Linear Extension Diameter of Downset Lattices of 2-Dimensional Posets

Stefan Felsner and Mareike Massow
Technische Universität Berlin
{felsner,massow}@math.tu-berlin.de

Abstract

The linear extension diameter of a finite poset $\mathcal{P}$ is the maximum distance between a pair of linear extensions of $\mathcal{P}$, where the distance between two linear extensions is the number of pairs of elements of $\mathcal{P}$ appearing in different orders in the two linear extensions. We prove a formula for the linear extension diameter of the Boolean Lattice and characterize the diametral pairs of linear extensions. For the more general case of a downset lattice $\mathcal{D}_{\mathcal{P}}$ of a 2-dimensional poset $\mathcal{P}$, we characterize the diametral pairs of linear extensions of $\mathcal{D}_{\mathcal{P}}$ and show how to compute the linear extension diameter of $\mathcal{D}_{\mathcal{P}}$ in time polynomial in $|\mathcal{P}|$.

1 Introduction

With a finite poset $\mathcal{P}$ consider its *linear extension graph* $G(\mathcal{P})$, which has the linear extensions of $\mathcal{P}$ as vertices, with two of them being adjacent exactly if they differ only in a single adjacent transposition. The linear extension graph was originally defined in [12]. It is implicitly used in investigations around finite posets and sorting problems, see e.g. [4]. Explicit research has been initiated in [13], see also [11].

The *linear extension diameter* of $\mathcal{P}$, denoted by $\text{led}(\mathcal{P})$, is the diameter of $G(\mathcal{P})$, see [6]. It equals the maximum number of pairs of elements of $\mathcal{P}$ that can be in different orders in two linear extensions of $\mathcal{P}$. A *diametral pair of linear extensions* of $\mathcal{P}$ is a diametral pair of $G(\mathcal{P})$.

A realizer of $\mathcal{P}$ is a set $\mathcal{R}$ of linear extensions of $\mathcal{P}$ such that the comparabilities of $\mathcal{P}$ are exactly the intersection of the comparabilities of the linear extensions in $\mathcal{R}$ (cf. [15]). The dimension of a poset $\mathcal{P}$ is the minimum size of a realizer. If $\mathcal{P}$ is 2-dimensional, i.e., if it has a realizer $\mathcal{R} = \{L_1, L_2\}$, then every incomparable pair $x||y$ of elements of $\mathcal{P}$ appears in different orders in L_1 and L_2. It follows that $\mathcal{P}$ is 2-dimensional exactly if $\text{led}(\mathcal{P})$ equals the number of incomparable pairs of $\mathcal{P}$. Figure 1 shows a six-element poset $\mathcal{P}$ (called the *chevron*) with its linear extension graph. Note that $\mathcal{P}$ has seven incomparable pairs, but the diameter of its linear extension graph is only six. Hence, the dimension of $\mathcal{P}$ must be at least three.

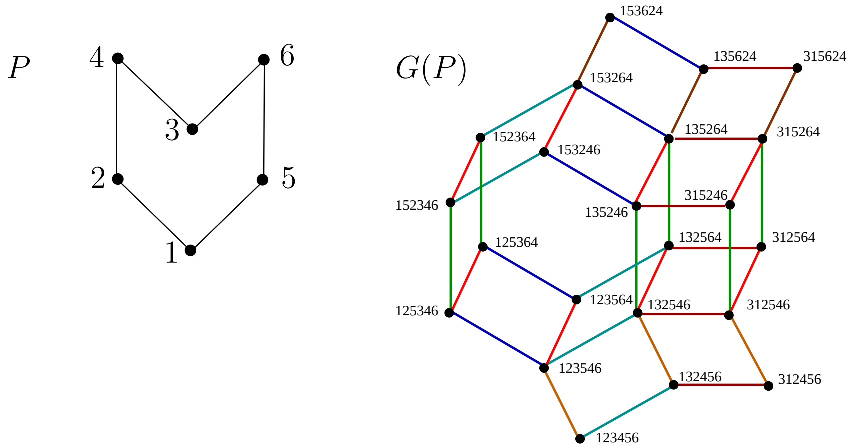

Figure 1: The chevron and its linear extension graph.

A diametral pair L_1, L_2 can be used to obtain an optimal drawing of $\mathcal{P}$: Use L_1 and L_2 on the two coordinate axes to get a position in the plane for each element of $\mathcal{P}$. Since the number of incomparable pairs which appear in different orders in L_1 and L_2 is maximized, the resulting drawing has a minimal number of pairs x, y which are comparable in the dominance order, i.e., $x_i < y_i$ for all coordinates i, but incomparable in $\mathcal{P}$. See Figure 2 for an example.

In [3] it was shown that, given a poset $\mathcal{P}$, it is NP-complete in general to determine the linear extension diameter of $\mathcal{P}$. In Section 2 we prove a formula for the linear extension diameter of the Boolean lattice, and characterize the diametral pairs of linear extensions of the Boolean lattice. In Section 3 we characterize the diametral pairs of linear extensions for the more general class of downset lattices of 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 2 are contained in the results of Section 3. Section 2 provides an easier acess to our proof techniques. However, since Section 3 is largely self-contained, the self-confident reader may proceed there right away.

2 Boolean Lattices

Let B_n denote the n-dimensional Boolean lattice, that is, the poset on all subsets of $[n]$, ordered by inclusion. In this section we prove the following conjecture from [6].

Conjecture 1 (Felsner, Reuter '99)**.** *$led(B_n) = 2^{2n-2} - (n+1) \cdot 2^{n-2}$.*

We will also characterize the diametral pairs of linear extensions of B_n. Central to our investigations is a generalization of the reverse lexicographic order. Before defining it, 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 $\sigma^{-1}(i) < \sigma^{-1}(j)$. We define σ-larger analogously. The σ-maximum (σ-minimum) of a finite set S is the element which is σ-largest (σ-smallest) in S. For example, if $\sigma = 2413$, then $4 <_\sigma 1$, and $\max_\sigma\{3,4\} = 3$.

Definition 1. *Define a relation on the subsets of $[n]$ by setting*

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

*for a pair $S, T \subseteq [n]$. We call this relation the σ-*revlex order*.*

In the lemma below we prove that the σ-revlex order defines a linear extension of B_n. We denote this linear extension by L_σ. By $\overline{\sigma}$ we denote the reverse of a permutation σ. In Theorem 9 we prove that the pairs $L_\sigma, L_{\overline{\sigma}}$ are exactly the diametral pairs of linear extensions of B_n. Figure 2 shows the drawing of B_4 resulting from $L_{\text{id}}, L_{\overline{\text{id}}}$.

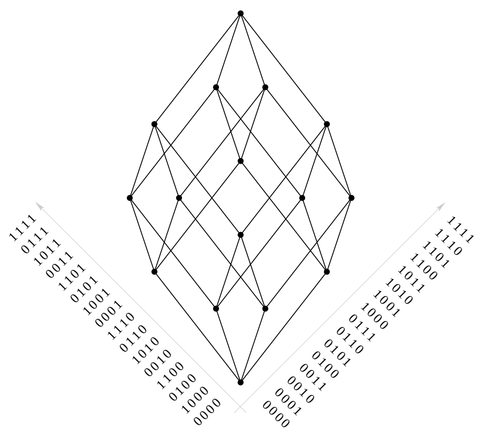

Figure 2: The drawing of B_4 based on the diametral pair $L_{\text{id}}, L_{\overline{\text{id}}}$.

Lemma 2. *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 $S \triangle T \subseteq T$, thus $S <_\sigma T$.

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 $\max_\sigma(A \triangle B) = b \in B$, $\max_\sigma(B \triangle C) = c \in C$ and $\max_\sigma(C \triangle A) = a \in A$.

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 $a = \min_\sigma\{a, b, c\}$. Then by definition of a we know that every element in C which is σ-larger than a is also in A. Hence $c \in A$. Since $c \notin B$, it follows that $c \in A \triangle B$, and thus $\max_\sigma(A \triangle B) = b >_\sigma c$. Now by definition of c, every element in B which is σ-larger than c is also in C. Thus $b \in C$. But we also know that $b \notin A$, and hence $b \in C \triangle A$. But this is a contradiction since $b >_\sigma a = \max_\sigma(C \triangle A)$. □

Definition 3. *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 $S \triangle T = D$ 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 with a subset $X \subseteq D$ 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 $T = X^c \cup I$. We obtain the following useful lemma:

Lemma 4. *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 $|D| = d$.*

Recall that the distance between two linear extensions of a poset $\mathcal{P}$ is the number of pairs of elements of $\mathcal{P}$ appearing in different orders in the two linear extensions. We call such a pair a *reversal.* Note that reversals are unordered pairs, whereas 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 1, was proved inductively in [6]. Here, we give a more combinatorial proof.

Proposition 5. *Given a permutation σ of $[n]$, the distance between L_σ and $L_{\overline{\sigma}}$ as linear extensions of B_n is*

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

Proof. We need to count the number of unordered pairs of subsets of $[n]$ appearing in different orders in L_σ and $L_{\overline{\sigma}}$. We claim that an equivalence class $\mathcal{C}_{D,I}$ contributes exactly 2^{d-2} reversals between L_σ and $L_{\overline{\sigma}}$ if $d \geq 2$, and none if $d < 2$.

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

$$S <_{\overline{\sigma}} T \iff \min_\sigma(S \triangle T) \in T.$$

Let $\mathcal{C}_{D,I}$ be an equvialence class as in Definition 3. 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 let us 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 $L_{\overline{\sigma}}$ if and only if exactly one of the elements $\min_\sigma D$ and $\max_\sigma D$ is contained in X. Since we want to count (S, T) and (T, S) only once, 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 $L_{\overline{\sigma}}$ as claimed.

How many reversals does this yield in total? Each set $D \subseteq [n]$ forms a class together with each set $I \subseteq [n] \setminus D$. Since there are 2^{n-d} choices for I, each D with $d \geq 2$ accounts for $2^{n-d} \cdot 2^{d-2} = 2^{n-2}$ reversals. There are $2^n - (n+1)$ possibilities to choose a set $D \subseteq [n]$ with $d \geq 2$. Thus the distance between L_σ and $L_{\overline{\sigma}}$ is

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

□

For proving the other half of the conjecture, we will need Kleitman's Lemma (see [8] or [1]):

Lemma 6 (Kleitman's Lemma)**.** *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 its subsets are also in $\mathcal{F}_i$. Then the following formula holds:*

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

Theorem 7. *led*$(B_n) = 2^{2n-2} - (n+1) \cdot 2^{n-2}$.

Proof. Proposition 5 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 3. 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 $|\mathcal{F}_1 \triangle \mathcal{F}_2|$.

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 $Y \in \mathcal{F}_i$ 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. 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

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

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 same calculation as in the proof of Proposition 5 to show that the total number of reversals is at most $2^{2n-2} - (n+1) \cdot 2^{n-2}$. □

In the above two proofs we have shown the following fact, which we state explicitely for later reference:

Fact 8. *If $L, \overline{L}$ 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 $\overline{L}$.*

We have shown that for every permutation σ of $[n]$, the linear extensions L_σ and $L_{\overline{\sigma}}$ form 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 9. *The diametral pairs of the Boolean lattice are unique up to isomorphism. More precisely, if $L, \overline{L}$ 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 $\overline{L} = L_{\overline{\sigma}}$.*

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

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 $\overline{L}$. Hence the permutation defining the order of the atoms in $\overline{L}$ is $\overline{\sigma}$.

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, $L^{k-1} = L^{k-1}_\sigma$, and that all pairs of sets in $\overline{L}^{k-1}$ are in $\overline{\sigma}$-revlex order, that is, $\overline{L}^{k-1} = L^{k-1}_{\overline{\sigma}}$. For the induction step, we will first show that each set of size k is in the desired order in L^k and $\overline{L}^k$ with all sets of strictly smaller cardinality, then that it is in the desired order with all sets of equal cardinality.

So 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 needs 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 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 $A' || B$. 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 of $A'||B$, 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 $\overline{L}$. Thus the pair (X, X^c) does not yield a reversal between L and $\overline{L}$, 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 $\overline{L}$. By Lemma 8, 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 $\overline{L}$. 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 $\overline{L}$, 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 $\max_{\overline{\sigma}}(B \triangle A'') = a \in A''$. So we have $B < A''$ in $\overline{L}$ by induction and thus $B < A$ in $\overline{L}$ 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. $\triangle$

Because $\max_\sigma(A \triangle B) = b \in B$, 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 $< k$ in L^k. By reversing the roles of L and $\overline{L}$ we see that A also has to be in $\overline{\sigma}$-revlex order in $\overline{L}$ 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, this means that 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 $a_i = \min_\sigma(A_j \triangle \{a_i\}) = \max_{\overline{\sigma}}(A_j \triangle \{a_i\})$. Thus by induction, $A_j < a_i < A_i$ in $\overline{L}$, that is, A_i and A_j belong to different slots in $\overline{L}$. Now since the class $\mathcal{C}_{D,I}$ containing (A_i, A_j) needs to contribute $2^{2-2} = 1$ reversal between L and $\overline{L}$, 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.

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

Another conjecture in [6] suggested a connection between diametral pairs and reversed critical pairs. A *critical pair* of a poset $\mathcal{P}$ is an ordered pair (x, y) of elements of $\mathcal{P}$ such that all elements smaller than x are also smaller than y in $\mathcal{P}$, and all elements larger than y are also larger than x in $\mathcal{P}$. Critical pairs appear as an important ingredient in the dimension theory of posets, see for example [15]. A critical pair (x, y) is *reversed* in a linear extension L of $\mathcal{P}$ if $x > y$ in L.

It was conjectured in [6] that in every diametral pair of linear extensions of $\mathcal{P}$, at least one of the two linear extensions reverses a critical pair of elements of $\mathcal{P}$. In [3] it was shown that the conjecture is false in general, but that almost all posets have the stronger property of being *diametrally reversing*, which means that *every* linear extension contained in a diametral pair reverses a critical pair. Still for Boolean lattices it remained open even whether they have the weaker property.

With the help of Theorem 9, we can now settle this question:

Corollary 10. *Boolean lattices are diametrally reversing.*

Proof. Let L be a linear extension of B_n that is contained in a diametral pair. Then by Theorem 9 we know that $L = L_\sigma$ for some permutation σ of $[n]$. It can be checked that the n atom-coatom pairs $(\{i\}, [n] \setminus i)$ for $i \in [n]$ are critical pairs of B_n (in fact, these are the only ones). Now let $i = \max_\sigma[n]$. Then $\max_\sigma(\{i\}\triangle[n] \setminus i) = i$. Thus $[n] \setminus i < \{i\}$ in L_σ, which means that L_σ reverses this critical pair. Hence every linear extension of B_n contained in a diametral pair reverses a critical pair of B_n. □

3 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. Now let $\mathcal{P}$ be an arbitrary poset. Denote by $\mathcal{D}_\mathcal{P}$ the downset lattice of $\mathcal{P}$, that is, the poset on all downsets of $\mathcal{P}$, ordered by inclusion. In Section 3.1 we give an upper bound for the linear extension diameter of $\mathcal{D}_\mathcal{P}$ and show that it is tight if $\mathcal{P}$ is 2-dimensional. We also characterize the diametral pairs of linear extensions of $\mathcal{D}_\mathcal{P}$ for 2-dimensional $\mathcal{P}$. For the proofs, we make use of the main ideas from the last section. In Section 3.2 we show how to compute the linear extension diameter of $\mathcal{D}_\mathcal{P}$ in time polynomial in $|\mathcal{P}|$.

We can use our results to obtain optimal drawings of $\mathcal{D}_{\mathcal{P}}$ as described in the introduction, see Figure 3 for an example.

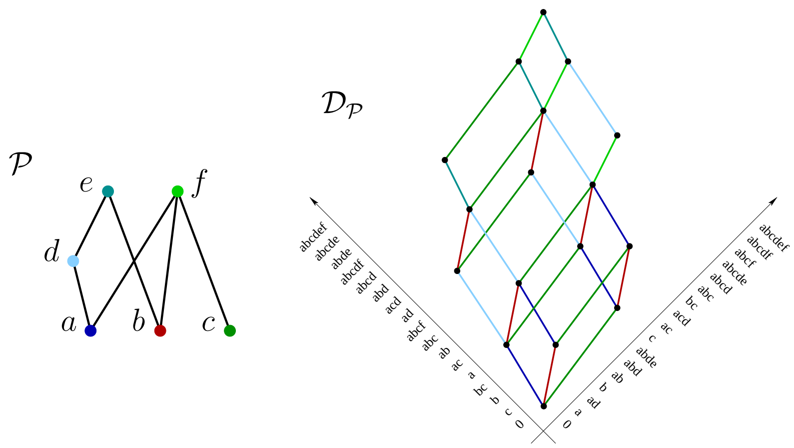

Figure 3: A drawing of a downset lattice based on a diametral pair.

In this section we make fundamental use of the canonical bijection between downsets and antichains of a poset. We frequently switch back and forth between the two viewpoints, concentrating on the one or the other in our proofs. We write $A^{\downarrow}$ to refer to the downset generated by an antichain A. We write $\mathrm{Max}(\mathcal{A})$ to refer to the antichain of maxima of a downset $\mathcal{A}$.

3.1 Characterizing Diametral Pairs

This subsection is organized as follows: We first prove an upper bound for the linear extension diameter of arbitrary downset lattices, using a generalization of the equivalence classes from Definition 3. Then we show that for downset lattices of 2-dimensional posets, a generalization of the revlex orders attains this bound. Finally we prove that these are the only pairs of linear extensions attaining the bound.

Definition 11. *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 $D = A \triangle B$ 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 12. *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 $D = A \triangle B$. First observe that since A and B are antichains, $\mathcal{P}[D]$ is a poset of height at most 2 (see Figure 3.1). Thus all elements of D belong to the antichain $\mathrm{Max}(D)$ of maxima of $\mathcal{P}[D]$ or to the antichain $\mathrm{Min}(D)$ 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. Note that in the second case the minima also belong to $B^{\downarrow}$, and in the first case also to $A^{\downarrow}$.

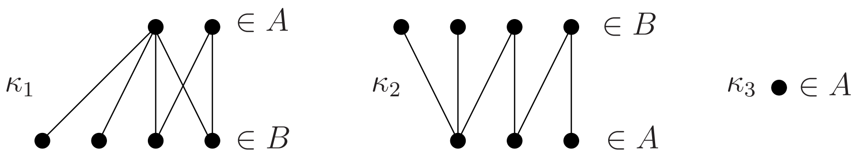

Figure 4: An example for $\mathcal{P}[D]$ with three components. The assignment of minima and maxima to A and B specifies one of the eight pairs in the class.

We can get 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 respectively B 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

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

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

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

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. Also, given a pair $(A, B) \in \mathcal{C}_{D,I}$ it is easy to construct $K \subseteq \mathcal{K}$ with $(A, B) = (A_K, B_K)$. Thus our map is a bijection as claimed. □

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 13. *Let $\mathcal{D}_{\mathcal{P}}$ be the downset lattice of a 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 $\mathcal{P}[A \triangle B]$ 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 11. 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 with Lemma 12 we know that $\mathcal{C}_{D,I}$ consists of the two pairs (A, B) and (B, A), where $A = \mathrm{Max}(D) \cup I$ and $B = (D \setminus \mathrm{Max}(D)) \cup I$. 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 (unordered) reversals that each class can contribute is at most a quarter of the number of pairs that it contains.

Recall the notions from Lemma 12. 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) 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 $|\mathcal{F}_1 \triangle \mathcal{F}_2|$.

Let $K \in \mathcal{F}_i$, and $K' \subset K$. We claim that $K' \in F_i$. 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 $A^{\downarrow}_{K'} < A^{\downarrow}_{K} < B^{\downarrow}_{K} < B^{\downarrow}_{K'}$ 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 7. 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

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

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. □

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 a particular type of linear extension of $\mathcal{D}_{\mathcal{P}}$. 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. In analogy to Definition 1, we have:

Definition 14. *We define a relation $<_\sigma$ on the pairs $\{A^{\downarrow}, B^{\downarrow}\}$ of downsets of $\mathcal{P}$ as follows:*

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

We call this relation the σ-revlex order.

Lemma 15. *Let $\mathcal{P}$ be a poset and σ a linear extension of $\mathcal{P}$. The relation of being in σ-revlex order defines a linear extension L_σ of $\mathcal{D}_{\mathcal{P}}$.*

Proof. Since $\max_\sigma(A^{\downarrow} \triangle B^{\downarrow}) = \max_\sigma(A \triangle B)$, we can 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 1) to the downsets of $\mathcal{P}$. In Lemma 2, 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 $\mathcal{P}$. □

Note that for the special case $\mathcal{P} = B_n$, the above lemma was already shown in [7].

Now let σ be a linear extension of $\mathcal{P}$ which is contained in a diametral pair of $\mathcal{P}$. Since $\mathcal{P}$ is 2-dimensional, σ has a unique partner $\overline{\sigma}$ with which it forms a realizer, i.e., a diametral pair. Thus all incomparable pairs of $\mathcal{P}$ are reversals between σ and $\overline{\sigma}$. We will frequently use the following characterization, which is implicite already in [5], cf. also [10].

Lemma 16. *Let $\mathcal{P}$ be a poset of dimension 2. A linear extensions π of $\mathcal{P}$ is contained in a realizer of cardinality 2 exactly if it is* non-separating, *that is, for $u, v \in \mathcal{P}$ with $u < v$ in $\mathcal{P}$, there is no element $x \in \mathcal{P}$ with $x \,||\, u, v$ in $\mathcal{P}$ and $u < x < w$ in π.*

In particular, σ and $\overline{\sigma}$ are non-separating. If a linear extension is not non-separating, we call it *separating.*

Theorem 17. *Let $\mathcal{P}$ be a 2-dimensional poset, and let $\sigma, \overline{\sigma}$ be a diametral pair of linear extensions of $\mathcal{P}$. Then $L_\sigma, L_{\overline{\sigma}}$ is a diametral pair of linear extensions of $\mathcal{D}_\mathcal{P}$.*

Proof. We will show that for each class $\mathcal{C}_{D,I}$, the pair $L_\sigma, L_{\overline{\sigma}}$ realizes the maximum possible number of reversals. Fix an arbitrary class $\mathcal{C}_{D,I}$. We have seen 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 $L_{\overline{\sigma}}$ exactly if one of the two elements $\max_\sigma(A \triangle B)$ and $\max_{\overline{\sigma}}(A \triangle B)$ 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 $\max_\sigma(A \triangle B) \in A$ and $\max_{\overline{\sigma}}(A \triangle B) \in B$.

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

Suppose for contradiction that $\max_\sigma \mathrm{Max}(D)$ and $\min_\sigma \mathrm{Max}(D)$ both lie in the component κ of $\mathcal{P}[D]$. Since we assumed that $\mathcal{P}[D]$ has at least two components, we can choose an element $x \in \mathrm{Max}(D)$ which is not contained in κ and thus has no relation to any element of κ. Then $\min_\sigma \mathrm{Max}(D) < x < \max_\sigma \mathrm{Max}(D)$ 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 $\min_\sigma \mathrm{Max}(D) \in \kappa_1$ and $\max_\sigma \mathrm{Max}(D) \in \kappa_2$. 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 $x || u, v$ in $\mathcal{P}$, and this is a contradiction because σ is non-separating.

We have shown that $\max_\sigma(A \triangle B)$ and $\max_{\overline{\sigma}}(A \triangle B)$ lie in different components of $\mathcal{P}[D]$, say, $\max_\sigma(A \triangle B) \in \kappa$ and $\max_{\overline{\sigma}}(A \triangle B) \in \lambda$ 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 $\max_\sigma(A \triangle B) \in A_K$ and $\max_{\overline{\sigma}}(A \triangle B) \in B_K$. Thus the pair (A_K, B_K) contributes a reversal between L_σ and $L_{\overline{\sigma}}$. 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 $L_{\overline{\sigma}}$. We have seen in Theorem 13 that this is maximal. Therefore L_σ and $L_{\overline{\sigma}}$ are a diametral pair of linear extensions of $\mathcal{D}_\mathcal{P}$. □

We have now proved the fact 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 18. *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 $\mathcal{P}[A \triangle B]$ has at least two connected components.*

Note that the proof of Conjecture 1, thus, the formula for the linear extension diameter of the Boolean lattice, follows from this corollary. A special case of the proof of Theorem 22 yields $\text{led}(B_n) = 2^{2n-2} - (n+1)2^{n-2}$.

In the next theorem we characterize the diametral pairs of linear extensions of downset lattices of 2-dimensional posets.

Theorem 19. *Let $\mathcal{P}$ be a 2-dimensional poset, and let $L, \overline{L}$ 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 diametral pair $\sigma, \overline{\sigma}$ of linear extensions of $\mathcal{P}$, and we have $L = L_\sigma$ and $\overline{L} = L_{\overline{\sigma}}$.*

Proof. We again use the equivalence classes from Definition 11. 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 $L, \overline{L}$ 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 $\overline{L}$. 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 $\overline{L}$ by $\overline{\sigma}$. We have seen that every incomparable pair of elements of $\mathcal{P}$ must be a reversal between σ and $\overline{\sigma}$. Hence, $\sigma, \overline{\sigma}$ is a diametral pair of linear extensions 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 $\overline{L} = L_{\overline{\sigma}}$, in analogy to the proof of Theorem 9. More precisely, we show that each downset of cardinality k is in σ-revlex order in L and in $\overline{\sigma}$-revlex order in $\overline{L}$ 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 $\overline{L}$. 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 $\overline{L}^{k-1} = L_{\overline{\sigma}}^{k-1}$. 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 $\overline{L}^k$ 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 $\overline{L}^k$.

So 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 $\max_{\overline{\sigma}}(A^\downarrow \triangle B^\downarrow) = \max_{\overline{\sigma}}(A \triangle B)$. Using these we see that $\{A^\downarrow, B^\downarrow\}$ needs to be a reversal between L and $\overline{L}$. From the order of $A^\downarrow$ and $B^\downarrow$ in $\overline{L}$ we can finally deduce that $A^\downarrow < B^\downarrow$ in L^k.

We have $\tilde{A}^\downarrow = A^\downarrow \backslash a$ 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 $\tilde{A}^\downarrow$ is σ-smaller than the element of $A^\downarrow$ that is missing in $\hat{A}^\downarrow$. Thus we can conclude that $a = \min_\sigma A$. Since the antichain A is completely reversed between σ and $\overline{\sigma}$, it follows that $a = \max_{\overline{\sigma}} A$.

Now observe that since $\tilde{A}^\downarrow < B^\downarrow$ in L^{k-1} and $|\tilde{A}^\downarrow| = k-1$ we have $\tilde{A}^\downarrow \,||\, B^\downarrow$. 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 $\mathcal{P}$. Observe that 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 $\max_{\overline{\sigma}}(A^\downarrow \triangle B^\downarrow) = a$.

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, $a \not< b$ in $\mathcal{P}$ because otherwise $a \in B^\downarrow \setminus \tilde{A}^\downarrow$ and thus $a = b$, a contradiction. It follows that $a || b$ 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 $a||a'$ in $\mathcal{P}$, then $a' \in A \setminus B$ and with $a = \min_\sigma A$ it follows that $a <_\sigma a'$. 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 $b \,||\, a', a$ and $a' < a$ in $\mathcal{P}$. This means that σ is a separating linear extension. But since σ is contained in the realizer $\sigma, \overline{\sigma}$ of $\mathcal{P}$, this is a contradiction.

We have shown that $a <_\sigma b$. We knew already that $\max_\sigma(\tilde{A}^\downarrow \triangle B^\downarrow) = b$, and because a is the only element in $A^\downarrow \setminus \tilde{A}^\downarrow$, we can conclude that

$\max_\sigma(A^\downarrow \triangle B^\downarrow) = b$. Also, since $a||b$ in $\mathcal{P}$ we now know that $a >_{\overline{\sigma}} b$. Because $a = \max_{\overline{\sigma}} A$ and $B^\downarrow \setminus A^\downarrow = \{b\}$, we have $\max_{\overline{\sigma}}(A^\downarrow \triangle B^\downarrow) = a$.

Let us now consider the class $\mathcal{C}_{D,I}$ defined by $D = A\triangle B$ and $I = A \cap B$. The elements a and b lie in different components of $\mathcal{P}[D]$, because $B\setminus A = \{b\}$ and $a||b$ 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 $\max_\sigma(A_K \triangle (b \cup I)) = b$, 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, $\max_{\overline{\sigma}}(A_K \triangle (a \cup I) = a$ holds, so we have $A_K^\downarrow < (a \cup I)^\downarrow$ in $\overline{L}$ by induction and thus $A_K^\downarrow < (a \cup I)^\downarrow < (X_K^c \cup I)^\downarrow = B_K^\downarrow$ in $\overline{L}$.

It follows that $A_K^\downarrow$ is down in L and $\overline{L}$ 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 $\overline{L}$, 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 $\overline{L}$. From the fact we remarked after Theorem 17 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 $\overline{L}$: For $(I \cup a)^\downarrow \subset A^\downarrow$, we have $\max_{\overline{\sigma}}(B \triangle (I \cup a)) = a$, so we have $B^\downarrow < (I \cup a)^\downarrow$ in $\overline{L}$ by induction and thus $B^\downarrow < A^\downarrow$ in $\overline{L}$ 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 $\max_\sigma(A\triangle B) = b \in B$, 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 $\overline{L}$, we obtain that $A^\downarrow$ is in $\overline{\sigma}$-revlex order in $\overline{L}^k$ 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, $A_i^{\downarrow} \triangle A_j^{\downarrow} = \{a_i, a_j\}$. We have $a_i = A_i^{\downarrow} \setminus \tilde{A}^{\downarrow} = \min_\sigma A_i = \max_{\overline{\sigma}} A_i$ and $a_j = A_j^{\downarrow} \setminus \tilde{A}^{\downarrow} = \min_\sigma A_j = \max_{\overline{\sigma}} A_j$. Note that a_i and a_j have to be incomparable in $\mathcal{P}$, and assume that $a_i <_\sigma a_j$, thus $a_j <_{\overline{\sigma}} a_i$. Then for the pair $\{A_j^{\downarrow}, a_i^{\downarrow}\}$ we know $\max_{\overline{\sigma}}(A_j^{\downarrow} \triangle a_i^{\downarrow}) = a_i$. Hence by induction, $A_j^{\downarrow} < a_i^{\downarrow} < A_i^{\downarrow}$ in $\overline{L}$. But since the equivalence class containing $(A_i^{\downarrow}, A_j^{\downarrow})$ needs to contribute $2^{2-2} = 1$ reversal between L and $\overline{L}$, 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 of $\max_\sigma A_i \triangle A_j = a_j$, this means that $A_i^{\downarrow}$ and $A_j^{\downarrow}$ are in revlex order in L.

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

3.2 Computing the Linear Extension Diameter

It is NP-complete to compute the linear extension diameter of a general poset, see [3]. That is, the linear extension diameter of a general poset $\mathcal{P}$ cannot be computed in time polynomial in $|\mathcal{P}|$ (unless P = NP). But with the results of the previous section, the problem is tractable if the 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 a well-known fact about distributive lattices: From a downset 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 $\mathcal{D}$. Finding a minimal realizer $\sigma, \overline{\sigma}$ of the 2-dimensional poset $\mathcal{P}$ amounts to finding a transitive orientation of its incomparability graph (cf. e.g. [10]). This can be done in time linear in $|\mathcal{P}|$ [9]. Then $\sigma, \overline{\sigma}$ is a diametral pair of linear extensions of $\mathcal{P}$. With the definition of the σ-revlex order we can compute L_σ and $L_{\overline{\sigma}}$, and this is a diametral pair of linear extensions of $\mathcal{D}$ by Theorem 17. We know by Theorem 19 that all diametral pairs arise this way. The linear extension diameter of $\mathcal{D}$ can now be computed by just checking for all pairs of elements of $\mathcal{D}$ whether they form a reversal between L_σ and $L_{\overline{\sigma}}$.

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, so $|\mathcal{D}_\mathcal{P}|$ is 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 [14] or [10]. We give a proof in the lemma below since the methods we use for the following theorem rely on the same ideas.

Lemma 20. *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 $x_i || x_j$, 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

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

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. □

Theorem 21. *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 18, 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. For a pair (A, B) of antichains of $\mathcal{P}$, we set $D = A \triangle B$ 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 $\alpha = a(\mathcal{P})^2$. 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 want to count each A exactly $|A|$ times. We want to refine the ideas of the proof of Lemma 20 to keep track of the sizes of the antichains. Hence 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 $i = 1, 2, \ldots, n$ as follows: The first entry of each $s(x_i)$ is 1, counting the antichain $\{x_i\}$. For the other entries we have

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

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 take a look at the structure of $\mathcal{P}[D \cup I]$, see Figure 3.2.

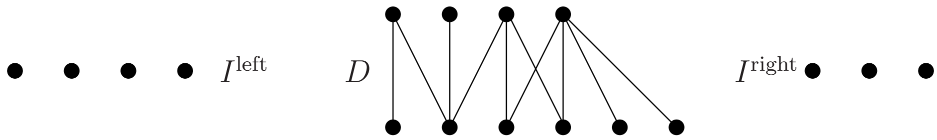

Figure 5: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in δ.

We know that $\mathcal{P}[I]$ is an antichain and $\mathcal{P}[D]$ consists of two antichains, $\mathrm{Max}(D)$ and $\mathrm{Min}(D)$, 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: I^{left}, living on the first interval, and I^{right}, living on the third interval.

Now define $\delta(k, \ell)$ as the number of pairs counted for δ which fulfill $\max_\sigma \mathrm{Min}(D) = x_k$ and $\max_\sigma \mathrm{Max}(D) = x_\ell$. 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 $\mathrm{Min}(D)$ and the antichain I^{left}. By definition we have $\max_\sigma \mathrm{Min}(D) = x_k$. Suppose that $\min_\sigma \mathrm{Min}(D) = x_i$ as in Figure 3.2.

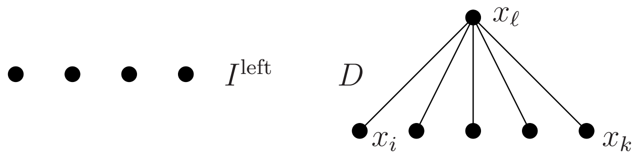

Figure 6: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in $\delta_1(k, \ell)$.

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 $\mathrm{Min}(D)$, 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 $\mathcal{P}^{\text{left}}_{i,\ell}$ be the poset induced by the elements $x \in \mathcal{P}$ with $x \in [x_1, x_i)$ and $x||x_\ell$. We claim that the sets which can be chosen as I^{left} are exactly the antichains of $\mathcal{P}^{\text{left}}_{i,\ell}$.

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 $\mathcal{P}^{\text{left}}_{i,\ell}$ as claimed. Altogether we have

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

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 7).

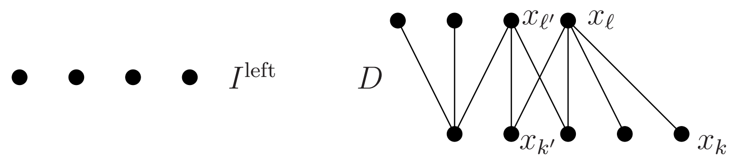

Figure 7: $\mathcal{P}[D \cup I]$ for a pair of antichains counted in $\delta_2(k, \ell)$.

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 $x_{\ell'}||x_\ell$. In general, $\mathcal{P}[D]$ is not connected after deletion of x_ℓ. But since $\mathcal{P}[D]$ was connected originally, $\mathrm{Min}(D)$

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 $\mathrm{Min}(D)$. With the same reasoning as for $\delta_1(k,\ell)$, these form exactly the antichains in $\mathcal{P}_{k',k,\ell}$. So we have

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

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, it cannot be larger than any element of D either: If $x > y$ for some element $y \in \mathrm{Min}(D)$, 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:

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

What is the overall running time for the computation of $\mathrm{led}(\mathcal{D}_\mathcal{P})$? From Lemma 20 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 $r = 1 \ldots n$ and $i = 1 \ldots n$. 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 20 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 $\mathrm{led}(D_\mathcal{P})$ can be computed in time $O(n^5)$. □

In the previous theorem we showed how to compute $\mathrm{led}(D_\mathcal{P})$ for 2-dimensional $\mathcal{P}$, 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. These lattices $\mathcal{D}_\mathcal{P}$ are also known as the factor lattices of integers: If $\mathcal{P} = C_1 \cup \ldots \cup C_w$ with $|C_i| = \ell_i$, we can associate each chain with a prime number p_i, and 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 22. *If $\mathcal{P} = C_1 \cup \ldots \cup C_w$ with $|C_i| = \ell_i$ is a disjoint union of chains, then the linear extension diameter of $\mathcal{D}_\mathcal{P}$ equals*

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

Proof. From Corollary 18 we know that $\operatorname{led}(\mathcal{D}_\mathcal{P})$ 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 substract 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 $\prod_{i=1}^{\omega} (\ell_i + 1)$ antichains, and this is also the number of pairs of antichains differing on zero chains. The number of all pairs of antichains is thus $(\prod_{i=1}^{\omega} (\ell_i + 1))^2$. 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. □

4 Open Problems

It is NP-hard to compute the linear extension diameter of a general poset, see [3]. 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 arbitrary distributive lattices?

Question 1. *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?*

It is also interesting to look at subposets of the Boolean lattice. Consider for example the subposet of B_5 induced by the sets of cardinality 2 and 3. It turns out that the revlex linear extensions do not form a diametral pair of linear extensions of this poset.

Question 2. *Can we construct the diametral pairs of linear extensions of a poset induced by two levels of B_n?*

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 [2] 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. Here, $\text{inc}(\mathcal{P})$ denotes the number of incomparable pairs of $\mathcal{P}$.
Acknowledgement. We thank Kolja Knauer for fruitful discussions.

References

[1] Ian Anderson. *Combinatorics of finite sets*. Oxford Clarendon Press, 1987.

[2] Graham Brightwell, 2009. Personal communication.

[3] Graham Brightwell and Mareike Massow. Diametral pairs of linear extensions. submitted; available at `http://arxiv.org/abs/0809.1828`.

[4] Russ Bubley and Martin E. Dyer. Faster random generation of linear extensions. *Discrete Math.*, 201:81–88, 1999.

[5] Ben Dushnik and E. W. Miller. Partially ordered sets. *Amer. J. Math.*, 63:600–610, 1941.

[6] Stefan Felsner and Klaus Reuter. The linear extension diameter of a poset. *SIAM J. Disc. Math.*, 12(3):360–373, 1999.

[7] Serkan Hoşten and Walter D. Morris, Jr. The order dimension of the complete graph. *Discrete Math*, 201:133–139, 1998.

[8] Daniel J. Kleitman. Families of non-disjoint subsets. *J. Combinat. Theory*, 1:153–155, 1966.

[9] Ross M. Mcconnell and Jeremy P. Spinrad. Modular decomposition and transitive orientation. *Discrete Math*, 201:189–241, 1997.

[10] Rolf H. Möhring. Computationally tractable classes of ordered sets. In Ivan Rival, editor, *Algorithms and Order*, volume 255 of *NATO ASI Series C*. Kluwer Academic Publishers, 1989.

[11] Michael Naatz. The graph of linear extensions revisited. *SIAM J. Disc. Math.*, 13(3):354–369, 2000.

[12] Gara Pruesse and Frank Ruskey. Generating the linear extensions of certain posets by transpositions. *SIAM J. Discrete Mathematics*, 4:413–422, 1991.

[13] Klaus Reuter. Linear extensions of a poset as abstract convex sets. *Hamburger Beiträge zur Mathematik*, Heft 56, 1996.

[14] George Steiner. Single machine scheduling with precedence constraints of dimension 2. *Math. Oper. Res.*, 9(2):248–259, 1984.

[15] William T. Trotter. *Combinatorics and partially ordered sets*. The Johns Hopkins University Press, 1992.