Navigating Posets with Few Maps

Stefan Felsner
Technische Universität Berlin, Germany

Jędrzej Hodor
Theoretical Computer Science Department, Faculty of Mathematics and Computer Science and Doctoral School of Exact and Natural Sciences, Jagiellonian University, Kraków, Poland

Giacomo Ortali
Università Telematica Internazionale Uninettuno, Rome, Italy

Alexander Wolff
Universität Würzburg, Germany

Abstract

We study two new parameters for finite posets motivated by the problem of efficiently determining the set of successors of a given element. A *plane map* of a poset $P = (X, \leq)$ is an injective mapping of X into the Cartesian plane $\mathbb{R}^2$. Given two different points a and b in the plane, we say that b *dominates* a if $a < b$ coordinatewise. We say that an element x of P is *tight* in a plane map μ if the following holds: $x < y$ in P if and only if $\mu(y)$ dominates $\mu(x)$. Note that, by definition, every 2-dimensional poset admits a map such that every element of the poset is tight. For any poset P, we define the *mapability* of P, $\mathrm{dmap}(P)$, to be the maximum number of elements that are tight in a single map, and we define the *atlas thickness* of P, $\mathrm{at}(P)$, to be the size of the smallest collection of maps such that every element is tight in at least one map of the collection.

We relate these parameters to the classical notions of dimension and width: for every poset P, we show that $\dim(P) \leq 2\,\mathrm{at}(P) \leq \mathrm{width}(P)$. On the other hand, there exists a sequence of posets $(P_n)_{n\geq 1}$ such that the atlas thickness of P_n is doubly exponential in the dimension of P_n.

On the computational side, we prove that it is NP-complete, for a given poset P, to compute the mapability of P and to decide whether $\mathrm{at}(P) \leq 2$. In contrast to the latter, we show that computing the mapability of a poset is fixed-parameter tractable with respect to the natural parameter.

2012 ACM Subject Classification Mathematics of computing → Graph theory

Keywords and phrases Poset, dimension, width, atlas thickness, mapability, NP-hard, fixed-parameter tractable

Category Track 1, Long Paper

Acknowledgements This work was started at the workshop Homonolo 2024 in Nová Louka. We thank the organizers and the other participants; in particular Bartosz Podkanowicz and Felix Schröder. J. Hodor is supported by the National Science Center of Poland under grant UMO-2022/47/B/ST6/02837 within the OPUS 24 program.

Generative AI Declaration Generative AI was not used in the preparation of this article.

1 Introduction

In combinatorics and in computer science there is interest in compact encodings of posets such that the presence of an order relation between two elements can be tested efficiently. We propose and study a new idea for such an encoding. The encoding consists of an atlas of pages which show the elements of the poset on plane maps. We aim for a small number of pages such that, for each element x, there is a page which shows all relations $x \leq y$.

We assume familiarity with basic notions of the theory of *partially ordered sets* (posets) as used, e.g., in West's book *Combinatorial Mathematics* [13]. Due to their relevance to our topic, we nevertheless recall the definition of dimension and some related concepts. A *realizer* of a poset P is a collection $\mathcal{R}$ of linear extensions of P such that, for every incomparable pair of elements x and y in P, there are L and L' in $\mathcal{R}$ such that $x <_L y$ and $y <_{L'} x$. This is equivalent to saying that P is the intersection of the linear extensions in $\mathcal{R}$. The *dimension* of a poset P, denoted $\dim(P)$, is the minimum size of a realizer of P. It is well know that $\dim(P)$ also equals the minimum k such that P is a suborder of $\mathbb{R}^k$ with the *product order*, also known as *dominance order*, where $x \leq_{\text{dom}} y$ if and only if $x_i \leq y_i$ for each $i \in \{1, \ldots, k\}$. In a poset P, a *chain* is a subset of P of pairwise comparable elements (that is, they are linearly ordered), whereas an *antichain* is a subset of P of pairwise incomparable elements. The *width* of a poset P, $\text{width}(P)$, is the size of a largest antichain in P.

A common motivation for studying the dimension of posets is that a realizer of size k for a poset of size n provides a data structure of size $\tilde{\mathcal{O}}(nk)$ that allows to determine the comparability status of a pair (x, y) with k lookup operations. To improve the size of the data structure, Nešetřil and Pudlák [12] introduced the concept of *Boolean dimension*. This provides a data structure that in general is of smaller size and requires fewer lookup operations but may need a long computation to determine the comparability status of a pair.

In this paper, we go the other direction. We are willing to use more storage but for a given element x of a poset P, we want to find the set $\text{succ}_P[x] = \{y \in P : x \leq_P y\}$ quickly.

Here is a first approach. Let P be a poset, and let $\mathcal{T}$ be a set of permutations of its elements such that, for every x in P, there is a π in $\mathcal{T}$ such that $\text{succ}_P[x] = \{y \in P : x \leq_\pi y\}$. In this case we say that x is *tight* in π. How many permutations are needed such that all elements of P are tight in at least one of them? The question has a simple answer. Observe that if x and y are incomparable, then at most one of x and y can be tight in a single permutation. Hence, the tight elements of a permutation form a chain and chains of tight elements of all permutations in $\mathcal{T}$ form a chain cover of P. From Dilworth's theorem [13, p. 546], we therefore get the lower bound $|\mathcal{T}| \geq \text{width}(P)$. Since, for every chain C of P, there is a linear extension L such that all elements of C are tight in L, we see that $\text{width}(P)$ is the minimum size of a family $\mathcal{T}$ and that indeed it is enough to use linear extensions in the family. A drawback of this approach is that it would use n linear extensions just to get all elements of a size-n antichain tight.

This motivates our second approach. A *(plane) map* for a poset P is an injective mapping μ of the elements of P into the Cartesian plane $\mathbb{R}^2$.

An element x of P is *tight* on a map μ if $\text{succ}_P[x] = \{y \in P : \mu(x) \leq_{\text{dom}} \mu(y)\}$. For example, the element b of the poset Q shown in Figure 1a is tight on the map in Figure 1c but not on the map in Figure 1b. We are interested in the following two parameters of a poset P:

- The *(dominance) mapability* of P, denoted $\text{dmap}(P)$, is the maximum number of tight elements on a single plane map for P.
- Let an *atlas* for P be a collection of plane maps for P such that every element of P is tight on at least one of the maps. Then the *atlas thickness* of P, denoted $\text{at}(P)$, is the minimum size of an atlas for P.

We refer to Figure 1. Figure 1a depicts a poset Q. Figures 1b and 1c show a collection of two maps where every element of Q is tight on at least one map, since elements $\{a, c, ab, ac, bc\}$ are tight in the former and b is tight in the latter. As we are going to observe later, Q is the crown poset, which has dimension 3. Hence, it is not possible to find a map where more than five elements of Q are tight, or a collection of maps for Q with fewer than two maps. This implies that P has mapability 5 and atlas thickness 2.

Note that, by definition, every 2-dimensional poset admits a single map such that every element of the poset is tight.

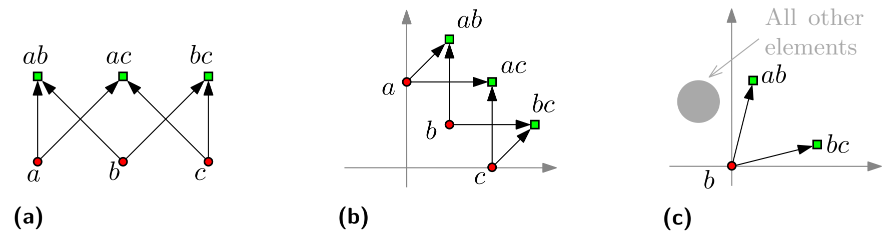

Figure 1 (a) The so-called crown poset Q. (b) A map of Q where all elements but b are tight. (c) A map of Q where b is tight.

Now that we have introduced the notion of an atlas, we briefly return to our motivation. Given an atlas A for P and an element x of P, we can use a simple look-up table to determine the map in A on which x is tight. If, for each map μ in A, we store the points in $\mu(P)$ in a separate priority search tree T_μ, we can query T_μ with the quadrant $[\mu(x)_1, \infty] \times [\mu(x)_2, \infty]$ (see paragraph "Notation" on page 4) to determine the set $\mathrm{succ}_P[x]$ in $O(|\mathrm{succ}_P[x]| + \log |P|)$ time [11]. The total space consumption of the priority search trees is $O(\mathrm{at}(P) \cdot |P|)$; the total preprocessing time for setting up the trees is their space consumption times a factor of $O(\log |P|)$.

While our definition requires each map in A to define locations for every element in P, this is not needed if just want to solve the above query problem. For example, in the map depicted in Figure 1c, there is no need to store the points in the gray circle. Hence, the total space consumption is just $O(\sum_{x \in P} \mathrm{succ}_P[x])$, which is the number of pairs of comparable elements in P.

Our contribution. We relate the atlas thickness of a poset to its dimension and width; see Section 2 for the proof of the following theorem.

▶ **Theorem 1.** *For every poset P,* $\dim(P) \leq 2\,\mathrm{at}(P) \leq \mathrm{width}(P)$.

The inequalities in Theorem 1 are tight, however, there are posets whose atlas thickness is much larger than their dimension. Note also that antichains have atlas thickness 1 and arbitrary width. The next two statements are proved in Section 3. Theorem 2 is witnessed by the family of standard examples, and Theorem 3 is witnessed by the family of the incidence posets of complete graphs.

▶ **Theorem 2.** *For every positive integer n, there is a poset P_n with* $\dim(P_n) = 2\,\mathrm{at}(P_n) = 2n$.

▶ **Theorem 3.** *There exists a sequence of posets $(P_n)_{n \geq 1}$ such that* $\mathrm{at}(P_n) = 2^{2^{\Omega(\dim(P_n))}}$.

The problem of checking whether a given poset has dimension at most k is NP-complete already for $k = 3$ [14]. Perhaps unsurprisingly, computing the parameters that we have introduced is also computationally hard; see Section 4.

▶ **Theorem 4.** *It is* NP*-complete to test whether the atlas thickness of a poset is at most* 2.

▶ **Theorem 5.** *It is* NP*-complete to compute the mapability of a poset.*

Theorem 4 shows that computing the atlas thickness of a poset is paraNP-hard with respect to the natural parameter. This is in contrast to the mapability of a poset, which turns out to be fixed-parameter tractable (FPT) with respect to the natural parameter; see Section 5. We prove that the size of each poset is bounded in terms of its mapability.

▶ **Theorem 6.** *There exists a computable function f such that, for every poset P, we have* $|P| \leq f(\mathrm{dmap}(P))$.

This directly yields an FPT-algorithm with respect to the mapability of the given poset.

▶ **Corollary 7.** *There exists a computable function g and an algorithm that, given a poset P, returns the value* $\mathrm{dmap}(P)$ *in time* $g(\mathrm{dmap}(P))$.

Related work. Our work follows a recent trend in graph drawing to visualize a complex combinatorial or geometric object (say, a graph) by means of a collection of drawings, each of which includes all parts (say, edges), but highlights only a subset of them. For example, Hliněný and Masařík [8] introduced the notion of an *uncrossed collection* of a graph, that is, a collection of straight-line drawings of the graph such that every edge of the graph is uncrossed in at least one drawing in the collection. They considered two optimization problems; namely minimizing the size of the collection (the uncrossed number) and minimizing the total number of crossings in the collection (the uncrossed crossing number). They showed the latter problem NP-hard but in FPT with respect to the natural parameter. Similarly, Antić, Liotta, Masařík, Ortali, Pfretzschner, Stumpf, Wolff, and Zink [1] introduced *unbent collections* for plane 4-graphs, that is, for embedded planar graphs of vertex degree at most 4. Such a collection contains orthogonal drawings of the given graph such that every edge of the graph is straight in at least one drawing in the collection. They also considered the corresponding optimization problems of computing, for a given plane 4-graph, the unbent number (which they proved to be always at most 3) and the unbent bend number (for which they showed NP-hardness and gave a 3-approximation algorithm).

Less related are storyplans [3,7] and graph stories [2], where the drawings in the collection show only parts of the graph (with a prescribed property such as planarity) and the order of the drawings is relevant.

Notation. For a positive integer k, we use $[k]$ as shorthand for $\{1, 2, \ldots, k\}$. For a poset P, we let $\mathrm{Max}(P)$ denote the set of all maximal elements of P, and we let $|P|$ denote the number of elements of P. For a map μ and an element x of a poset P, we write $\mu(x)_1$ and $\mu(x)_2$ for the first and the second coordinates of $\mu(x) \in \mathbb{R}^2$, respectively.

For a poset P, for each element x of P, we write $\mathrm{succ}_P[x] = \{y \in P : x \leq_P y\}$ and $\mathrm{succ}_P(x) = \{y \in P : x <_P y\}$ to denote, respectively, the *closed* and the *open* set of *successors* of x.

The *incidence poset* of a graph G, denoted by P_G, is the poset with ground set $V(G) \cup E(G)$ such that, for every two elements x and y of P_G, we have $x < y$ if and only if $x \in V(G)$, $y \in E(G)$, and x is incident to y in G. For an example, see Figure 2.

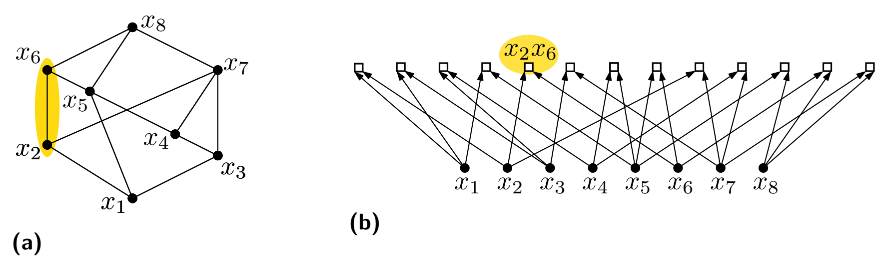

Figure 2 (a) A graph G and (b) the incidence poset P_G of G, represented as a DAG.

The pieces of text in gray explain statements that may be obvious for readers who feel at home with posets.

2 Comparision Between Parameters

The following lemma, which may be of independent interest, is the main tool for the proof of Theorem 1.

▶ **Lemma 8.** *For every poset P and for every map μ for P, there exists a map μ' for P such that each element of P that is tight on μ is tight on μ', and the two linear orders of the elements of P given by the orders on the coordinates of μ' are linear extensions of P.*

Proof. Let $T = \{t_1, \ldots, t_k\}$ be the set of elements of P that are tight on μ. Note that some elements of P might have the same value under $\mu(\cdot)_1$ or $\mu(\cdot)_2$. Let $x_1, \ldots, x_k$ be any linear order on the elements of T extending the order given by $\mu(\cdot)_1$, and let $y_1, \ldots, y_k$ any linear order on the elements of T extending $\mu(\cdot)_2$. For each element p of P with $p \notin T$, let

$$i(p) = \max(\{i \in [k] : x_i < p\} \cup \{0\}) \quad \text{and} \quad j(p) = \max(\{j \in [k] : y_j < p\} \cup \{0\}).$$

In other words, $x_{i(p)}$ is the largest element of T with respect to $\mu(\cdot)_1$ that is below p in P, and $y_{j(p)}$ is defined analogously. Next, for each $i \in [k] \cup \{0\}$, let

$$X_i = \{p \in P \setminus T : i(p) = i\} \quad \text{and} \quad Y_i = \{p \in P \setminus T : j(p) = i\}.$$

For each $i \in [k] \cup \{0\}$, let K_i be an arbitrary linear extension of the subposet of P induced by the elements of X_i, and let L_i be an arbitrary linear extension of the subposet of P induced by the elements of Y_i. Finally, let μ' be defined so that the linear order given by $\mu'(\cdot)_1$ is the concatenation $K = K_0 x_1 K_1 x_2 K_2 \ldots K_{k-1} x_k K_k$, and the linear order given by $\mu'(\cdot)_2$ is the concatenation $L = L_0 y_1 L_1 y_2 L_2 \ldots L_{k-1} y_k L_k$.

To complete the proof, we need to verify that the elements of T are tight on μ', and that L and K are linear extensions of P. An element $p \notin T$ can be pushed left and down. By definition of $i(p)$ and $j(p)$, however, it will not pass any $t \in T$ with $t < p$ in P. Hence, elements of T remain dominated by all their successors in P which means they remain tight.

More formally, let $t \in T$ be such that $t = x_i = y_j$. We will show that $\mathrm{succ}_P[t] = \{p \in P : \mu'(t) \leq \mu'(p)\}$ First, we have $\{p \in P : \mu(t) \leq \mu(p)\} \cap T = \{p \in P : \mu'(t) \leq \mu'(p)\} \cap T$. Next, let p be an element of P with $p \notin T$. If $p \in \mathrm{succ}_P(t)$, i.e., $t < p$ in P, then $i \leq i(p)$ and $j \leq j(p)$, hence, $t < p$ in both L and K, and so, $\mu'(t) \leq \mu'(p)$. If $\mu'(t) \leq \mu'(p)$, then $i \leq i(p)$ and $j \leq j(p)$, hence, since the elements of T are tight on μ, we have $\mu(t)_1 \leq \mu(x_{i(p)})_1 < \mu(p)_1$ and $\mu(t)_2 \leq \mu(y_{j(p)})_2 < \mu(p)_2$. It follows that $t \leq p$ in P, as desired.

Note that K and L are linear extensions of P. Indeed, let p and p' be distinct elements of P with $p < p'$, we claim that $p < p'$ in K. If $p, p' \in T$, then this is clear since the elements of T are tight on μ. If $p = x_i \in T$ and $p' \notin T$, then $i \leq i(p')$, hence, $p < p'$ in K. If $p \notin T$ and $p' = x_i \in T$, then $i(p) < i$ as otherwise $p' \leq x_{i(p)} < p$ in P, which is a contradiction. From $i(p) < i$ we then get $p < p'$ in K. Finally, assume that $p, p' \notin T$. By transitivity $i(p) \leq i(p')$ follows from $p < p'$. If $i(p) = i(p')$, then since $K_{i(p)}$ is a linear extension of the subposet of P induced by the elements of $X_{i(p)}$, $p < p'$ in K. If $i(p) < i(p')$, then clearly $p < p'$ in K. This completes the proof that K is a linear extension of P. The proof that L is a linear extension of P is symmetric. ◀

▶ **Corollary 9.** *For every poset P, $\dim(P) \leq 2\,\mathrm{at}(P)$.*

Proof. Let P be a poset. By Lemma 8, there exists a collection of maps $\mathcal{P}$ for P of size $\operatorname{at}(P)$ such that every element of P is tight on at least one of the maps, and for every map the two linear orders of the elements of P given by the orders on the coordinates of the map are linear extensions of P. We claim that the collection $\mathcal{L}$ of $2\operatorname{at}(P)$ mentioned linear extensions of P is a realizer of P.

Let x and y be two incomparable elements in P. To conclude, we show that there is a linear extension in $\mathcal{L}$ such that $x < y$ in L. Let $\mu \in \mathcal{P}$ be a map on which y is tight. In particular, $\mu(y) \not\leq \mu(x)$, and so, there is an $i \in [2]$ such that $\mu(x)_i < \mu(y)_i$. This completes the proof. ◀

▶ **Lemma 10.** *For every poset P and for every two chains C and D in P, there exists a map μ for P such that all the elements of C and D are tight on μ.*

Proof. Let P be a poset and let $C = c_1 < \cdots < c_k$ and $D = d_1 < \cdots < d_\ell$ be chains in P. Similarly as in the proof of Lemma 8, for each element p of P, we define

$$i(p) = \max(\{i \in [k] : c_i \leq p\} \cup \{0\}) \quad \text{and} \quad j(p) = \max(\{j \in [\ell] : d_j \leq p\} \cup \{0\}).$$

Next, for each $i \in [k] \cup \{0\}$ and $j \in [\ell] \cup \{0\}$, let

$$C_i = \{p \in P \setminus C : i(p) = i\} \quad \text{and} \quad D_j = \{p \in P \setminus D : j(p) = j\}.$$

For each $i \in [k] \cup \{0\}$, let K_i be an arbitrary linear extension of the subposet of P induced by the elements in C_i, and for each $j \in [\ell] \cup \{0\}$, let L_j be an arbitrary linear extension of the subposet of P induced by the elements in D_j. Finally, let μ be defined so that the linear order given by $\mu(\cdot)_1$ is the concatenation $K = K_0 c_1 K_1 c_2 K_2 \ldots K_{k-1} c_k K_k$, and the linear order given by $\mu(\cdot)_2$ is the concatenation $L = L_0 d_1 L_1 d_2 L_2 \ldots L_{\ell-1} d_\ell L_\ell$.

To complete the proof, we need to verify that the elements of $C \cup D$ are tight on μ. The orders L and K are linear extensions of P by definition. The elements of C are tight in K and the elements of D are tight in L. Since the order induced by the plane map μ is the intersection of K and L, all elements of $C \cup D$ are tight in μ.

More formally, let $i \in [k]$ and we will show that $\operatorname{succ}_P[c_i] = \{p \in P : \mu(c_i) \leq \mu(p)\}$. The proof for the elements of D is symmetric. First, let $p \in \operatorname{succ}_P[c_i]$. We claim that $\mu(c_i) \leq \mu(p)$. This is clear when $p = c_i$, hence, assume otherwise. Since $c_i < p$ in P, we have $i \leq i(p)$ and $c_i < p$ in L. This gives $\mu(c_i)_1 < \mu(p)_1$. By definition, $d_{j(p)} \leq p$ in P. It follows that $d_{j(p)} \not\leq c_i$ in P. In particular, as D is a chain, $j(c_i) < j(p)$ and $c_i < p$ in K. This gives $\mu(c_i)_2 < \mu(p)_2$ and $\mu(c_i) < \mu(p)$, as desired. Next, let p be an element of P with $\mu(c_i) \leq \mu(p)$. In particular, $\mu(c_i)_1 \leq \mu(p)_1$ and $c_i < p$ in K. Since C is a chain, this implies $c_i < p$ in P and $p \in \operatorname{succ}_P[c_i]$. This completes the proof. ◀

▶ **Theorem 1.** *For every poset P, $\dim(P) \leq 2\operatorname{at}(P) \leq \operatorname{width}(P)$.*

Proof. Corollary 9 implies the first inequality. By Dilworth's theorem, a poset P can be covered by a family of $\operatorname{width}(P)$ chains, hence, Lemma 10 implies the second inequality of the statement. ◀

3 The Constructions

In this section, we prove Theorems 2 and 3. In the core of the arguments, we use the following 6-element poset Q defined on the elements $\{a, b, c, ab, bc, ac\}$ where a, b, c are minimal elements, ab, bc, ac are maximal elements, and the order relations are: $a < ab$, $a < ac$, $b < ab$, $b < bc$,

$c < ac$, $c < bc$; see Figure 1(i). Among order theorists, the poset Q is known as a *crown* or *cycle* and also as the *standard example* S_3. It is well known that the dimension of Q is 3. This fact clearly implies Lemma 11 below. For completeness, we give a direct proof of the lemma anyway.

▶ **Lemma 11.** *There does not exist a map of Q where all of a, b, c are tight.*

Proof. Suppose to the contrary that μ is a map of Q where all of a, b, c are tight. Since they are all tight and they are pairwise incomparable in Q, we may assume that $\mu(a)_1 < \mu(b)_1 < \mu(c)_1$ and $\mu(c)_2 < \mu(b)_2 < \mu(a)_2$. Since $a < ac$ and $c < ac$ in Q, we have $\mu(c)_1 < \mu(ac)_1$ and $\mu(a)_2 < \mu(ac)_2$. It follows that $\mu(b)_1 < \mu(ac)_1$ and $\mu(b)_2 < \mu(ac)_2$, hence, $\mu(b) < \mu(ac)$. However, b and ac are incomparable in Q, which is a contradiction with b being tight in the considered map; see Figure 1(ii), where b is not tight. ◀

▶ **Theorem 2.** *For every positive integer n, there is a poset P_n with $\dim(P_n) = 2\,\mathrm{at}(P_n) = 2n$.*

Proof. For every positive integer m, let S_m be the poset consisting of minimal elements $a_1, \ldots, a_m$ and maximal elements $b_1, \ldots, b_m$ so that for all $i, j \in [m]$, we have $a_i < b_j$ if and only if $i \neq j$. In other words, S_m is the *standard example* of order m. Standard examples are the canonical posets forcing dimension to be large, and they date back to the foundational paper of Dushnik and Miller [5] that initiated the study of dimension of posets. It is well-known that the dimension of the standard example of order m is exactly m.

We show that the poset S_{2n} witnesses the statement for all positive integers n. We have $\dim(S_{2n}) = 2n$. It suffices to show that $n \leq \mathrm{at}(S_{2n})$. We claim that there is no triple $i, j, k \in [2n]$ such that a_i, a_j, a_k are all tight on a single map of S_n. Indeed, the subposet of S_n formed by $a_i, a_j, a_k, b_i, b_j, b_k$ is isomorphic to the poset Q, hence, the claim follows from Lemma 11. This implies that $n \leq \mathrm{at}(S_{2n})$, as desired. ◀

▶ **Theorem 3.** *There exists a sequence of posets $(P_n)_{n \geq 1}$ such that* $\mathrm{at}(P_n) = 2^{2^{\Omega(\dim(P_n))}}$.

Proof. For every positive integer m, let P_m be the incidence poset of the complete graph on m vertices, K_m. The dimension of such posets is well understood. In fact, in the 1970's, Spencer proved that $\dim(P_m) = \mathcal{O}(\log \log m)$, see [13, p.580] and [9] for more precise estimates of $\dim(P_m)$. To complete the proof, we show that for every positive integer m, $\mathrm{at}(P_{2m}) \geq m$.

Let m be a positive integer. We claim that there is no triple u, v, w of the vertices of K_m such that u, v, w are all tight on a single map of P_{2m}. Indeed, the subposet of P_{2m} formed by u, v, w, uv, vw, uw is isomorphic to the poset Q. Hence, the claim follows from Lemma 11. This implies that $m \leq \mathrm{at}(P_{2m})$, as desired. ◀

4 Complexity

In this section, we prove that it is NP-complete to compute the atlas thickness or the dominance mapability of a given poset (Theorems 4 and 5). The main idea is to consider incidence posets of graphs and to show that a subset of tight vertices in such a poset corresponds to a linear forest in this graph (see Lemmas 13 and 14).

In the case of the atlas thickness of a poset (Theorem 4), we will reduce from the problem of partitioning the vertex set of a graph into two sets such that each induces a linear forest. The NP-hardness of this problem was observed by Chaplick, Fleszar, Lipp, Ravsky, Verbitsky, and Wolff [4, Theorem 12]. It follows from a general result of Farrugia [6], who showed that, for any additive induced-hereditary properties Π, it is NP-hard to decide whether the vertex set of a graph can be partitioned into two sets that have property Π.

In the case of the dominance mapability of a poset (Theorem 5), we will reduce from the problem of finding a subset of k vertices of a given graph that induces a linear forest. The NP-hardness of this problem follows from another general result proved by Lewis and Yannakakis [10], who showed that removing the minimum number of vertices from a graph G such that the remaining induced subgraph has a property Π is NP-hard if Π is non-trivial and hereditary.

The crucial ingredient for the proof is the well-known characterization of incidence posets of dimension at most 2, which is stated in the following proposition. We provide a direct proof, which will help us to show Lemma 13.

▶ **Proposition 12.** *The dimension of the incidence poset P_G of a graph G is at most 2 if and only if G is a linear forest.*

Proof. The incidence order of a linear forest is an induced subposet of the incidence poset of a path, hence, it is sufficient to show that the incidence poset of a path is 2-dimensional. Let G be a path with vertices $a_1, \ldots, a_{\ell+1}$ and edges $e_i = a_i a_{i+1}$ for $i = 1, \ldots, \ell$, then the following two linear extensions of P_G are a realizer of size two: $K = a_1, a_2, e_1, a_3, e_2, \ldots, a_{\ell+1}, e_\ell$ and $L = a_{\ell+1}, a_\ell, e_\ell, \ldots, a_2, e_2, a_1, e_1$.

If G is not a forest, then G contains a cycle C. It is well known that the dimension of P_C is 3, hence $\dim(P_G) \geq 3$ in this case. If G is a forest but not a linear forest, then G contains a 3-star. The incidence poset of a 3-star is a poset on 7 elements called *spider*. It is well known that the dimension of the spider is 3, hence again $\dim(P_G) \geq 3$. Cycle posets and the spider are indeed 3-irreducible posets, i.e., they are 3-dimensional and all their subposets are 2-dimensional, cf. [13, p. 573] and the references given there. ◀

▶ **Lemma 13.** *Let G be a graph and let $U \subseteq V(G)$. If $G[U]$ is a linear forest, then there exists a map of P_G such that the elements in $U \cup E(G)$ are tight.*

Proof. In the proof of Proposition 12, we have shown how to construct two linear orders K and L that form a realizer of $P_{G[U]}$. Edges of G that are incident to just one element u in U can be placed in K and L just above u. All the remaining elements of P_G can be put at the beginning of K and at the end of L. Figure 3 illustrates the construction. ◀

▶ **Lemma 14.** *Let G be a graph, let $U \subseteq V(G)$, and let μ be a map of P_G. If all the elements of U are tight on μ, then $G[U]$ is a linear forest.*

Proof. This follows directly from Proposition 12. ◀

▶ **Theorem 4.** *It is NP-complete to test whether the atlas thickness of a poset is at most 2.*

Proof. The problem is clearly in NP. As announced at the beginning of this section, we infer the NP-hardness by a reduction from the problem of partitioning the vertex set of a graph into two sets such that each induces a linear forest, which is NP-hard [4, Theorem 12]. We show that the vertices of a graph G can be partitioned into two sets inducing linear forests if and only if $\mathrm{at}(G) \leq 2$.

First, given such a partition U_1, U_2 of $V(G)$, Lemma 13 implies that there are two maps μ_1 and μ_2 of P_G such that $U_i \cup E(G)$ is tight in μ_i for each $i \in [2]$. It follows that $\mathrm{at}(P_G) \leq 2$.

On the other hand, suppose that we are given two maps μ_1 and μ_2 of P_G such that every element of P_G is tight either on μ_1 or on μ_2. By Lemma 14, the vertices of G that are tight on μ_1 induce a linear forest in G, and so do the vertices that are tight on μ_2. This represents the required partition. ◀

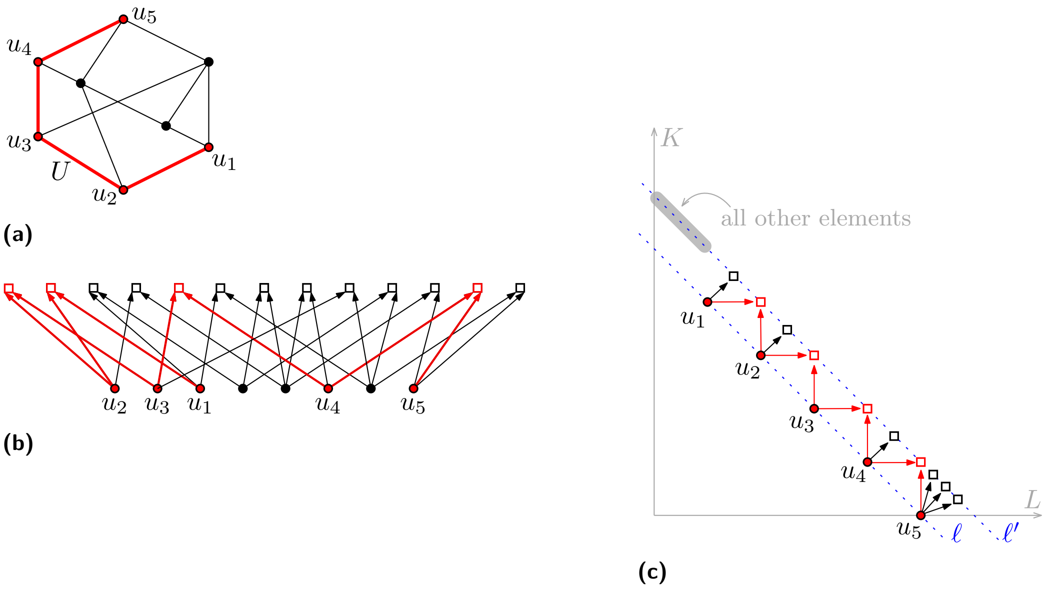

Figure 3 Illustration for the proof of Lemma 13: (a) a graph G with a path $G[U]$ highlighted in red; (b) the incidence poset P_G; (c) the map μ. All the vertices of U are placed on the line ℓ in an order compatible with the order of U. The edges of type u_iu_{i+1} (red squares) are placed on ℓ' between u_i and u_{i+1}. Each other edge (black square) is placed on ℓ' between the red squares $u_{i-1}u_i$ and u_iu_{i+1} if it is incident to u_i Finally, all the other elements are placed on the side (here at the top of ℓ').

▶ **Theorem 5.** *It is* NP*-complete to compute the mapability of a poset.*

Proof. The problem is clearly in NP. As announced at the beginning of this section, we infer NP-hardness by a reduction from the NP-hard problem of finding a subset of k vertices of a given graph that induces a linear forest. We show that a graph G contains a subset of at most k vertices that induces a linear forest in G if and only if $\mathrm{dmap}(P_G) \le k + |E(G)|$.

Given such a subset, Lemma 13 immediately implies that there exists a map of P_G where at least $k + |E(G)|$ elements are tight, and so, $\mathrm{dmap}(P_G) \le k + |E(G)|$.

On the other hand, suppose that we are given a map of P_G where at least $k + |E(G)|$ elements are tight. At least k of these elements are in $V(G)$ by the pigeonhole principle. It follows by Lemma 14 that this subset of vertices induces a linear forest in G, as desired. ◀

5 An FPT-Algorithm for Dominance Mapability

In this section, we prove Theorem 6, which implies that, given a poset P, computing its mapability $\mathrm{dmap}(P)$ is FPT with respect to the natural parameter (see Corollary 7). Namely, we will upperbound the size of a poset in terms of its mapability. The next statement is a direct corollary of Lemma 10 applied to a maximum chain and a singleton.

▶ **Corollary 15.** *For every poset P of height h that is not a chain, we have $h \le \mathrm{dmap}(P) - 1$.*

A subset M of P is an *up-module* if any two elements of M have the same successors, i.e., $\mathrm{succ}_P(u) = \mathrm{succ}_P(v)$ for all $u, v \in M$. Note that an up-module must be an antichain. If M is an up-module, then we write $\mathrm{succ}_P(M)$ to denote the set of successors of the elements of M.

▶ **Lemma 16.** *Let P be a poset. If M is an up-module of P, then $|M| \le \mathrm{dmap}(P)$.*

Proof. Let μ be a map where some element $v \in M$ is tight. By Lemma 8 we can assume that the coordinate projections of μ are linear extensions of P. We can also assume that all coordinates are integers. Let ℓ the line of slope -1 through $\mu(v)$ and let s be a segment of length 1 on ℓ which contains $\mu(v)$. Note that our assumption about μ imply that $\mathrm{succ}_P(M) = \mathrm{succ}_P(v) = \{w \in P : x <_{\mathrm{dom}} \mu(w)\}$ for every point x on s. Hence, if we place the elements of M at distinct points of s we create a map where all elements of M are tight. ◀

Observe that the set $\mathrm{Max}(P)$ is an up-module. Hence, Lemma 16 implies the following statement.

▶ **Corollary 17.** *For every poset P, we have $|\mathrm{Max}(P)| \leq \mathrm{dmap}(P)$.*

Given two positive integers a and k, ${}^k a$ is a tower of a's of height k. For example, ${}^3 2 = 2^{2^2} = 16$. More formally, ${}^1 a = a$ and, for any positive integer i, ${}^i a = a^{({}^{i-1}a)}$. We extend this notation and write ${}^k_b a$ to denote a tower of total height k where the top element is b and the remaining $(k-1)$ elements are a's. For example, ${}^3_4 2 = 2^{2^4} = 2^{16}$.

We prove the following theorem with respect to the function $f(d) = {}^d_d 4$.

▶ **Theorem 6.** *There exists a computable function f such that, for every poset P, we have $|P| \leq f(\mathrm{dmap}(P))$.*

Proof. Let h be the height of P, and let $d = \mathrm{dmap}(P)$. The *canonical antichain partition* of P is the partition of the elements of P into h antichains. Deviating from the standard definition we construct the canonical antichain partition top-down instead of bottom-up. Let $A_1 = \mathrm{Max}(P)$ and, for each $k \in [h] \setminus \{1\}$, let A_k be the set of maximal elements of the induced poset $P \setminus \bigcup_{j=1}^{k-1} A_j$. We will identify a function f such that, for every $k \in [h]$, $|A_k| \leq f(k)$. The function f will grow very fast. In particular, for every $k \geq 2$, we will have

$$f(k) \geq 3 \sum_{j=1}^{k-1} f(j) \quad \text{and} \quad d \leq 2^{(2/3) f(k-1)}. \tag{$\star$}$$

Since $A_1 = \mathrm{Max}(P)$, we obtain from Corollary 17 that $|A_1| \leq d$. We set $f(1) = d$. Note that this yields $|A_1| \leq f(1)$ and fulfills the right inequality in ($\star$) for every $k \geq 2$. (For $k = 2$, this is due to the fact that $x \leq 2^{2x/3}$ holds for every $x > 0$. For $k > 2$, this is due to the fact that f grows.)

Now consider $k \in [h] \setminus \{1\}$, and let $\mathcal{M}$ be a partition of A_k into maximal up-modules. By Lemma 16, for every $M \in \mathcal{M}$, we have $|M| \leq d$. In particular, $|A_k| \leq d \cdot |\mathcal{M}|$. A module $M \in \mathcal{M}$ is determined by $\mathrm{succ}_P(M)$, which is a subset of $\bigcup_{j=1}^{k-1} A_j$. Due to ($\star$), i.e., the fast growth of f, we have

$$\left| \bigcup_{j=1}^{k-1} A_j \right| \leq \sum_{j=1}^{k-1} f(j) = f(k-1) + \sum_{j=1}^{k-2} f(j) \leq \frac{4}{3} f(k-1).$$

This implies that there are at most $2^{(4/3)f(k-1)}$ choices for $\mathrm{succ}_P(M)$. Again using ($\star$), we get

$$|A_k| \leq d \cdot 2^{(4/3)f(k-1)} \leq 2^{(2/3)f(k-1)} 2^{(4/3)f(k-1)} = 2^{2f(k-1)} = 4^{f(k-1)}.$$

If we now set $f(k) = {}^k_d 4$, then ($\star$) is fulfilled for $k \geq 2$, and $|A_k| \leq f(k)$ for $k \in [h]$ (including $k = 1$). Using ($\star$) and Corollary 15, we finally get

$$|P| \leq \sum_{j=1}^{h} f(j) \leq \frac{1}{3} f(h+1) \leq f(d) = {}^d_d 4.$$ ◀

6 Open Problems and Future Directions

We have introduced the new concepts of mappability and atlas thickness, and we conclude by highlighting several open problems for future investigation.

1. If we embed into $\mathbb{R}$ instead of $\mathbb{R}^2$, we just get the width of the poset. Is it interesting to consider higher dimensions?
2. We have proved that, for every poset P, $\dim(P) \leq 2\,\mathrm{at}(P)$ and that there exist posets (incidence posets of complete graphs) with the atlas thickness huge relative to their dimension. However, we did not find a family of posets with constant dimension and unbounded atlas thickness. On the other hand, it is well-known that the Boolean dimension of incidence posets of complete graphs is bounded, thus our family of posets *separates* Boolean dimension and atlas thickness. It would be interesting to understand if dimension and atlas thickness are functionally equivalent or not.
3. The following class of parameters of posets may be worth studying: For a given poset P and positive integer k, find the minimum $t(k)$ such that there exists a family $\mathcal{R}$ of linear extensions with:
   - the size of $\mathcal{R}$ is at most $k \cdot \dim(P)$,
   - for each v in P, there is a subfamily $R(v) \subseteq \mathcal{R}$ of size at most $t(k)$ such that $\mathrm{succ}_P[v] = \{w \in P \colon v \leq_L w \text{ for every } L \in R(v)\}$ (we could say that v is tight in $R(v)$).

   The motivation for the definition is that we obtain a data structure that allows us to answer comparison queries ($<$, $\|$, $>$) with space requirement $O(k \cdot \dim(P) \cdot n)$ and query time $O(t(k))$.

References

1. Todor Antić, Giuseppe Liotta, Tomáš Masařík, Giacomo Ortali, Matthias Pfretzschner, Peter Stumpf, Alexander Wolff, and Johannes Zink. Unbent collections of orthogonal drawings. In Henning Fernau and Philipp Kindermann, editors, *Graph-Theoretic Concepts Comput. Sci. (WG)*, volume 16124 of *LNCS*, pages 18–32. Springer, 2025. doi:10.1007/978-3-032-11835-6_2.
2. Giuseppe Di Battista, Walter Didimo, Luca Grilli, Fabrizio Grosso, Giacomo Ortali, Maurizio Patrignani, and Alessandra Tappini. Small point-sets supporting graph stories. *J. Graph Algorithms Appl.*, 27(8):651–677, 2023. doi:10.7155/JGAA.00639.
3. Carla Binucci, Emilio Di Giacomo, William J. Lenhart, Giuseppe Liotta, Fabrizio Montecchiani, Martin Nöllenburg, and Antonios Symvonis. On the complexity of the storyplan problem. *J. Comput. Syst. Sci.*, 139:103466, 2024. doi:10.1016/J.JCSS.2023.103466.
4. Steven Chaplick, Krzysztof Fleszar, Fabian Lipp, Alexander Ravsky, Oleg Verbitsky, and Alexander Wolff. The complexity of drawing graphs on few lines and few planes. *J. Graph Algorithms Appl.*, 27(6):459–488, 2023. doi:10.7155/jgaa.00630.
5. Ben Dushnik and Edwin W. Miller. Partially ordered sets. *Amer. J. Math.*, 63(3):600–610, 1941. doi:10.2307/2371374.
6. Alastair Farrugia. Vertex-partitioning into fixed additive induced-hereditary properties is NP-hard. *Electr. J. Comb.*, 11(1):R46:1–9, 2004. doi:10.37236/1799.
7. Jiří Fiala, Oksana Firman, Giuseppe Liotta, Alexander Wolff, and Johannes Zink. Outerplanar and forest storyplans. In Henning Fernau, Serge Gaspers, and Ralf Klasing, editors, *Theory & Practice Comput. Sci. (SOFSEM)*, volume 14519 of *LNCS*, pages 211–225. Springer, 2024. doi:10.1007/978-3-031-52113-3_15.

8 Petr Hliněný and Tomáš Masařík. Minimizing an uncrossed collection of drawings. In Michael Bekos and Markus Chimani, editors, *Graph Drawing & Network Vis. (GD)*, volume 14466 of *LNCS*, pages 110–123. Springer, 2023. URL: https://arxiv.org/abs/2306.09550, doi:10.1007/978-3-031-49272-3_8.

9 Serkan Hoşten and Walter D. Morris. The order dimension of the complete graph. *Discrete Math.*, 201(1):133–139, 1999. doi:10.1016/S0012-365X(98)00315-X.

10 John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is NP-complete. *J. Comput. Syst. Sci.*, 20(2):219–230, 1980. doi:10.1016/0022-0000(80)90060-4.

11 Edward M. McCreight. Priority search trees. *SIAM J. Comput.*, 14(2):257–276, 1985. doi:10.1137/0214021.

12 Jaroslav Nešetřil and Pavel Pudlák. A note on Boolean dimension of posets. In *Irregularities of partitions (Fertőd, 1986)*, volume 8 of *Algorithms Combin. Study Res. Texts*, pages 137–140. Springer, 1989. doi:10.1007/978-3-642-61324-1_12.

13 Douglas B. West. *Combinatorial Mathematics*. Cambridge University Press, 2021.

14 Mihalis Yannakakis. The complexity of the partial order dimension problem. *SIAM J. Algebr. Discrete Meth.*, 3(3):351–358, 1982. doi:10.1137/0603036.