

Two examples of Wilf-collapse

Michael Albert¹Vít Jelínek²Michal Opler²¹ University of Otago, Dunedin, New Zealand² Charles University, Prague, Czechiareceived 18th Dec. 2019, revised 20th May 2021, accepted 21st May 2021.

Two permutation classes, the X-class and subpermutations of the increasing oscillation are shown to exhibit an exponential Wilf-collapse. This means that the number of distinct enumerations of principal subclasses of each of these classes grows much more slowly than the class itself whereas *a priori*, based only on symmetries of the class, there is no reason to expect this. The underlying cause of the collapse in both cases is the ability to apply some form of local symmetry which, combined with a greedy algorithm for detecting patterns in these classes, yields a Wilf-collapse.

Keywords: permutations, Wilf-equivalence, generating functions, growth rates

1 Introduction

The coincidences that arise when two unrelated, or peripherally related collections of combinatorial structures have the same generating function have always attracted interest. In such cases one frequently seeks an explanation of this coincidence which might be an explicit size-preserving bijection between the classes, or parallel structural descriptions of the classes which frequently amount to an implicit description of such a bijection. In this paper we are concerned with certain families of such coincidences which are called *Wilf-equivalence*. These arise when the two classes in question both occur as subsets of some universal class carrying a partial order (always interpreted as containment of structures, and presumed to satisfy the obvious conditions that such an order should have). Specifically if the universe is \mathcal{U} and it carries a partial order \leq , then, to each $\alpha \in \mathcal{U}$ we associate the subclass of structures *avoiding* α , i.e.,

$$\text{Av}(\alpha) = \{\tau \in \mathcal{U} : \alpha \not\leq \tau\}.$$

We say that α and β are *Wilf-equivalent* (in \mathcal{U}) if $\text{Av}(\alpha)$ and $\text{Av}(\beta)$ have the same generating function.

More particularly, we are interested in situations where Wilf-equivalence is the rule, rather than the exception. For two structures to be Wilf-equivalent they must have the same size (since, for any α the lowest order term of the generating function for $\text{Av}(\alpha)$ that differs from that of \mathcal{U} is the term of degree $|\alpha|$ (where $|\alpha|$ denotes the size of α). If, as $n \rightarrow \infty$ the number, w_n , the number of equivalence classes of the Wilf-equivalence relation on structures of size n in \mathcal{U} is small relative to the total number, u_n of structures of size n in \mathcal{U} then we say that \mathcal{U} exhibits a *Wilf-collapse*. More specifically, a Wilf-collapse occurs if $w_n = o(u_n)$. If, for some constant $c < 1$, $w_n = o(c^n u_n)$ then we say that \mathcal{U} exhibits an *exponential Wilf-collapse*.

At the risk of some confusion we will refer to the equivalence classes of the Wilf-equivalence relation in a universe \mathcal{U} as *Wilf-equivalence classes*.

Here we will demonstrate an exponential Wilf-collapse in two permutation classes, \mathcal{X} and \mathcal{SIO} (defined below). Our results complement those of another paper by the same authors, [3], which includes some sufficient conditions to establish exponential Wilf-collapse. The independent interest of these two particular examples is that neither \mathcal{X} nor \mathcal{SIO} satisfies those conditions. However, both admit a natural representation of structures as words over an infinite alphabet which combines with a greedy algorithm to decide the containment relation to allow “local” applications of symmetries in creating Wilf-equivalent structures. It is in exploring this central thesis with new examples that we believe the value of this paper lies.

1.1 Permutation classes

In this subsection we give the briefest possible introduction to the basics of permutation classes as required for the remainder of the paper. For a much more comprehensive introduction we recommend [11].

Permutations of size n are bijections from $[n] = \{1, 2, \dots, n\}$ to itself, usually written in one line notation as $\pi = \pi_1\pi_2 \dots \pi_n$. A permutation π of size k is *contained in* another permutation τ if there is a subsequence of k elements in τ whose relative ordering is the same as that of π . For instance 321 is contained in 3714526 because the subsequence 742 is in decreasing order. In this situation we write $\pi \leq \tau$ and also say that π is *involved in* τ . If $\pi \not\leq \tau$ then we say that τ *avoids* π .

The *symmetries* of the universe \mathcal{S} of all permutations are the order-automorphisms of \mathcal{S} . Concretely, there are eight of them generated by:

reverse which simply reverses the one-line representation of each permutation,

complement which, in a permutation of size n replaces each value a in its one line notation by $n + 1 - a$,
and

inverse the normal inverse operation on bijections.

These are most naturally understood as operating on the set of points (i, π_i) representing a permutation by reflection through a vertical, horizontal, or diagonal axis (and compositions of such reflections of course).

The *sum*, $\sigma = \alpha \oplus \beta$, of two permutations α and β is the permutation of size $|\alpha| + |\beta|$ where:

$$\sigma(i) = \begin{cases} \alpha(i) & \text{for } 1 \leq i \leq |\alpha| \\ |\alpha| + \beta(i - |\alpha|) & \text{for } |\alpha| < i \leq |\alpha| + |\beta| \end{cases}$$

That is, “place α to the left of and below β ”. If we change “below” to “above”, then we will have defined the *skew-sum*, $\alpha \ominus \beta$. A permutation is *sum-* (resp. *skew-*) *indecomposable* if it cannot be written as a sum (skew-sum) of two permutations both of smaller size.

A *permutation class* (or just *class*), \mathcal{C} is a collection of permutations downward-closed under containment, i.e., if $\tau \in \mathcal{C}$ and $\pi \leq \tau$, then $\pi \in \mathcal{C}$. When we speak about Wilf-collapse it will always be within a universe that is some permutation class (specifically, the classes \mathcal{X} and \mathcal{SIO} introduced below).

A class is *sum-closed* if $\alpha, \beta \in \mathcal{C}$ implies that $\alpha \oplus \beta \in \mathcal{C}$. As above we use $\text{Av}(\pi)$ to denote the set of permutations avoiding π (possibly in a given class when the context is clear), and also $\text{Inv}(\pi)$ to denote the set of permutations involving π .

We repeat now the definitions of Wilf-equivalence and Wilf-collapse in the context of permutation classes in particular. Throughout, we consider a particular class \mathcal{C} to be given as our “universe” and all permutations mentioned are assumed to belong to \mathcal{C} and all concepts are relative to \mathcal{C} . For instance, $\text{Av}(\alpha) = \mathcal{C} \cap \{\beta : \alpha \not\prec \beta\}$.

Two permutations $\alpha, \beta \in \mathcal{C}$ are *Wilf-equivalent* if the generating functions of $\text{Av}(\alpha)$ and $\text{Av}(\beta)$ are the same. Equivalent conditions include:

- the generating functions of $\text{Inv}(\alpha)$ and $\text{Inv}(\beta)$ are the same,
- there is a size preserving bijection between $\text{Av}(\alpha)$ and $\text{Av}(\beta)$,
- there is a size preserving bijection between $\text{Inv}(\alpha)$ and $\text{Inv}(\beta)$, and
- there is a size preserving bijection from \mathcal{C} to itself that sends $\text{Av}(\alpha)$ to $\text{Av}(\beta)$.

In different contexts we have observed any one of these equivalent conditions to be the most natural one to prove. In \mathcal{X} it will turn out to be easiest to compute (and compare) generating functions for $\text{Inv}(\alpha)$ and $\text{Inv}(\beta)$, while in \mathcal{SIO} we demonstrate the existence of implicitly defined bijections.

Let c_n denote the number of permutations in \mathcal{C} of size n and w_n the number of Wilf-equivalence classes among those. Then we say that \mathcal{C} exhibits a *Wilf-collapse* if $w_n = o(c_n)$, and an *exponential Wilf-collapse* if $w_n = o(r^n c_n)$ for some $0 < r < 1$.

In [3] we demonstrated some sufficient conditions for Wilf-collapse in sum-closed classes. One of these, Theorem 4.1, was:

Theorem. *Any supercritical sum-closed class, \mathcal{C} , that contains an incompatible pair has an exponential Wilf collapse, unless \mathcal{C} is the class of increasing permutations.*

A few words in the statement of that theorem require further explanation: \mathcal{C} is supercritical if the radius of convergence of the generating function of its sum-indecomposables is strictly greater than the radius of convergence of its generating function, and an incompatible pair is a pair of sum-indecomposable permutations a and b of \mathcal{C} such that neither $a \oplus b$ nor $b \oplus a$ is a subpermutation of any sum-indecomposable element of \mathcal{C} .

The examples we consider in this paper have an exponential Wilf collapse. However, the first one, \mathcal{X} , is not sum-closed and the second, \mathcal{SIO} is sum-closed and super-critical but contains no incompatible pair.

2 The \mathcal{X} -class

The class $\mathcal{X} = \text{Av}(2143, 2413, 3142, 3412)$ can be defined recursively as 1, together with all those permutations that start or finish with their maximum or minimum and that, when such an element is deleted, belong to \mathcal{X} . Alternatively, we can view it as the closure of $\{1\}$ under the four operations:

$$\pi \mapsto 1 \oplus \pi, \pi \mapsto 1 \ominus \pi, \pi \mapsto \pi \oplus 1, \pi \mapsto \pi \ominus 1,$$

or as the set of permutations that can be drawn on the diagonals of an axis-aligned square. The class \mathcal{X} is closed under all the symmetries of \mathcal{S} . We will not explicitly state symmetrical variations of results below, but will feel free to use them as required.

Plainly, \mathcal{X} is not sum- or skew-sum-closed and so the results of [3] cannot provide a Wilf-collapse. But, as we shall see there is a natural way of representing the elements of \mathcal{X} as words over an (infinite)




Fig. 1: A crossed and negative element of \mathcal{X} whose word representation is $(-3, -2)(2, 0)(-1, -1)(4, 3)(-3, -2)(3)$

alphabet \mathcal{A} and also a greedy method for determining when $\pi \leq \tau$ for $\pi, \tau \in \mathcal{X}$. Together, these provide a factorisation theorem for the generating functions of $\mathcal{X} \cap \text{Inv}(\pi)$ in terms of the letters that occur in the word representing π . This is more than sufficient for an exponential Wilf-collapse.

The type of $\pi \in \mathcal{X}$ is either *monotone* (if it is monotone increasing or decreasing) or *crossed* (otherwise). If π has size at least two and begins with its minimum or ends with its maximum we say it is *positive* (or *has positive sign*) while if it begins with its maximum or ends with its minimum it is *negative*. Note that, according to these definitions, the permutation 1 has no sign but every other permutation in \mathcal{X} has a uniquely defined sign. For a non-negative integer a we also use a to denote the monotone increasing permutation of size a , and $-a$ to denote the monotone decreasing permutation of size a .

If $\pi \in \mathcal{X}$ of size n is crossed and positive then, written in one-line notation, π either begins with 1 or ends with n . In fact, for some $a, b \geq 0$ with $0 < a + b < n - 1$, π begins with $12 \dots a$ but not $12 \dots (a + 1)$ and ends with $(n - b + 1)(n - b + 2) \dots n$ but not $(n - b)(n - b + 1) \dots n$. That is, for this unique pair (a, b) and some negative $\theta \in \mathcal{X}$, $\pi = a \oplus \theta \oplus b$. The conclusion of this discussion is:

Observation 2.1. If $\pi \in \mathcal{X}$ is crossed and positive then there exists a unique negative $\theta \in \mathcal{X}$ and a pair of non-negative integers a, b with $a + b > 0$ such that $\pi = a \oplus \theta \oplus b$.

Let \mathcal{A}^+ be the set of all ordered pairs of non-negative integers except $(0, 0)$ and let \mathcal{A}^- be the set of all ordered pairs of non-positive integers except $(0, 0)$. In either case, define the size of a pair, (a, b) , to be $|a| + |b|$. Finally let \mathcal{M} be the set of all integers of absolute value at least 2 and define the size of an element of \mathcal{M} to be its absolute value. The *sign* of an element of \mathcal{A}^+ , \mathcal{A}^- , or \mathcal{M} is defined in the obvious way. Define the language \mathcal{L} to consist of:

- The symbol 1, of size 1, and
- all words $w_1 w_2 \dots w_n$ for $n \geq 1$ satisfying the following conditions:
 - $w_n \in \mathcal{M}$,
 - $w_1, w_2, \dots, w_{n-1} \in \mathcal{A}^+ \cup \mathcal{A}^-$, and
 - the signs of w_1, w_2, \dots, w_n alternate.

The size of a word covered by the second case is defined to be the sum of the sizes of its letters.

Define $\Phi : \mathcal{L} \rightarrow \mathcal{X}$ recursively as follows:

$$\begin{aligned} \Phi(1) &= 1 \\ \Phi(a) &= a && \text{if } a \in \mathcal{M} \\ \Phi((a, b)v) &= a \oplus \Phi(v) \oplus b && \text{if } (a, b) \in \mathcal{A}^+ \\ \Phi((a, b)v) &= a \ominus \Phi(v) \ominus b && \text{if } (a, b) \in \mathcal{A}^- \end{aligned}$$

See Figure 2 for an example of the action of Φ which will also be helpful in understanding the observation below.

Observation 2.2. The function Φ is a size-preserving bijection.

Proof: Clearly, Φ is size-preserving. So, let $\omega \in \mathcal{X}$ be given. We provide an inductive proof that there is an element $w \in \mathcal{L}$ with $\Phi(w) = \omega$, and establish at the same time that w is uniquely determined by ω which shows that Φ is a bijection.

The base cases are where ω is monotone. Since the signs of words in \mathcal{L} having two or more letters alternate, no such word maps to a monotone permutation. On the other hand $\Phi(a) = a$ for $|a| \geq 2$ and $\Phi(1) = 1$ so each monotone permutation is the image of a unique element of \mathcal{L} .

If ω is positive and crossed (the negative case follows by symmetry of course) then, according to Observation 2.1, there is a unique pair, (a, b) , of non-negative integers such that $\omega = a \oplus \nu \oplus b$ where ν is negative. By induction there is a unique $v \in \mathcal{L}$ whose first symbol is negative such that $\nu = \Phi(v)$. Then $\theta = \Phi((a, b)v)$. Conversely if $\Phi(w') = \omega$ and $w' = w'_1 w'_2 \dots w'_m$, then we must have $m \geq 2$ since ω is crossed. If $w'_1 = (a', b')$ then, since: $\Phi(w') = a' \oplus \Phi(w'_2 \dots w'_m) \oplus b'$, $\Phi(w'_2 \dots w'_m)$ is negative, and Observation 2.1, it follows that $(a', b') = (a, b)$ and $\Phi(w'_2 \dots w'_m) = \nu$. By induction, $w' = w$ so Φ is one-to-one. \square

Since Φ is a size-preserving bijection we freely identify \mathcal{L} and \mathcal{X} . However, we would like to have an internal characterisation in \mathcal{L} of the ordering that corresponds to the subpermutation ordering in \mathcal{X} under this identification.

This is simple enough when we are looking for monotone subsequences of elements of \mathcal{X} . Given $a \geq 2$ and $t \in \mathcal{L}$, to check whether $a \leq \Phi(t)$ all we need to do is add up all the positive numbers occurring in letters of t (in either coordinate, and in the element of \mathcal{M} if it's positive) and add one more if the letter from \mathcal{M} occurring in t is negative. If the answer is $\geq a$ then $a \leq \Phi(t)$, and not otherwise. So, it remains to consider the crossed case.

In \mathcal{A}^+ and \mathcal{A}^- we will take sums coordinatewise, i.e., $(a, b) + (c, d) = (a + c, b + d)$ in either case (but sums of letters of different types are not defined). The ordering on \mathcal{A}^+ and \mathcal{A}^- is also to be taken coordinatewise but with respect to absolute value, i.e., $(a, b) \leq (c, d)$ if and only if $|a| \leq |c|$ and $|b| \leq |d|$. Let $w = w_1 w_2 \dots w_k$ be any word over $\mathcal{A}^+ \cup \mathcal{A}^-$ (only the case of alternating signs will be of interest to us but for the following definition it's not required). The *positive content* of w is the sum of all its positive letters, and the *negative content* of w is the sum of all its negative letters.

Given $(a, b) \in \mathcal{A}^+$ and $w = w_1 w_2 \dots w_k w_{k+1} \in \mathcal{L}$, define the (a, b) -*prefix* of w to be the minimal prefix (if any) of w whose positive content is greater than or equal to (a, b) and define the (a, b) -*remainder*

of w to be the remaining suffix after the (a, b) -prefix is removed – both are undefined (denoted \perp) if w has no prefix whose positive content is at least (a, b) .

The next proposition captures the essence of the argument that we can recognise involvement between crossed permutations and other elements of \mathcal{X} greedily – the uncrossed case has already been remarked upon above.

Proposition 2.3. *Let negative $\tau, \theta \in \mathcal{X}$, $(a, b) \in \mathcal{A}^+$ and (c, d) an arbitrary pair of non-negative integers be given. Then*

$$a \oplus \theta \oplus b \leq c \oplus \tau \oplus d$$

if and only if

$$\max(0, a - c) \oplus \theta \oplus \max(0, b - d) \leq \tau.$$

Moreover, if τ is monotone, then $a \oplus \theta \oplus b \not\leq \tau$, while if $\tau = c \oplus \nu \oplus d$ where ν is positive, then $a \oplus \theta \oplus b \leq \tau$ if and only if $a \oplus \theta \oplus b \leq \nu$.

Proof: The first part follows immediately from the observation that $1 \oplus \alpha \leq 1 \oplus \beta$ if and only if $\alpha \leq \beta$, and a symmetric variant for $\alpha \oplus 1$. Similarly, the non-trivial half of the second part follows from the observation that if α begins with its minimum or ends with its maximum, then $\alpha \leq 1 \oplus \beta$ if and only if $\alpha \leq \beta$ (and again, a symmetric variant for $\beta \oplus 1$). \square

Corollary 2.4. *Let $w = (a, b)v \in \mathcal{L}$ and let $t \in \mathcal{L}$. Then $\Phi(w) \leq \Phi(t)$ if and only if the (a, b) -prefix of t is defined and $\Phi(v) \leq \Phi(s)$ where s is the (a, b) -suffix of t .*

For $(a, b) \in \mathcal{A}^+ \cup \mathcal{A}^-$ set $P(a, b)$ to be the set of all words $w_1 w_2 \dots w_k$ over the alphabet $\mathcal{A}^+ \cup \mathcal{A}^-$ that satisfy:

- w_1 has the same sign as (a, b) ,
- the w_i alternate in sign,
- the (sign of w_1) content of $w_1 w_2 \dots w_k$ is at least (a, b) , and
- the (sign of w_1) content of $w_1 w_2 \dots w_{k-1}$ is *not* at least (a, b) .

Except for the sign restriction on w_1 , the elements of $P(a, b)$ are exactly the candidates for being an (a, b) -prefix. Note that it is a consequence of the last two conditions that the sign of w_k agrees with that of w_1 (i.e. k is odd). Let $F_{(a,b)}(x)$ be the generating function for the set $P(a, b)$ enumerated by weight.

For a letter $a \in \mathcal{M}$, let $M_a(x)$ be the generating function for the set of all words whose corresponding permutations contain a .

Applying Corollary 2.4 inductively we see that if $w = w_1 w_2 \dots w_n \in \mathcal{L}$, with w_1 positive and $t \in \mathcal{L}$ then $\Phi(w) \leq \Phi(t)$ if and only if t can be factored as the concatenation of: an optional negative letter, followed by elements of $P(w_1)$, $P(w_2)$, through $P(w_{n-1})$ followed by a word that represents a permutation having a monotone subsequence w_n . Therefore, the generating function for the permutations containing $\Phi(w)$ is:

$$\frac{1}{(1-x)^2} \left(\prod_{i=1}^{n-1} F_{w_i}(x) \right) M_{w_n}(x).$$

By obvious symmetries $F_{(a,b)} = F_{(b,a)} = F_{(-a,-b)} = F_{(-b,-a)}$, and $M_a = M_{-a}$ so we obtain:

Theorem 2.5. *If $w = w_1 w_2 \dots w_n$ and $v = v_1 v_2 \dots v_n$ belong to \mathcal{L} , and there is a permutation $\rho : [n-1] \rightarrow [n-1]$ such that w_i and $v_{\rho(i)}$ differ only in their sign and/or the order of their coordinates, and if $w_n = \pm v_n$ then the permutations $\Phi(w)$ and $\Phi(v)$ are Wilf-equivalent in \mathcal{X} .*

The generating function for \mathcal{X} (including the empty permutation) is

$$\frac{1 - 2x}{1 - 4x + 2x^2}$$

and the growth rate of the class is therefore $2 + \sqrt{2}$.

On the other hand, Theorem 2.5 implies that the number of Wilf-equivalence classes in \mathcal{X} for structures of size $N \geq 2$ is dominated by the number of structures of size N in another universe. In this universe, structures are made up of a multiset of atoms (a, b) with $0 \leq a \leq b$ of weight $a + b$ and a single atom $m \geq 2$ of weight m . The size of a structure is the sum of the weights of its elements. Effectively then those structures correspond to coloured partitions where there are $\lceil (i+1)/2 \rceil$ different colours for parts of size i corresponding to the pairs (a, b) with $0 \leq a \leq b$ and $a + b = i$, together with a single part of size at least two. The generating function for this universe is:

$$\frac{x^2}{1-x} \prod_{i=1}^{\infty} \left(\frac{1}{1-x^i} \right)^{\lceil (i+1)/2 \rceil}.$$

After taking logarithms it is easy to verify that this converges for all $0 \leq x < 1$. By Pringsheim's theorem, the radius of convergence of the generating function for Wilf-equivalence classes in \mathcal{X} is 1 (it certainly can't be greater than 1). In particular, the number of such Wilf-equivalence classes is $o(c^n)$ for all $c > 1$.

Summarising:

Theorem 2.6. *Let $x_n = |\mathcal{X}_n|$ and let w_n be the number of Wilf-equivalence classes in \mathcal{X}_n . Then,*

$$x_n^{1/n} \rightarrow 2 + \sqrt{2} \quad \text{and} \quad w_n^{1/n} \rightarrow 1 \quad \text{as} \quad n \rightarrow \infty.$$

The rules for Wilf-equivalence in \mathcal{X} given by Theorem 2.5 are by no means complete. For instance, it is easy to check by considering avoidance that if θ and τ are Wilf-equivalent in \mathcal{X} and neither ends with its maximum, then $1 \oplus \theta$ and $1 \oplus \tau$ are also Wilf-equivalent (see Figure 2 for a ‘‘proof without words’’). So, for instance 1243 and 1342 are Wilf-equivalent in \mathcal{X} .

3 Sum closure of increasing oscillations

We now consider an important class \mathcal{SIO} (Subpermutations of the Increasing Oscillation) to which the previous conditions do not apply and yet we can still demonstrate an exponential Wilf-collapse. This class consists of all finite permutations whose pattern can be found in the *infinite increasing oscillation*:

$$2, 4, 1, 6, 3, 8, 5, 10, 7, \dots$$

This sequence (and various minor modifications of it) arises in a wide variety of applications ranging from the recreational e.g., the theory of juggling [8] to the more serious in that it provides those sequences that are most difficult to sort by block reversals [6]. The class \mathcal{SIO} also plays a central role in the study of growth rates of permutation classes providing a fundamental building block for the construction of




Fig. 2: A proof by pictures that, if θ and τ are Wilf-equivalent in \mathcal{X} and neither ends with its maximum, then $1 \oplus \theta$ and $1 \oplus \tau$ are also Wilf-equivalent. The large box represents an element of either $\text{Av}(1 \oplus \theta)$ or $\text{Av}(1 \oplus \tau)$. If the element indicated by a \bullet is not present, then the permutation avoids both. If it is present, then the central region must avoid θ (respectively, τ) and the bijection witnessing their Wilf-equivalence can be applied to it.

intervals in the set of achievable growth rates and other threshold phenomena [7, 9, 10]. The class SIO is invariant under four of the symmetries for permutations in general: the identity, reverse-complement, inverse, and the composition of reverse-complement with inverse.

The *inversion graph* of a sequence or permutation π is the graph on the indices of the sequence where, for $i < j$, i and j are adjacent if $\pi_i > \pi_j$, i.e., if the pair (i, j) forms an inversion in the sequence. The inversion graph of the infinite increasing oscillation is a (one-ended) infinite path and the finite permutations in the class SIO have sum decompositions whose components have inversion graphs that are paths (possibly consisting only of a single element). For each path of length at least three there are two permutations whose inversion graph is that path and both of these occur in SIO . So the alphabet \mathcal{A} of sum-indecomposables for the class SIO consists of: a symbol a representing the singleton permutation (a), a symbol b representing 21 , and for $k \geq 3$ two symbols w_k and m_k representing the two different permutations of size k whose inversion graph is a path. The first few of these are shown explicitly in Figure 3 along with a convenient symbolic representation of each.

In the sum-indecomposables of size at least 3 we define the *Start* of that sum indecomposable to be “Up” (\nearrow) if the initial element of the permutation corresponds to an interior vertex of the inversion graph, and “Down” (\searrow) if it is a terminal vertex. The corresponding definition for *Finish* is that it is \nearrow if the final element of the permutation corresponds to an interior vertex of the inversion graph, and \searrow if it is a terminal vertex. Start and Finish are “Undefined” (\perp) for the permutations 1 or 21 .

Since every permutation in SIO is uniquely represented as a sum of sum-indecomposables and all such sums belong to SIO we freely identify SIO with words over the alphabet \mathcal{A} . We further define the Start of a word to be the Start of its first letter (possibly undefined) and its Finish to be the Finish of its terminal letter. The *type* of a word is the pair consisting of its Start and Finish.

The symbolic representation of permutations in SIO makes it particularly easy to identify when one permutation is involved in another – more specifically when a word is involved in a single letter. The basic idea is that you must try to pack the individual parts of the word into the letter (and may do so greedily from left to right) using the following conventions:










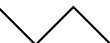

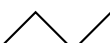

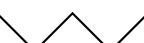



Sum indecomposable	Diagram	Symbol	Letter	Start	Finish
1			a	⊥	⊥
21			b	⊥	⊥
231			w ₃	↘	↗
312			m ₃	↗	↘
2413			w ₄	↘	↘
3142			m ₄	↗	↗
24153			w ₅	↘	↗
31524			m ₅	↗	↘

Fig. 3: The sum indecomposables of size at most 5 in the class SZO together with their diagrams, symbolic representation, representative letters of the alphabet \mathcal{A} and start and finish type (defined in text).

- The letter a can be matched to any single vertex,
- The letter b can be matched to any edge,
- Any letter w_k or m_k can be matched to a connected sequence of $k - 1$ edges that starts with an edge of the appropriate slope.

All this is subject to the condition that the parts of the symbol matching consecutive letters cannot be connected by an edge. For instance the relationship $w_3am_4b \leq m_{16}$ (and implicitly $w_3am_4b \not\leq m_{14}$) is shown by the thicker edges (and dot) below:



The symbolic viewpoint also makes it easy to understand the effect on SIO of the symmetries: reverse-complement, inverse, and their composition. Respectively, they correspond to: rotation by 180° , reflection in a horizontal axis, and reflection in a vertical axis.

Observation 3.1. The class SIO is supercritical since the generating function for its sum-closed permutations is:

$$x + x^2 + 2x^3/(1 - x).$$

However, it has no incompatible pairs since for any pair of plus-indecomposable permutations $\pi, \theta \in SIO$, $\pi \oplus \theta$ occurs as a pattern in a sum-indecomposable of size at most $|\pi| + |\theta| + 2$.

Define an equivalence relation \sim on \mathcal{A}^* by $X \sim Y$ if and only if there is a size and type-preserving bijection $\Phi : SIO \rightarrow SIO$ such that the image of $SIO \cap Av(X)$ is $SIO \cap Av(Y)$ (where again as usual we make no distinction between the words X and Y and the permutations they represent). In particular if $X \sim Y$ then $X \equiv_{SIO} Y$. So, if we can demonstrate a significant collapse for the \sim equivalence relation on SIO then we will have demonstrated a Wilf-collapse in SIO as well.


Lemma 3.2. *Let $X \in \mathcal{A}^*$ and a symmetry σ be given such that both the Start and Finish of X are defined and equal the Start and Finish of $\sigma(X)$. Then $X \sim \sigma(X)$*

Proof: The result would be trivial if we only required a size-preserving bijection since σ is such a bijection between $SIO \cap Av(X)$ and $SIO \cap Av(\sigma(X))$. However, the required construction is not that much more complex. Without loss of generality suppose that the type of X is (\nearrow, \searrow) (all other cases can be proven in the same way). In this case σ must be the symmetry whose action on symbols is reflection in a vertical axis. This symmetry has the effect of exchanging and inverting the Start and Finish of a word.

If $W \in \{a, b, w_3\}^*$ just set $\Phi(W) = W$ (since the Start and Finish of X are \nearrow and \searrow respectively, all such W avoid both X and $\sigma(X)$). Otherwise let $W = PES$ where P is the maximal prefix of W in $\{a, b, w_3\}^*$ and S is the maximal suffix in $\{a, b, w_3\}^*$ (either might be empty). The word E (the “essential” part of W) has a defined Start and Finish. If $\Phi(E)$ is defined (having the same type and size as E and behaving correctly in the sense that $X \leq E$ if and only if $\sigma(X) \leq \Phi(E)$) then we can set $\Phi(W) = P\Phi(E)S$ since this preserves type and size and $X \leq W$ if and only if $X \leq E$.

So let a word E in \mathcal{A}^* having defined Start and Finish be given that neither begins nor ends with w_3 . Define the word E' whose diagram is obtained from the diagram of E by deleting the first edge if the Start of E is not \nearrow and the last edge if the Finish of E is not \searrow . Because E does not begin or end with w_3 the Start of E' is \nearrow and its Finish is \searrow . Note also that $X \leq E$ if and only if $X \leq E'$. Finally, to define $\Phi(E)$ just take $\sigma(E')$ and reattach the edges that were deleted. For example, if $E = w_7am_5$ then

the following series of diagrams shows the construction of $\Phi(E) = w_6aw_6$ (the dots on the endpoint of the edge that is deleted/added are for emphasis only):



The following table shows how to adjust the construction for each possible type of X – the main point is that we remove the maximal unusable prefix (with respect to including X) and the maximal unusable suffix (ditto) and then operate on the remaining essential part.

Type of X	Prefix	Suffix
(\nearrow, \nearrow)	$\{a, b, w_3\}^*$	$\{a, b, m_3\}^*$
(\nearrow, \searrow)	$\{a, b, w_3\}^*$	$\{a, b, w_3\}^*$
(\searrow, \nearrow)	$\{a, b, m_3\}^*$	$\{a, b, m_3\}^*$
(\searrow, \searrow)	$\{a, b, m_3\}^*$	$\{a, b, w_3\}^*$

□

Now comes the main ingredient of the collapse – a substitution principle:

Proposition 3.3. *Suppose that $X \sim Y$ and that $P, S \in \mathcal{A}^*$. Then $PXS \sim PYS$.*

Proof: Let $\Phi : SIO \rightarrow SIO$ be a size and type-preserving bijection whose existence is guaranteed by the condition that $X \sim Y$. We will construct a size and type-preserving bijection $\Psi : SIO \rightarrow SIO$ which witnesses $PXS \sim PYS$.

To this end, let $W \in \mathcal{A}^*$ be given. If $PS \not\leq W$ set $\Psi(W) = W$. So suppose that $PS \leq W$. Take a leftmost embedding of P into W and a rightmost embedding of S into W . If P is non-empty then the embedding of P ends in some letter x and there might remain some “usable” part of x (beginning with the next vertex but one after the end of P) call it x' (if P is empty then x will just be the first character of W and $x' = x$). Likewise S starts in some letter y and there might be some usable part of y before the beginning of S , call it y' . Suppose first that $x \neq y$ (not just as characters but as letters in the word W) and write $W = AxMyB$. Consider $W' = x'My'$. Now $PXS \leq W$ if and only if $X \leq W'$. So $PXS \leq W$ if and only if $Y \leq \Phi(W')$. Since the Start and Finish of $\Phi(W')$ are the same as of W' we can reattach the part of x before W' and the part of y after W' to $\Phi(W')$ in the same way as they were originally attached to W . Define $\Psi(W)$ to be the concatenation of A , this reattachment, and B . If P is non-empty then Ax is non-empty and the Start of $\Psi(W)$ is the same as the Start of Ax and hence of W . On the other hand if P is empty then the Start of W is just the Start of x' (and so of W'), and once again the Start of $\Psi(W)$ is the same as that of W . Similar considerations show that the Finish of $\Psi(W)$ is the same as the Finish of W and complete this case.

However, the case when x any y are the same character is essentially the same, but W' just consists of a single letter which is the usable part of the letter in which P ends and S begins. □

Theorem 3.4. *The class SIO has an exponential Wilf-collapse.*

Proof: As noted previously, the generating function for the plus indecomposable members of SIO is

$$A(t) = t + t^2 + \frac{2t^3}{1-t}$$

which is rational and of radius of convergence 1 and therefore the class SIO is super-critical.

Proposition 3.10 of [3] implies that for some constant $c > 0$, in all but an exponentially small proportion of the words in \mathcal{A}^* of weight n there are at least cn disjoint occurrences of the factor w_3m_4 . However $w_3m_4 \sim w_4m_3$ so the \sim (and hence Wilf) equivalence classes of all such words contain at least 2^{cn} elements, which demonstrates an exponential Wilf-collapse. \square

4 Conclusion

This paper provides more evidence for a thesis that has previously been explored in [1, 5, 3]. Put most broadly:

If the elements of a class \mathcal{C} admit a natural representation as words over some weighted alphabet in such a way that containment can be recognised at the word level in a greedy fashion, then global symmetries of the class can often be applied locally to provide Wilf-equivalent structures and this frequently implies a Wilf collapse.

However, we have no other general criteria that yield Wilf-collapse.

The class $\text{Av}(321)$ provides a testing ground for this thesis for a couple of reasons. Results in [4] show that there is a greedy form of testing for containment within this class. On the other hand, results in [2] also give a natural (or at least arguably natural) encoding of the elements of the class in the form of words. This encoding implies a restricted form for the generating functions of the classes $\text{Av}(321, \pi)$, namely that they must be rational. These are exactly the classes where we would hope to see many coincidences between the generating functions if a Wilf-collapse exists in $\text{Av}(321)$. However, the two results do not fit together well enough to allow the sort of manipulations that we have seen here, so this case remains an intriguingly open one.

4.1 Acknowledgement

The authors are grateful to an anonymous referee for suggestions that have greatly improved the presentation and clarity of some of the material.

References

- [1] Michael Albert and Mathilde Bouvel. A general theory of Wilf-equivalence for Catalan structures. *Electron. J. Combin.*, 22(4):Paper 4.45, 29, 2015.
- [2] Michael Albert, Robert Brignall, Nik Ruškuc, and Vincent Vatter. Rationality for subclasses of 321-avoiding permutations. *European Journal of Combinatorics*, 78:44 – 72, 2019.
- [3] Michael Albert, Vít Jelínek, and Michal Opler. Wilf collapse in permutation classes. *ArXiv:1909.13348*, 2019.

- [4] Michael Albert, Marie-Louise Lackner, Martin Lackner, and Vincent Vatter. The complexity of pattern matching for 321-avoiding and skew-merged permutations. *Discrete Math. Theor. Comput. Sci.*, 18(2):Paper No. 11, 17, 2016.
- [5] Michael Albert and Jinge Li. Wilf-collapse in permutation classes having two basis elements of size three. *ArXiv:1710.04107*, October 2017.
- [6] V. Bafna and P. Pevzner. Genome rearrangements and sorting by reversals. *SIAM Journal on Computing*, 25(2):272–289, 1996.
- [7] David Bevan. Intervals of permutation class growth rates. *Combinatorica*, 38(2):279–303, 2018.
- [8] Joe Buhler, David Eisenbud, Ron Graham, and Colin Wright. Juggling drops and descents. *The American Mathematical Monthly*, 101(6):507–519, 1994.
- [9] Jay Pantone and Vincent Vatter. Growth rates of permutation classes: categorization up to the uncountability threshold. *ArXiv:1605.04289*, May 2016.
- [10] Vincent Vatter. Permutation classes of every growth rate above 2.48188. *Mathematika*, 56(1):182–192, 2010.
- [11] Vincent Vatter. Permutation classes. In *Handbook of enumerative combinatorics*, Discrete Math. Appl. (Boca Raton), pages 753–833. CRC Press, Boca Raton, FL, 2015.