

Unsplittable classes of separable permutations

Michael Albert

Department of Computer Science
University of Otago
New Zealand

michael.albert@cs.otago.ac.nz

Vít Jelínek*

Computer Science Institute
Faculty of Mathematics and Physics
Charles University in Prague
Czech Republic

jelinek@iuuk.mff.cuni.cz

Submitted: May 5, 2016; Accepted: Jun 1, 2016; Published: Jun 10, 2016

Mathematics Subject Classifications: 05A05

Abstract

A permutation class is splittable if it is contained in the merge of two of its proper subclasses. We characterise the unsplittable subclasses of the class of separable permutations both structurally and in terms of their bases.

1 Introduction

In recent years one of the main areas of investigation within the study of pattern-avoiding permutations, or permutation classes has been to develop structural characterisations for some classes in terms of simpler ones. Such a characterisation can have many useful consequences including: exact or approximate enumerations, determination of the type (e.g., rational, algebraic) of the generating function of a class, and algorithms for testing membership in such classes. The survey articles of Vatter [13] and Brignall [7] are perhaps the best introduction to this endeavour.

Among the constructions used to build new classes from old one of the most natural to consider is that of *merging* permutations (or permutation classes). We delay formal definitions to the next section but briefly a permutation π is a merge of two permutations σ and τ if its elements can be partitioned into two sets which are isomorphic (in the sense of relative ordering of corresponding pairs of points) to σ and τ respectively. However, this construction has been relatively ignored in the literature. Merges been used in [4, 5, 8] to establish bounds on the growth rates of certain permutation classes. A restricted version of the construction was used in [2] to show that certain collections of classes have equal growth rates. Those classes which are obtained by merging the classes consisting of

*Supported by project 16-01602Y of the Czech Science Foundation. This work also received financial support from the Neuron Foundation for Support of Science.

monotone increasing or monotone decreasing permutations were all shown to be finitely based in [12] and also played a role in [1] (as evidence for a conjecture concerning the lengths of longest pattern avoiding subsequences in random permutations). Apparently the first paper where merging and its dual, splitting, played a central role was [11].

Call a class *splittable* if it is contained in the merge of two of its proper subclasses. Then a natural question to ask is: what are the essential building blocks with respect to this notion? That is, can we characterise unsplittable classes? The previously mentioned paper [11] began the study of this and related notions and this paper continues it in a somewhat more limited context. Namely, aside from some initial general observations, we work within the class **Sep** of separable permutations. This class of permutations is the natural analog of the graph class consisting of the *cographs* – those graphs constructible from the one vertex graph by complementation and disjoint union. As a tool to help characterise the unsplittable subclasses of **Sep** we investigate an interesting monoid structure on the so-called representable subclasses of **Sep** – this structure is key to showing that the representable classes and the unsplittable ones coincide.

We view this work as representing the first steps of the investigation of splittability within well-behaved classes.

2 Terminology, notation and basic observations

For a non-negative integer n , let $[n] = \{1, 2, \dots, n\}$. A *permutation* (of $[n]$) is a bijective map $\pi: [n] \rightarrow [n]$. The set of all permutations of $[n]$ is denoted \mathcal{S}_n and the set of all permutations is denoted \mathcal{S} . If $\pi \in \mathcal{S}_n$ then we say that the *size* of π is n and write $|\pi| = n$. We frequently represent permutations in *one line notation* as the sequence of their values $\pi_1\pi_2\cdots\pi_n$ (where $\pi_i = \pi(i)$). Since a permutation is also a set of points in the plane, i.e., the set of points (i, π_i) for $i \in [n]$, we can speak of relationships among these points in the horizontal direction thought of as position (using words like before, after, to the left of, or to the right of), or in the vertical direction thought of as value (using words like above, below, greater or lesser). For instance, the set of elements lying above the leftmost point in π consists of the subset (or subsequence in the one line version) of all those points (j, π_j) with $\pi_1 < \pi_j$.

The geometric viewpoint of permutations as sets of points in the plane makes it clear that there are eight symmetries corresponding to the symmetries of the square that can be applied to permutations. In one line notation, these are generated by reversal, complementation (which, for $\pi \in \mathcal{S}_n$ replaces each value i by $n + 1 - i$), and ordinary functional inverse. All of our results respect these symmetries which frequently reduces the number of cases that need to be considered.

Given $\pi \in \mathcal{S}_n$ and a k -element subset, Σ , of the permutation π there is a unique permutation $\sigma \in \mathcal{S}_k$ such that there exists a correspondence between Σ and σ that preserves both the left to right, or positional, ordering, and also the bottom to top, or value, ordering. We say that σ is the *pattern* of the elements Σ . For instance, the pattern of the elements of 256143 occurring in the first, third and sixth positions (i.e., corresponding to the subsequence 263) is the permutation 132. If $\sigma \in \mathcal{S}_k$ is the pattern of some subset of

$$\pi \oplus \sigma = \begin{array}{|c|} \hline \sigma \\ \hline \pi \\ \hline \end{array} \qquad \pi \ominus \sigma = \begin{array}{|c|} \hline \pi \\ \hline \sigma \\ \hline \end{array}$$

Figure 1: The sum and skew sum operations.




Figure 2: 479832156 is an inflation of 2413, specifically 2413[1, 132, 321, 12].

π then we say that σ is *contained in* π and write $\sigma \leq \pi$. This relationship is clearly a partial order on permutations. When we wish to draw attention to a particular subset of π whose pattern is σ we sometimes say that σ is a *subpermutation* of π .

A *permutation class* or *class of permutations*, \mathcal{C} , is a set of permutations hereditarily closed downwards with respect to containment. That is, if $\pi \in \mathcal{C}$ and $\sigma \leq \pi$ then $\sigma \in \mathcal{C}$. Note that our definition permits empty permutations, and the empty permutation belongs to every non-empty permutation class. One way to characterise a permutation class is through a set of *forbidden permutations* or *obstructions*. Given a set F of permutations we can define the permutation class of permutations *avoiding* F :

$$\text{Av}(F) = \{\pi \in \mathcal{S} : \forall \sigma \in F \sigma \not\leq \pi\}.$$

For any permutation class \mathcal{C} we define the *basis* of \mathcal{C} , $\text{Ba}(\mathcal{C})$, to be the set of minimal elements of $\mathcal{S} \setminus \mathcal{C}$ (with respect to \leq). Then $\mathcal{C} = \text{Av}(\text{Ba}(\mathcal{C}))$ and also if F is an antichain in the ordering on permutations, then $F = \text{Ba}(\text{Av}(F))$.

There are various operations that apply to permutations and by extension to permutation classes. The *sum*, $\pi \oplus \sigma$, of two permutations π and σ is the permutation obtained by placing a copy of σ above and to the right of a copy of π . Formally, if $\pi \in \mathcal{S}_n$ and $\sigma \in \mathcal{S}_k$, then $\pi \oplus \sigma \in \mathcal{S}_{n+k}$ and $(\pi \oplus \sigma)(i) = \pi(i)$ for $i \leq n$, while $(\pi \oplus \sigma)(i) = n + \sigma(i - n)$ for $i > n$. The *skew-sum* $\pi \ominus \sigma$ is similarly defined, except σ is placed below and to the right of π . See Figure 1. If a permutation, π , can be written as a sum of two non-empty permutations then we say that π is *sum-decomposable*. If not, then of course we say that π is *sum-indecomposable*. *Skew-decomposable* and *skew-indecomposable* are defined similarly. Clearly, sum and skew-sum are associative (but not commutative) operations on permutations, and every permutation can be uniquely represented as a sum of one or more sum-indecomposable permutations or as a skew-sum of skew-indecomposable permutations. A permutation can be both sum- and skew-indecomposable (for instance 2413) but every sum-decomposable permutation is skew-indecomposable and vice versa.

If \mathcal{C} and \mathcal{D} are permutation classes then $\mathcal{C} \oplus \mathcal{D}$ (resp. $\mathcal{C} \ominus \mathcal{D}$) is defined to be the set of all permutations which are the sum (skew-sum) of a permutation in \mathcal{C} with one in \mathcal{D} . It is routine to verify that $\mathcal{C} \oplus \mathcal{D}$ ($\mathcal{C} \ominus \mathcal{D}$) is also a permutation class.

Sums and skew-sums are a particular case of a more general operation called *inflation*. An *interval* in a permutation is a contiguous set of positions on which the values also form a contiguous set. We say that an interval is *proper* if it is not the whole permutation, and *non-trivial* if it is not a singleton. A permutation whose only proper intervals are singletons is called *simple*. For instance, in the permutation 2756143 the second through fourth elements form an interval (positions from 2 to 4, and values from 5 to 7), and the last two elements also form an interval. On the other hand, the permutation 2753164 is simple.

Given $\sigma \in \mathcal{S}_k$ and permutations $\tau_1, \tau_2, \dots, \tau_k$ the *inflation*, $\pi = \sigma[\tau_1, \tau_2, \dots, \tau_k]$ is a permutation of size $\sum_{i=1}^k |\tau_i|$ obtained from σ by replacing each element σ_i by an interval whose pattern is τ_i where the relationships by position and value between the intervals corresponding to σ_i and σ_j are the same as those between σ_i and σ_j themselves. For instance $2756143 = 2413[1, 312, 1, 21]$, and see also Figure 2. Observe that $\sigma \oplus \tau = 12[\sigma, \tau]$ while $\sigma \ominus \tau = 21[\sigma, \tau]$. For permutation classes \mathcal{C} and \mathcal{D} we define $\mathcal{C}[\mathcal{D}]$ to be the set of all permutations obtained by inflating a permutation from \mathcal{C} with permutations from \mathcal{D} . Again it is clear that $\mathcal{C}[\mathcal{D}]$ is a permutation class.

Some particular permutation classes that we will be using repeatedly are:

- $\mathbb{1}$, which contains only the permutation 1 (and the empty permutation - we omit to mention this in what follows),
- \mathbb{I} , which contains only the increasing permutations,
- \mathbb{D} , which contains only the decreasing permutations,
- $\textcircled{132} = \text{Av}(132)$, together with its symmetries $\textcircled{213} = \text{Av}(213)$, $\textcircled{231} = \text{Av}(231)$ and $\textcircled{312} = \text{Av}(312)$.
- \mathbb{L} , the class of *layered* permutations, which is the closure of \mathbb{D} under \oplus .

Note that $\mathbb{1} = \text{Av}(12, 21)$, $\mathbb{I} = \text{Av}(21)$, $\mathbb{D} = \text{Av}(12)$ and $\mathbb{L} = \text{Av}(231, 312) = \mathbb{I}[\mathbb{D}]$.

A permutation class, \mathcal{C} , is said to be *sum-closed* if for all $\sigma, \tau \in \mathcal{C}$ we have $\sigma \oplus \tau \in \mathcal{C}$. *Skew-closed* is defined similarly.

Observation 1. A class \mathcal{C} is sum-closed if and only if $\mathcal{C} = \mathbb{I}[\mathcal{C}]$ and skew-closed if and only if $\mathcal{C} = \mathbb{D}[\mathcal{C}]$.

The following trivial observation can be surprisingly useful in practice:

Observation 2. Let \mathcal{C} be a permutation class and \mathcal{D} a non-empty subclass of \mathcal{C} . Then $\mathcal{C} = \mathcal{C}[\mathcal{D}]$ if and only if no basis element of \mathcal{C} belongs to $\mathcal{C}[\mathcal{D}]$, and $\mathcal{C} = \mathbb{D}[\mathcal{C}]$ if and only if no basis element of \mathcal{C} belongs to $\mathbb{D}[\mathcal{C}]$.

Some of the consequences of this observation are:

- \mathcal{C} is \oplus -closed, i.e., $\mathcal{C} = \mathbb{I}[\mathcal{C}]$ if and only if all the basis elements of \mathcal{C} are sum-indecomposable.

- $\mathcal{C} = \mathcal{C}[\mathcal{C}]$ if and only if all the basis elements of \mathcal{C} are simple.
- $\mathcal{C} = \mathcal{C}[\mathbb{I}]$ if and only if no basis element of \mathcal{C} has a non-trivial interval in \mathbb{I} (such permutations are sometimes called *plus-irreducible*).

2.1 Separable permutations

Permutation classes whose bases consist only of simple permutations are closed under inflation. We have already encountered three such classes: $\mathbb{1}$, \mathbb{I} and \mathbb{D} . Any other class of this type must contain both 12 and 21, and if we begin from the class $\{\emptyset, 1, 12, 21\}$ and then form its closure under inflation we obtain the class **Sep** of *separable* permutations. **Sep** is perhaps most naturally defined recursively as the smallest non-empty class having the property that if $\sigma, \tau \in \mathbf{Sep}$, then both $\sigma \oplus \tau, \sigma \ominus \tau \in \mathbf{Sep}$. Every non-singleton permutation in **Sep** is either sum- or skew-decomposable and its basis is $\{2413, 3142\}$ (this is easily verified but see Bose et al. [6]).

To every permutation $\pi \in \mathbf{Sep}$ we associate a *decomposition tree* $T(\pi)$. The leaves of $T(\pi)$ correspond to the elements of π and there are two types of internal nodes: \oplus -nodes and \ominus -nodes. The (obvious) correspondence between trees and permutations may be defined recursively: the singleton permutation is represented by a tree consisting of a single leaf. If π is a sum-decomposable permutation of the form $\pi = \pi_1 \oplus \pi_2 \oplus \cdots \oplus \pi_k$ with π_i being sum-indecomposable, then $T(\pi)$ has a root which is a \oplus -node with k children, whose subtrees are $T(\pi_1), \dots, T(\pi_k)$. Skew-decomposable permutations are handled analogously. Note that in a decomposition tree every child of a \oplus -node is a \ominus -node or a leaf, and every child of a \ominus -node is a \oplus -node or a leaf.

The *depth* of a node w in a tree T is the number of edges on the path from w to the root. The depth of T is then the maximum of the depths of the leaves of T . A *bottom leaf* of T is a leaf whose depth is equal to the depth of T . The depth of a separable permutation is the depth of its decomposition tree. The *least common ancestor* of two distinct leaves in a tree T is the unique internal node of maximum depth which is a common ancestor of the two leaves.

We say that a tree T is *slim* if each of its internal nodes has exactly two children and at least one of these children is a leaf. Note that in a slim tree T of depth d , there is exactly one internal node of depth i for any $i \in \{0, \dots, d-1\}$. A *slim* permutation is a separable permutation whose decomposition tree is slim.

2.2 Merging permutations and splitting classes

The last construction we need to introduce for permutations is the *merge*. Given two permutations σ and τ of sizes k and ℓ respectively the *merges* of σ and τ are those permutations of size $k + \ell$ whose elements can be partitioned into a set of k elements whose pattern is σ and a set of ℓ elements whose pattern is τ . For example, all five non-decreasing permutations of size 3 are merges of 12 with 1, and all permutations of size 4 except 1234, 2143, 3412, and 4321 are merges of 12 with 21. We write the set of merges of σ with τ as $\sigma \odot \tau$, and use the same notation for permutation classes. Again it is clear

that the merge of two permutation classes is itself a permutation class. At the level of permutation classes, \odot is an associative and commutative operation. We frequently think of the partition corresponding to a merge as being given by a red-blue colouring of the elements of a permutation. In that context, to prove $\pi \notin \mathcal{C} \odot \mathcal{D}$ requires showing that for every red-blue colouring of π either the red elements do not belong to \mathcal{C} (i.e., contain a basis element of \mathcal{C}) or the blue elements do not belong to \mathcal{D} . No proper permutation class is closed under the merge operation since by merging n copies of 1 we can produce every permutation in \mathcal{S}_n .

We say that a permutation class \mathcal{C} is *splittable*, if there is a sequence $(\mathcal{C}_i)_{i=1}^k$ of proper subclasses of \mathcal{C} such that $\mathcal{C} \subseteq \odot_{i=1}^k \mathcal{C}_i$. If X is a set of permutation classes, we say that \mathcal{C} is X -splittable and write $\mathcal{C} \in \mathbf{Split}(X)$ if there is a finite sequence $(\mathcal{C}_i)_{i=1}^k$ of classes each belonging to X such that \mathcal{C} is $(\mathcal{C}_i)_{i=1}^k$ -splittable. In the case where X is a singleton, say $X = \{\mathcal{D}\}$ then we often simply say that \mathcal{C} is \mathcal{D} -splittable.

A class is called *atomic* if it is not the union of a finite collection (hence two) of its proper subclasses. Every non-atomic class is splittable since if \mathcal{C} is contained in the union of two of its proper subclasses then it is also contained in the merge of those same subclasses as $\mathcal{X} \odot \mathcal{Y}$ always contains $\mathcal{X} \cup \mathcal{Y}$. Equivalently, every unsplittable class must be atomic. The main property of atomic classes that we need is contained in the following observation (due in a more general context to Fraïssé [9]) that atomic classes satisfy the *joint embedding property*.

Lemma 3. *A class \mathcal{C} is atomic if and only if for all $\pi, \sigma \in \mathcal{C}$ there exists $\tau \in \mathcal{C}$ with $\pi, \sigma \leq \tau$.*

Proof. Suppose that there exist $\pi, \sigma \in \mathcal{C}$ having no common extension in \mathcal{C} . Let $\mathcal{D} = \mathcal{C} \cap \text{Av}(\pi)$ and $\mathcal{E} = \mathcal{C} \cap \text{Av}(\sigma)$. Then both \mathcal{D} and \mathcal{E} are proper subclasses of \mathcal{C} but $\mathcal{C} = \mathcal{D} \cup \mathcal{E}$ so \mathcal{C} is not atomic. Conversely, if \mathcal{C} is not atomic, say $\mathcal{C} = \mathcal{D} \cup \mathcal{E}$ where \mathcal{D} and \mathcal{E} are proper subclasses of \mathcal{C} then we choose $\pi \in \mathcal{C} \setminus \mathcal{D}$ and $\sigma \in \mathcal{C} \setminus \mathcal{E}$. Any common extension of π and σ belongs to neither \mathcal{D} nor \mathcal{E} , hence not to \mathcal{C} either. \square

The following lemma provides us with a convenient characterisation of splittability:

Lemma 4. *A class \mathcal{C} is splittable if and only if there exist permutations $\pi, \pi' \in \mathcal{C}$ such that every $\sigma \in \mathcal{C}$ has a red-blue colouring whose red part avoids π and whose blue part avoids π' .*

Proof. Suppose that \mathcal{C} is splittable and choose a finite sequence $(\mathcal{C}_i)_{i=1}^k$ of proper subclasses of \mathcal{C} of minimum possible length such that $\mathcal{C} \subseteq \mathcal{C}_1 \odot \mathcal{C}_2 \odot \cdots \odot \mathcal{C}_k$. Then $k = 2$ since if $k > 2$ we could take $\mathcal{D} = \mathcal{C} \cap (\mathcal{C}_2 \odot \cdots \odot \mathcal{C}_k)$ a proper subclass of \mathcal{C} (by the assumption of minimality) but then $\mathcal{C} \subseteq \mathcal{C}_1 \odot \mathcal{D}$. So $\mathcal{C} \subseteq \mathcal{D} \odot \mathcal{E}$ for some proper subclasses \mathcal{D} and \mathcal{E} of \mathcal{C} . Now we can simply choose $\pi \in \mathcal{C} \setminus \mathcal{D}$ and $\pi' \in \mathcal{C} \setminus \mathcal{E}$.

Conversely, suppose that such π and π' exist and set $\mathcal{D} = \mathcal{C} \cap \text{Av}(\pi)$ and $\mathcal{E} = \mathcal{C} \cap \text{Av}(\pi')$. Then \mathcal{D} and \mathcal{E} are proper subclasses of \mathcal{C} and, by the given condition, $\mathcal{C} \subseteq \mathcal{D} \odot \mathcal{E}$, hence \mathcal{C} is splittable. \square

Now we can provide a supply of unsplittable classes.

Lemma 5. *If $\mathcal{C} = \mathcal{C}[\mathcal{C}]$, i.e., \mathcal{C} is closed under inflation, then \mathcal{C} is unsplittable.*

Proof. Let π and π' be any two permutations in \mathcal{C} and set $\tau = \pi[\pi', \dots, \pi']$. Any red-blue colouring of τ contains either an entirely blue interval of pattern π' or a red point in each such interval, and hence its red elements contain π . So \mathcal{C} cannot be splittable. \square

Lemma 6. *If \mathcal{X} and \mathcal{Y} are unsplittable classes, then so is $\mathcal{X}[\mathcal{Y}]$.*

Proof. Suppose that $\mathcal{X}[\mathcal{Y}]$ is splittable. That means there are permutations $\pi, \pi' \in \mathcal{X}[\mathcal{Y}]$ such that each element of $\mathcal{X}[\mathcal{Y}]$ has a red-blue coloring with no red π and no blue π' .

Since π is in $\mathcal{X}[\mathcal{Y}]$, it can be expressed as $\pi = \sigma[\tau_1, \dots, \tau_k]$ for some $\sigma \in \mathcal{X}$ and $\tau_1, \dots, \tau_k \in \mathcal{Y}$. Since \mathcal{Y} is unsplittable and therefore atomic, there is a permutation $\tau \in \mathcal{Y}$ which contains all the τ_1, \dots, τ_k . So $\pi \leq \sigma[\tau, \tau, \dots, \tau]$. Therefore, any red-blue colouring of an element of $\mathcal{X}[\mathcal{Y}]$ that has no red π also has no red $\sigma[\tau, \tau, \dots, \tau]$ and so we may assume without loss of generality that π is of the form $\sigma[\tau, \tau, \dots, \tau]$. Similarly, we may assume that π' is of the form $\sigma'[\tau', \tau', \dots, \tau']$ for some $\sigma' \in \mathcal{X}$ and $\tau' \in \mathcal{Y}$.

Since \mathcal{X} is unsplittable, it contains a permutation σ^+ whose every red-blue coloring has a red copy of σ or a blue copy of σ' . Likewise, \mathcal{Y} has a permutation τ^+ containing a red τ or a blue τ' in any red-blue coloring.

Consider the permutation $\pi^+ = \sigma^+[\tau^+, \tau^+, \dots, \tau^+] \in \mathcal{X}[\mathcal{Y}]$. We claim that any red-blue coloring of π^+ has a red π or a blue π' . Fix a red-blue coloring c of π^+ . Use the coloring c to define a red-blue coloring d of σ^+ as follows: an element σ_i^+ is red in d if the τ^+ -copy in π^+ obtained by inflating σ_i^+ has a red copy of τ in the coloring c ; similarly, σ_i^+ is blue in d if the corresponding τ^+ -copy has a blue copy of τ' . The coloring d of σ^+ has either a red copy of σ or a blue copy of σ' . In the first case, we find a red copy of $\sigma[\tau, \dots, \tau] = \pi$ in π^+ , while in the other case we find a blue π' . \square

Observation 7. *Let \mathcal{X} be a permutation class, and let \mathcal{Y} be a class which is $\{\mathcal{Y}_1, \dots, \mathcal{Y}_k\}$ -splittable for some $\mathcal{Y}_1, \dots, \mathcal{Y}_k$. Then $\mathcal{X}[\mathcal{Y}]$ is $\{\mathcal{X}[\mathcal{Y}_1], \dots, \mathcal{X}[\mathcal{Y}_k]\}$ -splittable, and $\mathcal{Y}[\mathcal{X}]$ is $\{\mathcal{Y}_1[\mathcal{X}], \dots, \mathcal{Y}_k[\mathcal{X}]\}$ -splittable.*

We need two results from [11], whose proofs we omit.

Lemma 8 ([11, Proposition 3.2]). *For any three permutations α, β, γ ,*

$$\text{Av}(\alpha \oplus \beta \oplus \gamma) \subseteq \text{Av}(\alpha \oplus \beta) \odot \text{Av}(\beta \oplus \gamma).$$

Lemma 9 ([11, Theorem 3.15]). *Let π be a sum-indecomposable permutation. If $\text{Av}(\pi)$ is $\{\text{Av}(\pi_1), \text{Av}(\pi_2), \dots, \text{Av}(\pi_k)\}$ -splittable for a set $\{\pi_1, \dots, \pi_k\}$ of sum-indecomposable permutations, then $\text{Av}(1 \oplus \pi)$ is $\{\text{Av}(1 \oplus \pi_1), \text{Av}(1 \oplus \pi_2), \dots, \text{Av}(1 \oplus \pi_k)\}$ -splittable.*

These results allow us to show that, when avoiding separable permutations, in order to obtain an unsplittable class all the basis elements must be slim.

Lemma 10. *Let π be a separable sum-decomposable permutation of depth d . Then $\text{Av}(\pi)$ is splittable over the set of classes $\text{Av}(\sigma)$ where σ ranges over the slim sum-decomposable subpermutations of π .*

Proof. Proceed by induction on d . The case $d = 0$ is trivial, as the only permutation of depth 0 is the singleton permutation. Suppose π has depth $d \geq 1$. We may write π as

$$\pi = \sigma_1 \oplus \sigma_2 \oplus \cdots \oplus \sigma_m$$

where the σ_i are sum-indecomposable permutations of depth at most $d - 1$. Applying Lemma 8 repeatedly we get

$$\begin{aligned} \text{Av}(\pi) &= \text{Av}(\sigma_1 \oplus \sigma_2 \oplus \cdots \oplus \sigma_m) \\ &\subseteq \text{Av}(\sigma_1 \oplus 1 \oplus \sigma_2 \oplus 1 \oplus \cdots \oplus 1 \oplus \sigma_m) \\ &\in \mathfrak{Split}(\text{Av}(\sigma_1 \oplus 1), \text{Av}(1 \oplus \sigma_2), \text{Av}(\sigma_2 \oplus 1), \dots, \text{Av}(1 \oplus \sigma_m)). \end{aligned}$$

Choose any $\sigma \in \{\sigma_1, \dots, \sigma_m\}$. By induction we have

$$\text{Av}(\sigma) \in \mathfrak{Split}(\text{Av}(\tau_1), \dots, \text{Av}(\tau_\ell)),$$

where τ_1, \dots, τ_ℓ are the slim skew-decomposable subpermutations of σ . By Lemma 9, we then get

$$\begin{aligned} \text{Av}(1 \oplus \sigma) &\in \mathfrak{Split}(\text{Av}(1 \oplus \tau_1), \dots, \text{Av}(1 \oplus \tau_\ell)), \quad \text{and} \\ \text{Av}(\sigma \oplus 1) &\in \mathfrak{Split}(\text{Av}(\tau_1 \oplus 1), \dots, \text{Av}(\tau_\ell \oplus 1)). \end{aligned}$$

The permutations $1 \oplus \tau_i$ are slim subpermutations of π except perhaps when $\sigma = \sigma_1$ in which case we do not need to split $\text{Av}(1 \oplus \sigma)$. An analogous claim holds for $\tau_i \oplus 1$. This completes the proof. \square

Of course an analogous statement to Lemma 10 holds for skew-decomposables as well by symmetry.

3 Representable subclasses of Sep

Our goal is to describe all the unsplitable subclasses of **Sep**. We will show that they are precisely the classes obtained by iterated inflations of **I**, **D**, $\textcircled{213}$, $\textcircled{312}$, $\textcircled{132}$, and $\textcircled{231}$ (recall that $\textcircled{213}$ denotes the class $\text{Av}(213)$ with $\textcircled{312}$, $\textcircled{132}$ and $\textcircled{231}$ its corresponding symmetries).

We will use $\mathcal{X}\mathcal{Y}$ as a shorthand for $\mathcal{X}[\mathcal{Y}]$. Note that $\mathcal{X}[\mathbb{1}] = \mathbb{1}[\mathcal{X}] = \mathcal{X}$ for any class \mathcal{X} . Note also that $\mathcal{X}(\mathcal{Y}\mathcal{Z}) = (\mathcal{X}\mathcal{Y})\mathcal{Z}$ for any three classes \mathcal{X} , \mathcal{Y} and \mathcal{Z} . Thus, permutation classes form a monoid with respect to the inflation operation, and $\mathbb{1}$ is its unit element. By Lemma 6 the unsplitable proper subclasses of **Sep** form a submonoid of this monoid, which will be denoted by \mathcal{U}_S .

Consider now the alphabet $\mathcal{A} = \{\mathbf{I}, \mathbf{D}, \textcircled{213}, \textcircled{312}, \textcircled{132}, \textcircled{231}\}$. Let \mathcal{A}^* be the set of words over \mathcal{A} . Each word $w \in \mathcal{A}^*$ can be seen as a description of an iterated inflation of a sequence of permutation classes, and in particular, each word represents a permutation class from \mathcal{U}_S . Thus, for example, the word **ID** represents the class **L** of layered permutations. We take the empty word $\lambda \in \mathcal{A}^*$ to represent the class $\mathbb{1}$. The subclasses of **Sep** that can be represented in this way will be called *representable* and again by Lemma 6 form a submonoid of \mathcal{U}_S .

which we will denote by \mathcal{R} . Before proceeding to the main result in the next section (that $\mathcal{U}_5 = \mathcal{R}$) we will investigate the structure of and relations between representable classes more closely.

The interpretation of a word $w \in \mathcal{A}^*$ as an element of \mathcal{R} defines a monoid homomorphism from \mathcal{A}^* to \mathcal{R} . We will frequently blur the distinction between \mathcal{A}^* and \mathcal{R} as the appropriate interpretation will generally be clear by context. When the distinction needs to be made explicit, we shall write $\text{cl}(w)$ for the class corresponding to $w \in \mathcal{A}^*$. The first thing to note is that the map $\text{cl}: \mathcal{A}^* \rightarrow \mathcal{R}$ is not an isomorphism. Specifically, using Observations 1 and 2 (and symmetry):

Observation 11. *The following identities hold in the monoid \mathcal{R} :*

$$\begin{array}{lll} \mathbb{I} = \mathbb{I}, & \mathbb{I} \textcircled{512} = \textcircled{512}, & \textcircled{512} \mathbb{D} = \textcircled{512}, \\ \mathbb{D} \mathbb{D} = \mathbb{D}, & \mathbb{I} \textcircled{231} = \textcircled{231}, & \textcircled{231} \mathbb{D} = \textcircled{231}, \\ & \mathbb{D} \textcircled{132} = \textcircled{132}, & \textcircled{132} \mathbb{I} = \textcircled{132}, \\ & \mathbb{D} \textcircled{213} = \textcircled{213}, & \textcircled{213} \mathbb{I} = \textcircled{213}. \end{array}$$

Let \sim denote the congruence on \mathcal{A}^* generated by these identities. Specifically, \sim is the transitive closure of the relation defined between two words in \mathcal{A}^* if the second can be obtained from the first by replacing either two consecutive symbols from the left hand side of one of the identities by the one on the right, or vice versa. Considering only the process of replacing pairs of symbols by single ones we can interpret the identities as a family of rewrite rules (such as $\mathbb{I} \textcircled{231} \rightarrow \textcircled{231}$) whose iterated application will allow us to produce a *reduced word* (one which allows no further applications of such rules) in each \sim -equivalence class. In fact it is easily seen that there is a unique reduced word in each such class since the rewrite rules are *locally confluent*. That is, if we begin with a word $w \in \mathcal{A}^*$ and apply two different rewrite rules to obtain w' and w'' , then either $w' = w''$ or there is a rewrite rule that applies to w' and another one that applies to w'' , both leading to the same word w''' . This is obvious if the rules that produced w' and w'' share no common character in w (since then in some sense the other application is still available). If they do share a common character then sometimes they produce the same word, as for instance $x \textcircled{512} \mathbb{D} \textcircled{132} y \rightarrow x \textcircled{512} \textcircled{132} y$ either by rewriting $\textcircled{512} \mathbb{D} \rightarrow \textcircled{512}$ or by rewriting $\mathbb{D} \textcircled{132} \rightarrow \textcircled{132}$, or the other rule is still available as in

$$x \mathbb{I} \textcircled{231} \mathbb{D} y \rightarrow x \textcircled{231} \mathbb{D} y \rightarrow x \textcircled{231} y \quad \text{or} \quad x \mathbb{I} \textcircled{231} \mathbb{D} y \rightarrow x \mathbb{I} \textcircled{231} y \rightarrow x \textcircled{231} y.$$

We wish to show that two distinct reduced words from \mathcal{A}^* represent distinct permutation classes. In fact, we prove a more general result, showing that set inclusion among permutation classes corresponds to a natural order relation on words.

Observation 12. *The following inclusions hold among permutation classes:*

$$\begin{array}{ll} \mathbb{I} \mathbb{D} \subseteq \textcircled{231}, & \mathbb{D} \mathbb{I} \subseteq \textcircled{132}, \\ \mathbb{I} \mathbb{D} \subseteq \textcircled{512}, & \mathbb{D} \mathbb{I} \subseteq \textcircled{213}. \end{array}$$

Let us say that a word $x \in \mathcal{A}^*$ is *dominated* by a word $y \in \mathcal{A}^*$ if x can be obtained from y by a (possibly empty) sequence of steps where in each step we either

- erase an arbitrary letter,
- replace an occurrence of the letter $\textcircled{231}$ or $\textcircled{312}$ by two consecutive letters ID, or
- replace an occurrence of the letter $\textcircled{132}$ or $\textcircled{213}$ by two consecutive letters DI.

Domination is a partial order on \mathcal{A}^* . This is clear since it is transitive by definition, and it contains no cycles since each single step either reduces the number of characters from the set $\{\textcircled{132}, \textcircled{213}, \textcircled{231}, \textcircled{312}\}$ in a word or leaves that the same but reduces the total length of the word. We will show that the restriction of domination to reduced words corresponds to the inclusion order of representable permutation classes. Before we prove this, we will need several auxiliary claims dealing with the structure of inflations of permutation classes.

Lemma 13. *Let $\mathcal{X} = \text{cl}(x)$ be a permutation class represented by a nonempty reduced word x and let $x_1 \in \mathcal{A}$ be the first symbol of x . Then the class \mathcal{X} is sum-closed if and only if x_1 is one of I, $\textcircled{312}$ or $\textcircled{231}$, i.e., x_1 is itself sum-closed. Symmetrically, \mathcal{X} is skew-closed if and only if x_1 is one of D, $\textcircled{132}$ or $\textcircled{213}$.*

Proof. If x_1 is equal to I, $\textcircled{231}$ or $\textcircled{312}$ then $\text{lx} \sim x$ so $\text{cl}(\text{lx}) = \text{cl}(x) = \mathcal{X}$, i.e., \mathcal{X} is sum-closed. Correspondingly if x_1 is equal to D, $\textcircled{132}$ or $\textcircled{213}$ then \mathcal{X} is skew-closed. The converse follows in each case because no proper subclass of **Sep** other than $\mathbb{1}$ can be both sum- and skew-closed. \square

Lemma 14. *Let $\mathcal{X} = \text{cl}(x)$ be a permutation class represented by a nonempty reduced word x and let $x_1 \in \mathcal{A}$ be the first symbol of x . If $x_1 \neq \text{I}$ then for every permutation $\pi \in \mathcal{X}$ there is a sum-indecomposable permutation $\sigma \in \mathcal{X}$ containing π .*

Proof. Let $x = x_1 x^-$ and suppose first that $x_1 = \textcircled{312}$, and $\pi \in \mathcal{X}$. Write $\pi = \sigma[\rho_1, \rho_2, \dots, \rho_k]$ for some $\sigma \in \textcircled{312}$ and $\rho_1, \rho_2, \dots, \rho_k \in \text{cl}(x^-)$. Then $\pi \ominus 1 = (\sigma \ominus 1)[\rho_1, \rho_2, \dots, \rho_k, 1]$ is in \mathcal{X} since $\sigma \ominus 1 \in \textcircled{312}$. If x_1 is neither I nor $\textcircled{312}$, a similar argument shows that for every $\pi \in \mathcal{X}$ we also have $1 \ominus \pi \in \mathcal{X}$. \square

Lemma 15. *Let \mathcal{X} be a permutation class represented by a nonempty reduced word $x = \text{lx}^-$. Let $\mathcal{X}^- = \text{cl}(x^-)$. A permutation π belongs to \mathcal{X}^- if and only if π is contained in a sum-indecomposable permutation σ belonging to \mathcal{X} .*

Proof. The result is trivial if x^- is the empty word. Otherwise, since x is reduced we know that the first character of x^- is not equal to I.

Suppose that π belongs to \mathcal{X}^- . Then, applying Lemma 14 to the class \mathcal{X}^- , we conclude that π is contained in a sum-indecomposable permutation σ belonging to \mathcal{X}^- (and hence also to \mathcal{X}).

Conversely, suppose π is contained in a sum-indecomposable permutation $\sigma \in \mathcal{X}$. Since \mathcal{X} is the sum closure of \mathcal{X}^- the sum-indecomposable permutations of \mathcal{X} all belong to \mathcal{X}^- , so $\sigma \in \mathcal{X}^-$ and hence $\pi \in \mathcal{X}^-$. \square

Lemma 16. Let \mathcal{X} be a permutation class represented by a nonempty reduced word $x = \textcircled{213}x^-$. Let $\mathcal{X}^- = \text{cl}(x^-)$. A permutation π belongs to \mathcal{X}^- if and only if π is contained in a sum-indecomposable permutation σ such that $\sigma \oplus 1 \in \mathcal{X}$.

Proof. The result is trivial if x^- is the empty word. Otherwise, since x is reduced, the first character of x^- is not equal to 1.

Suppose that π belongs to \mathcal{X}^- . By Lemma 14, there is a sum-indecomposable permutation $\sigma \in \mathcal{X}^-$ containing π , and then clearly $\sigma \oplus 1 = 12[\sigma, 1]$ belongs to \mathcal{X} .

Conversely, suppose that π is contained in a sum-indecomposable permutation σ such that $\sigma \oplus 1$ belongs to \mathcal{X} . We claim that σ , and hence also π , belongs to \mathcal{X}^- . Since $\sigma \oplus 1$ is in $\mathcal{X} = \textcircled{213}\mathcal{X}^-$, it can be obtained by inflating a 213-avoiding permutation ρ by elements of \mathcal{X}^- . As $\sigma \oplus 1$ ends with its maximum element, so must ρ , and as $\rho \in \textcircled{213}$, ρ must be increasing. But σ is sum-indecomposable, so the only possibility is that $\rho = 12$ and $\sigma \oplus 1 = 12[\sigma, 1]$. Hence $\sigma \in \mathcal{X}^-$ as claimed. \square

Theorem 17. Let $x, y \in \mathcal{A}^*$ be two reduced words. Then x is dominated by y if and only if $\text{cl}(x)$ is a subclass of $\text{cl}(y)$.

Proof. Clearly, if x is dominated by y , then $\text{cl}(x)$ is a subclass of $\text{cl}(y)$.

Let us prove the converse. Let $\mathcal{X} = \text{cl}(x)$ and $\mathcal{Y} = \text{cl}(y)$ be the permutation classes represented by x and y , and suppose that $\mathcal{X} \subseteq \mathcal{Y}$. Our goal is to show that x is dominated by y .

Let $x_1x_2 \cdots x_n$ be the sequence of symbols of x , and $y_1y_2 \cdots y_m$ the sequence of symbols of y . We proceed by induction on $m+n$. Note that if x or y is the empty word, the proof is trivial, so let us assume that both x and y are nonempty.

Let x^- be the word $x_2 \cdots x_n$, let y^- be the word $y_2 \cdots y_m$, let \mathcal{X}^- be the class $\text{cl}(x^-)$ and \mathcal{Y}^- the class $\text{cl}(y^-)$. Note that we clearly have $\mathcal{X}^- \subseteq \mathcal{X}$ and $\mathcal{Y}^- \subseteq \mathcal{Y}$.

We now distinguish several cases, based on the symbols x_1 and y_1 . We will restrict ourselves to the situations when y_1 is equal to 1 or to $\textcircled{213}$, as the remaining cases are symmetric to these two.

Case 1: $x_1 = 1$ and $y_1 = 1$. By Lemma 15, $\pi \in \mathcal{X}^-$ if and only if π is contained in some sum-indecomposable element of \mathcal{X} . But then since $\mathcal{X} \subseteq \mathcal{Y}$, π is contained in a sum indecomposable element of \mathcal{Y} and hence, by the same lemma, π belongs to \mathcal{Y}^- . So $\mathcal{X}^- \subseteq \mathcal{Y}^-$. Therefore, by induction, x^- is dominated by y^- and hence x is dominated by y .

Case 2: $x_1 \neq 1$ and $y_1 = 1$. In this case, we claim that $\mathcal{X} \subseteq \mathcal{Y}^-$. To see this, choose $\pi \in \mathcal{X}$. By Lemma 14, there is a sum-indecomposable permutation $\sigma \in \mathcal{X}$ containing π . Consequently, σ belongs to \mathcal{Y} as well, and by Lemma 15, π belongs to \mathcal{Y}^- . By induction, y^- dominates x and therefore y dominates x as well.

Case 3: $x_1 = \textcircled{213}$ and $y_1 = \textcircled{213}$. Applying Lemma 16 to \mathcal{X} and \mathcal{Y} , we conclude that $\mathcal{X}^- \subseteq \mathcal{Y}^-$. Hence x^- is dominated by y^- , and x is dominated by y .

Case 4: $x_1 = \mathbf{D}$, $x_2 \neq 1$ and $y_1 = \textcircled{213}$. If x has length 1, then clearly x is dominated by y , so assume that x^- is nonempty. Since x is reduced, either $x_2 = \textcircled{231}$ or $x_2 = \textcircled{312}$, and in particular \mathcal{X}^- is sum-closed. We wish to show that $\mathcal{X}^- \subseteq \mathcal{Y}^-$. Let $\pi \in \mathcal{X}^-$. By

Lemma 14 there is a sum-indecomposable permutation $\sigma \in \mathcal{X}^-$ that contains π . Since \mathcal{X}^- is sum-closed, $\sigma \oplus 1$ belongs to \mathcal{X}^- and therefore also to \mathcal{Y} , and so by Lemma 16, π belongs to \mathcal{Y}^- . By induction, x^- is dominated by y^- , hence x is dominated by y .

Case 5: $x_1 = \mathbf{D}$, $x_2 = \mathbf{I}$ and $y_1 = \textcircled{213}$. If x has length two we are done, so assume this is not the case. Consider the word $x^{--} = x_3 \cdots x_n$ and the class $\mathcal{X}^{--} = \text{cl}(x^{--})$. Since x is reduced, \mathcal{X}^{--} is skew-closed. Let us prove that $\mathcal{X}^{--} \subseteq \mathcal{Y}^-$. Let $\pi \in \mathcal{X}^{--}$. The permutation $\sigma = \pi \ominus 1$ is in \mathcal{X}^{--} , since \mathcal{X}^{--} is skew-closed. Therefore, $\sigma \oplus 1 \in \mathcal{X}^-$, and hence $\sigma \oplus 1 \in \mathcal{Y}$. By Lemma 16, $\mathcal{X}^{--} \subseteq \mathcal{Y}^-$, hence x^{--} is dominated by y^- , and x is dominated by y .

Case 6: $x_1 = \mathbf{I}$ and $y_1 = \textcircled{213}$. In this case, we may consider the class $\mathcal{X}^+ = \mathbf{D}\mathcal{X}$. Since \mathcal{Y} is skew-closed, \mathcal{X}^+ is a subclass of \mathcal{Y} , and by the argument from the previous case, we obtain that x^- is dominated by y^- , and x by y (the inductive hypothesis still applies since in the previous case it was applied to x^{--} which is two characters shorter than x).

Case 7: $x_1 \in \{\textcircled{23}, \textcircled{312}, \textcircled{132}\}$ and $y_1 = \textcircled{213}$. We show that $\mathcal{X} \subseteq \mathcal{Y}^-$. Let $\pi \in \mathcal{X}$. By Lemma 14, there is a sum-indecomposable permutation $\sigma \in \mathcal{X}$ containing π . Moreover, our assumptions on x_1 guarantee that $\sigma \oplus 1$ is also in \mathcal{X} , and therefore in \mathcal{Y} . By Lemma 16 $\pi \in \mathcal{Y}^-$, hence x is dominated by y^- , and therefore also by y .

This completes the proof of the theorem. \square

Corollary 18. *If \mathcal{X} and \mathcal{Y} are permutation classes represented by reduced words x and y , respectively, then $\mathcal{X} = \mathcal{Y}$ if and only if $x = y$.*

Proof. If $\mathcal{X} = \mathcal{Y}$, then by Theorem 17 x is dominated by y and y is dominated by x , hence $x = y$. The other direction is trivial. \square

3.1 Bases of representable classes

We can determine the basis of a representable class easily from the reduced word that represents it. Though we do not require this result for the remainder of the paper, we record the procedure for doing so here.

For a set F of permutations (generally always a subset of **Sep**) let $\text{Av}_S(F)$ be the set of F -avoiding separable permutations. Equivalently $\text{Av}_S(F) = \text{Av}(F \cup \{2413, 3142\})$. For a subclass \mathcal{X} of separable permutations, let $\text{Bas}_S(\mathcal{X})$ denote the minimal separable permutations not belonging to \mathcal{X} . With this notation we always have $\mathcal{X} = \text{Av}_S(\text{Bas}_S(\mathcal{X}))$. We call the elements of $\text{Bas}_S(\mathcal{X})$ the *minimal separable obstructions* for \mathcal{X} .

We will prove several general lemmas that will allow us to determine the minimal separable obstructions of any representable permutation class. For a set of permutations F , let $F \oplus 1$ denote the set $\{\pi \oplus 1; \pi \in F\}$, with $1 \oplus F$, $1 \ominus F$ and $F \ominus 1$ defined analogously.

Note that each of the following results has a number of symmetric variations that we do not explicitly specify.

Lemma 19. *Let F be a set of skew-decomposable permutations, and let $\mathcal{X} = \text{Av}_S(F)$. Then $\textcircled{213} \mathcal{X} = \text{Av}_S(F \oplus 1)$.*

Proof. Since the elements of F are skew-decomposable, the class \mathcal{X} is sum-closed. No permutation in $F \oplus 1$ can belong to $\textcircled{213}\mathcal{X}$ since then F would belong to \mathcal{X} by Lemma 16. Therefore $\textcircled{213}\mathcal{X} \subseteq \text{Av}_5(F \oplus 1)$. To prove the opposite inclusion let $\pi \in \text{Av}_5(F \oplus 1)$ and proceed by induction on $|\pi|$ to show that $\pi \in \textcircled{213}\mathcal{X}$. For $\pi = 1$ this is clear. Suppose that π is a skew sum of the form $\pi = \pi_1 \ominus \pi_2$. Then by induction both π_1 and π_2 are in $\textcircled{213}\mathcal{X}$, and since $\textcircled{213}\mathcal{X}$ is skew-closed we are done. Finally suppose that $\pi = \pi_1 \oplus \pi_2$. Then π_1 is in $\text{Av}_5(F) = \mathcal{X}$, and π_2 is an inflation of a permutation $\rho \in \textcircled{213}$ by elements of \mathcal{X} by induction. It follows that π is an inflation of $1 \oplus \rho \in \textcircled{213}$ by elements of \mathcal{X} as claimed. \square

Corollary 20. *If \mathcal{X} is a sum-closed permutation class, then $\text{Bas}_5(\textcircled{213}\mathcal{X}) = \text{Bas}_5(\mathcal{X}) \oplus 1$. If \mathcal{X} is an arbitrary permutation class (not necessarily sum-closed), then $\text{Bas}_5(\textcircled{213}\mathcal{X}) = \text{Bas}_5(\textcircled{213}\text{Id}\mathcal{X}) = \text{Bas}_5(\text{Id}\mathcal{X}) \oplus 1$.*

Lemma 21. *Let F be a set of skew-decomposable permutations, and let $\mathcal{X} = \text{Av}_5(F)$. Then $\text{D}\mathcal{X} = \text{Av}_5((F \oplus 1) \cup (1 \oplus F))$.*

Proof. We again easily see that the elements of $\text{D}\mathcal{X}$ must avoid all the permutations in $1 \oplus F$ and $F \oplus 1$.

Let $\pi \in \text{Av}_5((F \oplus 1) \cup (1 \oplus F))$, and write π as a skew sum $\pi = \pi_1 \ominus \pi_2 \ominus \cdots \ominus \pi_k$ for $k \geq 1$, where each π_i is either equal to 1 or sum-decomposable. We claim that each π_i is in \mathcal{X} : indeed, if π_i is sum-decomposable, each of its summands must avoid F , and since \mathcal{X} is sum-closed, we get $\pi_i \in \mathcal{X} = \text{Av}_5(F)$. Thus, we obtain $\pi \in \text{D}\mathcal{X}$ as claimed. \square

Corollary 22. *If \mathcal{X} is a sum-closed permutation class, then $\text{Bas}_5(\text{D}\mathcal{X}) = (\text{Bas}_5(\mathcal{X}) \oplus 1) \cup (1 \oplus \text{Bas}_5(\mathcal{X}))$.*

The results above are sufficient to characterise the minimal obstructions of a representable class, simply by working through the reduced word for a class from right to left (and using the known obstructions for the basic case). For example, consider the class $\mathcal{C} = \textcircled{213}\text{D}\textcircled{231}$.

$$\begin{aligned} \text{Bas}_5(\textcircled{231}) &= \{231\}, \\ \text{Bas}_5(\text{D}\textcircled{231}) &= \{1342, 2314\} && \text{(Lemma 21),} \\ \text{Bas}_5(\textcircled{213}\text{D}\textcircled{231}) &= \text{Bas}_5(\text{Id}\textcircled{231}) \oplus 1 && \text{(Corollary 20),} \\ \text{Bas}_5(\text{Id}\textcircled{231}) &= (\text{Bas}_5(\text{D}\textcircled{231}) \ominus 1) \cup (1 \ominus \text{Bas}_5(\text{D}\textcircled{231})) && \text{(Symmetry of Lemma 21),} \\ &= \{51342, 12453, 52314, 13425\}, \\ \text{Bas}_5(\textcircled{213}\text{D}\textcircled{231}) &= \{513426, 124536, 523146, 134256\}. \end{aligned}$$

Corollary 23. *Let \mathcal{X} be a representable class. Then the elements of $\text{Bas}_5(\mathcal{X})$ are slim separable permutations, all of the same size.*

Corollary 24. *For every slim permutation π , the permutation class $\text{Av}_5(\pi)$ is representable. Moreover, if π has size at least 3, $\text{Av}_5(\pi)$ has a representation all of whose symbols belong to $\{\textcircled{213}, \textcircled{132}, \textcircled{312}, \textcircled{231}\}$. Conversely, for every representable class, $\mathcal{C} = \text{cl}(c)$, defined by a word, c , over the alphabet $\{\textcircled{213}, \textcircled{132}, \textcircled{312}, \textcircled{231}\}$, $\text{Bas}_5(\mathcal{C})$ is a single slim permutation whose length is two more than the number of characters in c .*

4 Unsplittable subclasses of the separable permutations

By Proposition 2.7 of [11] and its symmetries, the classes $\textcircled{132}$, $\textcircled{213}$, $\textcircled{231}$, and $\textcircled{312}$ are all unsplittable. Combined with Lemma 6 this implies that all the representable classes are unsplittable. The aim of this section is to prove the converse - that any non-representable class is splittable. The only proper subclasses of \mathbf{D} and \mathbf{I} are finite and hence splittable so the first non-trivial result of this type is considered in the next lemma. Recall that $\mathbf{L} = \mathbf{ID}$ is the class of layered permutations.

Lemma 25. *Every proper subclass of $\textcircled{231}$ is \mathbf{L} -splittable.*

Proof. Suppose to the contrary that there exist proper subclasses of $\textcircled{231}$ which are not \mathbf{L} -splittable. It is well known (see Atkinson et al. [3, Corollary 2.6]) that the ordering of \mathbf{Sep} under permutation containment is a well quasi-order and, as a consequence the containment relation between subclasses of \mathbf{Sep} is well-founded. So, among the non \mathbf{L} -splittable proper subclasses of $\textcircled{231}$ we may choose one, \mathcal{C} , minimal with respect to containment.

Choose a minimal permutation $\sigma \in \textcircled{231} \setminus \mathcal{C}$ (i.e., a basis element of \mathcal{C} relative to $\textcircled{231}$). If σ is not slim, then \mathcal{C} is splittable, and since all proper subclasses of \mathcal{C} are \mathbf{L} -splittable, so is \mathcal{C} , a contradiction.

Suppose that $\sigma = \rho \oplus 1$, and let $\pi \in \mathcal{C}$. The elements of π less than its final element avoid ρ , while those greater than or equal to its final element are decreasing (since $\pi \in \textcircled{231}$). Hence $\pi \in \mathbf{D} \odot \mathbf{Av}(\rho)$ so \mathcal{C} is $\{\mathbf{D}, \mathbf{Av}(\rho)\}$ -splittable, and once again we have a contradiction. Similarly, if $\sigma = 1 \oplus \rho$ then the elements up to and including π 's minimum are decreasing, and those after avoid ρ , so again we have $\{\mathbf{D}, \mathbf{Av}(\rho)\}$ -splittability and a contradiction.

If $\sigma = \rho \ominus 1$ then, because $\sigma \in \textcircled{231}$ and σ is slim, $\sigma = 21$ and $\mathcal{C} \subseteq \mathbf{I}$, clearly a contradiction.

Finally suppose that $\sigma = 1 \ominus \rho$, where $\rho = 1 \oplus \tau$ or $\rho = \tau \oplus 1$. Let $\pi \in \mathcal{C}$ be sum-indecomposable. Then, since $\pi \in \textcircled{231}$, π begins with its maximum element. The remainder of π avoids ρ so, by the previous argument, it belongs to $\mathbf{D} \odot \mathbf{Av}(\tau)$. Therefore, $\pi \in \mathbf{D} \odot \mathbf{Av}(\tau)$ as well. An arbitrary permutation in \mathcal{C} is the sum of its sum-indecomposable parts and is therefore in $\mathbf{L} \odot \mathbf{I} \mathbf{Av}(\tau)$. But $\mathcal{C} \not\subseteq \mathbf{I} \mathbf{Av}(\tau)$ since $1 \ominus \tau \in \mathcal{C}$. Therefore, by assumption $\mathcal{C} \cap \mathbf{I} \mathbf{Av}(\tau)$ is \mathbf{L} -splittable and hence so is \mathcal{C} our final contradiction. \square

The plan from here is to continue essentially inductively. To proceed we need another technical lemma.

Lemma 26. *Let \mathcal{X} and \mathcal{Y} be permutation classes. Then $(\mathbf{D}\mathcal{X}) \cap (\mathbf{D}\mathcal{Y}) = \mathbf{D}(\mathcal{X} \cap \mathcal{Y})$ and $(\textcircled{213}\mathcal{X}) \cap (\textcircled{213}\mathcal{Y}) = \textcircled{213}(\mathcal{X} \cap \mathcal{Y})$.*

Proof. Clearly, if π is a permutation from $\mathbf{D}(\mathcal{X} \cap \mathcal{Y})$, then π belongs to both $\mathbf{D}\mathcal{X}$ and $\mathbf{D}\mathcal{Y}$. To prove the converse, let $\pi \in (\mathbf{D}\mathcal{X}) \cap (\mathbf{D}\mathcal{Y})$. Write π as $\pi = \pi_1 \ominus \pi_2 \ominus \cdots \ominus \pi_k$, where each π_i is a skew-indecomposable permutation. Then each π_i must belong to $\mathcal{X} \cap \mathcal{Y}$, and so π is in $\mathbf{D}(\mathcal{X} \cap \mathcal{Y})$, as claimed. This proves the first identity.

To prove the second one, we again easily observe that $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$ is a subclass of $(\textcircled{213}\mathcal{X}) \cap (\textcircled{213}\mathcal{Y})$.

We now proceed by induction on the length of a permutation $\pi \in (\textcircled{213}\mathcal{X}) \cap (\textcircled{213}\mathcal{Y})$ to show that it belongs to $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$. The case $\pi = 1$ is trivial. Moreover, if π is neither sum-decomposable nor skew-decomposable, then it belongs to $\mathcal{X} \cap \mathcal{Y}$ and we are done. Suppose that $\pi = \pi_1 \oplus \pi_2$, for some π_1 and π_2 . By induction, both π_1 and π_2 belong to $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$, and so does π itself, since $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$ is skew-closed.

Finally, suppose that π has the form $\pi_1 \oplus \pi_2$ where π_1 is sum-indecomposable. Since π is in $\textcircled{213}\mathcal{X}$, it was obtained by inflating a permutation $\rho \in \textcircled{213}$ by elements of \mathcal{X} . Note that the permutation π_1 was obtained by inflating a single element of ρ , since otherwise ρ would contain the pattern 213. In particular, π_1 belongs to \mathcal{X} , and by an analogous argument, it belongs to \mathcal{Y} as well. By induction, π_2 is in $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$, and in particular it can be obtained by inflating a permutation $\tau \in \textcircled{213}$ by elements of $\mathcal{X} \cap \mathcal{Y}$. Then π itself can be obtained by inflating the permutation $1 \oplus \tau \in \textcircled{213}$ by elements of $\mathcal{X} \cap \mathcal{Y}$, and so π belongs to $\textcircled{213}(\mathcal{X} \cap \mathcal{Y})$, as claimed. \square

The essential part of the remainder of the proof is to show that any proper subclass of a representable class \mathcal{Y} has a splitting into representable proper subclasses of \mathcal{Y} . The key step of the proof is presented below as a lemma.

Lemma 27. *Let \mathcal{Y} be a representable class represented by the reduced word $y = y_1 y_2 \cdots y_n$, and let $\sigma \in \mathcal{Y}$ be a slim permutation. Let \mathcal{X} be the class $\mathcal{Y} \cap \text{Av}_S(\sigma)$. Then there exist representable classes $\mathcal{Y}_1, \dots, \mathcal{Y}_k$ that are proper subclasses of \mathcal{Y} , and such that \mathcal{X} is $\{\mathcal{Y}_1, \dots, \mathcal{Y}_k\}$ -splittable.*

Proof. We proceed by induction, firstly over the possible values of n , i.e., the length of y , and secondly for a given n , over the size of σ . Note that the case of $n \leq 1$ is already handled by Lemma 25 and the preceding remark. Also the case $|\sigma| = 2$ is trivial for any n .

Suppose now that $n \geq 2$ and $|\sigma| \geq 3$. Since σ is slim it has one of the four possible forms $1 \oplus \rho$, $\rho \oplus 1$, $1 \ominus \rho$ or $\rho \ominus 1$ for a slim permutation ρ .

Let $\mathcal{Y}^- = \text{cl}(y_2 \cdots y_n)$. By symmetry, we may assume that y_1 is either D or $\textcircled{213}$. We will treat the two cases separately.

First, suppose that $y_1 = \text{D}$. Since y is reduced, we know that \mathcal{Y}^- is sum-closed. By Corollary 22, we know that the minimal separable obstructions of \mathcal{Y} are precisely the permutations of the form $1 \oplus \tau$ and $\tau \oplus 1$ where τ is a minimal separable obstruction of \mathcal{Y}^- .

Our first goal will be to give a splitting of \mathcal{X} into classes $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m$ where each \mathcal{X}_i is of one of four types:

- (1) the class $\mathbb{1}$,
- (2) the class \mathcal{Y}^- ,
- (3) the class $\mathcal{Y} \cap \text{Av}_S(\rho)$, or
- (4) a class of the form $\text{D}\mathcal{Z}$, where \mathcal{Z} is a proper subclass of \mathcal{Y}^- .

To obtain such splitting, let a permutation $\pi \in \mathcal{X}$ be given and distinguish cases based on the structure of σ .

If σ is of the form $1 \ominus \rho$, we split π into three parts as follows: the first part is the singleton permutation consisting of the leftmost element π_1 of π , the second part consists of all those elements lying above π_1 , and the third part consists of all the elements lying below π_1 .

Observe that the second part forms a permutation from \mathcal{Y}^- – indeed, if the second part contained any minimal separable obstruction τ of \mathcal{Y}^- , then π would contain $1 \oplus \tau$, which is an obstruction of \mathcal{Y} , contradicting the assumption $\pi \in \mathcal{X} \subseteq \mathcal{Y}$.

Similarly, the third part of the splitting avoids the permutation ρ , since π avoids $\sigma = 1 \ominus \rho$. Thus we can merge any permutation in \mathcal{X} from $\mathbb{1}$, \mathcal{Y}^- and $\mathcal{Y} \cap \text{Av}_S(\rho)$ which are classes of types (1), (2) and (3), respectively.

An analogous argument can be made in the case when σ has the form $\rho \ominus 1$. Fix the index i such that $\pi_i = 1$ and split π into three parts: the first part is the singleton π_i , the second part is formed by the elements to the left of π_i , and the third part are the remaining elements. We again see that the second part avoids ρ , and the third part is in \mathcal{Y}^- .

Suppose now that σ has the form $1 \oplus \rho$ or $\rho \oplus 1$. Then σ belongs not just to \mathcal{Y} but actually also to \mathcal{Y}^- . Therefore, \mathcal{X} is a subclass of $\text{D}(\mathcal{Y}^- \cap \text{Av}_S(\sigma))$, and in particular, \mathcal{X} itself is a class of type (4).

We now need to argue that the classes of types (1) through (4) can be split into representable subclasses of \mathcal{Y} . For types (1) and (2), this is obvious, and for the class of type (3), it follows by induction since ρ is shorter than σ .

Consider now a class of type (4), that is, a class equal to $\text{D}\mathcal{Z}$ with \mathcal{Z} a proper subclass of \mathcal{Y}^- . We may assume, without loss of generality, that \mathcal{Z} is equal to $\mathcal{Y}^- \cap \text{Av}_S(\tau)$ for some $\tau \in \mathcal{Y}^-$. Since the word y^- is shorter than y , we know by the inductive hypothesis that \mathcal{Z} is $\{\mathcal{Y}_1^-, \dots, \mathcal{Y}_k^-\}$ -splittable, with $\mathcal{Y}_1^-, \dots, \mathcal{Y}_k^-$ being representable proper subclasses of \mathcal{Y}^- . By Observation 7, $\text{D}\mathcal{Z}$ is $\{\text{D}\mathcal{Y}_1^-, \dots, \text{D}\mathcal{Y}_k^-\}$ -splittable. Moreover, since the word y^- strictly dominates the reduced representation of any of the classes \mathcal{Y}_i^- , we see that y strictly dominates the reduced representation of $\text{D}\mathcal{Y}_i^-$, and therefore all the classes $\text{D}\mathcal{Y}_i^-$ are proper subclasses of \mathcal{Y} by Theorem 17. We thus conclude that any subclass of \mathcal{Y} of type (4) has a splitting into representable proper subclasses of \mathcal{Y} .

This completes the induction step for the case when $y_1 = \text{D}$. Let us deal with the case of $y_1 = \text{213}$. The proof of this case follows the same basic structure as the proof of the case $y_1 = \text{D}$. By Corollary 20, $\text{Bas}(\mathcal{Y}) = \text{Bas}(\text{I}\mathcal{Y}^-) \oplus 1$.

This time, we will first show that the class \mathcal{X} is $\{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m\}$ -splittable where each \mathcal{X}_i is of one of these four types:

- (1) the singleton class $\mathbb{1}$,
- (2) the class $\text{DI}\mathcal{Y}^-$,
- (3) the class $\mathcal{Y} \cap \text{Av}_S(\rho)$, or
- (4) a class of the form $\text{213}\mathcal{Z}$, where \mathcal{Z} is a proper subclass of \mathcal{Y}^- .

To see this, let $\pi \in \mathcal{X}$ and distinguish cases based on the structure of σ .

If σ is of the form $1 \ominus \rho$, we let π_i be the largest element of π , and split π into three parts: the elements to the left of π_i , the element π_i itself, and the elements to the right of π_i . The first part then belongs to $\mathbf{I}\mathcal{Y}^-$, the second to $\mathbf{1}$, and the third to $\mathcal{Y} \cap \text{Av}_S(\rho)$.

Similarly, if σ equals $\rho \ominus 1$, we split π into the elements less than the rightmost element, the rightmost element itself, and the elements greater than the rightmost element. The three parts then belong to $\mathbf{I}\mathcal{Y}^-$, $\mathbf{1}$ and $\mathcal{Y} \cap \text{Av}_S(\rho)$, respectively.

Suppose that σ equals $1 \oplus \rho$. Write π as $\pi = \pi_1 \ominus \pi_2 \ominus \cdots \ominus \pi_\ell$, where all the π_i are skew-indecomposable. Since σ (and hence also ρ) is slim, we know that ρ has one of the two forms $\rho = 1 \ominus \eta$ or $\rho = \eta \ominus 1$, for a skew-indecomposable permutation η . Suppose that ρ equals $\eta \ominus 1$, the other case being analogous. We can then partition each permutation π_i into three parts: the elements smaller than the rightmost one (which belong to $\mathbf{I}\mathcal{Y}^-$), the rightmost element itself, and the elements larger than the rightmost one (which avoid η). Thus, the permutation π itself can be merged from three permutations, belonging respectively to $\mathbf{DI}\mathcal{Y}^-$, \mathbf{D} , and $\mathcal{Y} \cap \text{Av}_S(\eta)$, where we use the fact that $\mathbf{D} \text{Av}_S(\eta) = \text{Av}_S(\eta)$, since η is skew-indecomposable.

Suppose now that $\sigma = \rho \oplus 1$. Then $\text{Av}_S(\sigma) = \textcircled{213} \text{Av}_S(\rho)$ by Lemma 19. Therefore, by Lemma 26, we have the identity

$$\mathcal{Y} \cap \text{Av}_S(\sigma) = \textcircled{213} \mathcal{Y}^- \cap \textcircled{213} \text{Av}_S(\rho) = \textcircled{213} (\mathcal{Y}^- \cap \text{Av}_S(\rho)).$$

Let us argue that $\mathcal{Y}^- \cap \text{Av}_S(\rho)$ is a proper subclass of \mathcal{Y}^- , i.e., that ρ is in \mathcal{Y}^- . If ρ is not in \mathcal{Y}^- , then ρ cannot be in $\mathbf{I}\mathcal{Y}^-$ either, since ρ is skew-decomposable. Thus ρ contains a minimal separable obstruction τ of $\mathbf{I}\mathcal{Y}^-$, and therefore σ contains the minimal separable obstruction $\tau \oplus 1$ of \mathcal{Y} , which is impossible, since σ is in \mathcal{Y} .

What remains is to argue that each class of one of the four types above has a splitting into proper representable subclasses of \mathcal{Y} . This is clear for classes of types (1) and (2), and it follows from the inductive hypothesis for type (3). For type (4), the argument is analogous to the corresponding part of the argument in the case of $y_1 = \mathbf{D}$, solved previously. \square

We are finally ready to prove our main result, which describes all the unsplittable proper subclasses of **Sep**.

Theorem 28. *A proper subclass of **Sep** is unsplittable if and only if it is representable.*

Proof. We already know that every representable class is unsplittable. To prove the converse, let us first observe that every unsplittable proper subclass \mathcal{X} of **Sep** is a subclass of a representable class \mathcal{Y} . To see this, choose π to be a minimal separable obstruction of \mathcal{X} . We know that π is slim, otherwise \mathcal{X} would be splittable by Lemma 10. Thus, \mathcal{X} is a subclass of the class $\mathcal{Y} = \text{Av}_S(\pi)$, which is representable by Corollary 24.

From Theorem 17 the ordering of representable classes by containment is a strict refinement of the subword ordering on a subset of the free monoid \mathcal{A}^* . Higman's Lemma [10] implies that the subword ordering on \mathcal{A}^* is a well quasi-order and so we may deduce that

the ordering of the representable classes by inclusion is also a well quasi-order. In particular, for any unsplittable class \mathcal{X} there is a (not necessarily unique) minimal representable superclass.

Suppose now that \mathcal{X} is an unsplittable class and assume, for contradiction, that \mathcal{X} is not representable. Let \mathcal{Y} be a minimal representable superclass of \mathcal{X} and let σ be a minimal separable obstruction of \mathcal{X} that belongs to \mathcal{Y} . Note that σ is a slim permutation, otherwise \mathcal{X} would be splittable by Lemma 10.

By Lemma 27, the class $\mathcal{Y} \cap \text{Av}_5(\sigma)$, and therefore also its subclass \mathcal{X} , is $\{\mathcal{Y}_1, \dots, \mathcal{Y}_k\}$ -splittable, for some representable proper subclasses $\mathcal{Y}_1, \dots, \mathcal{Y}_k$ of \mathcal{Y} . Since \mathcal{Y} was a minimal representable superclass of \mathcal{X} , none of the \mathcal{Y}_i contains \mathcal{X} . Therefore \mathcal{X} is splittable, contrary to our assumption. \square

Theorem 29. *Let \mathcal{X} be a proper subclass of **Sep** and let A be the set of representable subclasses of \mathcal{X} . Then \mathcal{X} is A -splittable.*

Proof. Let \mathcal{X} be a proper subclass of **Sep** and let $\text{Bas}(\mathcal{X})$ be its minimal separable obstructions. Then $\text{Bas}(\mathcal{X})$ is finite since, as we noted earlier, **Sep** is well quasi-ordered by permutation containment (see [3, Corollary 2.6]). We may assume without loss of generality that all the elements of $\text{Bas}(\mathcal{X})$ are slim, for otherwise we could use Lemma 10 to show that \mathcal{X} can be split into its subclasses whose minimal separable obstructions are slim.

As argued in the proof of Theorem 28 we may choose a minimal representable superclass, \mathcal{Y} , of \mathcal{X} . By Theorem 17, \mathcal{Y} has only finitely many proper representable subclasses. We will proceed by induction on the number of representable subclasses of \mathcal{Y} . If \mathcal{Y} has no representable subclass beyond itself then $\mathcal{Y} = \mathbb{1}$ and the claim of the theorem is trivial.

Suppose $\mathcal{Y} \neq \mathbb{1}$. If $\mathcal{X} = \mathcal{Y}$, the claim is again trivial. If \mathcal{X} is a proper subclass of \mathcal{Y} , then Lemma 27 shows that \mathcal{X} is $\{\mathcal{Y}_1, \dots, \mathcal{Y}_k\}$ -splittable, where each \mathcal{Y}_i is a representable proper subclass of \mathcal{Y} . In particular, \mathcal{X} is also $\{\mathcal{X} \cap \mathcal{Y}_1, \dots, \mathcal{X} \cap \mathcal{Y}_k\}$ -splittable. Each class $\mathcal{X} \cap \mathcal{Y}_i$ is contained in the representable class \mathcal{Y}_i which has fewer representable subclasses than \mathcal{Y} . Hence we may apply induction to show that each $\mathcal{X} \cap \mathcal{Y}_i$ is splittable over the set of its representable subclasses. It follows that \mathcal{X} is splittable over its representable subclasses as well. \square

As an easy consequence, we obtain the following result.

Theorem 30. *Let A be a set of representable permutation classes. Let B be the set of minimal representable classes that are not contained in any class in A . Then B is finite and a proper subclass \mathcal{X} of **Sep** is A -splittable if and only if \mathcal{X} does not contain any element of B as a subclass.*

Proof. Since B is an antichain in the inclusion order of representable classes, B is finite. Let \mathcal{X} be a proper subclass of **Sep**. If \mathcal{X} is A -splittable, then every representable subclass of \mathcal{X} is A -splittable as well and, since representable classes are unsplittable, every representable subclass of \mathcal{X} is contained in an element of A . In particular, no subclass of \mathcal{X} is in B .

Conversely, if \mathcal{X} has no element of B as subclass, then all the representable subclasses of \mathcal{X} are subclasses of elements of A , and so \mathcal{X} is A -splittable by Theorem 29. \square

5 Concluding remarks

In the preceding sections we have answered the question: *what characterises the unsplittable subclasses of Sep?* This invites the question: *given a collection, A , of unsplittable subclasses of Sep, can one characterise the A -splittable classes?* The A -splittable subclasses of Sep are characterised by Theorem 30. It is a natural problem to extend such classification to more general classes.

For certain specific cases of A , the characterisation in Theorem 30 remains valid even without the assumption that \mathcal{X} is a subclass of Sep. For instance, it can be shown that an arbitrary class \mathcal{X} is 1-splittable if and only if it does not contain I and D as subclasses, it is I-splittable if and only if it does not contain D, it is {I, D}-splittable if and only if it does not contain ID and DI (see Vatter [14]), it is $\{\textcircled{231}, \textcircled{312}\}$ -splittable if it does not contain DI, etc.

Unfortunately, the assumption that \mathcal{X} is a subclass of Sep in Theorem 30 cannot in general be omitted. For instance, it can be shown that the class $\mathcal{X} = \text{Av}(1423)$, which is not a subclass of Sep, cannot be split over its representable subclasses. In fact, it can be shown that this class is not even Sep-splittable.

To extend our results from separable permutations to more general cases, we first need to get a better understanding of the unsplittable classes, especially those that cannot be obtained by inflations of smaller unsplittable classes. Our results show that the inflation monoid is closely related to splittability. Therefore, it seems natural to investigate splittability in substitution-closed classes particularly those that contain only finitely many simple permutations. Since the shortest non-monotone simple permutations are 2413 and 3142 this motivates the following question.

Question. *What are the unsplittable subclasses of the the closure under subpermutations and inflation of $\{2413\}$ (or of $\{2413, 3142\}$)?*

We believe that a better understanding of the consequences of splittability in permutation classes, and a more general understanding of the characteristics of unsplittable permutation classes could be a key element in the program of systematically investigating permutation classes by structural means.

Acknowledgement

We thank Vincent Vatter who provided the tikz code for Figures 1 and 2. We also thank a referee who made a number of useful suggestions for improvements of the presentation. The second author gratefully acknowledges the hospitality of the Computer Science Department of the University of Otago, where most of the research presented in this paper was conducted.

References

- [1] Michael H. Albert. On the length of the longest subsequence avoiding an arbitrary pattern in a random permutation. *Random Structures Algorithms*, 31(2):227–238, 2007.
- [2] Michael H. Albert, M. D. Atkinson, Robert Brignall, Nik Ruškuc, Rebecca Smith, and Julian West. Growth rates for subclasses of $\text{Av}(321)$. *Electron. J. Combin.*, 17:#R141, 16 pp., 2010.
- [3] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Partially well-ordered closed sets of permutations. *Order*, 19(2):101–113, 2002.
- [4] M. Bóna. On the best upper bound for permutations avoiding a pattern of a given length. [arXiv:1209.2404](https://arxiv.org/abs/1209.2404), September 2012.
- [5] Miklós Bóna. A new upper bound for 1324-avoiding permutations. *Combin. Probab. Comput.*, 23:717–724, 2014.
- [6] P. Bose, J. F. Buss, and A. Lubiw. Pattern matching for permutations. *Information Processing Letters*, 65(5):277–283, 1998.
- [7] Robert Brignall. A survey of simple permutations. In Steve Linton, Nik Ruškuc, and Vincent Vatter, editors, *Permutation Patterns*, volume 376 of *London Mathematical Society Lecture Note Series*, pages 41–65. Cambridge University Press, Cambridge, England, 2010.
- [8] Anders Claesson, Vít Jelínek, and Einar Steingrímsson. Upper bounds for the Stanley–Wilf limit of 1324 and other layered patterns. *J. Combin. Theory Ser. A*, 119:1680–1691, 2012.
- [9] Roland Fraïssé. *Theory of relations*, volume 145 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, Netherlands, 2000.
- [10] Graham Higman. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc. (3)*, 2:326–336, 1952.
- [11] Vít Jelínek and Pavel Valtr. Splittings and Ramsey properties of permutation classes. *Adv. in Appl. Math.*, 63:41–67, 2015.
- [12] André E. Kézdy, Hunter S. Snevily, and Chi Wang. Partitioning permutations into increasing and decreasing subsequences. *J. Combin. Theory Ser. A*, 73(2):353–359, 1996.
- [13] Vincent Vatter. Permutation classes. In *Handbook of enumerative combinatorics*, Discrete Math. Appl. (Boca Raton), pages 753–833. CRC Press, Boca Raton, FL, 2015.
- [14] Vincent Vatter. An Erdős–Hajnal analogue for permutation classes. *Discrete Math. Theor. Comput. Sci.*, 18(2):5 pp., 2016.