A Complexity Dichotomy for Permutation Pattern Matching on Grid Classes

Vít Jelínek
Computer Science Institute, Charles University, Prague, Czech Republic
jelinek@iuuk.mff.cuni.cz

Michal Opler
Computer Science Institute, Charles University, Prague, Czech Republic
opler@iuuk.mff.cuni.cz

Jakub Pekárek
Department of Applied Mathematics, Charles University, Prague, Czech Republic
pekarej@kam.mff.cuni.cz

Abstract

Permutation Pattern Matching (PPM) is the problem of deciding for a given pair of permutations π and τ whether the pattern π is contained in the text τ. Bose, Buss and Lubiw showed that PPM is NP-complete. In view of this result, it is natural to ask how the situation changes when we restrict the pattern π to a fixed permutation class $\mathcal{C}$; this is known as the $\mathcal{C}$-Pattern PPM problem. There have been several results in this direction, namely the work of Jelínek and Kynčl who completely resolved the hardness of $\mathcal{C}$-Pattern PPM when $\mathcal{C}$ is taken to be the class of σ-avoiding permutations for some σ.

Grid classes are special kind of permutation classes, consisting of permutations admitting a grid-like decomposition into simpler building blocks. Of particular interest are the so-called *monotone grid classes*, in which each building block is a monotone sequence. Recently, it has been discovered that grid classes, especially the monotone ones, play a fundamental role in the understanding of the structure of general permutation classes. This motivates us to study the hardness of $\mathcal{C}$-Pattern PPM for a (monotone) grid class $\mathcal{C}$.

We provide a complexity dichotomy for $\mathcal{C}$-Pattern PPM when $\mathcal{C}$ is taken to be a monotone grid class. Specifically, we show that the problem is polynomial-time solvable if a certain graph associated with $\mathcal{C}$, called the cell graph, is a forest, and it is NP-complete otherwise. We further generalize our results to grid classes whose blocks belong to classes of bounded grid-width. We show that the $\mathcal{C}$-Pattern PPM for such a grid class $\mathcal{C}$ is polynomial-time solvable if the cell graph of $\mathcal{C}$ avoids a cycle or a certain special type of path, and it is NP-complete otherwise.

2012 ACM Subject Classification Mathematics of computing → Permutations and combinations; Theory of computation → Pattern matching; Theory of computation → Problems, reductions and completeness

Keywords and phrases permutations, pattern matching, grid classes

Digital Object Identifier 10.4230/LIPIcs.MFCS.2020.52

Related Version `http://arxiv.org/abs/2008.04593`

Funding Supported by the GAUK project 1766318.
Vít Jelínek: Supported by project 18-19158S of the Czech Science Foundation.
Michal Opler: Supported by project 18-19158S of the Czech Science Foundation.

1 Introduction

A *permutation* is a sequence $\pi = \pi_1, \pi_2, \ldots, \pi_n$ in which each number from the set $[n] = \{1, 2, \ldots, n\}$ appears exactly once. We then say that a permutation $\pi = \pi_1, \ldots, \pi_n$ *contains* a permutation $\sigma = \sigma_1, \ldots, \sigma_k$, if π has a subsequence of length k whose elements have the same relative order as the elements of σ (see Section 2 for a more formal definition). If π does not contain σ, we say that π *avoids* σ.

45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020).
Editors: Javier Esparza and Daniel Král'; Article No. 52; pp. 52:1–52:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

An essential algorithmic problem involving permutations is Permutation Pattern Matching (PPM): Given a permutation π ("pattern") of size k and τ ("text") of size n, does τ contain π? Bose, Buss and Lubiw [10] have shown that PPM is NP-complete. This result has motivated a study of various variants of PPM , in particular to obtain the best possible runtime dependence on k. Guillemot and Marx [15] provided the break-through result in this direction by establishing the fixed-parameter tractability of PPM in terms of the pattern length with algorithm running in $2^{O(k^2 \log k)} \cdot n$ time. The first phase of their algorithm finds a suitable decomposition that relies on the proof of Stanley-Wilf conjecture given by Marcus and Tardos [21]. Subsequently, Fox [14] refined the results by Marcus and Tardos and thereby reduced the complexity of the algorithm to $2^{O(k^2)} \cdot n$.

Several algorithms also exist whose runtimes depend on different parameters than the length of π. Bruner and Lackner [12] described an algorithm for PPM with run time $O(1.79^{\mathrm{run}(\tau)} \cdot kn)$ where $\mathrm{run}(\tau)$ is the number of consecutive monotone sequences needed to obtain τ via concatenation. Ahal and Rabinovich [1] designed an algorithm for PPM that runs in time $n^{O(\mathrm{tw}(G_\pi))}$ where G_π is a certain graph associated to the pattern π and $\mathrm{tw}(G_\pi)$ denotes the treewidth of G_π. Later Jelínek, Opler and Valtr [19] introduced a related parameter $\mathrm{gw}(\pi)$, called the *grid-width* of π, and showed that $\mathrm{gw}(\pi)$ is equivalent to $\mathrm{tw}(G_\pi)$ up to a constant and thereby implying an algorithm for PPM running in time $O(n^{O(\mathrm{gw}(\pi))})$.

A *permutation class* is a set of permutations $\mathcal{C}$ such that for every $\sigma \in \mathcal{C}$, every pattern contained in σ belongs to $\mathcal{C}$ as well. For a permutation σ, we let $\mathrm{Av}(\sigma)$ denoted the class of permutations avoiding σ. Different approach to tackling the hardness of PPM is to restrict the choice of pattern to a particular permutation class $\mathcal{C}$.

$\mathcal{C}$-Pattern Permutation Pattern Matching ($\mathcal{C}$-Pattern PPM)	
Input:	A pattern $\pi \in \mathcal{C}$ of size k and a permutation τ of size n.
Question:	Does τ contain π?

Notice that Av(21)-Pattern PPM simply reduces to finding the longest increasing subsequence of τ, which is a well-known problem and can be solved in time $O(n \log \log n)$ [20]. Bose, Buss and Lubiw [10] showed that $\mathcal{C}$-Pattern PPM can be solved in polynomial time if $\mathcal{C}$ is the class of the so-called separable permutations. Further improvements were given by Ibarra [17], Albert et al. [3] and by Yugandhar and Saxena [28]. Recently, Jelínek and Kynčl [18] completely resolved the hardness of $\mathcal{C}$-Pattern PPM for classes avoiding a single pattern. They proved that $\mathrm{Av}(\alpha)$-Pattern PPM is polynomial-time solvable for $\alpha \in \{1, 12, 21, 132, 213, 231, 312\}$ and NP-complete otherwise.

Lately, a new type of permutation classes has been gaining a lot of attention. The *grid class* of a matrix $\mathcal{M}$ whose entries are permutation classes, denoted by $\mathrm{Grid}(\mathcal{M})$, is a class of permutations admitting a grid-like decomposition into blocks that belong to the classes $\mathcal{M}_{i,j}$. If moreover $\mathcal{M}$ contains only $\mathrm{Av}(21), \mathrm{Av}(12)$ and $\emptyset$ we say that $\mathrm{Grid}(\mathcal{M})$ is a *monotone grid class*. To each matrix $\mathcal{M}$ we also associate a graph $G_{\mathcal{M}}$, called the *cell graph* of $\mathcal{M}$. We postpone their full definitions to Section 2.

Monotone grid classes were introduced partly by Atkinson, Murphy and Ruškuc [6] and in full by Murphy and Vatter [22] who showed that a monotone grid class $\mathrm{Grid}(\mathcal{M})$ is partially well-ordered if and only if $G_{\mathcal{M}}$ is a forest. Brignall [11] later extended their results to a large portion of general grid classes. General grid classes themselves were introduced by Vatter [25] in his paper investigating the growth rates of permutation classes. Since then the grid classes played a central role in most consequent works on growth rates of permutations classes [4, 27]. Bevan [8, 9] tied the growth rates of grid classes to algebraic graph theory.

Given the prominent role of grid classes in recent developments of the permutation pattern research, it is only natural to investigate the hardness of searching patterns that belong to a grid class. Neou, Rizzi and Vialette [24] designed a polynomial-time algorithm solving $\mathcal{C}$-Pattern PPM when $\mathcal{C}$ is the class of the so-called wedge permutations, which can be also described as a monotone grid class. Consequently, Neou [23] asks at the end of his thesis about the hardness of Grid($\mathcal{M}$)-Pattern PPM for a monotone grid class Grid($\mathcal{M}$).

Our results. In Section 3, we answer the question of Neou by proving that for a monotone grid class Grid($\mathcal{M}$), the problem Grid($\mathcal{M}$)-Pattern PPM is polynomial-time solvable if $G_{\mathcal{M}}$ is a forest, NP-complete otherwise.

In Section 4, we further extend our results to matrices whose every entry is a permutation class of bounded grid-width. We prove that for such grid classes, Grid($\mathcal{M}$)-Pattern PPM is polynomial-time solvable if $G_{\mathcal{M}}$ is a forest which does not contain a certain type of path, and NP-complete otherwise.

2 Preliminaries

A *permutation of length* n is a sequence in which each element of the set $[n] = \{1, 2, \ldots, n\}$ appears exactly once. When writing out short permutations explicitly, we shall omit all punctuation and write, e.g., 15342 for the permutation $1, 5, 3, 4, 2$. The *permutation diagram* of π is the set of points $\{(i, \pi_i);\ i \in [n]\}$ in the plane. Note that we use Cartesian coordinates, that is, the first row of the diagram is at the bottom. We blur the distinction between permutations and their permutation diagrams, e.g. we shall refer to the point set π.

For a point p in the plane, we denote its first coordinate as $p.x$, and its second coordinate as $p.y$. A subset S of a permutation diagram is *isomorphic* to a subset R of a permutation diagram if there is a bijection $f : R \to S$ such that for any pair of points $p \neq q$ of R we have $f(p).x < f(q).x$ if and only if $p.x < q.x$, and $f(p).y < f(p).y$ if and only if $p.y < q.y$. A permutation π of length n *contains* a permutation σ of length k, if there is a subset of π isomorphic to the permutation diagram of σ. Such a subset is then an *occurrence* (or a *copy*) of σ in π. If π does not contain σ, we say that π *avoids* σ.

A *permutation class* is a set $\mathcal{C}$ of permutations with the property that if π is in $\mathcal{C}$, then all the permutations contained in π are in $\mathcal{C}$ as well. For a permutation σ, we let $\mathrm{Av}(\sigma)$ denote the class of σ-avoiding permutations. We shall sometimes use the symbols ⧄ and ⧅ as short-hands for the class of increasing permutations Av(21) and the class of decreasing permutations Av(12)

For a permutation π of length n the *reverse* of π is the permutation $\pi_n, \pi_{n-1}, \ldots, \pi_1$, the *complement* of π is the permutation $n+1-\pi_1, n+1-\pi_2, \ldots, n+1-\pi_n$, and the *inverse* of π is the permutation $\sigma = \sigma_1, \ldots, \sigma_n$ satisfying $\pi_i = j \iff \sigma_j = i$. We let π^r, π^c and π^{-1} denote the reverse, complement and inverse of π, respectively. Similarly, for a permutation class $\mathcal{C}$, we let $\mathcal{C}^r$ denote the set $\{\pi^r;\ \pi \in \mathcal{C}\}$, and similarly for $\mathcal{C}^c$ and $\mathcal{C}^{-1}$. Note that $\mathcal{C}^r$, $\mathcal{C}^c$ and $\mathcal{C}^{-1}$ are again permutation classes.

A permutation π of length n is a *horizontal alternation* if all the even entries of π precede all the odd entries of π, i.e. there are no indices $i < j$ such that π_i is odd and π_j is even. A permutation π is a *vertical alternation* if π^{-1} is a horizontal alternation.

Griddings and grid classes

A $k \times \ell$ matrix $\mathcal{M}$ whose entries are permutation classes is called a *gridding matrix.* Moreover, if the entries of $\mathcal{M}$ belong to the set $\{⧄, ⧅, \emptyset\}$ then we say that $\mathcal{M}$ is a *monotone gridding matrix.* A $k \times \ell$*-gridding* of a permutation π of length n are two sequences of (possibly empty)

$$\mathcal{M} = \begin{pmatrix} \mathrm{Av}(132) & & \mathrm{Av}(21) \\ & \mathrm{Av}(321) & \mathrm{Av}(12) \end{pmatrix}$$

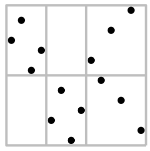

Figure 1 A gridding matrix $\mathcal{M}$ on the left and an $\mathcal{M}$-gridded permutation on the right. Empty entries of $\mathcal{M}$ are omitted and the edges of $G_{\mathcal{M}}$ are displayed inside $\mathcal{M}$.

disjoint integer intervals $I_1 < I_2 < \cdots < I_k$ and $J_1 < J_2 < \cdots < J_\ell$ such that both $\bigcup_{i=1}^k I_i$ and $\bigcup_{j=1}^\ell J_j$ are equal to $[n]$. We call the set $I_i \times J_j$ an (i, j)*-cell* of π. An $\mathcal{M}$*-gridding* of a permutation π is a $k \times \ell$-gridding such that the restriction of π to the (i, j)-cell is isomorphic to a permutation from the class $\mathcal{M}_{i,j}$. If π possesses an $\mathcal{M}$-gridding, then π is said to be $\mathcal{M}$*-griddable* and π equipped with a fixed $\mathcal{M}$-gridding is called an $\mathcal{M}$*-gridded permutation*. We let $\mathrm{Grid}(\mathcal{M})$ be the class of $\mathcal{M}$-griddable permutations. Note that for consistency with our Cartesian numbering convention, we number the rows of a matrix from bottom to top.

The *cell graph* of the gridding matrix $\mathcal{M}$ is the graph $G_{\mathcal{M}}$ whose vertices are the cells of $\mathcal{M}$ that contain an infinite class, with two vertices being adjacent if they share a row or a column of $\mathcal{M}$ and all cells between them are empty. A *proper turning path* in $G_{\mathcal{M}}$ is a path P such that no three consecutive cells of P share the same row or column. See Figure 1.

Grid-width

An *interval family* $\mathcal{I}$ is a set of pairwise disjoint integer intervals. The *intervalicity* of a set $A \subseteq [n]$, denoted by $\mathrm{int}(A)$, is the size of the smallest interval family whose union is equal to A. For a point set S in the plane, let $\Pi_x(S)$ denote its projection on the x-axis and similarly $\Pi_y(S)$ its projection on the y-axis. For a subset S of the permutation diagram, the *grid-complexity* of S is the maximum of $\mathrm{int}(\Pi_x(S))$ and $\mathrm{int}(\Pi_y(S))$.

A *grid tree* of a permutation π of length n is a rooted binary tree T with n leaves, each leaf being labeled by a distinct point of the permutation diagram. Let π_v^T denote the point set of the labels on the leaves in the subtree of T rooted in v. The *grid-width* of a vertex v in T is the grid-complexity of π_v^T, and the *grid-with* of T, denoted by $\mathrm{gw}^T(\pi)$, is the maximum grid-width of a vertex of T. Finally, the *grid-width* of a permutation π, denoted by $\mathrm{gw}(\pi)$, is the minimum of $\mathrm{gw}^T(\pi)$ over all grid trees T of π.

We also consider a linear version of this parameter. We say that a rooted binary tree T is a *caterpillar* if each vertex is either a leaf or has at least one leaf as a child. The *path-width* of a permutation π, denoted by $\mathrm{pw}(\pi)$, is the minimum of $\mathrm{gw}^T(\pi)$ over all caterpillar grid trees T of π.

We now provide a useful alternative definition of path-width. For permutations π and σ of length n, the *path-width of π in σ-ordering*, denoted by $\mathrm{pw}^\sigma(\pi)$ is the maximum grid-complexity attained by a set $\{(\sigma_1, \pi_{\sigma_1}), \dots, (\sigma_i, \pi_{\sigma_i})\}$ for some $i \in [n]$.

▶ **Lemma 2.1.** *A permutation π of length n has path-width p if and only if the minimum value of $\mathrm{pw}^\sigma(\pi)$ over all permutations σ of length n is exactly p.*

Ahal and Rabinovich [1] designed an algorithm for PPM that runs in time $n^{O(\mathrm{tw}(G_\pi))}$ where G_π is a certain graph associated to the pattern π and $\mathrm{tw}(G_\pi)$ denotes the treewidth of G_π. The following theorem follows by combining this algorithm with the result of Jelínek et al. [19] who showed that up to a constant, $\mathrm{gw}(\pi)$ is equivalent to $\mathrm{tw}(G_\pi)$.

▶ **Theorem 2.2** (Ahal and Rabinovich [1], Jelínek et al. [19]). *Let π be a permutation of length k and τ a permutation of length n. The problem whether τ contains π can be solved in time $n^{O(\mathrm{gw}(\pi))}$.*

Importantly, Theorem 2.2 implies that $\mathcal{C}$-Pattern PPM is decidable in polynomial time whenever the class $\mathcal{C}$ has bounded grid-width. In fact, we obtain all the polynomial-time solvable cases of $\mathcal{C}$-Pattern PPM in this paper via showing that $\mathcal{C}$ has bounded grid-width.

3 Monotone grid classes

This section is dedicated to proving that the complexity of Grid($\mathcal{M}$)-Pattern PPM is for a monotone gridding matrix $\mathcal{M}$ determined by whether $G_{\mathcal{M}}$ contains a cycle.

▶ **Theorem 3.1.** *For a monotone gridding matrix $\mathcal{M}$ one of the following holds:*

- *Either $G_{\mathcal{M}}$ is a forest, Grid($\mathcal{M}$) has bounded path-width and Grid($\mathcal{M}$)-Pattern PPM can be decided in polynomial time, or*
- *$G_{\mathcal{M}}$ contains a cycle, Grid($\mathcal{M}$) has unbounded grid-width and Grid($\mathcal{M}$)-Pattern PPM is NP-complete.*

A *consistent orientation* of a $k \times \ell$ monotone gridding matrix $\mathcal{M}$ is a pair of functions (c, r) such that $c : [k] \to \{-1, 1\}$, $r : [\ell] \to \{-1, 1\}$ and for every $i \in [k]$, $j \in [\ell]$ the value $c(i)r(j)$ is positive if $\mathcal{M}_{i,j} = ⧄$ and negative if $\mathcal{M}_{i,j} = ⧅$.

It was observed by Albert et al. [4] that we can without loss of generality assume that $\mathcal{M}$ has a consistent orientation. Let $\mathcal{M}$ be a $k \times \ell$ monotone gridding matrix and q a positive integer. The *refinement* $\mathcal{M}^{\times q}$ of $\mathcal{M}$ is the $qk \times q\ell$ matrix obtained from $\mathcal{M}$ by replacing each ⧄-entry by a $q \times q$ diagonal matrix with all the non-empty entries equal to ⧄, each ⧅-entry by a $q \times q$ anti-diagonal matrix with all the non-empty entries equal to ⧅ and each empty entry by a $q \times q$ empty matrix. It is easy to see that Grid($\mathcal{M}^{\times q}$) is a subclass of Grid($\mathcal{M}$). Moreover, if $G_{\mathcal{M}}$ is a forest then Grid($\mathcal{M}^{\times q}$) = Grid($\mathcal{M}$) and $G_{\mathcal{M}^{\times q}}$ is a forest as well.

▶ **Lemma 3.2** (Albert et al. [4, Proposition 4.1]). *For every monotone gridding matrix $\mathcal{M}$, the refinement $\mathcal{M}^{\times 2}$ admits a consistent orientation.*

We remark that Albert et al. [4] use a slightly different way of defining a permutation class from a given gridding matrix: specifically, their classes only contain permutations in which the entries represented by each cell of the gridding can be placed on a segment with slope $+1$ or -1. However, as Lemma 3.2 is a claim about gridding matrices and not about permutation classes, we can use it here for our purposes.

Now we provide bounds on the width parameters of monotone grid classes depending on the structure of their cell graphs.

▶ **Proposition 3.3.** *Let $\mathcal{M}$ be a $k \times \ell$ monotone gridding matrix that has consistent orientation. If $G_{\mathcal{M}}$ is a forest then every permutation of Grid($\mathcal{M}$) has path-width at most $\max(k, \ell)$.*

Proof. Let π be a permutation from Grid($\mathcal{M}$) of length n together with an $\mathcal{M}$-gridding $I_1^0 < \cdots < I_k^0$ and $J_1^0 < \cdots < J_\ell^0$. We fix a consistent orientation (c, r) for the matrix $\mathcal{M}$.

For $I \subseteq I_i^0$, the *extremal point* of I is the rightmost point of π restricted to $I \times [n]$ if $c(i)$ is positive, the leftmost point otherwise. Similarly for $J \subseteq J_j^0$, the extremal point of J is the topmost point of π restricted to $[n] \times J$ if $r(j)$ is positive, the bottommost point otherwise. Importantly, the definition of consistent orientation guarantees that if $I \times J$ contains the extremal points of both I and J, then these two points must actually be the same point.

We now construct an ordering σ of length n and a sequence of interval families $(\mathcal{I}^m, \mathcal{J}^m)$ for $m \in [n]$ such that for every m

- $\mathcal{I}^m$ contains k (possibly empty) intervals $I_1^m < \dots < I_k^m$ and $\mathcal{J}^m$ contains ℓ (possibly empty) intervals $J_1^m < \dots < J_\ell^m$,
- $I_s^m \subseteq I_s^0$ and $J_t^m \subseteq J_t^0$ for every s and t, and
- for $P = \{(\sigma_{m+1}, \pi_{\sigma_{m+1}}), (\sigma_{m+2}, \pi_{\sigma_{m+2}}), \dots, (\sigma_n, \pi_{\sigma_n})\}$ we have $\Pi_x(P) = \bigcup \mathcal{I}^m$ and $\Pi_y(P) = \bigcup \mathcal{J}^m$.

From the third condition then follows that $\mathrm{pw}^{\sigma^R}(\pi) \leq \max(k, \ell)$ which proves the proposition.

Suppose that we have already defined the sequences up to m. Let G^m be an auxiliary oriented bipartite graph whose vertices are $s_1, \dots, s_k, t_1, \dots, t_\ell$, and there is an edge (s_i, t_j) whenever the extremal point of I_i^m lies in the j-th row of the $\mathcal{M}$-gridding and an edge (t_j, s_i) whenever the extremal point of J_j^m lies in the i-th column of the $\mathcal{M}$-gridding.

Note that every vertex corresponding to a non-empty row has exactly one outgoing edge to a non-empty column and vice versa. Therefore, we can find a cycle in G^m by starting in arbitrary non-empty column and following the outgoing edges. However, observe that there cannot be a cycle of length larger than 2 as such cycle would imply a cycle in the graph $G_{\mathcal{M}}$. Thus, there exists a column i and a row j such that the extremal point of I_i^m lies in J_j^m and vice versa. As we observed, this must be the same point p. We set σ_{m+1} to be $p.x$ and we define the interval families $\mathcal{I}^{m+1}$, $\mathcal{J}^{m+1}$ by removing respective coordinates of p from $\mathcal{I}^m$ and $\mathcal{J}^m$. Observe that $\mathcal{I}^{m+1}$, $\mathcal{J}^{m+1}$ are well-defined because p was the extremal point of both intervals I_i^m and J_j^m. ◀

▶ **Lemma 3.4.** *Let $\mathcal{M}$ be a monotone gridding matrix such that $G_{\mathcal{M}}$ contains a path of length k. Then there exists a permutation $\pi \in Grid(\mathcal{M})$ with grid-width at least $\frac{k}{4}$.*

Proof. Without loss of generality, we assume that $G_{\mathcal{M}}$ is a path of length k, i.e. all the other entries are empty. We construct an $\mathcal{M}$-gridded permutation π of length k^3 such that $\mathrm{gw}(\pi) \geq \frac{k}{4}$ in the following way. The permutation π is a union of point sets $B_1, \dots, B_k$, called blocks, such that each B_i is contained in the cell of the $\mathcal{M}$-gridding corresponding to the i-th vertex of the path. Moreover for every i, the point set $B_i \cup B_{i+1}$ forms a vertical or horizontal alternation, depending on whether their respective cells share the same row or column. If the relation between point sets B_i and B_j for $|i - j| \geq 2$ is not fully determined by the position of their respective cells (they share the same row or column) they can be interleaved in arbitrary way as long as $B_k \cup B_{k+1}$ forms an alternation for every k.

For a contradiction, suppose that $\mathrm{gw}(\pi) < \frac{k}{4}$ and fix the corresponding grid tree T. By standard arguments there is a vertex $v \in T$ such that the subtree of v contains at least $\frac{n}{3}$ and at most $\frac{2n}{3}$ leaves. Let S be the subset of permutation diagram π defined by these leaves.

We first show that each block of π contains both a point in S and a point outside of S. Let the density of the block B_i be the ratio $\frac{|S \cap B_i|}{|B_i|}$, denoted by d_i. We claim that the densities of consecutive blocks cannot differ too much, in particular that the difference $|d_i - d_{i+1}|$ is at most $\frac{1}{4k}$. Without loss of generality assume that B_i and B_{i+1} share the same column and that $d_{i+1} > d_i$. If the density of B_i and B_{i+1} differed by at least $\frac{1}{4k}$, there would be at least $k/4$ more points in B_{i+1}. Since $B_i \cup B_{i+1}$ forms a vertical alternation, we could pair each point of B_i with the nearest point to the right of it in B_{i+1}. Then at least $k/4$ of these pairs would consist of a point of B_{i+1} in S and a point of B_i outside of S. The intervalicity of $\Pi_x(S \cap (B_i \cup B_{i+1}))$ would thus be at least $k/4$ and that is a contradiction.

Now suppose that block B_i is fully contained in S, i.e. the density d_i is equal to 1. Since the differences in densities of consecutive blocks cannot be larger than $\frac{1}{4k}$, every block has density at least $1 - k\frac{1}{4k} = \frac{3}{4}$. But then S must contain at least $\frac{3}{4}n$ of points, which is a contradiction. For the same reason, there cannot be any block whose density is equal to 0.

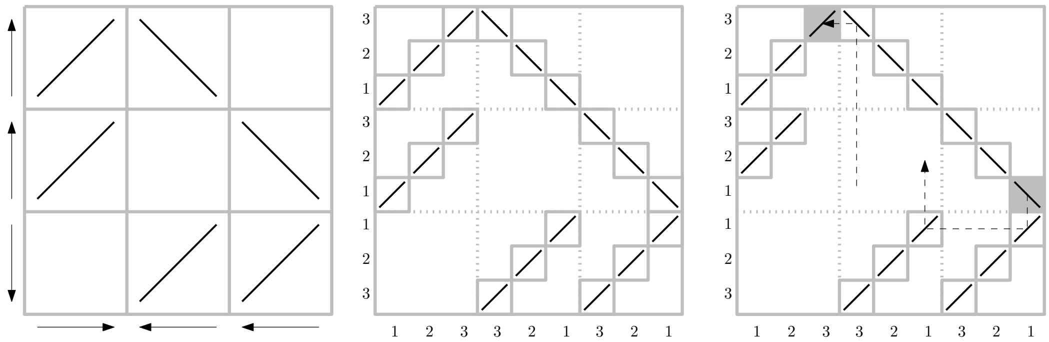

Figure 2 The matrix modification process in the proof of Lemma 3.5. A gridding matrix $\mathcal{M}$ with a consistent orientation on the left, its refinement $\mathcal{M}^{\times 3}$ with labeled rows and columns in the middle and the gridding matrix $\mathcal{M}_3$ on the right. The endpoints of path in $G_{\mathcal{M}_3}$ are highlighted.

Therefore, every block B_i contains both a point from S and a point outside of S. The number of rows and columns of $\mathcal{M}$ spanned by the path is at least k since every step on the path introduces either a new row or a new column. Let us therefore, without loss of generality, assume that the path spans at least $\frac{k}{2}$ columns. Since each of the columns contains point both from S and outside of S, any consecutive interval of $\Pi_x(S)$ cannot intersect more than 2 of these columns and the grid-complexity of S is at least $\frac{k}{4}$. ◀

We proceed to show that whenever $G_{\mathcal{M}}$ contains a cycle, Grid($\mathcal{M}$) contains grid subclasses with arbitrarily long paths in their cell graph. In fact, all the properties we show about cyclic grid classes rely structurally only on the existence of these long paths.

▶ **Lemma 3.5.** *Let $\mathcal{M}$ be a monotone gridding matrix such that $G_{\mathcal{M}}$ contains a cycle. For every $p \geq 1$, there is a gridding matrix $\mathcal{M}_p$ such that Grid($\mathcal{M}$) contains Grid($\mathcal{M}_p$) as a subclass and moreover $G_{\mathcal{M}_p}$ is a proper turning path of length at least p. Furthermore, given $\mathcal{M}$ and the integer p we can compute $\mathcal{M}_p$ in polynomial time.*

Proof. Let C be a cycle in $G_{\mathcal{M}}$. We can without loss of generality assume that C is proper turning and that every cell outside of C is actually empty, otherwise we could replace all the cells of $\mathcal{M}$ that do not correspond to the turns of C by empty cells.

The proof is illustrated in Figure 2. Fix a consistent orientation (c, r) of $\mathcal{M}$ and recall the definition of the refinement $\mathcal{M}^{\times p}$. We proceed by labeling the rows and columns of $\mathcal{M}^{\times p}$ using the set $[p]$. The p-tuple of columns created from the i-th column of $\mathcal{M}$ is labeled $1, 2, \ldots, p$ from left to right if $c(i)$ is positive, and right to left otherwise. Similarly, the p-tuple of rows created from the j-th row of $\mathcal{M}'$ is labeled $1, 2, \ldots, p$ from bottom to top if $r(j)$ is positive, and top to bottom otherwise. The *characteristic of a cell* in $\mathcal{M}^{\times p}$ is the pair of labels given to its column and row. The consistent orientation guarantees that each non-empty cell in $\mathcal{M}^{\times p}$ has a characteristic of form (s, s) for some $s \in [p]$. Therefore, $\mathcal{M}^{\times p}$ consists exactly of p components, each being a copy of $\mathcal{M}$.

The (i, j)-*block* of $\mathcal{M}^{\times p}$ is the $p \times p$ submatrix corresponding to the (i, j)-cell of $\mathcal{M}$. We pick an arbitrary non-empty cell (i, j) of $\mathcal{M}$ and obtain a matrix $\mathcal{M}_p$ by replacing (i, j)-block in $\mathcal{M}^{\times p}$ with the matrix whose only non-empty entries have the characteristic $(s, s + 1)$ for all $s \in [p - 1]$ and are of the same type as $\mathcal{M}_{i,j}$. Grid($\mathcal{M}_p$) is a subclass of Grid($\mathcal{M}$) since the modified (i, j)-block corresponds to shifting the original (anti-)diagonal matrix by one row either up or down, depending on the orientation of the j-th row of $\mathcal{M}$.

Observe that we connected all the p copies of $\mathcal{M}$ into a single long path. Moreover, we in fact described an algorithm how to compute $\mathcal{M}_p$ in polynomial time. ◀

Combining Lemmas 3.4 and 3.5, we directly obtain the following corollary.

▶ **Corollary 3.6.** *Let $\mathcal{M}$ be a monotone gridding matrix such that $G_{\mathcal{M}}$ contains a cycle. Then Grid($\mathcal{M}$) has unbounded grid-width.* ◀

In order to state and prove our hardness result that contrasts Proposition 3.3, we need to introduce few definitions. Let $\mathcal{C}$ and $\mathcal{D}$ be permutation classes and k a positive integer. The $(\mathcal{C}, \mathcal{D})$*-staircase of k steps* is the $(k+1) \times k$ gridding matrix $\mathrm{St}^k(\mathcal{C}, \mathcal{D})$ such that the (i, i)-cell contains $\mathcal{C}$, the $(i+1, i)$-cell contains $\mathcal{D}$ for every $i \in [k]$, and every other cell is empty. The staircase classes have been studied by Albert et al. [2] in the context of determining the growth rates of certain permutation classes.

Let $\mathcal{M}, \mathcal{N}$ be $k \times \ell$ gridding matrices, let π be an $\mathcal{M}$-gridded permutation and let τ an $\mathcal{N}$-gridded permutation. We say that τ contains a *grid-preserving copy* of π, if there is an occurrence of π in τ such that the elements from the (i, j)-cell of the $\mathcal{M}$-gridding of π are mapped to elements in the (i, j)-cell of the $\mathcal{N}$-gridding of τ for every i and j.

The final missing piece toward the Theorem 3.1 is showing that whenever $G_{\mathcal{M}}$ contains a cycle, Grid($\mathcal{M}$)-Pattern PPM is NP-complete. To that end, we need to inspect the results proved by Jelínek and Kynčl [18]. We remark that the following theorem describes only an intermediate step of their proof showing that Av(321)-Pattern PPM is NP-complete.

▶ **Theorem 3.7** (Jelínek and Kynčl [18]). *Let Φ be a 3-CNF formula with v variables and c clauses. There is a polynomial time algorithm that outputs a $St^{2c+1}(\boxslash, \boxslash)$-gridded permutation π and a $St^{2c+1}(Av(321), \boxslash)$-gridded permutation τ such that Φ is satisfiable if and only if there is a grid-preserving copy of π in τ. Additionally, the longest increasing subsequences of the $(1,1)$-cell of π and the $(1,1)$-cell of τ are both of length $2v$.*

Let us modify π and τ such that any embedding that maps the $(1, 1)$- cell of π to the $(1, 1)$-cell of τ must already be grid-preserving.

The *lane of k steps* is the $\mathrm{St}^k(\boxslash, \boxslash)$-gridded permutation such that each non-empty cell of the staircase contains exactly 2 points and two neighboring cells in the same row form a copy of 1423 while two neighboring cells in the same column form a copy of 1342. The intuition here is that as the lane intersects two adjacent cells in the same row, the two elements in the left cell (corresponding to 1 and 4 of the pattern 1423) are "sandwiching" the two elements of the right cell from above and from below. Similarly, for two cells in the same column, the bottom part of the lane sandwiches the top part from the left and from the right.

For a $\mathrm{St}^k(\mathcal{C}, \mathcal{D})$-gridded permutation π, the result of *confining* π is the following gridded permutation π'. Let $\mathcal{N}$ be the gridding matrix obtained by replacing every non-empty entry of $\mathrm{St}^k(\mathcal{C}, \mathcal{D})$ with a 3×3 diagonal matrix whose (1,1)-cell and (3,3)-cell contain $\boxslash$ and the $(2, 2)$-cell is equal to the original entry. Observe that $\mathcal{N}$ consists of 3 components, namely a copy of $\mathrm{St}^k(\mathcal{C}, \mathcal{D})$ "sandwiched" between two copies of $\mathrm{St}^k(\boxslash, \boxslash)$. We obtain π' by placing π in the middle copy and two lanes of k steps in the outer copies. Finally, we unify each 3×3 block of the $\mathcal{N}$-gridded π' into a single cell. See the left part of Figure 3.

Let π and τ be the gridded permutations from Theorem 3.7 and π', τ' their confined versions. Suppose there is an occurrence of π' in τ' that maps the $(1, 1)$-cell of π' to the $(1, 1)$-cell of τ'. It must map the first two points of the two lanes in π' to the first two points of the lanes in τ' since the longest increasing subsequences in the $(1, 1)$-cells of π and τ are of the same length. But then the whole lanes from π' map to the respective lanes in τ'; this is because the pair of elements in a given cell of the lane in π' (or τ') sandwiches exactly two other elements of π' (or τ'), namely the two elements of the same line belonging to the following cell. This then forces the elements of each non-empty (i, j)-cell of π' to map to the (i, j)-cell of τ'.

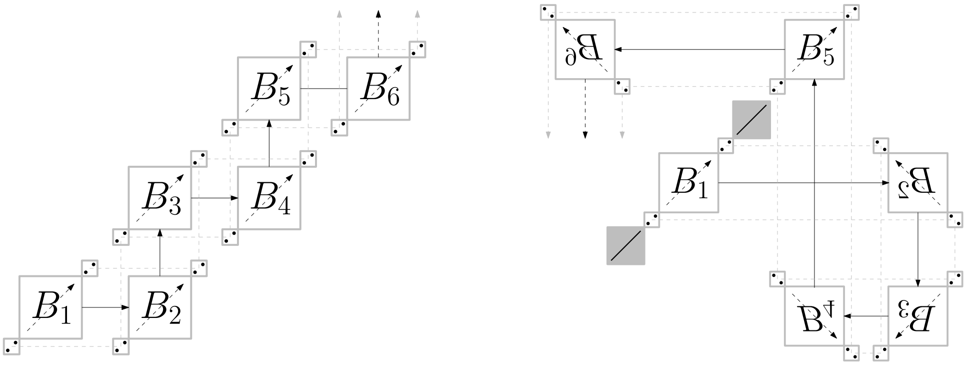

Figure 3 Situation in the proof of Proposition 3.8. The confined permutation π on the left, the permutation $\pi^\star$ obtained by adding anchors (highlighted) to the (f,g)-transform π' on the right.

▶ **Proposition 3.8.** *Let $\mathcal{M}$ be a monotone gridding matrix such that $G_{\mathcal{M}}$ contains a cycle. Then* Grid$(\mathcal{M})$-Pattern *PPM is NP-complete.*

Proof. We describe the reduction from 3-SAT to Grid$(\mathcal{M})$-Pattern PPM. Let Φ be a given 3-CNF formula with v variables and c clauses. Using Theorem 3.7, we compute gridded permutations π and τ such that Φ is satisfiable if and only if there is a grid-preserving copy of π in τ. As we have shown, we can assume that any occurrence of π in τ that maps the $(1,1)$-cell of π to the $(1,1)$-cell of τ must be grid-preserving. Furthermore, we obtain a monotone gridding matrix $\mathcal{M}'$ such that Grid$(\mathcal{M}')$ is a subclass of Grid$(\mathcal{M})$ and $G_{\mathcal{M}'}$ is a proper turning path P of length $4c+2$ by application of Lemma 3.5. Without loss of generality we may assume that the first two cells of this path occupy the same row.

First, we aim to construct gridded permutations π' and τ' such that $\pi' \in$ Grid$(\mathcal{M}')$ and there is a grid-preserving copy of π' in τ' if and only if there is a grid-preserving copy of π in τ. We essentially aim to generalize the "twirl" operation used by Jelínek and Kynčl [18]. A *signed permutation of length k* is a permutation of length k in which each entry is additionally equipped with a sign. Let $\mathcal{Q}$ be an arbitrary $k \times \ell$ gridding matrix and let f be a signed permutation of length k and g be a signed permutation of length ℓ. The (f,g)*-transform* of $\mathcal{Q}$ is a gridding matrix $\mathcal{Q}'$ such that $\mathcal{Q}'_{i,j} = (\mathcal{Q}_{|f(i)|,|g(j)|})^o$ where the operation o is identity if both $f(i)$ and $g(j)$ are positive, reversal if only $f(i)$ is negative, complement if only $g(j)$ is negative and reverse complement if both $f(i)$ and $g(j)$ are negative. In other words, $\mathcal{Q}'$ is created by permuting the rows and columns of $\mathcal{Q}$ according to the permutations f and g while also flipping around those with negative signs.

Analogously, we can define a transformation of a $\mathcal{Q}$-gridded permutation into a $\mathcal{Q}'$-gridded one. The (f,g)*-transform* of a $\mathcal{Q}$-gridded permutation σ is the gridded permutation σ' obtained by permuting the columns of the gridding according to f, rows according to g, replacing the point set of the i-th column with its reversal if $f(i)$ is negative and then replacing the point set of the j-th row with its complement if $g(j)$ is negative.

For any fixed f and g, there is indeed a grid-preserving occurrence of π in τ if and only if there is a grid-preserving occurrence of the (f,g)-transform of π in the (f,g)-transform of τ. Let us now define f and g such that the (f,g)-transform of $\mathrm{St}^{2c+1}(⍁,⍁)$ is exactly the gridding matrix $\mathcal{M}'$. Let (c,r) be a consistent orientation of $\mathcal{M}'$, let $s_1,\ldots,s_{2c+2}$ be the sequence of column indices in the order as visited by the path P and let $t_1,\ldots,t_{2c+1}$ be the sequence of row indices in the order as visited by P. We set $f(s_i) = c(i)i$ and $g(t_j) = r(j)j$. Observe that the i-th cell of the path in the staircase is mapped by the (f,g)-transform precisely to the i-th cell of P. Moreover, the consistent orientation guarantees that the type of the entry obtained via (f,g)-transform agrees with the type of the entry in $\mathcal{M}'$. Let $\mathcal{N}$ be the (f,g)-transform of $\mathrm{St}^{2c+1}(\mathrm{Av}(321),\boxslash)$.

We set π' to be the (f,g)-transform of π and τ' to be the (f,g)-transform of τ. As we already argued, π' is an $\mathcal{M}'$-gridded permutation, τ' is an $\mathcal{N}$-gridded permutation and there is a grid-preserving occurrence of π' in τ' if and only if there is a grid-preserving occurrence of π in τ. Moreover it is still true that if there is an occurrence of π' in τ' that maps the (s_1,t_1)-cell to the (s_1,t_1)-cell of τ' then it must be grid-preserving. By the very same argument as before, in any map of the (s_1,t_1)-cell of π' the beginning of the (f,g)-transformed lanes must map to the beginnings of the (f,g)-transformed lanes in τ'. As before, this forces the whole mapping to be grid-preserving.

To summarise the argument so far, we used Theorem 3.7 to reduce the NP-complete 3-SAT problem to the problem of deciding whether a $\mathrm{St}^{2c+1}(\boxslash,\boxslash)$-gridded permutation π has a grid-preserving occurrence in a $\mathrm{St}^{2c+1}(\mathrm{Av}(321),\boxslash)$-gridded permutation τ. Then, by means of the (f,g)-transform, we showed that this is equivalent to finding a grid-preserving occurrence of an $\mathcal{M}'$-gridded permutation π' in an $\mathcal{N}$-gridded permutation τ'.

As the final step of our argument, we transform π' and τ' into permutations $\pi^\star \in \mathrm{Grid}(\mathcal{M}')$ and $\tau^\star \in \mathrm{Grid}(\mathcal{N})$ (that are no longer gridded) such that $\tau^\star$ contains $\pi^\star$ if and only if there is a grid-preserving copy of π' in τ'. This will imply that $\textsc{Grid}(\mathcal{M}')$-$\textsc{Pattern}$ PPM is NP-complete, and therefore $\textsc{Grid}(\mathcal{M})$-$\textsc{Pattern}$ PPM is NP-complete as well.

Let p be the length of the longest monotone subpermutation of τ'. Suppose that the (s_1,t_1)-cell of $\mathcal{M}'$ contains $\boxslash$, the other case being symmetric. We obtain $\pi^\star$ by inserting two increasing sequences of length $p+1$ into π', one directly to the left and below the (s_1,t_1)-cell, called a lower anchor, and one directly above and to the right of the (s_1,t_1)-cell, called an upper anchor. We perform the same modification on τ' to obtain $\tau^\star$. See Figure 3.

Fix any embedding of $\pi^\star$ into $\tau^\star$. Since there is no increasing subsequence of length $p+1$ in τ', at least $p+2$ of the points from the anchors in $\pi^\star$ must map to the anchors of $\tau^\star$. In particular, there is a point of the upper anchor in $\pi^\star$ mapped to the upper anchor in $\tau^\star$ and the same holds for the lower anchors. This implies that the whole copy of the initial block of π' inside $\pi^\star$ maps to the initial block of $\tau^\star$. We claim that everything except perhaps its rightmost point and its leftmost point (that belong to the lanes) in fact maps to the initial block of τ' inside $\tau^\star$. Suppose for a contradiction that the two leftmost points q and r of the initial block of π', i.e. the beginning of one of the lanes, map both to the lower anchor. The second block of the lane must lie in the vertical interval between q and r either to the right of both of them or to the left of both of them. But for any two points of the lower anchor in $\tau^\star$ there are no such points to map q and r to. The same is true for the beginning of the second lane. Furthermore, we can then modify the embedding to map the leftmost, respectively rightmost, point of π' to the leftmost, respectively rightmost, point of τ'. By the previous arguments, such an embedding must be grid-preserving which concludes the proof. ◀

The first part of Theorem 3.1 follows by combining Proposition 3.3 with Theorem 2.2 and the second part follows from Corollary 3.6 and Proposition 3.8.

4 General grid classes

In this section, we generalize the results of Theorem 3.1 to any gridding matrix whose every entry has bounded grid-width. Note that a "bumper-ended" path is a certain kind of path in $G_{\mathcal{M}}$ whose definition we provide later.

▶ **Theorem 4.1.** *Let $\mathcal{M}$ be a gridding matrix such that every entry of $\mathcal{M}$ has bounded grid-width. Then one of the following holds:*
- *Either $G_{\mathcal{M}}$ is a forest that avoids a bumper-ended path, Grid($\mathcal{M}$) has bounded grid-width and* Grid($\mathcal{M}$)-Pattern *PPM can be decided in polynomial time, or*
- *$G_{\mathcal{M}}$ contains a bumper-ended path or a cycle, Grid($\mathcal{M}$) has unbounded grid-width and* Grid($\mathcal{M}$)-Pattern *PPM is NP-complete.*

Unlike monotone grid classes, general grid classes may contain finite entries. However, we can ignore them without affecting the properties we are interested in. Let $\mathcal{M}'$ be the gridding matrix obtained by removing all finite entries from a gridding matrix $\mathcal{M}$. The NP-completeness of Grid($\mathcal{M}'$)-Pattern PPM trivially implies the NP-completeness of Grid($\mathcal{M}$)-Pattern PPM. Moreover, Grid($\mathcal{M}'$) has bounded grid-width if and only if Grid($\mathcal{M}$) has bounded grid-width since inserting a constant number of points into a permutation increases its grid-width at most by a constant. Thus, from now on we suppose that $\mathcal{M}$ contains only infinite entries.

One natural way to generalize the notion of monotone classes is to consider classes that have bounded path-width in left-to-right ordering or in the bottom-to-top ordering. For permutation π, the *horizontal path-width* is $\mathrm{pw}^\sigma(\pi)$ where $\sigma_i = i$, and the *vertical path-width* is $\mathrm{pw}^\sigma(\pi)$ where $\sigma_i = \pi_i^{-1}$. The horizontal path-width was introduced independently by Ahal and Rabinovich [1] and Albert et al. [3] in the context of designing permutation pattern matching algorithms. Moreover, there is a connection to the so-called insertion-encodable classes which appear often in the area of permutation classes enumeration, see e.g. [5, 26, 7].

A *horizontal monotone juxtaposition* is a monotone grid class Grid($\mathcal{C}\ \mathcal{D}$) where both $\mathcal{C}$ and $\mathcal{D}$ are non-empty. Similarly, a *vertical monotone juxtaposition* is a monotone grid class Grid$\left(\frac{\mathcal{C}}{\mathcal{D}}\right)$. The following two lemmas are stated for the horizontal path-width, but their symmetric versions hold for the vertical path-width. The first lemma have been observed in different form by Albert et al. [5] when studying regular insertion-encodable classes.

▶ **Lemma 4.2.** *For a permutation class $\mathcal{C}$ the following are equivalent:*
(a) *$\mathcal{C}$ has unbounded horizontal path-width,*
(b) *$\mathcal{C}$ contains arbitrarily large horizontal alternations, and*
(c) *$\mathcal{C}$ contains a horizontal monotone juxtaposition as a subclass.*

▶ **Lemma 4.3.** *Let π be a permutation from class $\mathcal{C}$ with bounded horizontal path-width and let S be a subset of π such that $\mathrm{int}(\Pi_x(S)) = k$. Then $\mathrm{int}(\Pi_y(S)) \le \alpha k$ where the constant α depends only on $\mathcal{C}$.*

An ordered pair (p, q) of vertices in $G_{\mathcal{M}}$ is a *bumper* if either $\mathcal{M}_q$ has unbounded horizontal path-width and shares the same column with $\mathcal{M}_p$, or if $\mathcal{M}_q$ has unbounded vertical path-width and shares the same row with $\mathcal{M}_p$. A *bumper-ended path* is a path $P = p_1, \ldots, p_k$ in $G_{\mathcal{M}}$ such that both (p_2, p_1) and (p_{k-1}, p_k) are bumpers.

▶ **Lemma 4.4.** *If $G_{\mathcal{M}}$ contains a bumper-ended path then Grid($\mathcal{M}$) has unbounded grid-width and the problem* Grid($\mathcal{M}$)-Pattern *PPM is NP-complete.*

Proof. We aim to show that Grid($\mathcal{M}$) must contain a cyclic monotone grid class as its subclass. The proof is illustrated in Figure 4. Consider the bumper-ended path $p_1, p_2, \ldots, p_k$. Let us assume that both M_{p_1} and $\mathcal{M}_{p_k}$ have unbounded horizontal path-width as the other cases can be proved in an analogous way. Each of the infinite classes $\mathcal{M}_{p_i}$ contains a monotone subclass $\mathcal{C}_i$ due to the Erdős–Szekeres theorem [13]. Moreover, the classes $\mathcal{M}_{p_1}$ and $\mathcal{M}_{p_k}$ contain some monotone juxtaposition due to Lemma 4.2. Let Grid($\mathcal{C}_1\ \mathcal{D}_1$) be the

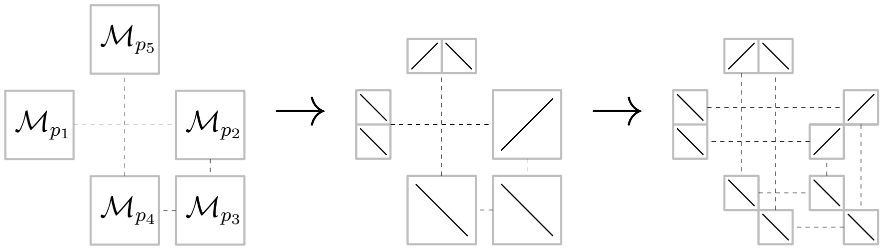

Figure 4 Transforming a bumper-ended path to cycle in the proof of Lemma 4.4.

juxtaposition contained in $\mathcal{M}_{p_1}$ and Grid($\mathcal{C}_k$ $\mathcal{D}_k$) the juxtaposition contained in $\mathcal{M}_{p_k}$. We define the monotone gridding matrix $\mathcal{M}'$ by replacing every entry of $\mathcal{M}$ with the following 2×2 matrix:

- entry $\mathcal{M}_{p_i}$ for i between 2 and $k-1$ is replaced with $\mathcal{C}_i^{\times 2}$
- entry $\mathcal{M}_{p_t}$ for $t \in \{1, k\}$ is replaced with $\left(\begin{smallmatrix} \emptyset & \emptyset \\ \mathcal{C}_t & \mathcal{D}_t \end{smallmatrix}\right)$, and
- every other entry is replaced with an empty 2×2 matrix.

Clearly, Grid($\mathcal{M}'$) is still a subclass of Grid($\mathcal{M}$). The (i,j)-block of $\mathcal{M}'$ is the 2×2 submatrix obtained from the (i,j)-cell in $\mathcal{M}$. If we forget about the blocks of p_1 and p_k we are left with two disjoint copies of the original path. Adding back the blocks connects the endpoints of both paths together and creates a cycle since p_2 shares the same column with p_1 and p_{k-1} shares the same column with p_k. Thus, *Grid*($\mathcal{M}'$) is a monotone grid subclass of Grid($\mathcal{M}$) whose cell graph contains a cycle and the claim follows from Proposition 3.8. ◀

▶ **Proposition 4.5.** *If $G_{\mathcal{M}}$ is a forest that avoids a bumper-ended path then Grid($\mathcal{M}$) has bounded grid-width and* Grid($\mathcal{M}$) *can be decided in polynomial time.*

Sketch of proof. We can without loss of generality assume that $G_{\mathcal{M}}$ is connected. Otherwise we could decompose each connected component individually.

We aim to transform $G_{\mathcal{M}}$ into a rooted tree, such that no path in $G_{\mathcal{M}}$ from root ends with a bumper. We pick an arbitrary non-empty cell of $\mathcal{M}$ as the root r. Suppose that there are vertices $v_1, \ldots, v_s$ such that the rv_i-path ends with a bumper. Select as the new root the vertex v_i which is furthest from r. Assume that there exists a vertex w such that the $v_i w$-path ends with a bumper. If w lied in the same subtree of v_i as r then the wv_i-path would have to be bumper-ended. Therefore, w and r lie in different subtrees of v_i. The rw-path then ends with a bumper which implies that w is the vertex v_j for some $j \in [s]$. However, that is not possible since w is further away from r than v_i.

Let us transform $G_{\mathcal{M}}$ into a tree $T_{\mathcal{M}}$, also rooted in r. For a vertex $v \neq r$ in $G_{\mathcal{M}}$, we set the parent of v in $T_{\mathcal{M}}$ to be the furthest vertex on the vr-path in $G_{\mathcal{M}}$ that shares the same row or column with v. Whenever a vertex v shares the same column with its parent w in $T_{\mathcal{M}}$ then the entry $\mathcal{M}_v$ has bounded horizontal path-width and whenever v shares the same row with w, the entry $\mathcal{M}_v$ has bounded vertical path-width.

Let π be an $\mathcal{M}$-gridded permutation that contains point in every non-empty cell and let π_v be the subset of points contained in the v-cell. We define an auxiliary oriented graph G_π. Suppose that the vertex v of $T_{\mathcal{M}}$ shares the same column with its parent w. The parent of a point p in the v-cell is the nearest point in the w-cell to the right of p, and if there is no such point then the rightmost point in the w-cell. If v and w share the same row, then the parent of p is the nearest point in the w-cell above p, or if there is no such point then the topmost point in the w-cell. Let P be a subset of the permutation diagram of P. The point set $\overline{P}$ contains P and every point that lies in G_π in a subtree of some point $p \in P$.

For a vertex v in $G_{\mathcal{M}}$, let α_v denote the constant obtained from Lemma 4.3 if v shares a column with its parent in $T_{\mathcal{M}}$, or from its "vertical" version otherwise. We inductively define a function h as follows. Let $h(u) = 1$ for any leaf u of $T_{\mathcal{M}}$ and for every other vertex v with children $w_1, \ldots, w_k$, set $h(v) = 1 + \sum_{i=1}^{j} \alpha_{w_i} h(w_i)$. Suppose that v is a vertex in $T_{\mathcal{M}}$ whose children lie in the same column and that S is a subset of π that forms one consecutive vertical interval in the column of v. We prove by induction that the set $\overline{S}$ has grid-complexity at most $h(v)$ when ignoring the row occupied by v. Let w_i be a child of v in $T_{\mathcal{M}}$. Due to Lemma 4.3, the points of $\overline{S}$ form at most α_{w_i} horizontal intervals inside w_i. Applying inductively the "vertical" version of the same claim on each of them bounds grid-complexity of $\overline{S}$ by $\sum_{i=1}^{j} \alpha_{w_i} h(w_i)$ when ignoring the row and column occupied by v. We obtain the promised bound by adding extra 1 for the column occupied by v.

Let p be a point of π_r. Using the arguments of previous paragraph, we construct a grid tree of the point set $\overline{\{p\}}$ with grid-width at most $h(r) + 1$. Let T_r be the optimum grid tree of π_r with grid-width at most g. A grid tree $T^\star$ of the whole π is obtained by replacing the leaf containing point p in T_r with T_p. The proof is concluded by arguing that for every subset S of π_r with grid-complexity at most g, the grid-complexity of $\overline{S}$ is at most $2gh(r)$. ◀

Finally, if $G_{\mathcal{M}}$ contains a cycle then Grid($\mathcal{M}$) contains a monotone grid subclass Grid($\mathcal{M}'$) where $\mathcal{M}'$ is obtained by replacing every infinite class in $\mathcal{M}$ by its monotone subclass. Applying Proposition 3.8 then wraps up the proof of Theorem 4.1.

5 Concluding remarks and open problems

The $\mathcal{C}$-PPM is the problem of determining whether a permutation $\pi \in \mathcal{C}$ is contained in a permutation $\tau \in \mathcal{C}$. Even though Av(321)-Pattern PPM is NP-complete, Av(321)-PPM can be decided in polynomial time [16]. This leads to the natural question of whether there are any such grid classes.

▶ **Open problem 1.** *Is there any (monotone) gridding matrix $\mathcal{M}$ such that Grid($\mathcal{M}$)-Pattern PPM is NP-complete and Grid($\mathcal{M}$)-PPM can be decided in polynomial time?*

We showed that if $\mathcal{C}$ contains arbitrarily large connected grid class (large with respect to the size of gridding matrix) as a subclass then $\mathcal{C}$ has unbounded grid-width and $\mathcal{C}$-Pattern PPM is NP-complete. In fact in all the cases when $\mathcal{C}$ is known to have unbounded grid-width, $\mathcal{C}$ contains arbitrarily large connected monotone grid class as its subclass. We may therefore ask whether this property precisely characterizes the classes with unbounded grid-width.

▶ **Open problem 2.** *Does every class $\mathcal{C}$ with unbounded grid-width contain arbitrarily large connected monotone grid subclasses?*

References

1 Shlomo Ahal and Yuri Rabinovich. On the complexity of the sub-permutation problem. In *Algorithm theory—SWAT 2000 (Bergen)*, volume 1851 of *Lecture Notes in Comput. Sci.*, pages 490–503. Springer, Berlin, 2000. doi:10.1007/3-540-44985-X_41.

2 Michael Albert, Jay Pantone, and Vincent Vatter. On the growth of merges and staircases of permutation classes. *Rocky Mountain J. Math.*, 49(2):355–367, 2019. doi:10.1216/RMJ-2019-49-2-355.

3 Michael H. Albert, Robert E. L. Aldred, Mike D. Atkinson, and Derek A. Holton. Algorithms for pattern involvement in permutations. In *Algorithms and computation (Christchurch, 2001)*, volume 2223 of *Lecture Notes in Comput. Sci.*, pages 355–366. Springer, Berlin, 2001. doi:10.1007/3-540-45678-3_31.

4 Michael H. Albert, M. D. Atkinson, Mathilde Bouvel, Nik Ruškuc, and Vincent Vatter. Geometric grid classes of permutations. *Trans. Amer. Math. Soc.*, 365(11):5859–5881, 2013. doi:10.1090/S0002-9947-2013-05804-7.

5 Michael H. Albert, Steve Linton, and Nik Ruškuc. The insertion encoding of permutations. *Electron. J. Combin.*, 12:Research Paper 47, 31, 2005. URL: http://www.combinatorics.org/Volume_12/Abstracts/v12i1r47.html.

6 M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Partially well-ordered closed sets of permutations. *Order*, 19(2):101–113, 2002. doi:10.1023/A:1016500300436.

7 Christian Bean, Bjarki Gudmundsson, and Henning Ulfarsson. Automatic discovery of structural rules of permutation classes. *Math. Comp.*, 88(318):1967–1990, 2019. doi:10.1090/mcom/3386.

8 David Bevan. Growth rates of geometric grid classes of permutations. *Electron. J. Combin.*, 21(4):Paper 4.51, 17, 2014.

9 David Bevan. Growth rates of permutation grid classes, tours on graphs, and the spectral radius. *Trans. Amer. Math. Soc.*, 367(8):5863–5889, 2015. doi:10.1090/S0002-9947-2015-06280-1.

10 Prosenjit Bose, Jonathan F. Buss, and Anna Lubiw. Pattern matching for permutations. *Inform. Process. Lett.*, 65(5):277–283, 1998. doi:10.1016/S0020-0190(97)00209-3.

11 Robert Brignall. Grid classes and partial well order. *J. Combin. Theory Ser. A*, 119(1):99–116, 2012. doi:10.1016/j.jcta.2011.08.005.

12 Marie-Louise Bruner and Martin Lackner. A fast algorithm for permutation pattern matching based on alternating runs. *Algorithmica*, 75(1):84–117, 2016. doi:10.1007/s00453-015-0013-y.

13 P. Erdös and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935. URL: http://www.numdam.org/item?id=CM_1935__2__463_0.

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

15 Sylvain Guillemot and Dániel Marx. Finding small patterns in permutations in linear time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 82–101. ACM, New York, 2014. doi:10.1137/1.9781611973402.7.

16 Sylvain Guillemot and Stéphane Vialette. Pattern matching for 321-avoiding permutations. In *Algorithms and computation*, volume 5878 of *Lecture Notes in Comput. Sci.*, pages 1064–1073. Springer, Berlin, 2009. doi:10.1007/978-3-642-10631-6_107.

17 Louis Ibarra. Finding pattern matchings for permutations. *Inf. Process. Lett.*, 61(6):293–295, 1997. doi:10.1016/S0020-0190(97)00029-X.

18 Vít Jelínek and Jan Kynčl. Hardness of permutation pattern matching. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 378–396. SIAM, Philadelphia, PA, 2017. doi:10.1137/1.9781611974782.24.

19 Vít Jelínek, Michal Opler, and Pavel Valtr. Generalized coloring of permutations. In *26th European Symposium on Algorithms*, volume 112 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 50, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2018.

20 Erkki Mäkinen. On the longest upsequence problem for permutations. *Int. J. Comput. Math.*, 77(1):45–53, 2001. doi:10.1080/00207160108805049.

21 Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 107(1):153–160, 2004. doi:10.1016/j.jcta.2004.04.002.

22 Maximillian M. Murphy and Vincent R. Vatter. Profile classes and partial well-order for permutations. *Electron. J. Combin.*, 9(2):Research paper 17, 30, 2002/03. Permutation patterns (Otago, 2003). URL: http://www.combinatorics.org/Volume_9/Abstracts/v9i2r17.html.

23 Both Emerite Neou. *Permutation pattern matching.* PhD thesis, Université Paris-Est ; Université de Vérone (Italie), 2017. URL: https://pastel.archives-ouvertes.fr/tel-01866721.

24 Both Emerite Neou, Romeo Rizzi, and Stéphane Vialette. Permutation pattern matching in (213, 231)-avoiding permutations. *Discret. Math. Theor. Comput. Sci.*, 18(2), 2016. URL: http://dmtcs.episciences.org/3199.

25 Vincent Vatter. Small permutation classes. *Proc. Lond. Math. Soc. (3)*, 103(5):879–921, 2011. doi:10.1112/plms/pdr017.

26 Vincent Vatter. Finding regular insertion encodings for permutation classes. *J. Symbolic Comput.*, 47(3):259–265, 2012. doi:10.1016/j.jsc.2011.11.002.

27 Vincent Vatter. Growth rates of permutation classes: from countable to uncountable. *Proc. Lond. Math. Soc. (3)*, 119(4):960–997, 2019. doi:10.1112/plms.12250.

28 V. Yugandhar and Sanjeev Saxena. Parallel algorithms for separable permutations. *Discret. Appl. Math.*, 146(3):343–364, 2005. doi:10.1016/j.dam.2004.10.004.

A Appendix

Due to the page limit, we had to omit some proofs of our results from the main paper. We provide them here.

Proof of Lemma 2.1. Suppose that $\mathrm{pw}(\pi) = p$ as witnessed by a caterpillar grid tree T. Observe that all leaves of T except for the deepest pair lie in different depths. Define σ to simply order the labels by the depths of their leaves, defining arbitrarily the order of the two deepest leaves. Then every set $\{(\sigma_1, \pi_{\sigma_1}), \dots, (\sigma_i, \pi_{\sigma_i})\}$ corresponds exactly to the set π_v^T for some vertex v of T.

In order to prove the other direction, we define a sequence of caterpillar trees $T_1, \dots, T_n$ in the following way. Let T_1 be a single vertex labeled by $(\sigma_1, \pi_{\sigma_1})$. For $i > 1$, let T_i be the binary rooted tree with left child a leaf labeled by $(\sigma_i, \pi_{\sigma_i})$ and right child the tree T_{i-1}. The tree T_n is a caterpillar grid tree of π and for every inner vertex v the set π_v^T is equal to $\{(\sigma_1, \pi_{\sigma_1}), \dots, (\sigma_i, \pi_{\sigma_i})\}$ for some i. The claim follows. ◀

Proof of Lemma 4.2. Suppose (a) holds and for every integer k there is a permutation $\pi^k \in \mathcal{C}$ such that the horizontal path-width of π^k is at least k. Then there is i such that the set $\{\pi_1^k, \dots, \pi_i^k\}$ has intervalicity at least k. Each pair of neighboring intervals is separated by π_j^k for some $j > i$. Therefore, π^k contains a horizontal alternation of size at least $2k-1$ which proves (b). On the other hand, a horizontal alternation of size $2k$ must have a horizontal path-width at least k and thus (b) implies (a).

Now suppose that (b) holds. A *monotone horizontal alternation* is a horizontal alternation whose set of odd entries and set of even entries both form monotone sequences. We claim that every horizontal alternation π of size $2k^4$ contains a monotone horizontal alternation of size $2k$. By applying the Erdős–Szekeres theorem [13] on the odd entries of π we obtain a horizontal alternation π' of size at least $2k^2$ whose odd entries form a monotone sequence. Applying the Erdős–Szekeres theorem again on even entries of π' yields a monotone horizontal alternation π'' of size at least $2k$. Therefore, $\mathcal{C}$ contains arbitrarily large monotone horizontal alternations. There are only four possible types of such alternations depending on the type of the monotone sequences. Therefore, $\mathcal{C}$ also contains arbitrarily large alternations belonging to a horizontal monotone juxtaposition $\mathrm{Grid}(\mathcal{D}_1\ \mathcal{D}_2)$ for some choice of $\mathcal{D}_1$, $\mathcal{D}_2$. Since every $\sigma \in \mathrm{Grid}(\mathcal{D}_1\ \mathcal{D}_2)$ is contained in a sufficiently large monotone alternation, in fact $\mathcal{C}$ must contain the whole class $\mathrm{Grid}(\mathcal{D}_1\ \mathcal{D}_2)$ as a subclass.

On the other hand, if (c) holds then $\mathcal{C}$ contains arbitrarily large monotone horizontal alternations which trivially implies (b). ◀

Proof of Lemma 4.3. Due to Lemma 4.2, there exists an l such that $\mathcal{C}$ does not contain any vertical alternation of size l. Let $\mathcal{I}$ be the interval family of size k such that $\bigcup \mathcal{I} = \Pi_x(S)$ and let I be a single interval of $\mathcal{I}$. Let us denote by S_I the subset of S such that $\Pi_x(S_I) = I$ and let $\mathcal{J}$ be the smallest interval family such that $\Pi_y(S_I) = \bigcup \mathcal{J}$. We claim that $\mathcal{J}$ contains at most $2l-1$ intervals. For a contradiction suppose that the size of $\mathcal{J}$ is at least $2l$. Then between each pair of consecutive intervals of $\mathcal{J}$ there is a value j such that π_j^{-1} lies outside the interval I. There is at least $2l-1$ gaps between intervals of $\mathcal{J}$ and therefore by the pigeon-hole principle either l of them contain a point to the right of I or at least l of the gaps contain a point to the left of I. Either way we have obtained a horizontal alternation of size l, which is a contradiction.

For each interval $I \in \mathcal{I}$, we showed that the intervalicity of $\Pi_y(S_i)$ is at most $2l-1$ and thus the intervalicity of $\Pi_y(S)$ is at most $k(2l-1)$. ◀

Full proof of Proposition 4.5. First, suppose that $G_\mathcal{M}$ contains more than one component. In that case, let $\mathcal{M}_1$ be a connected component of $\mathcal{M}$ and $\mathcal{M}_2$ the remaining part of $\mathcal{M}$. An $\mathcal{M}$-gridded permutation π can be split into π_1 and π_2 where π_i is the $\mathcal{M}_i$-gridded subpermutation of π consisting of the rows and columns of $\mathcal{M}_i$. Let T_i be the optimal grid tree of π_i. We define a grid tree T of π by taking a root vertex with children T_1 and T_2. The grid-complexity of any vertex in T_1 or T_2 has increased at most by $\max(k, \ell)$ where k and ℓ are the dimensions of $\mathcal{M}$. Therefore $\mathrm{gw}(\pi) \leq \max(\mathrm{gw}(\pi_1), \mathrm{gw}(\pi_2)) + \max(k, \ell)$. Applying this argument inductively shows that $\mathrm{Grid}(\mathcal{M})$ has bounded grid-width if and only if the grid-width of $\mathrm{Grid}(\mathcal{M}')$ is bounded for every connected component $\mathcal{M}'$ of $\mathcal{M}$. In the rest of the proof we assume that $G_\mathcal{M}$ is a tree.

We aim to transform the tree $G_\mathcal{M}$ into a rooted tree, such that no path in $G_\mathcal{M}$ from root ends with a bumper. We pick an arbitrary non-empty cell of $\mathcal{M}$ as the root r. Suppose that there are vertices $v_1, \dots, v_s$ such that the rv_i-path ends with a bumper. Select as the new root the vertex v_i which is furthest from r. Assume that there exists a vertex w such that the $v_i w$-path ends with a bumper. If w lied in the same subtree of v_i as r then the wv_i-path would have to be bumper-ended. Therefore, w and r lie in different subtrees of v_i. The rw-path then ends with a bumper which implies that w is the vertex v_j for some $j \in [s]$. However, that is not possible since w is further away from r than v_i and we picked i such that v_i is the vertex furthest from r among $v_1, \dots, v_s$.

Let us transform $G_\mathcal{M}$ into a slightly different tree $T_\mathcal{M}$ rooted in r. For a vertex $v \neq r$ in $G_\mathcal{M}$, we set the parent of v in $T_\mathcal{M}$ to be the furthest vertex on the vr-path in $G_\mathcal{M}$ that shares the same row or column with v. Observe that whenever a vertex v shares the same column with its parent w in $T_\mathcal{M}$ then the entry $\mathcal{M}_v$ has bounded horizontal path-width and whenever v shares the same row with w, the entry $\mathcal{M}_v$ has bounded vertical path-width.The *dominant cell* of a row, or a column, is the cell v such that every other cell in the row, or column, are its children in $T_\mathcal{M}$.

Let π be an $\mathcal{M}$-gridded permutation and let us assume that it contains point in every non-empty cell and let π_v be the subset of points contained in the v-cell. We define an auxiliary oriented graph G_π on the points of π such that every connected component is a tree rooted in some point of the r-cell. Suppose that the vertex v of $T_\mathcal{M}$ shares the same column with its parent w. The parent of a point p in the v-cell is the nearest point in the w-cell to the right of p, and if there is no such point (p lies to the right of all the points of the w-cell) then the rightmost point in the w-cell. If v and w share the same row, then the parent of p is the nearest point in the w-cell above p, or if there is no such point (p lies above all the points of the w-cell) then the topmost point in the w-cell. Let P be a subset of the permutation diagram of P. The point set $\overline{P}$ contains P and every point that lies in G_π in a subtree of some point $p \in P$.

Let v be a non-empty cell such that π_v contains m points. The permutation π_v^* is the standardized version of π_v, i.e. the point set inside $[m] \times [m]$ that is isomorphic to π_v. The construction of the graph G_π guarantees the following property and its symmetric version.

$\triangleright$ Observation A.1. Let v be the dominant cell of its row. Let S be a subset of π_v, let S^* be the corresponding subset of the standardized π_v^* and let S' be the set containing S and all its children in G_π that lie in the same row. Then $\operatorname{int}(\Pi_y(S^*)) = \operatorname{int}(\Pi_y(S'))$.

We inductively define a function h that we will later serve as an upper bound for grid-width of any $\pi \in \operatorname{Grid}(\mathcal{M})$. Let $h(u) = 1$ for any leaf u of $T_\mathcal{M}$ and for every other vertex v with children $w_1, \dots, w_k$, we set $h(v) = 1 + \sum_{i=1}^{j} \alpha_{w_i} h(w_i)$ where α_{w_i} is the constant obtained as follows. If v shares a column with its children then α_{w_i} is the constant from Lemma 4.3 applied on the class $\mathcal{M}_{w_i}$, otherwise it is the constant from the "vertical" version of Lemma 4.3 applied on the class $\mathcal{M}_{w_i}$. We state only one of the symmetric version of the following two claims. However, we are proving both of them simultaneously by induction.

$\triangleright$ Claim A.2. Let I be a subinterval of the i-th column of the $\mathcal{M}$-gridding of π and let v be the dominant cell of the i-th column. Let S be the set of points of π such that $\Pi_x(S) = I$ and let $S_v = S \cap \pi_v$. Then $\operatorname{int}(\Pi_x(\overline{(S \setminus S_v)} \cup S_v)) \le h(v)$ and $\operatorname{int}(\Pi_y(\overline{S \setminus S_v})) \le h(v) - 1$.

First, suppose that v is the only cell in its column, i.e. v is a leaf of $T_\mathcal{M}$. Then $\operatorname{int}(\Pi_x(\overline{S})) \le 1$ and $\operatorname{int}(\Pi_y(\overline{S'})) = 0$ since S' is an empty set.

Otherwise, let S_w be the set of $S \cap \pi_w$ for every non-empty cell $w \neq v$ in the same column. Observe that w is a child of v in $T_\mathcal{M}$ and that $\overline{S \setminus S_v}$ is a disjoint union of the sets $\overline{S_w}$. Due to Lemma 4.3, the points of S_w can be split into at most α_w subsets S_w^i, each satisfying the requirements of Claim A.1. Let R_w^i be the set S_w^i together with all its children in G_π that lie in the same row. By Observation A.1, each $\Pi_y(R_w^i)$ is a single interval, let us denote it by J_w^i. Using symmetric version of Claim A.2 on each J_w^i shows that $\operatorname{int}(\Pi_x(\overline{R_w^i \setminus S_w^i})) \le h(w) - 1$ and $\operatorname{int}(\Pi_y(\overline{(R_w^i \setminus S_w^i)} \cup S_w^i)) \le h(w)$. Therefore, $\operatorname{int}(\Pi_x(\overline{(S \setminus S_v)} \cup S_v))$ is at most $1 + \sum_w \alpha_w h(w) = h(v)$, where the extra 1 comes from counting the interval I, and $\operatorname{int}(\Pi_y(\overline{S \setminus S_v}))$ is at most $\sum_w \alpha_w (h(w) - 1) \le h(v) - 1$.

$\triangleright$ Claim A.3. Let I be a subinterval of the i-th column of the $\mathcal{M}$-gridding of π and let v be the dominant cell of the i-th column. Let S be the set of points of π such that $\Pi_x(S) = I$ and let $S_v = S \cap \pi_v$. If $\operatorname{int}(\Pi_y(S_v)) \le 1$ then there is a grid tree T_I of $\overline{(S \setminus S_v)} \cup S_v$ such that $\operatorname{gw}^{T_I}(\overline{(S \setminus S_v)} \cup S_v) \le h(v)$.

First, we remark that $\overline{(S \setminus S_v)} \cup S_v$ has bounded grid-complexity. It follows from Claim A.2 that $\operatorname{int}(\Pi_x(\overline{(S \setminus S_v)} \cup S_v))$ is at most $h(h)$ and $\operatorname{int}(\Pi_y(\overline{(S \setminus S_v)}))$ is at most $h(v) - 1$. It is sufficient to combine this with the assumption $\operatorname{int}(\Pi_y(S_v)) \le 1$.

Let p be the leftmost point of S and let $I' = I \setminus \{p\}$. We can use Claim A.3 inductively on I' to obtain a grid tree $T_{I'}$. First, suppose that p lies in the v-cell. We define T_I to be the grid tree whose root has as children the point p and the grid tree $T_{I'}$. We already checked the grid-complexity of the root of T_I, the rest follows inductively.

The other option is that p lies in an w-cell of $\mathcal{M}$ such that w is a child of v in $T_\mathcal{M}$. Let S' be the points of $\overline{\{p\}}$ that lie in the same row as w. Observe that $\operatorname{int}(\Pi_y(S')) = 1$ and let $J = \Pi_y(S')$. Therefore, we can obtain inductively a grid tree T_J of $\overline{\{p\}}$ using the symmetric version of Claim A.3 on J (it contains only a single point in the w-cell). We define T_I to be the grid tree whose root has as children the grid tree T_J and the grid tree $T_{I'}$. The grid-with of T_J is at most $h(w)$, while the grid-width of $T_{I'}$ and the grid-complexity of the root of T_S are both at most $h(v)$ which concludes the proof of Claim A.3.

For a point p of π_r, we define a grid tree T_p of the point set $\overline{\{p\}}$ as follows. Let S_x be the children of p in G_π that lie in the same column as the r-cell and let S_y be the children of p in G_π that lie in the same row as the r-cell. Both $\Pi_x(S_x \cup \{p\})$ and $\Pi_y(S_y \cup \{p\})$ are a single intervals due to Observation A.1, let us denote them by I and J. We can obtain grid trees T_x of $\overline{S_x} \cup \{p\}$ and T_y of $\overline{S_y} \cup \{p\}$ by applying Claim A.3 on I and J, both with grid-width at most $h(r)$. Set T_p to be tree whose root has children T_x and T'_y where T'_y is obtained from T_y by removing the leaf of p. The grid-width of T'_y is at most $h(r) + 1$ and thus the grid-width of T_p is also at most $h(r) + 1$.

Finally, let T_r be the optimum grid tree of π^*_r whose grid-width is at most g. A grid tree $T^\star$ of the whole permutation π is obtained by taking T_r and replacing the leaf containing point p with the tree T_p. We claim that $\mathrm{gw}^{T^\star}(\pi)$ is at most $2gh(r)$. The tree $T^\star$ contains every point of π and we showed that the grid-width of any vertex contained in a copy of some T_p is at most $h(r) + 1$. Let s be an inner vertex of $T^\star$ that belonged to T_r and let S be the associated subset of π_r. The subset of π contained in the subtree of s in $T^\star$ is precisely $\overline{S}$. Applying Observation A.1, we see that S together with its neighbors in G_π spans g consecutive intervals in the row and column dominated by the r-cell. Thus, we get that the grid-complexity of $\overline{S}$ is at most $2gh(r)$ by applying Claim A.2 individually on each of these $2g$ intervals. ◀