

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/163771>

Copyright and reuse:


This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



A Parametric Approach to Hereditary Classes

by

Bogdan Alecu

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Mathematics Institute

June 2021

Contents

Acknowledgments	iv
Declarations	vi
Abstract	vii
Chapter 1 Introduction	1
Chapter 2 Preliminaries	3
2.1 A short disclaimer	3
2.2 Graphs	4
2.3 Partial orders	6
2.4 Hereditary classes	6
2.5 Graph parameters	7
2.6 Bipartite graphs	9
2.7 Cographs	10
2.8 Intersection graphs	11
2.9 Clique-width	11
2.10 A short glossary of notation	12
Chapter 3 Lettericity and geometric griddability	14
3.1 Letter graphs and griddability: preliminaries	14
3.1.1 Letter graphs and lettericity	14
3.1.2 Monotone and geometric griddability	18
3.1.3 Permutation graphs	27
3.1.4 Hypergraphs	28
3.2 Characterisation, recognition and an intriguing connection	29
3.2.1 Characterisation and recognition of 3-letter graphs	30
3.2.2 Geometric griddability versus bounded lettericity	38

3.3	Bounded lettericity implies geometric griddability	42
3.3.1	Locally ordered hypergraphs	43
3.3.2	Showing bounded lettericity implies geometric griddability . .	47
3.4	Further directions: characterising geometric griddability	55
3.4.1	Chain circuits	58
3.4.2	Local consistency: a LOH perspective	66
3.4.3	A hierarchy of structure: the missing pieces	69
Chapter 4 Functionality		81
4.1	Background and motivation	81
4.1.1	The (labelled) speed of hereditary classes	81
4.1.2	Implicit graph representations	83
4.1.3	Functionality: a very short history	86
4.2	Graph functionality	88
4.2.1	Graphs of small functionality	88
4.2.2	Graphs of large functionality	96
4.3	Further directions of research	99
Chapter 5 The micro-world of cographs		106
5.1	Well-quasi-order and beyond: a prelude to our analysis	106
5.2	A hierarchy of parameters	110
5.2.1	Results that are immediate, or follow quickly from known facts	112
5.2.2	Co-chromatic number	114
5.2.3	Lettericity	115
5.2.4	Boxicity	117
5.2.5	H -index	118
5.2.6	Achromatic number	120
5.2.7	Contiguity	121
5.3	Conclusions and further directions	124
Chapter 6 Linear clique-width of bi-cographs		127
6.1	Preliminaries and some basic results	127
6.1.1	Linear clique-width	127
6.1.2	Bi-cographs	130
6.2	Unboundedness of linear clique-width	133
6.3	Minimality and uniqueness	135

Chapter 7	Bipartite permutation graphs	140
7.1	Lettericity of bipartite permutation graphs	141
7.2	Universal graphs within the class of bipartite permutation graphs . .	143
Chapter 8	Quasi-chain graphs	149
8.1	The structure of quasi-chain graphs	150
8.2	Lettericity and wqo of quasi-chain graphs	154
Chapter 9	Miscellanea	161
9.1	Linear Ramsey numbers versus co-chromatic number	161
9.2	The girth of matched k -partite graphs	164
9.3	Of two trees and 2-trees	169

Acknowledgments

This research was generously funded by the EPSRC through a Doctoral Training Partnership. This thesis was typeset with L^AT_EX 2_ε.¹

To say that this thesis would not have been possible without Prof. Vadim Lozin would be an understatement. In addition to being an accomplished mathematician, my advisor is an exceptional mentor. I am deeply grateful to him for his thoughtful guidance; for his unwavering support; for his faith in me throughout my doctoral studies; and for taking me with him on a beautiful mathematical journey.

Of course, no journey is the same without the people we meet along the way. In that regard, I am particularly fortunate. I would like to extend my thanks to all of my co-authors, and especially to Dr Aistis Atminas and Dr Viktor Zamaraev, for their unconditional camaraderie, and for their invaluable advice in oh-so-many matters. In addition to them, I would like to thank my fellow PhD students, for many delightful discussions, mathematical and otherwise, and for their friendship and support. It would be impossible for me to enumerate all of them here, but I would like to mention in particular Alex Wendland, Yani Pehova and Anna Skorobogatova.

More generally, I wish to express my gratitude towards the people who have kindled my passion for mathematics throughout the years. This includes my high-school teachers, Alain Lorée and Mihaela Bădescu; my undergraduate tutor, Dr David Wood; my summer project supervisor, Dr Dmitriy Rumynin; my Master's thesis advisor, Prof. David Loeffler; and the staff of the Warwick mathematics de-

¹L^AT_EX 2_ε is an extension of L^AT_EX. L^AT_EX is a collection of macros for T_EX. T_EX is a trademark of the American Mathematical Society.

partment, who have created a stimulating and enjoyable working atmosphere. To this list, I would like to add Dr Igor Razgon and Prof. Oleg Pikhurko, my thesis examiners. In addition to providing helpful feedback, they have offered me inspiring discussions concerning further directions of research.

I am also grateful to the many other wonderful people in my life, without whom I wouldn't have made it this far. Once more, I cannot hope to make this list even remotely exhaustive, but I would like to mention my friends from undergrad, James S, Sher, Will, Tariq and Marvin, as well as my flatmates, James P, Andrea and Sarvesh.

I would also like to thank my family for their support, and for taking care of me: my parents, Corina and Mihai; my brother, Călin; my grandparents, Florica and Marin; and my stepparents, Conrad and Cami.

Finally, from the bottom of my heart, I wish to thank Judith, my partner and best friend, for her love and encouragement, and for making every day brighter.

Declarations

This thesis is my own work, conducted under the supervision of my doctoral advisor, Prof. Vadim Lozin. It has not been submitted for another degree. In addition to this:

- Section 3.2 is based on [Ale+20b], which is joint work with Vadim Lozin, Dominique de Werra and Viktor Zamaraev.
- Section 4.2 is based on [AAL19], which is joint work with Aistis Atminas and Vadim Lozin; Theorem 105 will appear in [Ale+21a], which is joint work with the same two authors, and additionally with V. E. Alekseev and V. Zamaraev.
- Chapter 5 is based on [ALW21], which is joint work with Vadim Lozin and Dominique de Werra.
- Chapter 6 is based on [Ale+20a], which is joint work with Mamadou Kanté, Vadim Lozin and Viktor Zamaraev.
- Chapter 7 is based on [ALM21], which is joint work with Vadim Lozin and Dmitriy Malyshev.
- Chapter 8 is based on [Ale+21b], which is joint work with Aistis Atminas, Vadim Lozin and Dmitriy Malyshev.
- Section 9.1 is based on a result which appeared in [Ale+21c], which is joint work with Aistis Atminas, Vadim Lozin and Viktor Zamaraev.
- Section 9.3 is based on joint work with Viktor Zamaraev.

Abstract

The “minimal class approach” consists of studying downwards-closed properties of hereditary graph classes (such as boundedness of a certain parameter within the class) by identifying the minimal obstructions to those properties. In this thesis, we look at various hereditary classes through this lens. In practice, this often amounts to analysing the structure of those classes by characterising boundedness of certain graph parameters within them. However, there is more to it than this: while adopting the minimal class viewpoint, we encounter a variety of interesting notions and problems – some more loosely related to the approach than others. The thesis compiles the author’s work in the ensuing research directions.

Chapter 1

Introduction

A thesis is often written like a novel: it starts with a (hopefully) gripping exposition, then the reader is guided through a (hopefully) exciting narrative, and at the end, they are offered a (hopefully) satisfying resolution. To manage the reader's expectation, we warn them right away that the current work does not completely abide by that formula. Perhaps a more apt analogy for it would be a collection of short stories. The common setting of those stories is the world of *structural graph theory*.

Graph theory as a whole is a particularly active area of mathematics. The discipline undoubtedly owes some of its popularity to its many applications, but this does not stop it from exhibiting large amounts of beautiful theory. Broadly speaking, its structural branch attempts to express, characterise, and more generally understand various graph phenomena in terms of the configurations that occur in graphs. As a tongue-in-cheek example, a structural graph theorist might characterise trees as acyclic connected graphs, and remark that they always have a vertex of degree 1. In contrast, an extremal graph theorist might be interested in the fact that any graph with n vertices and at least n edges must contain a cycle; an enumerative combinatorist will derive Cayley's formula, which states that there are n^{n-2} labelled trees on n vertices; an algorithmic graph theorist will note that trees make for especially nice data structures; an algebraic one will point out that the Tutte polynomials of forests are particularly simple; and the list could go on.

This is, of course, a parodic oversimplification. In actuality, all of those fields, and several others, overlap significantly. As a consequence, the boundaries between the various areas inevitably become blurry. This thesis is no exception; if the author had to put labels on it, he might say that our study will feature algorithmic, enumerative, and occasionally extremal and algebraic aspects. Several

times, we will even venture outside of graph theory, and into the neighbouring world of permutation patterns.

In spite of (or, perhaps, elaborating on) our short story analogy, the reader must not think that the work is disjointed. Indeed, it is tightly bound together by a few common themes, and by many “recurring characters” – notions which make their appearance across several of our short stories. Because of this, we do not give an in-depth discussion of the literature and background at this stage. Instead, we will address these matters in the individual chapters where they are relevant. Nevertheless, in the next chapter, we will give a quick overview of the common themes permeating the work, together with a “character list” in which we introduce some of the recurring notions. But first, let us go through the outline of the thesis.

Chapter 2 lays out the common preliminaries.

Chapter 3 concerns a graph parameter called *lettericity*, originally introduced in [Pet02] to investigate well-quasi-orderability by induced subgraphs. Its study leads us into the world of permutations, where we reveal some intricate structural connections between lettericity and a notion called *geometric gridability* [Alb+13].

Chapter 4 introduces a new graph parameter that we call *functionality*. Its original motivation is enumerative, but its study leads to many questions that are interesting in their own right.

In each of Chapters 5 to 8, we restrict ourselves to individual classes, and investigate various structural and parametric problems within those classes. More specifically:

Chapter 5 is devoted to the class of cographs, for which we exhibit an intriguing hierarchy of graph parameters.

Chapter 6 concerns the bipartite analogues of cographs, for which we investigate the boundedness of linear clique-width.

In Chapter 7, we examine the class of bipartite permutation graphs, and produce some results on lettericity within the class, as well as bounds on the size of certain universal constructions.

In Chapter 8, we look at the class of so-called quasi-chain graphs, which is an extension of the well-studied class of bipartite chain graphs. We provide a structural characterisation of quasi-chain graphs, and once more, we study lettericity within the class.

Finally, Chapter 9 compiles together a few loose results which are interesting enough to be mentioned, but perhaps not consequential enough to deserve their own chapters.

Chapter 2

Preliminaries

In this chapter, we introduce the basic notation and terminology that we will (strive to) consistently use throughout the thesis. Among other things, we also present some of the “recurring characters” we mentioned, as well as the main concepts that will guide our study. Since most of what we discuss here appears in many introductory courses, textbooks or in the combinatorial folklore, we assume the reader will have some familiarity with the notions. As such, we will not dwell too long on the explanations, and we will omit illustrations and examples most of the time. Without further ado, let us begin.

2.1 A short disclaimer

Our notation will be as standard as possible. In particular, we believe it is safe for the reader to assume that any piece of notation which is not explicitly defined means exactly what they think it means. Nevertheless, we provide here a very short list of symbols which have different variants in the literature, and thus may present some ambiguity; we hope that we did not miss anything essential from this list.

The symbol “ $:=$ ” means “is defined to be”.

$\mathbb{N} := \{0, 1, 2, \dots\}$.

For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$ (and in particular, $[0] := \emptyset$).

The subset relation is denoted by “ \subseteq ”, and the proper subset relation is denoted by “ \subsetneq ”.

The disjoint union of two sets is denoted by \uplus .

Given sets A and B , the difference of A and B , denoted $A \setminus B$ or $A - B$, consists of all elements of A which are not elements of B .

Given a set A , we denote by A^k the set of ordered k -tuples of elements from A . We denote by $\binom{A}{k}$ the set of (unordered) k -subsets of A (that is, subsets of A with k elements).

Given two functions $f, g : \mathbb{N} \rightarrow \mathbb{R}$, we write $f = O(g)$ if there exists $C > 0$ and $N \in \mathbb{N}$ such that, for all $n \geq N$, $f(n) \leq Cg(n)$. We write $f = \Omega(g)$ if $g = O(f)$, and $f = \Theta(g)$ if $f = O(g)$ and $f = \Omega(g)$.

Given two functions $f, g : \mathbb{N} \rightarrow \mathbb{R}$, we write $f = o(g)$ if $f(n)/g(n) \rightarrow 0$ as $n \rightarrow \infty$, and $f \sim g$ if $f(n)/g(n) \rightarrow 1$ as $n \rightarrow \infty$.

2.2 Graphs

We will use very standard graph terminology, and avoid giving an exhaustive list of definitions for all the terms – the reader is invited to consult, for instance, Diestel’s book [Die00] for more details. Let us briefly summarise our particular implementation of the terminology.

An *undirected graph* G is a pair (V, E) , where V is a set whose elements we call *vertices*, and $E \subseteq \binom{V}{2}$ is a set whose elements we call *edges*. We note that this definition does not allow loops or multiple edges. Such graphs are also referred to as *simple* graphs. Unless otherwise specified, all graphs will be simple.

We say two vertices u and v are *adjacent*, or *neighbours*, if $\{u, v\} \in E$. We will often denote edges by simply concatenating the vertices they contain (that is, we may write uv instead of $\{u, v\}$). If $u \in e$ for a vertex u and edge e , we say u and e are *incident*. The *neighbourhood* of a vertex v in G is the set of neighbours of v , denoted $N_G(v)$ (the subscript is omitted when there is no ambiguity). The *complement* of a graph $G = (V, E)$ is the graph $\overline{G} := (V, \binom{V}{2} - E)$. Given two graphs G and H , their *disjoint union* $G + H$ is the graph with vertex set $V(G) \uplus V(H)$ and edge set $E(G) \uplus E(H)$. Their *join* $G \times H$ is the graph $\overline{\overline{G}} + \overline{\overline{H}}$ (that is, the graph with vertex set $V(G) \uplus V(H)$, and edge set $E(G) \uplus E(H) \uplus (V(G) \times V(H))$). The disjoint union of n copies of G is denoted by nG .

A *directed graph* (or *digraph* for short) consists, like in the undirected case, of a set V of vertices and a set E of edges, but this time, we require $E \subseteq V^2$ (that is, the edges are *ordered* pairs of vertices). Since we will sometimes work with undirected graphs and some auxiliary directed graphs simultaneously, we will call vertices “*nodes*” and edges “*arcs*” in the directed setting.

We will almost exclusively work with finite graphs. We will usually work with *unlabelled graphs*, in which vertices are only distinguished by the way they connect to one another. This is in contrast with *labelled graphs*, in which vertices come with labels (usually unique and from the set $[|V|]$).

Let G and H be two graphs. A function $\varphi : V(G) \rightarrow V(H)$ is a *graph isomorphism* if it is a bijection which preserves adjacency, in the sense that vertices $u, v \in V(G)$ are adjacent if and only if vertices $\varphi(u), \varphi(v) \in V(H)$ are adjacent. If such an isomorphism exists, G and H are said to be *isomorphic*. We will avoid discussing foundational issues, and abuse terminology by saying “graph” when we mean “isomorphism class of graphs” (for instance, we will talk about *the* (unique) complete graph on n vertices).

There are several natural order relations defined on graphs. The one we are most concerned with is the *induced subgraph* relation: for any subset $X \subseteq V(G)$, the *subgraph of G induced by X* , denoted $G[X]$, is the graph $(X, E(G) \cap \binom{X}{2})$. For a vertex $v \in V(G)$, we write $G - v := G[V(G) \setminus \{v\}]$. A graph H is an *induced subgraph* of G , written $H \leq_i G$, or simply $H \leq G$, if it is (isomorphic to) $G[X]$ for some $X \subseteq V(G)$. Put differently, H is induced in G if it can be obtained from G by deleting a set of vertices, together with all incident edges. In this case, we also say G *contains* H (as an induced subgraph).

There are of course other notions of containment: *subgraphs*, where in addition to vertex deletion, we also allow edge deletion; *spanning subgraphs*, where we only allow edge deletion; and *minors*, where we allow vertex deletion, edge deletion, and *edge contraction* (in which two adjacent vertices u and v are replaced with a single vertex w whose neighbourhood is the union $N(u) \cup N(v) \setminus \{u, v\}$).

A *clique* in G is a subset of vertices that are pairwise adjacent, and an *independent set*, or *stable set*, is a subset of vertices that are pairwise non-adjacent (that is, an independent set is a clique in \overline{G}). Cliques and independent sets together are sometimes referred to as *homogeneous sets*. Given a set $A \subseteq V(G)$ and a vertex $v \notin A$, we say v *dominates* A if v is adjacent to every vertex in A . A *dominating* vertex is a vertex v that dominates $V(G) \setminus \{v\}$. An *isolated* vertex is a vertex which is dominating in \overline{G} . Given two disjoint sets $A, B \subseteq V(G)$, we say A and B are *complete* to each other if every vertex in A dominates B (and vice-versa). We say A and B are *anticomplete* to each other if they are complete to each other in \overline{G} .

2.3 Partial orders

A *partial order* on a set X is a binary relation that is reflexive, antisymmetric and transitive. A set X together with a partial order \leq on X is called a *poset*. We write “ $x \leq y$ ” rather than “ $(x, y) \in \leq$ ”. A partial order \leq on X is said to be *total*, or *linear*, if any two elements of X are comparable (that is, if for any $x, y \in X$, we have either $x \leq y$ or $y \leq x$). Given a poset (X, \leq) , a *chain* is a set of pairwise comparable elements (that is, a subset of X totally ordered by \leq). An *ascending*, respectively *descending* chain is a (finite or infinite) sequence x_1, x_2, \dots with $x_1 \leq x_2 \leq \dots$, respectively $x_1 \geq x_2 \geq \dots$. An *antichain* is a set of pairwise incomparable elements.

A poset (X, \leq) is well-founded if it contains no infinite strictly descending chain. It is well-quasi-ordered (“wqo” for short) if it is well-founded, and it contains no infinite antichains. We note that all of the graph containment notions described above are well-founded, but not necessarily wqo (only the minor relation is wqo: this is the statement of Robertson and Seymour’s famous graph minor theorem [RS04]).

One can equivalently define well-quasi-orderings as follows:¹ an infinite sequence x_1, x_2, \dots in (X, \leq) is *good* if there exists $i < j$ with $x_i \leq x_j$ (the first “ $<$ ” denotes the usual order on \mathbb{N}). (X, \leq) is wqo if and only if every infinite sequence is good.

2.4 Hereditary classes

A graph *class*, graph *property*, or graph *family* is a collection of graphs closed under isomorphisms. Once more, we will avoid foundational issues, and identify isomorphic graphs, so that our “classes” can be treated like sets. A class is said to be *hereditary* if, in addition, it is closed under taking induced subgraphs (“downwards-closed”). Unless otherwise specified, every class we work with will be assumed to be hereditary.

An important principle in combinatorics is that downwards-closed classes of objects can be characterised by *minimal obstructions*.² Graphs are, of course, no exception: for a (hereditary) class \mathcal{X} of graphs, let $\text{Obs}(\mathcal{X})$ be the (possibly infinite) set of minimal graphs that are not in \mathcal{X} . In other words, any graph $G \in \text{Obs}(\mathcal{X})$ satisfies $G \notin \mathcal{X}$ and $G - v \in \mathcal{X}$ for any $v \in V(G)$; $\text{Obs}(\mathcal{X})$ consists of all graphs with this property. The set $\text{Obs}(\mathcal{X})$ is also called the set of *minimal forbidden induced subgraphs* for \mathcal{X} , sometimes also denoted by $\text{Forb}(\mathcal{X})$. Then, by standard theory, a

¹This equivalence and a number of others are more or less folklore; one of their early appearances in writing is in [Hig52].

²All we need to assume is well-foundedness of the relation – see Subsection 2.3

graph is in \mathcal{X} if and only if it does not contain any of the graphs from $\text{Obs}(\mathcal{X})$ (as induced subgraphs). We note that, directly from the minimality condition, $\text{Obs}(\mathcal{X})$ is an antichain.

Since every class admits a characterisation by minimal forbidden induced subgraphs, it can be useful to define classes via their sets of obstructions. To do this, for any set S of graphs, we write $\text{Free}(S)$ for the class of graphs that do not contain any graph in S as an induced subgraph. We also say a graph is S -free to mean that it is in $\text{Free}(S)$. We may omit the set brackets, or replace them with parentheses, when S consists of a small number of graphs (to produce expressions like “ $\text{Free}(G)$ ”, or “ (G, H) -free graphs”). As discussed above, for any class \mathcal{X} , we have $\mathcal{X} = \text{Free}(\text{Obs}(\mathcal{X}))$. We also have $S \supseteq \text{Obs}(\text{Free}(S))$ for any set S , with equality when S is an antichain.

Another way to specify classes of graphs is via the notion of *hereditary closure*: the hereditary closure of a set (or sequence) A of graphs consists of all graphs that are induced in some graph $G \in A$. Alternatively, the hereditary closure of a set A is the (inclusion-wise) smallest hereditary class containing all members of A .

A graph G is called *n -universal* for a class \mathcal{X} if G contains all n -vertex graphs from \mathcal{X} . It is called *proper* if, in addition, $G \in \mathcal{X}$. A *universal sequence* for \mathcal{X} can be defined in a number of ways; the most standard is, perhaps, as a sequence X_1, X_2, \dots , where X_n is n -universal for \mathcal{X} . Alternatively, it may simply be defined as a family X_1, X_2, \dots such that \mathcal{X} is contained in the hereditary closure of the X_i . The distinction will not matter to us, so we will use the term loosely to mean either of those things, as appropriate. The sequence is called *proper* if every member belongs to \mathcal{X} (in which case \mathcal{X} must be equal to the hereditary closure of the sequence).

2.5 Graph parameters

A *graph parameter* is a function that assigns to each graph a real number (in practice, the parameters we work with will often be natural number-valued). A parameter κ is *hereditary* if it does not increase when taking induced subgraphs (in other words, $G \leq H \implies \kappa(G) \leq \kappa(H)$). Like with classes, we will assume our parameters are hereditary, unless otherwise stated. Given a class \mathcal{X} and a parameter κ , we say κ is *bounded* in \mathcal{X} (or \mathcal{X} is *of bounded κ*) if there is a constant C such that for all $G \in \mathcal{X}$, $\kappa(G) \leq C$. If no such constant exists, we say κ is *unbounded* in \mathcal{X} .

Many of the problems that we study in this thesis are of the form “When is κ bounded in \mathcal{X} ?” for a specific parameter κ , and for a class \mathcal{X} belonging to a given

collection of classes. This collection is often going to be the set of subclasses of a fixed class that we will refer to as our *universe*.

A very fruitful approach to answering such questions is the so-called “minimal class approach”. Note that the set of classes in which a parameter κ is bounded is downwards-closed under inclusion. The idea is to then attempt to characterise boundedness of κ by producing a set $M(\kappa)$ of minimal classes in which κ is not bounded. Ideally, by analogy with the sets of obstructions for individual classes, we would then like to conclude that κ is bounded in a class from our universe if and only if it does not contain one of the “minimal obstacles” in $M(\kappa)$. This might not work in general, since the poset of classes we are considering might not be well-founded under inclusion.³ Nevertheless, in the many cases where it does work, the approach yields beautiful results, and valuable insight into the relevant problems.

It is worth noting that looking for minimal obstacles to the boundedness of parameters can be viewed as a Ramsey-type problem. Indeed, recall the special case of Ramsey’s classical theorem, which asserts the ubiquity of cliques and independent sets:

Theorem 1 (Ramsey’s theorem [Ram30]). *For any $p, q \in \mathbb{N}$, there exists $R = R(p, q) \in \mathbb{N}$ such that every graph on at least R vertices contains either a clique of size p , or an independent set of size q .*

This theorem can be restated as follows:

Theorem 2 (Ramsey’s theorem, minimal classes version). *Let $\kappa(G) := |V(G)|$ be the order of a graph. Then κ is bounded in a class \mathcal{X} if and only if \mathcal{X} does not contain the class \mathcal{K} of all cliques, or the class $\overline{\mathcal{K}}$ of all edgeless graphs.*

More generally, questions about boundedness of a parameter κ can be restated as: “What unavoidable structures arise when κ is large?”. As such, parametric questions can often be viewed as structural ones, and vice-versa.

We would also like to note that this point of view is not limited to examining the boundedness of parameters. Indeed, the same approach applies to any situation in which we are aiming to characterise downwards-closed sets of classes, such as classes in which a problem is not NP-hard, or classes which admit a universal construction of a certain size.

³In fact, it is another folklore result that the set of subclasses of a given class \mathcal{X} is well-founded under inclusion if and only if \mathcal{X} is wqo under induced subgraphs.

2.6 Bipartite graphs

A graph is called *k-partite* ($k \in \mathbb{N}$) if its vertex set can be partitioned into k independent sets (called *parts* in this context). The case where $k = 2$ is particularly interesting: those graphs are called *bipartite*. The same graph may admit several different partitions into two independent sets (so-called *bipartitions*) – in fact, a bipartite graph has a unique bipartition if and only if it is connected. To address this, bipartite graphs are often given with a fixed bipartition. In such cases, we write $G = (A, B, E)$ to emphasise that A and B are the two parts under consideration. In particular, $V(G) = A \uplus B$, and it is understood that any edge connects an element of A with an element of B .

When working with graphs with a distinguished bipartition, it is useful to think of them as *coloured* bipartite graphs. Given a graph (A, B, E) , we will call the vertices in A *black* and the vertices in B *white*. To make everything work smoothly in this setting, we need to tweak a bit our definition of isomorphism: two coloured bipartite graphs are isomorphic if there is an adjacency *and* colour-preserving bijection between them. As a consequence, other definitions need to be modified appropriately: for instance, induced subgraphs need to admit a colour-preserving embedding. We omit further details.

Keeping track of bipartitions in this way allows us to unambiguously define bipartite analogues of graph operations. More specifically, given two (coloured) bipartite graphs $G_1 = (A_1, B_1, E_1)$ and $G_2 = (A_2, B_2, E_2)$, their *bipartite disjoint union* is the graph $G_1 +_b G_2 := (A_1 \uplus A_2, B_1 \uplus B_2, E_1 \uplus E_2)$ (when there is no possibility of confusion, we may drop the $_b$ subscript). The *bipartite complement* of G is the graph $\widetilde{G} := (A, B, \{\{a, b\} : a \in A, b \in B\} \setminus E)$. The *bipartite join* of G_1 and G_2 is the graph $G_1 \times_b G_2 := \widetilde{G_1} +_b \widetilde{G_2}$. Finally, the *skew-join* of G_1 with G_2 is the graph $G_1 \rtimes G_2 := (A_1 \uplus A_2, B_1 \uplus B_2, E_1 \uplus E_2 \uplus \{\{a, b\} : a \in A_1, b \in B_2\})$.

We may also define a reflection operation on coloured bipartite graphs, which swaps the colours: the *reflection* of $G = (A, B, E)$ is the graph $G^R := (B, A, E)$. Additionally, from any coloured bipartite graph, we may obtain an *uncoloured* bipartite graph by “forgetting” the colouring – this is simply the graph $G = (A \cup B, E)$.

Also recall the standard result which characterises bipartite graphs as those graphs which do not contain any odd cycles.

2.7 Cographs

There are many tools that can be used to capture various features of the structure of graphs. One such tool is *modular decomposition*. Let G be a graph, and $u, v \in V(G)$. We say a vertex $w \neq u, v$ *distinguishes* u and v if it is adjacent to exactly one of u and v . A *module* of G is a subset A of $V(G)$ such that no vertex in $V(G) \setminus A$ distinguishes two vertices from A . As such, modules generalise connected components (and co-connected components, i.e., connected components of \overline{G}). \emptyset , singletons and $V(G)$ are *trivial* modules. A graph in which those are the only modules is called *prime*.

One can recursively decompose any graph into non-trivial modules. As is always the case with decompositions, one can record this information in a rooted tree: the *modular decomposition tree*. Modular decomposition has many structural and algorithmic uses, but we will never use it in its full generality, so we omit the details.

There is a class of graphs with particularly nice modular decompositions, called *cographs* (short for “complement reducible graphs”). This class can be defined as the smallest class containing the one vertex graph, and closed under disjoint unions and complements. Cographs have been rediscovered several times. They have been extensively studied, and admit many known characterisations (see, e.g., [BLS99] and the references therein). In particular, a graph G is a cograph if and only if it satisfies any of the following equivalent conditions:

- G is P_4 -free (where P_4 is the path on four vertices).
- For any induced subgraph $H \leq G$, either H or its complement is disconnected.
- For any induced subgraph $H \leq G$ with $|V(H)| \geq 2$, H contains a module of size 2 (the vertices in such a module are called *twins*; the literature sometimes distinguishes between *true* and *false* twins according to whether they are adjacent, respectively non-adjacent).
- G has clique-width at most 2.⁴

These properties imply that cographs have very nice properties with respect to modular decomposition, and those properties are captured in their representation by *cotrees*. A cotree is a rooted tree whose internal nodes (that is, the non-leaf nodes) are labelled with the numbers 0 and 1 (or equivalently, with “+” and “×”). From a cotree T , one can construct a cograph G_T as follows: the leaves of T correspond

⁴See Subsection 2.9 for a definition.

to single vertices. Inductively, for any node x of T with children y_1, \dots, y_s , the subtree T^x rooted at x corresponds to the graph G_{T^x} obtained by taking the disjoint union $\sum_i G_{T^{y_i}}$ or the join $\prod_i G_{T^{y_i}}$ according to whether x is labelled $0/+$ or $1/\times$ respectively.

Indeed, it is clear that every graph encoded by a cotree is a cograph, and it is not difficult to see that every cograph has a cotree representation. This follows from the fact that, for any cograph G , exactly one of G and its complement \overline{G} is connected. As such, even though cotree representations are in general not unique, we will abuse terminology and say *the* cotree T_G of a cograph G to refer to a special cotree. This cotree is obtained from G by starting with a root vertex labelled 0 or 1 according to whether G is disconnected or connected; the children of this root then recursively correspond to the connected or co-connected components of G respectively.⁵

2.8 Intersection graphs

Many classes of graphs are defined as the *intersection graphs* of other classes of objects. For example, an *interval graph* is a graph whose vertices correspond to intervals on the real line, and for which two vertices are adjacent whenever the corresponding intervals intersect. Another example is given by *line graphs*: given a graph G , its line graph $L(G)$ is the graph whose vertices are the edges of G , and for which two vertices are adjacent whenever the corresponding edges intersect. Given an intersection graph G , an *intersection model* for G is simply a set of the appropriate objects whose intersection graph is isomorphic to G .

One class of intersection graphs that we are particularly interested in is the class of *permutation graphs*. They are the intersection graphs of line segments between two parallel lines (although perhaps a more natural way to define them is as the *inversion graphs* of permutations – see Chapter 3).

2.9 Clique-width

Clique-width is a graph parameter whose boundedness has nice algorithmic consequences. While we will only use its linearised version (that we will define in Chapter 6), we briefly discuss clique-width here, since it is useful background knowledge.

The clique-width $\text{cw}(G)$ of a graph G is the smallest number of labels needed to construct G by means of the following operations:

1. Create a vertex v with label i ;

⁵This is effectively a special case of the modular decomposition tree.

2. Take the disjoint union of labelled graphs G and H ;
3. Add all edges between vertices labelled i and vertices labelled j for labels $i \neq j$;
4. Rename label i to label j .

Using those operations, one can produce *clique-width expressions* that can be represented by trees whose leaves correspond to vertex creation operations, and whose internal vertices correspond to one of the other operations. One then constructs graphs from those expressions in the obvious way.

Intuitively, clique-width can be obtained by minimising the number of labels needed to construct a graph G over all such expressions; the linearised version, *linear clique-width*, is what we obtain when we only allow expressions whose tree is “path-like”.

We note that there is a host of other width parameters in the literature, such as treewidth, branch-width, rank-width, and several others. We will occasionally use facts about those parameters (that we will state as we go along), but we will never need to work with their definitions, so we omit them.

2.10 A short glossary of notation

Some common graphs:

K_n – the *complete graph* $([n], \binom{[n]}{2})$.

C_n – the *cycle* $([n], \{12, 23, \dots, (n-1)n, n1\})$.

P_n – the *path* $([n], \{12, 23, \dots, (n-1)n\})$.

$K_{m,n}$ – the *complete bipartite graph* $\overline{K_m} \times \overline{K_n}$.

M_n – the *induced matching* nK_2 .

Sun_n – the graph obtained from C_n by adding a pendant vertex⁶ to each vertex.

H_n – the graph obtained from P_{n+1} by adding two pendant vertices to each end of the path.

$S_{k,l,m}$ – the tree with no vertices of degree 4 or more, a single vertex v of degree 3, and three leaves at distances k, l and m from v respectively.

⁶We say v is *pendant* to w if $N(v) = w$.

Some parameters:

χ – chromatic number (the smallest k such that $V(G)$ admits a partition into k independent sets).

z – co-chromatic number (the smallest k such that $V(G)$ admits a partition into k homogeneous sets).

ω – clique number (the size of the largest clique).

α – independence number (the size of the largest independent set).

cw – clique-width.

tw – treewidth.

Chapter 3

Lettericity and geometric griddability

In this chapter, we will study two notions – letter graph representations [Pet02], from the world of graphs, and geometric grid classes [Alb+13], from the world of permutations. We will see that, despite being defined in seemingly unrelated ways, the two notions have a close relationship between them. This stems from the fact that they capture the same structural data of their respective combinatorial objects: a partition of their elements into simple “bags”, and a linear ordering of the elements interacting nicely with this partition.

In Section 3.1, we introduce the necessary terminology and notation. Section 3.2 presents the results published in our paper [Ale+20b]. We prove a conjecture from that paper in Section 3.3, and finally, Section 3.4 provides a discussion about a further direction of research.

3.1 Letter graphs and griddability: preliminaries

3.1.1 Letter graphs and lettericity

The notion of letter graphs was introduced in [Pet02]. Our terminology differs only superficially from the one used there. We will need some basic notions from the theory of formal languages; rather than defining them separately in the most general setting possible, we will introduce the relevant definitions as we go along, and adapt them to our restricted setting.

Our starting point is a finite digraph $\mathcal{D} = (\Sigma, A)$ that we will call the *decoding digraph*, or simply *decoder*. We call Σ a (finite) *alphabet*, and we refer to its elements (that is, the vertices of \mathcal{D}) as *letters* (or *symbols*). Now let Σ^* be the set of finite

sequences of elements of Σ . We will refer to them as *words* (or *strings*) over Σ . The main idea is now to construct graphs from words in Σ^* by “decoding” them using \mathcal{D} . The intuition is that each of the indices $1, \dots, n$ of the word $w = w_1w_2 \dots w_n$ corresponds to a vertex, and their adjacency depends (in a straightforward way dictated by the arcs in \mathcal{D}) only on the relative order of the indices and on the symbols appearing at those indices. Formally, we have the following definition:

Definition 3. Let $\mathcal{D} = (\Sigma, A)$ be a decoder, and let $w = w_1w_2 \dots w_n \in \Sigma^*$. The *letter graph* $G(\mathcal{D}, w)$ is the finite simple graph defined by

- $V(G(\mathcal{D}, w)) = [n]$;
- $E(G(\mathcal{D}, w)) = \{\{i, j\} : (w_{\min(i,j)}, w_{\max(i,j)}) \in A\}$.

The map sending w to $G(\mathcal{D}, w)$ is called the *decoding map*.

Some examples are in order.

Example 4. In Figure 8.1, we show on the left a decoding digraph \mathcal{D} , and on the right the letter graph $G = G(\mathcal{D}, acdbad)$. Notice how, for each letter $l \in V(\mathcal{D})$, the set $\{i : w_i = l\}$ forms either a clique or an independent set, according to whether the loop (l, l) is in \mathcal{D} or not. Similarly, notice how, for two letters l_1, l_2 , the sets $A_s = \{i : w_i = l_s\} (s = 1, 2)$ are complete to each other if \mathcal{D} contains both arcs (l_1, l_2) and (l_2, l_1) , and anticomplete to each other if \mathcal{D} contains none of the two arcs. Finally, the least trivial situation is when \mathcal{D} contains exactly one of the arcs (l_1, l_2) . For instance, we note in the figure that \mathcal{D} has the arc (a, c) , but not the arc (c, a) . This tells us that in $G(\mathcal{D}, w)$, we connect each a to every c appearing after it, but not to any of the c s appearing before it.

Example 5. The simplest non-trivial example of a decoder is $\mathcal{D} = (\{a, b\}, \{(a, b)\})$ (that is, \mathcal{D} is a digraph with two vertices and a single directed arc between them). Graphs with this decoder are exactly the bipartite chain graphs; in particular, if $w = abab \dots ab$ is the concatenation of the word ab n times, the graph $G(\mathcal{D}, w)$ is the prime chain graph on $2n$ vertices (see Figure 3.2). The indices in the figure indicate the order in which the vertices appear in w .

Example 6. For our final example, we remark that any graph G has a letter graph representation $G \cong G(\mathcal{D}, w)$, if we put $V(\mathcal{D}) = V(G)$ and $E(\mathcal{D}) = \{(u, v), (v, u) : \{u, v\} \in E(G)\}$ (w can be any word containing each letter exactly once).

This last example also shows that the question of interest for this notion is not simply “can we represent a given graph as a letter graph?”. Instead, we




Figure 3.1: A letter graph




Figure 3.2: The prime chain graph on 10 vertices

want to investigate what happens when we fix a decoder, and consider all graphs representable with that particular decoder. Since, for a given size of the alphabet Σ , there are only finitely many possible decoders, this is more or less the same as studying what happens when we bound the number of letters. To this end, we have the following definitions:

Definition 7. Let G be a graph. The *lettericity* $\text{let}(G)$ of G is the smallest $n \in \mathbb{N}$ such that G is isomorphic to a letter graph over a decoder $\mathcal{D} = (\Sigma, A)$ with $|\Sigma| = n$.

For a decoder \mathcal{D} , we write $\mathcal{L}_{\mathcal{D}}$ for the class of graphs representable as letter graphs with that decoder, and call it the *class of letter graphs with decoder \mathcal{D}* . For a natural k , the *class of k -letter graphs* \mathcal{L}_k is the (finite) union $\bigcup_{\mathcal{D}: |V(\mathcal{D})|=k} \mathcal{L}_{\mathcal{D}}$.

Remark 8. The classes $\mathcal{L}_{\mathcal{D}}$ (and, as a consequence, \mathcal{L}_k) are hereditary. Indeed, it is easy to check that any induced subgraph H of $G \cong G(\mathcal{D}, w)$ can be written as $G(\mathcal{D}, w')$ where w' is obtained from w by deleting the entries not corresponding to vertices of H .

Remark 9. The definition immediately shows that not every class has bounded

lettericity, since lettericity is bounded below by co-chromatic number.

In [Pet02], Petkovšek characterises k -letter graphs as follows.

Proposition 10 ([Pet02], Proposition 1). *A graph G is a k -letter graph if and only if*

1. *there is a partition V_1, V_2, \dots, V_p of $V(G)$ with $p \leq k$ such that each V_i is either a clique or an independent set in G , and*
2. *there is a linear ordering L of $V(G)$ such that for each pair of distinct indices $1 \leq i, j \leq p$, the intersection of $E(G)$ with $V_i \times V_j$ is one of the following four types (where L is considered as a binary relation, i.e., as a set of pairs):*
 - i. $L \cap (V_i \times V_j)$;
 - ii. $L^{-1} \cap (V_i \times V_j)$;
 - iii. $V_i \times V_j$;
 - iv. \emptyset .

One of the main reasons the notion of letter graphs is interesting is that Σ^* comes with a natural partial order called *subword* (or *subsequence*) *embedding*, that interacts nicely with the induced subgraph partial order:

Definition 11. Let $w = w_1 w_2 \dots w_n$ and $w' = w'_1 w'_2 \dots w'_n$ be two words over an alphabet Σ . We say w is a *subword* (or *subsequence*) of w' (denoted $w \leq w'$) if $n \leq n'$, and there is an increasing injection $\iota : [n] \rightarrow [n']$ such that $w_i = w'_{\iota(i)}$ for $i = 1, \dots, n$. ι is called a *subword* or *subsequence embedding*.

Lemma 12. *Let $G \cong G(\mathcal{D}, w)$ and $G' \cong G(\mathcal{D}, w')$ for some decoder \mathcal{D} and words w, w' . If $w \leq w'$, then $G \leq_i G'$.¹*

Proof. Suppose $w \leq w'$, and let ι be the subword embedding. For $i < j$, we note that i is adjacent to j in G if and only if (w_i, w_j) is in the decoder. Similarly, $\iota(i)$ is adjacent to $\iota(j)$ in G' if and only if $(w'_{\iota(i)}, w'_{\iota(j)})$ is in the decoder. By construction, $(w_i, w_j) = (w'_{\iota(i)}, w'_{\iota(j)})$, so that i is adjacent to j if and only if $\iota(i)$ is adjacent to $\iota(j)$. But this is precisely the same as saying $\iota : V(G) \rightarrow V(G')$ is an induced subgraph embedding. \square

¹The converse of this lemma is the content of Remark 8, so that $w \leq w'$ if and only if $G \leq_i G'$.

This simple fact allows us to use order-theoretic results, namely Higman's Lemma [Hig52], on the classes $\mathcal{L}_{\mathcal{D}}$ (and in general, on classes of bounded lettericity).

Theorem 13 (Restricted version of Higman's Lemma, [Hig52], Theorems 1.2 and 4.3). *The subword relation defined above is a wqo when the alphabet is finite.*

Theorem 14 ([Pet02], Theorem 8). *The classes \mathcal{L}_k are wqo by the induced subgraph relation.*

Corollary 15. *Any class of bounded lettericity is wqo.*

Theorem 14 makes graph lettericity an important parameter when studying wqo of classes of graphs under the induced subgraph relation, since it provides non-trivial examples of wqo classes of graphs, and it also gives a useful method for proving certain classes are wqo. The theorem also provides an alternative argument that not all classes of graphs have bounded lettericity, since any class containing all cycles (or, indeed, any other infinite antichain) must have unbounded lettericity. Let us construct one explicit example of graphs of high lettericity:

Example 16. Let $n \in \mathbb{N}$. We have $\text{let}(nK_2) = n$. Indeed, it is easy to see that n letters are enough to represent the graph (just use one letter per edge). If we had $\text{let}(nK_2) < n$, then there would be 3 vertices with the same letter, say a . Denote their appearances in the word w representing nK_2 by a_1, a_2 and a_3 , so that a_2 lies between the other two in w . We note that no vertex can be adjacent to only the vertex corresponding to a_2 , which is a contradiction, since every vertex in nK_2 has degree 1.

Before moving on, we mention one more result shown in [Pet02].

Theorem 17 ([Pet02], Theorem 9). *For each k , the class \mathcal{L}_k is characterised by finitely many minimal forbidden induced subgraphs.*

Sketch of proof. Let S_k be the set of minimal forbidden induced subgraphs for \mathcal{L}_k . If $G \in S_k$, then for any vertex v of G , $G - v$ is in \mathcal{L}_k . It is not too difficult to see that this implies $G \in \mathcal{L}_{2k+1}$. The claim follows, since \mathcal{L}_{2k+1} is wqo and S_k is an antichain. \square

3.1.2 Monotone and geometric griddability

The study of permutations as combinatorial objects is a rich and rapidly developing area of research. Topics of interest include enumerative problems, and well-quasi-orderability. A detailed history of the field, albeit interesting, is outside our scope; we will instead only introduce the notions immediately relevant to us.

For our purposes, a permutation is a linear order on $[n]$ for some $n \in \mathbb{N}$ – in other words, a string in which every number in $[n]$ appears exactly once, such as “41325” or “7654321”. We will refer to the characters in the string as *digits* or *elements*. Permutations come with a natural partial order on them called *pattern containment*:

Definition 18. Let $w = w_1w_2 \dots w_t$ and $w' = w'_1w'_2 \dots w'_{t'}$ be two words in \mathbb{N}^* . We say w is *order-isomorphic* to w' if $t = t'$, and for all $1 \leq i, j \leq t$, $w_i \leq w_j$ if and only if $w'_i \leq w'_j$.

Now let σ and π be two permutations. We say σ is a *pattern* of π (or π *contains* σ as a pattern) if π contains a subsequence that is order-isomorphic to σ . If π contains no such subsequence, we say π *avoids* σ .

Example 19. The permutation 2713564 contains 1423 as a pattern. Indeed, the subsequence 2735 is order-isomorphic to 1423.

As another example, the permutations that avoid 21 as a pattern are exactly the increasing permutations 1, 12, 123, 1234, \dots

Pattern containment is analogous to the induced subgraph relation, and we can define *permutation classes* as sets of permutations closed under (isomorphisms and) pattern containment. By the same general theory as in the case of graphs, any permutation class \mathcal{X} can be characterised uniquely in terms of its set of minimal avoided patterns $\text{Av}(\mathcal{X})$, also known as the *basis* of \mathcal{X} .

Remark 20. Now is a good time to point out that in the study of permutations on the one hand, and graphs on the other, completely analogous concepts might have different terminology associated to them. This difference might be subtle – for instance, graph theorists usually use the word “hereditary” to specify when a graph class is closed under taking induced subgraphs, while in the permutation literature, classes are often closed under pattern containment from the definition. We will do our best to avoid any ambiguities caused by this, but the reader should be warned that, when we deem the risk of confusion to be low, we will liberally borrow from one field to refer to concepts from the other, like saying a graph “avoids” another (as an induced subgraph). Similarly, we might use more general terminology from standard combinatorial theory, like saying “minimal obstructions” (or “minimal obstacles”) to refer to either minimal forbidden induced subgraphs for a graph class, or to minimal avoided patterns for a permutation class.

We may identify a permutation π on $[n]$ with its *plot*, the set of points $\{(i, \pi(i)) : 1 \leq i \leq n\}$ in the plane. More generally, [Alb+13] describes a rigorous

framework for this geometric perspective on permutations. We do not need the full generality of their framework², but the gist of it is as follows: call a set of points in the plane *independent* if no two points lie on the same vertical or horizontal line. We may define a permutation as an equivalence class of finite independent sets of points, where two such sets are equivalent if, roughly speaking, we can get from one of them to the other by vertical and horizontal stretching or shrinking.

As an example, Figure 3.3 illustrates the plot of 614253 (axes are omitted), which is a representative for its equivalence class. The only thing that matters is the relationship between the vertical and horizontal orderings of the six points. More concretely, if we label the points in increasing order from the bottom to the top, then reading the labels from left to right yields 614253. The full equivalence class consist of exactly the (independent) sets of points with this property.




Figure 3.3: Geometric representation of $\pi = 614253$.

We will now talk about two tools used to study permutation classes: monotone and geometric griddability. The notion of monotone griddability was developed over several papers, by successively generalising previous notions. Some of the steps that led to the definition that we have today can be found in [Atk99; AMR02; MV02]. The definitions we give here are more or less equivalent to the ones introduced in [HV06].

²Indeed, [Alb+13] and to some degree [HV06] present everything with an added level of formalism. This has the benefit of making the geometric theory of permutations and the tools we are about to describe fairly robust, but it does so at the price of brevity. Since our focus is not on permutations themselves, but rather on their relationship to graphs (and, as we will see, to another combinatorial structure capturing some of their order properties), we will take some shortcuts along the way. Our aim here is to give the minimum amount of rigour necessary for developing an intuition in working with those tools; for the reader's peace of mind, we stress that everything we discuss in this subsection could be done carefully and in more detail.

Let $s, t \in \mathbb{N}$. An $s \times t$ *gridding* Γ is a set of $s + 1$ vertical and $t + 1$ horizontal lines in the plane. This partitions the rectangle in the plane defined by the extremal lines into st regions that we will call the *cells* of the gridding. The cells are labelled Z_{ij} , where the first index counts from left to right, and the second from bottom to top.³

Definition 21. Let π be a permutation and $M = (\alpha_{ij})$ an $s \times t$ matrix with entries in $\{0, \pm 1\}$. We say π is *monotonically griddable by M* (or just “griddable” for short) if there exists an $s \times t$ gridding Γ such that:

- If $\alpha_{ij} = 0$, then $\pi \cap Z_{ij} = \emptyset$.
- If $\alpha_{ij} = 1$, then $\pi \cap Z_{ij}$ is increasing.
- If $\alpha_{ij} = -1$, then $\pi \cap Z_{ij}$ is decreasing.

We say such a Γ is a *monotone gridding* of π by M – see Figure 3.4 for an example.




Figure 3.4: A monotone gridding of 614253 by $\begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix}$.

We note that monotone griddability is well-defined: if the plot of a permutation π is griddable by some matrix, then so is every other representative for π (we may simply stretch the gridding with the plot to get to any independent set of points in the equivalence class of π). We also note that if π is griddable by a matrix M , then so is any subpattern σ of π . This motivates the following definitions:

³In particular, whenever we use matrices, we will follow the same (non-standard) indexing convention: an $s \times t$ matrix $M = (a_{ij})$ has s columns and t rows; the indices i, j count the entries of M from left to right, and from bottom to top respectively.

Definition 22. Let M be a $0/\pm 1$ matrix. The *grid class* of M , denoted $\text{Grid}(M)$, is the class of permutations monotonically griddable by M . A class \mathcal{X} of permutations is called *monotonically griddable* if $\mathcal{X} \subseteq \text{Grid}(M)$ for some fixed $0/\pm 1$ matrix M .

Huczynska and Vatter [HV06] give a characterisation of monotone griddable classes in terms of minimal non-griddable classes. To state it, we first need a definition:

Definition 23. Let $\pi \in S_m$ and $\sigma \in S_n$. We define their *direct sum* $\pi \oplus \sigma$ by

$$(\pi \oplus \sigma)(i) = \begin{cases} \pi(i) & \text{if } i \in [m], \\ \sigma(i - m) + m & \text{if } i \in [m + n] \setminus [m], \end{cases}$$

and similarly their *skew sum* $\pi \ominus \sigma$ by

$$(\pi \ominus \sigma)(i) = \begin{cases} \pi(i) + n & \text{if } i \in [m], \\ \sigma(i - m) & \text{if } i \in [m + n] \setminus [m]. \end{cases}$$

Figure 3.5 illustrates the geometric meaning of the direct and skew sums.




Figure 3.5: Direct and skew sum of two permutations

Theorem 24 ([HV06], Theorem 2.5). *A permutation class is monotone griddable if and only if it does not contain arbitrarily long direct sums of 21 or skew sums of 12.*

In other words, a class \mathcal{X} of permutations is monotone griddable if and only if it does not contain the class of all (subpatterns of) direct sums of 21 or the class of all (subpatterns of) skew sums of 12.

We now discuss the second, stronger notion of griddability that we mentioned, introduced in [Alb+13] and called *geometric griddability*. The definition is very similar to that of monotone griddability, where we start with a $0/\pm 1$ matrix M and a gridding whose cells correspond to entries of M . However, instead of simply requiring that π is monotone in the cells of the gridding, we put the stronger condition that the entries of π in each cell lie on one of the diagonals.

Definition 25. Let π be a permutation and $M = (\alpha_{ij})$ an $s \times t$ matrix with entries in $\{0, \pm 1\}$. We say π is *geometrically griddable by M* if there exists an $s \times t$ gridding Γ such that:

- If $\alpha_{ij} = 0$, then $\pi \cap Z_{ij} = \emptyset$.
- If $\alpha_{ij} = 1$, then $\pi \cap Z_{ij}$ lies on the main diagonal⁴ of Z_{ij} .
- If $\alpha_{ij} = -1$, then $\pi \cap Z_{ij}$ lies on the antidiagonal of Z_{ij} .

We say such a Γ is a *geometric gridding* of π by M – see Figure 3.6 for an example. The union of the diagonals and antidiagonals on which the entries of π may lie is called the *standard figure* of M .

Remark 26. Without loss of generality, we may apply some normalisations so that the cells in any monotone or geometric gridding correspond to unit squares, with the bottom left corner of the bottom left cell at $(0, 0)$. Indeed, since we allow vertical and horizontal stretching, those normalisations do not affect the permutations geometrically griddable by M . In particular, we will assume the standard figure is subject to those normalisations.

Similarly to (monotone) griddability, we define geometric griddability of classes as follows:

Definition 27. Let M be a $0/\pm 1$ matrix. The *geometric grid class* of M , denoted $\text{Geom}(M)$, is the class of permutations geometrically griddable by M . A class \mathcal{X} of permutations is called *geometrically griddable* if $\mathcal{X} \subseteq \text{Geom}(M)$ for some fixed $0/\pm 1$ matrix M .

It is clear from the definition that any permutation geometrically griddable by a matrix is monotonically griddable by that matrix. Concisely, for any $0/\pm 1$ matrix M , $\text{Geom}(M) \subseteq \text{Grid}(M)$. Is the converse true? As one might expect, the answer is in general negative:

⁴That is, the straight line segment connecting the bottom left corner with the top right one.




Figure 3.6: A geometric gridding of 614253 by $\begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix}$.

Example 28. Let $\pi = 2413$ and $M = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$. Then $\pi \in \text{Grid}(M)$, but $\pi \notin \text{Geom}(M)$. That $\pi \in \text{Grid}(M)$ is easy to see – we can grid it with one element per cell. To see that $\pi \notin \text{Geom}(M)$, one can derive a contradiction by first noting (via simple case analysis) that there must be one element per cell, then seeing that, going around in a clockwise cycle starting at say 2, the distance from each element to the centre of the figure should strictly decrease, leaving no place to put 1 (see Figure 3.7).

Does it ever happen that $\text{Grid}(M)$ and $\text{Geom}(M)$ coincide? The answer is “yes”, and the matrices M for which this is the case are characterised in [Alb+13]. To state this characterisation, we need the notion of *cell graph* of a matrix M : for a matrix M , the vertices of the cell graph are the non-zero entries of M , and two vertices are adjacent if the corresponding entries share a row or a column, and all entries between them are 0 (see Figure 3.8).

The full characterisation says that $\text{Grid}(M) = \text{Geom}(M)$ if and only if the cell graph of M is a forest, and is an immediate consequence of the following:

Theorem 29 ([Alb+13], Theorem 3.2). *If the cell graph of M is a forest, then $\text{Grid}(M) = \text{Geom}(M)$.*

Theorem 30 ([Alb+13], Theorem 6.1). *Every geometrically griddable class is wqo.*

Theorem 31 ([MV02], Theorem 2.2). *$\text{Grid}(M)$ is wqo if and only if the cell graph of M is a forest.*





Figure 3.7: An attempt to geometrically grid 2413.

$$\begin{pmatrix} -1 & 1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & 0 & -1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}$$

(a) The matrix M



(b) Its cell graph.

Figure 3.8: A matrix and its cell graph.

Theorem 29 is shown by induction, the main insight being that we have a lot of freedom to move the entries of a permutation lying in a leaf of the cell graph.

Theorems 30 and 31 together imply the converse of Theorem 29. It is worth noting that Theorem 30 is (besides enumerative results) one of the big reasons geometrically griddable classes are interesting. It is also the first hint suggesting that the notions of geometric griddability and graph lettericity are related (compare with Theorem 14).

The last item we discuss in this subsection is the proof of Theorem 30. The proof, which is a straightforward generalisation of work from [VW11], consists of defining a finite alphabet Σ depending on M , then producing an order preserving, surjective map $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$. Σ^* is wqo by Higman's lemma, and Theorem 30 then immediately follows. To describe the map requires a bit of preparation; we start

with a result which, despite being a simple technicality, proves to be very useful. We need a quick definition:

Definition 32. We say an $s \times t$, $0/\pm 1$ matrix $M = (\alpha_{ij})$ is a *partial multiplication matrix* if there exist column and row signs $c_1, \dots, c_s, r_1, \dots, r_t \in \{\pm 1\}$ such that α_{ij} is either 0 or the product $c_i r_j$.

Proposition 33 ([Alb+13], Proposition 4.2). *Every geometric grid class is the geometric grid class of a partial multiplication matrix.*

Sketch of proof. We define a refinement $M^{\times k}$ of a matrix M by replacing each entry with a $k \times k$ matrix. 0s and 1s are replaced with the 0 and identity matrices respectively, and -1 is replaced with a matrix with -1 s on the antidiagonal, and 0s everywhere else. It is not difficult to see that $\text{Geom}(M) = \text{Geom}(M^{\times k})$ for any k , and that $M^{\times 2}$ is always a partial multiplication matrix. \square

From now on, we will assume unless otherwise specified that we are working with partial multiplication matrices; the column and row signs will often be represented by arrows to the left of and above the standard figure of the matrix – as in Figure 3.9. Which arrow direction corresponds to which sign is immaterial, as long as it is consistent. The partial multiplication matrix condition says exactly that those arrows can be chosen to always “agree” with the diagonals in the cells. This choice of signs also yields a distinguished corner for each cell, as indicated in the figure by the large black dots.

Let us now describe the map φ . We define the *cell alphabet* of the matrix $M = (\alpha_{ij})$ as the set $\Sigma := \{a_{kl} : \alpha_{kl} \neq 0\}$, that is, the letters a_{kl} correspond to non-zero entries of M . From a word $w = w_1 w_2 \dots w_t \in \Sigma^*$, we construct a permutation $\pi \in \text{Geom}(M)$ as follows: choose a set $0 < d_1 < d_2 \dots < d_t < 1$ of distances, then for $1 \leq i \leq t$, if $w_i = a_{kl}$, place a point x_i on the diagonal of cell Z_{kl} at infinity-norm distance d_i of the distinguished corner of that cell. Figure 3.9 illustrates this with $M = \begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix}$; then $\Sigma = \{a_{11}, a_{12}, a_{21}, a_{22}\}$. If we let $w = a_{12} a_{11} a_{21} a_{22} a_{12} a_{21}$, then choosing $d_i = \frac{i}{7}$ for $i = 1, \dots, 6$ produces the permutation 614253. The circled numbers are the indices i corresponding to each entry – one can interpret them as the order in which we insert the elements of the permutation into the picture.

It is not difficult to see that φ is well defined, in that its value does not depend on the choice of distances. Moreover, it is onto and order preserving ([Alb+13], Proposition 5.3). φ is not injective, since changing the relative order in which we add elements from *independent cells* (that is, cells that do not share a row or a




Figure 3.9: $\varphi(a_{12}a_{11}a_{21}a_{22}a_{12}a_{21}) = 614253$.

column) does not alter the permutation.⁵

3.1.3 Permutation graphs

Permutations can be related to graphs via the notion of permutation graphs. To a permutation π on $[n]$ we associate its *inversion graph* G_π , whose vertex set is $[n]$ and whose edges are the pairs $\{i, j\}$ that are *inverted* by π , in the sense that $(i - j)(\pi(i) - \pi(j)) < 0$. The class of permutation graphs consists of all graphs that are the inversion graph of some permutation. It has several known characterisations (for instance, as intersection graphs of line segments between two parallel lines, as graphs that are simultaneously comparability and co-comparability [DM41] or as comparability graphs of a poset of order dimension at most two [BFR72]), including a minimal forbidden induced subgraph one [Gal67].

The mapping from permutations to their permutation graphs is not injective, since for instance, both 2413 and 3142 have P_4 as their permutation graph. However, prime graphs with respect to modular decomposition have only two permutation representations that are inverse to each-other (in the function sense) [Gal67]. We

⁵In fact, as shown in [Alb+13], for any permutation π , $\varphi^{-1}(\pi)$ is an equivalence class of words where we are allowed to swap pairs of consecutive letters corresponding to independent cells. Such an equivalence class is called a *trace*, and the map φ could be made bijective by defining it instead on the so-called *trace monoid* Σ^* modulo this equivalence relation. Those objects have been studied relatively thoroughly, e.g., in [Die90]. We do not need to concern ourselves with these facts for the time being.

also note that, if σ contains π as a pattern, then G_σ contains G_π as an induced subgraph.

The inversion graph of a permutation is particularly easy to read from its geometric plot: two vertices are adjacent if and only if one appears to the bottom right of the other (see Figure 3.10).




Figure 3.10: Permutation graph of 614253.

This observation allows us to easily translate properties of (classes of) permutations into properties of the corresponding (classes of) permutation graphs. For instance, any increasing, respectively decreasing sequence in a permutation π corresponds to an independent set, respectively a clique in G_π . Similarly, if π admits a gridding by say $M = \begin{pmatrix} 1 & 1 \end{pmatrix}$, then G_π is a bipartite chain graph (the entries of π in each of the cells correspond to an independent set in G_π , and it is not difficult to see the neighbourhoods of the vertices in each part form a chain under inclusion).

In general, if a class \mathcal{X} is monotonically griddable by a matrix M , then there exists a k such that in the corresponding graph class $\mathcal{G}_\mathcal{X}$, any graph G admits a partition into at most k bags, where each bag is a clique or an independent set. Those bags correspond to the non-empty cells of the gridding of the permutations by M . Moreover, between any two bags we have either no edges, all possible edges, or a bipartite chain graph according to the relative positions of the corresponding cells.

3.1.4 Hypergraphs

We only use hypergraphs superficially, and we therefore keep this subsection short. A *hypergraph* \mathcal{H} is a pair $\mathcal{H} = (X, E)$, where X is a set and E is a set of non-empty subsets of X called *hyperedges*. Due to the extreme generality of this definition,

there are several valid ways to make sense of a subhypergraph and of hypergraph isomorphism. We will use the following:

Definition 34. Let $\mathcal{H} = (X, E)$ be a hypergraph. Write $X = \{x_i : i \in I\}$, $E = \{e_i : i \in J\}$, where I and J are finite index sets.

Given a set $A \subseteq X$, the *subhypergraph* $\mathcal{H}[A]$ induced by A is defined as

$$\mathcal{H}[A] = (A, \{e \cap A : e \in E \text{ and } e \cap A \neq \emptyset\}).$$

Given a subset $L \subseteq J$ of the hyperedge index set, the *partial hypergraph* generated by L is the hypergraph $(X, \{e_i : i \in L\})$. One may think of subhypergraphs as induced subgraphs, and partial hypergraphs as spanning subgraphs.

Suppose $\mathcal{H} = (X, E)$ and $\mathcal{J} = (Y, F)$ are two hypergraphs using the same hyperedge index set J . We say \mathcal{H} is *isomorphic* to \mathcal{J} if there exists a bijection $\psi : X \rightarrow Y$ and a permutation π of J such that $\psi(e_i) = f_{\pi(i)}$ for all $i \in J$. We say \mathcal{H} is *strongly isomorphic* to \mathcal{J} , written $\mathcal{H} \cong \mathcal{J}$, if the permutation π above is the identity.

Finally, a hypergraph is called *downwards closed* if every subset of a hyperedge is a hyperedge.

3.2 Characterisation, recognition and an intriguing connection

The results presented in this section are the product of joint work together with Vadim Lozin, Dominique de Werra and Viktor Zamaraev. The work was published in [Ale+20b].

In Subsection 3.2.1, we conduct a case study of the class of letter graphs with a fixed decoder \mathcal{D} . We investigate the problem of characterisation and recognition for graphs from that particular class in order to obtain an idea of how we may approach this problem in general.

In Subsection 3.2.2, we use the map $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$ described in Subsection 3.1.2 to show that the class $\mathcal{G}_{\text{Geom}(M)}$ of permutation graphs has bounded lettericity, and we show a converse statement when the class of permutation graphs has lettericity at most 2.

3.2.1 Characterisation and recognition of 3-letter graphs

In [Pet02], Petkovšek characterises and enumerates the classes of 2-letter graphs with each possible decoder. In this subsection, as a “proof of concept”, we analyse the class of 3-letter graphs with decoder $\{(a, b), (b, c), (c, a)\}$. We provide a structural characterisation, a recognition algorithm, and a minimal forbidden induced subgraph characterisation for graphs in this class.

Characterisation of 3-letter graphs over the decoder $\{(a, b), (b, c), (c, a)\}$

To characterise 3-letter graphs, we need a few observations about 2-letter graphs. Let $G = (V, E)$ be a graph and A an independent set in G . We will say that a linear order (a_1, a_2, \dots, a_k) of the vertices of A is

- *increasing* if $i < j$ implies $N(a_i) \subseteq N(a_j)$,
- *decreasing* if $i < j$ implies $N(a_i) \supseteq N(a_j)$,
- *monotone* if it is either increasing or decreasing.

By definition, each part of a chain graph (i.e., a $2K_2$ -free bipartite graph) admits a monotone ordering. Let $G = (A \cup B, E)$ be a chain graph given together with a bipartition $V(G) = A \cup B$ of its vertices into two independent sets. We fix an order of the parts (A is first and B is second), a decreasing order for A , an increasing order for B , and call G a *properly ordered graph*. This notion suggests an easy way of representing a $2K_2$ -free bipartite graph as a 2-letter graph.

Let $G = (A \cup B, E)$ be a properly ordered $2K_2$ -free bipartite graph. To represent G as a 2-letter graph, we fix the alphabet $\Sigma = \{a, b\}$ and the decoder $\mathcal{P} = \{(a, b)\}$. The word ω representing G can be constructed as follows. To each vertex of A we assign letter a and to each vertex of B we assign letter b . The a letters will appear in ω in the order in which the corresponding vertices appear in A and the b letters will appear in ω in the order in which the corresponding vertices appear in B . The rule defining the relative positions of a vertices with respect to b vertices can be described in two different ways as follows:

- R_1 an a vertex is located between the last b non-neighbour (if any) and the first b neighbour (if any),
- R_2 a b vertex is located between the last a neighbour (if any) and the first a non-neighbour (if any).

It is not difficult to see that both rules R_1 and R_2 define the same word and this word represents G .

Now we turn to 3-letter graphs. Let $G = (A \cup B \cup C, E)$ be a graph whose vertex set is partitioned into three independent sets A, B, C such that

- (a) $G[A \cup B]$, $G[B \cup C]$ and $G[C \cup A]$ are $2K_2$ -free bipartite graphs,
- (b) there are no three vertices $a \in A, b \in B, c \in C$ inducing either a triangle K_3 or an anti-triangle \overline{K}_3 .

We call any graph satisfying (a) and (b) *nice*. Our goal is to show that a graph G is a 3-letter graph over the decoder $\{(a, b), (b, c), (c, a)\}$ if and only if it is nice. First, we prove the following lemma.

Lemma 35. *Let $G = (A \cup B \cup C, E)$ be a nice graph. Then each of the independent sets A, B and C admits a linear ordering such that all three bipartite graphs $G[A \cup B]$, $G[B \cup C]$ and $G[C \cup A]$ are properly ordered.*

Proof. We start with a proper order of $G[A \cup B]$, in which case the order of B is increasing with respect to A . Let us show that the same order of B is decreasing with respect to C .

Consider two vertices b_i and b_j of B with $i < j$, i.e., b_i precedes b_j in the linear order of B and hence $N(b_i) \cap A \subseteq N(b_j) \cap A$. To show that the linear order of B is decreasing with respect to C , assume the contrary: b_j has a neighbour $c \in C$ non-adjacent to b_i . Without loss of generality, we may suppose that the inclusion $N(b_i) \cap A \subseteq N(b_j) \cap A$ is proper, since we can, if necessary, reorder all vertices with equal neighbourhoods in A decreasingly with respect to their neighbourhoods in C , which keeps the graph $G[A \cup B]$ properly ordered. According to this assumption, b_j must have a neighbour $a \in A$ non-adjacent to b_i . But then either a, b_j, c induce a triangle K_3 (if a is adjacent to c) or a, b_i, c induce an anti-triangle \overline{K}_3 (if a is not adjacent to c). A contradiction in both cases shows that the linear order of B is decreasing with respect to C .

Similar arguments show that the order of A which is decreasing with respect to B is increasing with respect to C . Now we fix a linear order of C which is increasing with respect to B and conclude, as before, that it is decreasing with respect to A . In this way, we obtain a proper order for all three graphs $G[A \cup B]$, $G[B \cup C]$ and $G[C \cup A]$ (notice, in the last graph C is the first part and A is the second). \square

Theorem 36. *A graph G is a 3-letter graph over the decoder $\{(a, b), (b, c), (c, a)\}$ if and only if it is nice.*

Proof. If G is a 3-letter graph over the decoder $\{(a, b), (b, c), (c, a)\}$, then obviously V_a (the set of vertices labelled by a), V_b and V_c are independent sets and condition (a) of the definition of nice graphs is valid for G . To show that (b) is valid, assume G contains a triangle induced by letters a, b, c . Then b must appear after a in the word representing G , and c must appear after b . But then c appears after a , in which case a is not adjacent to c , a contradiction. Similarly, an anti-triangle a, b, c is not possible and hence G is nice.

Suppose now that $G = (A \cup B \cup C, E)$ is nice. According to Lemma 35, we may assume that A, B and C are ordered in such a way that each of the three bipartite graphs $G[A \cup B]$, $G[B \cup C]$ and $G[C \cup A]$ is properly ordered.

We start by representing the graph $G[A \cup B]$ by a word ω with two letters a, b according to rules R_1 or R_2 . To complete the construction, we need to place the c vertices

- among the a vertices according to rule R_1 , i.e., every c vertex must be located between the last a non-neighbour $a_{l_{nn}}$ (if any) and the first a neighbour a_{f_n} (if any),
- among the b vertices according to rule R_2 , i.e., every c vertex must be located between the last b neighbour b_{l_n} (if any) and the first b non-neighbour $b_{f_{nn}}$ (if any).

This is always possible, unless

- either a_{f_n} precedes b_{l_n} in ω , in which case a_{f_n} is adjacent to b_{l_n} and hence a_{f_n}, b_{l_n}, c induce a triangle K_3 ,
- or $b_{f_{nn}}$ precedes $a_{l_{nn}}$ in ω , in which case $b_{f_{nn}}$ is not adjacent to $a_{l_{nn}}$ and hence $a_{l_{nn}}, b_{f_{nn}}, c$ induce an anti-triangle \overline{K}_3 .

A contradiction in both cases shows that ω can be extended to a word representing G . □

Note that this theorem can be viewed as a specialisation of Theorem 10.

Recognition of 3-letter graphs over the decoder $\{(a, b), (b, c), (c, a)\}$

In this section, we show how we can determine whether a graph G can be represented as a 3-letter graph over the cyclic decoder $\{(a, b), (b, c), (c, a)\}$. In order to do that, we will assume that G does indeed have such a representation ω , and derive various properties of ω .

If G has a twin v for a vertex u (i.e., $N(v) = N(u)$), then any word representing $G - v$ can be extended to a word representing G by assigning to v the same letter as to u and placing v next to u . This observation shows that we may assume without loss of generality that

- G is twin-free.

Due to the cyclic symmetry of the decoder, we may also assume without loss of generality that

- the last letter of ω is c .

Then

- the first letter is not a , since otherwise the first and the last vertices are twins.

Assume that the first letter of ω is b . Then according to the decoder

- (b1) no vertex between the first b and the last c is adjacent to both of them,
- (b2) every vertex non-adjacent to the first b and non-adjacent to the last c must be labelled by a ,
- (b3) every vertex non-adjacent to the first b and adjacent to the last c must be labelled by b ,
- (b4) every vertex adjacent to the first b and non-adjacent to the last c must be labelled by c .

If the first letter is c , we have instead that:

- (c1) no vertex between the first c and the last c is adjacent to both of them,
- (c2) every vertex adjacent to the first c and non-adjacent to the last c must be labelled by a ,
- (c3) every vertex non-adjacent to the first c and adjacent to the last c must be labelled by b ,
- (c4) every vertex non-adjacent to the first c and non-adjacent to the last c must be labelled by c .

This discussion, together with conditions (a) and (b) from the previous subsection, leads to the following recognition algorithm:

Algorithm: Recognition of 3-letter graphs over the cyclic decoder $\{(a, b), (b, c), (c, a)\}$

Input: A graph G

Output: **true** if G is a 3-letter graph over the cyclic decoder, **false** otherwise

```

1: set  $G' := G$ 
2: while  $G'$  has a pair of twins do
3:   remove one of the twins from  $G'$ 
4: for each ordered pair  $(u, v)$  of distinct vertices in  $G'$  do
5:   if  $u$  and  $v$  have no common neighbours then
6:     if  $u$  and  $v$  are adjacent then
7:       set  $A := \overline{N(u)} \cup \overline{N(v)}$ 
8:       set  $B := \overline{N(u)} \cup N(v) \cup \{u\}$ 
9:       set  $C := N(u) \cup \overline{N(v)} \cup \{v\}$ 
10:    else
11:      set  $A := N(u) \cup \overline{N(v)}$ 
12:      set  $B := \overline{N(u)} \cup N(v)$ 
13:      set  $C := \overline{N(u)} \cup \overline{N(v)} \cup \{u, v\}$ 
14:    if  $G[A \cup B]$ ,  $G[B \cup C]$  and  $G[C \cup A]$  are  $2K_2$ -free bipartite graphs and
      there are no vertices  $a \in A$ ,  $b \in B$ ,  $c \in C$  inducing either a triangle or an
      anti-triangle then
15:      return true
16: return false

```

Theorem 37. *The 3-letter graphs over the decoder $\{(a, b), (b, c), (c, a)\}$ can be recognized in polynomial time.*

Proof. It is easy to see that the above algorithm terminates, and correctness follows from the above discussion together with Theorem 36.

To determine its complexity, let n be the number of vertices of the graph and m the number of edges. Note first that the ‘while’ loop at line 2 takes $O(n^3)$ time per iteration and iterates at most n times.

Lines 6 to 14 can be implemented in linear time. It takes $O(n + m)$ time to recognize chain graphs (see, e.g., [HK07]), and the condition on triangles and anti-triangles can be checked in $O(n^{2.376})$ time (finding triangles can be reduced to matrix multiplication, see, e.g., [IR77], and we can use for instance the Coppersmith-Winograd algorithm after some straightforward preprocessing in order to detect the appropriate kinds of triangles or anti-triangles). Finally, the ‘for’ loop at line 5 is iterated through at most n^2 times. This gives $O(n^{4.376})$ time complexity overall. \square

We remark that the above algorithm can be made constructive without an

increase in complexity. We just need to make two modifications as follows: first, we use Theorem 36 to produce a word representing the twin-free graph G' . Second, we record the twins we removed to obtain G' from G and once we have a word for G' , we use the recorded information to obtain a word for G .

Minimal forbidden induced subgraphs for 3-letter graphs over the decoder $\{(a, b), (b, c), (c, a)\}$

To determine the list of minimal forbidden induced subgraphs for our class, we will rely on our earlier characterisation of graphs in this class as “nice” (Theorem 36). We start with a preparatory result.

Lemma 38. *Let G be a graph and let H_1 and H_2 be nice subgraphs of G with disjoint vertex sets $V(H_1) = A_1 \cup B_1 \cup C_1$ and $V(H_2) = A_2 \cup B_2 \cup C_2$. If the subgraphs induced*

- *by $A_1 \cup B_2$, $B_1 \cup C_2$ and $C_1 \cup A_2$ are complete bipartite,*
- *by $A_1 \cup A_2$, $A_1 \cup C_2$, $B_1 \cup A_2$, $B_1 \cup B_2$, $C_1 \cup B_2$ and $C_1 \cup C_2$ are edgeless,*

then the subgraph induced by $V(H_1) \cup V(H_2) = (A_1 \cup A_2) \cup (B_1 \cup B_2) \cup (C_1 \cup C_2)$ is nice.

Proof. By assumption, $A_1 \cup A_2$ and $B_1 \cup B_2$ are independent sets. Let us show that these two sets induce a chain graph. First, it is not difficult to see that $A_1 \cup (B_1 \cup B_2)$ induces a chain graph, because $G[A_1 \cup B_1]$ is a chain graph and $G[A_1 \cup B_2]$ is complete bipartite. Similar arguments show that $A_2 \cup (B_1 \cup B_2)$, $(A_1 \cup A_2) \cup B_1$ and $(A_1 \cup A_2) \cup B_2$ all induce chain graphs. Therefore, if the subgraph of G induced by $A_1 \cup A_2$ and $B_1 \cup B_2$ contains an induced $2K_2$, then this $2K_2$ contains exactly one vertex in each of the four sets, which is impossible. This contradiction shows that the subgraph of G induced by $A_1 \cup A_2$ and $B_1 \cup B_2$ is a chain graph.

By symmetry, $(B_1 \cup B_2) \cup (C_1 \cup C_2)$ and $(C_1 \cup C_2) \cup (A_1 \cup A_2)$ also induce chain graphs. It remains to show that no 3 vertices $a \in A_1 \cup A_2$, $b \in B_1 \cup B_2$, $c \in C_1 \cup C_2$ induce a triangle or an anti-triangle. Since H_1 and H_2 are nice, we may assume without loss of generality that two of the vertices belong to H_1 and one to H_2 . Also, due to the symmetry of the decoder, we may assume that $a \in A_1$, $b \in B_1$, and $c \in C_2$. Then a, b, c induce neither a triangle (since a is not adjacent to c) nor an anti-triangle (since b is adjacent to c). \square

We are now ready to prove the characterisation in terms of minimal forbidden induced subgraphs.

Theorem 39. *A graph G is a 3-letter graph over the decoder $\{(a, b), (b, c), (c, a)\}$ if and only if it is $(K_3, 2K_2 + K_1, C_5 + K_1, C_6)$ -free.*

Proof. For the “only if” direction, it is straightforward to check that none of the four graphs in our list is nice and that they are minimal with that property. For the “if” direction, we split the analysis into two cases.

Assume first that G is $2K_2$ -free. If, in addition, it is C_5 -free, then G is $2K_2$ -free bipartite, i.e., a chain graph (since it has no $2K_2$, K_3 , C_5 , and the absence of $2K_2$ forbids longer odd cycles), hence it is nice, with one of the 3 sets being empty. So suppose G has an induced C_5 . Label its vertices clockwise by v_1, \dots, v_5 (whenever indices are added in this proof, the addition will be modulo 5). Any vertex of G not in the C_5

- has to be adjacent to at least one vertex in the C_5 , since otherwise an induced $C_5 + K_1$ arises,
- cannot have a single neighbour in the C_5 , since otherwise an induced $2K_2$ can be easily found,
- cannot be adjacent to 3 or more vertices or to 2 consecutive vertices in the C_5 , since G is K_3 -free.

Hence the vertices of G can be partitioned into 5 sets V_1, \dots, V_5 such that the vertices in V_i are adjacent to exactly v_{i-1} and v_{i+1} in the C_5 (note $v_i \in V_i$ for $i = 1, \dots, 5$). Each V_i is an independent set (since they share a common neighbour, and triangles are forbidden), and adjacency between them is easy to determine:

- if $u_i \in V_i$, $u_{i+1} \in V_{i+1}$, then u_i and u_{i+1} are adjacent, since otherwise $u_i, v_{i-1}, u_{i+1}, v_{i+2}$ induce a $2K_2$,
- if $u_i \in V_i$, $u_{i+2} \in V_{i+2}$, then u_i and u_{i+2} are non-adjacent, since otherwise u_i, v_{i+1}, u_{i+2} induce a triangle.

This determines all adjacencies in G , and it is easy to check that G is nice, e.g., with partition $(V_1 \cup V_4), (V_2 \cup V_5), V_3$.

Now we turn to the case when G contains an induced $2K_2$. We denote one of the edges of the $2K_2$ by uw and partition the vertices of G into three subsets as follows (observe that there are no vertices adjacent to both u and w , since triangles are forbidden):

U is the set of vertices adjacent to w (u belongs to U). Since triangles are forbidden, U is an independent set.

W is the set of vertices adjacent to u (w belongs to W). Since triangles are forbidden, W is an independent set.

X is the set of vertices adjacent neither to u nor to w . The subgraph induced by X must be $K_2 + K_1$ -free, since otherwise an induced copy of $2K_2 + K_1$ would arise. It is not difficult to see that the $(K_2 + K_1, K_3)$ -free graphs that are not edgeless are precisely the complete bipartite graphs. Therefore, the vertices of X can be split into two independent sets with all possible edges between them. We call these independent sets C_1 and A_2 (this notation is chosen for consistency with Lemma 38) and observe that each of them is non-empty, because X contains the other edge of the K_2 .

Since G is K_3 -free, no vertex of G can have neighbours in both C_1 and A_2 . Thus W can be partitioned into three subsets as follows:

A_1 is the vertices of W that do have neighbours in C_1 (and hence have no neighbours in A_2),

C_2 is the vertices of W that do have neighbours in A_2 (and hence have no neighbours in C_1),

W' is the set of remaining vertices of W , i.e., those that have neighbours neither in C_1 nor in A_2 .

We partition U into three subsets in a similar way:

B_1 is the vertices of U that do have neighbours in C_1 (and hence have no neighbours in A_2),

B_2 is the vertices of U that do have neighbours in A_2 (and hence have no neighbours in C_1),

U' is the set of remaining vertices of U , i.e., those that have neighbours neither in C_1 nor in A_2 .

We note that

- *Every vertex of A_1 is adjacent to every vertex of B_2 .* Indeed, if $a_1 \in A_1$ is not adjacent to $b_2 \in B_2$, then u, w, a_1, b_2 together with a neighbour of a_1 in C_1 and a neighbour of b_2 in A_2 induce a C_6 .
- *Every vertex of B_1 is adjacent to every vertex of C_2 by similar arguments.*

- *Every vertex of U' is adjacent to every vertex of W' .* Indeed, if $u' \in U'$ is not adjacent to $w' \in W'$, then u', w', w together with any two vertices $c_1 \in C_1$ and $a_2 \in A_2$ induce a $2K_2 + K_1$.
- *Every vertex of W' is adjacent to every vertex in B_1 or to every vertex in B_2 .* Indeed, if a vertex $w' \in W'$ has a non-neighbour $b_1 \in B_1$ and a non-neighbour $b_2 \in B_2$, then w', w, b_1, b_2 together with a neighbour of b_1 in C_1 and a neighbour of b_2 in A_2 induce a $C_5 + K_1$.
- *Every vertex of U' is adjacent to every vertex in A_1 or to every vertex in C_2 by similar arguments.*

The above sequence of claims shows that we can move the vertices from W' to either A_1 or C_2 and those from U' to either B_1 or B_2 in such a way that the two subgraphs $G[A_1 \cup B_2]$ and $G[B_1 \cup C_2]$ are complete bipartite.

To sum up, we have partitioned G into independent sets A_1, A_2, B_1, B_2, C_1 , and C_2 , such that $G[A_1 \cup B_2]$, $G[B_1 \cup C_2]$ and $G[C_1 \cup A_2]$ are complete bipartite, while $G[A_1 \cup A_2]$, $G[A_1 \cup C_2]$, $G[B_1 \cup A_2]$, $G[B_1 \cup B_2]$, $G[C_1 \cup B_2]$ and $G[C_1 \cup C_2]$ are edgeless. To apply Lemma 38 it remains to show that $G[A_1 \cup B_1 \cup C_1]$ and $G[A_2 \cup B_2 \cup C_2]$ are nice.

Because of the $2K_2 + K_1$ -freeness, the subgraph induced by the set of non-neighbours of any vertex is $2K_2$ -free. Therefore, each of $G[A_1 \cup B_1]$, $G[B_1 \cup C_1]$ and $G[C_1 \cup A_1]$ is $2K_2$ -free, since they are induced by non-neighbours of a_2, u, w , respectively (where a_2 is an arbitrary vertex in A_2 , which exists because A_2 is not empty). We do not need to worry about triangles, since they are forbidden anyway. Finally, if there was an anti-triangle induced by $a_1 \in A_1, b_1 \in B_1, c_1 \in C_1$, then together with u and any vertex $a_2 \in A_2$ they would induce a $2K_2 + K_1$. This shows that $G[A_1 \cup B_1 \cup C_1]$ is nice. The other subgraph is treated analogously. Therefore, by Lemma 38 G is nice. \square

3.2.2 Geometric griddability versus bounded lettericity

We first show that geometric griddability of a class \mathcal{X} of permutations implies bounded lettericity of the corresponding class $\mathcal{G}_{\mathcal{X}}$ of permutation graphs. This uses the map φ defined at the end of Subsection 3.1.2.

Theorem 40. *Let \mathcal{X} be a class of permutations that is geometrically griddable by a partial multiplication matrix M . Then the corresponding class $\mathcal{G}_{\mathcal{X}}$ of permutation graphs has bounded lettericity.*

Proof. It suffices to prove this in the case $\mathcal{X} = \text{Geom}(M)$. For this, we let Σ be the cell alphabet of $M = (\alpha_{ij})$, and consider the surjective mapping $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$ defined earlier. To get $|\Sigma|$ -letter representations for the permutation graphs of the permutations in $\text{Geom}(M)$, we carefully construct a decoder \mathcal{D} with vertex set Σ . For each permutation $\pi \in \text{Geom}(M)$, we then pick any element $w \in \varphi^{-1}(\pi)$. It is then routine to check that the letter graph $G(\mathcal{D}, w)$ is in fact isomorphic to G_π . It remains to show how the decoder \mathcal{D} is constructed.

As in Figure 3.10, we observe that two points x_i and x_j of a permutation $\pi \in \text{Geom}(M)$ correspond to a pair of adjacent vertices in G_π if and only if one of them lies to the left and above the second one in the plane. Therefore, if

- $\alpha_{ij} = 1$, then the points lying in the cell Z_{ij} form an independent set in the permutation graph of π . Therefore, we do not include the pair (a_{ij}, a_{ij}) in \mathcal{D} .
- $\alpha_{ij} = -1$, then the points lying in the cell Z_{ij} form a clique in the permutation graph of π . Therefore, we include the pair (a_{ij}, a_{ij}) in \mathcal{D} .
- two cells Z_{ij} and Z_{kl} are independent with $i < k$ and $j < l$, then no point of Z_{ij} is adjacent to any point of Z_{kl} in the permutation graph of π . Therefore, we include neither (a_{ij}, a_{kl}) nor (a_{kl}, a_{ij}) in \mathcal{D} .
- two cells Z_{ij} and Z_{kl} are independent with $i < k$ and $j > l$, then every point of Z_{ij} is adjacent to every point of Z_{kl} in the permutation graph of π . Therefore, we include both pairs (a_{ij}, a_{kl}) and (a_{kl}, a_{ij}) in \mathcal{D} .
- two cells Z_{ij} and Z_{kl} share a column, i.e., $i = k$, then we look at the sign (direction) c_i associated with this column and the relative position of the two cells within the column.
 - If $c_i = 1$ (i.e., the column is oriented from left to right) and $j > l$ (the first of the two cells is above the second one), then only the pair (a_{ij}, a_{il}) is included in \mathcal{D} .
 - If $c_i = 1$ and $j < l$, then only the pair (a_{il}, a_{ij}) is included in \mathcal{D} .
 - If $c_i = -1$ (i.e., the column is oriented from right to left) and $j > l$ (the first of the two cells is above the second one), then only the pair (a_{il}, a_{ij}) is included in \mathcal{D} .
 - If $c_i = -1$ and $j < l$, then only the pair (a_{ij}, a_{il}) is included in \mathcal{D} .
- two cells Z_{ij} and Z_{kl} share a row, i.e., $j = l$, then we look at the sign (direction) r_j associated with this row and the relative position of the two cells within the row.

- If $r_j = 1$ (i.e., the row is oriented from bottom to top) and $i < k$ (the first of the two cells is to the left of the second one), then only the pair (a_{kj}, a_{ij}) is included in \mathcal{D} .
- If $r_j = 1$ and $i > k$, then only the pair (a_{ij}, a_{kj}) is included in \mathcal{D} .
- If $r_j = -1$ (i.e., the row is oriented from top to bottom) and $i < k$, then only the pair (a_{ij}, a_{kj}) is included in \mathcal{D} .
- If $r_j = -1$ and $i > k$, then only the pair (a_{kj}, a_{ij}) is included in \mathcal{D} .

□

This shows that geometrical griddability of a permutation class implies bounded lettericity of the corresponding class of permutation graphs. In [Ale+20b], we conjectured the converse – that if the class of permutation graphs corresponding to a class of permutations has bounded lettericity, then that class of permutations must be geometrically griddable. This converse direction proved to be a bit more tricky than expected; we will show the general case in Section 3.3. Before that, let us show it in the case where the lettericity is (at most) 2.⁶

Theorem 41. *Let \mathcal{X} be a class of permutations and $\mathcal{G}_{\mathcal{X}}$ the corresponding class of permutation graphs. If $\mathcal{G}_{\mathcal{X}}$ is a class of 2-letter graphs, then \mathcal{X} is geometrically griddable.*

Proof. Let $\Sigma = \{a, b\}$, and fix a decoder \mathcal{D} . Consider a graph $G_{\pi} \in \mathcal{G}_{\mathcal{X}}$ and represent it by a word over Σ with the decoder \mathcal{D} .

Assume first that \mathcal{D} contains either both of (a, b) and (b, a) , or none of them. Then we have either all possible edges between the set of vertices of G_{π} labelled by a and the set of vertices of G_{π} labelled by b , or none of them. It is then not difficult to see that in the first case, \mathcal{X} is contained in the geometric grid class of the matrix on the left, and in the second case, \mathcal{X} is contained in the geometric grid class of the matrix on the right:

$$\begin{pmatrix} m_a & 0 & 0 \\ 0 & m_b & 0 \\ 0 & 0 & m_a \end{pmatrix} \text{ and } \begin{pmatrix} 0 & 0 & m_b \\ 0 & m_a & 0 \\ m_b & 0 & 0 \end{pmatrix},$$

⁶As a warning to the reader prone to disappointment, we mention that the ad-hoc proof for lettericity 2 that we are about to present does not seem to generalise nicely. We have chosen not to omit it from the text because first, it provides a chronologically accurate account of our progress on the problem, and second, it gives a concrete, hands-on example that might help the reader gain more familiarity with the notions in play. The reader who is not looking for those things may skip it without fear of losing insight into the general case.

where m_a (resp. m_b) denotes either 1 if $(a, a) \notin \mathcal{D}$ (resp. $(b, b) \notin \mathcal{D}$) or -1 if $(a, a) \in \mathcal{D}$ (resp. $(b, b) \in \mathcal{D}$).

Now suppose only one of (a, b) and (b, a) is in \mathcal{D} . Without loss of generality assume it is (a, b) , since the other case is similar.

If only one of (a, a) and (b, b) is in \mathcal{D} , then G_π is a threshold graph. In this case, π can be placed in the figure of

$$\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$$

known as the \times -figure. Indeed, according to Proposition 5.6.1 in [Wat07], a permutation can be placed in the \times -figure if and only if it avoids 2143, 3412, 2413 and 3142. The first two of these permutations have permutation graphs $2K_2$ and C_4 , while the last two both have P_4 as their permutation graph. Since a graph is threshold if and only if it is $(P_4, C_4, 2K_2)$ -free, we conclude that π can be placed in the \times -figure, since G_π is threshold.

The cases when either both or none of (a, a) and (b, b) belong to \mathcal{D} are complementary to each other. Therefore, we may assume without loss of generality that none of them belongs to \mathcal{D} . Then G_π is a chain graph, and hence it is K_3 and $2K_2$ -free. Hence π avoids 321 and 2143. It is known [Atk99] that the class of permutations avoiding 321 and 2143 is the union of two classes: the class A_1 avoiding 321, 2143 and 3142, and the class A_2 avoiding 321, 2143 and 2413.

A short case analysis shows that any permutation in A_1 can have at most one drop (i.e., two consecutive elements such that the first one is larger than the second one), hence it can be placed in the figure of $\begin{pmatrix} 1 & 1 \end{pmatrix}$. Similarly, any permutation in A_2 consists of two increasing subsequences such that all the elements of one of them are greater than every element of the other, hence it can be placed in the figure of $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$.

The difference between the two classes can be illustrated as follows. For the class A_1 , the word representing G_π as a 2-letter graph can be read at the top of the diagram representing π (see the left diagram in Figure 3.11), while for the class A_2 , this word can be read at the bottom of the diagram (see the right diagram in Figure 3.11).

□




Figure 3.11: The diagrams of two permutations π such that G_π is the graph of the word $ababab$ with $\mathcal{D} = \{(a, b)\}$.

3.3 Bounded lettericity implies geometric griddability

While the implication from geometric griddability to bounded lettericity is relatively straightforward, a proof for the converse direction had eluded us for some time. One of the reasons behind this is that it is not clear how to construct all permutations with a fixed permutation graph $G \cong G(\mathcal{D}, w)$ by using the given letter graph representation of G .⁷

In order to overcome those difficulties, we had to adjust our perspective slightly. This enabled us to develop a framework which allows a conceptually easy (if a bit messy) proof that bounded lettericity implies geometric griddability. Our framework is inspired by the ideas described in [VW11, Section 3] – however, the scope of [VW11] is limited to certain monotone classes of permutations. Our present work provides a substantially more elaborate iteration of those ideas, and adapts them to a more general setting that includes all monotone classes of permutations, as well as graph classes of bounded lettericity. Our emerging point of view also shows some potential in tackling further problems regarding geometric griddability, such as characterising it via minimal obstacles. We recognise that our perspective and the new definitions that come with it are not yet completely optimised: one might rightfully question in places why we have made certain choices over others

⁷Enumerating the permutations with a given permutation graph without any further conditions should be feasible – we did not look too much into it. The problem is that if we do it without using the letter graph representation of the permutation graph, we are effectively ignoring the information that is supposed to imply geometric griddability in the first place.

concerning notation and terminology. We ask the reader to indulge us in this regard for the time being. The reader deserves, of course, a more concrete explanation of why we believe this new perspective is the “right tool for the job”; we postpone that discussion until the beginning of Subsection 3.4.2, when the nature of the “tool” and of the “job” will be clearer.

3.3.1 Locally ordered hypergraphs

The main insight in this subsection is that letter graph expressions on the one hand, and monotone gridings of permutations by a partial multiplication matrix on the other, share a common structure. This structure can be expressed in terms of the existence of “local orders” on appropriate subsets of the elements. Let us illustrate what we mean by this:

Example 42. Let M be a $0/\pm 1$ $s \times t$ partial multiplication gridding matrix, and let π be a permutation together with a monotone M -gridding. Write r_1, \dots, r_s and c_1, \dots, c_t for the rows and columns of M . Then for each r_i and c_j , we get a linear order \leq_{r_i} and \leq_{c_j} respectively on the elements from that row or column. It is given by the order in which the elements of π appear according to the direction corresponding to r_i or c_j (we will refer to this order as the *matrix order*, or the order *induced by M*). Since the matrix is monotone, \leq_{r_i} and \leq_{c_j} must agree on the points of π appearing in cell A_{ij} . For instance, in Figure 3.9, we have:

- $1 \leq_{r_1} 2 \leq_{r_1} 3$;
- $6 \leq_{r_2} 5 \leq_{r_2} 4$;
- $6 \leq_{c_1} 1 \leq_{c_1} 4$;
- $2 \leq_{c_2} 5 \leq_{c_2} 3$.

Example 43. Let Ω be a finite alphabet, and \mathcal{D} be a decoding digraph. Let w be a word with entries in Ω . For any $a, b \in \Omega$, we get a linear order \leq_{ab} defined on the set $\{x \in V(G(\mathcal{D}, w)) : w_x = a \text{ or } w_x = b\}$. The order is simply the one induced from w . In particular, for any fixed letter a and any other letter l , all orders \leq_{al} agree on the letter class of a (that is, on the set $\{x \in V(G(\mathcal{D}, w)) : w_x = a\}$).

Those two examples are similar, but they have different roles. The first example illustrates the local order structure coming from a (not necessarily geometrical) gridding matrix M . There is no guarantee that there is a global linear order *on the whole permutation* that restricts to each of the \leq_{r_i} and \leq_{c_j} on the appropriate

subsets.⁸ The second example shows instead a situation in which such a global order already exists, and the local orders described simply come from it. In this light, we can imagine the two examples as “before and after” pictures of what we are trying to achieve. Indeed, the main theme throughout this section will consist of being given some orders defined locally, and studying when and how they can be “glued together” into a globally defined linear order. As it turns out, in the case of permutations, this is possible exactly when the M -gridding can be made geometrical. Before we show this, we start with some definitions whose aim is to abstract the structure described above.

Definition 44. A *locally ordered hypergraph* (“LOH” for short) \mathcal{H} is a hypergraph (X, E) with no isolated vertices, where any hyperedge $e \in E$ has a linear order \leq_e on its elements, which we call the *local order* of e . In addition, we have the following

local consistency condition: for any edges $e, e' \in E$, \leq_e and $\leq_{e'}$ agree on $e \cap e'$.

We further define an equivalence relation on the elements of X : $x \sim y$ if they have the same sets of incident edges. The equivalence classes are called the *cells* of \mathcal{H} .

To accommodate the structure coming from griddings by non-partial multiplication matrices, we also propose the following definition:

Definition 45. A *semi-LOH* (“sLOH” for short) is like a LOH, except the local consistency condition is replaced with the following slightly weaker one:

local semi-consistency condition: on any cell $A_{e_{i_1}, \dots, e_{i_r}}$, any two linear orders induced by $\leq_{e_{i_1}}, \dots, \leq_{e_{i_r}}$ either agree, or are reverses of each-other.

In particular, any LOH is also a sLOH.

Definition 46. A LOH \mathcal{H} has the *global consistency property* (or is *globally consistent*) if there exists a linear order \leq on its vertices which restricts to \leq_e on each hyperedge.

An alternative way of defining global consistency is as follows: construct a directed graph on X , with arcs (x, y) for any elements $x \neq y$ with $x \leq_e y$ for some e . Call this digraph the *conflict graph* $\text{conf}(\mathcal{H})$ of \mathcal{H} . Then using topological sorting, \mathcal{H} is globally consistent if and only if its conflict graph is acyclic. We note that

⁸In the particular case from Figure 3.9, since the gridding is actually geometrical, such an order does in fact exist. Indeed, the distance-from-distinguished-corner order given by $6 \leq 1 \leq 2 \leq 5 \leq 4 \leq 3$ is an example.

conflict graphs can also be constructed in the same way for sLOHs (although some subsets might simply have all possible arcs between them).⁹

Definition 47. Let \mathcal{H} and \mathcal{K} be LOHs with labelled hyperedges. A *LOH isomorphism* $\phi : \mathcal{H} \rightarrow \mathcal{K}$ is a strong hypergraph isomorphism that preserves the local orders. We say \mathcal{H} and \mathcal{K} are *isomorphic as LOHs* (written $\mathcal{H} \cong \mathcal{K}$) if there is a LOH isomorphism between them.

As illustrated in Example 42, a permutation π with a gridding by a partial multiplication matrix M gives rise to a LOH $\mathcal{H}_M(\pi)$, each of whose hyperedges consists of the entries of π lying in a row or a column of M . The local orders \leq_e are given by the directions associated with the row or column corresponding to e , and the cells of $\mathcal{H}_M(\pi)$ correspond to the cells of M containing elements of π . Notice that $\mathcal{H}_M(\pi)$ depends on π , M , and the gridding of π by M .

We would now like to show that, as expected, permutations with griddings by a fixed matrix M are uniquely determined by their LOHs. For this, let π be a permutation together with a gridding consistent with M . Let us consider the elements in row i of the gridding. Denote them by x_1, \dots, x_l , ordered in the direction associated with that row. Denote by ρ_π^i the sequence of indices of the columns of x_1, \dots, x_l in that order. Construct sequences γ_π^i for the columns analogously. As an example, letting $\pi = 614253$ with the gridding from Figure 3.9, we have $\rho_\pi^1 = (1, 2, 2)$, $\rho_\pi^2 = (1, 2, 1)$, $\gamma_\pi^1 = (2, 1, 2)$ and $\gamma_\pi^2 = (1, 2, 1)$ (recall the rows are indexed from bottom to top).

Lemma 48. *Let π_1, π_2 be two permutations with griddings compatible with an $s \times t$ partial multiplication matrix M . Suppose π_1 and π_2 have the same number of elements in each cell of the gridding, and that $\rho_{\pi_1}^i = \rho_{\pi_2}^i$ and $\gamma_{\pi_1}^j = \gamma_{\pi_2}^j$ for all $1 \leq i \leq s$ and $1 \leq j \leq t$. Then π_1 and π_2 are the same permutation.*

Proof. The proof is by induction on the number of elements of the permutations. The base case is trivial. Let x_1, x_2 be the topmost elements in π_1 and π_2 respectively (say they lie in row i of the gridding). The fact that $\rho_{\pi_1}^i = \rho_{\pi_2}^i$ implies x_1 and x_2 are in the same column j , and therefore in the same cell (i, j) . Let π'_1 and π'_2 be the permutations we get by removing x_1 and x_2 from π_1 and π_2 respectively. These new permutations satisfy the conditions of the lemma. Indeed, the first condition is clearly satisfied, while for the second, we note that, for $\alpha = 1, 2$, the sequence $\rho_{\pi'_\alpha}^i$ is $\rho_{\pi_\alpha}^i$ with either the first or the last entry removed (depending on the direction

⁹We will forget about sLOHs for now – they are only briefly useful in the next subsection as “placeholders”, until we show that a certain sLOH is, in fact, a LOH.

of row i), while $\gamma_{\pi'_\alpha}^j$ is $\gamma_{\pi_\alpha}^j$ with either the first or the last occurrence of i removed (depending on the direction of column j). Hence the permutations with one point removed are equal; but x_1 and x_2 also have the same position in the horizontal order (the one corresponding to either the first or last occurrence of i in $\gamma_{\pi_\alpha}^j$). Hence π_1 and π_2 themselves must be equal. \square

Lemma 49. *Let π_1 and π_2 be two permutations with griddings compatible with a partial multiplication matrix M . Suppose $\mathcal{H}_M(\pi_1) \cong \mathcal{H}_M(\pi_2)$ are strongly isomorphic, with the hyperedge labellings coming from the rows and columns of M . Then $\pi_1 = \pi_2$.*

Proof. Note that, due to the strong isomorphism of the LOHs, the permutations have the same number of elements in each cell, and that $\mathcal{H}_M(\pi_\alpha)$ ($\alpha = 1, 2$) uniquely determines the $\rho_{\pi_\alpha}^i$ and $\gamma_{\pi_\alpha}^j$. Indeed, $\rho_{\pi_\alpha}^i$ can be obtained from $\mathcal{H}_M(\pi_\alpha)$ as the sequence of column hyperedges to which the vertices in the i th row hyperedge belong, in the order given by \leq_{r_i} , and similarly for the $\gamma_{\pi_\alpha}^j$. By Lemma 48, we are done. \square

We are now ready to show that geometric griddability of a permutation is the same as global consistency of its LOH.

Lemma 50. *Let M be a partial multiplication gridding matrix. Suppose π is a permutation monotonically griddable by M . Then $\mathcal{H}_M(\pi)$ is globally consistent if and only if π is geometrically griddable by M .*

Proof. If π is geometrically griddable by M , then any preimage of π by the map φ described in Subsection 3.1.2 gives the required linear order on the LOH.

For the converse, let \leq be a global order on $\mathcal{H}_M(\pi)$ that agrees with all the local orders \leq_e . We construct a new permutation π' as follows: we start by putting points in the standard figure of M one by one, in the order given by \leq . The points are placed on the main diagonals, in the same cells as the corresponding points from π , with increasing distance from the distinguished corners of their respective cell. Clearly, π' is a permutation geometrically gridded by M . If we show that it is actually equal to π , we are done. By Lemma 49, it suffices to show that $\mathcal{H}_M(\pi') \cong \mathcal{H}_M(\pi)$. But this follows by construction: the LOHs are strongly isomorphic as hypergraphs, and the orders on the hyperedges are given by the \leq_e for π , and by the corresponding subsets of \leq for π' , which by assumption agree. \square

3.3.2 Showing bounded lettericity implies geometric griddability

At the heart of our result is the following fact: if a permutation π has a monotone gridding by a (not necessarily partial multiplication) matrix M and a permutation graph G_π of lettericity k , then π is geometrically griddable by a partial multiplication matrix M' whose size only depends on the size of M and on k . The proof is somewhat technical and notation-heavy. However, the concept behind it is very transparent, so we will begin by describing it.

The permutation π has two sLOHs associated with it: one coming from its gridding by M , as in Example 42, and one LOH coming from the letter graph representation of G_π , as in Example 43. As noted in the previous section, the letter graph LOH is globally consistent by assumption; Lemma 50 tells us that what we are really after is global consistency of the matrix sLOH.

The problem is that those two sLOHs will in general differ. There is, however, a straightforward solution: we take a “common refinement” of them. The idea is to modify in tandem the matrix and the decoder so that, while still representing the same permutation, the information coming from the two refined sLOHs now coincides (and in particular, the new matrix sLOH is now a LOH, i.e., the gridding is by a partial multiplication matrix). This allows us to conclude that the new matrix LOH is, as required, globally consistent.

In practice, this is implemented as follows: in the gridding of π by M , we add some vertical and horizontal lines originating from the letter classes of G_π . This produces a gridding by a new matrix M' , which we use in turn to split the old letter classes and produce a new letter graph representation of G_π . After some further modifications, the resulting LOHs agree, up to flipping the orders on connected components of the cell graph of the new matrix. Since the size of M' will depend, by construction, only on M and k , this turns out to be enough for our purposes.

Notation and set-up

Our proof has a lot of moving parts. We therefore devote this sub-subsection to setting up the necessary notation and language, and to establishing some preliminary results.

Throughout the rest of this section, π will denote a fixed permutation, with permutation graph G_π . We will assume that π admits a gridding by a fixed $s \times t$ matrix M . We will also assume that $G_\pi \cong G(\mathcal{D}, w)$, where w is a fixed word over some alphabet Ω of size k , and \mathcal{D} is a decoder. As described above, from these objects associated with π , we will produce new ones. Those new objects will be denoted by

the addition of a prime symbol to the old objects' notation – for instance, the new matrix by which we will grid π will be denoted by M' .

Given a letter graph representation $G(\mathcal{E}, z)$ of G_π , for a letter a , we denote by $\text{let}(\mathcal{E}, z, a)$ the letter class of a , i.e., the vertices of G_π (and, by extension, the elements of π) with letter a in z . When the letter graph representation for G_π that we are using is clear from the context, we will just write $\text{let}(a)$ instead of $\text{let}(\mathcal{E}, z, a)$.

For a cell Z in a gridding of π by a matrix, we will denote by π_Z the set of points of π lying in Z .

The next remark shows that we may work with a somewhat simplified set-up.

Remark 51. We may assume that each letter of Ω occurs in at most one cell of M . Indeed, note that if a letter a occurs in more than one cell of the gridding by M , we may replace it with letters a_1, a_2, \dots, a_r where each a_i appears in a single cell, and r is bounded above by the total number of cells. Hence the size of the new alphabet is bounded above by kst , which is a function of s, t and k . The decoder can easily be modified accordingly.

In view of this remark, we will write $\text{cel}_N(a)$ for the (from now on by assumption unique) cell of a matrix N in which letter a appears; if the matrix N is clear from context, we will simply write $\text{cel}(a)$. Moreover, it might happen that a letter has an empty letter class in some letter graph representation; we will say that such a letter is *empty*.

Remark 52. We assume without loss of generality that in the letter graph representation $G(\mathcal{D}, w)$, there are no empty letters in Ω .

Finally, we will call two letters a and b :

- *N-cellmates* if $\text{cel}_N(a) = \text{cel}_N(b)$;
- *N-collinear* if $\text{cel}_N(a) \neq \text{cel}_N(b)$, but $\text{cel}_N(a)$ and $\text{cel}_N(b)$ share a row or a column;
- *N-unrelated* otherwise.

As usual, we will omit the N when the matrix is clear from context.

We call a pair of letters a, b in a decoder \mathcal{E} *symmetric* if either none or both of (a, b) and (b, a) are in \mathcal{E} , and *asymmetric* otherwise. For two distinct letters a and b , their letter classes (and, by extension, the letters themselves) are *complete* to each other if every vertex in $\text{let}(a)$ is adjacent to every vertex in $\text{let}(b)$, and the letter classes are *anticomplete* to each other if there are no edges between the vertices in $\text{let}(a)$ and the vertices in $\text{let}(b)$.

Remark 53. Without loss of generality, we may further simplify the set-up by assuming that the pair a, b is symmetric in \mathcal{D} whenever $\text{let}(a)$ and $\text{let}(b)$ are complete or anticomplete to each other in $G(\mathcal{D}, w)$. Conversely, if $\text{let}(a)$ and $\text{let}(b)$ are not complete or anticomplete to each other, it must be that the pair a, b is asymmetric. In addition, the same assumptions can be made about any letter graph representation $G(\mathcal{E}, z)$ of G_π .

As described at the beginning of this subsection, our plan is to show that, after a suitable refinement, the letter graph LOH and the matrix LOH are more or less the same. As such, we need to introduce some language that enables us to describe this statement.

Given a letter representation $G(\mathcal{E}, z)$ of G_π , we will denote by \leq_z the linear order on the elements of π coming from z . We will refer to this order as a *word order*. Given a gridding of π by a partial multiplication matrix N , we will denote by \leq_N the relation obtained by taking the union of the row and column orders in the matrix. Note that this relation is antisymmetric on π , but it is *not* an order relation in general; however, abusing notation, we will refer to it as a *local matrix order*. Given an order relation \preceq its *reverse* \preceq^r is defined by $x \preceq^r y$ if and only if $y \preceq x$.

Definition 54. Suppose π has an N -gridding and G_π has a letter graph representation $G(\mathcal{E}, z)$. Let $a \neq b$ be letters in \mathcal{E} . We say that $\{a, b\}$ is (N, z) -*forward* if

$$\leq_N \cap (\text{let}(a) \times \text{let}(b)) = \leq_z \cap (\text{let}(a) \times \text{let}(b))$$

and (N, z) -*backward* if

$$\leq_N \cap (\text{let}(a) \times \text{let}(b)) = \leq_z^r \cap (\text{let}(a) \times \text{let}(b)).$$

In other words, a pair of distinct letters is forward if the local matrix order agrees with the word order when comparing elements from the two letter classes, and backward if the local matrix order agrees with the reverse word order.

Similarly, we say the individual letter a is (N, z) -*forward* (respectively (N, z) -*backward*) if the local matrix order agrees with the word order (respectively the reverse word order) on the letter class $\text{let}(a)$.

When the matrix and letter graph representation are clear from context, we may drop the “ (N, z) -”.

A priori, a given letter or pair of letters does not need to be forward or backward – the matrix order and the word order might agree in some parts of

the letter classes, but not others. Of course, as one might expect, the nature of the problem gives us some control over when that can happen, as shown in the following lemma.

Lemma 55. *Let a, b be distinct asymmetric letters in a letter graph representation $G(\mathcal{E}, z)$ of G_π . Then the pair $\{a, b\}$ is either (M, z) -forward or (M, z) -backward.*

Proof. By our Remark 53, it suffices to restrict ourselves to M -collinear pairs of letters (since all other pairs of letters are symmetric). Assume without loss of generality that the cells of a and b share a column in M , with the cell of a above the cell of b , and that $(a, b) \in \mathcal{E}$. Moreover, suppose that the direction associated with the column is left to right. The $\text{let}(b)$ -neighbours in G_π of a point $x \in \text{let}(a)$ are all the points in $\text{let}(b)$ lying to the right of x , in other words, all the points in $\text{let}(b)$ succeeding x in the column's order. Those points are also exactly the points $y \in \text{let}(b)$ such that $y \geq_z x$. □

The refinement of the matrix and decoder

Let us now describe the construction of the new matrix M' in which we are planning to geometrically grid π . This construction is done graphically, by adding some vertical and horizontal lines to the M -gridding of π . For each letter a , we add 4 lines to the figure: a horizontal line just before¹⁰ the first element of $\text{let}(a)$ in the row order, a horizontal line just after the last element of $\text{let}(a)$ in the row order, a vertical line just before the first element of $\text{let}(a)$ in the column order, and a vertical line just after the last element of $\text{let}(a)$ in the column order. The lines thus induced by a letter a split the row of $\text{cel}(a)$ into three horizontal stripes that we will denote by R_a^1, R_a^2 and R_a^3 , and the column of $\text{cel}(a)$ into three vertical stripes C_a^1, C_a^2 and C_a^3 (with indices increasing in the row or column's direction – see Figure 3.12 for an illustration).

Notation. For any letter b sharing a row with a , we write $\text{let}(b)_a^t := \text{let}(b) \cap R_a^t$ for any $1 \leq t \leq 3$. Similarly, for any letter b sharing a column with a , we write $\text{let}(b)_a^t := \text{let}(b) \cap C_a^t$. We note that this is well-defined if b is a cellmate of a , since $\text{let}(b) \cap R_a^t = \text{let}(b) \cap C_a^t$ in that case. In particular, we have $\text{let}(a)_a^t = \text{let}(a)$ if $t = 2$, and \emptyset otherwise.

Adding those lines for every letter indeed produces, in the obvious way, a gridding of π by a larger matrix M' , whose size depends only on M and k , and

¹⁰By “just before”, we mean before, but after any element preceding it.




Figure 3.12: Lines induced by the letter class of a (drawn in red)

whose row and column orders are inherited from the ones in M . As stated at the beginning of the subsection, this geometric operation now induces a refinement of the decoder. Specifically, the lines we added to the figure suggest to us another alphabet Ω' , decoder \mathcal{D}' and word w' representing G_π as a letter graph, obtained by splitting the original letter classes. Let us describe explicitly how this new letter graph description is constructed.

Every cell Z of the M -gridding has several new lines going through it, some vertical and some horizontal. This set of lines induces a partition of the elements of π_Z into non-empty intervals with respect to the matrix order. Similarly, for each letter a with $\text{cel}(a) = Z$, the lines partition $\text{let}(a)$ into non-empty intervals. This is illustrated in Figure 3.13, where the numbers indicate the cells of M' containing those successive intervals, in increasing order.

For each letter $a \in \Omega$ with $\text{cel}(a) = Z$, the new alphabet Ω' contains letters a_1, a_2, \dots where $\text{let}(a_i)$ is the i th interval in the partition described above. The new word w' is obtained from w in the obvious way, by replacing appearances of the old letters with the appropriate new ones.¹¹ It is clear that, for any new letter a_i ,

¹¹This is, indeed, the obvious candidate for the new word w' . We will work with this word for




Figure 3.13: Lines through a cell Z

$\text{cel}_{M'}(a_i)$ is unique and thus well-defined. Moreover, by construction, we have the following:

Remark 56. For each pair of M -collinear letters or M -cellmates a, b , each letter class $\text{let}(a_i)$ in the partition of $\text{let}(a)$ into intervals is contained in one of the sets $\text{let}(a)_b^t$ ($t = 1, 2$ or 3). In addition, if a_i and b_j are M' -collinear or M' -cellmates, then $\text{let}(a_i) \subseteq \text{let}(a)_b^t$ and $\text{let}(b_j) \subseteq \text{let}(b)_b^t$ for the *same* t , which must thus be equal to 2.

The new decoder \mathcal{D}' is obtained as follows. Any pair of M' -cellmates or M' -unrelated letters is complete or anticomplete. We make these pairs symmetric in \mathcal{D}' in the obvious way. We now turn to pairs of M' -collinear letters. We note, from Remark 56, that an M' -collinear pair a_i, b_j cannot originate from a pair a, b that is complete or anticomplete, since $\text{let}(a_i) \subseteq \text{let}(a)_b^2$, and thus there are a 's between the first and last b 's. By Remark 53, the pair a, b is asymmetric; we then make the pair a_i, b_j asymmetric, letting $\{a_i, b_j\}$ inherit its orientation in \mathcal{D}' from the orientation of $\{a, b\}$ in \mathcal{D} .

now, however, we remark here that we will perform some minor modifications to it in Lemma 59, without changing the letter graph represented by it.

Remark 57. By Lemma 55, it follows that any pair of M' -collinear new letters is either (M, w') -forward or (M, w') -backward. Since the local matrix order on M' is inherited from the one on M , any such pair is in fact either (M', w') -forward or (M', w') -backward. It is not difficult to see that forwardness or backwardness of the new pair is inherited from the old pair.

Proof of the main result

In this sub-subsection, we complete the proof of the result. Our first step is to build up on Remark 57 and show that forwardness and backwardness satisfy a transitivity-like property in the new decoder \mathcal{D}' and matrix M' .

Lemma 58. *Let $a_i, b_j, c_l \in \Omega'$ be three letters such that a_i, b_j are M' -collinear and b_j, c_l are M' -collinear. Then a_i, b_j and b_j, c_l are either both (M', w') -forward or both (M', w') -backward.*

Proof. Suppose without loss of generality that a_i, b_j is (M', w') -forward. We need to show that the pair b_j, c_l is also (M', w') -forward. As in the end of the last section, we note that a, b and b, c are each asymmetric, and from Remark 57, a, b is (M, w) -forward.

We claim that w must contain a subword $bacb$ or $bcab$. Indeed, by Remark 56, $\text{let}(a_i) \subseteq \text{let}(a)_b^2$ and $\text{let}(c_l) \subseteq \text{let}(c)_b^2$, hence a and c each occur at least once between the first and last occurrences of b in the local matrix order, and thus in w . Suppose $bacb$ is a subword of w (the case for $bcab$ is identical), and write it as $b'acb''$ to distinguish between the initial and final copy of b . Note that forwardness of a, b implies that $b' \leq_M a$ and $a \leq_M b''$. Since the local matrix order restricts to a partial order on the row or column containing a and b , we obtain $b' \leq_M b''$. By symmetry, backwardness of b, c would imply $b'' \leq_M b'$. Since $b' \neq b''$, that is impossible, hence b, c is (M, w) -forward, from which b_j, c_l is (M', w') -forward as required. □

Lemma 58 immediately implies that for a given letter a_i , the collinear pairs containing a_i are either all (M', w') -forward or all (M', w') -backward. Call such a letter (M', w') -quasi-forward or (M', w') -quasi-backward respectively. A letter with no letters collinear to it is vacuously both. A posteriori, this implies that M' must be a partial multiplication matrix. Indeed, there can be no cell where the row and column orders are reverses of each other, since they must both either agree with the forward or the reverse of the order from w' . It follows that the sLOH coming from the gridding of π by M' is in fact a LOH.

Note that the cellmates of a_i are either all quasi-forward or quasi-backward (if there is a letter collinear to a_i , this follows from Lemma 58, otherwise it follows by vacuity). We thus extend this definition to cells: call a non-empty cell (M', w') -*quasi-forward* if the letters in it are quasi-forward, and (M', w') -*quasi-backward* if it is not quasi-forward. Thus every cell is either quasi-forward or quasi-backward.

Quasi-forwardness and quasi-backwardness are *almost* the properties we require of our cells. Indeed, it would be great if we could strengthen those properties in the obvious way, defined as follows: a cell Z is *forward* if $\leq_{M'}$ and $\leq_{w'}$ agree on Z , and *backward* if $\leq_{M'}$ and $\leq_{w'}^r$ agree on Z . As it stands, cells do not have to be forward or backward. However, the next lemma shows that we may modify w' slightly so that the required strengthening holds.

Lemma 59. *We may change w' without affecting the letter graph $G(\mathcal{D}, w')$ ¹² such that any quasi-forward cell is forward, and any quasi-backward cell is backward.*

Proof. Without loss of generality, let Z be a quasi-forward cell. We show that, after the appropriate modifications, Z is forward.

Let C be the set of letters collinear to letters in Z , and let w'_1, w'_2, \dots be the successive non-empty intervals of w' appearing strictly between elements of the set $\bigcup_{c_l \in C} \text{let}(c_l)$. Suppose $x, y \in \pi_Z$ such that $x \in w'_s, y \in w'_t$ with $s < t$. Then, a copy of a letter collinear to the letters of x and y appears between x and y , and hence, as in the proof of Lemma 58, quasi-forwardness of Z implies $x \leq_{M'} y$ if and only if $x \leq_{w'} y$.

If x and y are in the same interval w_s , then no letter collinear to the letters of x and y appears between them. In other words, all letters between them form symmetric pairs with the letters of x and y . In this case, we may swap x and y in w' , without changing the letter graph $G(\mathcal{D}', w')$. In general, we can permute the elements of π_Z appearing in each interval as necessary to ensure that $\leq_{M'}$ and $\leq_{w'}$ agree on π_Z in each interval. Repeating this for each cell of M' produces a modified word w' with the desired properties. \square

Using Lemma 59, we assume for the remainder of the section that w' is such that every quasi-forward cell is forward, and every quasi-backward cell is backward. We are now ready to finish our proof:

¹²In particular, these changes do not retroactively interfere with the previous lemmas. This can be checked rigorously by working in a suitable trace monoid over Ω ; in this terminology, what we have shown so far depends only on the trace of w' , and in the current lemma, we merely replace w' with another word from the same trace. For simplicity, we omit the details.

Theorem 60. *Let \mathcal{X} be a class of permutations, and let $\mathcal{G}_{\mathcal{X}}$ be the corresponding class of permutation graphs. If $\mathcal{G}_{\mathcal{X}}$ has lettericity bounded by some constant k , then \mathcal{X} is geometrically griddable.*

Proof. First, we claim that \mathcal{X} is monotonically griddable. Indeed, note that the classes of matchings and co-matchings have unbounded lettericity, hence $\mathcal{G}_{\mathcal{X}} \subseteq \text{Free}(nK_2, \overline{nK_2})$ for some n . This means that \mathcal{X} cannot contain arbitrarily large skew sums of 12 or direct sums of 21, which from [HV06] implies that \mathcal{X} is monotonically griddable by a matrix N : we take this as our fixed matrix M used throughout the section.

We now claim that each permutation in \mathcal{X} is geometrically griddable by some matrix of bounded size. For a fixed permutation π , this matrix is the matrix M' we have constructed in Section 3.3.2. To see that π is indeed geometrically griddable by M' , we use Lemma 50. It thus suffices to show that $\mathcal{H}_{M'}(\pi)$ is globally consistent.

This is a consequence of Lemmas 58 and 59. Together, they immediately imply that on any connected component of the cell graph of M' , $\leq_{M'}$ is a subset of either $\leq_{w'}$ or $\leq_{w'}^r$. Connected components of the cell graph induce connected components of the conflict graph $\text{conf}(\mathcal{H}_{M'}(\pi))$, hence we see that the conflict graph is acyclic as required.

Since there is a finite number of such potential matrices M' of size bounded in terms of s, t and k , \mathcal{X} is geometrically griddable by any matrix containing all matrices of that size as submatrices. \square

3.4 Further directions: characterising geometric griddability

The authors of [Alb+13] leave determining “the precise border between griddability and geometric griddability” as an open problem. One way to achieve this would be by describing minimal obstacles to geometric griddability analogous to the ones for griddability described in Theorem 24. Theorems 40 and 60 show that conceptually, looking for obstacles to geometric griddability is the same as looking for obstacles to bounded lettericity within the class of permutation graphs.¹³ One may then ask how much more difficult it would be to identify minimal obstacles to bounded lettericity *in general*, that is, without restricting ourselves to permutation graphs.

So far, the question is not yet settled: minimal classes of unbounded lettericity have only been identified in very specific settings. Several of those special

¹³In practice, one would have to carefully translate a prospective result from one language to the other, but there is no reason to expect significant difficulties there.

cases will appear later in this thesis; we will identify the obstructions within various classes, namely:

- cographs, in Section 5.2;
- bipartite permutation graphs, in Section 7.1;
- the so-called quasi-chain graphs, in Section 8.2.

In addition to those, the author is only aware of one other such result appearing in the literature. This result is the content of [FV21], in which Ferguson and Vatter identify the obstacles to bounded lettericity among classes with finitely many prime graphs.

In the rest of the current section, we will attempt to move one step closer to the general solution by investigating lettericity in a setting modelled after monotone grid classes of matrices whose cell graph is a cycle. Subsection 3.4.1 is devoted to understanding lettericity within this universe. For the sake of concreteness, we will first present our result – a complete list of minimal classes in this setting – in a purely graph-theoretic language. In Subsection 3.4.2, we examine a tentative formulation of those notions and results in a more abstract LOH setting. Finally, in Subsection 3.4.3, we take a step back and observe the structural hierarchy of graph and permutation classes emerging from our discussion.

Before doing any of that, a good place to start is the following proposition, which provides a sanity check that bounded lettericity can, in fact, be characterised via minimal classes.

Proposition 61. *Let \mathcal{X} be a hereditary class of unbounded lettericity. Then there exists a hereditary class $\mathcal{X}' \subseteq \mathcal{X}$ of unbounded lettericity such that the class $\mathcal{X}' \cap \text{Free}(G)$ has bounded lettericity for any $G \in \mathcal{X}'$.*

Proof. If \mathcal{X} is minimal of unbounded lettericity, we are done. Otherwise, \mathcal{X} contains a graph G such that $\mathcal{X} \cap \text{Free}(G)$ still has unbounded lettericity. Pick a graph $G_0 \in \mathcal{X}_0 := \mathcal{X}$ with this property that has the minimum possible number of vertices, and put $\mathcal{X}_1 := \mathcal{X}_0 \cap \text{Free}(G_0)$. Repeat the process for as long as possible, putting $\mathcal{X}_{k+1} := \mathcal{X}_k \cap \text{Free}(G_k)$, where $G_k \in \mathcal{X}_k$ is a minimum graph such that \mathcal{X}_{k+1} has unbounded lettericity. There are two cases:

- The process stops at some k . This means we have found a subclass $\mathcal{X}_k = \mathcal{X} \cap \text{Free}(G_0, \dots, G_{k-1})$ of unbounded lettericity such that forbidding any further $G \in \mathcal{X}_k$ yields a class of bounded lettericity, and \mathcal{X}_k is the minimal class we were looking for.

- Otherwise, the process goes on forever, and we get an infinite strictly descending chain $\mathcal{X}_0 \supsetneq \mathcal{X}_1 \supsetneq \dots$ of classes, all of which have unbounded lettericity. Let \mathcal{X}_{lim} be their intersection, that is, $\mathcal{X}_{\text{lim}} = \mathcal{X} \cap \text{Free}(G_0, G_1, \dots)$.

Note that this limit class cannot have bounded lettericity. Indeed, suppose \mathcal{X}_{lim} has lettericity at most t . The G_i s are by construction incomparable via the induced subgraph relation (since at each stage, they are chosen to be minimal), and from [Pet02], they all have lettericity bounded above by $2t + 1$ (since for any $v \in G_i, i \in \mathbb{N}, G_i - v$ is in \mathcal{X}_{lim}). This is a contradiction, since classes of bounded lettericity are well-quasi-ordered (Theorem 14).

Moreover, \mathcal{X}_{lim} is minimal of unbounded lettericity. To see this, note first that, by construction, $|G_i| \leq |G_j|$ for $i \leq j$, and $|G_n| \rightarrow \infty$ as $n \rightarrow \infty$ (since there are only finitely many graphs of a given size). Suppose we can forbid $G \in \mathcal{X}_{\text{lim}}$ with $|G| = k$, and we are still left with a class of unbounded lettericity. But then, by construction, G would have appeared in the sequence (G_i) before any graphs of size at least $k + 1$, contradicting $G \in \mathcal{X}_{\text{lim}} = \mathcal{X} \cap \text{Free}(G_1, G_2, \dots)$.

□

With this in mind, we can now ask: what are the minimal classes of unbounded lettericity? One such class, as we have already seen in Example 16, consists of induced matchings.

Remark 62. By symmetry, one sees that the complements of matchings are also an example. In fact, one can produce more examples by taking various complements that correspond intuitively to making changes in the decoder. For instance, let us consider a letter graph expression for, say, a clique co-matched to an independent set. By at most doubling the number of letters, we can obtain an expression that uses each letter in either the clique or the independent set (but not both); making appropriate changes to the decoder, the same word will then express a matching between two independent sets (that is, an induced matching). From here, it is easy to see that the class of graphs consisting of a clique co-matched to an independent set is another minimal class of unbounded lettericity. In the cases we study, one is usually safe to assume that, if a class is minimal of unbounded lettericity, so are the classes obtained from it by complementing homogeneous sets, or edges between homogeneous sets. We will not dwell on the details, since they are immaterial to our discussion. In the remainder of the chapter, we will use the word “complements” to refer to all of the classes obtained from a given class via such complementations.

Of course, this list is far from complete. Let us present some different kinds of obstacles.

3.4.1 Chain circuits

We will work in a locally consistent setting, modelled after monotonically griddable classes by a partial multiplication matrix.¹⁴ Keeping in mind that Remark 62 applies, we will simplify our setting as much as possible in order to illustrate a new type of obstacle to bounded lettericity. In the universe of permutations, this type of obstacle lies between monotone and geometric griddability; compare this with matchings and their complements, which lie “before” monotone griddability. We start with the following definition:

Definition 63. Let $k \geq 3$. A k -chain circuit (k -CC for short) is a graph whose vertex set can be partitioned into k independent sets (or “bags”) A_1, \dots, A_k with indices modulo k , such that:

- the only edges appear between consecutive A_i ;
- the edges between consecutive A_i induce chain graphs;
- if we order the vertices from A_i in decreasing order with respect to their neighbourhoods in A_{i+1} , then that order is increasing with respect to the neighbourhoods in A_{i-1} .


We will refer to $\{A_i : 1 \leq i \leq k\}$ as the *chain circuit partition*, and we will assume the vertices in each bag are ordered as described above. If v comes before w in the order, we write $v \preceq w$, and we say v is “to the left” of w (and w is “to the right” of v). Given a chain circuit, we can also define its *chain circuit complement* by complementing all edges between consecutive bags. It is clear that the graph thus obtained is still a chain circuit.

Remark 64. This definition is a generalisation of the “nice” graphs studied in Section 3.2.


Among k -chain circuits, we look at two subclasses of graphs. First, we introduce some notation:

¹⁴In fact, the obstacles we will describe have a natural interpretation in the setting of permutations. They are structurally similar to the infinite antichains produced in [MV02] and depicted in Figure 2 of the paper.

Notation 65. Let $C_{k,l}$ denote the chain circuit obtained by taking a union of l k -cycles, with the j -th cycle labelled $v_{1,j}, v_{2,j}, \dots, v_{k,j}$ (the first index is modulo k), and by adding edges between vertices $v_{i,m}$ and $v_{i+1,n}$ whenever $m < n$. See Figure 3.14 for two representations of the graph $C_{4,4}$.



(a) Bags in a cycle with chain graphs between them.



(b) Vertical cycles with matchings between them

Figure 3.14: Two representations of the chain circuit $C_{4,4}$.

The classes we are interested in are the class \mathcal{C}_k of all graphs $C_{k,l}$ and their induced subgraphs, and the class $\tilde{\mathcal{C}}_k$ of their CC-complements. Our first result is that \mathcal{C}_k and $\tilde{\mathcal{C}}_k$ have unbounded lettericity:

Theorem 66. \mathcal{C}_k and $\tilde{\mathcal{C}}_k$ have unbounded lettericity.

Proof. Without loss of generality, we may restrict ourselves to letter graph expressions that use each letter in only one of the k bags, since the minimum number of letters across all letter graph expressions is at most a factor of k away from the minimum number of letters in expressions with this property. By Remark 62, it suffices to prove the statement for \mathcal{C}_k . Suppose, for a contradiction, that the lettericity is bounded. If this is the case, then for any $t \in \mathbb{N}$, we can find $N \in \mathbb{N}$ such that t cycles in some word representation of $C_{k,N}$ are represented by the same subword. In particular, those t cycles form an induced $C_{k,t}$ using only k letters. Since this can be done for any t , this shows the lettericity would in fact be bounded by at most k , and we should be able to express any $C_{k,i}$ with one letter per bag. Some quick case analysis shows the only way this would be possible is, up to symmetry, with a

cyclic decoder (a_i, a_{i+1}) (modulo k). However, it is impossible to represent even a single cycle C_k in this way. \square

We now want to show that \mathcal{C}_k and $\widetilde{\mathcal{C}}_k$ are minimal of unbounded lettericity. In fact, we show a stronger result: in the universe of k -CCs, by forbidding any two chain circuits from \mathcal{C}_k and $\widetilde{\mathcal{C}}_k$ respectively, we obtain a class of bounded lettericity.¹⁵ Another way to state this is that \mathcal{C}_k and $\widetilde{\mathcal{C}}_k$ are the only minimal classes of unbounded lettericity among k -CCs. Because of how those classes are defined, it suffices to show this in the case where we forbid $C_{k,i}$ and $\widetilde{C}_{k,j}$ for some $i, j \geq 1$. We will do the proof by induction on $i + j$.

We will use, like in the previous section, a conflict graph. Rather than going through a LOH construction (which we will return to in the next subsection), let us for now define this conflict graph directly from the k -CC.

Definition 67. Let G be a k -chain circuit with CC-partition $\{A_i : 1 \leq i \leq k\}$. The *conflict graph* $\text{conf}(G)$ of G is the directed graph with vertex set $V(G)$, and arcs (v, w) whenever

- $v \in A_i, w \in A_{i+1}$ for some i (modulo k), and $\{v, w\} \in E(G)$, or
- $w \in A_i, v \in A_{i+1}$ for some i (modulo k), and $\{v, w\} \notin E(G)$.

The conflict graph of G gives us a way of describing obstacles to representing G with one letter per bag and a cyclic decoder. Specifically, $\text{conf}(G)$ has an arc from v to w exactly when the entry of v is forced to appear before the entry of w in such a representation. In particular, if such a representation of G does exist, $\text{conf}(G)$ is acyclic.

In fact, as one expects in view of Lemma 50 the converse is also true:

Lemma 68. *Suppose that G is a k -chain circuit and that $\text{conf}(G)$ is acyclic. Then there is a word w on letters a_1, \dots, a_k such that $G = G_{\mathcal{D}^c}(w)$, where \mathcal{D}^c is the cyclic decoder $\{(a_i, a_{i+1}) : i \in \mathbb{Z}/k\mathbb{Z}\}$.*

Proof. By standard results, $\text{conf}(G)$ admits a topological ordering w , i.e., a linear ordering of the vertices such that if (x, y) is an arc, then x comes before y in w . It is routine to check that with a cyclic decoder, w represents G . \square

We are now ready to prove the result. We start with a base case for our induction.

¹⁵Here, by “forbidding chain circuits” we mean that no copy of the underlying graphs appears as an induced subgraph that respects the cyclic ordering of the bags. We will skip the details for now.

Lemma 69. *The subclass of k -chain circuits obtained by forbidding $C_{k,1}$ and $\widetilde{C}_{k,1}$ has bounded lettericity.*

Proof. From Lemma 68, it is enough to show that if G avoids the two graphs, then its conflict graph is acyclic. Suppose not, and find a shortest directed cycle $v_1, \dots, v_t = v_1$, with $v_r \in A_{i_r}$ (i.e., the i_r are the indices of the bags successively visited by the cycle). Let us note a few facts about the sequence i_r .

- (i) $i_{r+1} - i_r = \pm 1 \pmod k$ for $1 \leq r < t$.
- (ii) We may assume without loss of generality that $i_1 = 1$. Indeed, we can permute the labels of the bags cyclically to make sure that this is the case. We may also assume that $i_2 = 2$, since otherwise we can work in \widetilde{G} instead. Indeed, $\text{conf}(\widetilde{G})$ is just $\text{conf}(G)$ with all arcs reversed, and G avoids $C_{k,1}$ and $\widetilde{C}_{k,1}$ if and only if its CC-complement does.
- (iii) For any $j \in \mathbb{Z}/k\mathbb{Z}$, j and $j+1$ appear consecutively at most once (and similarly for $j+1$ and j). Indeed, suppose the cycle visits bags j and $j+1$ in that order twice, through vertices $v \in A_j, w \in A_{j+1}$ the first time, and $v' \in A_j, w' \in A_{j+1}$ the second time. Since $G[A_1 \cup A_2]$ induces a chain graph, and we know $v \sim w$ and $v' \sim w'$ in $\text{conf}(G)$, we must have $v \sim w'$ or $v' \sim w$. In either case, we have a “shortcut” through our cycle, which shows it is not minimal, contrary to our assumption. An analogous argument shows the statement for $j+1$ and j .
- (iv) For any $j \in \mathbb{Z}/k\mathbb{Z}$, j and $j+1$ appear consecutively, or $j+1$ and j do. Indeed, suppose that there is a j such that the cycle has no edge between A_j and A_{j+1} . We may assume, after changing our choice and doing some relabelling if necessary, that $j = k$, and that the cycle does pass through A_1 . Let v_1 be the leftmost vertex of the cycle in A_1 . Label its position in A_1 by 0. Further, label by 0 the position of the leftmost neighbour $v_2 \in A_2$ of v_1 . Proceeding similarly, label by 0 the position of the leftmost neighbour $v_i \in A_i$ of v_{i-1} , for $i \leq k$. Note that if $v \in A_i$ ($1 \leq i < k$) has a non-negative label, then so do all its neighbours in A_{i+1} by construction, and if $v \in A_i$ ($1 < i \leq k$) has a non-negative label, then all its non-neighbours in A_{i-1} have (strictly) positive labels. This means that in our set-up, the cycle cannot actually return to v_1 .

The above observations imply that the sequence of i_r is (up to starting the cycle at another point, working with \widetilde{G} instead of G and relabelling bags if necessary) either:

- $1, 2, \dots, k, 1$. In this case, G contains a $C_{k,1}$ (or a $\widetilde{C}_{k,1}$ if we were working with \widetilde{G} , as described in (ii)).
- $1, 2, \dots, k, 1, k, \dots, 2, 1$, i.e., our cycle goes around the chain circuit, but instead of reaching its starting point v_1 , it reaches another vertex $v' \in A_1$ before looping back around. However, this is impossible: v' must be to the right of v_1 (otherwise the neighbour of v' preceding it in the cycle is also adjacent to v_1 and we can find a shorter cycle), and we can use the same indexing argument as in (iv) to conclude that the cycle cannot return to v_1 .

□

Remark 70. One cannot help but notice a vague and superficial similarity between the arguments of the above proof and certain exercises from basic homotopy theory, such as determining the winding number of loops. This is perhaps a stretch of the imagination, but it would be interesting (and outside our current scope) to look into whether there is something underlying this similarity. If the reader is willing to suspend their disbelief, let us dream together for a moment: what if there is an illuminating topological setting which gives elegant, satisfying interpretations and proofs for all the LOH-related phenomena we are attempting to formalise and explain? It is, of course, possible, that questions related to this in a not necessarily obvious way were already asked and answered – if the reader happens to have any insight or interest in this matter, they are kindly invited to contact the author.

Theorem 71. *Among k -chain circuits, the classes \mathcal{C}_k and $\widetilde{\mathcal{C}}_k$ are the only minimal classes of unbounded lettericity.*

Proof. We have shown the base case in Lemma 69. We need to show how the induction step works. Suppose thus that we have a k -chain circuit G with chain partition $\{A_1, \dots, A_k\}$, and with no $C_{k,p}$ and $\widetilde{C}_{k,q}$ for some $p, q \in \mathbb{N}$. We are going to split our chain circuit in a way which allows us to use the inductive hypothesis. We may assume without loss of generality that G has a $C_{k,1}$ as an induced subgraph (if not, consider its CC-complement and use Remark 62; if its CC-complement also has no $C_{k,1}$, then we are in the base case). Starting with any vertex $v_i \in A_i$ of the cycle, we colour the edge to its leftmost neighbour in A_{i+1} blue. We repeat this process with that leftmost neighbour, and keep doing this until we reach a vertex we have visited before. The process terminates, since any visited vertex in A_i has a neighbour in A_{i+1} (specifically, the appropriate vertex in the cycle we started with). We thus obtain a blue cycle C_b . Similarly, construct a red cycle C_r by starting with

v_i and colouring in red the edge to its rightmost neighbour in A_{i-1} , and like before, repeating the process. See Figure 3.15 for an illustration.

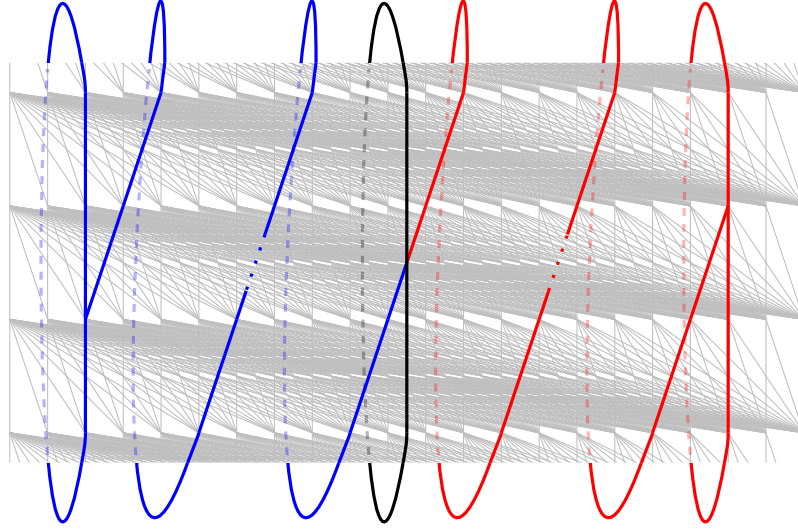


Figure 3.15: The red and blue spirals

Now let G^L be the induced subgraph strictly to the left of the blue cycle, let G^R be the subgraph strictly to the right of the red cycle, and let G^M be the middle subgraph (including the red and blue cycles). That is, for each A_i , G^L contains the vertices of A_i strictly to the left of the vertex in $C_b \cap A_i$, G^R contains the vertices strictly to the right of the vertex in $C_r \cap A_i$, and G^M contains the remaining vertices. Write $G_i^L := G^L \cap A_i$ (extending the notation to G^R and G^M in the obvious way).

Now notice that by construction, each of G^L and G^R is $C_{k,p-1}$ -free, since a $C_{k,p-1}$ in G^L together with the blue cycle would give a $C_{k,p}$, and similarly for a $C_{k,p-1}$ in G^R together with the red cycle. Hence the inductive hypothesis applies, and there is a c (depending on p) such that G^L and G^R can each be represented by a word with c letters.

Moreover, the edges between the three parts are easy to describe: for any i , we have no edges between G_i^R and $G_{i+1}^M \cup G_{i+1}^L$, because of how we constructed our red cycle, and no edges between G_i^L and $G_{i-1}^M \cup G_{i-1}^R$ because of how we constructed the blue cycle. We also have all possible edges between G_i^L and $G_{i+1}^M \cup G_{i+1}^R$, as well as all possible edges between G_i^R and $G_{i-1}^M \cup G_{i-1}^L$ because of the properties of chain circuits. Given these structural features, it is clear that if we have words representing each of G^L , G^M and G^R with c_1, c_2 and c_3 letters respectively, we can construct a word representing G using $c_1 + c_2 + c_3$ letters with a carefully chosen

decoder.

To prove the theorem, it remains to show that we can express G^M using a number of letters that only depends on p . Before we do this, note that, although the blue and red cycles are not uniquely defined (we might get different red and blue cycles if we choose a different starting cycle), the partition we get at the end into G^L, G^M and G^R still satisfies the properties we have described so far. In particular, we may assume without loss of generality that the cycle we start with is given by vertices v_1, \dots, v_k with $v_i \in A_i$, such that v_{i+1} is the leftmost neighbour of v_i for i modulo k (in other words, it *is* the blue cycle).

Now start with $v_k \in C_b$ and consider the sequence of vertices that we get by repeatedly taking the rightmost vertex in the previous bag modulo k . Label the vertices of this sequence as $v_{i,j}$, where i is the bag of the vertex, and j counts the number of times the sequence has visited that bag after this vertex (so v_k becomes $v_{k,1}$, and the sequence continues with $v_{k-1,1}, \dots, v_{1,1}, v_{k,2}, \dots$). One can picture this sequence as a spiral winding around the chain circuit, ending at the red cycle. Note that this spiral is indeed winding “to the right”, more precisely, for all i , if $j_2 > j_1$, v_{i,j_2} is to the right of v_{i,j_1} in the usual CC-ordering (or the two vertices are equal). To see this, observe that if a vertex of the sequence is to the right of a cycle (i.e., to the right of the vertex of the cycle lying in the corresponding bag), so are all following vertices.

Put $S_0 := C_b$, and let S_j ($j > 0$) be the cycles induced by vertices $v_{i,j}$ (note that those are indeed cycles, since the chain circuit property together with the above observation implies that $v_{k,j}$ is adjacent to $v_{1,j}$). Write G_j^M for the set of vertices strictly between S_j and S_{j+1} . Finally, write $A_{i,j}$ for $A_i \cap G_j^M$.

The following statements hold:

- $V(G^M) = \bigcup_{\substack{1 \leq i \leq k \\ j \geq 0}} A_{i,j} \cup \bigcup_{j \geq 0} S_j$.
- The number of disjoint S_j is bounded as a function of p . To see why, assume that $S_1, S_2, \dots, S_{2r+1}$ are all disjoint for some $r \in \mathbb{N}$. One can check that, by construction, $S_1, S_3, \dots, S_{2r-1}, S_{2r+1}$ induce a $C_{k,r}$. As $C_{k,p}$ is forbidden, it follows that $r < p$. Since the sequence $(v_{i,j})$ becomes periodic once it repeats a vertex, it follows that

$$|\{v_{i,j} : i \geq 1, j \geq 0\}| < k(2p + 1).$$

From [Pet02], we know that if the lettericity of a graph is l , then adding a vertex produces a graph of lettericity at most $2l + 1$. In view of the second statement

above, $\bigcup_{j \geq 0} S_j$ has size at most $k(2p+2)$, so it suffices to show that $G^M \setminus \bigcup_{j \geq 0} S_j$ can be expressed using a bounded number of letters. We claim that this can be done using one letter $a_{i,j}$ for each non-empty set $A_{i,j}$. This is enough, since from the above discussion, $|\{(i,j) : A_{i,j} \neq \emptyset\}| < k(2p+2)$. To see how one can construct a word and decoder representing $G^M \setminus \bigcup_{j \geq 0} S_j$, let us examine the edges between those sets. Let A_{i_1, j_1} and A_{i_2, j_2} be two such sets. If $i_1 - i_2 \not\equiv \pm 1 \pmod k$, then there are no edges between A_{i_1, j_1} and A_{i_2, j_2} , so write i for i_1 and assume without loss of generality that $i_2 = i + 1 \pmod k$. All of the following claims follow straightforwardly from our construction of the $A_{i,j}$ and the properties of chain circuits.

- If $i \neq k$, we have:
 - No edges between A_{i, j_1} and A_{i+1, j_2} when $j_2 < j_1$.
 - All possible edges A_{i, j_1} and A_{i+1, j_2} when $j_2 > j_1$
- For $i = k$, we have:
 - No edges between A_{k, j_1} and A_{1, j_2} when $j_2 < j_1 - 1$
 - All possible edges between A_{k, j_1} and A_{1, j_2} when $j_2 \geq j_1$.

The key is now to notice that all of the non-trivial adjacencies appear between consecutive bags in the following sequence that “spirals around” the chain circuit (some of the bags might be empty):

$$A_{k,0}, A_{k-1,0}, \dots, A_{1,0}, A_{k,1}, A_{k-1,1}, \dots, A_{1,1}, A_{k,2}, \dots$$

One helpful way of conceptualising this is by thinking of the red edges as “impermeable” to edges crossing them from top right to bottom left, and to non-edges crossing them from top left to bottom right.

Any two consecutive bags in the above sequence induce chain graphs. As described in Section 3.2, we can realise the subgraph consisting of the edges between consecutive bags by using one letter $a_{i,j}$ for each bag $A_{i,j}$; the decoder then contains the (ordered) pairs of letters corresponding to consecutive bags. The construction of the word can be done inductively: if we have a word describing the subgraph up to a certain bag $A_{i,j}$ in the sequence, we can add the letters corresponding to vertices from the next bag by placing them carefully among the letters of bag $A_{i,j}$ (see Subsubsection 3.2.1 for more details). To make the word represent $G^M \setminus \bigcup_{j \geq 0} S_j$, all we need to do is add pairs $(a_{i_1, j_1}, a_{i_2, j_2})$ and $(a_{i_2, j_2}, a_{i_1, j_1})$ to the decoder whenever the corresponding bags have all possible edges between them.

□

3.4.2 Local consistency: a LOH perspective

Let us start by giving, a posteriori, a short defense for the introduction of LOHs. First, note that they provide a language that is strictly more general than monotonically gridded permutations. For instance, one can associate LOHs to chain circuits in a straightforward way; while k -CCs with an even k have analogues in the world of monotonically gridded permutations, those with an odd k do not, since any cycle in the cell graph of a matrix has even length. However, for our purposes, there is no need to treat k -CCs differently based on the parity of k .

Second, and perhaps most importantly, note that LOHs provide a much cleaner environment to work in than the one given by either gridded permutations or lettericity. Indeed, gridding matrices and decoders come with a lot of superfluous information. We almost exclusively care only about bags with non-trivial relationships between them: collinear cells, or pairs of letters with exactly one arc between them in the decoder. We usually do not even care which one of the two arcs appears in the decoder, or whether the row or column's signs are 1 or -1. This information is not useful during most proofs; it is in fact just a burden that needs to be carried around, and that often obfuscates the real intuition behind the arguments. It thus makes sense to separate the useful information from the rest, which is what LOHs attempt to do.

In summary, the problems of characterising bounded lettericity and geometric griddability concern the passage from a locally consistent to a globally consistent regime.¹⁶ LOHs abstract the notion of local consistency, and thus provide what we feel is a more natural setting for those problems. We believe it would be very beneficial to the subject to develop a general theory of LOHs. Such a theory would not only help make the link between lettericity and geometric griddability more transparent and reliable, but it would also, in all likelihood, provide a more elegant formulation for several existing results, such as the antichain constructions in [MV02].

With all of this being said, it is natural to ask how the statements we have shown about the lettericity of chain circuits translate into statements about the geometric griddability of certain classes of permutations. And indeed, one finds that Theorem 71 can be used to produce a characterisation of obstacles to geometric griddability within classes that are monotonically griddable by a matrix with one cycle in its cell graph. Rather than stating that characterisation explicitly and

¹⁶For lettericity, the story is a bit more complicated – see the next subsection.

proving it directly (which would be tedious), let us try to find a more general analogue to Theorem 71 in the language of LOHs. The optimal formulation of such a statement is not yet clear, but let us nevertheless see what we can do. To produce the “correct” statement, we need a few ingredients:

- First, we would like a description of the end state we are trying to achieve (that is, the notion that we are trying to characterise). For graphs, this is bounded lettericity, and for permutations, it is geometric griddability. Naturally, the corresponding notion for LOHs is global consistency.
- We need to make the scope of such a statement precise. The holy grail would be a statement that characterises global consistency among all (classes of) LOHs.¹⁷ However, it might be necessary or helpful to restrict ourselves to certain kinds of LOHs, for instance the ones coming from monotone gridding by matrices with one cycle in the cell graph. The cell graph construction itself does not seem to generalise nicely to LOHs, however the row-column intersection graph [VW11, p. 2] easily generalises to the line graph of the LOH. It would therefore make sense to specify the scope of what we want to show by putting conditions on said line graph.
- We need some way to express the obstacles to achieving the desired end state. The obvious candidate for LOHs is by specifying forbidden structures in the conflict graphs.¹⁸ The most straightforward first choice is forbidding a “cycles-in-a-chain” structure modelled on the chain circuits $C_{k,l}$. It is not yet clear whether we want to forbid those as induced subgraphs, or more restrictively, as subgraphs; for a direct (non-generalised) analogue of Theorem 71, this distinction does not seem to matter.
- Finally, we need some way of determining whether we have reached the end state. What we mean by this is the following: for chain circuits, we have bounded lettericity if we are able to split each bag into a bounded number of letter classes using (a bounded number of) red and blue edges. For monotonically gridded permutations, we have geometric griddability if we are able to split the gridding via (a bounded number of) horizontal and vertical lines into a geometric gridding by a matrix of bounded size.

¹⁷To talk about classes of LOHs, we would need to define sub-LOH containment – this can be done in the obvious way in terms of subhypergraph containment.

¹⁸An alternative would be to describe obstacles directly in terms of sub-LOHs, but the chances are that such a description would just be the conflict graph one in disguise.

With this in mind, we propose the following definition:¹⁹

Definition 72. Let $\mathcal{H} = (X, E)$ be a LOH with local hyperedge orders \leq_e , and let $x \in X$. The *split* of \mathcal{H} at x is the LOH $\mathcal{H}^x = (X \setminus \{x\}, E')$, where

$$E' := \{e \in E : x \notin e\} \cup \{e \cap \{y \in e : y <_e x\} : x \in e\} \cup \{e \cap \{y \in e : y >_e x\} : x \in e\}$$

(we only consider non-empty sets in the union). The local orders on the new hyperedges are inherited from \mathcal{H} .

The intuition behind splitting a LOH at an element x is that we remove all of the comparisons between elements smaller than x and elements larger than x . In other words, splitting a LOH brings us closer to global consistency, since we are deleting arcs in the conflict graph. In the monotone gridded setting, it is analogous to adding the horizontal and vertical lines through x to the gridding – just as the row and column of the gridding to which x belongs get split into two, so do all hyperedges of the LOH containing x . Indeed, this is also the role of the red and blue edges from the proof of Theorem 71. We can thus define a parameter for LOHs measuring how far they are from being globally consistent:

Definition 73. The *global inconsistency* of a LOH is the smallest number of splits needed to make the LOH globally consistent.

Now our question becomes: what conditions do we need to put on the LOHs in some given class to guarantee that their global inconsistency is bounded? Let us formulate an analogue of Theorem 71 in this language.

Definition 74. A LOH is *k-cyclic* if its line graph is a cycle on k vertices.

Notation 75. Let $C_{k,l}$ be as in Notation 65, enhanced with an orientation of the edges from $v_{i,j}$ to $v_{i+1,j'}$ for all appropriate i, j, j' .

Theorem 76. Let $k, l \in \mathbb{N}$ be fixed. Let \mathcal{X} be a class of *k-cyclic* LOHs such that their conflict graphs avoid $C_{k,l}$ as an induced subgraph. Then \mathcal{X} has bounded global inconsistency.

This statement can be proved in a way completely analogous to Theorem 71; we can define a way of gluing together LOHs such that one of them is “smaller” than the other – this allows us to use an induction argument. The red and blue edge construction can then be reproduced by defining a notion of “infimum” and

¹⁹To make this definition smoother, one could try to work with an alternative definition of LOHs that makes them downwards closed. Let us stick with the original definition for now.

“supremum” for elements in some hyperedge of the LOH (strictly speaking, we need to be a bit more careful than that and make sure the infimum and supremum are in the appropriate cells). The heart of the argument stays the same, so we skip the details.

While Theorem 76 and its implications for the study of lettericity and geometric griddability is a step in the right direction, much remains to be done in order to fully understand the subject. We will elaborate on this in the next subsection.

3.4.3 A hierarchy of structure: the missing pieces

As we approach the end of the first chapter, a big picture begins to emerge – one of its main features is the transition from local to global consistency. Our understanding of this transition is incomplete, but quickly growing; this is however not the whole story. There is another transition at play here, namely the one that brought us from disorder to local consistency in the first place. In the universe of permutations, this initial transition is in big part understood: its description is given by Huczynska and Vatter in Theorem 24. In the universe of graphs, the situation is more complicated. Our first instinct might be to forbid the graph analogues to the obstacles from Theorem 24 and see what happens. Those analogues are induced matchings and their complements.²⁰ This problem has already been studied and to a large degree solved by Atminas in [Atm17]; the paper provides an impressive general result characterising the structure of graphs without star forests and their complements. A special case of that result concerns graphs without large matchings and complements. Atminas shows that those graphs can be partitioned into a bounded number of homogeneous bags (the bound only depends on the size of the matchings and complements we are forbidding) such that the edges between any pair of bags form a chain graph. However, this is not yet the locally consistent regime of LOHs: the “nice” orders from the chain graphs need not agree on each bag. This raises a second question of what lies between this intermediate regime and local consistency. We summarise this discussion in Figure 3.16, and leave two open problems that we then examine in more detail. But first, let us give an official parametric definition of local consistency in the universe of graphs (and use this chance to introduce another parameter):

Definition 77. Let G be a graph. A *chain partition* of G is a partition of $V(G)$ into homogeneous sets A_1, \dots, A_t such that, for each $1 \leq i, j \leq t$, the subgraph consisting of A_i, A_j and the edges between them is $2K_2$ -free. We define the *chain*

²⁰We are using “complements” in the loose sense of Remark 62.

co-chromatic number $\gamma(G)$ of G as the smallest natural number t such that $V(G)$ admits a chain partition into t bags.

A *locally consistent chain partition* of G is a partition of $V(G)$ into totally ordered homogeneous sets A_1, \dots, A_t such that, for each $1 \leq i, j \leq t$, the subgraph consisting of A_i, A_j and the edges between them is a properly ordered chain graph in the sense of Sub-subsection 3.2.1 (i.e., vertices in A_i are arranged in decreasing order with respect to their neighbourhoods in A_j and vertices in A_j are arranged in increasing order with respect to their neighbourhoods in A_i , or vice-versa). We define the *locally consistent chain co-chromatic number* $\lambda(G)$ of G as the smallest natural number t such that $V(G)$ admits a locally consistent chain partition into t bags. A class of graphs is *locally consistent* if λ is bounded in it.

Remark 78. We may define a parameter λ' slightly stronger than λ by only requiring the ordering of the vertices in A_i to be either increasing or decreasing with respect to their neighbourhoods in A_j – call such a chain partition *locally semi-consistent*. In Sub-subsection 3.4.3, we will give an example of a class which has bounded λ' and unbounded λ . In the world of permutations, bounded λ' is analogous to monotone griddability; the previously defined semi-LOHs are an analogue to LOHs for the study of bounded λ' .

Note that $\gamma(G) \leq \lambda'(G) \leq \lambda(G) \leq \text{let}(G)$ for any graph G . Also note that, like with chain circuits, we may attach LOHs to graphs of bounded λ in a straightforward way – in this sense, locally consistent classes of graphs are analogous to monotonically griddable classes of permutations by a partial multiplication matrix.

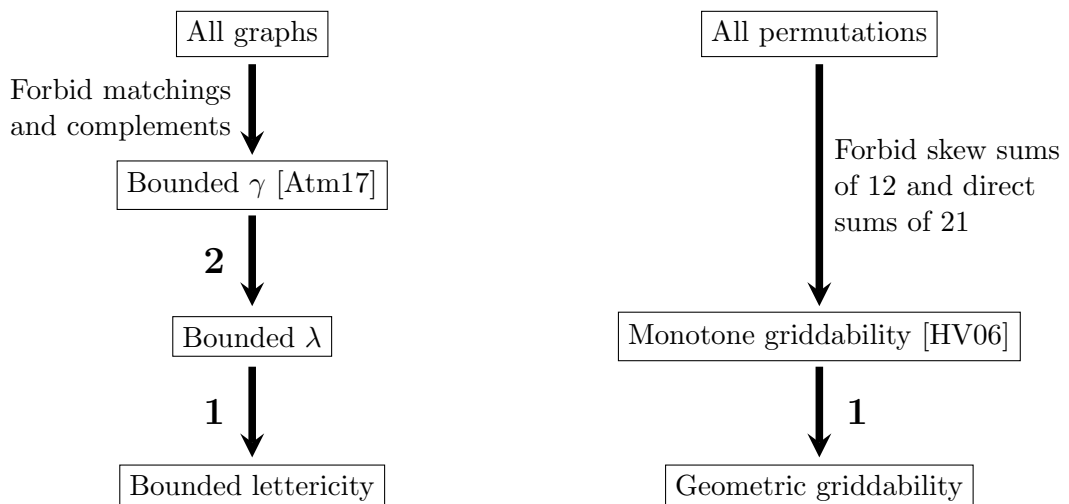


Figure 3.16: A hierarchy of structure

Open problem 79. What lies between local and global consistency? This question has several formulations that are more or less equivalent in spirit:

- What are the minimal classes of permutations that are monotonically griddable by a partial multiplication matrix, but not geometrically griddable?
- What are the minimal classes of unbounded lettericity, but bounded λ ?
- What are the obstacles to bounded global inconsistency in the conflict graphs of LOHs?

Open problem 80. Other than matchings and complements, what are the obstacles to local consistency in the universe of graphs? In other words, what are the minimal classes of unbounded λ , but bounded γ ?

Between local and global consistency

Theorem 76 provides an answer to Problem 79, but only in a restricted setting. One full answer to the problem would be a generalisation of the theorem to all LOHs with a bounded number of hyperedges, rather than just the k -cyclic ones. The key question then becomes: are there other obstacles to bounded global inconsistency than the cycles-in-a-chain structure? What do they look like?

At this point, our research on the problem enters the realm of speculation, and the notions we are working with start to lose their sharpness. We believe there is nevertheless some value in presenting a few loose ideas and tricks that show potential. One of the current strategies for tackling the question consists of trying to show there are essentially no other types of obstacles by using Theorem 76 as a black box, and applying it repeatedly. The general intuition is the following: working in a LOH with at most t hyperedges, we would try to show that any obstacle to global consistency must occur inside a k -cyclic LOH constructed from the original one. If we manage to bound the size of k in terms of t , there will be a bounded number of cyclic LOHs containing these obstacles. If indeed, all obstacles do come from cyclic LOHs, then by forbidding cycles-in-a-chain, we should be able to use Theorem 76 on each of those cyclic LOHs, which would yield a bound on the global inconsistency only depending on t .

Let us discuss some specifics in order to illustrate these ideas. The easiest unsolved instance of the problem consists of two chain circuits with a common cell; a

concrete example of this is the monotone grid class of the matrix $\begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}$. Identifying all obstacles to bounded global inconsistency (i.e., to geometric griddability) in this class will be a big step forward towards our goal.²¹ There are two “simple cycles” in the cell graph of the matrix: the one given by the four entries in the top left, and the one given by the four entries in the bottom right. Call them A and B respectively. Using Theorem 76, if we forbid $C_{4,p}$ ($p \in \mathbb{N}$) in the conflict graph, we are able to “clean out” all circuits that lie entirely in one of A or B . Any other obstacle to global consistency must somehow involve both A and B .

So what kinds of obstacles are there that can involve both A and B ? It turns out that our cycles-in-a-chain make another appearance here. Imagine working in a “virtual” chain circuit that starts in the middle cell of the matrix, goes around A once, then around B ; the middle cell thus occurs twice as a bag, but we treat those occurrences as disjoint copies. One may then construct new cycles-in-a-chain structures such as the one depicted in Figure 3.17: we go around cycle A (represented by red arcs), arrive to the left of where we started, then go around cycle B (represented by blue arcs) to get back to the start. In order to eliminate those structures, we may use Theorem 76 applied to this virtual chain circuit, provided $C_{8,p}$ is forbidden as a *not necessarily induced* subgraph of the conflict graph (while the two occurrences of the middle cell are independent in the virtual chain circuit, in the actual conflict graph, the two sets are complete to each other, with all arcs oriented in the same direction).

We note that this kind of obstacle is still, in some sense, periodic. What makes the problem messier at this stage is the possibility that aperiodic obstacles exist. One can think of them by analogy with the aperiodic fundamental antichains described in [MV02, Section 5]. The general idea of the construction presented there is that we produce an aperiodic word on A and B , and go around the cycles in the order indicated by the word.²² Since geometrically griddable classes are wqo, we had better be able to destroy all but a finite number of the elements in such an antichain by forbidding the appropriate graphs. Could it be that those graphs are just our old cycles-in-a-chain, or is some fundamentally different kind of obstacle hiding somewhere in there?

²¹The author is optimistic that it is only a matter of time before progress is made in this direction; either we will manage to show that forbidding the cycles-in-a-chain obstacles is enough to guarantee bounded global inconsistency in this class, or we will find another type of obstacle.

²²The construction in [MV02] is slightly more involved than that, but the details do not matter for this discussion.




Figure 3.17: A “figure 8” obstacle

Let us try to start answering that question. Sticking with the above example, consider a walk in the conflict graph. When we start in a bag and go around one of the simple cycles in the line graph of the LOH, we end up somewhere in the original bag. By applying Theorem 76 on the (boundedly many) simple cycles, we gain some control of where we end up: for instance, by construction, in the split LOH, we may not return to where we started by just walking repeatedly around the same simple cycle in the original line graph. However, as in Figure 3.17, we might be able to return to the start (and thus produce a circuit in the conflict graph) by walking around different cycles of the line graph in succession. This means that, in order to eliminate all circuits, we need to apply Theorem 76 to the “virtual chain circuits” coming from certain concatenations of simple cycles in the line graph.²³ If we manage to show that it suffices to do it only for a bounded number of such concatenations, we are done.

One way to go about this would be by stating and proving some kind of Ramsey-type result: if our (not necessarily chordless) circuit in the original conflict graph is long enough, the sequence of cells visited is going to have some repeated subsequences. We would want to use this in order to find a certain structure that would have been destroyed by our bounded number of applications of Theorem 76.²⁴

²³Additionally, we need to distinguish between the two possible directions in the simple cycles.

²⁴In fact, this is how we originally found the example from Figure 17: a certain number of

Indeed, it would be enough to show that any circuit above a certain size in the original conflict graph had an arc that was eliminated by one of a bounded number of splits.

One final item that we would like to mention in connection with this problem is scalability. Suppose we have completely solved the problem in the above instance: how do we generalise the solution? We do not yet know, but we have a perspective that might prove helpful in that regard. It begins from the noteworthy remark that there is a finite set of “obvious candidates” for a global order – in the world of permutations, they are the ones coming from spanning trees of the cell graph. Indeed, if the cell graph is already a tree, the unique candidate actually happens to work (Theorem 31). In the k -cyclic setting, it is possible to reinterpret the problem as follows: fix a spanning path of the cell graph; from this, using the methods from [VW11], we may produce in a systematic way a total order on the elements of the permutation. All arcs of the conflict graph between consecutive cells along the path are “forward”, in the sense that they agree with the order. The only “backward” arcs may only occur between the cells corresponding to the ends of the path (i.e., along the unique edge of the cell graph not belonging to the spanning tree). It should be possible to rewrite Theorem 76 in this setting, where the focus is on controlling those backward arcs. In the general case, we would be looking at the order given by a spanning tree,²⁵ and the backward arcs would appear along the edges not in it. While the fundamental difficulties we encountered before are still there, it is possible that a spanning tree/backward arc-focused perspective gives a cleaner interpretation of the problem, and thus a better path towards a solution.

Between bounded γ and local consistency

So far, we have seen two fundamentally different kinds of obstacles to bounded lettericity. First, we saw the matchings and their complements, which are the only minimal classes of unbounded γ .²⁶ Then, we saw the cycles-in-a-chain constructions that live between local and global consistency – we are trying to determine whether they are the only minimal classes of unbounded lettericity but bounded λ .

There is one place left to look for minimal classes of unbounded lettericity:

“backwards paths” like the red ones from the figure guarantees by simple arguments that certain configurations of their endpoints are unavoidable; one such configuration leads to the given obstacle.

²⁵Could the choice of the spanning tree matter? Perhaps we would have something to gain by considering multiple spanning trees simultaneously.

²⁶Strictly speaking, we have not yet shown unboundedness of γ in those classes, or their minimality with respect to this property, but this is not difficult; the difficult part of the statement is what was proved in [Atm17], namely that those are the *only* minimal classes.

among classes of unbounded λ , but bounded γ . Of course, we need to start by making sure that those kinds of classes actually exist within the universe of all graphs (as opposed to the universe of permutation graphs, where the two parameters seem to be equivalent). To do this, let us start by trying to construct an example as simple as possible.

Consider a graph G whose vertex set consists of three independent sets A , B and C on n vertices each, and assume $G[A \cup B]$ and $G[B \cup C]$ are prime chain graphs, while $G[A \cup C]$ is edgeless. In the scope of the current discussion, we will call graphs with this structure *linked chain graphs*. Let us order the vertices in B in increasing order with respect to their neighbourhoods in A , and label them by $1, \dots, n$. Since $G[B \cup C]$ is a prime chain graph, there is a unique permutation $\pi \in S_n$ such that the ordering $\pi(1), \dots, \pi(n)$ has decreasing neighbourhoods in C . We call π the *linking permutation* of G . See Figure 3.18 for an illustration.




Figure 3.18: A graph with linking permutation 614253

It is clear that every linked chain graph G has $\gamma(G) \leq 3$, and that G is uniquely determined by its linking permutation. Can we construct a sequence of permutations π_n such that the corresponding sequence of linked chain graphs G_n has unbounded λ ? Note that if the sequence of permutation graphs G_{π_n} has bounded chromatic number (that is, if each π can be partitioned into a fixed number of increasing subsequences), say at most t , then $\lambda(G_n) \leq t+2$. Indeed, we may partition B into t smaller bags, each corresponding to one of the increasing subsequences, and one easily checks that this new partition is locally consistent.

A similar argument applies if G_{π_n} has bounded clique cover number. If

G_{π_n} has bounded co-chromatic number,²⁷ one can obtain a locally semi-consistent partition into a bounded number of bags (that is, the graphs have bounded λ' , but not necessarily bounded λ).

We will start by looking for permutations whose family of permutation graphs has unbounded co-chromatic number. Such families are easy to construct. For instance, we may let π_n be the permutation on n^2 elements given by the concatenation $w_1 w_2 \dots w_n$, where w_i lists the elements of $\{x : 1 \leq x \leq n^2 \text{ and } x \equiv i \pmod n\}$ in decreasing order (see Figure 3.19). For the sequence $(\pi_n)_{n \geq 1}$, one checks that the size of the maximum homogeneous set in the corresponding permutation graphs is sublinear, which immediately implies unbounded co-chromatic number.




Figure 3.19: The permutation π_7

To show that the corresponding family of linked chain graphs has indeed unbounded λ , and in fact unbounded λ' , we first need a corollary to van der Waerden's theorem on arithmetic progressions. Let us start by recalling the theorem:

Theorem 81 (van der Waerden's Theorem [Wae27]). *For any $p, k \in \mathbb{N}$, there exists a number $N \in \mathbb{N}$ such that if $[N]$ is coloured with p different colours, then there are at least k integers in arithmetic progression whose elements are the same colour.*

Corollary 82. *For any $p, k \in \mathbb{N}$, there exists a number $N \in \mathbb{N}$ such that if $[N] \times [N]$ is coloured with p different colours, then there are two arithmetic progressions $X = \{x_1, \dots, x_s\}$ and $Y = \{y_1, \dots, y_t\}$ of length at least k such that $X \times Y$ is monochromatic.*

Proof. We first use van der Waerden's theorem to find a number N_1 such that any colouring of $[N_1]$ with p colours contains a monochromatic arithmetic progression of

²⁷Unbounded co-chromatic number implies both the chromatic number and clique cover number are unbounded, but not vice-versa – [Zve00] provides a characterisation of the graphs for which the converse implication holds.

length at least k . Note that there are at most $\binom{N_1}{2}$ possibilities for that arithmetic progression (it is uniquely determined by its first two terms). We then use van der Waerden's theorem a second time to find a number N_2 such that any colouring of $[N_2]$ with $p\binom{N_1}{2}$ colours contains an arithmetic progression of length at least k .

Suppose now that $[N_1] \times [N_2]$ is coloured with p colours. By choice of N_1 , for each $1 \leq i \leq N_2$, the set $[N_1] \times i$ contains a monochromatic arithmetic progression X' of length at least k . Colour i with the pair (X', c) , where c is the colour of X' . This gives a colouring of $[N_2]$ with (at most) $p\binom{N_1}{2}$ colours, and hence there is a monochromatic arithmetic progression Y of length at least k . Suppose its colour is (X, c) ; then by construction, $X \times Y$ is monochromatic with colour c . \square

We apply Corollary 82 to the plots of the permutations π_n . Indeed, the corollary directly implies the following: for any p and t , there exists N such that whenever we partition the plot of π_N into at most p pieces, one of the pieces contains π_t as a subpattern. With this, we are ready to prove our result.

Theorem 83. *Let π_n be the permutation on n^2 elements given by the concatenation $w_1 w_2 \dots w_n$, where w_i lists the elements of $\{x : 1 \leq x \leq n^2 \text{ and } x \equiv i \pmod{n}\}$ in decreasing order. Let G_n be the linked chain graph with linking permutation π_n . Then $(\lambda'(G_n))_{n \geq 1}$ is unbounded.*

Proof. The proof is similar in concept to the one of Theorem 66: we show that if all G_n have locally semi-consistent partitions with a bounded number p of bags,²⁸ then in fact, they have locally semi-consistent partitions with only 3 bags. In other words, the partition of the linked chain graphs into sets A, B, C should already be a locally semi-consistent partition. This is clearly not the case for $n \geq 2$, as one can find three vertices in B whose A -neighbourhoods are in strictly decreasing order, but whose C -neighbourhoods are in neither increasing nor decreasing order (and so, no ordering of B will be locally semi-consistent).

Suppose now that every G_n has a locally semi-consistent partition into p bags. Label the vertices of A by x_1, \dots, x_{n^2} in decreasing order of their B -neighbourhoods; label the vertices of B by y_1, \dots, y_{n^2} in increasing order of their A -neighbourhoods (so that $y_{\pi_n(1)}, \dots, y_{\pi_n(n^2)}$ is decreasing with respect to the neighbourhoods in C); finally, label the vertices of C by z_1, \dots, z_{n^2} in increasing order of their B -neighbourhoods.

²⁸By replacing p by $3p$ if necessary, we may assume that each bag is contained entirely in A, B or C .

Let $A_1, \dots, A_{p_1}, B_1, \dots, B_{p_2}, C_1, \dots, C_{p_3}$ be the parts lying in A, B and C respectively (so that $p = p_1 + p_2 + p_3$). We use the labelling described above to obtain, from the partitions of A, B and C , three partitions of $[n^2]$. Define:

- $A'_i := \{j \in [n^2] : x_j \in A_i\}$;
- $B'_i := \{j \in [n^2] : y_j \in B_i\}$;
- $C'_i := \{j \in [n^2] : z_{\pi_n^{-1}(j)} \in C_i\}$ (take note of the π_n^{-1} in the index).

Now consider the common refinement of the three partitions $(A'_i), (B'_i)$ and (C'_i) . This is a partition D'_1, \dots, D'_r of $[n^2]$, with $r \leq p^3$. From this refinement, we construct a new partition (D_i) of $V(G_n)$ by putting $D_i := \{x_j : j \in D'_i\} \cup \{y_j : j \in D'_i\} \cup \{z_{\pi_n^{-1}(j)} : j \in D'_i\}$. By construction, for each i , there exist i_1, i_2, i_3 such that $D_i \subseteq A_{i_1} \cup B_{i_2} \cup C_{i_3}$. This means that the induced subgraph $G_n[D_i]$ has $\lambda(G_n[D_i]) = 3$, since the partition $(A_i) \cup (B_i) \cup (C_i)$ of $V(G)$ is by assumption locally semi-consistent.

We claim $G_n[D_i]$ is, in fact, a linked chain graph whose linking permutation is the subpattern of π_n induced by the indices in D'_i . To see this, note that the neighbourhood of y_j in A is the interval x_1, \dots, x_j , while its neighbourhood in C is the interval $z_{\pi_n^{-1}(j)}, \dots, z_{n^2}$. Thus by construction, any vertex $y_j \in B \cap D_i$ has its rightmost neighbour from A and its leftmost neighbour from C also in D_i . From this, writing $B \cap D_i = \{y_{j_1}, \dots, y_{j_s}\}$ with $j_1 < \dots < j_s$, it is easy to see that this ordering of the vertices in $B \cap D_i$ is strictly increasing with respect to the neighbourhoods in $A \cap D_i$, and the ordering $y_{\pi_n(j_1)}, \dots, y_{\pi_n(j_t)}$ is strictly decreasing with respect to the neighbourhoods in $C \cap D_i$. But π_n reorders the y_{j_l} in exactly the same way as the the subpattern of π_n induced by the indices in D'_i reorders $\{1, \dots, t\}$, and so our claim follows.

Finally, to arrive at our contradiction, we use Corollary 82 (and the paragraph after it): the D'_i are a partition of $[n^2]$, and thus of the plot of π_n , into at most p^3 parts. It follows that, for any fixed t , if n is large enough, one of those parts, say the i th, will contain π_t as a pattern. By removing the appropriate vertices of $G_n[D_i]$, we can produce a locally semi-consistent partition of G_t into 3 bags, which is not possible for $t \geq 2$, as discussed at the beginning of the proof. □

This construction shows that, as far as local (semi-)consistency is concerned, things can go wrong even when there is a single bag where two chain graphs meet. A first step to understanding λ is an exhaustive analysis of this setting; one would

perhaps aim for a clean characterisation of bounded λ among linked chain graphs in terms of some conditions on the linking permutations.

In the general setting of bounded γ , we would be dealing with several bags A_1, \dots, A_p ; it would make sense to attempt to generalise the notion of “linking permutation” by associating to each bag A_i a set of permutations $\pi_{j \rightarrow k}^i$, which intuitively describe how we need to permute the bag A_i to get from an ordering that makes $G[A_i \cup A_j]$ properly ordered to one that makes $G[A_i \cup A_k]$ properly ordered. Of course, there are several complications here:

1. To properly order $G[A_i \cup A_j]$, we would need to re-order A_i and A_j in tandem. Working with linking permutations restricted to one bag does not seem fit for this job.
2. The chain graphs between pairs of bags will not, in general, be prime, so the permutations $\pi_{j \rightarrow k}^i$ will not be uniquely defined. How do we decide which ones to use?
3. The linking permutations alone do not suffice to characterise boundedness of λ ; one must also take into account the interactions between them. As an example, one can construct graphs of bounded λ' but unbounded λ by arranging bags in a cycle with properly ordered prime chain graphs between successive bags, except for one pair of bags where we “twist” the ordering in one of the bags – see Figure 3.20 for an illustration. That those graphs have unbounded λ is essentially a (simpler) variant of Theorem 83 to show that more bags do not help, together with the remark that, by uniqueness of the good orderings between bags, we may not make the partition in the figure locally consistent by simply reordering. In this example, the linking permutations are deceptively simple: they are either the identity, or its reverse. This shows that we must track, in some way, what happens between different bags – a tool like the cell graph of gridding matrices might be useful here.
4. Assuming we obtain a characterisation of bounded λ in terms of conditions on the linking permutations and their interactions, how do we transform it into a minimal class characterisation?

Dealing with those difficulties is a good subject for future research. We remark that the first point in the list can be mitigated by using local semi-consistency: we may refine our search for minimal obstacles by first looking at what happens when λ' is bounded. Once we understand the boundary between λ and λ' (that is, the boundary between monotone griddable classes, and monotone griddable classes by




Figure 3.20: Graphs of bounded λ' but unbounded λ

a partial multiplication matrix), we may look for obstacles to bounded λ' : in this setting, it makes sense to talk about linking permutations defined on individual bags. Something worth investigating is whether it is possible and meaningful to generalise the definition of local semi-consistency (which is like local consistency, except the local orders may be reverses of each other) by specifying various other conditions on the local orders in each bag (for example, we could allow them to be cyclic permutations of the vertices).

Chapter 4

Functionality

Graphs of bounded degeneracy can be constructed from a single vertex by adding vertices of small degree one at a time. Cographs can be constructed from a single vertex by adding twin vertices¹ one at a time (for various characterisations of cographs, together with references, see, e.g., [BLS99]). What do these examples have in common? In both cases, there is a vertex whose neighbourhood is in some sense “easy to describe”: we only need a bounded number of other vertices to express it. In this chapter, we define and study a new graph parameter called *functionality*, which attempts to capture that intuition.

In Section 4.1, we define functionality and discuss the motivation behind it; in doing so, we place our definition in the broader context of graph enumeration and representation. In Section 4.2, we present our results from [AAL19], where we originally introduced the parameter. Finally, in Section 4.3, we provide some further results and insights into it, and propose a few open questions emerging from our study.

4.1 Background and motivation

4.1.1 The (labelled) speed of hereditary classes

Let \mathcal{X} be a hereditary class. Write \mathcal{X}_n for the set of labelled n -vertex graphs in \mathcal{X} . The sequence $|\mathcal{X}_n|$ is called the *speed*, or the *growth rate* of \mathcal{X} ; we say \mathcal{X} is logarithmic, linear, exponential, \dots if its speed is logarithmic, linear, exponential, \dots in n . Numerous results on this subject appear in the literature – the interested reader is invited to consult, for instance, [BBW00] and the references therein. We summarise here only the background directly relevant to us.

¹We allow both true and false twins, that is, the twin vertices may be adjacent or non-adjacent.

Clearly, any class has speed at most $2^{\binom{n}{2}}$, since that is the total number of labelled graphs on n vertices. Interestingly, not every asymptotic behaviour between constant speed and this upper limit actually occurs. In fact, there are strong restrictions on what the growth rates can be. In an effort to characterise the possible speeds, [Ale92; Ale97] and independently [SZ94] describe a hierarchy of distinct layers of classes, with gaps between them. For instance, it turns out that any sublinear class is, in fact, constant – there are no classes with, say, logarithmic growth.

Alekseev [Ale97] provides a characterisation of the bottom four layers in the hierarchy via minimal classes not belonging to those layers, as well as a structural description for the bottom three – the constant, polynomial and exponential layers. Those three layers have a relatively simple structure; for instance, if a class belongs to the exponential layer, then graphs from it can be partitioned into a bounded number (that is, the bound depends only on the class) of homogeneous sets that are pairwise either complete or anticomplete to each other. The fourth layer is where things get interesting. This is the so-called *factorial* layer, in which classes have speed $2^{\Theta(n \log n)}$. Many well-studied classes lie in this layer: permutation and interval graphs [Spi03], line graphs [LMZ12], all proper minor-closed classes (such as forests) [Nor+06] to name just a few.

Using Alekseev’s characterisation, a class is superexponential if and only if it contains at least one of nine minimal classes. Three of those minimal classes are subclasses of bipartite graphs: induced matchings, their bipartite complements and chain graphs. The other six can be obtained from those three by replacing one, respectively both of the parts with cliques. This tells us that it is relatively easy to determine when a class of graphs is *at least* factorial. Determining whether a class is *at most* factorial can be a bit more tricky; there are, however, a variety of useful techniques to that end, many of which are discussed in [Spi03]. One general approach is to attempt representing graphs from that class as “space efficiently” as possible: if we manage to do it using only $O(n \log n)$ bits for each graph, then the class must be at most factorial. It goes without saying that the motivation behind the study of graph representations extends far beyond purely enumerative purposes – Spinrad’s monograph [Spi03] surveys the topic in detail.

Remark 84. We emphasize that the speed is defined in terms of *labelled* graphs. In particular, if the number of unlabelled graphs in a class is factorial, it follows that its speed is also factorial (since there are at most $n!$ possible labellings for an n -vertex graph). The converse need not be true: there are classes with exponentially many unlabelled graphs, but factorially many labelled ones. An example is given by

graphs of bounded lettericity. Indeed, their definition via words immediately yields an exponential upper bound for unlabelled graphs; a factorial lower bound for the labelled graphs follows immediately from Alekseev’s minimal class characterisation (as chain graphs have bounded lettericity). In the context of functionality, we will require the ability to refer to vertices using short labels, hence we will assume, throughout this chapter, that graphs come with a labelling of the vertices by the numbers 1 to n .

4.1.2 Implicit graph representations

There is one particularly interesting kind of graph representation, namely *implicit representation* (also known as *adjacency labelling scheme*, or just *labelling scheme*), introduced in [KNR88] and independently in [Mul88]. The idea is to store the graphs’ adjacency information locally: each vertex is assigned a label of size $O(\log n)$ in such a way that adjacency of a pair of vertices can be determined solely from reading their labels. Most sources then put some computational constraints on *how* adjacency is determined from the labels; it is usually expected that adjacency is at the very least computable (see, e.g., [Mul88]), and we often impose time polynomial in the size of the labels (see, e.g., [KNR88]). Not all graph classes admit implicit representations; it is clear from the prescribed label size that any class with an implicit representation must be at most factorial.

We will avoid giving formal definitions for implicit representations, since we do not need them for our discussion. After going through a few examples, there should in practice be no ambiguity as to what we (in our limited scope) mean when saying a class “has an implicit representation”. The archetypal example is given by interval graphs:

Example 85. Let G be an interval graph on n vertices. Starting with an interval intersection model for G , we number the endpoints of the intervals in increasing order of their appearance on the real line (see Figure 4.1 for an illustration; we may assume without loss of generality that no two intervals share an endpoint). We then label each vertex with the two numbers of the corresponding interval’s endpoints² – those are two integers between 1 and $2n$, so the label size is indeed $O(\log n)$. Adjacency can be readily determined from those labels: x and y are non-adjacent if and only if the largest of the numbers stored at x is smaller than the smallest of the numbers stored at y , or vice-versa.

²How we depict these labels is entirely a matter of implementation. We may simply think of them as pairs (i, j) with $i < j$.




Figure 4.1: Implicit representation of interval graphs

We provide a second less standard, yet unsurprising example.

Example 86. Graphs of bounded lettericity also have implicit representations. Let us describe a labelling scheme for graphs of lettericity at most k . Given such a graph, we start by finding a k -letter graph expression for it; for each vertex, we record its position in the expression ($O(\log n)$ bits), and its letter. It is clear that the $O(\log n)$ bits stored at a pair of vertices are enough to determine whether or not the vertices are adjacent – the decoder can be considered part of the adjacency computing algorithm or, if we want to keep avoiding digressions into formal definitions, we may simply store it at each vertex (since it is just a constant amount of information).

As one might expect, the big question concerning implicit representations is: which classes actually admit them? This question takes the form of the *Implicit Graph Conjecture*.

Conjecture 87 (Implicit Graph Conjecture [KNR88] – the usual formulation). Every factorial hereditary class of graphs admits an implicit representation with polynomially computable adjacency.

As of the writing of Spinrad’s monograph, it was not known whether the additional computational restrictions affected which classes are representable [Spi03, p. 22]. This has recently changed: in his thesis [Cha17], Chandoo undertakes a valiant effort to streamline some of the aspects surrounding implicit representations. Among other things, he develops a complexity theory for implicit representation, which yields a strict hierarchy of graph classes based on the time complexity of the algorithm required to determine adjacency from the labels. The hierarchy does not necessarily remain strict when staying in the universe of *hereditary* classes³ – in Chandoo’s own words, “The graph classes used to demonstrate these separations [in complexity] are far removed from any natural graph class” [Cha17, p. 2].

³Such a result would have implied that certain factorial hereditary classes do not have implicit representations with polynomially computable adjacency.

We remark that there exists a purely graph-theoretical, extremal-flavoured version of the conjecture, which circumvents the technicalities regarding computational complexity. This version is weaker than the usual statement of the Implicit Graph Conjecture; it relates implicit representations with the existence of small vertex-induced universal graphs [Mul88; KNR88; Spi03; Cha17]:

Conjecture 88 (Implicit Graph Conjecture – a weaker, graph-theoretical variant). Let \mathcal{X} be a factorial hereditary class. Then there exists a vertex-induced universal graph for \mathcal{X}_n of size polynomial in n .

Without going too much into detail, the idea behind the (pseudo-)equivalence of the two versions is as follows: given an implicit representation for a class \mathcal{X} with labels of size $c \log n$, we may treat all strings on $c \log n$ bits as the vertices of our universal graph, with adjacency determined by the appropriate procedure. Conversely, a universal graph suggests a (not necessarily computable) way of determining adjacency: the label of a vertex simply describes its embedding into the universal graph.

The author is somewhat mystified by the Implicit Graph Conjecture. Succinctly put, it sounds too good to be true – why should one expect to always find a nice representation for factorial hereditary classes, especially when many adjacency labelling schemes seem to use very ad-hoc, class-specific constructions?⁴ Indeed, to highlight the unpredictability of implicit representations, we put forward the case of cographs. While implicit representations for them are known, we are not aware of a direct and transparent way to obtain one from their emblematic features (such as cotrees). Instead, we need to go the long way around: we can represent them as permutation (or circle) graphs and use their intersection models, or we can use the fact that their clique-width is bounded.

On the other hand, a counterexample also seems out of reach, though there are a few candidates, i.e., factorial hereditary classes for which an implicit representation is not known. Among the more prominent candidates are some classes of intersection graphs, such as disk graphs [Spi03, p. 53].

Remark 89. In the context of implicit representations, we could distinguish between the labels produced by the adjacency labelling scheme and the “names” we give the vertices when working with labelled graphs. However, those names take

⁴For many known classes admitting implicit representations, the representations seem to emerge directly from known characterisations of the class (or can be obtained via simple reductions – see [Atm+15] and [Cha17, sec. 3.3]). There are, of course, more sophisticated examples, like that of classes of bounded clique-width [Spi03, p. 165], but proving the conjecture seems unachievable with the current techniques.

only $O(\log n)$ bits anyway, so whether or not we consider them part of the adjacency labels is immaterial.

With all of this in mind, we are finally ready to start talking about functionality.

4.1.3 Functionality: a very short history

In [All09], Allen investigates the speed of monogenic classes of bipartite graphs – that is, classes of bipartite graphs defined by a unique minimal forbidden induced subgraph. He provides an almost complete description of the set of graphs G for which the class of G -free bipartite graphs has at most factorial speed.⁵ One of the arguments used in the enumeration was that in some of the classes, all graphs had a pair of vertices such that the symmetric difference of their neighbourhoods was small (that is, bounded by a constant). Later, that argument was generalised in [Atm+15] using the notion of *functional vertices*. Let G be a graph with adjacency matrix A , and denote by $A(v, w)$ the entry of the matrix corresponding to vertices v and w .

Definition 90 ([Atm+15]). We say a vertex $y \in V(G)$ is a *function of the set* $X = \{x_1, \dots, x_k\} \subseteq V(G) \setminus \{y\}$ (or simply a function of the vertices x_1, \dots, x_k) if there exists a Boolean function $f : \{0, 1\}^k \rightarrow \{0, 1\}$ of k variables such that for any vertex $z \in V(G) \setminus (X \cup \{y\})$, we have $A(y, z) = f(A(x_1, z), \dots, A(x_k, z))$. Abusing notation, we will write $y = f(x_1, \dots, x_k)$ to mean that y is a function of x_1, \dots, x_k , and the Boolean function in the definition is f .

In other words, y is a function of X if the adjacency of every other vertex to y can be determined from its adjacency to X . In yet other words, y is a function of X if each of the 2^k sets of the form $\{z \in V(G) \setminus (X \cup \{y\}) : N(z) \cap X = A\}$ (where A runs over subsets of X) is either complete or anticomplete to y .

We give two simple and natural examples:

Example 91. Vacuously, any vertex y is a function of $V(G) \setminus \{y\}$ (any function works). Less vacuously (but still somewhat trivially), any vertex y is a function of $N(y)$ or of $\overline{N(y)}$ – the function f is constantly 0 or 1 respectively.

Example 92. The first non-trivial example is given by twin vertices: if x and y are twins (adjacent or non-adjacent), then each of them is a function of (the singleton containing) the other – the function f is the identity. More generally, for any two

⁵The only remaining open case, that of P_7 -free bipartite graphs, is settled later in [LZ17].

vertices x and y , any of them is a function of the set containing the other and the symmetric difference of their neighbourhoods.

The initial motivation behind the definition was that, if each graph in some hereditary class \mathcal{X} has a vertex which is a function of a set of bounded size (say by k), then \mathcal{X} is at most factorial [Atm+15, Theorem 2]. The proof consists of constructing inductively a “functional representation” of size $O(n \log n)$ bits for each graph in \mathcal{X} . Indeed, given a graph $G \in \mathcal{X}$, we find a vertex y which is a function of at most k other vertices. We then remove y from G , and record:

- the label of y ;
- the labels of the vertices $x_1, \dots, x_{k'}$ ($k' \leq k$) of which y is a function;
- a function f such that $y = f(x_1, \dots, x_{k'})$;
- the adjacency of y to $x_1, \dots, x_{k'}$.

We then repeat the process with $G' := G - y$, which must contain another vertex which is a function of at most k vertices. We keep doing this until we run out of vertices. It is clear that G can be recovered from the recorded information, and that at each step, $O(\log n)$ bits are recorded. We note that this is by no means an implicit representation, since we cannot recover adjacency of two vertices solely from the bits that we recorded when we removed them.⁶

Despite its origins as a purely enumerative tool, we felt the notion of functional vertices deserved to be studied in its own right. To this end, we propose in [AAL19] the following definitions:

Definition 93. Let G be a graph, and y a vertex of G . The *functionality* $\text{fun}_G(y)$ of y is the minimum size of a set X such that y is a function of X . We omit the subscript when the graph is clear from context. The functionality of G is

$$\text{fun}(G) := \max_{H \subseteq G} \min_{y \in V(H)} \text{fun}_H(y),$$

where the maximum is taken over induced subgraphs of G .

This makes functionality into a hereditary graph parameter akin to degeneracy. Indeed, degeneracy captures the fact that every (induced) subgraph of G

⁶If we try to use the above functional representation to naively construct an implicit one by recursively producing labels for the vertices, we very quickly run into difficulties: the size of the labels snowballs uncontrollably due to the adjacency information described in the fourth bullet point.

has a vertex of small degree, while functionality captures the fact that every induced subgraph of G has a vertex whose neighbourhood can be expressed in terms of few other vertices. In fact, functionality strictly generalises degeneracy: it follows from Example 91 (and from the above definition) that $\text{fun}(G) \leq \text{degen}(G)$ for any graph G , and so every class of bounded degeneracy is of bounded functionality. The converse fails, e.g., for cliques.

By Example 92, $\text{fun}(y)$ is small if there is a vertex x that is “almost twin” to y , in the sense that the size of the symmetric difference of their neighbourhoods is small. It turns out this special case of small functionality occurs often enough to deserve its own parameter:

Definition 94. Let G be a graph, and let x, y be vertices of G . Abusing notation, we will say “the symmetric difference of x and y ” to mean the symmetric difference of their neighbourhoods, excluding x and y themselves. We define $\text{sd}(x, y)$ to be the size of the symmetric difference of x and y . Abusing notation once more, the *symmetric difference* of G is

$$\text{sd}(G) := \max_{H \subseteq G} \min_{x, y \in V(H)} \text{sd}(x, y),$$

where the maximum is taken over induced subgraphs of G .

As discussed above, $\text{fun}(G) \leq \text{sd}(G) + 1$, so that bounded symmetric difference implies bounded functionality.⁷

The next section follows our work from [AAL19], in which we examine the behaviour of symmetric difference and functionality on various classes, and begin to place them into the vast hierarchy of graph parameters.

4.2 Graph functionality

4.2.1 Graphs of small functionality

Our first result compares clique-width to functionality. As it turns out, graphs of bounded clique-width have bounded functionality, and in fact even bounded symmetric difference:

Theorem 95. *For any graph G , $\text{sd}(G) \leq 2 \text{cw}(G) - 2$.*

Proof. Since clique-width is hereditary, it suffices to show that any graph of clique-width k has a pair of vertices with symmetric difference at most $2k - 2$. Let G be

⁷We shall see in the next section that this implication is strict.

a graph of clique-width k and let T be a rooted tree corresponding to a k -clique-width expression that describes G . For a node v of the rooted tree T , let T^v be the subtree of T induced by the node v and all its descendants. We can choose v in such a way that T^v has more than k leaves, and neither of the two children of v has this property (if no such v exists, we are done, since G has at most k vertices). Since T^v has more than k leaves, at least two of them, say x and y , have the same label at node v . On the other hand, T^v has at most $2k$ leaves by the choice of v . Therefore, G contains at most $2k - 2$ vertices that distinguish x and y , since the two vertices are not distinguished outside of T^v . In other words, $\text{sd}(x, y) \leq 2k - 2$ as required. \square

What about the converse? Does every class of bounded symmetric difference also have bounded clique-width? The answer is no. A counterexample is given by unit interval graphs: clique-width is unbounded in them [GR00], and symmetric difference is bounded, as we show in the next theorem.

Theorem 96. *The symmetric difference of unit interval graphs is at most 1.*

Proof. Since the class of unit interval graphs is hereditary, it suffices to show that each unit interval graph has a pair of vertices with symmetric difference at most 1.

Let G be a unit interval graph with n vertices and assume without loss of generality that G has no isolated vertices (by adding isolated vertices to a graph, we do not increase its symmetric difference). Take a unit interval representation for $G = (V, E)$ with the interval endpoints all distinct. We label the vertices v_1, \dots, v_n in the order in which they appear on the real line (from left to right), and denote the endpoints of interval I_i corresponding to vertex v_i by $a_i < b_i$. We will bound

$$S := \sum_{i=1}^{n-1} \text{sd}(v_i, v_{i+1}).$$

Note that any neighbour of v_i which is not a neighbour of v_{i+1} needs to have its right endpoint between a_i and a_{i+1} . Similarly, any neighbour of v_{i+1} but not of v_i needs to have its left endpoint between b_i and b_{i+1} . In other words, $\text{sd}(v_i, v_{i+1})$ is bounded above by the number of endpoints in $(a_i, a_{i+1}) \cup (b_i, b_{i+1})$ (we say bounded above and not equal, since it might happen that b_i lies between a_i and a_{i+1} , without contributing to the symmetric difference).

The key is now to note that any endpoint can be counted at most once in the whole sum S , since all (a_i, a_{i+1}) are disjoint (and the same applies to the (b_i, b_{i+1})), and the a 's are only counted when they appear between b 's (and vice-versa). In fact,

a_1 and b_n are never counted in S , and if a_2 is between b_1 and b_2 , then v_1 must be isolated, so a_2 is not counted either. The sum is thus at most $2n - 3$. Since it has $n - 1$ terms, one of the terms, say $\text{sd}(v_t, v_{t+1})$, must be at most 1, as required. \square

We have just shown that symmetric difference is strictly stronger than clique-width. We now turn to the gap between functionality and symmetric difference. We know from the last section that functionality is stronger than symmetric difference; is this comparison also strict? The answer is yes: an example of a class with bounded functionality, but unbounded symmetric difference is given by none other than permutation graphs. To show this, we will once more use the geometric perspective described in Subsection 3.1.2. We supplement that perspective by defining, in the obvious way, a notion of “geometric neighbourhood” of a point (illustrated in Figure 4.2).

Definition 97. The *geometric neighbourhood* of a point x is the union of two regions in the plane: the one above and to the left of x , and the one below and to its right.

It is clear from this definition that the set of points of π lying in the geometric neighbourhood of x is precisely the set of neighbours of vertex x in the permutation graph of π .




Figure 4.2: Representation of $\pi = 614253$, with the geometric neighbourhood of 4 shaded

We start by showing that permutation graphs have bounded functionality. The main idea of the proof is to find a point whose geometric neighbourhood can be neatly “approximated” using the neighbourhoods of points near it.

Theorem 98. *The functionality of permutation graphs is at most 8.*

Proof. Since the class of permutation graphs is hereditary, it suffices to show that every permutation graph contains a vertex of functionality at most 8. Let G be a permutation graph corresponding to a permutation π . The proof will be given in two steps: first, we show that if there is a vertex with a certain property in G (yet to be specified), then this vertex is a function of 4 other vertices. Second, we show how to find vertices that are “close enough” to having that property.

Step 1: Consider the plot of π . Among any 3 horizontally consecutive points, one is vertically between the two others. We call such a point *vertical middle* (in the permutation from Figure 4.2, the vertical middle points are 4, 2 and 3). Similarly, among any 3 vertically consecutive points, one is horizontally between the two others, and we call this point *horizontal middle* (in Figure 4.2, the horizontally middle points are 2, 5 and 4).

Now let us suppose that π has a point x that is simultaneously a horizontal and a vertical middle point. Then x is part of a triple x, b, t (not necessarily in that order) of horizontally consecutive points, where b is the bottom point (the lowest in the triple) and t is the top point (the highest in the triple). Also, x is part of a triple x, l, r (not necessarily in that order) of vertically consecutive points, where l is the leftmost and r is the rightmost point in the triple (see Figure 4.3 (a) for an illustration).

In general, x can be at any of the 9 intersection points of pairs of 3 consecutive vertical and horizontal lines, i.e., x is somewhere in X (see Figure 4.3 (b)). We also have $l \in L$, $r \in R$, $t \in T$ and $b \in B$ for the surrounding points (see Figure 4.3 (b)). The important thing to note is that, since the points are consecutive, those are the *only* points of the permutation lying in the shaded area $X \cup L \cup R \cup T \cup B$. Any point different from x, l, r, t, b lies in one of Q_1, Q_2, Q_3 or Q_4 .

It is not difficult to see that the geometric neighbourhood corresponding to $(N(r) \cap N(b)) \cup (N(l) \cap N(t))$ (see Figure 4.3 (a)) will always contain Q_2 and Q_4 , and will never intersect Q_1 or Q_3 . Therefore, the function that describes how x depends on $\{l, r, t, b\}$ can be written as follows:

$$f(x_r, x_b, x_l, x_t) = x_r x_b \vee x_l x_t,$$

where x_r, x_b, x_l, x_t are Boolean variables corresponding to points r, b, l, t , respec-

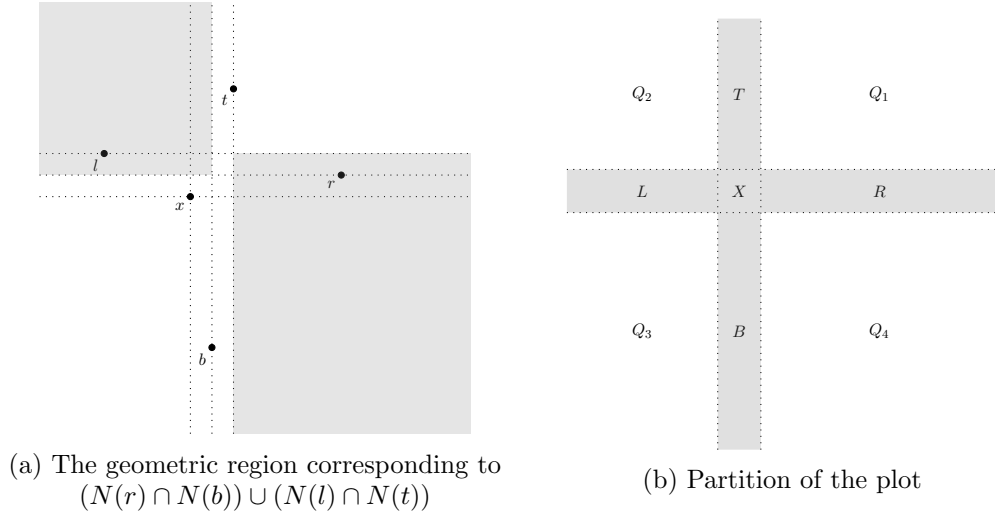


Figure 4.3: A middle point x and its four surrounding points

tively, and the Boolean AND is simply denoted by juxtaposition of the variables. In other words, a vertex $y \notin \{x, l, r, t, b\}$ is adjacent to x if and only if

$$f(A(y, r), A(y, b), A(y, l), A(y, t)) = 1.$$

Step 2: Let us relax the simultaneous middle point condition to the following one: amongst every 5 vertically (respectively horizontally) consecutive points, call the middle three *weak horizontal* (respectively *vertical*) *middle points*. For instance, in Figure 4.2, the weak horizontal middle points are 4, 2, 5 and 3, and those same points are also the weak vertical ones. Note that if the number of points is divisible by 5, at least $\frac{3}{5}$ of them are weak vertical and at least $\frac{3}{5}$ of them are weak horizontal middle points. Using this observation it is not hard to deduce that if there are at least 13 points, then more than half of them are weak vertical and more than half of them are weak horizontal middle points, and so there must exist a point x that is simultaneously both. We only need to deal with this case, as the functionality of any graph on at most 12 vertices is at most 6 (every vertex has at most 6 neighbours or non-neighbours).

Now x is simultaneously a weak vertical and a weak horizontal middle point, and so there must exist consecutive horizontal, respectively vertical quintuples l, x, m_1, m_2, r and t, x, m_3, m_4, b (not necessarily in that order) for which x is a weak middle point (and m_1, m_2, m_3 and m_4 are the other weak middle points in their respective quintuples). By removing m_1, m_2, m_3 and m_4 from the graph, we find ourselves in the configuration of Step 1 and conclude that x is a function of

$\{l, r, t, b\}$ in the reduced graph. Therefore, in the original graph, x is a function of $\{l, r, t, b, m_1, m_2, m_3, m_4\}$, concluding the proof. \square

We next show that permutation graphs have unbounded symmetric difference.

Theorem 99. *For any $t \in \mathbb{N}$, there is a permutation graph G with $\text{sd}(G) \geq t$.*

Proof. Given a permutation graph G corresponding to a permutation π and two vertices x_1, x_2 of G , the symmetric difference of their neighbourhoods can be represented geometrically as an area in the plane (see Figure 4.4). More precisely, a vertex different from x_1 and x_2 lies in the symmetric difference of their neighbourhoods if and only if the corresponding point of the plot of π lies in the shaded area.




Figure 4.4: Geometric symmetric difference of two points x_1 and x_2

In order to prove the theorem, it suffices, for each $t \in \mathbb{N}$, to exhibit a set S_t of points in the plane (with no two on the same vertical or horizontal line) such that for any pair $x_1, x_2 \in S_t$, there are at least t other points of S_t lying in the geometric symmetric difference of x_1 and x_2 . Indeed, such a set can be interpreted as the plot of a permutation, and in the corresponding permutation graph, the symmetric difference of the neighbourhoods of any pair of vertices is at least t .

We construct sets S_t in the following way (see Figure 4.5 for an example – incidentally, the corresponding permutations are the same ones we used in Subsubsection 3.4.3 for their large co-chromatic number):

- start with all the points with integer coordinates between 0 and t inclusive;

- apply to those points the counterclockwise rotation about the origin sending the vector $(1, 0)$ to the unit vector with direction $(t + 1, 1)$ (applying this rotation ensures none of the points share a horizontal or a vertical line).




Figure 4.5: The set S_6

To see that these sets have indeed the desired property, let $x_1, x_2 \in S_t$. For simplicity, we will use the coordinates of the points before the rotation. Suppose $x_1 = (a_1, b_1)$ and $x_2 = (a_2, b_2)$. There are four possible cases (after switching x_1 and x_2 if necessary):

- If $a_1 = a_2$ and $b_1 < b_2$, then the t points $(k, b_2), (l, b_1)$ with $k < a_1 < l$ are in the symmetric difference.
- Similarly, if $b_1 = b_2$ and $a_1 < a_2$, then the t points $(a_1, k), (a_2, l)$ with $k < b_1 < l$ are in the symmetric difference.
- If $a_1 < a_2$ and $b_1 < b_2$, the following points all lie in the symmetric difference of x_1 and x_2 :
 - (1) Points (a_1, k) with $k < b_1$ (in the bottom region).
 - (2) Points (a_1, k) with $b_1 < k \leq b_2$ (in the left region).
 - (3) Points (a_2, k) with $b_2 < k$ (in the top region).
 - (4) Points (a_2, k) with $b_1 \leq k < b_2$ (in the right region).

In particular, (1) and (3) account for at least $b_1 + t - b_2$ points, while (2) and (4) account for $2(b_2 - b_1)$ others. We conclude that in total, at least $t + (b_2 - b_1) > t$ points lie in the symmetric difference of x_1 and x_2 .

- If $a_1 < a_2$ and $b_1 > b_2$, the following points all lie in the symmetric difference of x_1 and x_2 :

- (1) Points (k, b_2) with $a_1 \leq k < a_2$ (in the bottom region).
- (2) Points (k, b_1) with $k < a_1$ (in the left region).
- (3) Points (k, b_1) with $a_1 < k \leq a_2$ (in the top region).
- (4) Points (k, b_2) with $a_2 < k$ (in the right region).

Summing up, we find once more at least t points in the symmetric difference. □

Remark 100. We note that, in conjunction with Theorem 95, this result yields an alternative (and in our opinion, ever so slightly easier) proof of the already known fact [GR00] that permutation graphs have unbounded clique-width.

The final class that we consider in this subsection is the class of line graphs (the line graph of a graph G is the intersection graph of its edges). This is another class of unbounded clique-width (see, e.g., [GW07]).

Theorem 101. *The functionality of line graphs is at most 6.*

Proof. Let G be a graph and H be the line graph of G . Since the class of line graphs is hereditary, it suffices to prove that H has a vertex of functionality at most 6. We will prove a stronger result showing that *every* vertex of H has functionality at most 6.

Let x be a vertex in H , i.e., an edge in G . We denote the two endpoints of this edge in G by a and b . Assume first that both the degree of a and the degree of b are at least 4. Let $Y = \{y_1, y_2, y_3\}$ be a set of any three edges of G different from x that are incident to a , and let $Z = \{z_1, z_2, z_3\}$ be a set of any three edges of G different from x that are incident to b .

We claim that a vertex $v \notin \{x\} \cup Y \cup Z$ is adjacent to x in H if and only if it is adjacent to every vertex in Y or to every vertex in Z . Indeed, if v is adjacent to x in H , then the edge v intersects the edge x in G . If the intersection consists of a , then v is adjacent to every vertex in Y in the graph H , and if the intersection consists of b , then v is adjacent to every vertex in Z in the graph H . Conversely, let v be adjacent to every vertex in Y , then v must intersect the edges y_1, y_2, y_3 in G at vertex a , in which case v is adjacent to x in H . Similarly, if v is adjacent to every vertex in Z , then v intersects the edges z_1, z_2, z_3 in G at vertex b and hence v is adjacent to x in H .

Therefore, in the case when both a and b have degree at least 4 in G , the function that describes how x depends on $\{y_1, y_2, y_3, z_1, z_2, z_3\}$ in the graph H can be written as follows: $f(y_1, y_2, y_3, z_1, z_2, z_3) = y_1 y_2 y_3 \vee z_1 z_2 z_3$.

If the degree of a is less than 4, we include in Y all the edges of G distinct from x which are incident to a (if there are any) and remove the term $y_1 y_2 y_3$ from the function. Similarly, if the degree of b is less than 4, we include in Z all the edges of G distinct from x which are incident to b (if there are any) and remove the term $z_1 z_2 z_3$ from the function. If both terms have been removed, the function is defined to be identically 0, i.e., no vertices are adjacent to x in H , except for those in $Y \cup Z$. \square

One may hope to generalise this result and show bounded functionality for the line graphs of k -uniform hypergraphs (that is, intersection graphs of families of k -subsets). It turns out that this is a substantially more difficult task; in [AAL19], we only manage to solve it for $k = 3$, with an upper bound of 462 (we omit the proof here).

4.2.2 Graphs of large functionality

In the previous section, we saw several examples of classes with bounded functionality. It is natural to now ask what classes have unbounded functionality. The first obvious example consists of all superfactorial classes. Indeed, this follows immediately from the discussion in Subsection 4.1.3 about functional vertices and functional representations. This allows us to compare functionality with another significant graph parameter, VC-dimension:

Definition 102. A set system (X, S) consists of a set X and a family S of subsets of X . A subset $A \subseteq X$ is *shattered* if for every subset $B \subseteq A$ there is a set $C \in S$ such that $B = A \cap C$. The VC-dimension of (X, S) is the cardinality of a largest shattered subset of X .

The VC-dimension of a graph $G = (V, E)$ was defined in [Alo+06] as the VC-dimension of the set system (V, S) , where S the family of closed neighbourhoods of vertices of G , i.e., $S = \{N[v] : v \in V(G)\}$. We denote the VC-dimension of G by $\text{vc}(G)$.

It is shown in [Loz18] that the only minimal classes of unbounded VC-dimension are bipartite, co-bipartite and split graphs. Since all of those classes are superfactorial ($\Omega(2^{n^2/4})$ graphs on n vertices), bounded functionality implies bounded VC-dimension.

Of course, the fact that, say, the class of bipartite graphs (or even less helpfully, the class all graphs) has unbounded functionality does not actually give us much insight into which individual graphs have high functionality. In particular, it would be interesting to construct specific graphs in which every vertex has high functionality. Another interesting item to find would be a factorial class of graphs of unbounded functionality, that is, a class which has unbounded functionality for non-trivial reasons. The example of hypercubes fulfills both of these criteria (and also shows that VC-dimension is *strictly* stronger than functionality).

Definition 103. Let $V_n = \{0, 1\}^n$ be the set of binary sequences of length n and let $v, w \in V_n$. The Hamming distance $d(v, w)$ between v and w is the number of positions in which the two sequences differ. The *hypercube* Q_n is the graph with vertex set $V_n = \{0, 1\}^n$, in which two vertices are adjacent if and only if the Hamming distance between them equals 1.

Theorem 104. *Functionality of the hypercube Q_n is at least $(n - 1)/3$.*

Proof. By symmetry, it suffices to show that the vertex $v = 00\dots 0 \in V_n$ has functionality at least $(n - 1)/3$. Suppose v is a function of vertices in a set $S \subseteq V_n \setminus \{v\}$. To provide a lower bound on the size of S , and hence a lower bound on the functionality of v , for each $i = 1, 2, \dots, n$ consider the set $S_i = \{w \in S : d(w, v) = i\}$, i.e., the set of all binary sequences in S that contain exactly i 1s. Also, consider the following set:

$$I = \{i \in \{1, 2, \dots, n\} : \exists z = z_1 z_2 \dots z_n \in S_1 \cup S_2 \cup S_3 \text{ with } z_i = 1\}.$$

Suppose $|I| \leq n - 2$. Then there exist two positions i and j such that for any sequence $z = z_1 z_2 \dots z_n \in S_1 \cup S_2 \cup S_3$, we have $z_i = 0$ and $z_j = 0$. Consider the following two vertices:

- $u = u_1 u_2 \dots u_n$ with $u_k = 1$ if and only if $k = i$,
- $w = w_1 w_2 \dots w_n$ with $w_k = 1$ if and only if $k = i$ or $k = j$.

We claim that u and w are not adjacent to any vertex $z \in S$. First, it is not hard to see that for any $z \in S_1 \cup S_2 \cup S_3$ we have $d(z, u) \geq 2$ and $d(z, w) \geq 2$. Indeed, any $z \in S_1 \cup S_2 \cup S_3$ differs from u and w in position i , i.e., $z_i = 0$ and $u_i = w_i = 1$, and there must exist a $k \neq i, j$ with $z_k = 1$ and $u_k = w_k = 0$. Also, it is not difficult to see that $d(z, u) \geq 2$ and $d(z, w) \geq 2$ for any vertex $z \in S \setminus (S_1 \cup S_2 \cup S_3)$, because any such z has at least four 1s, while u and w have at most two 1s. Therefore, by definition, u and w are not adjacent to any vertex in S .

We see that the assumption that $|I| \leq n-2$ leads to the conclusion that there are two vertices $u, w \in Q_n \setminus (S \cup \{v\})$ which are non-adjacent to any vertex in S , but have different adjacencies to v . This contradicts the fact that v is a function of the vertices in S . So, we must conclude that I has size at least $n-1$. As each vertex in $S_1 \cup S_2 \cup S_3$ has at most three 1s, we conclude that $S_1 \cup S_2 \cup S_3$ must contain at least $|I|/3 = (n-1)/3$ vertices. This completes the proof of the theorem. \square

As it turns out, the hereditary closure of hypercubes (that is, the class containing all hypercubes and their induced subgraphs) has factorial speed of growth. This result is due to appear in [Ale+21a].

Theorem 105 ([Ale+21a]). *The hereditary closure of hypercubes is a factorial class.*

Proof. Let \mathcal{Q} be the hereditary closure of the hypercubes Q_n . It is not hard to see that this class contains the class of trees, therefore \mathcal{Q} has at least factorial speed of growth. It remains to show that it has at most factorial speed of growth, and for this, it suffices to restrict our attention to connected graphs in \mathcal{Q} . Indeed, if we show that the set of connected graphs in \mathcal{Q} has at most factorial speed, then prime graphs (a subset of the connected graphs) have at most factorial speed, which implies at most factorial growth of \mathcal{Q} [Atm+15, Theorem 1].

Let $G \in \mathcal{Q}$ be a connected graph on n vertices. By definition of \mathcal{Q} , G embeds into Q_m for some m . We claim that, in fact, G embeds into Q_{n-1} . If $m < n$, this is clear. Otherwise, using an embedding into Q_m , each vertex of G corresponds to an m -digit binary sequence. For two adjacent vertices, the sequences differ in exactly one position. From this, it follows inductively that the n vertices of G all agree in at least $m - (n-1)$ positions. The coordinates on which they agree can simply be removed; this produces an embedding of G into Q_{n-1} .

Now write x_1, \dots, x_n for the vertices of G , and let $\varphi : V(G) \rightarrow \{0, 1\}^{n-1}$ be an embedding of G into Q_{n-1} . Without loss of generality, we may assume that $\varphi(x_1) = (0, 0, \dots, 0)$. Consider a spanning tree of G rooted at x_1 , and for $i = 2, 3, \dots, n$, define $p(i)$ to be the index of the parent of x_i . Since x_i and $x_{p(i)}$ are connected, $\varphi(x_i)$ and $\varphi(x_{p(i)})$ differ exactly in one position: we denote it by $d(i)$.

We claim that G can be restored from the sequence $p(2), d(2), p(3), d(3), \dots, p(n), d(n)$. Indeed, this information is enough to identify all children of x_1 , then determine their images via φ ; proceeding inductively, one obtains the embeddings of all vertices of G into Q_{n-1} . Finally, we note that the above sequence uses $O(n \log n)$ bits (since it consists of $2(n-1)$ integers of size at most n). \square

4.3 Further directions of research

Functionality is a very new parameter, and as such, we are still in the process of trying to understand it and to place it in the general graph-theoretic landscape. We recognise that, at the moment, we do not have a concrete, practical utility for it. That is not to say that such a niche cannot exist!⁸ In many applied situations concerning graph representation, however, it seems that whatever functional representations do, implicit representations do better.

Nevertheless, the study of functionality provides, at the very least, a new source of counterexample candidates to the Implicit Graph Conjecture. Indeed, assuming the conjecture is true, each class of bounded functionality should have an implicit representation. We note that functional representations seem more difficult to obtain than implicit ones. This is what we would expect; after all, implicit representations give us the freedom to individually change each of the $\Theta(\log n)$ bits for each vertex, while functional ones use somewhat clunky blocks of information that effectively describe how a vertex' neighbourhood relates to the neighbourhoods of other vertices. Despite this, it is not at all obvious how to go from bounded functionality to implicit representations. As a sidenote, using recent results from [Har19], it can be shown that the hereditary closure \mathcal{Q} of hypercubes (a factorial class of *unbounded* functionality) does admit an implicit representation.⁹ In general, we believe it would be interesting to further examine the relationship between functionality and implicit representations in conjunction with the Implicit Graph Conjecture, and try to give constructive answers for questions such as:

Open problem 106. Does every class of bounded functionality admit an implicit representation?

With all that said, we firmly believe that functionality deserves to be more thoroughly studied in its own right, forgetting its origin and ties to representation

⁸For instance, representing a vertex' adjacency as a function could be of interest in the area of graph learning and graph mining, since it makes graphs amenable to the techniques of Logical Analysis of Data [Lej+19]. Moreover, a parameter similar to symmetric difference called *twin-width* has recently been introduced independently in a series of preprints starting with [Bon+20b]. It is easy to see that bounded twin-width implies bounded symmetric difference, but it is not clear whether the converse implication holds; at any rate, the authors of [Bon+20b] propose algorithmic applications for twin-width. Later in the preprint series [Bon+20a], they also show that classes of bounded twin-width admit implicit representations.

⁹We remark that this is a non-trivial result: the labels coming from the hypercubes themselves may be too large for some of the graphs in \mathcal{Q} to constitute an implicit representation. Equivalently, we have seen in the proof of Theorem 105 that (connected) n -vertex graphs from \mathcal{Q} are in fact embeddable into Q_{n-1} ; what the result is really saying is that it is possible to embed all n -vertex graphs from \mathcal{Q} into some fixed graph of size *polynomial* in n .

and enumeration. There are several reasons for that – let us go through a few of them.

First, functionality has the potential to provide us with unexpected insight into various graph classes. For instance, the fact that cographs always have twin vertices has been known for some time [BLS99], but who would have guessed that permutation graphs always have a vertex whose neighbourhood can be neatly described in terms of at most 8 (perhaps even at most 4, if we are looking to optimise Theorem 98) other vertices? While this might seem like an arcane piece of knowledge at this point in time, it is not inconceivable that such a fact might prove itself useful in the future, when deriving other results about permutation graphs. There are many classes that are in a similar situation. Perhaps functionality could even be used to systematically obtain characterisations for them as the smallest classes closed under certain operations (such as adding a vertex with specific functional relations). We leave this as an open problem:

Open problem 107. Determine whether the class of permutation graphs can be characterised as the smallest hereditary class closed under the addition of vertices whose neighbourhood is of the form $(N(r) \cap N(b)) \cup (N(l) \cap N(t))$, where r, b, l, t are pre-existing vertices. Do similar characterisations hold for other well-studied classes?

Second, besides shedding light onto individual classes, functionality can also provide insight into various parameters. For example, Theorem 95 tells us that bounded clique-width implies the existence of vertices that are “almost twins”. What similar results can be derived from boundedness of other parameters? Could these results be used to gain a better understanding of when these parameters are bounded? In this parametric direction, we also point out that functionality appears unusually strong for a parameter whose boundedness implies factorial speed. In particular, we do not know of a “naturally defined” parameter whose boundedness implies factorial speed, and is strictly implied by the boundedness of functionality. This raises the further questions of whether the existence of implicit representations within factorial hereditary classes can be characterised via boundedness of a certain hereditary parameter,¹⁰ and whether factorial speed itself can be characterised in that way:¹¹

¹⁰Chandoo [Cha17] studies a variation on this problem in the setting of formal languages; however, in his complexity-focused analysis, he does not restrict the problem to factorial hereditary classes, and he also allows non-hereditary parameters. Moreover, the parameters he does find do not have particularly enlightening descriptions.

¹¹In [Loz18], such parameters are identified for other jumps in the speed hierarchy. For instance,

Open problem 108. Are there parameters whose boundedness lies “between” bounded functionality and the existence of implicit representations? What about between functionality and factorial speed of growth?

Open problem 109. Does the existence of implicit representations among factorial hereditary classes admit a parametric characterisation? What about factorial speed of growth?

Third, and somewhat more subjectively, we feel that the study of functionality comes with a very well-balanced mixture of challenge and intrigue – so much so that we are surprised we did not find similar studies in the literature.¹² Plainly put, studying functionality is fun! The reader does not have to blindly believe us. Instead, they are cordially invited to pick their favourite factorial class or their favourite parameter, and try to determine how it relates to functionality. As a small sample of specific classes for which the boundedness of functionality is not yet known, we suggest the following (some are motivated by results from the previous section, and some by their relevance in studying implicit representations):

Open problem 110. For each of the following classes, determine whether or not they have bounded functionality/symmetric difference:

- interval graphs;
- k -interval graphs;
- graphs of bounded boxicity (see [Rob69]);
- line graphs of k -regular hypergraphs;
- comparability graphs of partial orders of dimension k ;
- (unit) disk intersection graphs;
- line segment intersection graphs.

Hopefully, the above discussion will have convinced the reader that functionality is worthy of further consideration, if for no other reason than the sheer number

a class is subfactorial if and only if it has bounded *neighbourhood diversity*. That being said, there is no obvious reason this should be the case here; after all, factorial speed is a global property of the class, while a parametric description would express it as a local property of each graph. We note this kind of global-local interplay is also featured in the Implicit Graph Conjecture, hence a negative answer here would certainly be of great interest.

¹²The author dreads the day when it is brought to his attention that the ideas we have presented here appear in a foreign operational research journal from the 90s, using very different terminology...

of interesting questions engendered by its study so far. To finish this chapter, we present one final problem on the topic, together with our progress so far, and some possible avenues towards a solution. The problem is as follows:

Open problem 111. What is, asymptotically, the minimum number of vertices among graphs of functionality at least k ? Equivalently: what is, asymptotically, the maximum value of functionality among graphs on at most n vertices?

Write $m(k)$ for the minimum number of vertices in a graph of functionality at least k . It is clear that $m(k)$ is realised by a graph where every vertex has functionality at least k – by analogy with degeneracy, we call such a graph a *functional k -core*. Our aim is to find asymptotic bounds for $m(k)$. From the example of hypercubes it follows that $m(k) \leq 2 \cdot 8^k$ for all k . But what about lower bounds?

One immediately notices that $m(k) > 2k$, since in a graph of size $2k$, every vertex has either degree or co-degree smaller than k . But this leaves a very wide gap of possible growth speeds for $m(k)$. Ideally, we would like to be able to say at the very least whether the growth is polynomial or exponential. The following is the only non-trivial lower bound we have managed to obtain so far:

Theorem 112. $m(k) = \Omega(k^\alpha)$ for any $1 \leq \alpha < 2$.

Notation 113. Given two vertices u, v we will denote by $\iota(u, v)$ the number of vertices other than u and v that have identical adjacency to u and v , and by $\delta(u, v)$ the number of vertices other than u and v that have different adjacency to u and v .

Proof. To prove the theorem, we need a bit of setup. We consider functionality from a slightly different perspective: to say that vertex y is *not* a function of vertices x_1, \dots, x_k is the same as saying that there are two vertices w and w' that have the same adjacency to each of the x_i s, but different adjacencies to y . In that sense, the pair $\{w, w'\}$ provides a witness that y is not a function of x_1, \dots, x_k . Thus, for a graph $G = (V, E)$ with $|V| = n$ to be minimal of functionality at least $k + 1$, every pair (S, y) with $S \subseteq V, |S| = k$ and $y \in V \setminus S$ needs such a witness $\{w, w'\}$ with $w, w' \notin S \cup \{y\}$. The main idea of the proof consists of showing that if $n = O(k^\alpha)$ for some $1 \leq \alpha < 2$, then for large k , there will not be enough witnesses in some sense.

We need to be careful though: a “brute force” counting argument does not appear to work. Indeed, there are $n \binom{n-1}{k}$ pairs (S, y) , and there are

$$\sum_{\{w, w'\} \subseteq V} \binom{\iota(w, w')}{k} \delta(w, w')$$

“witness slots” (we are counting for how many different (S, y) each pair of vertices can be a witness). While there are graphs where the two numbers are equal (for example C_5), the expression on the right is difficult to control even for n linear in k . Indeed, the problem is that many vertices will be witnesses to the *same* (S, y) , and this brute force method of counting does not account for that, meaning this approach loses a lot of power.

We need to do something a bit more sophisticated instead. For notational reasons, we will assume from now on that our graph has $n + 1$ vertices. If we look from the perspective of a single vertex y , the only possible pairs witnessing y is not a function of S (for any given set S of k vertices) consist of a neighbour of y and a non-neighbour of y . The number of such pairs is bounded above by $n^2/4$ (the bound is achieved when the vertex has degree $n/2$). Moreover, we know that each k -subset of $V \setminus \{y\}$ must have the same adjacency to at least one of those pairs. In other words, for each pair (u, v) of vertices, let $I_{u,v}$ be the set of vertices that have the same adjacency to u and v ; then the sets $I_{u,v}$ cover all k subsets of $V \setminus \{y\}$, in the sense that every k -subset is contained in some $I_{u,v}$.

Now note that we can bound the size $\iota(u, v)$ of the $I_{u,v}$ s above: if we assume that every vertex of G has functionality at least $k + 1$, this means that for any two vertices u and v , $\delta(u, v)$ is at least k . From this, $\iota(u, v)$ is at most $n - 1 - k$. So our discussion from last paragraph implies that there is way of covering all k -subsets of an n -set with at most $n^2/4$ subsets of size at most $n - 1 - k$. There is already theory regarding such coverings (see e.g. [Sch64]). Let $C(n, t, k)$ denote the minimum number of t -subsets of an n -set required to cover all k -subsets (the names of the variables have been changed from [Sch64] to suit our purposes). Our discussion so far showed that, if some vertex in a graph G with $n + 1$ vertices has functionality at least $k + 1$, then

$$C(n, n - 1 - k, k) \leq \frac{n^2}{4}.$$

Using the so-called Schönheim bound [Sch64], the left hand side of the above inequality is bounded below by

$$\left\lceil \frac{n}{n - 1 - k} \left\lceil \frac{n - 1}{n - 2 - k} \cdots \left\lceil \frac{n - k + 1}{n - 2k} \right\rceil \right\rceil \right\rceil \geq \left(\frac{n}{n - (k + 1)} \right)^k.$$

Putting everything together, we have

$$\left(\frac{n}{n - (k + 1)} \right)^k \leq \frac{n^2}{4}.$$

As functions of n , the left hand side is decreasing (tending to 1), while the right hand side is increasing. So if we assume that $n \leq \beta(k+1)^\alpha$ for some $1 \leq \alpha < 2$ and $\beta > 0$, we may substitute for n and write $l := k+1$ to obtain

$$\left(\frac{\beta l^\alpha}{\beta l^\alpha - l}\right)^{l-1} \leq \frac{\beta^2 l^{2\alpha}}{4}.$$

Rewrite the left hand side as

$$\left(\frac{\beta l^{\alpha-1}}{\beta l^{\alpha-1} - 1}\right)^{l-1} = \left(1 + \frac{1}{\beta l^{\alpha-1} - 1}\right)^{l-1}.$$

Taking logarithms of both sides of the inequality yields

$$(l-1) \log\left(1 + \frac{1}{\beta l^{\alpha-1} - 1}\right) \leq 2\alpha \log l + 2 \log \beta - \log 4.$$

Using the expansion $\log(1+x) = x - \frac{x^2}{2} + O(x^3)$, we obtain:

$$(l-1) \left(\frac{1}{\beta l^{\alpha-1} - 1} + \frac{1}{2(\beta l^{\alpha-1} - 1)^2} + O((\beta l^{\alpha-1})^{-3})\right) = O(\log l).$$

The leading term on the left hand side has order $l^{2-\alpha}$, thus the above fails when $\alpha < 2$. The theorem follows. □

We note that the above proof only assumes the existence of a single vertex of high functionality. In order to improve on the bound, a good place to start would be trying to gain a better understanding of functional k -cores (and use the fact that *all* vertices have high functionality).

There is one more line of reasoning that might prove helpful in solving this and other problems. Suppose y is a function of a set $X = \{x, z_1, \dots, z_k\}$, where $\{z_1, \dots, z_k\}$ is the symmetric difference of x and y . What we are really saying is that y is a twin of x , as long as we ignore the vertices z_1, \dots, z_k . In this sense, y is a “simple” function of X , since its adjacency outside X only depends on a single vertex in X . It might make sense to try and formalise this idea; we propose the following definition:

Definition 114. Let $f : \{0, 1\}^k \rightarrow \{0, 1\}$ be a Boolean function. The f -degeneracy of a graph G is the minimum t such that in every induced subgraph $H \subseteq G$, we may remove t vertices to obtain a graph H' in which there exist vertices y, x_1, \dots, x_k with $y = f(x_1, \dots, x_k)$.

Degeneracy can be thought of the special case where f is identically 0 (abusing notation, we may put $k = 0$ in the above definition); symmetric difference is the case when $f : \{0, 1\} \rightarrow \{0, 1\}$ is the identity. We note that if $f : \{0, 1\}^k \rightarrow \{0, 1\}$, the functionality of G is bounded above by its f -degeneracy plus k . There are many directions in which we can go from here; it would be natural, for instance, to define F -degeneracy for a set of functions. We could also look into how properties of Boolean functions (such as reducibility as a direct product) translate into properties of the corresponding parameters (and classes where those parameters are bounded). We leave this as an open-ended direction for further research.

Chapter 5

The micro-world of cographs

In this chapter, we present our joint work with Vadim Lozin and Dominique de Werra [ALW21]. We take our magnifying glass out of the drawer and focus on the class of cographs. This class has relatively simple structure, as highlighted by the cotree description of cographs, and by myriad of other known characterisations (see the Introduction). Nevertheless, cographs constitute a rich and complex world in which many popular graph parameters jump to infinity on very specific subclasses. An example of this phenomenon is given by linear clique-width: a sophisticated approach shows that there are exactly two minimal subclasses of cographs where it is unbounded [BKV17]. Our analysis reveals several additional instances of this phenomenon.

5.1 Well-quasi-order and beyond: a prelude to our analysis

In general, it is possible for a hereditary parameter κ to be unbounded in a hereditary class \mathcal{X} , without there being a *minimal* subclass $\mathcal{Y} \subseteq \mathcal{X}$ in which κ is unbounded. We can easily construct examples: for instance, if we let $\kappa(G)$ be the largest k such that C_k is induced in G and assume κ is unbounded in \mathcal{X} , then \mathcal{X} must contain infinitely many cycles. Forbidding any one cycle C_i does not bound κ in $\mathcal{X} \cap \text{Free}(C_i)$, and hence no class in which κ is unbounded is actually minimal with that property.

What if we were to restrict ourselves to the universe of cographs? If a parameter κ is unbounded in a subclass \mathcal{X} of cographs, can we infer the existence of a minimal class $\mathcal{Y} \subseteq \mathcal{X}$ where κ is unbounded? Indeed, we can! The class of cographs is well-quasi-ordered by the induced subgraph relation [Dam90], and it is a very standard result that wqo of a class \mathcal{X} under induced subgraphs is equivalent to well-

foundedness of the set of subclasses of \mathcal{X} under inclusion. For completeness (and as a warm-up exercise), let us prove this fact now – our claim about the existence of minimal classes will then follow immediately from it.

Proposition 115. *Let \mathcal{Z} be a hereditary class. \mathcal{Z} is wqo under induced subgraphs if and only if the set of hereditary subclasses of \mathcal{Z} is well-founded under inclusion.*

Proof. For the “only if” direction, suppose that \mathcal{Z} is wqo, and let $\mathcal{Z}_1 \supseteq \mathcal{Z}_2 \supseteq \dots$ be a descending chain of classes. Write $\mathcal{Z}_* := \bigcap_{i=1}^{\infty} \mathcal{Z}_i$. Let T be the set of minimal forbidden induced subgraphs for \mathcal{Z}_* , and let $S := T \cap \mathcal{Z}$, so that $\mathcal{Z}_* = \mathcal{Z} \cap \text{Free}(S)$. Then S is an antichain in \mathcal{Z} , and thus finite. Each element in S stops appearing at some point in the chain, hence there must exist a k such that $\mathcal{Z}_k \subseteq \mathcal{Z} \cap \text{Free}(S) = \mathcal{Z}_*$. This shows the chain is eventually stationary.

For the “if” direction, we show the contrapositive: if \mathcal{Z} is not wqo, then there exists an infinite antichain G_1, G_2, \dots of graphs in \mathcal{Z} . Putting $\mathcal{Z}_i := \mathcal{Z} \cap \text{Free}(G_1, \dots, G_i)$, we obtain an infinite, strictly descending chain of subclasses of \mathcal{Z} . □

Corollary 116. *For any parameter κ , if \mathcal{X} is a class of cographs of unbounded κ , then \mathcal{X} contains a class \mathcal{Y} that is minimal of unbounded κ .*

Proof. If $\mathcal{X}_0 := \mathcal{X}$ itself is minimal, we are done. Otherwise, find $G_0 \in \mathcal{X}_0$ such that $\mathcal{X}_1 := \mathcal{X}_0 \cap \text{Free}(G_0)$ has unbounded κ . In general, if \mathcal{X}_i is not minimal, find $G_i \in \mathcal{X}_i$ such that $\mathcal{X}_{i+1} := \mathcal{X}_i \cap \text{Free}(G_i)$ has unbounded κ . Since $\mathcal{X}_0 \supsetneq \mathcal{X}_1 \supsetneq \dots$ is a strictly descending chain, it must be finite by the previous proposition. Thus the process must terminate with some minimal class \mathcal{X}_k . □

For the rest of the chapter, we will forget that there exist any graphs other than cographs. In particular, unless stated otherwise, whenever we say “class”, we will mean “subclass of cographs”. Moreover, when we say a parameter κ_1 is *stronger* than a parameter κ_2 , we will mean that boundedness of κ_2 in a subclass \mathcal{X} of cographs implies boundedness of κ_1 in \mathcal{X} . The above discussion implies that (un)boundedness of any parameter κ can be uniquely characterised in terms of a set $M(\kappa)$ of minimal classes where κ is unbounded.¹ Among other things, this makes comparisons between parameters particularly straightforward, as shown in the following lemma:

¹One may in fact identify a parameter κ with the set $\mathcal{B}(\kappa)$ of classes in which κ is bounded. This set is downwards closed under inclusion. From this point of view, the theory of parameters is completely analogous to the theory of hereditary classes: the set $M(\kappa)$ corresponds to the set of minimal forbidden induced subgraphs. This perspective also legitimises seemingly incomplete turns of phrase such as “characterising κ via minimal classes”.

Lemma 117. *Parameter κ_1 is stronger than κ_2 if and only if for every minimal hereditary class $\mathcal{F}_1 \in M(\kappa_1)$, there is a minimal hereditary class $\mathcal{F}_2 \in M(\kappa_2)$ such that $\mathcal{F}_2 \subseteq \mathcal{F}_1$.*

Proof. To prove the “if” direction, assume that for every minimal hereditary class \mathcal{F}_1 where κ_1 is unbounded, there is a minimal hereditary class \mathcal{F}_2 where κ_2 is unbounded such that $\mathcal{F}_2 \subseteq \mathcal{F}_1$. Now let \mathcal{X} be a κ_2 -bounded hereditary class. Since any minimal class where κ_1 is unbounded contains a class where κ_2 is unbounded, it follows \mathcal{X} cannot contain any minimal class where κ_1 is unbounded, and so by Corollary 116, \mathcal{X} is κ_1 -bounded, showing κ_1 is stronger than κ_2 .

Conversely, suppose κ_1 is stronger than κ_2 , and let \mathcal{F}_1 be a minimal hereditary class where κ_1 is unbounded. Since κ_1 is stronger, κ_2 is also unbounded in \mathcal{F}_1 , and by Corollary 116, \mathcal{F}_1 contains a minimal class \mathcal{F}_2 where κ_2 is unbounded, as required. \square

It turns out that we can even say a bit more about the sets $M(p)$, namely that they must be finite. The reason for this is that the set of subclasses of cographs under inclusion is not just well-founded, but it is actually itself well-quasi-ordered! This is a surprisingly non-trivial fact tied to the notion of *better-quasi-ordering* (“bqo”, for short). This notion has a fascinating history dating back to Nash-Williams in the 1960s [Nas65]. However, even defining it properly would be a long digression outside of the scope of this chapter. We will therefore content ourselves with a short bulleted list attempting to cover the main ideas and motivation behind the definition, as well as the facts that are relevant to this chapter. The reader wishing to dive deeper into this alluring rabbit hole may consult the first part of [AH07] and the references therein for a gentle, algorithmic algebra-flavoured introduction to the topic.²

Here is our rough understanding of bqo and the main facts about it that are important to us. We stress that this is only a very diluted, incomplete and non-rigorous introduction to the notion, and we certainly make no claim that it is the best or only way to conceptualise it:

- The notion of wqo is of great practical importance, but it lacks certain desirable algebraic properties. Specifically, it does not “lift” to the power set of a quasi-order. This can be formulated in a general, order-theoretic setting; for our modest graph-theoretic purposes, what this means is that there might exist

²Another very good and much more comprehensive introduction can be found in [FBS20]; that manuscript is however still unfinished.

a class \mathcal{Z} of graphs (not necessarily cographs) which is wqo under induced subgraphs, but whose set of subclasses is not wqo under inclusion.

- The stronger notion of bqo “fixes” this problem: in our setting, if a class of graphs is bqo under induced subgraphs, its set of subclasses is also bqo under inclusion. More explicitly, we may think of bqo as follows: if a class \mathcal{X} is bqo under induced subgraphs then it is wqo under induced subgraphs, and its set of (downwards-closed) subclasses is wqo under inclusion, and the set of downwards-closed sets of subclasses is also wqo under inclusion,³ and so on, to infinity and beyond.⁴
- The “definition” above is fundamentally algebraic; it can be rephrased in terms of fixed points of a certain function, and there is much theory in that direction. There is an equivalent definition that is purely combinatorial. Usual wqo can be defined by the absence of a certain structure in our poset: a so-called bad sequence in which there are no elements $a_i \leq a_j$ with $i < j$. A sequence is a function from \mathbb{N} to our poset; one can generalise this by defining an *array* to be a function from a *block* to our poset. Without going too much into detail, bqo can be expressed as the absence of bad arrays indexed by appropriately defined blocks.
- It turns out that proving bqo of a certain poset is not always as difficult as the ostensibly obscure definition might lead us to believe. In particular, any finite poset is trivially bqo, and bqo is preserved by certain constructions. To see what we mean by “certain constructions”, we look at Kruskal’s famous Tree Theorem: if (X, \leq) is wqo, then the set of finite trees labelled by X is wqo under tree embedding. In fact, the original motivation for introducing the notion of bqo was to extend Kruskal’s theorem to infinite trees: in [Nas65], Nash-Williams proves that infinite (in addition to finite) trees are actually bqo. Later, Laver [Lav71, Theorem 2.2] extended this to *labelled* trees, provided the set of labels is bqo.
- We are close to showing (or more accurately, compiling old results from which it follows that) cographs are bqo under induced subgraphs. A map $f : (X, \leq) \rightarrow (Y, \preceq)$ is called a *quasi-embedding* if, for all $a, b \in X$, $f(a) \preceq f(b) \implies a \leq b$.

³Remark that, as described in a previous footnote, parameters can be thought as downwards-closed sets of subclasses, and their inclusion is just the reverse of the “stronger than” relation. In other words, in a bqo class, parameters are wqo by their relative strength!

⁴This is no exaggeration. The definition of bqo requires wqo to hold after any finite number of applications of a certain power set construction, but also after a transfinite number of them – a step that the author of this thesis is not yet completely comfortable with.

b. It is immediate from the definitions that, if there exists a quasi-embedding $X \rightarrow Y$, and Y is wqo, then X is wqo. This remains true when replacing “wqo” with “bqo” (see, e.g., [AH07], Lemma 5.3). Damaschke’s proof that cographs are wqo [Dam90] consist of exhibiting a quasi-embedding from cographs to the set of finite trees labelled by 4 labels. But as seen in the previous bullet point, this set is bqo!

The additional strength of bqo (as opposed to just wqo) can appear subtle at first, but it is in fact very effective, and allows us to derive concrete results about cographs. For a poset (X, \leq) , we denote by $\mathcal{L}(X)$ the set of downwards-closed sets of X . The strength of bqo can be summarised with the following proposition (see, e.g., [AH07]).

Proposition 118. *Suppose (X, \leq) is bqo. Then (X, \leq) is wqo, and $(\mathcal{L}(X), \subseteq)$ is bqo.*

As an immediate consequence, we draw the following conclusions.

Corollary 119. *The set of hereditary subclasses of cographs is wqo by inclusion.*

Corollary 120. *The set of parameters is wqo by their strength in the class of cographs.*

In particular, we note that, for any parameter κ , the set $M(\kappa)$ of minimal subclasses of cographs in which κ is unbounded is an antichain, and hence finite as claimed. We are now ready to begin our analysis.

5.2 A hierarchy of parameters

Let us start by presenting straight away the outcome of our study, which we summarise in the following Hasse diagram, in Figure 5.1:

The parameters, which we will define later in the section where necessary, are compared by their strength (with stronger parameters higher up in the diagram). Each of them is given with the set $M(\kappa)$ of minimal classes where it is unbounded. The featured classes are as follows ($\overline{\mathcal{X}}$ denotes the class of complements of graphs in \mathcal{X}):

\mathcal{Q} the class of *quasi-threshold graphs*, i.e., (P_4, C_4) -free graphs (see, e.g., [YCC96]).

\mathcal{T} the class of *threshold graphs*. This is the class of $(P_4, C_4, 2K_2)$ -free graphs, i.e., the intersection of \mathcal{Q} and $\overline{\mathcal{Q}}$.

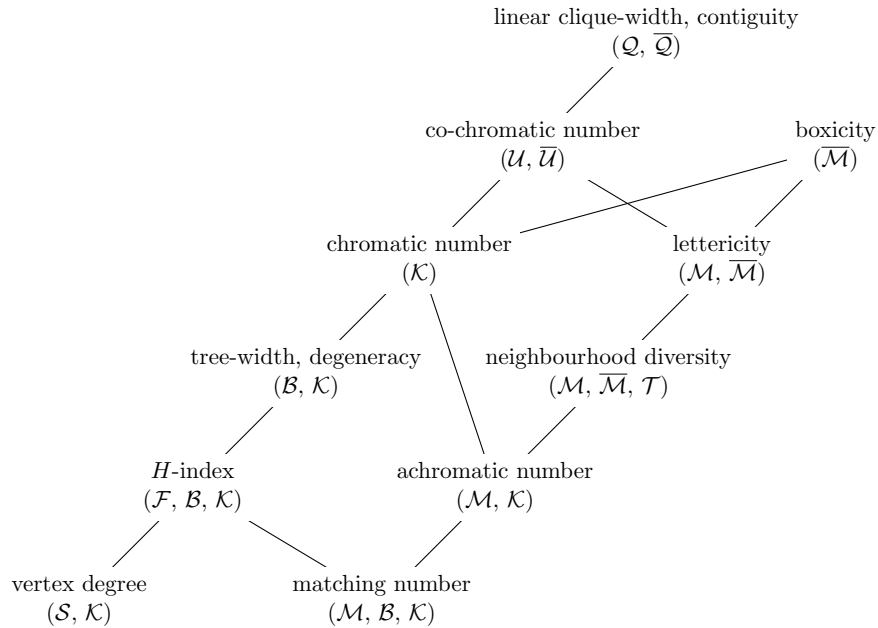


Figure 5.1: A Hasse diagram of graph parameters within the universe of cographs

\mathcal{U} the class of P_3 -free graphs, i.e., graphs every connected component of which is a clique.

\mathcal{K} the class of complete graphs.

\mathcal{F} the class of star forests, i.e., graphs every connected component of which is a star. This is the class of (P_4, C_4, K_3) -free graphs, i.e., the class of bipartite graphs in \mathcal{Q} .

\mathcal{M} the class of graphs of vertex degree at most 1. This is the class of (P_3, K_3) -free graphs, i.e., the class of bipartite graphs in \mathcal{U} .

\mathcal{B} the class of complete bipartite graphs (an edgeless graph is counted as complete bipartite with one part being empty). This is the class of (\overline{P}_3, K_3) -free graphs, i.e., the class of bipartite graphs in $\overline{\mathcal{U}}$.

\mathcal{S} the class of *stars*, i.e., graphs of the form $K_{1,n}$ and their induced subgraphs.

In the rest of this section, we derive, for each parameter, the minimal class characterisations shown in Figure 5.1 (definitions for the parameters, or quick reminders thereof, will be provided as necessary). We start with the parameters for which such a characterisation is easy to obtain from known results, then consider the remaining parameters one by one.

5.2.1 Results that are immediate, or follow quickly from known facts

Directly from Ramsey's Theorem we derive the following conclusion.

Proposition 121. *The class \mathcal{K} of complete graphs and the class of \mathcal{S} of stars are the only two minimal hereditary classes of graphs of unbounded maximum vertex degree.*

To report more results, we denote by

$\alpha(G)$ the *independence number* of G , i.e., the size of a maximum independent set in G ,

$\omega(G)$ the *clique number* of G , i.e., the size of a maximum clique in G ,

$\chi(G)$ the *chromatic number* of G , i.e., the minimum number of subsets in a partition of $V(G)$ such that each subset is an independent set,

$y(G)$ the *clique partition* (also known as *clique cover*) *number*, i.e., the minimum number of subsets in a partition of $V(G)$ such that each subset is a clique.

Clearly, the class \mathcal{K} of complete graphs is the only minimal hereditary class of unbounded clique number, i.e., by forbidding a complete graph we obtain a class of bounded clique number. Additionally, it is not difficult to see that \mathcal{K} is a minimal hereditary class of unbounded chromatic number. However, it is in general not the only minimal hereditary class of unbounded chromatic number. In other words, forbidding a complete graph does not guarantee a bound on the chromatic number. Moreover, as shown by Erdős [Erd59], chromatic number is unbounded even in the class of (C_3, C_4, \dots, C_k) -free graphs for any value of k , which means that in the universe of all hereditary classes, chromatic number cannot be characterized by means of minimal classes where this parameter is unbounded. On the other hand, when we restrict ourselves to cographs such a characterisation is possible by the discussion from the previous section. To obtain it, we note that cographs are perfect, and hence $\omega(G) = \chi(G)$ for any cograph G . As a result, we obtain:

Proposition 122. *The class \mathcal{K} of complete graphs is the only minimal hereditary subclass of cographs of unbounded clique number and chromatic number.*

The *degeneracy* of a graph G is the smallest value of k such that every induced subgraph of G has a vertex of degree at most k . It is known (and easily seen; see, e.g., [KBH01]) that tree-width is bounded below by degeneracy. Moreover, it is not

difficult to show that the class \mathcal{K} of complete graphs and the class of \mathcal{B} of complete bipartite graphs are minimal hereditary classes of unbounded degeneracy. Similarly to chromatic number, in the universe of all hereditary classes, neither degeneracy nor tree-width admit a characterization in terms of minimal classes where these parameters are unbounded. On the other hand, we claim that in the universe of cographs, those two minimal classes are the only ones:

Proposition 123. *The class \mathcal{K} of complete graphs and the class of \mathcal{B} of complete bipartite graphs are the only two minimal hereditary subclasses of cographs of unbounded degeneracy and tree-width.*

Proof. To prove the claim, it suffices to show that for any s and p , the tree-width of $(P_4, K_s, K_{p,p})$ -free graphs is bounded by a constant. For this, we refer the reader to the following result from [ALR12]: for every t, p, s , there exists a $z = z(t, p, s)$ such that every graph with a (not necessarily induced) path of length at least z contains either an induced P_t or an induced $K_{p,p}$ or a clique of size s . From this result it follows that $(P_4, K_s, K_{p,p})$ -free graphs do not contain (not necessarily induced) paths of length $z(4, p, s)$. It is well known (see, e.g., [FL89]) that graphs of bounded path number (the length of a longest path) have bounded tree-width. \square

The *matching number* of a graph G is the size of a maximum matching in G . The following result was proved in [DDL13].

Lemma 124. *For any natural numbers s, t and p , there is a number $N(s, t, p)$ such that every graph with a matching of size at least $N(s, t, p)$ contains either a clique K_s or an induced bi-clique $K_{t,t}$ or an induced matching pK_2 .*

A natural corollary from this result is the following characterization of the matching number in terms of minimal hereditary classes where this parameter is unbounded.

Theorem 125. *\mathcal{M}, \mathcal{B} and \mathcal{K} are the only three minimal hereditary classes of graphs of unbounded matching number.*

The *vertex cover number* of a graph G is the size of a minimum vertex cover in G . It is well known that the vertex cover number is never smaller than the matching number and never larger than twice the matching number. Therefore, the characterization of matching number given in Theorem 125 applies to the vertex cover number as well.

Theorem 126. *\mathcal{M}, \mathcal{B} and \mathcal{K} are the only three minimal hereditary classes of graphs of unbounded vertex cover number.*

The *neighbourhood diversity* of a graph was introduced in [Lam12] and can be defined as follows.

Definition 127. Let us say that two vertices x and y are *similar* if there is no vertex z distinguishing them (i.e., if there is no vertex z adjacent to exactly one of x and y). Vertex similarity is an equivalence relation. We denote by $nd(G)$ the number of similarity classes in G and call it the *neighbourhood diversity* of G .

Neighbourhood diversity was characterized in [Loz18] by means of nine minimal hereditary classes of graphs where this parameter is unbounded. Six of these minimal classes contain a P_4 . Therefore, when restricted to cographs, neighbourhood diversity can be characterized by three minimal classes as follows.

Theorem 128. \mathcal{M} , $\overline{\mathcal{M}}$, and \mathcal{T} are the only three minimal hereditary subclasses of cographs of unbounded neighbourhood diversity.

5.2.2 Co-chromatic number

The *co-chromatic number* of G , denoted $z(G)$, is the minimum number of subsets in a partition of $V(G)$ such that each subset is either a clique or an independent set [EGS90]. It is not difficult to see that the co-chromatic number can be arbitrarily large in the class of P_3 -free graphs, where each graph is a disjoint union of cliques. Therefore, it is also unbounded in the complements of P_3 -free graphs, also known as complete multipartite graphs. In what follows, we show that these are the only two minimal subclasses of cographs of unbounded co-chromatic number.

Lemma 129. Let n, m, t be positive integers with $t \geq 2$. If G is a $(nK_t, \overline{mK_t})$ -free cograph, then $z(G) \leq 2^{m+n-1}(t-1)$.

Proof. Call a partition of $V(G)$ *good* if it contains at least $t-1$ cliques and $t-1$ independent sets (empty sets in the partition may count as either). We prove by induction on $m+n$ that G admits a good partition into $2^{m+n-1}(t-1)$ sets, each of which is a clique or an independent set.

If $m+n=2$ ($n=m=1$), then G is K_t -free. Hence $\chi(G) = \omega(G) \leq t-1$; we add empty sets to the partition until we reach $2(t-1)$ sets in total. This makes the partition good, and we have proved the basis for the induction. In general, put $G' := G$. We are in one of the following three cases:

- (a) $G' = G_1 + G_2$, and both G_1 and G_2 are K_t -free, OR $G' = G_1 \times G_2$, and both G_1 and G_2 are $\overline{K_t}$ -free.

- (b) $G' = G_1 + G_2$, and both G_1 and G_2 contain a K_t , OR $G' = G_1 \times G_2$, and both G_1 and G_2 contain a $\overline{K_t}$.
- (c) $G' = G_1 + G_2$, G_1 contains a K_t and G_2 is K_t -free, OR $G' = G_1 \times G_2$, G_1 contains a $\overline{K_t}$ and G_2 is $\overline{K_t}$ -free.

As long as we are in case (c), iteratively put $G' := G_1$. We end up with a graph G' in either case (a) or (b). Note first that any good partition of G' extends to a good partition of G without increasing the number of sets. Indeed, at each step, G_2 was either K_t -free and anticomplete to the rest of the graph or $\overline{K_t}$ -free and complete to the rest of the graph. The disjoint union of all K_t -free G_2 s is again K_t -free and hence can be partitioned into at most $t - 1$ independent sets, and we take the union of each of these sets with one of the independent sets in the good partition of G' injectively. Similarly, the join of the $\overline{K_t}$ -free G_2 s can be partitioned into at most $t - 1$ cliques, each of which we join to one of the cliques in the good partition of G' injectively.

Now, if G' is in case (a), then G' is K_t -free or $\overline{K_t}$ -free and we act like in the base case to obtain a good partition of G' (and therefore of G) in $2(t - 1)$ sets. If G' is in case (c), then G_1 and G_2 are both either $(n - 1)K_t$ -free or $(m - 1)\overline{K_t}$ -free. In either case, the inductive hypothesis applies, and we have a good partition of G' of size at most

$$2^{m+n-2}(t - 1) + 2^{m+n-2}(t - 1) = 2^{m+n-1}(t - 1).$$

Like before, this extends to a partition of G , concluding the proof. \square

Lemma 129 naturally leads to the following conclusion.

Theorem 130. *The class \mathcal{U} of P_3 -free graphs and the class $\overline{\mathcal{U}}$ of $\overline{P_3}$ -free graphs are the only two minimal hereditary subclasses of cographs of unbounded co-chromatic number.*

5.2.3 Lettericity

A definition of lettericity can be found in Subsection 3.1.1. As a quick and relevant refresher, we give the following example (compare with Example 5). Consider the alphabet $\Sigma = \{a, b\}$ and the decoder $\mathcal{D} = \{(a, a), (a, b)\}$. Then the word $ababababab$ describes the graph represented in Figure 5.2. This graph can be constructed from a single vertex by means of two operations: adding a dominating vertex (corresponds to adding letter a as a prefix) or adding an isolated vertex (corresponds to adding

letter b as a prefix). The class of all graphs that can be constructed by means of these two operations coincides with the class \mathcal{T} of threshold graphs defined at the beginning of this section as $(2K_2, C_4, P_4)$ -free graphs [MP95]. The above discussion shows that a graph is threshold if and only if it is a letter graph over the alphabet $\Sigma = \{a, b\}$ with the decoder $\mathcal{D} = \{(a, a), (a, b)\}$.




Figure 5.2: The letter graph of the word $ababababab$ (the oval represents a clique). We use indices to indicate in which order the a -letters and the b -letters appear in the word.

Now recall (from Example 16) that the class \mathcal{M} of induced matchings has unbounded lettericity; since $\text{let}(G) = \text{let}(\overline{G})$ for any graph G , the same is true for the class $\overline{\mathcal{M}}$.

Theorem 131. \mathcal{M} and $\overline{\mathcal{M}}$ are the only two minimal hereditary subclasses of cographs of unbounded lettericity.

Proof. To prove the theorem, we will show that for any natural numbers $p, t \geq 2$, the lettericity of a $(P_4, pK_2, \overline{tK_2})$ -free graph G is at most 2^{p+t-3} . This will be shown by induction on $p + t$. Moreover, we will show that G can be represented with a decoder \mathcal{D} containing a source letter, i.e., a letter a such that $(a, b) \in \mathcal{D}$ for any letter b , and a sink letter, i.e., a letter b such that $(b, a) \notin \mathcal{D}$ for any letter a .

If $p = t = 2$, then G is a threshold graph and its lettericity is at most 2, because any threshold graph can be represented over the decoder $\mathcal{D} = \{(a, a), (a, b)\}$. In this decoder, a is a source letter and b is a sink letter.

Assume that every $(P_4, pK_2, \overline{tK_2})$ -free graph with $p+t \leq k$ can be represented as a letter graph over an alphabet of at most 2^{p+t-3} letters with a decoder containing a source vertex a and a sink vertex b . Consider now a $(P_4, pK_2, \overline{tK_2})$ -free graph G with $p + t = k + 1$.

The presence of source and sink letters in the decoder allows us to assume that G has neither dominating nor isolated vertices. Indeed, if v is dominating, then a word for G can be constructed from a word for $G - v$ by adding a source letter as a prefix, and if v is isolated, then a word for G can be constructed from a word for $G - v$ by adding a sink letter as a prefix. Therefore, in the rest of the proof we assume that G has neither isolated nor dominating vertices.

Case 1: G is disconnected. Denote by G_1 a connected component of G and by G_2 the rest of the graph. Observe that each of G_1 and G_2 contains a K_2 , since otherwise G has an isolated vertex. Therefore, each of G_1 and G_2 is $(p-1)K_2$ -free and hence we can apply induction to each of G_1 and G_2 . In other words, G_1 can be represented by a word w_1 over an alphabet Σ_1 of size at most 2^{p+t-4} with a decoder \mathcal{D}_1 containing a source vertex a_1 and a sink vertex b_1 , and G_2 can be represented by a word w_2 over an alphabet Σ_2 of size at most 2^{p+t-4} with a decoder \mathcal{D}_2 containing a source vertex a_2 and a sink vertex b_2 (we assume that A_1 and A_2 are disjoint). Then the word $w = w_1w_2$ represents G over the alphabet $\Sigma_1 \cup \Sigma_2$ of size at most 2^{p+t-3} with the decoder $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2$. In this decoder, b_2 is a sink letter. To guarantee the presence of a source letter, we add to \mathcal{D} the pair (a_2, c) for every $c \in \Sigma_1$. This extension transforms a_2 into a source letter and does not change the graph represented by the word w , since every letter from Σ_1 appears in w before any appearance of a_2 .

Case 2: G is connected. In this case, \overline{G} is disconnected and $(P_4, tK_2, \overline{pK_2})$ -free. A similar argument as above gives a representation for \overline{G} with at most 2^{p+t-3} letters, and complementing the corresponding decoder produces one for G (note that when doing that, sink letters become source letters and vice-versa).

□

5.2.4 Boxicity

The *boxicity* $\text{box}(G)$ of a graph G is the minimum dimension in which G can be represented as an intersection graph of hyper-rectangles. Equivalently, it is the smallest number of interval graphs on the same set of vertices whose intersection is G . The next lemma was shown in [Rob69]; we give here a proof for the sake of completeness.

Lemma 132. $\text{box}(\overline{nK_2}) = n$.

Proof. To see that $\text{box}(\overline{nK_2}) \leq n$, note that K_{2n} without an edge is an interval graph, and $\overline{nK_2}$ is the intersection of n such graphs. Conversely, note that two different matched non-edges in $\overline{nK_2}$ cannot belong to the same interval graph (since the corresponding four vertices would induce a C_4 , which is not an interval graph). Hence we need at least n interval graphs to obtain $\overline{nK_2}$ as an intersection. □

Lemma 133. *Let G_1 and G_2 be two graphs. Then*

$$\text{box}(G_1 + G_2) \leq \max(\text{box}(G_1), \text{box}(G_2)) \text{ and } \text{box}(G_1 \times G_2) \leq \text{box}(G_1) + \text{box}(G_2).$$

Moreover, if G_2 is a clique, then $\text{box}(G_1 \times G_2) = \text{box}(G_1)$.

Proof. Suppose $G_1 = \bigcap_{i=1}^s A_i$ and $G_2 = \bigcap_{i=1}^t B_i$ where the A_i and B_i are interval graphs, and assume without loss of generality that $s \geq t$. Put $C_i = A_i + B_i$ for $1 \leq i \leq t$ and $C_i = A_i + K_{|G_2|}$ for $t < i \leq s$. Put $D_i = A_i \times K_{|G_2|}$ for $1 \leq i \leq s$ and $D_i = K_{|G_1|} \times B_{i-s}$ for $s < i \leq s+t$.

The C_i and D_i are interval graphs, and with the obvious labellings of C_i and D_i , we have $G_1 + G_2 = \bigcap_{i=1}^s C_i$ and $G_1 \times G_2 = \bigcap_{i=1}^{s+t} D_i$.

For the final claim, if $G_2 = K_{|G_2|}$ is a clique, then $G_1 \times G_2 = \bigcap_{i=1}^s (A_i \times K_{|G_2|})$, and each of those is an interval graph. \square

Theorem 134. $\overline{\mathcal{M}}$ is the only minimal hereditary subclass of cographs of unbounded boxicity.

Proof. Let $n \geq 2$. We prove by induction on n that $(P_4, \overline{nK_2})$ -free graphs have boxicity at most 2^{n-2} . The result is true for $n = 2$, since (P_4, C_4) -free graphs are known to be interval graphs (see, e.g., [BLS99]).

For the induction step, suppose the result is true for some $n \geq 2$, and let G be a cograph that is $\overline{(n+1)K_2}$ -free. By Lemma 133, we may assume that G is connected, and in particular that $G = G_1 \times G_2$ where neither of the cographs G_1 or G_2 is a clique. But then G_1 and G_2 each have a $\overline{K_2}$, and so they are both $\overline{nK_2}$ -free. The induction hypothesis applies, and another application of Lemma 133 gives us that $\text{box}(G) \leq \text{box}(G_1) + \text{box}(G_2) \leq 2^{n-2} + 2^{n-2} = 2^{n-1}$ as required. \square

5.2.5 H -index

The H -index $h(G)$ of a graph G is the largest $k \geq 0$ such that G has k vertices of degree at least k . This parameter is important in the study of dynamic algorithms [ES12]. Clearly, H -index is unbounded for cographs, since it is unbounded for complete graphs. To characterize this parameter in terms of minimal subclasses of cographs with unbounded H -index, we start with an auxiliary lemma.

Lemma 135. Let G_1, \dots, G_t be graphs. Then

$$h\left(\sum_{i=1}^t G_i\right) \leq \sum_{i=1}^t h(G_i), \text{ and } h(G_1 \times G_2) \leq \min(h(G_1) + |V(G_2)|, h(G_2) + |V(G_1)|).$$

Proof. For the first bound, note that for any j , $1 + \sum_i h(G_i) > h(G_j)$. In particular, by definition of the H -index, each G_j has at most $h(G_j)$ vertices of degree $1 +$

$\sum_i h(G_i)$ or more, and so $\sum_j G_j$ has at most $\sum_j h(G_j)$ vertices of degree at least $1 + \sum_i h(G_i)$, from which the claim follows.

For the other bound, note that $G_1 \times G_2$ has at most $|V(G_2)|$ vertices of degree at least $h(G_1) + |V(G_2)| + 1$ coming from G_2 , and at most $h(G_1)$ coming from G_1 , since⁵ $\deg_{G_1 \times G_2}(v) = \deg_{G_1}(v) + |V(G_2)|$ for any $v \in G_1$, and G_1 does not have more than $h(G_1)$ vertices of degree $h(G_1) + 1$. By definition of the H -index, we obtain that $h(G_1 \times G_2) \leq h(G_1) + |V(G_2)|$, and the claim follows by symmetry. \square

Theorem 136. \mathcal{K} , \mathcal{B} and the class \mathcal{F} of star forests are the only minimal hereditary subclasses of cographs of unbounded H -index.

Proof. One can check that those are, indeed, minimal hereditary classes of unbounded H -index. To see they are the only ones, let $p, q, r, s \geq 1$. We will show by induction on $p + r$ that if G avoids K_p , $K_{q,q}$ and $rK_{1,s}$, then the H -index of G is bounded by a constant $H(p, q, r, s)$. For the base case, note that if $p = 1$, this is trivial, and if $r = 1$, then G is $(K_p, K_{1,s})$ -free and therefore the maximum vertex degree in G is bounded by $R(p, s)$. This in turn implies that $h(G) \leq R(p, s)$. We may thus assume $p, r \geq 2$.

If $G = G_1 \times G_2$ is a join of non-empty graphs, then not both G_1 and G_2 have more than $R(p, q)$ vertices. Indeed, if both do, then either one of them contains a clique of size p , which is forbidden, or they both have independent sets of size q , which again cannot happen since $K_{q,q}$ is forbidden. Without loss of generality, we may assume that $|V(G_2)| \leq R(p, q)$. In this case, by Lemma 135, $h(G) \leq h(G_1) + R(p, q)$. Since $|V(G_2)| \geq 1$, G_1 is K_{p-1} -free, so by the induction hypothesis, $h(G_1)$ is bounded by $H(p-1, q, r, s)$.

If $G = \sum_{i=1}^t G_i$ is a union of connected graphs, we may write $G = G_1 + \dots + G_l + G'$, where G_1, \dots, G_l each have a $K_{1,s}$, and G' is $K_{1,s}$ -free (we may have $l = 0$). Since K_p and $K_{1,s}$ are forbidden for G' , the maximum vertex degree, and hence the H -index of G' , is bounded by $R(p, s)$. Moreover, if $l \geq 2$ and so two of the components of G do have a $K_{1,s}$, then we may write G as the union of two graphs that are $(r-1)K_{1,s}$ -free, and by Lemma 135, $h(G) \leq 2H(p, q, r-1, s)$. Finally, if only one component has a $K_{1,s}$, then that component is a join of non-empty graphs and we obtain, again by Lemma 135 and from the previous paragraph, $h(G) \leq H(p-1, q, r, s) + R(p, q) + R(p, s)$.

⁵When a vertex v appears in more than one graph, we write $\deg_G(v)$ for the degree of v in graph G .

Combining the above, we obtain

$$H(p, q, r, s) \leq \max(H(p-1, q, r, s) + R(p, q) + R(p, s), 2H(p, q, r-1, s)).$$

□

5.2.6 Achromatic number

A *complete k -colouring* is a partition of G into k independent sets (the “colour classes”) such that any two independent sets in the partition have at least one edge between them. The *achromatic number* $\psi(G)$ of a graph G is the maximum number k such that G admits a complete k -colouring. Computing this parameter is a difficult task even for cographs and interval graphs [Bod89].

Note that the class \mathcal{K} of complete graphs and the class \mathcal{M} of matchings have unbounded achromatic number. Indeed, this is clear for complete graphs, and we note that $\binom{n}{2}K_2$ admits a complete n -colouring where each edge of the matching joins two of the colour classes. We claim that among cographs, those are the only minimal classes of unbounded achromatic number. To show this, we start with a short lemma.

Lemma 137. *Let $r, s \in \mathbb{N}$. The class of (K_r, sK_2, P_4) -free graphs has bounded neighbourhood diversity.*

Proof. From Theorem 128, the only minimal subclasses of cographs where neighbourhood diversity is unbounded are \mathcal{M} , $\overline{\mathcal{M}}$ and \mathcal{T} . K_r belongs to both $\overline{\mathcal{M}}$ and \mathcal{T} , while sK_2 belongs to \mathcal{M} . □

We are now ready to prove the main result of this subsection.

Theorem 138. *\mathcal{K} and \mathcal{M} are the only minimal hereditary subclasses of cographs of unbounded achromatic number.*

Proof. It suffices to show that for any $r, s \in \mathbb{N}$, the class of (K_r, sK_2, P_4) -free graphs has bounded achromatic number. Let G be a graph in this class. By Lemma 137, the class has bounded neighbourhood diversity. In other words, there is a constant k (independent of G) such that the vertex set of G can be partitioned into k similarity classes, each similarity class being a clique or an independent set. Moreover, since the size of cliques is bounded by r , we may further assume that each of these similarity classes is an independent set. Let G' be the quotient of G by this partition, i.e., the graph whose vertices are the independent sets, with two vertices being adjacent if and only if the corresponding sets are complete to each other.

Now consider a t -colouring of G , and interpret the colours as vertices of the complete graph K_t . From each edge e of G' , we obtain a complete bipartite subgraph of K_t as follows: if the edge e in G' joins independent sets A_1 and A_2 , then the two sets are complete to each other, so the sets of colours $I_1, I_2 \subseteq V(K_t)$ appearing in A_1 and A_2 respectively are disjoint. The complete bipartite graph B^e corresponding to e has I_1 and I_2 as its parts. With this set-up, the t -colouring is complete if and only if the edges of the graphs $B^e_{e \in E(G')}$ cover the edges of K_t . From [FH96], we need at least $\lceil \log_2(t) \rceil$ complete bipartite graphs to cover K_t . It follows that $t \leq 2^{|E(G')|} \leq 2^{\binom{k}{2}}$, as required. □

5.2.7 Contiguity

The notion of contiguity was introduced in [Gol+95] and was motivated by the need of compact representations of graphs in computer memory. One approach to achieving this goal is finding a linear order of the vertices in which the neighbourhood of each vertex forms an interval. Not every graph admits such an ordering, in which case one can relax this requirement by looking for an ordering in which the neighbourhood of each vertex can be split into at most k intervals. The minimum value of k which allows a graph G to be represented in this way is the *contiguity* of G , denoted $\text{cont}(G)$.

In [CG14], it was shown that the maximum contiguity of n -vertex cographs is $\Theta(\log n)$, implying that this parameter is unbounded in the class of cographs. In what follows, we identify two minimal hereditary subclasses of cographs of unbounded contiguity.

Lemma 139. *Contiguity is unbounded in the class \mathcal{Q} of (P_4, C_4) -free graphs and in the class of their complements.*

Proof. Let G be a graph and v a vertex of G . In a linear order of $V(G)$, the number of intervals representing the neighbourhood of v differs from the number of intervals representing the non-neighbourhood of v by at most 1. Therefore, the contiguity is bounded in a class X of graphs if and only if it is bounded in the class of complements of graphs in X . Thus, it suffices to prove the lemma only for (P_4, C_4) -free graphs (quasi-threshold graphs).

Every quasi-threshold graph can be recursively constructed from one-vertex graphs by applying one of the following two operations [YCC96]: disjoint union of two quasi-threshold graphs G and H , denoted $G + H$, and addition of a dominating vertex v to a quasi-threshold graph G , denoted $v \times G$.

Let G be a quasi-threshold graph of contiguity k . In particular, for any linear order L of $V(G)$, there exists a vertex u whose neighbourhood consists of at least k intervals in L . To prove the lemma, we will show that the contiguity of the graph $H = v \times (G + G + G)$ is strictly greater than k .

Let L be an arbitrary linear order of $V(H)$, and consider the order L^v that we obtain by restricting L to $H - v$, as well as the orders L_1, L_2 and L_3 that we obtain by further restricting L^v to the vertices of each of the three copies of G . Find vertices $u_1, u_2, u_3 \in V(H)$ belonging to each of the copies of G such that in its respective copy, the neighbourhood of u_i consists of at least k intervals in L_i . Since L_i is a restriction of L^v , the neighbourhood of u_i in $H - v$ still consists of at least k intervals in L^v (the number of intervals cannot increase when removing vertices).

Now, the neighbourhood of u_i in H consists of those at least k intervals in L^v , together with v . Note that v can only be adjacent to (or inside) at most one of these intervals. Moreover, since the u_i have disjoint neighbourhoods in $H - v$, v cannot be adjacent to intervals coming from all three neighbourhoods. In other words, there is an $i \in \{1, 2, 3\}$ such that u_i has a neighbourhood consisting of at least $k + 1$ intervals in L (one of which consists only of v). Since L was arbitrary, this shows the contiguity of H is at least $k + 1$, as required. □

Lemma 140. *For any pair of graphs $H \in \text{Free}(P_4, C_4)$ and $K \in \text{Free}(P_4, 2K_2)$, there is a constant $c(H, K)$ such that the contiguity of (P_4, H, K) -free graphs is at most $c(H, K)$.*

Proof. We prove the lemma by induction on $|V(H)| + |V(K)|$. For the basis of the induction we observe that if one of H and K consists of two vertices, then the statement is obvious.

Now assume that both H and K contain more than two vertices and let G be a (P_4, H, K) -free graph. Below we analyse various cases depending on the structure of H and K . Our analysis is based on the following observations (the first one can be derived by restricting orders like in the previous lemma, and the second immediately follows by a double complementation argument):

- (a) if G is disconnected and G_1, \dots, G_p are the components of G , then $\text{cont}(G) = \max_i \text{cont}(G_i)$;
- (b) if G is connected and G_1, \dots, G_p are the co-components (components of the complement) of G , then $\text{cont}(G) \leq \max_i \text{cont}(G_i) + 2$.

Assume first that H contains a dominating vertex v and let $H' = H - v$. By the inductive assumption, there is a constant $c(H', K)$ bounding the contiguity of (P_4, H', K) -free graphs. If G is connected, then each co-component of G is H' -free and hence by (b), $\text{cont}(G) \leq c(H', K) + 2$. If G is disconnected, then as in the previous sentence, the contiguity of each component of G is at most $c(H', K) + 2$ and hence by (a), the contiguity of G is at most $c(H', K) + 2$.

If K contains an isolated vertex, then the arguments are similar. Therefore, in the rest of the proof we assume that H is disconnected and K is the complement of a disconnected graph. We represent H as $H' + H''$, where H' is a component of H and H'' is the rest of the graph. Similarly, we represent $K = K' \times K''$, where K' is a co-component of K and K'' is the rest of the graph.

Assume without loss of generality that G is disconnected. If each of the components of $G'_0 := G$ is H' -free, then by the inductive assumption the contiguity of each component, and hence of G'_0 , is at most $c(H', K)$. Suppose now that one of the components of G'_0 contains H' as an induced subgraph. Denote that component by G'_1 , and the rest of the graph by G_1 . Note that each of the components of G_1 is H'' -free, and hence, by (a), G_1 has contiguity at most $c(H'', K)$. Applying similar arguments to G'_1 , we see that either all of its co-components are K' -free, or it can be expressed as the join of two graphs G'_2 and G_2 such that G'_2 is disconnected and contains K' as an induced subgraph, and G_2 has contiguity bounded by a constant depending on H and one of K', K'' .

Continue in this way for as long as possible. We produce two sequences G_i and G'_i such that $G'_i = G'_{i+1} \star G_{i+1}$, where \star stands for $+$ when i is even and \times when i is odd, G'_i is connected and contains H' when i is odd/disconnected and contains K' when i is even, and all G_i have contiguity uniformly bounded by some constant depending only on H and K . Since $|G'_i|$ strictly decreases as i increases, there exists a k such that every component or co-component of G'_k (according to whether k is even or odd respectively) is H' , respectively K' -free. Put $G_{k+1} := G'_k$.

Assuming without loss of generality that k is even, we have, by construction, that $G = G_1 + (G_2 \times (G_3 + \dots (G_k \times G_{k+1})))$, and each G_i has contiguity bounded by, e.g., $c'(H, K) := \max(c(H, K'), c(H, K''), c(H', K), c(H'', K)) + 2$.

Let $L_i, 1 \leq i \leq k+1$, be a linear order on the vertices of G_i that witnesses a contiguity of at most $c'(H, K)$, and consider the linear order on $V(G)$ given by the concatenation $L := L_1 L_3 \dots L_{k+1} L_k \dots L_4 L_2$. We claim that this order witnesses a contiguity of at most $c'(H, K) + 2$ for G . Indeed, the neighbourhood in G of any vertex $v \in G_i$ consists of its neighbours in G_i , together with some of the G_j , as follows:

- If i is even, the neighbourhood outside of G_i of v consists of $\bigcup_{j>i} V(G_j) \cup \bigcup_{\substack{j<i \\ j \text{ even}}} V(G_j)$.
- If i is odd, the neighbourhood outside of G_i of v consists of $\bigcup_{\substack{j<i \\ j \text{ even}}} V(G_j)$.

Note that each of the indexed unions above corresponds to an interval in L . Thus the neighbourhood of v consists of at most $c(H, K) := c'(H, K) + 2$ intervals in L , as required. □

Combining the two lemmas above we obtain the main result of this subsection as follows.

Theorem 141. *The class \mathcal{Q} of quasi-threshold graphs and the class of their complements are the only two minimal hereditary subclasses of cographs of unbounded contiguity.*

5.3 Conclusions and further directions

We first point out that, despite its relatively narrow scope (namely its restriction to the class of cographs), this study provides us with valuable insight into the general behaviour of various parameters. For instance, in [Ale+21a], we show that H -index is characterised by the same minimal classes *in the universe of all graphs*. Such a result might have been deemed too ambitious to approach, had we not known what happens in the cograph case. It would be interesting to look further into this: is it possible to formulate general conditions on a parameter which guarantee that its boundedness can be understood by only looking a few well-behaved classes (such as cographs)? What would those classes be?

Our study could also be continued by investigating many other interesting parameters that are unbounded in the class of cographs. Such examples include Dilworth number [Gho19], distinguishing number [AB20], shrub-depth [Gan+19], rank [CHY08], metric dimension [VW19], etc. We remark that not many “interesting” subclasses of cographs occur in our analysis. More concretely, we note that in our hierarchy from Figure 5.1, we have 13 (more or less arbitrarily chosen) different parameters, and the union of their corresponding sets $M(\kappa)$ only contains 11 different classes (8 if we count up to complements). In addition, some of the extra parameters suggested above (namely Dilworth number, distinguishing number and shrub-depth,

but potentially the others as well) can be characterised without extending this list any further. What makes those classes special?

It is not difficult to see that any class \mathcal{X} appearing in some $M(\kappa)$ must be *atomic*, in the sense that it cannot be written as the union of two proper subclasses. This property is equivalent to the *joint embedding property*, whereby if \mathcal{X} contains G and H , then it must contain a graph containing both G and H as induced subgraphs (Fraïssé [Fra54] studied these notions, albeit in a more general setting). Conversely, for any atomic class \mathcal{X} , one can artificially cook up a parameter $\kappa_{\mathcal{X}}$ with $M(\kappa_{\mathcal{X}}) = \{\mathcal{X}\}$ (just let $\kappa_{\mathcal{X}}(G)$ be the largest n such that G contains every n -vertex graph in \mathcal{X} as an induced subgraph). However, even when restricting our search to atomic classes, only a select few seem to occur when studying “natural” parameters. It is difficult, for instance, to imagine a parameter occurring “in the wild” for which the class of star forests of degree at most 7 is a minimal class. Understanding this phenomenon is a challenging research problem: how can we formalise what makes a parameter κ “natural”, and how is that “naturalness” reflected in $M(\kappa)$?

A third research direction concerns algorithmic problems. Just like with boundedness of parameters, the set of classes where a certain problem can be solved efficiently is also downwards closed under inclusion. Hence within the universe of cographs, it can be characterised via (finitely many) minimal classes not belonging to it. Examples of problems that are NP-complete in the whole class include ACHROMATIC NUMBER [Bod89], HARMONIOUS COLOURING [K A07], k -PATH PARTITION [K A07] and INDUCED SUBGRAPH ISOMORPHISM [BHH12]. It is known that each of these problems is still NP-complete in the class of quasi-threshold graphs; what are the minimal classes with this property?

Finally, one more series of questions stems from our observation that cographs are bqo. It appears that bqo properties under the induced subgraph relation have not yet been studied in depth. In particular, as far as the authors are aware, many fundamental questions in this area remain unanswered, the most immediate being: is every wqo class of graphs in fact bqo? We note that this is not the case for quasi-orders in general. For instance, the so-called *Rado structure*⁶ [Rad54] is a wqo, but its power set is not wqo under inclusion. In fact, this structure universal with this property, in the sense that it injects isomorphically into any quasi-order that is ω -good, but not ω^2 -good ([Rad54; Lav76; Jan99; FBS20] show various reformulations

⁶For brevity, we skip the construction; a very transparent and well-illustrated description of the structure can be found in [Jan99]. The point is that we are dealing with a very explicit, relatively simple quasi-order.

and variations on this statement).⁷ This suggests a very concrete and accessible first question in this direction: does there exist a Rado structure of graphs under induced subgraphs? We also note that bqo of graphs under the *minor* relation is already an open problem (see, e.g., [DK05]).

⁷Up to some details that we will ignore, what this means for our purposes is that any class of graphs that is wqo under induced subgraphs but whose set of subclasses is not wqo under inclusion must contain a copy of the Rado structure.

Chapter 6

Linear clique-width of bi-cographs

We now move our magnifying glass over the bipartite analogues of cographs, the so-called bi-cographs, and present our joint work with Mamadou Kanté, Vadim Lozin and Viktor Zamaraev [Ale+20a]. We once more expand on Brignall, Korpelainen and Vatter’s work from [BKV17], but in a different direction. This time, instead of staying in the class of cographs and studying other parameters, we study the same parameter as them, namely linear clique-width, but in the bipartite setting.¹ Moreover, our approach to identifying minimal classes of unbounded linear clique-width is different from theirs; in particular we do not make use of well-quasi-orderability.

6.1 Preliminaries and some basic results

6.1.1 Linear clique-width

Linear clique-width is the “linearised version” of clique-width: it is what we obtain if we require our clique-width expressions to be “path-like”, instead of general trees. More explicitly, it is like clique-width, except we do not allow disjoint unions – instead, we have to introduce the vertices one by one:

Definition 142. The *linear clique-width* of a graph G , denoted by $\text{lcw}(G)$, is the smallest number of labels needed to construct G by means of the following three operations:

¹We will identify two minimal classes of unbounded linear clique-width among bi-cographs. We remark that in [Ale+20a], we go a bit further and identify two additional boundary classes for linear clique-width among all bipartite graphs, and conjecture that the list (which consists of four classes) is complete. However, for the sake of brevity, we do not do that here.

- add a new vertex with label i (we denote this operation simply by i),
- add all edges between vertices labelled i and vertices labelled k , for $i \neq k$ (denoted by $i \times k$),
- relabel vertices labelled i to k (denoted by $i \rightarrow k$).

A *linear clique-width expression* A for a graph G is an ordered sequence of these three operations that constructs G .

Example 143. The following sequence constructs a path P_k with three different labels, showing that $\text{lcw}(P_k) \leq 3$ for any value of k :

$$1, 2, 1 \times 2, (3, 2 \times 3, 2 \rightarrow 1, 3 \rightarrow 2)^{k-2}. \quad (6.1)$$

This example can be extended to $S_{2,2,2}$ -free trees, that is, caterpillars, without increasing the number of labels.

We will be working with coloured bipartite graphs. The linear clique-width of a *coloured* bipartite graph is simply the linear clique-width of the underlying uncoloured graph. In this setting, it will be helpful to consider an auxiliary parameter that we will call *bipartite linear clique-width*.

Definition 144. The *bipartite linear clique-width* (or bi-linear clique-width, for short) of a coloured bipartite graph G , denoted by $\text{blcw}(G)$, is the minimum number of labels necessary to construct G (as an uncoloured graph) via a linear clique-width expression, but only allowing any given label to be used for either black or white vertices (we will call those labels black or white respectively).

The following is clear from the definition because it suffices to copy each label, and to reserve one copy for black, and the other for white vertices.

Remark 145. Let G' be a colouring of a bipartite graph G . Then, $\text{lcw}(G) \leq \text{blcw}(G') \leq 2 \cdot \text{lcw}(G)$.

It follows that, in a class of coloured bipartite graphs, linear clique-width is bounded if and only if bi-linear clique-width is bounded. The benefit of this auxiliary parameter is that it is well-behaved when taking bipartite complements, unions or joins, as we show in the next two lemmas:

Lemma 146. *If G is a coloured bipartite graph, then $\text{blcw}(\tilde{G}) \leq \text{blcw}(G) + 1$.*

Proof. Let A' be an expression using $\text{blcw}(G)$ labels such that each label is either black or white. We claim that we can modify A' to find a linear clique-width expression that uses $\text{blcw}(G)+1$ labels, in which vertices are connected to all of their already constructed neighbours immediately as they are inserted. Indeed, say that a new vertex v is inserted in A' with label l . Whether an already constructed vertex w is a neighbour of v only depends on its label, so we can say that the set of already constructed neighbours of v is a union $\bigcup_{k \in \Lambda} \{w : w \text{ has label } k\}$, where Λ is a set of labels. In A' , the label l might already be in use, so if we tried to connect v to all its already constructed neighbours right away, we might inadvertently add some extra edges (that do not appear in G) to the already constructed graph, between vertices labelled l and some other vertices. However, using a new, reserved label to insert v allows us to go around this. We can immediately connect it to all of its neighbours without changing the already constructed graph, and afterwards change the reserved label to the original label used for inserting v in A' . Proceeding inductively allows us to modify A' to an expression giving G with the desired properties.

A linear clique-width expression for \tilde{G} can be obtained from this modified expression by instead connecting newly inserted vertices to their non-neighbours in G of opposite colour that have already been inserted. \square

Lemma 147. *If G_1, \dots, G_r are coloured bipartite graphs with bi-linear clique-widths at most k_1, \dots, k_r respectively, then their disjoint union $\sum_{i=1}^r G_i$ and their join $\prod_{i=1}^r G_i$ have bi-linear clique-width at most $\max\{k_1, k_2 + 2, k_3 + 2, \dots, k_{r-1} + 2, k_r + 2\}$.*

Proof. We will prove the statement for the join $\prod_{i=1}^r G_i$. The case of the union is similar and we omit the details.

First, we construct G_1 using labels $1, 2, \dots, k_1$ in such a way that no vertices of different colour ever receive the same label, then we relabel all black vertices to 1 and all white vertices to 2. Next, we construct G_2 using labels $3, 4, \dots, k_2 + 2$ (which are now unused). To construct the bipartite join, we then connect vertices labelled by 1 to all vertices labelled by white labels, except 2, and we connect vertices labelled 2 to all vertices labelled by black labels, except 1. Finally, we relabel all black vertices to 1 and all white vertices to 2. In this way we construct $G_1 \times G_2$ using at most $\max\{k_1, k_2 + 2\}$ labels where in the end all black vertices are labelled 1 and all white vertices are labelled 2. Proceeding in the same way with G_3, G_4, \dots, G_r we will construct the join $\prod_{i=1}^r G_i$ with at most $\max\{k_1, k_2 + 2, k_3 + 2, \dots, k_{r-1} + 2, k_r + 2\}$ labels. \square

6.1.2 Bi-cographs

The main object of study in this chapter is the class of bi-complement reducible graphs, introduced in [GV97].

Definition 148. A *coloured bi-complement reducible graph* (or *coloured bi-cograph* for short) is a coloured bipartite graph defined recursively as follows:

- (i) A graph on a single black or white vertex is a coloured bi-cograph.
- (ii) If G_1, G_2 are coloured bi-cographs, then so is their disjoint union $G_1 + G_2$.
- (iii) If G is a coloured bi-cograph, then so is its bipartite complement \tilde{G} .

A *bi-cograph* is a bipartite graph obtained from a coloured bi-cograph by forgetting the colouring.

It is not difficult to see that (iii) in the above definition could be replaced by:

- (iii') If G_1, G_2 are coloured bi-cographs, then so is their bipartite join $G_1 \times G_2$.

In [GV97], the authors also provide an induced subgraph characterisation for (uncoloured) bi-cographs:

Proposition 149. A bipartite graph is a bi-cograph if and only if it is $(P_7, S_{1,2,3}, \text{Sun}_4)$ -free.

By analogy with quasi-threshold graphs (see Subsection 5.2.7 for a characterisation/possible definition), we introduce the class of bi-quasi-threshold graphs:

Definition 150. A *coloured bi-quasi-threshold graph* is a coloured bipartite graph defined inductively as follows:

- (i) A graph on a single black or white vertex is a coloured bi-quasi-threshold graph.
- (ii) If G_1, G_2 are coloured bi-quasi-threshold, then so is their disjoint union $G_1 + G_2$.
- (iii) If G is a coloured bi-quasi-threshold graph, then the bipartite join of G with a single black vertex is a coloured bi-quasi-threshold graph.

Remark 151. Note the asymmetry in this definition: we do not allow white dominating vertices while constructing a coloured bi-quasi-threshold graph. However, once we have finished constructing it, we can forget the colouring to obtain an (uncoloured) bi-quasi-threshold graph.

Throughout the remainder of the chapter, we denote by \mathcal{L} the class of all (uncoloured) bi-quasi-threshold graphs, and by $\tilde{\mathcal{L}}$ the class of uncoloured bipartite graphs obtained from bipartite complements of coloured bi-quasi-threshold graphs.

The following lemma provides a characterisation of coloured bi-quasi-threshold graphs, that we then use to obtain a characterisation of bi-quasi-threshold graphs.

Lemma 152. *The following are equivalent for a coloured bipartite graph G :*

- (a) G is a coloured bi-quasi-threshold graph;
- (b) G contains no induced P_5 with white centre;
- (c) any two black vertices of G have either comparable or disjoint neighbourhoods.

Proof. (a) \Rightarrow (b): A P_5 with white centre is not in \mathcal{L} , since it is not a disjoint union, and it does not have a black dominating vertex. Moreover, from the definition, \mathcal{L} is hereditary, hence no graph in \mathcal{L} contains a P_5 with white centre.

(b) \Rightarrow (c): If two black vertices x and y have incomparable and non-disjoint neighbourhoods, then x, y together with a private neighbour of each and with a common neighbour induce a P_5 with white centre.

(c) \Rightarrow (a): We want to show that, assuming (c), either G is disconnected, or it has a black dominating vertex (then use induction, and the fact that the condition (c) is hereditary). Suppose G is connected, and let b be a black vertex with a maximal (under set inclusion) neighbourhood. Let w be a white vertex non-adjacent to b (which exists because we assume that b is not a black dominating vertex), and consider a shortest path P from b to w (which exists, since G is connected). Write its vertices as $b = b_0, w_1, b_1, \dots, b_{k-1}, w_k = w$ (where the vertices w_i are white, and the vertices b_i are black). If $k > 1$, then w_1 is a common neighbour to b and b_1 , hence by (c) and maximality of the neighbourhood of b , $N(b_1) \subseteq N(b)$. In particular, w_2 and b are adjacent, and we have a shorter path between b and w , contradicting the choice of P . This shows $k = 1$, i.e., b and w are in fact adjacent, so b must be a dominating vertex. \square

We can now give a forbidden induced subgraph characterisation of the class \mathcal{L} .

Theorem 153. *A bipartite graph G is bi-quasi-threshold if and only if G is $(P_6, C_6, \text{domino}, \text{Sun}_4)$ -free.*

Proof. The “only if” direction comes from the fact that any colouring of one of the four graphs in black and white contains a P_5 with white centre, hence by the previous lemma, none of the four graphs is bi-quasi-threshold.

Conversely, suppose G is $(P_6, C_6, \text{domino}, \text{Sun}_4)$ -free. We show that there is a colouring of G in black and white such that there is no P_5 with white centre. This is clear if G is P_5 -free, so assume it is not. Without loss of generality, we can assume, in addition, that G is connected. Now find a P_5 induced by a, b, c, d and e such that the neighbourhood of its middle vertex is maximal among all P_5 's. Denote by S the part of G containing a, c and e , and by T the other part.

Let $x \in T$ be a neighbour of a . Then x is not a neighbour of e , otherwise a, b, c, d, e, x induce either a C_6 or a domino (depending on whether c and x are adjacent). Additionally, x must be a neighbour of c , otherwise the six vertices induce a P_6 . With this in mind, let B be the set of neighbours of a and c , let D be the set of neighbours of e and c (in particular, $b \in B$ and $d \in D$), and let N_c be the set of neighbours of c that are not neighbours of a or e .

Suppose now that a vertex y in T is a non-neighbour of c (i.e., $y \notin B \cup N_c \cup D$, and then y is also a non-neighbour of a and e). Find a path from y to c . Such a path must pass through $B \cup N_c \cup D = N(c)$. Let c' be the vertex of the path just before $N(c)$, and assume without loss of generality that y is adjacent to c' . Let $z' \in N(c)$ be a neighbour of c' .

If c' has a non-neighbour b' in B , then a, b, c, z', c', y is a P_6 , contradicting that G is P_6 -free. Symmetrically, if c' has a non-neighbour d' in D , then y, c', z', c, d', e is a P_6 , contradicting again that G is P_6 -free. Therefore, $B \cup D \subseteq N(c')$. We also have that $N_c \subseteq N(c')$, otherwise if $z \in N_c \setminus N(c')$, then $G[\{a, b, c, d, e, z, c', y\}]$ induces a Sun_4 . We can therefore conclude that $N(c) \subset N(c')$, contradicting the choice of the P_5 a, b, c, d, e because a, b, c', d, e is also a P_5 .

We can thus conclude that S has a vertex c dominating T . If there was another P_5 with its centre in T , then that other P_5 cannot contain c and would induce with c a domino. Hence if we colour S black and T white, we obtain a colouring of G with no P_5 's with white centre, and by Lemma 152, we can conclude that G is bi-quasi-threshold. \square

6.2 Unboundedness of linear clique-width

The main result of this section is the unboundedness of the linear clique-width of \mathcal{L} and $\tilde{\mathcal{L}}$. To prove the result we will use an auxiliary graph parameter which bounds linear clique-width below.

Let G be a graph and A a linear clique-width expression for G . Note that A defines a linear order of the vertex set of G , i.e., a permutation π in the symmetric group $\mathcal{S}(V(G))$. Let us denote by $S_{\pi,i}$ the set consisting of the first i elements of the permutation, and by A_i the maximal prefix of A containing only vertices from this set. If two vertices in $S_{\pi,i}$ have different neighbourhoods outside of the set, then they must have different labels in A_i , since otherwise in the rest of the expression we would not be able to add a neighbour to one of them without adding it to the other. Therefore, denoting by $\mu_{\pi,i}(G)$ the size of the set $\{N(x) \cap (V(G) \setminus S_{\pi,i}) \mid x \in S_{\pi,i}\}$, we conclude that A uses at least

$$\mu_{\pi}(G) := \max_i \mu_{\pi,i}(G)$$

different labels to construct G . As a result, the linear clique-width is bounded from below by²

$$\mu(G) := \min_{\pi \in \mathcal{S}(V(G))} \mu_{\pi}(G).$$

Therefore, to prove the main result of the section, it suffices to show that $\mu(G)$ is unbounded in the classes under consideration. In order to do that, we will need a technical lemma describing the behaviour of $\mu(G)$ in some situations.

We introduce some notation for the coming part. Given a graph G and a linear order π of its vertices, we will write $v < w$ if v appears before w in the order, and $v < S$ if v appears before every vertex of a set S . Notice that the order on a graph induces an order on all of its subgraphs in the obvious way.

Every $i \in \{1, \dots, n\}$ corresponds to a *cut in G with respect to π* , which separates the first i vertices in π from the rest of $V(G)$. It will be useful to mark cuts for which $\mu_{\pi,i}(G)$ is large. We will insert symbols α, β, \dots into our ordered list of vertices to mark such cuts. If α marks a cut with $\mu_{\pi,i}(G) \geq t$, then a set of t vertices in $S_{\pi,i}$ with pairwise different neighbourhoods outside of $S_{\pi,i}$ will be called a *diversity witness of size t* for α . The largest t such that there exists a diversity witness of size t for α will be called the *diversity* of (the cut at) α .

²This parameter, or variations thereon, appear throughout the literature in the study of clique-width, NLC-width, and their linear variants; for example, a clique-width version of it was used in [LR07a], while in [HMP12], the parameter is called *groupnumber*. [HMP12] gives further relevant references.

For a coloured bipartite graph G , we define $\mu(G)$ as $\mu(G')$ with G' obtained from G by forgetting the colouring.

Let H be a connected coloured bi-quasi-threshold graph with $\mu(H) = t \geq 2$. Since H is connected and has at least two vertices, it contains both white and black vertices. Let $G = v \times (H + H + H)$ for a black vertex v , and label the vertices of the three copies of H by $A = \{a_i : 1 \leq i \leq n\}$, $B = \{b_i : 1 \leq i \leq n\}$, and $C = \{c_i : 1 \leq i \leq n\}$, respectively.

Lemma 154. $\mu(G) \geq t + 1$.

Proof. To prove the lemma, we fix an arbitrary permutation π of $V(G)$ and show that $\mu_\pi(G) \geq t + 1$. Let α , β , and γ be the three cuts of diversity of at least t in the three copies of H with respect to the restrictions of π into A , B , and C , respectively. Without loss of generality we assume that $\alpha \leq \beta \leq \gamma$ in π . Let $B' \subset B$ be a diversity witness of size t for β in B , i.e., $B' < \beta$, $|B'| = t$, and the vertices of B' have pairwise different neighbourhoods in the subset of B to the right of β .

Assume first that a vertex a of A appears after β . Since $\mu(H) \geq 2$, there exist vertices of A before α (and in particular before β). Therefore, since H is connected, there must be an edge $a_i a_j$ such that $a_i < \beta < a_j$. Since, by the definition of G , none of the vertices in B' is adjacent to a_j , the set $B' \cup \{a_i\}$ is a diversity witness of size $t + 1$ for β , i.e., $\mu_\pi(G) \geq t + 1$. This conclusion allows us to assume, from now on, that

- $A < \beta$ and, by a similar argument, $\beta < C$ (we need $t \geq 2$ to make sure we do indeed have vertices of C after γ , and hence after β).

Suppose $v < \beta$. Since C has at least one white vertex, v has a neighbour in C and hence $B' \cup \{v\}$ is a diversity witness of size $t + 1$ for β , i.e., $\mu_\pi(G) \geq t + 1$. Therefore, in the rest of the proof we assume that

- $v > \beta$.

Assume B' contains a vertex b_i with no neighbour $b_j > \beta$ in B (observe that if such a vertex exists, then it is unique in B' , since otherwise B' is not a diversity witness). If b_i is white, then for any black vertex $a_k \in A$, the set $B' \cup \{a_k\}$ is a diversity witness of size $t + 1$ for β , because b_i is adjacent to v , while a_k is not (by the definition of G), and every vertex of B' different from b_i has a neighbour to the right of β , while a_k does not. Similarly, if b_i is black, then for any white vertex $a_k \in A$, the set $B' \cup \{a_k\}$ is a diversity witness of size $t + 1$ for β in G , because b_i is not adjacent to v , while a_k is, and every vertex of B' different from b_i has a neighbour in B to the right of β , while a_k does not. In both cases, we have $\mu_\pi(G) \geq t + 1$.

The above discussion allows us to assume that every vertex of B' has a neighbour in the subset of B to the right of β . Then for any vertex $a_k \in A$, the set $B' \cup \{a_k\}$ is a diversity witness of size $t + 1$ for β in G , since a_k has no neighbours in B , i.e., $\mu_\pi(G) \geq t + 1$. \square

Theorem 155. *Linear clique-width is unbounded in the classes \mathcal{L} and $\tilde{\mathcal{L}}$.*

Proof. Let $G_2 \simeq P_4$ given with any colouring. It is easy to see that G_2 is a connected coloured bi-quasi-threshold graph with $\mu(G_2) \geq 2$. Defining $G_k = v \times (G_{k-1} + G_{k-1} + G_{k-1})$ for $k > 2$, we conclude by Lemma 154 that G_k is a connected coloured bi-quasi-threshold graph with $\mu(G_k) \geq k$. Therefore, for each k , the class \mathcal{L} contains a graph of linear clique-width at least k . For the class $\tilde{\mathcal{L}}$, a similar conclusion follows from Lemma 146 and Remark 145. \square

6.3 Minimality and uniqueness

The goal of this section is to show that the two classes of unbounded linear clique-width identified in the previous section are *minimal* hereditary classes where this parameter is unbounded. Moreover, we prove a more general result showing that the classes \mathcal{L} and $\tilde{\mathcal{L}}$ are the *only two minimal hereditary classes* of bi-cographs where the linear clique-width is unbounded.³

For a coloured bipartite graph $G = (B, W, E)$, it will sometimes be useful to work with the coloured graph we obtain by swapping the colours. To this end, recall that the *reflection* G^R of G is the coloured bipartite graph (W, B, E) .

In order to prove the main result of the section, we will use the notion of bi-cotrees, the bipartite analogues of cotrees also defined in [GV97]:

Definition 156. Let G be a coloured bi-cograph. The *bi-cotree* T_G of G is the rooted labelled tree constructed as follows:

- Start with the root, which corresponds to G .
- For any internal node, label it by 0 if the corresponding subgraph is disconnected, and by 1 if it is not (in which case, its bipartite complement is disconnected). The children of the node then correspond to connected components, respectively bi-co-components.

³Note the similarity with the results from [BKV17]: among cographs, linear clique-width is characterised by quasi-threshold graphs and their complements. Among bi-cographs, it is characterised by bi-quasi-threshold graphs and their bipartite complements.

- For any leaf, label it by 0 if the corresponding vertex in the graph is white, and by 1 if the corresponding vertex in the graph is black.

Remark 157. The construction given here only defines “the (unique) bi-cotree of a coloured bi-cograph”. As the authors of [GV97] point out, if we instead define bi-cotrees as labelled trees giving sets of instructions for constructing a coloured bi-cograph, then different bi-cotrees might yield the same bi-cograph. This is in contrast to the usual cotrees of cographs, and happens because there exist disconnected coloured bipartite graphs whose bipartite complement is also disconnected.

Note that the bi-cotree T_{G^R} of G^R is obtained from the bi-cotree T_G of G by changing the labels of the leaves from 0 to 1 and from 1 to 0.

Over the next few lemmas, we will be talking about the presence of certain trees in the bi-cotrees of coloured bi-cographs. We will use the following definition of tree containment:

Definition 158. Let S and T be two rooted trees. We say S is *contained* or *appears* in a tree T , if there is an embedding $\phi : S \hookrightarrow T$ with the following properties:

- If $x, y \in V(S)$ and x is an ancestor of y , then $\phi(x)$ is an ancestor of $\phi(y)$.
- If $x, y, z \in V(S)$ and x is the lowest common ancestor of y and z , then $\phi(x)$ is the lowest common ancestor of $\phi(y)$ and $\phi(z)$.
- If S and T are labelled, then ϕ preserves labels.

We say that S is contained in T *internally* if no vertex of S is mapped to a leaf of T .

Our proof strategy for minimality and uniqueness is as follows: we first show that the bi-cotrees of coloured bi-cographs of large bi-linear clique-width must contain large perfect binary trees. We then show that, in particular, certain labelled perfect binary trees must appear in those bi-cotrees in a very specific way. Finally, we show that the latter implies that a family of coloured bi-cographs of unbounded bi-linear clique-width contains either colourings of graphs in \mathcal{L} or in $\tilde{\mathcal{L}}$.

For $h \in \mathbb{N}$, let B_h denote the unlabelled perfect binary tree of height h (i.e., the binary tree where every internal node has two children, and all leaves have the same depth h). Let $B_{h,0}$ and $B_{h,1}$ further denote perfect binary trees of height h , with all vertices labelled 0 or 1 respectively.

Lemma 159. *Let G be a coloured bi-cograph and h be a natural number. If T_G does not contain B_h , then the bi-linear clique-width of G is at most $2h$.*

Proof. We prove the lemma by induction on h . The result holds for $h = 1$, since forbidding B_1 means no node has two children, and G is trivial. Suppose the statement holds for some $h \geq 1$. We will prove that the bi-linear clique-width of any graph G whose bi-cotree T_G does not contain B_{h+1} is at most $2h + 2$. We proceed by induction on the height of the bi-cotree. Clearly, the statement holds for any graph with the bi-cotree of height at most $h - 1$, as in this case the bi-cotree does not contain B_h , and the bi-linear clique-width of the graph is at most $2h \leq 2h + 2$ by the induction hypothesis for h . Assume now that the statement holds of any graph with bi-cotree of height at most $r \geq 0$ and suppose that the height of T_G is $r + 1$. Let x be the root of T_G . Then at most one of the subtrees rooted at the children of x contains B_h , otherwise T_G would contain a B_{h+1} . If none of those subtrees do, we are done, since by the inductive hypothesis for h , the subgraphs corresponding to each child of x have bi-linear clique-width at most $2h$, and their join or disjoint union can be constructed using two additional labels. Otherwise, let b be the bi-linear clique-width of G , let x_1 be the child whose induced subtree contains a B_h , and b_1 the bi-linear clique-width of the graph corresponding to x_1 . Then, by Lemma 147, we have $b \leq \max\{b_1, 2h + 2\}$. Since the height of the tree rooted at x_1 is at most $r - 1$, by the inductive hypothesis for r we have $b_1 \leq 2h + 2$, which implies $b \leq 2h + 2$, and hence the lemma. \square

We next consider $B_{h,0}$ and $B_{h,1}$. We first show that if one of those trees appears in a certain way in the bi-cotree of a coloured bi-cograph, then that coloured bi-cograph contains either all coloured bi-quasi-threshold graphs up to a certain size, or their bipartite complements.

Definition 160. Let G be a coloured bi-cograph, h be a natural number, and $i \in \{0, 1\}$. We say $B_{h,i}$ is *meaningfully embedded* in T_G , if the following hold:

- $B_{h,i}$ is internally contained in T_G with embedding ϕ .
- Let x be a node in $B_{h,i}$ and let y be a child of x . Let P be the path in T_G between $\phi(x)$ and $\phi(y)$. Then there exists a vertex z on P labelled by $1 - i$ such that the subtree rooted at z excluding the branch containing $\phi(y)$ has a leaf in T_G corresponding to a black vertex (i.e., a leaf labelled 1).

Lemma 161. *Let G be a coloured bi-cograph, let h be a natural number, and $i \in \{0, 1\}$. Furthermore, suppose that $B_{h,i}$ is meaningfully embedded in T_G . Then if $i = 0$, G contains all coloured bi-quasi-threshold graphs on at most h vertices as induced subgraphs, and if $i = 1$, G contains bipartite complements of coloured bi-quasi-threshold graphs on at most h vertices as induced subgraphs.*

Proof. We assume that $i = 0$. The case of $i = 1$ is analogous and we omit the details.

We will prove by induction on h that G contains every coloured bi-quasi-threshold graph on at most h vertices as a coloured induced subgraph. If $B_{1,0}$ is meaningfully embedded in T_G , write x for the embedding of the root and y_1, y_2 for the embeddings of its two children. By definition, there is a vertex labelled 1 on the path between x and y_1 , and that vertex is not a leaf. It follows that G has at least one edge, and hence it contains both a black and a white vertex, i.e., it contains every coloured bi-quasi-threshold graph on 1 vertex as a coloured induced subgraph.

Assuming the statement holds for some $h \geq 1$, suppose that $B_{h+1,0}$ is meaningfully embedded in T_G . Like before, write x for the embedding of the root, and write y_1, y_2 for the embeddings of its two children. Each of the subtrees of T_G rooted at y_1 and y_2 have a meaningfully embedded $B_{h,0}$, so the corresponding induced subgraphs of G contain all coloured bi-quasi-threshold graphs on h vertices. Since x is labelled 0, G contains the disjoint union of any two such subgraphs, and the second condition in the definition of meaningful embeddings implies that G contains the join of any such subgraph with a single black vertex. The recursive construction of bi-quasi-threshold graphs then implies, as required, that G contains every coloured bi-quasi-threshold graph on $h + 1$ vertices. \square

The next two lemmas give a Ramsey type result on the presence of large meaningfully embedded $B_{h,i}$.

Lemma 162. *Let $r \geq 1$. There exists $n = n(r) \in \mathbb{N}$ such that any red-blue colouring of B_n contains internally a monochromatic B_r .*

Proof. We will show by induction on r that the recursion $n(r + 1) = n(r) + r + 2$, $n(1) = 3$, defines a desired function.

To prove the base case $r = 1$, let x be the root of a coloured B_3 , y_1, y_2 its children, and z_j ($1 \leq j \leq 4$) its grandchildren (z_1 and z_2 are the children of y_1 , and z_3 and z_4 are the children of y_2). All of those nodes are internal, since they all have descendants (the nodes on the last level of the B_3). Without loss of generality, we may assume that x is red. If a vertex in $\{y_1, z_1, z_2\}$ and a vertex in $\{y_2, z_3, z_4\}$ are also red, we are done, so assume not. Then in one of the two triples, all vertices are blue, and we are also done.

For the induction step, assume that for some $r \geq 1$ any red-blue colouring of $B_{n(r)}$ contains internally a monochromatic B_r , and consider $B_{n(r+1)} = B_{n(r)+r+2}$. By the induction hypothesis, the top $n(r)$ levels of the $B_{n(r+1)}$ contain without loss of generality a red internal B_r . The leaves of the red B_r are embedded at level at

most $n(r) - 1$, so their children are at level at most $n(r)$, and the subtrees rooted at those children have height at least $r + 2$. Either those subtrees each contain a red internal node, in which case we have a red internal B_{r+1} , or there is one such subtree with all internal nodes being blue, in which case we have a blue internal B_{r+1} . \square

Lemma 163. *Let $r \geq 1$, and let G be a coloured bi-cograph. There exists $m = m(r) \in \mathbb{N}$ such that if T_G contains B_m , then either T_G or T_{G^R} contains a meaningfully embedded $B_{r,0}$ or $B_{r,1}$.*

Proof. The proof consists of two applications of Lemma 162. First, we colour red the nodes of T_G labelled 0 and blue the nodes labelled 1. Since containment is transitive, this guarantees that if T_G contains $B_{n(r)}$, then it contains an internal monochromatic B_r , i.e., a $B_{r,i}$ for some $i \in \{0, 1\}$.

For the second application of the lemma, we start with an internal copy of $B_{n(r),i}$ in T_G . It follows from the definition of the bi-cotree, that for any internal node labelled by i , any of its internal children is labelled by $1 - i$. This implies that for a node $x \in B_{n(r),i}$ and a child y of x , the path P between $\phi(x)$ and $\phi(y)$ contains at least one node labelled $1 - i$, and for every such node, the subtree rooted at it has leaves outside the branch containing $\phi(y)$. Now, for each non-root node $y \in B_{n(r),i}$ with parent x , pick a vertex z labelled $1 - i$ on the path between $\phi(x)$ and $\phi(y)$; if at least one of the leaves in the tree rooted at z excluding the branch containing $\phi(y)$ is black, colour y red. Otherwise colour it blue. Colour the root arbitrarily. It can be checked that a red internal B_r corresponds to a $B_{r,i}$ meaningfully embedded in T_G , while a blue internal B_r corresponds to a $B_{r,i}$ meaningfully embedded in T_{G^R} .

Thus putting $m(r) = n(n(r))$ completes the proof. \square

We are ready to prove the main result of the section.

Theorem 164. *Let $H \in \mathcal{L}$ and $K \in \tilde{\mathcal{L}}$. The class of (H, K) -free bi-cographs has bounded linear clique-width.*

Proof. Suppose that G is an (H, K) -free bi-cograph, and let $r = \max\{|V(H)|, |V(K)|\}$. Let G' be a colouring of G and let $T_{G'}$ be the bi-cotree of G' . By Lemma 161, there is no meaningfully embedded $B_{r,i}$ in $T_{G'}$ or in $T_{G'^R}$. By Lemma 163, $T_{G'}$ contains no $B_{m(r)}$. Finally, by Lemma 159, $\text{blcw}(G') \leq 2m(r)$, and by Remark 145 we can conclude that $\text{lcw}(G) \leq 2m(r)$. \square

Chapter 7

Bipartite permutation graphs

In this chapter, we shift our gaze to the class of bipartite permutation graphs. We present our joint work with Vadim Lozin and Dmitriy Malyshev [ALM21]. In that paper, we examine several properties within the class of bipartite permutation graphs: well-quasi-orderability, boundedness of certain parameters, complexity of certain problems and the existence of universal graphs of a certain size. In each case, we attempt to characterise those properties via minimal obstacles (given as minimal classes where the property does not hold). For brevity, we include here only a selection of the results. In Section 7.1, we provide an upper bound for the lettericity of a bipartite permutation graph, proving a conjecture from [Teh+20]. In Section 7.2, we discuss universal bipartite permutation graphs and make some progress towards a question asked in [Atm+13].

As one might expect, the class \mathcal{BP} of bipartite permutation graphs, first studied in [SBS87], is the intersection of the classes of bipartite graphs and permutation graphs. They can be described as the inversion graphs of permutations avoiding 321 as a pattern (see, e.g., [Wat07, Section 5.7.2]). The main result we will use regarding this class is the universal construction described by Lozin and Rudolf in [LR07b], and presented in Figure 7.1. We emphasize that this figure contains two representations of the same graph. In most of our considerations the square representation is preferable; we denote a graph of this form with n rows and n columns by $H_{n,n}$. It is shown in [LR07b] that $H_{n,n}$ is a bipartite permutation graph which contains every n -vertex bipartite permutation graph as an induced subgraph.




Figure 7.1: Universal bipartite permutation graph $H_{n,n}$ for $n = 4$.

7.1 Lettericity of bipartite permutation graphs

While our main result from this section is stated in the (now familiar) language of lettericity, the question it answers was formulated in a slightly different terminology: the one of Parikh word representability.

Parikh word representable graphs were introduced in [BM16] as follows. Let $A = \{a_1 < a_2 < \dots < a_k\}$ be an ordered alphabet, and let $w = w_1 w_2 \dots w_n$ be a word over A . The *Parikh graph* of w has $\{1, 2, \dots, n\}$ as its vertex set and two vertices $i < j$ are adjacent if and only if there is $p \in \{1, 2, \dots, k-1\}$, such that $w_i = a_p$ and $w_j = a_{p+1}$. The reader might recognise those graphs as letter graphs with decoder $\{(a_1, a_2), (a_2, a_3), \dots, (a_{k-1}, a_k)\}$ for some $k \in \mathbb{N}$.

In [Teh+20], it was shown that the class of Parikh word representable graphs coincides with the class of bipartite permutation graphs (an alternative proof of this result can be obtained by comparing the universal graph from Figure 7.1 with the definition of lettericity). Moreover, it was conjectured that every bipartite permutation graph with n vertices admits a Parikh word representation over an alphabet of $\lfloor \frac{n}{2} \rfloor + 1$ letters. In this section, we prove the conjecture which, in the language of lettericity, is equivalent to the following theorem:

Theorem 165. *Let G be a bipartite permutation graph with n vertices. Then G has lettericity bounded above by $\lfloor \frac{n}{2} \rfloor + 1$. More precisely, it admits a letter graph representation with $k \leq \frac{n}{2} + 1$ letters a_1, \dots, a_k and decoder $\{(a_1, a_2), (a_2, a_3), \dots, (a_{k-1}, a_k)\}$.*

Proof. We first deal with the case when G is connected, and assume $n \geq 2$. We know that G can be expressed as a letter graph on letters a_1, \dots, a_n with decoder $\{(a_i, a_{i+1}) : 1 \leq i \leq n-1\}$.

Among all expressions $w_1 w_2 \dots w_n$ with that decoder, writing $l(j)$ for the index of the letter in position j of the word, pick one that minimises $\sum_{j=1}^n l(j)$ (i.e., an

expression that minimises the sum of the indices of the letters in w). Let us state some properties of this expression:

- a_1 appears at least once in w . If not, we can shift all indices down by 1.
- The last letter is a_2 . Indeed, the last letter cannot be a_1 , since that would mean G has an isolated vertex. If the last letter is a_t , for some $t \geq 3$, we can remove it, and add an a_{t-2} at the beginning of w : this new expression still represents G , but the sum of indices is smaller.
- Denote by r the largest index occurring in w . Let $2 < t \leq r$, and let $w_j = a_t$ be the rightmost appearance of a_t in w . Then there is an a_{t-1} to the right of w_j in w . Note that this holds for $t = r$ – otherwise, as before, we may replace this rightmost a_r with an a_{r-2} in the beginning of the word. Recursively, for $r > t > 2$, the rightmost a_t is thus to the right of the rightmost a_{t+1} , and must have an a_{t-1} to its right by the same argument as above. It follows that the rightmost occurrence of a_t has to its right at least one a_i for each $2 \leq i < t$, and no a_i for $i > t$.
- Let $2 \leq t \leq r$, and let $w_j = a_t$ be the rightmost occurrence of a_t . Then there is an a_{t-1} to the left of w_j in w . Indeed, since G is connected, the vertex j has a neighbour. But as shown above, there are no a_{t+1} s to the right of w_j , hence that neighbour must be an a_{i-1} to its left.

The above discussion implies that w uses letters $1, r$ at least once, and letters $2, \dots, r-1$ at least twice (since for $2 \leq t < r$, letter a_t appears after the rightmost a_{t+1} by the third bullet point, and before the rightmost a_{t+1} by the fourth). Since G has n vertices, this implies $r \leq \frac{n}{2} + 1$.

If G is disconnected, writing G_i , $1 \leq i \leq s$ for its connected components, we can produce words $w(G_i)$ as above for each G_i , where G_1 uses letters 1 to a_{r_1} , G_2 uses letters a_{r_1} to a_{r_2} , and so on. We then obtain a word representing G by concatenating $w(G_s), w(G_{s-1}), \dots, w(G_1)$ in that order. The resulting word uses once more each letter twice, except for possibly the first and the last one. \square

The upper bound in Theorem 165 is tight and attained on graphs of vertex degree at most 1 (see [Teh+20] for arguments given in the terminology of Parikh word representability or [ALW21] for arguments given in the terminology of lettericity). In the rest of this section we show that, within the universe of bipartite permutation graphs, the class of graphs of vertex degree at most 1 is the only obstacle for bounded lettericity, i.e., it is the unique minimal subclass of bipartite permutation graphs of unbounded lettericity.

Theorem 166. *For each p , there is an $f(p)$ such that the lettericity of pK_2 -free bipartite permutation graphs is at most $f(p)$.*

Proof. Let G be a pK_2 -free bipartite permutation graph. Then G has at most $p - 1$ nontrivial connected components, i.e., components of size at least 2. Each component can be embedded, as an induced subgraph, into the universal graph $H_{n,n}$ with at most $3p - 2$ rows, since any connected induced subgraph of the universal graph occupying at least $3p - 1$ rows contains an induced P_{3p-1} and, hence, an induced pK_2 . Therefore, any component of G requires at most $3p - 2$ letters to represent it. Altogether, we need at most $(p - 1)(3p - 2) + 1$ letters to represent G . \square

7.2 Universal graphs within the class of bipartite permutation graphs

As shown in [LR07b], the class of bipartite permutation graphs contains a universal element of quadratic order, i.e., a graph with n^2 vertices that contains all n -vertex bipartite permutation graphs as induced subgraphs. On the other hand, for the class of chain graphs, we have an n -universal graph on $2n$ vertices, i.e., a universal graph of linear order. This raises many questions regarding the growth rates of order-optimal universal graphs for subclasses of bipartite permutation graphs. One of the most immediate questions is identifying a boundary separating classes with a universal graph of linear order from classes where the smallest universal graph is super-linear. In this section, we show that the class of star forests is a minimal hereditary class with a super-linear universal graph.

Before we present the result for star forests, let us observe that in general not every hereditary class \mathcal{X} contains a universal graph, and even if it does, an optimal universal construction for \mathcal{X} does not necessarily belong to \mathcal{X} . In order to circumvent these difficulties (and to ensure downwards closure of the set of classes with, say, linear universal graphs), we will only consider universal constructions consisting of bipartite permutation graphs. In other words, whenever we look for universal constructions for some class $\mathcal{X} \subseteq \mathcal{BP}$, we allow *any* bipartite permutation graphs, and *only* bipartite permutation graphs, to appear in our universal constructions.

In the following lemma we first describe a star forest of order $O(n \cdot \log(n))$ containing all n -vertex star forests as induced subgraphs, and then we show that this construction is (asymptotically) order-optimal in the universe of all bipartite permutation graphs. To simplify notation, throughout this section we denote the

star $K_{1,n}$ by S_n .

Lemma 167. *The minimum number of vertices in a bipartite permutation graph containing all n -vertex star forests is $\Theta(n \cdot \log(n))$.*

Proof. Let F^* be the star forest $S_{\lfloor \frac{n}{1} \rfloor} + S_{\lfloor \frac{n}{2} \rfloor} + \dots + S_{\lfloor \frac{n}{n-1} \rfloor} + S_{\lfloor \frac{n}{n} \rfloor}$. It is a bipartite permutation graph, and it has n connected components and $\sum_{i=1}^n (\lfloor \frac{n}{i} \rfloor + 1)$ vertices.

As $\lfloor x \rfloor < x$, for any x , F^* has $O(\sum_{i=1}^n \frac{n}{i})$ vertices. Recall that the n -th harmonic number $\sum_{i=1}^n \frac{1}{i}$ is equal to $\ln(n) + \gamma + \epsilon_n$, where $\gamma = 0.577\dots$ is the Euler–Mascheroni constant and ϵ_n tends to 0 with n tending to infinity. Therefore, F^* has $O(n \cdot \log(n))$ vertices.

Let us show that F^* is a universal graph for n -vertex star forests. Indeed, let $F = S_{a_1} + S_{a_2} + \dots + S_{a_p}$ be an n -vertex star forest, where $a_1 \geq a_2 \geq \dots \geq a_p$. Clearly, $i \cdot a_i \leq a_1 + \dots + a_i < n$, for any $1 \leq i \leq p \leq n$. Hence, $a_i < \frac{n}{i}$ and $a_i \leq \lfloor \frac{n}{i} \rfloor$, as a_i is an integer. Therefore, for any i , S_{a_i} is an induced subgraph of $S_{\lfloor \frac{n}{i} \rfloor}$. Thus, F is an induced subgraph of F^* .

To prove a lower bound, let H be a bipartite permutation graph containing all n -vertex star forests. It can be embedded, as an induced subgraph, into $H_{n',n'}$ for some n' (see Figure 7.1). Now let n_1, n_2, \dots be a list in non-increasing order of the numbers of vertices of H embedded in each row of $H_{n',n'}$ (so that n_1 is the number of vertices of H in a row of $H_{n',n'}$ with the most vertices of H , and so on). We show that $n_i \geq \frac{1}{2} (\lfloor \frac{n}{10i} \rfloor - 1)$ for any $1 \leq i \leq \lfloor \frac{n}{20} \rfloor$, implying that the graph H has $\Omega(n \cdot \log(n))$ vertices.

Let $t \in \{10i : i \in \mathbb{N}\} \cap \{1, \dots, \lfloor \frac{n}{2} \rfloor\}$. By F_t , we denote the star forest with t connected components, each isomorphic to $S_{\lfloor \frac{n}{t} \rfloor - 1}$. For any t , the graph F_t is an induced subgraph of H and hence F_t must embed into $H_{n',n'}$. Since any two consecutive rows of $H_{n',n'}$ induce a chain graph, i.e., a $2K_2$ -free graph, any row of $H_{n',n'}$ contains the centres of at most two stars of F_t . Each star intersects at most 3 consecutive rows of $H_{n',n'}$, hence a star can intersect the same row as at most 9 other stars. It is therefore possible to find $\frac{t}{10}$ stars $S_{\lfloor \frac{n}{t} \rfloor - 1}$ in F_t such that no two of them intersect the same row of $H_{n',n'}$, and thus there are at least $\frac{t}{10}$ pairwise distinct rows in $H_{n',n'}$, each of which contains at least $\frac{1}{2} (\lfloor \frac{n}{t} \rfloor - 1)$ vertices of H . It follows that $n_{t/10} \geq \frac{1}{2} (\lfloor \frac{n}{t} \rfloor - 1)$ or, changing indices, that $n_i \geq \frac{1}{2} (\lfloor \frac{n}{10i} \rfloor - 1)$ for any $1 \leq i \leq \lfloor \frac{n}{20} \rfloor$ as required. \square

Theorem 168. *The class of star forests is a minimal hereditary class that does not admit a universal bipartite permutation graph of linear order.*

Proof. Let \mathcal{X} be any proper hereditary subclass of the class \mathcal{SF} of star forests. Then, $\mathcal{X} \subseteq \mathcal{SF} \cap \text{Free}(kS_k)$ for some k . Therefore, every graph in \mathcal{X} consists of at most $k-1$ stars with at least k leaves and arbitrarily many stars with at most $k-1$ leaves. But then $(k-1)S_n + nS_{k-1}$ is an n -universal graph for \mathcal{X} of linear order. \square

The class of star forests is not the only obstruction to admitting a universal graph of linear order. To see this, we show that the class of $3S_6$ -free bipartite permutation graphs requires a super-linear universal graph.

Lemma 169. *Suppose H is a bipartite permutation graph containing all n -vertex $3S_6$ -free bipartite permutation graphs as induced subgraphs. Then $|V(H)| = \Omega(n^{3/2})$.*

Proof. To prove the statement, we will show there must be $\Omega(n^3)$ pairs of vertices in H , from which we immediately get

$$|V(H)|^2 \geq \binom{|V(H)|}{2} \geq \Omega(n^3)$$

and so $|V(H)| = \Omega(n^{3/2})$.


We know H can be embedded as an induced subgraph into a universal graph $H_{n',n'}$ for some n' . The main idea is to construct a structure that is “rigid”, in the sense that we can guarantee the distance (within the structure) between certain vertices is not much greater than the distance in $H_{n',n'}$ between the embeddings of those vertices. To this end, we use a result due to Ferguson [Fer20], that states the lettericity of the path P_s , for $s \geq 3$, is precisely $\lfloor \frac{s+4}{3} \rfloor$.¹ In our language, it follows that any embedding of a chordless path P_s into $H_{n',n'}$ uses at least $\lfloor \frac{s+4}{3} \rfloor$ layers (rows).

Since edges appear only between successive layers, the set of layers used by an embedding of the path is an interval. Moreover, the ends of a chordless path do not appear more than one layer away from the extremal layers in the interval, otherwise a $2K_2$ forms between two of the layers. This implies that the distance between the ends of a chordless path P_s must be at least $\lfloor \frac{s+4}{3} \rfloor - 3 = \lfloor \frac{s-5}{3} \rfloor$.

For each t , we construct in two steps a graph Q_t as depicted in Figure 7.2.

It is easy to see that Q_t is a bipartite permutation graph, since we can embed the rows in the figure into successive layers of the universal graph. Moreover, writing d_G for the distance in a graph G , we have (using the triangle inequality and the

¹The bounds from [Pet02] are sufficient for the proof, but make it a bit messier.



Step 1: Start with a chordless path on $3t + 7$ vertices.

Step 2: Add a chordless path on $t + 3$ vertices, connecting it to the previous path as above.

Figure 7.2: The graph Q_t for $t = 3$

above discussion, and assuming for now that Q_t embeds into H),

$$d_H(x, y) \geq d_H(p, q) - 2 \geq t - 2$$

and

$$d_H(x, y) \leq d_{Q_t}(x, y) = t + 2.$$

We now construct bipartite permutation graphs $R_{n,t}$ from Q_t by replacing x and y with independent sets X, Y of twins of size $\lfloor \frac{n-4t-8}{2} \rfloor$ each, with the same adjacencies as x and y respectively (we only construct those $R_{n,t}$ for which the above quantity is positive). We note that $|V(R_{n,t})| \leq n$ by construction, and $R_{n,t}$ is easily seen to be $3S_6$ -free, so that $R_{n,t}$ is an induced subgraph of H . In addition, like with the original x and y , each new pair $x \in X$ and $y \in Y$ has $|d_H(x, y) - t| \leq 2$.

For $3 \leq t \leq \lfloor \frac{n}{6} \rfloor - 2$, we have $|X| = |Y| \geq \lfloor \frac{n}{6} \rfloor$. In particular, each choice of $t \in I := \{3, 4, \dots, \lfloor \frac{n}{6} \rfloor - 2\} \cap \{3 + 5i : i \in \mathbb{N}\}$ witnesses the existence in H of $|X||Y| \geq \lfloor \frac{n}{6} \rfloor^2$ pairs of vertices, and since the pairs' distance ranges for different $t \in I$ do not overlap, the sets of pairs are disjoint. Hence H must contain in total at least

$$|I| \left\lfloor \frac{n}{6} \right\rfloor^2 \geq \frac{1}{5} \left(\left\lfloor \frac{n}{6} \right\rfloor - 5 \right) \left\lfloor \frac{n}{6} \right\rfloor^2 = \Omega(n^3)$$

pairs, as claimed. \square

Lemma 169 shows indeed that there are other obstructions to a linear universal graph, but it is not yet clear what those obstructions are. For instance, the existence of a linear universal graph is a non-trivial question even for S_t -free graphs, i.e., bipartite permutation graphs of maximum degree at most $t - 1$. In fact, it is not clear whether every class with super-linear universal graphs contains a minimal such class. We leave the continuation of this study as an open problem.

Open problem 170. Characterise the family of hereditary subclasses of bipartite permutation graphs that admit a universal bipartite permutation graph of linear order.

We conclude the chapter with one more related open problem. Theorem 165 shows that the graph $H_{n,n}$ is not an optimal universal construction for the class of bipartite permutation graphs, because all graphs in this class can be embedded into $H_{n/2+1,n}$ as induced subgraphs. However, this construction is still quadratic. On the other hand, the following result provides an almost quadratic lower bound on the size of a universal graph. We note that this also yields a lower bound on the size of a shortest proper n -universal, 321-avoiding permutation – a problem asked in [Atm+13].

Theorem 171. *Suppose H is a bipartite permutation graph that contains all n -vertex bipartite permutation graphs as induced subgraphs. Then $|V(H)| = \Omega(n^\alpha)$ for any $\alpha < 2$.*

Proof. We show $|V(H)| = \Omega(n^{(2a-1)/a})$ for each $a \in \mathbb{N}$. This is a generalisation of Lemma 169, which deals with the case $a = 2$.

The proof of Lemma 169 generalises as follows. For $a \in \mathbb{N}$, we get $|V(H)| = \Omega(n^{2a-1/a})$ by counting a -sets of vertices. To do this, we associate to each a -set the $\binom{a}{2}$ -multiset consisting of distances between pairs of its vertices; we will refer to this $\binom{a}{2}$ -multiset as the “distance multiset (in H)” of the original a -tuple. In order to determine that two a -sets are distinct, it is enough to show they have distinct distance multisets.

We generalise the construction of the graphs $R_{n,t}$ to graphs $R_{n,T}$, where T is a set of $a - 1$ natural numbers, each at least 3. To construct $R_{n,T}$, we start with $Q_{\max(T)}$, but instead of inflating just the endpoints of the second path, we inflate the first vertex, then the $j + 3$ rd, for each $j \in T$. By putting an appropriate upper bound (linear in n) on the size of elements in T , say λn , we can arrange that each inflated set X_j has size linear in n , while ensuring $|V(R_{n,T})| \leq n$.

The set T can be viewed as a condition on the distance multiset in $R_{n,T}$ of certain a -sets: an a -set consisting of one vertex from each inflated set X_j has

$\{t + 2 : t \in T\}$ as a subset of its distance multiset in $R_{n,T}$. The distance multiset in H might differ from the one in $R_{n,T}$, but like before, rigidity of the structure ensures the two are within a small tolerance of each other. Therefore, as long as we are careful in choosing what sets T we consider, we can ensure that different choices of T will give rise to different distance multiset subsets in H . This is achieved by choosing, like before, $T \subseteq \{3, \dots, \lambda n\} \cap \{3 + 5i : i \in \mathbb{N}\}$.

One last hurdle is the following: in order to decide that two a -sets of vertices are distinct, we actually need to compare them via their whole distance multisets, not just via the $a - 1$ -subsets coming from the choice of T . We notice, however, that the same distance multiset can account (conservatively) for at most $\binom{a}{a-1}$ different choices of T .

Altogether, each choice of $T \subseteq \{3, \dots, \lambda n\} \cap \{3 + 5i : i \in \mathbb{N}\}$ witnesses the existence of $\Omega(n^a)$ a -sets of vertices in H , and each a -set is repeated overall at most a constant number of times. Since there are $\Omega(n^{a-1})$ choices for T , this shows $|V(H)|^a \geq \binom{|V(H)|}{a} = \Omega(n^{2a-1})$, from which $|V(H)| = \Omega(n^{(2a-1)/a})$ as required. \square

We conjecture that the optimal universal graph is, in fact, quadratic.

Conjecture 172. The minimum number of vertices in a bipartite permutation graph containing all n vertex bipartite permutation graphs is $\Omega(n^2)$.

Establishing the optimal constant would then be a problem analogous to the study of superpatterns from the world of permutations (see, for instance, [BDE14; EV18]²).

²We remark that the study of superpatterns is usually done in the universe of *all* permutations, while here we restrict ourselves to *bipartite* permutation graphs - this is the reason behind the apparent discrepancy between the upper bound from [BDE14] and our lower bound from Theorem 171.

Chapter 8

Quasi-chain graphs

The last class we examine before putting our magnifying glass back in the drawer is the class of so-called *quasi-chain graphs*. This is an extension of the more well-studied class of chain graphs. In [Ale+21b], we define quasi-chain graphs and prove various properties – some algorithmic, some structural and some parametric. As in the previous chapter, we only present here a selection of the results: in Section 8.1, we prove a structural theorem about the graphs. In Section 8.2, as is by now tradition, we study lettericity in the class.

Before that, let us begin by defining the class under investigation, and by providing a simple characterisation. We model our definition on the one for chain graphs (via neighbourhood inclusion), except we relax that property as follows. We say that a linear ordering (a_1, \dots, a_ℓ) of vertices is *good* if for all $i < j$, the neighbourhood of a_j contains at most 1 non-neighbour of a_i . We call a bipartite graph G a *quasi-chain graph* if the vertices in each part of its bipartition admit a good ordering. Alternatively, quasi-chain graphs are characterised as the (coloured) bipartite graphs that do not contain an unbalanced induced copy of $2P_3$.¹ Indeed, in the unbalanced bipartition, one of the parts does not admit a good ordering and hence quasi-chain graphs are free of unbalanced $2P_3$. On the other hand, if a bipartite graph G does not contain an unbalanced induced copy of $2P_3$, then by ordering the vertices in each part in a non-increasing order of their degrees we obtain a good ordering, i.e., G is a quasi-chain graph.

The class of quasi-chain graphs is substantially richer and more complex than the class of chain graphs. In particular, it is not well-quasi-ordered by induced subgraphs [KL11a] and the clique-width is not bounded in this class [LV08]. Additionally, in [Ale+21b], we establish a bijection f between the class of all permu-

¹Notice that $2P_3$ admits two bipartitions: one with parts of equal size (balanced) and the other with parts of different sizes (unbalanced).

tations and a subclass of quasi-chain graphs such that a permutation π contains a permutation ρ as a pattern if and only if the graph $f(\pi)$ contains the graph $f(\rho)$ as an induced subgraph. Together with the NP-completeness of the PATTERN MATCHING problem for permutations this implies the NP-completeness of the INDUCED SUBGRAPH ISOMORPHISM problem for quasi-chain graphs.

The order-preserving embedding of permutations into quasi-chain graphs also implies the existence of infinite antichains of quasi-chain graphs with respect to the induced subgraph relation and hence the unboundedness of lettericity. This motivates our study of lettericity in this class.

In spite of the more complex structure, the quasi-chain graphs inherit some attractive properties of chain graphs. Indeed, the structural characterisation that we show in Section 8.1 implies that quasi-chain graphs admit an implicit representation and that some algorithmic problems that are NP-complete for general bipartite graphs admit polynomial-time solutions when restricted to quasi-chain graphs (we omit the details here).

8.1 The structure of quasi-chain graphs

For two graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ on the same vertex set we denote by $G_1 \otimes G_2$ the graph $G = (V, E_1 \otimes E_2)$, where \otimes denotes the symmetric difference of two sets. The main result in this section is the following theorem.

Theorem 173. *If a bipartite graph $G = (A, B, E)$ is a quasi-chain graph, then $G = Z \otimes H$ for a chain graph Z and a graph H of vertex degree at most two such that $E(H) \cap E(Z)$ and $E(H) - E(Z)$ are matchings. Such a decomposition $G = Z \otimes H$ can be obtained in polynomial time.*

In the proof of this result, we use a word representation for our graphs, which builds the 2-letter graph representation of chain graphs. As a reminder, there is a bijective, order-preserving mapping between words over the alphabet $\{a, b\}$ (under the subword relation) and coloured chain graphs (under the coloured induced subgraph relation). This mapping sends a word w to the graph whose vertices are the entries of w , and we have edges between each a and each b appearing after it in w . See Figure 8.1 for an example (the indices of the letters indicate the order of their appearance in w).

We would like to extend this representation to graphs with the structure claimed in Theorem 173. To do so, we enhance the letter representation described above by introducing extra structure: *bottom edges* between pairs a, b with the a




Figure 8.1: The graph corresponding to the word $w = aababbab$

appearing before the b in w and *top edges* between pairs a, b with the a appearing after the b in w . We require, in addition, that the set of top edges forms a matching and the set of bottom edges forms a matching, and interpret the bottom edges as an instruction to remove the corresponding matching from the chain graph represented by w , and the top edges as an instruction to add the corresponding matching. We call such a word an *enhanced word*. For instance, $w' = \underline{a}ababbab$ is an enhanced word obtained from $w = aababbab$ by adding the bottom edge connecting the first a to the first b and the top edge connecting the second b to the last a .

Naturally, we may define an *enhanced subword containment* relation: for two enhanced words w and w' , w' is an enhanced subword of w if it can be obtained from w by removing entries. We may also extend the usual subword relation to enhanced words by saying w' is a subword of w if the underlying words obtained by removing the top and bottom edges are subwords of each other. We note that any enhanced subword is a subword, but not necessarily vice-versa: for instance, aab is a subword of $\underline{a}ab$, but not an enhanced subword.

If G is the graph described by an enhanced word w , we say w is an *enhanced letter representation* for G . In particular, $w' = \underline{a}ababbab$ is an enhanced letter representation of the graph obtained from the graph in Figure 8.1 by removing the edge a_1b_1 and adding the edge b_2a_4 . It is immediate from our discussion that Theorem 173 can be restated as follows.

Theorem 174. *Any quasi-chain graph admits an enhanced letter representation that can be found in polynomial time.*

Proof. At the core of our proof is an induction on the number of vertices of the quasi-chain graph G . The base case of the induction is trivial. To develop an inductive step, we prove the following claim.

Claim 175. *Let $G = (A, B, E)$ be a quasi-chain graph. Then either G or its bipartite complement has a vertex of degree at most 1.*

Proof. Let a_1, \dots, a_t be the vertices of A in a non-increasing order of their degrees. If a_1 has fewer than 2 non-neighbours, we are done (since a_1 then has degree at most one in the bipartite complement). Otherwise, let b, b' be two non-neighbours of a_1 . Note that b and b' have no common neighbour: if a was a common neighbour, then it would have two private neighbours with respect to a_1 ; since $2P_3$ s are forbidden, a would be adjacent to all but at most one of the neighbours of a_1 , from which $\deg(a) > \deg(a_1)$, contradicting our premise. But then at least one of b and b' has degree at most one, since otherwise an induced $2P_3$ appears. \square

Since the existence of enhanced letter representations is invariant under bipartite complementation and reflection (swapping the parts), we may assume, by reflecting and complementing if necessary, that $G = (A, B, E)$ has a vertex y of degree at most 1, and that $y \in B$.

Now our induction hypothesis says that $G' := G[A \cup (B - \{y\})]$ admits an enhanced letter representation w' . If y is isolated in G , we may always produce a representation w for G by adding b as a prefix to w' . The difficult case is when y has degree 1 in G . Even then, we may easily produce a representation for G by adding b as a prefix to w' and linking it with a top edge to (the letter corresponding to) the vertex x that y is pendant to, *provided that x does not already have an incident top edge in w'* . In the rest of the proof we show that G' admits an enhanced letter representation in which x is not incident to a top edge.

To show this, we first observe that the mapping from enhanced letter representations to graphs is not injective. As a very simple example, the enhanced words ab and \overline{ba} both represent the complete graph on two vertices, while ba and \underline{ab} both represent the edgeless graph on two vertices. In general, we may swap the above pairs when the two letters appear next to each other. We may also swap consecutive instances of the same letter, carrying over the top/bottom edges incident to them, e.g., we may go from \overline{baaaab} to $\overline{baaaaab}$ and vice-versa.

To prove the result, we assume, by contradiction, that in any enhanced letter representation of G' vertex x is incident to a top edge. Among all representations of G' , look at the ones that minimise the distance between x and its top-matched neighbour. Among those representations, pick one where the interval between x and its top-matched neighbour has the minimum number of bottom edges. Write w^* for this representation, and denote by y' the vertex top-matched to x . Given two letters α and β in w^* (two vertices in G'), we write $\alpha < \beta$ to indicate that α appears before β in the word, and denote by $\alpha - \beta$ the interval of letters (vertices) that appear *strictly between* α and β in w^* . In particular, $y' < x$, since $y' \in B$, $x \in A$ and they are top-matched. We now derive a number of conclusions about the interval $y' - x$.

- (1) *The interval $y' - x$ is not empty*, since otherwise we could remove the top edge by swapping y' and x , and due to its minimality, *this interval starts with an a , which we denote a^* , and ends with a b , which we denote b^* .*
- (2) *The interval $y' - x$ does not contain abb as an enhanced subword*, since otherwise the vertices corresponding to the abb together with the vertices x, y and y' induce a $2P_3$ in G .
- (3) *The interval $y' - x$ contains at most two bs* , which follows directly from (1) and (2).

To obtain a contradiction, we analyze the following two cases.

Case 1: a^ and b^* are not bottom-matched.* Then there is no b in the interval $a^* - b^*$. Indeed, if b' belongs to this interval, then, according to (2), a^* is bottom-matched to b' . However, this contradicts the choice of w^* , because, according to (3), this bottom edge can be removed by bringing a^* next to b' and swapping them. In a similar way, in the absence of a second b , any bottom edge can be removed from the interval $y' - x$, implying that this interval has no bottom edges.

We note that at least one of b^* and x must have a bottom-matched neighbour, since otherwise we could reduce the interval by swapping b^* and x and introducing the bottom edge between them. If x has a bottom-matched neighbour, then x, y, y' together with a^*, b^* and the bottom-matched neighbour of x induce a $2P_3$. Therefore, b^* has a bottom-matched neighbour a' with $a' < y'$.

We also note that at least one of a^* and b^* must have a top-matched neighbour, since otherwise we could bring a^* next to b^* , swap them by introducing a top edge, and then reduce the interval by swapping a^* and x . If b^* has a top-matched neighbour, then y', a', x together with b^*, a^* and the top-matched neighbour of b^* induce a $2P_3$. If a^* has a top-matched neighbour, then x, y, y' together with a^*, b^* and the top-matched neighbour of a^* induce another $2P_3$.

Case 2: a^ and b^* are bottom-matched.* Clearly, the interval $a^* - b^*$ is not empty, since otherwise we could remove the bottom edge by swapping a^* and b^* . Also, to avoid an easy reduction to Case 1, we conclude that the letter to the right of a^* is a b (we denote it by b°), and the letter to the left of b^* is an a (we denote it by a°).

We note that either a^* or b° is incident to a top edge, since otherwise we could swap them by introducing the top edge $\overline{b^\circ a^*}$ and then reduce the interval $y' - x$ by swapping y' and b° . Similarly, at least one of a° and b^* is incident to a top edge.

If a^* is incident to a top edge, then x, y, y' together with a^*, b° and a top-matched neighbour of a^* induce a $2P_3$. If a° is incident to a top edge, then x, y, y' together with a°, b^* and a top-matched neighbour of a° induce a $2P_3$. Therefore, b° is top-matched with a vertex a' and b^* is incident to a top edge. We can assume that $x < a'$, since otherwise we could remove the top edge between b° and a' by bringing them next to each other and swapping. But then a^*, b°, a' together with a°, b^* and a top-matched neighbour of b^* induce a $2P_3$.

A contradiction in all cases shows that G' admits an enhanced letter representation in which x is not incident to a top edge and completes the inductive step.

Our case analysis leads to a polynomial-time procedure for removing, if necessary, the top edge incident to x , which can be outlined as follows. The contradictions involving the appearance of a $2P_3$ concern cases that do not actually occur when we apply our procedure, so we ignore them. When a contradiction to the minimality in the construction of w^* appears in the case analysis, we repeatedly execute the operation that lead to the contradiction – we only need to iterate a linear number of times. We invariably arrive at the situation where y' and x appear next to each other, and we simply swap them to remove the top edge. \square


To conclude the section, we observe that the converse to Theorem 174 does not hold. In particular, $2P_3$ has 8 different enhanced letter graph representations (4 per colouring), up to moving the top/bottom edges between twin vertices.

8.2 Lettericity and wqo of quasi-chain graphs


It is shown in [KL11a] that quasi-chain graphs are not wqo under the induced subgraph relation (and indeed, this also follows directly from the order preserving embedding of permutations into quasi-chain graphs given in [Ale+21b]). To start this section, let us provide a simple, explicit example of an infinite antichain in the class.

Let Z_n be the universal chain graph on $2n$ vertices, with the labelling given in Figure 8.2a. Now let Q_n be the graph obtained from Z_n by deleting all edges of the form (a_i, b_{i+1}) (those edges form a matching), then adding a pendant vertex to each of a_1 and b_n , as shown in Figure 8.2b.

Lemma 176. *$(Q_k)_{k \geq 4}$ is an infinite antichain of quasi-chain graphs with respect to the induced subgraph relation.*



(a) The universal chain graph Z_6



(b) The graph Q_6 obtained from it

Figure 8.2: An infinite antichain of quasi-chain graphs

Proof. First, note that the graphs are indeed quasi-chain graphs. This follows from the fact that the ordering $a_1, a_2, \dots, a_n, a'_n$ is good (and, by symmetry, so is $b_n, b_{n-1}, \dots, b_1, b'_1$). Indeed, for $i < j$, a_j has at most one private neighbour with respect to a_i , namely b_j .

To see that the sequence $(Q_k)_{k \geq 4}$ is an antichain, let $4 \leq m \leq n$, and label the vertices of Q_m as in Figure 8.2b, and the vertices of Q_n by replacing a s with α s and b s with β s. Suppose $\iota : Q_m \rightarrow Q_n$ is an induced subgraph embedding. By symmetry and connectedness of Q_m , we may assume ι maps a -vertices to α -vertices and b -vertices to β -vertices, respectively.

Among ordered pairs of α -vertices with incomparable neighbourhoods, (α_1, α_2) is the *only* one where the first vertex has 3 private neighbours with respect to the second. This fact immediately forces $\iota(a_1) = \alpha_1$ and $\iota(a_2) = \alpha_2$. But then

$\iota(b_2) = \beta_2$, since β_2 is the only β -vertex non-adjacent to α_1 , implying that

$\iota(b_3) = \beta_3$, since otherwise the image of b_3 has no candidate neighbour for the image of a_3 , implying that b_1, b'_1 are mapped to β_1, β'_1 , implying that

$\iota(a_3) = \alpha_3$, since α_3 is the only neighbour of β_3 among not yet mapped vertices, implying that

$\iota(b_4) = \beta_4$, since β_4 is the only β -vertex non-adjacent to α_3 among not yet




Figure 8.3: The double-chain graph D_3

mapped vertices, etc.

Proceeding in this way, we conclude that $\iota(a_i) = \alpha_i$ and $\iota(b_i) = \beta_i$ for all $i \leq m$, which is possible only if $m = n$. \square

Knowing that the full class of quasi-chain graphs is not wqo, a natural question is to determine exactly what the obstacles to wqo are in this class. As a first step towards its solution, let us identify the minimal obstacles to unbounded lettericity appearing within this class. For the reader's ease, we recall here two facts about lettericity that we are going to use.

Fact 177. For any graph G and vertex x of G , $\text{let}(G) \leq 2\text{let}(G - x) + 1$ ([Pet02]).

Fact 178. Chain graphs have lettericity at most 2 (see the previous section).

As we have seen before, the class \mathcal{M} of graphs of degree at most one, and the class $\widetilde{\mathcal{M}}$ of their bipartite complements, are minimal of unbounded lettericity (and it is clear from the definition that they are quasi-chain graphs). We claim that, beside \mathcal{M} and $\widetilde{\mathcal{M}}$, there is only one more minimal class of unbounded lettericity among quasi-chain graphs, defined as follows. As before, let Z_n be the prime chain graph on $2n$ vertices illustrated in Figure 8.2a. We construct *double-chain* graphs D_n as follows: start with Z_{3n} , then like in the construction of Q_{3n} , delete all edges of the form (a_i, b_{i+1}) . Finally, delete all vertices whose index is divisible by 3. D_n can be thought of as Z_n , where we replace each vertical edge with a $2P_2$ – see Figure 8.3 for an illustration.

Let \mathcal{D} be the class containing, for each value of n , the graph D_n and all of their induced subgraphs. We note that the chain ordering inherited from the starting graph Z_{3n} is good in D_n , so that \mathcal{D} is indeed a subclass of quasi-chain graphs.²

²We remark that the class \mathcal{D} is, up to some complementations, the cycles-in-a-chain construction from Subsection 3.4.1. This gives an alternate proof for the next lemma.

Lemma 179. \mathcal{D} is a minimal hereditary class of unbounded lettericity.

Proof. We first show that any proper subclass of \mathcal{D} has bounded lettericity. Indeed, such a subclass is D_n -free for an appropriately large n , and any D_n -free graph G contains at most n copies of induced $2P_2$ s. This means we may remove at most $4n$ vertices from G to obtain a chain graph. Fact 178 and repeated application of Fact 177 gives a bound on $\text{let}(G)$ that only depends on n .

It remains to show that lettericity is unbounded in \mathcal{D} . To see this, suppose for a contradiction that the lettericity is bounded by k . The graph D_n consists of n copies of induced $2P_2$ s connected in a chainlike manner. Given a k -letter word w representing D_n , we consider the subwords of w representing each of the $2P_2$ s. In particular, by the pigeonhole principle, for any $t \in \mathbb{N}$, we may find an N large enough such that t of the $2P_2$ s in D_N are represented by the same subword. Those t copies of $2P_2$ s induce a copy of D_t in D_N whose letter graph representation only uses 4 letters; in particular, since any D_t has such a representation, we may assume $k \leq 4$. A similar argument shows that for each D_n there must exist a representation with letters a, b, c, d , where the four respective letter classes are (using the indexing from Figure 8.3) $A := \{a_i : i \equiv 1 \pmod{3}\}$, $B := \{b_i : i \equiv 1 \pmod{3}\}$, $C := \{a_i : i \equiv 2 \pmod{3}\}$ and $D := \{b_i : i \equiv 2 \pmod{3}\}$. Standard arguments show that, up to symmetry, the decoder for this representation must be $\{(a, b), (a, d), (c, b), (c, d)\}$. But even a single $2P_2$ cannot be expressed in this way – a contradiction. \square

We are ready for the main result of this section, which characterises classes of bounded lettericity among quasi-chain graphs. In the proof, given two vertex-disjoint bipartite graphs $G_1 = (A_1, B_1, E_1)$ and $G_2 = (A_2, B_2, E_2)$, we define the *skew-join* of G_1 with G_2 as the graph $(A_1 \cup A_2, B_1 \cup B_2, E_1 \cup E_2 \cup A_1 \times B_2)$.

Theorem 180. Let \mathcal{X} be a hereditary subclass of quasi-chain graphs. Then \mathcal{X} has bounded lettericity if and only if \mathcal{X} excludes at least one graph from each of $\mathcal{M}, \widetilde{\mathcal{M}}$ and \mathcal{D} .

Proof. The “only if” direction is clear, since $\mathcal{M}, \widetilde{\mathcal{M}}$ and \mathcal{D} all have unbounded lettericity. For the “if” direction, let \mathcal{X} be a hereditary subclass of quasi-chain graphs excluding a graph from each of the three classes. It suffices to show that the classes $\mathcal{X}_{s,t,n}$ of $(sP_2, t\widetilde{P}_2, D_n)$ -free quasi-chain graphs have bounded lettericity for all $s, t, n \in \mathbb{N}$, since \mathcal{X} is contained in such a class.

We prove the statement by induction on n . The statement is clearly true if $n = 1$ for all s, t , since $\mathcal{X}_{s,t,1}$ is a subclass of chain graphs, which have lettericity 2.

Now suppose $n \geq 1$, and let $G = (A, B, E) \in X_{s,t,n+1}$. By Theorem 173, $G = Z \otimes H$, where Z is a chain graph, and $E(H) \cap E(Z)$, $E(H) - E(Z)$ are both matchings.

Let a_1, \dots, a_k be the vertices of A listed in non-increasing order with respect to their neighbourhoods in Z . Each vertex a_i gives a partition of A into a “left” part $A_i^l = \{a_1, \dots, a_i\}$ and a “right” part $A_i^r = \{a_{i+1}, \dots, a_k\}$, and a partition of B into $B_i^l = B - N(a_i)$ and $B_i^r = N(a_i)$. This produces a cut of Z into two smaller chain graphs $Z_i^l := Z[A_i^l \cup B_i^l]$ and $Z_i^r := Z[A_i^r \cup B_i^r]$, and it is not difficult to see Z is the skew-join of Z_i^l with Z_i^r , since A_i^l is complete to B_i^r , while A_i^r is anticomplete to B_i^l . Similarly, we obtain a cut of G into quasi-chain graphs G_i^l and G_i^r . We will refer to those cuts as the cuts *induced by* a_i .

These cuts are very neat in the chain graph Z , but how do they look in the original quasi-chain graph G ? Specifically, where do induced $2P_2$ s in G appear with respect to these cuts? The first thing to note is that, for any given cut, the edges between A_i^r and B_i^l in G belong to $E(H) - E(Z)$, and thus induce a matching. Since G is sP_2 -free, there are at most $s - 1$ of them. Similarly, there are at most $t - 1$ non-edges in G between A_i^l and B_i^r . We call the (at most $2s + 2t - 4$) vertices incident to those edges or non-edges *i-dirty*. We call an induced $2P_2$ in G *i-bad* if it does not contain any *i-dirty* vertex (the reasoning being that the bad $2P_2$ s do not simply disappear when removing dirty vertices). We now claim that any *i-bad* $2P_2$ lies completely in G_i^l or in G_i^r (we call it *left i-bad* or *right i-bad* accordingly). To see that this is indeed the case, we simply note that any $2P_2$ with vertices in both G_i^l and G_i^r needs to have either a crossing edge between A^r and B^l , or a crossing non-edge between A^l and B^r . Finally, we call the cut induced by a_i *perfect* if there are no *i-bad* $2P_2$ s, *good* if there is both a left *i-bad* $2P_2$ and a right *i-bad* $2P_2$, and *bad* if it neither good nor perfect. There are three possible cases:

- i) *There is an i such that the cut induced by a_i is perfect.* In this case, we note that “cleaning the cut” by removing all *i-dirty* vertices from G yields a chain graph G' . But we have removed a bounded number of vertices, hence Fact 178 and repeated application of Fact 177 give an upper bound on the lettericity of G that only depends on s and t .
- ii) *There is an i such that the cut induced by a_i is good.* Then like before, cleaning the cut yields a quasi-chain graph G' which is a skew-join of the graphs $G'^l := G' \cap G^l$ and $G'^r := G' \cap G^r$. By construction, G'^l and G'^r each have a $2P_2$; since G (and hence G') is D_{n+1} -free, it follows that G'^l and G'^r are both D_n -free, and the inductive hypothesis applies. From the representations of G'^l

and G^r with a bounded number of letters, it is easy to construct one for their skew-join G' , then use that representation to construct one for G like in the previous case.

- iii) *Every cut is bad.* This means that each a_i has either a left or a right i -bad $2P_2$ (but not both). We note that a_1 must have a right 1-bad $2P_2$, while a_k must have a left k -bad $2P_2$. Moreover, if a $2P_2$ is left, respectively right i -bad, then it is left j -bad for any $j \geq i$, respectively right j -bad for any $j \leq i$. This implies that there is one specific i_0 such that a_1, \dots, a_{i_0} all have right bad $2P_2$ s, while a_{i_0+1}, \dots, a_k all have left bad $2P_2$ s. We claim that no $2P_2$ can be simultaneously i_0 - and $i_0 + 1$ -bad. Indeed, both vertices $a_{i_1}, a_{i_2} \in A$ of such a $2P_2$ would simultaneously need $i_1, i_2 > i_0$ and $i_1, i_2 \leq i_0 + 1$, which is impossible. It follows that cleaning both of the cuts induced by a_{i_0} and a_{i_0+1} leaves us with a chain graph, and we proceed as in the first case.

□

Theorem 180 gives us a characterisation of subclasses of quasi-chain graphs of bounded lettericity. All of those subclasses are wqo, but a wqo class need not have bounded lettericity – for instance, the minimal classes $\mathcal{M}, \widetilde{\mathcal{M}}$ and \mathcal{D} themselves are wqo. For \mathcal{M} and $\widetilde{\mathcal{M}}$, this is a special case of Theorem 2 from [KL11b]. As our last result from this chapter, let us now show the claim for \mathcal{D} .

Theorem 181. *\mathcal{D} is wqo by induced subgraphs.*

Proof. It suffices to produce an order-preserving surjection from a wqo poset (X, \leq) to \mathcal{D} ordered by the induced subgraph relation (this fact is standard – see, e.g., [VW11], Proposition 3.1).

Our poset X will be the set of words over a finite alphabet of incomparable letters, ordered under the subword relation – wqo of this poset is a special case of Higman’s Lemma. Note that a coloured $2P_2$ has, up to isomorphism, 9 distinct non-empty induced subgraphs. Consider an alphabet Ω consisting of incomparable letters A_1, \dots, A_9 , where each letter corresponds (arbitrarily) to one of those induced subgraphs. We define a map φ from the set Ω^* of words over Ω to graphs inductively, by defining $\varphi(A_i)$ to be the corresponding induced subgraph of $2P_2$, and $\varphi(A_i w')$ to be the skew-join of $\varphi(A_i)$ with $\varphi(w')$ (where $A_i w'$ denotes the concatenation of A_i with the word w').

We note that the image of any word of length n is an induced subgraph of D_n (see Figure 8.3), hence $\varphi(\Omega^*) \subseteq \mathcal{D}$. Since any induced subgraph of D_n can be

obtained in this way, φ is surjective. Finally, it is straightforward to check that φ is order-preserving. \square

Chapter 9

Miscellanea

9.1 Linear Ramsey numbers versus co-chromatic number

In this section, we provide a counterexample which appeared in [Ale+21c]. To give it some context, let us briefly mention two relevant conjectures concerning the sizes of homogeneous sets in graphs from certain classes. The first is the Erdős-Hajnal conjecture:

Conjecture 182 (The Erdős-Hajnal conjecture [EH89]). For any graph H , there exists a $\delta_H > 0$ such that all H -free graphs on n vertices have a homogeneous set of size $\Omega(n^{\delta_H})$.

The conjecture thus states that, when we forbid any graph as an induced subgraph, we should be able to find homogeneous sets of polynomial size. This is in contrast to the class of all graphs: the standard lower bounds on Ramsey numbers imply that in general, it is only possible to find homogeneous sets of logarithmic size. For a survey on this conjecture, the reader is invited to consult [Chu13].

A related conjecture is the so-called Gyárfás-Sumner conjecture, which in its original form [Gyá75; Sum81] states that forbidding a tree and a clique produces a class of bounded chromatic number. Chudnovsky and Seymour showed [CS14] that this conjecture is equivalent to the following:

Conjecture 183 (The Gyárfás-Sumner conjecture, alternative formulation). For a finite set S of graphs, the following are equivalent:

- $\text{Free}(S)$ has bounded co-chromatic number.
- S contains a disjoint union of cliques, a complete multipartite graph, a forest and the complement of a forest.

To see that Conjectures 182 and 183 do indeed share some similarities, we note that bounded co-chromatic number does in fact tell us something about the size of homogeneous sets. To make this precise, we first introduce some terminology. For a class \mathcal{X} , we consider the Ramsey numbers *restricted to* \mathcal{X} . More specifically, the number $R_{\mathcal{X}}(p, q)$ is the smallest natural n such that every graph *in* \mathcal{X} on at least n vertices contains either a clique of size p or an independent set of size q (thus the usual Ramsey numbers are the Ramsey numbers $R_{\mathcal{X}}$, with \mathcal{X} the class of all graphs). We say Ramsey numbers are *linear* in a class \mathcal{X} if $R_{\mathcal{X}}(p, q) = O(p + q)$. We note that this terminology is not standard in the literature. Working with our terminology has the advantage that exact values can be written explicitly (which we often do in [Ale+21c]); however, when studying asymptotics, the more popular perspective consists of looking for the size of the largest homogeneous set. In that terminology, a class \mathcal{X} has *linear homogeneous subgraphs* if any graph in \mathcal{X} on n vertices has a homogeneous set of size $\Theta(n)$. As a “dictionary” between the two terminologies, we give the following lemma:

Lemma 184. *Let X be a class of graphs. Then graphs in X have linear homogeneous subgraphs if and only if Ramsey numbers are linear in X . More generally, for any $0 < \delta \leq 1$, the following two statements are equivalent:*

- *There is a constant A such that $\max\{\alpha(G), \omega(G)\} \geq A \cdot |V(G)|^\delta$ for every $G \in X$.*
- *There is a constant B such that $R_X(p, q) \leq B(p + q)^{\frac{1}{\delta}}$.*

Proof. The second claim reduces to the first one when $\delta = 1$, so we just prove the stronger claim.

For the first implication, suppose there exists a constant A such that

$$\max\{\alpha(G), \omega(G)\} \geq A \cdot |V(G)|^\delta$$

for all $G \in X$. Let $H \in X$, let $p, q \in \mathbb{N}$, and suppose that $|V(H)| \geq \left(\frac{p+q}{A}\right)^{\frac{1}{\delta}}$. Then $\max\{\alpha(H), \omega(H)\} \geq A \cdot |V(H)|^\delta \geq p + q$, which means that H is guaranteed to have an independent set of size p or a clique of size q , and this proves the first direction (we can put, e.g., $B = A^{-\frac{1}{\delta}}$) in the statement of the lemma.

Conversely, suppose there exists a positive constant B such that for any $p, q \in \mathbb{N}$ and $G \in X$, if $|V(G)| \geq B(p + q)^{\frac{1}{\delta}}$, then G has an independent set of size p or a clique of size q . Let H be an arbitrary graph in X and let t be the largest integer such that $|V(H)| \geq 2^{\frac{1}{\delta}} B t^{\frac{1}{\delta}} = B(t + t)^{\frac{1}{\delta}}$. By the above assumption, H has a clique or an independent set of size t , i.e., $\max\{\alpha(H), \omega(H)\} \geq t$. Notice, by

definition of t , we have $|V(H)| \leq 2^{\frac{1}{\delta}} B(t+1)^{\frac{1}{\delta}}$, i.e., $|V(H)|^\delta \leq 2B^\delta(t+1)$. Hence if $t = 0$, then $|V(H)|^\delta \leq 2B^\delta$ and therefore $\max\{\alpha(H), \omega(H)\} \geq \frac{|V(H)|^\delta}{2B^\delta} \geq \frac{|V(H)|^\delta}{4B^\delta}$. On the other hand, if $t \geq 1$, then $|V(H)|^\delta \leq 2B^\delta(t+1) \leq 4B^\delta t$ and therefore $\max\{\alpha(H), \omega(H)\} \geq \frac{|V(H)|^\delta}{4B^\delta}$, and putting, e.g., $A = \frac{1}{4B^\delta}$ concludes the proof. \square

We now note that, if a class \mathcal{X} has co-chromatic number bounded by a constant C , then the largest homogeneous set in each graph has at least $\frac{1}{C}n$ vertices (where n denotes the number of vertices in the whole graph). In other words, bounded co-chromatic number immediately implies linear homogeneous sets. With a little bit more work (omitted here, but shown in [Ale+21c]), we obtain that, among classes defined by finitely many minimal forbidden induced subgraphs, those with linear Ramsey numbers must satisfy the second point in Conjecture 183. Thus in the universe of classes defined by finitely many minimal forbidden induced subgraphs, according to Conjecture 183, the three notions – bounded co-chromatic number, linear Ramsey numbers, and the avoidance of the four prescribed induced subgraphs – should coincide.

Our main result in this section is a counterexample showing that those notions do not coincide in the universe of all classes. Specifically, our example distinguishes between bounded co-chromatic number and linear Ramsey numbers in a class.

To construct this counterexample, we consider the Kneser graphs $KG_{a,b}$: the vertices are b -subsets of a set of size a , and two vertices are adjacent if and only if the corresponding subsets are disjoint. A well-known result due to Lovász says that, if $a \geq 2b$, then the chromatic number $\chi(KG_{a,b})$ is $a - 2b + 2$ [Lov78].

Let us denote by \mathcal{X} the hereditary closure of the family of Kneser graphs $KG_{3n,n}$, $n \in \mathbb{N}$, i.e., $\mathcal{X} = \{H : H \text{ is an induced subgraph of } KG_{3n,n}, \text{ for some } n \in \mathbb{N}\}$.

Theorem 185. *The class \mathcal{X} has linear Ramsey numbers and unbounded co-chromatic number.*

Proof. First, we note that by Lovász's result stated above, it follows that $\chi(KG_{3n,n}) = 3n - 2n + 2 = n + 2$. Also, it is not hard to see that the size of the biggest clique in $KG_{3n,n}$ is 3. It follows that the co-chromatic number of $KG_{3n,n}$ is at least $\frac{n+2}{3}$. As a result, the co-chromatic number is unbounded for this class.

Now consider any induced subgraph H of $KG_{3n,n}$. We will show that $\alpha(H) \geq \frac{|V(H)|}{3}$. Indeed, the vertices of the Kneser graph in this case are n -element subsets of $\{1, 2, \dots, 3n\}$. For each $i \in \{1, 2, \dots, 3n\}$ let V_i be the set of vertices of H

containing element i . Then, as each vertex is an n -element subset, it follows that $\sum_{n=1}^{3n} |V_i| = n \times |V(H)|$. Hence, by the Pigeonhole Principle, there is an i such that $|V_i| \geq \frac{|V(H)|}{3}$. As V_i is an independent set, it follows that $\alpha(H) \geq \frac{|V(H)|}{3}$. This implies that for any $H \in \mathcal{X}$ we have $|V(H)| \leq 3\alpha(H) \leq 3(\alpha(H) + \omega(H))$, and hence the Ramsey numbers are linear in the class \mathcal{X} . \square

An interesting direction of further research is to try and better understand the relationship between linear Ramsey numbers on the one hand, and bounded co-chromatic number on the other. For instance, it would be interesting to produce other examples like the one above – perhaps simpler ones. Alternatively, as an additional step in the Gyárfás-Sumner conjecture, we could try to ensure that the two notions coincide in the universe of classes defined by finitely many minimal forbidden induced subgraphs.¹

9.2 The girth of matched k -partite graphs

We start by introducing some terminology.

Definition 186. A *matched k -partite graph* is a k -partite graph in which the edges between any two parts induce a perfect matching (and in particular, all parts have the same size).

Fix labels from a common index set X for the vertices in each part. For $i < j$, the matching between sets A_i and A_j can be interpreted as a permutation $\pi_{i,j} \in S_X$, where $\pi_{i,j}$ sends the label of a vertex in A_i to the label of its matched vertex in A_j . We can extend this notation to any pair (i, j) by setting $\pi_{j,i} = \pi_{i,j}^{-1}$, and $\pi_{i,i} = 1_{S_X}$. This motivates the following definition.

Definition 187. Call a collection $(\pi_{i,j})_{1 \leq i, j \leq k}$ of permutations of X *symmetric* if $\pi_{j,i} = \pi_{i,j}^{-1}$, and $\pi_{i,i} = 1_{S_X}$.

From now on, we will denote a matched k -partite graph by (G, k, X, P) where G is a k -partite graph with k -partition $V(G) = \dot{\cup}_{i=1}^k A_i$, X is a common index set for the vertices in each bag, and $P = (\pi_{i,j})_{1 \leq i, j \leq k}$ is a symmetric collection of permutations of X describing the matchings.

¹A place to start would be to confirm that the class \mathcal{X} from Theorem 185 is not defined by finitely many minimal forbidden induced subgraphs; the author could not find relevant results in the literature.

Remark 188. We could have equivalently restricted ourselves to the “simpler” case $X = [n]$. The reason we allow more general indexing sets is that the additional structure of X will make some examples easier to write.

We would like to investigate the girth of matched k -partite graphs. In particular:

Question. Given $k, g \in \mathbb{N}$ with $g \geq 3$, does there exist a matched k -partite graph (G, k, X, P) of girth at least g ?

The answer to this question is trivially “yes” for $k = 1, 2$ (for an acyclic graph, we define its girth to be infinity). Intuitively, we may expect the answer to be “yes” in general, since increasing the size of X (i.e., the number of vertices in each part) gives us a significant amount of freedom. The purpose of this section is to show that this is indeed the case.

Note that any walk in (G, k, X, P) is uniquely determined by its starting vertex, and the sequence of independent sets it visits. We can thus identify walks with pairs $(x, w = i_1 i_2 \dots i_r)$, where x is the label of the starting vertex, and w is the sequence of indices of the visited sets (i_1 is the label of the set we start in). Since every set is independent, no two consecutive entries of the sequence are the same.

With this set-up, we see that only certain walks may represent cycles in G . Indeed, a walk $(x, w = i_1 i_2 \dots i_r)$ is closed if and only if $i_1 = i_r$, and x is a fixed point of the permutation π_w , defined as the composition $\pi_{i_{r-1}, i_r} \pi_{i_{r-2}, i_{r-1}} \dots \pi_{i_2, i_3} \pi_{i_1, i_2}$. And of course, any closed walk that does not visit any edge twice must contain a cycle. This motivates the following definition:

Definition 189. Call a walk in G (and the corresponding sequence) *potentially cyclic* if $i_1 = i_r$, and $i_s \neq i_{s+2}$ (indices modulo r) for any s .

With all of this in mind, to answer our original question, it suffices to answer the following one in the positive:

Question. Given $k, g \in \mathbb{N}$, is it possible to find a set X and a symmetric collection $P = (\pi_{i,j})_{1 \leq i, j \leq k}$ of permutations in S_X such that for any potentially cyclic sequence w of length at most g with entries in $[k]$, π_w has no fixed points?

Remark 190. Any cycle of length g can be represented by different potentially cyclic sequences (we can start our walk at any vertex in the cycle, and go in either direction). Define an equivalence relation $w \sim w'$, where two potentially cyclic

sequences are equivalent if they would represent the same cycle (i.e., if one of them can be obtained from the other by reversing the order if necessary, and by repeatedly applying the rule $i_1 i_2 \dots i_{r-1} i_1 \mapsto i_2 \dots i_{r-1} i_1 i_2$). If $w \sim w'$, then π_w has fixed points if and only if $\pi_{w'}$ has fixed points. Hence it suffices to check that each potentially cyclic sequence has an equivalent sequence whose corresponding permutation has no fixed points.

We begin by pointing out that triangles can always be avoided.

Lemma 191. *For any $k \in \mathbb{N}$, there exists a matched k -partite graph (G, k, X, P) of girth at least 4.*

Proof. If $k \leq 2$, this is trivial, so we may assume that $k \geq 3$. Put $X = \mathbb{Z}_{k-1}$, the integers modulo $k-1$. For $1 \leq i < j \leq k$, put $\pi_{i,j}(x) = x + i - 1$. Suppose G has a triangle. Without loss of generality, it is represented by a walk $(x, i_1 i_2 i_3 i_1)$ where $i_1 < i_2 < i_3$. The condition that x is fixed by the permutation associated to this walk amounts to saying that $x + i_1 - 1 + i_2 - 1 - i_1 + 1 = x \pmod{k-1}$, i.e., that $i_2 = 1$. But by assumption, $i_2 > i_1 \geq 1$, leading to a contradiction. \square

Now let us look at what happens when we fix k and try to avoid small cycles. We start with $k = 3$.

Lemma 192. *For any $g \geq 3$, there exists a matched 3-partite graph of girth at least g .*

Proof. Note any cycle in G can be represented by a walk $(x, 123123 \dots 1231)$. Put $X = \mathbb{Z}_{\lceil g/3 \rceil}$, and set $\pi_{1,2} = \pi_{1,3} = \pi_{2,3} : x \mapsto x + 1$. For any walk $(x, w = 1231)$, $\pi_w(x) = x + 1$. The fixed point condition tells us that the cycle is a concatenation of $\lceil g/3 \rceil$ such walks, which means its length is $3\lceil g/3 \rceil \geq g$, as required. \square

For the case $k = 3$, we managed to construct an example where all the permutations commute with each other (i.e., the subgroup of S_X generated by P is abelian). For $k \geq 4$, constructing examples is not as straightforward:

Lemma 193. *If $k \geq 4$ and the subgroup $\langle P \rangle$ of S_X generated by P is abelian, then (G, k, X, P) has a cycle of length at most 10.*

Proof. Let $w = 12432134231$. This is a potentially cyclic sequence of length 11, and since all elements of P commute, $\pi_w = 1_{S_X}$. In other words, a walk starting at any vertex in A_1 and visiting the bags in the order given by w satisfies the fixed point condition. Any such walk must contain a cycle of length at most 10. \square

Remark 194. In fact, the conclusion holds if we just assume there are 4 bags A_{i_s} ($1 \leq s \leq 4$) such that the subgroup generated by the permutations π_{i_s, i_t} with $1 \leq s, t \leq 4$ is abelian. This tells us that if we want to find a matched k -partite graph of high girth, we want to look for symmetric collections of permutations whose generated subgroups of S_X are not even “locally” abelian, in the sense that for any 4 bags, the permutations associated to the matchings between the bags cannot all commute with each other.

In order to construct examples for general k , we need to find a symmetric set of k^2 permutations of a certain set X such that no word of length less than g coming from a potentially cyclic sequence has any fixed points. We will borrow some ideas from group theory. The author is grateful to Alex Wendland for a fruitful discussion which pointed him towards the required group theoretic result. We need some standard definitions (see, e.g., [Coh89]):

Definition 195. The *free group* F_S over a set S of generators consists of all reduced words that can be built from members of S and their inverses (a reduced word is one that does not contain xx^{-1} as a subword for any $x \in S$). The group operation is word concatenation.

Definition 196. A group G is called *residually finite* if for any element x that is not the identity in G there is a homomorphism φ from G to a finite group such that $\varphi(x) \neq 1$.

The main result we will use is the following:

Proposition 197 ([Coh89], pp. 7, 11). *Free groups are residually finite.*

Proof. We follow the proof given in [Coh89]. Let $S = (x_i)_{i \in I}$ be a generating set, and let $w = x_{i_1}^{\epsilon_1} x_{i_2}^{\epsilon_2} \dots x_{i_n}^{\epsilon_n}$, where $\epsilon_i = \pm 1$, and if $i_r = i_{r+1}$ then $\epsilon_r = \epsilon_{r+1}$ (i.e., w is a reduced word in F_S). We give a homomorphism $f : F_S \rightarrow S_{n+1}$ that does not send w to 1. More precisely, we specify images for the x_i in such a way that $f(w)$ sends 1 to $n+1$. To achieve this, we want $f(x_{i_r})$ to send r to $r+1$ if $\epsilon_r = 1$, and $r+1$ to r if $\epsilon_r = -1$. A priori, there might be conflicts where $x_{i_r} = x_{i'_r} = x_i$, and $f(x_i)$ sends r to two different numbers. This can only happen if $x_{i_r} = x_{i_{r-1}}$, and $\epsilon_r = 1$ while $\epsilon_{r-1} = -1$. This is prevented by the fact that w is reduced. Note as well that so far, two different numbers are mapped to the same s if $x_{i_{s-1}} = x_{i_s} = x_i$, and $\epsilon_{s-1} = 1$ while $\epsilon_s = -1$. This is again impossible, since w is reduced. We can therefore extend each $f(x_i)$ to an element of S_{n+1} if x_i appears in w , and send the x_i not appearing in w to $1_{S_{n+1}}$. \square

In our setting, the proof of this theorem guarantees us that given a fixed potentially cyclic sequence w of length g with entries in $[k]$, we are able to find a symmetric collection of permutations $P \subseteq S_{g+1}$ labelled by $[k]^2$ such that the composition of permutations corresponding to w is not the identity. This is a step in the right direction. As a corollary to the theorem, we show that we can in fact do this simultaneously for any finite number of sequences.

Corollary 198. *Suppose x and y are two reduced non-identity elements in $G = F_S$ of lengths n_1 and n_2 respectively. There is a homomorphism φ from G to $S_{n_1+1} \times S_{n_2+1}$ such that $\varphi(x), \varphi(y) \neq 1$.*

Proof. From the proof of the theorem, we have homomorphisms $\varphi_1 : G \rightarrow S_{n_1+1}$ and $\varphi_2 : G \rightarrow S_{n_2+1}$ such that $\varphi_1(x) \neq 1$ and $\varphi_2(y) \neq 1$. Define $\varphi : G \rightarrow S_{n_1+1} \times S_{n_2+1}$ by $\varphi(z) := (\varphi_1(z), \varphi_2(z))$. φ has the desired property. \square

In other words, we can guarantee that none of the permutations coming from potentially cyclic walks of length less than g are the identity (since there are finitely many such walks for fixed k). The last step is going from this to guaranteeing that the permutations do not actually have *any* fixed points. This can be done via the following “trick”:

Remark 199. Suppose that a permutation π of X is not the identity. Then the regular action of π on S_X (given by $\pi \cdot \rho := \pi \circ \rho$) has no fixed points.

This leads us to the main result of this section:

Theorem 200. *Given $k, g \in \mathbb{N}$ with $g \geq 3$, there exists a matched k -partite graph (G, k, X, P) of girth at least g .*

Proof. To each potentially cyclic sequence $w = i_1 i_2 \dots i_r$ of length less than g with entries in $[k]$, we associate as before a word of permutations $\pi_w := \pi_{i_{r-1}, i_r} \pi_{i_{r-2}, i_{r-1}} \dots \pi_{i_2, i_3} \pi_{i_1, i_2}$. We need to find a set X , and a symmetric collection P of permutations of X labelled by $[k]^2$ such that none of the π_w have any fixed points. Set $X = \prod_{\pi_w} S_{|w|}$, where w runs over potentially cyclic sequences of length less than g , with entries in $[k]$. By the proposition, corollary and remark above, we can construct $P \subseteq S_X$ with the desired properties. \square

While the proof shows existence of matched k -partite graphs of high girth, the set X we have constructed gets very large very quickly (as a function of k and g). In other words, we need many vertices per bag to achieve high girth. An interesting question is how small we can make X given k and g .

Another potential further direction concerns the chromatic number of matched k -partite graphs. We note that it is not, in general, equal to the number of parts. An example is given by a matched k -partite graph with $k = 3$, $X = \mathbb{Z}_2$, and $\pi_{1,2} = \pi_{1,3} : x \mapsto x$, and $\pi_{2,3} : x \mapsto x + 1$. One can check that G is isomorphic to C_6 , with chromatic number 2. A natural follow-up question is: what can we say about the chromatic number given X and P ? In particular, is it possible to achieve simultaneously large chromatic number and large girth among matched k -partite graphs?

9.3 Of two trees and 2-trees

Let $n \in \mathbb{N}$, and let $T_1 = ([n], E_1)$ and $T_2 = ([n], E_2)$ be two trees with vertex set $[n]$. Let $\varphi \in S_n$ (the symmetric group on n elements), and for sets $U \subseteq [n]$ and $E \subseteq \binom{[n]}{2}$, write $\varphi(U) := \{\varphi(i) : i \in U\}$ and $\varphi(E) := \{\{\varphi(i), \varphi(j)\} : \{i, j\} \in E\}$ respectively. The *union of T_1 and T_2 along φ* is the graph $G_\varphi(T_1, T_2) := ([n], \varphi(E_1) \cup E_2)$. One way of visualising this is to think of T_1 and T_2 as trees on disjoint copies of $[n]$, and then “glue together” each vertex of T_1 and its image via φ in T_2 . See Figure 9.1 for an illustration.

One might hope that, given two trees T_1 and T_2 , the graph $G_\varphi(T_1, T_2)$ will behave nicely for an appropriate choice of φ . To describe this expected “niceness”, recall the definition of partial k -trees:

Definition 201. A k -tree is a graph that can be obtained by starting with K_k and repeatedly adding vertices and connecting them to a clique of size k . A *partial k -tree* is a (not necessarily induced) subgraph of a k -tree.

Note that k -trees are analogous to usual trees (which are the same as 1-trees): the vertices we add are generalisations of leaves.

Bodlaender’s survey [Bod98] and the references therein provide several equivalent characterisations of partial k -trees. The two that matter to us are the following:

- they are the graphs of treewidth at most k (though we will not discuss the usual definition of treewidth here);
- they are the subgraphs of chordal graphs² with maximum clique size $k + 1$.

Playing around with some simple examples might lead the reader to believe that pairs of trees can always be glued together in a nice way. For instance, gluing

²A chordal graph is a graph with no induced cycle on 4 or more vertices.




Figure 9.1: The union of two trees along different permutations

a star to any tree always produces a partial 2-tree. To see this, note that the union $G_\varphi(K_{1,n-1}, T)$ where T is a tree on n vertices is determined by the image of the centre of the star via φ . Let x be this image, and root T at it. We may now easily construct a 2-tree containing $G_\varphi(K_{1,n-1}, T)$: start with x and any of its children in T ; then add all of the other children of x , connecting each of them to a previous child and to x . Once we are done with the children of x , recursively add each successive level, connecting each vertex to its parent, and to x . It is clear that the graph constructed in this way is a 2-tree, and that it contains $G_\varphi(K_{1,n-1}, T)$.

Another simple argument shows that a path P_n can always be glued to a given tree T to produce a partial 2-tree (just root T , and glue the path in the order given by a depth-first search on T). Note, however, that this time, not every φ produces a partial 2-tree! For instance, P_4 can be glued to itself in a “bad” way (see Figure 9.2). Indeed, any chordal graph containing K_4 has clique size greater than 3, so the gluing from the figure does not produce a partial 2-tree. Nevertheless, it is tempting to hope that *some* φ should work given two fixed trees. This is, however, not the case. It turns out that there are pairs of trees for which no union produces a




Figure 9.2: $G_{2413}(P_4, P_4) = K_4$

partial 2-tree. In fact, it is even possible to construct trees for which any union has arbitrarily large clique-width, and thus arbitrarily large treewidth,³ but the proof is long, and we do not include it here.

Instead, in this section, we will investigate a subclass of trees which do admit partial 2-tree unions. A *caterpillar* is a tree in which all vertices are within distance 1 of a path. The main result from this section is the following:

Theorem 202. *Let T_1 and T_2 be two caterpillars on n vertices. Then there exists φ such that $G_\varphi(T_1, T_2)$ is a partial 2-tree.*

Proof. We start by (re-)labelling the caterpillars. For each of them, let x_1, x_2, \dots, x_r be the vertices in the central path (that is, the x_i are the non-leaf vertices, and consecutive x_i are adjacent). Order the vertices by concatenating, for each $i = 1, \dots, r$, the leaves pendant to x_i followed by x_i itself. Then assign increasing labels to the vertices, in this order. This is illustrated in Figure 9.3.

We now claim that that, with this labelling, the union of T_1 and T_2 along the identity is chordal and K_4 -free; from the equivalent characterisations given above, this is enough.

Let us show our claim. Note that the labelling of the caterpillars satisfies the following properties:

- (*) a vertex with label i has at most one neighbour with label $j > i$;
- (**) if two vertices with labels $i < j$ are adjacent, then for any s with $i \leq s < j$, the vertex with label s is adjacent to the vertex with label j .

Let us now colour the edges of T_1 red and the edges of T_2 black (like in the figures), and let $G = G_\varphi(T_1, T_2)$ be their union along the identity φ (in G , edges are allowed to be both red and black). We now note:

³Bounded treewidth implies bounded clique-width [LR07a].




Figure 9.3: Two relabelled caterpillars T_1 and T_2

- G is K_4 -free. Indeed, it cannot contain a K_4 , since property (*) would imply the vertex in the K_4 with the smallest label has at most one black edge and one red edge to vertices with greater labels, yielding a contradiction.
- G is chordal. To see this, consider a cycle in G , and let x be the vertex in the cycle with the smallest label. Let y and z be its neighbours in the cycle, and suppose that y has label smaller than z . Then by property (**), y must be adjacent to z , and thus the cycle either has length 3, or is not induced. \square

While there are examples whose unions always have large treewidth, it is not completely clear what causes this to happen. It is therefore interesting to see if adding additional restrictions causes trees to have unions with small (not necessarily 2) treewidth. For instance, we can conjecture:

Conjecture 203. Let T_1 and T_2 be two trees whose vertices lie within distance t from a central path. Then there exists $f(t)$ and a permutation φ such that $G_\varphi(T_1, T_2)$ has treewidth at most $f(t)$.

There is another interesting subclass of trees for which we can investigate this question. We may think of a caterpillar as a “path of stars”. A *spider* is a tree which has one vertex with degree at least 3, and all other vertices of degree at most 2. In other words, we may think of it as a “star of paths”.

Conjecture 204. Let T_1 and T_2 be two spiders. Then there exists φ such that $G_\varphi(T_1, T_2)$ is a partial 2-tree.

Bibliography

- [AAL19] B. Alecu, A. Atminas, and V. Lozin. “Graph Functionality”. *Lecture Notes in Comput. Sci.* 11789 (2019), pp. 135–147.
- [AB20] A. Atminas and R. Brignall. “Well-quasi-ordering and finite distinguishing number”. *J. Graph Theory* 95.1 (2020), pp. 5–26.
- [AH07] M. Aschenbrenner and R. Hemmecke. “Finiteness theorems in stochastic integer programming”. *Found. Comput. Math.* 7 (2007), pp. 183–227.
- [Alb+13] M. H. Albert, M. D. Atkinson, M. Bouvel, N. Ruškuc, and V. Vatter. “Geometric grid classes of permutations”. *Trans. Amer. Math. Soc.* 365 (2013), pp. 5859–5881.
- [Ale+20a] B. Alecu, M. Kanté, V. Lozin, and V. Zamaraev. “Between clique-width and linear clique-width of bipartite graphs”. *Discrete Math.* 343 (2020).
- [Ale+20b] B. Alecu, V. Lozin, D. de Werra, and V. Zamaraev. “Letter Graphs and Geometric Grid Classes of Permutations: Characterization and Recognition”. *Discrete Appl. Math.* 283 (2020), pp. 482–494.
- [Ale+21a] B. Alecu, V. E. Alekseev, A. Atminas, V. Lozin, and V. Zamaraev. “Graph parameters, implicit representations and factorial properties of graphs” (2021). In preparation.
- [Ale+21b] B. Alecu, A. Atminas, V. Lozin, and D. Malyshev. “Combinatorics and algorithms for quasi-chain graphs” (2021). Preprint, available at <https://arxiv.org/abs/2104.04471>. An extended abstract was accepted for IWOCA 2021, and will appear in LNCS.
- [Ale+21c] B. Alecu, A. Atminas, V. Lozin, and V. Zamaraev. “Graph classes with linear Ramsey numbers”. *Discrete Math.* 344.4 (2021).

- [Ale92] V. E. Alekseev. “Range of values of entropy of hereditary classes of graphs”. *Diskret. Mat.* 4.2 (1992). In Russian; translation in *Discrete Math. Appl.* 3 (1993), no. 2, 191–199., pp. 148–157.
- [Ale97] V. E. Alekseev. “On lower layers of the lattice of hereditary classes of graphs”. *Diskretn. Anal. Issled. Oper. Ser. 1* 4.1 (1997). In Russian., pp. 3–12.
- [All09] P. Allen. “Forbidden induced bipartite graphs”. *J. Graph Theory* 60.3 (2009), pp. 219–241.
- [ALM21] B. Alecu, V. Lozin, and D. Malyshev. “Critical properties of bipartite permutation graphs”. *European J. Combin.* (2021). Submitted. Preprint available at <https://arxiv.org/abs/2010.14467>.
- [Alo+06] N. Alon, G. Brightwell, H. Kierstead, A. Kostochka, and P. Winkler. “Dominating sets in k -majority tournaments”. *J. Combin. Theory Ser. B* 96 (2006), pp. 374–387.
- [ALR12] A. Atminas, V. V. Lozin, and I. Razgon. “Linear time algorithm for computing a small biclique in graphs without long induced paths”. *Lecture Notes in Computer Science* 7357 (2012), pp. 142–152.
- [ALW21] B. Alecu, V. Lozin, and D. de Werra. “The micro-world of cographs”. *Discrete Appl. Math.* (2021). Submitted, extended abstract available at https://doi.org/10.1007/978-3-030-48966-3_3.
- [AMR02] M. Atkinson, M. Murphy, and N. Ruškuc. “Partially Well-Ordered Closed Sets of Permutations”. *Order* 19 (2002), pp. 101–113.
- [Atk99] M. D. Atkinson. “Restricted permutations”. *Discrete Math.* 195 (1999), pp. 27–38.
- [Atm+13] A. Atminas, V. V. Lozin, S. Kitaev, and A. Valyuzhenich. “Universal graphs and universal permutations”. *Discrete Math., Algorithms and Applications* 5.4 (2013). 1350038.
- [Atm+15] A. Atminas, A. Collins, V. Lozin, and V. Zamaraev. “Implicit representations and factorial properties of graphs”. *Discrete Math.* 338 (2015), pp. 164–179.
- [Atm17] A. Atminas. “Classes of graphs without star forests and related graphs” (2017). Preprint, available at <https://arxiv.org/abs/1711.01483>.
- [BBW00] J. Balogh, B. Bollobás, and D. Weinreich. “The speed of hereditary properties of graphs”. *J. Combin. Theory Ser. B* 79 (2000), pp. 131–156.

- [BDE14] M. J. Bannister, W. E. Devanny, and D. Eppstein. “Small Superpatterns for Dominance Drawing”. *2014 Proceedings of the Meeting on Analytic Algorithmics and Combinatorics (ANALCO)*. 2014, pp. 92–103.
- [BFR72] K.A. Baker, P.C. Fishburn, and F.S. Roberts. “Partial Orders of Dimension 2”. *Networks* 2 (1972), pp. 11–28.
- [BHH12] R. Belmonte, P. Heggernes, and P. Van ’t Hof. “Edge contractions in subclasses of chordal graphs”. *Discrete Appl. Math.* 160.7–8 (2012), pp. 999–1010.
- [BKV17] R. Brignall, N. Korpelainen, and V. Vatter. “Linear clique-width for hereditary classes of cographs”. *Journal of Graph Theory* 84 (2017), pp. 501–511.
- [BLS99] A. Brandstädt, V. B. Le, and J. Spinrad. *Graph Classes: A Survey*. SIAM Monographs on Discrete Mathematics and Applications. Philadelphia: Society for Industrial and Applied Mathematics, 1999.
- [BM16] S. Bera and K. Mahalingam. “Structural properties of word representable graphs”. *Math. Comput. Sci.* 10 (2016), pp. 209–222.
- [Bod89] H. L. Bodlaender. “Achromatic number is NP-complete for cographs and interval graphs”. *Inform. Process. Lett.* 31 (1989), pp. 135–138.
- [Bod98] H. L. Bodlaender. “A partial k -arboretum of graphs with bounded treewidth”. *Theoret. Comput. Sci.* 209 (1998), pp. 1–45.
- [Bon+20a] E. Bonnet, Colin Geniet, E. Kim, S. Thomassé, and R. Watrigant. “Twin-width II: small classes” (2020). Preprint, available at <https://arxiv.org/abs/2006.09877>.
- [Bon+20b] E. Bonnet, E. Kim, S. Thomassé, and R. Watrigant. “Twin-width I: tractable FO model checking” (2020). Preprint, available at <https://arxiv.org/abs/2004.14789>.
- [CG14] C. Crespelle and P. Gambette. “(Nearly-)tight bounds on the contiguity and linearity of cographs”. *Theoretical Computer Science* 522 (2014), pp. 1–12.
- [Cha17] M. Chandoo. “Computational Complexity Aspects of Implicit Graph Representations”. Available at <https://www.repo.uni-hannover.de/handle/123456789/3601?show=full>. PhD thesis. Fakultät für Elektrotechnik und Informatik der Gottfried Wilhelm Leibniz Universität Hannover, 2017.

- [Chu13] M. Chudnovsky. “The Erdős-Hajnal Conjecture – A Survey”. *J. Graph Theory* 75.2 (2013), pp. 178–190.
- [CHY08] G. J. Chang, L.-H. Huang, and H.-G. Yeh. “On the rank of a cograph”. *Linear Algebra Appl.* 429.2–3 (2008), pp. 601–605.
- [Coh89] D. E. Cohen. *Combinatorial Group Theory: A Topological Approach*. Cambridge University Press, 1989.
- [CS14] M. Chudnovsky and P. Seymour. “Extending the Gyárfás-Sumner conjecture”. *J. Combin. Theory B* 105 (2014), pp. 11–16.
- [Dam90] P. Damaschke. “Induced subgraphs and well-quasi-ordering”. *J. Graph Theory* 14.4 (1990), pp. 427–435.
- [DDL13] K. Dabrowski, M. Demange, and V. V. Lozin. “New results on maximum induced matchings in bipartite graphs and beyond”. *Theoretical Computer Science* 478 (2013), pp. 33–40.
- [Die00] R. Diestel. *Graph Theory*. Springer-Verlag, 2000.
- [Die90] V. Diekert. “Combinatorics on traces”. *Lecture Notes in Computer Science* 454 (1990).
- [DK05] R. Diestel and D. Kühn. “Graph minor hierarchies”. *Discrete Appl. Math* 145 (2005), pp. 167–182.
- [DM41] B. Dushnik and E. Miller. “Partially Ordered Sets”. *American J. Math.* 63.3 (1941), pp. 600–610.
- [EGS90] P. Erdős, J. Gimbel, and H. J. Straight. “Chromatic number versus cochromatic number in graphs with bounded clique number”. *European J. Combinatorics* 11 (1990), pp. 235–240.
- [EH89] P. Erdős and A. Hajnal. “Ramsey-type theorems”. *Discrete Appl. Math.* 25 (1989), pp. 37–52.
- [Erd59] P. Erdős. “Graph theory and probability”. *Canad. J. Math.* 11 (1959), pp. 34–38.
- [ES12] D. Eppstein and E. S. Spiro. “The h -index of a graph and its application to dynamic subgraph statistics”. *J. Graph Algorithms and Applications* 16 (2012), pp. 543–567.
- [EV18] M. Engen and V. Vatter. “Containing all permutations” (2018). Preprint, available at <https://arxiv.org/abs/1810.08252v4>.

- [FBS20] T. Forster, N. Bowler, and M. Seisenberger. *An Introduction to WQO and BQO Theory*. Manuscript available at <https://www.dpmms.cam.ac.uk/~tf/BQObok.pdf>. 2020.
- [Fer20] R. Ferguson. “On the lettericity of paths” (2020). Preprint, available at <https://arxiv.org/abs/2007.03636v1>.
- [FH96] P. C. Fishburn and P. L. Hammer. “Bipartite dimensions and bipartite degrees of graphs”. *Discrete Math.* 160 (1996), pp. 127–148.
- [FL89] M. R. Fellows and M. A. Langston. “On search, decision and the efficiency of polynomial-time algorithms (extended abstract)”. *STOC*. 1989, pp. 501–512.
- [Fra54] R. Fraïssé. “Sur l’extension aux relations de quelques propriétés des ordres”. *Ann. Sci. Ecole Norm. Sup.* 71.3 (1954). In french, pp. 363–388.
- [FV21] R. Ferguson and V. Vatter. “Bounded lettericity for classes with finitely many prime graphs” (2021). Preprint, available at <https://arxiv.org/abs/2104.03055>.
- [Gal67] T. Gallai. “Transitiv orientierbare Graphen”. *Acta Math. Acad. Sci. Hung.* 18 (1967), pp. 25–66.
- [Gan+19] R. Ganian, P. Hliněný, J. Nešetřil, J. Obdržálek, and P. Ossona de Mendez. “Shrub-depth: capturing height of dense graphs”. *Logical Methods in Comp. Sci.* 15.1 (2019), 7:1–7:25.
- [Gho19] E. Ghorbani. “Cographs: eigenvalues and Dilworth number”. *Discrete Math.* 342.10 (2019), pp. 2797–2803.
- [Gol+95] P. Goldberg, M. Golumbic, H. Kaplan, and R. Shamir. “Four strikes against physical mapping of DNA”. *Journal of Computational Biology* 2.1 (1995), pp. 139–152.
- [GR00] M.C. Golumbic and U. Rotics. “On the clique-width of some perfect graph classes”. *Internat. J. Found. Comput. Sci.* 11.3 (2000), pp. 423–443.
- [GV97] V. Giakoumakis and J. Vanherpe. “Bi-complement Reducible Graphs”. *Advances in Appl. Math.* 18 (1997), pp. 389–402.
- [GW07] F. Gurski and E. Wanke. “Line graphs of bounded clique-width”. *Discrete Math.* 307 (2007), pp. 2734–2754.

- [Gyá75] A. Gyárfás. “On Ramsey covering-numbers”. *Infinite and Finite Sets* 2 (1975), pp. 801–816.
- [Har19] N. Harms. “Universal Communication, Universal Graphs, and Graph Labeling” (2019). Preprint, available at <https://arxiv.org/abs/1911.03757>.
- [Hig52] G. Higman. “Ordering by divisibility in abstract algebras”. *Proc. London Math. Soc.* s3-2.1 (1952), pp. 326–336.
- [HK07] P. Heggernes and D. Kratsch. “Linear-time certifying recognition algorithms and forbidden induced subgraphs”. *Nord. J. Comput.* 14 (2007), pp. 87–108.
- [HMP12] P. Heggernes, D. Meister, and C. Papadopoulos. “Characterising the linear clique-width of a class of graphs by forbidden induced subgraphs”. *Discrete Appl. Math.* 160 (2012), pp. 888–901.
- [HV06] S. Huczynska and V. Vatter. “Grid Classes and the Fibonacci Dichotomy for Restricted Permutations”. *Electron. J. Combin.* 13.R54 (2006).
- [IR77] A. Itai and M. Rodeh. “Finding a minimum circuit in a graph”. *STOC*. 1977.
- [Jan99] P. Jančar. “A note on well-quasi-orderings for powersets”. *Inform. Process. Lett.* 72.5–6 (1999), pp. 155–160.
- [K A07] S. D. Nikolopoulos K. Asdre. “NP-completeness results for some problems on subclasses of bipartite and chordal graphs”. *Theoret. Comput. Sci.* 381.1–3 (2007), pp. 248–259.
- [KBH01] A. Koster, H. Bodlaender, and S. van Hoesel. “Treewidth: Computational Experiments” (2001). ZIB-Report 01–38, available at <https://studylib.net/doc/14927407/treewidth--computational-experiments-konrad-zuse-zentrum-....>
- [KL11a] N. Korpelainen and V. Lozin. “Bipartite induced subgraphs and well-quasi-ordering”. *J. Graph Theory* 67 (2011), pp. 235–249.
- [KL11b] N. Korpelainen and V. Lozin. “Two forbidden induced subgraphs and well-quasi-ordering”. *Discrete Math.* 311.16 (2011), pp. 1813–1822.
- [KNR88] S. Kannan, M. Naor, and S. Rudich. “Implicit representation of graphs”. *STOC*. 1988, pp. 334–343.

- [Lam12] M. Lampis. “Algorithmic meta-theorems for restrictions of treewidth”. *Algorithmica* 64 (2012), pp. 19–37.
- [Lav71] R. Laver. “On Fraïssé’s Order Type Conjecture”. *Annals of Mathematics* 93.1 (1971), pp. 89–111.
- [Lav76] R. Laver. “Well-quasi-orderings and sets of finite sequences”. *Mathematical Proceedings of the Cambridge Phil. Soc.* 79.1 (1976), pp. 1–10.
- [Lej+19] M. Lejeune, V. Lozin, I. Lozina, A. Ragab, and S. Yacout. “Recent advances in the theory and practice of Logical Analysis of Data”. *European J. Oper. Res.* 275 (2019), pp. 1–15.
- [LMZ12] V. Lozin, C. Mayhill, and V. Zamaraev. “Locally bounded coverings and factorial properties of graphs”. *European J. Combin.* 33 (2012), pp. 534–543.
- [Lov78] L. Lovász. “Kneser’s conjecture, chromatic number, and homotopy”. *J. Combin. Theory A* 25.3 (1978), pp. 319–324.
- [Loz18] V. Lozin. “Graph Parameters and Ramsey Theory”. *Lecture Notes in Computer Science* 10765 (2018). Extended version available at <https://arxiv.org/abs/1608.07727v2>., pp. 185–194.
- [LR07a] V. Lozin and D. Rautenbach. “The relative clique-width of a graph”. *Journal of Combinatorial Theory B* 97 (2007), pp. 846–858.
- [LR07b] V. Lozin and G. Rudolf. “Minimal universal bipartite graphs”. *Ars Combinatoria* 84 (2007), pp. 345–356.
- [LV08] V. Lozin and J. Volz. “The clique-width of bipartite graphs in monogenic classes”. *Internat. J. Found. Comput. Sci.* 19 (2008), pp. 477–494.
- [LZ17] V. Lozin and V. Zamaraev. “The structure and the number of P_7 -free bipartite graphs”. *European J. Combinatorics* 65 (2017), pp. 143–153.
- [MP95] N. V. R. Mahadev and U. N. Peled. “Threshold graphs and related topics”. *Annals of Discrete Math.* 56 (1995), pp. xiv+543.
- [Mul88] J. H. Muller. “Local structure in graph classes”. PhD thesis. Georgia Institute of Technology, 1988.
- [MV02] M. Murphy and V. Vatter. “Profile classes and partial well-order for permutations”. *Electron. J. Combin.* 9.2 (2002).

- [Nas65] C. St. J. A. Nash-Williams. “On well-quasi-ordering infinite trees”. *Proc. Camb. Phil. Soc.* 61 (1965), pp. 697–720.
- [Nor+06] S. Norine, P. Seymour, R. Thomas, and P. Wollan. “Proper minor-closed families are small”. *J. Combin. Theory Ser. B* 96 (2006), pp. 754–757.
- [Pet02] M. Petkovšek. “Letter graphs and well-quasi-order by induced subgraphs”. *Discrete Math.* 244 (2002), pp. 375–388.
- [Rad54] R. Rado. “Partial well-ordering of sets of vectors”. *Mathematika* 1.2 (1954), pp. 89–95.
- [Ram30] F. P. Ramsey. “On a problem of formal logic”. *Proceedings of the London Mathematical Society* 30 (1930), pp. 264–286.
- [Rob69] F. S. Roberts. “On the boxicity and cubicity of a graph”. *Recent Progress in Combinatorics* (1969), pp. 301–310.
- [RS04] N. Robertson and P. Seymour. “Graph Minors. XX. Wagner’s conjecture”. *J. Combinatorial Theory B* 92 (2004), pp. 325–357.
- [SBS87] J. Spinrad, A. Brandstädt, and L. Stewart. “Bipartite permutation graphs”. *Discrete Appl. Math.* 18 (1987), pp. 279–292.
- [Sch64] J. Schönheim. “On coverings”. *Pacific J. Math.* 14.4 (1964), pp. 1405–1411.
- [Spi03] J.P. Spinrad. *Efficient graph representations*. Vol. 19. Fields Institute Monographs. American Mathematical Society, 2003.
- [Sum81] D. P. Sumner. “Subtrees of a graph and chromatic number”. In: *The Theory and Applications of Graphs* (1981), pp. 557–576.
- [SZ94] E. R. Scheinerman and J. Zito. “On the size of hereditary classes of graphs”. *J. Combin. Theory Ser. B* 61 (1994), pp. 16–39.
- [Teh+20] W. C. Teh, Z. C. Ng, M. Javaid, and Z. J. Chern. “Parikh word representability of bipartite permutation graphs”. *Discrete Appl. Math.* 282 (2020), pp. 208–221.
- [VW11] V. Vatter and S. Waton. “On partial well-order for monotone grid classes of permutations”. *Order* 28 (2011), pp. 193–199.
- [VW19] D. Vietz and E. Wanke. “The fault-tolerant metric dimension of cographs”. *Lecture Notes in Comput. Sci.* 11651 (2019), pp. 350–364.
- [Wae27] B. L. van der Waerden. “Beweis einer Baudetschen Vermutung”. *Nieuw Arch. Wisk.* 15 (1927). In German, pp. 212–216.

- [Wat07] S. D. Waton. “On permutation classes defined by token passing networks, gridding matrices and pictures: three flavours of involvement”. Available at <https://pdfs.semanticscholar.org/f45e/4ff6160425fbb0068f291be7a7d36ef64651.pdf>. PhD thesis. University of St Andrews, 2007.
- [YCC96] J.-H. Yan, J.-J. Chen, and G. J. Chang. “Quasi-threshold graphs”. *Discrete Appl. Math.* 69 (1996), pp. 247–255.
- [Zve00] I. Zverovich. “Perfect cochromatic graphs” (2000). Research report RRR 16-2000, archived version available at <https://web.archive.org/web/20160303173528/http://rutcor.rutgers.edu/pub/rrr/reports2000/16.ps>.