

Factoring Pattern-Free Permutations into Separable ones

Édouard Bonnet, Romain Bourneuf, Colin Geniet, and Stéphan Thomassé

Univ. Lyon, ENS de Lyon, UCBL, CNRS, LIP, France

Abstract

We show that for any permutation π there exists an integer k_π such that every permutation avoiding π as a pattern is a product of at most k_π separable permutations. In other words, every strict class \mathcal{C} of permutations is contained in a bounded power of the class of separable permutations. This factorisation can be computed in linear time, for any fixed π .

The central tool for our result is a notion of width of permutations, introduced by Guillemot and Marx [SODA '14] to efficiently detect patterns, and later generalised to graphs and matrices under the name of twin-width. Specifically, our factorisation is inspired by the decomposition used in the recent result that graphs with bounded twin-width are polynomially χ -bounded. As an application, we show that there is a fixed class \mathcal{C} of graphs of bounded twin-width such that every class of bounded twin-width is a first-order transduction of \mathcal{C} .

1 Introduction

Given a class \mathcal{C} of discrete structures, the arguably preeminent algorithmic task is the one of recognition: does the input belong to \mathcal{C} ? This problem is often tied with an effective *construction* of the class \mathcal{C} by performing elementary operations on some basic building blocks. For instance, totally unimodular matrices [27], minor-closed classes [25], and perfect graphs [13] can be constructed from simpler objects, respectively, network matrices, graphs embeddable on low-genus surfaces, and variants of bipartite graphs. In this paper we show that *strict* classes of permutations, that is, those avoiding a fixed pattern, can be constructed from separable permutations (the basic class) via a bounded number of compositions (the elementary operation).

Given a positive integer n , we denote by $[n]$ the set $\{1, 2, \dots, n\}$. Let $n \leq m$ be two integers, we say that a permutation $\pi \in \mathfrak{S}_n$ is a *pattern* of $\sigma \in \mathfrak{S}_m$ if there is an increasing function f from $[n]$ to $[m]$ such that $\pi(i) < \pi(j)$ if and only if $\sigma(f(i)) < \sigma(f(j))$ for all $i, j \in [n]$. Another way of characterizing patterns is to associate to a permutation $\sigma \in \mathfrak{S}_n$ its $n \times n$ matrix $A(\sigma) = (a_{ij})$ with $a_{ij} = 1$ if $j = \sigma(i)$, and $a_{ij} = 0$ otherwise. Observe that π is a pattern of σ if and only if $A(\pi)$ is a submatrix of $A(\sigma)$. For instance, 12345 is a pattern of σ if it contains an increasing subsequence of length five. A crucial achievement in permutation patterns is the Guillemot–Marx algorithm, which decides if a permutation π is a pattern of σ in time $f(\pi) \cdot |\sigma|$, where $|\sigma|$ is the size of σ .

Patterns readily offer a complexity notion for permutations: A permutation is “simple” if it does not contain a fixed small pattern. We will consider *classes* of permutations, which are assumed closed under taking patterns. The existence of a gap between the class of all permutations and any strict class is illustrated by the Marcus–Tardos theorem, answering the Stanley–Wilf conjecture: Every strict class of permutations has at most $2^{O(n)}$ permutations of size n , whereas the class of all permutations obviously has $n! = 2^{\Theta(n \log n)}$ such permutations. From an algorithmic perspective, sequences avoiding a fixed pattern can be comparison-sorted in almost linear time

$O(n \cdot 2^{(1+o(1))\alpha(n)})$ where α is the inverse Ackermann function [11, 21, 12], while linear algorithms when excluding some specific small patterns have been long known [20, 1]. Furthermore, as a generalisation of Guillemot–Marx algorithm, any property defined using first-order logic (FO) can be tested inside any strict permutation class \mathcal{C} in linear time [7]. For instance, given a fixed permutation τ , one can decide in linear time if an input n -permutation σ in \mathcal{C} is such that every pair of elements $i, j \in [n]$ is contained in a pattern τ of σ . Observe that even the existence of a linear-size positive certificate for this seemingly quadratic problem is far from obvious.

Strict classes of permutations are thus significantly simpler, both algorithmically and in terms of growth. The next question is to construct them from a basic class using some simple operations. In the case of permutations, possibly the most natural elementary operation is the product (or composition). Furthermore, the class of *separable permutations* is basic in several ways. It consists of those permutations whose permutation graph is a cograph; an elementary graph class (which coincides with graphs of twin-width 0). Like cographs have a natural auxiliary tree structure (the *cotrees*), separable permutations inherit their own tree structure, the so-called *separating tree* [9]. Separable permutations are originally themselves defined from the trivial permutation 1, by successively applying direct sums (setting two permutation matrices as diagonal blocks of the new permutation matrix) or skew sums (the same with antidiagonal blocks), or equivalently by closing $\{12, 21\}$ under substitution. They are well known to be the permutations avoiding the patterns 2413 and 3142 [9].

As our main result, we show:

Theorem 1.1. *For any pattern π , there exists $k_\pi = 2^{2^{O(|\pi|)}}$ such that every permutation avoiding π is a product of at most k_π separable permutations.*

Combining Theorem 1.1 with the definition of separable permutations, every permutation of $Av(\pi)$, the set of permutations avoiding the pattern π , can be built from the trivial permutation 1 via direct and skew sums, followed by a bounded-length product. Conversely, remark that for any c , the class of products of c separable permutations avoids some pattern, since it contains only $2^{O(cn)}$ permutations on n elements.

The proof of Theorem 1.1 is effective, and yields a fixed-parameter tractable (FPT) algorithm to compute the factorisation. With some more work, we show how to implement it in linear time.

1.1 Shortest separable decompositions

We call *separable index* of a permutation σ , denoted $\text{sep}(\sigma)$, the smallest k such that σ is the product of k separable permutations. Our main result then states that permutations avoiding a pattern have bounded separable index. Our upper bound is doubly exponential: we show that permutations avoiding a pattern π of length k have separable index at most $2^{2^{O(k)}}$. It is natural to ask to which extent this bound can be improved.

For specific patterns π , this bound may be very low: if π is the decreasing permutation of length k , then it is relatively simple to show that permutations avoiding π have separable index $O(\log k)$, see Lemma 2.2 below. However this situation is unusual: from lower bounds on the growth of pattern-avoiding classes due to Fox [17], it follows that for almost all patterns π of size k , the maximum separable index of permutations avoiding π is at least polynomial, more precisely at least $\Omega(k^{1/4-\epsilon})$ for any $\epsilon > 0$. This still leaves a huge gap between our double-exponential upper bound and the polynomial lower bound. In principle the following question could be met with a positive answer.

Question 1.2. Is the maximum separable index among permutations avoiding a pattern of size k polynomial in k ?

As an easier first step, one could try and bring the upper bound down to singly exponential.

From a different perspective, one may consider the problem of computing $\text{sep}(\sigma)$, given some permutation σ . Our result implies an FPT approximation. Given $k \in \mathbb{N}$, there is a pattern π that every product of k separable permutations avoids. Conversely, by our result, any permutation avoiding π is product of at most $f(k)$ separable permutations for some function f . Applying the pattern recognition algorithm of Guillemot and Marx [18], we can either detect that σ contains π and conclude that $\text{sep}(\sigma) > k$, or find that σ avoids π and conclude that $\text{sep}(\sigma) \leq f(k)$. Here, the size of π is exponential in k , hence the approximation function f is triple exponential in k . We ask whether this can be improved, possibly up to an exact algorithm:

Question 1.3. What is the parameterised complexity of computing $\text{sep}(\sigma)$, and of approximating it within a constant factor?

1.2 Applications to graphs

At a high level, our result states that permutations avoiding π can be decomposed into permutations avoiding a fixed pattern τ , where the length of the decomposition depends on π , but τ does not. Trading patterns with twin-width, we get results on graphs that convey a similar flavour of transformation into structures with universally bounded twin-width. For sparse graphs, this transformation is the *d-subdivision*, which consists of replacing every edge in a given graph by a path of length d .

Theorem 1.4. *There is a universal constant c and a function f such that for any graph G with twin-width k and no $K_{t,t}$ -subgraph, the $f(k,t)$ -subdivision of G has twin-width at most c .*

This can be extended to dense graphs (and more generally binary relational structures) using *first-order (FO) transductions*, which roughly speaking are transformations that can be described using first-order logic.

Theorem 1.5. *There is a fixed strict class \mathcal{C} of permutations such that for any class \mathcal{D} of binary structures, \mathcal{D} has bounded twin-width iff there exists an FO transduction Φ satisfying $\mathcal{D} \subseteq \Phi(\mathcal{C})$.*

We then say that \mathcal{D} is an FO transduction of \mathcal{C} . Theorem 1.5 nicely echoes similar characterizations for linear clique-width and clique-width, namely the fact that a graph class has bounded linear clique-width if and only if it is an FO transduction of a linear order, and bounded clique-width if and only if it is an FO transduction of a tree order [14]. Fittingly, these results can also be phrased in terms of pattern-avoiding permutation classes. Indeed a class has bounded linear clique-width if and only if it is an FO transduction of $\text{Av}(21)$, and bounded clique-width if and only if it is an FO transduction of $\text{Av}(231)$, or equivalently an FO transduction of the class of separable permutations (see for example [8, Proposition 8.1]).

Reconstructing a graph G from the d -subdivision of G with fixed d is a simple example of FO transduction, hence Theorem 1.4 is a special case of Theorem 1.5 utilizing a specific transduction. Theorems 1.4 and 1.5 are based on a representation of any factorisation $\sigma = \sigma_m \circ \dots \circ \sigma_1$ as a system of paths between ordered sets of vertices: each permutation σ_i is represented by a matching between two orders, which are joined together into paths of length m , see Figure 1 for an example. When the permutations σ_i are separable (or more generally avoid any fixed pattern), these path system representations have bounded twin-width, independently of the length of the factorisation. Furthermore, for any fixed length m , there is an FO transduction depending only on m that can reconstruct σ from the path system representation. We obtain Theorem 1.5 by combining this with the following result of Bonnet et al.




Figure 1: A factorisation of 163524 into three separable permutations, represented as a path system.

Theorem 1.6 ([8]). *For any class \mathcal{D} of binary structures with bounded twin-width, there exists a pattern avoiding class of permutations \mathcal{C} and an FO transduction Φ such that $\mathcal{D} \subseteq \Phi(\mathcal{C})$.*

Notice that the difference with Theorem 1.5 is that the class \mathcal{C} depends on \mathcal{D} .

1.3 Overview of the proof of Theorem 1.1

We see an n -element permutation σ as two linear orders $<, \prec$ over $[n]$.

Mixed minor free permutations. We fix a permutation π , and consider permutations σ in $\text{Av}(\pi)$, the class of all permutations avoiding π as a pattern. This class has bounded twin-width [18, 7], and it was further proven that there is a third linear order $<_{\text{tw}}$ over $[n]$ such that the adjacency matrix¹ of the binary structure $([n], <, \prec)$ has no c_π -mixed minor for some $c_\pi = 2^{O(|\pi|)}$, where a k -mixed minor of a matrix M is a k -by- k division of M into consecutive blocks of rows and columns such that each of the k^2 cells formed by the division has at least two distinct row vectors and at least two distinct column vectors. We then say that such cells are *mixed*.

It was cleverly observed that k -almost mixed minors is a better notion to conduct induction on [24], where one waives the mixedness condition off for the diagonal cells. Besides, one can remark that mixed minors and almost mixed minors are linearly tied, in the sense that if a matrix has a k -almost mixed minor it has a $\lfloor k/2 \rfloor$ -mixed minor (while a mixed minor is a fortiori an almost mixed minor).

If we summarize the situation after these opening moves, we have a permutation σ (in fact, two permutations whose product equals σ) whose matrix encoding has no k -almost mixed minor with $k = 2c_\pi$. We wish to decompose this permutation as a product of $g(k)$ permutations all of which have matrix encodings without $(k - 1)$ -almost mixed minor. If this scheme works while $k \geq 3$, we will have written our permutation as a product of $f(k)$ permutations without 2-almost mixed minor, which can be easily shown to be separable permutations. (Indeed the two excluded permutations 2413 and 3142 both admit a 2-mixed minor.) Theorem 1.1 will then come from an appropriate choice for the function g , which happens to make f single exponential.

Delayed substitutions. We decompose σ onto a tree T called *delayed structured tree*. This tree is inspired by the work of Pilipczuk and Sokolowski [24], and was formally introduced by Bourneuf and Thomassé [10] on their way to show that every graph class of bounded twin-width is polynomially χ -bounded (i.e. can be properly colored with a number of colors polynomial in their

¹At this point, we do not wish to concern the reader with the actual encoding.

maximum clique size). We now detail how to build an appropriate delayed structured tree in the particular case of permutations. This will serve as an informal definition of *delayed structured trees* and *delayed substitutions*, crucial elements of our proof.

Here it helps to renormalize (\prec, \prec_t) such that \prec is the natural order over $[n]$. This totally fixes the successor relation of \prec to an ordering eventually matching the left-to-right order of the leaves of T (as represented in Figure 2); leaves, which are in one-to-one correspondence with $[n]$. Let $a_1 < a_2 < \dots < a_n$ be the ordering of \prec . We recursively partition $\langle a_1, a_2, \dots, a_n \rangle$ in the following way. At the start, the root of T (the only node thus far) is labeled with the list $\langle a_1, a_2, \dots, a_n \rangle$.

If $\{a_i, a_{i+1}, \dots, a_j\}$ is a non-singleton interval along \prec (i.e. if $\langle a_i, a_{i+1}, \dots, a_j \rangle$ can be reordered as $j - i + 1$ consecutive integers), as it initially occurs with $i = 1$ and $j = n$ (when $n > 1$), we arbitrarily split $\langle a_i, a_{i+1}, \dots, a_j \rangle$ into two non-empty lists $\langle a_i, a_{i+1}, \dots, a_k \rangle$ and $\langle a_{k+1}, \dots, a_j \rangle$, each labeling a child of the node labeled by $\langle a_i, a_{i+1}, \dots, a_j \rangle$. Lists $\langle a_i \rangle$ of size 1 are given a unique child, a leaf labeled by a_i . Finally, the interesting case is when $\{a_i, a_{i+1}, \dots, a_j\}$ does not consist of a single interval of \prec , but several, say, I_1, \dots, I_s with $s \geq 2$. Then we cut $\langle a_i, a_{i+1}, \dots, a_j \rangle$ into $q + 1$ sublists $\langle a_i, a_{i+1}, \dots, a_{h_1} \rangle, \langle a_{h_1+1}, a_{h_1+2}, \dots, a_{h_2} \rangle, \dots, \langle a_{h_q+1}, a_{h_q+2}, \dots, a_j \rangle$ that are maximal for the property of being contained in a single I_ℓ .

Let us have a look, in Figure 2, at an example of delayed structured tree if σ is such that the ordering of \prec is 3, 1, 10, 5, 16, 9, 8, 4, 2, 6, 7, 11, 17, 15, 13, 12, 18, 14.




Figure 2: A well-chosen delayed structured tree of a permutation σ encoded as two linear orders \prec, \prec_t over $[18]$, where \prec is the left-to-right order on the leaves 3, 1, 10, 5, 16, 9, ... whereas \prec_t is the natural order 1, 2, 3, ... The permutations (\prec_t, \prec_t) over the grandchildren of a node t in a given colour are drawn in the same colour by the successor relation of \prec_t (\prec_t being the left-to-right order). If all these permutations are in \mathcal{C} , σ is said to be *obtained by delayed substitution from \mathcal{C}* .

that the list should be split in two (arbitrarily), and we accordingly chose to cut the first sublist after 11. Moving on to the left child of the root, the current set is $[1, 11] \cup \{16\}$ (two intervals), hence its three children labeled $\langle 3, 1, 10, 5 \rangle, \langle 16 \rangle, \langle 9, 8, 4, 2, 6, 7, 11 \rangle$.

Note that the figure also displays “local” linear orders (in different colours) over the grandchildren of a given node t (of that same colour). These orders \prec_t are naturally inherited from \prec , here

(among other possible options) by giving to a grandchild the label of its leftmost descendant. Let us focus on the blue order. Its successor relation defines the following ordering: 3 (leftmost descendant of the orange node), 9 (same with the purple node), 14, 15 (same with the yellow node), 16, 17, 18. As the left-to-right order $<$ on the leaves of T also naturally defines the left-to-right order $<_t$ on the grandchildren of t , every node t (with grandchildren) naturally defines the permutation $(<_t, \prec_t)$.

Finally a *delayed substitution from a class \mathcal{C} of permutations* is such a decomposition when every $(<_t, \prec_t)$ belongs to \mathcal{C} . We show (see Lemma 3.5) that any permutation obtained by delayed substitution from \mathcal{C} is the product of three permutations in \mathcal{C} , as long as the class \mathcal{C} is closed under some simple operations: substitution, taking patterns, and inverse. We then zoom in on the permutations $(<_t, \prec_t)$ and decompose them further.

Decomposing permutations of height two. What we gain here is that the permutations $(<_t, \prec_t)$ have a structured tree of height two. Interestingly the children of t naturally define a partition \mathcal{P} over the universe of $(<_t, \prec_t)$, that is, over the grandchildren of t . Each part of \mathcal{P} is by construction an interval along $<_t$. We look at how these parts behave and interlace w.r.t. \prec_t .

Two parts $P_1, P_2 \in \mathcal{P}$ are said *mixed* if there are $x_1, y_1 \in P_1$ and $x_2, y_2 \in P_2$ such that $x_1 \prec_t x_2 \prec_t y_1 \prec_t y_2$ or $x_2 \prec_t x_1 \prec_t y_2 \prec_t y_1$. This definition exactly matches the mixed cells in the definition of (almost) mixed minors. This brings us to a common theme in this context: colouring the conflict graph (red graph, or graph of the *mixed* relation) with a bounded number of colours to separately deal with large chunks without conflicting pairs. This is usually performed directly on a partition from a sequence witnessing low twin-width (see for instance [4, 3]), where the conflict graph of the partition has bounded maximum degree. In our case, like in [10], the subtlety is that \mathcal{P} does not come from a partition sequence, and its *mixed* graph may have large maximum degree.

Getting rid of the mixedness. Nevertheless, the mixed graph has no K_k (clique on k vertices) nor $K_{k,k}$ (biclique, with k vertices being fully adjacent to k other vertices) subgraphs, as they would entail the presence of a k -almost mixed minor. We can split the edge set of the *mixed* graph into a bounded-degeneracy graph and an *overlap graph* (also known as *circle graph*); the latter is degenerate as a consequence of excluding large clique and biclique subgraphs. Therefore the *mixed* graph is itself $h(k)$ -degenerate, for some polynomial function h , hence properly colourable with $h(k) + 1$ colours; see Lemma 4.4.

Once again we can zoom in on a more structured subset of our permutation, where no pairs of parts are mixed. This is because the so-called *shuffles*—permutation operations that can recreate $(<_t, \prec_t)$ from the $h(k) + 1$ “non-mixed” permutations—can be expressed as bounded products of separable permutations; see the more precise formulation of Lemma 2.2. We now want to decompose a permutation $(<', \prec')$ in the same setting as $(<_t, \prec_t)$ but with the additional assumption that its partition \mathcal{P} has no pair of mixed parts.

Endgame—the induction trick on the almost mixed minor. Our last goal is to show that for each $P \in \mathcal{P}$ the permutation restricted to the elements of P (at the level of the grandchildren of t) as well as the permutation over the parts of \mathcal{P} (at the level of the children of t) are simpler than the initial permutation σ . As announced, this takes the form of proving that these permutations have no $(k-1)$ -almost mixed minor. Once this is established, we are finally done by invoking Lemma 4.1, which decomposes $(<', \prec')$ into a bounded product involving the above-mentioned permutations and separable permutations.

We are left with the task of building a k -almost mixed minor of σ (a contradiction) from a supposed $(k-1)$ -almost mixed minor \mathcal{D} over the children of some node v of T . Let us recall that




Figure 3: Illustration of the induction trick. By construction of the children of v , v_1 and v'_1 have each a descendant (or simply child) x_1 and x'_1 , respectively, such that there is a y_1 outside the descendants of v in between x_1, x'_1 in the order \prec' . Assume R_1, \dots, R_5 is a 5-almost mixed minor on a choice of one descendant per child of v . For the sake of clarity, let us make the simplifying assumptions that the minor is symmetric ($C_i = R_i$ for every i), every y_i is to the left of v , and thus can be grouped together with **first**, the minimum along \prec' , in an interval $R_0 = C_0$ of \prec' , and **first** $\prec' x$ for every x descendant of v . It then holds that R_0 is mixed with each R_i (for $i \in [5]$) since for instance **first** $\prec' x_1 \prec' y_1 \prec' x'_1$ and **first** $\prec' x'_2 \prec' y_2 \prec' x_2$.

“starting a new sibling” in the construction of T means that the last element of the current child and the first element of the new child are in two distinct intervals I_a, I_b . Hence there is an element y outside the children of v such that $I_a \prec' y \prec' I_b$. We get one such y for every row part and every column part of \mathcal{D} . By leveraging again the *shuffle* trick, we can reduce to the case when all such vertices are to the left (by \prec') of the children of v (see Theorem 5.2). Then the minimum of \prec' together with these vertices y can make a k -th interval (in rows and columns) that is mixed with every previous part of \mathcal{D} , thereby yielding a k -almost mixed minor (see Figure 3).

1.4 Perspectives

We prove that for permutations, the minimum size of an avoided pattern, and the minimum length of a factorisation into separable permutations, are functionally equivalent. However, the bounds we give are huge, the dependency being exponential in one direction, and doubly exponential in the other, leaving open the question of finding tighter bounds. While avoided patterns have been extensively studied, there is to our best knowledge no work on the length of factorisations into separable permutations. A natural question is whether this new parameter can be computed efficiently.

In a different direction, a natural generalisation of our result would be to extend this factorisation to matrices. It was shown in [6] that a class \mathcal{M} of 0,1-matrices closed under taking submatrices has growth at most single-exponential if and only if \mathcal{M} has bounded twin-width. It is then natural to ask if every matrix of \mathcal{M} can be factorised into a bounded product (say, computed in \mathbb{F}_2) of some basic 0,1-matrices playing the role of separable permutations.

Finally, a crucial point of this work is perhaps the used strategy: To decompose a pattern-free permutation (\prec_1, \prec_2) , first find a third order \prec_3 compatible with both orders (i.e. such that (\prec_1, \prec_3) and (\prec_2, \prec_3) are k -mixed minor free), and then express both (\prec_1, \prec_3) and (\prec_2, \prec_3) as

delayed structured trees (see Figure 2). How delayed substitutions untangle the complexity of permutations and other binary structures is a novel and promising tool that, we believe, can be exploited in many other ways.

1.5 Organisation of the paper

In Section 2, we introduce a number of definitions, conventions, and simple results, which we use throughout this work. Notably we recall the definitions of ordered trees, substitutions and shuffles of permutations, degeneracy and proper colourings in the context of sparse circle graphs, and relevant properties of classes with bounded twin-width. Note that we will *not* actually need the definition of twin-width and of transductions (they can both be found for instance in [7]).

Section 3 focuses on delayed substitutions, detailing how to construct them, and how to decompose them into products and substitutions. Section 4 shows how to further decompose the permutations that appear in the delayed substitution, by studying their mixed graphs. The core of the proof of Theorem 1.1 is in Section 5, using lemmas from the previous sections to decompose k -almost mixed free permutations into $(k - 1)$ -almost mixed free ones. In Section 6, we explain how to implement Theorem 1.1 in linear time. Finally, Section 7 shows how to apply Theorem 1.1 to graphs, to prove Theorems 1.4 and 1.5.

2 Preliminaries

This section defines our notations and conventions, and introduces some important constructions on permutations. We denote by $[n]$ the interval of integers $\{1, \dots, n\}$.

2.1 Ordered trees

Trees are used throughout this work to describe decompositions of permutations. We work with rooted trees. Vertices of the trees are *nodes*. The *ancestors* of a node x are the nodes in the unique path from x to the root r , and the *parent* of x is the first node on this path. We will also speak of *descendants*, *children*, *grandchildren*, *siblings* (nodes with same parent), and *cousins* (non-sibling nodes with the same grandparent). The set of leaves of a tree T is denoted by $L(T)$. For any node $t \in T$, we denote by $T(t)$ the subtree rooted at t (i.e. consisting of the descendants of t), and by $L(t) \subseteq L(T)$ the set of leaves of $T(t)$.

Here an *ordered tree* $(T, <)$ is a rooted tree T equipped with a linear order $<$ on $L(T)$, such that for each node $t \in T$, the leaves $L(t)$ form an interval for $<$. In that case we also say that $<$ is *compatible* with T . It is natural to think of $<$ as a *left-to-right* order.

If $(T, <)$ is an ordered tree and $x, y \in T$ are not in an ancestor–descendant relationship, then $L(x), L(y)$ are disjoint, and each of them is an interval for $<$, hence either $L(x) < L(y)$, or $L(y) < L(x)$. We then naturally extend $<$ to x, y by $x < y$ in the former case, and $y < x$ in the latter. In particular, $<$ induces a linear order $<_t$ on the children of any given internal node t . This can be reversed. Given a tree T , and an order $<_t$ on the children of t for each internal node t , one can define $<$ compatible with T as follows: for any leaves x, y , we have $x < y$ iff $x' <_t y'$, where t is the last (that is, closest) common ancestor of x, y , and x' and y' are the children of t that are ancestors of x and y , respectively.

2.2 Permutations and biorders

Throughout the paper, we will represent permutations as *biorders*, i.e. the superposition of two linear orders on the same set. Precisely, a permutation $\sigma \in \mathfrak{S}_n$ is represented by the biorder $B_\sigma = ([n], <, <_\sigma)$, where $<$ is the natural order on $[n]$, and $x <_\sigma y$ iff $\sigma(x) < \sigma(y)$. Conversely, from a biorder $B = (X, <_1, <_2)$ on n elements, we define the associated permutation σ_B such that if $x \in X$ is in i th position for $<_1$, then it is in position $\sigma_B(i)$ for $<_2$. This is a bijection between permutations, and biorders considered up to isomorphism. In the proofs, this bijection will be left implicit, and we may for example say that a biorder belongs to some class of permutations.

The inverse of a permutation is obtained by swapping the roles of the two orders:

$$\sigma_{(X, <_1, <_2)}^{-1} = \sigma_{(X, <_2, <_1)}. \quad (1)$$

Similarly, factoring $(X, <_1, <_2)$ corresponds to choosing a third order $<_3$:

$$\sigma_{(X, <_3, <_2)} \circ \sigma_{(X, <_1, <_3)} = \sigma_{(X, <_1, <_2)}. \quad (2)$$

The notion of pattern has a very natural definition in terms of biorders. A pattern of $(X, <_1, <_2)$ is any induced restriction $(Y, <_1, <_2)$ for some subset $Y \subseteq X$.

2.2.1 Substitutions

Let us now introduce *substitutions* of permutations, which will be used throughout this work, and are a natural introduction to the more general delayed decompositions presented in Section 3. Substitutions are easily visualised through permutation matrices: the substitution of τ inside σ corresponds to replacing one of the 1-entries in the matrix of σ by the matrix of τ . In this work, we are interested in *iterated substitutions*—i.e. starting from a permutation σ in a class \mathcal{C} , and iteratively substituting other permutations of \mathcal{C} inside it—and we will consistently use substitution to mean iterated substitutions.

A precise description of an (iterated) substitution can be given by a tree as follows. Let T be a tree, and $<, \prec$ two linear orders on $L(T)$, both compatible with T . For every internal node $t \in T$, they induce orders $<_t, \prec_t$ on the children of t . Then the permutation $(<, \prec)$ is obtained by substitution of all the $\{(<_t, \prec_t)\}_{t \in T \setminus L(T)}$. In this case, we also say that $(<, \prec)$ is obtained by *substitution along T* . Conversely, any iterated substitution process can be expressed through such a tree.

For a class \mathcal{C} of permutations, let \mathcal{C}^* denote the *substitution closure* of \mathcal{C} , that is, σ is in \mathcal{C}^* if it is obtained by substitution of permutations in \mathcal{C} . We say that \mathcal{C} is *closed under substitution* if $\mathcal{C} = \mathcal{C}^*$. Note that $(\mathcal{C}^*)^* = \mathcal{C}^*$, meaning that the substitution closure is indeed closed under substitution. For example, the class \mathcal{S} of *separable* permutations is the substitution closure of the class $\{12, 21\}$ of permutations on two elements. It happens that \mathcal{S} is precisely the class of permutations excluding the patterns 3142 and 2413 [9]. This implies that the inverse of a separable permutation is also a separable permutation.

Lemma 2.1. *If \mathcal{C}, \mathcal{D} are substitution-closed classes of permutations, then so are $\mathcal{C} \circ \mathcal{D}$ and \mathcal{C}^{-1} .*

Proof. Let us prove more generally that for any classes \mathcal{C}, \mathcal{D} , the following holds:

$$(\mathcal{C}^{-1})^* = (\mathcal{C}^*)^{-1} \quad \text{and} \quad (\mathcal{C} \circ \mathcal{D})^* \subseteq \mathcal{C}^* \circ \mathcal{D}^*. \quad (3)$$

Let $\sigma \in (\mathcal{C}^{-1})^*$. Then, σ is obtained by substitution of $\{(<_t, \prec_t)\}_{t \in T \setminus L(T)}$ along a tree T , and for every internal node t the permutation $(<_t, \prec_t)$ is in \mathcal{C}^{-1} . Let $<, \prec$ be the corresponding linear orders

on $L(T)$. Then, $\sigma = (\prec, \prec)$, hence $\sigma^{-1} = (\succ, \prec)$. However, (\succ, \prec) is obtained by substitution of $\{(\prec_t, \prec_t)\}_{t \in T \setminus L(T)}$ along T , and for every internal node t , $(\prec_t, \prec_t) = (\prec_t, \prec_t)^{-1}$, so $(\prec_t, \prec_t) \in \mathcal{C}$. Thus $\sigma^{-1} \in \mathcal{C}^*$, and $\sigma \in (\mathcal{C}^*)^{-1}$. This proves that $(\mathcal{C}^{-1})^* \subseteq (\mathcal{C}^*)^{-1}$. Applying this result to \mathcal{C}^{-1} , we get $\mathcal{C}^* \subseteq ((\mathcal{C}^{-1})^*)^{-1}$. This immediately yields $(\mathcal{C}^*)^{-1} \subseteq (\mathcal{C}^{-1})^*$.

Let $\sigma \in (\mathcal{C} \circ \mathcal{D})^*$. Then, σ is obtained by substitution of $\{(\prec_t, \prec_t)\}_{t \in T \setminus L(T)}$ along a tree T , and for every internal node t the permutation (\prec_t, \prec_t) is in $\mathcal{C} \circ \mathcal{D}$. Thus, for every such t , there exists a linear order \prec'_t on the children of t such that $(\prec_t, \prec'_t) \in \mathcal{D}$ and $(\prec'_t, \prec_t) \in \mathcal{C}$. Let \prec, \prec' and \prec be the corresponding linear orders on $L(T)$. Then, $\sigma = (\prec, \prec) = (\prec', \prec) \circ (\prec, \prec')$. However, (\prec', \prec) is obtained by substitution of $\{(\prec'_t, \prec_t)\}_{t \in T \setminus L(T)}$ along T , so $(\prec', \prec) \in \mathcal{C}^*$. Similarly, $(\prec, \prec') \in \mathcal{D}^*$, thus $\sigma \in \mathcal{C}^* \circ \mathcal{D}^*$.

When \mathcal{C}, \mathcal{D} are closed under substitution, the inclusions (3) simplify to $(\mathcal{C}^{-1})^* = \mathcal{C}^{-1}$ and $(\mathcal{C} \circ \mathcal{D})^* \subseteq \mathcal{C} \circ \mathcal{D}$, proving the result. \square

2.2.2 Shuffles

A permutation $\sigma \in \mathfrak{S}_n$ is a k -shuffle if there is some partition $\bigsqcup_{i=1}^k X_i$ of its domain $[n]$ such that the restriction (in the sense of patterns) of σ to any X_i is the identity permutation. More generally, given a class \mathcal{C} of permutations, σ is a k -shuffle of \mathcal{C} if there is a partition $\bigsqcup_{i=1}^k X_i = [n]$ such that the restriction of σ to any X_i is in \mathcal{C} .

Let us say that a class \mathcal{C} of permutations is *closed under symmetry* if it is closed by conjugating with the decreasing permutation: if $\sigma \in \mathcal{C}$, then the permutation $i \mapsto n + 1 - \sigma(n + 1 - i)$ should also be in \mathcal{C} . In terms of biorders, this corresponds to replacing each of the two linear orders by its dual, i.e. its reverse order. Remark that if \mathcal{C}, \mathcal{D} are closed under symmetry, then so is $\mathcal{C} \circ \mathcal{D}$. For example, the class \mathcal{S} of separable permutations is closed under symmetry.

Lemma 2.2. *For a class \mathcal{C} of permutations closed under substitution and symmetry, any 2^k -shuffle of \mathcal{C} is in $\mathcal{S}^k \circ \mathcal{C} \circ \mathcal{S}^k$. In particular, any 2^k -shuffle is in \mathcal{S}^{2^k} .*

Proof. Let $\sigma = (X, \prec_1, \prec_2)$ be a 2^k -shuffle of \mathcal{C} . Then there is a partition of its domain $X = Y \uplus Z$ such that the restrictions of σ to both Y and Z are 2^{k-1} -shuffles of \mathcal{C} .

Define an intermediate linear order \prec_1 by

- $Y \prec_1 Z$,
- inside Y , \prec_1 coincides with \prec_1 , and
- inside Z , \prec_1 coincides with the dual of \prec_1 .

We claim that (X, \prec_1, \prec_1) is separable. Indeed, let x be the minimum of X for \prec_1 . If $x \in Y$, then x is the minimum of \prec_1 , while if $x \in Z$ it is the maximum of \prec_1 . Either way, we can separate X into $\{x\}$ and $X \setminus \{x\}$, and proceed by induction on the latter. This scheme is a decomposition of (X, \prec_1, \prec_1) as a separable permutation, with the specificity that the separating tree is a caterpillar. We define \prec_2 similarly with regards to \prec_2 , so that (X, \prec_2, \prec_2) is separable.

Consider now the permutation (X, \prec_1, \prec_2) . Since we have $Y \prec_i Z$ for both $i = 1, 2$, it suffices to consider this permutation restricted to either Y or Z . By construction, the restriction to Y is exactly (Y, \prec_1, \prec_2) , and the restriction to Z is (Z, \prec_1, \prec_2) up to symmetry. These two permutations are 2^{k-1} -shuffles of \mathcal{C} , hence by induction are in $\mathcal{S}^{k-1} \circ \mathcal{C} \circ \mathcal{S}^{k-1}$. Since this class is closed under symmetry and substitution, we obtain that (X, \prec_1, \prec_2) is also in $\mathcal{S}^{k-1} \circ \mathcal{C} \circ \mathcal{S}^{k-1}$, and conclude by composing with (X, \prec_1, \prec_1) and (X, \prec_2, \prec_2) .

Finally, the remark that 2^k -shuffles are in \mathcal{S}^{2^k} is obtained by applying the result with \mathcal{C} being the class of identity permutations. \square

2.3 Twin-width

Guillemot and Marx [18] introduced a width of permutations, that was later generalised to graphs, matrices, and binary structures by Bonnet et al. in [7] under the name of twin-width.

The twin-width of a graph $G = (V, E)$ is defined through *partition sequences* (equivalent to so-called *contraction sequences*), that is, sequences $\mathcal{P}_n, \dots, \mathcal{P}_1$ of partitions of V such that

1. $\mathcal{P}_n = \{\{x\} : x \in V\}$ is the partition into singletons,
2. $\mathcal{P}_1 = \{V\}$ is the trivial partition, and
3. \mathcal{P}_i is obtained from \mathcal{P}_{i+1} by merging two parts.

Two parts $X \neq Y$ in \mathcal{P}_i are *homogeneous* if there are either no edges or all possible edges between X and Y in G . If X, Y are not homogeneous (i.e. there is at least one edge and one non-edge between them), they are said to be *in error*. The *width* of the partition sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ is the maximum over $i \in [n]$ and $X \in \mathcal{P}_i$ of the number of parts of \mathcal{P}_i in error with X —called the *error degree of X* . Finally, the twin-width of G is the minimum width of a partition sequence. This definition readily generalises to *binary relational structures* (V, R_1, \dots, R_k) , consisting of a domain, or vertex set V , and a number of binary relations $R_i \subseteq V^2$. Remark in particular that biorders are binary relational structures. In this setting, parts $X, Y \in \mathcal{P}_i$ are *in error w.r.t. the relation R_j* if there are $x_1, x_2 \in X, y_1, y_2 \in Y$ such that $(x_1, y_1) \in R_j$ and $(x_2, y_2) \notin R_j$, or symmetrically $(y_1, x_1) \in R_j$ and $(y_2, x_2) \notin R_j$. The width of the partition sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ is now the maximum over all $j \in [k], i \in [n]$, and $X \in \mathcal{P}_i$ of the error degree of X in \mathcal{P}_i w.r.t. R_j .

As fundamental property, twin-width exhibits a duality with grid-like structures in matrices—or universal patterns of permutations. This result, based on the Marcus–Tardos theorem [22], appears in the seminal work of Guillemot and Marx, and can be restated as follows.

Theorem 2.3 ([18, Theorem 4.1]). *For any permutation π , there is $c_\pi = 2^{O(|\pi|)}$ such that if σ avoids π as pattern, then the biorder B_σ associated with σ has twin-width at most c_π .*

The bound $c_\pi = 2^{O(|\pi|)}$ in the restated theorem follows from a later improvement in the Marcus–Tardos constant due to Fox [17]. A *division \mathcal{D}* of a matrix consists of partitions \mathcal{R}, \mathcal{C} of its rows and columns respectively into *intervals*. It is a *k -division* if the partitions have k parts each. A *cell* of the division is the submatrix induced by $X \cap Y$ for some $X \in \mathcal{R}, Y \in \mathcal{C}$. A *k -grid* in a 0,1-matrix is a k -division in which every cell contains a 1-entry. Remark that the order of rows and columns is crucial in these definitions because divisions are required to be partitions into intervals.

The notion of grid is tightly linked to permutation patterns, and to Theorem 2.3. In order to generalise the latter to dense matrices, Bonnet et al. [7] introduced *mixed minors*. Say that a matrix is *horizontal* if all rows are constant, i.e. in each row, all the coefficients are the same (or equivalently, all columns of the matrix are equal). Symmetrically, it is *vertical* if all columns are constant, and finally it is *mixed* if it is neither horizontal nor vertical. Then, a *k -mixed minor* in a matrix is a k -division in which every cell is mixed. A matrix is said *k -mixed free* if it does not have a k -mixed minor. The generalisation of Theorem 2.3 can then be stated as follows [7, Theorem 5.8]: There is a function f such that if a graph G admits a k -mixed free adjacency matrix, then G has twin-width at most $f(k)$.

We will not need this result, but only its much simpler-to-prove converse.

Lemma 2.4 ([7, first part of Theorem 5.4]). *If (V, R_1, \dots, R_k) is a binary relational structure with twin-width t , then there exists a linear order $<$ on V such that the adjacency matrix of each relation R_i , ordered by $<$, is $(2t + 2)$ -mixed free.*

Pilipczuk and Sokołowski introduced the following subtle relaxation of mixed minors in [24]. Consider a k -division $\mathcal{D} = (\mathcal{R}, \mathcal{C})$ of a matrix M , and enumerate the parts as $\mathcal{R} = \{R_1, \dots, R_k\}$,

$\mathcal{C} = \{C_1, \dots, C_k\}$, following the order of rows and columns. We say that \mathcal{D} is a *k-almost mixed minor* if every cell of \mathcal{D} is mixed, except possibly the diagonal cells $R_i \cap C_i$ for $i \in [k]$. Clearly a *k-mixed minor* is also a *k-almost mixed minor*, and it is not hard to check that a *2k-almost mixed minor* yields a *k-mixed minor* by merging the parts by pairs. Thus these two notions are equivalent up to a factor of two. The point of almost mixed minors is to allow a proof by induction: Ideally, a structure without *k-almost mixed minors*, i.e. *k-almost mixed free*, is decomposed into smaller parts that have no $(k-1)$ -almost mixed minors. The same scheme is much harder to achieve using (regular) mixed minors.

2.4 Graphs, colouring, degeneracy

A *k-colouring* of a graph G is a map $c : V(G) \rightarrow [k]$ that assigns distinct colours to adjacent vertices. Equivalently, it is a partition into k colour classes $c^{-1}(1), \dots, c^{-1}(k)$, each of which is an independent set (i.e. a set of pairwise non-adjacent vertices).

A graph G is *k-degenerate* if G and all its subgraphs have a vertex of degree at most k . A well-known characterisation is that G is *k-degenerate* if and only if there is an acyclic orientation of its edges such that all vertices have out-degree at most k . Degenerate graphs can be efficiently coloured in the following sense.

Lemma 2.5 ([23]). *Any k-degenerate graph G can be $(k+1)$ -coloured by a greedy algorithm. Further, if the graph is given through adjacency lists, then this colouring can be computed in time $O(k|V(G)|) = O(|E(G)|)$.*

A closely related parameter is the *maximum edge density* of G , that is, the maximum, taken over all subgraphs H of G , of $|E(H)| / |V(H)|$. It is easy to see that *k-degenerate* graphs have maximum edge density at most k , and conversely graphs with maximum edge density k are *2k-degenerate*. If G admits an orientation (which may contain cycles) in which vertices have out-degree at most k , then its maximum edge density is at most k .

2.4.1 Circle graphs

Let \mathcal{I} be a family of intervals of some linear order. Two intervals A, B *overlap* if they intersect, but neither contains the other. The *overlap graph* of \mathcal{I} is the graph whose vertex set is \mathcal{I} , and where intervals are adjacent exactly when they overlap. Overlap graphs are also known as circle graphs, because they can be described by the intersections of a set of chords of a circle.

Let K_t denote the complete graph on t vertices, and $K_{t,t}$ the balanced complete bipartite graph on $2t$ vertices. We say that G is *H-subgraph-free* (resp. *H-free*) to mean that G does not have a subgraph (resp. induced subgraph) isomorphic to H . The following lemma shows that circle graphs excluding bicliques (and thus, cliques) as subgraphs are degenerate. Since we will use it for colouring, it is worth comparing it to a result of Davies and McCarty stating that K_t -free circle graphs have bounded chromatic number [16, 15]. Their hypothesis is weaker (arbitrary bicliques are allowed), and the $O(t \log t)$ bound in [15] is better, but this bound applies only to the chromatic number and not to the degeneracy, and the proof is far more complex.

Lemma 2.6. *Circle graphs that are K_t -free and $K_{t,t}$ -subgraph-free have maximum edge density at most $2(t-1)^2$.*

Proof. Let \mathcal{I} be a family of intervals whose overlap graph is K_t -free and $K_{t,t}$ -subgraph-free. We will define an orientation of the edges with out-degree at most $2(t-1)^2$.

Consider two overlapping intervals $A, B \in \mathcal{I}$. Relative to this edge, say that $C \in \mathcal{I}$ is a *private ancestor* of A if $A \subseteq C$ and $B \not\subseteq C$, and symmetrically with B . We consider that A and B are also private ancestors of themselves.

Claim 2.7. *Every private ancestor of A overlaps with every private ancestor of B .*

Proof. Let C, D be private ancestors of A, B respectively, with possibly $C = A$ or $D = B$. Pick $x \in a \cap B$ (which is non-empty as they overlap), and $y \in a \setminus D$ (which is non-empty by assumption), and symmetrically $z \in B \setminus C$. Then, because $A \subseteq C$ and $B \subseteq D$, we also have $x \in C \cap D$, $y \in C \setminus D$, and $z \in D \setminus C$, hence C and D overlap. \diamond

Thus, there is a biclique between A and its private ancestors on the one hand, and B and its private ancestors on the other. Since the overlap graph does not contain a $K_{t,t}$ subgraph, it follows that either A or B has less than t private ancestors. We orient the edge AB towards A if it has less than t private ancestors, and towards B otherwise.

Let us bound the out-degree of this orientation. Fix $I \in \mathcal{I}$, an interval with endpoints $x < y$, and consider $N^+(I)$ the set of $A \in \mathcal{I}$ with an edge oriented from I to A . Any $A \in N^+(I)$ contains either x or y , but not both. Let X be the set of intervals in $N^+(I)$ containing x , and consider the inclusion poset (X, \subseteq) .

Claim 2.8. *The poset (X, \subseteq) does not contain a chain of length t .*

Proof. If $A_1 \subsetneq \dots \subsetneq A_s$ are in X , then A_1 overlaps with I , and each A_i is a private ancestor of A_1 w.r.t. the edge $A_1 I$. From the orientation of the edge from I to A_1 , we know that A_1 has less than t private ancestors, hence $s < t$. \diamond

Claim 2.9. *The poset (X, \subseteq) does not contain an antichain of size t .*

Proof. If $A, B \in X$ are incomparable for inclusion, then they overlap because both contain x . Thus an antichain in the poset (X, \subseteq) is a clique in the overlap graph. \diamond

By Dilworth's theorem, it follows from the two claims that $|X| \leq (t-1)^2$. The same reasoning applies to intervals in $N^+(I)$ containing y , and we obtain $|N^+(I)| \leq 2(t-1)^2$. \square

Finally, for algorithmic purposes, let us sketch how to efficiently compute circle graphs from their interval representation.

Lemma 2.10. *Given \mathcal{I} a family of intervals in $[n]$, one can compute the overlap graph G of \mathcal{I} in time $O(n + |\mathcal{I}| + |E(G)|)$.*

Proof. We iterate for i from 1 to n , while maintaining the list L of intervals containing i , sorted by their right endpoint. In the i th iteration, we consider all intervals of the form $[i, j]$ sorted by *decreasing* values of j , and naively insert each of them in L . Suppose that $I = [i, j]$ is thus inserted into L in position k , after J_1, \dots, J_{k-1} . For $\ell < k$, if $J_\ell = [a_\ell, b_\ell]$, then it must be that $a_\ell < i \leq b_\ell < j$. It follows that I overlaps with J_1, \dots, J_{k-1} . Furthermore, any interval $[a, b]$ which overlaps with I and satisfies $a < i$ will be one of the J_ℓ . Thus, during the insertion of I , we can output all edges from I to intervals "to its left", and the cost of the insertion of I is proportional to the number of such edges. Finally, having inserted all intervals with left endpoint i , we remove from the beginning of L all intervals with right endpoint i , before proceeding with $i + 1$. \square

2.5 First-order transductions

For the sake of Theorem 1.5, let us introduce first-order logic, interpretations, and transductions.

A *relational signature* is a set $\Sigma = \{R_1, \dots, R_k\}$ of *relation symbols*, each with an associated *arity* $ar(R_i) \in \mathbb{N}$. A Σ -*structure* S is defined by a *domain* (or vertex set) $V(S)$, and, for each relation symbol $R \in \Sigma$ of arity $r := ar(R)$, a *realisation* $R(S) \subseteq V(S)^r$ of this relation. For example, (simple undirected) graphs are structures over the signature $\{E\}$ with $ar(E) = 2$ (with the specificity that the realisation of E is symmetric and irreflexive), and biorders are structures over the signature $\{<_1, <_2\}$, $ar(<_i) = 2$ (the realisations being total orders). When all relation symbols have arity 2, the signature and its structures are called *binary*. Recall from Section 2.3 that binary structures are the objects for which twin-width is defined.

An *FO formula* ϕ over the language of Σ can quantify on vertices, and test whether a relation $R \in \Sigma$ holds for a given tuple of vertices. For example, over the language of graphs, the formula

$$\phi(x, y) = E(x, y) \vee \exists z E(x, z) \wedge E(z, y)$$

expresses that x, y are at distance at most two. If ϕ is a formula over the language of Σ and S is a Σ -structure, then $S \models \phi$ denotes that ϕ is *satisfied* by S , which is defined in the obvious way.

Given two signatures Σ, Γ , an *FO interpretation* Φ from Σ to Γ is a map, defined in FO logic, from Σ -structures to Γ -structures. Precisely, Φ is described by giving

1. for each relation $R \in \Gamma$, a formula $\phi_R(x_1, \dots, x_r)$ over the language of Σ , with as many free variables as the arity $r := ar(R)$ of R , and
2. one last formula $\phi_{dom}(x)$ with one free variable, again over the language of Σ .

Given a Σ -structure S , its image $\Phi(S)$ is defined as follows.

1. The domain consists of vertices satisfying ϕ_{dom} , i.e.

$$V(\Phi(S)) = \{x \in V(S) : S \models \phi_{dom}(x)\}.$$

2. For each symbol $R \in \Gamma$, the realisation $R(\Phi(S))$ is described by ϕ_R , i.e. for $x_1, \dots, x_r \in V(\Phi(S))$,

$$(x_1, \dots, x_r) \in R(\Phi(S)) \iff S \models \phi_R(x_1, \dots, x_r).$$

Transductions are a non-deterministic generalisation of interpretations. Let Σ be a signature, and C_1, \dots, C_k be k unary relation symbols (i.e. with arity 1), disjoint from Σ . The k -*colouring* is the one-to-many operation which maps a Σ -structure S to all possible extensions of S as $(\Sigma \uplus \{C_1, \dots, C_k\})$ -structures S^+ , meaning that $V(S) = V(S^+)$ and $R(S) = R(S^+)$ for any $R \in \Sigma$, while the $C_i(S^+)$ are chosen to be arbitrary subsets of $V(S^+)$. An *FO transduction* is the composition of a k -colouring (with k fixed), followed by an FO interpretation.

It is folklore that interpretations and transductions can be composed.

Lemma 2.11. *If Φ, Ψ are FO transductions (resp. interpretations) from Σ to Γ and Γ to Δ respectively, then the composition $\Psi \circ \Phi$ is an FO transduction (resp. interpretation) from Σ to Δ .*

Finally, transductions preserve bounded twin-width in the following sense.

Theorem 2.12 ([7, Theorem 8.1]). *For any FO transduction Φ , there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for any binary structures S and $T \in \Phi(S)$, $tw(T) \leq f(tw(S))$.*

3 Delayed substitutions

In this section, we recall the definition of delayed decompositions from [10, Section 2], which generalise substitutions. Since we consider biorders and not arbitrary binary relations, we obtain some additional structure on these delayed decompositions. We show that they can be expressed as a bounded product of substitutions.

3.1 Definition

A *delayed structured tree* $(T, <, \{\prec_t\}_{t \in T})$ consists of an ordered tree $(T, <)$, equipped with, for each node $t \in T$, a linear order \prec_t on the *grandchildren* of t . This is analogous to the trees describing substitutions, except that \prec_t is defined on the grandchildren instead of the children, hence ‘delayed’. We add the technical requirement that each leaf is a single child (with no siblings), so that whenever $x \neq y$ are leaves, their closest ancestor is at distance at least 2.

The *realisation* of this delayed structured tree is the structure $(L(T), <, \prec)$, where for two leaves x, y , we have $x \prec y$ if and only if $x' \prec_t y'$, where t is the closest ancestor of x, y , and x', y' are the grandchildren of t which are ancestors of x, y respectively. In general, this binary relation \prec is not an order, as it might not be transitive. We call the delayed structured tree *well-formed* if \prec is a linear order, so that the realisation is a permutation on $L(T)$. We only consider well-formed delayed structured trees. Remark that in the realisation of the tree, we only use the linear order \prec_t between cousins; the order between siblings is irrelevant.

We say that the permutation $(<, \prec)$ is obtained by *delayed substitution* from the permutations $\{(<_t, \prec_t)\}_{t \in T}$, understood as permutations on the grandchildren of t . If for each $t \in T$ the permutation $(<_t, \prec_t)$ on its grandchildren is in a given class \mathcal{C} , we also say that the delayed structured tree T is *labelled* with permutations in \mathcal{C} , and that the permutation $(<, \prec)$ is obtained by *delayed substitution from \mathcal{C}* .

3.2 Distinguishability

The linear order \prec is defined on the leaves of T . We extend it to a partial order on all nodes of T where $x \prec y$ iff $L(x) \prec L(y)$, i.e. all descendants of x are strictly before all descendants of y for \prec .

Remark 3.1. *Let x, y be nodes in T , and t their least common ancestor. If t is at distance at least 2 of both x and y , then x, y are comparable by \prec , meaning either $x \prec y$ or $y \prec x$. In particular, if x, y are at the same level in T but are not siblings, then they are comparable by \prec .*

Being incomparable by \prec is not an equivalence relation, even when restricted to siblings: one may have three siblings x, y, z such that $x \prec z$, but $L(x), L(y)$, resp. $L(y), L(z)$ are interleaving for \prec . However, we can define an equivalence relation by considering how other nodes can separate siblings: let x, y be two children of a node t , and $v \in L(T) \setminus L(t)$. By Remark 3.1, v is comparable by \prec to both x and y . We say that x and y are *distinguished* by v if $x \prec v \prec y$ or $y \prec v \prec x$. One can check that if $v \in L(T) \setminus L(t)$ distinguishes x, y , it must be a descendant of a cousin of x, y . Finally, we call x, y *indistinguishable*, denoted $x \sim y$, if x, y are siblings with parent t , and no $v \in L(T) \setminus L(t)$ distinguishes x and y . See Figure 4 for an example.

Lemma 3.2. *Indistinguishability is an equivalence relation.*

Proof. From its definition, one can see that \sim is reflexive and symmetric. Let x, y, z be siblings with parent t , and suppose that v distinguishes x, z , say $x \prec v \prec z$. Since $v \notin L(t)$, y and v are comparable. Thus it either holds that $v \prec y$, in which case v distinguishes x, y , or $y \prec v$, and v distinguishes y, z . By contraposition, \sim is transitive. \square




Figure 4: Indistinguishability for the grandchildren of a node t . The values and arrows represent the linear order \prec_t , and colours represent equivalence classes of indistinguishability.

Lemma 3.3. *Let x, y be grandchildren of t such that $x \not\sim y$. Then $x \prec y$ if and only if $x \prec_t y$.*

Proof. Between cousins with grandparent t , the orders \prec and \prec_t coincide by definition of the realisation. Thus, consider x, y siblings distinguished by some v , say $x \prec v \prec y$. Let v' be the cousin of x, y that is the ancestor of v . Now x and v' are comparable by \prec , and $v' \prec x$ is impossible as it would imply $v \prec x$, hence we have $x \prec v'$. As noted at the beginning of the proof, this implies $x \prec_t v'$. The same reasoning gives $v' \prec_t y$, hence $x \prec_t y$, i.e. \prec and \prec_t coincide on the pair x, y as desired. \square

A crucial property of indistinguishability is that it is compatible with \prec as follows.

Lemma 3.4. *Let A be an equivalence class of \sim . Then, among the leaves $L(T)$, the subset $L(A) := \bigcup_{x \in A} L(x)$ is an interval for \prec .*

Proof. Let A be an equivalence class of \sim . All elements of A are siblings. Consider t their parent. Let $x, z \in L(A)$ and $y \in L(T) \setminus L(A)$ be leaves, and suppose for a contradiction that $x \prec y \prec z$. Let x', z' be the ancestors in A of x, z respectively. We have two cases to consider.

1. If y is not a descendant of t , then x' and z' are comparable by \prec with y , and it must be that $x' \prec y \prec z'$, contradicting that x', z' are indistinguishable.
2. Otherwise, y is a descendant of some child y' of t , which is distinguishable from x' and from z' . Thus there exist $v_1, v_2 \notin L(t)$ such that $x' \prec v_1 \prec y'$ and $y' \prec v_2 \prec z'$. But then we also have $x' \prec v_1 \prec z'$, a contradiction. \square

3.3 Factoring delayed substitutions

We will now prove that any delayed substitution can be decomposed into products of substitutions.

Lemma 3.5. *Let \mathcal{C} be a class of permutations closed under substitution, taking patterns and inverse. Then any permutation obtained by delayed substitution from \mathcal{C} is in \mathcal{C}^3 .*

Proof. Let $(X, <, \prec)$ be a biorder, realised by a delayed structured tree $(T, <, \{\prec_t\}_{t \in T})$ with leaves $L(T) = X$, labelled in \mathcal{C} . The linear order $<$ is compatible with T , while \prec is realised by the delayed substitution.

Let us define an intermediate linear order $<'$ on X . For each internal node t of T , and for each child x of t , choose an arbitrary descendant $f_t(x) \in X$ of x . Then, on the children of t , define $x <'_t y$ if and only if $f_t(x) \prec f_t(y)$. These local orders extend to a linear order $<'$ on X , which by construction is compatible with T .

Claim 3.6. *For each $t \in T$, the permutation (\prec, \prec'_t) on the children of t is in \mathcal{C} .*

Proof. Let A be the set of children of t . For each $x \in A$, let $g(x)$ be the sole child of x which is an ancestor of $f_t(x)$. Thus $g(x)$ is a grandchild of t . For $x \neq y \in A$, we have $g(x) \prec g(y)$ if and only if $x \prec y$, by compatibility of \prec with T . Furthermore, the closest ancestor of $f_t(x), f_t(y)$ is t , hence $f_t(x) \prec f_t(y)$ if and only if $g(x) \prec_t g(y)$, by definition of \prec . Thus the permutation (\prec, \prec'_t) on A is equal to the permutation (\prec, \prec_t) on $g(A)$, which is a pattern of the permutation (\prec, \prec_t) on all grandchildren of t . The latter is in \mathcal{C} by hypothesis. \diamond

Since \mathcal{C} is closed under substitution, it follows that the permutation (X, \prec, \prec') is in \mathcal{C} .

We now quotient the tree T by indistinguishability. Since \sim can only identify siblings, the quotient $T' = T/\sim$ naturally has a tree structure. Furthermore, since the leaves of T are required to be single children, no two leaves are identified, hence the set of leaves of T' is exactly X .

Claim 3.7. *The tree T' is compatible with both \prec' and \prec .*

Proof. Lemma 3.4 precisely proves that T' is compatible with \prec : each node of T' , which is an equivalence class of \sim , corresponds to an interval of (X, \prec) .

Since \prec' is compatible with T , to show that it is also compatible with $T' = T/\sim$, it is sufficient to consider only the children of a given $t \in T$. That is, it is enough to prove that among the children of t , any equivalence class A for indistinguishability is an interval for \prec' . Suppose for a contradiction that there are children $x \prec' y \prec' z$ of t with $x, z \in A$ and $y \notin A$. Since $x \not\prec y$, it must be that $x \prec y$ or $y \prec x$, but the latter is impossible as it would imply $y \prec' x$. Thus $x \prec y$, and similarly $y \prec z$, contradicting Lemma 3.4. \diamond

Thus the permutation (X, \prec', \prec) is obtained by substitution along the tree T' . It only remains to show that this structured tree is labelled with permutations in \mathcal{C}^2 .

Fix a node $\bar{s} \in T'$, which in T is an equivalence class of \sim . Let $t \in T$ be the parent of the nodes of \bar{s} (only siblings can be in the same equivalence class for \sim). The children in T' of \bar{s} are equivalence classes of grandchildren of t in T . Denote by $\bar{x}_1, \dots, \bar{x}_k$ these children, where the representatives $R = \{x_1, \dots, x_k\}$ are grandchildren of t . Recall that $L(x_i) \subseteq X$ is the set of leaves descendant of x_i . The subsets $L(x_1), \dots, L(x_k)$ are non-interleaved for \prec and \prec' because these linear orders are compatible with T . Furthermore, they are non-interleaved for \prec because the x_i are pairwise distinguished. Thus R is equipped with the three linear orders \prec, \prec', \prec . In the structured tree T' which realises (X, \prec', \prec) , \bar{s} is labelled with (R, \prec', \prec) , hence we only need the following to conclude.

Claim 3.8. *The permutation (R, \prec', \prec) is in \mathcal{C}^2 .*

Proof. We decompose the permutation (R, \prec', \prec) into (R, \prec', \prec) and (R, \prec, \prec) . This may at first seem counterproductive, as we go back to the initial permutation (\prec, \prec) , but the point is that we now only consider a small subset R , rather than all X .

We already know that (X, \prec, \prec') is in \mathcal{C} , hence so is (R, \prec', \prec) , which is a pattern of its inverse. Furthermore, the x_i are pairwise distinguished grandchildren of t , hence the linear orders \prec and \prec_t coincide on R by Lemma 3.3. Thus (R, \prec, \prec) is a pattern of (\prec, \prec_t) , which is in \mathcal{C} . \diamond

As (X, \prec, \prec) can be factorised into $(X, \prec, \prec') \in \mathcal{C}$ and $(X, \prec', \prec) \in \mathcal{C}^2$, it is in \mathcal{C}^3 . \square

3.4 Constructing delayed structured trees

Finally, we show how to express any permutation as a delayed substitution.

Lemma 3.9. *Let $\sigma = (X, <, \prec)$ be a biorder. There is a structured tree $(T, <, \{\prec_t\}_{t \in T})$ whose realisation is σ , that satisfies the following.*

1. *For any node $t \in T$, the linear order \prec_t on its grandchildren is obtained by a choice of representatives. That is, if A is the set of grandchildren of t , then there is a mapping $f : a \rightarrow X$, with $f(a)$ descendant of a , such that for $a, b \in A$, $a \prec_t b$ if and only if $f(a) \prec f(b)$.*
2. *If x, y are consecutive siblings along $<$, then x and y are distinguished for the linear order \prec , except possibly if x and y do not have any other sibling.*

In the previous statement, Condition 1 is a technical requirement ensuring that the permutations labelling T are patterns of σ , so that hypotheses on σ can also be applied to the former. Condition 2 is crucial for the induction on almost mixed minors in Section 5. Informally, it ensures that the subpermutations $\{(<, \prec_t)\}_{t \in T}$ labelling T are strictly simpler—in the sense of almost mixed minors—than the global permutation $(X, <, \prec)$.

Proof of Lemma 3.9. We construct the tree T inductively starting from the root, choosing for each internal node t the interval (for $<$) $L(t) \subseteq X$ that, at the end of the construction, will be the set of leaves descendant of t . We maintain the condition that whenever t' is a child of t and $x \in X \setminus L(t)$, then x does not split $L(t')$, i.e. either $x \prec L(t')$ or $x \succ L(t')$.

Initially, there is only the root r with $L(r) = X$. If t and $L(t)$ have already been constructed, then we create the children of t by the following rules.

- If $L(t) = \{x\}$ is a singleton, we add a leaf x as the sole child of t , and the construction stops there for this branch.
- If $L(t)$ has size at least 2, and is an interval of (X, \prec) (in addition to being one for $<$), then we add two children u, v to t , and split $L(t)$ arbitrarily into two intervals $L(u), L(v)$ for $<$. Remark that this case occurs at least for the root r .

If $x \notin L(t)$, then x does not split $L(t)$, hence a fortiori it splits neither $L(u)$ nor $L(v)$, as required. On the other hand, this means that u and v will be indistinguishable in T , which is why condition 2 is waived for nodes with only two children.

- Otherwise, enumerate $L(t)$ as $x_1 < \dots < x_k$. Say that a subset $A \subseteq L(t)$ is a *local module* of $L(t)$ if no $y \notin L(t)$ splits A (for \prec). We greedily partition $L(t)$ into local modules A_1, \dots, A_l : A_1 is $\{x_1, \dots, x_{i_1}\}$ with i_1 chosen maximal such that A_1 is a local module, then A_2 is $\{x_{i_1+1}, \dots, x_{i_2}\}$ with again i_2 maximal, etc.

By construction, no element $y \notin L(t)$ can split any A_i . Further, A_i, A_{i+1} are distinguished (in fact, the last element of A_i is distinguished from the first element of A_{i+1}), as otherwise A_i could have been extended. Also, the number l of local modules in this partition is at least 2, as otherwise $L(t)$ would be an interval of (X, \prec) and we would fall in the previous case.

With this tree constructed, consider a node $t \in T$, and u, v two grandchildren of t that are not siblings. Then the condition maintained during the construction ensures that no $x \in L(u)$ splits $L(v)$, and symmetrically no $y \in L(v)$ splits $L(u)$. It follows that either $L(u) \prec L(v)$ or $L(v) \prec L(u)$. In short, the order \prec is well defined between cousins in T .

Now, define the linear order \prec_t on the grandchildren through an arbitrary choice of representatives, as required in condition 1 of the lemma. Then for cousins u, v , we have $u \prec_t v$ if and only if $u \prec v$. In turn, this implies that the realisation of $(T, <, \{\prec_t\}_{t \in T})$ is the permutation $(X, <, \prec)$ from which we started. Condition 1 is satisfied by definition of \prec_t , and condition 2 was verified during the construction of T . \square

4 Partitions and mixity

In this section, we introduce tools to further decompose a single level of a delayed structured tree. Their purpose is similar to the right module partitions of [10, Section 4], but once again working with biorders gives stronger results and simpler proofs.

4.1 Definitions

Let (X, \prec) be a linear order, and \mathcal{P} a partition of X . For any subsets $X_1, X_2 \in \mathcal{P}$, we distinguish three cases:

1. X_1, X_2 do not interleave, i.e. $X_1 \prec X_2$ or $X_2 \prec X_1$.
2. Restricted to $X_1 \cup X_2$, X_1 is an interval which is said to *split* X_2 in two, written $X_1 \sqsubset X_2$, or vice versa.
3. Neither X_1 nor X_2 is an interval in $X_1 \cup X_2$, and we say that X_1, X_2 are *mixed*.

These three cases correspond to zones in the adjacency matrix that are respectively constant, vertical or horizontal, and mixed.

Assume now that no pair of subsets in \mathcal{P} is mixed—we then say that \mathcal{P} is *non-mixed*. For any $X_1 \in \mathcal{P}$, let \overline{X}_1 denote the interval closure of X_1 , i.e. the interval of (X, \prec) between the minimum and maximum of X_1 . Then we have $X_1 \prec X_2$ if and only if $\overline{X}_1 \prec \overline{X}_2$, and $X_1 \sqsubset X_2$ if and only if $\overline{X}_1 \subset \overline{X}_2$. In particular, either $\overline{X}_1, \overline{X}_2$ are disjoint, or one contains the other. Thus (X, \mathcal{P}) induces a laminar family, meaning that subsets in \mathcal{P} are the nodes of a rooted forest F , where X_1 is a descendant of X_2 if and only if $X_1 \sqsubset X_2$. Furthermore, the order \prec is defined for any pair of subsets $X_1, X_2 \in \mathcal{P}$ that are not in an ancestor–descendant relationship. Thus \prec gives to F the structure of an ordered forest: each connected component of F is an ordered tree, and furthermore \prec orders the components.

4.2 Non-mixed partitions

We now consider a biorder $(X, <, \prec)$, and a partition \mathcal{P} of X into intervals for $<$. In later proofs, such a structure will arise from delayed substitutions as follows: Given a node t in a delayed structured tree, X is the set of grandchildren of t , the partition \mathcal{P} is the one given by the children of t , $<$ is the linear order compatible with the tree, while \prec is the linear order on the grandchildren of t given as part of the delayed structured tree. Our goal is to decompose this biorder into simpler permutations.

Lemma 4.1. *Let \mathcal{C} be a substitution-closed class of permutations. Let $(X, <, \prec)$ be a biorder, and \mathcal{P} a partition of X into intervals of $<$. Suppose that \mathcal{P} is non-mixed w.r.t. \prec . Finally, assume that*

1. *for each part $P \in \mathcal{P}$, $(P, <, \prec)$ is in \mathcal{C} , and*




Figure 5: The tree T , ordered by \prec (left) and \prec' (right) respectively. Leaves are elements of X , and their numbering is according to \prec . Internal nodes correspond to parts in \mathcal{P} , with each colour representing a part of \mathcal{P} . The transversal R is indicated by circled leaves, and decides how \prec' orders the parts.

2. there exists a transversal R of \mathcal{P} (i.e. a choice of a single element in each part of \mathcal{P}) such that the permutation $(R, <, \prec)$ is in \mathcal{C} .

Then $(X, <, \prec)$ is in $\mathcal{S}^2 \circ \mathcal{C}$, where \mathcal{S} denotes the class of separable permutations.

Proof. Let us define an intermediate linear order \prec' on X as follows:

- each part $P \in \mathcal{P}$ is an interval of (X, \prec') ,
- inside each part P , the linear order \prec' coincides with \prec , and
- between parts $P_1, P_2 \in \mathcal{P}$, the order is given by the representatives in R , i.e. if x_i is the sole element in $P_i \cap R$, then $P_1 \prec' P_2$ if and only if $x_1 \prec x_2$.

Claim 4.2. *The permutation $(X, <, \prec')$ is in \mathcal{C} .*

Proof. The permutation inside a given part $P \in \mathcal{P}$ is $(P, <, \prec)$, which is known to be in \mathcal{C} . Furthermore, the permutation on the set \mathcal{P} is isomorphic to $(R, <, \prec)$, which is also in \mathcal{C} . Therefore $(X, <, \prec')$ can be expressed as a substitution from permutations in \mathcal{C} , with a structured tree of depth 2. \diamond

We now focus on (X, \prec', \prec) , and will prove that this permutation is obtained by substitution from 2-shuffles. Such permutations are in \mathcal{S}^2 by Lemma 2.2 and since \mathcal{S} is closed under substitution.

Recall the structure on (X, \prec) and \mathcal{P} described in Section 4.1: Since the partition \mathcal{P} is non-mixed w.r.t. \prec , the *splitting* partial order \sqsubset (related to \prec) gives \mathcal{P} the structure of a rooted forest F , which furthermore is compatible with the linear order \prec . We construct a tree T from F by

- adding a new root, whose children are the roots of the connected components of F , and
- for each $P \in \mathcal{P}$ node of F , adding each element $x \in P$ as a new child of P .

The leaves of T are exactly the elements of X , and we will describe (X, \prec', \prec) as a substitution along T . See Figure 5 for an illustration.

Claim 4.3. *The linear orders \prec' and \prec are compatible with T .*

Proof. Consider a node t of T , other than the leaves and the root for which it is straightforward that $L(t)$ is an interval of $<'$ and of $<$. Thus t corresponds to some part $P \in \mathcal{P}$ (and we now identify t to P). Then it is simple to check that $L(P)$, the set of leaves descending from P in T , is exactly $\bigcup_{P' \sqsubseteq P} P'$, i.e. the union of parts which are descendants of P in F . We claim that this is an interval for both $<'$ and $<$.

Suppose otherwise towards a contradiction. Then there are $u, w \in L(P)$ and $v \notin L(P)$ such that either $u < v < w$ or $u <' v <' w$. Let U, V, W be the parts of \mathcal{P} containing u, v, w respectively. We have $U, W \sqsubseteq P$, and $V \not\sqsubseteq P$. Recall that \overline{P} denotes the interval between the minimum and maximum of P for $<$. Then we have $U, W \subseteq \overline{P}$, while V is disjoint from \overline{P} . This implies $u, w \in \overline{P}$ but $v \notin \overline{P}$, hence it cannot be that $u < v < w$ since \overline{P} is an interval of $<$.

For $<'$, we also need to consider the representative $v' \in V \cap R$. For the same reasons as above, we have either $v' < U \cup W$, or $U \cup W < v'$, which implies $V <' U \cup W$ or $U \cup W <' V$ respectively. Either way, $u <' v <' w$ is impossible. \diamond

To conclude the proof, we only need to show that for any internal node P of T , the permutation $(X, <', <)$ restricted to its children is a 2-shuffle. We partition the children of P in two categories: leaves (which are in X), and internal nodes (which are in \mathcal{P}). We claim that $(X, <', <)$ restricted to either of these categories is the identity.

1. The leaves that are children of P are exactly the elements of P , and inside P the linear orders $<'$ and $<$ coincide by construction.
2. The parts of \mathcal{P} that are children of P in T are exactly the children of P in F . We know that they are totally ordered under $<$, i.e. if $A, B \in \mathcal{P}$ are children of P , then either $A < B$ or $B < A$. By construction of $<'$, if $A < B$ (resp. $B < A$) then $A <' B$ (resp. $B <' A$). It follows from these two remarks that $<'$ and $<$ coincide on the children of P in F . \square

4.3 Separating mixed parts

The previous lemma shows how to decompose a permutation that is non-mixed w.r.t. a given partition. We will now generalise it to permutations with few pairs of mixed parts. Let $(X, <)$ be a linear order and \mathcal{P} a partition of X . The associated *mixed graph* is the graph with vertex set \mathcal{P} , and where two parts $A, B \in \mathcal{P}$ are adjacent when they are mixed.

Lemma 4.4. *Let G be the mixed graph of a partition, and assume that it is K_t -free and $K_{t,t}$ -subgraph-free. Then G is $(4t^2 - 1)$ -degenerate.*

Proof. Let $G = (\mathcal{P}, E)$ be the mixed graph of a partition \mathcal{P} of a linear order $(X, <)$. We partition the edges of G into $E = E_1 \uplus E_2$ as follows:

1. An edge $AB \in E$ is in E_1 if the interval closures satisfy $\overline{A} \subset \overline{B}$ or $\overline{B} \subset \overline{A}$.
2. The remaining edges are in E_2 . Thus, a pair $(A, B) \in \mathcal{P}$ is an edge in E_2 if and only if \overline{A} and \overline{B} overlap.

The graph (\mathcal{P}, E_2) is precisely the overlap graph of the interval closures of parts of \mathcal{P} . Thus by Lemma 2.6, its maximum edge density is at most $2(t-1)^2$. We will show that (\mathcal{P}, E_1) is $(t-1)$ -degenerate, hence has maximum edge density at most $t-1$. Then, their union (\mathcal{P}, E) has maximum edge density at most $t-1 + 2(t-1)^2 \leq 2t^2 - 1$, hence it is $(4t^2 - 1)$ -degenerate as desired.

To this end, we define an acyclic orientation of (\mathcal{P}, E_1) for which the out-degree is at most $t-1$. If AB is an edge in E_1 and $\overline{A} \subset \overline{B}$, we orient it from A to B . We claim that for any part $A \in \mathcal{P}$,




Figure 6: Representation of six parts of \mathcal{P} . Each row is a part P , the dots being elements of P , and the line being the interval \overline{P} . The left-to-right order is \prec . The B_i s are out-neighbours of A in the graph (\mathcal{P}, E_1) , and $y_i \in B_i \cap \overline{A}$ witnesses that A and B_i are mixed. For any $i \neq j$, y_i, y_j together with the endpoints of $\overline{B_i}, \overline{B_j}$ ensure that B_i and B_j are mixed. The part C satisfies $\overline{A} \subset \overline{C}$, but is not mixed with A as $C \cap \overline{A} = \emptyset$. The part D is mixed with A because \overline{A} and \overline{D} overlap, hence $AD \in E_2$.

the out-degree of A in (\mathcal{P}, E_1) is at most $t - 1$. Indeed, let B_1, \dots, B_k be the out-neighbours of A , so that each B_i is mixed with A , and $\overline{A} \subset \overline{B_i}$ (see Figure 6). We will show that A, B_1, \dots, B_k is a clique in (\mathcal{P}, E) , which implies the claim as this graph has no clique of size t . To this end, we need to show that any two $B_i \neq B_j$ are mixed. Without loss of generality, assume that $\overline{B_i} \not\subset \overline{B_j}$, hence there is some $x \in B_i$ outside of $\overline{B_j}$, say $x < B_j$. On the other hand, since B_i is mixed with A , there must be $y \in B_i \cap \overline{A}$, and thus $y \in B_i \cap \overline{B_j}$. Thus we have two points in B_i , one inside the interval closure of B_j and one outside. This implies that B_i and B_j are mixed. \square

Combining this with Lemma 4.1, we obtain the following, where (as throughout the paper) the log function is in base 2.

Lemma 4.5. *Let \mathcal{C} be a class of permutations closed under substitution, symmetry, and taking patterns. Let $(X, <, \prec)$ be a biorder, and \mathcal{P} a partition of X into intervals of $<$, satisfying the following:*

1. *The mixed graph of \mathcal{P} for the linear order (X, \prec) does not contain K_t or $K_{t,t}$ subgraphs.*
2. *For each part $P \in \mathcal{P}$, the permutation $(P, <, \prec)$ is in \mathcal{C} .*
3. *There exists a transversal R of \mathcal{P} (i.e. a choice of a single element in each part of \mathcal{P}) such that the permutation $(R, <, \prec)$ is in \mathcal{C} .*

Then the permutation $(X, <, \prec)$ is in $\mathcal{S}^{k+2} \circ \mathcal{C} \circ \mathcal{S}^k$, with $k = \lceil 2 \log t \rceil + 2$.

Proof. Using Lemmas 4.4 and 2.5, the mixed graph of \mathcal{P} is $4t^2$ -colourable. Fix a proper $4t^2$ -colouring, and for any fixed colour $c \in [4t^2]$, let \mathcal{P}_c be the set of parts of colour c , and $X_c = \bigcup_{P \in \mathcal{P}_c} P$ the points contained therein. Since the colouring is proper, no two parts of \mathcal{P}_c are mixed. Thus the restricted biorder $(X_c, <, \prec)$ with the partition \mathcal{P}_c satisfies the conditions of Lemma 4.1, and the permutation $(X_c, <, \prec)$ is in $\mathcal{S}^2 \circ \mathcal{C}$. Hence, $(X, <, \prec)$ is a $4t^2$ -shuffle of permutations in $\mathcal{S}^2 \circ \mathcal{C}$. We conclude by applying Lemma 2.2 with $k = \lceil \log(4t^2) \rceil = \lceil 2 \log t \rceil + 2$. \square

5 Reducing the size of almost mixed minors

Given a biorder $\sigma = (X, <, \prec)$ define its *adjacency matrix* M_σ as the 0,1-matrix whose rows and columns are X ordered by $<$, and with a 1 in position (x, y) if and only if $x \prec y$, i.e. if and only if $\sigma(x) < \sigma(y)$. Equivalently, M_σ is obtained by applying σ^{-1} to both the rows and the

columns of the full upper triangular matrix. This should not be confused with the more usual sparse permutation matrix, with a 1 in positions of the form $(x, \sigma(x))$. We say that σ is k -almost mixed free as shorthand for M_σ being k -almost mixed free.

Using the tools from the previous sections, we will show how to factorize k -almost mixed free permutations, by induction on k . The base case is the following simple remark.

Lemma 5.1. *Any 2-almost mixed free permutation is separable.*

Proof. A permutation that is not separable contains either 3142 or 2413 as pattern, both containing 2-almost mixed minors in their adjacency matrices (they in fact also contain 2-mixed minors). See Figure 7 for the adjacency matrix of 3142, the matrix of 2413 is its transpose. \square

		3	1	4	2
3	0	0	1	0	
1	1	0	1	1	
4	0	0	0	0	
2	1	0	1	0	

Figure 7: Adjacency matrix of the permutation $\sigma = 3142$, which contains a 2-mixed minor represented by dashed lines. The matrix has a 1 at the intersection of the i th row and j th column if and only if $\sigma(i) < \sigma(j)$.

Let \mathcal{A}_k denote the class of k -almost mixed free permutations, and recall that \mathcal{S} is the class of separable permutations.

Theorem 5.2. *If $\mathcal{A}_{k-1} \subseteq \mathcal{S}^r$ for some $r \in \mathbb{N}$, then $\mathcal{A}_k \subseteq \mathcal{S}^s$ for*

$$s = 3r + 12 \lceil \log k \rceil + 28.$$

Proof. Assume that $\mathcal{A}_{k-1} \subseteq \mathcal{S}^r$, and consider $\sigma = (X, <, \prec) \in \mathcal{A}_k$ a k -almost-mixed free permutation. Denote by **first** and **last** the minimum and maximum of X w.r.t. $<$. For the linear order \prec , **first** and **last** split X into three intervals as

$$X_1 \prec \mathbf{first} \prec X_2 \prec \mathbf{last} \prec X_3,$$

up to swapping **first** and **last**. We consider each X_i independently, before recombining X_1, X_2, X_3 , and $\{\mathbf{first}, \mathbf{last}\}$. Let $X' = X_i$ be one of these intervals, and $\sigma' = (X', <, \prec)$ the restricted permutation.

Let $(T, <, \{\prec_t\}_{t \in T})$ be a delayed structured tree for σ' obtained by Lemma 3.9. Consider any internal node $t \in T$, let A be its set of children, and consider a transversal $S \subseteq X'$ of A , i.e. all elements of S are descendants of t , and each child $v \in A$ has exactly one descendant in S .

Claim 5.3. *The transversal S admits a partition into two parts $S = S_L \uplus S_R$ such that the permutation $(<, \prec)$ restricted to each of S_L, S_R is $(k-1)$ -almost mixed free.*

Proof. Enumerate the children of t as v_1, \dots, v_ℓ in the linear order $<$, and let x_i be the descendant of v_i in S . We can assume $\ell > 2$ as the claim is otherwise trivial, hence condition 2 of Lemma 3.9 gives that for all $i \in \{1, \dots, \ell-1\}$, v_i and v_{i+1} are distinguishable. Thus, we can choose some $y_i \in X' \setminus L(t)$ such that $v_i \prec y_i \prec v_{i+1}$, or $v_{i+1} \prec y_i \prec v_i$.

If $y_i < L(t)$ (resp. $y_i > L(t)$) we say that v_i is *split to the left* (resp. *to the right*). Thus every v_i except v_ℓ is split either to the left or to the right. We partition S accordingly, so that S_L (resp. S_R) contains x_i only if v_i is split to the left (resp. to the right). We place x_ℓ in either S_L or S_R arbitrarily. We will prove that $(S_L, <, \prec)$ is $(k-1)$ -almost mixed free, the case of S_R being symmetrical.

Suppose for a contradiction that $(S_L, <, \prec)$ has a $(k-1)$ -almost mixed minor. It is given by two partitions of $(S_L, <)$ into $(k-1)$ intervals, say

$$\mathcal{R} = \{R_1 < \cdots < R_{k-1}\} \quad \text{and} \quad \mathcal{C} = \{C_1 < \cdots < C_{k-1}\}$$

such that for every $i \neq j \in [k-1]$, R_i and C_j are mixed in the linear order \prec . Note in particular that all R_i and C_j have size at least 2. Starting from \mathcal{R} and \mathcal{C} , we build two new partitions $\mathcal{R}' = \{R'_0 < \cdots < R'_{k-1}\}$ and $\mathcal{C}' = \{C'_0 < \cdots < C'_{k-1}\}$, which will form a k -almost mixed minor of (a pattern of) σ , as follows:

- Initially set $R'_0 = \{\mathbf{first}\}$, $C'_0 = \{\mathbf{first}\}$, and $R'_i = R_i$ and $C'_i = C_i$ for $i \geq 1$.
- For each $i \geq 1$, consider x_j the minimum of R_i w.r.t. $<$. Since R_i has size at least 2, we have $j < \ell$, and either $x_{j+1} \in R_i$, or x_{j+1} is not in S_L . We add x_{j+1} to R'_i in the latter case, and in both cases we add y_j to C'_0 . Recall that $y_j < L(t)$ since v_j is split left.
- The same operation is applied to \mathcal{C}' , adding y_j to R'_0 , and x_{j+1} to C'_i if needed.

After this modification, all elements of R'_0 are smaller than $L(t)$ for $<$, hence \mathcal{R}' is still a partition of a subset of X into intervals of $<$, and similarly with \mathcal{C}' . The parts which were originally in \mathcal{R}, \mathcal{C} have only increased, hence for $i \neq j$, $i, j \geq 1$, R'_i and C'_j are mixed. We claim that R'_0 is mixed with C'_i for any $i > 0$, and symmetrically for C'_0, R'_i , which implies that the new $\mathcal{R}', \mathcal{C}'$ form a k -almost mixed minor in (a submatrix of) M_σ , a contradiction.

Indeed, let x_j be the smallest element of R_i , so that $y_j \in C'_0$ distinguishes x_j and x_{j+1} , i.e. either $x_j \prec y_j \prec x_{j+1}$ or $x_{j+1} \prec y_j \prec x_j$. On the other hand, by construction of X' , \mathbf{first} is not interleaved with x_j, x_{j+1} (or any elements of X'), i.e. either $\mathbf{first} \prec x_j, x_{j+1}$ or $x_j, x_{j+1} \prec \mathbf{first}$. The above implies that $\{\mathbf{first}, y_j\}$ and $\{x_j, x_{j+1}\}$ are mixed. The former is contained in C'_0 , while the latter is contained in R'_i , proving the claim.

The same reasoning also holds for S_R , using \mathbf{last} instead of \mathbf{first} and adding parts R'_t and C'_t instead of R'_0 and C'_0 . \diamond

It follows from Claim 5.3 that for any transversal S of the children of a node t , the permutation $(S, <, \prec)$ is a 2-shuffle of permutations in \mathcal{A}_{k-1} , hence is in \mathcal{S}^{r+2} by Lemma 2.2.

Fix now a node $t \in T$, and let A be the set of grandchildren of t . Consider the permutation $(A, <, \prec_t)$, and the partition \mathcal{P} of A induced by the children of t . Recall that \mathcal{S}^{r+2} is closed under substitution, inverse and taking patterns. Let us check that the three conditions of Lemma 4.5 are satisfied.

1. If the mixed graph of \mathcal{P} for the linear order (A, \prec_t) contains either K_k or $K_{k,k}$ as a subgraph, then the corresponding parts form a k -almost mixed minor in the matrix of (A, \prec_t) ordered by $<$. We know that the linear order \prec_t is defined through a choice of representatives, hence $(A, <, \prec_t)$ is a pattern of $(X, <, \prec)$, meaning that the latter would also have a k -almost mixed minor, a contradiction.
2. Fix a part $P \in \mathcal{P}$. The elements of P are the children of some child t' of t , thus the restricted permutation $(P, <, \prec_t)$ is isomorphic to $(S, <, \prec)$ for some transversal S of P . By Claim 5.3 applied on t' , this permutation is in \mathcal{S}^{r+2} .

3. Similarly, applying Claim 5.3 this time to t shows that for any transversal S of \mathcal{P} , the permutation $(S, <, \prec_t)$ is in \mathcal{S}^{r+2} .

It follows from Lemma 4.5 that the permutation $(A, <, \prec_t)$ is in $\mathcal{S}^{r'}$ with

$$\begin{aligned} r' &= (r + 2) + 2([2 \log k] + 2) + 2 \\ &\leq r + 4 \lceil \log k \rceil + 8. \end{aligned}$$

Next, we apply Lemma 3.5 to the delayed structured tree T and conclude that $(X', <, \prec)$ is in $\mathcal{S}^{3r'}$. Recall that X' was one of three intervals of (X, \prec) defined by **first** and **last**. Combining the permutations on these three intervals and on **{first, last}**, we finally find that $\sigma = (X, <, \prec)$ is a 4-shuffle of permutations in $\mathcal{S}^{3r'}$. By Lemma 2.2, this implies that σ is in \mathcal{S}^s for

$$\begin{aligned} s &= 3r' + 4 \\ &\leq 3r + 12 \lceil \log k \rceil + 28. \end{aligned} \quad \square$$

Recursively applying Theorem 5.2 gives the following bound on the length of factorisations of k -almost mixed free permutations.

Corollary 5.4. *Any k -almost mixed free permutation factorises into a product of at most $4 \cdot 3^k$ separable permutations.*

Proof. Lemma 5.1 and Theorem 5.2 give that k -almost mixed free permutations are product of at most $f(k)$ separable permutations for any function f satisfying

$$\begin{aligned} f(2) &\geq 1 \\ \text{and } f(k) &\geq 3f(k-1) + 12 \lceil \log k \rceil + 28. \end{aligned}$$

This is satisfied for

$$f(k) = 4 \cdot 3^k - 6 \lceil \log k \rceil - 23.$$

Indeed, we have

$$\begin{aligned} f(2) &= 4 \cdot 3^2 - 6 \lceil \log 2 \rceil - 23 = 7. \\ \text{and } f(k) &= 4 \cdot 3^k - 6 \lceil \log k \rceil - 23 \\ &\geq 4 \cdot 3^k - 18(1 + \lceil \log(k-1) \rceil) + 12 \lceil \log k \rceil - 23 \\ &= 3(4 \cdot 3^{k-1} - 6 \lceil \log(k-1) \rceil - 23) - 18 + 12 \lceil \log k \rceil + 2 \cdot 23 \\ &= 3f(k-1) + 12 \lceil \log k \rceil + 28. \end{aligned} \quad \square$$

The last step is to show that pattern-avoiding permutations decompose into almost-mixed free ones. Here, it is worth noting that permutations avoiding a fixed pattern may contain arbitrarily large mixed minors. This is the case of 2-shuffles: if $<$ is the usual order on $[n]$, and \prec orders all odd integers before the even ones, then the adjacency matrix of \prec ordered by $<$ contains an $(n/2)$ -mixed minor. Thus the permutation $\sigma = (<, \prec)$ has arbitrary mixed minors despite avoiding the pattern 321. (In this specific example, the inverse permutation $\sigma^{-1} = (\prec, <)$ is 4-mixed free, but one could combine σ and σ^{-1} to obtain an example where the inverse also has arbitrary mixed minors.)

In general, we need to introduce a third linear order to obtain a k -mixed free structure.

Lemma 5.5. *For any pattern π , there exists $k = 2^{O(|\pi|)}$ such that any permutation σ avoiding π factorises as*

$$\sigma = \sigma_2 \circ \sigma_1^{-1}$$

where σ_1, σ_2 are k -almost mixed free.

Proof. Let $(X, <_1, <_2)$ be a biorder representing σ . If σ avoids π , then by Theorem 2.3, $(X, <_1, <_2)$ has twin-width $t = 2^{O(|\pi|)}$. Thus, by Lemma 2.4, there exists a third linear order $<_3$ on X such that the adjacency matrix of either $<_1$ or $<_2$ ordered by $<_3$ is $(2t + 2)$ -mixed free. It follows that these adjacency matrices ordered by $<_3$ are k -almost mixed free for $k = 4t + 4$. Defining the permutations $\sigma_1 = (X, <_3, <_1)$ and $\sigma_2 = (X, <_3, <_2)$, we obtain the result. \square

Our main result immediately follows from Corollary 5.4 and Lemma 5.5.

Theorem 1.1. *For any pattern π , there exists $k_\pi = 2^{2^{O(|\pi|)}}$ such that every permutation avoiding π is a product of at most k_π separable permutations.*

6 Algorithmic implementation

The proofs of the previous sections are all effective, and yield an FPT algorithm to compute the factorisation given by Theorem 1.1. A relatively naive implementation should run in quadratic time, for any fixed excluded pattern. In this section, we explain how to improve this to get a linear complexity focusing on the few sections of the proof for which this is not straightforward: computing the delayed structured tree in Lemma 3.9, finding a proper colouring of the mixed graph in Lemma 4.5, and finding a mixed-minor free order in Lemma 5.5.

We work in the RAM model. We assume that a permutation σ is given as an array containing $[n]$ in permuted order. If σ is the biorder $(X, <_1, <_2)$, it means that

1. we can sort X according to either $<_1$ or $<_2$ in linear time, and
2. after this preprocessing, we can in constant time compute successors and predecessors for either of the two linear orders $<_1, <_2$ (in addition to being able to compare two given elements).

For example, the linear algorithm would *not* apply if $<_2$ was given as a comparison oracle, or as an array of arbitrary numbers. Note that Guillemot–Marx algorithm for pattern detection [18] makes the same assumption, and we use it in the form of Theorem 2.3.

6.1 Delayed substitutions

We first describe how to compute a delayed structured tree for a permutation $\sigma = (X, <, \prec)$. We see σ as an array containing values from 1 to n , so that $<$ is the left-to-right order on the array, and \prec compares the values. Thus we will prefer to say ‘left’ and ‘right’ when referring to $<$, and ‘greater’ or ‘smaller’ when referring to \prec .

We use as subroutine the classical problem of finding extrema in a given interval of an array.

Theorem 6.1 ([19, 26, 2]). *Given an array A , one can, after a linear-time preprocessing, find in constant time the minimum and maximum (and their positions) in any interval in A .*

Given $x, y \in X$, let $D(x, y)$ be the interval between x and y in the linear order \prec . This is the set of elements that *distinguish* x and y .

Lemma 6.2. *Given a permutation σ , one can compute in time $O(|\sigma|)$ a delayed structured tree for σ subject to the restrictions of Lemma 3.9.*

Proof. For all intents and purposes, the proof of Lemma 3.9 is already an algorithm. The issue is to implement it in linear time. Let us recall the important basic operation in the proof of Lemma 3.9: given $Y \subseteq X$ an interval for $<$, we want to partition it into some intervals $Y_1 < \dots < Y_k$ such that

- no $z \in X \setminus Y$ distinguishes two elements of the same Y_i (i.e. Y_i is a local module), and
- Y_i and Y_{i+1} are distinguished by some $z \in X \setminus Y$ (i.e. the Y_i are maximal).

The delayed structured tree is obtained by repetitively applying the former operation, except when Y is simultaneously an interval for $<$ and for \prec , in which case we do an arbitrary split. To construct the tree in linear time, it is sufficient to implement this splitting operation in time $O(k)$, where k is the number of parts produced by the splitting.

To this end, we define some auxiliary arrays. Let x_1, \dots, x_n be the elements of X ordered left to right. For each pair of consecutive x_i, x_{i+1} , let l_i and r_i be the leftmost and rightmost elements of $D(x_i, x_{i+1})$. Remark that this definition is independent of Y . Let L be the array l_1, \dots, l_{n-1} , and R , the array r_1, \dots, r_{n-1} .

Claim 6.3. *One can compute all l_i and r_i in linear time.*

Proof. This is exactly Theorem 6.1, but on the inverse permutation σ^{-1} . Indeed, l_i and r_i are the minimum and maximum w.r.t. $<$ of the interval between x_i and x_{i+1} for \prec . \diamond

We now want to generalise this claim to any interval of X . Given $Y \subseteq X$ an interval for $<$, let $l(Y)$ and $r(Y)$ be the leftmost and rightmost elements in X that distinguish some pair of elements in Y . Remark that if z distinguishes two elements of Y , then it must also distinguish two consecutive elements of Y . It follows that if $Y = \{x_i, \dots, x_j\}$, then $l(Y)$ is the leftmost element among l_i, \dots, l_{j-1} , and $r(Y)$ is the rightmost element among r_i, \dots, r_{j-1} .

Claim 6.4. *After a linear-time preprocessing, one can answer the following query in constant time: Given an interval Y of $<$,*

1. *find $l(Y)$ and $r(Y)$, and*
2. *find two consecutive $x_i, x_{i+1} \in Y$ distinguished by $l(Y)$, or by $r(Y)$.*

Proof. Apply Theorem 6.1 to the auxiliary arrays L and R . \diamond

We are now ready to solve the problem of splitting an interval Y into local modules. Using Claim 6.4, we compute $l(Y)$ and $r(Y)$. If both are contained in Y , then Y is simultaneously an interval of $<$ and of \prec , and there is nothing to do. Otherwise, if for instance $l(Y) \notin Y$, we also obtain from Claim 6.4 two consecutive x_i, x_{i+1} distinguished by $l(Y)$. We split Y into Y_1, Y_2 containing elements left of x_i , and right of x_{i+1} (inclusive). The above only takes constant time. We then recurse in Y_1 and Y_2 . If for instance $l(Y_1)$ and $r(Y_1)$ are both inside Y (but not necessarily in Y_1), then Y_1 is a local module of Y , and we do not need to split it further. Otherwise, we obtain a new position at which Y_1 can be split, and we continue.

This construction stops once Y has been partitioned into local modules Y_1, \dots, Y_k . The number of steps, and hence the complexity, is proportional to the number of parts k . Finally, each split between consecutive Y_i, Y_{i+1} is explicitly given by some element outside Y distinguishing the rightmost element of Y_i from the leftmost element of Y_{i+1} . Thus the sets Y_i are also maximal as desired. \square

6.2 Mixed graphs

Next, we explain how to colour the mixed graphs of Lemma 4.5 in linear time, at the price of a somewhat worse bound than the one given by Lemma 4.4.

Lemma 6.5. *Given a linear order $(X, <)$, and a partition \mathcal{P} of X whose mixed graph has no K_t or $K_{t,t}$ subgraph, one can compute a $4t^3$ -colouring of the mixed graph of \mathcal{P} in time $O(t^3 |X|)$.*

Proof. We know by Lemma 4.4 that the mixed graph of \mathcal{P} is $(4t^2 - 1)$ -degenerate. Thus, using Lemma 2.5, we could compute a colouring in linear time if we were given the mixed graph. The only problem is to compute the latter. We set $n = |X|$. Recall the partition of the edges of the mixed graph used in Lemma 4.4: E_1 contains the edges AB such that the interval closures of the parts $A, B \in \mathcal{P}$ satisfy $\bar{A} \subseteq \bar{B}$ or $\bar{B} \subseteq \bar{A}$, and the remaining edges are in E_2 .

The graph (\mathcal{P}, E_2) is the overlap graph of the interval closures of parts of \mathcal{P} . By Lemma 2.6, its maximum edge density is at most $2(t - 1)^2$. In particular, the number of edges in E_2 is $O(n \cdot t^2)$, hence by Lemmas 2.5 and 2.10, we can compute (\mathcal{P}, E_2) and find a $4t^2$ -colouring, all in time $O(n \cdot t^2)$.

Let us now focus on the edges of E_1 . We will not try to compute E_1 entirely. Instead, we consider a colour class $\mathcal{P}' \subseteq \mathcal{P}$ of (\mathcal{P}, E_2) —hence a subset of parts satisfying that no pair $A, B \in \mathcal{P}'$ overlaps—and we will find a t -colouring of the subgraph (\mathcal{P}', E_1) .

The interval closures of parts in \mathcal{P}' do not overlap, hence they are a laminar family, or equivalently a rooted forest F , where the ancestor relation is inclusion. We add as leaves of F all the singletons $\{x\}$, $x \in X$, whose parent is the smallest interval \bar{A} containing x with $A \in \mathcal{P}'$, if any. Using a stack and iterating over \mathcal{P}' ordered by left endpoints, one can compute the edges of F in time $O(|X| + |\mathcal{P}'|)$.

Consider now an edge $AB \in E_1$ with $\bar{A} \subseteq \bar{B}$. For this edge to exist, i.e. for A and B to be mixed, there must be some $b \in B \cap \bar{A}$. We say that b *witnesses* the edge AB . Remark that this edge may be witnessed by many different elements of B . Now suppose that B and $b \in B$ are fixed, and consider all the $A \in \mathcal{P}'$ such that there is an edge AB witnessed by b . These are exactly the $A \in \mathcal{P}'$ satisfying $b \in \bar{A} \subseteq \bar{B}$, that is, the nodes on the path from b to B in F . Therefore, to compute all edges of E_1 going down (in F) from B , it suffices to consider all elements $b \in B$, and take the union of all paths from such a b to B in F . It is simple to compute this in time linear in $|B|$ plus the number of edges going down from B . Repeating this process for every $B \in \mathcal{P}'$ allows to compute the restriction of E_1 to \mathcal{P}' in time linear in $|X|$ plus the number of edges. Finally, we know from the proof of Lemma 4.4 that E_1 is $(t - 1)$ -degenerate. Thus the number of edges above is $O(t |\mathcal{P}'|)$, and we can compute a t -colouring of (\mathcal{P}', E_1) in time $O(t |\mathcal{P}'|)$.

Thus, we have computed a $4t^2$ -colouring of (\mathcal{P}, E_2) , and inside each colour class \mathcal{P}' , a t -colouring of (\mathcal{P}', E_1) . Combining these yields a $4t^3$ -colouring of the mixed graph $(\mathcal{P}, E_1 \cup E_2)$. \square

6.3 Mixed-minor free orders

Let us finally comment on the algorithmic aspect of Lemma 5.5, since it is the part of our proof which is not self-contained.

Firstly, Theorem 2.3, which is a key part of Guillemot–Marx algorithm for pattern recognition, can be implemented efficiently: For any fixed pattern π , there is an algorithm which given any permutation σ , either finds an instance of π as pattern of σ , or finds a partition sequence for σ of width $t = 2^{O(|\pi|)}$, in linear time $f(|\pi|) \cdot O(|\sigma|)$. In our case, σ is assumed to avoid a fixed pattern π , hence we obtain a partition sequence for σ in linear time.

We then need to retrieve the order $<$ given by Lemma 2.4 for which the adjacency matrix of σ (as biorder) is $(2t + 2)$ -mixed free. Consider a partition sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ for a binary

structure (V, R_1, \dots, R_k) . An ordering $<$ of V is *compatible* with this partition sequence if for any $i \in [n]$ and part $X \in \mathcal{P}_i$, X is an interval of $(V, <)$. Equivalently, one can represent the partition sequence as a tree T with leaves V , whose internal nodes represent each merge of two parts, and the order $<$ is compatible with the partition sequence if and only if it is compatible with T . There always are orders compatible with any given partition sequence, and these are the orders used in Lemma 2.4. That is, a more precise statement of Lemma 2.4 is (cf. [7, Theorem 5.4]): if $\mathcal{P}_n, \dots, \mathcal{P}_1$ is a partition sequence of width t for (V, R_1, \dots, R_k) , and $<$ is any ordering of V compatible with it, then the adjacency matrix of R_i ordered by $<$ is $(2t+2)$ -mixed free. Given a partition sequence, it is simple to compute a compatible ordering in linear time, by constructing the tree associated with the partition sequence, and traversing the latter.

Combining these two arguments, we obtain a linear algorithm to compute the decomposition given by Lemma 5.5 of pattern-avoiding permutations into k -mixed free permutations.

7 Factoring structures of bounded twin-width

In this final section, we show how to adapt our result to graphs of bounded twin-width in order to describe them by structures with universally bounded twin-width. To this end, we first introduce a representation of products of permutations as ordered graphs.

7.1 Products of permutations as path systems

Let $\sigma_1, \dots, \sigma_m$ be permutations of a linear order $(X, <)$. The *path system representation* of the product $\sigma_m \circ \dots \circ \sigma_1$ is defined as follows. First take $m+1$ copies X_0, \dots, X_m of $(X, <)$, and for each $i \in [0, m]$ and $x \in X$, let x_i denote the copy of x in X_i . Then, for each $i \in [m]$, $x \in X$, add an edge between x_{i-1} and $\sigma_i(x)_i$. See Figure 1 in the introduction for an example. Finally, we fix a *canonical linear order* on the vertices, namely $X_0 < \dots < X_m$, while keeping the order inside X for each copy.

We will prove the following crucial property: If the permutations σ_i have bounded twin-width, then the path system representation has twin-width bounded independently of the length m . The arguments are exactly the same as in [5, Proposition 6.5], we reproduce them for the sake of completeness. Here, we use the standard sparse permutation matrices (and not the dense biorder adjacency matrices of Section 5): the matrix M_σ of a permutation σ has a 1 at position $(i, \sigma(i))$ for each i . Recall that a k -grid in a 0,1-matrix is a k -division in which every zone contains a 1, and that a class \mathcal{C} of permutations avoids some pattern if and only if there is a k such that matrices of permutations in \mathcal{C} have no k -grid.

Lemma 7.1. *Let $\sigma_1, \dots, \sigma_m$ be permutations whose adjacency matrices have no r -grid. Then the adjacency matrix of the path system representation of $\sigma_m \circ \dots \circ \sigma_1$, has no $(3r+2)$ -grid when ordered according to the canonical linear order.*

Proof. Let G be the path system representing $\sigma_m \circ \dots \circ \sigma_1$, with vertices $X_0 \uplus \dots \uplus X_m$ as above. The adjacency matrix of G consists of a double diagonal of blocks corresponding to the adjacency matrices of X_{i-1} against X_i for each i . The latter is exactly the permutation matrix M_{σ_i} , or its transpose. See Figure 8.

Consider now an l -grid in this matrix induced by a division \mathcal{R}, \mathcal{C} .

Claim 7.2. *There exists $i \in [m]$ such that every part of \mathcal{R} intersects $X_{i-1} \cup X_i \cup X_{i+1}$.*

Proof. Let C_1 be the first part of \mathcal{C} , and consider i minimal such that $C_1 \subseteq X_0 \cup \dots \cup X_i$. Then there is no edge between C_1 and X_j for $j > i+1$, hence every $R \in \mathcal{R}$ must intersect $X_0 \cup \dots \cup X_{i+1}$.




Figure 8: The adjacency matrix of the path system representation of $\sigma_m \circ \cdots \circ \sigma_1$.

Symmetrically, if C_l is the last part of \mathcal{C} and j is maximal such that $C_l \subseteq X_j \cup \cdots \cup X_m$, we find that any $R \in \mathcal{R}$ must intersect $X_{j-1} \cup \cdots \cup X_m$. Thus any $R \in \mathcal{R}$ must intersect $X_{j-1} \cup X_j \cup \cdots \cup X_i \cup X_{i+1}$. But necessarily $i \leq j$, hence any part $R \in \mathcal{R}$ must intersect $X_{i-1} \cup X_i \cup X_{i+1}$. \diamond

Claim 7.3. *There exists $j \in [m]$ such that at least $\frac{l-4}{3}$ parts of \mathcal{R} are contained in X_j .*

Proof. By Claim 7.2, any part of \mathcal{R} intersects one of X_{i-1}, X_i, X_{i+1} . Excluding four parts at the borders of the latter, at least $l - 4$ parts of \mathcal{R} are contained in one of X_{i-1}, X_i, X_{i+1} . By the pigeonhole principle, at least $\frac{l-4}{3}$ parts are contained in the same of these three subsets. \diamond

Naturally, Claims 7.2 and 7.3 also hold for columns. Thus we obtain two sets X_i, X_j which contain at least $\frac{l-4}{3}$ parts of \mathcal{R} and \mathcal{C} respectively. This gives a $(\frac{l-4}{3})$ -grid in one of the blocks of the adjacency matrix, which are r -grid free by hypothesis. It follows that $l \leq 3r + 1$, i.e. the adjacency matrix is $(3r + 2)$ -grid free. \square

7.2 Subdivisions of sparse graphs

We first consider sparse graphs. A class \mathcal{C} of graphs is said to have *bounded sparse twin-width* if it has bounded twin-width, and it excludes a biclique $K_{t,t}$ as a subgraph for some t . Bounded sparse twin-width classes were studied in [5, section 7]. They are characterised by excluding grids in adjacency matrices.

Theorem 7.4 ([5, Theorem 2.12]). *Let G be a $K_{t,t}$ -subgraph-free graph of twin-width k . Then there is some linear ordering $<$ of the vertices of G for which the adjacency matrix of G is $f(k, t)$ -grid free, for some function f .*

Theorem 1.4. *There is a universal constant c and a function f such that for any graph G with twin-width k and no $K_{t,t}$ -subgraph, the $f(k, t)$ -subdivision of G has twin-width at most c .*

Proof. Let $G = (V, E)$ be $K_{t,t}$ -subgraph-free with twin-width k , and, applying Theorem 7.4, consider a linear ordering $<$ on V for which the adjacency matrix is r -grid free, with r function of k, t only. Fix an arbitrary orientation of the edges of G , and denote by \vec{E} the set of oriented edges: for each $uv \in E$, exactly one of (u, v) or (v, u) is in \vec{E} . If $\vec{e} = (u, v) \in \vec{E}$, we denote by $s(\vec{e}) = u$ its starting point and by $t(\vec{e}) = v$ its endpoint.

We define two lexicographic orders on \vec{E} : $<_s$ orders first by starting points (ordered by $<$), and then by endpoints, while $<_t$ orders first by endpoints, then by starting points. We consider the permutation $\sigma = (\vec{E}, <_s, <_t)$. Let M_σ be the sparse permutation matrix of σ : the columns are \vec{E} ordered by $<_s$, while the rows are \vec{E} ordered by $<_t$, and for each $\vec{e} \in \vec{E}$, there is a 1 at the intersection of the row and the column corresponding to \vec{e} .

Claim 7.5. *The matrix M_σ is $3r$ -grid free.*

Proof. Let \mathcal{C}, \mathcal{R} be partitions of \vec{E} into intervals of $<_s$ and $<_t$ respectively, inducing a $3r$ -grid, i.e. for each $C \in \mathcal{C}, R \in \mathcal{R}$, there is some $\vec{e} \in C \cap R$. Consider \mathcal{P}_s the partition of \vec{E} which groups edges with the same starting point, i.e. $\mathcal{P}_s = \{s^{-1}(v) : v \in V\}$. This is also a partition of \vec{E} into intervals of $<_s$. Remark that inside each $P \in \mathcal{P}_s$, the orders $<_s$ and $<_t$ coincide, hence the matrix M_σ restricted to the columns in P has no 2-grid. It follows that there cannot be two distinct parts $C, C' \in \mathcal{C}$ such that $C, C' \subseteq P$, as $\mathcal{R}, \{C, C'\}$ would yield a 2-grid using only columns of P . Therefore, it is impossible to have more than three parts of \mathcal{C} intersecting the same part $P \in \mathcal{P}_s$. Naturally, the same applies to \mathcal{R} and the partition by endpoints $\mathcal{P}_t := \{t^{-1}(v) : v \in V\}$.

We then pick every third part of \mathcal{C} and of \mathcal{R} , yielding subsets $\mathcal{C}' \subset \mathcal{C}$ and $\mathcal{R}' \subset \mathcal{R}$ such that each part $P \in \mathcal{P}_s$ (resp. \mathcal{P}_t) intersects at most one part of \mathcal{C}' (resp. \mathcal{R}'). The families of intervals $\mathcal{C}', \mathcal{R}'$ have size r , and induce an r -grid in the matrix M_σ . By projecting on starting points and endpoints respectively, we obtain $\mathcal{C}'' := \{s(C) : C \in \mathcal{C}'\}$ and $\mathcal{R}'' := \{t(R) : R \in \mathcal{R}'\}$, two families of r disjoint intervals of $(V, <)$. For any $s(C) \in \mathcal{C}'', t(R) \in \mathcal{R}''$, there is an edge $\vec{e} \in C \cap R$ hence $s(\vec{e}) \in s(C)$ is adjacent to $t(\vec{e}) \in t(R)$. This proves that $\mathcal{C}'', \mathcal{R}''$ define an r -grid in the adjacency matrix of G , a contradiction. \diamond

Claim 7.5 implies that σ avoids a pattern of size $(3r)^2$, hence by Theorem 1.1 we obtain a factorisation of σ into $m = 2^{2^{O(r^2)}}$ separable permutations, as $\sigma = \sigma_m \circ \dots \circ \sigma_1$. Consider the path system representation of this factorisation, with vertex set $X_0 \uplus \dots \uplus X_m$ as in the previous section. The sets X_0, X_m ordered canonically are in bijection with $(\vec{E}, <_s)$ and $(\vec{E}, <_t)$ respectively, and for each $\vec{e} \in \vec{E}$, there is a path joining the copy of \vec{e} in X_0 to its copy in X_m . We now add the vertices V of G to this structure, and for each edge $\vec{e} \in \vec{E}$, we connect the copy of \vec{e} in X_0 to $s(\vec{e})$, and the copy of \vec{e} in X_m to $t(\vec{e})$. Thus, for each edge $\vec{e} \in \vec{E}$, a path on m vertices joins $s(\vec{e})$ to $t(\vec{e})$, hence the resulting graph is the $(m+1)$ -subdivision of G . Furthermore, for each $v \in V$, the neighbourhood of v inside X_0 , resp. X_m , is an interval of the canonical linear order.

To bound the twin-width of this structure, we consider the following linear order on the vertices: V is ordered by $<$, the path system representation of σ is ordered by the canonical linear order, and we place all vertices of V before the rest. The adjacency matrix of this graph then consists of

1. the adjacency matrix of the path system representation, which is 11-grid free using Lemma 7.1, because separable permutations have no 3-grid, and
2. the adjacency matrix of V against X_0 and X_m , which consists of two increasing sequences (one for X_0 , one for X_m), and can be seen to be 3-grid free.

This implies that the adjacency matrix of the $(m+1)$ -subdivision of G is 15-grid free. \square

7.3 Transducing structures of bounded twin-width

Finally, let us generalise the results of the former section by replacing subdivisions with first-order transductions. We first show that permutations of twin-width t can be encoded into structures of twin-width bounded by a constant c , where the decoding is an FO transduction which depends on t , but c is independent of t .

Lemma 7.6. *There is a universal constant c , and for any pattern π there is an FO interpretation Φ such that for any permutation σ avoiding π , there is a structure S of twin-width at most c such that $\Phi(S) = \sigma$.*

Proof. The structure $S = (V, E, <)$ is the path system representation of the decomposition of σ into separable permutations obtained by Theorem 1.1. It is a binary relational structure, with two relations: the edges E , and the canonical order $<$. By Lemma 7.1, since separable permutations do not contain 3-grids, the adjacency matrix of S has no 11-grid. This gives a universal bound on its twin-width.

Let us show, given the path system representation S of the factorisation $\sigma = \sigma_m \circ \dots \circ \sigma_1$, how to reconstruct the biorder σ using an FO interpretation. As previously, the sets of starting points and endpoints of paths in S are denoted by X_0 and X_m respectively. Remark that the vertices in $X_0 \cup X_m$ are exactly the ones with degree 1 in the edge relation E , and that any other vertex $x \notin X_0 \cup X_m$ satisfies $X_0 < x < X_m$. This allows to test whether a given vertex is in X_0 , or in X_m . We construct a biorder $(X_0, <, \prec)$ isomorphic to σ , where $<$ is the already given canonical order, and \prec is defined as follows. Consider $x_1, x_2 \in X_0$, and denote by $y_i \in X_m$ the other endpoint of the path starting from x_i . Then $x_1 \prec x_2$ if and only if $y_1 < y_2$.

The vertex y_i is the only one in X_m connected by a path of length m to x_i . For fixed m , ‘being connected by a path of length m ’ can be expressed by an FO formula, thus allowing, given x_1, x_2 , to identify y_1, y_2 . It follows that \prec can be defined by an FO formula, hence there is an FO interpretation Φ which maps S to σ . This interpretation only depends on the length m of the factorisation, which itself is function of the excluded pattern π , as desired. \square

Theorem 1.5. *There is a fixed strict class \mathcal{C} of permutations such that for any class \mathcal{D} of binary structures, \mathcal{D} has bounded twin-width iff there exists an FO transduction Φ satisfying $\mathcal{D} \subseteq \Phi(\mathcal{C})$.*

Proof. Firstly, the ‘if’ part of the claim follows from the fact that the (yet to be defined) class \mathcal{C} avoids a pattern, hence has bounded twin-width by Theorem 2.3, and transductions preserve bounded twin-width by Theorem 2.12.

For the other direction, given the class \mathcal{D} , we first apply Theorem 1.6 to obtain a class \mathcal{C}_3 of permutations with bounded twin-width, and Φ_3 an FO transduction such that $\mathcal{D} \subseteq \Phi_3(\mathcal{C}_3)$. Next we apply Lemma 7.6 to obtain a class \mathcal{C}_2 of structures with twin-width bounded by a universal constant c , and Φ_2 an FO transduction such that $\mathcal{C}_3 \subseteq \Phi_2(\mathcal{C}_2)$. Note that \mathcal{C}_2 is independent of \mathcal{D} . Finally, we apply Theorem 1.6 a second time to obtain a class \mathcal{C} of permutations with twin-width bounded by a universal constant c' (function of c only), and Φ_1 an FO transduction such that $\mathcal{C}_2 \subseteq \Phi_1(\mathcal{C})$. We conclude by composing the transductions as $\mathcal{C} \xrightarrow{\Phi_1} \mathcal{C}_2 \xrightarrow{\Phi_2} \mathcal{C}_3 \xrightarrow{\Phi_3} \mathcal{D}$. \square

Acknowledgments

The authors would like to thank Patrice Ossona de Mendez for interesting discussions regarding the subject of Section 7. This work was supported by the ANR projects TWIN-WIDTH (ANR-21-CE48-0014) and Digraphs (ANR-19-CE48-0013).

References

- [1] David Arthur. Fast sorting and pattern-avoiding permutations. In *2007 Proceedings of the Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pages 169–174, 2007.

- [2] Michael A. Bender and Martín Farach-Colton. The lca problem revisited. In Gaston H. Gonnet and Alfredo Viola, editors, *LATIN 2000: Theoretical Informatics*, pages 88–94. Springer Berlin Heidelberg, 2000.
- [3] Pierre Bergé, Édouard Bonnet, Hugues Déprés, and Rémi Watrigant. Approximating highly inapproximable problems on graphs of bounded twin-width. In Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté, editors, *STACS 2023, Hamburg, Germany*, pages 10:1–10:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [4] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set, min dominating set, and coloring. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *ICALP 2021, Glasgow, Scotland (Virtual Conference)*, pages 35:1–35:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [5] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. *Combinatorial Theory*, 2(2), 2022.
- [6] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Torunczyk. Twin-width IV: ordered graphs and matrices. In Stefano Leonardi and Anupam Gupta, editors, *STOC 2022, Rome, Italy*, pages 924–937. ACM, 2022.
- [7] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022.
- [8] Édouard Bonnet, Jaroslav Nešetřil, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. Twin-width and permutations. *CoRR*, abs/2102.06880, 2021.
- [9] Prosenjit Bose, Jonathan F. Buss, and Anna Lubiw. Pattern matching for permutations. *Inf. Process. Lett.*, 65(5):277–283, 1998.
- [10] Romain Bourneuf and Stéphan Thomassé. Bounded twin-width graphs are polynomially χ -bounded. *arXiv preprint arXiv:2303.11231*, 2023.
- [11] Parinya Chalermsook, Mayank Goswami, László Kozma, Kurt Mehlhorn, and Thatchaphol Saranurak. Pattern-avoiding access in binary search trees. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 410–423, 2015.
- [12] Parinya Chalermsook, Seth Pettie, and Sorrachai Yingchareonthawornchai. Sorting pattern-avoiding permutations via 0-1 matrices forbidding product patterns. *CoRR*, abs/2307.02294, 2023.
- [13] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Annals of mathematics*, pages 51–229, 2006.
- [14] Thomas Colcombet. a combinatorial theorem for trees. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland*, pages 901–912. Springer, 2007.
- [15] James Davies. Improved bounds for colouring circle graphs. *Proceedings of the American Mathematical Society*, 150(12):5121–5135, 2022.
- [16] James Davies and Rose McCarty. Circle graphs are quadratically χ -bounded. *Bulletin of the London Mathematical Society*, 53(3):673–679, 2021.

- [17] Jacob Fox. Stanley-Wilf limits are typically exponential. *arXiv preprint arXiv:1310.8378*, 2013.
- [18] 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, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 82–101, 2014.
- [19] Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, 13(2):338–355, 1984.
- [20] Donald E Knuth. The art of computer programming, vol 1: Fundamental. *Algorithms. Reading, MA: Addison-Wesley*, 1968.
- [21] László Kozma and Thatchaphol Saranurak. Smooth heaps and a dual view of self-adjusting data structures. In *STOC 2018, Los Angeles, CA, USA*, page 801–814. ACM, 2018.
- [22] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Comb. Theory, Ser. A*, 107(1):153–160, 2004.
- [23] David W. Matula and Leland L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *J. ACM*, 30(3):417–427, 1983.
- [24] Michał Pilipczuk and Marek Sokółowski. Graphs of bounded twin-width are quasi-polynomially χ -bounded. *Journal of Combinatorial Theory, Series B*, 161:382–406, 2023.
- [25] Neil Robertson and Paul D. Seymour. Graph minors. XVI. excluding a non-planar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003.
- [26] Baruch Schieber and Uzi Vishkin. On finding lowest common ancestors: Simplification and parallelization. *SIAM Journal on Computing*, 17(6):1253–1262, 1988.
- [27] Paul D. Seymour. Decomposition of regular matroids. *J. Comb. Theory, Ser. B*, 28(3):305–359, 1980.