Cover Graphs and Order Dimension

vorgelegt von

M. Sc.–Mathematiker

Veit Wiechert

aus Lübz

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

Doktor der Naturwissenschaften
– Dr. rer. nat. –

vorgelegte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Jörg Liesen
Berichter: Prof. Dr. Stefan Felsner
Prof. Dr. Tibor Szabó
Dr. Piotr Micek

Eidesstattliche Versicherung

Veit Wiechert

Hiermit versichere ich an Eides statt, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe. Alle Ausführungen, die anderen veröffentlichten oder nicht veröffentlichten Schriften wörtlich oder sinngemäß entnommen wurden, habe ich kenntlich gemacht. In Fällen von Co-Autorenschaft ist die Darstellung des Eigenanteils gemäß der beigefügten Anlage "Vorveröffentlichungen" zutreffend .

Die Arbeit hat in gleicher oder ähnlicher Fassung noch keiner anderen Prüfungsbehörde vorgelegen.

Titel der schriftlichen Arbeit:
Cover Graphs and Order Dimension

Verfasser:
Veit Wiechert

Betreuer:
Stefan Felsner

Ort, Datum, Unterschrift

Acknowledgments

This dissertation is the result of three years of research within the *Discrete Mathematics Group* at TU Berlin. I like to express my greatest gratitude to my advisor Stefan Felsner for introducing me to the world of order dimension and for every kind of support. I enjoyed big freedom under his supervision and was always free to choose my favorite research problem. Stefan is also the reason why his working group is known for its relaxed and productive atmosphere. I also like to thank my (former) colleagues from this group for the fantastic time here: Andrea Hoffkamp, Kolja Knauer, Daniel Heldt, Torsten Ueckerdt, Irina Mustaţă, Nieke Aerts, Udo Hoffmann, Linda Kleist, and Torsten Mütze. Special thanks go to my office mate Linda for the many funny moments we shared over the last six months.

Another important person of my academic career is Tom Trotter. He is pushing and contributing to *dimension theory* for many decades now. Still, he seems to be as enthusiastic about posets and their dimension as he was from the beginning on. I am grateful to him for sharing this enthusiasm in the numerous discussions we had. I also would like to thank him for giving me the opportunity to speak at the minisymposium he organized around the SIAM Conference on Discrete Mathematics in Halifax in 2012. This event plays a special role in my career as it strongly influenced my research interests.

I would like to thank my colleagues and friends Gwenaël Joret and Piotr Micek for many exciting research sessions. Especially the week we spent together in Brussels in 2015 was a wonderful experience. We worked an uncountable number of hours on two of our papers; needless to say that I learned a lot about writing mathematics from Gwen and Piotr. While his 2-years stay at our working group, Piotr introduced me to many interesting math problems. I am grateful to him for our regular morning sessions, which led to several joint publications.

Last and not least, I very much want to thank Johanna and our daughter Helena for the beautiful time besides the academic world.

Thank you all,

Veit

Contents

Introduction and Overview

The goal of this dissertation is to study the various connections between the dimension of posets and graph theoretic properties of their cover graphs.

The *dimension* of a poset P is the least number of linear extensions of P whose intersection gives rise to P. This parameter was introduced in 1941 by Dushnik and Miller [17] and turned out to be one of the most important concepts when studying the combinatorics of posets.

Computing the dimension is a classical computationally hard problem. The decision problem whether the dimension of a poset is at most k, for some fixed k, was posed as one of the twelve problems on Garey's and Johnson's famous list of problems with unknown complexity in 1979. While it is decidable in polynomial time whether a poset has dimension at most 2, the problem becomes NP-complete once k is at least 3, as shown by Yannakakis in 1982 [77]. The problem even remains NP-complete if only posets of height 2 are under consideration (Felsner et al. [25]). Giving good approximations is hard as well: Chalermsook et al. [10] argued that it is unlikely that there exists a polynomial-time algorithm that approximates the dimension within a factor of $\mathcal{O}(n^{1-\epsilon})$, where n denotes the number of elements in the poset.

Despite all these hardness results, we want to understand this parameter and try to find sufficient conditions and witnesses of large-dimensional posets. On the other hand, we want to determine properties that are *necessarily* satisfied by them. A number of such properties are known; for instance, Hiraguchi showed that every n-dimensional poset ($n \geqslant 3$) has to contain at least $2n$ elements [36]. Moreover, when n is large, then some of the poset elements must be involved in many comparabilities (Rödl, Trotter [67]; Füredi, Kahn [30]). Another such property is a simple consequence of a classic theorem of Dilworth [15]: the width of a poset, which is the maximum size of a contained antichain, is at least as large as the dimension of the poset. This fact can be rephrased as follows: Large-dimensional posets have to be *wide*. The situation is different for the *height* of a poset, which is maximum size of a chain in the poset. The so-called *standard example* of order n has dimension n, while its height is only 2; Figure 1 (left) illustrates such an example of order 4. Therefore, large-dimensional posets need not be *tall*.

In the following we discuss structural graph properties that are necessarily satisfied in cover graphs of posets with large dimension.

Cover Graphs. Posets are usually visualized by their *Hasse diagrams*, in which every cover relation is represented by a y-monotone curve. The abstract undirected graph behind the Hasse diagram is the *cover graph* of the poset. Another graph that we associate with a poset is its *comparability graph*, where two elements are joined by an edge if they are comparable. Those two graphs behave differently in

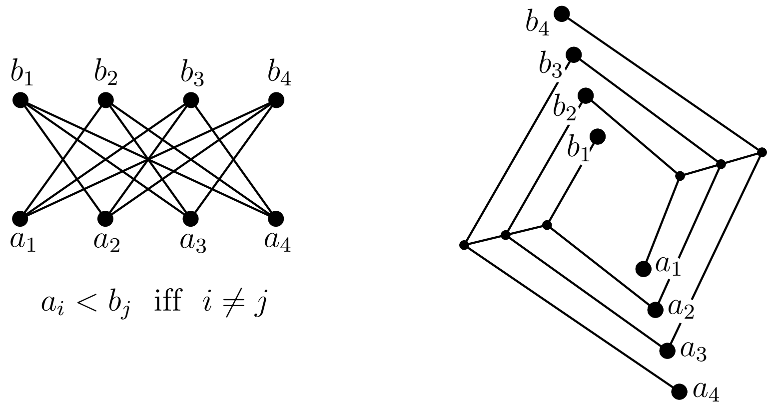

Figure 1: Standard example (left), and Kelly's example containing it as an induced subposet (right).

many aspects: While width, height, and also dimension are invariant on posets with the same comparability graph, this is not true for cover graphs. The dimension of posets with the same cover graph may even differ by an arbitrary number. However, there are structural properties of graphs G that yield valuable information about the dimension of (all) posets whose cover graphs are isomorphic to G. Another fundamental difference between these two graph concepts is given by the computational complexity of their corresponding decision problems: It is well known that deciding whether a graph is a comparability graph can be done in polynomial time. On the other hand, testing whether a graph is a cover graph is an NP-complete problem as shown by Nešetřil and Rödl [58], and Brightwell [9].

Why studying cover graphs? First of all, these are the ones that we use to visualize posets and hence they encode the topology of posets. Another important reason to consider cover graphs is that they are usually sparser than comparability graphs, and this sparsity may admit the existence of an underlying topological structure that can be used to give good estimations on the dimension.

An example where such a structure can be exploited is given by the class of planar posets. A poset is *planar* if its Hasse diagram can be drawn without any edge crossings. In 1972, Baker, Fishburn, and Roberts proved that planar posets with a single minimal element and a single maximal element have dimension at most 2 [2]. The first result in this context that is using only the structural properties of cover graphs is due to Moore and Trotter; they show that posets with trees as cover graphs have dimension at most 3 [75].

Dimension versus Height. These early results led to the question whether the dimension of planar posets is bounded by a constant. It was answered in the negative by Kelly in 1981, who constructed an explicit family of planar posets that contain arbitrarily large standard examples and hence have unbounded dimension; see Figure 1 (right). After Kelly's discovery, researchers in the field lost interest in studying connections between the topological structure of cover graphs and dimension for decades. Probably, the main reason here was that Kelly's result broke a possible analogy between order dimension and chromatic number: While planar graphs have chromatic number at most 4, there is no such constant bound on the dimension of planar posets. This interest got renewed only after Felsner, Li, and Trotter [24] discovered in 2010 that the dimension of height-2 posets with planar cover graphs is at most 4. The question arose whether such a constant bound also holds when the height is bounded by values larger than 2. In other words, do large-dimensional posets with planar cover graphs have to be *tall*?

In 2014, this question was answered in the affirmative by Streib and Trotter [71], who set the starting point for a series of papers in this direction with their contribution. Joret et al. [38] continued this line of research and showed that bounded height posets whose cover graphs have bounded tree-width also have bounded dimension. This was further generalized by Walczak [76] to the case that cover graphs exclude a fixed graph as a topological minor.

Bounded degree graphs, planar graphs, graphs of bounded tree-width, and more generally graphs excluding a fixed graph as a topological minor have in common that their degeneracy is bounded and hence contain only linear many edges. The following well-known construction shows that one cannot extend Walczak's result to 2-degenerated cover graphs: The inclusion order on the 1-element and 2-element subsets of $\{1, \ldots, n\}$ is a height-2 poset with dimension at least $\log \log n$, as already shown by Dushnik and Miller [17]. The cover graph of this example though is isomorphic to the 1-subdivision of the complete graph on n vertices, which is 2-degenerated. At first glance, this observation seems to close the subject on the "dimension versus height" problem. However, as we explain now there is another natural question arising at this point.

Over the last decade, Nešetřil and Ossona de Mendez [55] carried out a thorough study of sparse graph classes. They particularly introduced graph classes with *bounded expansion*, which is a general notion for uniform sparseness in graphs. Those classes include many other natural sparse graph classes, among them graphs of bounded degree, graphs with bounded tree-width, and more generally graphs that exclude a fixed graph as a topological minor. Moreover, they properly extend the other classes as there exist bounded expansion classes that contain every graph as a topological minor (e.g. k-planar graphs).

Over the years, it turned out that the concept of bounded expansion is a very robust notion with many seemingly unrelated characterizations [56, 62]. As a consequence, bounded expansion classes form the natural boundary in many different aspects. Therefore, let us ask the obvious question whether order dimension yields yet another aspect here: Do bounded expansion classes form the limit in the hierarchy of sparse graphs until we can guarantee a bounding function on the dimension as discussed above?

Aim of the Thesis and Results. Continuing the line of research being shaped by the aforementioned results, it is the aim of this thesis to give precise answers to the following general two questions:

1) *Do large-dimensional posets with somewhat sparse cover graphs have to be tall?*
2) *And if so, how tall?*

Using the notions developed by Nešetřil and Ossona de Mendez, we determine the exact type of sparsity in cover graphs for which the answer to the first question is 'yes'. More precisely, we prove the following theorem, which is joint work with Gwenaël Joret and Piotr Micek.

Theorem ([42])**.** *Posets of bounded height whose cover graphs belong to some fixed class with bounded expansion have bounded dimension.*

The proof of this theorem uses *weak coloring numbers*, which were introduced by Kierstead and Yang [45]. They can be used to characterize bounded expansion classes [78] and have been proven to be helpful in many different settings. As a consequence, our proof does not rely on topological arguments that can be derived from structure theorems, but rather on combinatorial density arguments.

We also show that the above theorem cannot be pushed any further by requiring the cover graphs to be *nowhere dense*, which is an even more general concept of uniform sparseness in graphs than bounded expansion. Therefore, we can give further support for the robustness of this sparsity concept by establishing a surprising connection to order dimension.

A big part of the research contained in this dissertation deals with the second general question from above. That is to say, given a concrete sparsity property of graphs (such as being planar, having bounded tree-width, et cetera), we want to estimate the maximal dimension of posets whose cover graphs have this property by a function in height. Our main contribution in this context is the linear bound on the dimension of planar posets. It is an improvement upon exponential bounds and is best possible up to a constant factor as witnessed by Kelly's examples.

The illustration in Figure 2 summarizes the results contained in this thesis. It particularly shows the current best lower and upper bounds on the maximal dimension of posets of height at most h with cover graphs in the respective graph classes. Bounds drawn within a frame that has a star at one of its corners are proven in this thesis. Let us also remark that the asymptotic bounds are calculated with respect to the poset height h, and hence other parameters like g in the bounded genus case are treated as constants (so they might vanish in the $\mathcal{O}$-notation).

Overview. Let us continue with an overview of the single chapters and the results proven therein. For background material and motivation of these results, we refer the reader to the introductory part of the individual chapters.

In Chapter 1 we start with introducing elementary poset terminology and notations. This includes concepts like dimension, alternating cycles, and cover graphs. Moreover, we cover the needed graph theoretic notions that might not be contained in a basic course on graph theory.

In Chapter 2 we present some essential techniques and tools that are applied numerous times throughout the thesis. For example, we explain the method of unfolding a poset in a breadth-first search manner, which was first introduced by Streib and Trotter [71]. Furthermore, we describe several distinct applications of poset unfoldings which will be applied in the respective chapters.

In Chapter 3 we prove the aforementioned linear upper bound on the dimension of planar posets. This result is complemented by a construction of a new family of planar posets with unbounded dimension. It yields a better lower bound on the maximal dimension of planar posets than Kelly's construction.

We continue with posets whose cover graphs have bounded tree-depth, path-width, or tree-width in Chapter 4. We start to consider the cases where we can achieve constant upper bounds on the dimension, meaning that they are independent of the poset's height. This holds when the cover graphs have bounded tree-depth or path-width at most 2. Then, we prove a polynomial upper bound on the dimension in the bounded path-width case. This contrasts the situation when the cover graph only has bounded tree-width as in this case the dimension can be exponential. We conclude the chapter by several lower bound constructions witnessing that the proved upper bounds of the chapter are essentially best possible.

In Chapter 5 we give a new proof of Walczak's result. Our argument is entirely combinatorial and avoids sophisticated tools such as graph structural decomposition theorems. We develop a new approach for establishing upper bounds on the dimension of posets. The key idea is to iteratively unfold large-dimensional posets and to use the initial parts of these unfoldings to find a large topological minor in the cover graphs. We can also use this approach for $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude a fixed graph as a topological minor and show that they do not contain arbitrarily large standard examples.

In Chapter 6 we deal with the broadest class of cover graphs in the thesis. We prove the aforementioned upper bound on the dimension of bounded height posets whose cover graphs belong to some fixed class with bounded expansion. Moreover, we show that such a result cannot hold when cover graphs are required to be nowhere dense. We finish Chapter 6 by applying the developed ideas to a seemingly unrelated problem. We show that graphs with no K_t-minor have boxicity $\mathcal{O}(t^2 \log t)$.

The dissertation ends with our conclusions; we discuss new and exciting open problems that are directly related to our results.

Let us add some remarks about the ordering of the chapters. Chapters 4, 3, 5, and 6 can be read independently. Their order is chosen according to the hierarchy on the sparsity properties studied in the respective cases. The tools of Chapter 2 are applied in Chapters 4, 3, and 5. Therefore, the reader may skip Chapters 2–5 and directly jump to Chapter 6.

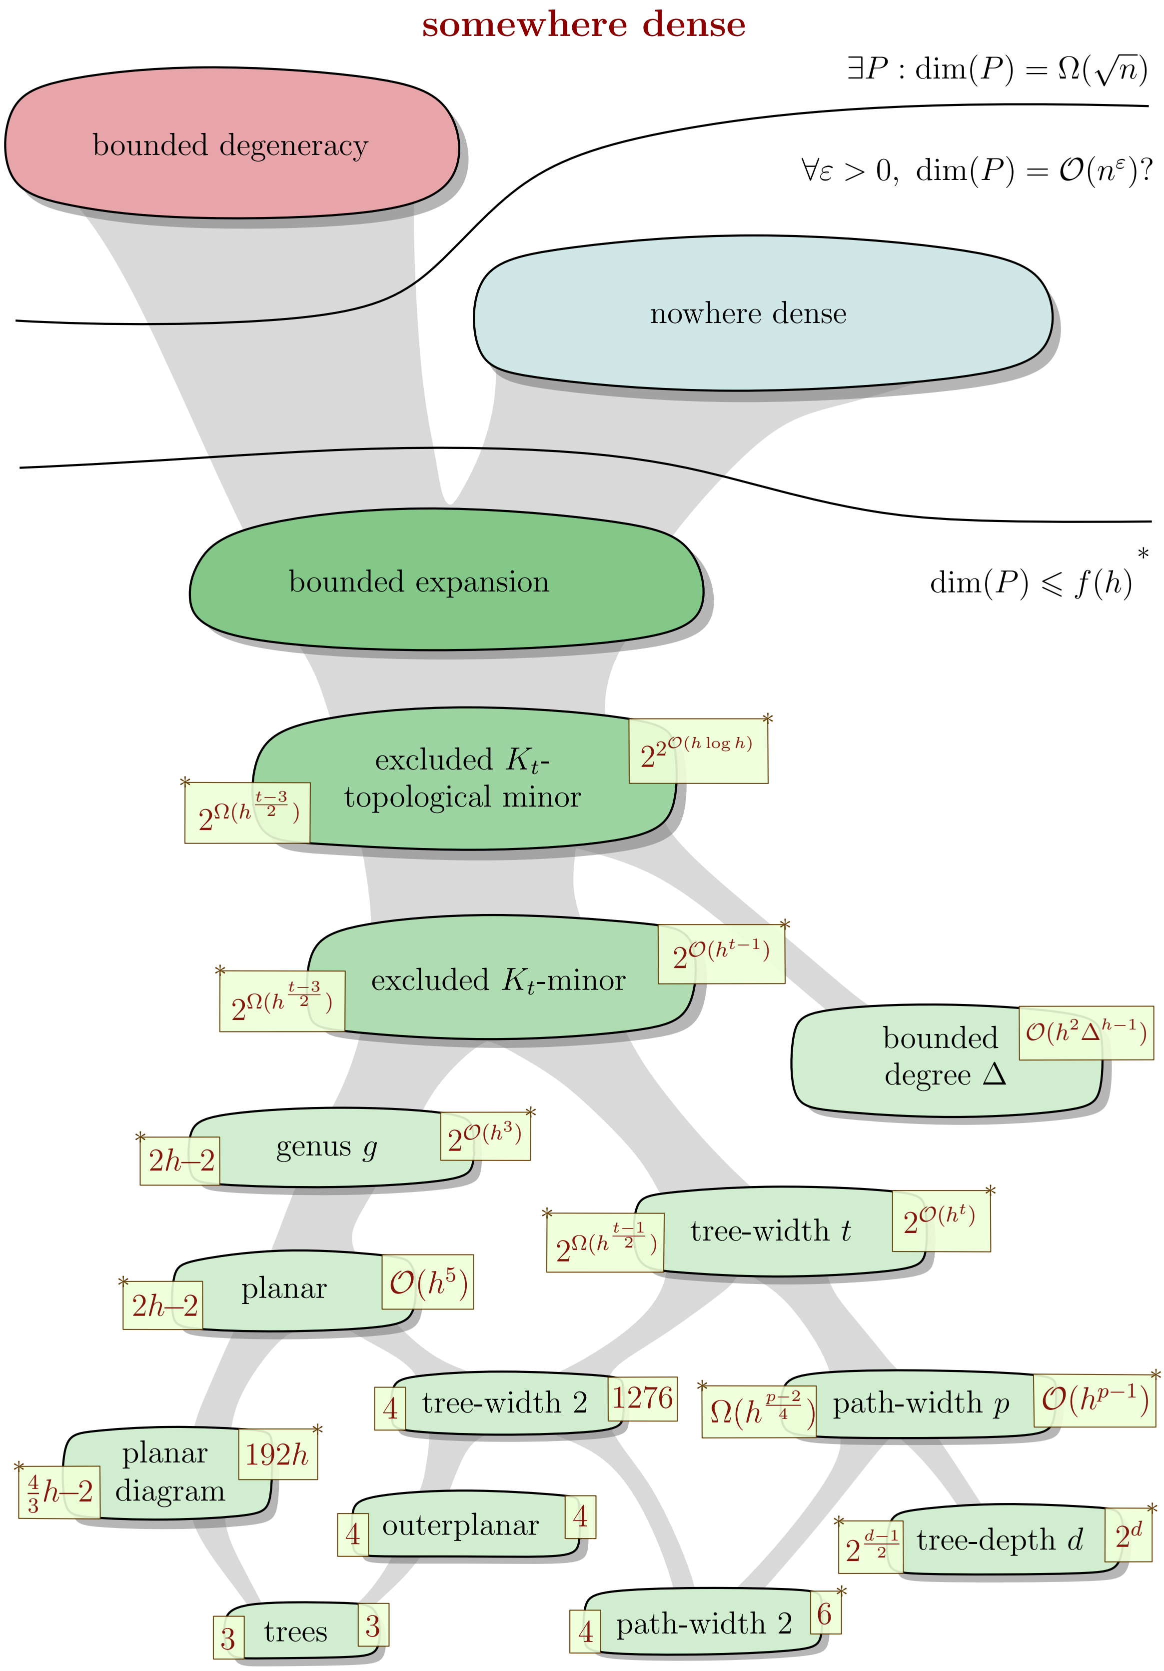

Figure 2: Dimension of posets of height at most h with cover graphs in the respective graph classes.

Chapter 1

Preliminaries

The purpose of this chapter is to introduce elementary concepts that are related to the topic of this thesis. We begin with a concise overview of the basic notions for partially ordered sets, which we simply call *posets*. Then, we introduce more carefully the concept of the dimension of a poset. In particular, some standard observations from dimension theory will be stated and explained. We conclude the chapter with some basics about graph minors.

Let us quickly present some abbreviations that we use throughout the thesis. If $n \geqslant 1$ is an integer, then we denote the set $\{1, \ldots, n\}$ by $[n]$. In the same spirit, for integers n, m such that $m \leqslant n$, we let $[m, n]$ denote the set $\{m, m+1, \ldots, n\}$. The set of subsets of $[n]$ is denoted by $2^{[n]}$.

1.1 Posets, Cover Graphs, and Dimension

All posets in this thesis are finite. A poset $P = (X, \leqslant_P)$ is a pair consisting of a *ground set* X and a binary relation $\leqslant_P$ on X that is reflexive, antisymmetric, and transitive. Whenever applicable, we avoid to specify a symbol for the ground set of P and simply write '$x \in P$' for $x \in X$. Similarly, we do not always provide a symbol for P's relation; in these cases we write '$x \leqslant y$ in P' to express that $x \leqslant_P y$. If we want to emphasize that a comparability $x \leqslant_P y$ involves two distinct elements x and y, then we write '$x <_P y$' or '$x < y$ in P'.

We say that elements $x, y \in X$ are *comparable* in P if $x \leqslant_P y$ or $y \leqslant_P x$ holds. Otherwise, x and y are *incomparable* in P. A set C of elements in P is called a *chain* in P if the elements of C are pairwise comparable in P. If the set of all elements in P forms a chain, then we also say that P itself is a chain. A set A of elements of P is called an *antichain* in P if the elements of A are pairwise incomparable in P. In the case where any two distinct elements of P are incomparable, we also say that P itself is an antichain.

The *height* of P, which we denote by $\text{height}(P)$, is defined to be the maximum size of a chain in P. Similarly, the *width* of P is defined as the maximum size of an antichain in P, and denoted by $\text{width}(P)$. An element $x \in X$ is *minimal* in P if there is no $y \in X$ such that $y <_P x$. Dually, $x \in X$ is *maximal* in P if there is no $y \in X$ such that $x <_P y$. We denote the set of minimal elements in P by $\text{Min}(P)$ and the set of maximal elements in P by $\text{Max}(P)$.

Given a subset $Y \subseteq X$ of elements in P, the *upset* $\text{U}_P(Y)$ of Y in P is the set of all elements $x \in X$ for which there is $y \in Y$ such that $y \leqslant_P x$. Similarly, we define

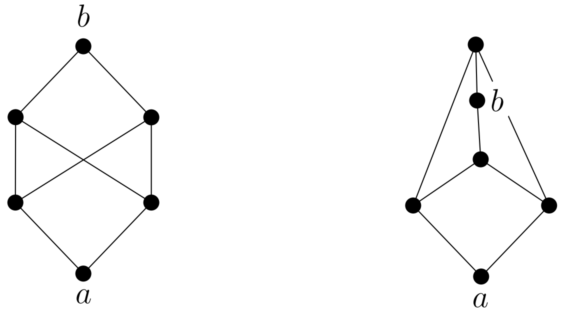

Figure 1.1: The diagram of a non-planar poset (left) and a planar embedding of its cover graph (right).

the *downset* $\mathrm{D}_P(Y)$ of Y in P as the set of all elements $x \in X$ for which there is $y \in Y$ such that $x \leqslant_P y$. In the case of one-element sets we use the abbreviations $\mathrm{U}_P(y)$ and $\mathrm{D}_P(y)$ for $\mathrm{U}_P(\{y\})$ and $\mathrm{D}_P(\{y\})$, respectively. Whenever the poset P is clear from the context, we may drop the subscript 'P'.

A poset $Q = (Y, \leqslant_Q)$ is an *induced subposet* of $P = (X, \leqslant_P)$, if $Y \subseteq X$ and for any $x, y \in Y$ it holds that $x \leqslant_Q y$ if and only if $x \leqslant_P y$. Note that Q can also be seen as the restriction of P to the elements of Y, which we also denote by $P|_Y$. If P does not contain Q as an induced subposet, then P is *Q-free*.

For readers not familiar with these introduced notions, we recommend the book *Ordered Sets* of Bernd Schröder [70], which serves as a standard reference in order theory.

Cover Graphs. Let $P = (X, \leqslant_P)$ be a poset. We say that $x <_P y$ is a *cover relation* in P if there is no $z \in X$ such that $x <_P z <_P y$ (so $x <_P y$ cannot be deduced by transitivity). In this case, element x *is covered by* element y in P. The *cover graph* of P, which we denote by $\mathrm{cover}(P)$, is the graph with vertex set X and with an edge between any distinct vertices x, y if and only if $x <_P y$ or $y <_P x$ is a cover relation in P. The poset P is *connected* if the cover graph of P is connected. An induced subposet Q of P is a *component* of P if the elements of Q build a component in the cover graph of P.

A *covering chain* in P is a sequence $x_1, \ldots, x_k$ of elements of P such that $x_i <_P x_{i+1}$ is a cover relation in P for each $i \in [k-1]$. Note that in this case $x_1, \ldots, x_k$ is also a path in the cover graph of P. Therefore, we also say that this sequence is a covering chain *of* $\mathrm{cover}(P)$.

The usual way to visualize P is to draw its *Hasse diagram* (short: *diagram*): Elements are drawn as distinct points in the plane and each cover relation $a < b$ in P is represented by a curve from a to b going upwards. That is to say, the diagram is a drawing of P's cover graphs with a special monotonicity property on the curves representing edges. If the diagram of P can be drawn in a planar way, then P is said to be a *planar poset*. Note that planar posets have planar cover graphs. However, the converse is not necessarily true: For every height $h \geqslant 3$, there exists a non-planar poset with a planar cover graph; see Figure 1.1 for an example.

Convex Subposets and Convex Hulls. We introduce now a quite recent notion for posets that will be useful in many places of the thesis. Let $P = (X, \leqslant)$ be a poset. We say that an induced subposet Q of P is *convex*, if $x, y \in Q$ and $x \leqslant z \leqslant y$ in P imply that z is an element of Q. Given a subset $Y \subseteq X$, the *convex hull*

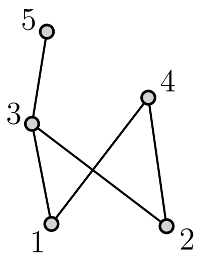

$1 < 2 < 3 < 4 < 5$
$1 < 2 < 4 < 3 < 5$
$1 < 2 < 3 < 5 < 4$
$2 < 1 < 3 < 4 < 5$
$2 < 1 < 4 < 3 < 5$
$2 < 1 < 3 < 5 < 4$

Figure 1.2: A poset on five elements and its six linear extensions. The second and the six-th linear extension form a realizer of the poset.

of Y in P is the subposet of P consisting of all points $x \in X$ for which there are $y, z \in Y$ such that $y \leqslant x \leqslant z$ in P. We denote this subposet by $\mathrm{conv}_P(Y)$ (we will drop the subscript if the poset P is clear from the context). Note that $\mathrm{conv}_P(Y)$ is convex and that it is equivalently defined by taking the smallest convex subposet of P containing Y.

We state the following important property of convex subposets for emphasis.

Proposition. *Let Q be a convex subposet of P. Then the cover graph of Q is an induced subgraph of the cover graph of P.*

Linear Extensions and Dimension. Let $P = (X, \leqslant_P)$ be a poset. We denote by $\mathrm{Inc}(P)$ the set of ordered pairs (x, y) such that x and y are incomparable in P. If Y, Z are subsets of the ground set X, then $\mathrm{Inc}(Y, Z)$ denotes the set of pairs $(y, z) \in \mathrm{Inc}(P)$ such that $y \in Y$ and $z \in Z$.

A linear order L on the elements of X is a *linear extension* of P, if for any $x, y \in X$ we have that $x \leqslant_P y$ implies $x \leqslant y$ in L. A set of linear extensions $\{L_1, \ldots, L_t\}$ of P is a *realizer* of P if for any $x, y \in X$ we have that $x \leqslant y$ in P if and only if $x \leqslant y$ in L_i for each $i \in [t]$; in other words, the intersection of $L_1, \ldots, L_i$ gives rise to P; see Figure 1.2 for an example.

Note that the set of all linear extensions of P forms a realizer of P; this follows from the basic fact that for each incomparable pair (x, y) of P there are linear extensions L, L' of P such that $x < y$ in L while $y < x$ in L'. Observe also that a realizer can be seen as a way to encode a poset in an efficient way (provided the realizer is small).

This all motivates the following definition, which can be seen as a measure for the complexity of a poset: The *dimension* $\dim(P)$ of P is the least number t such that P has a realizer of size t. For example, a poset is 1-dimensional if and only if it is a chain.

Dushnik and Miller introduced the notion of a poset's dimension, but, to be honest, it is not very intuitive at first sight. Therefore, let us have a look at a geometric interpretation of dimension. Equivalently, and sometimes called the *Ore-dimension* of a poset, $\dim(P)$ is the least d such that there is an order-preserving embedding of P into $\mathbb{R}^d$, where $\mathbb{R}^d$ is equipped with the product order $\leqslant_d$ (that is, $x \leqslant_d y$ if and only if $x_i \leqslant y_i$ for all $i \in [d]$). More precisely, it is the least d such that P is an induced subposet of $(\mathbb{R}^d, \leqslant_d)$. Throughout the paper we will not work with the geometric version, even though there are examples where it is useful; see for instance Figure 1.3 illustrating an order-preserving embedding of the *Boolean lattice* of order three (which we will define in a moment) into $\mathbb{R}^3$.

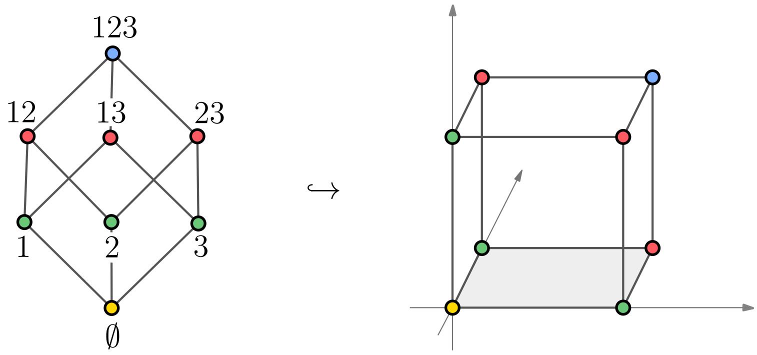

Figure 1.3: An order-preserving embedding of the Boolean lattice of order three into $\mathbb{R}^3$ using the vertices of the cube.

Let us now discuss some standard facts about the dimension of posets. Dimension is a monotone property, in the sense that if Q is an induced subposet of P then $\dim(Q) \leqslant \dim(P)$. Indeed, any realizer of P restricted to the elements of Q yields a realizer of Q. Another simple observation is the following one.

Proposition. *Every poset P with n elements has dimension at most n.*

Proof. For each element x in P, we define a linear extension L_x as follows. We let L'_x be an arbitrary linear extension of the downset $\mathrm{D}(x)$, and we let L''_x be an arbitrary linear extension of the induced subposet of P consisting of all elements in P that are not contained in the downset $\mathrm{D}(x)$. Then we set $L_x := L'_x < L''_x$, so L_x can be seen as the concatenation of these two linear extensions. Observe that L_x is indeed a linear extension of P.

We claim that $\mathcal{R} = \{L_x \mid x \in P\}$ is a realizer of P. Clearly, if $y \leqslant z$ in P then we also have $y \leqslant z$ in L_x for all x in P by the definition of a linear extension. It remains to consider the case that y and z are incomparable elements of P. Then we observe that $y < z$ in L_y and $z < y$ in L_z. Thus, y and z are also incomparable in the intersection of L_y and L_z. This shows that $\mathcal{R}$ is a realizer of P and implies $\dim(P) \leqslant n$. □

Hiraguchi [36] proved a stronger statement: for all $n \geqslant 4$, posets with n elements have dimension at most $n/2$. His inequality cannot be strengthened as is witnessed by the so-called standard examples. The *standard example* S_n of order n is the height-2 poset with n minimal elements $a_1, \ldots, a_n$ and n maximal elements $b_1, \ldots, b_n$, such that for any $i, j \in [n]$ we have $a_i < b_j$ in S_n if and only if $i \neq j$; see again Figure 1 (left) from the introduction.

Proposition (Dushnik and Miller [17])**.** *For all $n \geqslant 2$, the standard example S_n has dimension n.*

Proof. We only show $\dim(S_n) \geqslant n$ as the other direction is straightforward. Suppose that $\mathcal{R}$ is a realizer of S_n and let $i \in [n]$. Then there exists a linear extension $L \in \mathcal{R}$ such that $b_i < a_i$ in L. As L is a linear extension of S_n this even implies $a_j < b_i < a_i < b_j$ in L for every $j \in [n] \setminus \{i\}$. Therefore, each linear extension of the realizer 'reverses' at most one incomparable pair of the form (a_i, b_i). This implies $|\mathcal{R}| \geqslant n$. □

Standard examples play a key role in the lower bound constructions of this thesis. In a sense they are the "cliques" for the dimension problem.

A famous poset supporting that dimension is a natural parameter is given by the Boolean lattice. The Boolean lattice of order n, denoted by B_n, is the poset in which the subsets of $[n]$ are ordered by inclusion. The embedding illustrated in Figure 1.3 shows that the dimension of B_n is at most n. Since the 1-element and $(n-1)$-element sets of B_n induce the standard example S_n, we deduce that the dimension of B_n is exactly n.

Reversible Sets and Alternating Cycles. In many proofs of the thesis we work with the observation that the dimension problem can be rephrased as a hyper graph coloring problem. Concepts like *reversible sets* and *alternating cycles* play an important role for this characterization and hence also for our proofs. We explain them now.

We say that a linear extension L of P *reverses* an incomparable pair $(x, y) \in \mathrm{Inc}(P)$ if we have $y < x$ in L. Note for example that if $\mathcal{R}$ is a realizer of P, then for each $(x, y) \in \mathrm{Inc}(P)$ there must be linear extension $L \in \mathcal{R}$ that is reversing (x, y). Given a set $I \subseteq \mathrm{Inc}(P)$, we call I *reversible* if there exists a linear extension of P that reverses each pair of I.

A standard obstruction for a set to be reversible is the following structure: An *alternating cycle* of P is a sequence $C = (x_1, y_1), \ldots, (x_k, y_k)$ of incomparable pairs of P such that $k \geqslant 2$ and $x_i \leqslant_P y_{i+1}$ holds cyclically for each $i \in [k]$ (meaning that we also have $y_k \leqslant_P x_1$). We call C a *strict alternating cycle* if for any $i, j \in [k]$, we have that $x_i \leqslant_P y_j$ holds if and only if $j = i + 1$. Note that if C is strict then the elements $x_1, \ldots, x_k$ are pairwise distinct and form an antichain, and this holds analogously for $y_1, \ldots, y_k$.

It is easy to see that the pairs of an alternating cycles $(x_1, y_1), \ldots, (x_k, y_k)$ are not reversible. Indeed, otherwise there exists a linear extension L reversing all of those pairs, implying that $y_i < x_i \leqslant y_{i+1}$ holds in L for all $i \in [k]$ (cyclically), which is not possible in a linear order. The converse is also true:

Proposition (Moore, Trotter [75]). *A set $I \subseteq \mathrm{Inc}(P)$ is reversible if and only if I contains no alternating cycle.*

Proof. We have already seen the forward direction. So let us prove the backward implication now. We need to show that if $I \in \mathrm{Inc}(P)$ is not reversible, then I contains an alternating cycle.

Let H be the directed graph on the ground set of P, where we put an edge from y to x if we either have $y < x$ in P or $(x, y) \in I$. This definition suggests a partition of the edges in H into P-edges and I-edges.

First we observe that H has to contain a directed cycle: Otherwise, any topological ordering of the vertices corresponds to a linear extension of P that is reversing all pairs of I, which is not possible as I is not reversible by assumption. Let C be such a directed cycle. Clearly, C has to contain an I-edge as the poset P is acyclic. In fact, it has to contain at least two I-edges, since otherwise the two endpoints of the unique I-edge would be at the same time incomparable in P (by the definition of an I-edge) and comparable in P (by transitivity of P and the directed path of P-edges in C). So let $(y_1, x_1), \ldots, (y_k, x_k)$ be the I-edges of C appearing in that cyclic order on C. For each $i \in [k]$ there is a directed path from x_i to y_{i+1} in C

consisting of P-edges (this path possibly contains only one vertex, in which case $x_i = y_{i+1}$), which implies $x_i \leqslant y_{i+1}$ in P. We conclude that $(x_1, y_1), \dots, (x_k, y_k)$ is an alternating cycle in I. This completes the proof of the backward implication. □

Let us note at this point that the above proposition remains true if we rephrase it with strict alternating cycles. This proposition is standard in dimension theory and we will also use it for almost all of our proofs that are establishing upper bounds on the dimension. Therefore, we will apply it without referencing in what follows.

Consider now the following hyper graph $\mathcal{H}_P$ associated with a poset P. The vertex set of $\mathcal{H}_P$ is $\mathrm{Inc}(P)$, and edges of $\mathcal{H}_P$ correspond to sets of incomparable pairs forming an alternating cycle of P. Then the chromatic number of $\mathcal{H}_P$ is equal to the dimension of P. To see this, observe first that each color class in a proper coloring of $\mathcal{H}_P$ forms a reversible set by the proposition above, implying that $\dim(P) \leqslant \chi(\mathcal{H}_P)$. Conversely, the set of incomparable pairs being reversed by a linear extension in a realizer of P forms a reversible set by definition. Therefore, a realizer of P naturally induces a proper coloring of $\mathcal{H}_P$, which yields $\dim(P) \geqslant \chi(\mathcal{H}_P)$. This shows $\dim(P) = \chi(\mathcal{H}_P)$ and turns the dimension problem into a coloring problem. We state this observation in a slightly different form again for emphasis.

Proposition. *For every P that is not a chain, we have that the dimension of P is the least number t for which there are reversible sets $I_1, \dots, I_t$ such that* $\mathrm{Inc}(P) = I_1 \cup \cdots \cup I_t$.

Motivated by this proposition we make the following definitions. Given a non-empty set $I \subseteq \mathrm{Inc}(P)$ of incomparable pairs, we let $\dim_P(I)$ denote the least integer t such that I can be partitioned into t reversible sets. In the special case of I being the empty set, we use the convention that $\dim_P(I)$ is equal to 1. Observe that our last proposition implies that $\dim(P)$ and $\dim_P(\mathrm{Inc}(P))$ are equal for all posets P. For subsets X, Y of elements of P, we use the abbreviation $\dim_P(X, Y)$ to denote the value of $\dim_P(\mathrm{Inc}(X, Y))$. We may again drop the subscript 'P' from the newly introduced notions if the poset is clear from the context.

Let us emphasize here once more the important fact that the dimension problem is equivalent to partitioning incomparable pairs into reversible sets. We use this characterization throughout the thesis as we usually apply the following strategy to bound the dimension of a poset from above: First, we partition the incomparable pairs of P into a 'small' number of sets, and second, we show that all these sets are reversible by proving that they contain no alternating cycle.

We have been discussing the dimension of posets and various related concepts in this section. However, the material presented here still does not serve as a careful introduction to this topic. For further examples and explanation regarding dimension, the reader may consult Trotter's book [74].

1.2 Graph Minors

All graphs in this thesis are finite, simple, and undirected. We assume the reader to be familiar with basic notions from graph theory, such as *cliques, independent sets, paths, cycles, degeneracy, chromatic number*. (As a standard textbook introducing

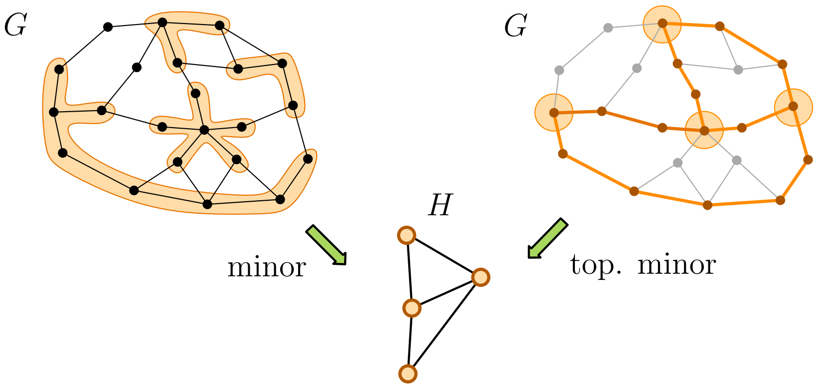

Figure 1.4: Graph H is a minor of G (left) and also a topological minor of G (right). Vertices covered by disks on the right-hand side are the principal vertices of the H-subdivision.

these notions we recommend Diestel's book [14].) Since this dissertation deals with posets whose cover graphs have certain graph structural properties which might not be part of basic graph theory course, we will introduce notions like *graph minor, path-width, tree-width, bounded expansion classes*, et cetera. Most of these parameters and concepts are defined in the chapters where they are needed for the first time. In this section we define (topological) graph minors as those play an important role throughout the thesis.

Before we do so, let us introduce some notation and abbreviations. Given a graph G, we denote by the $V(G)$ and $E(G)$ the vertex set and the edge set of G, respectively. For a subset $X \subseteq V(G)$ we denote by $G[X]$ the subgraph of G induced by vertices in X. The complete graph on n vertices is denoted by K_n. For positive integers n, m we let $K_{n,m}$ denote the complete bipartite graph with n vertices on one side and m vertices on the other side.

A graph H is a *minor* of G if H can be obtained from G by applying a sequence of the following three operations: removing a vertex, removing an edge, and contracting an edge. Equivalently, H is a minor of G if H can be obtained by first specifying a set of disjoint connected subgraphs of G, then contracting those subgraphs to a single vertex each, and finally by possibly removing vertices and edges of the resulting graph. The graph G *excludes* H as a minor if H is not a minor of G; in this case we also say that G is H-*minor-free*.

A class of graphs $\mathcal{C}$ is *minor-closed* if every minor of a graph in $\mathcal{C}$ also belongs to $\mathcal{C}$. We call a graph class *proper* if it does not contain all finite graphs. Clearly, every proper minor-closed graph class excludes some graph as a minor (and hence an infinite list of finite graphs as well). Surprisingly, such graph classes can be characterized by a *finite* list of excluded minors as implied by the *graph minor theorem* that was proven by Robertson and Seymour [65] in their fundamental series of papers on *Graph Minors*. More precisely, if $\mathcal{C}$ is a proper minor-closed graph class, then there exist graphs $H_1, \ldots, H_m$ such that

$$\mathcal{C} = \{G \mid G \text{ excludes } H_i \text{ as a minor, for all } i \in [m]\}.$$

For instance, planar graphs form a minor-closed graph class since contracting an edge in a planar graph does not destroy planarity, and this class is proper as the complete graph on five vertices is not planar. By the previously mentioned *graph* *minor theorem* there exists a forbidden minor characterization of planar graphs, and this is given by Wagner's theorem that a graph is planar if and only if it does not contain K_5 or $K_{3,3}$ as a minor.

We aim to define topological minors now. By *subdividing an edge* e in a graph G we mean the operation of replacing e by a path on three vertices that has the same endpoints as e. Informally, this can be seen as if we "subdivide" the edge e in a drawing of G by putting a new vertex in the middle of e. A graph G' is a *subdivision* of G if G' can be obtained from G by a sequence of edge subdivisions; in this case we say that the vertices of G' that we associate with the vertices of G are *principal* in G'. Then we call H a *topological minor* of G if G contains a subgraph that is isomorphic to a subdivision of H. Figure 1.4 illustrates the two graph minor notions.

Observe that if H is a topological minor of G, then H is also an ordinary minor of G. Indeed, this follows from the fact that any subdivision of H can be turned into H by contracting along subdivided edges.

The *r-subdivision* of a graph G is the graph that we obtain if we replace all the edges in G by internally disjoint paths of length at most $r + 1$. (We note at this point that we follow the convention that the *length* of a path is given by the number of its edges.) Moreover, an $\leqslant r$-subdivision of G is a graph that can be obtained by replacing the edges in G by internally disjoint paths of length at most $r + 1$.

Chapter 2

Essential Tools for Dimension

Many proofs that are establishing upper bounds on the dimension of posets start with somewhat standard reductions and make use of known tools. In this chapter we introduce those reductions and tools that are necessary for the type of problems we solve in the dissertation.

A classic example in this context is the reduction of the dimension problem to the 'min-max dimension' of a poset. Once this trick has been applied, it will be enough to find small realizers that are reversing incomparable pairs consisting of a minimal element and a maximal element. We discuss this *Min-Max Reduction* in Section 2.1.

In Section 2.2 we present the concept of unfolding a poset. It is a relatively new concept and was first introduced by Streib and Trotter [71]; subsequently it was used in several other works [39, 51, 42]. It has already been proven to be helpful and we give further support for its importance by applying it a number of times in the thesis.

Unfolding a poset can be seen as a performing a *breadth first search* (BFS) on the minimal and maximal elements of a poset. Similar to the fact that there is BFS layer witnessing at least half of the chromatic number of a graph, there is 'min-max layer' of the unfolding that is witnessing at least half of the dimension of a poset. An important new idea developed in this thesis is to use unfoldings in an iterative manner: first unfold the poset, and then proceed with unfolding the 'heavy' layer. For instance, we use this idea to find large clique minors in cover graphs of large-dimensional posets with bounded height in Chapter 5.

We also introduce two direct applications of the *Unfolding Lemma* in Section 2.2: the *Global Min Support Reduction* and the *Extended Unfolding Lemma*. In the first one we win an element in the poset that lies above every minimal element. It is applicable whenever cover graphs belong to some minor-closed graph class, which will be the case in Chapter 4. The second application is similar to the first one, with the difference that it is designed to find large cliques as a topological minor in cover graphs (see Chapter 5).

We conclude this chapter with Section 2.3, where we introduce *zig-zag paths* with respect to some fixed unfolding of a poset. Those paths connect elements of the 'heavy' layer with the element from which the unfolding was started. Moreover, their union forms a tree. This tree structure, somewhat surrounding the poset outside the 'heavy layer', will be used in Chapter 3 to bound the dimension of planar posets.

Let us remark that it is not necessary to go through all the tools of this chapter before reading specific proofs of subsequent chapters. For instance, none of the tools is used in Chapter 6. In Chapter 3, we only apply the *Unfolding Lemma* and make use of *zig-zag paths*. And in Chapters 4 and 5, we need the *Global Min Support Reduction* and the *Extended Unfolding Lemma*, respectively.

2.1 Min-Max Reduction

Our first tool of the chapter allows us to focus on incomparable pairs consisting of a minimal and a maximal element.

Lemma 2.1.1 (Min-Max Reduction)**.** *Let P be a poset of height at most h. Then there exists a poset P' containing P as an induced subposet such that*

(i) *P' has height at most h,*
(ii) $\dim(P) \leqslant \dim_{P'}(\mathrm{Min}(P'), \mathrm{Max}(P'))$,
(iii) $\mathrm{cover}(P')$ *can be obtained by attaching degree-*1 *vertices to* $\mathrm{cover}(P)$.

Proof. We construct P' by extending P in two steps. In the first step, we introduce for each non-minimal element $x \in P - \mathrm{Min}(P)$ a new minimal element x', which is such that for every $y \in P$ we have $x' < y$ in P if and only if $y \in \mathrm{U}[x]$. Note that $x' < x$ is the only cover relation involving x' in the resulting poset.

In the second step, we introduce for each non-maximal element $x \in P - \mathrm{Max}(P)$ a maximal element x'', which is such that for every $y \in P$ we have $y < x'$ in P if and only if $y \in \mathrm{D}[x]$. Moreover, for each $a, b \in P$ we set $a' < b''$ if and only if $a \leqslant b$ in P (to guarantee transitivity). Similarly as before, the only cover relation involving x'' is $x < x''$. Figure 2.1 illustrates the result after the two construction steps.

Let P' be the resulting poset for which we show conditions (i)-(iii) now. The height of P' clearly agrees with the height of P, hence (i) holds. Since each newly introduced element is involved in exactly one cover relation where the other element is from P, we deduce that (iii) holds.

To show (ii), let $t = \dim_{P'}(\mathrm{Min}(P'), \mathrm{Max}(P'))$. Then there exist linear extensions $L_1, \ldots, L_t$ of P' that reverse each pair of $\mathrm{Inc}(\mathrm{Min}(P'), \mathrm{Max}(P'))$. We claim that the restriction of these linear extensions to elements of P yields a realizer of P. To verify this, we need to show that each incomparable pair $(x, y) \in \mathrm{Inc}(P)$ is reversed in at least one linear extension of $L_1, \ldots, L_t$. For convenience, if $x \in \mathrm{Min}(P)$ (so x' does not exist) then we also write x' for x, and similarly, if $y \in \mathrm{Max}(P)$ (so y'' does not exist), then we also write y'' for y. With this convention, it follows that $(x', y'') \in \mathrm{Inc}(\mathrm{Min}(P'), \mathrm{Max}(P'))$ and thus there is $i \in [t]$ such that $y'' < x'$ in L_i. Together with the comparabilities in P we conclude that $y \leqslant y'' < x' \leqslant x$ in L_i,

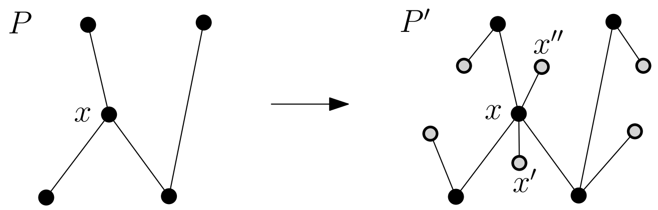

Figure 2.1: A poset P and the construction of P' as done in Lemma 2.1.1.

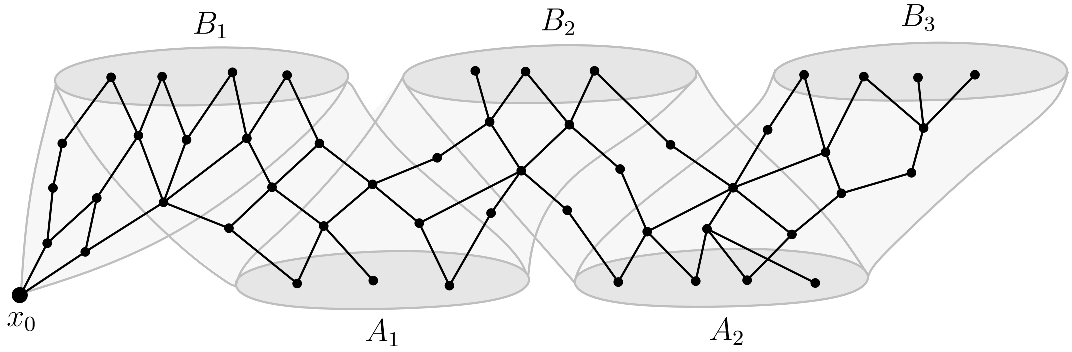

Figure 2.2: Illustration of an unfolding of P starting from a_0.

and hence (x, y) is reversed in L_i, as claimed. This shows (ii) and completes the proof. □

Observe already here that we can apply this lemma whenever we work with bounded height posets whose cover graphs are sparse; indeed, adding degree-1 vertices does not turn a graph into a dense one.

2.2 Unfolding Posets and Cores

We continue with the second major tool of the chapter. Let P be a connected poset with at least two elements and let $A := \mathrm{Min}(P)$ and $B := \mathrm{Max}(P)$. Note that A and B are disjoint. Given a fixed element $x_0 \in A \cup B$ we can *unfold* the poset P starting from x_0 in the following way. If $x_0 \in A$ then set $A_0 := \{x_0\}$ and $B_1 := \{b \in B \mid a_0 \leqslant b \text{ in } P\}$. If $x_0 \in B$ then set $A_0 := \emptyset$ and $B_1 := \{x_0\}$. In both cases, define the following sets for $i = 1, 2, \ldots$

$$A_i := \Big\{a \in A - \bigcup_{0 \leqslant j < i} A_j \mid \text{ there is } b \in B_i \text{ with } a \leqslant b \text{ in } P\Big\},$$
$$B_{i+1} := \Big\{b \in B - \bigcup_{1 \leqslant j < i+1} B_j \mid \text{ there is } a \in A_i \text{ with } a \leqslant b \text{ in } P\Big\}.$$

Since P is connected, the sets $A_0, A_1, \ldots$ partition A, and the sets $B_1, B_2 \ldots$ partition B. Let us emphasize that A_0 is possibly empty. Note that at some point all sets in the sequence $A_0, B_1, A_1, B_2 \ldots$ are empty. A prefix of the sequence $A_0, B_1, A_1, B_2 \ldots$ containing all non-empty sets is called an *unfolding* of P from x_0. See Figure 2.2 for an illustration. Observe that an unfolding of P corresponds to a BFS in the comparability graph of P when restricted to the elements of $\mathrm{Min}(P)$ and $\mathrm{Max}(P)$.

Let us point out the following simple but important observation:

If $x \in \mathrm{U}(A_i)$ with $i \geqslant 1$, then x is in at least one of the two downsets $\mathrm{D}(B_i)$ and $\mathrm{D}(B_{i+1})$, possibly both, but in no other downset $\mathrm{D}(B_j)$ with $j \neq i, i+1$. If $x \in \mathrm{U}(A_0)$ then x is included only in $\mathrm{D}(B_1)$.

Dually, if $x \in \mathrm{D}(B_i)$ with $i \geqslant 1$, then x is in at least one of the two upsets $\mathrm{U}(A_{i-1})$ and $\mathrm{U}(A_i)$, possibly both, but in no other upset $\mathrm{U}(A_j)$ with $j \neq i-1, i$. (2.1)

The following lemma plays a key role in all applications of the unfolding concept. Intuitively, it says that every unfolding contains a local part that roughly witnesses the dimension of the poset (more precisely, $\dim(A, B)$), up to a factor 2. (Note that this is quite analogous to the behavior of chromatic number with respect to some BFS.)

Lemma 2.2.1 (Unfolding lemma)**.** *Let P be a connected poset, let $A := \mathrm{Min}(P)$, $B := \mathrm{Max}(P)$, and suppose that $\dim(A, B) \geqslant 2$. Consider the sequence $A_0, B_1, \ldots, A_{m-1}, B_m$ obtained by unfolding P from some element $x_0 \in A \cup B$. Then there exists an index i such that*

$$\dim(A_i, B_i) \geqslant \dim(A, B)/2$$
or
$$\dim(A_i, B_{i+1}) \geqslant \dim(A, B)/2.$$

Proof. Let

$$t = \max\{\dim(A_i, B_{i+1}) \mid i = 0, \ldots, m-1\},$$
$$t' = \max\{\dim(A_i, B_i) \mid i = 1, \ldots, m-1\}.$$

We aim to show $\dim(A, B) \leqslant t + t'$, which implies the lemma. To do so, we use linear extensions reversing 'local' incomparable pairs and combine those to obtain linear extensions of P.

For each $i \in \{0, \ldots, m-1\}$, we fix t linear extensions $L_1^i, \ldots, L_t^i$ of $P|_{A_i \cup B_{i+1}}$ that reverse all pairs from the set $\mathrm{Inc}(A_i, B_{i+1})$ (those exist by the definition of t). Next, we combine these to obtain t linear extensions of P in the following way. For each $j \in [t]$, we concatenate the linear extensions $L_j^{m-1}, \ldots, L_j^0$ as follows:

$$L_j^{m-1} < \cdots < L_j^1 < L_j^0. \tag{2.2}$$

Observe that this yields a linear extension of the poset $P|_{A \cup B}$. Therefore, the linear order in (2.2) can be extended to a linear extension L_j of P. Then it is not hard to see that $L_1, \ldots, L_t$ reverse all pairs $(a, b) \in \mathrm{Inc}(A, B)$ with $a \in A_i$ and $b \in B_j$ such that $j \geqslant i + 1$.

Similarly, we construct t' linear extensions that reverse the remaining incomparable pairs. For each $i \in [m-1]$, we fix t' linear extensions $L_1^i, \ldots, L_{t'}^i$ of $P|_{A_i \cup B_i}$ reversing all pairs from $\mathrm{Inc}(A_i, B_i)$. For $j \in [t']$, we see that

$$L_j^1 < L_j^2 < \cdots < L_j^{m-1}$$

is a linear extension of $P|_{A\cup B}$ (unless $x_0 \in A_0$, in which case we put x_0 at the first position of the linear extension), allowing us to extend it to a linear extension L'_j of P. Now, $L'_1 \ldots, L'_{t'}$ reverse all pairs $(a, b) \in \mathrm{Inc}(A, B)$ with $a \in A_i$ and $b \in B_j$ such that $j \leqslant i$. We conclude that $L_1, \ldots, L_t, L'_1 \ldots, L'_{t'}$ all together reverse every pair of $\mathrm{Inc}(A, B)$, and hence we have $\dim(A, B) \leqslant t + t'$. □

In the following, suppose that we have $\dim(A_i, B_j) \geqslant \dim(A, B)/2$ for $j \in \{i, i+1\}$ as in the lemma. Assume further that $\dim(A, B) \geqslant 6$, implying $\dim(A_i, B_j) \geqslant 3$. Consider the convex hull $\mathrm{conv}_P(A_i \cup B_j)$. Recall that it is not necessarily connected. However, it is a well-known fact that the dimension of a poset is witnessed by that of one of its components, unless the poset is the union of at least two chains, in which case it has dimension 2. The same statement remains true for the variant of dimension under consideration (see e.g. [39, Observation 5]). Since $\dim(A_i, B_j) > 2$ by our assumption, this means that there exists a component P' of $\mathrm{conv}_P(A_i \cup B_j)$ with $\mathrm{Min}(P') \subseteq A_i$, $\mathrm{Max}(P') \subseteq B_j$, and such that

$$\dim(\mathrm{Min}(P'), \mathrm{Max}(P')) = \dim(A_i, B_j).$$

The poset P' is then said to be a *core* of P with respect to x_0, or simply an x_0-*core* of P for brief. Clearly, P' is a convex subposet of P. We attribute a type to P' depending on whether $j = i$ or $j = i + 1$: If $j = i$ then P' is *left-facing*, and if $j = i + 1$ then P' is *right-facing*. For example, under the assumption that $\mathrm{conv}(A_2, B_2)$ in Figure 2.6 on the left is a core, $\mathrm{conv}(A_2, B_2)$ would be left-facing.

To summarize, when P is a connected poset with $\dim(\mathrm{Min}(P), \mathrm{Max}(P)) \geqslant 6$, then we may consider a core P' of P, and such a core has the following properties:

- P' is connected;
- P' is a convex subposet of P with element set $\mathrm{U}_P(\mathrm{Min}(P')) \cap \mathrm{D}_P(\mathrm{Max}(P'))$;
- $\dim(\mathrm{Min}(P'), \mathrm{Max}(P')) \geqslant \dim(\mathrm{Min}(P), \mathrm{Max}(P))/2$, and
- P' is either left-facing or right-facing.

The *Unfolding Lemma* plays a key role in several proofs included in this thesis. To see its power, consider the following useful idea. Starting with some poset P, we first unfold P, take the core of this unfolding, and then we continue with the same operations on that core. In other words, we iteratively apply the *Unfolding Lemma*, and if we start with P being large-dimensional then this allows us to perform this step many times. We end up with a core that is connected to many 'unfoldings' whose vertices are somewhat floating around the core. This is because the vertices of a core are connected to the vertex from which we started the unfolding via a path 'through' the unfolding sequence. In Chapter 5 we will use this floating material to specify a large complete graph as a topological minor in the cover graph of P.

We continue with a first direct application of the *Unfolding Lemma*. It is useful whenever we work with posets whose cover graphs are contained in a minor-closed graph class.

Lemma 2.2.2 (Global Min Support Reduction). *For every poset P of height at most h there exists a poset Q such that*

(i) *Q has height at most h,*
(ii) $\mathrm{cover}(Q)$ *is a minor of* $\mathrm{cover}(P)$*,*
(iii) $\dim_P(\mathrm{Min}(P), \mathrm{Max}(P)) \leqslant 2 \dim_Q(\mathrm{Min}(Q), \mathrm{Max}(Q))$,

(iv) *there is* $x_0 \in \mathrm{Max}(Q)$ *such that* $x_0 \geqslant a$ *in* Q *for all* $a \in \mathrm{Min}(Q)$.

Proof. If P itself satisfies item (iv), then we can take P as Q and are done. Similarly, if the dual P^d of P satisfies item (iv), then P^d can be chosen as Q and we are done as well.

So suppose that neither P nor its dual P^d fulfill item (iv). It is not hard to see that this implies $\dim(\mathrm{Min}(P), \mathrm{Max}(P)) \geqslant 2$. This allows us to apply the *Unfolding Lemma* (Lemma 2.2.1): Let $A_0, B_1, \ldots, A_{m-1}, B_m$ be a sequence obtained by unfolding P, then there is $i \in \{0, \ldots, m\}$ and $j \in \{i, i+1\}$ such that

$$\dim_P(A_i, B_j) \geqslant \dim_P(\mathrm{Min}(P), \mathrm{Max}(P))/2. \tag{2.3}$$

We distinguish between the cases whether $j = i$ or $j = i + 1$. First, we consider the case $j = i + 1$ which can be seen as the case where the core of the unfolding is right-facing. We define Q to be the poset that is obtained from $\mathrm{conv}(A_i \cup B_{i+1})$ by adding to it an extra element q, which is such that $q > x$ holds in Q if and only if $x \in \mathrm{D}_P(B_i)$; see Figure 2.3 for an illustration of this definition. It follows that $q > a$ in Q for all $a \in A_i = \mathrm{Min}(Q)$, and hence item (iv) is satisfied by Q. Observe that the cover graph of Q is an induced subgraph of $\mathrm{cover}(P)$ with an extra vertex q linked to some of the other vertices. Here, q can be seen as the result of the contraction of the connected set $\bigcup_{1 \leqslant k \leqslant i} D(B_k) - \mathrm{conv}(A_i \cup B_{i+1})$ plus the deletion of some of the edges incident to the contracted vertex (see again Figure 2.3, dashed edges indicate deletions). The deletion step is necessary, as after the contraction it might be that some edges incident to q do not correspond to cover relations anymore. It follows that $\mathrm{cover}(Q)$ is a minor of $\mathrm{cover}(P)$, thus (ii) holds. Furthermore, by (2.3) we have

$$\dim_P(\mathrm{Min}(P), \mathrm{Max}(P)) \leqslant 2 \dim_Q(\mathrm{Min}(Q), \mathrm{Max}(Q)).$$

Therefore, Q also satisfies condition (iii). By the definition of Q we also have $\mathrm{height}(Q) \leqslant \mathrm{height}(P) \leqslant h$, which yields item (i). This concludes the case $j = i+1$.

Second, we consider the case that $j = i$, meaning that the core of the unfolding is left-facing. We aim to reduce this case to the previous one. To do so, we take the dual poset P^d of P and unfold it starting from the same element as before. It is easy to see that we obtain exactly the same sequence as before (with a possible 'shift', but the sets are the same), with the only difference that maximal and minimal elements are exchanged now. Moreover, since the core of the original unfolding is left-facing, it follows that the dual of the core is a right-facing core in the dual unfolding (recall that the dimension is invariant under taking the dual). Therefore, applying the operations of the first case to P^d we find a poset Q satisfying (i)-(iv) with respect to P^d. Since the cover graphs of P and P^d are isomorphic, the poset Q fulfills the conditions of the lemma also with respect to P. This concludes the proof. □

Note that by the arguments applied in the proof it also follows that there exists a dual version of the *Global Min Support Reduction*; here, condition (iv) is replaced by the existence of an element $x_0 \in \mathrm{Min}(Q)$ such that $x_0 \leqslant b$ in Q for all $b \in \mathrm{Max}(Q)$.

We proceed with the second application of the *Unfolding Lemma*. It will be useful in the case that the cover graph of a poset P excludes a fixed graph as a topological minor. In contrast to the previous application, this means that we are not allowed to contract the portion of an unfolding of P that lies before the core.

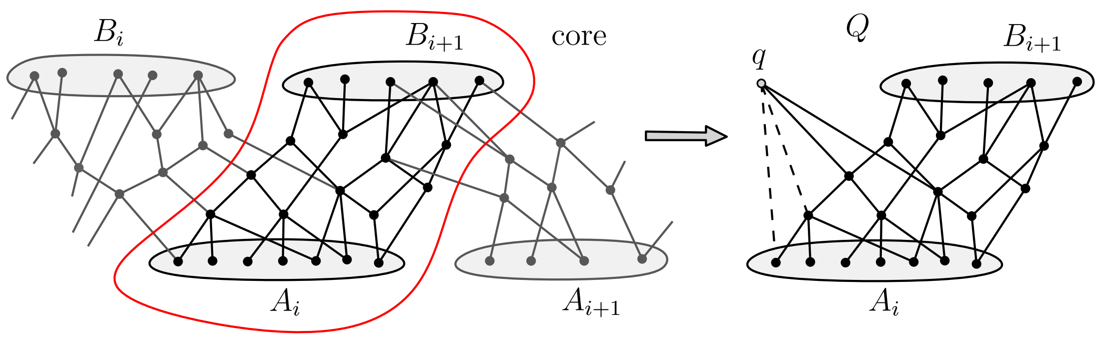

Figure 2.3: Definition of Q in the case that $\dim(A_i, B_{i+1})$ is at least $\dim(\mathrm{Min}(P), \mathrm{Max}(P))/2$.

However, this portion might still be used to connect two vertices of the core by a path. We will use this idea in Chapter 5, where we aim to construct a topological minor after iteratively unfolding a large-dimensional poset.

Lemma 2.2.3 (Extended Unfolding Lemma). *For every poset P and sets $A \subseteq \mathrm{Min}(P)$, $B \subseteq \mathrm{Max}(P)$ with $\dim(A, B) \geqslant 3$, there are sets $A' \subseteq A$, $B' \subseteq B$, and $S \subseteq \mathrm{conv}_P(A \cup B)$ such that*

(i) *elements of S induce a connected subgraph of* $\mathrm{cover}(P)$*,*
(ii) $\dim_P(A', B') \geqslant \dim_P(A, B)/2$*,*
(iii) *it holds that either*

$$\text{a)}\quad A' \cap \mathrm{D}(S) = \emptyset \quad \textit{and} \quad B' \subseteq \mathrm{U}(S),$$
$$\textit{or}$$
$$\text{b)}\quad B' \cap \mathrm{U}(S) = \emptyset \quad \textit{and} \quad A' \subseteq \mathrm{D}(S).$$

Proof. First, using $\dim(A, B) \geqslant 3$ we see that $\mathrm{conv}_P(A \cup B)$ cannot be a union of chains (which is 2-dimensional). It is well known that in this case the dimension of $\mathrm{conv}_P(A \cup B)$ is witnessed by the dimension of one of its components, and this holds as well for the 'min-max dimension'. Therefore, there are sets $A'' \subseteq A$ and $B'' \subseteq B$ such that $\dim_P(A'', B'') = \dim_P(A, B)$ and the poset $Q := \mathrm{conv}_P(A'' \cup B'')$ is connected.

Now choose $a_0 \in A''$ arbitrarily and let $A_0, B_1, \ldots, A_{m-1}, B_m$ be a sequence obtained by unfolding Q from a_0. By the *Unfolding Lemma* (Lemma 2.2.1) there is $\ell \in [m-1]$ such that

$$\dim_Q(A_\ell, B_\ell) \geqslant \dim_Q(A'', B'')/2 \quad \text{or} \quad \dim_Q(A_\ell, B_{\ell+1}) \geqslant \dim_Q(A'', B'')/2.$$

In the first case we will find a set S fulfilling part a) of item (iii), and in the second case we will find a set S satisfying part b) of item (iii).

Suppose first that $\dim_Q(A_\ell, B_\ell) \geqslant \dim_Q(A'', B'')/2$. Then we are in the situation that the core of the unfolding is contained in $\mathrm{conv}_Q(A_\ell \cup B_\ell)$ and that it is left-facing. Since Q is an induced subposet of P, the inequality $\dim_P(A_\ell, B_\ell) \geqslant \dim_P(A'', B'')/2$ holds as well. Set $A' := A_\ell$ and $B' := B_\ell$.

Next, we use the portion of the unfolding lying before the core to define the set S. We let

$$S = \mathrm{conv}_Q\Big(\bigcup_{i \in [0, \ell-1]} A_i \cup \bigcup_{i \in [1, \ell-1]} B_i\Big). \tag{2.4}$$

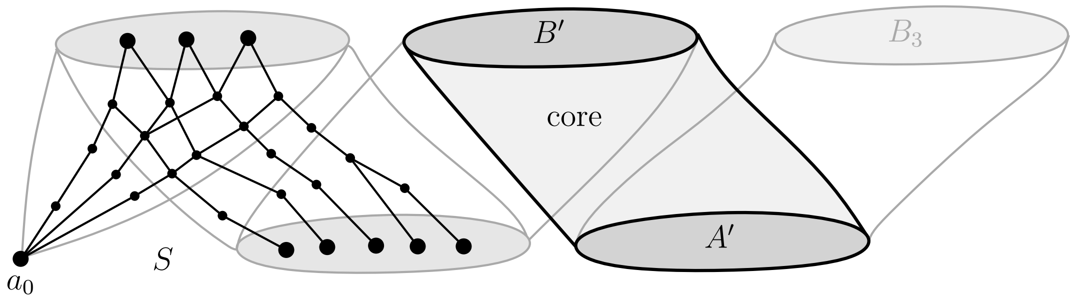

Figure 2.4: Illustration of Lemma 2.2.3. The core of the unfolding is contained in $\mathrm{conv}(A' \cup B')$ and the portion lying before is used to define the set S.

See Figure 2.4 for an illustration of this definition. By the way we unfold, for every $x \in S$ there is path in $\mathrm{cover}(P)$ from x to a_0 that contains only elements of S. Therefore, the elements of S induce a connected subgraph of $\mathrm{cover}(P)$, which shows item (i).

Now we show that $A' \cap \mathrm{D}(S) = \emptyset$ holds. Suppose for a contradiction that this is not true. As S and A' are disjoint this means that there are $a' \in A'$ and $s \in S$ such that $a' < s$ in P. By the definition of S (see (2.4)) this implies $s \leqslant b$ for some $b \in B_i$ with $i \in [1, \ell - 1]$, and therefore $s \in \mathrm{D}(B_i) \cap \mathrm{U}(A_\ell)$. However, this is a contradiction as by property (2.1) the intersection $\mathrm{D}(B_i) \cap \mathrm{U}(A_\ell)$ must be empty whenever $i \leqslant \ell - 1$.

So we indeed have $A' \cap \mathrm{D}(S) = \emptyset$. Since $B' = B_\ell \subseteq \mathrm{U}(A_{\ell-1})$ (by the definition of an unfolding) and $A_{\ell-1} \subseteq S$, we deduce that $B' \subseteq \mathrm{U}(S)$. This shows that part a) of item (iii) holds and completes the case of a left-facing core.

We are left with the case that $\dim_Q(A_\ell, B_{\ell+1}) \geqslant \dim_Q(A'', B'')/2$. Here, we set $A' = A_\ell$ and $B' = B_{\ell+1}$. Similarly to the previous case, we let

$$S = \mathrm{conv}_Q\Big(\bigcup_{i \in [0, \ell-1]} A_i \cup \bigcup_{i \in [1, \ell]} B_i \Big).$$

The properties required for A', B', and S now follow along the same lines as for the first case. This concludes the proof. □

2.3 Zig-Zag Paths

In the last section of this chapter we introduce zig-zag paths with respect to some fixed unfolding. Their union will turn out to be a tree, which makes them applicable whenever we work with a planar embedding of a cover graph as this embedding induces a natural left-right notion on the elements of the tree.

Let P be a connected poset and let $A := \mathrm{Min}(P)$ and $B := \mathrm{Max}(P)$. Choose some $x_0 \in A \cup B$ and let $A_0, B_1, \ldots, A_{m-1}, B_m$ be the sequence obtained by unfolding P starting from x_0. We will define for each element $x \in P$ a corresponding zig-zag path connecting x with x_0 in the cover graph of P.

To do so, we first need to introduce some notations. For $x \in P - \{x_0\}$, let $\alpha(x)$ denote the smallest index $i \geqslant 0$ such that $x \in \mathrm{U}(A_i)$, and $\beta(x)$ the smallest index $j \geqslant 1$ such that $x \in \mathrm{D}(B_j)$. (For instance, in Figure 2.5 we have $\alpha(x) = 2, \beta(x) = 3$ and

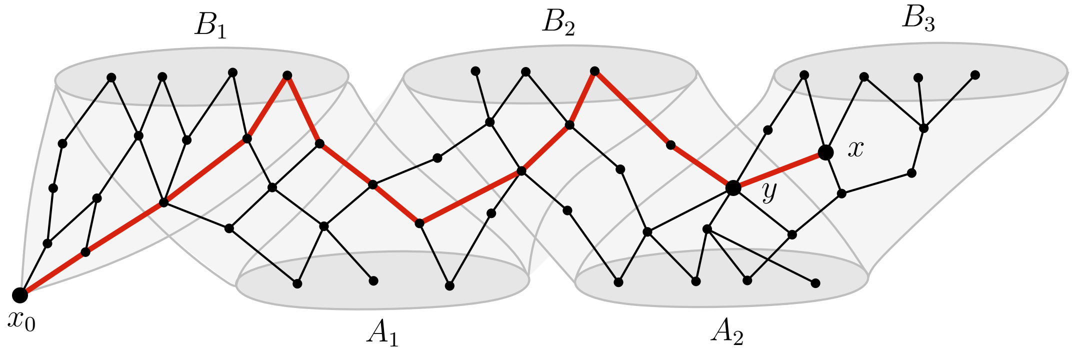

Figure 2.5: Unfolding P from x_0 and the zig-zag path $Z_P^{x_0}(x)$ of x.

$\alpha(y) = 2, \beta(y) = 2$.) Note that $x \in \text{conv}(A_{\alpha(x)}, B_{\beta(x)})$ and $\beta(x) \in \{\alpha(x), \alpha(x) + 1\}$ by (2.1).

Next, we associate to each element $x \in P - \{x_0\}$ a *parent*: If $\beta(x) = \alpha(x) = i$, then note that $x \notin B_i$ since $\alpha(b) = i - 1$ for all elements $b \in B_i$. Thus, there exist $y \in P$ and $b \in B_i$ such that $x < y \leqslant b$ in P and $x < y$ is a cover relation in P, and we choose (arbitrarily) one such element y to be the parent of x. If $\beta(x) = \alpha(x)+1 = i+1$ (as in Figure 2.5) then $x \notin A_i$ since $\beta(a) = i$ for all elements $a \in A_i$. Thus, there exist $y \in P$ and $a \in A_i$ such that $a \leqslant y < x$ in P and $y < x$ is a cover relation. We choose one such element y to be the parent of x.

We write parent(x) to denote the parent of x. Observe that when $\alpha(x) = \beta(x) = i$, then by definition parent$(x) \in \text{D}(B_i)$ and hence $\alpha(\text{parent}(x)) \leqslant i = \alpha(x)$ by (2.1). Also since parent$(x) \in \text{D}(B_i)$, we have $\beta(\text{parent}(x)) \leqslant i = \beta(x)$. Similarly, when $\alpha(x) = i$ and $\beta(x) = i + 1$, then by definition parent$(x) \in \text{U}(A_i)$ and hence $\alpha(\text{parent}(x)) \leqslant i = \alpha(x)$. And since parent$(x) \in \text{U}(A_i)$, we also have $\beta(\text{parent}(x)) \leqslant i + 1 = \beta(x)$ by (2.1).

To summarize, we have:

$$\alpha(\text{parent}(x)) \leqslant \alpha(x) \text{ and } \beta(\text{parent}(x)) \leqslant \beta(x), \tag{2.5}$$

for every $x \neq x_0$ in P.

Let T be the spanning subgraph of cover(P) obtained by only keeping edges connecting an element and its designated parent.

Claim 2.3.1. *The graph T is a tree.*

Proof. To see this, it is convenient to orient each edge $\{x, \text{parent}(x)\}$ of T towards parent(x). Then every element of P is incident to exactly one outgoing edge, except for x_0 which is a sink. Thus, to show that T is a tree it is enough to show that there is no directed cycle in this orientation of T. Arguing by contradiction, suppose that $x_1x_2\dots x_k$ is a directed cycle. Then $\alpha(x_1) \leqslant \cdots \leqslant \alpha(x_k) \leqslant \alpha(x_1)$ and $\beta(x_1) \leqslant \cdots \leqslant \beta(x_k) \leqslant \beta(x_1)$ by (2.5). Thus all these inequalities hold with equality. By (2.1) we know that $\beta(x_i) \in \{\alpha(x_i), \alpha(x_i) + 1\}$, for all $i \in \{1, \dots, k\}$. Suppose first that $\beta(x_i) = \alpha(x_i)$, for all $i \in \{1, \dots, k\}$. Then $x_i < \text{parent}(x_i) = x_{i+1}$ in P for each $i \in \{1, \dots, k\}$ (cyclically), which is a contradiction. Now assume that $\beta(x_i) = \alpha(x_i) + 1$, for all $i \in \{1, \dots, k\}$. Then $x_i > \text{parent}(x_i) = x_{i+1}$ in P for each $i \in \{1, \dots, k\}$ (cyclically), which is again a contradiction. This completes the proof. □

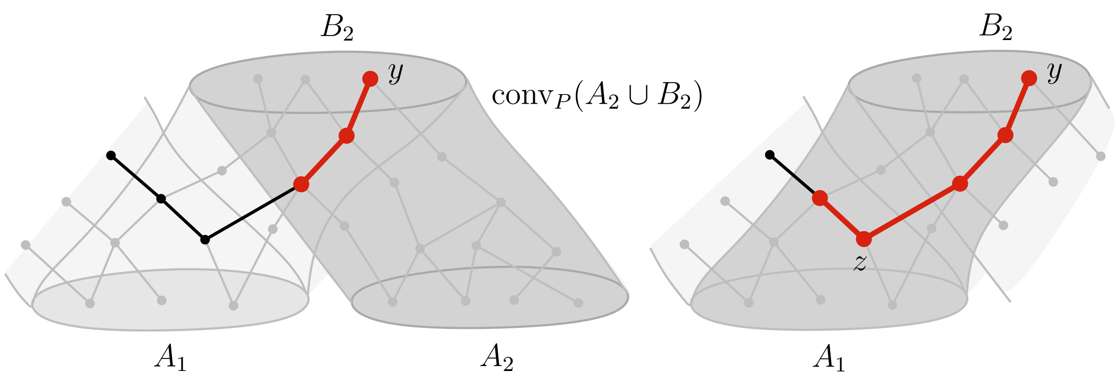

Figure 2.6: Illustration of Claim 2.3.2.

For each $x \in P$ let $Z_P^{x_0}(x)$ denote the unique path in T that connects x_0 and x, which we call the *zig-zag path* of x. (We drop the subscript P when the poset is clear from the context.) See Figure 2.5 for an illustration.

The following claim describes the shape of the initial part of the zig-zag path $Z^{x_0}(y)$ starting from y. It is easy to see that when y belongs to all three sets $\mathrm{U}(A_{i-1})$, $\mathrm{U}(A_i)$, and $\mathrm{D}(B_i)$, then the zig-zag path goes down from y in P. We will see that the trace of the zig-zag path $Z^{x_0}(y)$ in $\mathrm{U}(A_i)$ is a chain, while its trace in $\mathrm{D}(B_i)$ has a unique minimal element, see Figure 2.6 for an illustration.

Claim 2.3.2. *If an element y lies in all three sets $\mathrm{U}(A_{i-1})$, $\mathrm{U}(A_i)$, and $\mathrm{D}(B_i)$ for some $i \in \{1, \ldots, k\}$, then*

- (i) *the elements of $Z^{x_0}(y) \cap \mathrm{U}(A_i)$ form a chain contained in $\mathrm{U}(A_{i-1})$, and y is the maximal element of that chain;*
- (ii) *the elements of $Z^{x_0}(y) \cap \mathrm{D}(B_i)$ induce a subposet of P with a unique minimal element.*

Proof. Since y belongs to $\mathrm{U}(A_{i-1})$, $\mathrm{U}(A_i)$, and $\mathrm{D}(B_i)$, we see that $\alpha(y) = i-1$ and $\beta(y) = i$ by (2.1).

Now start in y and keep walking along the zig-zag path $Z^{x_0}(y)$ towards x_0, as long as the parent of the current element x is below x in P, or we reach x_0. Let z be the element we stop at. If $z = x_0$ then the whole zig-zag path $Z^{x_0}(y)$ is simply a chain in P with y being its maximal element, thus (i) holds. Property (ii) holds as well as the whole chain lies in $D(B_1)$ and has a unique minimal element, namely x_0.

So we assume that $z \neq x_0$. In this case, we keep walking along the zig-zag path $Z^{x_0}(y)$ towards x_0, as long as the parent of the current element x is above x in P, or we reach x_0. Let z' be the new element we stop at.

Observe that $z < \mathrm{parent}(z)$ in P. Since $z \leqslant y$ in P and $y \in \mathrm{D}(B_i)$, we have $z \in D(B_i)$. Clearly, z is below (in P) all the other elements of the zig-zag path $Z^{x_0}(y)$ that we traversed. We have $\alpha(z) \leqslant \alpha(y) = i-1$, by (2.5). We also have $\alpha(z) \geqslant \alpha(y)$ since $z \leqslant y$ in P and thus $\mathrm{D}(z) \subseteq \mathrm{D}(y)$. Hence $\alpha(z) = \alpha(y) = i-1$. Since $z < \mathrm{parent}(z)$ in P, we must have $\beta(z) = \alpha(z) = i-1$ by the definition of parents.

It follows that $\beta(x) \leqslant \beta(z) = i-1$ for all x on the zig-zag path $Z^{x_0}(y)$ appearing after z (towards x_0) as well, by (2.5). For such elements x we cannot have $x \in U(A_i)$ by (2.1). Hence $Z^{x_0}(y) \cap \mathrm{U}(A_i)$ is a subset of the y–z portion of the zig-zag path $Z^{x_0}(y)$. Thus it is a chain with y as its maximal element. This proves (i).

For the proof of (ii), we are going to show that z is the unique minimal element in the subposet of P induced by $Z^{x_0}(y) \cap \mathrm{D}(B_i)$. Recall that $z \in \mathrm{D}(B_i)$.

If $z' = x_0$ then all elements on the zig-zag path $Z^{x_0}(y)$ are greater or equal to z in the poset P. Thus (ii) holds.

Now, assume that $z' \neq x_0$. Thus $z' > \mathrm{parent}(z')$ in P. We have $\beta(z') \leqslant \beta(z)$, by (2.5). We also have $\beta(z') \geqslant \beta(z)$ since $z' > z$ in P and thus $\mathrm{U}(z') \subseteq \mathrm{U}(z)$. Hence $\beta(z') = \beta(z) = i - 1$. Since $z' > \mathrm{parent}(z')$ in P, we must have $\alpha(z') = \beta(z') - 1 = i - 2$ by the definition of parents.

It follows that $\alpha(x) = \alpha(z') \leqslant i-2$ for all x on the zig-zag path $Z^{x_0}(y)$ appearing after z' (towards x_0) as well, by (2.5). For such elements x we cannot have $x \in D(B_i)$ by (2.1). Hence $Z^{x_0}(y) \cap \mathrm{D}(B_i)$ is a subset of the y–z' portion of the zig-zag path $Z^{x_0}(y)$. Clearly, z is the unique minimal element in the subposet induced by that subset. This completes the proof of (ii). □

Chapter 3

Planar Posets

The roots of dimension problems on planar posets can be traced back to the early 1970's, or even some years before. The book *Lattice Theory* by G. Birkhoff [5] contains an exercise (page 32, Exercise 7 (c), with credits to J. Zilber), where the reader is asked to show that a lattice has a planar diagram if and only if its cocomparability graph is a comparability graph. The latter property was already known to be equivalent to having dimension at most 2 (see [17]), so this exercise was revealing the first connection between planarity and dimension: A lattice has a planar diagram if and only if it has dimension at most 2. In 1972, Baker, Fishburn, and Roberts [2] extended this observation using the so-called Dedekind-MacNeille completions. As a corollary of their results they obtained that every planar poset with a single minimal and a single maximal element has dimension at most 2. Moore and Trotter [75] followed this line of research and showed in 1977 that planar posets with a single minimal element have dimension at most 3. This early research on the dimension of planar posets stopped in 1981 by Kelly's construction of planar posets with unbounded dimension (see again Figure 1 from the introduction).

Let us continue now with pointing out some fundamental differences between the two concepts of planar posets and posets whose cover graphs are planar. Clearly, every planar poset has a planar cover graph, but the converse need not be true as we have already shown in Chapter 1 with the example in Figure 1.1. Regarding posets with planar cover graphs, there is also no analog to the previously mentioned result by Baker, Fishburn, and Roberts: There are posets with planar cover graphs that have a single minimal element and a single maximal element but unbounded dimension, as follows by a construction of Streib and Trotter [71] (see Figure 3.1 and its caption). Another fundamental difference between these two concepts is witnessed by the computational complexity of their corresponding recognition problems: While one can test in linear time whether the cover graph of a given poset is planar, Garg and Tamassia [31] proved that it is NP-complete to decide whether a poset is planar. On posets of height 2 though, those two concepts agree; it is a non-trivial fact that every poset of height 2 has a planar cover graph if and only if it is planar [52, 4]

We proceed with the general theme of this thesis and consider the question whether large-dimensional planar posets have to be tall. As already discussed in the introduction of the thesis this is indeed the case, and therefore we add the question: "How tall?".

It was first proven by Streib and Trotter [71] that the dimension of planar posets

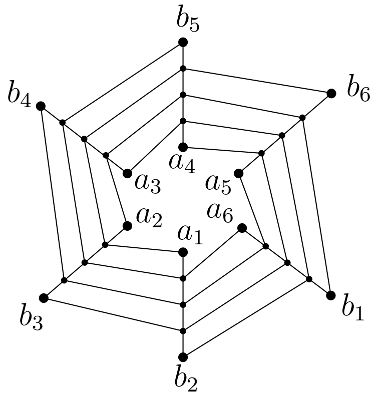

Figure 3.1: A planar cover graph of a poset that contains the standard example S_6. Note that this construction is extendable in an obvious fashion. Moreover, we can add two elements without destroying planarity so that the resulting poset has only one minimal and one maximal element: the minimal one is placed in the center of the 'spider net' and connected it to all the a_i's, and the maximal one is placed in the outer face and connected to all the b_i's.

is bounded by their height. They made no effort in computing the bounding function as this must be enormous due to their repeated applications of Ramsey arguments. A better upper bound is implicitly proven by Joret et al. [38]; they show a double exponential upper bound in height, which can also be deduced by a work of Micek and myself [51].

So far, all the aforementioned results were obtained in the quite general settings when cover graphs are planar, have bounded tree-width, or exclude a fixed graph as a minor. Consequently, no proof of these results uses the special properties of a planar diagram. In joint work with Gwenaël Joret and Piotr Micek, we make use of these properties and establish a linear bound on the dimension of planar posets. In other words, we can show that large-dimensional planar posets are very tall, meaning that those posets contain a chain that is of linear size with respect to the poset's dimension. This is the main result of this chapter.

Theorem 3.0.1 ([40]). *If P is a planar poset of height at most h, then*

$$\dim(P) \leqslant 192h + 96.$$

The linear bound in the theorem is optimal up to a constant factor, as shown by Kelly's examples that have dimension at least $h+1$. Up to now, these examples still provide the best known lower bound on the maximum dimension of a planar poset of height h. The second contribution of us in this chapter is a slight improvement of this lower bound:

Theorem 3.0.2 ([40]). *For every $h \geqslant 1$, there is a planar poset P of height h with*

$$\dim(P) \geqslant (4/3)h - 2.$$

As mentioned before, the upper bound of Streib and Trotter [71] on dimension in terms of height holds in fact in a more general setting, that of posets with planar cover graphs. It is not known whether the dimension of posets with planar cover graphs is bounded by a linear function of their height (or any polynomial function for that matter). We also present a slightly better construction in that less restrictive setting:

Theorem 3.0.3 ([40])**.** *For every $h \geqslant 1$, there is a poset of height h with a planar cover graph and dimension at least $2h - 2$.*

The rest of the chapter is organized as follows. First, we prove a lemma in Section 3.1 that is at the heart of our proof of Theorem 3.0.1. Informally, this lemma states that a linear bound holds for planar posets that have a special element supporting all maximal elements of the poset. Its proof uses heavily the planarity of the diagram (as opposed to merely using the planarity of the cover graph). Second, we finish the proof of Theorem 3.0.1 in Section 3.2, by reducing the case of general planar posets to the case covered by the lemma. Finally, we describe in Section 3.3 the two lower bound constructions giving us Theorem 3.0.2 and Theorem 3.0.3.

All results of this chapter were obtained jointly with Piotr Micek and Gwenaël Joret (see [40]).

3.1 Planar Posets with Support

In this section, we prove that if a planar poset has an element that is supporting all of its maximal elements (meaning that it is below all maximal elements) then its 'min-max' dimension is bounded by a linear function of its height. This result will serve as a main lemma in the proof of our main theorem of the chapter. We remark that the proof of the lemma uses heavily the drawing of the diagram in the plane and does not extend to the case of a poset with a planar cover graph.

Lemma 3.1.1. *Let P be a planar poset of height h and let $B \subseteq \mathrm{Max}(P)$. If there is an element x_0 such that $x_0 \leqslant b$ in P for all $b \in B$, then $\dim(\mathrm{Min}(P), B) \leqslant 6h + 3$.*

The proof of Lemma 3.1.1 is split into a number of steps. The general strategy is to partition the set of $\mathrm{Inc}(\mathrm{Min}(P), B)$ into a number (eventually $6h+3$) of reversible subsets. Here is a brief outline of the proof: First, we deal with some incomparable pairs that can easily be reversed using a bounded number of linear extensions (three). The remaining set of incomparable pairs has a natural partition into those pairs (a, b) such that a is 'to the left' of b and those such that a is 'to the right' of b (see Claim 3.1.1). Using symmetry, we can then focus on one of these two types of incomparable pairs. At that point, we study how alternating cycles must look like given the properties of the incomparable pairs garnered so far. We establish various technical properties of these alternating cycles, which enable us to show that if the dimension is large, then there is a large 'nested structure' in the diagram of P, with some maximal element b buried deep inside and the special element x_0 drawn outside the structure. This is where the existence of x_0 is used: A path in the diagram witnessing the relation $x_0 \leqslant b$ in P has to go through the whole nested structure. This path will then meet a number of disjoint curves from the diagram, and by planarity the path can only enter such a curve in an element of the poset. This way, we deduce that the path under consideration contains many elements from P, and therefore that the height of P is large.

Let us now turn to the proof of Lemma 3.1.1. Fix a planar drawing of the diagram of P. Let $A := \mathrm{Min}(P)$ and $I_1 := \mathrm{Inc}(A, B)$. We use the standard coordinate system where each point in the plane is characterized by an x-coordinate and a y-coordinate. We may assume without loss of generality that no two elements of P have the same y-coordinate in the drawing.

Some easy Cases

Given two distinct elements $a, b \in P$, we say that a is *drawn below* (*above*) b if the y-coordinate of a is less than (greater than, respectively) that of b.

Let I_1' be the set of all incomparable pairs $(a, b) \in I_1$ such that a is drawn above b. Observe that if we order the elements of P by increasing order of their y-coordinates in the drawing, then we obtain a linear extension of P (by the definition of a diagram). This linear extension reverses all pairs in I_1'. Let $I_2 := I_1 - I_1'$ be the set of remaining incomparable pairs in I_1, that is, those pairs $(a, b) \in I_1$ such that a is drawn below b. Clearly,

$$\dim(I_1) \leqslant \dim(I_2) + \dim(I_1') \leqslant \dim(I_2) + 1,$$

thus we can restrict our attention to pairs in I_2.

Before pursuing further, let us introduce some terminology. A curve in the plane that can be oriented so that it is strictly increasing on the y-coordinate is said to be *y-increasing*. If a y-increasing curve γ is completely contained in the drawing of P then we call γ a *walk* in the diagram of P. Every covering chain witnessing $a \leqslant b$ in P corresponds to a walk from a to b. Conversely, every two elements of P contained in the same walk are comparable in P. For an element p of P, we define the *p-line* as the horizontal line in the plane through (the image of) p. When p and q are two elements of P, we say that p *sees* the q-line if there is a walk containing p that intersects the q-line.

Let I_2' be the set of all pairs $(a, b) \in I_2$ such that a does not see the b-line. It turns out that the set I_2' is reversible: If not, then I_2' contains an alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$. We may assume that b_1 is drawn below all other b_i's. Given that a_1 is drawn below b_1 (since $(a_1, b_1) \in I_2$) and b_2 is drawn above b_1, any covering chain witnessing the relation $a_1 \leqslant b_2$ in P crosses the b_1-line. Thus a_1 sees the b_1-line, contradicting $(a_1, b_1) \in I_2'$.

Let I_2'' be a set of all pairs $(a, b) \in I_2$ such that b does not see the a-line. A dual argument to the one above shows that I_2'' is reversible. Let $I_3 := I_2 - (I_2' \cup I_2'')$, that is, I_3 is the set of pairs $(a, b) \in I_2$ such that a sees the b-line and b sees the a-line. We have

$$\dim(I_1) \leqslant \dim(I_2) + 1 \leqslant \dim(I_3) + \dim(I_2') + \dim(I_2'') + 1 \leqslant \dim(I_3) + 3.$$

Thus, to prove Lemma 3.1.1, it remains to partition I_3 into at most $6h$ reversible sets.

Given two distinct elements $p, q \in P$, we say that p *sees the left side (right side)* of q if there is a walk containing p and intersecting the q-line to the left (right) of q. Note that in general p could possibly see both the left and the right side of q. However, this cannot happen with elements of pairs in I_3:

Claim 3.1.1. *For each $(a, b) \in I_3$, either*

(i) *a sees only the left side of b and b sees only the right side of a, or*
(ii) *a sees only the right side of b and b sees only the left side of a.*

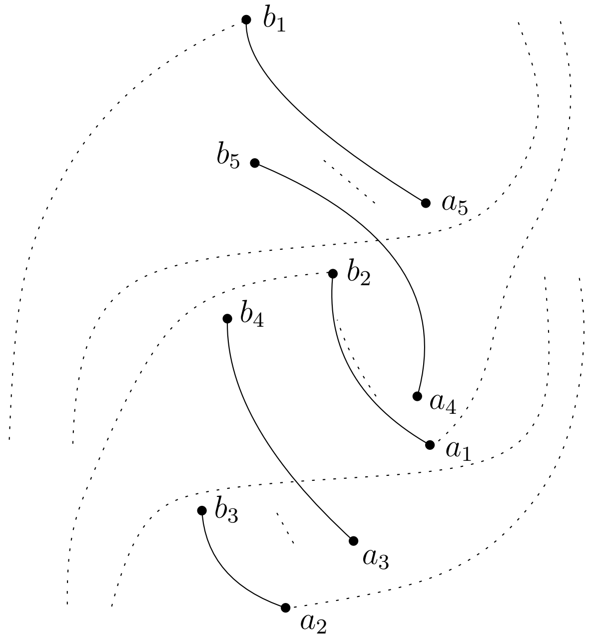

Figure 3.2: An alternating cycle of size 5 with all pairs in I_4^{sep}. Solid line are walks witnessing comparabilities of the cycle. Dotted lines are walks witnessing that for all i: (1) a_i sees the right side of b_i; (2) b_i sees the left side of a_i; (3) the pair (a_i, b_i) is separated.

Proof. Let $(a, b) \in I_3$. Thus a sees the b-line and b sees the a-line. We may assume without loss of generality that a sees the left side of b, the other case being symmetric. Let γ_a be a walk witnessing the fact that a sees the left side of b. First, we show that b cannot see the left side of a. Suppose for a contradiction that it does, and let γ_b be a walk witnessing this. Clearly, γ_a and γ_b intersect. Since the drawing of the diagram of P under consideration is planar, the intersection of γ_a and γ_b contains the image of an element $p \in P$. Since there is a walk containing a and p, we see that $a \leqslant p$ in P, and similarly $p \leqslant b$ in P since there is a walk containing p and b. We conclude that $a \leqslant b$ in P, a contradiction. Thus, b cannot see the left side of a, and hence only sees its right side.

A symmetric argument shows that a cannot see the right side of b, and therefore a only sees its left side. □

We partition the set of pairs $(a, b) \in I_3$ into two sets I_3' and I_3'', depending on whether (a, b) satisfies (i) or (ii) in the claim above. It is enough to show that we can partition one of the two sets into at most $3h$ reversible sets, as we would obtain the same result for the other set by symmetric arguments (that is, by exchanging the notion of left and right). We focus on the set I_3'', and thus aim to prove that $\dim(I_3'') \leqslant 3h$. For convenience, let $I_4 := I_3''$.

We are going to partition pairs in I_4 according to whether they admit a 'separator': Say that a walk σ is a *separator* for a pair $(a, b) \in I_4$ if σ starts on the a-line to the left of a, ends on the b-line to the right of b, and every element of P that appears on σ is incomparable to both a and b. Then let I_4^{sep} be the set of pairs in I_4 having a separator, and let $I_4^{\text{no-sep}}$ be the set of pairs in I_4 with no separator. See Figure 3.2 for an example of an alternating cycle with all pairs in I_4^{sep}.

Say that an incomparable pair $(a, b) \in I_4$ is *dangerous* if a is drawn below x_0

and a sees the left side of x_0. Note if (a, b) is *not* dangerous, then either a is drawn above x_0, or a is drawn below x_0 but then a only sees the right side of x_0 (recall that there is walk from a to the b-line).

We consider dangerous and non-dangerous pairs in I_4^{sep} and $I_4^{\text{no-sep}}$. This defines four subsets of incomparable pairs. A first observation is that one of these sets is empty:

Claim 3.1.2. *There are no dangerous pairs in I_4^{sep}.*

Proof. Arguing by contradiction, suppose that $(a, b) \in I_4^{\text{sep}}$ is dangerous. Let σ denote a walk that separates the pair. Thus a is to the right of σ and b is to its left. Since a is drawn below x_0, the walk σ intersects the x_0-line. Consider a walk β from a to the x_0-line with x_0 to its right, which exists since (a, b) is dangerous. By the definition of a separator, β cannot intersect σ, and hence the top endpoint of β is to right of σ. It follows that x_0 is to the right of σ. However, this implies that any walk witnessing the relation $x_0 \leqslant b$ in P must intersect σ, since b is to the left of σ. This contradicts the fact that σ separates the pair (a, b). □

The plan for the rest of the proof is to partition into at most h reversible sets each of the remaining three sets of incomparable pairs, i.e. non-dangerous pairs in I_4^{sep}, non-dangerous pairs in $I_4^{\text{no-sep}}$, and dangerous pairs in $I_4^{\text{no-sep}}$. Altogether, this proves that $\dim(I_4) \leqslant 3h$, as desired. To do so, we first make a little detour: In the next few pages we introduce the key notion of a 'wall' and establish several useful properties of walls. These walls will help us getting a better understanding of strict alternating cycles in I_4^{sep} and in $I_4^{\text{no-sep}}$, which in turn will help us to partition the three sets mentioned above into at most h reversible sets each.

Walls and i-Walls in the Plane

Given a point p in the plane and a walk γ that intersects the p-line in a point q distinct from p, we say that p is *to the left (right) of* γ if p is to the left (right, respectively) of q on the p-line. A set W of walks is a *wall* if there is a walk $\gamma \in W$ such that every walk $\gamma' \in W$ distinct from γ has the property that its topmost point is to the right of some walk in W. Note that in this case the walk γ is uniquely defined; we call it the *root* of the wall W. Observe also that the topmost point of the root walk has the maximum y-coordinate among all points in walks of W. We begin by showing an easy property of walls.

Claim 3.1.3. *Let W be a wall and let γ be a walk that is disjoint from every walk in W. If the bottommost point of γ is to the right of some walk in W, then every point of γ is either to the right of some walk in W, or above all the walks in W.*

Proof. Assume that the bottommost point of γ is to the right of some walk in W. We are going to walk along γ starting at its bottommost point and going upwards. We claim that at all times the current point satisfies the property that it is either to the right of some walk in W, or above all the walks in W. This is true at the beginning, as by our assumption the bottommost point of γ is to the right of some walk, say δ, in W. Since γ does not intersect δ, walking along γ we have the curve δ to the left, until one of the two curves ends. If γ stops first then every point of γ is to the right of some walk in W, as desired. If δ stops first, then either δ is the root

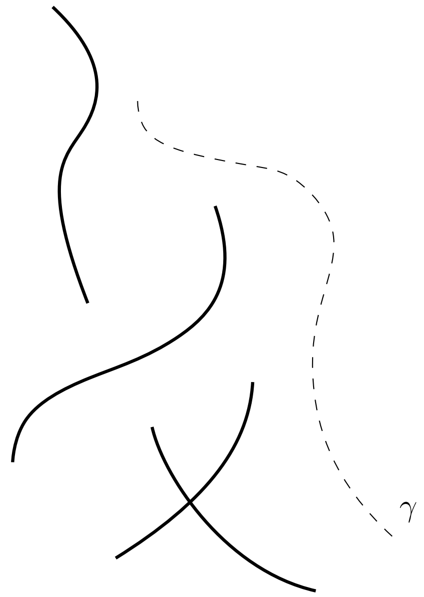

Figure 3.3: A wall (thick curves) and a walk γ that is disjoint from every walk in the wall.

walk of W and therefore all points of γ above the current point are above all walks in W, or δ is not the root walk and by definition of a wall there is another walk δ' in W to the left of δ's topmost point, and hence to the left of the current point of γ. Continuing in this way, we see that every point of γ has the desired property. This completes the proof. □

Consider a strict alternating cycle C in I_4 consisting of the pairs $(a_1, b_1), \ldots, (a_k, b_k)$ such that b_1 is drawn above all other b_i's. For $i \in \{2, \ldots, k\}$, we say that a wall W is an *i-wall for the cycle C* if for each element $p \in P$ that is included in some walk of W, there exists $\ell \in \{1, \ldots, i-1\}$ such that $a_\ell \leqslant p$ in P or $p \leqslant b_{\ell+1}$ in P. Note that if W is an i-wall for C then W is also a j-wall for C, for every $j \in \{i+1, \ldots, k\}$.

Say that a point p is *to the left of a wall W* if p is to the left of all walks in W intersecting the p-line, and there is at least one such walk.

Claim 3.1.4. *Let C denote a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ in I_4 such that b_1 is drawn topmost among the b_i's. If W is an i-wall for C for some $i \in \{2, \ldots, k\}$ and b_1 is to the left of W, then none of $a_i, \ldots, a_k$, $b_i, \ldots, b_k$ is to the right of some walk in W.*

Proof. We will prove the following two implications:

(i) If W is an i-wall for C with $i \in \{2, \ldots, k\}$ such that b_1 is to the left of W and a_i to the right of some walk in W, then b_{i+1} is also to the right of some walk in W. (Indices are taken cyclically, as always.)

(ii) If W is an i-wall for C with $i \in \{2, \ldots, k\}$ such that b_1 is to the left of W and b_i is to the right of some walk in W, then there is an i-wall W' for C with b_1 to the left of W' and a_i to the right of some walk in W'.

Note that these two statements together imply that if at least one of $a_i, \ldots, a_k$, $b_i, \ldots, b_k$ is to the right of a walk from an i-wall W for C with $i \in \{2, \ldots, k\}$ and with b_1 to the left of W, then there is a k-wall W' for C with a_k to the right of one of its walks and with b_1 to the left of W'. Then applying statement (i) again we get that b_1 lies also to the right of some walk of W', which is not possible. Hence, to establish our observation, it is enough to prove these two statements, which we do now.

For the proof of the first implication, suppose that we have an i-wall W for C with $i \in \{2, \ldots, k\}$ and with b_1 to the left of W and a_i to the right of some walk in W. Consider a walk γ_i witnessing the relation $a_i \leqslant b_{i+1}$ in P. Note that γ_i cannot intersect any walk in W. Indeed, otherwise their intersection would have some element p of P in common (by the planarity of the diagram), and we would have $a_i \leqslant p \leqslant b_{i+1}$ in P, which together with an extra comparability of the form $a_\ell \leqslant p$ in P or $p \leqslant b_{\ell+1}$ in P for some $\ell \in \{1, \ldots, i-1\}$ contradicts the fact that C is a strict alternating cycle. Since a_i is to the right of some walk in W, by Claim 3.1.3 we conclude that the topmost point of γ_i, namely b_{i+1}, is to the right of some walk in W, as desired. (Here we used that b_{i+1} is not above b_1 in the drawing; note that this argument applies even in the special case $i = k$.)

For the proof of the second implication, suppose that we have an i-wall W for C with $i \in \{2, \ldots, k\}$ with b_1 to the left of W and b_i to the right of some walk in W. Recall that $(a_i, b_i) \in I_4$, so there is a walk β from b_i going downwards to the a_i-line and intersecting it to the left of a_i. Let $W' := W \cup \{\beta\}$. Since b_i is the topmost point of β, given that W is an i-wall for C it should be clear that W' is also an i-wall for C and that a_i is to the right of $\beta \in W'$, as desired. This concludes the proof. □

Alternating Cycles with Special Pairs

Our next goal is to use walls and their properties to show that each strict alternating cycle in I_4 has at least one 'special pair': A pair (a_j, b_j) of a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ is said to be *special* if

(i) a_j is drawn below a_{j+1} and b_{j+1} is drawn below b_j;
(ii) a_{j+1} is to the right of every walk from a_j to b_{j+1}, and to the left of every walk from a_j to the b_j-line.

See Figure 3.4 for an illustration.

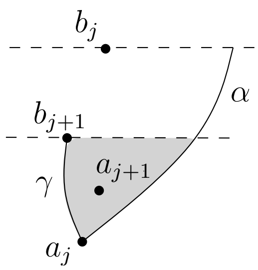

Figure 3.4: A special pair (a_j, b_j).

First we show that strict alternating cycles in $I_4^{\text{no-sep}}$ have a special pair.

Claim 3.1.5. *Let C denote a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ in $I_4^{\text{no-sep}}$ with b_1 drawn above all other b_i's. Then the pair (a_1, b_1) is special.*

Proof. Arguing by contradiction, suppose that (a_1, b_1) is not special. Since b_1 is drawn above b_2 and the pair (a_1, b_1) is not special, we can find a walk γ from a_1 to b_2 witnessing the relation $a_1 \leqslant b_2$ in P, and a walk α from a_1 to the b_1-line (hitting that line to the right of b_1) such that either

(a) a_2 is to the right of α, or
(b) a_2 is to the left of γ, or
(c) a_2 is drawn below a_1.

Let us start with an easy consequence of Claim 3.1.4. Since $\{\alpha\}$ is a 2-wall for C with b_1 to its left, we obtain in particular that a_2 and b_2 do not lie to the right of α. This already rules out the case (a), and it remains to find a contradiction when (b) or (c) hold.

Now that b_2 is not to the right of α, it has to be to the left of α as it is drawn above a_1, the bottommost point of α. This is used in the following partitioning of the plane. Let D be the curve obtained by starting in b_2 and going downwards along γ until the first intersection point r with α, at which point we switch to α and go upwards until its topmost point, which we denote q. (Note that the intersection of α and γ is not empty since a_1 belongs to both curves.) Extend D to the left by adding the horizontal half-line starting at b_2 going to the left, and to the right by adding the horizontal half-line starting at q going to the right. Note that D is not self-intersecting by our previous observation. Now let D' be the horizontal half-line starting at r going to the right. The removal of $D \cup D'$ defines three regions of the plane. We call the region consisting of all points between the r-line and the q-line that are to the right of α the *right* region. The remaining two regions are referred to as the *top* and *bottom* regions, in the natural way (see Figure 3.5).

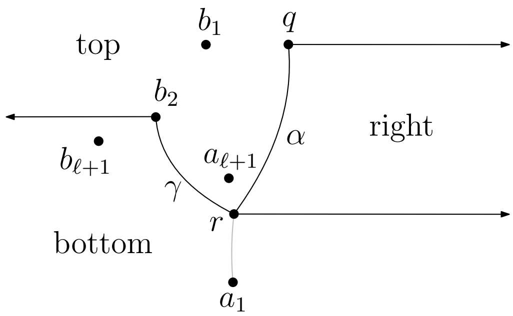

Figure 3.5: Illustration of the proof of Claim 3.1.5. The three regions and the placement of $a_{\ell+1}$, and $b_{\ell+1}$.

Clearly, neither a_1 nor b_1 is contained in the right region (recall that our regions do not include points from $D \cup D'$). Together with Claim 3.1.4 applied to the 2-wall $\{\alpha\}$, we obtain that:

None of $a_1, \ldots, a_k, b_1, \ldots, b_k$ lies in the right region. (3.1)

It follows from (3.1) that

a_2 is contained in the bottom region (3.2)

in both cases (b) and (c). We will show that this leads to a contradiction. (Whether we are in case (b) or (c) will not be used in the rest of the proof.)

Let ℓ be the largest integer such that $2 \leqslant \ell \leqslant k$ and a_i is in the bottom region for every $i \in \{2, \ldots, \ell\}$. Such an integer exists by (3.2). For each $i \in \{2, \ldots, \ell\}$, let β_i be a walk witnessing that b_i sees the left side of a_i (so a_i lies to the right of β_i). We claim that:

$$\text{The set } B_i := \{\beta_2, \ldots, \beta_i\} \text{ is an } i\text{-wall for } C \text{ with } \beta_2 \text{ being its root walk, for each } i \in \{2, \ldots, \ell\}. \text{ Moreover, } b_3, \ldots, b_{\ell+1} \text{ all lie in the bottom region.} \tag{3.3}$$

First we show that, if B_i ($i \in \{2, \ldots, \ell\}$) is an i-wall for C with β_2 as the root walk, then b_{i+1} is to the right of some walk in B_i and lies in the bottom region. To do so we consider a walk γ_i witnessing the comparability $a_i \leqslant b_{i+1}$ in P. We aim to show now that γ_i does not intersect the curve D.

The walk γ_i cannot intersect any walk from B_i, nor γ nor α. Indeed, if γ_i did intersect one of these walks, then there would be an element of P lying in their intersection (by the planarity of the diagram), and this would imply a non-existing comparability in the strict alternating cycle C. Since the bottommost point of γ_i (i.e. a_i) is to the right of β_i, we deduce from Claim 3.1.3 applied to B_i that γ_i cannot contain a point that is to the left of b_2 on the b_2-line (recall that b_2 is the topmost point of the root β_2). Therefore, γ_i cannot intersect the horizontal half-line starting at b_2 and going left. Clearly, γ_i is disjoint from the horizontal half-line starting at q and going right (since b_{i+1} is not drawn above b_1; in fact, b_{i+1} is drawn below b_1 unless $i = k$, in which case it is to the left of q). Thus, γ_i does not intersect the curve D.

Since a_i lies in the bottom region, we conclude that b_{i+1} lies in the bottom or in the right region. But by property (3.1) b_{i+1} cannot be in the right region, and hence it is in the bottom region, as claimed. Lastly, since a_i is to the right of the walk $\beta_i \in B_i$, using Claim 3.1.3 with the walk γ_i we deduce that b_{i+1} is also to the right of some walk in B_i, as desired.

Now we are ready to prove (3.3) by induction on i. The base case $i = 2$ is immediate since $B_2 = \{\beta_2\}$ is a 2-wall for C. For the inductive step, assume $i \geqslant 3$, and let us show that B_i has the desired property. By the induction hypothesis, B_{i-1} is an $(i-1)$-wall for C with β_2 as the root walk. With the observation from the previous paragraph this implies that b_i lies in the bottom region and to the right of some walk in B_{i-1}. This directly yields that $B_i = B_{i-1} \cup \{\beta_i\}$ is an i-wall for C rooted at β_2. This concludes the proof of (3.3).

Observe that a corollary of (3.3) is that $\ell < k$ (if $\ell = k$ then $b_{k+1} = b_1$ would lie in the bottom region by (3.3), which is clearly not the case).

We are now ready to get a final contradiction for the case under consideration, namely, that all pairs in our alternating cycle are in $I_4^{\text{no-sep}}$. Recall that, by definition of $I_4^{\text{no-sep}}$, none of $(a_1, b_1), \ldots, (a_k, b_k)$ admits a separator. Using the properties established above, we now exhibit a separator for the pair $(a_{\ell+1}, b_{\ell+1})$, which will be the desired contradiction.

First, note that $b_{\ell+1}$ is in the bottom region (by (3.3)) but that $a_{\ell+1}$ is not (by definition of ℓ). By (3.1), $a_{\ell+1}$ is not in the right region either. Since $\ell < k$, we also know that $a_{\ell+1} \neq a_1$, and hence that $a_{\ell+1}$ is not on $D \cup D'$. This implies that $a_{\ell+1}$ is contained in the top region. In particular, $a_{\ell+1}$ lies above the r-line and consequently so does $b_{\ell+1}$ (as $(a_{\ell+1}, b_{\ell+1}) \in I_2$). Since $b_{\ell+1}$ is in the bottom region and is above the r-line, $b_{\ell+1}$ must be to the left of γ. The element $a_{\ell+1}$, on the other hand, lies to the right of γ as this is the only place occupied by the top region below the position of $b_{\ell+1}$.

Finally, every element of P appearing on γ is incomparable to both $a_{\ell+1}$ and $b_{\ell+1}$ in P, as otherwise this would imply a non-existing comparability in C. This shows that the pair $(a_{\ell+1}, b_{\ell+1})$ is separated by γ, which is the contradiction we were looking for. $\square$

Let us consider strict alternating cycles in I_4^{sep} now. We start with a useful observation about i-walls with respect to such cycles.

Claim 3.1.6. *Let C denote a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ in I_4^{sep} such that b_1 is drawn above all other b_i's. Let W be an i-wall for C with $i \in \{2, \ldots, k\}$ such that b_1 is to the left of W and a_{i-1} belongs to some walk in W. Then for each walk λ from a_i to some point drawn below b_1 such that λ intersects the a_{i-1}-line and is disjoint from all walks in W, we have that the set $W \cup \{\lambda\}$ is an $(i+1)$-wall for C with b_1 to its left.*

Proof. Let λ be an arbitrary walk as in the statement of the claim and let p be the intersection point of λ with the a_{i-1}-line. (See Figure 3.6 for an illustration of the current situation and upcoming arguments.) First, we show that p is to the right of a_{i-1}. Consider a walk δ_i that separates the pair (a_i, b_i), and let γ_{i-1} be a walk witnessing $a_{i-1} \leqslant b_i$ in P. Observe that both λ and γ_{i-1} must be disjoint from δ_i, by the definition of a separator. Since γ_{i-1} ends in b_i, thus to the left of δ_i, we obtain that the whole walk γ_{i-1} stays to the left of δ_i. Similarly, as λ starts in a_i, thus to the right of δ_i, we get that λ stays to the right of δ_i. That is to say, δ_i separates γ_{i-1} and λ in the diagram.

Now, since all three walks intersect the a_{i-1}-line we conclude that γ_{i-1} intersects the a_{i-1}-line to the left of λ. In other words, p lies to the right of a_{i-1} on the a_{i-1}-line, and in particular to the right of some walk in W by assumption. Next, consider the portion λ' of λ that lies on or above the p-line. Since λ is disjoint from walks in W and since the bottommost point of λ' (which is p) is to the right of some walk in W, we deduce by Claim 3.1.3 that the topmost point of λ' is either to the right of some walk in W, or above all walks in W. However, by our assumptions the latter cannot happen as the b_1-line is hit by some walk in W but not by λ. Therefore, the topmost point of λ is to the right of some walk in W, and it follows that $W \cup \{\lambda\}$ is an $(i+1)$-wall for C with b_1 to its left. $\square$

In the next claim we show that strict alternating cycles in I_4^{sep} have a special pair.

Claim 3.1.7. *Let C denote a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ in I_4^{sep}. Then at least one of the pairs is special.*

Proof. We may assume that b_1 is drawn above all other b_i's. Let j be the smallest index in $\{1, \ldots, k\}$ such that a_{j+1} is drawn above a_j. (Such an index clearly exists.) We will show that (a_j, b_j) is a special pair of C.

Let γ be a walk from a_j to b_{j+1}, and let α be a walk from a_j to the b_j-line. We have to show that b_{j+1} is drawn below b_j, and that a_{j+1} is to the right of γ and to the left of α, as illustrated in Figure 3.7 (left).

First let us quickly deal with the $j = 1$ case. In this case, elements a_1, a_2, b_2, b_1 appear in this order in the drawing, from bottom to top. Also, we already know that

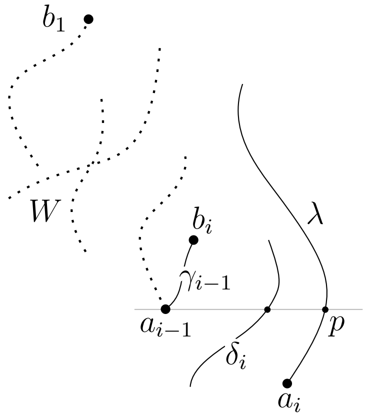

Figure 3.6: Situation in the proof of Claim 3.1.6. Dotted lines indicate walks in W.

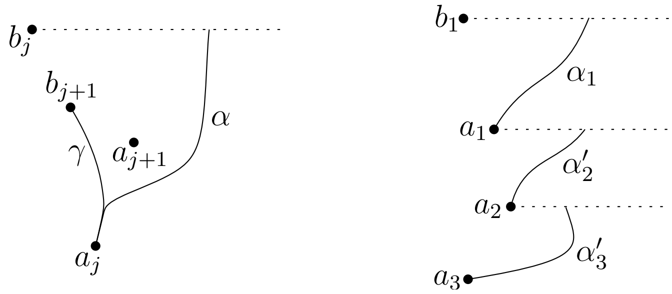

Figure 3.7: Positions of a_{j+1} and b_{j+1} (left), and the 4-wall A_3 (right).

a_2 is to the right of γ, since b_2 only sees the left side of a_2. Furthermore, $\{\alpha\}$ is a 2-wall for C with b_1 to its left. Hence, a_2 cannot be to the right of α by Claim 3.1.4, and therefore must be to the left of α. This proves the statement for $j = 1$.

Next, assume $j \geqslant 2$. We set up some walls as follows. For each $i \in \{1, \dots, j\}$, let α_i denote a walk from a_i to the b_i-line. Also, if $i \geqslant 2$, let α'_i denote the portion of α_i that goes from a_i to the a_{i-1}-line. Define $A_i := \{\alpha_1\} \cup \{\alpha'_2, \dots, \alpha'_i\}$, for each $i \in \{1, \dots, j\}$. See Figure 3.7 (right) for an illustration of how the walks in A_3 might look like.

We claim that A_i is an $(i+1)$-wall for C with b_1 to the left of A_i, which we prove by induction on i. For the base case $i = 1$, it is clear that $A_1 = \{\alpha_1\}$ is a 2-wall for C with b_1 to its left (since α_1 hits the b_1-line to the right of b_1). Now assume $i \geqslant 2$ for the inductive step. By induction, $A_{i-1} = \{\alpha_1, \alpha'_2, \dots, \alpha'_{i-1}\}$ is an i-wall for C with b_1 to the left of A_{i-1}. Applying Claim 3.1.6 to the i-wall A_{i-1} and the walk α'_i, which is clearly disjoint from all walks in A_{i-1}, we obtain that $A_{i-1} \cup \{\alpha'_i\} = A_i$ is an $(i + 1)$-wall for C with b_1 to its left, as desired.

Let us point out the following consequence of Claim 3.1.6 when applied to the j-wall A_{j-1}:

> If λ is a walk from a_j to some point drawn below b_1 such that λ hits the a_{j-1}-line and is disjoint from all walks in A_{j-1}, then a_{j+1} cannot be to the right of λ. (3.4)

Indeed, if a_{j+1} is to the right of λ, then a_{j+1} is to the right of a walk from the set $A_{j-1} \cup \{\lambda\}$, which is a $(j + 1)$-wall with b_1 to its left by Claim 3.1.6. However, this is not possible by Claim 3.1.4.

We may now conclude the proof. Recall that a_{j+1} is drawn above a_j. We already know that a_{j+1} is to the right of γ, since b_{j+1} only sees the left side of a_{j+1}. We need to show that b_{j+1} is drawn below b_j, and that a_{j+1} lies to the left of α.

Note that γ is disjoint from all walks in the j-wall A_{j-1} as each intersection would imply a non-existing comparability within the strict alternating cycle C. Thus if γ intersects the a_{j-1}-line, then property (3.4) applies to γ and yields that a_{j+1} is to the left of γ, which is not true. Thus, γ does not intersect the a_{j-1}-line. In particular, b_{j+1} is drawn below a_{j-1}, which is itself drawn below b_j. Hence, b_{j+1} is drawn below b_j.

It only remains to show that a_{j+1} is to the left of α. Let α' be the portion of α ranging from a_j to the a_{j-1}-line. Recall that a_{j+1} is drawn above a_j and below b_{j+1} (which itself is drawn below a_{j-1}). Since a_{j+1} cannot be on α' (because C is a strict alternating cycle), we deduce that a_{j+1} is either to the left of α' or to the right of α'. However, property (3.4) applies to α' and hence a_{j+1} must lie to the left of α'. Thus in particular, a_{j+1} is to the left of α. This completes the proof of the claim. □

Back to the Main Proof

We are now ready to go back to our main objective, which is to show that the following three sets can be partitioned into at most h reversible sets each: Non-dangerous pairs in I_4^{sep}, non-dangerous pairs in $I_4^{\text{no-sep}}$, and dangerous pairs in $I_4^{\text{no-sep}}$. (Recall that an incomparable pair $(a, b) \in I_4$ is said to be dangerous if a is drawn below x_0 and a sees the left side of x_0.) To do so, we define an auxiliary directed graph on each of these three sets: Suppose that J denotes one of these three sets. Define a directed graph G with vertex set J as follows. Given two distinct pairs $(a, b), (a', b') \in J$, we put an edge from (a, b) to (a', b') in G if there exists a strict alternating cycle $(a_1, b_1), \ldots, (a_k, b_k)$ in J and an index $j \in \{1, \ldots, k\}$ with $(a, b) = (a_j, b_j)$ and $(a', b') = (a_{j+1}, b_{j+1})$ such that (a_j, b_j) is special. In this case, we say that the alternating cycle is a *witness* for the edge $((a, b), (a', b'))$. Observe that when there is an edge from (a, b) to (a', b') in G then a is drawn below a'. This implies in particular that G has no directed cycle. Define the *chromatic number* $\chi(G)$ of G as the chromatic number of its underlying undirected graph. Since every strict alternating cycle in J has at least one special pair (c.f. Claims 3.1.5 and 3.1.7), we have

$$\dim(J) \leqslant \chi(G).$$

Thus it is enough to show that $\chi(G) \leqslant h$. We will do so by showing that every directed path in G has at most h vertices. This is done in the next two sections; the proofs will be different depending on whether the pairs in J are dangerous or not. In the meantime, let us make some observations on directed paths in G that hold independently of the type of pairs in J.

Let $(a^1, b^1), \ldots, (a^\ell, b^\ell)$ denote any directed path in G. (To avoid confusion, we use superscripts when considering directed paths and subscripts when considering alternating cycles.) Note that the a^i's go up in the diagram, while the b^i's go down: Indeed, for each $i \in \{1, \ldots, \ell - 1\}$ the two pairs (a^i, b^i) and (a^{i+1}, b^{i+1}) appear consecutively in that order in some strict alternating cycle in J where (a^i, b^i) is special, and thus a^{i+1} is drawn above a^i and b^{i+1} is drawn below b^i.

Now, for each $i \in \{1, \ldots, \ell - 1\}$, let γ^i denote an arbitrary walk from a^i to b^{i+1} witnessing the relation $a^i \leqslant b^{i+1}$ in P. The following observation will be useful in

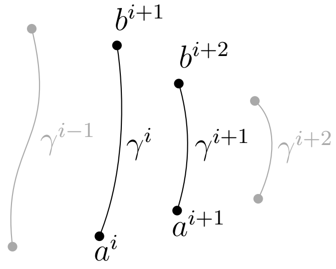

Figure 3.8: How the γ^i's look like.

our proofs:

$$\text{For each } i \in \{1,\ldots,\ell-2\}, \text{ the walk } \gamma^{i+1} \text{ is completely to the right of } \gamma^i. \quad (3.5)$$

By 'completely to the right', we mean that every point of γ^{i+1} is to the right of γ^i (as in Figure 3.8). To see this, recall once again that a^i is drawn below a^{i+1} and that b^{i+2} is drawn below b^{i+1}. Moreover, γ^i and γ^{i+1} cannot intersect because a^{i+1} and b^{i+1} are incomparable in P. Finally, a^{i+1} is to the right of γ^i, as follows from the fact that (a^i, b^i) is a special pair of some alternating cycle in J where $(a^i, b^i), (a^{i+1}, b^{i+1})$ appear consecutively in that order. Altogether, this implies that γ^{i+1} is completely to the right of γ^i.

We continue our study of directed paths in G in the next two sections.

Non-dangerous Pairs

In this section we consider pairs in I_4^{sep} and in $I_4^{\text{no-sep}}$ that are not dangerous. Our approach is independent of whether the pairs have separators or not, thus in this section we fix J as the subset of non-dangerous pairs in one of these two sets. We show that $\chi(G) \leqslant h$, where G denotes the directed graph on J defined earlier:

Claim 3.1.8. *Every directed path in G has at most h vertices, and hence* $\dim(J) \leqslant \chi(G) \leqslant h$.

Proof. Let $(a^1, b^1), \ldots, (a^\ell, b^\ell)$ denote any directed path in G. For $i \in \{1, \ldots, \ell-1\}$, let γ^i be a walk witnessing the relation $a^i \leqslant b^{i+1}$ in P that is 'leftmost' in the diagram among all such walks: For every point p of γ^i and every walk β witnessing that relation, either p is also on β, or p is to the left of β. (A little thought shows that γ^i is well defined, and uniquely defined.)

Let q be the largest index in $\{1, \ldots, \ell-1\}$ such that a^q is drawn below x_0 if there is such an index, otherwise set $q := 0$. To illustrate the usefulness of property (3.5), we use it to show that $h \geqslant q+1$. Assuming $q > 0$ (otherwise the claim is vacuous), we first note that x_0 is to the left of γ^1, since a^1 is drawn below x_0 and sees only its right side. (This is where we use that (a^1, b^1) is not dangerous.) Fix some walk α witnessing the relation $x_0 \leqslant b^\ell$ in P. Since a^i is below the x_0-line for each $i = 1, \ldots, q$, we directly deduce from (3.5) that α must intersect each of $\gamma^1, \ldots, \gamma^q$; see Figure 3.9. Given that these walks are disjoint, and that each intersection contains at least one element from P, counting x_0 we conclude that the height of P is at least $q+1$.

Thus we are already done with the proof if $q = \ell - 1$. Assume from now on that $q < \ell - 1$ (recall that possibly $q = 0$). Ideally, we would like to argue that

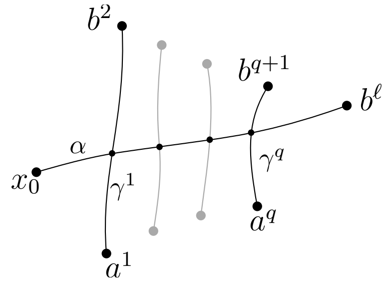

Figure 3.9: Each walk from x_0 to b^ℓ intersects the q walks $\gamma^1, \ldots, \gamma^q$.

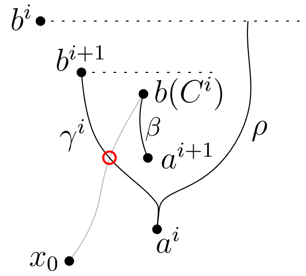

Figure 3.10: Situation if we assume that $b(C^i)$ is drawn below b^{i+1}. Considering any walk from x_0 to $b(C^i)$, we see a forbidden comparability in the poset.

α must intersect all of $\gamma^{q+1}, \ldots, \gamma^{\ell-1}$ as well, which would be enough to conclude the proof. However, this is not necessarily true: For instance, α could avoid γ^{q+1} by passing under a^{q+1} and then going to its right. In the rest of the proof, we will enrich our current structure—i.e. the walks $\gamma^1, \ldots, \gamma^{\ell-1}$—in a way that will allow us to conclude that α intersects at least $\ell - 1 - q$ disjoint walks, that are moreover disjoint from $\gamma^1, \ldots, \gamma^q$. This will show that α contains at least $q+1+(\ell-1-q) = \ell$ elements of P, and hence that $h \geqslant \ell$, as desired.

For each $i \in \{q+1, \ldots, \ell-1\}$, choose some strict alternating cycle C^i in J witnessing the edge $((a^i, b^i), (a^{i+1}, b^{i+1}))$. Thus $(a^i, b^i), (a^{i+1}, b^{i+1})$ appear consecutively in that order in C^i, and (a^i, b^i) is a special pair of C^i. The element b of the pair (a, b) appearing just after (a^{i+1}, b^{i+1}) in C^i will be important for our purposes, let us denote it by $b(C^i)$. (Note that $b(C^i) = b^i$ in case C^i is of length 2.) We show:

$$b(C^i) \text{ is drawn above } b^{i+1} \text{ for each } i \in \{q+1, \ldots, \ell-1\}. \tag{3.6}$$

To prove this, let $i \in \{q+1, \ldots, \ell-1\}$ and consider the alternating cycle C^i. First, if C^i has length 2, then $b(C^i) = b^i$, which is indeed drawn above b^{i+1}. Next, assume that C^i has length at least 3. Arguing by contradiction, suppose that $b(C^i)$ is drawn below b^{i+1}. Let ρ denote a walk starting in a^i and hitting the b^i-line to the right of b^i. Since (a^i, b^i) is a special pair of C^i, we know that a^{i+1} lies to the right of γ^i and to the left of ρ. We claim that $b(C^i)$ also lies to the right of γ^i and to the left of ρ, as depicted in Figure 3.10. Since $b(C^i)$ is drawn below b^{i+1} and b^i, both walks γ^i and ρ hit the $b(C^i)$-line. Let β denote a walk witnessing the comparability $a^{i+1} \leqslant b(C^i)$ in P. Now observe that both γ^i and ρ are disjoint from β as otherwise this would imply $a^i \leqslant b(C^i)$ in P, contradicting the fact that the alternating cycle C^i is strict (here we use that C^i has length at least 3). Given that a^{i+1} is to the right of γ^i and to the left of ρ, we deduce from this discussion that β is completely to the right of γ^i, and completely to the left of ρ. In particular, $b(C^i)$ is to the right of γ^i

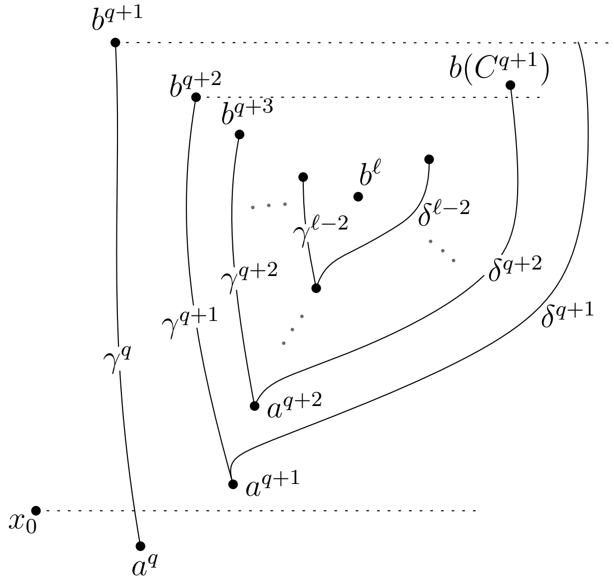

Figure 3.11: Nested structure built by the walks $\gamma^{q+1}, \ldots, \gamma^{\ell-2}$ and $\delta^{q+1}, \ldots, \delta^{\ell-2}$.

and to the left of ρ, as claimed. Since x_0 is drawn below a^i, this implies that any walk witnessing the comparability $x_0 \leqslant b(C^i)$ in P must intersect at least one of γ^i and ρ (by planarity of the diagram). It follows that $a^i \leqslant b(C^i)$ in P, contradicting the assumption that C^i is strict. This concludes the proof of (3.6).

We define the following additional walks: Let δ^{q+1} denote a walk from a^{q+1} to the b^{q+1}-line, and for each $i \in \{q+2, \ldots, \ell\}$, let δ^i denote a walk from a^i to $b(C^{i-1})$ witnessing the relation $a^i \leqslant b(C^{i-1})$ in P. (We remark that δ^{q+1} is defined slightly differently from $\delta^{q+2}, \ldots, \delta^\ell$; fortunately, this will not cause any complication in the arguments below.) In the remaining part of the proof, the key step will be to show that the different walks create together a 'nested structure' around b^ℓ with x_0 completely outside of it, as illustrated in Figure 3.11. Using this nested structure, we will easily conclude that the walk α intersects $\ell - 1 - q$ pairwise disjoint walks, each disjoint from $\gamma^1, \ldots, \gamma^q$ and from x_0, as desired.

For each $i \in \{q+1, \ldots, \ell\}$, let $\tilde{\delta}^i$ denote the portion of δ^i from a^i to the b^i-line. Thus $\tilde{\delta}^{q+1} = \delta^{q+1}$, and for $i \geqslant q+2$ the walk $\tilde{\delta}^i$ is a proper prefix of δ^i since $b(C^{i-1})$ is drawn above b^i by (3.6).

Let us point out some properties of the pair of walks γ^i and $\tilde{\delta}^{i+1}$ for $i \in \{q+1, \ldots, \ell-1\}$: First, γ^i and $\tilde{\delta}^{i+1}$ are disjoint as otherwise we would deduce that $a^{i+1} \leqslant b^{i+1}$ in P. Recall also that a^{i+1} lies to the right of γ^i, implying that $\tilde{\delta}^{i+1}$ starts to the right of γ^i. By the disjointedness of the two walks, this in turn implies that $\tilde{\delta}^{i+1}$ is completely to the right of γ^i (here we use that both walks have their top endpoints on the b^{i+1}-line).

Building on the observations above, we show:

(i) $\tilde{\delta}^{i+1}$ is completely to the left of $\tilde{\delta}^i$, for each $i \in \{q+1, \ldots, \ell-1\}$,

(ii) γ^i and $\tilde{\delta}^j$ are disjoint, for all $i \in \{1, \ldots, \ell-1\}$ and $j \in \{q+1, \ldots, \ell\}$ such that $i \neq j$.

Let us first prove (i). Since $\tilde{\delta}^i$ starts in a^i and ends on the b^i-line, we see that $\tilde{\delta}^i$ starts below the whole walk $\tilde{\delta}^{i+1}$ and ends above $\tilde{\delta}^{i+1}$. Using that (a^i, b^i) is a special pair of C^i, we also see that a^{i+1} is to the left of δ^i. Thus, to prove (i) it is enough to show that δ^i and δ^{i+1} are disjoint. Recall that δ^{i+1} witnesses the comparability $a^{i+1} \leqslant b(C^i)$ in P, since $i+1 \geqslant q+2$. If δ^i and δ^{i+1} intersect, then we deduce that $a^i \leqslant b(C^i)$ in P, in contradiction with the fact that C^i is a strict alternating cycle. So δ^i and δ^{i+1} must be disjoint.

Next we show (ii). First we consider the case $j \leqslant i-1$. By (i), $\tilde{\delta}^i$ is completely to the left of $\tilde{\delta}^{i-1}$. Iterating, we see that $\tilde{\delta}^i$ is completely to the left of $\tilde{\delta}^j$. On the other hand, every point p in $\tilde{\delta}^i$ is either also on γ^i, or is to the right of γ^i. This is because of the 'leftmost' property of γ^i: The walk $\tilde{\delta}^i$ starts in a^i and hits the b^{i+1}-line to the right of b^{i+1}; if some point of $\tilde{\delta}^i$ were to the left of γ^i then there would be a way to 're-route' a part of γ^i to the left using the appropriate part of $\tilde{\delta}^i$, contradicting the fact that γ^i has been chosen as the leftmost walk from a^i to b^{i+1}. We deduce that γ^i is completely to the left of $\tilde{\delta}^j$, and thus in particular that the two walks are disjoint.

Next we consider the case $j \geqslant i+1$. Here we use the fact that γ^j is completely to the right of γ^i (c.f. (3.5)): Since every point of γ^j is either on $\tilde{\delta}^j$ or to the left of it (by the leftmost property of γ^j, exactly as in the previous paragraph), we conclude that $\tilde{\delta}^j$ is also completely to the right of γ^i. In particular, they are disjoint. This completes the proof of (ii).

We are finally ready to show that $\ell \leqslant h$. Recall that we proved at the beginning of the proof that the walk α intersects each of $\gamma^1, \dots, \gamma^q$. Counting x_0, these intersections already single out $q+1$ elements of P on α. Using the walks we defined above, we now identify $\ell - q - 1$ extra elements of P on α, implying that P has height at least ℓ.

Let $i \in \{q+1, \dots, \ell-2\}$. The walks γ^i and $\tilde{\delta}^i$ both start in a^i. Recall also that γ^i hits the b^ℓ-line to the left of b^ℓ. Since $\tilde{\delta}^\ell$ hits the b^ℓ-line to the right of b^ℓ, we deduce from (i) that the same is true for $\tilde{\delta}^i$. Recall also that x_0 is drawn below a^i, since $i > q$. From these observations it follows that the walk α has to intersect the union U_i of the two walks γ^i and δ^i in an element of P. Using (ii) we see that U_i is disjoint from U_j for all $j \in \{q+1, \dots, \ell-2\}$ with $j \neq i$, and from $\gamma^1, \dots, \gamma^q$ as well.

Hence considering the intersection of α with U_i for $i = q+1, \dots, \ell-2$ we identify $\ell - q - 2$ 'new' elements of P on α. Finally, we can get an extra one by observing that element b^ℓ (the top endpoint of α) has not been counted so far. This concludes the proof. □

Dangerous Pairs

It only remains to partition the set J of dangerous pairs in $I_4^{\text{no-sep}}$ into reversible sets. As in the previous section, let us show that $\chi(G) \leqslant h$ for the directed graph G we defined on J:

Claim 3.1.9. *Every directed path in G has at most h vertices, and hence* $\dim(J) \leqslant \chi(G) \leqslant h$.

Proof. Let $(a^1, b^1), \dots, (a^\ell, b^\ell)$ denote any directed path in G. For each $i \in \{1, \dots, \ell-1\}$, we let C^i denote a strict alternating cycle in J that is witnessing the edge $((a^i, b^i), (a^{i+1}, b^{i+1}))$. Note that thanks to Claim 3.1.5 we may (and do) choose C^i so that b^i is drawn topmost among all elements of P appearing in C^i. In this proof, the element a of the pair (a, b) appearing just before (a^i, b^i) in C^i will play a special role, let us denote it by $a(C^i)$. (Observe that $a(C^i) = a^{i+1}$ in case C^i is of length 2.) Let δ^i denote a walk witnessing the relation $a(C^i) \leqslant b^i$ in P.

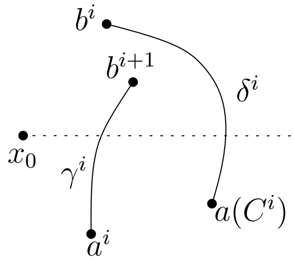

Figure 3.12: x_0 and b^{i+1} are to the left of δ^i.

A key property of the walks $\delta^1, \dots, \delta^{\ell-1}$ is the following:

$$x_0 \text{ and } b^{i+1} \text{ are both to the left of } \delta^i \text{ for each } i \in \{1, \dots, \ell-1\}. \tag{3.7}$$

Figure 3.12 illustrates this situation. To show this, let $i \in \{1, \dots, \ell-1\}$ and let $(c_1, d_1), \dots, (c_k, d_k)$ denote the pairs forming the cycle C^i in order, with $(c_1, d_1) = (a^i, b^i)$, $(c_2, d_2) = (a^{i+1}, b^{i+1})$, and $c_k = a(C^i)$.

Recall that for each $j \in \{2, \dots, k\}$ the element d_j is drawn above c_k (since the x_0-line lies in between c_k and d_j) and below d_1 (by our choice of C^i). In particular, d_j is either to the left or to the right of δ^i. (Note that d_j cannot be on δ^i since $j \neq 1$.) We will prove that all of $d_2, \dots, d_k$ are to the left of δ^i, and thus in particular $d_2 = b^{i+1}$ is.

We start by showing that d_k is to the left of δ^i. Suppose not. Let β be a walk witnessing the fact that d_k sees the left side of c_k. Since β's bottom endpoint is to the left of δ^i and its top endpoint d_k is to the right of δ^i, it follows that these two walks intersect (by planarity of the diagram). However, this implies $c_k \leqslant d_k$ in P, a contradiction. Hence d_k is to the left of δ^i as claimed.

Next, we observe that x_0 and d_j are on the same side of δ^i for each $j \in \{2, \dots, k\}$. For if x_0 and d_j are on different sides of δ^i, then any walk witnessing the relation $x_0 \leqslant d_j$ in P intersects δ^i, which in turn implies that $c_k \leqslant d_j$ in P, in contradiction with the fact that C^i is a strict alternating cycle. (Notice that the fact that $d_1 = b^i$ is drawn above d_j is crucial here.) Combining this with the fact that d_k is to the left of δ^i, we deduce that x_0 and $d_2, \dots, d_k$ are all to the left of δ^i, which establishes (3.7).

For $i \in \{1, \dots, \ell-1\}$, let γ^i be an arbitrary walk witnessing the relation $a^i \leqslant b^{i+1}$ in P, let $\tilde{\delta}^i$ denote the portion of δ^i extending from b^i to the x_0-line, and let similarly $\tilde{\gamma}^i$ denote the portion of γ^i from b^{i+1} to the x_0-line. Using (3.7), we aim to show that these walks build a 'nested structure' as depicted in Figure 3.13. To do so, we need to prove:

(i) $\tilde{\delta}^{i+1}$ is completely to the left of $\tilde{\delta}^i$, for each $i \in \{1, \dots, \ell-2\}$,

(ii) $\tilde{\gamma}^i$ is completely to the left of $\tilde{\delta}^i$, for each $i \in \{1, \dots, \ell-1\}$,

(iii) $\gamma^1, \dots, \gamma^{\ell-1}, \delta^{\ell-1}, \dots, \delta^1$ cross the x_0-line in this order from left to right,

(iv) x_0 is to the right of $\gamma^{\ell-2}$ and to the left of $\delta^{\ell-1}$.

In order to motivate these four properties, let us use them to show that $h \geqslant \ell$. For $i \in \{2, \dots, \ell-1\}$, let U_i denote the union of $\tilde{\gamma}^{i-1}$ and $\tilde{\delta}^i$, both of which have b^i as top endpoint. Recalling that γ^i is completely to the right of γ^{i-1} (c.f. (3.5)), it follows from (i), (ii), and (iii) that $U_2, \dots, U_{\ell-1}$ are pairwise disjoint, and moreover

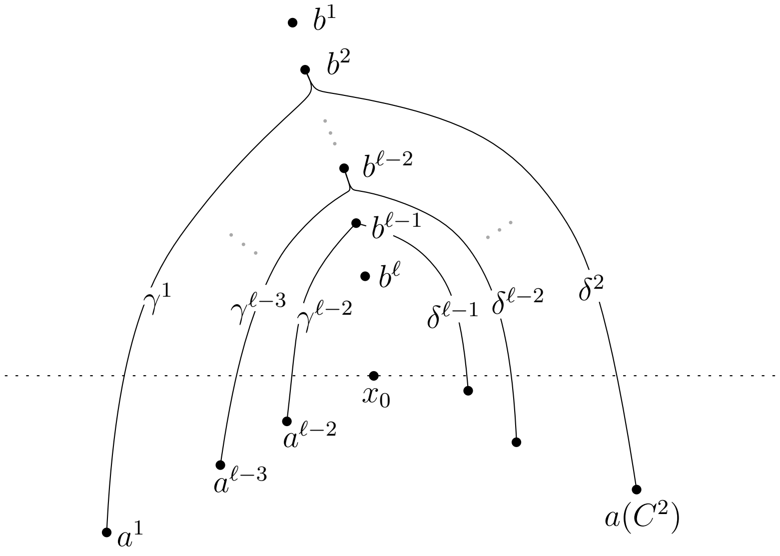

Figure 3.13: Nested structure formed by the walks $\gamma^1, \ldots, \gamma^{\ell-2}$ and $\delta^2, \ldots, \delta^{\ell-1}$.

that they form a nested structure: Say that a point p is *inside* U_i if p is to the right of $\tilde{\gamma}^{i-1}$ and to the left of δ^i. Then every point inside U_i is also inside $U_{i-1}, \ldots, U_2$. Since x_0 is inside $U_{\ell-1}$ by (iv), this implies that any walk witnessing the relation $x_0 \leqslant b^1$ in P intersects each of $U_2, \ldots, U_{\ell-1}$; see Figure 3.13 illustrating this situation. Counting x_0 and b^1, this shows that P has height at least ℓ, as desired.

Thus it only remains to prove the four properties above. To show (i), first note that δ^i and δ^{i+1} must be disjoint, for otherwise we would have $a(C^i) \leqslant b^{i+1}$ in P, contradicting the fact that C^i is a strict alternating cycle. Thus, using that b^{i+1} is to the left of $\tilde{\delta}^i$ by (3.7), we deduce that $\tilde{\delta}^{i+1}$ is completely to the left of $\tilde{\delta}^i$.

Property (ii) is a consequence of (3.7): b^{i+1} is to the left of δ^i by (3.7), and γ^i and δ^i cannot intersect as this would imply $a^i \leqslant b^i$ in P.

Property (iii) follows directly from (3.5), (i), and (ii).

Finally, to establish (iv), we only need to show that x_0 is to the right of $\gamma^{\ell-2}$ since we already know that x_0 is to the left of $\delta^{\ell-1}$ by (3.7). Arguing by contradiction, suppose that x_0 is to the left of $\gamma^{\ell-2}$. Here we exploit the fact that pairs in J are dangerous: Consider a walk β witnessing the fact that $a^{\ell-1}$ sees the left side of x_0. Since $(a^{\ell-2}, b^{\ell-2})$ is a special pair of $C^{\ell-2}$, the element $a^{\ell-1}$ is drawn above $a^{\ell-2}$ and is moreover to the right of $\gamma^{\ell-2}$; this situation is illustrated in Figure 3.14. Thus we deduce that β intersects $\gamma^{\ell-2}$. However, this implies that $a^{\ell-1} \leqslant b^{\ell-1}$ in P, a contradiction. This concludes the proof. □

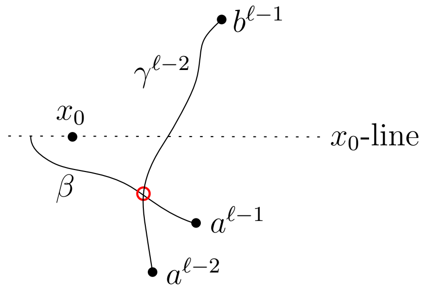

Figure 3.14: Situation if we assume that x_0 is to the left of $\gamma^{\ell-2}$. Considering a walk from $a^{\ell-1}$ to the left side of x_0, we obtain a forbidden comparability in P.

Hence, we conclude that $\dim(I_4) \leqslant 3h$ as claimed. This finishes the proof of Lemma 3.1.1.

3.2 A Linear Upper Bound on the Dimension

The whole section is devoted to a proof of the linear upper bound in Theorem 3.0.1.

Let P be a planar poset of height h. First, we apply the *Min-Max Reduction* (Lemma 2.1.1) to P and get a poset P' of height h such that $\dim(P) \leqslant \dim(\operatorname{Min}(P'), \operatorname{Max}(P'))$, and such that the cover graph of P' can be obtained by attaching degree-1 vertices to the cover graph of P. It is not hard to see that from the last condition and the planarity of P it follows that P' is also planar. In what follows we consider only the poset P', so with a slight abuse of notation let us simply write P for P' from now on. Our aim is to show that

$$\dim(\operatorname{Min}(P), \operatorname{Max}(P)) \leqslant 192h + 96,$$

which implies our main theorem.

Clearly, we may assume that $\dim(\operatorname{Min}(P), \operatorname{Max}(P)) > 192 + 96 = 288$, as otherwise we are done. As noted before (see for instance Chapter 2) we may assume that P is connected.

Fix a planar diagram of P and let G denote the cover graph of P embedded in the plane according to this drawing. For simplicity, we may assume without loss of generality that no two elements of P have the same y-coordinate in the drawing. We remark however that it is not essential to the arguments developed in this section that the diagram itself can be drawn in planar way, just that the cover graph can be. Planarity of the diagram itself will be needed only when invoking Lemma 3.1.1.

The following straightforward observation will be used several times implicitly in the proof: If Q is a convex subposet of P then our fixed planar diagram of P induces a planar diagram of Q.

Unfolding the Poset

Given a connected convex subposet Q of P, let $a(Q)$ denote the element of Q with smallest y-coordinate in the drawing. Clearly, $a(Q)$ is a minimal element of Q. Similarly, let $b(Q)$ denote the element of Q with largest y-coordinate in the drawing, which is a maximal element of Q.

Using the terminology introduced in Section 2.2, we iteratively unfold the poset P three times: Let $Q_0 := P$. For $i = 1, 2, 3$ let $c_{i-1} := a(Q_{i-1})$ and let Q_i be a c_{i-1}-core of Q_{i-1}. Note that

$$\begin{aligned}\dim(\operatorname{Min}(Q_2), \operatorname{Max}(Q_2)) &\geqslant \dim(\operatorname{Min}(Q_1), \operatorname{Max}(Q_1))/2 \\ &\geqslant \dim(\operatorname{Min}(P), \operatorname{Max}(P))/4 \geqslant 6,\end{aligned}$$

thus these cores can be defined (see the notes after the proof of the Unfolding Lemma in Chapter 2).

At least two of the cores Q_1, Q_2, Q_3 have the same type (left-facing or right-facing). Say this is the case for indices i and j, with $i < j$. We would like to focus on the left-facing case in the rest of the proof. We will be able to do so thanks to the following trick: Say Q_i and Q_j are right-facing. Then we turn our attention

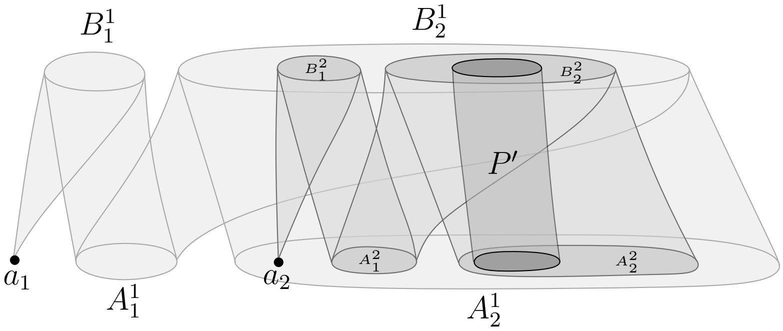

Figure 3.15: Unfolding P^1 and P^2.

to the dual poset P^d of P and its drawing obtained from that of P by flipping it. For $k = 1, 2, 3$, unfold Q^d_{k-1} from $c_{k-1} = b(Q^d_{k-1})$; notice that Q^d_k is a core of the unfolding, as follows from these two observations:

(i) if $\{x\}, B_1, A_1, \ldots, A_{m-1}, B_m$ is an unfolding of a poset R with $x \in \mathrm{Min}(R)$, then $\emptyset, \{x\}, B_1, A_1, \ldots, A_{m-1}, B_m$ is an unfolding of the dual poset R^d;

(ii) given a poset R and $A \subseteq \mathrm{Min}(R)$, $B \subseteq \mathrm{Max}(R)$ we have $\dim_R(A, B) = \dim_{R^d}(B, A)$.

In other words, we simply mirror the three unfoldings we did before. The key observation is that the type of the core Q^d_k is the opposite of that of Q_k. In particular, Q^d_i and Q^d_j are left-facing. We may then work with P^d and Q^d_i, Q^d_j instead of P and Q_i, Q_j.

To summarize, going to the dual poset if necessary, we may assume at this point that Q_i and Q_j are both left-facing, and that these two cores were obtained by unfolding Q_{i-1} and Q_{j-1} from c_{i-1} and c_{j-1} respectively. Observe that c_{i-1} is either the bottommost or the topmost point in the drawing of Q_{i-1}, and the same holds for c_{j-1} with respect to Q_{j-1}.

Let $x_1 := c_{i-1}$ and $x_2 := c_{j-1}$. (Thus, either $x_1, x_2 \in A$ or $x_1, x_2 \in B$.) Let $P^1 := Q_{i-1}$ and $P^2 := Q_{j-1}$. Also, let $P^3 := Q_j$ and $A^3 := \mathrm{Min}(P^3)$, $B^3 := \mathrm{Max}(P^3)$.

By the Unfolding Lemma (Lemma 2.2.1) and the definition of cores we have:

$$\dim(A^3, B^3) \geqslant \dim(\mathrm{Min}(P), \mathrm{Max}(P))/8.$$

Our goal is to show $\dim(A^3, B^3) \leqslant 24h + 12$, which implies our main theorem.

Before pursuing further, let us emphasize that P^3, P^2, P^1 form an increasing sequence of *convex* subposets of P. In particular, they all are planar posets, and the drawing of P induces in a natural way drawings of their respective diagrams. It is perhaps also good to recall that we do not need to specify whether $\dim(A^3, B^3)$ is to be understood with respect to P^3, P^2, P^1, or P as this is the same quantity.

Let $A^1_0, B^1_1, A^1_1, \ldots, A^1_{m_1-1}, B^1_{m_1}$ and $A^2_0, B^2_1, A^2_1, \ldots, A^2_{m_2-1}, B^2_{m_2}$ be sequences obtained by unfolding P^1 and P^2 from x_1 and x_2, respectively. It follows from our left-facing assumption that there are indices $k \geqslant 1$ and $\ell \geqslant 1$ such that P^2 is contained in $\mathrm{conv}_{P^1}(A^1_k \cup B^1_k)$ and P^3 is contained in $\mathrm{conv}_{P^2}(A^2_\ell \cup B^2_\ell)$. Figure 3.15 illustrates this in the case $k = \ell = 2$.

Let us point out that x_1 is not in P^2. This can be seen as follows: First, note that either $A^1_0 = \{x_1\}$, or $A^1_0 = \emptyset$ and $B^1_1 = \{x_1\}$. In the first case, it is clear that x_1 is not in P^2 simply because x_1 will not be included in $\mathrm{D}(B^1_k) \cap \mathrm{U}(A^1_k)$. In the second case, x_1 is not in P^2 because $\mathrm{Inc}(A^1_1, B^1_1) = \emptyset$, thus $\dim(A^1_1, B^1_1) = 1$, and hence $k \geqslant 2$ in that case. The same observation holds for x_2 with respect to P^3. We record these two facts, for emphasis:

$$x_1 \notin P^2 \quad \text{and} \quad x_2 \notin P^3. \tag{3.8}$$

Red Tree and Blue Tree

Let G^3 denote the cover graph of P^3 embedded in the plane according to our fixed drawing of P. Since neither x_1 nor x_2 is in P^3 (c.f. (3.8)) and both elements have either smallest or largest y-coordinates among elements of P^1 and P^2, respectively, we see that x_1 and x_2 are each drawn either below all elements of G^3 or above all of them. In particular, x_1 and x_2 are drawn in the outer face of G^3.

We will define two specific subgraphs of G, the cover graph of P. Both subgraphs will be trees and will have the property that all their internal nodes are drawn in the outer face of G^3 and all their leaves are drawn on the boundary of that outer face. The first tree is rooted at x_1 and will be colored in red, while the second one is rooted at x_2 and will be colored in blue. See Figure 3.16 for an illustration of how the trees will look like. To define the trees we will use zig-zag paths (which have been introduced in Section 2.3) with respect to our fixed unfoldings of P^1 and P^2. Given an element $v \in P^3$, for simplicity we denote by $Z^1(v)$ and $Z^2(v)$ the zig-zag paths $Z^{x_1}_{P^1}(v)$ and $Z^{x_2}_{P^2}(v)$, respectively.

Now, consider an element $b \in B^3$ and its zig-zag path $Z^1(b)$ in P^1 connecting b to x_1. Let x be the element of P^3 on that path that is closest to x_1. Recall that x_1 is not in P^3 by (3.8), thus $x \neq x_1$. Since P^3 is a core of P^2, and in particular the element set of P^3 is $\mathrm{U}_{P^2}(A^3) \cap \mathrm{D}_{P^2}(B^3)$, there is an element $a \in A^3$ such that $a \leqslant b$ in P^3. This shows that $b \in \mathrm{U}_{P^1}(A^1_k)$, since $a \in A^3 \subseteq A^1_k$. We also know that $b \in B^3 \subseteq B^1_k \subseteq \mathrm{U}_{P^1}(A^1_{k-1})$ (by definition of unfolding). Hence, we may apply Claim 2.3.2(i) on the poset P^1 and element b with respect to the unfolding sequence $A^1_0, B^1_1, A^1_1, \ldots, A^1_{m_1-1}, B^1_{m_1}$ of P^1. By this claim, we know that $Z^1(b) \cap \mathrm{U}_{P^1}(A^1_k)$ is a chain C with the topmost element being b and $C \subseteq \mathrm{U}_{P^1}(A^1_{k-1})$. Observe in particular that x is in this chain C.

No element of the zig-zag path $Z^1(x)$ is in P^3 except for x, by the choice of x. Thus, since x_1 is drawn in the outer face of G^3, it follows that the whole path $Z^1(x)$ is drawn in that outer face, except for its endpoint x which is on its boundary. We call x the *red exit point* of b. Since x is in P^3 and in the chain $C \subseteq \mathrm{U}_{P^1}(A^1_{k-1})$ we obtain the following obvious relations, which we record for future reference:

$$x \in \mathrm{D}_{P^1}(B^1_k), \qquad x \in \mathrm{U}_{P^1}(A^1_k), \quad \text{and} \quad x \in \mathrm{U}_{P^1}(A^1_{k-1}). \tag{3.9}$$

Let X be the set of red exit points of elements in B^3. Let $T^1 := \bigcup_{x \in X} Z^1(x)$. Then T^1 is a tree with X as set of leaves, as follows from Claim 2.3.1. We color the vertices and edges of T^1 in red and refer to T^1 as the *red tree*. Note that in our drawing of cover(P), the red tree is drawn in the outer face of G^3 with its leaves on the boundary, as in Figure 3.16.

By considering P^2 instead of P^1 in the definition of red exit points, we analogously define *blue exit points* for elements in B^3: Given $b \in B^3$, the blue exit point of b is the element y of P^3 on the zig-zag path $Z^2(b)$ that is closest to x_2. This element y has the following properties:

$$y \in \mathrm{D}_{P^2}(B^2_\ell), \qquad y \in \mathrm{U}_{P^2}(A^2_\ell), \quad \text{and} \quad y \in \mathrm{U}_{P^2}(A^2_{\ell-1}). \tag{3.10}$$

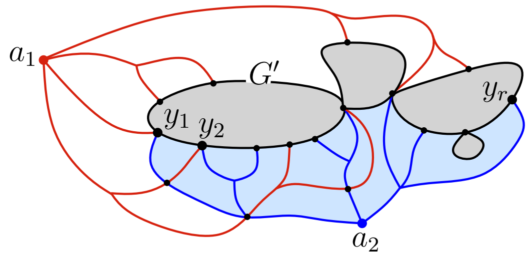

Figure 3.16: The graph G^3 together with the red and blue trees. The part outside G^3 is not drawn according to our fixed embedding of the diagram but has instead been redrawn freely for clarity, relying only on the planarity of the cover graph.

We let Y denote the set of blue exit points of elements in B^3, and let $T^2 := \bigcup_{y \in Y} Z^2(y)$ denote the tree defined by the union of the zig-zag paths of these exit points. Y is thus the set of its leaves, and the tree is drawn in the outer face of G^3, except for its leaves which are on its boundary. We color the vertices and edges of T^2 in blue, and refer to T^2 as the *blue tree*.

Recall that we wish to show that $\dim(A^3, B^3) \leqslant 24h + 12$. Before pursuing further, we pause to remark that Lemma 3.1.1 already shows that $\dim(A^3, B^3) \leqslant |Y|(6h+3)$: For each element of B^3, consider the corresponding blue exit point. Partition B^3 into $|Y|$ sets according to these blue exit points, and apply Lemma 3.1.1 on P^3 with each of these subsets of B^3 and the corresponding blue exit point in Y (playing the role of x_0 in that lemma). Thus, we may assume that $|Y| \geqslant 5$, as otherwise we are done.

Using the Red Tree and the Blue Tree

In this section, we use our newly defined trees to partition a certain subset of $\mathrm{Inc}(A^3, B^3)$ into a bounded number of reversible sets. As a result, the remaining incomparable pairs will have some extra structure, which we will then exploit in the next section.

Let H denote the plane graph obtained by taking the union of G^3 and the blue tree. The edges on the boundary of the outer face of G^3 that are *not* on the boundary of the outer face of H together define a facial walk along the boundary of G^3 corresponding to the part of the boundary that is "inside" H; let $\vec{F}_{\text{in}}$ denote this facial walk oriented counter-clockwise. Recalling that Y is the set of leaves of the blue tree, we see that $\vec{F}_{\text{in}}$ begins at some vertex from Y and ends at another vertex from Y. Enumerate the elements of Y as $y_1, \ldots, y_r$ in order of their appearance in the facial walk $\vec{F}_{\text{in}}$.

Let R denote the bounded region of the plane defined by the union of the paths $Z^2(y_1)$ and $Z^2(y_r)$, and the facial walk $\vec{F}_{\text{in}}$. (This is the blue-shaded region in Figure 3.16.) Clearly, x_1 is not in R since x_1 is the element of P^1 that is drawn topmost or bottommost. Let W denote the union of $\{y_1, y_r\}$ with the set of red exit points $x \in X$ that are such that the path $Z^1(x)$ leaves x inside R. (For instance, in Figure 3.16 there are five elements in W.)

Our aim is to partition the set $\mathrm{Inc}(A^3, B^3 \cap \mathrm{U}_{P^2}(W))$ into at most $12h + 6$ reversible sets. To do so, the two zig-zag paths $Z^2(y_1)$ and $Z^2(y_r)$ will play an important role. Recall that $Y \subseteq \mathrm{D}_{P^2}(B^2_\ell)$, $Y \subseteq \mathrm{U}_{P^2}(A^2_\ell)$, and $Y \subseteq \mathrm{U}_{P^2}(A^2_{\ell-1})$ by (3.10). In particular, we may apply Claim 2.3.2(ii) on P^2 with element y_1 and unfolding sequence $A^2_0, B^2_1, A^2_1, \ldots, A^2_{m_2-1}, B^2_{m_2}$. By this claim, there is an element z_1 of P^2 such that $Z^2(y_1) \cap \mathrm{D}_{P^2}(B^2_\ell) \subseteq \mathrm{U}_{P^2}(z_1)$. Doing the same with element y_r instead of y_1, we find an element z_r of P^2 such that $Z^2(y_r) \cap \mathrm{D}_{P^2}(B^2_\ell) \subseteq \mathrm{U}_{P^2}(z_r)$. Let us repeat the properties of z_1 and z_r, for future reference:

$$Z^2(y_1) \cap \mathrm{D}_{P^2}(B^2_\ell) \subseteq \mathrm{U}_{P^2}(z_1) \quad \text{and} \quad Z^2(y_r) \cap \mathrm{D}_{P^2}(B^2_\ell) \subseteq \mathrm{U}_{P^2}(z_r). \tag{3.11}$$

Or in words: In P^2, the element z_1 is smaller or equal to each element of the zig-zag path $Z^2(y_1)$ that is in the downset of B^2_ℓ, and the same holds for z_r with respect to y_r.

The key observation that we will use to partition $\mathrm{Inc}(A^3, B^3 \cap \mathrm{U}_{P^2}(W))$ into a small number of reversible sets is that the union of the upsets of z_1 and z_2 contains all of W:

Claim 3.2.1. $W \subseteq \mathrm{U}_{P^2}(\{z_1, z_r\})$.

Proof. It follows from (3.11) that $y_1, y_r \in \mathrm{U}_{P^2}(\{z_1, z_r\})$, so let us consider an element $x \in W - \{y_1, y_r\}$. First we observe that the zig-zag path $Z^1(x)$ leaves its endpoint x inside the region R and ends in x_1, which is not R. Hence, $Z^1(x)$ intersects $Z^2(y_1)$ or $Z^2(y_r)$ in an element distinct from x. Let y denote such an element.

If $y \in \mathrm{D}_{P^2}(B^2_\ell)$, then we deduce from (3.11) that $x \in \mathrm{U}_{P^2}(\{z_1, z_r\})$ since y is in at least one of $Z^2(y_1)$ and $Z^2(y_r)$. Thus it is enough to show that $y \in \mathrm{D}_{P^2}(B^2_\ell)$, which we do now.

Since x is a red exit point, we have $x \in \mathrm{D}_{P^1}(B^1_k)$, $x \in \mathrm{U}_{P^1}(A^1_k)$, and $x \in \mathrm{U}_{P^1}(A^1_{k-1})$ by (3.9). Note also that $y \in \mathrm{D}_{P^1}(B^1_k)$ and $y \in \mathrm{U}_{P^1}(A^1_k)$ since $y \in P^2$. Using Claim 2.3.2(i) on P^1 and zig-zag path $Z^1(x)$, we then see that $y \leqslant x$ in P^1. This relation also holds in P^2 since x and y are both in P^2 (and P^2 is an induced subposet of P^1). Recalling that $x \in \mathrm{D}_{P^3}(B^3) \subseteq \mathrm{D}_{P^2}(B^2_\ell)$, we deduce that $y \in \mathrm{D}_{P^2}(B^2_\ell)$, as desired. □

A consequence of the above claim is that we can partition $B^3 \cap \mathrm{U}_{P^2}(W)$ into B^3_1 and B^3_r in such a way that $z_1 \leqslant b$ in P^2 for all $b \in B^3_1$ and $z_r \leqslant b$ in P^2 for all $b \in B^3_r$. We may now apply Lemma 3.1.1 on the planar poset P^2 with the set $B^3_1 \subseteq \mathrm{Max}(P^2)$ and element z_1 (playing the role of x_0 in that lemma), and deduce that $\mathrm{Inc}(\mathrm{Min}(P^2), B^3_1)$ can be partitioned into at most $6h+3$ reversible sets. In particular, we obtain a partition of $\mathrm{Inc}(A^3, B^3_1)$ into at most $6h+3$ reversible sets. Proceeding similarly with the set B^3_r and element z_r yields a partition of $\mathrm{Inc}(A^3, B^3_r)$ into at most $6h+3$ reversible sets. Therefore, we have a partition of $\mathrm{Inc}(A^3, B^3 \cap \mathrm{U}_{P^2}(W))$ into at most $12h+6$ reversible sets, as desired.

Finishing the Proof

It remains to deal with incomparable pairs $(a, b) \in \mathrm{Inc}(A^3, B^3)$ such that $b \in B^3 - \mathrm{U}_{P^2}(W)$. We remark that $B^3 - \mathrm{U}_{P^2}(W) = B^3 - \mathrm{U}_{P^3}(W)$, since W is a set of elements of P^3 and P^3 is a convex subposet of P^2. Let $I' := \mathrm{Inc}(A^3, B^3 - \mathrm{U}_{P^3}(W))$, and let us assume that I' is not empty (otherwise we are done). Let $X' := X - W$ and $Y' := Y - W$, both of which are non empty. For each pair $(a, b) \in I'$ we have that $b \in \mathrm{U}_{P^3}(X')$ and $b \in \mathrm{U}_{P^3}(Y')$; furthermore, since $b \notin \mathrm{U}_{P^3}(W)$ we also know that $b \notin \mathrm{U}_{P^3}(y_1)$ and $b \notin \mathrm{U}_{P^3}(y_r)$. In this section we show that I' can be partitioned into at most $12h + 6$ reversible sets as well, which will thus conclude the proof of our main theorem.

Observe that for each $x \in X'$ the zig-zag path $Z^1(x)$ leaves its endpoint x outside the region R, while for each $y \in Y'$ the path $Z^2(y)$ leaves y inside R. Let F^{ac} denote the facial walk on the outer face boundary of G^3 going from y_1 to y_r in anticlockwise order that spans all the edges that are in R but no more. Let F^c denote the 'opposite' facial walk going in clockwise order from y_1 to y_r spanning all the edges of the outer face boundary that are not in R (and no more). Enumerate the elements of X' as $x'_1, \dots, x'_s$ in the order of first appearance in the walk F^c. (Some elements could appear multiple times, in which case they are cutvertices of G^3.) Enumerate similarly the elements of Y' as $y'_1, \dots, y'_t$ in order of their first appearance in F^{ac}.

Let us point out that the two sets X' and Y' are not necessarily disjoint. As is easily checked, every element that appears in both sets is a cutvertex of G^3. We note that in the rest of the proof we focus exclusively on the poset P^3 and will not use P^1, P^2, nor our colored trees anymore. Instead, we will use the two facial walks F^{ac} and F^c to help us reverse pairs in I'.

For each element $z \in \mathrm{U}_{P^3}(X')$ choose a corresponding element $x \in X'$ with $x \leqslant z$ in P^3 and a path in G^3 witnessing that relation, which we call an *X'-path*. We may and do assume that all the choices of X'-paths have been made in such a way that the union of all the X'-paths forms a forest consisting of exactly one tree per element $x \in X'$. (Thus every two X'-paths that intersect do so in a common prefix of the two paths starting in their endpoint in X'.) Replacing X' with Y' in the above definition, we similarly define Y'-paths for all elements $z \in \mathrm{U}_{P^3}(Y')$, whose union is a forest consisting of exactly one tree per element $y \in Y'$.

Let B^* be the subset of $B^3 - \mathrm{U}_{P^3}(W)$ consisting of the b's appearing in pairs $(a, b) \in I'$. For each $b \in B^*$, we will define a path $Q(b)$ in G^3 that separates y_1 from y_r in G^3. A key property of $Q(b)$ will be that all its elements are contained in the downset of b. To define $Q(b)$ we first need to specify its *midpoint* $m(b)$: Choose arbitrarily one element from the set $\mathrm{D}_{P^3}(b) \cap \mathrm{U}_{P^3}(X') \cap \mathrm{U}_{P^3}(Y')$ that is minimal in the subposet it induces. Note that that set is not empty since it contains at least b. Next, we let $Q^{X'}(b)$ denote the X'-path of the midpoint $m(b)$, and let $Q^{Y'}(b)$ denote its Y'-path. The path $Q(b)$ is then defined as the union of these two paths. Note that this union is indeed a path because $Q^{X'}(b)$ and $Q^{Y'}(b)$ share $m(b)$ as one of their two endpoints but have no other element in common (otherwise $m(b)$ would not be minimal in $\mathrm{D}_{P^3}(b) \cap \mathrm{U}_{P^3}(X') \cap \mathrm{U}_{P^3}(Y')$). Finally, let $x(b)$ denote the endpoint of $Q^{X'}(b)$ that is in X', and define $y(b)$ similarly with respect to Y'.

Recall that $b \notin \mathrm{U}_{P^3}(y_1)$ and $b \notin \mathrm{U}_{P^3}(y_r)$. Hence, $Q(b)$ avoids both y_1 and y_r. We remark that, since X' and Y' could share some elements, the path $Q(b)$ could possibly consist of only its midpoint $m(b)$ and nothing else. In that case, $m(b)$ belongs to both X' and Y' and is a cutvertex of G^3. (In Figure 3.17, this could be the cutvertex on the right for instance.) This somewhat degenerated case will not cause any difficulties in the proof.

A key observation is that not only the path $Q(b)$ avoids both y_1 and y_r, it moreover separates y_1 from y_r in G^3, as illustrated in Figure 3.17. This is clear if the outer face of G^3 bounds a cycle of G^3, because $y_1, x(b), y_r, y(b)$ appear then in that order on the cycle (clockwise). More generally, even if the outer face boundary contains cutvertices, one can observe that all occurrences of $x(b)$ in the facial walk F^c appear after all occurrences of y_1 and before all occurrences of y_r in that walk

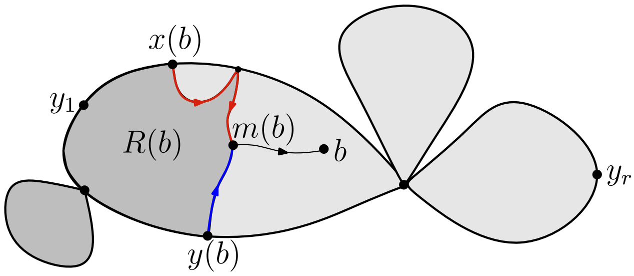

Figure 3.17: The X'-path $Q^{X'}(b)$ (in red), the Y'-path $Q^{Y'}(b)$ (in blue), and their union $Q(b)$.

(for if not, $Q^{X'}(b)$ would contain y_1 or y_r). Similarly, all occurrences of $y(b)$ in the facial walk F^{ac} appear after all occurrences of y_1 and before all occurrences of y_r in that walk (otherwise, $Q^{Y'}(b)$ would contain y_1 or y_r). Hence we see that $Q(b)$ does indeed separate y_1 from y_r in G^3. (Again, in a degenerated case $Q(b)$ is just a cutvertex of G^3 separating y_1 from y_r, which is perfectly fine for our purposes.)

As illustrated in Figure 3.17, the path $Q(b)$ defines a corresponding region $R(b)$ of the plane, which we define precisely now. Let $Q^c(b)$ denote the prefix of the facial walk F^c obtained by stopping it at the first occurrence of $x(b)$. Similarly, let $Q^{ac}(b)$ denote the prefix of the facial walk F^{ac} obtained by stopping it at the first occurrence of $y(b)$. Define $R(b)$ as the region of the plane bounded by the union of $Q(b)$, $Q^c(b)$, and $Q^{ac}(b)$ (see Figure 3.17 again). Note that the path $Q(b)$ separates every vertex of G^3 that is in $R(b)$ from every vertex not in $R(b)$. This will be used in the proof of the following claim.

Claim 3.2.2. *Suppose that b and b' are two elements in B^* with b' drawn outside $R(b)$. Then $R(b) \subseteq R(b')$. Moreover, equality holds if and only if $m(b) = m(b')$.*

Proof. First observe that if $m(b') = m(b)$, then $Q(b) = Q(b')$ and thus $R(b) = R(b')$. So assume $m(b') \neq m(b)$. We need to prove that $R(b)$ is a strict subset of $R(b')$. To do so, we first show that $m(b')$ is not in $R(b)$.

Arguing by contradiction, we suppose that it is contained in $R(b)$. Consider a covering chain Q of G' witnessing the comparability $m(b') \leqslant b'$ in P'. Since b' is not in $R(b)$ but so is $m(b')$, the path $Q(b)$ separates $m(b')$ from b' in G'. Thus, Q intersects $Q(b)$. Let p be an element in their intersection. Clearly, $m(b') \leqslant p \leqslant b'$ in P'. It follows that p is in the upsets of both X' and Y' (since this is true for $m(b')$). This in turn implies that $p = m(b)$, because by the definition of $m(b)$ no other element from $Q(b)$ lies in both upsets (otherwise this would contradict the minimality of $m(b)$). However, since $m(b') \neq m(b)$, we deduce that $m(b') < p = m(b)$ in P', which contradicts the minimality of $m(b)$ (since $m(b')$ is in the upsets of X' and Y', and also in the downset of b). Hence, we conclude that $m(b')$ is not in $R(b)$, as claimed.

For the rest of the proof we aim to show that the whole boundary of $R(b)$ is contained in $R(b')$. This would imply $R(b) \subseteq R(b')$, and this inclusion would be strict since $m(b')$ is not in $R(b)$.

So all we need to show is that $Q^c(b)$, $Q^{ac}(b)$, and $Q(b)$ are contained in $R(b')$. We start to prove $Q^c(b) \subseteq Q^c(b')$. If $x(b')$ is not on $Q^c(b)$, then $Q^c(b')$ (which ends in $x(b')$) is a longer facial walk starting from y_1 than $Q^c(b)$, implying that $Q^c(b) \subseteq Q^c(b')$. If $x(b') = x(b)$ then we clearly have $Q^c(b') = Q^c(b)$. So suppose we have $x(b') \in Q^c(b)$ and $x(b') \neq x(b)$. Since $m(b')$ is outside $R(b)$ but $x(b')$ is inside, we deduce that $Q^{X'}(b')$ has to intersect $Q(b)$ in this case. However, it cannot intersect $Q^{Y'}(b) - \{m(b)\}$ as no element from this set is in the upset of X', and it also cannot intersect $Q^{X'}(b)$ as this implies $x(b) = x(b')$ by our definition of X'-paths. Thus, we indeed have $Q^c(b) \subseteq Q^c(b')$. Symmetrically, we also obtain $Q^{ac}(b) \subseteq Q^{ac}(b')$.

It remains to show $Q(b) \subseteq R(b')$. We just observed that the two paths $Q^{X'}(b')$ and $Q^{Y'}(b)$ cannot intersect. Similarly, the two paths $Q^{Y'}(b')$ and $Q^{X'}(b)$ do not intersect either. Moreover, recall that if $Q^{X'}(b)$ and $Q^{X'}(b')$ intersect, then they do so in a common prefix starting from their common endpoint in X', and thus the rest of $Q^{X'}(b')$ lies outside of $R(b)$ (since $m(b')$ is outside that region). The two paths $Q^{Y'}(b)$ and $Q^{Y'}(b')$ behave similarly. Since $Q(b')$ is a path connecting vertices that lie on the outer face cycle, all this shows that $Q(b)$ must be contained in $R(b')$. Therefore, we conclude that $R(b) \subsetneq R(b')$ and the claim is proven. □

We will also need a dual version of the above claim. The proof is roughly identical up to exchanging the notions of inside $R(b)$ and outside $R(b)$, and is thus left to the reader.

Claim 3.2.3. *Suppose that b and b' are two elements in B^* with b' drawn inside $R(b)$. Then $R(b') \subseteq R(b)$. Moreover, equality holds if and only if $m(b) = m(b')$.*

We are now ready to partition I' into a bounded number of reversible sets. For each $x \in P^3$ let $B_x := \{b \in B^* \mid m(b) = x\}$ and $I_x := \{(a, b) \in I' \mid b \in B_x\}$. Since $B_x \subseteq \mathrm{U}_{P^3}(x)$ we may apply our main lemma, Lemma 3.1.1, on P^3 with the set B_x and element x to obtain a partition of I_x into at most $6h + 3$ reversible sets. Let $I_x^1, \ldots, I_x^{6h+3}$ denote such a partition (with possibly some empty sets).

Next, for each $i \in \{1, \ldots, 6h+3\}$ let $I^i := \bigcup_{b \in B^*} I^i_{m(b)}$. Observe that $I^1, \ldots, I^{6h+3}$ is a partition of I' (again, with possibly some empty sets). For each i we split I^i into two sets $I^{i,\text{in}}$ and $I^{i,\text{out}}$. The set $I^{i,\text{in}}$ contains the pairs $(a, b) \in I^i$ with a drawn in the region $R(b)$, while $I^{i,\text{out}}$ contains the remaining one. We show that both subsets are reversible, starting with $I^{i,\text{out}}$:

Claim 3.2.4. *$I^{i,\text{out}}$ is reversible for each $i \in \{1, \ldots, 6h + 3\}$.*

Proof. Arguing by contradiction, suppose that $I^{i,\text{out}}$ is not reversible for some $i \in \{1, \ldots, 6h+3\}$. Then there is a sequence of pairs $(a_1, b_1), \ldots, (a_k, b_k)$ in $I^{i,\text{out}}$ forming an alternating cycle in P^3. Consider an index $j \in \{1, \ldots, k\}$ and let Q be a path witnessing the relation $a_j \leqslant b_{j+1}$ in P^3 (indices are taken cyclically as always). This path cannot intersect $Q(b_j)$ as otherwise a_j and b_j would be comparable in P^3. Since a_j is not in $R(b_j)$ by assumption, it follows that no element from Q is in $R(b_j)$. In particular b_{j+1} is not in $R(b_j)$. By Claim 3.2.2, we have that $R(b_j) \subseteq R(b_{j+1})$.

Now, the inclusion $R(b_j) \subseteq R(b_{j+1})$ holds for all $j \in \{1, \ldots, k\}$ (cyclically), which implies that $R(b_1) = \cdots = R(b_k)$. By Claim 3.2.2 we must then have $m(b_1) = \cdots = m(b_k)$, and hence all pairs of our alternating cycle are contained in the set $I^i_{m(b_1)}$. However, the latter set is reversible, a contradiction. □

By mirroring the above proof but working inside $R(b_j)$, reversing the relation $R(b_j) \subseteq R(b_{j+1})$, and invoking Claim 3.2.3 instead of Claim 3.2.2, we similarly obtain that the sets $I^{i,\text{in}}$ are reversible:

Claim 3.2.5. *$I^{i,\text{in}}$ is reversible for each $i \in \{1, \ldots, 6h + 3\}$.*

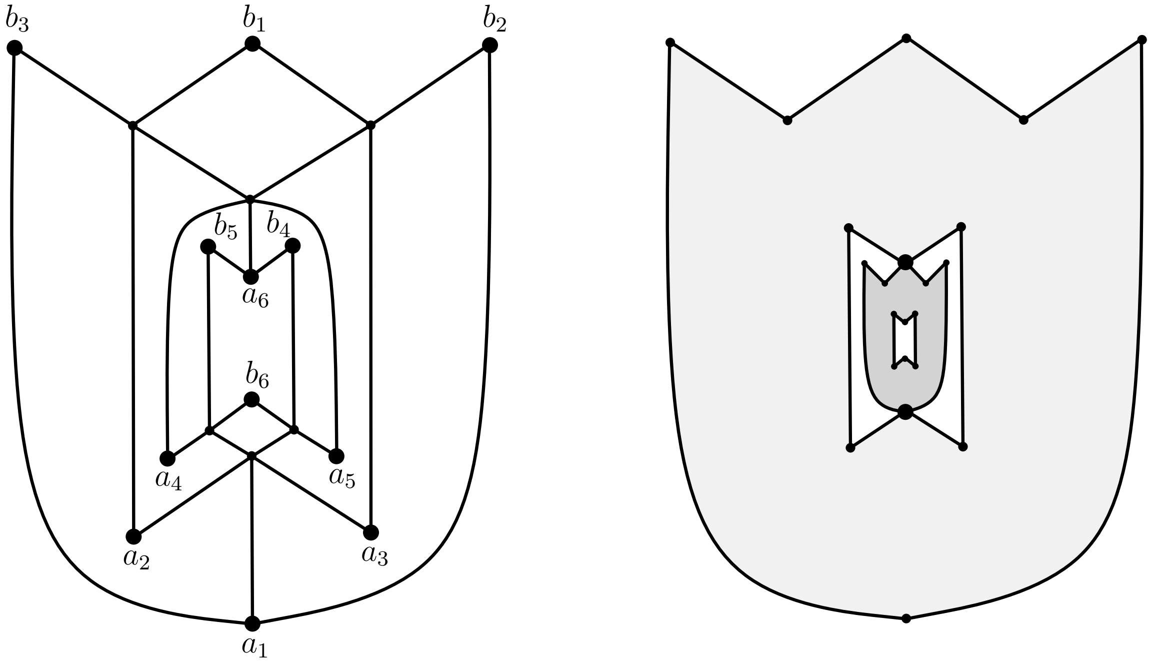

Figure 3.18: Iterative construction of planar posets with arbitrarily large dimension.

We conclude that I' can be partitioned into at most $12h + 6$ reversible sets, as claimed at the beginning of this section. Putting everything together, we deduce that

$$\dim(P) \leqslant 8 \dim(A^3, B^3) \leqslant 8(24h + 12) = 192h + 96,$$

which concludes the proof of our theorem.

3.3 New Lower Bound Constructions

As mentioned at the beginning of this chapter, the original construction of Kelly [43] shows that planar posets of height h can have dimension at least $h + 1$. It was even suggested in [71] (perhaps provocatively) that there might be some constant c such that the dimension of planar posets is at most $h + c$. Theorem 3.0.2 shows that such a constant cannot exist. To prove Theorem 3.0.2, we show the following statement that clearly implies it.

Theorem 3.3.1. *For every $h \geqslant 1$ with $h \equiv 1 \mod 3$, there is a planar poset P of height h with*

$$\dim(P) \geqslant (4/3)h + 2/3.$$

Proof. If $h = 1$ then it suffices to take an antichain of size at least 2, which has dimension 2. For larger heights h with $h \equiv 1 \mod 3$, we give an inductive construction that contains the standard example of size $(4/3)h + 2/3$ as an induced subposet, thus showing that the dimension is at least $(4/3)h + 2/3$. For the base case $h = 4$, start with the planar poset P in Figure 3.18 on the left. With the provided labeling it is easy to see that it contains the standard example S_6 as an induced subposet.

Next, using this base case, we describe how to build the desired poset for $h = 7$. Observe that in the diagram of P the maximal element b_6 is drawn below the minimal element a_6, and that there is some free space between them. Thus, we can take a

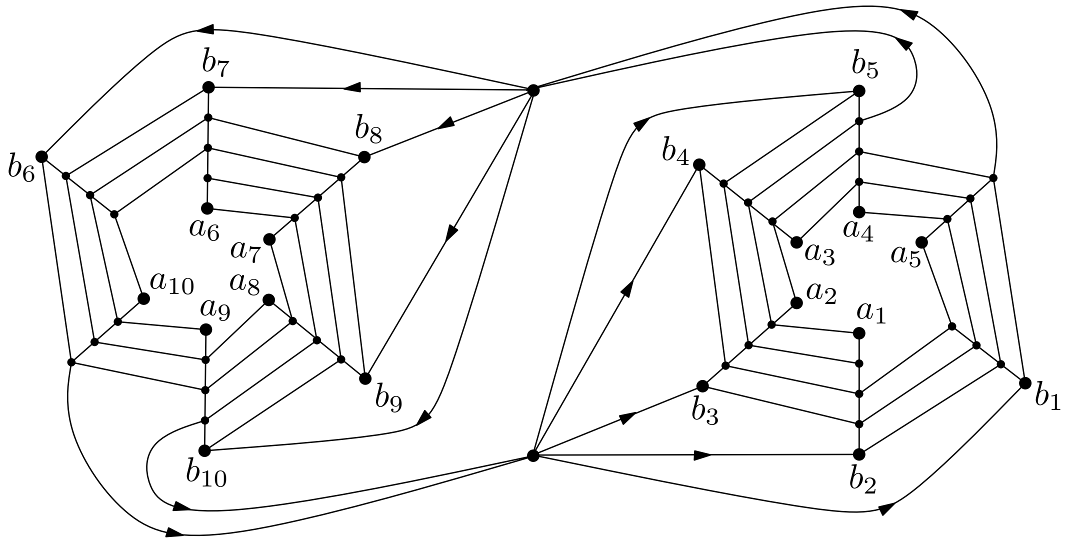

Figure 3.19: Construction of posets with planar cover graphs and large dimension.

small copy P' of P and insert P' in that space. We do it in such a way that the element a_1 of P' is identified with the element b_6 of P, and the element b_1 of P' is identified with element a_6 of P, see the right of Figure 3.18 for an illustration. The resulting poset has height 7 and the labeled elements of P and P', except for the ones that we identified, form a standard example of order 10.

It is not hard see that this copy-pasting procedure gives rise to an iterative construction, where in each step we increase the height by 3 and the order of the standard example under consideration by 4, as desired. □

Next, we show that there are posets of height h with planar cover graphs and dimension at least $2h - 2$, as stated in Theorem 3.0.3.

Proof of Theorem 3.0.3. Consider the construction of a poset with a planar cover graph illustrated in Figure 3.19. Directed edges indicate the precise cover relations. Undirected edges are part of one of the two induced subposets that resemble a spider net (c.f. Figure 3.1 in the introduction of this chapter). These edges are oriented away from the centers of the spider nets. The resulting poset has height 6 and contains a standard example of order 10 as indicated by the labeled elements. Clearly, we can extend this example in such a way that the poset has height h and contains a standard example of order $2h - 2$, and hence has dimension at least $2h - 2$. □

Chapter 4

Tree-width, Path-width, and Tree-depth

Bounding the tree-width, path-width, or tree-depth of a poset's cover graph can impose an upper bound on the dimension of that poset. As a prime example for this statement serves a classical result of Trotter and Moore [75], which can be formulated as follows: Posets with cover graphs of tree-width at most 1 have dimension at most 3. Such a constant upper bound on the dimension does not exist once we only require cover graphs to have tree-width at most 3. This fact is witnessed by Kelly's construction, which we already presented in the introduction (see Figure 1). However, as shown by Joret et al. [38], the dimension of posets whose cover graphs have bounded tree-width is bounded from above by a function in their height.

The purpose of this chapter is to follow this line of research and to discuss and prove the results depicted in the following table, which collects current best lower and upper bounds on the dimension of posets whose cover graphs have bounded tree-width (tw), path-width (pw), or tree-depth (td). (The parameter h refers to the height of the poset; asymptotic formulas in columns two and four are calculated with respect to h, so t and p are treated as constants.)

graph property	tw $\leqslant 2$	tw $\leqslant t$	pw $\leqslant 2$	pw $\leqslant p$	td $\leqslant d$
dim lower bound	4	$2^{\Omega(h^{\lfloor (t-1)/2 \rfloor})}$	4	$\Omega(h^{(p-2)/4})$	$2^{\lfloor d/2 \rfloor}$
dim upper bound	1276	$2^{\mathcal{O}(h^t)}$	6	$\mathcal{O}(h^{p-1})$	2^d

The chapter is organized as follows. Subsequent to this overview we introduce the necessary graph theoretic concepts, including tree- and path-decompositions and the various width-parameters from the chapter title. In Sections 4.1 and 4.2 we establish the upper bounds from the table above that are independent of the poset's height. First, in Section 4.1, we show that posets with cover graphs of tree-depth at most d have dimension at most 2^d. And second, in Section 4.2, we prove that the dimension of posets with cover graphs of path-width at most 2 is at most 6. Here, we also discuss how one can achieve a constant upper bound in the more general case of cover graphs with tree-width at most 2 (but we do not give a full proof). In Section 4.3 we establish a polynomial upper bound in height when the path-width of the cover graph is bounded. We conclude this chapter with Section 4.4, where we construct poset families that give us the lower bounds from the table above. The upper bound of the second column is not proven in this chapter but in Chapter 6. There, it will follow as a corollary from a result using concepts that do not fit very well here.

Tree- and Path-Decompositions. Let $G = (V, E)$ be a graph. A *tree-decomposition* of G is a pair $(T, \{\mathcal{B}_t \mid t \in V(T)\})$, consisting of a tree T and a family of non-empty subsets of V, such that

- for each edge $uv \in E$, there is a node $t \in V(T)$ such that $u, v \in \mathcal{B}_t$,
- for every $v \in V$, the nodes of $\{t \in V(T) \mid v \in \mathcal{B}_t\}$ form a non-empty subtree of T.

The sets $\mathcal{B}_v$ are the *bags* of the tree-decomposition. The maximum size of a bag minus one is the *width* of the tree-decomposition. Then the *tree-width* of G is defined to be the minimum width of a tree-decomposition of G, and denoted by $\operatorname{tw}(G)$.

A*path-decomposition* of G is simply a tree-decomposition in which the underlying tree is a path. However, for our arguments it will be convenient to work with the following equivalent definition. A *path-decomposition* $\mathcal{P}$ of G is a sequence $\mathcal{B}_1, \ldots, \mathcal{B}_\ell$ of non-empty subsets of V, such that

- the union $\bigcup_{1 \leqslant i \leqslant \ell} \mathcal{B}_i$ is equal to V,
- for each edge $uv \in E$, there is $i \in [\ell]$ with $u, v \in \mathcal{B}_i$, and
- for any $i, j, k \in [\ell]$ such that $i < j < k$, we have $\mathcal{B}_i \cap \mathcal{B}_k \subseteq \mathcal{B}_j$.

The sets $\mathcal{B}_i$ are called the *bags* of $\mathcal{P}$. The *width* of $\mathcal{P}$ is the maximum size of a bag of $\mathcal{P}$ minus one. Then, the *path-width* of a graph G is defined as the minimum width of a path-decomposition of G, and denoted by $\operatorname{pw}(G)$. For instance, paths (or more generally, caterpillars) with at least two vertices have path-width 1. Figure 4.1 illustrates a path-decomposition of the cover graph of Kelly's example. It has width 3 and this cannot be improved once the order of the example is large enough.

Another equivalent view on path-decompositions can be obtained via interval representations. Observe that for each vertex of the graph G, the bags in a path-decomposition that contain this vertex form a continuous interval in the sequence of all bags. This gives rise to an interval representation of a supergraph whose largest clique is of size one more than the width of the path-decomposition.

Planar graphs can have arbitrarily large tree-width as is witnessed by grids. On the other hand, graphs of tree-width at most 3 need not be planar. Therefore, planar graphs and graphs of bounded tree-width are independent graph classes. They have in common that they are proper minor-closed: Contracting an edge in a planar graph clearly leaves again a planar graph; and doing the same operation in a graph with bounded tree-width, we see that identifying the two contracted vertices with a new vertex in a tree-decomposition yields a decomposition of smaller or equal width. This implies that graphs of tree-width at most t are characterized by a finite list of excluded minors [65]. Moreover, it is easy to see that those graphs exclude K_{t+2} as a minor.

Let us now introduce the third graph parameter of the chapter title. The *tree-depth* of a graph G, which we denote by $\operatorname{td}(G)$, is the least number d for which there is a rooted forest F of depth d on the vertices of G, such that every edge of G connects two nodes that have an ancestor-descendant relationship in F. Notice that if G is connected, then F has to be a tree. Intuitively, while tree-width measures how tree-like a graph is, tree-depth measures how far a graph is from being a star.

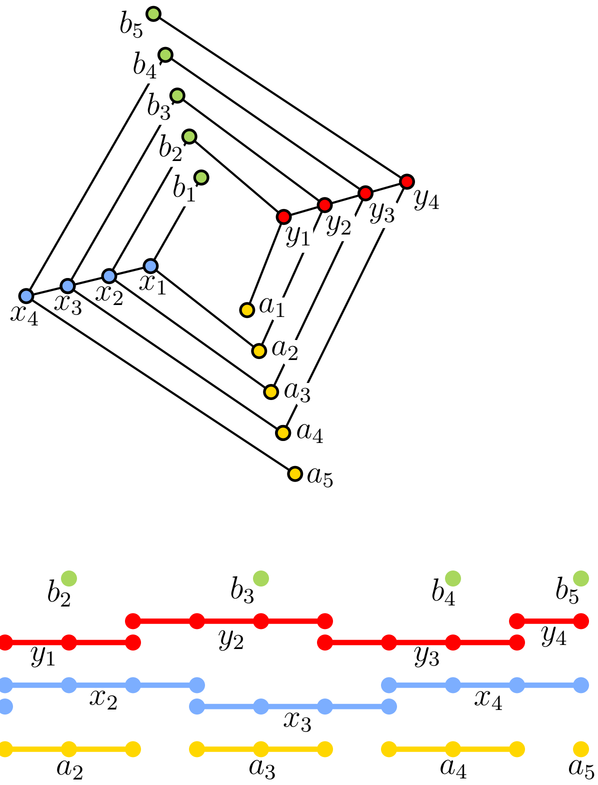

Figure 4.1: The cover graph of Kelly's example has path-width at most 3. Points that lie on a vertical line in the bottom figure build a bag of the path-decomposition. For instance, the first bag is $\{a_1, x_1, y_1, b_1\}$ and the second one is $\{a_2, x_1, x_2, y_1\}$.

Figure 4.2 depicts two graphs with underlying rooted forests that are witnessing the tree-depth of these graphs.

We explain now how the three introduced parameters are related to each other. Clearly, since every path-decomposition is also a tree-decomposition, we see that $\text{tw}(G) \leqslant \text{pw}(G)$ for all graphs G. Next, given a forest F of depth d that is witnessing the tree-depth of a graph G, we sketch how to use F to prove $\text{pw}(G) \leqslant d - 1$. As described earlier, it is enough to build an interval representation of a super graph of G whose maximum clique size is at most d. We start by introducing an interval I_x for each root x in F, and we let those be pairwise disjoint. Next we recurse in parallel on each tree of F. That is, if T_x is a tree with root x in F, then we proceed similarly with the forest $T_x - \{x\}$ (it has depth less than d) and place the resulting interval representation inside I_x. This procedure yields an interval representation of a supergraph of G with clique-size at most d.

We summarize that for all graphs G, we have

$$\text{tw}(G) \leqslant \text{pw}(G) \leqslant \text{td}(G) - 1.$$

4.1 Cover Graphs with Bounded Tree-depth

In this section we consider posets whose cover graphs have bounded tree-depth. The following theorem shows that the dimension of such posets is bounded from above by a constant.

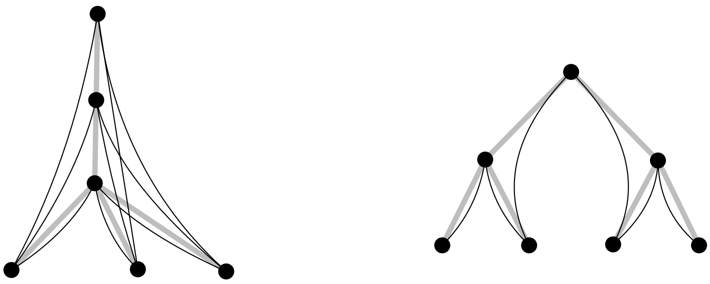

Figure 4.2: The tree-depth of the complete bipartite graph $K_{3,3}$ is 4 (left), while the tree-depth of the path P_7 is 3 (right).

Theorem 4.1.1. *For every poset P whose cover graph has tree-depth at most d, we have*

$$\dim(P) \leqslant 2^d.$$

Proof. Let P be a poset and let G be its cover graph. Let F be a rooted forest of depth d that is witnessing the tree-depth of G. If P is the union of chains, then $\dim(P) \leqslant 2$ and hence the theorem trivially holds. Otherwise, it is well known that the dimension of P is witnessed by a single component of it, implying that we may assume that P is connected. With this assumption it follows that F consists of a single tree T.

Let r be the root of T. We perform a depth-first search on T starting in r and we write $x \prec y$ if the node x is explored before the node y during the search. For each element y in P there exists a unique path $y_1, y_2, \ldots, y_\ell$ in T such that $r = y_1$ and $y_\ell = y$. Clearly, we have $\ell \leqslant d$ by the tree-depth bound. We define a function $\sigma(y)$ that saves the relation between y and the other elements from this unique path in P. That is, we let $\sigma(y) = (\sigma_1(y), \ldots, \sigma_{\ell-1}(y))$, where for $i \in [\ell - 1]$ we set

$$\sigma_i(y) = \begin{cases} 1, & \text{if } y_i \leqslant y \text{ in } P, \\ 0, & \text{otherwise.} \end{cases}$$

We fill up $\sigma(y)$ to a vector of length $d - 1$ by appending zeros to its end. Clearly, there are at most 2^{d-1} possible values for $\sigma(y)$.

Next, we define a signature $\sigma(x, y)$ for every incomparable pair $(x, y) \in \mathrm{Inc}(P)$. We set

$$\sigma(x, y) := (\sigma(y), \mathbb{1}_{x \prec y}),$$

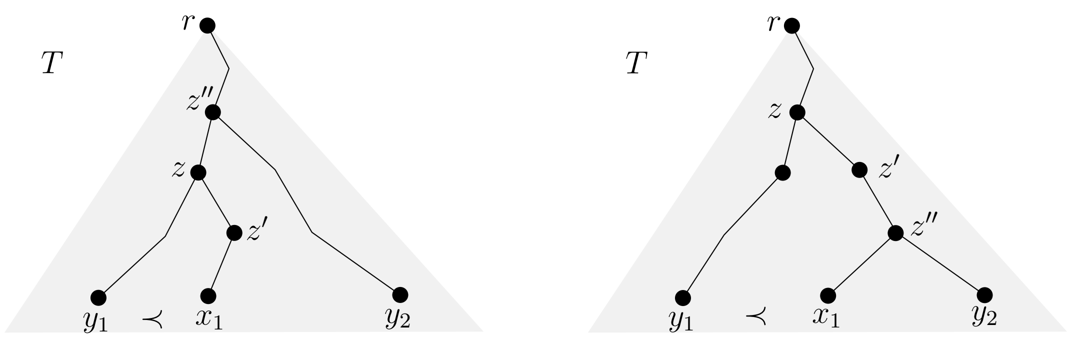

Figure 4.3: The two cases in the proof of Theorem 4.1.1. In both cases we have $x_1 \leqslant z'' \leqslant y_2$ in P.

where $\mathbb{1}_{x \prec y}$ evaluates to 1, if $x \prec y$, and to 0, otherwise.

We claim that for each possible signature σ, the set of incomparable pairs $(x, y) \in \mathrm{Inc}(P)$ satisfying $\sigma(x, y) = \sigma$ is reversible.

Suppose for a contradiction that this is not true for signature σ. Then there exists a strict alternating cycle $(x_1, y_1), \ldots, (x_k, y_k)$ in P such that $\sigma(x_i, y_i) = \sigma$ for each $i \in [k]$. We may assume that $\mathbb{1}_{x_i \prec y_i} = 0$ for every $i \in [k]$ as the other case is symmetric. Thus we have $y_i \prec x_i$ for all $i \in [k]$. We may also assume without loss of generality that y_1 is maximal among the y_i's with respect to $\prec$.

Starting in x_1 and going upwards along tree edges towards the root, we let z denote the first common ancestor of x_1 and y_1 in T. Since $y_1 \prec x_1$, we cannot have that x_1 is an ancestor of y_1 in T, and hence $z \neq x_1$. Let z' be the neighbor of z on the path from z to x_1 in T. Note that

$$y_1 \prec z' \tag{4.1}$$

as otherwise by the properties of a depth-first search ordering we would obtain $x_1 \prec y_1$, which is not true.

Next, consider a covering chain Q in G witnessing that $x_1 \leqslant y_2$ in P. Since edges of G can only connect nodes in T that are in an ancestor-descendant relationship, there is a vertex z'' in Q that is an ancestor of all other vertices in Q. Clearly, we have $x_1 \leqslant z'' \leqslant y_2$ in P. Note also that z, z', and z'' all lie on the path from x_1 to r in T. We distinguish two cases now.

First, suppose that z'' is an ancestor of z in T or $z'' = z$ (see Figure 4.3 (left) for such a situation). Then z'' is an ancestor of both y_1 and y_2 in T. It follows that there is an index j such that z'' is the j-th vertex on the path from r to y_1 and also the j-th vertex on the path from r to y_2 in T. Since all pairs of our cycle have the same signature, we have $\sigma(y_1) = \sigma(y_2)$. As $z'' \leqslant y_2$ in P, this implies $\sigma_j(y_1) = \sigma_j(y_2) = 1$ and hence $z'' \leqslant y_1$ in P. However, this yields $x_1 \leqslant y_1$ in P, which is a contradiction.

Second, suppose that z'' is not an ancestor of z and $z'' \neq z$ (see Figure 4.3 (right) for this case). Then z' must be an ancestor of z'' in T or $z' = z''$. By the choice of z'' it follows that z' is also an ancestor of y_2 in T. However, with (4.1) this implies $y_1 \prec z' \preceq y_2$, contradicting the choice of y_1 to be maximal with respect to $\prec$. This completes the proof of our subclaim.

We conclude that our signature function $\sigma(x, y)$ induces a partition of $\mathrm{Inc}(P)$ into reversible sets. Since the function ranges over at most $2 \cdot 2^{d-1} = 2^d$ possible values, we deduce $\dim(P) \leqslant 2^d$. □

Note that in the previous theorem we did not make an assumption on the height of the posets. However, graphs of tree-depth at most d contain no paths of length larger than 2^d. Therefore, posets whose cover graphs have bounded tree-depth actually have bounded height.

In Section 4.4 we show that the exponential upper bound 2^d in the previous theorem is essentially tight. We provide a construction yielding a lower bound of $2^{\lfloor d/2 \rfloor}$ there.

4.2 Cover Graphs with Tree- or Path-width 2

Recall from the introduction that there are graph properties of cover graphs that impose a constant bound on the dimension. For instance, if P is poset whose cover

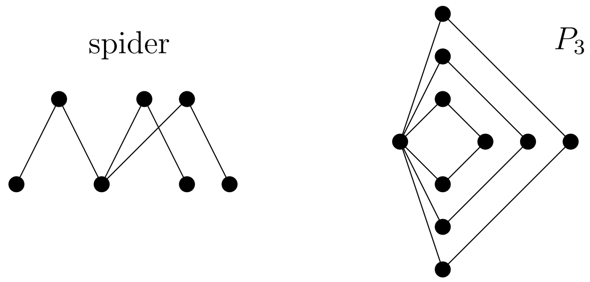

Figure 4.4: Left: 3-dimensional 'spider'-poset. Right: three 'diamonds' glued together at a common element yield the poset P_3. Extending it in the obvious fashion we have $\dim(P_n) \geqslant 4$ for all $n \geqslant 17$ [26]. Moreover, P_n has an outerplanar cover graph for all $n \geqslant 1$.

graph is a forest, then the dimension of P is at most 3 [75]. Forests form a subclass of outerplanar graphs, and Felsner, Trotter, and myself [26] showed that posets with outerplanar cover graphs have dimension at most 4. Let us note that both results are best possible as witnessed by the examples in Figure 4.4. Biró, Keller, and Young proved that posets with cover graphs of path-width at most 2 have dimension at most 17 [8].

The situation is different for posets whose cover graphs are planar or have path-width at most 3. Kelly's examples satisfy both conditions (see again Figure 4.1 illustrating this fact) but have arbitrarily large dimension. Observe that the three aforementioned results yielding constant upper bounds on the dimension all deal with cover graphs of tree-width at most 2: Forests, outerplanar graphs, and graphs of path-width 2. Consequently, it was asked by several authors in the field if there is a constant upper bound on the dimension of posets with cover graphs of tree-width 2. In joint work with Gwenaël Joret, Piotr Micek, Tom Trotter, and Ruidong Wang, we answered this question in the affirmative.

Theorem 4.2.1 ([39])**.** *For every poset P whose cover graph has tree-width at most 2, we have*

$$\dim(P) \leqslant 1276.$$

Although the research on this problem was a big part of my PhD studies, I decided to not include the proof of it into this thesis. The main reason is that the proof of this theorem is very long, technical, and admittedly not very insightful.

Nevertheless, let me give at least some basic ideas for the proof of this theorem. Given a poset whose cover graph has tree-width 2, we start with some standard tools in dimension theory. First, we apply the *Min-Max Reduction* (Lemma 2.1.1), which allows us to focus on incomparable pairs consisting of a minimal element and a maximal element. Second, we use the dual version of the *Global Min Support Reduction* (Lemma 2.2.2, which is an application of the *Unfolding Lemma*), to obtain the additional assumption that there exists a minimal element that is below all the maximal elements in the poset. From here it is then a long way of deriving properties of incomparable pairs with respect to a fixed tree-decomposition of the poset's cover graph. Many cases have to be considered until the final partition of incomparable pairs into reversible sets is done.

We do not believe that the bound of Theorem 4.2.1 is best possible. In fact, the correct answer is probably much closer to the best lower bound that we have.

This lower bound is 4 and comes again from the example on the right-hand side of Figure 4.4.

We continue with posets whose cover graphs have path-width at most 2 now. As mentioned before, Biró et al. [8] proved that they have dimension at most 17. We improve upon this bound with the following theorem.

Theorem 4.2.2. *For every poset P whose cover graph has path-width at most 2, we have*

$$\dim(P) \leqslant 6.$$

Our proof of this theorem does not directly rely on the existence a fixed path-decomposition of width 2. Instead, we use the fact that graphs of path-width at most 2 are 'almost' outerplanar. All we need to do is then to apply and extend the ideas developed in [26], where Felsner, Trotter, and myself show that posets with outerplanar cover graphs have dimension at most 4. As a consequence, we obtain a slightly more general result than Theorem 4.2.2.

Theorem 4.2.3. *Let P be a poset with a cover graph that can be obtained from an outerplanar graph G by subdividing each chord at most once (with respect to some fixed embedding of G witnessing that it is outerplanar). Then*

$$\dim(P) \leqslant 6.$$

Let us briefly discuss why this theorem implies Theorem 4.2.2. By $\mathcal{G}$ we denote the class of graphs that can be obtained from an outerplanar graph by subdividing each chord at most once (with respect to some fixed embedding witnessing outerplanarity). Barát et al. [3] and then independently Biró et al. [8] characterized the 2-connected graphs of path-width at most 2 as so-called *parallel nearly outerplanar graphs*. This characterization is a bit technical and therefore it is enough to note here that parallel nearly outerplanar graphs particularly belong to the class $\mathcal{G}$. With Lemma 3.3 of [8], where it is described how the 2-connected blocks intersect, it then follows that every graph of path-width at most 2 is a graph of $\mathcal{G}$.

Before we turn to the proof of Theorem 4.2.3, let us first introduce some useful notations. Let $G \in \mathcal{G}$. Then there exists an outerplanar graph G' and a planar embedding of G' with all vertices lying at the outer face such that G can be obtained from G' by subdividing some of the chords of G'. Since G' is outerplanar, there exists a linear order L on the vertices of G' such that no two edges cross with respect to L. That is, there exist no four vertices u, v, x, y in $V(G')$ such that $u < x < v < y$ in L and $uv, xy \in E(G')$. The linear order L yields a so-called *1-page embedding*: The vertices of G are placed on a horizontal line from left to right according to their ordering in L, and edges are drawn as half-circles above the line without mutual crossings.

We extend this 1-page embedding to a planar embedding of G now. For each vertex $s \in V(G) \setminus V(G')$, we consider the half-circle that corresponds to the chord of G' being subdivided by s, and we place a point representing s in the middle of that half-circle. We call the resulting drawing an *augmented 1-page embedding* of G. Figure 4.5 shows an augmented 1-page embedding of a graph from $\mathcal{G}$.

Suppose now that G is the cover graph of a poset P. We call the vertices being linearly ordered by L as *line points*, and for simplicity we identify L also with the

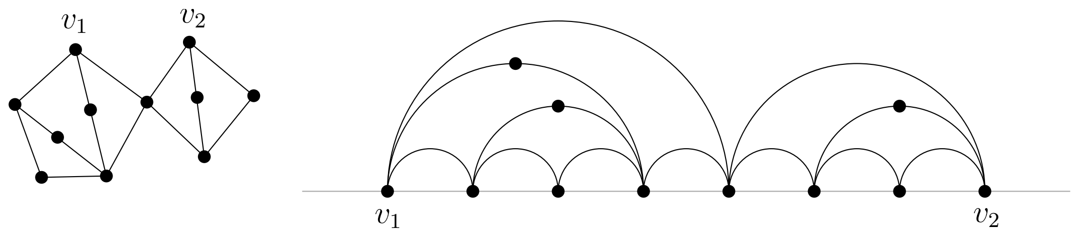

Figure 4.5: A graph from $\mathcal{G}$ and an augmented 1-page embedding of it.

fixed horizontal line of the embedding. The vertices in $S := V(G) \setminus L$ are the *subdivision points* of P.

Next, we want to associate with each element of P some special line points. For each line point x, we let ℓ_x be the leftmost element from the downset $\mathrm{D}(x)$ on the line L. Similarly, we let r_x be the rightmost element from the downset $\mathrm{D}(x)$ on L. For each subdivision point $y \in S$, let s_y and t_y be the adjacent line points such that $s_y < t_y$ in L. Then we define ℓ_y to be the leftmost element of $\mathrm{D}(y) \cup \{s_y\}$ on the line L. (Note that ℓ_y need not be in the downset of y, but in any case it is comparable to y in P.) Analogously, we define r_y to be the rightmost element of $\mathrm{D}(y) \cup \{t_y\}$ on the line L.

We continue with introducing a 'left-right' notion on elements of P. This can be naturally achieved if we only have line points. However, with the subdivision points it becomes a bit more subtle. For distinct elements $a, b \in P$, we write $a \prec b$ and say that *a lies to the left of b* in the augmented 1-page embedding if

(i) $a, b \in L$ and $a < b$ in L, or
(ii) $a \in L$, $b \in S$ and $a \leqslant s_b$ in L, or
(iii) $a \in S$, $b \in L$ and $t_a \leqslant b$ in L, or
(iv) $a, b \in S$ and $t_a \leqslant s_b$ in L.

In Figure 4.6 the four situations in which a lies to the left of b are illustrated. Furthermore, if $a \prec b \prec c$ or $c \prec b \prec a$ in the embedding, then b lies *between* a

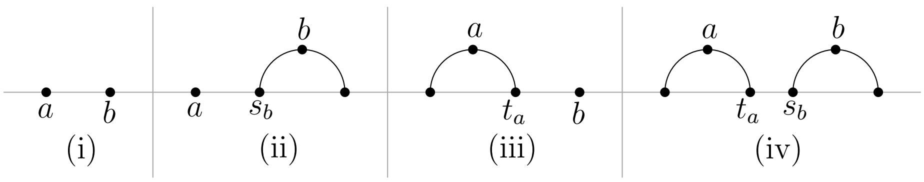

Figure 4.6: The four case in which a lies to the left of b ($a \prec b$).

and c. An element $a \in S$ *encloses* another element $b \neq s_a, t_a$ if the closed bounded region, which is defined by the cover edges as_a, at_a, and the horizontal line segment connecting s_a and t_a, contains b.

We turn to the proof of Theorem 4.2.3 now.

Proof of Theorem 4.2.3. Let P be a poset with a cover graph G as in the statement. As described in the previous paragraphs, we find an augmented 1-page embedding of G. We use the same notation as before, so in particular L denotes the linear order on the line points and S denotes the set of subdivision points.

We aim to partition the set of incomparable pairs $\mathrm{Inc}(P)$ into six reversible sets. The first one is defined as follows.

$$\mathcal{O}_1 := \{(x, y) \in \mathrm{Inc}(P) \mid \ell_y \prec x \prec y\}.$$

Claim. *The set $\mathcal{O}_1$ is reversible.*

Proof. Suppose to the contrary that $\mathcal{O}_1$ is not reversible. Then there exists a strict alternating cycle $C = (a_1, b_1), \ldots, (a_k, b_k)$ in $\mathcal{O}_1$. Observe that we may assume without loss of generality that b_1 has the following extremal property: for each $i \in \{2, \ldots, k\}$, the element b_i does not lie to the left of b_1 and it also does not enclose b_1.

We distinguish three cases now. First, let us suppose that both a_1 and b_1 are line points. By the choice of b_1 it follows that $b_1 \prec b_2$, and together with the definition of $\mathcal{O}_1$ we obtain

$$\ell_{b_1} \prec a_1 \prec b_1 \prec b_2. \tag{4.2}$$

Let Q and Q' be covering chains witnessing that $\ell_{b_1} \leqslant b_1$ and $a_1 \leqslant b_2$ hold in P, respectively. If b_2 is a line point, then all elements in (4.2) are line points. However, in this case the two curves representing Q and Q' in the 1-page embedding clearly have to intersect in a common point $z \in P$. As Q and Q' are chains, we particularly deduce that $a_1 \leqslant z \leqslant b_1$ in P, which is a contradiction.

Hence b_2 is a subdivision point. Notice that in this case Q' has to contain either s_{b_2} or t_{b_2}, since these are the only adjacent vertices of b_2 in G. Let $c \in \{s_{b_2}, t_{b_2}\}$ denote this element on Q'. Then we have $a_1 \leqslant c$ in P and thus $b_1 \neq c$. Moreover, by $b_1 \prec b_2$ this implies $\ell_{b_1} \prec a_1 \prec b_1 \prec c$. As only line points are involved, we deduce with the same argument as before that $a_1 \leqslant b_1$ has to hold in P, which is again a contradiction. This concludes the case that both a_1 and b_1 are line points.

Suppose now that $a_1 \in L$ and $b_1 \in S$. Since $\ell_{b_1} \prec a_1 \prec b_1$, we particularly obtain that $\ell_{b_1} < a_1 \leqslant s_{b_1}$ in L and hence $\ell_{b_1} \neq s_{b_1}$. By the definition of ℓ_{b_1} this implies $\ell_{b_1} \in \mathrm{D}(b_1)$. Moreover, as b_1 is only with s_{b_1} and t_{b_1} in a cover relation, there is $z \in \{s_{b_1}, t_{b_1}\}$ such that $\ell_{b_1} \leqslant z \leqslant b_1$ in P. Clearly, we must have $a_1 \neq z$ and hence $a_1 \prec z$ holds.

Let us consider the possible positions of b_2 in the drawing with respect to z now. First, we note that $b_2 \neq s_{b_1}, t_{b_1}$ (in particular $b_2 \neq z$) as otherwise b_1 and b_2 have to be comparable, which contradicts to the fact that C is strict. If $z \prec b_2$ then we are in a similar situation as in the first case since inequality 4.2 holds now with b_1 replaced by the line point z. So following previous arguments we deduce that $a_1 \leqslant z$ in P. However, this is a contradiction as this also implies $a_1 \leqslant z \leqslant b_1$ in P.

Next, observe that b_2 cannot enclose z as this would imply that b_2 encloses b_1, which is not true by our choice of b_1. Therefore, it remains to consider the subcase $b_2 \prec z$. Again by the choice of b_1, we deduce that b_1 encloses b_2 and $z = t_{b_1}$. Clearly, as a_1 is not enclosed by b_1, any covering chain from a_1 to b_2 has to contain s_{b_1} or t_{b_1}. So there is $z' \in \{s_{b_1}, t_{b_1}\}$ such that $a_1 \leqslant z' \leqslant b_2$ in P. Since b_1 and z' are comparable in P, we have $b_1 \leqslant z'$ or $b_1 \geqslant z'$ in P. In the first case we conclude $a_k \leqslant b_1 \leqslant z' \leqslant b_2$ in P, which contradicts to the fact that C is a strict alternating cycle, and in the second case it follows that $a_1 \leqslant z' \leqslant b_1$ in P, which is also a contradiction. This completes the case that $a_1 \in L$ and $b_1 \in S$.

We are left with the case that $a_1 \in S$ and $b_1 \in L \cup S$. Since $a_1 \prec b_1$, we have that $a_1 \neq b_2$ by our choice of b_1. Hence there is $y \in \{s_{a_1}, t_{a_1}\}$ such that $a_1 < y \leqslant b_2$ in P. By the definition of $\mathcal{O}_1$ we know that $\ell_{b_1} \prec a_1$, and together with $\ell_{b_1} \neq y$ (as otherwise $a_1 \leqslant \ell_{b_1} \leqslant b_1$ in P) it follows that $\ell_{b_1} \prec y$. By similar reasons, $a_1 \prec b_1$ implies $y \prec b_1$. Putting these inequalities together we obtain $\ell_{b_1} \prec y \prec b_1$. Observe also that y and b_1 have to be incomparable as C is strict.

Therefore, if we apply the arguments from the previous two cases with a_1 replaced by y, we end up with a contradiction or we conclude that $y \leqslant b_1$ in P. Of course, the latter is also a contradiction as it makes a_1 and b_1 comparable in P. This completes the last of the three main cases and finishes the proof. □

We proceed with a set of incomparable pairs that corresponds to the symmetric analogue of $\mathcal{O}_1$. We let

$$\mathcal{O}_2 := \{(x, y) \in \mathrm{Inc}(P) \mid y \prec x \prec r_y\}.$$

Clearly, the set $\mathcal{O}_2$ is reversible as follows from symmetric arguments of the previous proof.

Next, we consider the set

$$\mathcal{O}_3 := \{(x, y) \in \mathrm{Inc}(P) \mid x \prec y \text{ and } \ell_y \not\prec x\}.$$

Claim. *The set $\mathcal{O}_3$ is reversible.*

Proof. Arguing by contradiction, suppose that $\mathcal{O}_3$ is not reversible and hence contains a strict alternating cycle $C = (a_1, b_1), \dots, (a_k, b_k)$. We may assume that b_1 is extremal in the sense that $\ell_{b_i} \not\prec \ell_{b_1}$ for all $i \in [k]$.

Suppose first that a_1 is a line point. Since $\ell_{b_1} \not\prec a_1$ holds by the definition of $\mathcal{O}_3$, this implies $a_1 \prec \ell_{b_1}$ (we have $a_1 \neq \ell_{b_1}$ as otherwise a_1 and b_1 are comparable in P). However, a_1 is in the downset of b_2 and hence it follows that $\ell_{b_2} \preceq a_1 \prec \ell_{b_1}$, a contradiction to the choice of b_1.

Therefore, we have $a_1 \in S$. Since $\ell_{b_1} \not\prec a_1$ it follows that $s_{a_1} \prec \ell_{b_1}$. We also have $a_1 \neq b_2$ as otherwise $\ell_{b_2} = \ell_{a_1} \preceq s_{a_1} \prec \ell_{b_1}$ in L, which is contradicting our choice of b_1. Thus there is $z \in \{s_{a_1}, t_{a_1}\}$ such that $a_1 < z \leqslant b_2$ in P. By our assumption on b_1 we must have $z = t_{a_1}$, and hence

$$a_1 \leqslant t_{a_1} \leqslant b_2$$

in P. We also see that $\ell_{b_2} \preceq t_{a_1}$, and together with $\ell_{b_2} \not\prec \ell_{b_1}$ this implies $\ell_{b_1} \preceq t_{a_1}$. Recall now that b_1 is comparable to ℓ_{b_1}. So we have $\ell_{b_1} \neq t_{a_1}$ as otherwise b_1 and t_{a_1} are comparable in P, which would imply that b_1 is comparable to a_1 or b_2, contradicting to C being strict. Therefore,

$$s_{a_1} \prec \ell_{b_1} \prec t_{a_1}$$

and a_1 encloses ℓ_{b_1} bot not b_1. It follows that any covering chain Q from ℓ_{b_1} to b_1 has to contain s_{a_1} or t_{a_1}. However, Q cannot contain s_{a_1} as by $s_{a_1} \prec \ell_{b_1}$ we know that s_{a_1} is not in the downset of b_1. And Q also cannot contain t_{a_1} since otherwise $a_1 \leqslant t_{a_1} \leqslant b_1$ in P. Thus, we end up with a contradiction. This completes the proof of the claim. □

Using symmetric arguments to the ones from the previous proof, we also obtain that the following set is reversible:

$$\mathcal{O}_4 := \{(x, y) \in \mathrm{Inc}(P) \mid y \prec x \text{ and } x \not\prec r_y\}.$$

Let us note at this point that incomparable pairs (x, y) of P such that x and y are comparable in the $\prec$-relation are contained in exactly one of the sets $\mathcal{O}_1, \ldots, \mathcal{O}_4$. For instance, if $x \prec y$ then (x, y) is either contained in $\mathcal{O}_1$ or in $\mathcal{O}_3$, depending on whether $\ell_y \prec x$ holds or not.

We continue with the remaining incomparable pairs, that is, with those being incomparable in the $\prec$-relation. Each of them is contained in one of the following sets:

$$\mathcal{O}_5 := \{(x, y) \in \mathrm{Inc}(P) \mid y \text{ encloses } x\},$$
$$\mathcal{O}_6 := \{(x, y) \in \mathrm{Inc}(P) \mid x \text{ encloses } y\}.$$

We show now that both sets are reversible.

Claim. *The set $\mathcal{O}_5$ is reversible.*

Proof. Suppose that $\mathcal{O}_5$ is not reversible and hence contains a strict alternating cycle $C = (a_1, b_1), \ldots, (a_k, b_k)$. Note that each b_i ($1 \leqslant i \leqslant k$) is in S as it encloses a_i. We may additionally assume that b_1 is minimal among the b_i's, meaning that b_1 does not enclose any other b_i ($2 \leqslant i \leqslant k$).

Since b_1 encloses a_1 but not b_2, any covering chain from a_1 to b_2 has to contain s_{b_1} or t_{b_1}. It follows that there is $z \in \{s_{b_1}, t_{b_1}\}$ such that $a_1 \leqslant z \leqslant b_2$ in P. However, as b_1 is comparable to z we conclude that b_1 is comparable to at least one of a_1 and b_2. This clearly cannot hold in a strict alternating cycle and finishes the proof. □

Claim. *The set $\mathcal{O}_6$ is reversible.*

Proof. Suppose that $\mathcal{O}_6$ is not reversible and hence contains a strict alternating cycle $C = (a_1, b_1), \ldots, (a_k, b_k)$. We may assume that a_1 is not enclosed by any other a_i ($2 \leqslant i \leqslant k$).

Since a_2 encloses b_2 but not a_1, any covering chain from a_1 to b_2 has to contain s_{a_2} or t_{a_2}. Thus, there is $r \in \{s_{a_2}, t_{a_2}\}$ such that $a_1 \leqslant r \leqslant b_2$ in P. However, as a_2 and r are comparable we conclude that a_2 is comparable to a_1 or b_2. This cannot hold in a strict alternating cycle and completes the proof. □

Since

$$\mathrm{Inc}(P) = \mathcal{O}_1 \cup \cdots \cup \mathcal{O}_6,$$

we conclude that $\dim(P) \leqslant 6$. □

We do not know if the upper bound of Theorem 4.2.2 (or even Theorem 4.2.3) is best possible. The example on the right-hand side of Figure 4.4 yields a lower bound of 4 as it is straightforward to verify that its cover graph has path-width at most 2. We tend to believe that the best upper bound is also "4".

4.3 Cover Graphs with Path-width at most p

The dimension of posets with cover graphs of bounded path-width is bounded by their height. This result was first obtained by Joret et al. in [38], where they actually prove the stronger result for bounded tree-width. Their bounds are double exponential in height, and as we will prove in Section 4.4, exponential growth is possible in the setting of bounded tree-width.

In this section we show that bounding the path-width of cover graphs is much more restrictive than bounding only the tree-width. More precisely, we show that the dimension of posets with cover graphs of bounded path-width is bounded by a polynomial function in their height. This is unpublished joint work with Piotr Micek and Gwenaël Joret.

Theorem 4.3.1. *Let $p \geqslant 1$. For every poset P of height at most h with a cover graph of path-width at most p, we have*

$$\dim(P) = \mathcal{O}(h^{p-1}).$$

Let P be a poset of height h whose cover graph has path-width at most p. The case $p \leqslant 1$ is easy, since in this case the cover graph is a forest and hence the dimension is at most 3 by the classic Moore-Trotter result [75].

So suppose $p \geqslant 2$ now. Then, we continue with the *Min-Max Reduction* described in Chapter 2. That is, we apply Lemma 2.1.1 to P and obtain a poset P' such that

- P is an induced subposet of P',
- the height of P' is at most h,
- $\dim(P) \leqslant \dim(\mathrm{Min}(P'), \mathrm{Max}(P'))$, and
- the elements in $P' \setminus P$ have degree 1 in the cover graph of P'.

We aim to achieve a polynomial upper bound on $\dim(\mathrm{Min}(P'), \mathrm{Max}(P'))$. To do so, we first show that the cover graph of P' has path-width at most $p + 1$.

Let $\mathcal{P}$ be a path-decomposition witnessing that the cover graph of P has path-width at most p. Let $x \in P' \setminus P$ and consider the unique neighbor $y \in P$ of x in the cover graph of P'. We choose a bag $\mathcal{B}$ of $\mathcal{P}$ that is containing element y and make a copy $\mathcal{B}^x$ of it. Next, we insert x to $\mathcal{B}^x$ and add $\mathcal{B}^x$ to the path-decomposition by putting it next to the bag $\mathcal{B}$. Realizing this step for all elements of $P' \setminus P$, it is easy to see this yields a path-decomposition of width at most $p + 1$ for the cover graph of P'.

With these observations at hand, Theorem 4.3.1 is implied by the following theorem.

Theorem 4.3.2. *Let $p \geqslant 2$. For every poset P of height at most h with a cover graph of path-width at most p, we have*

$$\dim(\mathrm{Min}(P), \mathrm{Max}(P)) = \mathcal{O}(h^{p-2}).$$

Before we give a proof of this theorem, let us first show a useful lemma.

Lemma 4.3.1. *Let P be poset and $\mathcal{B}_1, \ldots, \mathcal{B}_\ell$ be a path-decomposition of width p of the cover graph of P. Fix $j \in [\ell]$ and let $X := \bigcup_{i \leqslant j} \mathcal{B}_i$ and $Y := \bigcup_{i \geqslant j} \mathcal{B}_j$. Then we have*

$$\dim(X, Y) \leqslant 2^{p+1}$$

Proof. To establish the upper bound on $\dim(X, Y)$, we partition $\mathrm{Inc}(X, Y)$ into 2^{p+1} reversible sets.

Let $q \leqslant p + 1$ be the size of the bag $\mathcal{B}_j$ and denote its elements by $z_1, \ldots, z_q$. With each $x \in X$ we associate a vector $\tau(x)$ that saves the comparability status to the elements of $\mathcal{B}_j$. That is to say, we let $\tau(x) := (\tau_i(x))_{i \in [q]}$, where

$$\tau_i(x) = \begin{cases} 1, & \text{if } x \leqslant z_i \text{ in } P \\ 0, & \text{otherwise.} \end{cases}$$

We claim that for each such possible vector $\tau \in \{0, 1\}^q$, the set $S(\tau)$ of incomparable pairs $(x, y) \in \mathrm{Inc}(X, Y)$ satisfying $\tau(x) = \tau$ is reversible.

Suppose to the contrary that this is not true for some $\tau \in \{0, 1\}^q$. Then there exists a strict alternating cycle $(x_1, y_1), \ldots, (x_k, y_k)$ in $S(\tau)$. Since $x_1 \leqslant y_2$ in P, there is a covering chain

$$x_1 = c_1 < c_2 < \cdots < c_r = y_2 \;\text{ in } P.$$

Let $s \in [r]$ be maximal such that $c_s \in X$ (note that this is well-defined as $c_1 \in X$). We aim to show $c_s \in \mathcal{B}_j$ now.

Observe that by the definition of a path-decomposition it follows that $X \cap Y = \mathcal{B}_j$. As $c_s \in X$ holds, it is therefore enough to show that $c_s \in Y$. If $s = r$, then $c_s = y_2 \in Y$. And if $s < r$, then $c_s c_{s+1}$ is an edge of the cover graph of P and consequently there is an index $t \in [\ell]$ such that $c_s, c_{s+1} \in \mathcal{B}_t$. By the maximality of s we deduce $t > j$, and thus we obtain $c_s \in \mathcal{B}_t \subseteq Y$, as desired.

So we indeed have $c_s \in \mathcal{B}_j$. In particular, there is an index $m \in [q]$ such that $c_s = z_m$. As we have $x_1 \leqslant z_m$ in P, this implies $\tau_m(x_1) = 1$. However, all pairs $(x, y) \in S(\tau)$ satisfy $\tau(x) = \tau$, and hence we have $\tau_m(x_2) = \tau_m = \tau_m(x_1) = 1$. This implies $x_2 \leqslant z_m = c_s \leqslant y_2$ in P, which is a contradiction. Therefore, $S(\tau)$ is reversible.

Finally, observe that there are at most $2^q \leqslant 2^{p+1}$ possible vectors τ, implying that we can partition $\mathrm{Inc}(X, Y)$ into at most 2^{p+1} reversible sets. This completes the proof. □

Proof of Theorem 4.3.2. The proof will go by induction on $p \geqslant 2$. So let P be a poset as in the statement of the theorem. In the base case $p = 2$ we have an upper bound of 6 on the dimension of P by Theorem 4.2.2. In particular, $\dim(\mathrm{Min}(P), \mathrm{Max}(P))$ is bounded by 6.

For the induction step suppose $p \geqslant 3$ now. First, we apply the dual version of our *Global Min Support Reduction* (Lemma 2.2.2) to P. We obtain a poset Q of height at most h such that

(i) $\mathrm{cover}(Q)$ is a minor of $\mathrm{cover}(P)$,
(ii) $\dim(\mathrm{Min}(P), \mathrm{Max}(P)) \leqslant 2 \dim(\mathrm{Min}(Q), \mathrm{Max}(Q))$,
(iii) there is $a_0 \in \mathrm{Min}(Q)$ such that $a_0 \leqslant b$ in Q for all $b \in \mathrm{Max}(Q)$.

Graphs of bounded path-width form a minor-closed class of graphs. Therefore, by (i) we deduce that the cover graph of Q has path-width at most p. In the following we aim to bound $\dim(\mathrm{Min}(Q), \mathrm{Max}(Q))$, which by (ii) will also yield a bound on $\dim(\mathrm{Min}(P), \mathrm{Max}(P))$.

For convenience, we set $A = \mathrm{Min}(Q)$ and $B = \mathrm{Max}(Q)$. Let $\mathcal{B}_1, \ldots, \mathcal{B}_m$ be a path-decomposition of width p of the cover graph of Q. Let $j \in [m]$ be such that

$a_0 \in \mathcal{B}_j$, where a_0 is the special minimal element of (iii). Next, we partition the minimal elements and maximal elements of Q into a respective 'left' and 'right' set. That is, we define

$$A_r := A \cap \bigcup_{i \geqslant j} \mathcal{B}_i \quad \text{and} \quad A_\ell := A \setminus A_r,$$
$$B_r := B \cap \bigcup_{i \geqslant j} \mathcal{B}_i \quad \text{and} \quad B_\ell := B \setminus B_r.$$

This yields the following partition of the 'min-max' incomparable pairs in Q:

$$\operatorname{Inc}(A, B) = \operatorname{Inc}(A_\ell, B_r) \cup \operatorname{Inc}(A_r, B_\ell) \cup \operatorname{Inc}(A_\ell, B_\ell) \cup \operatorname{Inc}(A_r, B_r).$$

Therefore,

$$\dim(A, B) \leqslant \dim(A_\ell, B_r) + \dim(A_r, B_\ell) + \dim(A_\ell, B_\ell) + \dim(A_r, B_r).$$

Observe that A_ℓ is contained in the first j bags of the path-decomposition. Thus, by Lemma 4.3.1 we have $\dim(A_\ell, B_r) \leqslant 2^{p+1}$, and a symmetric argument also shows $\dim(A_r, B_\ell) \leqslant 2^{p+1}$.

We aim to bound $\dim(A_r, B_r)$ now. If B_r is empty, then $\dim(A_r, B_r) = 1$ and this value is clearly good enough for our purposes. Otherwise, we let k ($j \leqslant k \leqslant m$) be maximal such that $\mathcal{B}_k$ contains some element of B_r, and let $b_r \in \mathcal{B}_k \cap B_r$ be such one. Then we denote by Q_r the subposet of Q that is induced by the elements in $\bigcup_{i=j}^{k} \mathcal{B}_i$. Clearly, we have that $A_r \subseteq \operatorname{Min}(Q_r)$ and $B_r \subseteq \operatorname{Max}(Q_r)$, and therefore also

$$\dim(A_r, B_r) \leqslant \dim(\operatorname{Min}(Q_r), \operatorname{Max}(Q_r)).$$

However, observe that the cover graph of Q_r is not necessarily a subgraph of the cover graph of Q. This is because the removal of elements from Q may create new cover relations. We will overcome this issue with the following claim.

Claim. *The sequence $\mathcal{P}' = \mathcal{B}_j, \ldots, \mathcal{B}_k$ is a path-decomposition of the cover graph of Q_r.*

Proof. Suppose that $x < y$ is a cover relation in Q_r. We need to show that there is a bag of $\mathcal{P}'$ that contains x and y. Since Q_r is an induced subposet of Q, we also have $x < y$ in Q. We consider two cases now.

First, suppose that $x < y$ is also a cover relation in Q. Then there is a bag $\mathcal{B}_i$ of $\mathcal{P}$ containing x and y. If $j \leqslant i \leqslant k$, then $\mathcal{B}_i$ is a bag of $\mathcal{P}'$ and we are done already. If $i < j$, then by the fact that x and y are contained in some bag of $\mathcal{P}'$ we deduce that $x, y \in \mathcal{B}_j$ (here we also used the properties of a path-decomposition). Finally, if $k < i$, then we similarly obtain that $x, y \in \mathcal{B}_k$. This completes the first case.

Second, we suppose that $x < y$ is not a cover relation in Q. Then there exists a covering chain $x < z_1 < \cdots < z_s < y$ in Q such that $z_1, \ldots, z_s \in Q \setminus Q_r$. Since $z_1 \notin Q_r$, we have that z_1 is either contained in $\bigcup_{1 \leqslant i < j} \mathcal{B}_i$ or in $\bigcup_{k < i \leqslant m} \mathcal{B}_i$ (both is not possible by the properties of a path-decomposition). Without loss of generality we may assume that the first case applies. Combining the facts that $z_1, \ldots, z_s$ are not contained in Q_r and that they form a path in cover(Q), we see that $z_1, \ldots, z_s \in \bigcup_{1 \leqslant i < j} \mathcal{B}_i$ (and these are the only bags that possibly contain some z_i). Moreover, since xz_1 and $z_s y$ are edges in cover(Q), we also deduce that $x, y \in \bigcup_{1 \leqslant i < j} \mathcal{B}_i$. Together with $x, y \in Q_r$ this then implies $x, y \in \mathcal{B}_j$. This completes the second case and hence $\mathcal{P}'$ is indeed a path-decomposition of width at most p of cover(Q_r). □

Recall that a_0 and b_r are elements of Q_r such that $a_0 \in \mathcal{B}_j$, $b_r \in \mathcal{B}_k$, and $a_0 \leqslant b_r$ in Q_r. Thus there is a covering chain

$$a_0 = c_0 < c_1 < \cdots < c_t < b_r$$

in Q_r with $t \leqslant h-2$. Note that each bag of $\mathcal{P}'$ contains at least one element of this chain, since $a_0, c_1, \ldots, c_t, b_r$, viewed as a path in cover(Q_r), connects the first bag of $\mathcal{P}'$ with the last bag of $\mathcal{P}'$.

Next, we partition the set $B_r \setminus \{b_r\}$ into $t+1$ sets $B_{r,0}, \ldots, B_{r,t}$, where $B_{r,i}$ consists of all $b \in B_r \setminus \{b_r\}$ for which i $(0 \leqslant i \leqslant t)$ is maximal such that $c_i \leqslant b$ holds in Q_r. Note that this is indeed a partition as a_0 is below each maximal element of $B_r \setminus \{b_r\}$. This also induces the following partition of $\text{Inc}(A_r, B_r)$:

$$\text{Inc}(A_r, B_r) = \text{Inc}(A_r, B_{r,0}) \cup \cdots \cup \text{Inc}(A_r, B_{r,t}) \cup \text{Inc}(A_r, \{b_r\}).$$

Clearly, this also means that

$$\dim(A_r, B_r) \leqslant \dim(A_r, B_{r,0}) + \cdots \dim(A_r, B_{r,t}) + 1. \tag{4.3}$$

We aim to bound $\dim(A_r, B_{r,i})$ by some function in $\mathcal{O}(h^{p-3})$ for each $i \in \{0, \ldots, t\}$ now.

So we arbitrarily fix $i \in \{0, \ldots, t\}$ for the moment. First, we get rid of somewhat redundant elements. That is, we let $A_{r,i}$ be the set of elements $a \in A_r$ which are involved in an incomparable pair of $\text{Inc}(A_r, B_{r,i})$. Then we still have $\text{Inc}(A_{r,i}, B_{r,i}) = \text{Inc}(A_r, B_{r,i})$ and hence $\dim(A_{r,i}, B_{r,i}) = \dim(A_r, B_{r,i})$. Let $Q_{r,i}$ denote the convex subposet $\text{conv}(A_{r,i} \cup B_{r,i})$ of Q_r. Recall that by convexity, the cover graph of $Q_{r,i}$ is an induced subgraph of the cover graph of Q_r. Moreover, it is easy to see that

$$\dim_{Q_r}(A_{r,i}, B_{r,i}) = \dim_{Q_{r,i}}(A_{r,i}, B_{r,i}).$$

It follows that restricting $\mathcal{P}'$ to the elements of $Q_{r,i}$ yields a path-decomposition $\mathcal{P}'_i$ of the cover graph of $Q_{r,i}$.

Next, we show that the width of $\mathcal{P}'_i$ is at most $p-1$. This can be achieved by proving that $Q_{r,i}$ does not contain any of the elements $a_0, c_1, \ldots, c_t, b_r$. (Recall that each bag of $\mathcal{P}'$ contains at least one element of $a_0, c_1, \ldots, c_t, b_r$.) It does not contain a_0 as in $A_{r,i}$ there are only elements that are incomparable to at least one element of $B_{r,i}$. It also does not contain b_r as this element is not in $B_{r,i}$ by definition.

So suppose for a contradiction that $c \in \{c_1, \ldots, c_t\}$ is an element of $Q_{r,i}$. Then there are $a \in A_{r,i}$ and $b \in B_{r,i}$ such that $a \leqslant c \leqslant b$ in $Q_{r,i}$ by the definition of a convex hull. By the definition of the set $B_{r,i}$, we even obtain $a \leqslant c \leqslant c_i \leqslant b$ in $Q_{r,i}$. However, $a \leqslant c_i$ in $Q_{r,i}$ implies that a is below each element of $B_{r,i}$, which is a contradiction as elements of $A_{r,i}$ do not have this property.

We conclude that $\mathcal{P}'_i$ indeed has width at most $p-1$, which allows us to apply the induction hypothesis to the poset $Q_{r,i}$. It follows that there exists a function $f(h) \in \mathcal{O}(h^{p-3})$ such that $\dim(A_{r,i}, B_{r,i}) \leqslant f(h)$ for every $i \in \{0, \ldots, t\}$. With (4.3) and $t \leqslant h-2$ this implies

$$\dim(A_r, B_r) \leqslant (h-1)f(h) + 1.$$

A symmetric argument shows $\dim(A_\ell, B_\ell) \leqslant (h-1)f(h) + 1$. Finally, coming back to our starting poset P we conclude that

$$\dim(\text{Min}(P), \text{Max}(P)) \leqslant 2\dim(A, B) \leqslant 4(h-1)f(h) + 2^{p+3} = \mathcal{O}(h^{p-2}).$$

This completes the proof. □

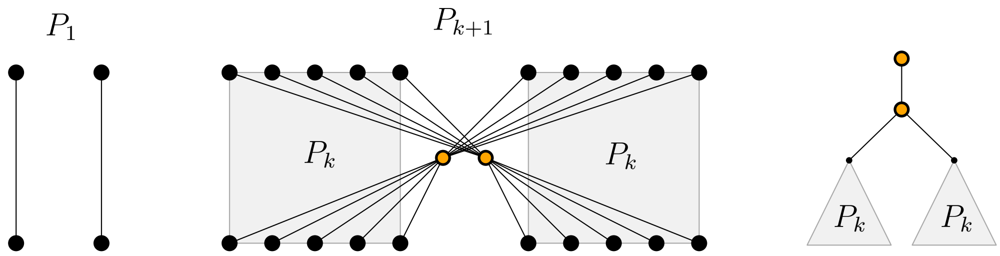

Figure 4.7: Inductive construction of the family $(P_k)_{k\geqslant 1}$. P_1 is simply the standard example S_2. For $k \geqslant 2$, the poset P_{k+1} consists of two copies of P_k and two special elements (drawn in orange). The first special point covers all minimal elements of the first P_k-copy and is covered by all maximal elements of the second P_k-copy. For the other special point it is defined the other way around. The figure on the right-hand side shows that the tree-depth grows by at most 2 in each step.

4.4 Lower Bound Constructions

In this section we provide constructions that establish the lower bounds from the table at the beginning of this chapter. The lower bound of 4 on the maximal dimension of posets whose cover graphs have path- or tree-width 2 is given by the example on the right-hand side of Figure 4.4 (the 'diamonds' that share a single element); see [26] for details.

Cover Graphs of Bounded Tree-depth. Let us begin now with the lower bound in the case that cover graphs have bounded tree-depth. That is, we show that for every $d \geqslant 1$ there exists a poset with dimension $2^{\lfloor d/2 \rfloor}$ that has a cover graph of tree-depth at most d. The following construction yields this bound (except for the trivial case $d = 1$).

Theorem 4.4.1. *There is a family of height-3 posets $(P_k)_{k\geqslant 1}$ such that for every $k \geqslant 1$ we have that $\dim(P_k) \geqslant 2^k$ and the cover graph of P_k has tree-depth at most $2k$.*

Proof. Consider the inductive construction in Figure 4.7. It is easy to see that P_k contains the standard example S_{2^k} as an induced subposet, which yields the lower bound on the dimension. Figure 4.7 also shows that tree-depth of cover(P_{k+1}) is at most two more than the tree-depth of cover(P_k). Indeed, in the illustrated rooted tree the two special points are in an ancestor-descendant relationship with all other elements, implying that every incident edge is witnessed by such a relation. As the cover graph of P_1 has tree-depth 2, we conclude that the cover graph of P_k has tree-depth at most $2k$. □

Cover Graphs of Bounded Path-width. We continue with posets whose cover graphs have bounded path-width. The following theorem is the outcome of several joint sessions with Gwenaël Joret, Piotr Micek, Stefan Felsner, and Tom Trotter.

Theorem 4.4.2. *Let $p \geqslant 3$ be fixed. Then for every integer $h \geqslant (p+3)/2$ there exists a poset P of height at most h whose cover graph has path-width at most p, and*

$$\dim(P) = \Omega(h^{(p-2)/4}).$$

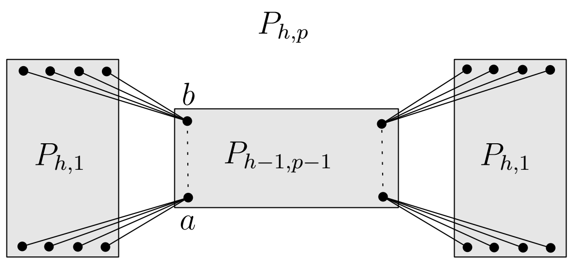

Figure 4.8: Recursive construction of $P_{h,p}$.

Before we provide the construction that establishes this theorem, let us introduce a new notation first. Let P be a poset that contains the standard example S_k as an induced subposet. Then there are elements $a_1, \ldots, a_k$ and $b_1, \ldots, b_k$ in P such that for any $i, j \in [k]$, we have that $a_i \leqslant b_j$ holds in P if and only if $i \neq j$. We call the incomparable pairs (a_i, b_i) in P the *vertical pairs* of S_k in P.

The following theorem will imply Theorem 4.4.2.

Theorem 4.4.3. *Let p, h be integers such that $h \geqslant p \geqslant 1$. Then there exists a poset $P_{h,p}$ and a path-decomposition $\mathcal{P}_{h,p}$ of its cover graph, such that*

(i) *the height of $P_{h,p}$ is $2h$,*
(ii) *the minimal and maximal elements of $P_{h,p}$ induce a standard example S_k of size $k = \prod_{i=1}^{p}(2h - 2i + 3)$.*
(iii) *for each vertical pair (a, b) of S_k in $P_{h,p}$, there is a bag of $\mathcal{P}_{h,p}$ containing both a and b.*
(iv) *$\mathcal{P}_{h,p}$ has width $4p - 1$.*

Proof. We prove the theorem by induction on p. In the base case $p = 1$ we define $P_{h,1}$ to be Kelly's example of order $2h + 1$. It is important to note that we take the *original* example from [43], which can be obtained by removing the degree-1 vertices in the example of Figure 4.1. It has height $2h$, and its minimal and maximal elements induce a standard example of size $2h + 1$; this is exactly what we need for item (ii). Figure 4.1 illustrates a path-decomposition of width 3 for the modern version of Kelly's example, and hence it also induces one for $P_{1,h}$ that satisfies item (iv). It is easy to check that this decomposition also fulfills (iii).

Suppose $h \geqslant p > 1$ now. We aim to construct $P_{h,p}$ by extending the poset $P_{h-1,p-1}$. Let S_k be the standard example of size $k = \prod_{i=1}^{p-1}(2(h-1) - 2i + 3)$ that is induced by the minimal elements and maximal elements of $P_{h-1,p-1}$. For each vertical pair (a, b) of S_k in $P_{h-1,p-1}$, we introduce a copy of $P_{h,1}$ and set $x < a$ for every $x \in \text{Min}(P_{h,1})$ and $b < y$ for every $y \in \text{Max}(P_{h,1})$; see Figure 4.8 for an illustration of this operation. The transitive closure of this construction yields the poset $P_{h,p}$.

We need to show items (i)-(iv) for $P_{h,p}$ now. Item (i) is immediate as we extend the height of $P_{h-1,p-1}$ by two units in the construction and as $P_{h-1,p-1}$ itself has height $2h - 2$ (by (i)). It is straightforward to verify that the minimal elements and

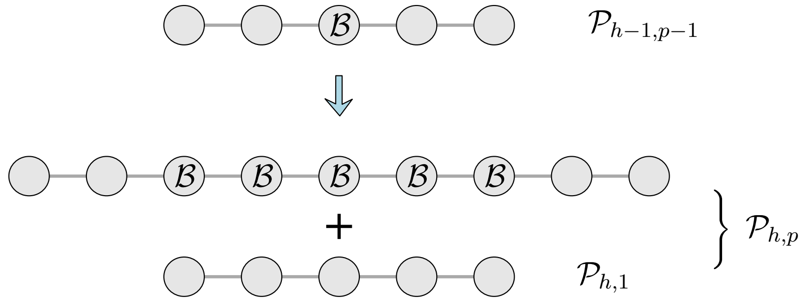

Figure 4.9: Construction of $\mathcal{P}_{h,p}$.

maximal elements of $P_{h,p}$ induce a standard example of size

$$\begin{aligned} k \cdot (2h+1) &= \prod_{i=1}^{p-1} (2(h-1) - 2i + 3) \cdot (2h+1) \\ &= \prod_{i=2}^{p} (2h - 2i + 3) \cdot (2h+1) = \prod_{i=1}^{p} (2h - 2i + 3). \end{aligned}$$

This shows item (ii) and it remains to build a path-decomposition $\mathcal{P}_{h,p}$ of the cover of $P_{h,p}$ with the desired properties.

For each vertical pair (a, b) in S_k of $P_{h-1,p-1}$, we extend the path-decomposition $\mathcal{P}_{h-1,p-1}$ in the following way. Consider a bag $\mathcal{B}$ of $\mathcal{P}_{h-1,p-1}$ containing both a and b, which exists by (iii). Now we artificially blow up $\mathcal{P}_{h-1,p-1}$ by taking many copies of $\mathcal{B}$ and putting them next to each other. That is to say, the copies of $\mathcal{B}$ build an "interval" in this extended path-decomposition. We take as many copies of $\mathcal{B}$ as there are bags in the path decomposition $\mathcal{P}_{h,1}$ of the cover graph of $P_{h,1}$. Next, we combine the sequence of $\mathcal{B}$'s with $\mathcal{P}_{h,1}$ by putting them onto each other as illustrated in Figure 4.9. From another perspective, the elements of $\mathcal{B}$ are added to each bag of $\mathcal{P}_{h,1}$ in this construction step. This particularly holds for the two elements a and b.

It is not hard to see that we can realize this operation simultaneously for each vertical pair of S_k in $P_{h-1,p-1}$ so that each bag of the starting decomposition $\mathcal{P}_{h-1,p-1}$ is joined with at most one bag of a copy of $\mathcal{P}_{h,1}$. So let $\mathcal{P}_{h,p}$ be a path-decomposition realized in this way.

Clearly, $\mathcal{P}_{h,p}$ is a path-decomposition of the cover graph of $P_{h,p}$. Moreover, since $\mathcal{P}_{h,1}$ and $\mathcal{P}_{h-1,p-1}$ have width 3 and $4(p-1)-1$, respectively, we obtain that $\mathcal{P}_{h,p}$ has width $3 + (4(p-1)-1) + 1 = 4p - 1$, which shows (iv) for $\mathcal{P}_{h,p}$. Finally, item (iii) holds for $\mathcal{P}_{h,p}$ as it already holds for each copy of $\mathcal{P}_{h,1}$ that we use in the construction. This completes the proof. $\square$

Proof of Theorem 4.4.2. Let $p \geqslant 3$ be fixed and suppose that $h \geqslant (p+3)/2$. Let q be an integer such that $p - 3 \leqslant q \leqslant p$ and $q \equiv 3 \mod 4$. Then the poset $P_{\lfloor h/2 \rfloor, \frac{q+1}{4}}$ from Theorem 4.4.3 has height at most h, it has a cover graph of path-width at most $q \leqslant p$, and it contains a standard example of size at least

$$\prod_{i=1}^{\frac{q+1}{4}} (2\lfloor h/2 \rfloor - 2i + 3) = \Omega(h^{\frac{q+1}{4}}) = \Omega(h^{\frac{p-2}{4}}).$$

Since the maximal size of a standard example contained in a poset is a lower bound on the dimension, the theorem follows. □

Cover Graphs of Bounded Tree-width. Our last construction of this section establishes an exponential lower bound on the dimension for posets whose cover graphs have bounded tree-width. This is joint work with Gwenaël Joret and Piotr Micek.

Theorem 4.4.4. *Let $t \geqslant 3$ be fixed. For each $h \geqslant 4$, there exists a poset P of height at most h whose cover graph has tree-width at most t, and*

$$\dim(P) \geqslant 2^{\Omega(h^{\lfloor (t-1)/2 \rfloor})}.$$

This theorem will be implied by the following slightly more technical theorem. We will apply similar ideas as for the previous construction. In particular, we use again the notion of vertical pairs with respect to some fixed induced standard example.

Theorem 4.4.5. *For each $h, t \geqslant 1$ there exists a poset $P_{h,t}$ and a tree-decomposition $\mathcal{T}$ of its cover graph, such that*

(i) *$P_{h,t}$ has height at most $2h$,*
(ii) *the minimal and maximal elements of $P_{h,t}$ induce the standard example of size $2^{\binom{h+t-1}{t}}$,*
(iii) *the width of $\mathcal{T}$ is at most $2t+1$,*
(iv) *for each vertical pair (a,b) in $P_{h,t}$ there is a bag in $\mathcal{T}$ containing both a and b.*

Proof. We prove the theorem by induction on h and t. Let us first show the boundary cases in which $h = 1$ or $t = 1$, which serve as the base cases for the induction. If $h = 1$, then it is easy to see that for all $t \geqslant 1$ the standard example S_2 fulfills the desired conditions. (In fact, even S_{2t+1} does the job so that we could prove a larger bound in (ii), but the gain is negligible.)

So suppose that $t = 1$. For the definition of $P_{h,1}$ we use the inductive construction shown on the left of Figure 4.10. That is, given the poset $P_{h-1,1}$, for each vertical pair (a,b) in $P_{h-1,1}$ we introduce four elements x_1, x_2, y_1, y_2 forming a standard example of size 2, and we put a above x_1 and x_2, and b below y_1 and y_2. The poset $P_{h,1}$ is then obtained by taking the transitive closure of this construction. It is easy to see that the height of $P_{h,1}$ increases by two compared to $P_{h-1,1}$, which implies (i). It is also not hard to see, that the number of minimal (resp. maximal) elements grows by a factor of two. Since we have 2^{h-1} vertical pairs in $P_{h-1,1}$, the minimal and maximal elements of $P_{h,1}$ induce the standard example of size 2^h, which yields (ii).

Now, let $\mathcal{T}$ be a tree-decomposition of the cover graph of $P_{h-1,1}$ satisfying (iii) and (iv). Given a vertical pair (a,b) of $P_{h-1,1}$, let $\mathcal{B}$ be a bag that contains both a and b. If x_1, x_2, y_1, y_2 are as above, then we can extend the tree-decomposition at $\mathcal{B}$ as shown in Figure 4.11. Clearly, if this extension is done for each vertical pair, then the resulting tree-decomposition satisfies (iii) and (iv) with respect to $P_{h,1}$.

Next, suppose that $h \geqslant 2$ and $t \geqslant 2$. Consider a copy of $P_{h-1,t}$ and its vertical pairs $(a_1,b_1), \ldots, (a_\ell,b_\ell)$. Similarly to the previous construction, we introduce for each vertical pair (a_i,b_i) a copy P^i of $P_{h,t-1}$, and we put a_i above x, for all $x \in$

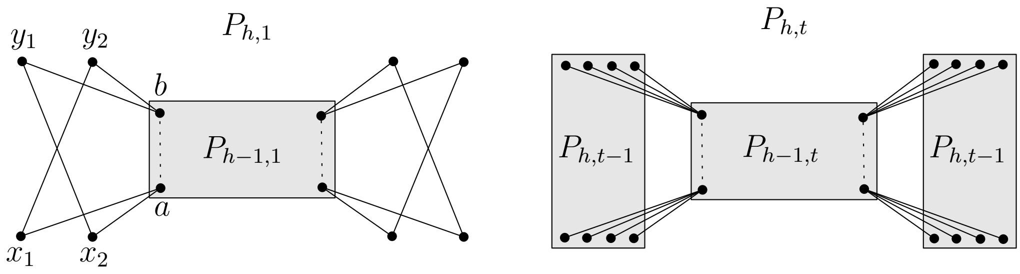

Figure 4.10: Inductive construction of $P_{h,t}$.

$\mathrm{Min}(P_i)$, and b_i below y, for all $y \in \mathrm{Max}(P_i)$ (see Figure 4.10 on the right). Then $P_{h,t}$ is obtained by taking the transitive closure of this construction.

Clearly, the height of $P_{h,t}$ is equal to h. Now observe that the minimal and maximal elements of $P_{h,t}$ induce a standard example, and by the induction hypothesis it has size

$$2^{\binom{h+t-2}{t}} \cdot 2^{\binom{h+t-2}{t-1}} = 2^{\binom{h+t-1}{t}}.$$

Let $\mathcal{T}'$ be a tree-decomposition obtained by applying the induction hypothesis to $P_{h-1,t}$. Again, we will extend $\mathcal{T}'$ in a clever way. To do so, let (a_i, b_i) be an arbitrary vertical pair of $P_{h-1,t}$. Using (iv) we know that there is a bag $\mathcal{B}$ of $\mathcal{T}$ containing both a_i and b_i. Now let $\mathcal{T}_i$ be a tree-decomposition of the cover graph of P_i, such that (iii) and (iv) are satisfied. Then we add an edge between the node of $\mathcal{T}'$ corresponding to $\mathcal{B}$ and an arbitrary node of $\mathcal{T}_i$. We put a_i and b_i into every bag of $\mathcal{T}_i$.

Let $\mathcal{T}$ be the tree-decomposition that we obtain by repeating the last step for every vertical pair of $P_{h-1,t}$. Using the induction hypothesis, it is straightforward to verify that $\mathcal{T}$ has width at most $2t + 1$ and that for each vertical pair (a, b) of $P_{h,t}$ there is a bag of $\mathcal{T}$ containing both a and b. □

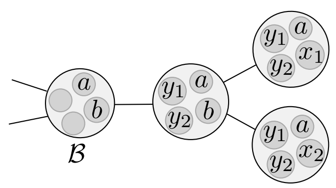

Figure 4.11: Extending the tree-decomposition.

Proof of Theorem 4.4.4. Let $t \geqslant 3$ and $h \geqslant 4$ as in the assumptions of the statement. Then we set $h' := \lfloor h/2 \rfloor$ and $t' = \lfloor (t-1)/2 \rfloor$. With these values, the poset $P_{h',t'}$ from Theorem 4.4.5 has height at most $2h' \leqslant h$ and its cover graph has tree-width at most $2t' + 1 \leqslant t$. Moreover,

$$\dim(P_{h',t'}) \geqslant 2^{\binom{h'+t'-1}{t'}} = 2^{\Omega(h'^{t'})} = 2^{\Omega(h^{\lfloor (t-1)/2 \rfloor})}.$$

(Regarding the last equality, recall that the asymptotics are taken with respect to h and that the parameter t is treated as a constant.) □

We conclude the chapter with a simple consequence of Theorem 4.4.4. It is well known that graphs of tree-width at most t contain no K_{t+2} minor. Therefore, the previous construction also yields the following result.

Corollary 4.4.1. *For each $t \geqslant 5$ and each $h \geqslant 4$, there exists a poset P of height at most h whose cover graph excludes K_t as a minor, such that it has dimension*

$$\dim(P) = 2^{\Omega(h^{\lfloor (t-3)/2 \rfloor})}.$$

Chapter 5

Cover Graphs with Excluded Topological Minor

In previous chapters we have already studied posets whose cover graphs belong to some proper minor-closed graph class. This study particularly captured the cases when cover graphs are planar, have bounded path-width, or even have bounded tree-width. It is not hard to see that in all these cases there exists a graph that cannot be a minor of those cover graphs: Planar graphs have no K_5-minor, and graphs of path- or tree-width at most t exclude K_{t+2} as a minor.

In this chapter we want to study posets whose cover graphs exclude a fixed graph as a minor (and even more general, as a topological minor). Walczak [76] first proved that their dimension is bounded from above by a function in their height.

Theorem 5.0.1 (Walczak [76])**.** *Posets of height at most h whose cover graphs exclude K_n as a topological minor have dimension bounded in terms of h and n.*

Walczak's argument is based on structural decomposition theorems due to Robertson and Seymour [66], and Grohe and Marx [34]. Although the proof is not long, it is not very applicable to readers not familiar with these structure theorems. Moreover, the upper bound on the dimension must be enormous (it is not calculated in the paper).

In Section 5.1, we give an alternative proof of Walczak's result, which is joint work with Piotr Micek [51]. Our argument is entirely combinatorial and does not rely on the decomposition theorems mentioned above. We are able to state explicit upper bounds on the dimension, which are double exponential in the poset height. Admittedly, we can show even better bounds nowadays; in Chapter 6 we derive single exponential upper bounds, which are essentially best possible because of the lower bound constructions from the previous chapter. Nevertheless, we believe that our developed tools and ideas of Section 5.1 are important and useful. This claim is supported by the fact that we establish a framework for proving upper bounds on the dimension that can also be used in several other settings. For instance, we apply the same framework a second time on so-called $(\mathbf{k}+\mathbf{k})$-free posets in Section 5.2. An idea that was used for the first time in this framework is the iterated application of the *Unfolding Lemma*. Subsequently, it was also used in the context of planar posets [40] (see Chapter 3) and posets whose cover graphs belong to some class with bounded expansion [42].

We continue our study of posets whose cover graphs exclude a fixed graph as a topological minor in Section 5.2. This time though, we do not require the posets to have bounded height, which can also be seen as excluding a chain as an induced subposet. Instead, we only forbid the poset P consisting of two incomparable chains of size k. This means that no element of one chain is comparable to any element of the other chain in P. We denote the poset P by $\mathbf{k}+\mathbf{k}$ from now on.

Over the last few years, a number of nice results [16, 23, 49] emerged pointing out that problems difficult for the class of all posets might be tractable for $(\mathbf{k}+\mathbf{k})$-free posets. For instance, any on-line algorithm that is ask to build a realizer for posets of width at most w can be forced to use arbitrarily many linear extensions (even when $w = 3$) [46]. For $(\mathbf{k}+\mathbf{k})$-free posets however, Felsner et al. [23] devised an on-line algorithm that constructs realizers of bounded size.

Motivated by results of this type and by the simple observation that Kelly's family of planar posets is not $(\mathbf{k}+\mathbf{k})$-free for any $k \geqslant 1$, it was asked by several colleagues whether the dimension of $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude a fixed graph as a minor is bounded (e.g. see [76, 51, 37]). Quite recently, Howard et al. [37] gave support towards a positive resolution of this question as they proved that the dimension of $(\mathbf{k}+\mathbf{k})$-free posets with planar cover graphs is bounded from above by a function in k.

We will not answer the aforementioned question in Section 5.2. However, we give further evidence that the correct answer is "yes" by showing that those posets contain only standard examples of bounded size. This is again joint work with Piotr Micek.

Theorem 5.0.2 ([51]). *The $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude K_n as a topological minor contain only standard examples of size bounded in terms of k and n.*

Let us note here that in the conclusions of this thesis we discuss some interesting consequences of Theorem 5.0.2. Moreover, we elaborate on future research directions regarding large-dimensional $(\mathbf{k}+\mathbf{k})$-free posets there.

5.1 First Case: Posets of Bounded Height

The aim of this section is to give a proof of Theorem 5.0.1: Posets P of height at most h whose cover graphs exclude K_n as a topological minor have dimension bounded in terms of h and n.

Let us begin with an outline of the proof. Given a poset P with large dimension but bounded height, we are going to construct a subdivision of K_n in the cover graph of P. First, in a preprocessing step, we apply the *Min-Max Reduction* to P to ensure that the dimension of P is witnessed by its minimal and maximal elements. Then we will run through two phases. In Phase 1 we are going to set up two collections of disjoint sets in P such that the elements of each set induce a connected subgraph of P's cover graph. We start with two empty collections and then we iteratively apply the *Unfolding Lemma* to get new sets for the collections. At the end of the Phase 1, one of the two collections will be large enough. This collection is refined in Phase 2, where we fix n elements of P that will be the principal vertices of a K_n-subdivision in the cover graph of P. We conclude Phase 2 with at least $\binom{n}{2}$ sets remaining in our collection. Finally, we use the sets of the collection to connect the fixed vertices. This will yield a subdivision of K_n.

We continue with the detailed proof now. Let us omit the trivial cases of the theorem and assume $n \geqslant 3$ and $h \geqslant 2$. Let P be a poset with $\text{height}(P) \leqslant h$ and

$$\dim(P) > n^L, \text{ where } L = 2 \cdot \binom{M+h}{h} - 1 \text{ and } M = \binom{n}{2}^{h^n}.$$

We start by applying the *Min-Max Reduction* (Lemma 2.1.1) to P and obtain a poset P' of height at most h such that $\dim(\text{Min}(P'), \text{Max}(P')) \geqslant \dim(P) > n^L$. Note that by item (iii) of Lemma 2.1.1 we also have that $\text{cover}(P')$ can be obtained by attaching degree-1 vertices to $\text{cover}(P)$. Therefore, if we find a subdivision of K_n in the cover graph of P', then this subdivision also exists in the cover graph of P (recall that $n \geqslant 3$). So all we need to show is that $\dim(\text{Min}(P'), \text{Max}(P')) > n^L$ implies the existence of a K_n topological minor in $\text{cover}(P')$. For convenience, from now on we write P instead of P'.

The Data Structure of Phase 1 and its Invariants

During Phase 1 we maintain an additional structure $(A, B, \mathcal{C}, \mathcal{D})$ while running a loop. After the i-th loop iteration we have the following invariants:

(a) $A \subseteq \text{Min}(P)$, $B \subseteq \text{Max}(P)$, and $\dim(A, B) > n^{L-i}$,

(b) $\mathcal{C}$ is a collection of pairwise disjoint subsets of P with $|\mathcal{C}| \leqslant M$ and
- (b.1) C is connected in $\text{cover}(P)$, for every $C \in \mathcal{C}$,
- (b.2) $A \cap \text{D}(C) = \emptyset$ and $B \subseteq \text{U}(C)$, for every $C \in \mathcal{C}$,
- (b.3) $\text{D}(c) \cap C = \emptyset$ for each singleton c in $\mathcal{C}$ and each $C \in \mathcal{C} - \{c\}$,

(c) $\mathcal{D}$ is a collection of pairwise disjoint subsets of P with $|\mathcal{D}| \leqslant M$ and
- (c.1) D is connected in $\text{cover}(P)$, for every $D \in \mathcal{D}$,
- (c.2) $A \subseteq \text{D}(D)$ and $B \cap \text{U}(D) = \emptyset$, for every $D \in \mathcal{D}$,
- (c.3) $\text{U}(d) \cap D = \emptyset$ for each singleton d in $\mathcal{D}$ and each $D \in \mathcal{D} - \{d\}$.

We also have a measure of quality of the maintained structure. For each $C \in \mathcal{C}$ with $|C| > 1$, we let $\text{value}(C) = h$, and if $C = \{c\}$, we let $\text{value}(C)$ be the size of the longest covering chain in $\text{cover}(P)$ from c to any $b \in B$. For each $D \in \mathcal{D}$ with $|D| > 1$, we let $\text{value}(D) = h$, and if $D = \{d\}$, we let $\text{value}(D)$ be the size of the longest covering chain in $\text{cover}(P)$ from any $a \in A$ to d. Since the height of P is at most h, $\text{value}(X) \leqslant h$ holds for every $X \in \mathcal{C} \cup \mathcal{D}$. Note also that $\text{value}(C)$ and $\text{value}(D)$ depend on the current sets A and B within the structure.

Next, we define the value of a collection $\mathcal{X}$ of subsets of P, denoted by $\text{value}(\mathcal{X})$, to be the sequence of size M sorted in a non-decreasing order with one entry $\text{value}(X)$ for each $X \in \mathcal{X}$, and with $M - |\mathcal{X}|$ positions filled with '$h+1$' values. Note that $\text{value}(\mathcal{C})$ and $\text{value}(\mathcal{D})$ are sequences of length M with sorted values from the set $\{1, \ldots, h+1\}$. Therefore, there are at most $\binom{M+h}{h}$ possible values for $\mathcal{C}$ and $\mathcal{D}$, respectively. We write $\text{value}(\mathcal{X}') < \text{value}(\mathcal{X})$, if there is an index $j \in [M]$ such that the first $j-1$ entries of $\text{value}(\mathcal{X}')$ and $\text{value}(\mathcal{X})$ are the same, and the j-th entry of $\text{value}(\mathcal{X}')$ is smaller than the j-th entry of $\text{value}(\mathcal{X})$ (which is also known as the *lexicographic order*).

During Phase 1 the values of the collections $\mathcal{C}$ and $\mathcal{D}$ will decrease. In such a case we say that the quality of our maintained structure is improving. Intuitively, a small value of $\mathcal{C}$ is good since then the sets in $\mathcal{C}$ are somehow close to all elements in B, which makes it easier to construct a topological minor in the cover graph.

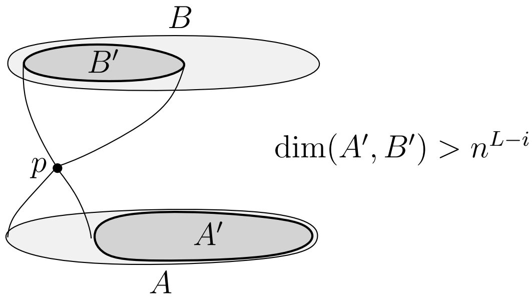

Figure 5.1: Definition of A' and B' in the case of a "yes"-answer for question (Q1).

Phase 1: Updating the Data Structure

We set up the initial structure as follows: $A = \mathrm{Min}(P)$, $B = \mathrm{Max}(P)$, and $\mathcal{C}$ and $\mathcal{D}$ are empty. Clearly, conditions (a)-(c) hold for $i = 0$. Note that we start with value($\mathcal{C}$) and value($\mathcal{D}$) being the sequence with M entries of '$h+1$'.

Now we run a loop to improve the quality of the data structure. In each iteration we ask up to three questions about the current structure $(A, B, \mathcal{C}, \mathcal{D})$. If we get only negative answers, then the loop will terminate and Phase 1 is done. If we get a positive answer to one of the questions, then we finish the iteration by updating the structure to $(A', B', \mathcal{C}', \mathcal{D}')$ that is satisfying conditions (a)-(c) and additionally

$$\text{value}(\mathcal{C}') < \text{value}(\mathcal{C}) \text{ and } \text{value}(\mathcal{D}') \leqslant \text{value}(\mathcal{D}),$$
$$\text{or}$$
$$\text{value}(\mathcal{C}') \leqslant \text{value}(\mathcal{C}) \text{ and } \text{value}(\mathcal{D}') < \text{value}(\mathcal{D}).$$

Since the number of values for the collection $\mathcal{C}$ (and $\mathcal{D}$, resp.) is bounded by $\binom{M+h}{h}$, the quality of our structure can be improved at most $L-1 = 2\binom{M+h}{h}-2$ times. Thus there will be at most L iterations in total. (If we reach the L-th iteration, then the structure is not improvable anymore, implying that we necessarily get only negative answers to the three questions.)

Now we describe the i-th ($1 \leqslant i \leqslant L$) iteration in detail. Let $(A, B, \mathcal{C}, \mathcal{D})$ be the current structure. Hence it satisfies conditions (a)-(c) with respect to $i-1$. The iteration starts with the evaluation of the following question:

Is there an element $p \in P$ such that (Q1)

(i) $\dim(A, B \cap \mathrm{U}(p)) > n^{L-i}$, and
(ii) there is $C \in \mathcal{C}$ and $c \in C$ such that $c < p$ in P?

First, suppose that the answer is "yes" and fix such an element $p \in P$. In this case we finish the i-th iteration by updating the structure to $(A', B', \mathcal{C}', \mathcal{D}')$, where

$$\begin{aligned} A' &= A - \mathrm{D}(p), \\ B' &= B \cap \mathrm{U}(p), \\ \mathcal{C}' &= \mathcal{C} \cup \{p\} - \{C \in \mathcal{C} \mid p \in \mathrm{U}(C)\}, \\ \mathcal{D}' &= \mathcal{D}. \end{aligned}$$

See Figure 5.1 for a visualization of the sets A', B'.

Claim 5.1.1. *The structure* $(A', B', \mathcal{C}', \mathcal{D}')$ *satisfies the invariants* (a)-(c). *Moreover,* value($\mathcal{C}'$) < value($\mathcal{C}$) *and* value($\mathcal{D}'$) ⩽ value($\mathcal{D}$).

Proof. Clearly, $A' \subseteq A$ is a set of minimal elements and $B' \subseteq B$ is a set of maximal elements in P. Since the answer for (Q1) is "yes" and all elements in $A \cap \mathrm{D}(p)$ are below all elements in $B \cap \mathrm{U}(p)$, we have

$$\dim(A', B') = \dim(A, B \cap \mathrm{U}(p)) > n^{L-i},$$

so (a) holds.

Next, we show that $\mathcal{C}'$ satisfies all conditions of (b). The set $\{p\}$ cannot be contained in $\mathcal{C}$. Indeed, otherwise item (ii) of question (Q1) would yield a contradiction to invariant (b.3) for $\mathcal{C}$. Thus, $\{p\}$ is a new set in $\mathcal{C}'$ compared to $\mathcal{C}$. In order to prove that the sets in $\mathcal{C}'$ are pairwise disjoint we only need to argue that $p \notin C$ for every $C \in \mathcal{C}' - \{p\}$. But this follows immediately from the definition of $\mathcal{C}'$. In order to prove $|\mathcal{C}'| \leqslant M$, note that p witnesses the positive answer for (Q1), so there are $C \in \mathcal{C}$ and $c \in C$ with $c < p$ in P, and therefore $C \notin \mathcal{C}'$. This implies $|\mathcal{C}'| \leqslant |\mathcal{C}| \leqslant M$.

Item (b.1) trivially holds. Item (b.2) for old sets in $\mathcal{C}'$ follows immediately from the same invariant for $\mathcal{C}$, and for $\{p\}$ it follows from the definition of sets A' and B'. For item (b.3), observe first that $\mathrm{D}(p) \cap C = \emptyset$ for every $C \in \mathcal{C}' - \{p\}$ by the definition of $\mathcal{C}'$. We also have to argue that $p \notin \mathrm{D}(c')$ for each singleton c' in $\mathcal{C}' - \{p\} \subseteq \mathcal{C}$. Suppose to the contrary that $p \leqslant c'$ in P. Recall that there is $C \in \mathcal{C}$ and $c \in C$ with $c < p$. However, this implies $c < c'$ in P, contradicting (b.3) for $\mathcal{C}$. This completes the verification of (b) for $\mathcal{C}'$. Since $\mathcal{D}' = \mathcal{D}$ and $A' \subseteq A$, $B' \subseteq B$, condition (c) still holds.

Now we show that the quality of the new structure improved. We aim to prove $\mathrm{value}(\mathcal{C}') < \mathrm{value}(\mathcal{C})$ and $\mathrm{value}(\mathcal{D}') \leqslant \mathrm{value}(\mathcal{D})$. The values of sets in $\mathcal{C} \cap \mathcal{C}'$ cannot increase with the update as $B' \subseteq B$. Thus, to show $\mathrm{value}(\mathcal{C}') < \mathrm{value}(\mathcal{C})$ it is enough to argue that $\mathrm{value}(\{p\})$ is smaller than the value of each set removed from $\mathcal{C}$. So let $C \in \mathcal{C}$ such that $p \in \mathrm{U}(C)$. If $|C| > 1$ then we have $\mathrm{value}(C) = h$ by our definition. Note that there is no covering chain of size h starting from p, since otherwise p must be minimal in P, which is not true by item (ii) of question (Q1). Hence $\mathrm{value}(\{p\}) \leqslant h - 1 < \mathrm{value}(C)$. If $|C| = 1$, that is $C = \{c\}$ for some $c \in P$, then recall that p is not a singleton of $\mathcal{C}$ and hence $p \neq c$. Therefore, $p \in \mathrm{U}(C)$ implies $c < p$ in P, which in turn yields $\mathrm{value}(\{p\}) < \mathrm{value}(\{c\})$. Thus we indeed have $\mathrm{value}(\mathcal{C}') < \mathrm{value}(\mathcal{C})$.

Finally, $\mathrm{value}(\mathcal{D}') \leqslant \mathrm{value}(\mathcal{D})$ holds as $\mathcal{D}' = \mathcal{D}$ and $A' \subseteq A$. This completes the verification of the invariants for the updated structure $(A', B', \mathcal{C}', \mathcal{D}')$ in the case of a 'yes' answer to question (Q1). □

If the answer for question (Q1) is "no", then the procedure continues with a dual question:

$$\text{Is there an element } p \in P \text{ such that} \tag{Q2}$$

(i) $\dim(A \cap \mathrm{D}(p), B) > n^{L-i}$, and
(ii) there is $D \in \mathcal{D}$ and $d \in D$ such that $p < d$ in P?

If the answer for (Q2) is "yes", then we improve the current structure analogously to the "yes"-case of question (Q1). We finish the i-th iteration by updating the structure to $(A', B', \mathcal{C}', \mathcal{D}')$, where

$$\begin{aligned} A' &= A \cap \mathrm{D}(p), \\ B' &= B - \mathrm{U}(p), \\ \mathcal{C}' &= \mathcal{C}, \\ \mathcal{D}' &= \mathcal{D} \cup \{p\} - \{D \in \mathcal{D} \mid p \in \mathrm{D}(D)\}. \end{aligned}$$

The proof that this new structure satisfies conditions (a)-(c) and that it improves the quality goes dually to the one for question (Q1).

If the answers for questions (Q1) and (Q2) are both "no", then the procedure continues with a third question:

$$\text{Are } |\mathcal{C}| < M \text{ and } |\mathcal{D}| < M? \tag{Q3}$$

Again, we first deal with the "yes"-answer. In this case, we show how to find a new candidate set to extend $\mathcal{C}$ or $\mathcal{D}$. We apply the *Extended Unfolding Lemma* (Lemma 2.2.3) with respect to the sets A and B in P (recall that $\dim(A, B) > n^{L-(i-1)} \geqslant 3$) and obtain that there exist sets $A' \subseteq A$, $B' \subseteq B$, and $S \subseteq \mathrm{conv}(A \cup B)$ satisfying

(i) elements of S induce a connected subgraph of $\mathrm{cover}(P)$,
(ii) $\dim(A', B') \geqslant \dim(A, B)/2$,
(iii) it holds that either

$$\begin{aligned} &(1) \quad A' \cap \mathrm{D}(S) = \emptyset \quad \text{and} \quad B' \subseteq \mathrm{U}(S), \\ &\qquad\qquad\qquad\qquad \text{or} \\ &(2) \quad B' \cap \mathrm{U}(S) = \emptyset \quad \text{and} \quad A' \subseteq \mathrm{D}(S). \end{aligned}$$

First, we consider the case that part (1) of (iii) holds. In this case we finish the i-th iteration by updating the structure to $(A', B', \mathcal{C}', \mathcal{D}')$, where $\mathcal{C}' = \mathcal{C} \cup \{S\}$ and $\mathcal{D}' = \mathcal{D}$.

Claim 5.1.2. *The structure* $(A', B', \mathcal{C}', \mathcal{D}')$ *keeps the invariants* (a)-(c). *Moreover,* $\mathrm{value}(\mathcal{C}') < \mathrm{value}(\mathcal{C})$ *and* $\mathrm{value}(\mathcal{D}') \leqslant \mathrm{value}(\mathcal{D})$.

Proof. Clearly, $A' \subseteq A$ is a set of minimal and $B' \subseteq B$ is a set of maximal elements in P and

$$\dim(A', B') \geqslant \dim(A, B)/2 > n^{L-(i-1)}/2 > n^{L-i}$$

as $n > 2$, so (a) holds.

To argue that $\mathcal{C}'$ is a set of disjoint sets, we need to check whether S is disjoint from every $C \in \mathcal{C}$. This holds, since in particular $S \subseteq \mathrm{U}(A)$ and on the other hand $\mathrm{U}(A) \cap C = \emptyset$ (by (b.2)), for every $C \in \mathcal{C}$. Note also that all sets in $\mathcal{C}'$ are connected in $\mathrm{cover}(P)$; this follows from (b.1) for $\mathcal{C}$ and the fact that S itself is connected in $\mathrm{cover}(P)$. This proves (b.1) for $\mathcal{C}'$. Since $A' \subseteq A$ and $B' \subseteq B$, invariant (b.2) remains true for all $C \in \mathcal{C}$. The new set S has the required property in (b.2) explicitly. Therefore, invariant (b.2) holds for the whole collection $\mathcal{C}'$.

For invariant (b.3), let c be a singleton of $\mathcal{C}$ and suppose that there is $z \in \mathrm{D}(c) \cap S$. Since $S \subseteq \mathrm{U}(A) \cap \mathrm{D}(B)$, there exists $a \in A$ with $a \leqslant z$ in P. However, this implies $a \leqslant c$ in P and hence $\mathrm{D}(c) \cap A \neq \emptyset$, contradicting (b.2) for $\{c\} \in \mathcal{C}$. Therefore, $D(c) \cap S = \emptyset$ for each singleton c in $\mathcal{C}$. To complete the verification of (b.3) we also have to consider the case in which S contains only one element. So say we have $S = \{s\}$ and suppose to the contrary that $\mathrm{D}(s) \cap C \neq \emptyset$ for some $C \in \mathcal{C}$. Let $z \in \mathrm{D}(s) \cap C$ be an element in the intersection. By the disjointedness of the sets in $\mathcal{C}'$, it holds that $z \neq s$ and hence $z < s$ in P. Observe that s fulfills the conditions of question (Q1), contradicting the fact that this question was answered with "no". This establishes (b.3) for $\mathcal{C}'$.

Since $\mathcal{D}' = \mathcal{D}$, invariant (c) for $\mathcal{D}'$ is immediate. Finally, $\mathrm{value}(\mathcal{C}') < \mathrm{value}(\mathcal{C})$ and $\mathrm{value}(\mathcal{D}') \leqslant \mathrm{value}(\mathcal{D})$ hold as $\mathcal{C} \subsetneq \mathcal{C}'$ and $\mathcal{D} = \mathcal{D}'$. □

Let us now consider the case that (2) of (iii) holds. Then we finish the i-th iteration by updating the structure to $(A', B', \mathcal{C}', \mathcal{D}')$, where $\mathcal{C}' = \mathcal{C}$ and $\mathcal{D}' = \mathcal{D} \cup \{S\}$. The proof that the invariants are kept in this case goes along similar arguments as in the first case. Therefore, we omit the details here.

Finally, if the answer to all questions (Q1)-(Q3) is "no", then we stop iterating and Phase 1 is done.

Let $(A, B, \mathcal{C}, \mathcal{D})$ be the final data structure of Phase 1 and suppose it is obtained in the i-th loop iteration (so $i \in \{0, \ldots, L-1\}$). Thus, $(A, B, \mathcal{C}, \mathcal{D})$ satisfies the invariants (a)-(c) and the answers to questions (Q1)-(Q3) were "no" for $(A, B, \mathcal{C}, \mathcal{D})$ in the $(i+1)$-th loop iteration. The negative answer for (Q3) tells us that

$$|\mathcal{C}| = M \quad \text{or} \quad |\mathcal{D}| = M.$$

For the rest of the proof we assume that $|\mathcal{C}| = M$. (We may do this, since the other case can be treated dually. Indeed, observe that $\mathcal{C}$ and $\mathcal{D}$ satisfy properties that are dual to each other, and hence we could apply all of the following arguments to the dual poset of P in the case $|\mathcal{D}| = M$.)

In a moment we will start with Phase 2, which consists of a loop that has n iterations. In each iteration we find a new principal vertex for the final construction of a K_n subdivision. Simultaneously, we refine the collection $\mathcal{C}$ maintaining a large enough subcollection that interacts well with vertices already fixed.

Let us go more into detail now. It will be convenient to use the following definition. For a family $\mathcal{F}$ of subsets of P and an element $p \in P$, we define

$$\mathcal{F}^p = \{F \in \mathcal{F} \mid p \in \mathrm{U}(F)\}.$$

While running the loop of Phase 2, we maintain as an invariant a pair $(V, \mathcal{E})$ with $V \subseteq P$ and $\mathcal{E} \subseteq \mathcal{C}$, that is satisfying the following items after the j-th loop iteration:

- (d.1) $|V| = j$ and $|\mathcal{E}| \geqslant M^{(1/h)^j} = \binom{n}{2}^{h^{n-j}}$,
- (d.2) V is disjoint from every $C \in \mathcal{E}$, and
- (d.3) for every $v \in V$ and $C \in \mathcal{E}$, there is $x \in P$ such that x is covered by v in P and $\mathcal{E}^x = \{C\}$.

Condition (d.3) particularly tells us that for every $v \in V$ there exist a set of covering chains such that the following holds: For each $C \in \mathcal{E}$ there is a covering chain in the set starting in C and ending in v, and these covering chains are pairwise disjoint except for their common endpoint v.

Phase 2: Selecting the Principal Vertices

Before the first iteration we set up the pair $(V, \mathcal{E})$ with $V = \emptyset$ and $\mathcal{E} = \mathcal{C}$. Invariants (d.1)-(d.3) are satisfied for $j = 0$ vacuously.

Now we describe the j-th iteration of the loop ($1 \leqslant j \leqslant n$). Let $(V, \mathcal{E})$ be the pair satisfying the invariants after the $(j-1)$-th iteration. The main issue is to find a new vertex to put into V. We start to look for it from an appropriate vertex in B. We want to pick any vertex from $B - \bigcup_{v \in V} \mathrm{U}(v)$, so we need to argue that this set is not empty. By invariant (d.3), we get in particular that for every $v \in V$ there is $C \in \mathcal{C}$ and $c \in C$ such that $c < v$ in P. Since the answer to question (Q1) was "no" in Phase 1, we have $\dim(A, B \cap \mathrm{U}(v)) \leqslant n^{L-(i+1)}$ for every $v \in V$. Thus,

$$\begin{aligned}\dim\Big(A, B - \bigcup_{v \in V} \mathrm{U}(v)\Big) &\geqslant \dim(A, B) - \sum_{v \in V} \dim(A, B \cap \mathrm{U}(v)) \\ &> n^{L-i} - |V| \cdot n^{L-(i+1)} \\ &> n^{L-i} - n \cdot n^{L-(i+1)} = 0.\end{aligned}$$

It follows that the left-hand side is at least 2 and hence $B - \bigcup_{v \in V} \mathrm{U}(v)$ is not empty. Let b any element in this set.

Now starting from the element b we go down along cover relations in the poset P. Let $M_0 = M^{(1/h)^{j-1}}$. Recall that by $\mathcal{E} \subseteq \mathcal{C}$ and (b.2) for data structure $(A, B, \mathcal{C}, \mathcal{D})$ we know that each element in B lies in the upset of each set in $\mathcal{E}$. Thus we have $\mathcal{E}^b = \mathcal{E}$ and together with (d.1) we obtain $|\mathcal{E}^b| = |\mathcal{E}| \geqslant M_0$. Initially we set $v = b$, and as long as there is an element $x \in P$ such that x is covered by v in P and

$$|\mathcal{E}^x| > |\mathcal{E}^v| / M_0^{1/(h-1)}, \quad \text{we update } v = x. \tag{5.1}$$

Note that the process must stop as the height of v in P is decreasing in every move. Furthermore, v never goes down to a minimal element of P. Indeed, if $x < v$ and x is minimal in P, then at most $h - 2$ steps were done and hence $|\mathcal{E}^v| > |\mathcal{E}^b| / M_0^{(h-2)/(h-1)} \geqslant M_0^{1/(h-1)}$. However, $|\mathcal{E}^x| \leqslant 1 = M_0^0$ as all sets in $\mathcal{E}^x$ must contain x when x is minimal in P, and the sets in $\mathcal{E} \subseteq \mathcal{C}$ are pairwise disjoint (by (b)). Therefore, the inequality of (5.1) cannot hold strictly if x is minimal in P and we conclude that $v \notin \mathrm{Min}(P)$ at the end of the process.

Again, by invariant (b) there is at most one set in $\mathcal{E}$ containing v. If such a set C exists we define $\mathcal{E}_* = \mathcal{E} - \{C\}$, and otherwise we let $\mathcal{E}_* = \mathcal{E}$. In following we want to use the expression $\mathcal{E}_*^x$ for the set $(\mathcal{E}_*)^x$ (note that we have to specify it here as the expression is ambiguous).

Now consider the set X consisting of all elements that are covered by v in P. As no set in $\mathcal{E}_*$ contains v, we have $\mathcal{E}_*^v = \bigcup_{x \in X} \mathcal{E}_*^x$. We want to ignore somewhat redundant sets in this union, so we take a minimal subset X' of X such that

$$\mathcal{E}_*^v = \bigcup_{x \in X'} \mathcal{E}_*^x.$$

The minimality of X' allows us to fix for every $x \in X'$ a set

$$C^x \in \mathcal{E}_*^x - \bigcup_{y \in X' - \{x\}} \mathcal{E}_*^y. \tag{5.2}$$

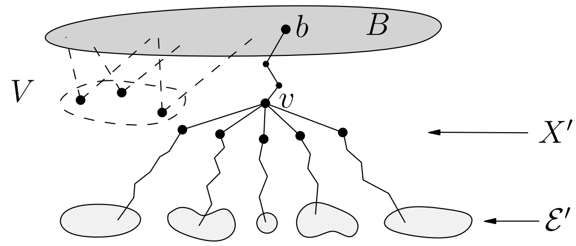

Figure 5.2: Element b in $B - \mathrm{U}(V)$ and a new element v and its cover relations downwards in P.

Finally, we update our maintained pair to $(V', \mathcal{E}')$, where

$$V' = V \cup \{v\} \quad \text{and} \quad \mathcal{E}' = \{C^x \mid x \in X'\}.$$

See Figure 5.2 for an illustration of the sets X', $\mathcal{E}'$ and their interaction. This finishes the j-th iteration of the loop.

Claim 5.1.3. *The pair* $(V', \mathcal{E}')$ *fulfills the invariants* (d.1)-(d.3).

Proof. First of all, we show that $v \notin V$ and therefore $|V'| = j$. Recall that we have chosen $b \in B$ such that $w \not\leqslant b$ in P, for every $w \in V$. On the other hand, by our procedure we have $v \leqslant b$ in P and hence $v \notin V$.

Now we aim to get the required lower bound for $|\mathcal{E}'|$. Since the sets C^x, C^y are distinct for distinct $x, y \in X'$, we have $|\mathcal{E}'| = |X'|$. Moreover,

$$\begin{aligned} |\mathcal{E}^v| \leqslant |\mathcal{E}^v_*| + 1 = \left| \bigcup_{x \in X'} \mathcal{E}^x_* \right| + 1 \leqslant \sum_{x \in X'} |\mathcal{E}^x_*| + 1 \leqslant \sum_{x \in X'} |\mathcal{E}^x| + 1 \\ \leqslant |X'| \cdot |\mathcal{E}^v| / M_0^{1/(h-1)} + 1. \end{aligned}$$

Since $|\mathcal{E}^v| > M_0^{1/(h-1)} \geqslant 1$ and $M_0 \geqslant 2^{h(h-1)}$ (this follows from (d.1) and $n \geqslant 3$), we deduce that

$$|\mathcal{E}'| = |X'| \geqslant \frac{|\mathcal{E}^v| - 1}{|\mathcal{E}^v|} \cdot M_0^{1/(h-1)} \geqslant \frac{1}{2} M_0^{1/(h-1)} \geqslant M_0^{1/h} = M^{(1/h)^j}.$$

This proves (d.1).

Invariant (d.2) holds as V is disjoint from every set in $\mathcal{E} \supseteq \mathcal{E}'$ and as v is not contained in any set of $\mathcal{E}_* \supseteq \mathcal{E}'$.

It remains to verify (d.3) for $(V', \mathcal{E}')$. Since $\mathcal{E}' \subseteq \mathcal{E}$, we only need to check this condition for the new vertex v. Consider a set $C \in \mathcal{E}'$. By the definition of $\mathcal{E}'$ there is $x \in X'$ such that $C^x = C$. Recalling the way we defined C^x (see (5.2)), we obtain $\mathcal{E}'^x = \{C\}$. This completes the verification of the invariants for $(V', \mathcal{E}')$. □

For the rest of the proof let $(V, \mathcal{E})$ denote the pair satisfying the invariants (d.1)-(d.3) after the n-th iteration of Phase 2. We are now ready to connect the vertices of V, which serve as the principal vertices of the K_n topological minor, by internally disjoint paths.

Final Step: Connecting the Principal Vertices by Paths

By (d.1) we have

$$|V| = n \quad \text{and} \quad |\mathcal{E}| \geqslant \binom{n}{2}.$$

Since $|\mathcal{E}| \geqslant \binom{n}{2}$, for every pair of distinct elements $v_1, v_2 \in V$ we can fix a unique set $C_{v_1v_2} \in \mathcal{E}$.

By invariant (d.3) there are cover relations $x_1 < v_1$ and $x_2 < v_2$ in P such that $\mathcal{E}^{x_1} = \mathcal{E}^{x_2} = \{C_{v_1v_2}\}$. In particular, there are $c_1, c_2 \in C_{v_1v_2}$ such that $c_1 \leqslant x_1 < v_1$ and $c_2 \leqslant x_2 < v_2$ in P. Let

$$c_1 = y_1 < y_2 < \cdots < y_r = x_1 \quad \text{and} \quad c_2 = z_1 < z_2 < \cdots < z_s = x_2$$

be covering chains in P. Fix a path $Q_{v_1v_2}$ connecting v_1 and v_2 in cover(P) using only vertices from the sets $\{y_1, \ldots, y_r, v_1\}$, $\{z_1, \ldots, z_s, v_2\}$, and $C_{v_1v_2}$. Such a path exists since the elements of $C_{v_1v_2}$ induce a connected subgraph of cover(P) (by (b.1)).

We claim that the union of these paths forms a subdivision of K_n in cover(P). All we need to prove is that whenever there is $z \in Q_{v_1v_2} \cap Q_{v'_1v'_2}$ for distinct two-sets $\{v_1, v_2\}, \{v'_1, v'_2\} \subseteq V$, then z is an endpoint of both paths. So fix such an element z. By the construction of our paths, there are cover relations

$$\begin{aligned} x_1 < v_1 \quad &\text{and} \quad x_2 < v_2, \\ x'_1 < v'_1 \quad &\text{and} \quad x'_2 < v'_2 \end{aligned}$$

in P such that

$$\begin{aligned} \mathcal{E}^{x_1} = \mathcal{E}^{x_2} &= \{C_{v_1v_2}\}, \\ \mathcal{E}^{x'_1} = \mathcal{E}^{x'_2} &= \{C_{v'_1v'_2}\}. \end{aligned}$$

First, suppose that z is an endpoint of one path and an internal vertex of the other path. Without loss of generality we assume that $z = v_1$ (so it is an endpoint of $Q_{v_1v_2}$). By the definition of $Q_{v'_1v'_2}$ we have $z \leqslant x'_1$ in P, or $z \leqslant x'_2$ in P, or $z \in C_{v'_1v'_2}$. In the first case it follows that $\mathcal{E} = \mathcal{E}^{v_1} \subseteq \mathcal{E}^{x'_1} = \{C_{v'_1v'_2}\}$, which is a clear contradiction. The second case is similar. And the third one contradicts the fact that V is disjoint from $\mathcal{E}$ (by (d.2)).

So suppose that $z \in Q_{v_1v_2} \cap Q_{v'_1v'_2}$ is an internal vertex of both paths. The sets $C_{v_1v_2}$ and $C_{v'_1v'_2}$ cannot both contain z as they are disjoint (by (b)). Hence we may assume $z \notin C_{v_1v_2}$. By the definition of $Q_{v_1v_2}$ we then must have $z \leqslant x_1$ or $z \leqslant x_2$ in P. Say $z \leqslant x_1$ holds in P. Now observe that $z \in Q_{v'_1v'_2}$ implies that there is $c' \in C_{v'_1v'_2}$ with $c' \leqslant z$ in P. Hence $c' \leqslant x_1$ in P. However, from this we deduce $C_{v'_1v'_2} \in \mathcal{E}^{x_1} = \{C_{v_1,v_2}\}$, which is a contradiction. We conclude that both paths $Q_{v_1v_2}$ and $Q_{v'_1v'_2}$ are indeed internally disjoint.

As a consequence we established the existence of a subdivision of K_n in the cover graph of P. This completes the proof of Theorem 5.0.1.

5.2 Second Case: $(\mathbf{k}+\mathbf{k})$-free Posets

This section is devoted to a proof of Theorem 5.0.2: $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude K_n as a topological minor have dimension bounded in terms of k and n.

To prove this theorem, we are going to use the framework from the previous section. In particular, we run through two phases which will give us appropriate sets to construct a subdivision of K_n. Compared to the previous section, the structure with its invariants in Phase 1 is slightly different. For instance, the two collections will contain only singletons so that the structure becomes simpler as we can use two sets of elements instead. More importantly, we have to adjust the invariants (b.2) and (c.2) of the previous section to the unbounded height setting. This can be achieved by using the following notions of *close upsets* and *close downsets*. For an element $p \in P$ and an integer $k \geqslant 2$, we denote by $\mathrm{U}^{<k}(p)$ the set of elements $x \in \mathrm{U}(p)$ for which *every* covering chain from p to x is of size less than k. (Let us emphasize here that it is crucial for our proof to define it for every covering chain.) Similarly, we define $\mathrm{D}^{<k}(p)$ to be the set of elements $x \in \mathrm{D}(p)$ for which every covering chain from x to p is of size less than k. We will use these two new notions in the invariants and for the two new variants of questions (Q1) and (Q2). (Note that we could have used close upsets and downsets already in the previous section as $\mathrm{U}(p) = \mathrm{U}^{<h+1}(p)$ and $\mathrm{D}(p) = \mathrm{D}^{<h+1}(p)$ holds in posets of height at most h.)

The third difference to the previous Phase 1 is the way we will get new elements (after a "yes"-answer for the third question). It will be simpler and avoids the use of unfoldings. Phase 2 and the construction of the subdivision will then go along the same lines as in the first proof.

We continue with the detailed proof now. We omit the trivial cases and assume $n \geqslant 3$. Let P be a $(\mathbf{k}+\mathbf{k})$-free poset that contains a standard example S_m with

$$m > n^L, \text{ where } L = 2\binom{M+k-1}{k-1} - 1 \text{ and } M = \binom{n}{2}^{(k-1)^n}.$$

During Phase 1 we maintain an additional structure (A, B, C, D) while running a loop. After the i-th iteration step we have the following invariants:

(a) there is a standard example of size n^{L-i} in P with A and B being the sets of its minimal and maximal elements, respectively,

(b) C is an antichain in P with $|C| \leqslant M$ and

  (b.1) $A \cap \mathrm{D}(c) = \emptyset$ and $B \subseteq \mathrm{U}^{<k}(c)$, for every $c \in C$,

(c) D is an antichain in P with $|D| \leqslant M$ and

  (c.1) $B \cap \mathrm{U}(d) = \emptyset$ and $A \subseteq \mathrm{D}^{<k}(d)$, for every $d \in D$,

Again, we have a measure of quality of the maintained structure. For $c \in C$ let value(c) be the maximum size of a covering chain in cover(P) from c to any $b \in B$. Note that by invariant (b.1) we have value$(c) < k$ for every $c \in C$. Then we define value(C) in the same way as in the previous section, implying that value(C) is a non-decreasing sequence of length M with entries from the set $\{1, \ldots, k\}$. Dually, we define value(D). To compare the values we use the same lexicographic order as before. Note that the number of possible values for C and D, respectively, is bounded by $\binom{M+k-1}{k-1}$.

Phase 1: Updating the Data Structure

We set up the initial structure as follows. We fix a copy of a standard example of size m in P. Set A and B to be the set of its minimal elements and maximal elements, respectively. Set C and D to be the empty set. Clearly, conditions (a)-(c) hold for $i = 0$.

Now we run a loop to improve the quality of the data structure. In each loop iteration we ask up to three questions about the current structure (A, B, C, D). If we get only negative answers, then the loop terminates and Phase 1 is done. If we get a positive answer to one of the questions, then we finish the iteration by updating the structure to (A', B', C', D') that is satisfying conditions (a)-(c) and additionally

$$\text{value}(C') < \text{value}(C) \text{ and } \text{value}(D') \leqslant \text{value}(D),$$

$$\text{or}$$

$$\text{value}(C') \leqslant \text{value}(C) \text{ and } \text{value}(D') < \text{value}(D).$$

Since the number of possible values for the set C (and D, resp.) is bounded by $\binom{M+k-1}{k-1}$, there will be at most $L = 2\binom{M+k-1}{k-1} - 1$ iterations in total. (If we reach the L-th iteration, then the structure is not improvable anymore, implying that we necessarily get only negative answers to the three questions.)

Now we are going to describe the i-th $(1 \leqslant i \leqslant L)$ iteration in detail. Let (A, B, C, D) be the current structure. Hence it satisfies conditions (a)-(c) after the $(i-1)$-th step. The iteration starts with the evaluation of the following question:

$$\text{Is there an element } p \in P \text{ such that} \tag{Q1}$$

(i) $\left|B \cap \mathrm{U}^{<k}(p)\right| > n^{L-i}$, and
(ii) there is $c \in C$ such that $c < p$ in P?

First, suppose that the answer is "yes" and fix such an element $p \in P$. In this case we finish the i-th iteration by updating the structure to (A', B', C', D'), where

$$\begin{aligned}
B' &= B \cap \mathrm{U}^{<k}(p),\\
A' &= \{a \in A \mid a \text{ is incomparable to some } b \in B'\},\\
C' &= C \cup \{p\} - \{c \in C \mid c < p \text{ in } P\},\\
D' &= D.
\end{aligned}$$

Note that A' and B' induce a standard example of size larger than n^{L-i} and hence invariant (a) is satisfied. We skip the proof for the fact that (A', B', C', D') satisfies the invariants (b)-(c) and moreover, that $\text{value}(C') < \text{value}(C)$ and $\text{value}(D') \leqslant \text{value}(D)$. It follows along the same lines as in the argument for the analogue claim in the previous section.

If the answer for question (Q1) is "no", then the procedure continues with a dual question:

$$\text{Is there an element } p \in P \text{ such that} \tag{Q2}$$

(i) $\left|A \cap \mathrm{D}^{<k}(p)\right| > n^{L-i}$, and
(ii) there is $d \in D$ such that $p < d$ in P?

If the answer for (Q2) is "yes", then we finish the i-th iteration by updating the structure to (A', B', C', D'), where

$$\begin{aligned}
A' &= A \cap \mathrm{D}^{<k}(p),\\
B' &= \{b \in B \mid b \text{ is incomparable to some } a \in A'\},\\
C' &= C,\\
D' &= D \cup \{p\} - \{d \in D \mid p < d \text{ in } P\}.
\end{aligned}$$

The proof that this new structure satisfies conditions (a)-(c) and that it improves the quality is dual to the one of question (Q1).

If the answers for questions (Q1) and (Q2) are both "no", then the procedure continues with the third question:

$$\text{Are } |C| < M \text{ and } |D| < M? \tag{Q3}$$

We first deal with the "yes" answer. In this case, we are going to show how to find a new element to extend C or D. This part of the procedure is simpler than its analogue in the previous section. We do not need to unfold the poset, but instead make use of the structure of standard examples and the fact that P is $(\mathbf{k}+\mathbf{k})$-free.

Fix an incomparable pair $(a_0, b_0) \in \operatorname{Inc}(A, B)$. Consider the following two sets:

$$A_{\text{cl}} := A \cap \mathrm{D}^{<k}(b_0) \quad \text{and} \quad B_{\text{cl}} := B \cap \mathrm{U}^{<k}(a_0).$$

The key observation here is that for each incomparable pair $(a, b) \in \operatorname{Inc}(A, B) - \{(a_0, b_0)\}$, we have $a \in A_{\text{cl}}$ or $b \in B_{\text{cl}}$. Indeed, otherwise there are two covering chains in $\operatorname{cover}(P)$ of size at least k, one from a_0 to b and one from a to b_0. And since P is $(\mathbf{k}+\mathbf{k})$-free, we would deduce that $a_0 < b_0$ or $a < b$ in P, which is not true. As a consequence, we get

$$|A_{\text{cl}}| \geqslant (|A| - 1)/2 \ \text{ or } |B_{\text{cl}}| \geqslant (|B| - 1)/2. \tag{5.3}$$

Suppose first $|B_{\text{cl}}| \geqslant (|B| - 1)/2$. Then we finish the i-th iteration by updating the structure to (A', B', C', D'), where

$$\begin{aligned}
B' &= B_{\text{cl}},\\
A' &= \{a \in A \mid a \text{ is incomparable to some } b \in B'\},\\
C' &= C \cup \{a_0\},\\
D' &= D.
\end{aligned}$$

Claim 5.2.1. *The structure (A', B', C', D') keeps the invariants (a)-(c). Moreover,* $\operatorname{value}(C') < \operatorname{value}(C)$ *and* $\operatorname{value}(D') \leqslant \operatorname{value}(D)$.

Proof. Clearly, A' and B' form a standard example that is contained in the original standard example and its size is

$$|A'| = |B'| = |B_{\text{cl}}| \geqslant (|B| - 1)/2 \geqslant (n^{L-(i-1)} - 1)/2 \geqslant n^{L-i},$$

as $n \geqslant 3$. Thus, (a) holds.

Now we aim to prove that $C' = C \cup \{a_0\}$ is an antichain in P. Since C is an antichain in P (by (b)), all we need to prove is that a_0 is incomparable to all elements in C. To see this, note that if $a_0 < c$ in P for some $c \in C$ then this violates invariant (b.1) for (A, B, C, D). And if $c < a_0$ in P for some $c \in C$, then a_0 would be a witness for a "yes"-answer for question (Q1), which cannot be as we reached question (Q3) in this iteration. Hence C' is an antichain in P as required.

We need to check invariant (b.1) only for the new element a_0 in the collection as the invariant was satisfied before and as $A' \subseteq A$, $B' \subseteq B$. It follows immediately as a_0 is incomparable to all elements in A' and by the definition of B'. Invariant (c) holds as $D' = D$ and $A' \subseteq A$, $B' \subseteq B$.

Finally, $\operatorname{value}(C') < \operatorname{value}(C)$ and $\operatorname{value}(D') \leqslant \operatorname{value}(D)$ simply as $C \subsetneq C'$ and $D = D'$, respectively. □

If the other case of (5.3) holds, that is, we have$|A_{\mathrm{cl}}| \geqslant (|A|-1)/2$, then we finish the i-th iteration by updating the structure to (A', B', C', D'), where

$$
\begin{aligned}
A' &= A_{\mathrm{cl}}, \\
B' &= \{b \in B \mid b \text{ is incomparable to some } a \in A'\}, \\
C' &= C, \\
D' &= D \cup \{b_0\}.
\end{aligned}
$$

The proof that (A', B', C', D') keeps the invariants (a)-(c) and that value$(C') \leqslant$ value(C) and value$(D') <$ value(D) goes dually to the previous case.

Let (A, B, C, D) be the final data structure of Phase 1 and suppose that in the last iteration we received a negative answer to question (Q3) because $|C| = M$. Next, to specify the n principal vertices of the K_n subdivision in the cover graph of P, we apply Phase 2 of the previous section to the structure (A, B, C, D). To make it work, we only need to replace the sets C and D with the collections of their singletons. Moreover, we can also keep the invariants (d.1)-(d.3) for the maintained pair $(V, \mathcal{E})$, except that we have to replace the height parameter 'h' by '$k-1$' in invariant (d.1).

The main reason why we can apply Phase 2 so easily here is the following. In each iteration of the phase, a new principal vertex v is found by first starting in some element $b \in B$ and then going down along cover relations of the subposet conv$(C \cup B)$. Since conv$(C \cup B)$ is poset of height at most $k-1$ by (b.1), the same arguments as before apply (with respect to '$k-1$' instead of 'h').

Suppose we finish Phase 2 with the pair $(V, \mathcal{E})$ satisfying invariants (d.1)-(d.3) after n iterations. We may proceed now exactly in the same way as in the previous section and find a subdivision of K_n in the cover graph of P. This finishes the proof for Theorem 5.0.2.

Let us point out here that the subdivision of the K_n that we found in the previous proof is in a sense special. This is because the paths connecting principal vertices of the topological minor have length at most $2k-4$. Indeed, following the last construction steps of Section 5.1 we see that those paths consists of two covering chains of length at most $k-2$, which meet at the latest in a singleton of the collection C.

Therefore, we found a K_n topological minor of *bounded depth* in the cover graph of P in the proof of Theorem 5.0.2. This observation will allow us in Chapter 6 to extend Theorem 5.0.2 to broader graph classes, that is, to so-called *nowhere dense* graph classes.

Corollary 5.2.1. *For each $k \geqslant 1$ and $n \geqslant 1$ there is a constant $d(k,n)$ such that if P is a $(\mathbf{k}+\mathbf{k})$-free poset containing a standard example of size at least $d(k,n)$ as an induced subposet, then the cover graph of P contains a K_n-subdivision in which every path connecting principal vertices has length at most $2k-4$.*

Chapter 6

Sparsity and Dimension

In previous chapters we have seen that the dimension of posets can be bounded by a function of their height whenever their cover graphs have bounded degree, are planar, have bounded tree-width, or more generally exclude a fixed graph as a (topological) minor. Graphs from the mentioned classes have in common that they are sparse, in the sense that they contain only linear many edges. One might believe at this point that simply having a sparse cover graph is enough to guarantee a bounding function in terms of the height on the dimension. It is the main goal of this chapter to adress that question and to determine the exact type of combinatorial sparsity needed to ensure the existence of such a bounding function. Before we introduce this sparsity type let us recall some preliminary observations from the introduction.

Requiring cover graphs to contain only linear many edges is not enough. Indeed, adding a sufficient number of independent elements to large standard examples we easily obtain large-dimensional height-2 posets with fewer comparabilities than elements. This simple idea does not apply though when we require that also the subgraphs of a cover graph have linear many edges. (Note that this property also holds for graphs from the classes mentioned above). However, this condition is not enough either: The inclusion order on the 1-element and 2-element subsets of $\{1, \ldots, n\}$ yields a height-2 poset with dimension at least $\log \log n$, as already shown by Dushnik and Miller [17] by repeated applications of the Erdős-Szekeres theorem on monotone sequences. Yet, the cover graphs of these posets are complete graphs where each edge is subdivided once, which are 2-degenerate.

We deduce that 1-subdivisions of complete graphs are in a sense not sparse enough for our purposes. So what properties make those graphs dense? One way to answer this question is to consider also the edge densities of their minors. Clearly, complete graphs are minors of those graphs and have unbounded average degree. Moreover, already local contractions around the principal vertices yield the clique minors here.

It turned out that the idea of looking at densities of bounded depth minors, that is, minors that can be obtained by contracting disjoint subgraphs of bounded radius, leads to a fruitful theory of uniform sparsity in graphs. This theory is based on the work of Nešetřil and Ossona de Mendez [55]. Most notably, they introduced the notions of nowhere dense classes and classes with bounded expansion. We discuss these sparsity models in detail in Section 6.1, but let us give some first intuition here. In nowhere dense classes we have the property that bounded depth minors exclude at least some graph. The requirement for classes with bounded expansion is stronger; bounded depth minors have only linear many edges in this case.

Over the last decade, several seemingly unrelated characterizations of these two models have been found. For example, nowhere dense classes can be defined in terms of uniform quasi-wideness [13], an important concept in model theory, in terms of the so-called splitter game [33], and in terms of generalized coloring numbers [78]. As a consequence, many strong algorithmic applications have been discovered (e.g. see [33, 18]).

We add yet another reason for the robustness of this theory by establishing a sharp connection to the dimension of posets. We show that the notion of bounded expansion gives the exact sparsity model to guarantee bounding functions as discussed at the beginning of this chapter. More precisely, we prove the following theorem.

Theorem 6.0.1 ([42]). *For every class of graphs $\mathcal{C}$ with bounded expansion, posets of bounded height whose cover graphs belong to $\mathcal{C}$ have bounded dimension.*

We complement this theorem and show that it cannot be extended to nowhere dense classes with the following result.

Theorem 6.0.2 ([42]). *There exists a class of graphs $\mathcal{C}$ with locally bounded tree-width such that posets of height* 2 *with cover graphs in $\mathcal{C}$ have unbounded dimension.*

Both theorems are joint work with Gwenaël Joret and Piotr Micek. They were first presented at the Symposium on Discrete Algorithms 2016 (SODA 216); see [42] for the printed version published in the proceedings of the conference. For the proof of Theorem 6.0.1 in [42] we use two main tools. First, we once more work with unfolding posets in an iterative manner, and second, we use a characterization of bounded expansion classes in terms of *p-centered colorings*.

Although we believe that this is an insightful proof, we decided to not include it here in this chapter. Instead, we present a shorter and even more elegant proof that is based on weak coloring numbers. It avoids unfoldings and moreover leads to improved bounds on the dimension in several important cases. (Yet another proof of Theorem 6.0.1 by Joret, Micek, and myself can be found in [41].)

Chapter Overview. In Section 6.1 we introduce graph classes with bounded expansion and nowhere dense classes. We also consider posets whose cover graphs belong to some fixed nowhere dense class of graphs. More precisely, we first construct a family of posets that will establish a proof of Theorem 6.0.2, and second, we discuss a challenging conjecture in this context that would lead to a characterization of nowhere dense classes in terms of dimension.

In Section 6.2 we introduce weak coloring numbers and prove an upper bound on the dimension of posets in terms of those numbers. Combined with a characterization of bounded expansion classes due to Zhu this will yield a proof of Theorem 6.0.1. As a byproduct we obtain a single exponential bound in height on the dimension of posets whose cover graphs exclude a fixed graph as a minor. This is an improvement upon double exponential bounds that were known before (established in [38] for the bounded tree-width case and in [51] for the excluded minor case). Surprisingly, our derived bound is essentially best possible as witnessed by the lower bound constructions of Chapter 4. We conclude Section 6.2 by studying the dimension of height-2 posets with nowhere dense cover graphs. Again by using weak coloring number, we obtain much better bounds in this setting. For instance, in the case that cover graphs contain no K_t-minor we derive a polynomial bound in t for the dimension of height-2 posets, while there exists an exponential lower bound in t for height-4 posets (this follows from Theorem 4.4.5).

We finish the chapter with Section 6.3, where we apply our developed ideas to the concept of boxicity. We show a close connection between the boxicity of a graph and its weak 2-coloring number. This will imply an $\mathcal{O}(t^2 \log t)$ bound on the boxicity of K_t-minor-free graphs, which improves upon an $\mathcal{O}(t^4 \log^2 t)$ bound of Esperet and Joret [22].

6.1 Classes with Bounded Expansion and Nowhere Dense Classes

Let G be a graph. The *distance* between two vertices in G is the number of edges in a shortest path connecting them. If the two vertices lie in distinct components of G, then their distance is infinite. The set of all vertices at distance at most r from a vertex v in G is denoted by $N_G^r(v)$, and the subscript is omitted if G is clear from the context. The *radius* of a connected graph G is the least integer $r \geqslant 0$ for which there is a vertex $v \in V(G)$ such that $N_G^r(v) = V(G)$.

In Chapter 1 we introduced the concept of graph minors. We build on those and define bounded depth versions of them now. A graph H is a *depth-r minor* (also known as an *r-shallow minor*) of a graph G if there are disjoint connected subgraphs of radius at most r in G, such that H can be obtained from G by first contracting each of these subgraphs to a single vertex and then by possibly removing vertices and edges of the resulting graph. Figure 6.1 illustrates this definition on an example.

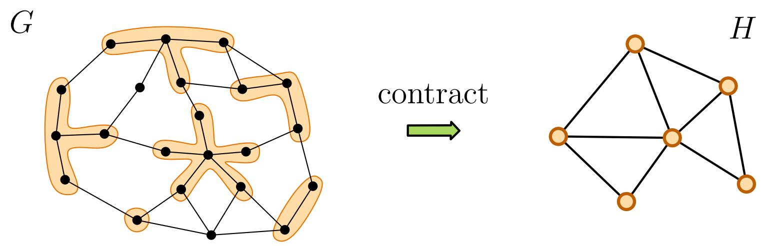

Figure 6.1: Graph H is a 1-shallow minor of G, since it can be obtained by contracting disjoint subgraphs of radius at most 1 in G.

The *greatest reduced average density* of *rank* r of a graph G, denoted by $\nabla_r(G)$, is defined as

$$\nabla_r(G) = \max\left\{\frac{|E(H)|}{|V(H)|} \mid H \text{ is a depth-}r\text{ minor of } G\right\}.$$

A class of graphs $\mathcal{C}$ has *bounded expansion* if there exists a function $f : \mathbb{R} \to \mathbb{R}$ such that $\nabla_r(G) \leqslant f(r)$ for every integer $r \geqslant 0$ and graph $G \in \mathcal{C}$. A class of graphs $\mathcal{C}$ is *nowhere dense* if for each integer $r \geqslant 0$ there exists a graph which is *not* a depth-r minor of any graph $G \in \mathcal{C}$. If a class is not nowhere dense, then we call it *somewhere dense*.

We discuss some examples of graph classes satisfying these types of sparsity now. First, we observe that every graph class with bounded expansion is also nowhere dense. This follows directly from the definition: r-shallow minors of graphs from a bounded expansion class have only linear many edges and hence large cliques are not among those. Next, we argue that every graph class $\mathcal{C}$ that excludes a fixed graph H as a minor has bounded expansion. It is a well-known fact that graphs with no H-minor have bounded average degree. This particularly means that there is a constant c such that $\nabla_0(G) \leqslant c$ for every $G \in \mathcal{C}$. Since every r-shallow minor of a graph in $\mathcal{C}$ also needs to exclude H as a minor (otherwise H is a minor of some graph in $\mathcal{C}$), we deduce that the bounding function for the greatest reduced average densities can be taken to be the constant function $f(r) \equiv c$ for the class $\mathcal{C}$. Graph classes that exclude a fixed graph as a topological minor also have bounded expansion, though in that case the bounding function might not be constant anymore, see [55].

Let us now give an example of a bounded expansion class that does not exclude a fixed graph as a topological minor. Consider the graphs for which there exists an embedding in the plane such that each edge is crossed at most k times. Those graphs are called k-planar graphs. Every graph G is a topological minor of some 1-planar graph: fix an arbitrary embedding of G in the plane and subdivide edges around every crossing so that the resulting graph is 1-planar. On the other hand, k-planar graphs are sparse; Pach and Tóth [60] proved that their average degree is at most $8\sqrt{k}$. Since bounded depth minors of k-planar graphs are k'-planar for some integer k', this shows that k-planar graphs have bounded expansion [57]. Further bounded expansion classes that do not exclude a fixed graph as a topological minor are given by the graphs with bounded book-thickness or by the graphs with bounded queue-number (see [57] for details). These examples coming from graph drawing show that the sparsity notion of bounded expansion goes beyong the concept of excluding a topological minor.

Nowhere dense classes of graphs form an even more general type of graph classes than those that have bounded expansion. For example, graphs with locally bounded tree-width are nowhere dense but not necessarily have bounded expansion. A class of graphs $\mathcal{C}$ has *locally bounded tree-width* if there exists a function $f : \mathbb{R} \to \mathbb{R}$ such that $\mathrm{tw}(G[N^r(v)]) \leqslant f(r)$ for every integer $r \geqslant 0$, graph $G \in \mathcal{C}$ and vertex $v \in V(G)$. An explicit family of graphs with locally bounded tree-width is given by the graphs G that satisfy $\Delta(G) \leqslant \mathrm{girth}(G)$, where $\Delta(G)$ and $\mathrm{girth}(G)$ denote the maximum degree and girth of G, respectively. (This is not difficult to check; or see [55] for a proof.) It is moreover an example of a nowhere dense class that does not have bounded expansion.

Posets with Cover Graphs in a Nowhere Dense Class. As mentioned in the introduction of this chapter, the statement of Theorem 6.0.1 cannot be pushed further towards nowhere dense classes. In fact, it already fails for classes with locally bounded tree-width, as we show now.

Our construction is based on the class of graphs G with $\Delta(G) \leqslant \mathrm{girth}(G)$. As mentioned above, this is a useful example of a class with locally bounded tree-width that does not have bounded expansion. (Indeed, this class is invoked several times in the textbook [55].) We will use in particular that the chromatic number of these graphs is unbounded. This is a well-known fact that can be shown in multiple ways; we can note for instance that the chromatic number of the n-vertex d-regular non-bipartite Ramanujan graphs with girth $\Omega_d(\log n)$ built by Lubotzky, Phillips, and Sarnak have chromatic number $\Omega(\sqrt{d})$ (see [50]).

For the proof of Theorem 6.0.2 we use a classic construction of height-2 posets that we associate with graphs. Given a graph G, the *adjacency poset* P_G of G is the poset with ground set $\{a_v \mid v \in V(G)\} \cup \{b_v \mid v \in V(G)\}$ such that, for every two distinct vertices $u, v \in V(G)$, we have $a_u \leqslant b_v$ in P_G if and only if $uv \in E(G)$. It is well known that $\dim(P_G) \geqslant \chi(G)$. This can be seen as follows. Fix a realizer $\{L_1, \ldots, L_d\}$ of P_G. For every vertex $v \in V(G)$, fix a number $\phi(v) = i$ such that $b_v < a_v$ in L_i. We claim that ϕ is a proper coloring of G. Consider any two adjacent vertices u and v in G and, say, $\phi(v) = i$. Then $a_u \leqslant b_v < a_v \leqslant b_u$ in L_i, which witnesses that $\phi(u) \neq i$. Therefore, $\dim(P_G) \geqslant \chi(G)$.

Proof of Theorem 6.0.2. Let $\mathcal{C}$ denote the class of graphs G satisfying $\Delta(G) \leqslant \text{girth}(G)$. The key observation about the class $\mathcal{C}$ in this context is that if $G \in \mathcal{C}$, then $\text{cover}(P_G) \in \mathcal{C}$. This can be seen as follows: First, clearly $\Delta(G) = \Delta(\text{cover}(P_G))$, so it is enough to show $\text{girth}(G) \leqslant \text{girth}(\text{cover}(P_G))$. To show the latter, we remark that if C is a cycle of $\text{cover}(P_G)$, then C naturally corresponds to a closed walk W in G of the same length. Moreover, every three consecutive vertices in that walk W are pairwise distinct, as follows from the adjacency poset construction. Hence, W contains a cycle, which is of length at most that of C. Therefore, $\text{girth}(G) \leqslant \text{girth}(\text{cover}(P_G))$, as claimed.

To summarize, graphs in $\mathcal{C}$ have unbounded chromatic number, implying that adjacency posets of these graphs have unbounded dimension. Yet, the cover graphs of these adjacency posets all belong to $\mathcal{C}$, a class with locally bounded tree-width. This concludes the proof. □

By Theorem 6.0.2 we know that the dimension of bounded height posets with nowhere dense cover graphs can grow with the poset size. A natural problem arising here is to determine the maximal growth rate. Is it possible that the dimension is polynomial in the poset size here? Analogously to the behaviour of many graph parameters on nowhere dense graphs (such as p-centered coloring number, weak r-coloring number, r-neighborhood complexity, etc.) we believe that the dimension of these posets is $\mathcal{O}(n^\epsilon)$, for each $\epsilon > 0$.

Conjecture 6.1.1. *Let $\mathcal{C}$ be a nowhere dense class of graphs. For every integer $h \geqslant 1$ and every $\epsilon > 0$, posets of height at most h whose cover graphs belong to $\mathcal{C}$ have dimension $\mathcal{O}(n^\epsilon)$.*

Let us argue now that the conjecture, if true, would yield a characterization of monotone nowhere dense classes in terms of order dimension. (Recall that a monotone class is defined to be closed under taking subgraphs.) Nešetřil, Ossona de Mendez, and Wood [56] proved that bounded expansion classes and nowhere dense classes can be equivalently defined by edge densities of r-shallow topological minors. A graph H is an *r-shallow topological minor* of a graph G if G contains a subgraph that is isomorphic to an $\leqslant r$-subdivision of H. Combining the characterization with a result of Dvořák [19], it follows that a class of graphs $\mathcal{C}$ is nowhere dense if and only if for each $r \geqslant 0$ there exists a graph that is not an r-shallow topological minor of any graph in $\mathcal{C}$ (see also [54]).

Therefore, Conjecture 6.1.1 is false for *every* monotone somewhere dense class $\mathcal{C}$: There exists an integer $r \geqslant 0$ (depending on $\mathcal{C}$) such that every graph is an r-shallow topological minor of some graph in $\mathcal{C}$. In particular, this holds for the cover graph of the standard example S_m. This means that for every m, the class $\mathcal{C}$ contains a graph G_m that is an $\leqslant r$-subdivision of $\text{cover}(S_m)$. Notice that G_m has at most

$rm^2 + 2m$ vertices. Now it is easy to see that G_m is also the cover graph of a poset P_m of height at most $r + 2$ that contains S_m as an induced subposet. Since P_m has at most $n := rm^2 + 2m$ elements we deduce that its dimension is $\Omega(\sqrt{n})$.

On the other hand, a family of bounded height posets with nowhere dense cover graphs and dimension $\Omega(\sqrt{n})$ clearly cannot exist if the above conjecture is true. Thus, we would indeed obtain a new characterization of monotone nowhere dense graphs with a positive resolution of Conjecture 6.1.1.

Let us give further support of this conjecture now. If posets contain only standard examples of bounded size as induced subposets, then this can be seen as an indication that their dimension is small. Of course, this does not hold in general as interval orders do not contain standard examples of size 2 while their dimension can be arbitrary large. However, with respect to the restricted setting of sparse cover graphs, there is a general believe that the maximal difference between the dimension and the largest size of an induced standard example is way smaller than in the general setting.

We show now that posets of bounded height with cover graphs from a fixed nowhere dense class can contain only standard examples of bounded size. In fact, we prove this statement for $\mathbf{k}+\mathbf{k}$-free posets. (Recall that this is more general as posets of height less than h do not contain $\mathbf{h}+\mathbf{h}$ as an induced subposet.)

Theorem 6.1.1. *Let $k \geqslant 1$ and let $\mathcal{C}$ be a nowhere dense class of graphs. Then there exists $d := d(k, \mathcal{C})$ such that $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs belong to $\mathcal{C}$ contain only standard examples of size at most d as induced subposets.*

Proof. Suppose that $\mathcal{P}$ is a family of $(\mathbf{k}+\mathbf{k})$-free posets that contain standard examples of unbounded size as induced subposets. All we need to show is that the cover graphs of posets in $\mathcal{P}$ form a somewhere dense class of graphs.

By Corollary 5.2.1 we know that for every $t \geqslant 1$ there is an integer $c(k, t)$, such that if P is a $(\mathbf{k}+\mathbf{k})$-free poset with a standard example of size at least $c(k, t)$ as an induced subposet, then the cover graph of P contains K_t as a depth-$(2k-4)$ minor. This implies that the cover graphs of posets in $\mathcal{P}$ contain every clique (and hence every graph) as a depth-$(2k-4)$ minor, and hence they form a somewhere dense class of graphs. □

We further discuss the question of standard examples versus dimension in the conclusions of this thesis. In Section 6.2 we give more support for Conjecture 6.1.1 by proving that it holds for posets of height at most 2.

6.2 Weak Coloring Numbers and Dimension

Weak coloring numbers were introduced by Kierstead and Yang [45] as a generalization of the degeneracy of a graph (which is also known as the *coloring number*). Originally, they were applied in the context of coloring games and marking games on graphs. A quite recent application of them for the sparsity theory of graphs was found by Zhu [78]. He proves characterizations of bounded expansion classes and nowhere dense classes in terms of weak coloring numbers, which form the basis in many proofs of problems dealing with sparse graphs. Further important applications were found in the algorithmic theory of sparse graphs. For example, weak coloring numbers play a key role in Dvořák constant-factor approximation algorithm for the dominating set problem on graphs with bounded expansion [18]. Moreover, they were used to find sparse neighborhood covers of nowhere dense graphs, which was a major step to prove the existence of almost linear time model-checking algorithms for first-order formulas; see Grohe et al. [33].

In this section we also want to use the powerful concept of weak coloring numbers to give strong bounds on the dimension of posets with sparse cover graphs. Before we state the main theorem of this section, let us provide some necessary definitions.

Let G be a graph and let $\Pi(G)$ denote the set of linear orders on $V(G)$. Fix some linear order $\pi \in \Pi(G)$ for the moment. We write $x <_\pi y$ if x is smaller than y in π. For a path Q in G, we let $\ell(Q, \pi)$ denote the minimal vertex of Q with respect to π. We write $\ell(Q)$ for $\ell(Q, \pi)$ if it is clear from the context that we are working with π. For a vertex v in G and integer $r \geqslant 0$, we say that $u \in V(G)$ is *weakly r-reachable from v* with respect to the order π, if there exists a path Q of length at most r from v to u in G such that $\ell(Q, \pi) = u$. By $\mathrm{WReach}_r^\pi[v]$ we denote the set of weakly r-reachable vertices from v with respect to π (note that this set contains v for all $r \geqslant 0$). Then the weak r-coloring number $\mathrm{wcol}_r(G)$ is defined as

$$\mathrm{wcol}_r(G) = \min_{\pi \in \Pi(G)} \max_{v \in V(G)} |\mathrm{WReach}_r^\pi[v]|.$$

Zhu [78] established a close connection between weak coloring numbers and densities of bounded depth minors, which can be used to characterize graph classes with bounded expansion.

Theorem 6.2.1 (Zhu [78])**.** *A class of graphs $\mathcal{C}$ has bounded expansion if and only if there exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that $\mathrm{wcol}_r(G) \leqslant f(r)$ for all $G \in \mathcal{C}$ and every $r \geqslant 0$.*

Zhu's observations also imply a characterization of nowhere dense classes of graphs: A class $\mathcal{C}$ is nowhere dense if and only if for each $r \geqslant 0$ and every $\epsilon > 0$, we have that $\mathrm{wcol}_r(G) = \mathcal{O}(n^\epsilon)$ for every n-vertex graph $G \in \mathcal{C}$.

We are ready to state and prove the main theorem of this section, which together with Zhu's characterization of bounded expansion classes implies Theorem 6.0.1.

Theorem 6.2.2. *For every poset P of height at most h whose cover graph G satisfies $\mathrm{wcol}_{3h-3}(G) \leqslant c$, we have*

$$\dim(P) \leqslant 4^c.$$

Proof. Let $P = (X, \leqslant)$ be a poset as in the statement and let G be its cover graph. By assumption, there is a linear order π on the elements of X such that

$$\mathrm{WReach}_{3h-3}^\pi[x] \leqslant c, \tag{6.1}$$

for each $x \in X$.

First, we greedily color the elements of X by going through π. We start with the smallest element in π, and when element x is about to be colored, meaning that all elements smaller than x in π have been colored yet, then we give x the smallest color $\Phi(x)$ from $\{1, \ldots, c\}$ that does not appear on elements of $\mathrm{WReach}_{3h-3}^\pi[x] \setminus \{x\}$. (Note that by (6.1) at least one color is available.) The resulting greedy coloring Φ has the following property.

Claim. *Let $x \in X$. If $y, z \in \mathrm{WReach}^{\pi}_{h-1}[x]$ are distinct, then $\Phi(y) \neq \Phi(z)$.*

Proof. We may assume that $y <_\pi z$. Since z is weakly $(h-1)$-reachable from x, there is a path $Q(x, z)$ from x to z of length at most $h-1$ in G that does not go below z in π. Moreover, there is a path $Q(x, y)$ from x to y of length at most $h-1$ in G witnessing that y is weakly $(h-1)$-reachable from x. Clearly, the union of the two paths $Q(x, z)$ and $Q(x, y)$ is connected in G and contains a path Q from z to y of length at most $2h-2$. As y is the minimal element of the union with respect to π, we deduce that Q is a witness that y is weakly $(2h-2)$-reachable from z with respect to π. By our coloring rule this implies $\Phi(y) \neq \Phi(z)$. □

In what follows, we need only consider those elements in $\mathrm{WReach}^{\pi}_{h-1}[x]$ that are weakly reachable from x via covering chains. Therefore, we introduce the *weakly reachable upset* and the *weakly reachable downset* of x:

$$\mathrm{WU}[x] := \{y \in \mathrm{U}(x) \mid \exists \text{ a covering chain } Q \text{ from } x \text{ to } y \text{ such that } \ell(Q) = y\},$$
$$\mathrm{WD}[x] := \{y \in \mathrm{D}(x) \mid \exists \text{ a covering chain } Q \text{ from } y \text{ to } x \text{ such that } \ell(Q) = y\}.$$

Note that $\mathrm{WU}[x] \subseteq \mathrm{WReach}^{\pi}_{h-1}[x]$ and by the previous claim this implies that the elements in $\mathrm{WU}[x]$ have pairwise different colors. Similarly, this holds for elements in $\mathrm{WD}[x] \subseteq \mathrm{WReach}^{\pi}_{h-1}[x]$. Therefore, given that $i \in \Phi(\mathrm{WU}[x])$, it is well-defined to denote by $\mathrm{wu}_i(x)$ the unique element in $\mathrm{WU}[x]$ with color i. We also let $\mathrm{wd}_i(x)$ denote the unique element in $\mathrm{WD}[x]$ with color i, given that $i \in \Phi(\mathrm{WD}[x])$.

Next, we assign a signature to each incomparable pair of P. As we will show, this yields a partition of $\mathrm{Inc}(P)$ into at most 4^c reversible sets. Given $(x, y) \in \mathrm{Inc}(P)$, we define a vector $\tau(x, y) = (\tau_i(x, y))_{i \in [c]}$, where

$$\tau_i(x, y) = \begin{cases} 1 & \text{if } i \in \Phi(\mathrm{WU}[x]) \cap \Phi(\mathrm{WD}[y]) \text{ and } \mathrm{wu}_i(x) <_\pi \mathrm{wd}_i(y), \\ 0 & \text{otherwise.} \end{cases}$$

So we have $\tau_i(x, y) = 1$ if both $\mathrm{WU}[x]$ and $\mathrm{WD}[y]$ contain an element with color i, and the corresponding unique elements in color i satisfy $\mathrm{wu}_i(x) <_\pi \mathrm{wd}_i(y)$. Now the signature $\sigma(x, y)$ of the pair (x, y) is a 3-tuple defined as follows:

$$\sigma(x, y) := (\Phi(\mathrm{WU}[x]), \Phi(\mathrm{WU}[x]) \cap \Phi(\mathrm{WD}[y]), \tau(x, y)).$$

By Σ we denote the set of all possible signatures $\sigma(x, y)$ for an incomparable pair (x, y) in P.

Claim. *Let $\sigma \in \Sigma$. Then the set of incomparable pairs (x, y) of P with $\sigma(x, y) = \sigma$ is reversible.*

Proof. Suppose for a contradiction that the statement is not true. Then there exists an alternating cycle $(x_1, y_1), \ldots, (x_k, y_k)$ such that $\sigma(x_j, y_j) = \sigma$ for all $j \in [k]$. If we write the fixed signature as $\sigma =: (A, B, \tau)$, then for each $j \in [k]$ it holds that $A = \Phi(\mathrm{WU}[x_j])$, $B = \Phi(\mathrm{WU}[x_j]) \cap \Phi(\mathrm{WD}[y_j])$, and $\tau = \tau(x_j, y_j)$.

For every $j \in [k]$, among the covering chains witnessing the comparability $x_j \leqslant y_{j+1}$ in P, let Q_j denote one with $\ell(Q_j)$ being minimal with respect to π. Let us write q_j for $\ell(Q_j)$ from now on. Clearly, for each $j \in [k]$ we have

$$x_j \leqslant q_j \leqslant y_{j+1} \text{ in } P \quad \text{and} \quad q_j \in \mathrm{WU}[x_j] \cap \mathrm{WD}[y_{j+1}]. \tag{6.2}$$

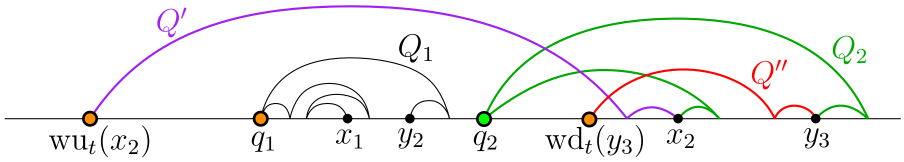

Figure 6.2: Elements with color t are drawn in orange. The union $Q' \cup Q'' \cup Q_2$ yields a path witnessing that $\mathrm{wu}_t(x_2)$ is weakly $(3h-3)$-reachable from $\mathrm{wd}_t(y_3)$.

From this we conclude $\Phi(q_j) \in \Phi(\mathrm{WU}[x_j]) = A = \Phi(\mathrm{WU}[x_{j+1}])$ and also $\Phi(q_j) \in \Phi(\mathrm{WD}[y_{j+1}])$. Hence $\Phi(q_j) \in B = \Phi(\mathrm{WU}[x_{j+1}]) \cap \Phi(\mathrm{WD}[y_{j+1}])$, for each $j \in [k]$. So if we denote by C the set of colors used on the q_j's, then we have $C \subseteq B$.

Without loss of generality we may assume that q_1 is minimal among the q_j's with respect to π. Let $t \in C$ be the color of q_1. Observe that by (6.2) this means

$$\mathrm{wu}_t(x_1) = q_1 = \mathrm{wd}_t(y_2).$$

We first show $\tau_t(x_2, y_2) = 0$ now. A symmetric argument will then give us $\tau_t(x_1, y_1) = 1$, which yields a clear contradiction to the fact that all pairs of the alternating cycle have the same signature σ.

To show $\tau_t(x_2, y_2) = 0$, let us first consider the element q_2. We have $q_1 \neq q_2$, as otherwise $x_2 \leqslant q_2 = q_1 \leqslant y_2$ in P by (6.2). Hence $q_1 <_\pi q_2$ by our assumption on q_1. If q_2 is also colored with t, then $\mathrm{wd}_t(y_2) = q_1 <_\pi q_2 = \mathrm{wu}_t(x_2)$ and we already obtain $\tau_t(x_2, y_2) = 0$.

So suppose that q_2 has color $t' \in C$ with $t' \neq t$. Assume for a contradiction that $\tau_t(x_2, y_2) = 1$, that is, we have

$$\mathrm{wu}_t(x_2) <_\pi \mathrm{wd}_t(y_2) = q_1 <_\pi q_2.$$

(See Figure 6.2 for such a situation and upcoming arguments.) From this it follows that $\mathrm{wd}_t(y_3) \neq \mathrm{wu}_t(x_2)$, as otherwise there is a covering chain witnessing $x_2 \leqslant y_3$ in P that is containing $\mathrm{wu}_t(x_2)$, which contradicts that Q_2 was chosen so that its element $\ell(Q_2) = q_2$ is as small as possible in π. (Recall at this point that in our alternating cycle we have $y_3 = y_1$ if $k = 2$.)

Next, we aim to show that $\mathrm{wd}_t(y_3)$ is weakly $(3h-3)$-reachable from $\mathrm{wu}_t(x_2)$ with respect to π, or the other way around. Since $\mathrm{wu}_t(x_2) \in \mathrm{WU}[x_2]$, there is a covering chain Q' of length at most $h-1$ from x_2 to $\mathrm{wu}_t(x_2)$ in G, such that $\ell(Q') = \mathrm{wu}_t(x_2)$. Similarly, there is a covering chain Q'' of length at most $h-1$ from $\mathrm{wd}_t(y_3)$ to y_3 in G, such that $\ell(Q'') = \mathrm{wd}_t(y_3)$. Observe that $Q_2 \cup Q' \cup Q''$, which we view as a union of paths in $\mathrm{cover}(P)$ now, contains a path Q of length at most $3h-3$ from $\mathrm{wu}_t(x_2)$ to $\mathrm{wd}_t(y_3)$. Moreover, as $\mathrm{wu}_t(x_2) <_\pi q_2 = \ell(Q_2)$ we obtain that either $\mathrm{wu}_t(x_2)$ or $\mathrm{wd}_t(y_3)$ is the leftmost element of Q. If $\ell(Q) = \mathrm{wu}_t(x_2)$, then Q is a witness for $\mathrm{wu}_t(x_2) \in \mathrm{WReach}^\pi_{3h-3}[\mathrm{wd}_t(y_3)]$. If, on the other hand, $\ell(Q) = \mathrm{wd}_t(y_3)$, then we get $\mathrm{wd}_t(y_3) \in \mathrm{WReach}^\pi_{3h-3}[\mathrm{wu}_t(x_2)]$. Since $\mathrm{wu}_t(x_2)$ and $\mathrm{wd}_t(y_3)$ have the same color, we obtain in both cases a contradiction to the properties of our initial greedy coloring.

This shows $\tau_t(x_2, y_2) = 0$. If we now start our argumentation with the element q_k instead of q_2, then a symmetric argument, where the roles of x's and y's are exchanged, shows $\tau_t(x_1, y_1) = 1$. This yields the aimed contradiction and completes the proof of the claim. □

It remains to bound the number of signatures in Σ. A standard counting approach yields $|\Sigma| \leqslant 2^c \cdot 2^c \cdot 2^c = 8^c$. However, we can do better. Let $\sigma = (A, B, \tau) \in \Sigma$ be a signature that is attained by some incomparable pair of P. Note that by definition we have $B \subseteq A \subseteq 2^{[c]}$. Moreover, for every color $i \in [c]$ we can only have $\tau_i = 1$ if $i \in B$. Therefore, for each color $i \in [c]$ exactly one of the following four options hold:

(1) $i \notin A$,
(2) $i \in A$ and $i \notin B$,
(3) $i \in B$ and $\tau_i = 1$,
(4) $i \in B$ and $\tau_i = 0$.

It follows that we can encode every possible signature of Σ by a vector of length c with entries from the set $\{1, 2, 3, 4\}$. Therefore, this shows $|\Sigma| \leqslant 4^c$ and concludes the proof of the theorem. □

Let us note that the exponential upper bound in the theorem cannot be improved to a polynomial one. Recall that in Theorem 4.4.1 we proved that there is a family of height-3 posets $(P_k)_{k\geqslant 1}$ such that $\dim(P_k) \geqslant 2^k$ and the cover graph of P_k has tree-depth at most $2k$. Since for each graph G and $r \geqslant 1$ it holds that $\mathrm{wcol}_r(G) \leqslant \mathrm{td}(G)$, the family $(P_k)_{k\geqslant 1}$ also yields an exponential lower bound on the dimension in terms of weak coloring numbers.

Applications of Theorem 6.2.2. We already mentioned that Theorem 6.2.2 combined with Zhu's characterization (Theorem 6.2.1) implies Theorem 6.0.1, the main result of this chapter. Hence we obtain a surprisingly short proof of this general result, especially when comparing it with other existing proofs showing weaker statements. It might be objected that Zhu's result serves as some kind of "black box" here. However, his proof of the characterization is also very short and elementary (one to two pages in [78]), implying that we could give a complete proof of Theorem 6.0.1 on less than four pages.

Let us discuss some further applications of Theorem 6.2.2 now. Here we want to focus again on posets whose cover graphs are contained in certain bounded expansion classes. Namely, such as graphs with bounded genus, bounded tree-width, and those that exclude a fixed graph as a (topological) minor.

We start with graphs of bounded genus (those are including planar graphs). Van den Heuvel et al. [35] show that for each graph G with genus g, we have

$$\mathrm{wcol}_r(G) \leqslant \left(2g + \binom{r+2}{2}\right) \cdot (2r+1).$$

Combining this inequality with Theorem 6.2.2, we obtain the following upper bound.

Corollary 6.2.1. *For every poset P of height at most h whose cover graph has genus g, we have*

$$\dim(P) \leqslant 4^{\left(2g+\binom{3h-1}{2}\right)\cdot(6h-5)} = 2^{\mathcal{O}(h^3)}.$$

In particular, we obtain a single exponential upper bound on the dimension of posets with planar cover graphs. Therefore, it is an improvement to the current best bound due to Joret et al. [38], which is double exponential in height. Meanwhile, Micek, Kozik, and Trotter announced to be able to establish a polynomial upper bound in the case of planar cover graphs. This is an impressive result and might even lead to polynomial bounds in the more general case of bounded genus cover graphs.

We continue with graphs of bounded tree-width, which have polynomial weak coloring numbers as shown by Grohe et al. [32]. More precisely, they show that graphs G of tree-width at most k satisfy

$$\operatorname{wcol}_r(G) \leqslant \binom{k+r}{k} = \mathcal{O}(r^k). \tag{6.3}$$

Regarding the dimension of posets with cover graphs of bounded tree-width, Joret et al. [38] proved an upper bound that is double exponentional in the height of the poset. Combining Theorem 6.2.2 with (6.3), we obtain a single exponential bound.

Corollary 6.2.2. *For every poset P of height at most h with a cover graph of tree-width at most k, we have*

$$\dim(P) \leqslant 4^{\binom{k+3h-3}{k}} = 2^{\mathcal{O}(h^k)}.$$

Surprisingly, this upper bound turns out to be essentially best possible. Recall that in Theorem 4.4.4 we showed that there are posets of height at most h with cover graphs of tree-width at most k and with dimension $2^{\Omega(h^{\lfloor (k-1)/2 \rfloor})}$.

Let us continue with the case that cover graphs exclude a fixed graph as a minor. We use the strong (and almost tight) theorem of van den Heuvel et al. [35] that each graph G excluding K_t as a minor satisfies

$$\operatorname{wcol}_r(G) \leqslant \binom{r+t-2}{t-2} \cdot (t-3)(2r+1) \in \mathcal{O}(r^{t-1}).$$

Together with Theorem 6.2.2, this yields the following improvement on existing double exponential bounds due to Micek and myself [51].

Corollary 6.2.3. *For every poset P of height at most h whose cover graph excludes K_t as a minor, we have*

$$\dim(P) \leqslant 4^{\binom{3h+t-5}{t-2} \cdot (t-3)(6h-5)} = 2^{\mathcal{O}(h^{t-1})}.$$

Again, this bound is essentially best possible by the same construction as in the bounded tree-width case (see Corollary 4.4.1).

Regarding graphs G that exclude K_t as a topological minor, it is implicitly proven in the work of Kreutzer et al. [48] that these graphs satisfy

$$\operatorname{wcol}_r(G) \leqslant 2^{\mathcal{O}(r \log r)}.$$

Combining this inequality with Theorem 6.2.2 we get a slight improvement upon the double exponential bounds derived in [51].

Corollary 6.2.4. *For every poset P of height at most h whose cover graph excludes K_t as a topological minor, we have*

$$\dim(P) \leqslant 2^{2^{\mathcal{O}(h \log h)}}.$$

As we have seen so far, Theorem 6.2.2 implies several essentially best possible upper bounds on the dimension in cases where the cover graphs belong to certain classes with bounded expansion. However, we do not obtain the bounds we are hoping for in the more general case of nowhere dense cover graphs (see Conjecture 6.1.1). Combining Theorem 6.2.2 with Zhu's characterization of nowhere dense classes we only get an exponential bound of $2^{\mathcal{O}(n^{\epsilon})}$, which trivially holds as the dimension of posets with n elements is at most n. For height-2 posets though, the desired bound holds as we show now.

Height-2 Posets with Sparse Cover Graphs. In this part of the section we focus on height-2 posets. We show that we can substantially improve the bound in Theorem 6.2.2 for those posets. This will again have some interesting applications, and one of those is for the case that cover graphs have no K_t-minor. To achieve better bounds in this particular case, we need to introduce a variant of the weak coloring number.

Let G be a graph and $r \geqslant 0$ be an integer. Then we denote by $\mathrm{wcol}^*_r(G)$ the least number c for which there exists a vertex coloring of G with c colors and a linear order π on $V(G)$, such that the following implication holds for any $u, v \in V(G)$: if u is weakly r-reachable from v with respect to π, then u and v have different colors.

Using a greedy coloring as done in the proof of Theorem 6.2.2, we immediately observe that for all $r \geqslant 0$ and every graph G we have

$$\mathrm{wcol}^*_r(G) \leqslant \mathrm{wcol}_r(G).$$

Notice also that we could prove a slightly stronger statement than Theorem 6.2.2 by requiring cover graphs G to satisfy only the weaker condition $\mathrm{wcol}^*_{3h-3}(G) \leqslant c$ there.

Let us continue with our next important result.

Theorem 6.2.3. *For every poset P of height 2 whose cover graph G satisfies $\mathrm{wcol}^*_2(G) \leqslant c$, we have*

$$\dim(P) \leqslant 2c.$$

Proof. Let P be a poset as in the statement of the theorem. We may assume that P contains no independent elements as those do not contribute to $\dim(P)$ except for some trivial cases. With this assumption the sets $\mathrm{Min}(P)$ and $\mathrm{Max}(P)$ are disjoint.

Fix a coloring ϕ of the elements of P that is using colors from the set $[c]$, and let π be a linear order of the elements of P such that ϕ and π are witnessing $\mathrm{wcol}^*_2(G) \leqslant c$.

For each color $i \in [c]$ we construct two linear extensions $L_{i,1}, L_{i,2}$ of P now. Let A_i denote the set of minimal elements with color i and let $\pi(A_i)$ be the restriction of π to the elements of A_i. Similarly, we define B_i and $\pi(B_i)$ with respect to maximal elements of P. (Note that A_i and B_i are disjoint by our initial assumption.)

We construct the linear extension $L_{i,1}$ as follows. We start with $\pi(B_i)$ and extend it by pushing minimal elements of $\mathrm{D}(B_i)$ as high as possible into $\pi(B_i)$. That is, if $a \in \mathrm{D}(B_i)$ then we place a directly before the smallest $b \in B_i$ in $\pi(B_i)$ for which we have $a < b$ in P. Notice that several distinct elements might be assigned to the same position in $\pi(B_i)$; in this case we arbitrarily order those elements. Let $\pi^*(B_i)$ be the resulting linear order.

Dually, we extend $\pi(A_i)$ to $\pi^*(A_i)$ by pushing maximal elements $b \in \mathrm{U}(A_i)$ as low as possible into $\pi(A_i)$. Next, let M_i be an arbitrary linear order on elements not contained in $\pi^*(B_i)$ or $\pi^*(A_i)$ so that comparabilities in P are respected. Then we define $L_{i,1}$ to be the following concatenation:

$$L_{i,1} := \pi^*(B_i) < M_i < \pi^*(A_i).$$

Note that there are no comparabilities between elements of A_i and B_i in P as otherwise we would see a monochromatic edge in the cover graph. Therefore, $L_{i,1}$ is indeed a linear extension of P.

Similarly, we build the linear order $L_{i,2}$. Instead of $\pi(B_i)$ and $\pi(A_i)$ we start with their dual linear orders $\mathrm{rev}(\pi(B_i))$ and $\mathrm{rev}(\pi(A_i))$, respectively, and then proceed with the same pushing operation as done before. If we denote by $\mathrm{rev}^*(\pi(B_i))$ and $\mathrm{rev}^*(\pi(A_i))$ the respective resulting orders, then we set

$$L_{i,2} := \mathrm{rev}^*(\pi(B_i)) < M_i < \mathrm{rev}^*(\pi(A_i)).$$

We claim that $R = \{L_{1,1}, L_{1,2}, \ldots, L_{c,1}, L_{c,2}\}$ is a realizer of P. It is easy to check that all incomparabilites between minimal elements are reversed in R, and that the same holds for incomparabilites between maximal elements. From this and the fact that P has no independent elements it also follows that all pairs $(b, a) \in \mathrm{Inc}(P)$ such that $b \in \mathrm{Max}(P)$ and $a \in \mathrm{Min}(P)$ are reversed in R (those pairs simply are not "critical pairs").

It remains to consider pairs $(a, b) \in \mathrm{Inc}(P)$ for which $a \in \mathrm{Min}(P)$ and $b \in \mathrm{Max}(P)$. Let (a, b) be such a pair and let i, j be the colors of a and b, respectively. If $i = j$ then the pair (a, b) is reversed in $L_{i,1}$. So assume $i \neq j$ from now on. We claim that the pair (a, b) is reversed in at least one of the linear extensions $L_{i,1}, L_{i,2}, L_{j,1}, L_{j,2}$.

Suppose for a contradiction that this is not true. Then we particularly have $a < b$ in $L_{j,1}$ and $L_{j,2}$, which means that we could not push a above b in $\pi(B_j)$ and also not above b in $\mathrm{rev}(\pi(B_j))$. This implies that there exist $b', b'' \in B_j$ such that $a \leqslant b'$ in P, $a \leqslant b''$ in P, and $b' < b < b''$ in π. Observe that this forces $a < b$ in π, since otherwise b' becomes weakly 2-reachable from b'' along the path b'', a, b', which contradicts the fact that b' and b'' received the same color j by our initial coloring ϕ.

Since we also have $a < b$ in $L_{i,1}$ and $L_{i,2}$, we know that we could not push b below a in $\pi(A_i)$ and also not below a in $\mathrm{rev}(\pi(A_i))$. This implies that there exist $a', a'' \in A_i$ such that $a' \leqslant b$ in P, $a'' \leqslant b$ in P, and $a' < a < a''$ in π. However, since $a < b$ in π we deduce that a' is weakly 2-reachable from a'' along the path a'', b, a', which yields a contradiction to the properties of the initial coloring ϕ.

We conclude that (a, b) is reversed by some of our constructed linear extensions, and therefore we have built a realizer of size $2c$. □

Let us discuss some direct consequences of the previous theorem. As noted before, the weak r-coloring numbers can be bounded on nowhere dense graph classes by $\mathcal{O}(n^\epsilon)$ for each $\epsilon > 0$ and each $r \geqslant 0$ (see Zhu [78]). Together with Theorem 6.2.3 this yields the following result.

Corollary 6.2.5. *Posets of height* 2 *with* n *elements and whose cover graphs belong to a nowhere dense class of graphs have dimension* $\mathcal{O}(n^\epsilon)$*, for every* $\epsilon > 0$.

Next, we discuss height-2 posets whose cover graphs contain no K_t-minor. It is proven implicitly in [53] that if every minor of a graph G has average degree at most d, then the *star-chromatic number* $\chi_s(G)$ is $\mathcal{O}(d^2)$. A closer look at the proof contained in the paper reveals that this bound also holds for $\text{wcol}_2^*(G)$. (See Theorem 2.1 in this paper; the authors do not state this observation, but it directly follows from the constructed conflict graph that contains edges between any two vertices which are weakly 2-reachable from each other.) Since graphs with no K_t-minor have average degree $\mathcal{O}(t\sqrt{\log t})$ [47, 72], we obtain the following result with Theorem 6.2.3.

Corollary 6.2.6. *For every height-2 poset P whose cover graph contains no K_t-minor, we have*

$$\dim(P) = \mathcal{O}(t^2 \log t).$$

Further improved upper bounds on the dimension of height-2 posets can be obtained by combining Theorem 6.2.3 with the aforementioned bounds on the weak r-coloring numbers in the special case of $r = 2$.

Corollary 6.2.7. *For every height-2 poset P whose cover graph has genus at most g, we have*

$$\dim(P) = 20g + 60.$$

Corollary 6.2.8. *For every height-2 poset P whose cover graph has tree-width at most t, we have*

$$\dim(P) \leqslant t^2 + 3t + 2.$$

6.3 Application: Boxicity of Graphs

In this section we make a short excursion and consider the concept of boxicity. We show how the ideas developed in the previous section can be used to bound the boxicity of graphs in terms of weak coloring numbers. As an application of this connection we will obtain the following theorem.

Theorem 6.3.1. *Graphs with no K_t-minor have boxicity $\mathcal{O}(t^2 \log t)$.*

A *d-box* is a the Cartesian product $[x_1, y_1] \times \ldots \times [x_d, y_d]$ of d closed intervals of the real line. The *intersection graph* of a collection S of d-boxes is the graph with vertex set S, such that two d-boxes of S are adjacent if and only if their intersection is non empty. The *boxicity* $\text{box}(G)$ of a graph G is defined to be the smallest integer $d \geqslant 1$ for which G is the intersection graph of a collection of d-boxes. A graph with boxicity 1 is called an *interval graph*.

For our proofs we will also use the following equivalent definition: The boxicity of a graph G is the least integer $d \geqslant 1$ for which G is the intersection of d interval graphs. (Recall that the intersection of ℓ graphs $G_1, \ldots, G_\ell$ defined on the same vertex set V is the graph $G = (V, E(G_1) \cap \cdots \cap E(G_\ell))$.)

Roberts [63] introduced the concept of boxicity in 1969. It has applications in the study of ecological [64] and social networks [28]. Outerplanar graphs have boxicity at most 2 [68], planar graphs have boxicity at most 3 [73], and graphs with Euler genus g have boxicity $\mathcal{O}(\sqrt{g} \log g)$ [21]. Graphs of tree-width at most k have $\text{box}(G) \leqslant k + 2$ [12]. For further aspects of this graph parameter we recommend [1, 12, 11] and the references therein.

As noted in [21], the boxicity is unbounded on graphs of bounded degeneracy: The 1-subdivision of the complete graph K_n (which is 2-degenerate) has $\mathrm{box}(K_n) = \Theta(\log\log n)$. This example also shows that the boxicity cannot be bounded from above in terms of the chromatic number. However, it can be bounded in terms of the *acyclic chromatic number* $\chi_a(G)$, which is the smallest integer k for which there exists a proper vertex coloring of G with k colors, such that every cycle in G receives at least three colors (equivalently, any two color classes induce a forest).

Theorem 6.3.2 (Esperet, Joret [22])**.** *For every graph G, we have* $\mathrm{box}(G) \leqslant \chi_a(G)(\chi_a(G) - 1)$.

A consequence of this theorem is that graphs belonging to some fixed class with bounded expansion have bounded boxicity. This follows for instance from Zhu's result (Theorem 6.2.1) and the simple fact that for every graph G we have $\chi_a(G) \leqslant \mathrm{wcol}_2(G)$. Combining this inequality with Theorem 6.3.2, we obtain that the boxicity is bounded by a quadratic function in the weak 2-coloring number.

We prove a linear bound now, which even holds for our newly introduced parameter $\mathrm{wcol}_2^*(G)$. (I am grateful to Louis Esperet for pointing out that the proof of Theorem 6.2.3 directly translates to the boxicity setting.)

Theorem 6.3.3. *For every graph G, we have* $\mathrm{box}(G) \leqslant 2\,\mathrm{wcol}_2^*(G)$.

Proof. Let G be a graph on n vertices and let $c := \mathrm{wcol}_2^*(G)$. By definition, there exist a linear order π on $V(G)$ and a vertex coloring ϕ with colors from the set $\{1, \dots, c\}$, such that whenever a vertex u is weakly 2-reachable from another vertex v with respect to π, then $\phi(u) \neq \phi(v)$.

We aim to show that G is the intersection of $2c$ interval graphs $I_1, \dots, I_{2c}$. We associate to each color $i \in [c]$ the two interval graphs I_{2i-1} and I_{2i}. So let us fix color i for the moment. We explicitly define the intervals representing the vertices of $V(G)$ in I_{2i-1} and I_{2i}, respectively. Let us consider the vertices $v_1, \dots, v_\ell$ that received color i by ϕ. By relabelling the vertices if needed, we may assume that $v_1 < \dots < v_\ell$ holds in π.

We start with I_{2i-1}. Here, we map v_j ($1 \leqslant j \leqslant \ell$) to the point $\{j\}$; and for every vertex u that is not colored with i, we first choose the minimal k ($1 \leqslant k \leqslant \ell$) such that u is adjacent to v_k, and then we map u to the interval $[k, n]$. Notice that I_{2i-1} is a supergraph of G.

We proceed with I_{2i}. Here, we reverse the order of the vertices with color i, that is, we map v_j ($1 \leqslant j \leqslant \ell$) to the point $\{\ell - j + 1\}$; and for every vertex u not colored with i, we first choose the maximmal k' ($1 \leqslant k' \leqslant \ell$) such that u is adjacent to $v_{k'}$, and then we map u to the interval $[\ell - k' + 1, n]$. Notice that I_{2i} is also a supergraph of G. In Figure 6.3 the two interval graphs I_{2i-1} and I_{2i} are illustrated by their induced box representation in dimensions $2i - 1$ and $2i$.

Next, we show that G is the intersection of $I_1, \dots, I_{2c}$. Since all involved interval graphs are supergraphs of G, we only need to show that for each pair of non-adjacent vertices $u, v \in V(G)$ there is an interval graph I_j ($1 \leqslant j \leqslant 2c$) in which the two vertices are mapped to disjoint intervals.

So suppose that u and v are non-adjacent in G. We may assume without loss of generality that $v < u$ in π. If both u and v have the same color i, then their intervals are distinct points in I_{2i-1} (and also in I_{2i}) and thus disjoint. So suppose that u and v have distinct colors i_1 and i_2, respectively. We assume for a contradiction

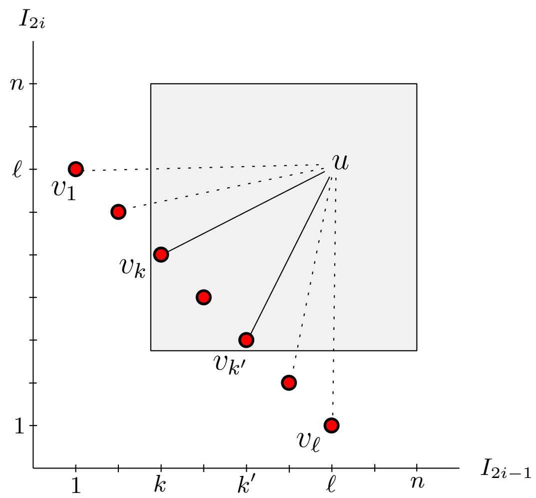

Figure 6.3: Illustration of I_{2i-1} and I_{2i} as the corresponding box representation. Vertices with color i are mapped to the red points. Projections onto the two axis yield the intervals representing the vertices.

that the intervals of u and v intersect in every interval graph $I_1, \dots, I_{2c}$. This holds in particular in I_{2i_2-1} and I_{2i_2} (where v is mapped to a point and u to an interval containing point $\{n\}$); and from this we deduce that there are vertices v' and v'' with color i_2 such that u is adjacent to both of them and $v' < v < v''$ in π. However, together with our assumption $v < u$ in π this implies that v' is weakly 2-reachable from v'' with resptect to π, as is witnessed by the path v', u, v''. This is a contradiction to the properties of the coloring ϕ.

We conclude that G is indeed the intersection of $I_1, \dots, I_{2c}$, and thus $\text{box}(G) \leqslant 2c$. □

We continue with some direct applications of this theorem. As done with our derived inequalities in the previous section, we can now check whether combining Theorem 6.3.3 with existing bounds on the weak 2-coloring numbers of graphs yields improved upper bounds. For instance, for bounded genus graphs we obtain an $\mathcal{O}(g)$ bound on the boxicity in this way. This is a fairly good bound, but does not beat the previously mentioned $\mathcal{O}(\sqrt{g}\log g)$ bound. We are in a similar situation with graphs of tree-width at most k as we cannot beat the $k+2$ bound by Chandran et al.

In the case of K_t-minor-free graphs though, we can achieve an improvement. By using their Theorem 6.3.2, Esperet and Joret [22] proved that the boxicity of graphs with no K_t-minor is $\mathcal{O}(t^4 \log^2 t)$. With our Theorem 6.3.3 and the observation made at the end of Section 6.2 that K_t-minor-free graphs G satisfy $\text{wcol}_2^*(G) = \mathcal{O}(t^2 \log t)$, we conclude that those graphs have boxicity $\mathcal{O}(t^2 \log t)$, which proves Theorem 6.3.1. From below the current best bound can be achieved with a probabilistic argument; it shows that there are K_t-minor-free graphs with boxicity $\Omega(t\sqrt{\log t})$ (see for instance [21]).

Let us conclude now this section with the following problem: Determine the asymptotic behaviour of the maximal boxicity of graphs with no K_t-minor.

Conclusions and Open Problems

It was the general objective of this dissertation to study properties of large-dimensional posets with somewhat sparse cover graphs. This was done with special emphasis upon the following two questions: Do these posets have to be tall; so do they contain large chains as induced subposets? And if so, how tall do they have to be?

At the beginning of my PhD research, the first question was already answered in the affirmative for posets whose cover graphs are planar [71], have bounded treewidth, or more generally exclude a fixed graph as a (topological) minor [76]. With respect to the second question, these results were a bit unsatisfactory as the bounds on the dimension in these results are huge due to the type of methods that have been applied. In this thesis, we were able to significantly improve those bounds in the respective cases and to show that some of them are essentially best possible. A particular strong improvement was obtained for planar posets, where the best upper bounds on their dimension have been exponential in height. Here, I was able to prove a linear upper bound and this one is best possible up to a constant factor.

Besides improving bounds, one of my major goals was to precisely answer the first question above, namely, to determine the exact type of sparsity that is needed in cover graphs to guarantee a bounding function in height on the dimension. I achieved that goal by showing that this property holds up to bounded expansion classes, and that it cannot be extended to nowhere dense graphs. Therefore, I particularly obtained yet another dichotomy result that is establishing a natural boundary between the two concepts of having bounded expansion and being nowhere dense. It is usually not easy to attract the attention of mathematicians from other fields with theorems about order dimension. With the new connections between order dimension and structural sparsity of graphs we have already won the interest of many researchers outside the field.

Shortly after its discovery, a colleague of mine suggested that this result closes the subject of "dimension versus height". For a moment I agreed, but meanwhile with a better understand of sparsity notions I have some objections. There are many exciting problems that naturally arise with our results and that might lead to characterizations of bounded expansion classes and nowhere dense classes in terms of order dimension.

We discuss those problems in the following. Moreover, we comment on several other research directions that are closely related to the results contained in this thesis.

Characterizing Sparsity Notions by Order Dimension

As mentioned above, we have established a natural boundary at bounded expansion classes with respect to order dimension. This does not mean though that we proved a characterization of this sparsity notion in terms of dimension. Several colleagues of mine and I expressed the belief that the following conjecture might be true:

Conjecture 1. *Let $\mathcal{C}$ be a monotone class of graphs. Then $\mathcal{C}$ has bounded expansion if and only if for all $h \geqslant 1$, posets of height at most h with cover graphs in $\mathcal{C}$ have bounded dimension.*

Such a characterization is desirable as it would be yet another one in the long list of seemingly unrelated characterizations for bounded expansion classes. Moreover, it would share the spirit of Schnyder's famous result that planar graphs are characterized by the dimension of their incidence posets [69].

Theorem 6.0.1 shows the forward implication of Conjecture 1, while the backward directions remains wide open. It seems to be a particularly hard problem, especially when comparing it to other rather simple characterizations of bounded expansion classes. Of course, together with my colleagues Gwenaël Joret and Piotr Micek I have tried to prove it. During this research we came up with the following problem, whose positive resolution would imply the backward direction of the conjecture above.

Problem 6.3.1. *Let $\mathcal{C}$ be a monotone class of graphs with unbounded degeneracy. Does it necessarily hold that the height-2 posets with cover graphs in $\mathcal{C}$ have unbounded dimension?*

We believe that this problem is interesting on its own, and therefore we posed it as an open problem at the *Order and Geometry* workshop in Gułtowy in September 2016. Many researcher got attracted and tried to tackle it, but everyone ended up finding interesting special cases that cannot be solved either. I would love to see a proof for this problem.

I also believe that a similar characterization should hold for nowhere dense classes of graphs:

Conjecture 2. *Let $\mathcal{C}$ be a monotone class of graphs. Then $\mathcal{C}$ is nowhere dense if and only if for every $h \geqslant 1$ and every $\epsilon > 0$ it holds that posets of height at most h with cover graphs in $\mathcal{C}$ have dimension $\mathcal{O}(n^\epsilon)$.*

In Section 6.1 we provided several reasons for the correctness of this conjecture. In particular, we observed that the backward direction holds. The forward direction would also show that order dimension (parameterized by height) behaves analogously to many graph parameters for which $\mathcal{O}(n^\epsilon)$–bounds hold on nowhere dense graphs. Due to some partial progress, I am optimistic that this conjecture will be settled within the year 2017.

Sparse Cover Graphs and Induced Subposets

A big line of research of this thesis was concerned with establishing structural properties of large-dimensional posets. By Theorem 6.0.1 we now know that large-dimensional posets with 'sparse' cover graphs need to be *tall*, meaning that they contain a large chain as an induced subposet. A natural question arising here is whether there are other posets being forced to appear as induced subposets. A quick look at Kelly's example (which is 'sparse') yields two further candidates: the first one is $\mathbf{k}+\mathbf{k}$ (two incomparable chains of size k), and the second one is the standard example S_k.

Posets with no two long incomparable chains. Excluding $\mathbf{k}+\mathbf{k}$ as an induced subposet is in a sense more restrictive that bounding the height of a poset. This is justified by the simple observation that posets of height at most $h-1$ do not contain $\mathbf{h}+\mathbf{h}$ as an induced subposet.

Let us consider some known results on $\mathbf{2}+\mathbf{2}$-free posets. In 1970, Fishburn [27] showed that those are characterized by the interval orders. It is well-known that so-called *canonical* interval orders on n elements have dimension $\Theta(\log\log n)$ [29]. Therefore, simply excluding $\mathbf{2}+\mathbf{2}$ as an induced subposet is not enough to bound the dimension. The situation might be different though once we require an additionally sparsity property on the cover graphs. For example, Howard et al. [37] proved that $(\mathbf{k}+\mathbf{k})$-free posets with planar cover graphs have bounded dimension, for every $k \geqslant 1$. This result led the authors to support the following conjecture of Micek and myself [51] (see also [38, 76], where it was posed as a problem).

Conjecture 3. *For all integers $k \geqslant 1$ and $t \geqslant 1$, the $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude K_t as a minor have bounded dimension.*

It was first noted in [51] that this conjecture holds for $k \leqslant 2$: Recall the classic result of Kierstead and Trotter [44] that interval orders of unbounded dimension contain canonical interval orders of unbounded size. Now, the conjecture follows in this case with the following non-trivial fact: If a poset P contains a canonical interval order of large enough size as an induced subposet, then the cover graph of P has a large clique as a minor.

For a short time I believed that the above conjecture might even hold in the setting of cover graphs with no K_t topological minor, or more generally for cover graphs that belong to some fixed class with bounded expansion. However, the following observation of mine destroyed that hope:

Take your favorite large-dimension interval order P and consider an interval representation of P with open intervals whose endpoints lie on integer coordinates. It is easy to see that we can realize it with the additional property that each integer coordinate is the endpoint of at most one interval. Next, we add all possible open intervals of length 1 with integer endpoints. Let P' be the interval order corresponding to the resulting interval representation. Clearly, P is an induced subposet of P' and thus $\dim(P) \leqslant \dim(P')$. Moreover, it is easy to see that the cover graph of P' has maximum degree 3 and hence excludes the star on five vertices as a topological minor.

This construction shows the following: There are $(\mathbf{2}+\mathbf{2})$-free posets whose cover graphs exclude a fixed graph as a topological minor and that have unbounded dimension. Therefore, for $(\mathbf{k}+\mathbf{k})$-free posets there is no such dimension-sparsity result as we have it for the bounded height case, where we established a natural boundary at bounded expansion classes.

Let me conclude the case of $(\mathbf{k}+\mathbf{k})$-free posets with another observation. I argued that Conjecture 3 does not hold with a weaker condition on cover graphs. It turns out that we also cannot ask for something less restrictive on the poset condition, meaning that the conjecture does not hold for $(\mathbf{k}+\mathbf{k}+\mathbf{k})$-free posets. This is witnessed by Kelly's examples as those are $(\mathbf{2}+\mathbf{2}+\mathbf{2})$-free.

Posets without large standard examples. Let me continue with the second conspicuous family of posets being contained in Kelly's construction: the standard examples. Standard examples are witnessing the dimension of Kelly's examples and form the basis in almost all of our lower bound constructions (e.g., those from Section 4.4). With respect to dimension, standard examples play in some sense the role that cliques play for chromatic number.

In the previous paragraph we already observed that S_2-free posets (those are exactly the $(\mathbf{2}+\mathbf{2})$-free posets) can have arbitrarily large dimension. Thus, simply excluding S_k as an induced subposet does not bound the dimension.

Similar to the case of $(\mathbf{k}+\mathbf{k})$-free posets, several colleagues of mine believe that an additional sparsity condition for cover graphs does impose a bound on the dimension of S_k-free posets. Therefore, let me ask: Do large-dimensional posets with 'sparse' cover graphs necessarily contain large standard examples? Recall from the previous paragraph on $(\mathbf{k}+\mathbf{k})$-free posets that the answer to this question is negative once we only require that the cover graphs exclude some fixed graph as a topological minor. However, the following beautiful conjecture might be true.

Conjecture 4. *For every integer $k \geqslant 1$, the dimension of S_k-free posets whose cover graphs exclude some fixed graph as a minor is bounded.*

Somehow expected but not obvious, Conjecture 4 is harder than Conjecture 3: Recall that $(\mathbf{k}+\mathbf{k})$-free posets whose cover graphs exclude some fixed graph as a minor are also S_ℓ-free, for some positive integer ℓ (see Theorem 5.0.2). Therefore, if the above conjecture is true, then this implies the correctness of Conjecture 3.

Another fruitful research direction regarding S_k-free posets is to study their maximal dimension with respect to their number of elements n. It was proven by Biró, Hamburger, and Pór [6] that for every fixed $k \geqslant 1$, the dimension of S_k-free posets is $o(n)$. In the other direction, S_k-free posets can have polynomial dimension in n: Already in 1959 Erdős implicitly proved that there are graphs of girth at least ℓ and chromatic number $\Omega(n^{1/\ell+o(1)})$ [20] (for an improvement see [59]). It follows that the adjacency posets of those graphs have dimension $\Omega(n^{1/\ell+o(1)})$. Moreover, if the girth ℓ is at least 7 then it is easy to see that the adjacency posets contain no S_3. Thus, we deduce that there are S_3-free posets with polynomial dimension $\Omega(n^{1/7+o(1)})$. (For S_2-free posets P such a result does not hold, since their dimension is $\mathcal{O}(\log(\text{height}(P))$ [61].)

If we add a sparsity condition on the cover graphs though, the dimension of S_k-free posets might be much smaller. Analogous to Conjecture 6.1.1 we ask:

Question. *Do S_k-free posets with nowhere dense cover graphs have dimension $\mathcal{O}(n^\epsilon)$, for every $\epsilon > 0$?*

For other aspects of S_k-free posets (e.g. the case that k is not fixed but close to n) we recommend [6, 7].

Cover Graphs that Impose Constant Bounds

Throughout the thesis we have seen that certain graph properties of the cover graph of a poset P can have a dramatic effect on the dimension of P. In some cases the dimension can even be bounded by constant. For instance, recall that if the cover graph of P is a tree then the dimension of P is at most 3. As we have discussed, absolute constant upper bounds also exist when the cover graph of P is outerplanar, has path-width 2, or more generally has tree-width at most 2.

All these graph properties define proper minor-closed graph classes. Interestingly, graphs of tree-width at most 2 are characterized by the graphs that exclude K_4 as a minor. This leads to the following open problem:

Question. *For which graphs H does there exist a constant c_H, such that for every poset P whose cover graph excludes H as a minor we have* $\dim(P) \leqslant c_H$*?*

Let $\mathcal{H}$ be the set of graphs H that admit such a constant c_H. Then we have for instance $K_4 \in \mathcal{H}$, while $K_5 \notin \mathcal{H}$ as witnessed by Kelly's examples. Can we characterize $\mathcal{H}$? Clearly, $\mathcal{H}$ is a proper minor-closed graph class. Thus, by the graph minor theorem of Robertson and Seymour it follows that there is finite list of graphs $H_1, \ldots, H_k$ such that

$$\mathcal{H} = \{G \mid G \text{ excludes } H_i \text{ as a minor, for all } i \in [k]\}.$$

So can we even characterize $\mathcal{H}$ by specifying the explicit list of graphs $H_1, \ldots, H_k$?

I believe that this is very nice question and that $\mathcal{H}$ might be a very natural graph class. However, it seems to be a very hard problem. Already proving $K_4 \in \mathcal{H}$ turned out to be tough; our current shortest proof of this result takes more than 40 pages (see [39]). On the other hand, there may be further exciting unknown connections from graph theory to dimension theory that might lead to short proofs and strong results.

Bibliography

[1] Abhijin Adiga, Diptendu Bhowmick, and L. Sunil Chandran. Boxicity and poset dimension. *SIAM J. Discrete Math.*, 25(4):1687–1698, 2011.

[2] K. A. Baker, P. C. Fishburn, and F. S. Roberts. Partial orders of dimension 2. *Networks*, 2:11–28, 1972.

[3] János Barát, Péter Hajnal, Yixun Lin, and Aifeng Yang. On the structure of graphs with path-width at most two. *Studia Sci. Math. Hungar.*, 49(2):211–222, 2012.

[4] Giuseppe Di Battista, Wei-Ping Liu, and Ivan Rival. Bipartite graphs, upward drawings, and planarity. *Information Processing Letters*, 36(6):317 – 322, 1990.

[5] Garrett Birkhoff. *Lattice theory.* Third edition. American Mathematical Society Colloquium Publications, Vol. XXV. American Mathematical Society, Providence, R.I., 1967.

[6] Csaba Biró, Peter Hamburger, and Attila Pór. Standard examples as subposets of posets. *Order*, 32(2):293–299, 2015.

[7] Csaba Biró, Peter Hamburger, Attila Pór, and William T. Trotter. Forcing posets with large dimension to contain large standard examples. *Graphs Combin.*, 32(3):861–880, 2016.

[8] Csaba Biró, Mitchel T. Keller, and Stephen J. Young. Posets with cover graph of pathwidth two have bounded dimension. *Order*, 33(2):195–212, 2016.

[9] Graham Brightwell. On the complexity of diagram testing. *Order*, 10(4):297–303, 1993.

[10] Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Graph products revisited: tight approximation hardness of induced matching, poset dimension and more. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1557–1576. SIAM, Philadelphia, PA, 2012.

[11] L. Sunil Chandran, Mathew C. Francis, and Naveen Sivadasan. Boxicity and maximum degree. *J. Combin. Theory Ser. B*, 98(2):443–445, 2008.

[12] L. Sunil Chandran and Naveen Sivadasan. Boxicity and treewidth. *J. Combin. Theory Ser. B*, 97(5):733–744, 2007.

[13] Anuj Dawar. Homomorphism preservation on quasi-wide classes. *J. Comput. System Sci.*, 76(5):324–332, 2010.

[14] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[15] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51(1):161–166, 1950.

[16] Vida Dujmović, Gwenaël Joret, and David R. Wood. An improved bound for first-fit on posets without two long incomparable chains. *SIAM J. Discrete Math.*, 26(3):1068–1075, 2012.

[17] Ben Dushnik and Edwin W. Miller. Partially ordered sets. *Amer. J. Math.*, 63:600–610, 1941.

[18] Zdeněk Dvořák. Constant-factor approximation of the domination number in sparse graphs. *European J. Combin.*, 34(5):833–840, 2013.

[19] Zdeněk Dvořák. Asymptotical structure of combinatorial objects. PhD thesis, Charles University, Faculty of Mathematics and Physics, 2007.

[20] P. Erdős. Graph theory and probability. *Canad. J. Math.*, 11:34–38, 1959.

[21] Louis Esperet. Boxicity and topological invariants. *European J. Combin.*, 51:495–499, 2016.

[22] Louis Esperet and Gwenaël Joret. Boxicity of graphs on surfaces. *Graphs Combin.*, 29(3):417–427, 2013.

[23] Stefan Felsner, Tomasz Krawczyk, and William T. Trotter. On-line dimension for posets excluding two long incomparable chains. *Order*, 30(1):1–12, 2013.

[24] Stefan Felsner, Ching Man Li, and William T. Trotter. Adjacency posets of planar graphs. *Discrete Math.*, 310(5):1097–1104, 2010.

[25] Stefan Felsner, Irina Mustaţă, and Martin Pergel. The complexity of the partial order dimension problem - closing the gap. to appear in SIDMA, https://arxiv.org/abs/1501.01147arXiv:1501.01147.

[26] Stefan Felsner, William T. Trotter, and Veit Wiechert. The Dimension of Posets with Planar Cover Graphs. *Graphs Combin.*, 31(4):927–939, 2015.

[27] Peter C. Fishburn. Intransitive indifference with unequal indifference intervals. *J. Mathematical Psychology*, 7:144–149, 1970.

[28] L. C. Freeman. Spheres, cubes and boxes: Graph dimensionality and network structure. *Social Networks*, 5(2):139–156, 1983.

[29] Z. Füredi, P. Hajnal, V. Rödl, and W. T. Trotter. Interval orders and shift graphs. In *Sets, graphs and numbers (Budapest, 1991)*, volume 60 of *Colloq. Math. Soc. János Bolyai*, pages 297–313. North-Holland, Amsterdam, 1992.

[30] Z. Füredi and J. Kahn. On the dimensions of ordered sets of bounded degree. *Order*, 3(1):15–20, 1986.

[31] Ashim Garg and Roberto Tamassia. On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.*, 31(2):601–625 (electronic), 2001.

[32] Martin Grohe, Stephan Kreutzer, Roman Rabinovich, Sebastian Siebertz, and Konstantinos Stavropoulos. *Colouring and Covering Nowhere Dense Graphs*, pages 325–338. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.

[33] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. In *STOC'14—Proceedings of the 2014 ACM Symposium on Theory of Computing*, pages 89–98. ACM, New York, 2014.

[34] Martin Grohe and Dániel Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. *SIAM J. Comput.*, 44(1):114–159, 2015.

[35] Jan van den Heuvel, Patrice Ossona de Mendez, Daniel Quiroz, Roman Rabinovich, and Sebastian Siebertz. On the generalised colouring numbers of graphs that exclude a fixed minor. arXiv:1602.09052, 2016.

[36] Toshio Hiraguchi. On the dimension of partially ordered sets. *Sci. Rep. Kanazawa Univ.*, 1:77–94, 1951.

[37] David M. Howard, Noah Streib, William T. Trotter, Ruidong Wang, and Bartosz Walczak. The dimension of posets with planar cover graphs excluding two long incomparable chains. arXiv:1608.08843.

[38] Gwenaël Joret, Piotr Micek, Kevin G. Milans, William T. Trotter, Bartosz Walczak, and Ruidong Wang. Tree-width and dimension. *Combinatorica*, 36(4):431–450, 2016.

[39] Gwenaël Joret, Piotr Micek, William T. Trotter, Ruidong Wang, and Veit Wiechert. On the dimension of posets with cover graphs of treewidth 2. *Order*, to appear. arXiv:1406.3397.

[40] Gwenaël Joret, Piotr Micek, and Veit Wiechert. Planar posets have dimension linear in height. Submitted.

[41] Gwenaël Joret, Piotr Micek, and Veit Wiechert. Sparsity and dimension. *Combinatorica*, to appear. arXiv:1507.01120v2.

[42] Gwenaël Joret, Piotr Micek, and Veit Wiechert. Sparsity and dimension. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 1804–1813, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics.

[43] David Kelly. On the dimension of partially ordered sets. *Discrete Math.*, 35:135–156, 1981.

[44] H. A. Kierstead and W. T. Trotter. Interval orders and dimension. *Discrete Math.*, 213(1-3):179–188, 2000. Selected topics in discrete mathematics (Warsaw, 1996).

[45] H. A. Kierstead and Daqing Yang. Orderings on graphs and game coloring number. *Order*, 20(3):255–264 (2004), 2003.

[46] Henry A. Kierstead, George F. McNulty, and William T. Trotter. A theory of recursive dimension for ordered sets. *Order*, 1(1):67–82, 1984.

[47] A. V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[48] Stephan Kreutzer, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. The generalised colouring numbers on classes of bounded expansion. *CoRR*, abs/1606.08972, 2016.

[49] Michał Lasoń, Piotr Micek, Noah Streib, William T. Trotter, and Bartosz Walczak. An extremal problem on crossing vectors. *J. Combin. Theory Ser. A*, 128:41–55, 2014.

[50] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[51] Piotr Micek and Veit Wiechert. Topological minors of cover graphs and dimension. *Journal of Graph Theory*, to appear. arXiv:1504.07388.

[52] John Isham Moore, Jr. *Graphs and Partially Ordered Sets*. ProQuest LLC, Ann Arbor, MI, 1975. Thesis (Ph.D.)–University of South Carolina.

[53] Jaroslav Nešetřil and Patrice Ossona de Mendez. Colorings and homomorphisms of minor closed classes. In *Discrete and computational geometry*, volume 25 of *Algorithms Combin.*, pages 651–664. Springer, Berlin, 2003.

[54] Jaroslav Nešetřil and Patrice Ossona de Mendez. On nowhere dense graphs. *European J. Combin.*, 32(4):600–617, 2011.

[55] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity*, volume 28 of *Algorithms and Combinatorics*. Springer, Heidelberg, 2012. Graphs, structures, and algorithms.

[56] Jaroslav Nešetřil, Patrice Ossona de Mendez, and David R. Wood. Characterisations and examples of graph classes with bounded expansion. *European J. Combin.*, 33(3):350–373, 2012.

[57] Jaroslav Nešetřil, Patrice Ossona de Mendez, and David R. Wood. Characterisations and examples of graph classes with bounded expansion. *European J. Combin.*, 33(3):350–373, 2012. arXiv:0902.3265.

[58] Jaroslav Nešetřil and Vojtěch Rödl. Complexity of diagrams. *Order*, 3(4):321–330, 1987.

[59] Deryk Osthus and Anusch Taraz. Random maximal H-free graphs. *Random Structures Algorithms*, 18(1):61–82, 2001.

[60] János Pach and Géza Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997.

[61] Issie Rabinovitch. *The Dimension-Theorey of Semiorders and Interval-orders.* ProQuest LLC, Ann Arbor, MI, 1973. Thesis (Ph.D.)–Dartmouth College.

[62] Felix Reidl, Fernando Sanchez Villaamil, and Konstantinos Stavropoulos. Characterising bounded expansion by neighbourhood complexity. arXiv:1603.09532.

[63] Fred S. Roberts. On the boxicity and cubicity of a graph. In *Recent Progress in Combinatorics (Proc. Third Waterloo Conf. on Combinatorics, 1968)*, pages 301–310. Academic Press, New York, 1969.

[64] Fred S. Roberts. *Discrete Mathematical Models with Applications to Social, Biological, and Environmental Problems.* Prentice Hall, Englewood Cliffs, New Jersey, 1976.

[65] Neil Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. *J. Combin. Theory Ser. B*, 92(2):325–357, 2004.

[66] Neil Robertson and P.D Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43 – 76, 2003.

[67] Vojtec Rödl and William T. Trotter. unpublished result.

[68] Edward R. Scheinerman. *Intersection Classes and Multiple Intersection Parameters of Graphs.* ProQuest LLC, Ann Arbor, MI, 1984. Thesis (Ph.D.)–Princeton University.

[69] Walter Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.

[70] Bernd Schröder. *Ordered sets.* Birkhäuser/Springer, second edition, 2016.

[71] Noah Streib and William T. Trotter. Dimension and height for posets with planar cover graphs. *European J. Combin.*, 35:474–489, 2014.

[72] Andrew Thomason. An extremal function for contractions of graphs. *Math. Proc. Cambridge Philos. Soc.*, 95(2):261–265, 1984.

[73] Carsten Thomassen. Interval representations of planar graphs. *J. Combin. Theory Ser. B*, 40(1):9–20, 1986.

[74] William T. Trotter. *Combinatorics and partially ordered sets. Dimension theory.* Johns Hopkins Series in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 1992.

[75] William T. Trotter, Jr. and John I. Moore, Jr. The dimension of planar posets. *J. Combinatorial Theory Ser. B*, 22(1):54–67, 1977.

[76] Bartosz Walczak. Minors and dimension. *J. Combin. Theory Ser. B*, 122:668–689, 2017.

[77] Mihalis Yannakakis. The complexity of the partial order dimension problem. *SIAM J. Algebraic Discrete Methods*, 3(3):351–358, 1982.

[78] Xuding Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Math.*, 309(18):5562–5568, 2009.

Index