


Mareike Massow

LINEAR EXTENSION GRAPHS AND LINEAR EXTENSION DIAMETER



Cuvillier Verlag Göttingen
Internationaler wissenschaftlicher Fachverlag

LINEAR EXTENSION GRAPHS AND LINEAR EXTENSION DIAMETER

vorgelegt von
Diplom-Mathematikerin
Mareike Massow
aus Paderborn

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

Doktor der Naturwissenschaften
– Dr. rer. nat. –

genehmigte Dissertation

Promotionsausschuss:

Vorsitzende: Prof. Dr. Petra Wittbold
Berichter: Prof. Dr. Stefan Felsner
Prof. Dr. Graham R. Brightwell

Tag der wissenschaftlichen Aussprache: 11. Dezember 2009

Berlin 2009

D 83

Bibliografische Information der Deutschen Nationalbibliothek

Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über <http://dnb.d-nb.de> abrufbar.

1. Aufl. - Göttingen: Cuvillier, 2010
Zugl.: (TU) Berlin (Univ.), Diss., 2009

978-3-86955-245-3

© CUVILLIER VERLAG, Göttingen 2010
Nonnenstieg 8, 37075 Göttingen
Telefon: 0551-54724-0
Telefax: 0551-54724-21
www.cuvillier.de

Alle Rechte vorbehalten. Ohne ausdrückliche Genehmigung des Verlages ist es nicht gestattet, das Buch oder Teile daraus auf fotomechanischem Weg (Fotokopie, Mikrokopie) zu vervielfältigen.

1. Auflage, 2010
Gedruckt auf säurefreiem Papier

978-3-86955-245-3

IT TURNS OUT THAT AN EERIE TYPE OF CHAOS
CAN LURK JUST BEHIND A FAÇADE OF ORDER
– AND YET, DEEP INSIDE THE CHAOS
LURKS AN EVEN EERIER TYPE OF ORDER.

D. HOFSTADTER

Preface

The time I spent working on this thesis has been very enjoyable for a number of reasons. The great atmosphere for Discrete Mathematics in Berlin is certainly one of them. It has been praised in so many prefaces that it is hard to find original words. Nonetheless, be assured – the rumors are true.

I much like to thank my advisor Stefan Felsner: For giving me the freedom to choose my problems, and for discussing them until they surrendered. I am also very thankful to Graham Brightwell, for hosting me at the LSE in London for a fruitful spring, and for being the second reviewer of my thesis.

My research was made possible by the support of the research training group “Methods for Discrete Structures”¹. Apart from financial support, it provided the opportunity to discuss with many people interested in my topic, in particular Felix Breuer, Holger Dell, Kolja Knauer, Christina Puhl, and Daria Schymura. More thanks go to Wiebke Höhn and Axel Werner for being there and keeping a chocolate supply.

It was a pleasure being part of the workgroup “Discrete Mathematics” at TU Berlin. I especially like to thank Torsten Ueckerdt for his algorithmic help, and Kolja Knauer for always being up to discussing wild conjectures with me (some of which even prove to be true, see Section 5.2).

I am unspeakably thankful to Aaron Dall for proofreading this thesis, for discussing language, and for everything beyond language.

Last and least: This thesis is proof that mathematical research without coffee (well, almost) is possible.

Mareike Massow
Berlin, October 2009

¹Research supported by DFG-GRK 1408 and the Berlin Mathematical School.

Contents

Contents	vii
Introduction	1
1 Notation and Tools	5
1.1 The Basics	5
1.1.1 Linear Extensions	7
1.1.2 Special Classes of Posets	8
1.2 Poset Dimension	9
1.3 Partial Cubes	13
1.4 Modular Decomposition	17
2 Properties of Linear Extension Graphs	23
2.1 Context and Previous Results	24
2.2 Color Classes	29
2.3 Adjacent Swap Colors	35
2.4 The Linear Extension Graph as an Invariant	40
3 Reconstructing Posets from Linear Extension Graphs	43
3.1 Answering Admissible Questions	45
3.2 Checking Adjacency of Swap Colors	47
3.3 Reconstructing the Comparability Graph	48
3.4 Orienting the Comparability Graph	51
3.4.1 The Parallel Case	54
3.4.2 The Series Case	54
3.4.3 The Prime Case	55
3.4.4 The Exceptional Case	57

3.5	Putting Things Together	57
3.6	Recognizing Linear Extension Graphs	58
3.6.1	Constructing a Candidate Poset	58
3.6.2	Finding a Start	60
3.6.3	Labeling Vertices with Linear Extensions	61
4	Complexity of Linear Extension Diameter	63
4.1	Previous Results	64
4.2	NP-Completeness for General Posets	67
4.3	Algorithm for Posets of Width 3	73
5	Linear Extension Diameter of Distributive Lattices	77
5.1	Boolean Lattices	78
5.1.1	Revlex Linear Extensions	78
5.1.2	Determining the Linear Extension Diameter	80
5.1.3	Characterizing Diametral Pairs	83
5.2	Downset Lattices of 2-Dimensional Posets	85
5.2.1	Revlex Pairs are Diametral Pairs	86
5.2.2	Diametral Pairs are Revlex Pairs	91
5.2.3	Computing the Linear Extension Diameter	94
5.3	Open Questions	100
6	Diametrically Reversing Posets	101
6.1	Not all Posets are Diametrically Reversing	103
6.2	Most Posets are Diametrically Reversing	107
6.2.1	Modules	107
6.2.2	Interval Orders	108
6.2.3	3-Layer Posets	112
	Bibliography	115
	Index	121

Introduction

A poset is a set equipped with a partial order relation. Since the ordering is only partial, there are usually many ways to extend it to a linear order. Each of them yields a linear extension of the poset.

We are interested in the set of all linear extensions of a given finite poset \mathcal{P} . To make the structure of this set tangible, we consider the linear extension graph $G(\mathcal{P})$. It has the linear extensions of \mathcal{P} as vertices, with two of them being adjacent if they differ exactly by one adjacent swap of elements. Figure 0.1 shows a simple example, the cover of this thesis a more complicated one.

This thesis investigates how properties of the poset \mathcal{P} are reflected in the linear extension graph $G(\mathcal{P})$, and vice versa. We place a special emphasis on the diameter of $G(\mathcal{P})$.




Figure 0.1: A poset and its linear extension graph with swap coloring.

Chapter 1

The function of this chapter is to set the stage for what happens in the later chapters. We recall basic notions of poset theory and specify our notation. We introduce partial cubes and some important properties of them. Also,

we present the basic results of Gallai's theory of modular decomposition and transitive orientations.

Chapter 2

In this chapter we introduce linear extension graphs. We present some of their history and previously known properties. Then we focus on the *swap coloring* of the edges of $G(\mathcal{P})$, in which each edge is colored with the pair of elements of \mathcal{P} swapped along it. Our main result of this chapter is that we can characterize in terms of the graph $G(\mathcal{P})$ which pairs of swap colors share an element of \mathcal{P} .

We also discuss which modifications of the poset \mathcal{P} leave the graph $G(\mathcal{P})$ invariant, and which do not. The results of this chapter provide the ingredients for the reconstruction procedure of the next chapter.

Chapter 3

Here we present a procedure which, given a linear extension graph G , reconstructs all posets \mathcal{P} with $G = G(\mathcal{P})$. We prove that if G cannot be decomposed into several non-trivial Cartesian factors, then \mathcal{P} is unique up to duality and the addition of global elements. The procedure makes fundamental use of Gallai's modular decomposition theory. In the last section, we show how to use the reconstruction procedure to recognize linear extension graphs.

Chapter 4

In this chapter we introduce the second part of the title of this thesis: The linear extension diameter of a poset. Given a poset \mathcal{P} , the linear extension diameter $\text{led}(\mathcal{P})$ is the maximum number of pairs of elements of \mathcal{P} which can appear in different orders in two linear extensions of \mathcal{P} . It can easily be seen that $\text{led}(\mathcal{P})$ equals the graph diameter of $G(\mathcal{P})$.

At the beginning of this chapter we present some previous results on the linear extension diameter. Then we prove that, given a poset \mathcal{P} and some $k \geq 2$, it is NP-complete to decide whether $\text{led}(\mathcal{P}) \geq k$. On the positive side, we show how to compute in polynomial time the linear extension diameter of a poset of width 3, using a dynamic programming approach.

The results of this chapter are joint work with Graham Brightwell. They are also contained in [7].

Chapter 5

In the first part of this chapter we prove a formula for the linear extension diameter of Boolean lattices which had been conjectured in [22]. Moreover, we characterize the diametral pairs of linear extensions of Boolean lattices. The proofs only use basic combinatorial arguments.

The cover of this thesis shows the linear extension graph of the 3-dimensional Boolean lattice. The linear extensions contained in diametral pairs are highlighted.

Boolean lattices are downset lattices of antichains. Now let $\mathcal{D}_{\mathcal{P}}$ be the downset lattice of an arbitrary 2-dimensional poset \mathcal{P} . In the second part of this chapter we characterize the diametral pairs of linear extensions of $\mathcal{D}_{\mathcal{P}}$. Furthermore, we show that we can compute the linear extension diameter of $\mathcal{D}_{\mathcal{P}}$ in time polynomial in $|\mathcal{P}|$. The proofs use characteristics of the 2-dimensional poset \mathcal{P} .

This chapter is joint work with Stefan Felsner. The results can also be found in [20].

Chapter 6

This chapter deals with a property of posets which we call diametrically reversing. A linear extension of a poset \mathcal{P} is *reversing* if it reverses some critical pair of elements of \mathcal{P} . A poset \mathcal{P} is *diametrically reversing* if every linear extension of \mathcal{P} which is part of a diametral pair of linear extensions of \mathcal{P} is reversing.

It follows from the results of Chapter 5 that Boolean lattices are diametrically reversing. We give an example of a poset \mathcal{P} such that no linear extension of \mathcal{P} contained in a diametral pair is reversing. This disproves a conjecture from [22]. On the other hand, we exhibit some classes of posets which are diametrically reversing, including interval orders and 3-layer posets. The last class shows that almost all posets are diametrically reversing.

The results of this chapter are joint work with Graham Brightwell. They can also be found in [7].

Chapter 1

Notation and Tools

This chapter sets the stage for the results and proofs we are going to present in this thesis. Readers who are familiar with the topics may skip it and consult it only as a reference when needed.

Section 1.1 contains basic definitions and notations concerning posets. Section 1.2 introduces the dimension of posets, with a special emphasis on critical pairs and on characterizations of 2-dimensional posets. In Section 1.3 we look at partial cubes and some of their properties. Section 1.4 presents the basics of Gallai's theory on modular decomposition and transitive orientation of comparability graphs.

1.1 The Basics

Throughout this thesis we assume basic graph theory to be known to the reader. We mainly use the notation from [13].

Let us start with the objects which are at the basis of this thesis: A partial order or *poset* \mathcal{P} is a set P equipped with a binary relation \leq on P , which is reflexive, antisymmetric, and transitive. Formally, the relation \leq is a subset of the ordered pairs $(x, y) \in P \times P$. In this thesis we will only consider *finite* posets, that is, the *ground set* P is finite.

A poset \mathcal{Q} is a *subposet* of \mathcal{P} if $Q \subseteq P$, and the relations of \mathcal{Q} are induced by the relations of \mathcal{P} , that is, for each pair x, y of elements of \mathcal{Q} we have $x \leq y$ in \mathcal{Q} exactly if $x \leq y$ in \mathcal{P} . By standard abuse of notation, we usually identify the ground set of a poset with the whole poset.

We write $x < y$ if $x \leq y$ and $x \neq y$. We also write $y \geq x$ if $x \leq y$ in \mathcal{P} , and $y > x$ if $y \geq x$ and $y \neq x$. Two elements of \mathcal{P} standing in the relation of \mathcal{P} are called *comparable*, otherwise they are *incomparable*. We write $x \sim y$ if $\{x, y\} \in \mathcal{P}$ is a comparable pair of elements, and $x \parallel y$ if they are incomparable. Comparable and incomparable pairs are unordered pairs, and we mostly write x, y or xy instead of $\{x, y\}$. We denote the set of incomparable pairs of \mathcal{P} by $\text{Inc}(\mathcal{P})$, and let $|\text{Inc}(\mathcal{P})| = \text{inc}(\mathcal{P})$.

If $x < y$, then the relation between x and y is a *cover relation* if there is no third element $w \in \mathcal{P}$ such that $x < w < y$. The *interval* $I(x, y) \subseteq \mathcal{P}$ consists of all elements w such that $x \leq w \leq y$.

A poset is a *chain* if it contains no incomparable pair of elements. In this case, the partial order is a *linear order*. A poset is an *antichain* if all of its pairs are incomparable. We denote the antichain on n elements by \mathcal{A}_n . The *height* of a poset \mathcal{P} is the number of elements of the longest chain appearing as a subset of \mathcal{P} . The *width* of a poset is the number of elements of the largest antichain appearing as a subset of \mathcal{P} .

The *Hasse diagram* of \mathcal{P} is a graph which has the elements of \mathcal{P} as vertices and draws an upward edge between $x, y \in \mathcal{P}$ for every cover relation $x < y$. That is, the graph is drawn undirected, but the edges have an implicit direction. Figure 1.1 shows the Hasse diagram of a poset called *Chevron*. All posets depicted in this thesis are shown by their Hasse diagrams.




Figure 1.1: The Hasse diagram of the Chevron.

The *comparability graph* $\text{Comp}(\mathcal{P})$ of a poset \mathcal{P} is the undirected graph whose vertex set is the ground set of \mathcal{P} , with two vertices being adjacent exactly if the corresponding elements are comparable in \mathcal{P} . A poset \mathcal{P} is called *connected* if $\text{Comp}(\mathcal{P})$ is connected, and the *connected components* of \mathcal{P} are the connected components of $\text{Comp}(\mathcal{P})$. A *comparability invariant* is a poset parameter that depends only on the comparability graph of the poset.

The *incomparability graph* $\text{Incomp}(\mathcal{P})$ of a poset \mathcal{P} is the undirected graph whose vertex set is the ground set of \mathcal{P} , with two vertices being adjacent exactly if the corresponding elements are incomparable in \mathcal{P} .

The *dual* \mathcal{P}^* of a poset \mathcal{P} is the poset we obtain from \mathcal{P} by reversing its relation, that is, by replacing $x \leq y$ by $y \geq x$ for all comparable pairs x, y of \mathcal{P} . We say that \mathcal{P} and \mathcal{P}^* differ in *direction*. This is the first non-standard term we define, and it will become important later.

An element $x \in \mathcal{P}$ with $x \leq y$ for all $y \in \mathcal{P}$ is called a *global minimum* of \mathcal{P} . Analogously, if $x \geq y$ for all $y \in \mathcal{P}$, then x is a *global maximum* of \mathcal{P} . If $x \sim y$ for all $y \in \mathcal{P}$ with $y \neq x$, then x is a *global element* of \mathcal{P} . We extend this notation to graphs and call a vertex v of a graph G a *global vertex* if v is adjacent to every other vertex of G .

If $x, y \in \mathcal{P}$ and $x < y$, then x is a *predecessor* of y , and y is a *successor* of x . We denote the set of predecessors of y by $\text{Pred}(y)$, and the set of successors of x by $\text{Succ}(x)$. If we need to specify the poset \mathcal{P} , we write $\text{Pred}_{\mathcal{P}}(y)$ and $\text{Succ}_{\mathcal{P}}(x)$. We define the *downset of a set* $A \subseteq \mathcal{P}$ as the set of all $x \in \mathcal{P}$ with $x \leq a$ for some $a \in A$, and denote it by A^\downarrow . The *upset of* A , denoted by A^\uparrow , is analogously defined. This means that all downsets and upsets in this thesis are closed.

There is a canonical bijection between the downsets of \mathcal{P} and the antichains of \mathcal{P} , since the maximal elements of every downset form an antichain, and conversely, each antichain $A \subseteq \mathcal{P}$ generates the downset $\{v \in \mathcal{P} : v \leq a \text{ for some } a \in A\}$.

1.1.1 Linear Extensions

As mentioned in the introduction, this thesis investigates the set of linear extensions of a poset \mathcal{P} . A *linear extension* L of \mathcal{P} is a linear order on the elements of \mathcal{P} , such that $x \leq y$ in \mathcal{P} implies $x \leq y$ in L for all $x, y \in \mathcal{P}$. We write a linear extension as follows:

$$L = x_1 x_2 \dots x_n,$$

which stands for $x_1 < x_2 < \dots < x_n$ in L . For example, $L = 123456$ and $L' = 315246$ are linear extensions of the Chevron (see Figure 1.1).

Every poset has a linear extension. The *generic algorithm* to build a linear extension L of \mathcal{P} consists of n simple steps: In step i , choose x_i from the minima of $\mathcal{P} - \{x_1, \dots, x_{i-1}\}$. After the n -th step, output $L = x_1 x_2 \dots x_n$. Specifications of the generic algorithm can be given by defining a *priority list* for choosing x_i in step i .

The number of linear extensions of a poset \mathcal{P} on n elements can be as large as $n!$, which is attained for $\mathcal{P} = \mathcal{A}_n$. Thus it may be exponential in $|\mathcal{P}|$. Brightwell and Winkler showed in [8] that counting the number of linear extensions of a given poset is $\#P$ -complete.

Let L and L' be two linear extensions of a poset \mathcal{P} . A pair $\{x, y\}$, or xy for short, of elements of \mathcal{P} such that $x < y$ in L and $x > y$ in L' is called a *reversal* between L and L' . The *distance* $\text{dist}(L, L')$ between L and L' is the number of reversals between L and L' . Clearly, only incomparable pairs of elements of \mathcal{P} can be reversals. Hence we have $\text{dist}(L, L') \leq \text{inc}(\mathcal{P})$. As for the two linear extensions of the chevron given above, there are four reversals between them, thus $\text{dist}(L, L') = 4$.

1.1.2 Special Classes of Posets

The *standard example* S_n is the poset of height two consisting of two antichains $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ such that $a_i < b_j$ in S_n exactly if $j \neq i$.

A poset \mathcal{P} is *graded* if it can be equipped with a rank function ρ from \mathcal{P} to the integer numbers such that $x \leq y$ implies $\rho(x) \leq \rho(y)$, and whenever $x < y$ is a cover relation, then $\rho(y) = \rho(x+1)$. We say that the elements with the same ρ -value form a *level* of \mathcal{P} . Each graded poset \mathcal{P} can be partitioned into levels A_1, \dots, A_r , such that each level A_i forms an antichain in \mathcal{P} , and cover relations only appear between adjacent levels A_i and A_{i+1} .

A *lattice* is a poset \mathcal{P} such that any two elements $x, y \in \mathcal{P}$ have a join and a meet in \mathcal{P} . The *join* $x \vee y$ of x, y is an element $z \in \mathcal{P}$ such that $z \geq x$ and $z \geq y$, and if there is another element $w \in \mathcal{P}$ fulfilling these two relations, then we have $w \geq z$. The *meet* $x \wedge y$ is defined analogously by replacing \geq with \leq . A *join-irreducible element* of \mathcal{P} is an element v which is not the join of two elements different from v .

A *distributive lattice* is a lattice such that $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ for every pair $x, y \in \mathcal{P}$. The *downset lattice* $\mathcal{D}_{\mathcal{P}}$ of a poset \mathcal{P} is the poset on the set of all downsets of \mathcal{P} , ordered by inclusion. Each downset lattice is a distributive lattice, with the least upper bound and the greatest lower bound of two downsets given by their union and their intersection, respectively. Birkhoff's famous representation theorem for distributive lattices [3] says that each distributive lattice \mathcal{D} is the downset lattice of the poset induced by the join-irreducible elements of \mathcal{D} (see e.g. [11]).

Each downset lattice is graded, with a rank function given by the number of elements in the downsets. The 1-element sets of a downset lattice $\mathcal{D}_{\mathcal{P}}$ are called the *atoms* of $\mathcal{D}_{\mathcal{P}}$. The downsets containing all but one element of \mathcal{P} are called the *coatoms* of $\mathcal{D}_{\mathcal{P}}$.

The *Boolean lattice* B_n is the poset on all subsets of $[n]$, ordered by inclusion. Here, we use the notation $[n] = \{1, 2, \dots, n\}$, which is hopefully on its way to being a ubiquitous standard. By definition, B_n is the downset lattice of the antichain \mathcal{A}_n . Hence, B_n is graded. Level i of B_n contains

all subsets of $[n]$ with i elements. The Hasse diagram of B_n equals the n -dimensional hypercube (see Section 1.3 for the definition).

1.2 Poset Dimension

The dimension of posets is a classic parameter which is discussed extensively in Trotter's classic book [59]. Here is the definition:

Definition 1.1 ([17]). *A set $\mathcal{R} = \{L_1, L_2, \dots, L_k\}$ of linear extensions of \mathcal{P} is a realizer of \mathcal{P} if for every incomparable pair $x, y \in \mathcal{P}$, there are $L_i, L_j \in \mathcal{R}$ with $x < y$ in L_i and $x > y$ in L_j . The dimension $\dim(\mathcal{P})$ of \mathcal{P} is the minimum cardinality of a realizer.*

Equivalently, $\dim(\mathcal{P})$ can be defined as the minimum k for which there are linear extensions L_1, \dots, L_k such that

$$\mathcal{P} = \bigcap_{i=1}^k L_i,$$

where the intersection is taken over the sets of relations of the L_i .

Another characterization of dimension uses an embedding of \mathcal{P} into \mathbb{Z}^d . For two vectors \mathbf{v}, \mathbf{w} in a d -dimensional vectorspace V we have $\mathbf{v} < \mathbf{w}$ in the *dominance relation* if the i -th coordinate of \mathbf{v} is smaller than the i -th coordinate of \mathbf{w} , for every $i = 1, \dots, d$. A realizer $\mathcal{R} = \{L_1, \dots, L_d\}$ of a poset \mathcal{P} yields an embedding of \mathcal{P} into \mathbb{Z}^d such that the order relation of \mathcal{P} equals the dominance relation on \mathbb{Z}^d . Let us call such an embedding *proper*. It can be obtained by assigning to each $v \in \mathcal{P}$ the vector in \mathbb{Z}^d whose i -th coordinate equals the position of v in L_i , for every $i = 1, \dots, d$. Then the dimension of a poset \mathcal{P} is the minimal d such that \mathcal{P} has a proper embedding into \mathbb{Z}^d .

Suppose we have a set \mathcal{R} of linear extensions of a poset \mathcal{P} and want to check whether it is a realizer of \mathcal{P} . Using Definition 1.1, we would have to go through all incomparable pairs of elements. It turns out that it suffices to check a special type of incomparable pairs, the *critical pairs*. See Figure 1.2 for illustration.

Definition 1.2. *A critical pair of \mathcal{P} is an ordered pair (x, y) of elements of \mathcal{P} such that $\text{Pred}(x) \subseteq \text{Pred}(y)$ and $\text{Succ}(y) \subseteq \text{Succ}(x)$.*

The canonical order of a critical pair (x, y) in a linear extension L of \mathcal{P} is $x < y$. If $y < x$ in L , then L reverses the critical pair (x, y) . If L reverses some critical pair of \mathcal{P} , it is reversing.




Figure 1.2: Elements x and y form the critical pair (x, y) .

The following characterization of critical pairs follows right from the definition: The pair (x, y) is a critical pair of \mathcal{P} if the addition of $x < y$ to the relations of \mathcal{P} does not transitively force any other additional relation (or, equivalently, if $y < x$ cannot be forced by adding any other relation).

We will see in the proof of the lemma below that every poset which has an incomparable pair also contains a critical pair. That is, the only posets without critical pairs are chains.

Lemma 1.3 ([49]). *A set \mathcal{R} of linear extensions of \mathcal{P} is a realizer of \mathcal{P} if and only if for every critical pair (x, y) of \mathcal{P} there is a linear extension in \mathcal{R} which reverses (x, y) .*

Proof. The “only if” direction follows immediately from the definition of a realizer. For the other direction, it suffices to show that for every ordered pair (a, b) of incomparable elements of \mathcal{P} we can find $L_i \in \mathcal{R}$ with $b < a$ in L_i . Given (a, b) , we check if a has a predecessor which is not a predecessor of b . If so, let us call it a_1 . Now given an a_i with $i \geq 1$, we check if a_i has a predecessor a_{i+1} which is not a predecessor of b . Since \mathcal{P} is finite, we will eventually arrive at an a_k such that $\text{Pred}(a_k) \subseteq \text{Pred}(b)$.

Next we check if b has a successor b_1 which is not a successor of a_k . We iterate this until arriving at a b_ℓ such that $\text{Succ}(b_\ell) \subseteq \text{Succ}(a_k)$. Observe that we have $\text{Pred}(a_k) \subseteq \text{Pred}(b_1) \subseteq \text{Pred}(b_\ell)$. It follows that (a_k, b_ℓ) is a critical pair of \mathcal{P} . Since \mathcal{R} is a realizer, there is a linear extension $L \in \mathcal{R}$ with $b_\ell < a_k$ in L . Then we have $b < b_\ell < a_k < a$ in L , which is what we wanted to show. \square

It is easy to see that for the standard examples we have $\dim(S_n) = n$. The critical pairs of S_n are exactly the pairs (a_i, b_i) for $i = 1, \dots, n$. Hence a collection of linear extensions is a realizer of S_n if, for every i , there is a linear extension in which a_i is above b_i . But by transitivity we can have $a_i > b_i$ for at most one i in each linear extension of S_n .

It was shown in [37] (and is folklore by now) that the dimension of the Boolean lattice B_n is n . To see this, observe that B_n contains S_n as a

subposet, induced by the atoms and coatoms. Clearly, the dimension of a poset cannot be smaller than the dimension of one of its subposets. Hence we have $\dim(B_n) \geq n$.

To show that $\dim(B_n)$ is exactly n , we characterize the critical pairs of B_n . We extend the well-known result to some subposets of B_n which we use later. For an atom $a \in B_n$, we set $a^c = [n] \setminus \{a\}$.

Lemma 1.4. *Let \mathcal{P} be a subposet of B_n which is induced by a set of subsets of $[n]$ containing each atom a and each coatom a^c of B_n . Then the critical pairs of \mathcal{P} are exactly the n pairs (a, a^c) .*

Proof. Suppose two subsets S and T of $[n]$ form a critical pair of \mathcal{P} . Then they are incomparable, so $S \not\subseteq T$. Therefore S contains an atom a which is not contained in T . Now if $S \neq a$, then a would be a predecessor of S which is not a predecessor of T , contradicting that (S, T) is a critical pair. Hence $S = a$. Since $a \notin T$, we know that $T \subseteq a^c$. But if T is a proper subset of the coatom a^c , then a^c forms a successor of T which is not a successor of S , a contradiction. Thus we have $T = a^c$, and the critical pairs of \mathcal{P} are as claimed. \square

It is now easy to deduce that $\dim(B_n) = n$: In order to build a realizer $\{L_1, L_2, \dots, L_n\}$ of B_n , it suffices to choose as L_a a linear extension of B_n in which $a^c < a$, for $a = 1, \dots, n$.

In general, computing the dimension of a poset is a difficult problem. Yannakakis [65] showed that deciding whether the dimension of a poset \mathcal{P} is at least t for some fixed $t \geq 3$ is an NP-complete problem. The only special case which is easy to compute is dimension 2. The class of 2-dimensional posets is well understood and a number of characterizations are available (see [17], [2] or the overview in [43]).

Since 2-dimensional posets will appear a number of times in this thesis, we want to prove two characterizations here. We need the following notions: A poset \mathcal{Q} is a *conjugate* of a poset \mathcal{P} if their ground sets coincide and every pair of distinct elements is comparable in exactly one of the two posets. A linear extension L of \mathcal{P} is called *separating* if there are elements $u, v, w \in \mathcal{P}$ such that $u < v < w$ in L , $u < w$ in \mathcal{P} , but $v \parallel u, w$ in \mathcal{P} . If a linear extension is not separating, it is *non-separating*.

Theorem 1.5 ([17]). *Let \mathcal{P} be a poset. Then the following properties are equivalent:*

- (1) \mathcal{P} is 2-dimensional.
- (2) \mathcal{P} has a conjugate.

(3) \mathcal{P} has a non-separating linear extension.

Proof. To show that (1) and (2) are equivalent, we start with a realizer L_1, L_2 of \mathcal{P} . Then by definition, we have $\mathcal{P} = L_1 \cap L_2$, where the intersection is taken over the relations. Now $L_1 \cap L_2^*$ defines a partial order \mathcal{Q} which is a conjugate of \mathcal{P} . On the other hand, if \mathcal{Q} is a conjugate of \mathcal{P} , then it is not difficult to see that $\mathcal{P} \cup \mathcal{Q}$ and $\mathcal{P} \cup \mathcal{Q}^*$ are linear extensions of \mathcal{P} . Together they build a realizer of \mathcal{P} .

For the equivalence of (2) and (3), let \mathcal{Q} be a conjugate of \mathcal{P} and let L' be a linear extension of \mathcal{Q} . Build a linear extension L of \mathcal{P} with the generic algorithm, using L' as a priority list. If we have $v \parallel u, w \in \mathcal{P}$, then because \mathcal{Q} is transitive, we have either $v < u, w$ or $v > u, w$ in \mathcal{Q} . It follows that L is non-separating. Conversely, if L is a non-separating linear extension of \mathcal{P} , then consider the set of relations which are in L , but not in \mathcal{P} . Since L is non-separating, this set of relations is transitively closed, and thus defines a conjugate of \mathcal{P} . \square

If L and L' form a realizer of a 2-dimensional poset \mathcal{P} , then all incomparable pairs of \mathcal{P} need to be reversals between L and L' , that is, $\text{dist}(L, L') = \text{inc}(\mathcal{P})$. Obviously, the converse also holds. It follows that L' is the unique partner of L in a minimum realizer of \mathcal{P} . Not all linear extensions of \mathcal{P} have a partner at distance $\text{inc}(\mathcal{P})$, but we can characterize them using the arguments from the previous proof.

Corollary 1.6. *Let \mathcal{P} be a 2-dimensional poset and let L be a linear extension of \mathcal{P} . Then L is contained in a minimum realizer of \mathcal{P} exactly if it is non-reversing.*

Proof. Let L be non-reversing. Then the relations which are in L , but not in \mathcal{P} , define a conjugate \mathcal{Q} of \mathcal{P} . Now $L = \mathcal{P} \cup \mathcal{Q}$ and $\mathcal{P} \cup \mathcal{Q}^*$ form a realizer of \mathcal{P} . Conversely, suppose that $\{L, L'\}$ is a minimum realizer of \mathcal{P} , and let $u, v, w \in \mathcal{P}$ with $u < w$ and $v \parallel u, w$ in \mathcal{P} . If $u < v < w$ in L , then we have $v < u$ and $v > w$ in L' , which is a contradiction. Thus L has to be non-separating. \square

It follows from the equivalence between properties (1) and (2) of Theorem 1.5 that a poset is 2-dimensional exactly if its incomparability graph is a comparability graph, cf. [2]. Gallai [25] gave a complete list of the minimal forbidden subgraphs of comparability graphs.

A poset \mathcal{P} is called *t-irreducible* if it has dimension t , but deleting any element from \mathcal{P} makes the dimension drop. Gallai's characterization of comparability graphs via forbidden subgraphs yields a list of all 3-irreducible posets. The Chevron (cf. Figure 1.1) is one of the smallest 3-irreducible posets. See Trotter's book [59] for a discussion of *t-irreducible* posets.

1.3 Partial Cubes

In this section, we introduce the graph class of partial cubes and present basic properties and results that will be useful later. Partial cubes were first mentioned by Firsov [23] and Graham and Pollak [29], and first characterized by Djoković [15]. Several papers with other characterizations and generalizations followed, see [30], [31], [64], [9] and more recently [47]. The class of partial cubes includes various important graph classes, e.g. trees, hyperplane arrangement graphs and median graphs. For overviews providing more references and connections to other fields, see [31], [33] and [12].

Definition 1.7. *The Hamming distance between two strings is the number of coordinates in which they differ. The r -dimensional hypercube \mathcal{Q}_r is the graph on all 0-1-strings of length r , in which two strings are adjacent if their Hamming distance is 1.*

Let G be a graph and H a subgraph of G . Then H is a convex subgraph of G if for every $u, v \in V(H)$, every shortest u - v -path in G is fully contained in H . Moreover, H is an isometric subgraph of G if for every $u, v \in V(H)$, at least one shortest u - v -path in G is fully contained in H .

A partial cube is an isometric subgraph of a hypercube. A Hamming labeling of a partial cube $G = (V, E)$ is a labeling of V with 0-1-strings such that the distance of two vertices in G equals the Hamming distance of the corresponding strings.

By definition, each partial cube has a Hamming labeling. The following proposition summarizes basic properties of hypercubes, see e.g. [33].

Proposition 1.8. *Let \mathcal{Q}_r be a hypercube. Then*

- (i) \mathcal{Q}_r is connected, bipartite, r -regular and has diameter r .
- (ii) $|V(\mathcal{Q}_r)| = 2^r$ and $|E(\mathcal{Q}_r)| = r2^{r-1}$.
- (iii) For any pair of vertices $u, v \in \mathcal{Q}_r$, the subgraph induced by the interval $I(u, v)$ is a hypercube whose dimension equals the Hamming distance of u and v .
- (iv) If G is a subgraph of \mathcal{Q}_r , then $|E(G)| \leq \frac{1}{2}|V(G)| \cdot \log_2(|V(G)|)$, with equality holding if and only if $G = \mathcal{Q}_s$ for some $s \leq r$.

The fact that \mathcal{Q}_r is bipartite carries over to all partial cubes. A bipartition is given by the parity of the number of 1s in a Hamming labeling of the vertices.

Definition 1.9 ([15]). *For an edge uv of a graph G , let W_{uv} be the set of vertices of G that are closer to u than to v in G . Two edges $e = xy$ and f of G stand in relation θ if f joins a vertex in W_{xy} with a vertex in W_{yx} . The relation θ is called the Djoković-Winkler relation.*

The definition of θ given above is due to Djoković [15]. Winkler [64] defined another relation θ' on the edges of a graph, in which two edges $e = vw$ and $f = xy$ are related if $d(x, v) + d(y, w) \neq d(x, w) + d(y, v)$, where d denotes the graph distance. It is not difficult to see that θ' and θ coincide on bipartite graphs, hence in particular on partial cubes. Since we will only use these two relations for partial cubes, we do not distinguish between them, and call $\theta = \theta'$ the Djoković-Winkler relation, following [33].

We will use the following characterizations of partial cubes:

Theorem 1.10 ([15], [64]). *Let $G = (V, E)$ be a connected graph. The following statements are equivalent:*

- (1) *G is a partial cube.*
- (2) *G is bipartite and for every edge uv of G , the sets W_{uv} and W_{vu} induce convex subgraphs of G .*
- (3) *G is bipartite and θ is an equivalence relation on E .*

Let $G = (V, E)$ be a partial cube. By the above theorem, the Djoković-Winkler relation is an equivalence relation on the edges of G . This yields a partition of E into *Djoković-Winkler classes* $\theta_1, \dots, \theta_r$. If $e \in \theta_k$, we write $\theta(e) = \theta_k$, abusing the notation. Here is a crucial lemma about the Djoković-Winkler classes:

Lemma 1.11. *Let G be a partial cube given with a Hamming labeling. Then the coordinates of the Hamming labeling are in bijection with the Djoković-Winkler classes of G . Moreover, the Hamming labeling is unique up to permutation of coordinates and addition of redundant coordinates.*

Proof. Consider a shortest path P in \mathcal{Q}_r between two vertices x and y which coincide in the i -th coordinate. It is easy to see that all vertices on P coincide with x and y in the i -th coordinate. Now let $e = xy$ be an edge of \mathcal{Q}_r such that x and y differ in the i -th coordinate. It follows that all vertices coinciding with x in the i -th coordinate are closer to x than to y , and vice versa. In other words, W_{xy} consists of all vertices which coincide with x in the i -th position, and W_{yx} consists of all vertices which coincide with y in the i -th position. The edges joining a vertex in W_{xy} to a vertex in W_{yx} are therefore exactly the edges corresponding to a change of the i -th

coordinate. This yields a bijection between the coordinates of the Hamming labeling and the Djoković-Winkler classes of G .

Assume that G is an isometric subgraph of Q_r , and let P be a shortest path in G . Then P is also a shortest path in Q_r . Thus the observation of the last paragraph carries over to P : If the endvertices of P coincide in a coordinate, then all vertices on P do. Consequently, if $e = xy$ is an edge of G such that the Hamming labeling of x and y differs in the i -th coordinate, then W_{xy} consists of all vertices which coincide with x in the i -th position, and W_{yx} consists of all vertices which coincide with y in the i -th position. Also, if $\theta(e) = \theta_k$, then the edges in θ_k are exactly the edges connecting two vertices differing in the i -th position.

Since the partition of the edges into Djoković-Winkler classes is clearly unique, it follows that the Hamming labeling is unique up to permutation of coordinates and addition of redundant coordinates. \square

The hypercube Q_r is the Cartesian product of r copies of K_2 . In [64] it was shown that the isometric embedding of a graph into a product of complete graphs is essentially unique. This generalizes the uniqueness of the Hamming labeling shown above.

Using Winkler's definition, it is easy to see that the Djoković-Winkler relation of a bipartite graph can be computed in polynomial time. With characterization (iii) of Theorem 1.10, this yields a recognition algorithm for partial cubes. The currently fastest recognition algorithm has been given by Eppstein [18].

Theorem 1.12 ([18]). *Let G be an undirected graph on n vertices. Then there is an algorithm which checks in time $O(n^2)$ whether G is a partial cube. The algorithm constructs the Djoković-Winkler classes of G , and uses them to obtain a valid Hamming labeling of G .*

Djoković [15] defines the *isometric dimension* $\dim_I(G)$ of a graph G as the smallest r such that G has an isometric embedding into the hypercube Q_r . He proved the following:

Theorem 1.13 ([15]). *If G is a partial cube, then $\dim_I(G)$ equals the number of Djoković-Winkler classes of G .*

With Eppstein's algorithm, it follows that the isometric dimension of a partial cube can be computed in quadratic time.

The following easy lemma is probably known to anybody interested in partial cubes, though the publications we are aware of only mention one direction. We give a proof for completeness.

Lemma 1.14. *Let P be a path in a partial cube G . Then P is a shortest path in G exactly if no two distinct edges on P belong to the same Djoković-Winkler class.*

Proof. Consider a Hamming labeling of G . By Lemma 1.11, the Djoković-Winkler classes are in bijection with the coordinates of the Hamming labeling. Thus it suffices to show that P is a shortest path in G exactly if no two distinct edges on P correspond to the same coordinate.

Let x and y be the two endvertices of P , and observe that every x - y -path needs to use at least one edge for every coordinate in which x and y differ. Let S be the set of these coordinates. If G is the full hypercube, then it is clearly true that P is a shortest path in G if and only if it contains exactly one edge for every coordinate in S . This carries over to an arbitrary partial cube G , since a shortest path in G cannot be shorter than $|S|$. \square

We call a cycle in a graph G an *isometric cycle* if it is an isometric subgraph of G . *Convex cycles* are defined analogously. Note that every convex cycle is an isometric cycle, but the converse does not hold.

We can now characterize the isometric cycles in partial cubes. Again, this characterization was probably known beforehand, but we are not aware of a publication.

Theorem 1.15. *Let G be a partial cube, and let C be a cycle in G . Then C is an isometric cycle exactly if, for any two edges e, f on C , the following two conditions are equivalent:*

(i) e and f are opposite edges of C .

(ii) $\theta(e) = \theta(f)$.

Proof. Assume that C is an isometric cycle of G . Let C contain two edges $e = xy$ and $f = uv$ such that $x \in W_{uv}$ as in Figure 1.3. If e and f are opposite edges in C , then $y \in W_{vu}$. Thus $\theta(e) = \theta(f)$. Now suppose that e and f are not opposite in C . Then there is a shortest y - v -path on C which contains e and f . By Lemma 1.14, the Djoković-Winkler classes of the edges on this path are pairwise different.

For the other direction, assume that (i) and (ii) are equivalent for C . Let x and v be vertices on C . Then there is an x - v -path on C such that no two distinct edges of P belong to the same Djoković-Winkler class. By Lemma 1.14, this path is a shortest path. Thus C is an isometric cycle. \square




Figure 1.3: Opposite and non-opposite edges e, f in a cycle C .

1.4 Modular Decomposition

In this section, we state results on the modular decomposition and transitive orientation of graphs from Gallai's classic paper [25]. We keep the presentation brief, only presenting the results (without proofs) which we need later in Section 3.4. We mainly follow Gallai's notation and the notation used in the translation in [39].

Definition 1.16. *An undirected graph $G = (V, E)$ is called a comparability graph if it is possible to assign a transitive orientation to the edges in E , such that whenever (u, v) and (v, w) are arcs, also (u, w) is an arc.*

It is easy to see that G is a comparability graph exactly if there is a poset \mathcal{P} such that $G = \text{Comp}(\mathcal{P})$, and every transitive orientation of G corresponds to such a poset.

Now let $G = (V, E)$ be an arbitrary graph. If $uv \in E$ and $uw \in E$, but $vw \notin E$, then orienting one of the two edges uv and uw forces the orientation of the other in a transitive orientation: The edges uv and uw must either both be directed towards their common endvertex, or both away from it. We say that uv forces uw . The forcing relation is reflexive and symmetric, and its transitive closure yields a partition of E . Let us call the classes of this partition the *forcing classes* of G . Then the choice of an orientation for one edge in a forcing class F forces the orientation of all other edges in F .

Two subsets V_1, V_2 of the vertices of a graph G will be called *completely adjacent* if all possible edges between V_1 and V_2 are present in G .

Theorem 1.17 ([25]). *Let $G = (V, E)$ be an undirected graph on at least two vertices. Then exactly one of the following cases holds:*

- (||) G is disconnected and has connected components G_1, G_2, \dots, G_t , $t \geq 2$. Then the forcing classes of G_1, G_2, \dots, G_t are exactly the forcing classes of G .
- (S) G^c is disconnected (thus G is connected) and has connected components $G_1^c, G_2^c, \dots, G_t^c$, $t \geq 2$. Let $M_i = V(G_i^c)$. Then for each pair i, j with $i \neq j$, it holds that M_i and M_j are completely adjacent in G , and the M_i - M_j -edges form one forcing class F_{ij} of G . The forcing classes of G different from the F_{ij} are exactly the forcing classes of all the $(G_i^c)^c$.
- (P) If G and G^c are both connected and have at least two vertices, then there exists a unique proper partition M_1, \dots, M_t , $t \geq 2$, of V with the following properties:
 - (a) For each pair $i \neq j$, if there is an M_i - M_j -edge in G , then M_i and M_j are completely adjacent.
 - (b) The edges whose endpoints are in different M_i form a unique forcing class F . Every vertex of G is incident with at least one edge from F .
 - (c) The forcing classes different from F are exactly the forcing classes of the graphs $G[M_i]$, $i = 1, \dots, t$.
 - (d) The partition of V into M_1, \dots, M_t is not a refinement of another partition with properties (a)–(c).

In each of the three cases of Theorem 1.17, we obtain a unique proper partition $\{M_1, \dots, M_t\}$ of V , the *canonical partition* of G . The M_i are the *vertex classes* of G . An edge of G with both endpoints in some M_i is called an *inner edge*, and the other edges are called *outer edges*. The edges of a given forcing class are either all inner or all outer edges. By definition, the canonical partition of the one-vertex graph K_1 is K_1 itself.

Consider the subgraphs $G[M_i]$ of G given by the canonical partition. We can now apply Theorem 1.17 to the $G[M_i]$. The vertex classes of the canonical partition of the $G[M_i]$ are the *second order vertex classes* of G . Recursively, we obtain vertex classes of order 3, 4, \dots of G . Since G is finite, there is an $s \in \mathbb{N}$ such that all vertex classes of order s are singletons.

We now present a recursion-free characterization of the vertex classes. It uses the concept of modules. Gallai uses the term *homogeneous set* instead of module.

Definition 1.18. Let $G = (V, E)$ be a graph and $M \subseteq V$. Then M is a module of G if every vertex of G which is not in M is either adjacent to all

vertices in M , or adjacent to no vertex in M . A module M is called *strong* if every other module of G is either disjoint to M or comparable to M under inclusion. A (strong) module M is called *quasi-maximal* if there is no other (strong) module M' with $M \subset M' \subset G$.

Clearly, all vertex classes with respect to the canonical partition form modules of G . Furthermore, we have:

Lemma 1.19. *Let M be a module and F a forcing class of a graph G . If F has one edge in M , then all edges of F are in M , and F is a forcing class of M . Conversely, the set of vertices induced by F form a module of G .*

Theorem 1.20. *If G has at least two vertices, the quasi-maximal modules of G are the vertex classes of the canonical partition of G . If G is not empty, then the collection of all strong modules of G is identical to the collection of all vertex classes of all orders of G .*

The *modular decomposition* is the collection of all strong modules of G , ordered by inclusion. The modular decomposition of G forms a tree, with the whole vertex set V as root and singletons as leaves. It is called the *modular decomposition tree*.

The number of modules of G can be exponential in general. However, the modular decomposition tree has at most $2|V| - 1$ nodes (cf. [42]), which can be proved by a straightforward induction. This makes it an efficient structure for storing G , which we will exploit in Chapter 3.

Definition 1.21. *The partition graph $G^\#$ of G is defined as follows: The vertices of $G^\#$ are the vertex classes of the canonical partition of G . Two vertices of $G^\#$ are adjacent exactly if there exists an edge between the corresponding vertex classes in G .*

Theorem 1.22. *Let G be a graph and $G^\#$ the partition graph of G . Then we have:*

- (1) *If G is not connected, then $G^\#$ has no edge.*
- (2) *If G^c is not connected, then $G^\#$ is a complete graph, and each edge of $G^\#$ forms its own forcing class in $G^\#$.*
- (3) *If G and G^c are connected and have at least two vertices, then $G^\#$ and $(G^c)^\#$ are connected. All edges of $G^\#$ belong to the same forcing class. Also, $G^\#$ contains an induced path on four vertices.*

A graph with exactly two transitive orientations (which are then the reverse of one another) is called *uniquely partially orderable*. The transition graph $G^\#$ in case (3) of the above theorem is uniquely partially orderable, because all edges belong to the same forcing class.

The following theorem characterizes all transitive orientations of a given graph.

Theorem 1.23. *Let G be a non-empty graph and let M^{r-1} be a module of order $r - 1$ of G ($r \geq 1$). Consider the canonical partition $\{M_1^r, \dots, M_k^r\}$ of M^{r-1} , the partition graph $G[M^{r-1}]^\#$, and two vertex classes M_i^r, M_j^r which are completely adjacent in G .*

1. *If G is a comparability graph then, for each transitive orientation of G , the M_i^r - M_j^r -edges are either all directed from M_i^r to M_j^r or all directed from M_j^r to M_i^r . Moreover, if we orient the edge $M_i^r M_j^r$ of $G[M^{r-1}]^\#$ in the same way, and repeat this for each edge of $G[M^{r-1}]^\#$, we obtain a transitive orientation of $G[M^{r-1}]^\#$.*
2. *Conversely, assume that all partition graphs $G[M^{r-1}]^\#$ are comparability graphs and choose a transitive orientation for each of them. For each vertex class M^{r-1} of any order $r - 1$, and for any two vertices M_i^r, M_j^r of the partition graph $G[M^{r-1}]^\#$, let us assign to all M_i^r - M_j^r -edges of G the orientation of the edge $M_i^r M_j^r$ in $G[M^{r-1}]^\#$. Then we obtain a transitive orientation of G .*

Gallai does not explicitly mention an algorithm to construct a transitive orientation of a comparability graph, but the above result is constructive and gives an outline of a possible algorithm. We will use this in Section 3.4. Gallai's paper [25] also gives a characterization of comparability graphs by means of forbidden subgraphs.

There has been a lot of research on modular decomposition and transitive orientations of a comparability graph after Gallai's paper. Algorithmic aspects were treated in Golumbic's book [28] and in Möhring's article [42].

A linear time algorithm for finding the modular decomposition of a graph has been given by McConnell and Spinrad [40]. The algorithm can also be extended to find a transitive orientation in linear time if the given graph is a comparability graph. However, it is very technical and hard to understand. In [41], they give a simpler algorithm which fulfills both tasks in time $O(n + m \log n)$.

We want to transfer some of the terms defined for comparability graphs to posets. The term of a module of a poset will appear frequently in this thesis:

Definition 1.24. Let \mathcal{P} be a poset. We call a subset M of the elements a module of \mathcal{P} if the elements in M cannot be distinguished from the outside. More precisely, M is a module if for any $x \in \mathcal{P} \setminus M$ we have either $x > m$ for all $m \in M$ or $x < m$ for all $m \in M$ or $x \parallel m$ for all $m \in M$.

If $M = \{x, y\}$ with $x \parallel y$, we say that xy is a twin of \mathcal{P} .

For the following, consider the comparability graph $\text{Comp}(\mathcal{P})$ of \mathcal{P} and the canonical partition M_1, \dots, M_t of $\text{Comp}(\mathcal{P})$, given by Theorem 1.17. By Theorem 1.23, each M_i is a module of \mathcal{P} . Then \mathcal{P} induces a transitive orientation of the partition graph $\text{Comp}(\mathcal{P})^\#$. Observe that if case (S) of Theorem 1.23 applies to $\text{Comp}(\mathcal{P})$, then this transitive orientation is a linear order of the modules M_1, \dots, M_t .

Definition 1.25. Let \mathcal{P} be a poset and let M_1, \dots, M_t be the canonical partition of $\text{Comp}(\mathcal{P})$. In view of the three cases of Theorem 1.17, we define:

If case (\parallel) applies to $\text{Comp}(\mathcal{P})$, then the M_i are parallel modules of \mathcal{P} , and \mathcal{P} is the parallel composition of $\mathcal{P}[M_1], \dots, \mathcal{P}[M_t]$.

If case (S) applies to $\text{Comp}(\mathcal{P})$, then the M_i are series modules of \mathcal{P} , and \mathcal{P} is the series composition of $\mathcal{P}[M_1], \dots, \mathcal{P}[M_t]$. Furthermore, if $M_1 < M_2 < \dots < M_t$ is the transitive orientation of $\text{Comp}(\mathcal{P})^\#$ induced by \mathcal{P} , then we write $\mathcal{P} = \mathcal{P}[M_1] * \dots * \mathcal{P}[M_t]$.

If case (P) applies to $\text{Comp}(\mathcal{P})$, then the M_i are prime modules of \mathcal{P} , and \mathcal{P} is the prime composition of $\mathcal{P}[M_1], \dots, \mathcal{P}[M_t]$.

Observe that $\mathcal{P} = \mathcal{P}_1 * \mathcal{P}_2$ if \mathcal{P} arises from placing \mathcal{P}_2 “on top of” \mathcal{P}_1 , that is, the relations of \mathcal{P} are obtained by keeping the relations within \mathcal{P}_1 and \mathcal{P}_2 unchanged, and setting $v_1 < v_2$ whenever $v_1 \in \mathcal{P}_1$ and $v_2 \in \mathcal{P}_2$.

As last result of this chapter, we present a characterization of comparability invariants (see Section 1.1 for the definition). Theorem 1.23 implies a convenient method to show that a given poset parameter is a comparability invariant. It was proved explicitly in [16].

Before stating it, we need some notation: Let \mathcal{P} and \mathcal{Q} be posets and let $x \in \mathcal{P}$. We denote by $\mathcal{P}_x^\mathcal{Q}$ the poset resulting from replacing x by \mathcal{Q} in \mathcal{P} . That is, $\mathcal{P}_x^\mathcal{Q}$ has $\mathcal{P} - x \cup \mathcal{Q}$ as ground set. The relation in $\mathcal{P}_x^\mathcal{Q}$ between two elements of $\mathcal{P} - x$ is the same as in \mathcal{P} , and the relation in $\mathcal{P}_x^\mathcal{Q}$ between an element v in $\mathcal{P} - x$ and an element in \mathcal{Q} is the same as the relation between v and x in \mathcal{P} . Then $\mathcal{P}_x^\mathcal{Q}$ arises from \mathcal{P} by inserting \mathcal{Q} for x in \mathcal{P} .

Theorem 1.26 ([16]). A poset parameter ρ is a comparability invariant exactly if it remains unchanged whenever a subposet induced by a module is replaced by its dual, that is, if $\rho(\mathcal{P}_x^\mathcal{Q}) = \rho(\mathcal{P}_x^{\mathcal{Q}^*})$ for all choices of posets \mathcal{P} , \mathcal{Q} and elements $x \in \mathcal{P}$.

Chapter 2

Properties of Linear Extension Graphs

Linear extension graphs play a central role in this thesis. In this chapter we explore their beautiful structure and exhibit connections to the underlying posets.

A linear extension graph is a graph on the set of linear extensions of a poset, where two linear extensions are adjacent exactly if they differ in one adjacent swap of elements. This yields a coloring of the edges with the corresponding swaps. Figure 2.1 shows the Chevron with its linear extension graph.




Figure 2.1: The Chevron and its linear extension graph with a swap coloring.

Linear extension graphs were originally defined by Pruesse and Ruskey in [48]. The first line of research on linear extension graphs was concerned with the existence of a Hamilton path; see also [52] and [61]. There has been subsequent research on other structural properties of $G(\mathcal{P})$, see [51], [50], [44], [45] and [22].

In Section 2.1 we present basic properties and previous results of linear extension graphs, and place them into a larger context. In Section 2.2 we concentrate on properties of the edge classes induced by the swap colors. This prepares Section 2.3, in which we will characterize which pairs of swap colors share an element. The last section of this chapter deals with the question which modifications of the poset \mathcal{P} leave the linear extension graph $G(\mathcal{P})$ invariant.

2.1 Context and Previous Results

Before pointing out the connection of linear extension graphs to larger known graph classes, we want to make our notation precise and present some basic features of linear extension graphs.

Definition 2.1. *The linear extension graph $G(\mathcal{P}) = (V, E)$ of a poset \mathcal{P} has as vertices the linear extensions of \mathcal{P} , with two of them being adjacent if they only differ in one adjacent transposition, or swap, of elements.*

Equivalently, for an edge $LL' \in E$ there is exactly one pair $xy \in \text{Inc}(\mathcal{P})$ with $x < y$ in L and $x > y$ in L' . This pair is the swap color of LL' . The set of all edges with the same swap color forms a color class of $G(\mathcal{P})$.

A graph G is a linear extension graph if there is an underlying poset \mathcal{P} such that $G = G(\mathcal{P})$.

Recall that a reversal between L and L' is a pair of elements of \mathcal{P} appearing in different orders in L and L' . Thus, each swap color corresponds to a reversal. By definition, two linear extensions of \mathcal{P} are adjacent in $G(\mathcal{P})$ exactly if there is only one reversal between them, i.e., if the distance between them is 1. The very first observation about linear extension graphs we want to prove is that this generalizes to higher distances: The distance between two linear extensions L, L' of \mathcal{P} equals the graph distance between the two corresponding vertices in $G(\mathcal{P})$.

Informally, this holds because changing L into L' is just the same as sorting the elements of \mathcal{P} into the linear order L' , starting with the linear order L . Thus we may, for example, use a selection sort algorithm (see e.g. [54]) to show that we never have to make a superfluous reversal in the sorting process.

For completeness, we give an explicit proof in the following lemma. A similar version of this lemma appeared in [44]. Note that we identify a vertex of $G(\mathcal{P})$ with the corresponding linear extension of \mathcal{P} .

Lemma 2.2. *Let L and L' be two linear extensions of a poset \mathcal{P} . Let T be a shortest L - L' -path in $G(\mathcal{P})$, and let S be the set of swap colors appearing on T . Then the swap colors in S are in bijection with the reversals between L and L' , and each swap color in S appears only once on T .*

Proof. An L - L' -path T in $G(\mathcal{P})$ with swap color set S corresponds to a sequence of swaps modifying L into L' . To get from L to L' by swaps, we need to swap every reversal between L and L' at least once. We will describe a process modifying L into L' which swaps only reversals, and every reversal only once. This corresponds to a shortest L - L' -path in $G(\mathcal{P})$, and it follows that every shortest L - L' -path in $G(\mathcal{P})$ has the desired properties.

Let $L = u_1 u_2 \dots u_n$ and $L' = v_1 v_2 \dots v_n$. Start with finding the position of v_1 in L , say, $v_1 = u_j$. Then u_j is a minimal element of \mathcal{P} , so we can swap it with u_{j-1} , then u_{j-2} , and so on, until we obtain a linear extension of \mathcal{P} which coincides with L' in the first element.

Now assume that we are given $L^i = u_1 u_2 \dots u_n$ which coincides with L' in the first i elements. Find the position of v_{i+1} in L^i , say, $v_{i+1} = u_j$. It follows that $j > i$, and u_j is a minimal element of $\mathcal{P} - \{u_1, \dots, u_i\}$. We can thus swap u_j with u_{j-1} , then u_{j-2} , and so on, until we arrive at a linear extension L^{i+1} of \mathcal{P} which coincides with L' in the first $i + 1$ elements. Inductively, we obtain $L^n = L'$.

Consider a pair of elements appearing in the same order in L and L' , that is, a pair $x, y \in \mathcal{P}$ such that y has higher u -index and higher v -index than x . Then x and y are never swapped in our process, and hence we swap only reversals. Also, all of our swaps take an element which appears in L above some element with higher v -index, and swap it below that element. Therefore no pair of elements is swapped twice. This means that each reversal is swapped exactly once. \square

Let $xy \in \text{Inc}(\mathcal{P})$ and set $G = G(\mathcal{P})$. We denote by W_{xy} the set of linear extensions of \mathcal{P} in which $x < y$, and by W_{yx} the set of linear extensions of \mathcal{P} in which $y < x$. Then the edges of swap color xy are exactly the edges connecting a linear extension in W_{xy} with a linear extension in W_{yx} . On the other hand, every path connecting a linear extension in W_{xy} with a linear extension in W_{yx} has to pass an edge with swap color xy . It follows that each color class is an edge cut of G .

Observe that the proof of the lemma above yields that each linear extension graph is connected. Now set $G_{xy} = G[W_{xy}]$ and $G_{yx} = G[W_{yx}]$.

Then G_{xy} is the linear extension graph of $\mathcal{P} \cup (x < y)$, and G_{yx} is the linear extension graph of $\mathcal{P} \cup (x > y)$. Therefore G_{xy} and G_{yx} are connected. Thus every color class cuts the graph G into exactly two components.

Lemma 2.3 ([51]). *Let \mathcal{P} be a poset and $xy \in \text{Inc}(\mathcal{P})$. Then G_{xy} and G_{yx} are convex subgraphs of $G(\mathcal{P})$.*

Proof. Let $L, L' \in W_{xy}$. We need to show that an arbitrary shortest L - L' -path T in $G(\mathcal{P})$ is contained in W_{xy} . By Lemma 2.2, all swap colors appearing on T are reversals between L and L' . Hence xy does not appear as a swap color on T . Therefore T lies fully in W_{xy} . \square

The above lemma is the crucial tool embedding linear extension graphs into the much larger class of partial cubes.

Theorem 2.4. *Linear extension graphs are partial cubes, and the color classes equal the Djoković-Winkler classes.*

Proof. We use the characterization from Theorem 1.10. Let G be a linear extension graph, and let LL' be an edge in G . Suppose L and L' differ by the swap of the elements x and y , and $x < y$ in L . By Lemma 2.2, we have $W_{LL'} = W_{xy}$, that is, the set of linear extensions which are closer to L than to L' in G is exactly the set of linear extensions in which $x < y$. Hence it follows from Lemma 2.3 that G is a partial cube.

The definition of the Djoković-Winkler relation and $W_{LL'} = W_{xy}$ yield that the color classes of G equal the Djoković-Winkler classes of G . \square

For the following, we want to fix some more notation concerning the swap colors of linear extension graphs. In view of the above theorem we use the style of the partial cube notation.

Definition 2.5. *Let $G = G(\mathcal{P})$ be a linear extension graph. We denote the color classes of $G(\mathcal{P})$ by $\theta_1, \theta_2, \dots, \theta_r$. The swap partition $\Theta(G)$ is the partition of $E(G)$ into the color classes $\theta_i, i = 1, \dots, r$.*

A swap coloring of G is a bijection c between $\Theta(G)$ and $\text{Inc}(\mathcal{P})$ which assigns to every color class the swap color of its edges. We usually denote the reverse mapping of c by θ . That is, if $\theta_i \in \Theta(G)$, and $c(\theta_i) = xy \in \text{Inc}(\mathcal{P})$, we write $\theta(xy) = \theta_i$. We extend this notation to the edges of G , i.e., if the swap color of e is xy , we write $c(e) = xy$, and if $e \in \theta_i$, we write $\theta(e) = \theta_i$.

If a graph G is a linear extension graph and we need to specify which underlying poset a swap coloring of G refers to, we say that c is a swap coloring of G with respect to \mathcal{P} .

If G is a linear extension graph, and thus a partial cube, then it has a unique partition of its edges into Djoković-Winkler classes. Hence it follows

from Theorem 2.4 that G has a unique swap partition $\Theta(G)$. Furthermore, as a partial cube, G has a Hamming labeling. By Lemma 1.11, this Hamming labeling is essentially unique, and partitioning the edges of G according to its coordinates again yields the swap partition $\Theta(G)$.

With Theorem 1.15, the following result now follows directly from Theorem 2.4.

Corollary 2.6 ([44]). *Let G be a linear extension graph with swap coloring c , and let C be a cycle in G . Then C is an isometric cycle exactly if for any two edges e, f on C , the following two conditions are equivalent:*

(i) e and f are opposite on C .

(ii) $c(e) = c(f)$. □

The fact that linear extension graphs are partial cubes is the background setting that we will use most in this thesis, but it is by far not the only larger context in which linear extension graphs can be viewed. In the remainder of this section, we want to point out some other connections. We only hint at the rich fields behind them. More connections and references can be found in Reuter's article [51].

The linear extension graph of the antichain, $G(\mathcal{A}_n)$, is the 1-skeleton of the *permutahedron* Π_{n-1} , see e.g. Ziegler's book on polytopes [66]. The permutahedron is a well-known polytope which is defined as the convex hull of all vectors that are obtained by permuting the coordinates of the vector $(1 \ 2 \ \dots \ n)^t$. In the following, we use the term permutahedron, and the notation Perm_n , to denote the 1-skeleton of Π_{n-1} .

The permutahedron Perm_n is also a hyperplane arrangement graph: The set of hyperplanes $H_{i,j} = \{x \in \mathbb{R}^n : x_i = x_j\}$ for $1 \leq i < j \leq n$ is called a *braid arrangement*. It is not difficult to see that Perm_n is isomorphic to the graph on all regions of the braid arrangements, where two regions are adjacent if they are separated by only one hyperplane. For an introduction to hyperplane arrangements, see e.g. Stanley's recent overview [56].

The braid arrangement is a Coxeter arrangement of type A_{n-1} . See the overview [24] by Fomin and Reading for an introduction to reflection groups and Coxeter arrangements. Hyperplane arrangement graphs are also the tope graphs of oriented matroids, see e.g. the book [4] by Björner et al..

A fact that we will use in the following is that hyperplane arrangement graphs are also partial cubes [46], thus, they interpolate between linear extension graphs and partial cubes.

In the last paragraphs we only considered $G(\mathcal{A}_n) = \text{Perm}_n$. Now let \mathcal{P} be an arbitrary poset on n elements, and consider the regions of the braid

arrangement again. For each pair $i < j$ in \mathcal{P} , let us restrict the hyperplane arrangement to the *halfspace* induced by $H_{i,j}$ in which $x_i < x_j$. Then $G(\mathcal{P})$ is the region graph of the modified braid arrangement. If we restrict the modified braid arrangement to the box in which all coordinates of vectors are at most 1, we obtain Stanley's *order polytope*, see [55]. This additional restriction does not change the region graph. It also follows from Stanley's results that $G(\mathcal{P})$ is the dual graph of the canonical triangulation of the order polytope.

Let us consider what happens to Perm_n if we restrict the braid arrangement to halfspaces. Since Perm_n is a hyperplane arrangement graph, each edge e of Perm_n corresponds to a hyperplane $H_{i,j}$. Equivalently, e has swap color ij . Restricting Perm_n to the vertices induced by one halfspace of $H_{i,j}$ is the same as restricting it to one side of the edge cut $\theta(ij)$, thus to W_{ij} or W_{ji} .

This correspondence sheds new light on the notion of *convex* subgraphs in hyperplane arrangement graphs, since we can now see the close connection to geometric convexity. We formulate the following lemma for the larger class of partial cubes. For an adjacent pair x, y of vertices of a partial cube G , we call the subgraph $G[W_{xy}]$ a *halfspace* of G , cf. [45].

Lemma 2.7 ([45]). *Let G be a partial cube. A subgraph of G is convex if and only if it is an intersection of halfspaces.* □

With this lemma, it can now be seen that linear extension graphs are in bijection with convex subgraphs of the permutahedron: If we start with the full graph Perm_n and an n -element poset \mathcal{P} , we take each comparable pair in \mathcal{P} and delete the halfspace of Perm_n where the corresponding relation is violated. This yields a convex subgraph of Perm_n . Conversely, starting with a convex subgraph we may consider it as intersection of halfspaces, and this yields the relations of the corresponding poset.

In fact, something even stronger is true. The convex subgraphs of Perm_n , ordered by inclusion, form a lattice $\text{Conv}(\text{Perm}_n)$. Now let us consider the set of all extensions of an n -element poset \mathcal{P} . If we order it by inclusion of the relations, it forms a poset $\text{Ext}(\mathcal{P})$. Let us add an artificial global maximum to it, corresponding to the complete relation $\mathcal{P} \times \mathcal{P}$, and denote the result by $\text{Ext}(\mathcal{P})^+$.

The following fundamental correspondance was first proved by Feldman in [19], and later rediscovered by Björner and Wachs [5] and by Reuter [51].

Theorem 2.8. *For an n -element poset \mathcal{P} , the lattice $\text{Conv}(\text{Perm}_n)$ is isomorphic to the dual of $\text{Ext}(\mathcal{P})^+$.* □

We want to mention one more fundamental property of linear extension graphs in this section:

Theorem 2.9 ([44]). *The cycle space of a linear extension graph is generated by 4-cycles and 6-cycles.*

Proof Outline. It is not difficult to show that the cycle space of a graph G is generated by the isometric cycles of G . In fact, it holds that every non-isometric cycle is a sum of *shorter* isometric cycles. Now let G be a linear extension graph. It can be proved that each isometric cycle of G is the sum of 4-cycles and 6-cycles, by induction on the length of the cycle.

The first crucial step is to show that for each cycle C of length at least 8 in G , there are two consecutive edges of C which are contained in a common isometric 4-cycle C' , or three consecutive edges of C which are contained in a common isometric 6-cycle C' . This can be deduced from Lemma 2.16.

Now let C be an isometric cycle of G . By induction we may assume that C and $C + C'$ have the same length. The second crucial step is to apply Corollary 2.6, which says that opposite edges of C' have the same swap color. But then, there are two opposite edges of $C + C'$ which do not have the same swap color. It follows by Corollary 2.6 that $C + C'$ is not an isometric cycle. Thus it is a sum of shorter isometric cycles, and we are done by induction. \square

The reader is invited to go back to Figure 2.1 and check that every cycle in $G(\mathcal{P})$ is composed of 4- and 6-cycles. It can also be observed that every 4-cycle in $G(\mathcal{P})$ corresponds to two disjoint swaps, and every convex 6-cycle in $G(\mathcal{P})$ corresponds to the permutations of the elements of an antichain \mathcal{A}_3 in \mathcal{P} . It will follow from Lemma 2.16 that this holds in general.

2.2 Color Classes

The color classes play a very important role for our analysis of linear extension graphs. In this section, we analyze the relation of two color classes in a linear extension graph. However, we start with looking at the relation of two edges in the same color class.

Proposition 2.10. *Let $G = G(\mathcal{P})$ be a linear extension graph with swap coloring c , and let $e, f \in E(G)$. Then we have $c(e) = c(f)$ exactly if there is a sequence $e_1 e_2 \dots e_k$ of edges with $e_1 = e$ and $e_k = f$ such that e_i and e_{i+1} are opposite edges in a convex 4-cycle or 6-cycle of G .*

Proof. Suppose that there is a sequence $e_1 e_2 \dots e_k$ of edges as described in the proposition. Successive edges in this sequence are opposite in a convex,

and thus isometric, 4-cycle or 6-cycle of G . Recall that edges which are opposite in an isometric cycle of G have the same color by Corollary 2.6. Thus $c(e_i) = c(e_{i+1})$ for $i = 1 \dots k - 1$, and hence $c(e) = c(f)$.

For the other direction, suppose that $c(e) = c(f) = xy \in \text{Inc}(\mathcal{P})$. Let $e = L_1L_2$ and $f = L'_1L'_2$, such that $L_1, L'_1 \in W_{xy}$ and $L_2, L'_2 \in W_{yx}$, see Figure 2.2. Our plan is to construct an L_1 - L'_1 -path T_1 in W_{xy} , and an L_2 - L'_2 -path T_2 in W_{yx} . The path T_2 mirrors T_1 , that is, the i -th vertex on T_2 differs from the i -th vertex on T_1 only in the order of x and y . Both paths stay as close as possible to the edge cut $\theta(yx)$, and hence to one another. More precisely, there are no two consecutive vertices on T_1 or T_2 which are not adjacent to an edge in $\theta(xy)$. This means that either the i -th vertex of T_1 is connected to the i -th vertex of T_2 via an edge in $\theta(xy)$, or this holds for the $(i - 1)$ -th vertex and the $(i + 1)$ -th vertex.

The edges connecting T_1 and T_2 form a sequence of edges in which consecutive edges are opposite in 4-cycles or 6-cycles. All 4-cycles of G are convex, since G is bipartite. Our construction also yields that all 6-cycles induced by the sequence are convex in G . Thus the edges connecting T_1 and T_2 form the sequence we are looking for.




Figure 2.2: The paths T_1 and T_2 and the edges in $\theta(xy)$ connecting them.

First let us construct T_1 by describing a process to modify L_1 into L'_1 while always keeping the invariant that there is at most one element between x and y . This invariant assures that T_1 stays as close as possible to the edge cut $\theta(yx)$. Assume that $L_1 = u_1u_2 \dots u_n$ and $L'_1 = v_1v_2 \dots v_n$, and that the first $i - 1$ elements of L_1 and L'_1 coincide. Find the position of v_i in L_1 , say, $v_i = u_j$. Then u_j is a minimum of $\mathcal{P} - \{u_1, u_2, \dots, u_{i-1}\}$.

If $x \neq u_j$, we perform swaps in L_1 bringing u_j to the i -th position of L_1 . That is, we first swap u_j below u_{j-1} , then below u_{j-2} , and so on until we finally swap it below u_{i+1} . This keeps the invariant intact.

If $x = u_j$, then $y = u_{j+1} = v_{i+1}$. Thus y is a minimal element of $\mathcal{P} - \{u_1, u_2, \dots, u_{i-1}, x\}$. In this case, we bring x and y down in L_1 together by alternately swapping x and y with some element. That is, we first

swap x below u_{j-1} , then we swap y below u_{j-1} , and so on, until we swap x below u_{i+1} and then y below u_{i+1} . This keeps our invariant intact. We arrive at a linear extension which coincides with L'_1 in the first i positions (at least). Iterating this process, we can modify L_1 into L'_1 and obtain the path T_1 with the desired properties.

Note that any element which appears between x and y in a linear extension L on T_1 is incomparable with both x and y , because it is above both x and y in L_1 and below both x and y in L'_1 , or below x and y in L_1 and above x and y in L'_1 . Thus we obtain another linear extension of \mathcal{P} by exchanging x and y in L . In this way we obtain the path T_2 with the desired properties. Note that all edges connecting T_1 and T_2 are in $\theta(uv)$. These edges form the sequence $e_1 \dots e_k$ we are looking for. If e_i and e_{i+1} are opposite in a 6-cycle C , then C corresponds to the six permutations of x , y and a third element v which is incomparable to x and y . It follows that C forms a convex Perm_3 in G . This concludes the proof. \square

From the above result it follows that the swap partition $\Theta(G)$ can be obtained greedily: Start with an arbitrary edge and assign it to the first color class. Then iteratively assign all edges that are opposite in a 4-cycle or 6-cycle to an already assigned edge to the same color class. If this process ends, repeat it starting with an unassigned edge and a new color class. However, since the color classes equal the Djoković-Winkler classes, Eppstein's partial cube recognition algorithm mentioned in Theorem 1.12 provides a more efficient way of obtaining $\Theta(G)$.

Below we define some relations between color classes of a linear extension graph. They play a central role in our proofs. See Figure 2.3 for illustration.

Definition 2.11. *Let G be a linear extension graph, and let $\theta_1, \theta_2 \in \Theta(G)$. We say that θ_1 and θ_2 are crossing if both components of $G - \theta_1$ contain edges of θ_2 . Otherwise, we say that θ_1 and θ_2 are parallel.*

The color classes θ_1 and θ_2 touch if they are parallel and there is a vertex incident to edges of both classes.

If θ_1 and θ_2 are parallel, then $\theta_i \in \Theta(G)$ lies between θ_1 and θ_2 if all edges of θ_1 are contained in one component of $G - \theta_i$, and all edges of θ_2 are contained in the other component of $G - \theta_i$.

Let c be a swap coloring of $G = G(\mathcal{P})$. We also use the notions crossing, parallel, lying between, and touching for the corresponding swap colors $c(\theta_1), c(\theta_2) \in \text{Inc}(\mathcal{P})$.

Let us first check that crossing is well-defined, that is, if both components of $G - \theta_1$ contain edges of θ_2 , then both components of $G - \theta_2$ contain edges of θ_1 . Let c be a swap coloring of $G = G(\mathcal{P})$, and let $c(\theta_1) = xv$




Figure 2.3: A linear extension graph with swap colors. The color classes $\theta(uy)$ and $\theta(xv)$ are crossing, while $\theta(uv)$ and $\theta(xy)$ are parallel. Moreover, $\theta(uy)$ and $\theta(xy)$ are touching, and $\theta(xv)$ lies between $\theta(xy)$ and $\theta(uv)$.

and $c(\theta_2) = uy$. Both W_{xv} and W_{vx} contain edges from θ_2 . Pick a vertex L in $W_{xv} \cup W_{uy}$ and a vertex L' in $W_{vx} \cup W_{uy}$. Any shortest $L-L'$ -path in G contains an edge from θ_1 by Lemma 2.2 and stays within W_{uy} by Lemma 2.3. Hence W_{uy} contains an edge from θ_1 . The same argument can be applied to show that W_{yu} contains an edge from θ_1 .

Note that any pair of color classes (and hence, of swap colors) is either crossing or parallel. If two color classes are parallel, it means that there is some transitive forcing between the elements appearing in the corresponding swap colors. Let us look at the example in Figure 2.3: If $x < y$ in a linear extension L of \mathcal{P} , then from $y < v$ in \mathcal{P} it follows by transitivity that $x < v$ in L . We say that $x < y$ forces $x < v$. On the other hand, no order of u, y forces an order of x, v . This is the reason that $\theta(xy)$ and $\theta(xv)$ are parallel in $G(\mathcal{P})$, while $\theta(uy)$ and $\theta(xv)$ are crossing.

In the following lemma, we characterize the relations between color classes given in Definition 2.11 in terms of the corresponding swap colors.

Lemma 2.12. *Let $G = G(\mathcal{P})$ be a linear extension graph with swap coloring c , and let $\theta_1, \theta_2 \in \Theta(G)$.*

- (1) *The color classes θ_1 and θ_2 are parallel exactly if $c(\theta_1) = uv$ and $c(\theta_2) = xy$ for $u, v, x, y \in \mathcal{P}$ such that $u \leq x$ and $y \leq v$ in \mathcal{P} .*
- (2) *Let $\theta_i \in \Theta(G)$ be a third color class of $G(\mathcal{P})$. Then θ_i lies between θ_1 and θ_2 exactly if $c(\theta_i) = ab$ for $a \in I(u, x)$ and $b \in I(y, v)$.*
- (3) *The color classes θ_1 and θ_2 touch exactly if either $y = v$ and $u < x$ is a cover relation in \mathcal{P} , or $u = x$ and $y < v$ is a cover relation in \mathcal{P} .*

Proof. To prove (1), assume that θ_1 and θ_2 are parallel. Then by definition there are elements $u, v, x, y \in \mathcal{P}$ with $c(\theta_1) = uv$ and $c(\theta_2) = xy$ such that $W_{uv} \subset W_{xy}$. Now $W_{uv} \subset W_{xy}$ is equivalent to saying that $x < y$ forces $u < v$. This is the case exactly if $u \leq x$ and $y \leq v$ in \mathcal{P} .

In order to prove (2) recall that, by definition, θ_i lies between θ_1 and θ_2 exactly if there are elements $a, b \in \mathcal{P}$ such that $c(\theta_i) = ab$ and $W_{uv} \subset W_{ab}$ and $W_{yx} \subset W_{ba}$. This is equivalent to the fact that $a < b$ forces $u < v$ and $b < a$ forces $y < x$. This in turn is equivalent to $u \leq a$ and $b \leq v$ and also $y \leq b$ and $a \leq x$ in \mathcal{P} , which simply means that $a \in I(u, x)$ and $b \in I(y, v)$.

It remains to prove (3). Let us first assume that θ_1 and θ_2 are touching and that $c(\theta_1) = uv$ and $c(\theta_2) = xy$ with $u \leq x$ and $y \leq v$ in \mathcal{P} as in the previous paragraphs. From $u \parallel v$ and $x \parallel y$ it follows that $x \parallel v$ and $u \parallel y$. If $u \neq x$ and $y \neq v$, then the color classes $\theta(xv)$ and $\theta(uy)$ lie between θ_1 and θ_2 by (2). But then θ_1 and θ_2 cannot touch, a contradiction. So we may assume that $y = v$. If $u < x$ is not a cover relation, then there is an element $a \in \mathcal{P}$ with $u < a < x$ in \mathcal{P} . But then the color class $\theta(av)$ lies between θ_1 and θ_2 , and again they cannot touch. This proves the first direction.

For the other direction, let us assume that $c(\theta_1) = uv$ and $c(\theta_2) = vx$ such that $u < x$ is a cover relation in \mathcal{P} . Then θ_1 and θ_2 are parallel by (1). Because $u < x$ is a cover relation, we can construct a linear extension L by first picking all elements in $\text{Pred}(u) \cup \text{Pred}(v)$, then picking uvx in this order, and then the remaining elements. Now L is incident to an edge with swap color uv and an edge with swap color vx . Thus, θ_1 and θ_2 touch in L . This concludes the proof. \square

We now present two lemmas characterizing the color classes corresponding to special types of incomparable pairs, see Figure 2.4 for illustration. We start with the case of a twin, see Definition 1.24.

Lemma 2.13. *Let $G(\mathcal{P})$ be a linear extension graph with swap coloring c . Let $\theta \in \Theta(G)$ with $c(\theta) = xy$. Then the pair xy is a twin of \mathcal{P} exactly if no color class of $G(\mathcal{P})$ is parallel to θ . In this case, G_{xy} and G_{yx} are isomorphic.*

Proof. If xy is a twin of \mathcal{P} , then by definition, every element larger (smaller) than x in \mathcal{P} is also larger (smaller) than y in \mathcal{P} , and vice versa. But by Lemma 2.12, a color uv is parallel to xy in $G(\mathcal{P})$ exactly if $u \geq x$ and $v \leq y$, where at least one of these relations is strict. Furthermore, observe that xy is a twin of \mathcal{P} exactly if x and y can be exchanged in every linear extension of \mathcal{P} . This yields an isomorphism between G_{xy} and G_{yx} . \square

Observe that G_{xy} and G_{yx} may be isomorphic even if xy is not a twin. As an example, consider the pair uy in the example of Figure 2.3. Furthermore, note that Lemma 2.13 shows that a linear extension graph can have different swap colorings corresponding to the same poset.

Now let us look at critical pairs, cf. Definition 1.2.




Figure 2.4: The poset \mathcal{P} contains the twin vw and the critical pairs (u, a) and (a, y) .

Lemma 2.14 ([51]). *Let $G = G(\mathcal{P})$ be a linear extension graph with swap coloring c . Let $\theta \in \Theta(G)$ with $c(\theta) = xy$. Then (x, y) is a critical pair of \mathcal{P} exactly if there is no color class of G which is completely contained in G_{yx} .*

Proof. Consider two incomparable pairs v, w and x, y of \mathcal{P} . We have $G_{vw} \subset G_{yx}$ exactly if $v < w$ forces $y < x$. In other words, there is no color class completely contained in G_{yx} exactly if there is no relation which forces $y < x$. We also know that (x, y) is a critical pair of \mathcal{P} exactly if $y < x$ cannot be forced by any other relation. This yields the result. \square

A sequence of pairwise parallel color classes looks like *stripes* in a linear extension class, see for examples the swap colors containing the element a in Figure 2.4. If we consider a maximal sequence of such stripes in a linear extension graph, then the above lemma tells us that the color classes at the two ends of that sequence correspond to critical pairs. In our example, these are the swap colors au and ay . They are reversed on the side of their color class which does not contain any other color class. This may hold for both sides of the color class – then the corresponding swap color is a twin and hence critical in both orders.

In the next proposition we come back to the relation between color classes. We characterize the way that two color classes can cross in a linear extension graph. Let θ_1 and θ_2 be two color classes of a linear extension graph G , and suppose there is a convex cycle C in G which contains edges from both color classes. Then by Corollary 2.6, we know that θ_1 and θ_2 are crossing. We say that they *cross in C* .

Proposition 2.15. *Let G be a linear extension graph and let θ_1 and θ_2 be a crossing pair of color classes of G . Then they cross in a convex 4-cycle or a convex 6-cycle of G .*

Proof. Let $G = G(\mathcal{P})$ and let c be a swap coloring of G . Let $c(\theta_1) = uv$ and choose $e \in \theta_2$ in G_{uv} and $f \in \theta_2$ in G_{vu} , see Figure 2.5. By Proposition 2.10, there is a sequence S of edges in θ_1 starting with e and ending with f , such that successive edges in the sequence are opposite in a convex 4- or 6-cycle.

Let e' be the last edge in S which is contained in G_{uv} and let f' be the first edge in S which is contained in G_{vu} . Then e' and f' are opposite edges in a convex 4- or 6-cycle C . By definition of e' and f' , a path from an endpoint of e' to an endpoint of f' has to pass an edge in θ_1 . In particular, the path contained in C has to contain an edge in θ_1 . Thus, θ_1 and θ_2 cross in C . \square




Figure 2.5: If θ_1 and θ_2 are crossing, they meet in a convex 4- or 6-cycle.

2.3 Adjacent Swap Colors

In this section, we analyze the relation of two swap colors by looking at the corresponding color classes. Recall that the incomparability graph $\text{Incomp}(\mathcal{P})$ is the undirected graph on the ground set of \mathcal{P} in which two vertices are adjacent if the two corresponding elements are incomparable in \mathcal{P} . Thus, the incomparable pairs of \mathcal{P} , and hence the swap colors of $G(\mathcal{P})$, correspond to the edges of $\text{Incomp}(\mathcal{P})$.

Let us call two swap colors adjacent if they correspond to adjacent edges in the incomparability graph. More precisely, if $G = G(\mathcal{P})$ has swap coloring c , and $\theta_1, \theta_2 \in \Theta(G)$, then $c(\theta_1)$ and $c(\theta_2)$ are *adjacent* if $c(\theta_1) \cap c(\theta_2) \neq \emptyset$.

Our aim in this section is to characterize adjacent swap colors in terms of the relation of their color classes in $G(\mathcal{P})$. If we pick two adjacent edges of $G(\mathcal{P})$, then we can easily characterize how their swap colors relate.

Lemma 2.16. *Let $G = G(\mathcal{P})$ be a linear extension graph with swap coloring c . Let e and f be two adjacent edges of G . Then the following holds:*

- (i) *The color classes $\theta(e)$ and $\theta(f)$ cross in a 4-cycle exactly if $c(e)$ and $c(f)$ are disjoint.*
- (ii) *The color classes $\theta(e)$ and $\theta(f)$ cross in a convex 6-cycle exactly if $c(e)$ and $c(f)$ share an element, and the remaining two elements are incomparable.*
- (iii) *The color classes $\theta(e)$ and $\theta(f)$ are parallel exactly if $c(e)$ and $c(f)$ share an element, and the remaining two elements form a cover relation of \mathcal{P} .*

Proof. To prove (i), let $c(e) = uv$ and $c(f) = xy$. We first assume that u, v, x, y are pairwise different elements of \mathcal{P} . Let L be the linear extension of \mathcal{P} in which e and f meet. Thus elements u and v are adjacent in L , as well as x and y . Since uv and xy are disjoint, there is a linear extension L' at distance 2 from L in which uv and xy are reversed. There are two L - L' -paths of length 2, one which swaps first uv and then xy , and one which swaps first xy and then uv . The first path starts with e , and the second path with f . These two paths form a 4-cycle in which $c(e)$ and $c(f)$ cross.

Now suppose that $\theta(uv)$ and $\theta(xy)$ cross in a 4-cycle C . Assume for contradiction that uv and xy are not disjoint, say, $xy = vy$. Since G is bipartite, C is isometric. By Corollary 2.6, every linear extension in C is incident to an edge of color uv in C and an edge of color vy in C . Let $L \in C$. Then element v has to sit between u and y in L . But if we swap uv in L to obtain $L' \in C$, the pair vy is not adjacent in L' . Thus L' does not have an incident edge of color vy , which is a contradiction.

To prove (ii), first assume that $c(e)$ and $c(f)$ cross in a convex 6-cycle C . If $c(e)$ and $c(f)$ are disjoint, we have seen above that the edges e and f are contained in a 4-cycle. But this is a contradiction since C is convex. Thus we may assume that $c(e) = uv$ and $c(f) = vw$. Now suppose that $u \sim w$ in \mathcal{P} , say, $u < w$. Then $v < u$ in a linear extension of \mathcal{P} forces $v < w$ by transitivity. Hence there is no edge of color vw in W_{vu} . It follows that $\theta(e)$ and $\theta(f)$ are parallel, a contradiction. Thus we have $u \parallel w$ in \mathcal{P} .

For the other direction, let $c(e) = uv$ and $c(f) = vw$, and $u \parallel w$ in \mathcal{P} . Let L be the linear extension of \mathcal{P} in which e and f meet. Then the three elements u, v, w must appear consecutively in L such that v sits between u and w . Let L' be the linear extension at distance 3 from L in which the order of u, v, w is reversed. There are two L - L' -paths of length 3. The

first swaps uv , then uw , and then vw ; the second swaps vw , then uw , and then uv . The first path thus starts with e , and the second path with f . These two paths form a convex 6-cycle in which $\theta(e)$ and $\theta(f)$ cross.

Now (iii) holds by complete exhaustion of the possible cases. If $\theta(e)$ and $\theta(f)$ do not cross in a 4-cycle nor in a 6-cycle, then by Proposition 2.15, they are parallel. Thus $\theta(e)$ and $\theta(f)$ touch, and we can use Lemma 2.12 to prove the desired equivalence. \square

Note that possibilities (i) and (ii) are disjoint. Since it follows from Proposition 2.15 that any two crossing color classes contain two adjacent edges, we now know that two crossing color classes either cross in a 4-cycle or in a convex 6-cycle of G , but not in both. They cross in a 4-cycle exactly if the corresponding colors are disjoint.

One could be led to think that two arbitrary colors are disjoint exactly if they cross in a 4-cycle. This was also claimed by Reuter in [50]. However, there is a case missing, which can be seen in the example in Figure 2.3: The swap colors uv and xy are parallel, but disjoint.

The following theorem characterizes adjacent colors in terms of their color classes.

Theorem 2.17. *Let $G = G(\mathcal{P})$ be a linear extension graph with swap coloring c and let $\theta, \theta' \in \Theta(G)$. Then $c(\theta)$ and $c(\theta')$ are adjacent exactly if one of the following two options holds:*

- (1) $c(\theta) = uv$ and $c(\theta') = vw$ for elements $u, v, w \in \mathcal{P}$ with $u \parallel w \iff \theta$ and θ' cross in a convex 6-cycle.
- (2) $c(\theta) = uv$ and $c(\theta') = vw$ for elements $u, v, w \in \mathcal{P}$ with $u \sim w \iff \theta$ and θ' are parallel, and no two color classes between them cross in a 4-cycle.

Proof. To show the equivalence (1), let us first assume that $c(\theta) = uv$ and $c(\theta') = vw$, and that u, v and w form an antichain in \mathcal{P} . Then we can build linear extensions of \mathcal{P} by first picking all elements which are predecessors of one of u, v, w , then picking these three elements in some order, and then the remaining elements of \mathcal{P} . This yields six linear extensions of \mathcal{P} which differ only in the order of u, v, w . Thus, these linear extension induce Perm_3 as convex subgraph of G , which is a convex 6-cycle C . The colors appearing on C are uv, vw and uw . By Corollary 2.6, the classes θ and θ' cross in C .

Conversely, if θ and θ' cross in a convex 6-cycle, then there are two adjacent edges of θ and θ' for which (ii) of Lemma 2.16 holds. This yields the desired properties.

To show (2), assume that $c(\theta) = uv$ and $c(\theta') = vw$ with $u < w$. Then we know from Lemma 2.12 that θ and θ' are parallel. The lemma also tells us that any color class between θ and θ' has swap color av for some $a \in I(u, w)$. Thus the swap colors of any two classes lying between θ and θ' share the element v . Now, by Lemma 2.16, if two color classes cross in a 4-cycle, their swap colors are disjoint. Therefore no pair of color classes between θ and θ' crosses in a 4-cycle.

To prove the converse we can apply Lemma 2.12 to the parallel color classes θ and θ' and assume that $c(\theta) = uv$ and $c(\theta') = xy$ with $u \leq x$ and $y \leq v$ in \mathcal{P} . We want to show that $u = x$ or $y = v$ has to hold. First observe that with $u||v$ and $x||y$ in \mathcal{P} it follows that $x||v$ and $u||y$ in \mathcal{P} . Thus there are two color classes θ_i, θ_j in G which have swap colors xv and uy , respectively. By Lemma 2.12, they both lie between θ and θ' . Assume for contradiction that $u \neq x$ and $y \neq v$. Then $c(\theta_i)$ and $c(\theta_j)$ are disjoint. Furthermore, θ_i and θ_j are crossing, since no choice of an order of x and v forces an order of u and y . As we noted after Lemma 2.16, any pair of crossing color classes with disjoint colors cross in a 4-cycle. Thus θ_i and θ_j cross in a 4-cycle, which is a contradiction. \square

We give a second characterization of adjacent swap colors which we find interesting in its own right. The key observation for the characterization is contained in the following lemma. For a set S of swap colors, let us set $\Theta_S = \{\theta \in \Theta(G) : c(\theta) \in S\}$.

Lemma 2.18. *Let $G = G(\mathcal{P})$ be a linear extension graph with swap partition $\Theta(G)$ and swap coloring c . Let $S \subset \text{Inc}(\mathcal{P})$ be a set of swap colors of G which all share an element $v \in \mathcal{P}$. Then it holds that $G[\Theta_S]$ is cycle-free.*

Proof. Assume for contradiction that C is a shortest cycle in $G[\Theta_S]$. Let L be the linear extension appearing in C in which v has the highest position that it obtains among all the linear extensions on C . Assume that $L = \dots xvy \dots$, i.e., $x < v$ and $v < y$ are cover relations in L . Since any color on C contains v , but v does not appear after y on C , the two edges incident to L in C must both have color xv . But this is a contradiction, since L cannot have two incident edges of the same color. \square

We need one more new notation: Let a θ - θ' -path be a path starting with an edge of the color class θ and ending with an edge of the color class θ' .

Theorem 2.19. *Let θ, θ' be two color classes of a linear extension graph $G = G(\mathcal{P})$ with swap coloring c . Then $c(\theta)$ and $c(\theta')$ are adjacent exactly if for any shortest θ - θ' -path T , the color classes appearing on T induce a forest in G .*

Proof. First let us assume that $c(\theta)$ and $c(\theta')$ share an element, say, $c(\theta) = uv$ and $c(\theta') = vw$. If $u \parallel w$ in \mathcal{P} , then θ and θ' cross in a 6-cycle by Theorem 2.17. Thus any shortest θ - θ' -path consists of just two edges, and the swap colors used on that path all share the element v . By Lemma 2.18, the induced graph $G[\theta \cup \theta']$ is a forest.

If $u \sim w$ in \mathcal{P} , then θ and θ' are parallel by Theorem 2.17. Let us assume that $u < w$ in \mathcal{P} . By Lemma 2.12, a color class lies between θ and θ' exactly if it has swap color av for some $a \in I(u, w)$. Let us denote the set of these swap colors av by S . We claim that a θ - θ' -path T is a shortest θ - θ' -path if and only if the swap colors used on T are exactly the swap colors in S .

Every color class lying between θ and θ' must clearly be used on each θ - θ' -path. So the set of swap colors used on T contains S . We will construct a θ - θ' -path using only the swap colors of S to prove our claim. To do so, we construct two linear extensions differing only in the position of v .

Let $D = \text{Pred}(w) \cup \text{Pred}(v)$, and $U = \text{Succ}(u) \cup \text{Succ}(v)$. Then we have $\mathcal{P} = D \cup I(u, w) \cup v \cup U$. Let L_X be a linear extension of $\mathcal{P}[X]$, for $X \in \{D, U, I(u, w)\}$. Observe that because of $v \parallel u$ and $v \parallel w$, we know that v is incomparable in \mathcal{P} to all elements of $I(u, w)$. Therefore $L = L_D v L_I L_U$ and $L' = L_D L_I v L_U$ are linear extensions of \mathcal{P} . The vertices L and L' of G are connected by a path which uses exactly the swap colors in S . This proves our claim.

All colors in S , and thus all swap colors on a shortest θ - θ' -path, share the element v . Hence it follows from Lemma 2.18 that the classes appearing on a shortest θ - θ' -path cannot induce a cycle in G . This proves the first direction of the theorem.

For the other direction we assume that the swap colors on any shortest θ - θ' -path in G induce a forest. Since G is connected, there is such a path T . Let $T = e_1 e_2 \dots e_k$ with $\theta(e_1) = \theta$, and $\theta(e_k) = \theta'$. Denote the set of swap colors used on T by S . Because $G[\Theta_S]$ is cycle-free, it does not contain 4-cycles. Hence it follows from part (i) of Lemma 2.16 that any pair of swap colors appearing consecutively on T shares an element. Let $c(e_1)$ and $c(e_2)$ share the element $b \in \mathcal{P}$. We prove inductively that all swap colors on T contain the element b .

Assume that the swap colors $c(e_1) \dots c(e_i)$ all contain the element b , and assume that $c(e_{i-1}) = ab$ and $c(e_i) = bc$. Consider the linear extension L in which e_{i-1} and e_i meet. We may assume that the sequence $\dots abcd \dots$ is contained in L . Then the sequence $\dots acbd \dots$ is contained in the linear extension in which e_i and e_{i+1} meet. Since $c(e_{i+1})$ needs to share an element with $c(e_i) = bc$, we have $c(e_{i+1}) = ac$ or $c(e_{i+1}) = bd$. If $c(e_{i+1}) = ac$, then S contains the three swap colors ab, bc and ac , for a 3-antichain $a, b, c \in \mathcal{P}$. But then it follows from part (ii) of Lemma 2.16 that these three swap

colors form a convex 6-cycle in G . Thus this cycle appears also in $G[\Theta_S]$, which is a contradiction. Therefore we have $c(e_{i+1}) = bd$, and it follows by induction that all colors in S contain the element b . \square

2.4 The Linear Extension Graph as an Invariant

In this section we consider the question which modifications of the poset leave the linear extension graph invariant.

Given a graph G , we say that a poset \mathcal{P} is G -compatible if $G = G(\mathcal{P})$. Can a graph G have different G -compatible posets? What immediately comes to mind is that we can reverse the direction of a poset without changing its linear extension graph: If we change a poset \mathcal{P} to \mathcal{P}^* , then the direction of any linear extension is also reversed, and the swaps stay the same. Therefore we have $G(\mathcal{P}) = G(\mathcal{P}^*)$ for any \mathcal{P} .

Starting from a poset \mathcal{P} , it is also easy to find an infinite family of posets which have the same linear extension graph as \mathcal{P} : If a global minimum is added to \mathcal{P} , then it needs to be the smallest element in any linear extension of \mathcal{P} . Thus it does not have any influence on the structure of the linear extension graph of \mathcal{P} . Hence we can add arbitrarily long chains of global minima to \mathcal{P} without changing its linear extension graph. If \mathcal{P} contains other *global elements*, that is, elements comparable to all other elements in \mathcal{P} , then these can also be deleted or replaced by arbitrarily long chains without effect on the linear extension graph.

Definition 2.20. *A G -compatible poset is essentially unique if it is unique up to its direction and the addition or deletion of global elements.*

More basic alterations which leave the linear extension graph invariant are possible if \mathcal{P} is a series composition. Recall that a series composition $\mathcal{P}_1 * \mathcal{P}_2$ results from putting \mathcal{P}_2 “on top of” \mathcal{P}_1 , see Definition 1.25.

The *Cartesian product* $G_1 \square G_2$ of two graphs G_1, G_2 is defined as follows: The vertices of $G_1 \square G_2$ are the pairs (v_1, v_2) with $v_1 \in V(G_1)$ and $v_2 \in V(G_2)$. Two such vertices (v_1, v_2) and (v'_1, v'_2) are adjacent if $v_1 = v'_1$ and $v_2 v'_2 \in E(G_2)$, or vice versa.

A slightly more general version of the following basic Proposition appeared in [44].

Proposition 2.21 ([50]). *If $\mathcal{P} = \mathcal{P}_1 * \mathcal{P}_2 * \dots * \mathcal{P}_k$ for posets $\mathcal{P}_1, \dots, \mathcal{P}_k$, then $G(\mathcal{P}) = G(\mathcal{P}_1) \square G(\mathcal{P}_2) \square \dots \square G(\mathcal{P}_k)$.*

Proof. Every linear extension L of \mathcal{P} is a concatenation of linear extensions of the \mathcal{P}_i . On the other hand, every concatenation of linear extensions

of the \mathcal{P}_i in the right order is a linear extension of \mathcal{P} . Moreover, a swap in L has to swap two elements of the same \mathcal{P}_i . That is, an edge in $G(\mathcal{P})$ corresponds to an edge in one of the $G(\mathcal{P}_i)$. It is now easy to see that $G(\mathcal{P})$ is the Cartesian product of the $G(\mathcal{P}_i)$. \square

Definition 2.22. *Let $\mathcal{P}, \mathcal{P}'$ be two posets. We say that \mathcal{P}' arises from series alteration of \mathcal{P} if $\mathcal{P} = \mathcal{P}_1 * \mathcal{P}_2 * \dots * \mathcal{P}_k$ and \mathcal{P}' arises from \mathcal{P} by changing the order of the \mathcal{P}_i and the direction of some of the \mathcal{P}_i in the series composition.*

The Cartesian product is clearly commutative, so we have the following corollary:

Corollary 2.23. *The linear extension graph is invariant under series alteration.* \square

The next question that comes up is how far we can go with modifying \mathcal{P} while keeping $G(\mathcal{P})$ invariant. Recall that the comparability graph $\text{Comp}(\mathcal{P})$ of a poset \mathcal{P} is the graph on the ground set of \mathcal{P} , with two elements being adjacent exactly if they are comparable in \mathcal{P} . Clearly, series alterations are a modification of \mathcal{P} which do not change the comparability graph.

Does $G(\mathcal{P})$ stay invariant under all poset modifications leaving the comparability graph invariant? It was first observed by Reuter [50] that the linear extension graph is not a comparability invariant. He gave the counterexample depicted in Figure 2.6.

Theorem 1.26 says that a poset parameter is a comparability invariant exactly if it remains the same whenever a subposet of \mathcal{P} induced by a module is replaced by its dual. In Reuter's example, \mathcal{P} and \mathcal{P}' differ in the direction of the subposet induced by the module $\{c, d, e\}$, but the linear extension graph does not remain invariant. We will see in Section 3.4 how the color classes of a linear extension graph G determine the direction of modular subposets in a G -compatible poset \mathcal{P} .

In the next chapter, we give a reconstruction procedure which proves that the series alteration and the addition of global elements are the only modifications leaving the linear extension graph invariant. Let us call a graph *Cartesian prime* if it is not the Cartesian product of several non-trivial Cartesian factors. Our procedure will imply the following theorem (see Section 3.5):

Theorem 2.24. *Let G be a linear extension graph. Then the G -compatible poset is essentially unique exactly if G is Cartesian prime.*




Figure 2.6: The linear extension graph is not a comparability invariant, since the two posets \mathcal{P} and \mathcal{P}' have the same comparability graph, but their linear extension graphs are not isomorphic.

Chapter 3

Reconstructing Posets from Linear Extension Graphs

We have seen in the last chapter that each linear extension graph G has many different G -compatible posets. However, the differences we saw were not very exciting. In this chapter we give a procedure to reconstruct all G -compatible posets of a linear extension graph G . We will see that there are no poset modifications leaving the linear extension graph invariant apart from the ones we saw. We will also show how to use our procedure to recognize linear extension graphs.

The first five sections of this chapter are devoted to the reconstruction procedure, while the last section explains the recognition. Let us fix a linear extension graph $G = (V, E)$ for all sections dealing with the reconstruction. We denote the set of G -compatible posets by \mathbf{P}_G .

Since a poset can have exponentially many linear extensions, G can be huge compared to a poset $\mathcal{P} \in \mathbf{P}_G$. But much of the information that it carries appears redundantly. It turns out that for the reconstruction we do not need to know the whole graph G . Recall that G comes with a unique swap partition $\Theta(G)$. For the reconstruction, we only need to know certain relations between the color classes in $\Theta(G)$.

Let us assume that we do not know G , but we know a visionary person, a G -medium, which can answer certain questions related to $\Theta(G)$. To start with, the G -medium tells us the number r of color classes of G . Recall that $r = \dim_I(G)$ by Theorem 1.13. Moreover, the G -medium can give answers to the following *admissible questions* about a pair $\theta_i, \theta_j \in \Theta(G)$:

1. Are θ_i and θ_j parallel or crossing?
2. In case θ_i and θ_j are crossing, do they cross in a 4-cycle or in a 6-cycle?
3. In case θ_i and θ_j are parallel, are they touching?
4. If θ_i and θ_j are parallel and θ_k is a third color class of G , does θ_k lie between θ_i and θ_j ?

The time we have to wait for an answer to an admissible question depends on what power our G -medium possesses. It is long if the G -medium only sees G , but very short if it can look beyond and see some $\mathcal{P} \in \mathbf{P}_G$. This is made precise in Section 3.1.

The first main aim of this chapter is to prove the following:

Theorem 3.1. *Let G be a linear extension graph. Given a G -medium, we can reconstruct the set \mathbf{P}_G in time $O(\dim_I(G)^4 \cdot q)$, where q is the time the G -medium needs to answer an admissible question.*

Recall that the addition of global elements to a poset $\mathcal{P} \in \mathbf{P}_G$ yields another G -compatible poset. Thus \mathbf{P}_G contains infinitely many posets. The reconstruction procedure outputs a minimal poset in \mathbf{P}_G and a concise description of how to obtain all other posets on \mathbf{P}_G .

We will prove Theorem 3.1 by giving a reconstruction procedure in several steps. Each section of this chapter is dedicated to one step of the procedure. In most steps we construct an object (e.g. a comparability graph or a transitive orientation) which corresponds to a G -compatible poset. These objects will also be called G -compatible. This will become clear from the context.

In Section 3.1, we reveal how to answer admissible questions. This includes showing how to obtain the color classes of G efficiently. In Section 3.2 we determine which pairs of color classes correspond to adjacent swap colors. Thus, we construct the unique line graph of a G -compatible incomparability graph. From this we obtain, in Section 3.3, the unique minimal G -compatible comparability graph $\text{Comp}(G)$. The most difficult part is to construct all G -compatible transitive orientations of $\text{Comp}(G)$. We do this in Section 3.4, thus reconstructing \mathbf{P}_G . In Section 3.5, we finish the proof of Theorem 3.1 by discussing the running time. Finally, in Section 3.6 we show how to use the reconstruction procedure to recognize whether an arbitrary given graph is a linear extension graph.

3.1 Answering Admissible Questions

In this section we first show that the G -medium can answer each admissible question in constant time if it sees a poset $\mathcal{P} \in \mathbf{P}_G$ and a swap coloring c of G with respect to \mathcal{P} . After that we show that if the G -medium sees G with the partition $\Theta(G)$, then it can answer each admissible question in time $O(|V| + |E|)$. We also show that $\Theta(G)$ can be obtained from G in time $O(|V|^2)$.

Let us first assume that our G -medium is given a poset $\mathcal{P} \in \mathbf{P}_G$. We want the given structure to distinguish between relations and cover relations of \mathcal{P} . If this is not the case, i.e., if only the relations of \mathcal{P} are known, it can be achieved in a preprocessing step: Consider the matrix M whose rows and columns are labeled with \mathcal{P} , such that position (u, v) has entry 1 exactly if $u < v$ in \mathcal{P} , and 0 otherwise. Now compute the matrix M^2 . Then the rules of matrix multiplications imply that the entry of M^2 at position (u, v) equals the number of elements $w \in \mathcal{P}$ such that $u < w$ and $w < v$ in \mathcal{P} . Thus, $u < v$ is a cover relation of \mathcal{P} exactly if position (u, v) is 1 in M , but 0 in M^2 . This preprocessing step can be done in time $O(|\mathcal{P}|^\alpha)$, where $\alpha \geq 2$ is the exponent of a matrix multiplication algorithm. Currently, the fastest known algorithm has $\alpha = 2.376$ [10].

So, in the following lemma, we assume that \mathcal{P} is given in a way which allows us to check for each pair $u, v \in \mathcal{P}$ in constant time whether $u = v$ or $u \parallel v$ or $u < v$ or $v < u$ in \mathcal{P} and, if one of the last two options holds, whether the relation is a cover relation.

Lemma 3.2. *Given a G -compatible poset \mathcal{P} and the swap coloring c of G with respect to \mathcal{P} , the answer to an admissible question can be found in constant time.*

Proof. Suppose that $c(\theta_i) = uv$ and $c(\theta_j) = xy$. By Lemma 2.12, the color classes θ_i and θ_j are parallel exactly if either $u \leq x$ and $y \leq v$, or $u \leq y$ and $x \leq v$, or $u \geq x$ and $y \geq v$, or $u \geq y$ and $x \geq v$. Since any pair of color classes is either crossing or parallel, they are crossing otherwise. Therefore Question 1 can be settled by checking the above four cases.

Assume that θ_i and θ_j are crossing. Then by Proposition 2.15, they either cross in a 4-cycle or in a 6-cycle, thus there are adjacent edges of the two classes. Therefore it follows from Lemma 2.16 that θ_i and θ_j cross in a 4-cycle exactly if uv and xy are disjoint. Otherwise, they cross in a 6-cycle. Hence Question 2 can be answered in constant time by looking at the elements uv and xy .

Questions 3 and 4 can be answered with the help again of Lemma 2.12. Suppose θ_i and θ_j are parallel because $u \leq x$ and $y \leq v$. Then θ_i and θ_j

touch exactly if either $y = v$ and $u < x$ is a cover relation in \mathcal{P} , or $u = x$ and $y < v$ is a cover relation in \mathcal{P} . Hence by assumption, Question 3 can be answered in constant time. Moreover, a third class θ_k with $c(\theta_k) = ab$ lies between θ_i and θ_j exactly if $a \in I(x, u)$ and $b \in I(v, y)$. By checking $x \leq a$ and $a \leq u$ and $v \leq b$ and $b \leq y$ in \mathcal{P} , Question 4 can be answered in constant time. \square

Now let us assume that the G -medium sees only the (unlabeled) linear extension graph G . Let us first show how to construct $\Theta(G)$.

Lemma 3.3. *Given a linear extension graph $G = (V, E)$, the swap partition $\Theta(G)$ can be obtained in time $O(|V|^2)$.*

Proof. By Theorem 2.4, the linear extension graph G is a partial cube, and the color classes of G equal its Djoković-Winkler classes. Thus the partial cube recognition algorithm from Theorem 1.12 constructs the color classes of G . The algorithm runs in time $O(|V|^2)$. \square

Now assume that we are given $G = (V, E)$ with the unique swap partition $\Theta(G) = \{\theta_1, \theta_2, \dots, \theta_r\}$ of E into color classes.

Lemma 3.4. *Given a linear extension graph $G = (V, E)$ and the swap partition $\Theta(G)$, the answer to an admissible question can be found in time $O(|V| + |E|)$.*

Proof. Consider two color classes $\theta_i, \theta_j \in \Theta(G)$. To check whether they are crossing, use the following procedure: Delete all edges in θ_i from G . We are left with two components. Now check if both components contain edges of θ_j . If so, then θ_i and θ_j are crossing, otherwise they are parallel. These steps can be performed in time $O(|V| + |E|)$ and answer Question 1.

To answer Question 2, build $G^{ij} = (V, \theta_i \cup \theta_j)$, the graph induced by the two given color classes. It follows from Theorem 1.13 that the components of G^{ij} are partial cubes of dimension 2. Thus each component is either a single edge, a path of length 2, or a 4-cycle. We then check if G^{ij} contains a 4-cycle. Since a crossing pair of color classes either crosses in a 4-cycle or in a 6-cycle, but not in both, this decides Question 2. Building the graph G^{ij} and searching its components for a 4-cycle can be done in time $O(|V| + |E|)$.

To answer Question 3, we can run through all the vertices of G and check at each vertex whether there is an incident edge from θ_i and an incident edge from θ_j . Thus Question 2 can be decided in time $O(|V| + |E|)$.

Finally, to answer Question 4, observe that in case that θ_i and θ_j are parallel, the graph $G - \theta_i - \theta_j$ has three components, cf. Figure 3.1. Let us denote the two components of $G - \theta_i$ by G_i^+ and G_i^- , such that G_i^+

contains edges of θ_j . Similarly, let us denote the two components of $G - \theta_j$ by G_j^+ and G_j^- , such that G_j^+ contains edges of θ_i . The three components of $G - \theta_i - \theta_j$ then are G_i^- , G_j^- and $G_i^+ \cap G_j^+$.

It follows from the definition that a color class θ_k lies between θ_i and θ_j exactly if all edges of θ_k are contained in $G_i^+ \cap G_j^+$. Thus to answer Question 4, we can build $G - \theta_i - \theta_j$ and check if there is no edge of θ_k contained in G_i^- and G_j^- . This can be done in time $O(|V| + |E|)$. \square




Figure 3.1: If θ_i and θ_j are parallel, then $G - \theta_i - \theta_j$ has three components.

In the following sections, we will just assume that our G -medium can answer admissible questions in time q , knowing that this can vary from constant time to $O(|V| + |E|)$.

3.2 Checking Adjacency of Swap Colors

At the start of our reconstruction procedure, the G -medium only told us the number r of color classes of G . Let us denote the swap color corresponding to θ_i by c_i . If $\mathcal{P} \in \mathbf{P}_G$, then there is a swap coloring c of G with respect to \mathcal{P} such that $c(\theta_i) = c_i$. However, so far we do not know anything about the c_i except which color class of G they correspond to.

If $\mathcal{P} \in \mathbf{P}_G$, then each c_i corresponds to an edge in the incomparability graph of \mathcal{P} . Recall that two swap colors are adjacent if the two corresponding incomparable pairs of \mathcal{P} share an element, that is, if they correspond to adjacent edges in $\text{Incomp}(\mathcal{P})$.

In this section, we want to determine which pairs c_i, c_j of swap colors of G are adjacent. In Theorem 2.17 we have seen that this is uniquely determined by G .

Proposition 3.5. *We can determine all adjacencies between swap colors of G in time $O(r^4 \cdot q)$.*

Proof. For each pair c_i, c_j of swap colors, we need to check whether one of the properties (1) and (2) of Theorem 2.17 holds. Consider the correspond-

ing color classes $\theta_i, \theta_j \in \Theta(G)$. It takes one admissible question to check whether they are crossing or parallel.

In case θ_i and θ_j are crossing, we ask Question 2 to check whether they cross in a 6-cycle. Then (1) is fulfilled exactly if the answer is yes. So we can check (1) with no more than two admissible questions.

In case θ_i and θ_j are parallel, we ask Question 4 for every color class $\theta_k \neq \theta_i, \theta_j$. In this way we determine the set $\Theta(G)_{ij}$ of color classes lying between θ_i and θ_j . After that, we use Questions 1 and 2 to find out whether any pair of color classes in $\Theta(G)_{ij}$ crosses in a 4-cycle. Hence we can check property (2) with $O(r^2)$ admissible questions.

There are $O(r^2)$ pairs of swap colors, so we conclude that the adjacencies can be determined in a total running time of $O(r^4 \cdot q)$. \square

3.3 Reconstructing the Comparability Graph

In the previous section, we have determined which swap colors of G are adjacent. Put differently, we have shown that there is a unique graph $X(G)$ which is the line graph of $\text{Incomp}(\mathcal{P})$ for some poset $\mathcal{P} \in \mathbf{P}_G$. Thus, $X(G)$ is the line graph of $\text{Incomp}(\mathcal{P})$ for *every* poset in \mathbf{P}_G . In this section, we show that there is a unique minimal graph $\text{Comp}(G)$ which is the comparability graph of a poset in \mathbf{P}_G . We also show that it comes with a unique swap coloring c , which we will make precise later. Moreover, we will see how to construct $\text{Comp}(G)$ and c .

Recall from Section 2.4 that G does not determine the number of global elements in a G -compatible poset \mathcal{P} . This is reflected in the fact that a global element of \mathcal{P} is not part of any edge of $\text{Incomp}(\mathcal{P})$.

Lemma 3.6. *Let $\Gamma(G)$ be the family of graphs which are the comparability graph of some poset in \mathbf{P}_G . Then $\Gamma(G)$ contains a unique minimal graph $\text{Comp}(G)$. Furthermore, $\Gamma(G)$ equals the family of all graphs which can be obtained from $\text{Comp}(G)$ by successively adding global vertices.*

Proof. We have seen that there is a unique graph $X(G)$ which is the line graph of $\text{Incomp}(\mathcal{P})$ for some $\mathcal{P} \in \mathbf{P}_G$. We first show that there is a unique graph $H(G)$ which equals $\text{Incomp}(\mathcal{P})$ for some $\mathcal{P} \in \mathbf{P}_G$.

If X is a simple connected graph, then it is a line graph of a simple graph H exactly if the edges of X can be decomposed into cliques such that each vertex appears in exactly two cliques. Such a decomposition is called a *Krausz decomposition* (see [62]). Then H is the graph which has the cliques of a Krausz decomposition as vertices, with two of them being adjacent if they share a vertex. The Krausz decomposition is unique

unless X is a triangle, which is the line graph of both K_3 and $K_{1,3}$. The line graph of a disconnected graph equals the union of the line graphs of its components. These results are easy to prove, see e.g. the discussion of line graphs in West's book [62]. The uniqueness result was originally proved by Whitney [63].

It follows that the incomparability graph of the posets in \mathbf{P}_G is unique unless $X(G)$ has a triangle component. To settle this exception, let us look at K_3 and $K_{1,3}$ as incomparability graphs. Clearly, the unique poset with K_3 as incomparability graph is the 3-antichain \mathcal{A}_3 . Now, $G(\mathcal{A}_3)$ has three color classes which are pairwise crossing. On the other hand, the unique poset \mathcal{Q} with $K_{1,3}$ as incomparability graph consists of a chain on three elements and an additional element incomparable to the chain. Then $G(\mathcal{Q})$ has three color classes which are pairwise parallel.

Observe that $\text{Incomp}(\mathcal{P})$ consists of different components exactly if \mathcal{P} is a series composition of several subposets. This carries over to the line graph of $\text{Incomp}(\mathcal{P})$. Thus if one component of $X(G)$ is a triangle, then the corresponding series component in a G -compatible poset is either a 3-antichain or a chain on three elements with an additional element. Since G has a unique swap partition, G uniquely determines which one of these two possibilities holds.

This shows that there is a unique graph $H(G)$ which equals $\text{Incomp}(\mathcal{P})$ for some $\mathcal{P} \in \mathbf{P}_G$. It is easy to see that the comparability graph $\text{Comp}(\mathcal{P})$ equals $\text{Incomp}(\mathcal{P})^c$ plus a set of global vertices corresponding to the set of global elements of \mathcal{P} . Thus the unique minimal graph in $\Gamma(G)$ is $H(G)^c$, which we denote by $\text{Comp}(G)$, and the characterization of $\Gamma(G)$ follows. \square

A weaker version of the above result was also stated in Reuter's article [50]. However, it was based on a flawed characterization of adjacent colors, as mentioned in Section 2.3.

Let H be a comparability graph. Then we say that a bijection c between the color classes of G and the non-adjacent pairs of vertices of H is a *G -compatible swap coloring for H* if there is a poset $\mathcal{P} \in \mathbf{P}_G$ with $H = \text{Comp}(\mathcal{P})$ such that c is the swap coloring of G with respect to \mathcal{P} .

Lemma 3.7. *The G -compatible swap coloring for $\text{Comp}(G)$ is unique up to automorphisms of $\text{Comp}(G)$.*

Proof. We use the notation of the previous proof. We identify the swap colors c_i of G with the vertices of $X(G)$. The cliques of the Krausz decomposition of $X(G)$ become the vertices of $H(G)$. For a non-triangle component C of $X(G)$, we assign $c_i \in V(X(G))$ to the unique edge of $H(G)$

connecting the two cliques in which it appears. By construction, this is the unique assignment preserving the adjacencies of the swap colors.

A triangle component C of $X(G)$ corresponds to a component \bar{C} of $H(G)$ which is either K_3 or $K_{1,3}$. Hence, relabeling the edges of \bar{C} is an automorphism of $H(G)$. We thus assign the three swap colors of C arbitrarily to the edges of \bar{C} .

It follows that the assignment of the c_i to edges of $H(G)$ preserving the adjacencies of the swap colors is unique up to automorphisms of $H(G)$. This is equivalent to the fact that there is a unique G -compatible swap coloring for $\text{Comp}(G)$ up to automorphisms of $\text{Comp}(G)$. \square

Recall that G has r swap colors, hence $X(G)$ has r vertices. Then $H(G) = \text{Comp}(G)^c$ has r edges and may have $\theta(r)$ vertices. Thus we assume for the following that $\text{Comp}(G)$ has $O(r)$ vertices.

Lemma 3.8. *Given the graph $X(G)$, we can reconstruct $\text{Comp}(G)$ and the G -compatible swap coloring for $\text{Comp}(G)$ in time $O(r^2 + r \cdot q)$.*

Proof. We consider the components of $X(G) = (V_X, E_X)$ separately. For a component C which is not a triangle, we use an algorithm by Lehot [38] which takes a line graph $X \neq K_3$ as input and reconstructs the unique graph H such that X is the line graph of H . The algorithm runs in time $O(|V(X)| + |E(X)|)$. It works by assigning the vertices of X to the cliques of the Krausz decomposition. Thus, it also yields the part of the G -compatible swap coloring corresponding to C .

For a triangle component C of $X(G)$, we pick $c_i, c_j \in V(C)$ arbitrarily. We pose Question 1 to our G -medium to find out whether θ_i and θ_j are crossing or parallel. By the proof of the previous lemma, this determines the corresponding component \bar{C} of $H(G)$. For obtaining c , we have seen that we can assign the swap colors of C arbitrarily to the three edges of \bar{C} .

Taking care of the non-triangle components of $X(G)$ takes at most $O(|E_X| + |V_X|) = O(r^2)$ time altogether. If $X(G)$ contains a linear number of triangle components, these take time $O(r \cdot q)$. Thus we can reconstruct $H(G)$ in time $O(r^2 + r \cdot q)$. From this, $\text{Comp}(G)$ can be obtained in $O(r)$ by taking the complement. \square

If the linear extension graph was a comparability invariant, then we would have completed our reconstruction procedure by now: The set of G -compatible posets would be given by the set of all transitive orientations of $\text{Comp}(G)$. However, we have seen in Section 2.4 that the linear extension graph is not a comparability invariant. Therefore we have to continue to find out which transitive orientations of $\text{Comp}(G)$ are G -compatible.

3.4 Orienting the Comparability Graph

In this section we show how to find all G -compatible transitive orientations of $\text{Comp}(G)$, that is, all transitive orientations yielding a poset $\mathcal{P} \in \mathbf{P}_G$. Since $\text{Comp}(G)$ is the minimal G -compatible comparability graph, we denote the family of these posets by \mathbf{P}_G^{\min} . By Lemma 3.6, the family \mathbf{P}_G of all G -compatible posets can be obtained from the family \mathbf{P}_G^{\min} by the addition of global elements to the posets in \mathbf{P}_G^{\min} . Hence this section contains the last (and most involved) step of our reconstruction procedure.

From the preceding section, we are also given a G -compatible swap coloring c for $\text{Comp}(G)$ which is unique up to automorphisms of $\text{Comp}(G)$. We are only interested in posets $\mathcal{P} \in \mathbf{P}_G^{\min}$ respecting c , that is, such that c is the swap coloring of G with respect to \mathcal{P} . It follows from Lemma 3.6 that posets in \mathbf{P}_G^{\min} with a different swap coloring correspond to automorphisms of the posets respecting c .

It will turn out that in most cases, the poset in \mathbf{P}_G^{\min} is essentially unique. In this section, essentially unique means unique up to automorphisms and duality. Recall that we use the notion *direction* for the difference between a poset \mathcal{P} and its dual \mathcal{P}^* .

Let $\text{Comp}(G) = (V_C, E_C)$ and recall that $|V_C| \in O(r)$. If $\text{Comp}(G)$ has exactly two transitive orientations, then $\text{Comp}(G)$ is uniquely partially orderable. In this case, we can use the algorithm of McConnell and Spinrad [40] to compute its transitive orientation in time $O(|V_C| + |E_C|)$. Hence we can complete the last step without the help of a G -medium.

In the general case, we compute the modular decomposition tree T of $\text{Comp}(G)$, see Section 1.4. This can be done with the algorithm of [40] in time $O(r)$. Each node of T corresponds to a module of $\text{Comp}(G)$. The root of T corresponds to V_C , and each leaf of T is a single vertex of $\text{Comp}(G)$. The nodes of T are labelled (\parallel), (S) or (P), according to which case of Theorem 1.17 applies to their canonical partition.

The case that the root of T is an (S)-node is the only exception where the poset in \mathbf{P}_G^{\min} is not essentially unique. We will deal with this case in the last subsection. The main work is to prove the following proposition:

Proposition 3.9. *Assume that we are given the minimal G -compatible comparability graph $\text{Comp}(G)$ and the G -compatible swap coloring c for $\text{Comp}(G)$, as well as the modular decomposition tree T of $\text{Comp}(G)$.*

If the root node of T is not an (S)-node, then \mathbf{P}_G^{\min} contains an essentially unique poset, and we can construct it in time $O(r^3q)$.

Proof. We construct a G -compatible transitive orientation of $\text{Comp}(G)$ by starting with the leaves of T and then working our way upwards.

For each type of node there will be a subsection explaining how to process it. Before going into details, we want to give a rough explanation of how the orientation procedure works. The details will then be explained in three subsections.

Because $\text{Comp}(G)$ is a comparability graph, each induced subgraph of it is also a comparability graph and thus has a transitive orientation. In particular, this holds for a subgraph induced by a node of T , which is a module M of G . We will see that each subgraph $G[M]$ has only one G -compatible transitive orientation (up to its direction). This yields a poset $\mathcal{P}(M)$ which we will also call G -compatible.

Note that a transitively oriented graph defines a poset and vice versa. In this proof, we often switch between the two concepts, in particular between a transitively oriented graph $G[M]$ and the corresponding poset $\mathcal{P}(M)$. We use notions for these objects interchangeably.

We start with processing the leaves. Because the leaf-nodes are single-vertex graphs, there is nothing to do to choose a transitive orientation. In each further step, we choose a node M of T such that we have already processed all children of M . The children of M in T are the vertex sets M_1, \dots, M_t of the canonical partition of M . Thus, for every child M_i of M we have already found the unique G -compatible transitive orientation of $G(M_i)$. This yields the posets $\mathcal{P}(M_1), \dots, \mathcal{P}(M_t)$.

Now we construct a G -compatible transitive orientation of the partition graph $G[M]^\#$ of $G[M]$. This yields a poset $\mathcal{P}(M^\#) = \mathcal{P}(G[M]^\#)$ which we want to use to obtain $\mathcal{P}(M)$. By Theorem 1.23, each transitive orientation of $G[M]$ arises from replacing every vertex of $G[M]^\#$ by the corresponding child of M in T and directing all edges between two children according to the edge between the corresponding vertices of $G[M]^\#$. The only thing we need to check is in which direction we insert the children of M . That means, to obtain $\mathcal{P}(M)$, we need to decide whether to replace an M_i in $\mathcal{P}(M^\#)$ by $\mathcal{P}(M_i)$ or by $\mathcal{P}(M_i)^*$.

We will see that the direction of the first inserted child forces the direction of all other children! In other words, once we fixed the direction of some $\mathcal{P}(M_i)$ in $\mathcal{P}(M)^\#$, we fixed the direction of $\mathcal{P}(M)$, and every other $\mathcal{P}(M_j)$ has only one direction which is compatible with G and the chosen direction of $\mathcal{P}(M_i)$. (This holds, of course, unless $\mathcal{P}(M_i)$ is self-dual. But then the other direction corresponds to an automorphism of $\text{Comp}(G)$, and will thus be ignored.)

This motivates the following definition: We say that $\mathcal{P}(M_i)$ and $\mathcal{P}(M_j)$ *interlock* if their directions are G -compatible, that is, if there is a G -compatible poset in which they both appear in the given direction. The crucial tool of this proof is the Interlocking Lemma stated below.

For the Interlocking Lemma, we need the notion of a *pivot* for a pair M_i, M_j , which is a vertex $v \in \text{Comp}(G)$ such that v is non-adjacent in $\text{Comp}(G)$ to all elements of M_i and M_j . Also, for each M_i we need a minimal element m_i^- and maximal element m_i^+ of $\mathcal{P}(M_i)$, such that $m_i^- \leq m_i^+$ in $\mathcal{P}(M_i)$. If $\mathcal{P}(M_i)$ has no relations, then m_i^- and m_i^+ coincide.

Interlocking Lemma. *Let M_i and M_j be two children of M , and $\mathcal{P}(M_i)$ and $\mathcal{P}(M_j)$ the corresponding G -compatible posets.*

- (i) *If $M_i \parallel M_j$ in $\mathcal{P}(M^\#)$, then it is uniquely determined by G in which directions $\mathcal{P}(M_i)$ and $\mathcal{P}(M_j)$ interlock. We can determine it in time $O(q)$.*
- (ii) *If $M_i \sim M_j$ in $\mathcal{P}(M^\#)$ and there exists a pivot v for the pair M_i, M_j , then it is uniquely determined by G whether $M_i \sim M_j$ is a cover relation of $\mathcal{P}(M^\#)$, and if so, in which directions $\mathcal{P}(M_i)$ and $\mathcal{P}(M_j)$ interlock. Given v , we can determine it in time $O(q)$.*

Proof. For case (i), observe that m_i^σ and $m_j^{\sigma'}$ for $\sigma, \sigma' \in \{+, -\}$ are non-adjacent in $\text{Comp}(G)$. Thus we may consider the color classes of G associated with $m_i^+m_j^-$ and $m_i^-m_j^+$ by c . If there is a poset $\mathcal{P} \in \mathbf{P}_G^{\min}$ respecting c which contains $\mathcal{P}(M_i)$ and $\mathcal{P}(M_j)$ in the given directions, then these two colors are parallel by Lemma 2.12. If not, these two colors are crossing, and we have to reverse the direction of $\mathcal{P}(M_j)$ (or $\mathcal{P}(M_i)$) to make them interlock. We can find out which case holds by posing Question 1 to the G -medium.

For case (ii), let $v \in \text{Comp}(G)$ be a pivot for M_i, M_j , and let $\sigma, \sigma' \in \{+, -\}$. It holds by Lemma 2.12 that $m_i^\sigma \sim m_j^{\sigma'}$ is a cover relation in a poset $\mathcal{P} \in \mathbf{P}_G^{\min}$ respecting c exactly if the colors $m_i^\sigma v$ and $m_j^{\sigma'} v$ touch in G . That is, $M_i \sim M_j$ is a cover relation of $\mathcal{P}(M^\#)$ exactly if there are $\sigma, \sigma' \in \{+, -\}$ such that $m_i^\sigma v$ and $m_j^{\sigma'} v$ are touching.

Thus, by posing Question 3 to our G -medium at most four times, we can find out whether $M_i \sim M_j$ is a cover relation of $\mathcal{P}(M^\#)$, and how M_i and M_j interlock. \triangle

Let us go back to the outline of our orientation procedure. Recall that in every step, we process a node M of the decomposition tree T , and want to turn $\mathcal{P}(M^\#)$ into $\mathcal{P}(M)$ by replacing every M_j by $\mathcal{P}(M_j)$. Using the Interlocking Lemma, we can find the unique direction of $\mathcal{P}(M_j)$ compatible with G and the chosen direction of the first $\mathcal{P}(M_i)$. In the three following subsections, we explain how to do this according to the type of the node M of T .

Thus we construct the unique G -compatible transitive orientation of $G[M]$, yielding a poset $\mathcal{P}(M)$. When we have processed the root of T , we know from Theorem 1.23 that we have constructed a transitive orientation of $\text{Comp}(G)$. In this way we obtain the essentially unique poset in \mathbf{P}_G^{\min} .

Note that it might seem more natural to process the decomposition tree T from the root downwards. Our method only works upwards because in every step, we need the fact that the children of the current node of T are already transitively oriented and thus correspond to a poset.

3.4.1 The Parallel Case

First, let us consider the case that the chosen node M of T is a (\parallel)-node.

Lemma 3.10. *If M is a (\parallel)-node, then the G -compatible poset $\mathcal{P}(M)$ is essentially unique and we can reconstruct it in time $O(r \cdot q)$.*

Proof. By Theorem 1.22, the graph $G[M]^\#$ has no edge. Thus we obtain a transitive orientation of $G[M]^\#$ and hence a poset $\mathcal{P}(M^\#)$ for free. We only need to check which directions we need to insert the posets $\mathcal{P}(M_i)$.

We pick the smallest i such that $G[M_i]$ contains at least one edge (that is, such that $\mathcal{P}(M_i)$ is not an antichain), and fix its direction arbitrarily. For $j = i + 1, \dots, t$, we want to match the direction of $\mathcal{P}(M_j)$ to the direction of $\mathcal{P}(M_i)$. By the Interlocking Lemma, this is uniquely determined by G and can be determined in time $O(q)$.

It follows that there is a unique G -compatible transitive orientation of $G[M]$. To compute it, we have to interlock each pair $M_i M_j$ for $j = i + 1, \dots, t$. There are at most $t - 1$ such pairs, each taking time q . The module M cannot have more children than there are vertices of $\text{Comp}(G)$, thus $t \in O(r)$. Hence a (\parallel)-node can be processed in time $O(r \cdot q)$. \triangle

3.4.2 The Series Case

Next, let us consider the case that M is an (S)-node of T .

Lemma 3.11. *If M is an (S)-node which is not the root node of T , then the G -compatible poset $\mathcal{P}(M)$ is essentially unique and we can reconstruct $\mathcal{P}(M)$ in time $O(r^2 \cdot q)$.*

Proof. It holds by Theorem 1.22 that $G[M]^\#$ is a complete graph, and each transitive orientation of $G[M]^\#$ can be found by choosing a linear order of the children M_1, \dots, M_t of M and directing the edges of $G[M]^\#$ from lower to higher indices in this order.

To find a G -compatible linear order of M_1, \dots, M_t , we want to use part (ii) of the Interlocking Lemma. Since M is not the root node of T , on the path from M to the root of T there is a node M' of T which is not labeled (S). Then M' has a child M'_1 which contains M , and a child M'_2 which is not adjacent to M'_1 in $G[M']^\#$. Any vertex $v \in M'_2$ is non-adjacent in $\text{Comp}(G)$ to all vertices of M . Thus, v can be used as a pivot vertex for all pairs M_i, M_j of children of M . Hence it follows with the Interlocking Lemma that the G -compatible linear order of M_1, \dots, M_t is unique up to direction.

To reconstruct this linear order, we can check for each pair M_i, M_j whether it is a cover relation in $\mathcal{P}(M^\#)$. Once we have done this for every pair, we choose a direction of the resulting linear order to obtain $\mathcal{P}(M^\#)$. Since the Interlocking Lemma also tells us in which direction adjacent pairs of $\mathcal{P}(M^\#)$ interlock, this yields $\mathcal{P}(M)$.

We can find the pivot vertex v in time $O(r)$ by searching the vertices of $\text{Comp}(G)$. Given v , the running time is $O(q)$ for every pair of children of M . Since M has $O(r)$ children, we can build $\mathcal{P}(M)$ in time $O(r^2 \cdot q)$. \triangle

3.4.3 The Prime Case

The most difficult case is when M is a (P)-node.

Lemma 3.12. *If M is a (P)-node, then the G -compatible poset $\mathcal{P}(M)$ is essentially unique, and we can reconstruct it in time $O(r^2 \cdot q)$.*

Proof. It was observed after Theorem 1.22 that if M is a (P)-node, then $G[M]^\#$ has exactly two transitive orientations, where one is the reverse of the other. Construct one of these transitive orientations with the algorithm from [40]. This yields a partial order $\mathcal{P}(M^\#)$ on M_1, \dots, M_t .

Now we have to choose directions of the $\mathcal{P}(M_i)$ which are compatible with G . After relabeling, we may assume that $M_1 M_2 \dots M_t$ is a linear extension of $\mathcal{P}(M^\#)$. We perform t steps, where step j serves to fix the direction of $\mathcal{P}(M_j)$. We start with fixing the direction of $\mathcal{P}(M_1)$ arbitrarily. In each later step j , consider M_j in the poset $\mathcal{P}(M^\#)$. If there is an index $i < j$ such that $M_i \parallel M_j$ in $\mathcal{P}(M^\#)$, then the direction of $\mathcal{P}(M_j)$ can be fixed using part (i) of the Interlocking Lemma. This case certainly applies if $j = t$, otherwise M_t would be a series module of M , which is a contradiction since M is a (P)-node.

In case we have $M_1, \dots, M_{j-1} < M_j$ in $\mathcal{P}(M^\#)$, we want to apply part (ii) of the Interlocking Lemma. If there is a vertex $v \in \text{Comp}(G)$ such that the vertices of M are non-adjacent to v in $\text{Comp}(G)$, then v can

serve as a pivot vertex for the pair M_1, M_j . This allows us to fix the direction of $\mathcal{P}(M_j)$. However, it may happen (if M is the root node of T) that no such v exists. In this case, we have to find a *private pivot* for each M_j .

Let S be the set of immediate predecessors of M_j in $\mathcal{P}(M^\#)$, that is, of predecessors forming a cover relation with M_j . Then by assumption we have $S^\downarrow = \{M_1, \dots, M_{j-1}\}$. If $M_k > M_i$ in $\mathcal{P}(M^\#)$ for all M_i, M_k with $M_i \in S$ and $k > j$, then S^\downarrow is a series module of $\mathcal{P}(M^\#)$. This is a contradiction since M is a prime node in T . Thus there is an $M_i \in S$ and an M_k with $k > j$ such that $M_i \parallel M_k$ in $\mathcal{P}(M^\#)$. It follows that $M_j \parallel M_k$ in $\mathcal{P}(M^\#)$. Hence we can use a vertex from M_k as a pivot for the pair M_i, M_j in order to fix the direction of $\mathcal{P}(M_j)$.

Thus in each step j , the direction of $\mathcal{P}(M_j)$ is uniquely determined. After fixing the direction of $\mathcal{P}(M_t)$, we have constructed a G -compatible poset $\mathcal{P}(M)$. We have also seen that it is essentially unique.

It remains to calculate the running time for the (P)-case. A transitive orientation of $G[M]^\#$ can be computed in time linear in $|M|$ with the algorithm of [40], and thus in $O(r)$. As for fixing the directions of the $\mathcal{P}(M_j)$, the worst case happens if we need to find a private pivot for each j . A private pivot can be found by searching $\text{Comp}(G)$ and hence in time $O(r)$. Thus it follows from the Interlocking Lemma that each step can be completed in time $O(r \cdot q)$. There are $t \in O(r)$ steps in total. Thus a (P)-module can be processed in time $O(r^2 \cdot q)$. \triangle

Completing the proof of Proposition 3.9.

We have seen in the three subsections above that we can process every node M of T to construct a G -compatible poset $\mathcal{P}(M)$. In every case, $\mathcal{P}(M)$ is essentially unique. Now it follows from Theorem 1.23 that \mathbf{P}_G^{\min} contains an essentially unique poset.

We complete the proof of Proposition 3.9 by checking the overall running time. Each node M_i of T has a unique father in T . When this father is processed, we need a minimal element m_i^- and a maximal element m_i^+ of $\mathcal{P}(M_i)$ such that $m_i^- \leq m_i^+$. Clearly, we can find m_i^- and m_i^+ in time linear in $|M_i|$. Recall that the modular decomposition tree has a linear number of nodes. Thus, the number of M_i we consider is in $O(r)$. Hence we can find all elements m_i^σ with $\sigma \in \{+, -\}$ in a total time of $O(r^2)$.

After this, every node of T can be processed in time $O(r^2 \cdot q)$. Thus the time to construct the G -compatible transitive orientation of $\text{Comp}(G)$ is bounded by $O(r^3 \cdot q)$. \square

3.4.4 The Exceptional Case

Now let us finally consider the case that the root of T is an (S)-node.

Lemma 3.13. *Suppose that M is the root node of T and that it is an (S)-node with children M_1, \dots, M_t . Then \mathbf{P}_G^{\min} consists of the posets formed by any series composition of the posets $\mathcal{P}(M_1), \dots, \mathcal{P}(M_t)$, each taken in arbitrary direction.*

Proof. If the root node M of T is a series node, this means by Definition 1.25 that a poset $\mathcal{P}(M)$, and thus a poset in \mathbf{P}_G^{\min} , is a series composition of the posets $\mathcal{P}(M_i)$. By Corollary 2.23, the linear extension graph is invariant under series alteration, that is, under a change of the order or the direction of the series modules. The result follows. \square

The lemma above yields a precise description of all G -compatible posets in the case that the root of T is an (S)-node. Thus in this case, we know the posets in \mathbf{P}_G^{\min} already after processing all the children of the root. There can be up to $2^t t!$ such posets (their description is obviously only implicit).

3.5 Putting Things Together

Completing the proof of Theorem 3.1.

Recall that once we know all posets in \mathbf{P}_G^{\min} , we obtain the set \mathbf{P}_G by adding an arbitrary number of global elements to each $\mathcal{P} \in \mathbf{P}_G^{\min}$.

To obtain the running time of our reconstruction procedure we need to look at the running times of the different steps. The step taking the longest time is to find the adjacencies of the swap colors, namely, $O(r^4 \cdot q)$. This can probably be improved.

We have thus shown that with a G -medium answering admissible questions in time q , we can construct the set \mathbf{P}_G of all G -compatible posets in time $O(r^4 \cdot q)$. This concludes the proof of Theorem 3.1. \square

The case that the root of T is an (S)-node makes a big difference for the set of G -compatible posets. Note that deciding whether or not this case holds can already be done by looking at G : It follows from Lemma 3.13 and Proposition 2.21 that the root of T is an (S)-node exactly if G is a Cartesian product of several non-trivial graphs. Hence the G -compatible poset is unique up to direction and the addition of global elements exactly if G is Cartesian prime. We have thus proved Theorem 2.24.

To finish the discussion of the reconstruction procedure, let us mention that Sabidussi [53] and, independently, Vizing [60] showed that every finite connected graph has a Cartesian factorization which is unique up to

order and isomorphism of the factors. The unique Cartesian factorization of a graph can be found in linear time with an algorithm of Imrich and Peterin [34].

Hence another approach to our reconstruction algorithm would be to first ask the G -medium to find the Cartesian factors G_1, \dots, G_f of G . Then each G_i would have an essentially unique G_i -compatible poset. The set of G -compatible posets then equals the series compositions of the G_i -compatible posets in any order and direction, completed by some global elements. However, if the G -medium has access to a G -compatible poset \mathcal{P} and can thus answer the admissible questions in constant time, we can avoid looking at the whole graph G . In this case it is not favorable to start with the Cartesian factorization.

3.6 Recognizing Linear Extension Graphs

In the last section of this chapter, we show that the recognition procedure developed in the previous sections can also be used to recognize whether an arbitrary given graph is a linear extension graph. Essentially, it works by applying the reconstruction procedure to the given graph G . If the procedure gets stuck at some point, we know that G is not a linear extension graph. Otherwise, we obtain a candidate \mathcal{P} for a poset in \mathbf{P}_G . We then try to match the linear extensions of \mathcal{P} and the vertices of G . If this is successful, we have a labeling of G with the linear extensions of \mathcal{P} . Thus $G = G(\mathcal{P})$ and hence G is a linear extension graph.

Theorem 3.14. *Whether a given graph $G = (V, E)$ is a linear extension graph can be recognized in time $O(|V|^2 + \dim_I(G)^4 \cdot |E|)$.*

Proof. We partition the proof into several parts. In Subsection 3.6.1, we construct the candidate \mathcal{P} for a G -compatible poset, after ensuring that \mathcal{P} is essentially unique. In Subsection 3.6.2, we choose a vertex $v \in G$ and construct the linear extension L_v of \mathcal{P} which corresponds to v if $G = G(\mathcal{P})$. In Subsection 3.6.3, we try to match the linear extensions of \mathcal{P} to the vertices of G , starting from L_v .

3.6.1 Constructing a Candidate Poset

In order to recognize if G is a linear extension graph, we cannot avoid looking at the whole graph G . Thus without increasing the running time we may start with applying the linear time algorithm by Imrich and Peterin [34] which decomposes G into its Cartesian factors. It follows from

Proposition 2.21 and Theorem 2.24 that G is a linear extension graph exactly if each of its Cartesian factors is a linear extension graph. So for the rest of this proof, we will assume that G is Cartesian prime.

Next, we check whether G is a partial cube using Eppstein's Algorithm from Theorem 1.12. If the answer is negative, then G is not a linear extension graph by Theorem 2.4. If the answer is positive, the algorithm also computes the swap partition $\Theta(G)$.

Now we follow the steps of our reconstruction procedure. In this case, our G -medium is just the graph G . We check adjacencies of the swap colors as shown in Section 3.2. After that, we build the unique candidate $\text{Comp}(G) = (V_C, E_C)$ for a minimal G -compatible comparability graph as shown in Section 3.3.

At this point, we test whether $\text{Comp}(G)$ is a comparability graph. This can be done with an algorithm by Golumbic [27]. If $\text{Comp}(G)$ is not a comparability graph, we know that G is not a linear extension graph.

If $\text{Comp}(G)$ is a comparability graph, then we continue with computing the modular decomposition tree T of $\text{Comp}(G)$. Since G is Cartesian prime, it follows from Proposition 2.21 that the root node of T is not a series node. Thus by Proposition 3.9, the candidate \mathcal{P} for a poset in \mathbf{P}_G is essentially unique.

We try to reconstruct \mathcal{P} by finding a G -compatible transitive orientation of $\text{Comp}(G)$ with the methods of Section 3.4. If some part of the orientation algorithm cannot be carried out, e.g., we do not find a pivot to apply the Interlocking Lemma, or we do not obtain a linear order of the modules in in Subsection 3.4.2, then G is not a linear extension graph. Otherwise, we obtain a poset \mathcal{P} , which is the essentially unique poset in \mathbf{P}_G^{\min} if G is a linear extension graph.

As for the running time, the partial cube recognition algorithm takes time $O(|V|^2)$. Recall that the number r of swap colors of G equals the isometric dimension $\dim_I(G)$ of G . Thus $|V_C|, |\mathcal{P}| \in O(\dim_I(G))$. Golumbic's recognition algorithm for comparability graphs runs in time $O(\Delta \cdot |E_C|)$, where Δ is the maximum degree of $\text{Comp}(G)$. Thus $\text{Comp}(G)$ can be tested in time $O(\dim_I(G)^3)$, which is dominated by the running time of the reconstruction.

By Lemma 3.4, we can answer each admissible question in time $O(|V| + |E|)$. Therefore it follows from Theorem 3.1 that we can reconstruct \mathcal{P} in time $O(\dim_I(G)^4 \cdot |E|)$. Altogether, the first part of the recognition algorithm takes time $O(|V|^2 + \dim_I(G)^4 \cdot |E|)$.

3.6.2 Finding a Start

If all the steps in Subsection 3.6.1 worked out, we now have a poset \mathcal{P} which is the essentially unique candidate for a poset in \mathbf{P}_G^{\min} . By our construction, it comes with a candidate c for a swap coloring of G with respect to \mathcal{P} which is unique up to automorphisms of \mathcal{P} . Thus, c is a bijection between $\Theta(G)$ and $\text{Inc}(\mathcal{P})$.

Let c^{-1} be the inverse map of c . Recall that for the linear extension graph $G(\mathcal{P})$, we denote the color class corresponding to $xy \in \text{Inc}(\mathcal{P})$ by $\theta(xy)$. If $G = G(\mathcal{P})$, then $c^{-1}(xy) = \theta(xy)$ for each $xy \in \text{Inc}(\mathcal{P})$.

Choose a vertex v of G . Our aim is to find a linear extension L_v of \mathcal{P} which corresponds to v in case $G = G(\mathcal{P})$. We know that L_v is unique up to automorphisms of \mathcal{P} . To find it, we consider the swap colors of the edges which are incident to v in G . Let v be incident to the edges e_1, \dots, e_k in G . Denote the color of e_i by $c_i = x_i y_i$. Then the pairs $x_i y_i$ are the *jumps* of \mathcal{P} . Our plan is to determine for each i which of the two relations, $x_i < y_i$ or $x_i > y_i$, is valid in L_v . Then we obtain L_v by adding these relations to \mathcal{P} , as is proved in the lemma below.

Lemma 3.15. *Let \mathcal{P} be a poset and $L = x_1 x_2 \dots x_n$ a linear extension of \mathcal{P} . Let x_{i_j}, x_{i_j+1} be the jumps in L , for $j \in J \subseteq \{1, \dots, n-1\}$. Then it holds that $\mathcal{P}^+ = \mathcal{P} + \{(x_{i_j}, x_{i_j+1}), j \in J\} = L$.*

Proof. After adding the relations of all jumps to \mathcal{P} , we have $x_i < x_{i+1}$ in \mathcal{P}^+ for any $i = 1, \dots, n-1$. Hence $\mathcal{P}^+ = L$. \triangle

Now we proceed as follows: For each i , consider the elements x_i and y_i in \mathcal{P} , and look for a pivot element which distinguishes them. Here, an element $v \in \mathcal{P}$ is a *pivot element* if it is comparable to one of x_i and y_i , and incomparable to the other. Suppose there is such a v , and suppose that $v \sim x_i$ and $v \parallel y_i$. We consider the two cases $v < x_i$ and $x_i < v$.

If $v < x_i$ in \mathcal{P} , then $x_i < y_i$ forces $v < y_i$ in a linear extension of \mathcal{P} . Thus, the color classes $\theta(x_i y_i)$ and $\theta(v y_i)$ are parallel in $G(\mathcal{P})$, cf. Figure 3.2. The class $\theta(v y_i)$ is contained in $G_{y_i x_i}$, that is, in the component of $G(\mathcal{P}) - \theta(x_i y_i)$ consisting of the linear extensions of \mathcal{P} in which $y_i < x_i$. To check whether $x_i < y_i$ or $x_i > y_i$ has to hold in L_v , we check on which side of $G - c^{-1}(x_i y_i)$ we find v and $c^{-1}(v y_i)$. Then, $x_i < y_i$ holds in L_v exactly if v lies on the other side of $c^{-1}(x_i y_i)$ as the class $c^{-1}(v y_i)$.

If $x_i < v$ in \mathcal{P} , then $y_i < x_i$ forces $y_i < v$ in a linear extension of \mathcal{P} . Again, the color classes $\theta(x_i y_i)$ and $\theta(v y_i)$ are parallel in $G(\mathcal{P})$. This time the class $\theta(v y_i)$ is contained in $W_{x_i y_i}$. Hence we have $x_i < y_i$ in L_v exactly if v lies on the same side of $c^{-1}(x_i y_i)$ as the class $c^{-1}(v y_i)$.




Figure 3.2: The position of a color class $\theta(vy_i)$ parallel to $\theta(x_iy_i)$ can be used to determine which side of $\theta(x_iy_i)$ contains the linear extensions with $x_i < y_i$.

If there is no pivot element to distinguish between x_i and y_i , then any other element $v \in \mathcal{P}$ is either larger than both these elements, or smaller than both, or incomparable to both. This means that x_iy_i is a twin of \mathcal{P} (cf. Definition 1.24). By Lemma 2.13, it follows that the two components of $G(\mathcal{P}) - \theta(x_iy_i)$ are isomorphic. Thus if $G = G(\mathcal{P})$, then the two components of $G - c^{-1}(x_iy_i)$ are isomorphic. In this case we can decide arbitrarily whether $x_i < y_i$ or $x_i > y_i$ in L_v .

As for the running time of this second part, there may be $O(\dim_I(G))$ edges incident to v in G , thus we may have to consider $O(\dim_I(G))$ pairs x_iy_i . For each pair, we can find a pivot element (or decide that there is none) in time $O(\dim_I(G))$, by checking for every $v \neq x_i, y_i$ its relations to x_i and y_i . Building the graph $G - c^{-1}(x_iy_i)$ and searching its two components for v and $c^{-1}(vy_i)$ can be done in time $O(|V| + |E|)$. Thus the second part of the recognition can be completed in time $O(\dim_I(G)^2 \cdot (|V| + |E|))$.

3.6.3 Labeling Vertices with Linear Extensions

Now that we have found a linear extension L_v of \mathcal{P} corresponding to v in case $G = G(\mathcal{P})$, our plan is to generate all linear extensions of \mathcal{P} and label the vertices of G with them in a breadth first search fashion. If L is a linear extension of \mathcal{P} and $x, y \in \text{Inc}(\mathcal{P})$ are adjacent in L , then we denote by L^{xy} the linear extension arising from L by swapping x and y .

We start with v and label it with L_v . For the following, we use a queue to keep track of the linear extensions we still need to process. Initially, the queue only contains L_v .

In every step, we pick the first linear extension L from the queue and find all incomparable pairs of \mathcal{P} which are adjacent in L . For each such incomparable pair xy , check if v has an incident edge $e \in c^{-1}(xy)$ in G . If the answer is no, then $G \neq G(\mathcal{P})$. If the answer is yes, then we try to label the neighbor v' of v along e with the linear extension L^{xy} . This can be done without problem if v' has not yet received a label.

If v' is already labeled with a linear extension L' , then we check if $L' = L^{xy}$. If not, then $G \neq G(\mathcal{P})$. If the labeling was successful, then add L^{xy} to the queue. When all adjacent incomparable pairs in L have been looked at, we delete L from the queue. Now we pick the first linear extension in the queue and process it in the same way as L .

When the queue is empty, we have processed all linear extensions of \mathcal{P} , because any linear extension of \mathcal{P} can be reached from L_v via a sequence of adjacent swaps. Now we check if there are still unlabeled vertices in G . If yes, then $G \neq G(\mathcal{P})$. If not, then we have correctly labeled the vertices of G with the linear extensions of \mathcal{P} , and thus it holds that $G = G(\mathcal{P})$. Hence, G is a linear extension graph.

Let us consider the running time of the above labeling procedure. In the case that $G = G(\mathcal{P})$, there are $|V(G)|$ linear extensions which need to be processed. For each one of them, we need to run through the $O(\dim_I(G))$ elements of \mathcal{P} to check which incomparable pairs are adjacent. On the other hand, as soon as the algorithm exhibits a swap between two linear extensions which is not compatible with G , it stops. Thus we do not perform more than $O(|V(G)|)$ steps. Hence the running time of the labeling procedure is $O(\dim_I(G) \cdot |V(G)|)$.

The overall running time of the recognition algorithm is clearly dominated by the $O(|V|^2 + \dim_I(G)^4 \cdot |E|)$ needed for the first part. This concludes the proof of Theorem 3.14. \square

To close this chapter, let us remark that even if $\dim_I(G)$ is expected to be small compared to the size of the whole graph G , it can be as large as $|E(G)|$ (e.g. if G is a path). Hence the running time of the reconstruction and the recognition procedures presented in this chapter can be very long. This also holds because linear extension graphs are typically exponentially large in the size of the underlying poset.

In the next chapters, we are interested in properties of $G(\mathcal{P})$ that we can determine in time polynomial in $|\mathcal{P}|$.

Chapter 4

Complexity of Linear Extension Diameter

In the preceding two chapters we analyzed the structure of linear extension graphs and the connections to properties of the corresponding posets. For the remainder of this thesis, we will mainly be interested in one particular parameter of the linear extension graph: its diameter. The diameter of $G(\mathcal{P})$ equals the *linear extension diameter* of \mathcal{P} .

The diameter of a graph is a parameter which is not very complicated to determine: We can use any shortest path algorithm to compute the longest shortest distance between two vertices in the graph. But since $G(\mathcal{P})$ is typically exponentially large in the size of \mathcal{P} , we are interested in results about the linear extension diameter of \mathcal{P} which we can obtain without looking at the whole graph $G(\mathcal{P})$. That is, we view the linear extension diameter as a parameter of the underlying poset. This comes out more clearly in the following definition:

Definition 4.1 ([22]). *The linear extension diameter of a poset \mathcal{P} , denoted by $\text{led}(\mathcal{P})$, is the maximum distance between two linear extensions of \mathcal{P} . A diametral pair of linear extensions of \mathcal{P} , or in short diametral pair of \mathcal{P} , is a pair of linear extensions of \mathcal{P} achieving this maximum distance. If L, L' is a diametral pair of \mathcal{P} , then L' is a diametral partner of L .*

Recall that the distance between two linear extensions L, L' of \mathcal{P} is the number of reversals between L and L' . Since only incomparable pairs of elements can form a reversal, we have $\text{dist}(L, L') \leq \text{inc}(\mathcal{P})$. It follows that $\text{led}(\mathcal{P}) \leq \text{inc}(\mathcal{P})$ for any poset \mathcal{P} .

A diametral pair L, L' can be used to obtain a drawing of \mathcal{P} which is in some sense optimal: Use L and L' on the two coordinate axes to get a position in the plane for each element of \mathcal{P} . Since the number of incomparable pairs of \mathcal{P} which appear in different orders in L and L' is maximized, the resulting drawing has a minimal number of pairs which are comparable in the dominance order, but incomparable in \mathcal{P} . Figure 4.1 illustrates this for the example of the Chevron.

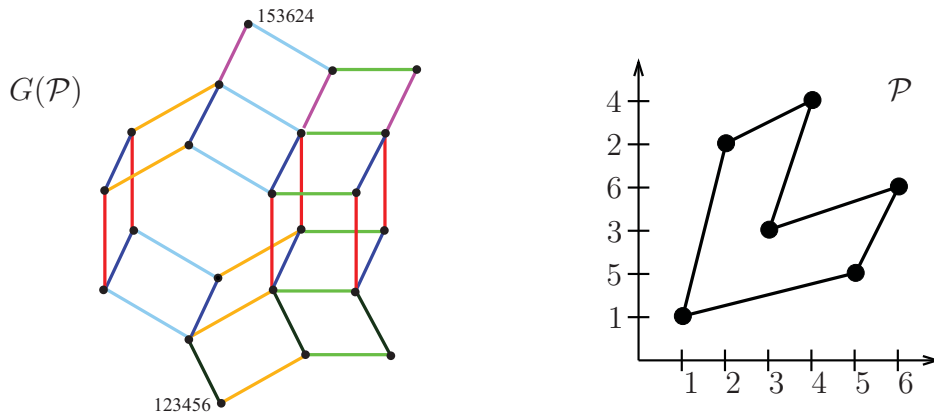


Figure 4.1: The linear extension graph of the Chevron with a diametral pair and the resulting drawing of the Chevron.

In Section 4.1, we revise some known results about the linear extension diameter. In Section 4.2 of this chapter we prove that, in general, computing the linear extension diameter in polynomial time is NP-complete. In the third and last section of this chapter, we prove that the linear extension diameter of posets of width 3 can be computed in polynomial time.

4.1 Previous Results

In this section we give a short overview of previously known result about the linear extension diameter. They have been proved by Felsner and Reuter in [22].

There are very few exact results about the linear extension diameter. Let us start with considering the easiest cases: If \mathcal{P} is a chain, then $G(\mathcal{P})$ consists of a single vertex, so its diameter is 1. If the poset is the antichain \mathcal{A}_n , then a diametral pair consists of two permutations of the elements such that one is the reverse of the other. Hence $\text{led}(\mathcal{A}_n) = \binom{n}{2}$.

If \mathcal{P} has dimension 2, then by definition it has a realizer $\mathcal{R} = \{L, L'\}$ such that every incomparable pair $x \parallel y$ of elements of \mathcal{P} appears in different orders in L and L' . So $\text{led}(\mathcal{P})$ equals the number of incomparable pairs of \mathcal{P} . The converse is also true: If we have $\text{led}(\mathcal{P}) = \text{inc}(\mathcal{P})$ for some poset, then a diametral pair of linear extensions of \mathcal{P} yields a realizer of \mathcal{P} , and thus $\text{dim}(\mathcal{P}) = 2$. We have thus proved the following:

Theorem 4.2 ([22]). *Let \mathcal{P} be a poset. Then $\text{led}(\mathcal{P}) = \text{inc}(\mathcal{P})$ holds exactly if $\text{dim}(\mathcal{P}) = 2$.* □

Note that the Chevron has seven incomparable pairs, but its linear extension diameter is only six. It follows that its dimension is at least 3.

Recall from Section 1.2 that the dimension of a poset \mathcal{P} can be characterized as the minimum d such that \mathcal{P} has a proper embedding into \mathbb{Z}^d , that is, an embedding such that the dominance order on \mathbb{Z}^d equals the order relation of \mathcal{P} . From the above theorem it follows that the drawing of \mathcal{P} induced by a diametral pair is a proper embedding into the plane exactly if \mathcal{P} is 2-dimensional.

Let us consider the Chevron again: The drawing on the right of Figure 4.1 is not a proper embedding into the plane, since the pair 1, 3 of elements is comparable in the dominance order, but not in the Chevron. However, this is the only pair which is not embedded properly.

Felsner and Reuter [22] give a number of lower bounds on the linear extension diameter, relating $\text{led}(\mathcal{P})$ to the width, dimension and fractional dimension of \mathcal{P} . They are shown by rather straightforward constructions of pairs of linear extensions. As for upper bounds, $\text{led}(\mathcal{P}) \leq \text{inc}(\mathcal{P})$ can be refined to $\text{led}(\mathcal{P}) \leq \text{inc}(\mathcal{P}) - (\text{dim}(\mathcal{P}) - 2)$. This bound is tight for the standard examples.

There are also some results in [22] about the behavior of $\text{led}(\mathcal{P})$ under removals of elements or relations. For an element $x \in \mathcal{P}$ denote by $\text{inc}(x)$ the number of elements of \mathcal{P} which are incomparable to x . Then it holds that $\text{led}(\mathcal{P}) \geq \text{led}(\mathcal{P} - x) \geq \text{led}(\mathcal{P}) - \text{inc}(x)$. For a cover relation $r = (x < y)$ of \mathcal{P} , we have $\text{led}(\mathcal{P}) \leq \text{led}(\mathcal{P} - r) \leq \text{led}(\mathcal{P}) + \min\{\text{inc}(x), \text{inc}(y)\} + 1$. Most of these inequalities are sharp, for which examples are given in [22]. Again, the inequalities are quite straightforward to prove. Since our considerations go into a different direction, we do not go into more details here.

A result that we will use concerns the situation of modules in \mathcal{P} . Recall that a module of \mathcal{P} is a subset of the elements of \mathcal{P} which cannot be distinguished from the outside, see Definition 1.24. For completeness, we provide a proof of the following lemma:

Lemma 4.3 ([22]). *For every poset \mathcal{P} and module M of \mathcal{P} , there is a diametral pair L, L' of \mathcal{P} such that the elements of M appear successively in L and L' .*

Proof. Consider a diametral pair L, L' of linear extensions of \mathcal{P} . Each element $x \in \mathcal{P}$ contributes a certain number of reversals to the distance of L and L' . Let x be an element of M which contributes a maximum number of reversals with the elements not in M . Now we move all the other elements of M to the position of x in L and L' without changing their internal order. Since M is a module, this cannot violate any relation of \mathcal{P} , thus the result will again be two linear extensions of \mathcal{P} . By the choice of x , their distance cannot have decreased. So we have constructed a diametral pair of \mathcal{P} such that the elements of M appear consecutively in both linear extensions. \square

Recall that a comparability invariant is a poset parameter that depends only on the comparability graph of the poset. We have seen in Section 2.4 that the linear extension graph is not a comparability invariant. With the above result, it can be shown easily that the linear extension diameter is indeed a comparability invariant.

Theorem 4.4 ([22]). *The linear extension diameter is a comparability invariant.*

Proof. By Theorem 1.26, it suffices to prove that the linear extension diameter remains unchanged whenever a subposet induced by a module is replaced by its dual, that is, if $\text{led}(\mathcal{P}_x^{\mathcal{Q}}) = \text{led}(\mathcal{P}_x^{\mathcal{Q}^*})$ for all choices of posets \mathcal{P}, \mathcal{Q} and an element $x \in \mathcal{P}$.

Since the ground set of \mathcal{Q} is a module in $\mathcal{P}_x^{\mathcal{Q}}$, by Lemma 4.3 there is a diametral pair L, L' of linear extensions of $\mathcal{P}_x^{\mathcal{Q}}$ in which the elements of \mathcal{Q} appear successively. If we reverse the order of the elements of \mathcal{Q} in L and L' , we obtain two linear extensions of $\mathcal{P}_x^{\mathcal{Q}^*}$. Their distance is the same as the distance between L and L' . Thus we have $\text{led}(\mathcal{P}_x^{\mathcal{Q}}) \leq \text{led}(\mathcal{P}_x^{\mathcal{Q}^*})$. With the same argument, the converse also holds. \square

One of the main results in [22] is an exact formula for the linear extension diameter of generalized crowns. To state it, we first give some definitions.

The *crown* C_n was defined in [2] as the height-2 poset consisting of two antichains $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ such that $a_i < b_j$ in S_n exactly if $j \in \{i, i + 1\}$. Here, the indices of the a_i and b_i are to be interpreted cyclically.

The *generalized crowns*, defined by Trotter [58], form a class of posets interpolating between the standard examples and the crowns. We use the

definition from [22]: The generalized crown C_n^k is a height two poset and has as ground set the two antichains a_0, \dots, a_{n-1} and b_0, \dots, b_{n-1} . The relations are defined by $a_i < b_j$ exactly if $i - \lfloor (k-1)/2 \rfloor \leq j \leq i + \lfloor k/2 \rfloor$. Again the indices are to be read cyclically. Note that $C_n^{n-1} = S_n$ and $C_n^2 = C_n$.

Theorem 4.5 ([22]). *For each $n \geq k \geq 2$, the linear extension diameter of the generalized crown C_n^k is given by*

$$\text{led}(C_n^k) = 2n(n - k + k(k - 1)).$$

For the class of Boolean lattices B_n , a formula for the linear extension diameter is conjectured in [22] and proved up to $n = 4$. We prove the conjecture for all n in Section 5.1.

Felsner and Reuter [22] also provide a heuristic method for the construction of two linear extensions with large distance. It is extremely simple: Start with one linear extension L . Now construct L' , the *complementary linear extension of L* , with the generic algorithm, using the reverse of L as priority order. Intuitively, there are many reversals between L and L' . In [22] it is shown that if this process is applied iteratively, it converges to a pair of linear extensions which form a *complementary pair*, that is, they are mutually complementary. The convergence is reached after $2h$ steps, where h is the height of the poset.

It is easy to see that each complementary pair L, L' is locally extreme in $G(\mathcal{P})$. That is, there is no neighbor of L' which has larger distance from L than L' , and no neighbor of L which has larger distance from L' than L . Unfortunately, complementary pairs do not have to be diametral, and diametral pairs do not have to be complementary. It is unknown, however, how far a complementary pair can be from being diametral.

4.2 NP-Completeness for General Posets

In this section we prove that it is NP-complete to determine the linear extension diameter of a general poset. More precisely, we consider the following decision problem:

LINEAR EXTENSION DIAMETER

Input: Finite poset \mathcal{P} , natural number k

Question: Are there two linear extensions of \mathcal{P} with distance at least k ?

For the hardness proof we use a reduction of the following problem:

BALANCED BIPARTITE INDEPENDENT SET

Input: Bipartite graph G , natural number k

Question: Is there an independent set of size $2k$, consisting of k vertices in each bipartition set?

This problem is NP-complete as the equivalent problem BALANCED COMPLETE BIPARTITE SUBGRAPH is NP-complete [26].

Theorem 4.6. *LINEAR EXTENSION DIAMETER is NP-complete.*

Proof. The problem is in NP, because the distance of two given linear extensions can be checked in time quadratic in the number of elements of \mathcal{P} , and thus two linear extensions at distance k are a certificate for a YES-instance.

To show that the problem is NP-hard, suppose that an instance of BALANCED BIPARTITE INDEPENDENT SET is given: A bipartite graph $G = (A \cup B, E)$ and $k \in \mathbb{N}$. In a preprocessing step, we take a disjoint union of two copies $G_1 = (A_1 \cup B_1, E_1)$ and $G_2 = (A_2 \cup B_2, E_2)$ of G and join all vertices of A_1 to all vertices of B_2 as well as all vertices of A_2 to all vertices of B_1 . We call the resulting graph $G' = (A' \cup B', E')$. Then G has a balanced independent set of size $2k$ exactly if G' has two disjoint balanced independent sets of size $2k$. A further convenient fact is that, after this preprocessing step, we know that $k \leq |A'|/2, |B'|/2$.

Now we build a poset \mathcal{P} starting from G' by designating the vertices of G' as *black elements* of \mathcal{P} . The sets A' and B' each form an antichain in \mathcal{P} , and an element of B' is larger than an element of A' exactly if they are adjacent in G' . Then we add the *green elements* $A_1, A_2, \dots, A_r, B_1, B_2, \dots, B_s$ and C and D with relations as shown in Figure 4.2.




Figure 4.2: For the hardness proof, we build a poset from a bipartite graph.

Finally we form \mathcal{P}^* from \mathcal{P} by replacing the green elements by long chains. Let n be the number of vertices of G' . Each A_i and each B_j is replaced by a chain of length $2n^4$, and C and D by chains of length $(2k-1)n^4$. All these chains form modules in \mathcal{P}^* , so we know by Lemma 4.3 that there is a diametral pair of linear extensions of \mathcal{P}^* in which they appear successively. Since we are only interested in the distance between the linear

extensions of a diametral pair, it suffices to consider such a diametral pair. This diametral pair corresponds to a pair of linear extensions of \mathcal{P} . The distance in \mathcal{P}^* between the two linear extensions can be thought of as a *weighted distance* between linear extensions of \mathcal{P} . From now on, we work in \mathcal{P} .

We will analyse what a weighted diametral pair of linear extensions of \mathcal{P} has to look like, and we will eventually see that its distance depends on the existence of a balanced independent set of G' . Recall that the distance between two linear extensions is the number of reversals between them. There are three types of pairs of elements of \mathcal{P} that can be reversed: First, the pairs consisting of two black elements. Every such black/black reversal adds 1 to the distance. These reversals form the *unit reversals*. Second, there are the reversals of a black element with a green element. A black/green reversal contributes $2n^4$ or $(2k-1)n^4$ unit reversals to the distance, depending on the type of the green element involved. Finally there are reversals between two green elements, that is, between two A_i or two B_j . Every such green/green reversal yields $4n^8$ unit reversals.

We saw that the contribution of the three types of reversals differs a lot. In fact, the total gain of all possible reversals of one type yields still less than one reversal of the next bigger type: There are n black elements, so the number of black/black reversals is at most $\binom{n}{2}$. This is $\theta(n^2)$, far less than the $\theta(n^4)$ that one black/green reversal yields. And how many of those can there be in total? Even if the black elements were incomparable to the whole rest of the poset, the black/green reversals could altogether contribute only $\theta(n^6)$ unit reversals, again far less than the $\theta(n^8)$ of a single green/green reversal.

The consequence of this is that we can analyse the linear extensions forming a weighted diametral pair of \mathcal{P} in three separate steps. First we check how the green elements have to be ordered in the linear extensions to yield a maximum number of green/green reversals. In the second step, we fill in the black elements in order to get a maximum number of black/green elements, knowing that these cannot influence the order of the green elements, because they cannot contribute enough. In the last step, we use the remaining freedom to order the black elements so that they add a maximum number of unit reversals. Note that the first two steps do not depend on the particular graph G' . We define the *base distance* to be the weighted distance we can achieve between two linear extensions of \mathcal{P} independently of the particular given graph G' ; put differently, it is the linear extension diameter of \mathcal{P}^* built from a complete bipartite graph G' . In the last step we will see that the existence of a balanced independent set determines how much $\text{led}(\mathcal{P}^*)$ exceeds the base distance.

We start with the first step. For this, it is enough to look at the poset \mathcal{P}' induced only by the green elements, so $\mathcal{P}' = \mathcal{P} \setminus G'$. It consists of an antichain of minima, formed by the A_i , the two elements C and D which are comparable to all other elements of \mathcal{P}' , and an antichain of maxima, formed by the B_j . See Figure 4.3 for illustration.




Figure 4.3: In the first step, we only consider the green part \mathcal{P}' of \mathcal{P} .

The poset \mathcal{P}' is two-dimensional, and the diametral pairs consist of two linear extensions L'_1, L'_2 such that the A_i and B_j come in an arbitrary order in L'_1 , and in the opposite order in L'_2 . Thus up to labelling the A_i and B_j , a diametral pair of linear extensions of \mathcal{P}' has the form

$$\begin{aligned} L'_1 &= A_1 A_2 \dots A_r C D B_1 B_2 \dots B_s \\ L'_2 &= A_r A_{r-1} \dots A_1 C D B_s B_{s-1} \dots B_1. \end{aligned}$$

Hence the green/green reversals in a diametral pair contribute a total number of $\binom{r}{2} + \binom{s}{2} \cdot 4n^8$ unit reversals to the base distance.

For the second step we need to take the elements of G' into account. We concentrate on the *lower half* $\bar{\mathcal{P}}$ of \mathcal{P} , induced by the elements of $A_1 \cup \dots \cup A_r \cup C \cup a_1 \cup \dots \cup a_r$ (see Figure 4.4).




Figure 4.4: In the second step, we look at the lower half $\bar{\mathcal{P}}$ of \mathcal{P} .

The subposet induced by the remaining elements is the *upper half*. The only incomparabilities between the upper and the lower half occur between two black elements, and we are not interested in black/black reversals in the second step. Other than that, the lower half and the upper half are symmetric; we only need to replace r by s .

Let us now look at the lower half. Each a_i is incomparable to C and to all A_j with $j \neq i$. One option to insert a_i into L'_1 and L'_2 yielding many reversals is to place it right above its predecessor A_i in both linear extensions, thus reversing with all A_j , $j \neq i$, but not with C . This is clearly best possible among all options which do not reverse a_i and C . If we want to reverse a_i with C , then the best option is to insert a_i above C (and thus also above all A_j) in one linear extension, and place it as low as possible in the other. Note for this option that a_i still has to be above A_i , and thus it cannot be reversed with the A_j which are below A_i in the other linear extension.

In the first option, a_i contributes $(r-1) \cdot 2n^4$ unit reversals. In the second option, we win $(2k-1)n^4$ unit reversals from C , but lose $2n^4 \cdot \min\{i-1, r-i\}$ unit reversals from the A_j . Which option is best therefore depends on i . We will consider three different cases. Recall for the following that $k \leq r/2$. We say that the elements between C and D are *in the middle* of a linear extension of \mathcal{P} .

If $i \leq k$, then $\min\{i-1, r-i\} = i-1$. Then we have $2n^4 \cdot (i-1) < (2k-1)n^4$. Hence it is best to choose the second option and insert a_i above C in L'_2 , thus, to put a_i in the middle of L'_2 .

If $i \geq r-k+1$, then $\min\{i-1, r-i\} = r-i$, and again $2n^4 \cdot (r-i) < (2k-1)n^4$. In this case it is best to insert a_i in the middle of L'_1 .

Finally, if $k+1 \leq i \leq r-k$, then $k \leq i-1$ and $k \leq r-i$. Thus $2k \leq 2 \cdot \min\{i-1, r-i\}$ and consequently $(2k-1)n^4 < 2n^4 \cdot \min\{i-1, r-i\}$. Therefore the first option is best in this case, that is, a_i will *not* be put in the middle of L'_1 or L'_2 .

It follows that exactly k of the a_i will be inserted in the middle of L'_1 and L'_2 , respectively. The analysis for the upper half can be made completely analogously, so k of the b_j end up in the middle of L'_1 and L'_2 , respectively. Up to reversals in the middle, we can now write down what a diametral pair L_1, L_2 of linear extensions of the whole poset \mathcal{P} looks like.

$$\begin{aligned} L_1 &= A_1 a_1 A_2 a_2 \dots A_r C a_{r-k+1} \dots a_{r-1} a_r b_1 b_2 \dots b_k D B_1 B_2 \dots b_{s-1} B_{s-1} b_s B_s \\ L_2 &= A_r a_r A_{r-1} a_{r-1} \dots A_1 C a_k \dots a_2 a_1 b_s b_{s-1} \dots b_{s-k+1} D B_s B_{s-1} \dots b_2 B_2 b_1 B_1 \end{aligned}$$

Note that the chosen order of the A_i and B_j determines which elements a_i and b_j end up in the middle of L_1 and L_2 .

The black/green reversals contributed by the lower half amount to the following number of unit reversals:

$$\begin{aligned} & r(r-1) \cdot 2n^4 \quad -2 \left(\sum_{i=1}^k (i-1) \right) \cdot 2n^4 \quad +2k \cdot (2k-1)n^4 \\ & = r(r-1) \cdot 2n^4 \quad -k(k+1) \cdot 2n^4 \quad +2k \cdot (2k-1)n^4. \end{aligned}$$

For the number of contributed black/green reversals of the upper half, we only have to replace r by s .

The third step is now easy. To analyse how many black/black reversals we can obtain, we first observe that it is clearly possible to completely reverse the order of the a_i as well as the order of the b_j . This adds another $\binom{r}{2} + \binom{s}{2}$ to the base distance.

Now we put our calculations together to find the base distance d between two diametral linear extensions of \mathcal{P} .

$$\begin{aligned} d &= \sum_{t=r,s} \binom{t}{2} \cdot 4n^8 + t(t-1) \cdot 2n^4 - k(k+1) \cdot 2n^4 + 2k \cdot (2k-1)n^4 + \binom{t}{2} \\ &= \left(\binom{r}{2} + \binom{s}{2} \right) \cdot (2n^4 + 1)^2 - 2k(k+1) \cdot 2n^4 + 4k \cdot (2k-1)n^4 \end{aligned}$$

For the only other possible black/black reversals, we have to check if some of the a_i can be brought above some of the b_j . This can only happen between the a_i and b_j in the middle of L_1 and L_2 . So there are at most $2k^2$ additional unit reversals to be won. This number can be obtained exactly if G' contains two disjoint balanced independent sets of size $2k$. Because if it does, then we can choose an order of the A_i and B_j so that the vertices of one independent set end up in the middle of L_1 , and the vertices of the other in the middle of L_2 , and then bring all the involved a_i above all the involved b_j . On the other hand, if we can win these additional $2k^2$ unit reversals, then clearly there have to be two disjoint balanced independent sets in G' .

Recall that G' has two disjoint balanced independent sets of size $2k$ exactly if G has one. So we conclude that G has a balanced independent set of size $2k$ exactly if \mathcal{P}^* has two linear extensions of distance at least $d + 2k^2$. Since it is clearly possible to build \mathcal{P}^* in time polynomial in the size of G , this proves the hardness of LINEAR EXTENSION DIAMETER. \square

Note that BALANCED BIPARTITE INDEPENDENT SET is fixed-parameter tractable, that is, there is an algorithm which solves it in time $f(k) \cdot \text{poly}(|V(G)|)$, where f is an arbitrary function depending only on k , and poly is a polynomial. This can be done by checking for all pairs of k -subsets of A and B whether they form an independent set in G .

It thus remains open whether LINEAR EXTENSION DIAMETER is still NP-hard if we fix the desired number of reversals between two linear extensions.

Open Question 1. Is the problem LINEAR EXTENSION DIAMETER fixed-parameter tractable?

4.3 Algorithm for Posets of Width 3

In this section we describe an algorithm that determines the linear extension diameter of posets of width 3, using a dynamic programming approach.

Before we go into the proof, recall the famous Dilworth's Theorem [14] which states that any poset of width k can be decomposed into k chains (see e.g. [59]).

Theorem 4.7. *The linear extension diameter of posets of width 3 can be computed in polynomial time.*

Proof. Let a poset \mathcal{P} of width 3 be given. By Dilworth's Theorem, \mathcal{P} has a decomposition into three chains. Such a decomposition can be found in linear time with an algorithm of Felsner, Raghavan and Spinrad [21]. Let \mathcal{P} have n elements, partitioned into the three chains A, B and C . A convenient consequence is that \mathcal{P} can have no more than n^3 downsets, because every downset is uniquely determined by an antichain of elements forming the maxima of the downset. Let \mathcal{P}_D denote the poset induced by a downset D of \mathcal{P} . Our approach is to calculate the linear extension diameter of \mathcal{P} dynamically by calculating it for every \mathcal{P}_D and re-using data in the process.

Let us analyse what the linear extensions of \mathcal{P} can look like. In this proof we always read linear extensions from top to bottom, so the initial segment of a linear extension consists of its topmost vertices, to find the i -th element we count from the top, and so on. The first element of a linear extension L can either be the top element of A , the top element of B , or the top element of C . Then there is an initial segment with only elements of that chain, until at a certain position the top element of a second chain appears.

For chains $V, W, X, Y \in \{A, B, C\}$ and $i, j \in \mathbb{N}$ with $i, j \geq 2$, we now define $\text{led}_D(VW, i, XY, j)$ as the maximum distance between two linear extensions L_1, L_2 of \mathcal{P}_D such that the following holds: the maximal element of L_1 belongs to chain V , the second chain appearing is chain W , and the first element of chain W appears at position i in L_1 ; the maximal element of L_2 belongs to chain X , the second chain appearing is chain Y , and the first element of chain Y appears at position j in L_2 . We call such a pair of linear extensions *relevant* for $\text{led}_D(VW, i, XY, j)$, see Figure 4.5 for illustration. If there is no relevant pair of linear extensions, we set $\text{led}_D(VW, i, XY, j) = -\infty$.

Our plan now is to calculate $\text{led}_D(VW, i, XY, j)$ for all choices of parameters D, V, W, X, Y, i, j . The number of values to compute is then bounded by $n^3 \cdot 6 \cdot 6 \cdot n \cdot n = O(n^5)$, and thus polynomial.




Figure 4.5: Scheme of a relevant pair of linear extensions for $\text{led}_D(VW, i, XY, j)$.

If \mathcal{P}_D consists of only one chain or a chain plus one element, then we can read off the desired value immediately. For all other cases we describe a recursive formula. So let us assume that \mathcal{P}_D contains elements from all three chains, or that there are two chains from which \mathcal{P}_D contains more than one element, and that we have computed all values for all downsets of \mathcal{P}_D . We want to compute $\text{led}_D(VW, i, XY, j)$ with $V, W, X, Y \in \{A, B, C\}$. We first choose a chain which appears in VW as well as in XY . Since there are only three chains available, it must exist. Suppose it is chain A . We will use the values for $D' = D - a_1$ for our recursion, where a_i is defined as the i -th element of chain A (counted from the top) in \mathcal{P}_D . There are three cases to consider, and each case will be divided into several possibilities.

First we consider the case where A is the chain appearing second in both L_1 and L_2 and the chains appearing first are different. We show a formula for $\text{led}_D(BA, i, CA, j)$; the formula for $\text{led}_D(CA, i, BA, j)$ is analogous. If $a_1 > b_{i-1}$ or $a_1 > c_{j-1}$, then $\text{led}_D(BA, i, CA, j) = -\infty$. Otherwise we know that the element a_1 contributes exactly $i - 1 + j - 1$ reversals to the distance between L_1 and L_2 . Hence to calculate $\text{led}_D(BA, i, CA, j)$, we can maximize over all distances between two linear extensions L'_1, L'_2 of D' in which we can reinsert a_1 to get two relevant linear extensions.

What do the relevant linear extensions L'_1 and L'_2 look like? Of course it holds that L'_1 starts with $i - 1$ elements of chain B , and L'_2 with $j - 1$ elements of chain C . So the first chains appearing in L'_1 and L'_2 are fixed. But after the initial segment, there are four possibilities. If the second appearing chain is again A in both L'_1 and L'_2 , then we are interested in the maximum distance over all possible positions of a_2 , so let us set $\alpha = \max_{r \geq i, p \geq j} \text{led}_{D'}(BA, r, CA, p)$. The distance may increase if the second chain in L'_1 is C , so we want to know $\beta = \max_{r \geq i, p \geq j} \text{led}_{D'}(BC, r, CA, p)$; or if the second chain in L'_2 is B , in which

case we are interested in $\gamma = \max_{r \geq i, p \geq j} \text{led}_{D'}(BA, r, CB, p)$. The last possibility we have to take into account is $\delta = \max_{r \geq i, p \geq j} \text{led}_{D'}(BC, r, CB, p)$. To compute $\text{led}_D(BA, i, CA, j)$, we have to maximize over all these possibilities. We obtain

$$\text{led}_D(BA, i, CA, j) = i + j - 2 + \max\{\alpha, \beta, \gamma, \delta\}.$$

In the second case, A appears first in one linear extension, say L_1 , and second in the other, L_2 . We will show a formula for $\text{led}_D(AB, i, CA, j)$. The formulae for $\text{led}_D(AC, i, BA, j)$, $\text{led}_D(AB, i, BA, j)$ and $\text{led}_D(AC, i, CA, j)$ are built analogously. Again, if $b_1 > a_{i-1}$ or $a_1 > c_{j-1}$, then $\text{led}_D(AB, i, CA, j) = -\infty$. Otherwise, we know that a_1 contributes $j - 1$ reversals to the distance between L_1 and L_2 . Let us distinguish the possibility $i > 2$ from the possibility $i = 2$. If $i > 2$, then we know that a relevant L'_1 starts with $i - 2$ elements of A and then lists b_1 . So we only need to maximize over the possible second chains in L'_2 . We obtain

$$\begin{aligned} \text{led}_D(AB, i, CA, j) &= j - 1 + \max\{\alpha, \beta\}, \quad \text{where} \\ \alpha &= \max_{p \geq j} \text{led}_{D'}(AB, i - 1, CA, p) \quad \text{and} \\ \beta &= \max_{p \geq j} \text{led}_{D'}(AB, i - 1, CB, p). \end{aligned}$$

If $i = 2$, then a relevant L'_1 starts with an element from B , and we maximize over all possible combinations of second chains in L'_1 and in L'_2 :

$$\begin{aligned} \text{led}_D(AB, 2, CA, j) &= j - 1 + \max\{\alpha, \beta, \gamma, \delta\}, \quad \text{where} \\ \alpha &= \max_{r \geq 2, p \geq j} \text{led}_{D'}(BA, r, CA, p), \quad \beta = \max_{r \geq 2, p \geq j} \text{led}_{D'}(BA, r, CB, p), \\ \gamma &= \max_{r \geq 2, p \geq j} \text{led}_{D'}(BC, r, CA, p), \quad \delta = \max_{r \geq 2, p \geq j} \text{led}_{D'}(BC, r, CB, p). \end{aligned}$$

In the third case, A is the first chain in both L_1 and L_2 . We will show how to compute $\text{led}_D(AB, i, AC, j)$; the other cases for the second chains work analogously. The principle is the same as in the first two cases, we only need to distinguish more possibilities now. First we observe that a_1 does not contribute any reversals here, but on the other hand we will always get a finite value for $\text{led}_D(AB, i, AC, j)$. If both $i > 2$ and $j > 2$, we know exactly which chains come first and second in the relevant linear extensions. Hence, nothing changes but the position of the first element of the second chain:

$$\text{led}_D(AB, i, AC, j) = \text{led}_{D'}(AB, i - 1, AC, j - 1).$$

If one of the two parameters i and j is exactly 2, then we need to distinguish two possibilities as before.

If $i = 2$ and $j > 2$, we have

$$\begin{aligned} \text{led}_D(AB, 2, AC, j) &= \max\{\alpha, \beta\}, \quad \text{where} \\ \alpha &= \max_{r \geq 2} \text{led}_{D'}(BA, r, AC, j - 1) \quad \text{and} \\ \beta &= \max_{r \geq 2} \text{led}_{D'}(BC, r, AC, j - 1). \end{aligned}$$

For $i > 2$ and $j = 2$, we obtain

$$\begin{aligned} \text{led}_D(AB, i, AC, 2) &= \max\{\alpha, \beta\}, \quad \text{where} \\ \alpha &= \max_{p \geq 2} \text{led}_{D'}(AB, i - 1, CA, p), \quad \text{and} \\ \beta &= \max_{p \geq 2} \text{led}_{D'}(AB, i - 1, CB, p). \end{aligned}$$

Finally, if $i = 2$ and $j = 2$, we again have to consider four possibilities:

$$\begin{aligned} \text{led}_D(AB, 2, AC, 2) &= \max\{\alpha, \beta, \gamma, \delta\}, \quad \text{where} \\ \alpha &= \max_{r \geq 2, p \geq 2} \text{led}_{D'}(BA, r, CA, p), \quad \beta = \max_{r \geq 2, p \geq 2} \text{led}_{D'}(BA, r, CB, p), \\ \gamma &= \max_{r \geq 2, p \geq 2} \text{led}_{D'}(BC, r, CA, p), \quad \delta = \max_{r \geq 2, p \geq 2} \text{led}_{D'}(BC, r, CB, p). \end{aligned}$$

Hence we can compute the desired values for all downsets in time polynomial in n . The linear extension diameter of the whole poset \mathcal{P} is now of course the maximum of the values for $\mathcal{P}_D = \mathcal{P}$ over all choices of parameters for $\text{led}_D(VW, i, XY, j)$. \square

However, since the number of linear extensions increases exponentially with the width, a large width makes the computation of the linear extension diameter more difficult. A natural question to ask is whether there is a polynomial time algorithm computing the linear extension diameter poset with arbitrary fixed width. We have found no way of generalizing the above proof to answer this. Our proof uses the special property of a set of size 3 that, whenever you chose a pair from this set, and then again a pair from this set, there is an element that you chose twice. This does not generalize to a set containing more elements.

Open Question 2. Is the problem of computing the linear extension diameter solvable in polynomial time for posets of fixed width?

Chapter 5

Linear Extension Diameter of Distributive Lattices

In the previous chapter we showed that it is NP-complete to determine the linear extension diameter of a poset, in general. For a poset \mathcal{P} of width 3, we have seen a dynamic programming approach to compute $\text{led}(\mathcal{P})$. However, we could not give a closed formula.

In this chapter, we consider the linear extension diameter of different classes of distributive lattices. In Section 5.1 we deal with Boolean lattices. We prove a formula for their linear extension diameter and characterize their diametral pairs of linear extensions. In Section 5.2 we consider the more general class of downset lattices of 2-dimensional posets. Again, we characterize their diametral pairs of linear extensions. The proofs are based on the same ideas as those for Boolean lattices, but they also make heavy use of the structure of the underlying 2-dimensional posets. We also show how to compute the linear extension diameter of the downset lattice of a 2-dimensional poset \mathcal{P} in time polynomial in $|\mathcal{P}|$.

Note that the results of Section 5.1 are contained in the results of Section 5.2. We nevertheless carry out the proofs for the Boolean lattices in detail since they provide an accessible introduction to the techniques used in the general case. However, since Section 5.2 is largely self-contained, the self-confident reader may proceed there directly.

5.1 Boolean Lattices

Recall that the n -dimensional Boolean lattice B_n is the poset on all subsets of $[n]$, ordered by inclusion. In this section, we prove a formula for the linear extension diameter of Boolean lattices, confirming Conjecture 5.1 below. Then we characterize the diametral pairs of linear extensions of B_n .

Conjecture 5.1 (Felsner, Reuter [22]). $led(B_n) = 2^{2n-2} - (n+1) \cdot 2^{n-2}$.

The formula in Conjecture 5.1 is the distance between two linear extensions of B_n , arising from a generalization of the reverse lexicographic order. Since these revlex linear extensions play a central role throughout the whole chapter, we devote a subsection to their introduction.

5.1.1 Revlex Linear Extensions

Before defining revlex linear extensions, let us clarify some notation: Let σ be a permutation of $[n]$. We say that i is σ -smaller than j , and write $i <_\sigma j$, if $\sigma^{-1}(i) < \sigma^{-1}(j)$. The σ -minimum $\min_\sigma S$ of a finite set S is the element which is σ -smallest in S . For example, if $\sigma = 2413$, then $4 <_\sigma 1$, and $\min_\sigma \{1, 2, 3\} = 2$. We define σ -larger and σ -maximum analogously.

Definition 5.2. The σ -revlex order $<_\sigma$ on the pairs $S, T \subseteq [n]$ is defined as follows:

$$S <_\sigma T \iff \max_\sigma(S \Delta T) \in T.$$

Lemma 5.3. The relation of being in σ -revlex order defines a linear extension L_σ of the Boolean lattice.

Proof. We need to show that the σ -revlex relation defines a linear order on the subsets of $[n]$ which respects the inclusion order. The last part is easy to see: For $S, T \subseteq [n]$ with $S \subseteq T$ we have $S \Delta T \subseteq T$, thus $S \leq_\sigma T$.

It is also clear that the relation is antisymmetric and total. So we only need to prove transitivity. Assume for contradiction that there are three sets $A, B, C \subseteq [n]$ with $A <_\sigma B$ and $B <_\sigma C$ and $C <_\sigma A$. Then $\max_\sigma(A \Delta B) = b \in B$, $\max_\sigma(B \Delta C) = c \in C$ and $\max_\sigma(C \Delta A) = a \in A$.

Note that from $b \in B$ and $c \notin B$ we have that $b \neq c$. Because the situation of the three elements a, b, c is symmetric, it follows that they are pairwise different. Assume that $c = \max_\sigma\{a, b, c\}$. Then by definition of b , we have $c \notin A \Delta B$, and by definition of a , we have $c \notin A \Delta C$. But now it follows that $c \notin B \Delta C$, which is a contradiction. \square

Definition 5.4. Let σ be a permutation on $[n]$. The linear extension of B_n given by the σ -revlex order is a revlex linear extension L_σ of B_n .

By $\bar{\sigma}$ we denote the reverse of σ . We call the pair $L_\sigma, L_{\bar{\sigma}}$ a revlex pair of linear extensions.

In [22], the linear extensions L_{id} and $L_{\bar{\text{id}}}$ are called the reverse lexicographic and the reverse antilexicographic order. As linear extensions of B_4 , they have the following form:

$$L_{\text{id}} = \emptyset \ 1 \ 2 \ 12 \ 3 \ 13 \ 23 \ 123 \ 4 \ 14 \ 24 \ 124 \ 34 \ 134 \ 234 \ 1234$$

$$L_{\bar{\text{id}}} = \emptyset \ 4 \ 3 \ 34 \ 2 \ 24 \ 23 \ 234 \ 1 \ 14 \ 13 \ 134 \ 12 \ 124 \ 123 \ 1234$$

Theorem 5.11 proves that the pairs $L_\sigma, L_{\bar{\sigma}}$ are exactly the diametral pairs of linear extensions of B_n . Recall that a diametral pair of linear extensions can be used to obtain an optimal drawing of a poset. Figure 5.1 shows the drawing of B_4 resulting from $L_{\text{id}}, L_{\bar{\text{id}}}$.




Figure 5.1: The drawing of B_4 based on the diametral pair $L_{\text{id}}, L_{\bar{\text{id}}}$. The subsets of $[4]$ are denoted by their characteristic vectors.

Revlex orders have an appealing intrinsically symmetric structure: It follows immediately from the definition that any subcube is traversed in the same manner. More precisely, consider a revlex linear extension L_σ of B_n . A k -dimensional *subcube* B_k of B_n is induced by a set S of subsets of $[n]$ such that $n - k$ atoms are fixed to appear either in all or in no subset in S , and only k atoms are *free*. Then the linear extension of B_k induced

by L_σ is again a revlex linear extension. It corresponds to the permutation induced by σ on the free atoms of the subcube.

Another property of the revlex linear extensions is that they are built in a depth first search fashion. In [22], this property is called *super-greedy*: In the generic algorithm for building linear extensions, x_i is chosen from $\text{Min}(\mathcal{P} - \{x_1, \dots, x_{i-1}\}) \cap \text{Succ}(x_j)$, where $j < i$ is maximal such that this set is non-empty.

We want to highlight one more property of revlex orders. It is related to their super-greediness and will be central in the next chapter. Recall that a linear extension is called reversing if it contains a critical pair of elements in the non-canonical order, see Definition 1.2.

Lemma 5.5. *Every revlex linear extension of B_n is reversing.*

Proof. Let σ be a permutation of $[n]$, consider the revlex linear extension L_σ of B_n . By Lemma 1.4, the critical pairs of B_n are exactly the n atom-coatom pairs (a, a^c) , where $a \in [n]$ and $a^c = [n] \setminus \{a\}$. Now let $a = \max_\sigma[n]$. Then $\max_\sigma(a \Delta a^c) = a$. Thus $a^c < a$ in L_σ , which means that L_σ reverses this critical pair. \square

5.1.2 Determining the Linear Extension Diameter

Results about Boolean lattices are usually proved by induction on the dimension. This is often done by partitioning B_n into two subcubes, e.g. by partitioning the subsets of $[n]$ into those which contain the atom n and those which do not. To prove Conjecture 5.1, we use a novel partition, defined as follows:

Definition 5.6. *For $D \subseteq [n]$ and $I \subseteq [n] \setminus D$ we define $\mathcal{C}_{D,I}$ as the set of all ordered pairs (S, T) of subsets of $[n]$ with $S \Delta T = D$ and $S \cap T = I$.*

Clearly, the sets $\mathcal{C}_{D,I}$ partition the ordered pairs of subsets of $[n]$ into equivalence classes. Note that these equivalence classes are in bijection with the intervals of B_n : The class $\mathcal{C}_{D,I}$ corresponds to the interval $[I, D \cup I]$, and it contains all pairs (S, T) with $S \cap T = I$ and $S \cup T = D \cup I$.

Another important observation is that we can associate a subset $X \subseteq D$ with the pair $(X \cup I, X^c \cup I) \in \mathcal{C}_{D,I}$, where $X^c = D \setminus X$. On the other hand, for each pair $(S, T) \in \mathcal{C}_{D,I}$ we have $S \setminus I =: X \subseteq D$ and $T = X^c \cup I$. This yields the following useful lemma:

Lemma 5.7. *The pairs of a class $\mathcal{C}_{D,I}$ are in bijection with the subsets of D . Each class contains 2^d ordered pairs, where $|D| = d$.* \square

Recall that the distance between two linear extensions is the number of reversals, that is, of *unordered* pairs of elements appearing in different orders in the two linear extensions. In contrast to this, the equivalence classes $\mathcal{C}_{D,I}$ contain *ordered pairs*. Therefore we will often have to switch between ordered and unordered pairs in our proofs.

The following proposition, settling the lower bound of Conjecture 5.1, was proved inductively in [22]. Here we give a more combinatorial proof.

Proposition 5.8 ([22]). *The distance between a revlex pair of linear extensions of B_n is*

$$2^{2^{n-2}} - (n+1) \cdot 2^{n-2}.$$

Proof. Let σ be a permutation of $[n]$ and let $L_\sigma, L_{\bar{\sigma}}$ be the corresponding revlex pair. We claim that an equivalence class $\mathcal{C}_{D,I}$ contributes exactly 2^{d-2} reversals between L_σ and $L_{\bar{\sigma}}$ if $d \geq 2$, and none if $d < 2$.

Observe that since $\bar{\sigma}$ is the reverse permutation of σ , the σ -minimum of a set equals its $\bar{\sigma}$ -maximum, and vice versa. Thus we have

$$S <_{\bar{\sigma}} T \iff \min_{\sigma}(S \Delta T) \in T.$$

Let $\mathcal{C}_{D,I}$ be an equivalence class as in Definition 5.6. If D is empty, then $\mathcal{C}_{D,I}$ contains only the pair (I, I) , and thus cannot contribute any reversal. If D consists of only one element, say, x , then $\mathcal{C}_{D,I}$ consists of the two pairs $(I, x \cup I)$ and $(x \cup I, I)$. Since $I \subset x \cup I$, the class $\mathcal{C}_{D,I}$ again cannot contribute any reversal.

So we may assume that D contains at least two elements, and hence $\max_{\sigma} D \neq \min_{\sigma} D$. Then a pair in $\mathcal{C}_{D,I}$ corresponding to some $X \subseteq D$ is a reversal between L_σ and $L_{\bar{\sigma}}$ if and only if exactly one of the elements $\min_{\sigma} D$ and $\max_{\sigma} D$ is contained in X . We want to count (S, T) and (T, S) only once, so let us count the sets $X \subseteq D$ with $\min_{\sigma} D \in X$ and $\max_{\sigma} D \notin X$. There are 2^{d-2} such sets, and thus $\mathcal{C}_{D,I}$ contributes 2^{d-2} reversals between L_σ and $L_{\bar{\sigma}}$, as claimed.

How many reversals does this yield in total? A pair (D, I) induces a class $\mathcal{C}_{D,I}$ exactly if $D \subseteq [n]$ and $I \subseteq [n] \setminus D$, and D is *large enough*, that is, $d \geq 2$. Let us count how many reversals a large enough D contributes. It has 2^{n-d} partners I with which it forms a class. Thus, D contributes $2^{n-d} \cdot 2^{d-2} = 2^{n-2}$ reversals. There are $2^n - (n+1)$ large enough sets $D \subseteq [n]$. Therefore the distance between L_σ and $L_{\bar{\sigma}}$ is

$$2^{n-2} \cdot (2^n - (n+1)) = 2^{2^{n-2}} - (n+1) \cdot 2^{n-2}. \quad \square$$

For proving the upper bound on $\text{led}(B_n)$, we need Kleitman's Lemma:

Kleitman's Lemma ([35]). *Let \mathcal{F}_1 and \mathcal{F}_2 be families of subsets of $[d]$ which are closed downwards, that is, for every set in \mathcal{F}_i all of its subsets are also in \mathcal{F}_i . Then the following formula holds:*

$$|\mathcal{F}_1| \cdot |\mathcal{F}_2| \leq 2^d |\mathcal{F}_1 \cap \mathcal{F}_2|.$$

Kleitman's lemma can be proved by a rather straightforward induction, see e.g. [1].

Theorem 5.9. $\text{led}(B_n) = 2^{2n-2} - (n+1) \cdot 2^{n-2}$.

Proof. Proposition 5.8 yields $\text{led}(B_n) \geq 2^{2n-2} - (n+1) \cdot 2^{n-2}$ since the distance between any pair of linear extensions is a lower bound for the linear extension diameter. To prove that this formula is also an upper bound, we will again use the equivalence classes from Definition 5.6. Given two linear extensions L_1, L_2 of B_n , we will show that each $\mathcal{C}_{D,I}$ can contribute at most 2^{d-2} reversals between L_1 and L_2 .

Let us fix a class $\mathcal{C}_{D,I}$. It will turn out that the set I actually plays no role for our argument, therefore we assume that $I = \emptyset$. The only thing this assumption changes is that for a set $X \subseteq D$, now (X, X^c) itself is a pair of $\mathcal{C}_{D,I}$. The reader is invited to check that the following argument goes through unchanged if each X is replaced by $X \cup I$ and each X^c by $X^c \cup I$.

We say that $X \subseteq D$ is *down* in a linear extension L if $X < X^c$ in L . Let \mathcal{F}_1 be the family of subsets of D which are down in L_1 , and \mathcal{F}_2 the family of subsets of D which are down in L_2 . A pair $(X, X^c) \in \mathcal{C}_{D,I}$ yields a reversal between L_1 and L_2 exactly if X is down in one L_i , but not in the other. Thus our aim is to find an upper bound on $|\mathcal{F}_1 \Delta \mathcal{F}_2|$.

The following key observation captures the essence of transitive forcing between the different pairs: If $X < X^c$ in L_i , and $Y \subseteq D$ is a subset of X , then $X^c \subseteq Y^c$, and hence, by transitivity, $Y < X < X^c < Y^c$ in L_i . Thus from $X \in \mathcal{F}_i$ it follows that $Y \in \mathcal{F}_i$ for every subset Y of X . This means that \mathcal{F}_1 and \mathcal{F}_2 each form a family of subsets of $[d]$ which is closed downwards. (Put differently, \mathcal{F}_1 and \mathcal{F}_2 form downsets in B_d .) Hence we can apply Kleitman's Lemma, which yields $|\mathcal{F}_1| \cdot |\mathcal{F}_2| \leq 2^d |\mathcal{F}_1 \cap \mathcal{F}_2|$.

We observe that for every L and every set $X \subseteq D$, exactly one of X and X^c is down in L . Hence we have $|\mathcal{F}_1| = |\mathcal{F}_2| = 2^{d-1}$. It follows that $|\mathcal{F}_1 \cap \mathcal{F}_2| \geq 2^{d-2}$.

Also, if X is down in both L_1 and L_2 , then X^c is down in neither. That is, $X \in \mathcal{F}_1 \cap \mathcal{F}_2 \iff X^c \in (\mathcal{F}_2 \cup \mathcal{F}_1)^c$, and thus $|\mathcal{F}_1 \cap \mathcal{F}_2| = |(\mathcal{F}_2 \cup \mathcal{F}_1)^c|$. From this we obtain

$$|\mathcal{F}_1 \Delta \mathcal{F}_2| = 2^n - |\mathcal{F}_1 \cap \mathcal{F}_2| - |(\mathcal{F}_2 \cup \mathcal{F}_1)^c| \leq 2^d - 2^{d-2} - 2^{d-2} = 2^{d-1}.$$

In $\mathcal{F}_1 \triangle \mathcal{F}_2$, every reversal is counted twice – once with the set that is down in L_1 and once with the set that is down in L_2 . Therefore the number of (unordered) reversals that $\mathcal{C}_{D,I}$ can contribute is at most 2^{d-2} . Now we can use the argument from the end of the proof of Proposition 5.8 to show that the total number of reversals is at most $2^{2n-2} - (n+1) \cdot 2^{n-2}$. \square

In the above two proofs we have shown the following fact, which we state explicitly here since it will be useful for the characterization of the diametral pairs of B_n .

Fact 5.10. If L, \bar{L} is a diametral pair of linear extensions of B_n , then each equivalence class $\mathcal{C}_{D,I}$ with $d \geq 2$ contributes exactly 2^{d-2} reversals between L and \bar{L} .

5.1.3 Characterizing Diametral Pairs

We have shown that every revlex pair of linear extensions is a diametral pair of the Boolean lattice. Thus we know $n!/2$ diametral pairs. The following theorem proves that these are in fact the only ones.

Theorem 5.11. *The diametral pairs of the Boolean lattice are unique up to isomorphism. More precisely, if L, \bar{L} is a diametral pair of linear extensions of B_n and σ is the order of the atoms in L , then $L = L_\sigma$ and $\bar{L} = L_{\bar{\sigma}}$.*

Proof. We show by induction on k that each set of cardinality k is in σ -revlex order in L and in $\bar{\sigma}$ -revlex order in \bar{L} with all sets of cardinality less or equal to k . To do so, we use Fact 5.10: Every equivalence class $\mathcal{C}_{D,I}$ with $d \geq 2$ contributes exactly 2^{d-2} reversals between the diametral linear extensions L and \bar{L} .

Recall that σ denotes the order of the atoms in L . For every pair $i <_\sigma j$ of atoms, consider the class $\mathcal{C}_{D,I}$ defined by $D = \{i, j\}$ and $I = \emptyset$. This class needs to contribute $2^{2-2} = 1$ reversal. Thus, i and j must appear in reversed order in \bar{L} . Hence the permutation defining the order of the atoms in \bar{L} is $\bar{\sigma}$.

Let L^k be the restriction of L to the sets of cardinality at most k . Our induction hypothesis is that all pairs of sets in L^{k-1} are in σ -revlex order, that is, $L^{k-1} = L_\sigma^{k-1}$, and that all pairs of sets in \bar{L}^{k-1} are in $\bar{\sigma}$ -revlex order, that is, $\bar{L}^{k-1} = L_{\bar{\sigma}}^{k-1}$. For the induction step, we will first show that each set of size k is in the desired order in L^k and \bar{L}^k with all sets of strictly smaller cardinality, then that it is in the desired order with all sets of equal cardinality.

We use the following expression: Let $X, Y, Z \subseteq [n]$ such that X and Z have cardinality $k-1$ and Y has cardinality k . If $X < Z$ is a cover relation

in L^{k-1} and $X < Y < Z$ in L^k , then we say that Y sits in the *slot* between X and Z in L^k .

Let A be a set of size k in B_n and let A' be the subset of A which is largest in L^{k-1} . Let B be the immediate successor of A' in L^{k-1} .

Claim. A has to sit in the slot between A' and B in L^k .

Note that $A' = A \setminus a$ for an atom $a \in [n]$. By induction, all subsets of A are in σ -revlex order in L . So we know that if A'' is a second subset of cardinality $k - 1$, then the element of A that is missing in A' is σ -smaller than the element of A that is missing in A'' . Therefore $a = \min_\sigma A$.

Now observe that since $A' < B$ in L^{k-1} and $|A'| = k - 1$ we have $A' \parallel B$. Again by induction we know that $\max_\sigma(A' \triangle B) = b \in B$. If there were $b' \in B \setminus A'$ with $b' \neq b$, then $A' < B \setminus b' < B$ in L^{k-1} , which is a contradiction to the choice of B . Therefore $B \setminus A' = \{b\}$.

Because $A' \parallel B$, there is an element $a' \in A' \setminus B$. We have $b >_\sigma a'$ and therefore also $b >_\sigma a$. Hence, $b = \max_\sigma(A \triangle B)$ and $a = \min_\sigma(A \triangle B)$.

Consider the class $\mathcal{C}_{D,I}$ defined by $D = A \triangle B$ and $I = A \cap B$. Note that $|D \cup I| = |A \cup B| = |A \cup \{b\}| = k + 1$. Choose a set $X \subseteq D \setminus \{a, b\}$. Then $|X \cup I| \leq k - 1$, thus we can apply the induction hypothesis to get $X \cup I < b \cup I$ in L , and with $b \in X^c$ it follows $X \cup I < X^c \cup I$ in L . Analogously, we have $X \cup I < a \cup I < X^c \cup I$ in \bar{L} . Therefore the pair (X, X^c) does not yield a reversal between L and \bar{L} , and neither does the pair (X^c, X) .

There are 2^{d-2} choices for X , hence we have found 2^{d-1} ordered pairs in $\mathcal{C}_{D,I}$ which do not yield reversals between L and \bar{L} . By Fact 5.10, the class $\mathcal{C}_{D,I}$ contributes exactly 2^{d-2} reversals. But the remaining 2^{d-1} ordered pairs in $\mathcal{C}_{D,I}$ can yield at most 2^{d-2} reversals, thus, they all have to be reversed between L and \bar{L} . It follows that all subsets Y of D containing exactly one of the two atoms a and b have to be down in exactly one of the two linear extensions.

In particular, we can choose $Y = A \setminus I$ to see that $\{A, B\}$ must be a reversal. In \bar{L} , we know the order of A and B : Set $A'' = I \cup a \subseteq A$, then $\min_\sigma(B \triangle A'') = a \in A''$. But this means $\max_{\bar{\sigma}}(B \triangle A'') = a \in A''$. So we have $B < A''$ in \bar{L} by induction and thus $B < A$ in \bar{L} by transitivity. Hence it follows that $A < B$ in L . This proves our claim that A has to sit in the slot between A' and B in L . \triangle

Recall that $\max_\sigma(A \triangle B) = b \in B$. Hence, by showing $A < B$ in L^k , we have shown that A is in σ -revlex order with B in L^k . Since the slot after A' is the lowest possible position for A in L^k , it follows by transitivity that A is in σ -revlex order with all sets of cardinality less than k in L^k . By

reversing the roles of L and \bar{L} we see that A also has to be in $\bar{\sigma}$ -revlex order in \bar{L} with all sets of smaller cardinality. Now we will show that all pairs of sets with equal cardinality k need to be in σ -revlex order in L .

Let us consider two sets $A_i, A_j \in B_n$ with cardinality k . If they are inserted into different slots in L^k , then their order in L^k equals their order in L_σ , thus they are in σ -revlex order. If they go into the same slot, then they have the same largest $(k-1)$ -subset A' in L . Thus $|A_i \triangle A_j| = 2$. We also know that $a_i := A_i \setminus A' = \min_\sigma A_i$ and $a_j := A_j \setminus A' = \min_\sigma A_j$. Assume without restriction that $a_i <_\sigma a_j$. Then for the pair of sets $\{A_j, \{a_i\}\}$, we know that $a_i = \min_\sigma(A_j \triangle \{a_i\}) = \max_{\bar{\sigma}}(A_j \triangle \{a_i\})$. By induction it follows that $A_j < a_i < A_i$ in \bar{L} , that is, A_i and A_j belong to different slots in \bar{L} . Now, since the class $\mathcal{C}_{D,I}$ containing (A_i, A_j) needs to contribute $2^{2-2} = 1$ reversal between L and \bar{L} , and (A_i, A_j) is the only incomparable pair in this class (except for the reversed pair), we know that we must have $A_i < A_j$ in L^k . Since $\max_\sigma(A_i \triangle A_j) = a_j$, this means that A_i and A_j are in σ -revlex order in L^k .

We can apply the same argument with the roles of L and \bar{L} reversed to show that all sets of cardinality k are in $\bar{\sigma}$ -revlex order in \bar{L} , thus $\bar{L}^k = L_\sigma^k$. By induction we conclude that $L = L_\sigma$ and $\bar{L} = L_{\bar{\sigma}}$. \square

5.2 Downset Lattices of 2-Dimensional Posets

The Boolean lattice B_n can be viewed as the distributive lattice of downsets of the n -element antichain, see Section 1.1. Now let \mathcal{P} be an arbitrary poset and consider its downset lattice $\mathcal{D}_\mathcal{P}$, that is, the poset on all downsets of \mathcal{P} , ordered by inclusion.

In Section 5.2.1, we give an upper bound for the linear extension diameter of $\mathcal{D}_\mathcal{P}$. We then define revlex linear extensions of $\mathcal{D}_\mathcal{P}$. We show that if \mathcal{P} is 2-dimensional, our upper bound on $\text{led}(\mathcal{D}_\mathcal{P})$ is attained by revlex pairs. In Subsection 5.2.2, we prove that for 2-dimensional \mathcal{P} , all diametral pairs of $\mathcal{D}_\mathcal{P}$ are revlex pairs. For the proofs, we make use of the main ideas from the last section. In Subsection 5.2.3 we show how to compute the linear extension diameter of $\mathcal{D}_\mathcal{P}$ in time polynomial in $|\mathcal{P}|$.

As described before, we can use our results to obtain optimal drawings of $\mathcal{D}_\mathcal{P}$. See Figure 5.2 for an example.

In this section we make fundamental use of the canonical bijection between the downsets and the antichains of a poset which was mentioned in Section 1.1. We frequently switch back and forth between the two viewpoints. We write A^\downarrow to refer to the downset generated by an antichain A . We write $\text{Max}(\mathcal{A})$ to refer to the antichain of maxima of a downset \mathcal{A} .




Figure 5.2: A drawing of a downset lattice based on a diametral pair.

5.2.1 Revlex Pairs are Diametral Pairs

To prove an upper bound for the linear extension diameter of downset lattices, we use a generalization of the equivalence classes from Definition 5.6.

Definition 5.12. Let \mathcal{P} be a poset, $D \subset \mathcal{P}$ and $I \subset \mathcal{P} \setminus D$. We define $\mathcal{C}_{D,I}$ as the set of all ordered pairs (A, B) of antichains of \mathcal{P} with $D = A \Delta B$ and $I = A \cap B$.

It is easy to see that the sets $\mathcal{C}_{D,I}$ partition the ordered pairs of antichains of \mathcal{P} into equivalence classes. Note that for a class $\mathcal{C}_{D,I}$, the sets D and I are disjoint. Furthermore, there is no relation in \mathcal{P} between any element of I and any element of D .

Lemma 5.13. Let $\mathcal{C}_{D,I}$ be an equivalence class as defined above. Let $\mathcal{P}[D]$ be the subposet of \mathcal{P} induced by the elements of D , and let \mathcal{K} be the set of connected components of $\mathcal{P}[D]$. Then the pairs in $\mathcal{C}_{D,I}$ are in bijection with the subsets of \mathcal{K} .

Proof. For a given equivalence class $\mathcal{C}_{D,I}$, let (A, B) be a pair in $\mathcal{C}_{D,I}$, thus we have $D = A \Delta B$. First observe that since A and B are antichains, $\mathcal{P}[D]$ is a poset of height at most 2 (see Figure 5.3). Thus all elements of D belong to the antichain $\text{Max}(D)$ of maxima of $\mathcal{P}[D]$ or to the antichain $\text{Min}(D)$ of minima of $\mathcal{P}[D]$.

Consequently, every connected component κ of $\mathcal{P}[D]$ is either a single element or has height 2. If κ is a single element, it belongs either to A or to B . If κ has height 2, there are again two possibilities: Either the maxima of κ all belong to A and its minima all belong to B , or the maxima of κ all belong to B and its minima all belong to A . In the first case, the minima of κ also belong to A^\perp , but not to A . In the second case, the minima of κ also belong to B^\perp , but not to B .




Figure 5.3: An example for $\mathcal{P}[D]$ with three components. The assignment of minima and maxima to A and B specifies one of the eight pairs in the class.

We obtain a different pair $(A', B') \in \mathcal{C}_{D,I}$ by switching the roles of A and B in one component of $\mathcal{P}[D]$. We can do this switch independently for each component, but we have to do it for the whole component to ensure that the elements belonging to A and B , respectively, still form an antichain.

Let $K \subseteq \mathcal{K}$ be a subset of the components of $\mathcal{P}[D]$. With K we associate a subset X_K of D by setting

$$X_K = \bigcup_{\kappa \in K} \text{Max}(\kappa) \cup \bigcup_{\kappa \notin K, |\kappa| > 1} \text{Min}(\kappa).$$

Let $X_K^c = D \setminus X_K$. Now define a map from the powerset of \mathcal{K} to $\mathcal{C}_{D,I}$ via

$$K \mapsto (A_K, B_K) = (X_K \cup I, X_K^c \cup I).$$

We claim that this defines a bijection. Recall that D and I are disjoint and that there are no relations between them. It follows that

$$A_K \Delta B_K = (X_K \cup I) \Delta (X_K^c \cup I) = X_K \cup X_K^c = D,$$

and

$$A_K \cap B_K = (X_K \cup I) \cap (X_K^c \cup I) = I.$$

Thus K is indeed mapped to a pair in $\mathcal{C}_{D,I}$.

If K and K' are two different subsets of \mathcal{K} , then X_K and $X_{K'}$ differ on at least one component of $\mathcal{P}[D]$, and hence $(A_K, B_K) \neq (A_{K'}, B_{K'})$. Therefore our map is injective. On the other hand, given a pair $(\mathcal{A}, \mathcal{B}) \in \mathcal{C}_{D,I}$, we have seen that for every component κ of $\mathcal{P}[D]$, the downset \mathcal{A} either contains all

maxima of κ or none of them. Now let $K_{(\mathcal{A}, \mathcal{B})}$ be the set of those components for which \mathcal{A} contains the maxima. Then $K_{(\mathcal{A}, \mathcal{B})}$ is mapped to $(\mathcal{A}, \mathcal{B})$. Hence our map is also surjective. \square

Note that the above lemma implies that there are exactly 2^d pairs in the class $\mathcal{C}_{D, I}$, where $d = |\mathcal{K}|$ denotes the number of connected components of $\mathcal{P}[D]$.

For the following, let us keep in mind that $\mathcal{C}_{D, I}$ contains ordered pairs, whereas reversals, constituting the distance between two linear extensions, are unordered pairs.

Theorem 5.14. *Let $\mathcal{D}_{\mathcal{P}}$ be the downset lattice of an arbitrary poset \mathcal{P} . The linear extension diameter of $\mathcal{D}_{\mathcal{P}}$ is bounded by a quarter of the number of pairs (A, B) of antichains of \mathcal{P} such that $\mathcal{P}[A \Delta B]$ has at least two connected components.*

Proof. Let L_1, L_2 be an arbitrary pair of linear extensions of $\mathcal{D}_{\mathcal{P}}$ and $\mathcal{C}_{D, I}$ an equivalence class as in Definition 5.12. First note that if D is empty, then the class $\mathcal{C}_{D, I}$ only consists of a single pair, namely, (I, I) . This class cannot contribute any reversal. Now let D be non-empty, and assume that $\mathcal{P}[D]$ is connected. Then, by Lemma 5.13, we know that $\mathcal{C}_{D, I}$ consists of the two pairs (A, B) and (B, A) , where $A = \text{Max}(D) \cup I$ and $B = (D \setminus \text{Max}(D)) \cup I$. But then we have $B^\downarrow \subset A^\downarrow$, that is, the two downsets form a comparable pair in $\mathcal{D}_{\mathcal{P}}$. Hence $\mathcal{C}_{D, I}$ cannot contribute any reversals if $\mathcal{P}[D]$ is connected.

Now let us assume that $\mathcal{P}[D]$ has at least two components. We claim that at most half of the pairs contained in $\mathcal{C}_{D, I}$ can be reversed between L_1 and L_2 . This means that the number of reversals that each class can contribute is at most a quarter of the number of pairs that it contains.

Recall the terminology from Lemma 5.13. For a pair $(A, B) \in \mathcal{C}_{D, I}$, let us call A^\downarrow *down in L_i* if $A^\downarrow < B^\downarrow$ in L_i , for $i = 1, 2$. The pairs in $\mathcal{C}_{D, I}$ are in bijection with the subsets of \mathcal{K} . Define the family \mathcal{F}_i as the family of subsets K of \mathcal{K} such that A_K^\downarrow is down in L_i .

A pair $\{A^\downarrow, B^\downarrow\}$ is a reversal between L_1 and L_2 exactly if A^\downarrow is down in one of the two linear extensions, but not in both. Put differently, a pair $(A_K, B_K) \in \mathcal{C}_{D, I}$ yields a reversal exactly if K is contained in one of the \mathcal{F}_i , but not in both. Thus we are interested in bounding $|\mathcal{F}_1 \Delta \mathcal{F}_2|$.

Let $K \in \mathcal{F}_i$, and $K' \subset K$. We claim that $K' \in \mathcal{F}_i$. Indeed, by definition of X_K we have $(X_{K'} \cup I)^\downarrow \subset (X_K \cup I)^\downarrow$, that is, $A_{K'}^\downarrow \subset A_K^\downarrow$. Analogously it holds $(X_K^c \cup I)^\downarrow \subset (X_{K'}^c \cup I)^\downarrow$ and hence $B_K^\downarrow \subset B_{K'}^\downarrow$. Thus from $A_K^\downarrow < B_K^\downarrow$ in L_i it follows that $A_{K'}^\downarrow < A_K^\downarrow < B_K^\downarrow < B_{K'}^\downarrow$ in L_i . Therefore the families \mathcal{F}_1 and \mathcal{F}_2 are closed downwards, and we can apply Kleitman's Lemma as

in the proof of Theorem 5.9. This yields $|\mathcal{F}_1 \cap \mathcal{F}_2| \geq 2^{-d} \cdot |\mathcal{F}_1| \cdot |\mathcal{F}_2|$, where $d = |\mathcal{K}|$.

Observe that if $(A, B) \in \mathcal{C}_{D,I}$ is associated with the set $K \subseteq \mathcal{K}$, then (B, A) is associated with the set $K^c = \mathcal{K} \setminus K$. Now in each L_i , either A^\downarrow is down or B^\downarrow is down. Thus $K \in \mathcal{F}_i$ exactly if $K^c \notin \mathcal{F}_i$. Hence $|\mathcal{F}_1| = |\mathcal{F}_2| = 2^{d-1}$. It follows that $|\mathcal{F}_1 \cap \mathcal{F}_2| \geq 2^{d-2}$. Similarly, if A^\downarrow is down in both L_i , then B^\downarrow is down in neither, and vice versa. This means that $K \in \mathcal{F}_1 \cap \mathcal{F}_2$ exactly if $K^c \in (\mathcal{F}_1 \cup \mathcal{F}_2)^c$. Therefore $|\mathcal{F}_1 \cap \mathcal{F}_2| = |(\mathcal{F}_1 \cup \mathcal{F}_2)^c|$. We conclude that

$$|\mathcal{F}_1 \Delta \mathcal{F}_2| = 2^d - |\mathcal{F}_1 \cap \mathcal{F}_2| - |(\mathcal{F}_1 \cup \mathcal{F}_2)^c| \leq 2^d - 2^{d-2} - 2^{d-2} = 2^{d-1}.$$

We have thus shown that each class $\mathcal{C}_{D,I}$ can contribute at most 2^{d-2} reversals between two linear extensions of $\mathcal{D}_{\mathcal{P}}$ as claimed. \square

Next we are going to prove that the bound of the above theorem is tight in the case of a 2-dimensional poset \mathcal{P} . To do so, we define revlex linear extensions of $\mathcal{D}_{\mathcal{P}}$. In analogy to Definition 5.2, we have:

Definition 5.15. *Let σ be a linear extension of \mathcal{P} . For a set $S \subseteq \mathcal{P}$, let $\max_{\sigma} S$ be the element of S which is largest in σ . The σ -revlex order $<_{\sigma}$ on the pairs $\{A^\downarrow, B^\downarrow\}$ of downsets of \mathcal{P} is defined as follows:*

$$A^\downarrow <_{\sigma} B^\downarrow \iff \max_{\sigma}(A \Delta B) \in B.$$

Lemma 5.16. *Let \mathcal{P} be a poset and let σ be a linear extension of \mathcal{P} . The relation of being in σ -revlex order defines a linear extension of $\mathcal{D}_{\mathcal{P}}$.*

Proof. Since $\max_{\sigma}(A \Delta B) = \max_{\sigma}(A^\downarrow \Delta B^\downarrow)$, we could equivalently define the relation $<_{\sigma}$ directly via the downsets. Thus this relation is just a restriction of the σ -revlex order on all subsets of \mathcal{P} (see Definition 5.2) to the downsets of \mathcal{P} . In Lemma 5.3, we proved that the σ -revlex order on all subsets is a linear order and that it extends the inclusion relation. This carries over to the restriction to the downsets of \mathcal{P} . \square

The above lemma was shown in [32] for the special case that \mathcal{P} equals B_n .

Definition 5.17. *Let \mathcal{P} be a 2-dimensional poset, and let σ be a linear extension of \mathcal{P} . We call the linear extension of $\mathcal{D}_{\mathcal{P}}$ which is given by the σ -revlex order a revlex linear extension of $\mathcal{D}_{\mathcal{P}}$.*

Now let σ be contained in a realizer of \mathcal{P} . We denote by $\bar{\sigma}$ the unique partner of σ in a realizer of \mathcal{P} . The pair $L_{\sigma}, L_{\bar{\sigma}}$ is a revlex pair of linear extensions of $\mathcal{D}_{\mathcal{P}}$.

Recall that all incomparable pairs of \mathcal{P} are reversals between σ and $\bar{\sigma}$. Thus σ and $\bar{\sigma}$ form a diametral pair of linear extensions of \mathcal{P} . Recall also that by Corollary 1.6, the linear extensions σ and $\bar{\sigma}$ are non-separating, that is, there are no three elements $u, v, w \in \mathcal{P}$ such that $u < w$ and $v \parallel u, w$ in \mathcal{P} , but $u < v < w$ in σ or $\bar{\sigma}$.

Theorem 5.18. *Let \mathcal{P} be a 2-dimensional poset, and let $\{\sigma, \bar{\sigma}\}$ be a realizer of \mathcal{P} . Then the revlex pair $L_\sigma, L_{\bar{\sigma}}$ is a diametral pair of linear extensions of $\mathcal{D}_{\mathcal{P}}$.*

Proof. We will show that for each class $\mathcal{C}_{D,I}$, a revlex pair realizes the maximum possible number of reversals. Fix an arbitrary class $\mathcal{C}_{D,I}$. We have seen in the proof of Theorem 5.14 that $\mathcal{C}_{D,I}$ cannot contribute any reversals if $\mathcal{P}[D]$ is empty or consists of only one component. Now let \mathcal{K} be the set of connected components of $\mathcal{P}[D]$ as before, and assume that $|\mathcal{K}| = d \geq 2$.

A pair $(A, B) \in \mathcal{C}_{D,I}$ yields a reversal between L_σ and $L_{\bar{\sigma}}$ exactly if one of the two elements $\max_\sigma(A\Delta B)$ and $\max_{\bar{\sigma}}(A\Delta B)$ is contained in A , and the other in B . Since (A, B) and (B, A) contribute only one reversal, let us count the number of pairs (A, B) with $\max_\sigma(A\Delta B) \in A$ and $\max_{\bar{\sigma}}(A\Delta B) \in B$.

Note that the σ -maximum of $A\Delta B$ belongs to the antichain $\text{Max}(D)$, which is completely reversed between σ and $\bar{\sigma}$. It follows that we have $\max_\sigma \text{Max}(D) = \min_{\bar{\sigma}} \text{Max}(D)$ and $\max_{\bar{\sigma}} \text{Max}(D) = \min_\sigma \text{Max}(D)$. We claim that $\max_\sigma \text{Max}(D)$ and $\min_\sigma \text{Max}(D)$ lie in different components of $\mathcal{P}[D]$.

Suppose for contradiction that there is a component κ of $\mathcal{P}[D]$ containing both $\max_\sigma \text{Max}(D)$ and $\min_\sigma \text{Max}(D)$. Since we assumed that $\mathcal{P}[D]$ has at least two components, we can choose an element $x \in \text{Max}(D)$ which is not contained in κ and thus has no relation to any element of κ . Then $\min_\sigma \text{Max}(D) < x < \max_\sigma \text{Max}(D)$ in σ . Denote by κ_1 the set of elements of κ which are σ -smaller than x , and by κ_2 the set of elements of κ which are σ -larger than x . Both sets are non-empty since $\min_\sigma \text{Max}(D) \in \kappa_1$ and $\max_\sigma \text{Max}(D) \in \kappa_2$. Since κ is a connected component of $\mathcal{P}[D]$, there are elements $u \in \kappa_1$ and $v \in \kappa_2$ with $u < v$ in \mathcal{P} . But then $u < x < v$ in σ with $x \parallel u, v$ in \mathcal{P} , and this is a contradiction because σ is non-separating.

We have shown that $\max_\sigma(A\Delta B)$ and $\max_{\bar{\sigma}}(A\Delta B)$ lie in different components of $\mathcal{P}[D]$, say, $\max_\sigma(A\Delta B) \in \kappa$ and $\max_{\bar{\sigma}}(A\Delta B) \in \lambda$ with $\kappa, \lambda \in \mathcal{K}$. Now let $K \subseteq \mathcal{K}$ be a set of components with $\kappa \in K$ and $\lambda \notin K$. Consider the pair $(A_K, B_K) \in \mathcal{C}_{D,I}$. By definition, we have $\max_\sigma(A\Delta B) \in A_K$ and $\max_{\bar{\sigma}}(A\Delta B) \in B_K$. Thus the pair (A_K, B_K) contributes a reversal between L_σ and $L_{\bar{\sigma}}$. There are 2^{d-2} possibilities of choosing K . Thus every class $\mathcal{C}_{D,I}$

with $d \geq 2$ contributes at least 2^{d-2} reversals between L_σ and $L_{\bar{\sigma}}$. We have seen in Theorem 5.14 that this is maximal. Therefore L_σ and $L_{\bar{\sigma}}$ are a diametral pair of linear extensions of $\mathcal{D}_\mathcal{P}$. \square

We have now proved that, for a 2-dimensional poset \mathcal{P} , every class $\mathcal{C}_{D,I}$ for which $\mathcal{P}[D]$ has at least two components contributes exactly 2^{d-2} reversals between a diametral pair of linear extensions of $\mathcal{D}_\mathcal{P}$. Since there are 2^d pairs in $\mathcal{C}_{D,I}$, we have the following corollary:

Corollary 5.19. *Let $\mathcal{D}_\mathcal{P}$ be the downset lattice of a 2-dimensional poset \mathcal{P} . The linear extension diameter of $\mathcal{D}_\mathcal{P}$ equals a quarter of the number of pairs (A, B) of antichains of \mathcal{P} such that $\mathcal{P}[A\Delta B]$ has at least two connected components.*

Note that the proof of Conjecture 5.1, that is, the proof that $\text{led}(B_n)$ equals $2^{2n-2} - (n+1)2^{n-2}$, follows from this corollary. This is implied more directly by Theorem 5.23.

5.2.2 Diametral Pairs are Revlex Pairs

We are now ready to characterize the diametral pairs of downset lattices of 2-dimensional posets.

Theorem 5.20. *Let \mathcal{P} be a 2-dimensional poset, and let L, \bar{L} be a diametral pair of linear extensions of $\mathcal{D}_\mathcal{P}$. Let σ be the linear extension of \mathcal{P} defined by the order of the downsets x^\downarrow for $x \in \mathcal{P}$ in L . Then σ is contained in a realizer $\{\sigma, \bar{\sigma}\}$ of \mathcal{P} , and we have $L = L_\sigma$ and $\bar{L} = L_{\bar{\sigma}}$.*

Proof. We again use the equivalence classes from Definition 5.12. For each incomparable pair $x, y \in \mathcal{P}$, consider the equivalence class $\mathcal{C}_{D,I}$ defined by $D = \{x, y\}$ and $I = \emptyset$. Then $\mathcal{P}[D]$ consists of two singletons. Since L, \bar{L} is a diametral pair, this class must contribute $2^{2-2} = 1$ reversal, and thus x^\downarrow and y^\downarrow must appear in opposite order in L and \bar{L} . Recall that σ is the linear extension of \mathcal{P} defined by the order of the downsets x^\downarrow in L . Let us denote the linear extension defined analogously for \bar{L} by $\bar{\sigma}$. We have seen that every incomparable pair of elements of \mathcal{P} must be a reversal between σ and $\bar{\sigma}$. Hence, $\{\sigma, \bar{\sigma}\}$ is a realizer of \mathcal{P} .

In the following, we use induction on the cardinality of the downsets of \mathcal{P} to show that $L = L_\sigma$ and $\bar{L} = L_{\bar{\sigma}}$, in analogy to the proof of Theorem 5.11. More precisely, we show that each downset of cardinality k is in σ -revlex order in L and in $\bar{\sigma}$ -revlex order in \bar{L} with all downsets of cardinality $\leq k$, by induction on k . We use the fact that every equivalence class $\mathcal{C}_{D,I}$ for which $\mathcal{P}[D]$ is disconnected contributes exactly 2^{d-2} reversals between L

and \bar{L} . Note that we have settled the base case already: All downsets of cardinality 1 are of the form x^\downarrow for some minimal element $x \in \mathcal{P}$, and we have shown in the previous paragraph that these behave as expected.

Let L^k be the restriction of L to the sets of cardinality at most k . Our induction hypothesis is that $L^{k-1} = L_\sigma^{k-1}$ and $\bar{L}^{k-1} = L_{\bar{\sigma}}^{k-1}$. We structure the induction step as follows: We first show that each set of size k is in the desired order in L^k and \bar{L}^k with all sets of smaller size. This will be the main part of the proof. Then we show that all pairs of sets of equal size k are in the desired order in L^k and \bar{L}^k .

We use the following expression: Let X, Y, Z be downsets of \mathcal{P} such that X and Z have cardinality $k-1$ and Y has cardinality k . If $X < Z$ is a cover relation in L^{k-1} and $X < Y < Z$ in L^k , then we say that Y sits in the *slot* between X and Z in L^k .

Let A^\downarrow be a downset of cardinality k of \mathcal{P} . Let \tilde{A}^\downarrow be the subset of A^\downarrow which is largest in L^{k-1} , and let B^\downarrow be its immediate successor in L^{k-1} .

Claim. A^\downarrow needs to sit in the slot between \tilde{A}^\downarrow and B^\downarrow in L^k .

Proving this claim requires some technical details. Here is an outline of what we are going to do: We first locate the elements $\max_\sigma(A^\downarrow \triangle B^\downarrow) = \max_\sigma(A \triangle B)$ and $\max_{\bar{\sigma}}(A^\downarrow \triangle B^\downarrow) = \max_{\bar{\sigma}}(A \triangle B)$. Using these we see that $\{A^\downarrow, B^\downarrow\}$ needs to be a reversal between L and \bar{L} . From the order of A^\downarrow and B^\downarrow in \bar{L} we can finally deduce that $A^\downarrow < B^\downarrow$ in L^k .

We have $\tilde{A}^\downarrow = A^\downarrow \setminus a$ for some $a \in A$. All subsets of A^\downarrow are in σ -revlex order in L by induction. So we know that if \hat{A}^\downarrow is a second subset of cardinality $k-1$ of A^\downarrow , then the element of A^\downarrow that is missing in \hat{A}^\downarrow is σ -smaller than the element of A^\downarrow that is missing in \tilde{A}^\downarrow . Thus we can conclude that $a = \min_\sigma A$. Since the antichain A is completely reversed between σ and $\bar{\sigma}$, it follows that $a = \max_{\bar{\sigma}} A$.

Now observe that since $\tilde{A}^\downarrow < B^\downarrow$ in L^{k-1} and $|\tilde{A}^\downarrow| = k-1$, we have $\tilde{A}^\downarrow \parallel B^\downarrow$. By induction we know that $\max_\sigma(\tilde{A}^\downarrow \triangle B^\downarrow) \in B^\downarrow$. Let b be the σ -smallest element of $B^\downarrow \setminus \tilde{A}^\downarrow$ which is σ -larger than all elements of $\tilde{A}^\downarrow \setminus B^\downarrow$. Then $(\tilde{A}^\downarrow \cap B^\downarrow) \cup b$ is a downset of \mathcal{P} . By induction, $\tilde{A}^\downarrow < (\tilde{A}^\downarrow \cap B^\downarrow) \cup b$ in L^{k-1} . Since $(\tilde{A}^\downarrow \cap B^\downarrow) \cup b \subseteq B^\downarrow$ we must have $(\tilde{A}^\downarrow \cap B^\downarrow) \cup b = B^\downarrow$ by the choice of B^\downarrow . Thus $B^\downarrow \setminus \tilde{A}^\downarrow = \{b\}$ and $\max_\sigma(\tilde{A}^\downarrow \triangle B^\downarrow) = b$. We will use the next three paragraphs to show that $\max_\sigma(A^\downarrow \triangle B^\downarrow) = b$ and $\max_{\bar{\sigma}}(A^\downarrow \triangle B^\downarrow) = a$.

Note that $B^\downarrow \not\subseteq A^\downarrow$ by the choice of \tilde{A}^\downarrow , and thus $B^\downarrow \setminus A^\downarrow = \{b\}$. Hence $b \not\leq a$ in \mathcal{P} . On the other hand, $a \not\leq b$ in \mathcal{P} because otherwise $a \in B^\downarrow \setminus \tilde{A}^\downarrow$ and thus $a = b$, a contradiction. It follows that $a \parallel b$ in \mathcal{P} .

Next let us show that $a <_\sigma b$. Because of $|\tilde{A}^\downarrow| = k - 1$, we know that $\tilde{A}^\downarrow \setminus B^\downarrow \neq \emptyset$. Let $a' \in \tilde{A} \setminus B$. If $a \parallel a'$ in \mathcal{P} , then $a' \in A \setminus B$ and with $a = \min_\sigma A$ it follows that $a <_\sigma a'$. Putting this together with $a' <_\sigma b$, we get $a <_\sigma b$. If a and a' are comparable, then $a' < a$ in \mathcal{P} . Now assume for contradiction that $b <_\sigma a$. Then we have $a' <_\sigma b <_\sigma a$, with $b \parallel a'$, a and $a' < a$ in \mathcal{P} . This means that σ is a separating linear extension. But since σ is contained in the realizer $\{\sigma, \bar{\sigma}\}$ of \mathcal{P} , this is a contradiction.

We have shown that $a <_\sigma b$. We knew already that $\max_\sigma(\tilde{A}^\downarrow \Delta B^\downarrow) = b$, and because a is the only element in $A^\downarrow \setminus \tilde{A}^\downarrow$, we can conclude that $\max_\sigma(A^\downarrow \Delta B^\downarrow) = b$. Also, since $a \parallel b$ in \mathcal{P} we now know that $a >_{\bar{\sigma}} b$. Because $a = \max_{\bar{\sigma}} A$ and $B^\downarrow \setminus A^\downarrow = \{b\}$, we have $\max_{\bar{\sigma}}(A^\downarrow \Delta B^\downarrow) = a$.

Let us now consider the class $\mathcal{C}_{D,I}$ defined by $D = A \Delta B$ and $I = A \cap B$. The elements a and b lie in different components of $\mathcal{P}[D]$, because $B \setminus A = \{b\}$ and $a \parallel b$ in \mathcal{P} . So we may assume that $a \in \alpha$ and $b \in \beta$, where α and β are different elements from the set \mathcal{K} of components of $\mathcal{P}[D]$. As before, set $d = |\mathcal{K}|$.

Observe that $|A^\downarrow \cup B^\downarrow| = |A^\downarrow \cup \{b\}| = k + 1$. Now choose a subset $K \subset \mathcal{K}$ with $\alpha, \beta \notin K$. For the corresponding downset $(X_K \cup I)^\downarrow = A_K^\downarrow$ we have $A_K^\downarrow \subseteq A^\downarrow \cup B^\downarrow$. Since $a, b \notin A_K^\downarrow$, we can apply the induction hypothesis to A_K^\downarrow . We can also apply it to the set $(b \cup I)^\downarrow$. It holds that $\max_\sigma(A_K \Delta (b \cup I)) = b$, so we have $A_K^\downarrow < (b \cup I)^\downarrow$ in L by induction. Let $X_K^c \cup I = B_K$. Then we have $(b \cup I)^\downarrow \subseteq B_K^\downarrow$, and thus $A_K^\downarrow < B_K^\downarrow$ in L by transitivity. Analogously, $\max_{\bar{\sigma}}(A_K \Delta (a \cup I)) = a$ holds. Hence, $A_K^\downarrow < (a \cup I)^\downarrow$ in \bar{L} by induction and thus $A_K^\downarrow < (a \cup I)^\downarrow < (X_K^c \cup I)^\downarrow = B_K^\downarrow$ in \bar{L} .

It follows that A_K^\downarrow is down in L and \bar{L} for every $K \subset \mathcal{K}$ with $\alpha, \beta \notin K$. Thus $(A_K^\downarrow, B_K^\downarrow)$ cannot yield a reversal between L and \bar{L} , and neither can $(B_K^\downarrow, A_K^\downarrow)$. There are 2^{d-2} possibilities to choose K . Thus we have exhibited $2 \cdot 2^{d-2}$ pairs in $\mathcal{C}_{D,I}$ which do not contribute a reversal between L and \bar{L} . From the fact we remarked after Theorem 5.18 it follows that all other pairs in $\mathcal{C}_{D,I}$ have to contribute reversals.

Consequently, all subsets K of \mathcal{K} containing exactly one of the two components α and β need to contribute a reversal, or equivalently, all A_K^\downarrow which contain exactly one of the two elements a, b need to be down in exactly one of the two linear extensions. In particular, our set A^\downarrow needs to be down relative to $\mathcal{C}_{D,I}$ in exactly one of the two linear extensions.

It turns out that A^\downarrow cannot be down in \bar{L} : For $(I \cup a)^\downarrow \subset A^\downarrow$, we have $\max_{\bar{\sigma}}(B \Delta (I \cup a)) = a$. So we have $B^\downarrow < (I \cup a)^\downarrow$ in \bar{L} by induction and thus $B^\downarrow < A^\downarrow$ in \bar{L} by transitivity. Hence it follows that $A^\downarrow < B^\downarrow$ in L . This proves our claim that A^\downarrow has to sit in the slot between \tilde{A}^\downarrow and B^\downarrow in L^k . \triangle

Because $\max_\sigma(A\Delta B) = b \in B$, and $A^\downarrow < B^\downarrow$ in L as shown in the claim, we now know that A^\downarrow is in σ -revlex order with B^\downarrow in L^k . Since the slot after \tilde{A}^\downarrow is the lowest possible position for A^\downarrow in L^k , it follows from the transitivity of the σ -revlex order that A^\downarrow is in σ -revlex order in L^k with all sets of smaller cardinality. By reversing the roles of L and \bar{L} , we obtain that A^\downarrow is in $\bar{\sigma}$ -revlex order in \bar{L}^k with all sets of smaller cardinality. Next we show that all pairs of sets with equal cardinality k are in σ -revlex order in L^k .

Let $A_i^\downarrow, A_j^\downarrow \in \mathcal{D}_{\mathcal{P}}$ be two downsets of the same cardinality k . If they are inserted into different slots in L , they are in σ -revlex order by transitivity. If they belong into the same slot, this means that they have the same largest $(k-1)$ -subset \tilde{A}^\downarrow in L . So their symmetric difference contains only two elements, say, $A_i^\downarrow \Delta A_j^\downarrow = \{a_i, a_j\}$. We have $a_i = A_i^\downarrow \setminus \tilde{A}^\downarrow = \min_\sigma A_i = \max_{\bar{\sigma}} A_i$ and $a_j = A_j^\downarrow \setminus \tilde{A}^\downarrow = \min_\sigma A_j = \max_{\bar{\sigma}} A_j$. Note that a_i and a_j have to be incomparable in \mathcal{P} , and assume that $a_i <_\sigma a_j$, thus $a_j <_{\bar{\sigma}} a_i$. Then for the pair $\{A_j^\downarrow, a_i^\downarrow\}$ we know $\max_{\bar{\sigma}}(A_j^\downarrow \Delta a_i^\downarrow) = a_i$. Hence, by induction, $A_j^\downarrow < a_i^\downarrow < A_i^\downarrow$ in \bar{L} . But since the equivalence class containing $(A_i^\downarrow, A_j^\downarrow)$ needs to contribute $2^{2-2} = 1$ reversal between L and \bar{L} , and this can only be the pair $\{A_i^\downarrow, A_j^\downarrow\}$, we know that we must have $A_i^\downarrow < A_j^\downarrow$ in L . Because $\max_\sigma A_i \Delta A_j = a_j$, this means that A_i^\downarrow and A_j^\downarrow are in revlex order in L .

We can apply the same argument (with the roles of L and \bar{L} reversed) to show that all pairs of downsets of equal cardinality k are in $\bar{\sigma}$ -revlex order in \bar{L} . By induction we conclude that $L = L_\sigma$ and $\bar{L} = L_{\bar{\sigma}}$. \square

5.2.3 Computing the Linear Extension Diameter

We have seen in Theorem 4.6 that it is NP-complete to compute the linear extension diameter of a general poset. That is, the linear extension diameter of a general poset \mathcal{P} cannot be computed in a running time polynomial in $|\mathcal{P}|$ (unless $P = NP$). But with the results of the previous section, the problem is tractable if the given poset is a downset lattice \mathcal{D} of a 2-dimensional poset.

In fact, we can construct any diametral pair of linear extensions of \mathcal{D} in time polynomial in $|\mathcal{D}|$. To see this, we use Birkhoff's representation theorem of distributive lattices, which says that from a distributive lattice \mathcal{D} one can obtain \mathcal{P} with $\mathcal{D} = \mathcal{D}_{\mathcal{P}}$ as the poset induced by the join-irreducible elements of \mathcal{D} (see Subsection 1.1).

By the characterization of 2-dimensional posets given in Theorem 1.5, finding a realizer $\{\sigma, \bar{\sigma}\}$ of \mathcal{P} amounts to finding a transitive orientation of

its incomparability graph. This can be done in time linear in $|\mathcal{P}|$ by the algorithm given in [40]. With the definition of the σ -revlex order we can compute L_σ and $L_{\bar{\sigma}}$. This is a diametral pair of linear extensions of \mathcal{D} by Theorem 5.18. We know by Theorem 5.20 that all diametral pairs arise in this way. The linear extension diameter of \mathcal{D} can now be computed by simply checking for all pairs of elements of \mathcal{D} whether they form a reversal between L_σ and $L_{\bar{\sigma}}$.

In this subsection we show that we can in fact do much better: For a 2-dimensional poset \mathcal{P} , we can compute the linear extension diameter of $\mathcal{D}_\mathcal{P}$ in time polynomial in $|\mathcal{P}|$. Note that in general, \mathcal{P} can have exponentially many downsets, e.g., if \mathcal{P} is an antichain or has only very few relations. Hence $|\mathcal{D}_\mathcal{P}|$ may be exponentially larger than $|\mathcal{P}|$.

For the proofs of this subsection, we mainly consider antichains instead of downsets, again using the canonical bijection between them. It is known that the antichains of a 2-dimensional poset can be counted in polynomial time, see [57] or [43]. We give a proof in the lemma below since the methods we use for proving the following theorem rely on the same ideas.

Lemma 5.21 ([57]). *Let \mathcal{P} be a 2-dimensional poset. Denote by $A(\mathcal{P})$ the set of antichains of \mathcal{P} and let $a(\mathcal{P}) = |A(\mathcal{P})|$. Then $a(\mathcal{P})$ can be computed in time $\mathcal{O}(|\mathcal{P}|^2)$.*

Proof. Let $\sigma = x_1x_2\dots x_n$ be a non-separating linear extension of \mathcal{P} . Denote by $A(x_i)$ the set of antichains of \mathcal{P} which contain x_i as σ -largest element, and let $a(x_i) = |A(x_i)|$. We will use a dynamic programming approach to compute $a(x_i)$ for all i .

To start with, we have $a(x_1) = 1$. Now suppose we have computed $a(x_j)$ for all $j < i$. The main observation is that for any $A \in A(x_j)$ with $j < i$ and $x_i \parallel x_j$, the set $x_i \cup A$ is again an antichain. This holds because any $x_k \in A$ with $x_k < x_i$ would yield a contradiction to σ being non-separating. Therefore we have

$$a(x_i) = 1 + \sum_{j < i, x_i \parallel x_j} a(x_j),$$

where the 1 accounts for the antichain $\{x_i\}$. Consequently, we obtain the number of all antichains of \mathcal{P} as $a(\mathcal{P}) = 1 + \sum_i a(x_i)$, where the 1 accounts for the empty set.

With the above formula, the evaluation of $a(x_i)$ can be done in linear time for each i . Thus $a(\mathcal{P})$ can be computed in quadratic time. \square

Theorem 5.22. *The linear extension diameter of the downset lattice $\mathcal{D}_\mathcal{P}$ of a 2-dimensional poset \mathcal{P} can be computed in time $\mathcal{O}(|\mathcal{P}|^5)$.*

Proof. From Corollary 5.19, we know that $\text{led}(\mathcal{D}_{\mathcal{P}})$ equals a quarter of the number of pairs (A, B) of antichains of \mathcal{P} such that $\mathcal{P}[A\Delta B]$ has at least two connected components. For a pair (A, B) of antichains of \mathcal{P} , we set $D = A\Delta B$ and $I = A \cap B$.

We will count four different classes of pairs of antichains. Let α be the number of *all* ordered pairs of antichains of \mathcal{P} . Let β be the number of pairs (A, B) with $D = \emptyset$, and γ the number of pairs with $|D| = 1$. Finally, let δ be the number of pairs such that $|D| > 1$ and $\mathcal{P}[D]$ is connected. Then $\text{led}(\mathcal{D}_{\mathcal{P}}) = \frac{1}{4}(\alpha - \beta - \gamma - \delta)$.

We have $\alpha = a(\mathcal{P})^2$. Moreover, the pairs we count for β are just the pairs (A, A) , so $\beta = a(\mathcal{P})$.

For the following, let $\sigma = x_1x_2 \dots x_n$ be a non-separating linear extension of \mathcal{P} . We denote by $[x_i, x_k]$ the set $\{x_i, x_{i+1}, \dots, x_k\}$, and by (x_i, x_k) the set $\{x_{i+1}, \dots, x_{k-1}\}$. We use analogous notions for “half-open intervals” of σ .

To obtain γ , we count the pairs $(A, A - x)$, where A is a non-empty antichain in \mathcal{P} , and x is an element of A . Thus we count each A exactly $|A|$ times. We want to refine the ideas of the proof of Lemma 5.21 to keep track of the sizes of the antichains. Therefore we define vectors $s(x_i)$, where $s_r(x_i)$ is the number of antichains of cardinality r in $A(x_i)$.

We can recursively compute $s(x_i)$ for $i = 1, 2, \dots, n$ as follows: The first entry of each $s(x_i)$ is 1, counting the antichain $\{x_i\}$. For the other entries we have

$$s_r(x_i) = \sum_{j < i: x_j \parallel x_i} s_{r-1}(x_j).$$

Then the number of pairs $(A, A - x)$ equals $\sum_r r \sum_i s_r(x_i)$. Now γ is twice this number, since we also need to count the pairs $(A - x, A)$.

The most difficult part is to compute δ , the number of pairs (A, B) of antichains of \mathcal{P} such that $|D| > 1$ and $\mathcal{P}[D]$ is connected. Let us look at the structure of $\mathcal{P}[D \cup I]$, see Figure 5.2.3.




Figure 5.4: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in δ .

We know that $\mathcal{P}[I]$ is an antichain and $\mathcal{P}[D]$ consists of two antichains, $\text{Max}(D)$ and $\text{Min}(D)$, which are both non-empty by definition of δ . We claim that each element of I is either σ -smaller than all elements of D ,

or σ -larger than all elements of D . Indeed, suppose there is an $x \in I$ and $u, v \in D$ with $u <_\sigma x <_\sigma v$. Then since $\mathcal{P}[D]$ is connected, we can find $u', v' \in D$ with $u' < v'$ in \mathcal{P} and $u' <_\sigma x <_\sigma v'$. But since there are no relations between x and the elements of D , this means that σ is separating. This contradiction proves our claim. Thus σ can be partitioned into three intervals, such that the elements of D are all contained in the middle interval, and I is split up into two parts: The first interval, I^{left} , and the third interval, I^{right} .

Now define $\delta(k, \ell)$ as the number of pairs counted for δ which fulfill $\max_\sigma \text{Min}(D) = x_k$ and $\max_\sigma \text{Max}(D) = x_\ell$. In addition, we require that $x_\ell = \max_\sigma \mathcal{A} \cup \mathcal{B}$, which means that I^{right} is empty. We split up $\delta(k, \ell)$ into the number $\delta_1(k, \ell)$ of pairs for which $\mathcal{P}[D]$ has only one maximum and the number $\delta_2(k, \ell)$ of pairs for which it has several.

To compute $\delta_1(k, \ell)$, we have to count the number of possibilities to choose the antichain $\text{Min}(D)$ and the antichain I^{left} . By definition we have $\max_\sigma \text{Min}(D) = x_k$. Suppose that $\min_\sigma \text{Min}(D) = x_i$ as in Figure 5.5.




Figure 5.5: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in $\delta_1(k, \ell)$.

Let us define $\mathcal{P}_{i,k,\ell}$ as the poset induced by the elements $x_j \in (x_i, x_k)$ with $x_j < x_\ell$ and $\{x_i, x_j, x_k\} \in A(\mathcal{P})$. Then to choose $\text{Min}(D)$, we have to choose an antichain in $\mathcal{P}_{i,k,\ell}$.

Once a set D is fixed, it remains to choose I^{left} to determine a pair of antichains counted in $\delta_1(k, \ell)$. Let $\mathcal{P}_{i,\ell}^{\text{left}}$ be the poset induced by the elements $x \in \mathcal{P}$ with $x \in [x_1, x_i)$ and $x \parallel x_\ell$. We claim that the sets which can be chosen as I^{left} are exactly the antichains of $\mathcal{P}_{i,\ell}^{\text{left}}$.

By definition, each element $x \in I^{\text{left}}$ is σ -smaller than x_i . We have to choose x so that it is incomparable to all elements in D . It is clear that $x \in [x_1, x_i)$ cannot be larger in \mathcal{P} than any element in $D \subseteq [x_i, x_\ell]$. Now if we choose x incomparable to x_ℓ , it cannot be smaller than any element in D , either. Thus to choose I , we have to choose an antichain in $\mathcal{P}_{i,\ell}^{\text{left}}$ as claimed.

Altogether we have

$$\delta_1(k, \ell) = \sum_{x_i \in [x_1, x_k], \{x_i, x_k\} \in A(\mathcal{P}), x_i < x_\ell} a(\mathcal{P}_{i,k,\ell}) \cdot a(\mathcal{P}_{i,\ell}^{\text{left}}).$$

It remains to compute $\delta_2(k, \ell)$, the number of pairs (A, B) counted in $\delta(k, \ell)$ for which $\mathcal{P}[D]$ has several maxima. We want to cut off the σ -largest maximum and (possibly) some minima and recursively use values $\delta_2(k', \ell')$ and $\delta_1(k', \ell')$ that we have calculated already (cf. Figure 5.6).




Figure 5.6: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in $\delta_2(k, \ell)$.

For a pair (A, B) counted in $\delta_2(k, \ell)$, the second largest maximum of $\mathcal{P}[D]$ in σ is an element $x_{\ell'} \in [x_1, x_{\ell})$ with $x_{\ell'} \parallel x_{\ell}$. In general, $\mathcal{P}[D]$ is not connected after deletion of x_{ℓ} . But since $\mathcal{P}[D]$ was connected originally, $\text{Min}(D)$ contains an element $x_{k'}$ with $x_{k'} < x_{\ell}$ and $x_{k'} < x_{\ell'}$. Let $x_{k'}$ be the σ -smallest such element. There can be more elements in $[x_{k'}, x_k]$ which are part of $\text{Min}(D)$. With the same reasoning as for $\delta_1(k, \ell)$, these form exactly the antichains in $\mathcal{P}_{k', k, \ell}$. So we have

$$\delta_2(k, \ell) = \sum_{x_{\ell'} \in [x_1, x_{\ell}), x_{\ell'} \parallel x_{\ell}} \sum_{x_{k'} \in [x_1, x_{\ell'}), x_{k'} < x_{\ell'}, x_{\ell}} (\delta_1(k', \ell') + \delta_2(k', \ell')) \cdot a(\mathcal{P}_{k', k, \ell}).$$

Recall that for the pairs counted in $\delta(k, \ell)$, we required that I^{right} is empty. So to compute δ , we have to weight every pair with the number of possible choices for I^{right} . Let $\mathcal{P}_{k, \ell}^{\text{right}}$ be the poset induced by the elements of \mathcal{P} which are in $(x_{\ell}, x_n]$ and incomparable to x_k in \mathcal{P} . We claim that the sets eligible for I^{right} are exactly the antichains of $\mathcal{P}_{k, \ell}^{\text{right}}$.

Each element $x \in I^{\text{right}}$ has to be incomparable to all elements of D . Since $x \in (x_{\ell}, x_n]$, it cannot be smaller than any element in D . If we choose x incomparable to x_k , then x cannot be larger than any element of D either: If $x > y$ for some element $y \in \text{Min}(D)$, then y has to be σ -smaller than x_k , which makes σ separating. Thus to choose I^{right} we have to choose an antichain in $\mathcal{P}_{k, \ell}^{\text{right}}$ as claimed. Hence we obtain δ as follows:

$$\delta = \sum_{k, \ell} (\delta_1(k, \ell) + \delta_2(k, \ell)) \cdot a(\mathcal{P}_{k, \ell}^{\text{right}}).$$

To finish the proof of the theorem, let us consider the overall running time for the computation of $\text{led}(\mathcal{D}_{\mathcal{P}})$. From Lemma 5.21 we know that the number of antichains of a poset can be computed in quadratic time. Thus

α and β can be determined in quadratic time. To compute γ we need to compute $s_r(x_i)$ for $r = 1, \dots, n$ and $i = 1, \dots, n$. For each value $s_r(x_i)$, our formula can be evaluated in linear time. Thus it takes $O(n^3)$ to determine γ .

For the computation of δ we first determine, in a preprocessing step, the values $a(\mathcal{P}_{i,k,\ell})$ for all triples i, k, ℓ . Given such a triple, we can build $\mathcal{P}_{i,k,\ell}$ in linear time, and then compute $a(\mathcal{P}_{i,k,\ell})$ in quadratic time using Lemma 5.21 again. Altogether, this can be done in $O(n^5)$. Similarly, we can determine $a(\mathcal{P}_{k,\ell}^{\text{left}})$ and $a(\mathcal{P}_{k,\ell}^{\text{right}})$ for all pairs k, ℓ in a preprocessing step taking time $O(n^4)$. Then for each pair k, ℓ , we can compute $\delta_1(k, \ell)$ and $\delta_2(k, \ell)$ in linear time. Thus it takes $O(n^3)$ to obtain all these values. In the end, we can put them together in quadratic time to obtain δ .

The overall running time is the maximum over all these separate steps. We conclude that $\text{led}(D_{\mathcal{P}})$ can be computed in time $O(n^5)$. \square

In the previous theorem we showed how to compute $\text{led}(D_{\mathcal{P}})$ for a 2-dimensional poset \mathcal{P} , but we could not give an explicit formula like the one we have for the Boolean lattice. This is possible for the special case where \mathcal{P} is a disjoint union of chains. Such a lattice $\mathcal{D}_{\mathcal{P}}$ is also known as a factor lattice of integers: If $\mathcal{P} = C_1 \cup \dots \cup C_w$ with $|C_i| = \ell_i$, we can associate each chain with a prime number p_i . Then $\mathcal{D}_{\mathcal{P}}$ is the lattice of all factors of $m = \prod_{i=1}^w p_i^{\ell_i}$, ordered by divisibility.

Theorem 5.23. *If $\mathcal{P} = C_1 \cup \dots \cup C_w$ with $|C_i| = \ell_i$ is a disjoint union of chains, then the linear extension diameter of $\mathcal{D}_{\mathcal{P}}$ equals*

$$\frac{1}{4} \cdot \left(\left(\prod_{i=1}^w (\ell_i + 1) \right)^2 - \sum_{k=1}^w (\ell_k + 1) \ell_k \cdot \prod_{i \neq k} (\ell_i + 1) - \prod_{i=1}^w (\ell_i + 1) \right).$$

Proof. From Corollary 5.19 we know that $\text{led}(\mathcal{D}_{\mathcal{P}})$ equals a quarter of the number of pairs (A, B) of antichains of \mathcal{P} such that $\mathcal{P}[A \triangle B]$ has at least two connected components. So we need to count the pairs of antichains of \mathcal{P} which differ on at least two of the C_i . We will count *all* pairs of antichains and subtract from it the number of pairs differing on zero or one chain.

To choose one antichain, we have $\ell_i + 1$ choices in each C_i . So \mathcal{P} contains exactly $\prod_{i=1}^w (\ell_i + 1)$ antichains, and this is also the number of pairs of antichains differing on zero chains. The number of all pairs of antichains is thus $(\prod_{i=1}^w (\ell_i + 1))^2$. The number of pairs of antichains which differ on one chain is the sum over k of all choices of two different elements in chain C_k and one element from each other chain. This yields the desired formula. \square

5.3 Open Questions

It is NP-hard to compute the linear extension diameter of a general poset, see [7]. For Boolean lattices and for downset lattices of 2-dimensional posets we can now construct the diametral pairs of linear extensions in polynomial time. Is this possible for more general classes?

Open Question 3. Is it possible to compute the linear extension diameter of an arbitrary distributive lattice in polynomial time, or even characterize its diametral pairs of linear extensions?

Since any distributive lattice is the downset lattice of some poset, we would only have to extend our theorem to downset lattices of posets with arbitrary dimension. Unfortunately, our method may fail already when the dimension of the underlying poset is 3.

We considered the 3-dimensional Chevron C . Instead of taking a realizer of C to construct a revlex pair of linear extensions of \mathcal{D}_C , we could only use a diametral pair of C . The resulting revlex pair turned out to not even be locally diametral. That is, there was an incomparable pair of elements of C which appeared in the same order in both of the revlex linear extensions – and in one of them they were even adjacent. So we could swap this pair and obtain two linear extensions of \mathcal{D}_C with larger distance.

Besides looking at larger classes of distributive lattices, another common way to generalize results about Boolean lattices is to look at subposets induced by levels.

Open Question 4. Can we construct the diametral pairs of linear extensions of a poset induced by two levels of B_n ?

Again, our methods do not answer this question. Consider for example the subposet of B_5 induced by the sets of cardinality 2 and 3. The revlex pairs (built by restricting revlex linear extensions of B_5) do not form diametral pairs of this poset.

Another natural question we were interested in asks whether there is a fixed fraction of the incomparable elements of a poset that can always be reversed between two linear extensions. Graham Brightwell [6] recently answered this question in the negative by constructing a family of random posets \mathcal{P} for which $\text{led}(\mathcal{P}) \in o(\text{inc}(\mathcal{P}))$ holds with high probability.

Open Question 5. Which properties of a family \mathcal{F} of posets ensure that there is a $c > 0$ such that $\text{led}(\mathcal{P})/\text{inc}(\mathcal{P}) > c$ for each $\mathcal{P} \in \mathcal{F}$?

Chapter 6

Diametrically Reversing Posets

In the previous chapter we characterized the diametral pairs of linear extensions of special types of posets. Since computing the linear extension diameter is an NP-complete problem, we cannot hope to achieve this in general. In this chapter, we are interested in necessary conditions for linear extensions to be part of a diametral pair. Let us call a linear extension which is part of a diametral pair a *diametral* linear extension.

By Theorem 5.11, we know what the diametral linear extensions of the Boolean lattice look like: They are exactly the revlex linear extensions, induced by the σ -revlex orders, see Definition 5.2. As mentioned in Subsection 5.1.1, the revlex linear extensions have many appealing properties. For example, by Lemma 5.5 each revlex linear extension L of B_n is reversing. That is, there is a critical pair of elements of B_n appearing in non-canonical order in L .

Let us consider the reversing linear extensions in a linear extension graph G , see Figure 6.1. We call a swap color a *critical color* if it corresponds to a critical pair. By Lemma 2.14, the critical colors are exactly the colors $c = ab$ such that there is one component of $G - \theta(c)$, say, G_{ba} , which does not completely contain any other color class. The vertices of G_{ba} are exactly the linear extensions reversing the critical pair (a, b) . That is, if we take a sequence of pairwise parallel color classes, then the critical colors are those at the two ends of such a sequence. Intuitively, it seems plausible that diametral linear extensions should also appear at the end of a sequence of parallel color classes.




Figure 6.1: The critical pairs of \mathcal{P} are (x, b) and (z, a) , so the critical colors of $G(\mathcal{P})$ are xb and za .

In the linear extension graph depicted in Figure 6.1, it looks as if the critical colors cut off the “extremal parts” of the graph. If someone would ask us to find diametral pairs of vertices in this graph, we would probably look for them among the reversing linear extensions. In the case of the depicted example, as in the case of Boolean lattices, this is enough, since all diametral linear extensions are reversing.

These somewhat hand-waving arguments are meant to motivate the following definition:

Definition 6.1. *A poset \mathcal{P} is diametrically reversing if every diametral linear extension of \mathcal{P} is reversing.*

Note that chains are not diametrically reversing, because they contain no critical pair. Intuition may lead us to think that all posets containing at least one incomparable pair are diametrically reversing. In Section 6.1 we will see that this is not true. However, in Section 6.2, we will exhibit several classes of posets which are indeed diametrically reversing, including posets of height 2, interval posets and 3-layer posets. From the last class it follows that almost all posets are diametrically reversing.

Felsner and Reuter [22] conjectured that a property which is a little weaker than being diametrically reversing holds for all posets:

Conjecture 6.2 ([22]). *Let L, L' be a diametral pair of a poset \mathcal{P} . Then at least one of the two linear extensions L, L' is reversing.*

In Section 6.1 we provide a non-trivial counterexample to this conjecture. The trivial fact that chains do not fulfill the conjecture turns out to be crucial for the construction of the counterexample.

6.1 Not all Posets are Diametrically Reversing

In this section we construct a counterexample to Conjecture 6.2. We will first give a nontrivial example of a poset which is not diametrically reversing, because it provides a simpler version of the construction. The idea in both examples is to replace some elements of a Boolean lattice by long chain modules. The chains function like a weight on the elements they replace, and we can therefore use them to manipulate the behavior of the diametral pairs. At the same time, the chains do not add any new critical pairs, as the following lemma shows.

Lemma 6.3. *Let \mathcal{P} be a finite poset, and let \mathcal{P}' arise from it by replacing each element x of \mathcal{P} by the chain $X = x_1 \leq x_2 \leq \dots \leq x_{k(x)}$. Then by mapping a pair (x, y) of elements of \mathcal{P} to the pair $(x_1, y_{k(y)})$ of elements of \mathcal{P}' , we obtain a bijection between the critical pairs of \mathcal{P} and the critical pairs of \mathcal{P}' .*

Proof. Observe that, since all the introduced chains are modules, we have $v < w$ in \mathcal{P} exactly if all elements of V are smaller than all elements of W in \mathcal{P}' , and $v \parallel w$ in \mathcal{P} exactly if all elements of V are incomparable to all elements of W in \mathcal{P}' .

If (x, y) is a critical pair of \mathcal{P} , then $x \parallel y$ in \mathcal{P} and hence $x_1 \parallel y_{k(y)}$ in \mathcal{P}' . We have

$$\text{Pred}_{\mathcal{P}}(x) \subseteq \text{Pred}_{\mathcal{P}}(y) \implies \text{Pred}_{\mathcal{P}'}(x_1) \subseteq \text{Pred}_{\mathcal{P}'}(y_{k(y)})$$

and

$$\text{Succ}_{\mathcal{P}}(y) \subseteq \text{Succ}_{\mathcal{P}}(x) \implies \text{Succ}_{\mathcal{P}'}(y_{k(y)}) \subseteq \text{Succ}_{\mathcal{P}'}(x_1),$$

and thus $(x_1, y_{k(y)})$ is a critical pair of \mathcal{P}' .

Now let (x_i, y_j) be a critical pair of \mathcal{P}' . Then we claim that $i = 1$ and $j = k(y)$ must hold. Since $x_i \parallel y_j$, all pairs of elements from the two chains are incomparable. Therefore each element in X other than x_1 has a predecessor which is not a predecessor of y_j , namely, x_1 . In the same way, each element in Y other than $y_{k(y)}$ has a successor which is not a successor of x_i , namely, $y_{k(y)}$. Hence all critical pairs of \mathcal{P}' have the form $(x_1, y_{k(y)})$. On the other hand, if $(x_1, y_{k(y)})$ is critical in \mathcal{P}' , then clearly (x, y) is critical in \mathcal{P} , and thus the defined mapping is a bijection. \square

The construction of our counterexamples starts with the Boolean lattice. As a reminder, Figure 6.2 shows the Hasse diagram of B_4 . Recall Lemma 1.4, which characterizes the critical pairs of subposets of B_n . We call the 2-, 3- and 4-element sets in B_n *doubles*, *triples* and *quadruples*, respectively.




Figure 6.2: The Boolean lattice B_4 .

Theorem 6.4. *Let B_4^* be the poset resulting from B_4 if the doubles are replaced with chains of length 3. Then B_4^* is not diametrically reversing.*

Proof. Let us first think of B_4^* as the poset obtained from B_4 by replacing the doubles with long chains. Let w be their length; later we will show that $w = 3$ suffices.

Note that the introduced chains all form modules in B_4^* . We are interested in pairs of linear extensions of B_4^* with large distance. By Lemma 4.3, we can always make the chain modules appear successively in a pair of linear extensions without lowering the distance. Therefore we will first look at such linear extensions. We denote the elements in such a linear extension as if they were elements of B_4 , treating the long chains as one element.

In analogy with the hardness proof (Theorem 4.6), the distance between two linear extensions of B_4^* can be thought of as a *weighted distance* between linear extensions of B_4 . Reversing two doubles yields w^2 *unit reversals*. If w is chosen large enough, then reversing as many doubles as possible has priority over all other reversals. Thus we may assume that in a diametral pair L_1, L_2 of B_4^* , the doubles appear in some order in L_1 and in the opposite order in L_2 .

Let L be a reversing linear extension of B_4^* . By Lemma 6.3 and Lemma 1.4 we know that the critical pairs of B_4^* are exactly the pairs (a, a^c) for $a = 1, 2, 3, 4$ and $a^c = [4] \setminus a$. Assume without restriction that $234 < 1$ in L . Then we know that $\{23, 24, 34\} < 234 < 1 < \{12, 13, 14\}$ in L . Let L' be a linear extension of B_4^* which has maximum distance to L . For large w we have $\{12, 13, 14\} < \{23, 24, 34\}$ in L' . Suppose for contradiction that L' reverses the critical pair $(a, [4] \setminus a)$. Then the three doubles containing a must appear in L' after the three doubles contained in $[4] \setminus a$. But the last three doubles in L' do not have an atom in common. It follows that L'

cannot be reversing. We conclude that if a diametral linear extension of B_4^* is reversing, then no diametral partner of it is.

With a closer analysis we can bound w . Here are two linear extensions of B_4 which reverse all pairs of doubles:

$$\begin{aligned} L_1 &= \emptyset 1 2 12 3 13 23 123 4 14 24 124 34 134 234 1234 \\ L_2 &= \emptyset 4 3 34 2 24 1 14 23 234 13 134 12 124 123 1234 \end{aligned}$$

Their weighted distance is $\binom{6}{2}w^2 + 14w + 13$. If we consider two reversing linear extensions, then their quadratic term is smaller, as shown above. For the linear term, we observe that the largest atom in a linear extension of B_4 can be larger than at most three doubles, and the second largest atom larger than at most one double. Analogously, the smallest coatom can be smaller than at most three doubles, the second smallest coatom smaller than at most one. Therefore the linear term is at most $16w$. The constant term is bounded by 14 by counting six reversals among atoms and among coatoms, respectively, and two reversals for the two critical pairs. Now we just have to choose w such that $\binom{6}{2}w^2 + 14w + 13 > (\binom{6}{2} - 1)w^2 + 16w + 14$, and this holds for $w = 3$.

So there is a gap between the distance of two reversing linear extensions of B_4^* and the distance of a diametral pair. This gap remains even if we do not insist that the chain modules appear successively, so we conclude that B_4^* is not diametrically reversing. \square

Note that B_4^* is a graded poset (see Subsection 1.1 for the definition), so not all graded posets are diametrically reversing. Now we refine the ideas of the above construction to disprove Conjecture 6.2.

Theorem 6.5. *There is a poset \mathcal{P}^* such that no diametral linear extension of \mathcal{P}^* is reversing.*

Proof. We consider the subposet \mathcal{P} of the six-dimensional Boolean lattice B_6 induced by all the atoms, the three doubles 12, 34, 56, no triples, the six quadruples 1235, 1246, 1345, 2346, 1356, 2456, and all the coatoms, see Figure 6.3. We replace the doubles and the quadruples in \mathcal{P} by chains of length w , which we will specify later. The resulting poset is our counterexample \mathcal{P}^* .

Again we are interested in pairs of linear extensions of \mathcal{P}^* with large distance, and we first look only at linear extensions in which the chain modules appear successively (cf. Lemma 4.3). We work with \mathcal{P} instead of \mathcal{P}^* , treating the chains of length w as single elements, called *red elements*, which have weight w . We make w so big that gaining a maximum number




Figure 6.3: The counterexample to Conjecture 6.2: the red elements represent long chains.

of red-red reversals, each contributing w^2 unit reversals, is more desirable than any other type of reversal. Let us call the poset induced by the red elements \mathcal{P}_{red} . A maximum number of red-red reversals is achieved by taking two linear extensions of \mathcal{P}^* which restricted to the red elements form a diametral pair of \mathcal{P}_{red} .

The poset \mathcal{P}_{red} is two-dimensional, so we have $\text{led}(\mathcal{P}_{\text{red}}) = \text{inc}(\mathcal{P}_{\text{red}}) = \binom{3}{2} + \binom{6}{2} + 3 \cdot 4 = 30$. Since we can extend a linear extension of \mathcal{P}_{red} to a linear extension of \mathcal{P}^* , we obtain $\text{led}(\mathcal{P}^*) > 30w^2$.

Which distance can a pair of linear extensions achieve in which at least one of them is reversing? Again by Lemmas 6.3 and 1.4 we know that the critical pairs of \mathcal{P}^* are exactly the pairs $(i, [6] - i)$, $i = 1, \dots, 6$. So let (L_1, L_2) be a pair of linear extensions of \mathcal{P}^* , and let L_1 reverse the critical pair $(6, 12345)$, say. Then for the red elements we have $\{12, 34, 1235, 1345\} < 12345 < 6 < \{56, 1246, 2346, 1356, 2456\}$ in L_1 . Assume that there is a diametral linear extension L of \mathcal{P}_{red} which respects these relations. Recall that a diametral pair of a 2-dimensional poset forms a realizer. Thus the partner of L has to fulfill $1246 < 34$ and $2346 < 12$ at the same time, which is impossible. Thus there are no more than 29 red-red reversals between L_1 and L_2 .

Let us bound the number of other reversals between L_1 and L_2 . Each atom of \mathcal{P}^* is incomparable to four red elements, and each coatom to five. If we add all possible reversals of atoms, all possible reversals of coatoms, and finally the reversed critical pair, we see that the distance between L_1 and L_2 is at most $29w^2 + 6 \cdot 4w + 6 \cdot 5w + 6! + 6! + 1 = 29w^2 + 54w + 1441$.

We choose w so big that $30w^2 > 29w^2 + 54w + 1441$, e.g., $w = 100$. Then there is a gap between the distance we can achieve when reversing a critical pair and the distance of a diametral pair. Again, we cannot gain anything by relaxing the condition that the chain modules have to appear

successively, so we conclude that every diametral pair of \mathcal{P}^* consists of two non-reversing linear extensions. \square

Note that our construction disproves the conjecture in a very strong sense: Not only have we shown that not every diametral pair of our example contains a reversing linear extension, but in fact no diametral pair at all does.

6.2 Most Posets are Diametrically Reversing

In this section we will present a number of classes of diametrically reversing posets. Consequently, all of these posets fulfill Conjecture 6.2. As a warm-up, we start with 2-dimensional posets.

Proposition 6.6. *Every poset of dimension 2 is diametrically reversing.*

Proof. If \mathcal{P} is a 2-dimensional poset, then the diametral pairs of linear extensions of \mathcal{P} are exactly the realizers of \mathcal{P} . By Lemma 1.3, a pair L_1, L_2 of linear extensions forms a realizer of \mathcal{P} if and only if every critical pair of \mathcal{P} is reversed in at least one of them. If a non-reversing diametral linear extension L would exist, then its diametral partner L' would have to reverse all critical pairs of \mathcal{P} . But in this case, L' alone would already be a realizer of \mathcal{P} . This is a contradiction to \mathcal{P} being 2-dimensional. \square

6.2.1 Modules

What can we deduce about the diametral linear extensions of a poset \mathcal{P} by looking at the modules of \mathcal{P} ? As an example, consider the case that \mathcal{P} is a series composition of smaller posets \mathcal{P}_i . Then the linear extension graph $G(\mathcal{P})$ is the Cartesian product of the $G(\mathcal{P}_i)$; see Proposition 2.21. Therefore it is easy to see that it suffices that one of the \mathcal{P}_i is diametrically reversing to make the whole poset \mathcal{P} diametrically reversing. It turns out that this is even true for general modules:

Proposition 6.7. *Let \mathcal{P} be a poset containing a module M . If $\mathcal{P}[M]$ is diametrically reversing, then so is \mathcal{P} .*

Proof. Let M be a module of \mathcal{P} , and assume that $\mathcal{P}[M]$ is diametrically reversing. By Lemma 4.3, there is a diametral pair of linear extensions of \mathcal{P} in which the elements of M appear successively. Restricting the two linear extensions of such a pair to M clearly yields a diametral pair of linear extensions of $\mathcal{P}[M]$.

If L_1 and L_2 form a diametral pair of \mathcal{P} in which the elements of M do not appear successively, they still have to contribute the same number of reversals. Hence the restrictions of L_1 and L_2 to M again yield a diametral pair of $\mathcal{P}[M]$. Thus, they both reverse a critical pair of $\mathcal{P}[M]$. Since M is a module, the critical pairs of $\mathcal{P}[M]$ stay critical in \mathcal{P} . We conclude that if every diametral linear extension of $\mathcal{P}[M]$ is reversing, then also every diametral linear extension of \mathcal{P} is reversing. \square

A special case of a module is a twin, that is, a pair xy of elements of \mathcal{P} having the same set of predecessors and the same set of successors (cf. Definition 1.24). Note that the two elements in a twin must be incomparable. Thus, both (x, y) and (y, x) form a critical pair of \mathcal{P} . So any linear extension of \mathcal{P} reverses one of these two critical pairs, and we have the following very useful result:

Corollary 6.8. *Every poset containing a twin is diametrically reversing.*

Proof. This result also follows immediately from Propositions 6.6 and 6.7 by setting $M = \mathcal{A}_2$. \square

6.2.2 Interval Orders

In this section we prove that interval orders are diametrically reversing.

Definition 6.9. *A poset is an interval order if its elements can be represented by intervals on the real line such that $u < v$ if and only if the interval representing u is completely left of the interval representing v . If this can be done only with intervals of length 1, we speak of a unit interval order.*

We may assume that the endpoints in an interval representation of an interval order \mathcal{P} are pairwise different. Then the left endpoints of the intervals yield a linear order of the elements of \mathcal{P} , the *left-endpoint-order*, and the right endpoints yield the *right-endpoint-order*. Clearly, the left-endpoint-order and the right-endpoint-order are linear extensions of \mathcal{P} . It is also easy to see that an interval order is a unit interval order exactly if it has an interval representation in which the left-endpoint-order equals the right-endpoint order.

Note that the critical pairs of an interval order correspond to pairs of intervals which do not contain each other, ordered from left to right.

In the following proposition, we prove that unit interval orders are diametrically reversing (except for chains). We need this result for the proof of the general case. The proof of the special case also contains basic ideas we will use repeatedly in later proofs.

In our proofs, we identify an element of an interval order with the interval representing it, so we may speak, e.g., of an incomparable pair of intervals.

Proposition 6.10. *Every unit interval order which is not a chain is diametrically reversing.*

Proof. Let \mathcal{P} be a unit interval order. We can assume that \mathcal{P} is not an antichain, otherwise we are done by Proposition 6.6. Consider a unit interval representation of \mathcal{P} . Then the left-endpoint-order is the same as the right-endpoint-order. As a consequence, each incomparable pair u, v is a critical pair (u, v) , where u is represented by the interval more to the left. So in fact there is only one non-reversing linear extension: the left-endpoint-order L . We will show that L cannot be diametral.

Cover the plane with a grid of vertical lines a, b, c, \dots of distance 1. We may assume that all intervals are closed on the right end and open on the left end. Then every interval intersects exactly one grid line. We place the grid in such a way that a hits the right endpoint of the leftmost interval. Let us denote the elements intersecting line a with a_1, a_2, \dots, a_k , ordered as in L , the elements intersecting line b with b_1, b_2, \dots , and so on. Note that we have $a_1 < b_1$ in \mathcal{P} , because if they would be incomparable, b_1 had to intersect line a , a contradiction. Also observe that if we have $a_k < b_1$ in \mathcal{P} , then the a_i form a series module of \mathcal{P} , and every pair $a_i a_j$ is a twin of \mathcal{P} . So by Corollary 6.8 we can assume that $a_k \parallel b_1$.

Now assume for contradiction that L forms a diametral pair with L' . In L , element a_k is adjacent to b_1 , hence we must have $b_1 < a_k$ in L' , otherwise we could increase the distance of the two linear extensions by exchanging a_k and b_1 in L .

Similarly, the order of the a_i can be chosen freely in L since they form an antichain of successive elements. Therefore the order of the a_i in L must be the reverse of their order in L' , otherwise we could construct a pair of linear extensions with larger distance.

So we conclude that in L' we must have $b_1 < a_k$ and $a_k < a_1$, which implies $b_1 < a_1$. But this is a contradiction because $a_1 < b_1$ in \mathcal{P} . \square

In the general case, we will use the same idea: We show that if, in a pair of linear extensions, one linear extension is not reversing, then we can always construct two linear extensions with larger distance.

For the proof we want to use a canonical representation of interval orders. We need the following well-known characterization, see e.g. [43].

Theorem 6.11. *A poset \mathcal{P} is an interval order if and only if the maximal antichains of \mathcal{P} can be linearly ordered such that, for each element $v \in \mathcal{P}$, the maximal antichains containing v occur consecutively.*

Proof. The “only if” direction is immediately clear when considering an interval representation. For the other direction, we take a linear order of the antichains as in the characterization and define $\ell(v)$ and $r(v)$ as the indices of the first and last antichains containing v . Then $(\ell(v), r(v))_{v \in \mathcal{P}}$ defines a representation of \mathcal{P} by open intervals with integer endpoints. \square

The representation induced by the linearly ordered antichains mentioned in the proof above will be our *canonical interval representation*. Figure 6.4 shows an example where we have drawn vertical lines to mark the integers which constitute left or right endpoints of intervals. An observation important for the proof is that no intervals start or end between the lines, but at every line except for the last one an interval starts, and at every line except for the first one an interval ends.




Figure 6.4: Example of an interval order in canonical representation.

Now we have enough tools at hand to go into the proof.

Theorem 6.12. *Every interval order which is not a chain is diametrically reversing.*

Proof. Let \mathcal{P} be an interval order given in canonical representation. Let L be a diametral linear extension of \mathcal{P} with diametral partner L' . Suppose for contradiction that L is non-reversing, that is, every critical pair (x, y) of \mathcal{P} appears in the canonical order $x < y$ in L . A non-reversing linear extension of the example in Figure 6.4 is given by $L = xewabcdyz$; the names of the elements are chosen as to illustrate the proof that follows. We will construct a pair of linear extensions with larger distance, contradicting the fact that L, L' is a diametral pair.

If there are no two intervals in \mathcal{P} such that one contains the other, then the left-endpoint-order equals the right-endpoint-order. Thus \mathcal{P} is a unit interval order and we are done by the previous proposition. So we may

assume that there are incomparable pairs of intervals containing each other; these are exactly the pairs which are not critical.

Let c be the first element in L which appears after an element a with $r(a) \geq r(c)$. (If there is no such c , then we apply the proof to L read backwards, exchanging left and right endpoints.) Note that since L is non-reversing, it follows that $\ell(a) < \ell(c)$. Let b be the element appearing immediately before c in L . Our plan is to move b down in L and up in L' , creating a pair of linear extensions with larger distance.

We want to place b *low* in L , that is, in a position such that all the elements before b in L are predecessors of b . To show that this is possible we claim that the predecessors of b form an initial segment of L . Indeed, if we consider an element d that is incomparable to b , then clearly it has a larger right endpoint than any predecessor e of b . Hence if d would appear before e in L , we would have chosen e as c , a contradiction. Therefore we can safely move b down in L and place it right after its first predecessor, obtaining a linear extension \bar{L} . Let S be the set of elements that b passes, that is, that are smaller than b in L and larger than b in \bar{L} .

Now we move b up in L' , that is, we move it just before its lowest successor in L' . We call the resulting linear extension \bar{L}' . First we show that the distance between \bar{L} and \bar{L}' is not smaller than the distance between L and L' . Note that b is low in \bar{L} and thus any element that b passes when moving up in L' only increases the distance between the two linear extensions. If we can ensure that in \bar{L}' , the element b is larger than all elements of S , then we have shown that the distance between the linear extensions did not decrease. But this can be done: By the choice of c , all elements which are smaller than b in L , in particular the elements of S , have a smaller right endpoint than b . Therefore every successor of b is also a successor of all elements in S . Hence if we move b up in L' and place it just before its lowest successor, we will have passed all elements of S .

It remains to show that the distance between \bar{L} and \bar{L}' is actually larger than the distance between L and L' . We will identify a pair of elements that is a reversal between the first pair of linear extensions, but not between the latter. First let us concentrate on the pair b, c . By the choice of c , we must have $r(b) \geq r(a) \geq r(c)$. Since $b < c$ in L , it follows that b and c are incomparable. If $\ell(b) \geq \ell(c)$, then (c, b) is a critical pair which is reversed in L , a contradiction. Thus we have $\ell(b) < \ell(c)$.

Now from the properties of the canonical interval representation it follows that there exists an element w which is incomparable to b but smaller than c . What is the order of b, c and w in L and L' ? In L we know that b and c appear adjacently, so w has to appear before both of them. In L' the order of b and c must be reversed, otherwise we could construct a pair of

linear extensions with larger distance immediately by changing their order in L . Hence in L' we have $w < c < b$, and thus the pair w, b is not reversed between L and L' . But b is low in \bar{L} , and therefore we have $b < w$ in \bar{L} . Therefore w and b are reversed between \bar{L} and \bar{L}' , and we have found the desired pair. This shows that the non-reversing linear extension L cannot be diametral. \square

6.2.3 3-Layer Posets

In this subsection we will prove that a class of posets covering the vast majority of all posets is diametrically reversing.

Definition 6.13. *A 3-layer poset is a poset in which each maximal chain has length 3, and in addition, each minimum is smaller than each maximum (cf. Figure 6.5). We call a 3-layer poset complete if each incomparable pair consists of two elements from the same layer.*




Figure 6.5: Scheme of a 3-layer poset. The sets A, B and C form antichains. All elements of A are smaller than all elements of C .

Theorem 6.14. *Every 3-layer poset which is not a chain is diametrically reversing.*

Proof. Let \mathcal{P} be a 3-layer poset consisting of a layer A of minima, a middle layer B , and a layer C of maxima. First we consider the easy case where there is only one element b in the middle layer. Since \mathcal{P} is graded, every minimum has a successor in the middle layer. In this case all minima are smaller than b . But this means in fact that every pair of minima is a twin, and by Corollary 6.8 we are done. Hence we can assume $|B| \geq 2$.

Observe that, if we have an incomparable pair (a, b) with $a \in A$ and $b \in B$, then it forms a critical pair automatically: $\text{Pred}(a) = \emptyset \subseteq \text{Pred}(b)$, and $\text{Succ}(b) \subseteq C \subseteq \text{Succ}(a)$. An analogous argument shows that every incomparable pair (b, c) with $b \in B$ and $c \in C$ is critical. Let us call a

linear extension *mixing* if it contains an element of B appearing before an element of A or an element of C appearing before an element of B . Thus we have shown that any mixing linear extension is reversing.

Now let us consider a non-mixing diametral linear extension L with diametral partner L' . We will first analyse L' . If L' is also a non-mixing linear extension, then \mathcal{P} must be a complete 3-layer poset, since in all other cases we can find two linear extensions with larger distance. But in a complete 3-layer poset, any pair of minima forms a twin, and so by Corollary 6.8 we are done. Thus we can assume that L' is a mixing linear extension.

The linear extension L' starts with an initial segment consisting of only minima, followed by a part where elements of the three layers are mixed, and it finishes with a final segment consisting of only maxima. We call the set of elements forming the initial segment $A' \subseteq A$, and the set forming the final segment $C' \subseteq C$. Note that since \mathcal{P} is graded, the partition of the elements into three layers is unique. In particular, A' and C' cannot be empty.

We label the elements of the three layers as $a_1, a_2, \dots, a_k, b_1, b_2, \dots, b_\ell$ and c_1, c_2, \dots, c_m in the order in which they appear in L' . Now we return to L . Since L is non-mixing, the elements of every level appear successively. So we are free to choose the layer orders in L and therefore choose them in a way which differs most from L' . Hence the elements of each layer appear in L exactly in the opposite order of their order in L' . Thus L has the form

$$L = a_k a_{k-1} \dots a_2 a_1 b_\ell b_{\ell-1} \dots b_2 b_1 c_m c_{m-1} \dots c_2 c_1.$$

Assuming that $A' = \{a_1, a_2, \dots, a_{k'}\}$ and $C' = \{c_{m'}, c_{m'+1}, \dots, c_m\}$ we know that L' looks like this:

$$L' = a_1 a_2 \dots a_{k'} b_1 X b_\ell c_{m'} \dots c_{m-1} c_m.$$

Here, X denotes the mixed part of L' consisting of elements of potentially all three layers.

We claim that (b_1, b_ℓ) is a critical pair. Since they are both elements of the middle layer, they are incomparable. The predecessors of b_1 are contained in A' , since these are the only elements smaller than b_1 in L' . We want to show that $A' \subseteq \text{Pred}(b_\ell)$.

In L , the set A' is found at the end of the sequence of minima, immediately before the element b_ℓ . Hence if a_1 and b_ℓ were incomparable, they could be exchanged in L to yield a pair of linear extensions with larger distance, a contradiction. So we have $a_1 < b_\ell$. We can extend this argument to show that any $a_i \in A'$ needs to be smaller than b_ℓ . This holds because

we can choose any order of the initial segment A' in L' . So we are free to turn any $a_i \in A'$ into a_1 by reordering the elements of A' in L' and L . Thus with the above argument, $a_i < b_\ell$ for any $a_i \in A'$. Hence we have $\text{Pred}(b_1) \subseteq A' \subseteq \text{Pred}(b_\ell)$.

It remains to show that $\text{Succ}(b_\ell) \subseteq \text{Succ}(b_1)$. The argument works analogously: From the position of b_ℓ in L' we deduce that $\text{Succ}(b_\ell) \subseteq C'$. In L , the element b_1 appears immediately before the elements of C' . If some $c_i \in C'$ were incomparable with b_1 , then we could construct a pair of linear extensions with larger distance by rearranging C' in L and L' and exchanging b_1 with c_i . This would be a contradiction to L, L' being a diametral pair, and thus $\text{Succ}(b_\ell) \subseteq \text{Succ}(b_1)$.

We have shown that (b_1, b_ℓ) is a critical pair of \mathcal{P} . In L , we have $b_\ell < b_1$, so we found a critical pair that is reversed in L . This shows that if a non-mixing linear extension of \mathcal{P} is diametral, then it is reversing. \square

Corollary 6.15. *Every poset of height 2 which is not a chain is diametrically reversing.*

Proof. A poset of height 2 can be thought of as 3-layer posets with $C = \emptyset$. We can use the proof of Theorem 6.14 with only very few changes.

We consider L and L' defined as before. The argument is the same until we define A' and C' . Here it holds that $C' = \emptyset$, since $C = \emptyset$. We then have

$$\begin{aligned} L &= a_k a_{k-1} \dots a_2 a_1 b_\ell b_{l-1} \dots b_2 b_1 \quad \text{and} \\ L' &= a_1 a_2 \dots a_{k'} b_1 X b_\ell, \end{aligned}$$

where the mixed segment X of L' consists of elements from A and B .

Now it can be shown exactly in the same way as in the previous proof that $\text{Pred}(b_1) \subseteq A' \subseteq \text{Pred}(b_\ell)$. Furthermore, $\text{Succ}(b_\ell) = \emptyset = \text{Succ}(b_1)$. Thus (b_1, b_ℓ) is a critical pair of \mathcal{P} which is reversed in L . \square

Kleitman and Rothschild [36] showed that almost all posets are 3-layer posets. More precisely, among the posets on n elements, the proportion of the number of 3-layer posets to the number of all posets tends to 1 as n tends to infinity. Since there is only one chain for each n , we can neglect these, and so with Theorem 6.14 we immediately obtain the following result:

Corollary 6.16. *Almost all posets are diametrically reversing.* \square

Of course, there are still open questions about which classes of posets are diametrically reversing. We close with the following:

Open Question 6. Are all posets of height 3 diametrically reversing?

Bibliography

- [1] I. ANDERSON, *Combinatorics of Finite Sets*, Oxford Clarendon Press, 1987. (82)
- [2] K. BAKER, P. FISHBURN, AND F. ROBERTS, *Partial orders of dimension 2*, *Networks*, 2 (1972), pp. 11–28. (11, 12, 66)
- [3] G. BIRKHOFF, *On the combination of subalgebras*, *Proc. Cambridge Philos. Soc.*, 29 (1933), pp. 441–464. (8)
- [4] A. BJÖRNER, M. LAS VERGNAS, B. STURMFELS, N. WHITE, AND G. M. ZIEGLER, *Oriented Matroids*, vol. 46 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, second ed., 1999. (27)
- [5] A. BJÖRNER AND M. L. WACHS, *Permutation statistics and linear extensions of posets*, *J. Combinatorial Theory Ser. A*, 58 (1991), pp. 85–114. (28)
- [6] G. BRIGHTWELL, 2009. Personal communication. (100)
- [7] G. BRIGHTWELL AND M. MASSOW, *Diametral pairs of linear extensions*. Submitted; available at <http://arxiv.org/abs/0809.1828>. (2, 3, 100)
- [8] G. BRIGHTWELL AND P. WINKLER, *Counting linear extensions*, *Order*, 8 (1991), pp. 225–242. (7)
- [9] V. D. CHEPOI, *Isometric subgraphs of hamming graphs and d-convexity*, *J. Cybernetics and Systems Analysis*, 24 (1988), pp. 6–11. (13)

- [10] D. COPPERSMITH AND S. WINOGRAD, *Matrix multiplication via arithmetic progressions*, J. Symb. Comput., 9 (1990), pp. 251–280. (45)
- [11] B. DAVEY AND H. PRIESTLEY, *Introduction to Lattices and Order*, Cambridge University Press, second ed., 2002. (8)
- [12] M. M. DEZA AND M. LAURENT, *Geometry of Cuts and Metrics*, Algorithms and Combinatorics, Springer, 1997 ed., 1997. (13)
- [13] R. DIESTEL, *Graph Theory*, Springer, third ed., 2005. (5)
- [14] R. P. DILWORTH, *A decomposition theorem for partially ordered sets*, Annals of Math., 51 (1950), pp. 161–166. (73)
- [15] D. Ž. DJOKOVIĆ, *Distance preserving subgraphs of hypercubes*, J. Combinatorial Theory Ser. B, 14 (1973), pp. 263–267. (13, 14, 15)
- [16] B. DREESEN, W. POGUNTKE, AND P. WINKLER, *Comparability invariance of the fixed point property*, Order, 2 (1985), pp. 269–274. (21)
- [17] B. DUSHNIK AND E. W. MILLER, *Partially ordered sets*, Amer. J. Math., 63 (1941), pp. 600–610. (9, 11)
- [18] D. EPPSTEIN, *Recognizing partial cubes in quadratic time*, in Proceedings of SODA '08, New York, 2008, ACM, pp. 1258–1266. (15)
- [19] J. FELDMAN, *Ordres et permutoèdres*, Mathématiques et sciences humaines, 28 (1969), pp. 27–38. (28)
- [20] S. FELSNER AND M. MASSOW, *Linear extension diameter of downset lattices of 2-dimensional posets*. Submitted; available at <http://arxiv.org/abs/0907.1793>. (3)
- [21] S. FELSNER, V. RAGHAVAN, AND J. SPINRAD, *Recognition algorithms for orders of small width and graphs of small Dilworth number*, Order, 20 (2003), pp. 351–364. (73)
- [22] S. FELSNER AND K. REUTER, *The linear extension diameter of a poset*, SIAM J. Discrete Math., 12 (1999), pp. 360–373. (3, 24, 63, 64, 65, 66, 67, 78, 79, 80, 81, 102)
- [23] V. V. FIRSOV, *Isometric embedding of a graph in a Boolean cube*, Cybernetics and Systems Analysis, 1 (1965), pp. 112–113. (13)

-
- [24] S. FOMIN AND N. READING, *Root systems and generalized associahedra*, in Geometric Combinatorics, vol. 13 of IAS / Park City Math. Ser., Amer. Math. Soc., 2007, pp. 63–131. (27)
- [25] T. GALLAI, *Transitiv orientierbare Graphen*, Acta Math. Acad. Sci. Hungar., 18 (1967), pp. 25–66. (12, 17, 20)
- [26] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, 1979. (68)
- [27] M. C. GOLUMBIC, *The complexity of comparability graph recognition and coloring*, Computing, 18 (1977), pp. 199–208. (59)
- [28] ———, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, 1980. (20)
- [29] R. L. GRAHAM AND H. O. POLLAK, *On the addressing problem for loop switching*, Bell System Tech. J., 50 (1971), pp. 2495–2519. (13)
- [30] ———, *On embedding graphs in squashed cube*, in Graph Theory and Applications, vol. 303 of Lecture Notes in Mathematics, Springer, 1972, pp. 99–110. (13)
- [31] R. L. GRAHAM AND P. M. WINKLER, *On isometric embeddings of graphs*, Trans. Amer. Math. Soc., 288 (1985), pp. 527–539. (13)
- [32] S. HOŞTEN AND W. D. MORRIS, JR., *The order dimension of the complete graph*, Discrete Math., 201 (1998), pp. 133–139. (89)
- [33] W. IMRICH AND S. KLAVŽAR, *Product Graphs*, Wiley, 2000. (13, 14)
- [34] W. IMRICH AND I. PETERIN, *Recognizing Cartesian products in linear time*, Discrete Math., 307 (2007), pp. 472–483. (58)
- [35] D. J. KLEITMAN, *Families of non-disjoint subsets*, J. Combinatorial Theory, 1 (1966), pp. 153–155. (82)
- [36] D. J. KLEITMAN AND B. L. ROTHSCHILD, *Asymptotic enumeration of partial orders on a finite set*, Trans. Amer. Math. Soc., 205 (1975), pp. 205–220. (114)
- [37] H. KOMM, *On the dimension of partially ordered sets*, Amer. J. of Math., 70 (1948), pp. 507–520. (10)
- [38] P. LEHOT, *An optimal algorithm to detect a line graph and output its root graph*, J. Assoc. Comp. Mach., 21 (1974), pp. 569–575. (50)

- [39] F. MAFFRAY AND M. PREISSMANN, *A translation of Gallai's paper: 'Transitiv orientierbare Graphen'*, in Perfect Graphs, J. L. R. Alfonsín and B. A. Reed, eds., Wiley, 2001, pp. 25–66. (17)
- [40] R. M. MCCONNELL AND J. P. SPINRAD, *Modular decomposition and transitive orientation*, Discrete Math., 201 (1999), pp. 189–241. (20, 51, 55, 56, 95)
- [41] —, *Ordered vertex partitioning*, Discrete Math. and Theor. Comp. Sci., 4 (2000), pp. 45–60. (20)
- [42] R. H. MÖHRING, *Algorithmic aspects of comparability graphs and interval graphs*, in Graphs and Order, I. Rival, ed., vol. 147 of NATO ASI Series C, D. Reidel Pub. Comp., 1985, pp. 41–101. (19, 20)
- [43] —, *Computationally tractable classes of ordered sets*, in Algorithms and Order, I. Rival, ed., vol. 255 of NATO ASI Series C, Kluwer Academic Publishers, 1989, pp. 105–193. (11, 95, 109)
- [44] M. NAATZ, *The graph of linear extensions revisited*, SIAM J. Discrete Math., 13 (2000), pp. 354–369. (24, 25, 27, 29, 40)
- [45] —, *Acyclic Orientations of Mixed Graphs*, PhD thesis, Technische Universität Berlin, 2001. (24, 28)
- [46] S. OVCHINNIKOV, *Media theory: Representations and examples*, Discrete Appl. Math., 156 (2008), pp. 1197–1219. (27)
- [47] —, *Partial cubes: structures, characterizations, and constructions*, Discrete Math., 308 (2008), pp. 5597–5621. (13)
- [48] G. PRUESSE AND F. RUSKEY, *Generating the linear extensions of certain posets by transpositions*, SIAM J. Discrete Math., 4 (1991), pp. 413–422. (24)
- [49] I. RABINOVITCH AND I. RIVAL, *The rank of a distributive lattice*, Discrete Math., 25 (1979). (10)
- [50] K. REUTER, *The comparability graph and the graph of linear extensions of a poset*, Hamburger Beiträge zur Mathematik, Heft 57 (1996). (24, 37, 40, 41, 49)
- [51] —, *Linear extensions of a poset as abstract convex sets*, Hamburger Beiträge zur Mathematik, Heft 56 (1996). (24, 26, 27, 28, 34)

-
- [52] F. RUSKEY AND C. D. SAVAGE, *Hamilton cycles that extend transposition matchings in Cayley graphs of S_n* , SIAM J. Discrete Math., 6 (1993), pp. 152–166. (24)
- [53] G. SABIDUSSI, *Graph multiplication*, Math. Z., 72 (1960), pp. 446–457. (57)
- [54] R. SEDGEWICK, *Algorithms in C++*, Addison-Wesley Professional, third ed., 1998. (24)
- [55] R. P. STANLEY, *Two poset polytopes*, Discrete Comp. Geom., 1 (1986), pp. 9–23. (28)
- [56] R. P. STANLEY, *An introduction to hyperplane arrangements*, in Geometric Combinatorics, vol. 13 of IAS/Park City Math. Ser., Amer. Math. Soc., 2007, pp. 389–496. (27)
- [57] G. STEINER, *Single machine scheduling with precedence constraints of dimension 2*, Math. Oper. Res., 9 (1984), pp. 248–259. (95)
- [58] W. T. TROTTER, *Dimension of the crown S_n^k* , Discrete Math., 8 (1974), pp. 85–103. (66)
- [59] —, *Combinatorics and Partially Ordered Sets*, The Johns Hopkins University Press, 1992. (9, 12, 73)
- [60] V. G. VIZING, *The Cartesian product of graphs*, Vyc. Sis., 9 (1963), pp. 30–43. (57)
- [61] D. B. WEST, *Generating linear extensions by adjacent transpositions*, J. Combinatorial Theory Ser. B, 58 (1993), pp. 58–64. (24)
- [62] —, *Introduction to Graph Theory (2nd Edition)*, Prentice Hall, 2001. (48, 49)
- [63] H. WHITNEY, *Congruent graphs and the connectivity of graphs*, vol. 54, 1932, pp. 150–168. (49)
- [64] P. M. WINKLER, *Isometric embedding in products of complete graphs*, Discrete Appl. Math., 7 (1984), pp. 221–225. (13, 14, 15)
- [65] M. YANNAKAKIS, *The complexity of the partial order dimension problem*, SIAM J. Algebraic Discrete Methods, 3 (1982), pp. 351–358. (11)
- [66] G. M. ZIEGLER, *Lectures on Polytopes*, Springer, 1995. (27)

Index

- 3-layer poset, 112
- σ -larger, 78
- σ -maximum, 78
- σ -minimum, 78
- σ -revlex order, 78, 89
- σ -smaller, 78
- 2-dimensional posets, 11

- adjacent swap colors, 35
- admissible questions, 43
- antichain, 6
- atom, 8

- Boolean lattice, 8
- braid arrangement, 27

- canonical bijection, 7
- canonical interval representation, 110
- canonical order of a critical pair, 9
- canonical partition, 18
- Cartesian product, 40
- Cartesian prime, 41
- chain, 6
- Chevron, 6
- coatom, 8
- comparability graph, 6, 17
- comparability invariant, 6
- comparable, 6
- complementary linear extension, 67
- complementary pair, 67
- complete 3-layer poset, 112
- completely adjacent, 17
- component of a poset, 6
- conjugate, 11
- connected poset, 6
- convex cycle, 16
- convex subgraph, 13
- cover relation, 6
- Coxeter arrangement, 27
- critical color, 101
- critical pair, 9
- crossing color classes, 31
- crown, 66

- diametral linear extension, 101
- diametral pair of linear extensions, 63
- diametral partner, 63
- diametrically reversing, 102
- dimension of a poset, 9
- direction, 7
- distance btw linear extensions, 8

- distributive lattice, 8
- Djoković-Winkler classes, 14
- Djoković-Winkler relation, 14
- dominance relation, 9
- doubles, 103
- downset, 7
- downset lattice, 8
- dual poset, 7

- essentially unique, 40

- forcing classes, 17
- forcing relation, 17, 32

- G -compatible poset, 40
- G -compatible swap coloring, 49
- G -medium, 43
- generalized crowns, 66
- generic algorithm, 7
- global element, 7
- global maximum, 7
- global minimum, 7
- global vertex, 7
- graded, 8
- ground set, 5

- halfspace, 28
- Hamming distance, 13
- Hamming labeling, 13
- Hasse diagram, 6
- height, 6
- hypercube, 13
- hyperplane arrangement, 27

- incomparability graph, 6
- incomparable, 6
- inner edges, 18
- inserting a module, 21
- interlock, 52
- interval, 6
- interval order, 108
- irreducible, 12

- isometric cycle, 16
- isometric dimension, 15
- isometric subgraph, 13

- join, 8
- join-irreducible, 8
- jump, 60

- Kleitman's Lemma, 81
- Krausz decomposition, 48

- lattice, 8
- left-endpoint-order, 108
- level, 8
- lie between, 31
- linear extension, 7
- linear extension diameter, 63
- linear extension graph, 24
- linear order, 6

- meet, 8
- modular decomposition, 19
- modular decomposition tree, 19
- module of a graph, 18
 - quasi-maximal, 19
 - strong, 19
- module of a poset, 21
 - parallel module, 21
 - prime module, 21
 - series module, 21

- non-separating, 11

- order polytope, 28
- outer edges, 18

- parallel color classes, 31
- parallel composition, 21
- partial cube, 13
- permutahedron, 27
- pivot, 53
- pivot element, 60
- predecessor, 7

prime composition, 21
priority list, 7
private pivot, 56
proper embedding, 9

quadruples, 103

realizer, 9
reversal, 8
reversing, 9
revlex linear extension, 79, 89
revlex order, 78
revlex pair, 79, 89
right-endpoint-order, 108

separating, 11
series alteration, 41
series composition, 21
slot, 84, 92
standard example, 8
stripes, 34
subposet, 5
successor, 7
swap, 24
swap color, 24
swap coloring, 26
swap partition, 26

touching color classes, 31
triples, 103
twin, 21

uniquely partially orderable, 20
unit interval order, 108
unit reversal, 69, 104
upset, 7

vertex classes, 18

weighted distance, 69, 104
width, 6

