Permutation Pattern Matching for Doubly Partially Ordered Patterns

Laurent Bulteau
LIGM, CNRS, Univ Gustave Eiffel, F77454 Marne-la-Vallée, France

Guillaume Fertin
Nantes Université, LS2N (UMR 6004), CNRS, Nantes, France

Vincent Jugé
LIGM, CNRS, Univ Gustave Eiffel, F77454 Marne-la-Vallée, France

Stéphane Vialette
LIGM, CNRS, Univ Gustave Eiffel, F77454 Marne-la-Vallée, France

Abstract

We study in this paper the Doubly Partially Ordered Pattern Matching (or DPOP Matching) problem, a natural extension of the Permutation Pattern Matching problem. Permutation Pattern Matching takes as input two permutations σ and π, and asks whether there exists an occurrence of σ in π; whereas DPOP Matching takes *two partial orders* P_V and P_P defined on the same set X and a permutation π, and asks whether there exist $|X|$ elements in π whose values (resp., positions) are in accordance with P_V (resp., P_P). Posets P_V and P_P aim at relaxing the conditions formerly imposed by the permutation σ, since σ yields a total order on both positions and values. Our problem being NP-hard in general (as Permutation Pattern Matching is), we consider restrictions on several parameters/properties of the input, e.g., bounding the size of the pattern, assuming symmetry of the posets (i.e., P_V and P_P are identical), assuming that one partial order is a total (resp., weak) order, bounding the length of the longest chain/anti-chain in the posets, or forbidding specific patterns in π. For each such restriction, we provide results which together give a(n almost) complete landscape for the algorithmic complexity of the problem.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Partial orders, Permutations, Pattern Matching, Algorithmic Complexity, Parameterized Complexity

Digital Object Identifier 10.4230/LIPIcs.CPM.2022.21

1 Preamble

Let us play the following little puzzle game. Among the selection of fifteen cities of the Czech Republic depicted in Figure 1 together with their geographic coordinates, find (if they exist) five cities, say **A**, **B**, **C**, **D** and **E**, such that:

- **A** and **C** are west of **D** and north of **B**,
- **E** is east of **B** and south of **A**,
- **D** is west of **B** and north of **A** and **C**.

It is assumed that no two cities have the same longitude or latitude. Notice that the game does not provide complete information as, for example, no information is provided about the relative positioning of **A** and **C** (and silence is tantamount to consent). We may assume that the information is minimal: requiring **C** is west of **B** is unnecessary since **C** is west of **D** and **D** is west of **B**. One solution is **A** = Praha, **B** = Brno, **C** = Plzeň, **D** = Liberec and **E** = Olomouc. Note that the solution is not unique, as **A** = Plzeň, **B** = Jindřichův Hradec, **C** = Cheb, **D** = Ústí nad Labem and **E** = Brno is another solution.

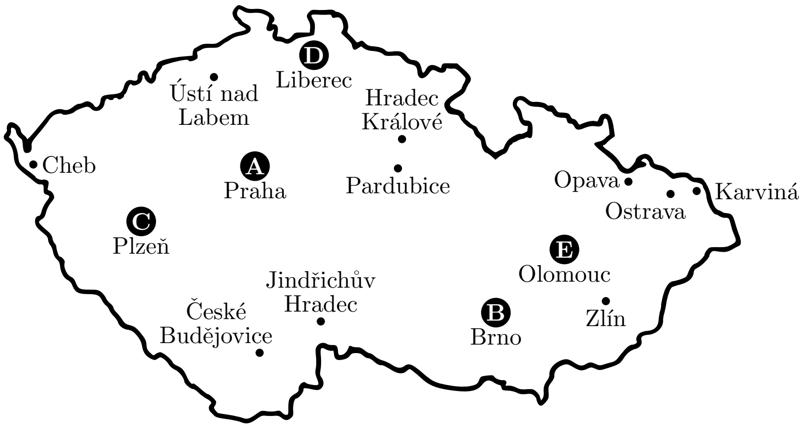

Figure 1 A Czech Republic map showing 15 of its cities. Their names and GPS coordinates, in increasing order of their longitudes, are: Cheb (50°4′N 12°22′E), Plzeň (49°44′N 13°22′E), Ústí nad Labem (50°39′N 14°1′E), Praha (50°5′N 14°25′E), České Budějovice (48°58′N 14°28′E), Jindřichův Hradec (49°9′N 15°0′E), Liberec (50°46′N 15°3′E), Pardubice (50°2′N 15°46′E), Hradec Králové (50°12′N 15°49′E), Brno (49°11′N 16°36′E), Olomouc (49°35′N 17°15′E), Zlín (49°13′N 17°40′), Opava (49°56′N 17°54′E), Ostrava (49°50′N 18°17′E) and Karviná (49°51′N 18°32′E).

We show that this puzzle game can be modeled as a permutation pattern matching problem for doubly partially ordered patterns. Let us first associate a permutation $\pi \in \mathfrak{S}(15)$ with the problem (see Figure 2). We sort the fifteen cities of the Czech Republic depicted in Figure 1 both by increasing longitude (E) and by increasing latitude (N), so that $\pi(i) = j$ if the i-th city going west to east is also the j-th city going south to north. In our example, the "Czech Republic permutation" is $\pi = 11\ 6\ 14\ 12\ 1\ 2\ 15\ 10\ 13\ 3\ 5\ 4\ 9\ 7\ 8$. For example, $\pi(2) = 6$ since Plzeň is the second city going west to east, and the sixth city going south to north. What is left is to define our pattern P: P is composed of two partially ordered sets on the variables $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}\}$ (see Figure 3): one partially ordered set (denoted P_{V} for *value poset* in the sequel) describes the south-to-north constraints and another partially ordered set (denoted P_{P} for *position poset* in the sequel) describes the west-to-east constraints.

2 Introduction

We say that a permutation σ *occurs* in another permutation π (or that π *contains* σ) if there exists a subsequence of elements of π that has the same relative order as σ. Otherwise, we say that π *avoids* σ. For example, π contains the permutation $\sigma = 123$ (resp., $\sigma = 321$) if it has an increasing (resp., a decreasing) subsequence of size 3. Similarly, $\sigma = 4312$ occurs in $\pi = 6152347$, as shown in 6152347, but the same $\pi = 6152347$ avoids $\sigma' = 2341$.

Deciding whether a permutation $\sigma \in \mathfrak{S}(k)$ occurs in some permutation $\pi \in \mathfrak{S}(n)$ is NP-complete [7], but is fixed-parameter tractable for the parameter k [15, 17]. Several exponential-time algorithms have been recently proposed [5, 16], improving upon [1, 10]. A vast literature is devoted to the case where both the pattern σ and the target π are restricted to a proper permutation class, e.g., 321-avoiding permutations [18, 2, 21], $(213, 231)$-avoiding permutations [26], $(2413, 3142)$-avoiding (a.k.a. *separable*) permutations [19, 25], and $(k \dots 1)$-avoiding permutations [11]. For more background on permutation patterns and pattern avoidance, we refer to [6] and [24].

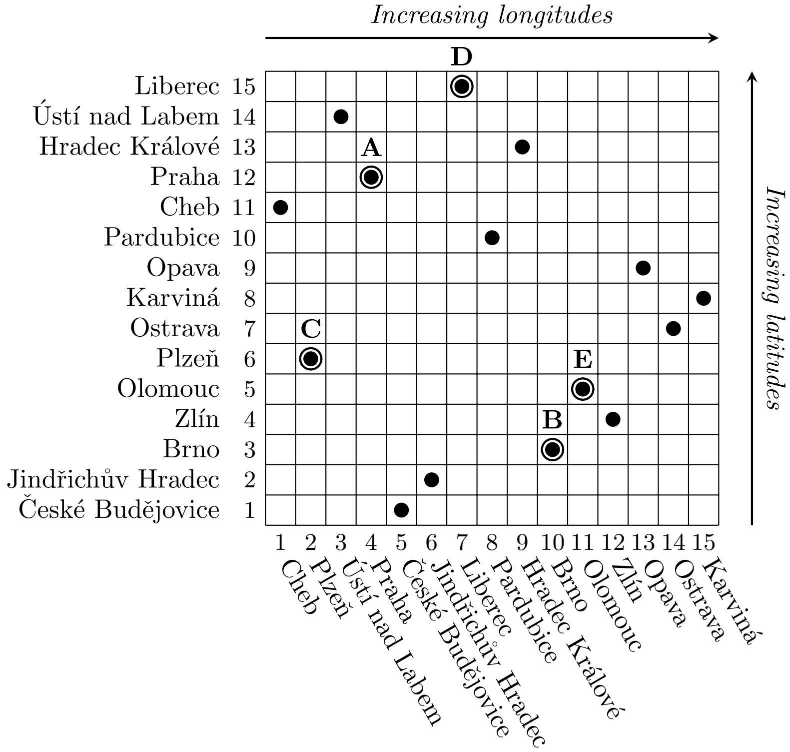

Figure 2 Permutation $\pi = 11\ 6\ 14\ 12\ 1\ 2\ 15\ 10\ 13\ 3\ 5\ 4\ 9\ 7\ 8$ corresponding to the map from Figure 1. The solution of our puzzle, depicted Figure 1, is also represented.

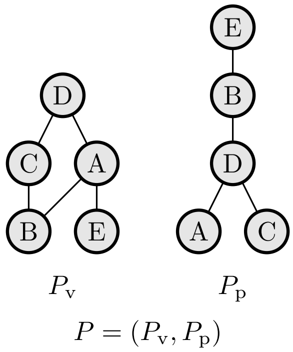

	Cheb	Plzeň	Ústí nad Labem	Praha	České Budějovice	Jindřichův Hradec	Liberec	Pardubice	Hradec Králové	Brno	Olomouc	Zlín	Opava	Ostrava	Karviná
$\pi =$	11	6	14	12	1	2	15	10	13	3	5	4	9	7	8
	.	**C**	.	**A**	.	.	**D**	.	.	**B**	**E**	.	.	.	.
	A	**C**	**D**	.	.	**B**	.	.	.	**E**	.	.	.	.	.
	C	.	.	**A**	.	.	.	.	**D**	.	**B**	**E**	.	.	.

Figure 3 A dpop $P = (P_{\mathrm{V}}, P_{\mathrm{P}})$ representing the pattern for our puzzle, together with three distinct occurrences. Note that partial orders are represented by Hasse diagrams, i.e., a bottom-up path in P_{V} (resp., P_{P}) implies a bottom-up (resp., left-right) relation in the occurrence of P in π.

In the last years, the notion of *pattern* has been generalized in several ways. A *vincular pattern* is a permutation in which some entries must occur consecutively [4]. *Consecutive patterns* are a special case of vincular patterns in which all entries need to be adjacent [13]. *Bivincular patterns* generalize classical patterns even further than vincular patterns by requiring that not only positions but also values of elements involved in a matching may be forced to be adjacent [8]. *Mesh patterns* (a further generalization of bivincular patterns) impose further restrictions on the relative positions of the entries in an occurrence of a pattern [9] and *boxed mesh patterns* are special cases of mesh patterns [3]. Strongly related to our approach are *partially ordered patterns* that are vincular patterns in which the relative order of some elements is not fixed [23]. The best general reference is [24].

In this paper, we consider a new generalization of classical patterns in which both the relative order and the relative positioning of some elements are not fixed. The idea is to allow the possibility for some elements to be incomparable in value (i.e., their relative order is unknown) and to go one step further by allowing the possibility for some elements to be incomparable in position (i.e., their relative positioning in the occurrence is unknown). Since the problem is clearly NP-hard (as it contains Permutation Pattern Matching as a sub-problem), our goal is to identify tractable cases when restrictions apply to the pattern and/or to the permutation.

The restrictions we consider here apply to the following parameters of the problem: size of the pattern; symmetry (i.e., same partial order in positions and values); one partial order is a total (resp., weak) order; size of the longest chain (resp., anti-chain) in the partial orders (height and width); forbidden patterns in π. On the positive side, we show that the FPT algorithm for Permutation Pattern Matching can be generalized to our setting (with the pattern size as parameter). We further give polynomial-time algorithms when the pattern is a symmetric disjoint union of a constant number of weak orders. Finally, we also provide polynomial-time algorithms when the pattern is symmetric and the permutation belongs to some restricted classes, such as $(123, 132)$-avoiding permutations. We complement these positive results with NP- or W[1]-hardness proofs in most of the remaining cases.

3 Definitions

Permutations and Patterns

A permutation σ is said to be *contained* in (or is a *sub-permutation* of) another permutation π, which we denote by $\sigma \preceq \pi$, if π has a (not necessarily contiguous) subsequence whose terms are order-isomorphic to σ. We also say that π *admits an occurrence* of the *pattern* σ. If no such subsequence exists, we say that π *avoids* σ (or is σ*-avoiding*). A permutation is *separable* if it avoids both 2413 and 3142. Permutation Pattern Matching is the problem of deciding whether a permutation is contained into another permutation.

For any non-negative integer n, we denote by $[n]$ the set $\{1, 2, \ldots, n\}$. When $n \geqslant 1$, we also note $\mathrm{ip}(n) = 1\,2\,\ldots\,n$ the *increasing permutation* of length n and $\mathrm{dp}(n) = n\,(n-1)\,\ldots\,1$ the *decreasing permutation* of length n. Let $\pi \in \mathfrak{S}(n)$. The *reverse* (resp., *complement*) of π is the permutation $\pi^{\mathrm{r}} = \pi(n)\pi(n-1)\ldots\pi(1)$ (resp., $\pi^{\mathrm{c}} = (n-\pi(1)+1)(n-\pi(2)+1)\ldots(n-\pi(n)+1)$). The *inverse* of π is the permutation $\pi^{-1} \in \mathfrak{S}(n)$ defined by $\pi^{-1}(j) = i$ if and only if $\pi(i) = j$. Given a permutation π of size m and a permutation σ of size n, the *skew sum* of π and σ is the permutation of size $m + n$ defined by

$$(\pi \ominus \sigma)(i) = \begin{cases} \pi(i) + n & \text{for } 1 \leqslant i \leqslant m, \\ \sigma(i - m) & \text{for } m + 1 \leqslant i \leqslant m + n, \end{cases}$$

and the *direct sum* of π and σ is the permutation of size $m + n$ defined by

$$(\pi \oplus \sigma)(i) = \begin{cases} \pi(i) & \text{for } 1 \leqslant i \leqslant m, \\ \sigma(i - m) + m & \text{for } m + 1 \leqslant i \leqslant m + n. \end{cases}$$

Orders

A relation $\leqslant$ is a *partial order* on a set X if it has:

- *reflexivity*: for all $x \in X$, $x \leqslant x$ (i.e., every element is related to itself);
- *transitivity*: for all $x, x', x'' \in X$, if $x \leqslant x'$ and $x' \leqslant x''$, then $x \leqslant x''$;
- *antisymmetry*: for all $x, x' \in X$, if $x \leqslant x'$ and $x' \leqslant x$, then $x = x'$ (i.e., no two distinct elements precede each other).

If $\leqslant$ has the following additional property, we say that it is a *weak* order on X:

- *transitivity of incomparability*: for all pairwise distinct $x, x', x'' \in X$, if x is incomparable with x' (i.e., neither $x \leqslant x'$ nor $x' \leqslant x$ is true) and if x' is incomparable with x'', then x is incomparable with x''.

Two subsets X_1, X_2 are *independent* if there is no $x_1 \in X_1$, $x_2 \in X_2$ such that $x_1 \leqslant x_2$ or $x_2 \leqslant x_1$. We say that a partial order is k*-weak* if there exists a partition of X into k pairwise independent sets $X_1, \ldots, X_k$ such that, for each i, the restriction of $\leqslant$ to X_i is a weak order (in other words, $\leqslant$ is the disjoint union of k weak partial orders).

Let $\mathcal{P} = (X, \leqslant)$ be a finite partially ordered set. A *chain* in $\mathcal{P}$ is a set of pairwise comparable elements (i.e., a totally ordered subset) and an *antichain* in $\mathcal{P}$ is a set of pairwise incomparable elements. The *partial order height* of $\mathcal{P}$, denoted by $\text{height}(\mathcal{P})$, is defined as the maximum cardinality of a chain in $\mathcal{P}$, and the *partial order width* of $\mathcal{P}$, denoted by $\text{width}(\mathcal{P})$, is defined as the maximum cardinality of an antichain in $\mathcal{P}$. By Dilworth Theorem, $\text{width}(\mathcal{P})$ is also the minimum number of chains in any partition of $\mathcal{P}$ into chains. The *dual* of $\mathcal{P}$ is the partial order $\mathcal{P}^\partial = (X, \leqslant^\partial)$ defined by letting $\leqslant^\partial$ be the converse relation of $\leqslant$, i.e., $x \leqslant^\partial x'$ if and only if $x' \leqslant x$. The dual of a partial order is a partial order and the dual of the dual of a relation is the original relation. A *total order* is a partial order in which any two elements are comparable, and a set equipped with a total order is a *totally ordered set*. A *linear extension* of a partial order is a total order that is compatible with the partial order. It will be convenient to represent a linear extension of a poset $\mathcal{P} = (X, \leqslant)$ as the mapping $\tau_\mathcal{P} : X \to [\,|X|\,]$ such that $\tau_\mathcal{P}(i) < \tau_\mathcal{P}(j)$ if $i < j$ in the linear extension.

A *doubly partially ordered pattern* (dpop) P is a pair, denoted by $P = (P_\text{v}, P_\text{p})$, of posets $P_\text{v} = (X, \leqslant_\text{v})$ and $P_\text{p} = (X, \leqslant_\text{p})$ defined over the same set X. We call P_v and P_p the *value poset* and the *position poset*, respectively. A dpop $P = (P_\text{v}, P_\text{p})$ is *symmetric* if $P_\text{v} = P_\text{p}$, *dual* if $P_\text{v} = {P_\text{p}}^\partial$, and *semi-total* if one of P_p or P_v is a total order. We let $\text{height}(P)$ and $\text{width}(P)$ stand for $\max\{\text{height}(P_\text{v}), \text{height}(P_\text{p})\}$ and $\max\{\text{width}(P_\text{v}), \text{width}(P_\text{p})\}$, respectively. Finally, the *size* of P is defined as the cardinality $|X|$ and is denoted by $|P|$.

▶ **Definition 1** (DPOP Matching). *Given a permutation $\pi \in \mathfrak{S}(n)$ and a dpop $P = (P_\text{v}, P_\text{p})$, an* occurrence *(or* mapping*) of P in π is an injective function $\varphi : X \to [n]$ such that:*

- *$\pi \circ \varphi$ is $\leqslant_\text{v}$-non-decreasing, i.e., for all $x, y \in X$, if $x \leqslant_\text{v} y$ then $\pi(\varphi(x)) \leqslant \pi(\varphi(y))$, and*
- *φ is $\leqslant_\text{p}$-non-decreasing, i.e., for all $x, y \in X$, if $x \leqslant_\text{p} y$ then $\varphi(x) \leqslant \varphi(y)$.*

The DPOP Matching problem consists in deciding whether P occurs in π.

First Observations

▶ **Observation 2.** *Permutation Pattern Matching is the special case of DPOP Matching where both $\leqslant_\text{v}$ and $\leqslant_\text{p}$ are total orders.*

We note that applying a vertical and/or horizontal symmetry on both pattern and permutation does not alter the existence of an occurrence.

▶ **Observation 3.** *Let $P = (P_\mathrm{v}, P_\mathrm{p})$ be a dpop and π be a permutation. The following statements are equivalent:*

1. $(P_\mathrm{v}, P_\mathrm{p})$ *occurs in* π*;*
2. $(P_\mathrm{v}, (P_\mathrm{p})^\partial)$ *occurs in* π^r*;*
3. $((P_\mathrm{v})^\partial, P_\mathrm{p})$ *occurs in* π^c*;*
4. $((P_\mathrm{v})^\partial, (P_\mathrm{p})^\partial)$ *occurs in* $\pi^{\mathrm{c\,r}}$*;*
5. $(P_\mathrm{p}, P_\mathrm{v})$ *occurs in* π^{-1}*.*

The following reformulation will prove useful.

▶ **Observation 4.** *Let $P = (P_\mathrm{v}, P_\mathrm{p})$ be a dpop with $P_\mathrm{v} = (X, \leqslant_\mathrm{v})$, $P_\mathrm{p} = (X, \leqslant_\mathrm{p})$ and $k = |X|$, and let $\pi \in \mathfrak{S}(n)$ be a permutation. The following statements are equivalent:*

- *P occurs in π.*
- *There exists a linear extension $\tau_\mathrm{v} : X \to [k]$ of P_v and a linear extension $\tau_\mathrm{p} : X \to [k]$ of P_p such that the permutation $\sigma \in \mathfrak{S}(k)$ defined by $\sigma(i) = \tau_\mathrm{v}(\tau_\mathrm{p}^{-1}(i))$ for $1 \leqslant i \leqslant k$ is contained in π.*

The rationale for the reformulation introduced in Observation 4 stems from the following corollary that sets the general context.

▶ **Corollary 5** ([17])**.** *DPOP Matching is* FPT *for the parameter $|P|$.*

Indeed, it is enough to guess two linear extensions $\tau_\mathrm{v} : X \to [k]$ of P_v and $\tau_\mathrm{p} : X \to [k]$ of P_p, and to check if the permutation $\sigma \in \mathfrak{S}(k)$ defined by $\sigma(i) = \tau_\mathrm{v}(\tau_\mathrm{p}^{-1}(i))$ for $1 \leqslant i \leqslant k$ is contained in π. There are $O(k!^2)$ pairs of such extensions and, for each of them, one can check whether σ occurs in π in $n\,2^{O(k^2 \log k)}$ time [17].

4 Semi-Total Patterns

In this section we focus on semi-total patterns, i.e., without loss of generality, on the case where P_p is a total order (up to symmetry by Observation 3). This case still contains Permutation Pattern Matching as a special case, and is thus NP-hard. We focus on small-height value partial orders, i.e., on dpops with constant height(P_v), and give an XP algorithm for weak orders (Proposition 6) and paraNP-hardness in general (Proposition 7).

▶ **Proposition 6.** *DPOP Matching is solvable in $O\left(n^{\mathrm{height}(P_\mathrm{v})}\right)$ time if P_v is a weak order and P_p is a total order.*

Proof. Let $\pi \in \mathfrak{S}(n)$ be a permutation and $P = (P_\mathrm{v}, P_\mathrm{p})$ be a dpop on some ground set X, where P_v is a weak order and P_p is a total order. Without loss of generality, we assume that X is the set $[k]$ and that P_p is the usual order on integers. For every $x \in X$, we abusively denote by height(x) the maximum cardinality of a chain with maximum element x in P_v. Finally, set $\ell = \mathrm{height}(P_\mathrm{v})$.

For any two distinct variables $x, y \in X$, we have $x <_\mathrm{v} y$ if and only if $\mathrm{height}(x) < \mathrm{height}(y)$. Thus, P occurs in π if and only there exists a sequence $0 = a_0 < a_1 < a_2 < \ldots < a_\ell = n$ such that w_σ is a subsequence of w_π, where $w_\pi \in [\ell]^n$ and $w_\sigma \in [\ell]^k$ are the two words defined by $w_\pi[i] = \min\{j \colon a_{j-1} < \pi(i) \leqslant a_j\}$ and $w_\sigma[i] = \mathrm{height}(i)$.

As for the running time, there exist $\binom{n-1}{\ell-1}$ distinct sequences $(a_i)_{0 \leqslant i \leqslant \ell}$ and deciding whether w_σ occurs in w_π as a subsequence is a linear-time procedure. ◀

▶ **Proposition 7.** *DPOP Matching is* NP*-complete even if* $\mathrm{height}(P_\mathrm{v}) = 2$*, P_p is a total order and π avoids* 1234*.*

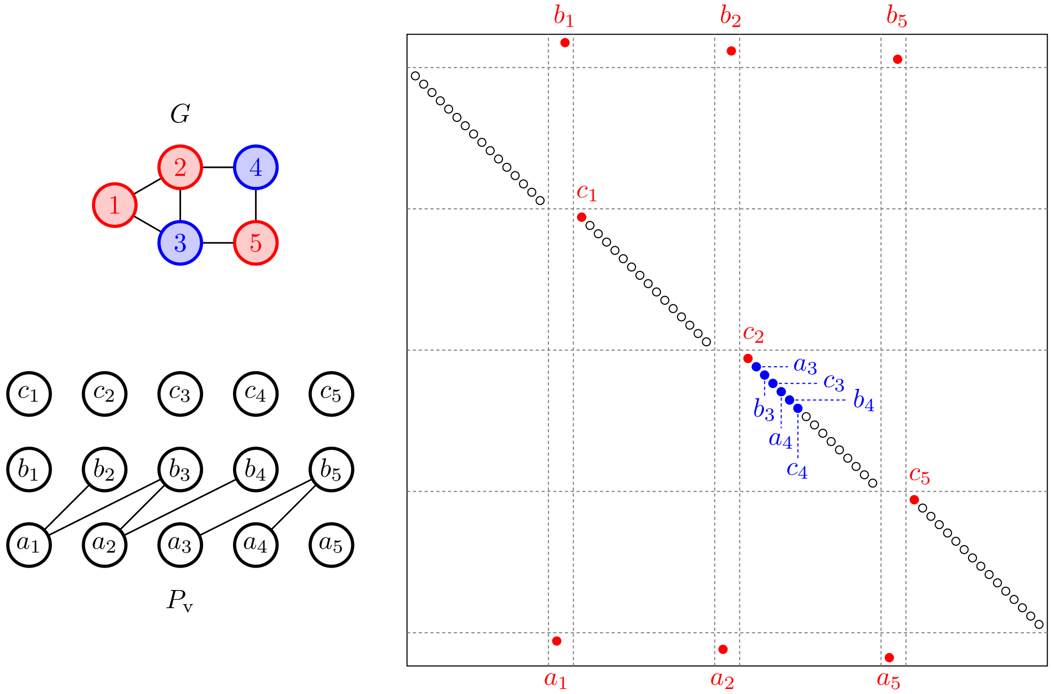

Figure 4 Top left: vertex cover (in red) in a 5-vertex graph G. Bottom left: partially ordered set P_{V} constructed from G. Right: an occurrence of P in π obtained from our size-3 vertex cover.

Proof. We perform a reduction from Vertex Cover, which is known to be NP-complete [22]. Let $G = (V, E)$ be a graph and let k be a positive integer. We identify V with the set $[n]$.

We construct a dpop $P = (P_{\mathrm{v}}, P_{\mathrm{p}})$, where $P_{\mathrm{v}} = (X, \leqslant_{\mathrm{v}})$ is a height-2 partial order and $P_{\mathrm{p}} = (X, \leqslant_{\mathrm{p}})$ a total order, as follows. We set $X = \{a_1, b_1, c_1, a_2, b_2, c_2, \ldots, a_n, b_n, c_n\}$, so that $|X| = 3n$. Then, we set $a_1 \leqslant_{\mathrm{p}} b_1 \leqslant_{\mathrm{p}} c_1 \leqslant_{\mathrm{p}} a_2 \leqslant_{\mathrm{p}} b_2 \leqslant_{\mathrm{p}} c_2 \leqslant_{\mathrm{p}} \cdots \leqslant_{\mathrm{p}} a_n \leqslant_{\mathrm{p}} b_n \leqslant_{\mathrm{p}} c_n$, which defines the total order $\leqslant_{\mathrm{p}}$. Finally, for each edge $\{i, j\}$ in E with $i < j$, we set $a_i \leqslant_{\mathrm{v}} b_j$; all other elements of X are pairwise incomparable by $\leqslant_{\mathrm{v}}$. This defines a partial order $\leqslant_{\mathrm{v}}$ such that $\mathrm{height}(\leqslant_{\mathrm{v}}) = 2$.

Write now $N = 3n + 3$ and $m = (k + 1)N - 2$, and define a permutation $\pi \in \mathfrak{S}(m)$ as follows:

- $\pi(iN + j) = (m + 1) - (iN + j + k)$ whenever $0 \leqslant i \leqslant k$ and $1 \leqslant j \leqslant N - 2$;
- $\pi(iN - 1) = (k + 1) - i$ and $\pi(iN) = (m + 1) - i$ whenever $1 \leqslant i \leqslant k$.

It is straightforward to check that π is 1234-avoiding. It is also easy to see how the construction, illustrated in Figure 4, can be accomplished in polynomial-time.

Let us see under which conditions an injective $\leqslant_{\mathrm{p}}$-non-decreasing function $\varphi \colon X \to [m]$ maps P into π. We say that a vertex i belongs to the *j-th gadget* if one of the integers $\varphi(a_i)$ or $\varphi(b_i)$ is equal to $jN - 1$ or to jN, i.e., if $\{\varphi(a_i), \varphi(b_i)\} \cap \{jN - 1, jN\} \neq \emptyset$. When two elements in the range of φ are consecutive, either they are integers $\varphi(a_i)$ and $\varphi(b_i)$ for a given i, or one of them is an integer $\varphi(c_i)$ for some i. Therefore, no two distinct vertices i and i' can belong to the same j-th gadget. Consequently, and since there are k gadgets, the set V' of vertices i that belong to some gadget is of size at most k.

Then, we define a notion of *height* as follows: for each element x of X, we set $\mathrm{height}(x) = 0$ if N divides $\varphi(x) + 1$, $\mathrm{height}(x) = 2$ if N divides $\varphi(x)$, and $\mathrm{height}(x) = 1$ otherwise. By construction, for all $x, y \in X$ such that $x \leqslant_{\mathrm{p}} y$, we have $\pi(\varphi(x)) \leqslant \pi(\varphi(y))$ if and only if x

is of smaller height than y. Therefore, if φ maps P into π, and for each relation $a_i \leqslant_{\mathrm{v}} b_j$, either a_i has height 0 or b_j has height 2. In particular, either i or j must belong to V', and therefore V' is a vertex cover of size at most k.

Conversely, provided that there exist vertices $v(1) < v(2) < \ldots < v(k)$ that form a vertex cover V', we construct an occurrence of P in π as follows. First, we abusively set $v(0) = 0$. Then, for all $i \in [k]$, we set $f(i) = jN + 3(i - v(j))$, where j is the largest integer such that $v(j) \leqslant i$. We set $\varphi(a_j) = f(j) - 1$, $\varphi(b_j) = f(j)$ and $\varphi(c_j) = f(j) + 1$.

By construction, we have $f(i) + 3 \leqslant f(i+1)$ for all i, and therefore φ is an injective $\leqslant_{\mathrm{p}}$-non-decreasing function. Moreover, for every $i \in [n]$, the elements a_i and b_i have heights 0 and 2 if $i \in V'$, and they have height 1 if $i \notin V'$. It follows that $\pi(\varphi(a_i)) \leqslant \pi(\varphi(b_j))$ whenever $a_i \leqslant_{\mathrm{v}} b_j$, i.e., that φ is an occurrence of P in π. ◀

5 Symmetric Patterns

This section is devoted to studying complexity issues of pattern matching for symmetric dpop (i.e., those dpops $P = (\mathcal{P}, \mathcal{P})$, whose value and position posets coincide). We further focus on two special cases, first when $\mathcal{P}$ has a bounded width, then when π is restricted to constrained pattern-avoiding classes of permutations.

5.1 Symmetric Pattern with Bounded Width

We first observe that the problem is polynomial for width 1 (Observation 8). We further prove W[1]-hardness for the parameter k when $\mathcal{P}$ is a disjoint union of k chains (Proposition 10). We complement this result with an XP algorithm for the slightly more general case where $\mathcal{P}$ is a disjoint union of weak orders (Proposition 11). Note that the existence of an XP algorithm for the width parameter remains open, and we conjecture that the problem is NP-hard even for constant width.

▶ **Observation 8.** *DPOP Matching is solvable in $O(n \log\log |P|)$ time for a symmetric dpop P of width 1 (i.e., a total symmetric dpop P).*

Proof. If P has width 1, then $P = (\mathcal{P}, \mathcal{P})$ for some total order $\mathcal{P} = (X, \preccurlyeq)$. In particular, we can write $X = \{x_1, \ldots, x_{|X|}\}$ with $x_i \prec x_j$ for $i < j$, and in any mapping $\phi : X \to [n]$, the elements $\pi_{\phi(x_1)}, \ldots, \pi_{\phi(x_{|X|})}$ must form an increasing subsequence of π. Conversely, any size-$|X|$ increasing subsequence of π can be used as an image for ϕ, so in this setting DPOP Matching corresponds to the longest increasing subsequence problem, which can be solved in $O(n \log\log |X|)$ time [12]. ◀

To simplify the exposition of our next result, we introduce a new problem that may be of independent interest. Given a positive integer k and a permutation $\pi \in \mathfrak{S}(kn)$, Balanced k-Increasing Coloring is the problem of deciding whether there exists a balanced k-coloring of π (i.e., a partition of $[kn]$ into k subsets of size exactly n) such that each color induces an increasing subsequence of π.

▶ **Proposition 9.** *Balanced k-Increasing Coloring for 312-avoiding permutations is W[1]-hard for the parameter k.*

Proof. We perform a reduction from Unary Bin Packing parameterized by the number of bins, which is known to be W[1]-hard [20]. In this version of Bin Packing, we are given a list of integers $s_1, s_2, \ldots, s_n$ encoded in unary, and two integers B and k. These integers are interpreted as item sizes, and the task is to decide whether the items can be partitioned

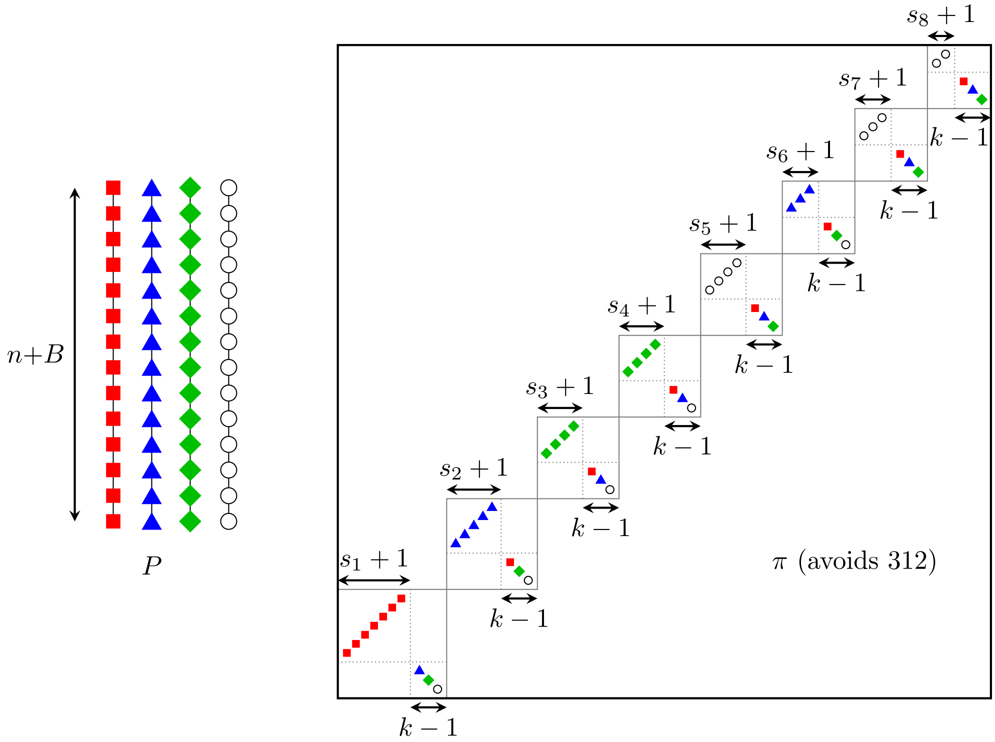

Figure 5 Reduction from UNARY BIN PACKING to BALANCED k-INCREASING COLORING for the list $6, 4, 3, 3, 3, 2, 2, 1$, which admits the partition $(\{6\}, \{4, 2\}, \{3, 3\}, \{3, 2, 1\})$, and the integers $B = 6$, $k = 4$. Left: dpop P that consists of k chains, each of length $n + B$. Right: 312-avoiding permutation π that consists of n blocks. Each color/shape induces an increasing subsequence of π.

into k susbets, each of total size B. We show that there is a reduction from UNARY BIN PACKING, parameterized by the number of bins, to BALANCED k-INCREASING COLORING, parameterized by the number of colors.

Consider an arbitrary instance of UNARY BIN PACKING containing n items with item sizes $S = \{s_1, s_2, \ldots, s_n\}$, and two integers B and k. Define $\pi \in \mathfrak{S}(kB + kn)$ by

$$\pi = \bigoplus_{i=1}^{n} (\mathrm{ip}(s_i + 1) \ominus \mathrm{dp}(k - 1)) .$$

Each pattern $\mathrm{ip}(s_i+1) \ominus \mathrm{dp}(k-1)$ is called the *i-th block* of π. See Figure 5 for an illustration. It is straightforward to check that π is 312-avoiding.

We claim that the n items $s_1, s_2, \ldots, s_n$ can be partitioned into k susbets, each of total size B, if and only if there exists a k-coloring of π such that each color induces an increasing pattern of length $B + n$.

Suppose first that the n items $s_1, s_2, \ldots, s_n$ can be partitioned into k susbets, each of total size B. Write $S = S_1 \cup S_2 \cup \cdots \cup S_k$ such a partition. Define a k-coloring of π as follows. Consider the i-th block $\mathrm{ip}(s_i + 1) \ominus \mathrm{dp}(k - 1)$ of π, and suppose that $s_i \in S_j$. Color the whole ascending pattern $\mathrm{ip}(s_i + 1)$ with color c_j and arbitrarily color the elements of the descending pattern $\mathrm{dp}(k - 1)$ with the remaining $k - 1$ colors (each element of $\mathrm{dp}(k - 1)$ is assigned to a distinct color). We claim that every color c_j induces an increasing pattern of length $B+n$ in π. First, it is clear that the above k-coloring induces increasing patterns only. As for the length of each induced increasing pattern, focus on any color c_j. We note that, in every block $\text{ip}(s_i+1) \ominus \text{dp}(k-1)$ of π, either the whole subpattern $\text{ip}(s_i+1)$ is colored with color c_j (if $s_i \in S_j$) or exactly one element of the subpattern $\text{dp}(k-1)$ is colored with color c_j (if $s_i \notin S_j$). It follows that the increasing pattern induced by color c_j in π has length $\sum_{s_i \in S_j}(s_i+1) + n - |S_i| = \sum_{s_i \in S_j} s_i + |S_i| + n - |S_i| = B + n$.

For the reverse direction, suppose now that there exists a k-coloring of π such that each color induces an increasing pattern of length $B+n$. Every block $\text{ip}(s_i+1) \ominus \text{dp}(k-1)$ requires at least k colors, as it contains a decreasing subpattern of length k. Therefore, the whole subpattern $\text{ip}(s_i+1)$ is colored with the same color. For every $j \leqslant k$, let S_j be the set of all s_i such that, in the i-th block $\text{ip}(s_i+1) \ominus \text{dp}(k-1)$, the subpattern $\text{ip}(s_i+1)$ is colored with color c_j. We have $B+n = \sum_{s_i \in S_j}(s_i+1) + n - |S_j| = \sum_{s_i \in S_j} s_i + |S_j| + n - |S_j|$, and hence $\sum_{s_i \in S_j} s_i = B$. Therefore, the n items $s_1, s_2, \ldots, s_n$ can be packed into k bins, each of capacity B. ◀

Most of the interest in Proposition 9 stems from the following proposition.

▶ **Proposition 10.** *DPOP MATCHING for symmetric dpop and* 312*-avoiding permutations is W[1]-hard for the parameter* $\text{width}(P)$.

Proof. We perform a reduction from BALANCED k-INCREASING COLORING, which is W[1]-hard for the parameter k. Let $\pi \in \mathfrak{S}(kn)$ for some positive integers k and n. We construct a symmetric dpop $P = (\mathcal{P}, \mathcal{P})$, where $\mathcal{P} = (X, \preccurlyeq)$, as follows: $X = [k] \times [n]$ and $(i,j) \preccurlyeq (i',j')$ if and only if $i = i'$ and $j \leqslant j'$. We claim that P occurs in π if and only π admits a k-coloring for which every color induces an increasing pattern of length n.

If π admits such a k-coloring into colors $c_1, c_2, \ldots, c_k$, the function $\varphi\colon X \to [kn]$ that maps each pair (i,j) to the j-th smallest position with color c_i is an occurrence of P in π.

Conversely, suppose that some injective function $\varphi : X \to [kn]$ is an occurrence of P in π. For each $i \leqslant k$, the set $\{i\} \times [n]$ forms a chain of $\preccurlyeq$, and therefore it is mapped to an increasing pattern of size n. Coloring this pattern in color c_i produces the desired k-coloring. ◀

We show now that the problem where P consists of k independent chains is XP for the parameter k. In fact, we generalize this result to k-weak partial orders (i.e., if P consists of k independent weak orders).

▶ **Proposition 11.** *DPOP MATCHING for k-weak symmetric dpop is* XP *with parameter k.*

Proof. Let P be a disjoint union of k weak symmetric dpops $P_1, P_2, \ldots, P_k$. For each dpop P_i, let $\preccurlyeq_i$ be a linear extension of P_i, and let $P_{i,1}, P_{i,2}, \ldots, P_{i,p_i}$ be the maximal antichains of P_i, ordered by $\preccurlyeq_i$. Finally, for each k-tuple $\mathbf{a} = (a_1, a_2, \ldots, a_k)$ of integers such that $a_i \leqslant |P_i|$, we denote by $P_{\mathbf{a}}$ the dpop obtained from P by removing the a_i $\preccurlyeq_i$-least elements of each dpop P_i, and by $P_{i,\mathbf{a}}^{\min}$ the set of $\preccurlyeq_i$-minimal elements of $P_{\mathbf{a}}$.¨

Then, given a permutation $\pi \in \mathfrak{S}(n)$, a k-tuple $\mathbf{I} = (I_1, \ldots, I_k)$ of intervals of $[n]$, a k-tuple $\mathbf{a}$ and an integer ℓ, a function $\varphi\colon P_{\mathbf{a}} \to \{\ell, \ell+1, \ldots, n\}$ is called a *partial matching* for $(\pi, \mathbf{I}, \mathbf{a}, \ell)$ if:

- φ and $\pi \circ \varphi$ are $\preccurlyeq_i$-non-decreasing for each i, and
- for each i, and each element x of $P_{\mathbf{a}}$, $\pi(x) \in I_i$ if and only if $x \in P_{i,\mathbf{a}}^{\min}$.

Before going further, we denote by $\mathbf{1}_i$ the k-tuple with one element 1 (in position i) and $k-1$ elements 0. We also denote by $<_i$ the partial order on tuples $\mathbf{I}$ of intervals, where $\mathbf{I} <_i \mathbf{I}'$ if $I_j = I'_j$ whenever $j \neq i$ and $x < x'$ whenever $x \in I_i$ and $x' \in I'_i$.

When $a_i = |P_i|$ for all i, such a partial matching exists for all permutations π, tuples of intervals $\mathbf{I}$ and integers ℓ. When $a_i = 0$ for all i and $\ell = 1$, and once π is fixed, such partial matchings coincide with (standard) matchings, and thus we are interested in checking whether a partial matching exists. Finally, for all tuples $\mathbf{I}$ and $\mathbf{a}$ and for all $\ell \leqslant n$, a partial matching φ for $(\pi, \mathbf{I}, \mathbf{a}, \ell)$ exists precisely when one of the following cases occur:

1. φ is a partial matching for $(\pi, \mathbf{I}, \mathbf{a}, \ell + 1)$, i.e., $\ell \notin \varphi(P_\mathbf{a})$;
2. there exists an integer $i \leqslant k$ for which the $\preccurlyeq_i$-least element of $P^{\min}_{i,\mathbf{a}}$, say x, is such that $\varphi(x) = \ell$ and $\pi(\ell) \in I_i$, and either
   - x is not the only element of $P^{\min}_{i,\mathbf{a}}$, and φ is a partial matching for $(\pi, \mathbf{I}, \mathbf{a} + \mathbf{1}_i, \ell + 1)$, or
   - x is the only element of $P^{\min}_{i,\mathbf{a}}$ and there exists a tuple $\mathbf{I}' >_i \mathbf{I}$ such that φ is a partial matching for $(\pi, \mathbf{I}', \mathbf{a} + \mathbf{1}_i, \ell + 1)$.

Consequently, we can compute by dynamic programming the list of triples $(\mathbf{I}, \mathbf{a}, \ell)$ such that there exists a partial matching for $(\pi, \mathbf{I}, \mathbf{a}, \ell)$: deciding whether adding a triple $(\mathbf{I}, \mathbf{a}, \ell)$ to the list simply requires to check which triples of the form $(\mathbf{I}', \mathbf{a}', \ell + 1)$ already belong to the list. Since there are less than n^{3k+1} triples, this provides us with a $\tilde{\mathcal{O}}(n^{6k+2})$ algorithm. ◀

5.2 Symmetric Pattern and Pattern-Avoiding π

In this final section, we consider restrictions on the *shape* of π, via pattern-avoiding restrictions. Our goal here is to identify tractable cases among classes of permutations avoiding one or more size-3 patterns. We give an almost complete dichotomy of polynomial/NP-hard cases, as shown in Table 1. Hardness results are proven in Proposition 12, and also apply to height-2 dpops. Polynomial cases are proven in Proposition 13 and apply to dpops of any height.

▶ **Proposition 12.** *DPOP Matching for height-2 symmetric dpop P and permutation π is NP-hard even if π is separable (it avoids 2413 and 3142) and one of the following restrictions occurs:*

1. *π is 123-avoiding;*
2. *π is $(132, 213)$-avoiding;*
3. *π is $(132, 321)$-avoiding;*
4. *π is $(231, 312)$-avoiding;*
5. *π is $(132, 312)$-avoiding;*
6. *π is $(213, 321)$-avoiding;*
7. *π is $(213, 312)$-avoiding;*
8. *π is $(132, 231)$-avoiding;*
9. *π is $(213, 231)$-avoiding.*

Proof. In each of the cases presented below, we define a symmetric dpop $P = (\mathcal{P}, \mathcal{P})$ for some partially ordered set $\mathcal{P} = (X, \preccurlyeq)$. Each time, we identify P with the partial order $\preccurlyeq$.

Table 1 Polynomial (green/light) and NP-hard (red/dark) cases for DPOP Matching with symmetric dpop and pattern-avoiding permutation π, for combinations of size-3 avoided patterns. For each case, see the referenced proposition and case for more details. Diagonal cases follow from any other hard case in the same row or column. For hard cases, the problem used for reduction is indicated as follows: BIC: Biclique, 3P: 3-Partition, BIN: Unary Bin Packing, BIS: Bisection.

π avoids … and …	123	132	213	231	312	321
123	12.1 3P					
132	13.2					
213	13.5	12.2 BIN				
231	13.1	12.8 BIC	12.9 BIC			
312	13.4	12.5 BIC	12.7 BIC	12.4 BIC		
321	13.3	12.3 BIS	12.6 BIS	(open)		

NP-hard cases

Poynomial cases

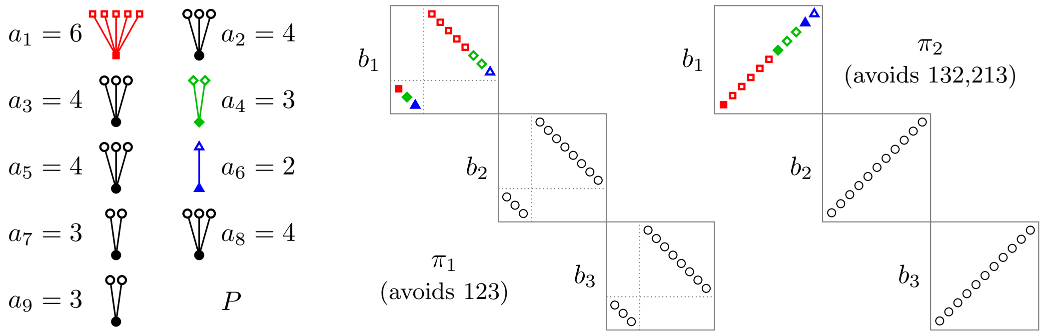

Figure 6 Reductions from 3-PARTITION and UNARY BIN PACKING to DPOP MATCHING on 123-avoiding and $(132, 213)$-avoiding permutations. Left: the height-2 dpop P used in both reductions. Right: the 123-avoiding and $(132, 213)$-avoiding permutations used in each reduction. The mapping of three subsets of P corresponding to a first bin gadget is highlighted in each figure.

Case 1: π is 123-avoiding and separable. We use a reduction from 3-PARTITION, as illustrated in Figure 6 with permutation π_1. Let (A, B) be an instance of 3-PARTITION, where A is a list of integers $a_1, a_2, \ldots, a_{3n}$ with sum nB, all being larger than 1.

For all $p \leqslant n$, we define a *bin gadget* b_p as the permutation $\mathrm{dp}(3) \oplus \mathrm{dp}(B-3)$: we see this gadget as consisting of two parts. Our permutation π is now defined by $\pi = \ominus_{p=1}^{n} b_p$. Then, our partial order $\preccurlyeq$ is defined on a set X of nB elements, noted $x_i, y_{i,2}, \ldots, y_{i,a_i}$ for each $i \leqslant 3n$, so that $x_i \preccurlyeq y_{i,j}$ for all i and j.

If P has an occurrence $\varphi \colon X \to [nB]$ in π, this occurrence is bijective. Moreover, each element x_i is sent to the bottom-left of $y_{i,2}$, and thus it must is mapped to the left part of some gadget, say $b_{f(i)}$. Each element $y_{i,j}$ must then be mapped to the right part of the same gadget. Now, for each $p \leqslant n$, the set $S_p = \{i \colon f(i) = p\}$ is of size 3, and exactly B elements of X are mapped to the gadget b_p, which means that $\sum_{i \in S_p} a_i = B$. Moreover, $S_1 \cup \ldots \cup S_p$ forms a partition of $[3n]$, hence it yields a 3-partition of A.

Conversely, given a partition $S_1 \cup \ldots \cup S_n$ of $[3n]$ such that $|S_p| = 3$ and $\sum_{i \in S_p} a_i = B$ for each p, we build an occurrence of P in π by mapping the three elements x_i (for $i \in S_p$) to the left part of b_p, and the $B-3$ elements $y_{i,j}$ (for $i \in X_p$) to the right part of b_p.

Case 2: π is $(132, 213)$-avoiding. We use a reduction from UNARY BIN PACKING, as illustrated in Figure 6 with permutation π_2. Given an instance (A, B, k) of UNARY BIN PACKING, where A is a list of integers $a_1, a_2, \ldots, a_n$ larger than 1, we use the same dpop as in Case 1: our partial order $\preccurlyeq$ is defined on a set X of nB elements, noted $x_i, y_{i,2}, \ldots, y_{i,a_i}$ for each $i \leqslant n$, so that $x_i \preccurlyeq y_{i,j}$ for all i and j. However, this time, our gadget b_p is the permutation $\mathrm{ip}(B)$, and our permutation π is again defined by $\pi = \ominus_{p=1}^{n} b_p$.

If P has an occurrence $\varphi \colon X \to [nB]$ in π, this occurrence is bijective. Each element x_i is sent to some gadget, say $b_{f(i)}$, and the elements $y_{i,j}$ must then be mapped to the same gadget. Now, for each $p \leqslant k$, let $S_p = \{i \colon f(i) = p\}$. Exactly B elements of X are mapped to the gadget b_p, which means that $\sum_{i \in S_p} a_i = B$. This means that (S, B, k) is a positive instance of the UNARY BIN PACKING problem.

Conversely, given a partition $S_1 \cup \ldots \cup S_k$ of $[n]$ such that $\sum_{i \in S_p} a_i = B$ for each i, we build an occurrence of P in π by mapping the B elements x_i and $y_{i,j}$ (for $i \in S_p$) to b_p.

Case 3: π is $(132, 321)$-avoiding. We use a reduction from BISECTION, as illustrated in Figure 7. Given a graph $G = (V, E)$ and an integer k, the BISECTION problem consists in deciding whether V admits a partition $V_1 \cup V_2$ such that $|V_1| = |V_2|$ and that *splits* at most k edges (i.e., at most k edges have one endpoint in V_1 and one endpoint in V_2).

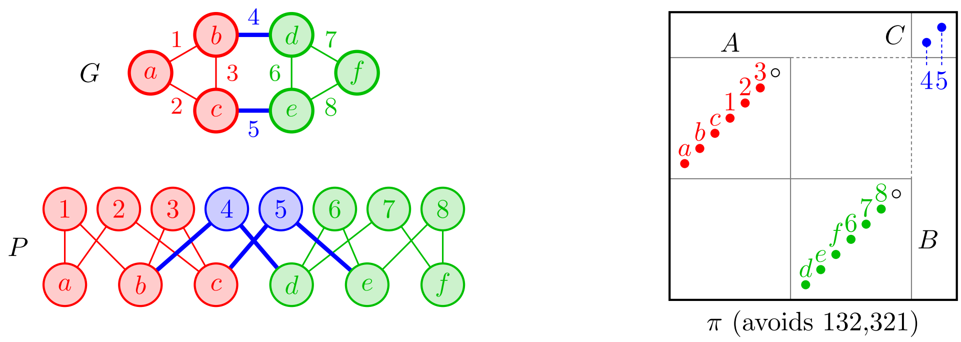

Figure 7 Simplified version (for $W = 1$ and $L = 7$) of the reduction from BISECTION to DPOP MATCHING on $(132, 321)$-avoiding permutations. Top left: a size-6 graph with a bisection $(\{a, b, c\}, \{d, e, f\})$ that splits $k = 2$ edges. Bottom left: height-2 dpop P (note that, in general, each element from the bottom line should appear W times and not just once). Right: permutation π; elements of P are mapped to A, B or C depending on whether they are colored in red, green or blue.

Our reduction is as follows. Let $n = |V|/2$, $m = |E|$, $W = m + k + 1$, and $L = nW + m$. Our permutation is defined by $\pi = (\mathrm{ip}(L) \ominus \mathrm{ip}(L)) \oplus \mathrm{ip}(k)$. These three parts of π are noted A, B and C, from left to right. Then, our partial order $\preccurlyeq$ is defined on a set X of $2nW + m$ elements: $2nW$ elements, noted $x_{v,i}$ for each $v \in V$ and $i \leqslant W$, and m elements, noted y_e for each $e \in E$. This order contains the relations $x_{v,i} \preccurlyeq y_e$ for which v is an endpoint of e.

Assume that there exists a mapping of P into π. For each $v \in V$, and since C has size $k < W$, at least one of the elements $x_{v,i}$ is mapped to A, in which case we say that v has type A, or to B, in which case v has type B. Then, each vertex has at least one type, and possibly both. We partition V into three sets V_A, V_B, V_{AB} containing the vertices of type A, B and both A and B, respectively. Moreover, for each $v \in V_A$, each element $x_{v,i}$ must be mapped either to A or to C: these two parts together contain $L + k$ elements, so $|V_A| \leqslant (L + k)/W = n + 1 - 1/W$, and $|V_A| \leqslant n$. Similarly, $|V_B| \leqslant n$.

We build a set V_1 as the union of V_A with $n - |V_A|$ vertices of V_{AB}, and V_2 as $V \setminus V_1$, so that $|V_1| = |V_2| = n$. Moreover, for every $v \in V_1$ (resp., $v \in V_2$), some element $x_{v,i}$, say $x_{v,1}$, is mapped to A (resp., to B). Then, each edge $e = (u, v)$ that is split by (V_1, V_2) must be mapped to a point above some point of A and to the right of some point of B. This means that y_e is mapped to C, and that (V_1, V_2) splits at most k edges, i.e., is a valid bisection.

Conversely, given a bisection (V_1, V_2) splitting at most k edges, we map P into π as follows: map elements $x_{v,i}$ for $v \in V_1$ (resp., V_2) to the first nW elements of A (resp., B), map elements y_e for which e is induced by V_1 (resp., V_2) to the following elements of A (resp. B), and finally map all elements y_e such that e is split by (V_1, V_2) into C. This mapping is an occurrence of P in π.

Case 4: π is $(231, 312)$-avoiding. We use a reduction from BICLIQUE, as illustrated in Figure 8 with permutation π_1. Given a bipartite graph $G = (V, E)$ and an integer k, the BICLIQUE problem consists in deciding whether V admits a complete bipartite subgraph $K_{k,k}$. If $V = A \cup B$ is a partition of V into two independent sets, adding independent vertices if needed allows us to assume that A and B have the same size n, and that no vertex in either side is fully connected to the other side.

Our permutation π is defined by $\pi = \mathrm{dp}(n - k) \oplus \mathrm{dp}(2k) \oplus \mathrm{dp}(n - k)$. These three parts of π are noted b_1, b_2 and b_3. Our partial order $\preccurlyeq$ is the order on V such that $x \preccurlyeq y$ whenever $x \in A$, $y \in B$ and $\{x, y\} \notin E$.

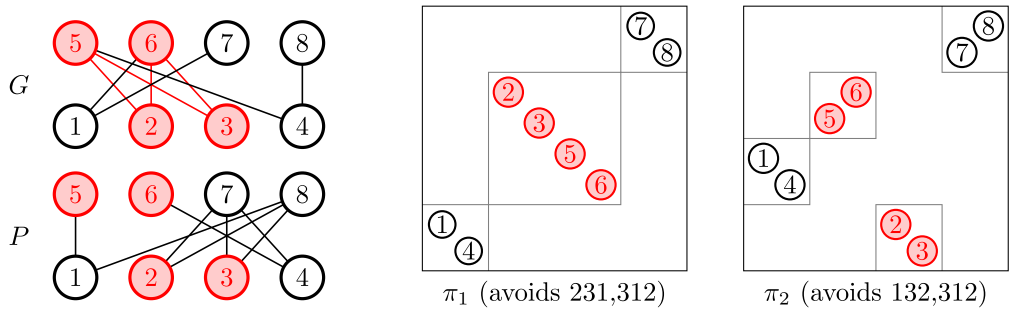

Figure 8 Reduction from BICLIQUE to DPOP MATCHING on $(231, 312)$-avoiding and $(213, 231)$-avoiding permutations. Left: a bipartite graph G with a $(2, 2)$ biclique and the corresponding height-2 dpop P, built from the complement of G. Right: permutations π_1 and π_2 with a mapping of the vertices 1 to 8, including the biclique vertices mapped into the central $2k$ positions.

Consider a mapping of P into π. For each element $x \in A$, there exists $y \in B$ such that $x \preccurlyeq y$, and therefore x cannot be mapped into b_3. Symmetrically, no element $y \in B$ may be mapped into b_1. Overall, since $|\pi| = |V| = 2n$, b_1 contains $n - k$ elements from A, b_3 contains $n - k$ elements from B, and b_2 contains a size-k subset A' of A and a size-k subset B' of B. No two elements $x \in A'$ and $y \in B'$ that are mapped into b_2 are comparable for $\preccurlyeq$, which means that $\{x, y\} \in E$ for each such pair, i.e., that (A', B') is a biclique.

Conversely, if G has a biclique (A', B'), we map all elements of $A \setminus A'$ into b_1, all elements of $A' \cup B'$ into b_2, and all elements of $B \setminus B'$ into b_3. This mapping satisfies all relations $x \preccurlyeq y$ with $x \in A$ and $y \in B$, except for $x \in A'$ and $y \in B'$, but indeed there is no such relation since (A', B') is a biclique.

Case 5: π is $(132, 312)$-avoiding. We also use a reduction from BICLIQUE, as illustrated in Figure 6 with permutation π_2. Our partial order $\preccurlyeq$ is the same as in Case 4, and our permutation π is defined by $\pi = ((\mathrm{dp}(n-k) \oplus \mathrm{ip}(k)) \ominus \mathrm{dp}(k)) \oplus \mathrm{ip}(n-k)$. These three parts of π are noted b_1, b_2, b_3 and b_4.

Consider a mapping of P into π. For each element $y \in B$, there exists $x \in A$ such that $x \preccurlyeq y$, and therefore y cannot be mapped into b_1 or b_3. Thus, and since $|B| = n$, the elements of B are mapped to b_2 or b_4, and the elements of A are mapped to b_1 or b_3. Hence, b_2 contains a size-k subset B' of B and b_3 contains a size-k subset A' of A. No element of b_2 is comparable to any element of b_3, and therefore (A', B') is a biclique.

Conversely, if G has a biclique (A', B'), we map all elements of $A \setminus A'$ into b_1, all elements of A' into b_3, all elements of B' into b_2 and all elements of $B \setminus B'$ into b_4. This mapping satisfies all relations $x \preccurlyeq y$ with $x \in A$ and $y \in B$, except for $x \in A'$ and $y \in B'$, but indeed there is no such relation since (A', B') is a biclique.

Cases 6–9: These cases are symmetric to Cases 3, 5, 5 and 7, respectively. Indeed, if (P, π) is an instance of DPOP MATCHING with P a height-2 symmetric dpop, $(P^{\partial}, \pi^{\mathrm{c\,r}})$ and (P, π^{-1}) are equivalent instances of DPOP MATCHING with height-2 symmetric dpops, and

Case 6: if π avoids 132 and 321 (Case 3), $\pi^{\mathrm{c\,r}}$ avoids $132^{\mathrm{c\,r}} = 213$ and $321^{\mathrm{c\,r}} = 321$;
Case 7: if π avoids 132 and 312 (Case 5), $\pi^{\mathrm{c\,r}}$ avoids $132^{\mathrm{c\,r}} = 213$ and $312^{\mathrm{c\,r}} = 312$;
Case 8: if π avoids 132 and 312 (Case 5), π^{-1} avoids $132^{-1} = 132$ and $312^{-1} = 231$;
Case 9: if π avoids 213 and 312 (Case 7), π^{-1} avoids $213^{-1} = 213$ and $312^{-1} = 231$. ◀

▶ **Proposition 13.** *DPOP* Matching *is in* P *for symmetric dpop P if one of the following restrictions on π occurs:*

1. *π is $(123, 231)$-avoiding;*
2. *π is $(123, 132)$-avoiding;*
3. *π is $(123, 321)$-avoiding;*
4. *π is $(123, 312)$-avoiding;*
5. *π is $(123, 213)$-avoiding.*

Proof. In each of the cases presented below, we are given a permutation π and a symmetric dpop $P = (\mathcal{P}, \mathcal{P})$ for some partially ordered set $\mathcal{P} = (X, \preccurlyeq)$. Each time, we identify P with the partial order $\preccurlyeq$.

Case 1: π is $(123, 231)$-avoiding. There exist integers k, ℓ and m, with sum n, such that $\pi = \mathrm{dp}(k) \ominus (\mathrm{dp}(\ell) \oplus \mathrm{dp}(m))$. These three parts of π are noted b_1, b_2 and b_3. Then, for every pair (u, v) such that $u \prec v$, we must map u into b_2 and v into b_3. Such values can be mapped greedily, since elements in b_2 are pairwise incompatible, as well as those in b_3. Thus, P can be mapped into π if and only if it has height at most 2, there are at most a elements that are lower bounds, and at most b elements that are upper bounds.

Note that, if P is not symmetric, the problem becomes NP-hard, since reversing the horizontal order of P_{p} and π transforms π into the $(132, 321)$-avoiding permutation of the NP-hard Case 3 in Proposition 12.

Case 2: π is $(123, 132)$-avoiding. The permutation π is a skew sum $\pi = \ominus_{p=1}^{k} d_p$ of patterns of the form $d_p = \mathrm{dp}(a_p) \oplus \mathrm{dp}(1)$ for some integer $a_p \geqslant 0$. Then, no two elements in X can share a strict lower bound, i.e., if $u \prec v$ and $u \prec w$ then $v = w$. Thus, P is of height at most 2, and there exists a partition $S_1 \cup \ldots \cup S_\ell$ of X in which each set S_i contains a distinguished element s_i, such that $x \preccurlyeq s_i$ if and only if $x \in S_i$. Up to reordering the patterns d_p and the sets S_i, which are pairwise incomparable, we assume that $a_1 \geqslant a_2 \geqslant \ldots \geqslant a_k$ and that $|S_1| \geqslant |S_2| \geqslant \ldots \geqslant |S_\ell|$. Let also m be the number of sets S_i with size at least 2.

Each set S_i must be mapped into a single pattern, say $d_{p(i)}$, and if $i \leqslant m$, i.e., if $|S_i| \geqslant 2$, the element s_i must be mapped to the unique top-right element of $d_{p(i)}$. Such a mapping exists if and only if $k \geqslant m$ and $a_i \geqslant |S_i| - 1$ for all $i \leqslant m$: we shall choose $p(i) = i$ and map greedily the elements of $S_i \setminus \{s_i\}$ to the bottom-left part of d_i. Finally, the elements of singleton sets S_i can be mapped to the remaining places in π.

Case 3: π is $(123, 321)$-avoiding. Erdős-Szekeres theorem [14] proves that $n \leqslant 4$.

Cases 4–5: These cases are symmetric to Cases 1 and 2, respectively. Indeed, if (P, π) is an instance of DPOP Matching with P a height-2 symmetric dpop, $(P^{\partial}, \pi^{\mathrm{c\,r}})$ and (P, π^{-1}) are equivalent instances of DPOP Matching with height-2 symmetric dpops, and

Case 4: if π avoids 123 and 231 (Case 1), π^{-1} avoids $123^{-1} = 123$ and $231^{-1} = 312$;

Case 5: if π avoids 123 and 132 (Case 2), $\pi^{\mathrm{c\,r}}$ avoids $123^{\mathrm{c\,r}} = 123$ and $132^{\mathrm{c\,r}} = 213$. ◀

6 Concluding Remarks

Some open complexity questions remain among the parameters we identified for DPOP Matching. For semi-total dpops, the complexity is open for constant width, and for most classes of pattern-avoiding permutations (although, according to Propositions 7 and 10, the problem is NP-hard when π avoids 1234 or 312, respectively). For symmetric dpops, it would be interesting to settle the complexity status of deciding whether a dpop occurs in a $(231, 321)$-avoiding or $(312, 321)$-avoiding permutation. In particular, for these cases, we conjecture that the problem becomes polynomial when $\mathrm{height}(P)$ is constant.

Regarding the original puzzle formulation of the problem, an interesting question is to generate instances that yield a unique solution, i.e., given a permutation π, find a dpop with a unique occurrence in π. This can be done by using a semi-total dpop (e.g., take X

with $|X| = |\pi|$, let P_p be a total order and P_v be an empty order), but one could try to minimize $|X|$ or the number of pairs of comparable elements in P (i.e., the number of clues) in order to have a unique solution.

References

1 Shlomo Ahal and Yuri Rabinovich. On complexity of the subpattern problem. *SIAM J. Discret. Math.*, 22(2):629–649, 2008.

2 Michael H. Albert, Marie-Louise Lackner, Martin Lackner, and Vincent Vatter. The complexity of pattern matching for 321-avoiding and skew-merged permutations. *Discret. Math. Theor. Comput. Sci.*, 18(2), 2016.

3 Sergey V. Avgustinovich, Sergey Kitaev, and Alexandr Valyuzhenich. Avoidance of boxed mesh patterns on permutations. *Discret. Appl. Math.*, 161(1-2):43–51, 2013.

4 Eric Babson and Einar Steingrímsson. Generalized permutation patterns and a classification of the mahonian statistics. *Séminaire Lotharingien de Combinatoire [electronic only]*, 44:B44b, 18 p.–B44b, 18 p., 2000. URL: `http://eudml.org/doc/120841`.

5 Benjamin Aram Berendsohn, László Kozma, and Dániel Marx. Finding and counting permutations via csps. *Algorithmica*, 83(8):2552–2577, 2021.

6 Miklós Bóna. *Combinatorics of Permutations, Second Edition*. Discrete mathematics and its applications. CRC Press, 2012.

7 Prosenjit Bose, Jonathan F. Buss, and Anna Lubiw. Pattern matching for permutations. *Inf. Process. Lett.*, 65(5):277–283, 1998.

8 Mireille Bousquet-Mélou, Anders Claesson, Mark Dukes, and Sergey Kitaev. (2+2)-free posets, ascent sequences and pattern avoiding permutations. *J. Comb. Theory, Ser. A*, 117(7):884–909, 2010.

9 Petter Brändén and Anders Claesson. Mesh patterns and the expansion of permutation statistics as sums of permutation patterns. *Electron. J. Comb.*, 18(2), 2011.

10 Marie-Louise Bruner and Martin Lackner. A fast algorithm for permutation pattern matching based on alternating runs. *Algorithmica*, 75(1):84–117, 2016.

11 Laurent Bulteau, Romeo Rizzi, and Stéphane Vialette. Pattern matching for k-track permutations. In Costas S. Iliopoulos, Hon Wai Leong, and Wing-Kin Sung, editors, *Combinatorial Algorithms - 29th International Workshop, IWOCA 2018, Singapore, July 16-19, 2018, Proceedings*, volume 10979 of *Lecture Notes in Computer Science*, pages 102–114. Springer, 2018.

12 Maxime Crochemore and Ely Porat. Fast computation of a longest increasing subsequence and application. *Inf. Comput.*, 208(9):1054–1059, 2010.

13 Sergi Elizalde and Marc Noy. Consecutive patterns in permutations. *Adv. Appl. Math.*, 30(1-2):110–125, 2003.

14 Paul Erdös and George Szekeres. A combinatorial problem in geometry. *Compositio mathematica*, 2:463–470, 1935.

15 Jacob Fox. Stanley-wilf limits are typically exponential. *CoRR*, abs/1310.8378, 2013. `arXiv:1310.8378`.

16 Paweł Gawrychowski and Mateusz Rzepecki. Faster exponential algorithm for permutation pattern matching. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 279–284. SIAM, 2022.

17 Sylvain Guillemot and Dániel Marx. Finding small patterns in permutations in linear time. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 82–101. SIAM, 2014.

18 Sylvain Guillemot and Stéphane Vialette. Pattern matching for 321-avoiding permutations. In Yingfei Dong, Ding-Zhu Du, and Oscar H. Ibarra, editors, *Algorithms and Computation, 20th International Symposium, ISAAC 2009, Honolulu, Hawaii, USA, December 16-18, 2009. Proceedings*, volume 5878 of *Lecture Notes in Computer Science*, pages 1064–1073. Springer, 2009.

19 Louis Ibarra. Finding pattern matchings for permutations. *Inf. Process. Lett.*, 61(6):293–295, 1997.

20 Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.*, 79(1):39–49, 2013.

21 Vít Jelínek and Jan Kynčl. Hardness of permutation pattern matching. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 378–396. SIAM, 2017.

22 Richard M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. Springer US, Boston, MA, 1972.

23 Sergey Kitaev. Introduction to partially ordered patterns. *Discret. Appl. Math.*, 155(8):929–944, 2007.

24 Sergey Kitaev. *Patterns in Permutations and Words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, 2011.

25 Both Emerite Neou, Romeo Rizzi, and Stéphane Vialette. Pattern matching for separable permutations. In Shunsuke Inenaga, Kunihiko Sadakane, and Tetsuya Sakai, editors, *String Processing and Information Retrieval - 23rd International Symposium, SPIRE 2016, Beppu, Japan, October 18-20, 2016, Proceedings*, volume 9954 of *Lecture Notes in Computer Science*, pages 260–272, 2016.

26 Both Emerite Neou, Romeo Rizzi, and Stéphane Vialette. Permutation pattern matching in (213, 231)-avoiding permutations. *Discret. Math. Theor. Comput. Sci.*, 18(2), 2016.