 102(4):900–923, 2012.

[BDHK06] Dietmar Berwanger, Anuj Dawar, Paul Hunter, and Stephan Kreutzer. DAG-width and parity games. In *STACS 2006. 23rd annual symposium on theoretical aspects of computer science, Marseille, France, February 23–25, 2006. Proceedings*, pages 524–536. Berlin: Springer, 2006.

[BDHT20] Marthe Bonamy, Oscar Defrain, Meike Hatzel, and Jocelyn Thiebaut. Avoidable paths in graphs. *Electronic Journal of Combinatorics*, 27(4):P4.46, 2020.

[Ber83] Alan A. Bertossi. Finding Hamiltonian circuits in proper interval graphs. *Information Processing Letters*, 17(2):97–101, 1983.

[BG02] Jørgen Bang-Jensen and Gregory Gutin. *Digraphs. Theory, algorithms and applications.* London: Springer, 2002.

[BG18] Jørgen Bang-Jensen and Gregory Gutin, editors. *Classes of directed graphs*. Cham: Springer, 2018.

[BJ00] Hans L. Bodlaender and Klaus Jansen. On the complexity of the maximum cut problem. *Nordic Journal of Computing*, 7(1):14–31, 2000.

[BKKM99] Hajo Broersma, Ton Kloks, Dieter Kratsch, and Haiko Müller. Independent sets in asteroidal triple-free graphs. *SIAM Journal on Discrete Mathematics*, 12(2):276–287, 1999.

[BKMT01] Vincent Bouchitté, Dieter Kratsch, Haiko Müller, and Ioan Todinca. On treewidth approximations. *Electronic Notes in Discrete Mathematics*, 8:98–101, 2001.

[BLS99] Andreas Brandstädt, Van Bang Le, and Jeremy P. Spinrad. *Graph classes: a survey*. SIAM, 1999.

[Bod88] Hans L. Bodlaender. Dynamic programming on graphs with bounded treewidth. Automata, languages and programming, Proc. 15th Int. Colloq., Tampere/Finn. 1988, Lect. Notes Comput. Sci. 317, 105-118, 1988.

[Bod96] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.

[Bod97] Hans L. Bodlaender. Treewidth: Algorithmic techniques and results. In *Mathematical foundations of computer science 1997. 22nd international symposium, MFCS '97, Bratislava, Slovakia, August 25–29, 1997. Proceedings*, pages 19–36. Berlin: Springer, 1997.

[Bod98] Hans L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2):1–45, 1998.

[Bod05] Hans L. Bodlaender. Discovering treewidth. In *SOFSEM 2005: Theory and practice of computer science. 31st conference on current trends in theory and practice of computer science, Liptovský Ján, Slovakia, January 22–28, 2005. Proceedings.*, pages 1–16. Berlin: Springer, 2005.

[BP11] Anne Berry and Romain Pogorelcnik. A simple algorithm to generate the minimal separators and the maximal cliques of a chordal graph. *Information Processing Letters*, 111(11):508–511, 2011.

[BRST91] Dan Bienstock, Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a forest. *J. Comb. Theory, Ser. B*, 52(2):274–283, 1991.

[BT97] Hans L. Bodlaender and Dimitrios M. Thilikos. Treewidth for graphs with small chordality. *Discrete Applied Mathematics*, 79(1-3):45–61, 1997.

[Bun74] Peter Buneman. A characterisation of rigid circuit graphs. *Discrete Math.*, 9:205–212, 1974.

[BW99] Kenneth P. Bogart and Douglas B. West. A short proof that "proper = unit". *Discrete Mathematics*, 1(201):21–23, 1999.

[BW12] Anne Berry and Annegret Wagler. Triangulation and clique separator decomposition of claw-free graphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 7–21. Springer, 2012.

[Cam89] Kathie Cameron. Induced matchings. *Discrete Appl. Math.*, 24(1-3):97–102, 1989.

[Cao16] Yixin Cao. Linear recognition of almost interval graphs. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1096–1115. SIAM, 2016.

[CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *J. ACM*, 63(5):65, 2016. Id/No 40.

[CDH13a] Derek G. Corneil, Barnaby Dalton, and Michel Habib. LDFS-based certifying algorithm for the minimum path cover problem on cocomparability graphs. *SIAM J. Comput.*, 42(3):792–807, 2013.

[CDH13b] Derek G. Corneil, Barnaby Dalton, and Michel Habib. LDFS-based certifying algorithm for the minimum path cover problem on cocomparability graphs. *SIAM Journal on Computing*, 42(3):792–807, 2013.

[CFK+15] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*. Cham: Springer, 2015.

[Chu14] Maria Chudnovsky. The Erdős-Hajnal conjecture. A survey. *J. Graph Theory*, 75(2):178–190, 2014.

[Chu15] Julia Chuzhoy. Excluded grid theorem: improved and simplified. In *Proceedings of the 47th annual ACM symposium on theory of computing, STOC '15, Portland, OR, USA, June 14–17, 2015*, pages 645–654. New York, NY: Association for Computing Machinery (ACM), 2015.

[CK08] Derek G. Corneil and Richard M. Krueger. A unified view of graph searching. *SIAM Journal on Discrete Mathematics*, 22(4):1259–1276, 2008.

[CLMS22] Victor Campos, Raul Lopes, Ana Karolinna Maia, and Ignasi Sau. Adapting the directed grid theorem into an FPT algorithm. *SIAM J. Discrete Math.*, 36(3):1887–1917, 2022.

[Cor04a] Derek G. Corneil. A simple 3-sweep LBFS algorithm for the recognition of unit interval graphs. *Discrete Appl. Math.*, 138(3):371–379, 2004.

[Cor04b] Derek G. Corneil. Lexicographic breadth first search–a survey. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 1–19. Springer, 2004.

[COS97] Derek G. Corneil, Stephan Olariu, and Lorna Stewart. Asteroidal triple-free graphs. *SIAM Journal on Discrete Mathematics*, 10(3):399–430, 1997.

[CRS02] Vašek Chvátal, Irena Rusu, and R. Sritharan. Dirac-type characterizations of graphs without long chordless cycles. *Discrete Math.*, 256(1-2):445–448, 2002.

[CRST06] Maria Chudnovsky, Neil Robertson, Paul D. Seymour, and Robin Thomas. The strong perfect graph theorem. *Ann. Math. (2)*, 164(1):51–229, 2006.

[CSS17] Maria Chudnovsky, Alex Scott, and Paul D. Seymour. Induced subgraphs of graphs with large chromatic number. III: Long holes. *Combinatorica*, 37(6):1057–1072, 2017.

[CSSS20] Maria Chudnovsky, Alex Scott, Paul D. Seymour, and Sophie Spirkl. Induced subgraphs of graphs with large chromatic number. VIII. Long odd holes. *J. Comb. Theory, Ser. B*, 140:84–97, 2020.

[CST03] Kathie Cameron, R. Sritharan, and Yingwen Tang. Finding a maximum induced matching in weakly chordal graphs. *Discrete Math.*, 266(1-3):133–142, 2003.

[CT21] Julia Chuzhoy and Zihan Tan. Towards tight(er) bounds for the excluded grid theorem. *J. Comb. Theory, Ser. B*, 146:219–265, 2021.

[DF13] Rodney G. Downey and Michael R. Fellows. *Fundamentals of parameterized complexity*. London: Springer, 2013.

[dFdMdSOS21] Celina M. H. de Figueiredo, Alexsander Andrade de Melo, Fabiano de S. Oliveira, and Ana Silva. Maximum cut on interval graphs of interval count four is np-complete. In Filippo Bonchi and Simon J. Puglisi, editors, *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021, Tallinn, Estonia*, volume 202 of *LIPIcs*, pages 38:1–38:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[DGH+21] Clément Dallard, Robert Ganian, Meike Hatzel, Matjaž Krnc, and Martin Milanič. Graphs with two moplexes. *Procedia Computer Science*, 195:248–256, 2021. Proceedings of the XI Latin and American Algorithms, Graphs and Optimization Symposium.

[Die17] Reinhard Diestel. *Graphentheorie*. Berlin: Springer, 2017.

[Die18] Reinhard Diestel. Abstract separation systems. *Order*, 35(1):157–170, 2018.

[Dir61] Gabriel A. Dirac. On rigid circuit graphs. *Abh. Math. Semin. Univ. Hamb.*, 25:71–76, 1961.

[DJGT99] Reinhard Diestel, Tommy R. Jensen, Konstantin Yu. Gorbunov, and Carsten Thomassen. Highly connected sets and the excluded grid theorem. *J. Comb. Theory, Ser. B*, 75(1):61–73, 1999.

[DO21] Reinhard Diestel and Sang-il Oum. Tangle-tree duality in abstract separation systems. *Adv. Math.*, 377:25, 2021. Id/No 107470.

[DS94] Jitender S. Deogun and George Steiner. Polynomial algorithms for hamiltonian cycle in cocomparability graphs. *SIAM J. Comput.*, 23(3):520–552, 1994.

[Duc84] Pierre Duchet. Classical perfect graphs. An introduction with emphasis on triangulated and interval graphs. Perfect graphs, Ann. Discrete Math. 21, 67-96 (1984)., 1984.

[Ege31] Jenö Egerváry. Über kombinatorische Eigenschaften von Matrizen. *Mat. Fiz. Lapok*, 38:16–28, 1931.

[EH89] Paul Erdős and András Hajnal. Ramsey-type theorems. *Discrete Appl. Math.*, 25(1-2):37–52, 1989.

[EHPS05] Elaine M. Eschen, Chính T. Hoàng, Mark D. T. Petrick, and R. Sritharan. Disjoint clique cutsets in graphs without long holes. *J. Graph Theory*, 48(4):277–298, 2005.

[Erd20] Joshua Erde. Directed path-decompositions. *SIAM J. Discrete Math.*, 34(1):415–430, 2020.

[FGK11] Jiří Fiala, Petr A. Golovach, and Jan Kratochvíl. Parameterized complexity of coloring problems: treewidth versus vertex cover. *Theor. Comput. Sci.*, 412(23):2513–2523, 2011.

[FKM04] Fedor V. Fomin, Dieter Kratsch, and Haiko Müller. Algorithms for graphs with small octopus. *Discrete Applied Mathematics*, 134(1-3):105–128, 2004.

[Flo97] Carsten Flotow. Graphs whose powers are chordal and graphs whose powers are interval graphs. *J. Graph Theory*, 24(4):323–330, 1997.

[Gar07] Frédéric Gardi. The Roberts characterization of proper and unit interval graphs. *Discrete Mathematics*, 307(22):2906–2908, 2007.

[Gav74] Fanica Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *J. Comb. Theory, Ser. B*, 16:47–56, 1974.

[GHW19] Archontia C. Giannopoulou, Meike Hatzel, and Sebastian Wiederrecht. Braces of perfect matching width 2. *arXiv preprint arXiv:1902.06307*, 2019.

[GKK+20] Archontia C. Giannopoulou, Ken-ichi Kawarabayashi, Stephan Kreutzer, O-joung Kwon, et al. The canonical directed tree decomposition and its applications to the directed disjoint paths problem. *arXiv preprint arXiv:2009.13184*, 2020.

[GKKK20] Archontia C. Giannopoulou, Ken-ichi Kawarabayashi, Stephan Kreutzer, and O-Joung Kwon. The directed flat wall theorem. In *Proceedings of the 31st annual ACM-SIAM symposium on discrete algorithms, SODA 2020, Salt Lake City, UT, USA, January 5–8, 2020*, pages 239–258. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM); New York, NY: Association for Computing Machinery (ACM), 2020.

[GKMV21] Vladimir Gurvich, Matjaž Krnc, Martin Milanič, and Mikhail Vyalyi. Shifting paths to avoidable ones. *J. Graph Theory*, 100(1):69–83, 2021.

[GKW21] Archontia C. Giannopoulou, Stephan Kreutzer, and Sebastian Wiederrecht. Excluding a planar matching minor in bipartite graphs. *arXiv preprint arXiv:2106.00703*, 2021.

[GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics: Study and Research Texts*. Springer-Verlag, Berlin, 1988.

[Gol80] Martin Charles Golumbic. Algorithmic graph theory and perfect graphs. Computer Science and Applied Mathematics; New York etc.: Academic Press. A Subsidiary of Harcourt Brace Jovanovich, Publishers. XX, p. 284, 1980.

[GPv12] Petr A. Golovach, Daniël Paulusma, and Erik Jan van Leeuwen. Induced disjoint paths in AT-free graphs. In *Algorithm theory – SWAT 2012. 13th Scandinavian symposium and workshops, Helsinki, Finland, July 4–6, 2012. Proceedings*, pages 153–164. Berlin: Springer, 2012.

[GW21] Archontia C. Giannopoulou and Sebastian Wiederrecht. A flat wall theorem for matching minors in bipartite graphs, 2021.

[Gya87] András Gyarfás. Problems from the world surrounding perfect graphs. *Zastosow. Mat.*, 19(3-4):413–441, 1987.

[Heg06] Pinar Heggernes. Minimal triangulations of graphs: A survey. *Discrete Mathematics*, 306(3):297–317, 2006.

[HKK19] Meike Hatzel, Ken-ichi Kawarabayashi, and Stephan Kreutzer. Polynomial planar directed grid theorem. In *Proceedings of the 30th annual ACM-SIAM symposium on discrete algorithms, SODA 2019,* *San Diego, CA, USA, January 6–9, 2019*, pages 1465–1484. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM); New York, NY: Association for Computing Machinery (ACM), 2019.

[HRW19] Meike Hatzel, Roman Rabinovich, and Sebastian Wiederrecht. Cyclewidth and the grid theorem for perfect matching width of bipartite graphs. In *Graph-theoretic concepts in computer science. 45th international workshop, WG 2019, Vall de Núria, Spain, June 19–21, 2019. Revised papers*, pages 53–65. Cham: Springer, 2019.

[HW18] Meike Hatzel and Sebastian Wiederrecht. On perfect linegraph squares. In *Graph-theoretic concepts in computer science. 44th international workshop, WG 2018, Cottbus, Germany, June 27–29, 2018. Proceedings*, pages 252–265. Cham: Springer, 2018.

[JN01] Jeannette Janssen and Lata Narayanan. Approximation algorithms for channel assignment with constraints. *Theor. Comput. Sci.*, 262(1-2):649–667, 2001.

[JRST01] Thor Johnson, Neil Robertson, Paul D. Seymour, and Robin Thomas. Directed tree-width. *J. Comb. Theory, Ser. B*, 82(1):138–154, 2001.

[JST18] Jisu Jeong, Sigve Hortemo Sæther, and Jan Arne Telle. Maximum matching width: new characterizations and a fast algorithm for dominating set. *Discrete Appl. Math.*, 248:114–124, 2018.

[KK08] Ken-ichi Kawarabayashi and Yusuke Kobayashi. The induced disjoint paths problem. In *Integer programming and combinatorial optimization. 13th international conference, IPCO 2008 Bertinoro, Italy, May 26–28, 2008. Proceedings*, pages 47–61. Berlin: Springer, 2008.

[KK15] Ken-ichi Kawarabayashi and Stephan Kreutzer. The directed grid theorem. In *Proceedings of the 47th annual ACM symposium on theory of computing, STOC '15, Portland, OR, USA, June 14–17, 2015*, pages 655–664. New York, NY: Association for Computing Machinery (ACM), 2015.

[KKKK13] Ken-ichi Kawarabayashi, Daniel Král', Marek Krčál, and Stephan Kreutzer. Packing directed cycles through a specified vertex set. In *Proceedings of the 24th annual ACM-SIAM symposium on discrete algorithms, SODA 2013, New Orleans, LA, USA, January 6–8, 2013*, pages 365–377. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM); New York, NY: Association for Computing Machinery (ACM), 2013.

[KKM97] Ton Kloks, Dieter Kratsch, and Haiko Müller. Asteroidal sets in graphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 229–241. Springer, 1997.

[KKM01] Ton Kloks, Dieter Kratsch, and Haiko Müller. On the structure of graphs with bounded asteroidal number. *Graphs and Combinatorics*, 17(2):295–306, 2001.

[KKTW01] Daniel Král', Jan Kratochvíl, Zsolt Tuza, and Gerhard J. Woeginger. Complexity of coloring graphs without forbidden induced subgraphs. In *Graph-theoretic concepts in computer science (Boltenhagen, 2001)*, volume 2204 of *Lecture Notes in Comput. Sci.*, pages 254–262. Springer, Berlin, 2001.

[KM98] P. Sreenivasa Kumar and C. E. Veni Madhavan. Minimal vertex separators of chordal graphs. *Discrete Applied Mathematics*, 89(1-3):155–168, 1998.

[KM07] Ken-ichi Kawarabayashi and Bojan Mohar. Some recent progress and applications in graph minor theory. *Graphs Comb.*, 23(1):1–46, 2007.

[KM12] Dieter Kratsch and Haiko Müller. Colouring AT-free graphs. In *Algorithms—ESA 2012*, volume 7501 of *Lecture Notes in Comput. Sci.*, pages 707–718. Springer, Heidelberg, 2012.

[KM14] Ekkehard Köhler and Lalla Mouatadid. Linear time LexDFS on cocomparability graphs. In *Algorithm theory—SWAT 2014*, volume 8503 of *Lecture Notes in Comput. Sci.*, pages 319–330. Springer, Cham, 2014.

[KM16] Ekkehard Köhler and Lalla Mouatadid. A linear time algorithm to compute a maximum weighted independent set on cocomparability graphs. *Information Processing Letters*, 116(6):391–395, 2016.

[KMN20] Jan Kratochvíl, Tomáš Masařík, and Jana Novotná. U-bubble model for mixed unit interval graphs and its applications: The maxcut problem revisited. In *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

[KMT08] Dieter Kratsch, Haiko Müller, and Ioan Todinca. Feedback vertex set on AT-free graphs. *Discrete Applied Mathematics*, 156(10):1936–1947, 2008.

[KO11] Stephan Kreutzer and Sebastian Ordyniak. Digraph decompositions and monotonicity in digraph searching. *Theor. Comput. Sci.*, 412(35):4688–4703, 2011.

[KO14] Stephan Kreutzer and Sebastian Ordyniak. Width-measures for directed graphs and algorithmic applications. In Matthias Dehmer and Frank Emmert-Streib, editors, *Quantitative Graph Theory: Mathematical Foundations and Applications*, pages 181–232. Springer, Chapman and Hall/CRC, 1 edition, 2014.

[KS93] Dieter Kratsch and Lorna Stewart. Domination on cocomparability graphs. *SIAM Journal on Discrete Mathematics*, 6(3):400–417, 1993.

[KS94] C. W. Ko and F. Bruce Shepherd. Adding an identity to a totally unimodular matrix. *LSE Operations Research Working Paper LSEOR.94.12*, 1994.

[KTW20] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph. *arXiv preprint arXiv:2010.12397*, 2020.

[Kö99] Ekkehard Köhler. *Graphs without asteroidal triples*. PhD dissertation, Technische Universität Berlin, 1999.

[Kő31] Denes Kőnig. Graphen und Matrices. *Mat. Fiz. Lapok*, 38:119, 1931.

[LB62] Cornelis Lekkerkerker and Johan Boland. Representation of a finite graph by a set of intervals on the real line. *Fundamenta Mathematicae*, 51(1):45–64, 1962.

[LB79] George S. Lueker and Kellogg S. Booth. A linear time algorithm for deciding interval graph isomorphism. *Journal of the ACM (JACM)*, 26(2):183–195, 1979.

[LdM15] Cláudio L. Lucchesi, Marcelo H. de Carvalho, and U. S. R. Murty. Thin edges in braces. *Electron. J. Comb.*, 22(4):research paper p4.14, 38, 2015.

[Lit75] Charles H. C. Little. A characterization of convertible (0,1)-matrices. *J. Comb. Theory, Ser. B*, 18:187–208, 1975.

[LMW98] In-Jen Lin, Terry McKee, and Douglas B. West. The leafage of a chordal graph. *Discussiones Mathematicae Graph Theory*, 18(1):23–48, 1998.

[Lov72] László Lovász. Normal hypergraphs and the perfect graph conjecture. *Discrete Mathematics*, 2(3):253–267, 1972.

[Lov87] László Lovász. Matching structure and the matching lattice. *J. Comb. Theory, Ser. B*, 43:187–222, 1987.

[LP09] László Lovász and Michael D. Plummer. *Matching theory*. Providence, RI: AMS Chelsea Publishing, 2009.

[LR92] Errol L. Lloyd and Subramanian Ramanathan. On the complexity of distance-2 coloring. In *Computing and Information, 1992. Proceedings. ICCI'92., Fourth International Conference on*, pages 71–74. IEEE, 1992.

[Mah02] Mohammad Mahdian. On the computational complexity of strong edge coloring. *Discrete Appl. Math.*, 118(3):239–248, 2002.

[McC00] William McCuaig. Even dicycles. *J. Graph Theory*, 35(1):46–68, 2000.

[Mei05] Daniel Meister. Recognition and computation of minimal triangulations for AT-free claw-free and co-comparability graphs. *Discrete Applied Mathematics*, 146(3):193–218, 2005.

[Men27] Karl Menger. Zur Allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927.

[NKGB00] Thyagarajan Nandagopal, Tae-Eun Kim, Xia Gao, and Vaduvur Bharghavan. Achieving mac layer fairness in wireless packet networks. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 87–98. ACM, 2000.

[Nor05] Serguei Norine. *Matching Structure and Pfaffian Orientations of Graphs*. PhD dissertation, Georgia Institute of Technology, 2005.

[NT07] Serguei Norine and Robin Thomas. Generating bricks. *J. Comb. Theory, Ser. B*, 97(5):769–817, 2007.

[Obd06] Jan Obdržálek. DAG-width – connectivity measure for directed graphs. In *Proceedings of the seventeenth annual ACM-SIAM symposium on discrete algorithms, SODA 2006, Miami, FL, January 22–24, 2006*, pages 814–821. New York, NY: Association for Computing Machinery (ACM); Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2006.

[OCF76] Tatsuo Ohtsuki, Lap Kit Cheung, and Toshio Fujisawa. Minimal triangulation of a graph and optimal pivoting order in a sparse matrix. *Journal of Mathematical Analysis and Applications*, 54(3):622–633, 1976.

[Plu80] Michael D. Plummer. On n-extendable graphs. *Discrete Math.*, 31:201–210, 1980.

[Plu86] Michael D. Plummer. Matching extension in bipartite graphs. Combinatorics, graph theory, and computing, Proc. 17th Southeast. Conf., Boca Raton/Fl., Congr. Numerantium 54, 245–258., 1986.

[PS97] Andreas Parra and Petra Scheffler. Characterizations and algorithmic applications of chordal graph embeddings. *Discrete Applied Mathematics*, 79(1-3):171–188, 1997.

[PW10] Arnaud Pêcher and Annegret K. Wagler. Clique and chromatic number of circular-perfect graphs. In *ISCO 2010. International symposium on combinatorial optimization. Papers based on the presentations at the symposium, Hammamet, Tunesia, March 24–26, 2010.*, pages 199–206. Amsterdam: Elsevier, 2010.

[Ram97] Subramanian Ramanathan. A unified framework and algorithm for (T/F/C) DMA channel assignment in wireless networks. In *INFOCOM'97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Driving the Information Revolution., Proceedings IEEE*, volume 2, pages 900–907. IEEE, 1997.

[Ree99] Bruce Reed. Introducing directed tree width. In *6th Twente workshop on graphs and combinatorial optimization. Univ. of Twente, Enschede, Netherlands, May 26–28. Extended abstracts*, page no pag. Amsterdam: Elsevier, 1999.

[RL93] Subramanian Ramanathan and Errol L. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE/ACM Transactions on networking*, 1(2):166–177, 1993.

[Rob69] Fred S. Roberts. Indifference graphs. *Proof techniques in graph theory*, pages 139–146, 1969.

[Ros70] Donald J. Rose. Triangulated graphs and the elimination process. *J. Math. Anal. Appl.*, 32:597–609, 1970.

[RS84] Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *J. Comb. Theory, Ser. B*, 36:49–64, 1984.

[RS86] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41:92–114, 1986.

[RS91] Neil Robertson and Paul D. Seymour. Graph minors. X: Obstructions to tree-decomposition. *J. Comb. Theory, Ser. B*, 52(2):153–190, 1991.

[RS93] Neil Robertson and Paul D. Seymour. Excluding a graph with one crossing. In *Graph structure theory. Proceedings of the AMS-IMS-SIAM joint summer research conference on graph minors held June 22 to July 5, 1991 at the University of Washington, Seattle, WA (USA)*, pages 669–675. Providence, RI: American Mathematical Society, 1993.

[RS99] Neil Robertson and Paul D. Seymour. Graph minors. XVII: Taming a vortex. *J. Comb. Theory, Ser. B*, 77(1):162–210, 1999.

[RS04a] Bert Randerath and Ingo Schiermeyer. Vertex colouring and forbidden subgraphs – a survey. *Graphs Comb.*, 20(1):1–40, 2004.

[RS04b] Neil Robertson and Paul D. Seymour. Graph minors. XX: Wagner's conjecture. *J. Comb. Theory, Ser. B*, 92(2):325–357, 2004.

[RST94] Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *J. Comb. Theory, Ser. B*, 62(2):323–348, 1994.

[RST08] Ivan Rapaport, Karol Suchan, and Ioan Todinca. Minimal proper interval completions. *Information Processing Letters*, 106(5):195–202, 2008.

[RTL76] Donald J. Rose, R. Endre Tarjan, and George S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.*, 5:266–283, 1976.

[Sei74] Dieter Seinsche. On a property of the class of n-colorable graphs. *J. Comb. Theory, Ser. B*, 16:191–193, 1974.

[Shi88] Yukio Shibata. On the tree representation of chordal graphs. *Journal of Graph Theory*, 12(3):421–428, 1988.

[Sim91] Klaus Simon. A new simple linear algorithm to recognize interval graphs. In *Computational geometry - methods, algorithms and applications. International workshop, CG, Bern, Switzerland, March 21-22. Proceedings*, pages 289–308. Berlin etc.: Springer-Verlag, 1991.

[SR19] Ingo Schiermeyer and Bert Randerath. Polynomial χ-binding functions and forbidden induced subgraphs: a survey. *Graphs Comb.*, 35(1):1–31, 2019.

[SS16] Alex Scott and Paul D. Seymour. Induced subgraphs of graphs with large chromatic number. I. Odd holes. *J. Comb. Theory, Ser. B*, 121:68–84, 2016.

[ST90] A. Satyanarayana and L. Tung. A characterization of partial 3-trees. *Networks*, 20(3):299–322, 1990.

[ST93] Paul D. Seymour and Robin Thomas. Graph searching and a min-max theorem for tree-width. *J. Comb. Theory, Ser. B*, 58(1):22–33, 1993.

[ST09] Karol Suchan and Ioan Todinca. Minimal interval completion through graph exploration. *Theoretical computer science*, 410(1):35–43, 2009.

[SV82] Larry J. Stockmeyer and Vijay V. Vazirani. NP-completeness of some generalizations of the maximum matching problem. *Inf. Process. Lett.*, 15:14–19, 1982.

[SW18] Robert Scheidweiler and Sebastian Wiederrecht. On chordal graph and line graph squares. *Discrete Appl. Math.*, 243:239–247, 2018.

[Tho91] Carsten Thomassen. Highly connected non-2-linked digraphs. *Combinatorica*, 11(4):393–395, 1991.

[Tro16] Nicolas Trotignon. LexBFS, Structure, and Algorithms. In *Encyclopedia of Algorithms*, pages 1103–1107. Springer New York, 2016.

[UTN05] Ryuhei Uehara, Seinosuke Toda, and Takayuki Nagoya. Graph isomorphism completeness for chordal bipartite graphs and strongly chordal graphs. *Discrete Applied Mathematics*, 145(3):479–482, 2005.

[Vat12] Martin Vatshelle. *New width parameters of graphs*. PhD dissertation, The University of Bergen, 2012.

[Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Math. Ann.*, 114:570–590, 1937.

[Wal78] James R. Walter. Representations of chordal graphs as subtrees of a tree. *Journal of Graph Theory*, 2(3):265–267, 1978.

[Wes96] Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, Inc., Upper Saddle River, NJ, 1996.

[Wie20] Sebastian Wiederrecht. Digraphs of directed treewidth one. *Discrete Math.*, 343(12):9, 2020. Id/No 112124.

[XLL13] Shou-Jun Xu, Xianyue Li, and Ronghua Liang. Moplex orderings generated by the LexDFS algorithm. *Discrete Applied Mathematics*, 161(13-14):2189–2195, 2013.

[YHF14] Zhengwu Yang, Hong Huo, and Tao Fang. Ordering strategies for removing redundant fill edges. In *International Advance Computing Conference (IACC)*, pages 744–753. IEEE, 2014.

Index

Schriftenreihe **Foundations of computing**
Hrsg.: Prof. Dr. Stephan Kreutzer, Prof. Dr. Uwe Nestmann, Prof. Dr. Rolf Niedermeier

ISSN 2199-5249 (print)
ISSN 2199-5257 (online)

01: Bevern, René van: Fixed-Parameter Linear-Time Algorithms for NP-hard Graph and Hypergraph Problems Arising in Industrial Applications. - 2014. - 225 S.
ISBN **978-3-7983-2705-4** (print) EUR **12,00**
ISBN **978-3-7983-2706-1** (online)

02: Nichterlein, André: Degree-Constrained Editing of Small-Degree Graphs. - 2015. - xiv, 225 S.
ISBN **978-3-7983-2705-4** (print) EUR **12,00**
ISBN **978-3-7983-2706-1** (online)

03: Bredereck, Robert: Multivariate Complexity Analysis of Team Management Problems. - 2015. - xix, 228 S.
ISBN **978-3-7983-2764-1** (print) EUR **12,00**
ISBN **978-3-7983-2765-8** (online)

04: Talmon, Nimrod: Algorithmic Aspects of Manipulation and Anonymization in Social Choice and Social Networks. - 2016. - xiv, 275 S.
ISBN **978-3-7983-2804-4** (print) EUR **13,00**
ISBN **978-3-7983-2805-1** (online)

05: Siebertz, Sebastian: Nowhere Dense Classes of Graphs. Characterisations and Algorithmic Meta-Theorems. - 2016. - xxii, 149 S.
ISBN **978-3-7983-2818-1** (print) EUR **11,00**
ISBN **978-3-7983-2819-8** (online)

06: Chen, Jiehua: Exploiting Structure in Computationally Hard Voting Problems. - 2016. - xxi, 255 S.
ISBN **978-3-7983-2825-9** (print) EUR **13,00**
ISBN **978-3-7983-2826-6** (online)

07: Arbach, Youssef: On the Foundations of dynamic coalitions. Modeling changes and evolution of workflows in healthcare scenarios - 2016. - xv, 171 S.
ISBN **978-3-7983-2856-3** (print) EUR **12,00**
ISBN **978-3-7983-2857-0** (online)

08: Sorge, Manuel: Be sparse! Be dense! Be robust! Elements of parameterized algorithmmics. - 2017. - xvi, 251 S.
ISBN 978-3-7983-2885-3 (print) EUR **13,00**
ISBN 978-3-7983-2886-0 (online)

09: Dittmann, Christoph: Parity games, separations, and the modal μ-calculus. - 2017. - x, 274 S.
ISBN **978-3-7983-2887-7** (print) EUR **13,00**
ISBN **978-3-7983-2888-4** (online)

10: Karcher, David S.: Event Structures with Higher-Order Dynamics. - 2019. - xix, 125 S.
ISBN **978-3-7983-2995-9** (print) EUR **11,00**
ISBN **978-3-7983-2996-6** (online)

11: Jungnickel, Tim: On the Feasibility of Multi-Leader Replication in the Early Tiers. - 2018. - xiv, 177 S.
ISBN **978-3-7983-3001-6** (print) EUR **13,00**
ISBN **978-3-7983-3002-3** (online)

12: Froese, Vincent: Fine-grained complexity analysis of some combinatorial data science problems. - 2018. - xiv, 166 S.
ISBN **978-3-7983-3003-0** (print) EUR **11,00**
ISBN **978-3-7983-3004-7** (online)

13: Molter, Hendrik: Classic graph problems made temporal – a parameterized complexity analysis. - 2020. - xii, 206 S.
ISBN **978-3-7983-3172-3** (print) EUR **12,00**
ISBN **978-3-7983-3173-0** (online)

14: Bentert, Matthias: Elements of Dynamic and 2-SAT Programming: Paths, Trees, and Cuts. - 2021. - xiv, 199 S.
ISBN **978-3-7983-3209-6** (print) EUR **12,00**
ISBN **978-3-7983-3210-2** (online)

15: Kaczmarczyk, Andrzej: Algorithmic Aspects of Resource Allocation and Multiwinner Voting: Theory and Experiments. - 2021. - xii, 228 S.
ISBN **978-3-7983-3215-7** (print) EUR **13,00**
ISBN **978-3-7983-3216-4** (online)

16: Wiederrecht, Sebastian: Matching minors in bipartite graphs. - 2022. - XVII, 463 S.
ISBN **978-3-7983-3252-2** (print) EUR **25,00**
ISBN **978-3-7983-3253-9** (online)

Universitätsverlag der TU Berlin

Dualities in graphs and digraphs

This thesis explores dualities in directed and undirected graphs, utilising width-parameters, obstructions, and substructures. The focus is primarily on the structure of directed graphs. We introduce new width-measures based on directed separations, which relate to DAG-width. Our results include a duality theorem for a tangle-like obstruction. We also introduce cyclewidth as a new width measure, which is equivalent to directed treewidth. By leveraging the connection between directed graphs and bipartite graphs with perfect matchings, we characterise digraphs with low cyclewidth. In addition, we present a new flat wall theorem for directed graphs, which could potentially serve as a better foundation for a directed structure theorem. Regarding undirected graphs, we present several results on induced subgraphs in the graphs themselves or the square graph of their linegraph. Our findings range from general statements about all graphs to the consideration of specific graph classes such as those with exactly two moplexes.

ISBN 978-3-7983-3291-1 (print)
ISBN 978-3-7983-3292-8 (online)

ISBN 978-3-7983-3291-1

https://verlag.tu-berlin.de