

Growth rates for subclasses of $\text{Av}(321)$

M. H. Albert¹, M. D. Atkinson¹, R. Brignall², N. Ruškuc³,
Rebecca Smith⁴, and J. West⁵

¹Department of Computer Science, University of Otago

²Department of Mathematics and Statistics, The Open University

³School of Mathematics and Statistics, University of St Andrews

⁴Department of Mathematics, SUNY Brockport

⁵Heilbronn Institute for Mathematical Research, University of Bristol

Submitted: Jan 15, 2010; Accepted: Sep 27, 2010; Published: Oct 22, 2010

Mathematics Subject Classification: 05A05, 05A16

Abstract

Pattern classes which avoid 321 and other patterns are shown to have the same growth rates as similar (but strictly larger) classes obtained by adding articulation points to any or all of the other patterns. The method of proof is to show that the elements of the latter classes can be represented as *bounded merges* of elements of the original class, and that the bounded merge construction does not change growth rates.

1 Introduction

A pattern class is, roughly, a collection of permutations that satisfy certain restrictions on the configurations of their elements (formal definitions can be found in the next section). Most commonly, the restrictions are expressed by prohibiting particular types of subsequence. For example, the collection of all permutations containing no descending subsequence of length 3 is a pattern class. More generally, if B is any set of such restrictions, we write $\text{Av}(B)$ to denote the pattern class they define. The study of such classes dates back at least to work of Knuth [7], or even further to the celebrated result of Erdős and Szekeres [5] that every permutation of length greater than ad must include either an ascending subsequence of length $a + 1$ or a descending one of length $d + 1$.

Initially, research into pattern classes focussed on enumeration – determining the number of permutations of length n in a given pattern class. An early result of this type [7] was that $\text{Av}(231)$ and $\text{Av}(321)$ are both enumerated by the Catalan sequence (and by easy symmetries so also is every class $\text{Av}(\alpha)$ with $|\alpha| = 3$). It was initially expected that

$\text{Av}(231)$ and $\text{Av}(321)$ would have further properties in common but these hopes have evaporated since the discovery [2] that $\text{Av}(231)$ contains only countably many subclasses whilst $\text{Av}(321)$ contains uncountably many. In fact $\text{Av}(231)$ is a very tractable class compared to $\text{Av}(321)$ and, in particular, there is an efficient algorithm [1] to enumerate any given subclass of it. By contrast comparatively few subclasses of $\text{Av}(321)$ have been enumerated exactly and so attention has turned to growth rate estimates.

Growth rates offer a way of approximating the number of permutations of a given length in a pattern class. They have become especially important since Marcus and Tardos [8] proved the Stanley-Wilf conjecture (that for every proper pattern class there is an exponential bound on the number of permutations of length n which it contains). This result implies that every proper pattern class \mathcal{C} has a *growth rate* defined to be the limit superior of the n^{th} root of the number of permutations in \mathcal{C} of length n . Growth rates have been investigated by Bóna [3, 4] who found bounds (relative to the size of the forbidden patterns) and established results on what form this growth rate might take. Recently, Vatter [10] has proven that every real number greater than 2.482 occurs as the growth rate of some pattern class. Because of these results and others, we shall investigate the growth rates of pattern subclasses of $\text{Av}(321)$ and we will be particularly interested in the case when distinct subclasses of $\text{Av}(321)$ have the same growth rate.

Consider a pattern class \mathcal{C} of the form $\text{Av}(321, X)$ where X is some arbitrary set of permutations. Consider also $\mathcal{C}' = \text{Av}(321, X')$ where X' is obtained from X by adding or removing “articulation points” (similar to the 3 of 21354) anywhere within the patterns of X . The main result of this paper is that \mathcal{C} and \mathcal{C}' have the same growth rate. In order to prove this result we introduce a number of new concepts and constructions, including the notions of k -rigidity, bounded merges, and staircase decompositions, which we discuss in some generality.

The structure of the remainder of this paper is as follows:

Section 2 introduces the formal definitions, and certain preliminary results concerning rigidity and growth rates.

Section 3 contains the proof of the main result, divided into two cases for clarity, using staircase decompositions.

Section 4 examines the distributive lattices of occurrences of 21 in a 321-avoiding permutation, and shows that every subdirect product of two chains can arise in this fashion.

Section 5 concludes the paper with some further remarks, and open problems.

2 Preliminaries

A permutation $\pi \in \mathcal{S}_n$ is a bijective map from $[n] = \{1, 2, \dots, n\}$ to itself, and is therefore a set of ordered pairs

$$\{(1, \pi(1)), (2, \pi(2)), \dots, (n, \pi(n))\}$$

(traditionally and more frequently written as the sequence $\pi(1)\pi(2)\cdots\pi(n)$). So, when we say $x \in \pi$ we are simply referring to some member of this set. However, it is frequently necessary to relate elements of π either by the values of their first or second coordinates. Normally, we think of the first coordinates as lying on a horizontal axis so words and phrases such as “precedes”, “follows”, “to the left of”, etc. refer to that ordering. Conversely words such as “larger”, “smaller”, “above” and “below” relate to the ordering of the second coordinate.

An *involvement* or *embedding* of a permutation α in π is a map $f : \alpha \rightarrow \pi$ that respects both these orderings. In other words x precedes (is larger than) y in α if and only if $f(x)$ precedes (is larger than) $f(y)$ in π . In particular an embedding is necessarily injective. The composition of embeddings is an embedding and so the relation “is involved in” is a partial order, which will be denoted \preceq . If a subset of π is the image of α under an embedding, then we say that the pattern of the subset is α . We say that $x \in \pi$ *occurs as an i in an embedding of α* (or just “as i in an α ”) if there is an embedding of α in π such that x is the image of the element of α whose *second* coordinate (i.e. value¹) is i . A *pattern class*, or simply *class* of permutations is a set of permutations closed downward under \preceq . Such a class, \mathcal{C} , can also be defined as the set of permutations which *avoid*, i.e. do not involve, any of the elements of some set B of permutations. In that case we write $\mathcal{C} = \text{Av}(B)$. If B is a \preceq -antichain, then it is called the *basis* of \mathcal{C} (note that, for any set B , the set of minimal elements of B is an antichain and forms the basis of $\text{Av}(B)$). We define the *growth rate* (sometimes called the *Stanley-Wilf limit*, or *upper growth rate*) of \mathcal{C} :

$$s(\mathcal{C}) = \limsup_{n \rightarrow \infty} |\mathcal{C} \cap \mathcal{S}_n|^{1/n}.$$

As noted in the introduction, Marcus and Tardos [8] proved that if \mathcal{C} is a proper pattern class, then $s(\mathcal{C}) < \infty$.

The *increasing* and *decreasing* permutations of length k are

$$\begin{aligned} \iota_k &= \{(1, 1), (2, 2), (3, 3), \dots, (k, k)\} \\ \delta_k &= \{(k, 1), (k-1, 2), (k-2, 3), \dots, (1, k)\} \end{aligned}$$

respectively. A subset of π is called increasing (respectively decreasing) if its pattern is some increasing (decreasing) permutation.

Throughout this paper, we are primarily concerned with permutations that can be written as the union of k increasing subsets for some fixed value of k . These permutations form a pattern class \mathcal{I}_k , whose basis is the single decreasing permutation δ_{k+1} . We say that a permutation $\pi \in \mathcal{I}_k$ is *k-rigid* if every element of π belongs to a subset whose pattern is δ_k .

Suppose that $\pi \in \mathcal{I}_k$. We can define a decomposition of π into increasing subsets C_1, C_2, \dots, C_k by defining, for $1 \leq t \leq k$:

$$C_t = \left\{ x \in \pi : \begin{array}{l} x \text{ occurs as the maximum of some} \\ \delta_t \text{ but not of any } \delta_{t+1} \end{array} \right\}.$$

¹Why value? Because, in the usual “one line” notation for permutations, it is easy to identify the element of value i , and not necessarily so easy to identify the element at position i .

This decomposition is the one produced by a greedy algorithm, which takes the elements of π in order from right to left, and adds each successive element x to the first C_j of which x is smaller than the current minimum. If $x \in \pi$ belongs to C_i then we say that the *rank* of x is i .

Lemma 1. *If $\pi \in \mathcal{I}_k$, and $x \in \pi$ occurs as an i in some δ_k , then the rank of x is i . Consequently, the position of x in all the δ_k to which it belongs is the same.*

Proof. Choose a δ_k in which x occurs as i , and write it in one line notation as AxB (so A is a decreasing sequence of length $k - i$ and B a decreasing sequence of length $i - 1$). Then x occurs as the maximum of the δ_i , xB . It cannot occur as the maximum of any δ_{i+1} , xC , because then AxC would be a δ_{k+1} in π . \square

It follows that if ρ is k -rigid, then any embedding of ρ in $\pi \in \mathcal{I}_k$ must preserve the ranks of the elements of ρ , as it preserves sets whose pattern is δ_k .

If two elements of a permutation coincide or form a 12 pattern, then it makes sense to speak of their *infimum* – it is simply the smaller and earlier of the two, and likewise their *supremum* which is the larger and later. If $f, g : \rho \rightarrow \pi$ are two embeddings of a k -rigid permutation into an element of \mathcal{I}_k , then for any $x \in \rho$, the ranks of $f(x)$ and $g(x)$ are the same. Therefore $f(x)$ and $g(x)$ occur in some increasing subset of π and hence their infimum and supremum are defined. In fact more is true:

Theorem 2. *Let $\pi \in \mathcal{I}_k$, ρ a k -rigid permutation, and two embeddings $f, g : \rho \rightarrow \pi$ be given. Then $I, S : \rho \rightarrow \pi$ defined for $x \in \rho$ by $I(x) = \inf(f(x), g(x))$, and $S(x) = \sup(f(x), g(x))$ are also embeddings of ρ in π . In particular, the embeddings of ρ in π form a distributive lattice.*

Proof. We give the argument for I only (that for S is similar). It suffices to show that for any $x, y \in \rho$ (without loss of generality, x preceding y), the pattern of $I(x)$ and $I(y)$ in π is the same as the pattern of x and y in ρ . But, this is essentially trivial. If the pattern of xy is 12 then $\inf(f(x), f(y)) = f(x)$ and $\inf(g(x), g(y)) = g(x)$. So, $\inf(f(x), g(x))$ must form a 12 pattern with $\inf(f(y), g(y))$. The case where xy has pattern 21 is just the same. \square

More generally, given two embeddings f and g of an arbitrary permutation α in an arbitrary permutation β such that the images $f(a)$ and $g(a)$ of any $a \in \alpha$ coincide or form a 12 pattern, the maps I and S defined in the theorem are also embeddings of α in β . We will defer a discussion of the distributive lattices mentioned in the theorem above to Section 4.

Applying the previous theorem repeatedly, we can take the infimum of *all* of the embeddings of a k -rigid permutation into an element $\pi \in \mathcal{I}_k$, thus obtaining:

Corollary 3. *Let $\pi \in \mathcal{I}_k$ and ρ a k -rigid permutation be given. If $\rho \preceq \pi$ then there is an embedding of ρ in π which simultaneously minimizes the position and value of every element of the image of ρ among all such embeddings.*

Naturally enough, we call the embedding whose existence is asserted by this corollary the *leftmost-bottommost* embedding of ρ in π .

A permutation π is called a *merge* of two permutations α and β if it can be written as the disjoint union of two sets, the first of which has pattern α and the second of which has pattern β . If \mathcal{A} and \mathcal{B} are pattern classes, then

$$\mathcal{M}(\mathcal{A}, \mathcal{B}) = \{\pi : \pi \text{ is a merge of some } \alpha \in \mathcal{A} \text{ and some } \beta \in \mathcal{B}\}$$

is also a permutation class, called the *merge* of \mathcal{A} and \mathcal{B} . For instance $\mathcal{M}(\mathcal{I}_s, \mathcal{I}_t) = \mathcal{I}_{s+t}$ for any s and t .

Let two permutations α and β be given, together with embeddings $a : \alpha \rightarrow \pi$, $b : \beta \rightarrow \pi$ that witness π being a merge of α and β (so the ranges of the embeddings are disjoint and their union is equal to π). For $x \in \pi$ define the *type of x* , $\text{tp}(x) = a$ if x is in the range of a and $\text{tp}(x) = b$ if it is in the range of b . For $1 \leq c < |\pi|$, if the types of $(c, \pi(c))$ and $(c+1, \pi(c+1))$ are different, then we say that there is a *type change by position* at c . Similarly, for $1 \leq r < |\pi|$, if the types of $(\pi^{-1}(r), r)$ and $(\pi^{-1}(r+1), r+1)$ are different, then we say that there is a *type change by value* at r .

Given a positive integer B and two permutation classes \mathcal{C} and \mathcal{D} we define the *B-bounded merge of \mathcal{C} and \mathcal{D}* :

$$\mathcal{M}_B(\mathcal{C}, \mathcal{D}) = \left\{ \pi : \begin{array}{l} \pi \text{ is a merge of some } \alpha \in \mathcal{C} \text{ and some } \beta \in \mathcal{D} \text{ having} \\ \text{at most } B \text{ type changes in total, either by position} \\ \text{or value} \end{array} \right\}$$

As the number of type changes cannot increase when we delete elements of a merge, $\mathcal{M}_B(\mathcal{C}, \mathcal{D})$ is also a permutation class.

Example 1. *The permutation*

$$\{(1, 1), (2, 2), (3, 3), (4, 7), (5, 8), (6, 9), (7, 4), (8, 5), (9, 6)\}$$

(123789456 in one line notation) lies in $\mathcal{M}_3(\mathcal{I}_1, \mathcal{I}_1)$ because of the subsequences 123789 and 456 and the type changes $(6, 9)$ to $(7, 4)$ by position and $(3, 3)$ to $(7, 4)$ and $(9, 6)$ to $(4, 7)$ by value.

Theorem 4. *Let a positive integer B and two permutation classes \mathcal{C} and \mathcal{D} be given. Then,*

$$\begin{aligned} s(\mathcal{M}(\mathcal{C}, \mathcal{D})) &\leq \left(\sqrt{s(\mathcal{C})} + \sqrt{s(\mathcal{D})} \right)^2, \text{ and} \\ s(\mathcal{M}_B(\mathcal{C}, \mathcal{D})) &= \max(s(\mathcal{C}), s(\mathcal{D})). \end{aligned}$$

Proof. Let $c_n = |\mathcal{C} \cap \mathcal{S}_n|$, $d_n = |\mathcal{D} \cap \mathcal{S}_n|$, $M_n = |\mathcal{M}(\mathcal{A}, \mathcal{B}) \cap \mathcal{S}_n|$ and $m_n = |\mathcal{M}_B(\mathcal{A}, \mathcal{B}) \cap \mathcal{S}_n|$. A merge of $\alpha \in \mathcal{A} \cap \mathcal{S}_k$ and $\beta \in \mathcal{B} \cap \mathcal{S}_{n-k}$ can be defined by independently choosing k (from n) positions and k values to hold the pattern α , while fitting the pattern β in the remaining positions and values. It follows that:

$$M_n \leq \sum_{k=0}^n \binom{n}{k}^2 c_k d_{n-k}.$$

So,

$$s(\mathcal{M}(\mathcal{C}, \mathcal{D})) \leq \limsup_{n \rightarrow \infty} \left(\sum_{k=0}^n \binom{n}{k}^2 c_k d_{n-k} \right)^{1/n}.$$

The similarity of the square root of each term in the sum to a term of the expansion of $(\sqrt{s(\mathcal{C})} + \sqrt{s(\mathcal{D})})^n$ is sufficient to establish the first of the results claimed in the theorem (an argument that goes back to [9]).

For the second result, in order to specify a B -bounded merge of length n we need to specify at most B positions and values at which a type change can occur, and then two permutations in \mathcal{C} and \mathcal{D} of suitable length. Additionally, $\mathcal{C}, \mathcal{D} \subseteq \mathcal{M}_B(\mathcal{C}, \mathcal{D})$. So (certainly for $n > 2B$):

$$\max(c_n, d_n) \leq m_n \leq \binom{n}{B}^2 \max\{c_k d_{n-k} : 0 \leq k \leq n\}.$$

Taking n^{th} roots throughout, and observing that $\binom{n}{B}^{2/n} \rightarrow 1$ establishes the second result. \square

Note that $s(\mathcal{I}_k) = k^2$, so the bound given by the first estimate is tight for $\mathcal{M}(\mathcal{I}_n, \mathcal{I}_m)$. For the remainder of this paper we will only be using the second of these estimates; that the growth rate of a bounded merge of two permutation classes is the maximum of their individual growth rates.

The *direct sum* $\alpha \oplus \beta$ of two permutations α and β is that merge of α with β in which the image of α occupies the first $|\alpha|$ places both by position and value. A permutation π is called *plus indecomposable* if $\pi \neq \alpha \oplus \beta$ for any pair of non-empty permutations α and β .

If $\pi \in \mathcal{I}_2$ is not 2-rigid, then, for some α and β , $\pi = \alpha \oplus 1 \oplus \beta$ since it must contain an element which has no larger preceding element, nor any smaller following element. Thus, all the preceding elements (of pattern α) are smaller than it and the following ones (of pattern β) are larger. Such an element is called an *articulation point* of π . Conversely, $\pi \in \mathcal{I}_2$ is 2-rigid exactly if $\pi = \alpha_1 \oplus \alpha_2 \oplus \cdots \oplus \alpha_k$ where $k \geq 1$ and each α_i is a plus indecomposable permutation of length at least 2.

Let $1^n = \iota_n$ be the direct sum of n copies of the singleton permutation. If $\pi \in \mathcal{I}_2$ is an arbitrary permutation then there is a unique sequence $\rho_1, \rho_2, \dots, \rho_c$ of plus indecomposable permutations all of length at least 2 such that:

$$\pi = 1^{m_0} \oplus \rho_1 \oplus 1^{m_1} \oplus \rho_2 \oplus \cdots \oplus 1^{m_{c-1}} \oplus \rho_c \oplus 1^{m_c}.$$

In this case, we define the *rigid reduction* of π

$$\text{red}(\pi) = \rho_1 \oplus \rho_2 \oplus \cdots \oplus \rho_c.$$

For example:

$$\text{red}(241357689) = 241365.$$

For a set X of permutations $\text{red}(X) = \{\text{red}(\pi) : \pi \in X\}$.

3 The main result

We now turn our attention almost exclusively to infinite subclasses of $\mathcal{I}_2 = \text{Av}(321)$ with the aim of proving:

Theorem 5. *Let X be any finite set of permutations. Then $\mathcal{I}_2 \cap \text{Av}(X)$ and $\mathcal{I}_2 \cap \text{Av}(\text{red}(X))$ have the same growth rate.*

This seems a surprising result as, *a priori*, the second class appears to be much smaller than the first one – consider for instance $\mathcal{I}_2 \cap \text{Av}(21\ 34\ 65\ 7)$ and $\mathcal{I}_2 \cap \text{Av}(2143)$. To prove it, some further preparation is required.

A *staircase decomposition* of a permutation $\pi \in \mathcal{I}_2$ is a partition $\alpha_1, \alpha_2, \dots, \alpha_k$ of π that has the following properties:

- The pattern of each α_i is increasing;
- For $j \geq 1$, α_{2j} lies entirely to the right of α_{2j-1} ;
- For $j \geq 1$, α_{2j+1} lies entirely above α_{2j} ;
- If $i - j \geq 2$ then α_i lies entirely above and to the right of α_j .

Figure 1 should make it clear why the term “staircase decomposition” was chosen. We refer to the individual constituents α_i of the staircase as its *blocks*.

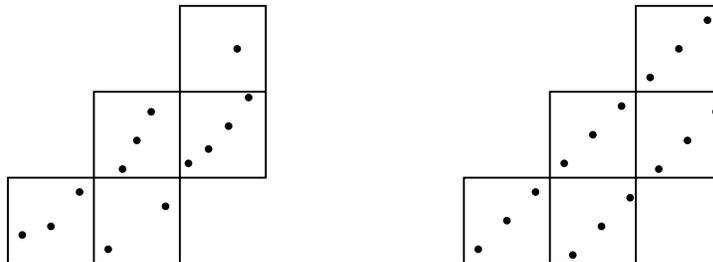


Figure 1: On the left, a staircase decomposition; and on the right, a generic staircase with five blocks of size three.

Every $\pi \in \mathcal{I}_2$ has a staircase decomposition. This can be constructed inductively by taking, for odd i , α_i to be the longest initial segment by position of $\pi \setminus \cup_{j < i} \alpha_j$ that has an increasing pattern; and for even i , α_i to be the longest initial segment by value of $\pi \setminus \cup_{j < i} \alpha_j$ that has an increasing pattern.

Let positive integers k and b be given. The *generic staircase with k blocks of size b* or (k, b) -generic staircase is that permutation π which has a staircase decomposition $\alpha_1, \alpha_2, \dots, \alpha_k$, where for each i , $|\alpha_i| = b$ and additionally:

- If $i \geq 1$ and $t \leq b$, then the t^{th} element of α_{2i} lies in value between the $(t - 1)^{\text{st}}$ and t^{th} elements of α_{2i-1} ;

- If $i \geq 1$ and $t \leq b$, then the t^{th} element of α_{2i+1} lies in position between the t^{th} and $(t + 1)^{\text{st}}$ elements of α_{2i} .

Figure 1 also illustrates an example of a generic staircase.

Proposition 6. *Every $\pi \in \mathcal{I}_2$ occurs as a pattern in some generic staircase.*

Proof. Let $\pi \in \mathcal{I}_2$ be given, and choose a staircase decomposition $\alpha_1, \alpha_2, \dots, \alpha_k$ of π . Consider the infinite set of points shown in Figure 2. The points in each of the line segments within a block are a translation of the set $D \cap (0, 1)$ where D is the set of dyadic rationals (rationals whose denominator is a power of 2) and therefore form a dense linear order without endpoints. Choose an arbitrary embedding of α_1 into the first block. Then, α_2 can be embedded into the second block in such a way that the pattern of $\alpha_1 \cup \alpha_2$ is preserved (simply because we have a dense linear order available here). Similarly, having embedded α_1 and α_2 , we can embed α_3 in the third block. Its relationship with the embedded copy of α_1 is fixed by the fourth condition in the definition of a staircase decomposition, and its proper relationship with the embedded copy of α_2 can be assured using the density again. Proceeding inductively we can find an embedding of π into this infinite set. Since π is finite, the range of this embedding is contained entirely among the points whose coordinates have a denominator at most 2^m for some m . Now reduce the infinite staircase to the finite set of points of this type. The result is not a generic staircase as some points share a common horizontal or vertical component. However, each odd numbered block can be shifted upwards by $1/2^{m+1}$ (or any suitably small amount) and each even numbered block leftwards by the same amount. This does not change the relationship of any pair of points that were previously on different horizontal or vertical lines (and in particular, the images of the points of π), and the resulting staircase is generic with k blocks of size $2^m - 1$. \square

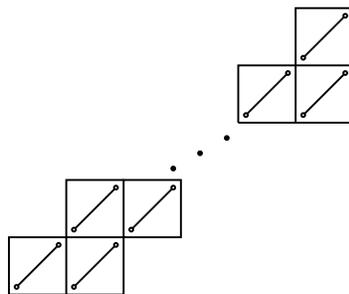


Figure 2: A staircase where each block is a dense linear order without endpoints.

The following technical proposition links together bounded merges and generic staircases. It shows that a 321-avoiding permutation that avoids a generic staircase is a

bounded merge of two increasing permutations where the parameters of the bounded merge are dependent on the parameters of the generic staircase. We use it in Propositions 8 and 9 to show that a permutation of $\text{Av}(321)$ that avoids some extra pattern other than 321 lies in a bounded merge of classes which avoid shorter (but related) patterns.

Proposition 7. *Let positive integers k and b be given. There is a positive integer B (depending only on k and b) such that for all $\pi \in \mathcal{I}_2$, either π contains a (k, b) -generic staircase, or π is a B -bounded merge of two permutations λ and β such that the image of λ contains all the elements preceding the minimum element of π , and the image of β contains all the elements less than the first element of π .*

Proof. The proof will show that the proposition is true with $B = (k + 2)(b + 1)/2$.

Let $\pi \in \mathcal{I}_2$ be given. Then there is a decomposition of π into a pair of intertwined staircases which is illustrated in Figure 3. In this decomposition consider the staircase that begins with the block λ_1 which consists of all the elements preceding the least element of π . If this staircase has fewer than k blocks then π is a k -bounded merge of two permutations having the requisite properties. So, suppose that at least k blocks occur in this staircase.

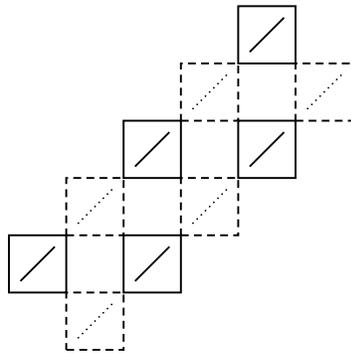


Figure 3: A general picture of intertwined staircases. The solid blocks represent λ_1, λ_2 etc.

Label the elements of these blocks in the following way:

- The elements of λ_1 are labeled with their values.
- For even $i > 1$, each element of λ_i is labeled with the largest label of an element of λ_{i-1} of smaller value.
- For odd $i > 1$, each element of λ_i is labeled with the largest label of an element of λ_{i-1} to its left.

Note that, within each block, if a label occurs in that block, then it labels an interval of elements in the block; and that together with all the elements of the preceding block

sharing the same label we obtain an interval by position or value within π according to whether the block is of odd or even index.

Our first claim is that if at least b labels occur in λ_k , then π contains a (k, b) -generic staircase. This is clear enough: simply choose a set of b labels that occur in λ_k and then, for each chosen label, in each λ_i for $1 \leq i \leq k$ take the first element carrying that label. The pattern of these elements is that of a (k, b) -generic staircase.

So, we assume henceforth that the set L of labels occurring in λ_k has fewer than b elements. Let C be its complement (in the set of labels occurring in λ_1). We claim that if we take λ to consist of all elements with labels in C together with all the elements of λ_1 , and take β to be the remaining elements of π , then the number of alternations between λ and β in the resulting merge is bounded by a function of k and b (independent of π).

Consider the elements of λ_1 through λ_k whose labels come from C (there are of course none in λ_k). They define a certain set of intervals by value and by position in π . If $x, y \in \lambda_i$ lie in different intervals, then they are separated by an element whose label is in L . Thus, using the note following the definition of labeling, the elements of C belonging to a vertical pair of blocks (λ_{2i} and λ_{2i+1}) project onto at most $|L| + 1$ intervals by position. Similarly, the elements of C in a horizontal pair of blocks project onto at most $|L| + 1$ intervals by value. So, within π the number of intervals determined by the elements with labels from C is bounded above by $k(|L| + 1)/2$ (whether we consider intervals by position or by value). Now add to this set of elements the remaining $|L|$ elements of λ_1 . This might increase the number of intervals by value, but not by more than the number of elements added. If anything, it decreases the number of intervals by position (since the entire block λ_1 is now included which forms a single interval by position). So, λ_1 together with elements whose labels come from C determine at most $k(b + 1)/2 + b$ intervals either by position or value. We set λ to be the pattern of this part, β the pattern of the remainder of π and then their merge has at most $1 + k(b + 1)/2 + b$ type changes. \square

We have all the tools required to prove Theorem 5 at this point, but it will still be helpful to approach it gently. The following proposition is not technically required in the main proof, but isolates half of the argument and, we hope, will make it easier to follow the full proof. It is also included for historical accuracy – this result was proved before the significance of rigid permutations in the main result was understood.

Proposition 8. *Let $X \subseteq \mathcal{I}_2$, $\beta \in \mathcal{I}_2 \cap \text{Av}(X)$ and suppose that $\mathcal{C} = \mathcal{I}_2 \cap \text{Av}(X) \cap \text{Av}(\beta)$ is an infinite class. Then, the growth rates of \mathcal{C} and $\mathcal{C}' = \mathcal{I}_2 \cap \text{Av}(X) \cap \text{Av}(1 \oplus \beta)$ are the same.*

Proof. Since $\mathcal{C} \subseteq \mathcal{C}'$ it is sufficient to show that $\mathcal{C}' \setminus \mathcal{C}$ is contained in some class (or indeed any set) whose growth rate is not greater than that of \mathcal{C} . So, let $\pi \in \mathcal{C}' \setminus \mathcal{C}$ be given. If π begins with its minimum, then it belongs to the class $\mathcal{C} \cup (1 \oplus \mathcal{C})$ and this class has the same growth rate as \mathcal{C} does. Otherwise, since π avoids $1 \oplus \beta$, and hence also some generic staircase, it must by Proposition 7 be a bounded merge of two permutations each avoiding $1 \oplus \beta$ and each beginning with their minimum elements. Since these permutations avoid $1 \oplus \beta$, their patterns after the first element must avoid β . So, in any case, π belongs to a bounded merge of the class $1 \oplus \mathcal{C}$ with itself. Thus $s(\mathcal{C}) = s(\mathcal{C}')$ as claimed. \square

Now we extend this proposition to a form from which Theorem 5 will follow by an easy inductive argument.

Proposition 9. *Let $X \subseteq \mathcal{I}_2$, $\alpha, \beta \in \mathcal{I}_2$ and suppose that α is 2-rigid, $\alpha \oplus \beta \in \mathcal{I}_2 \cap \text{Av}(X)$ and $\mathcal{C} = \mathcal{I}_2 \cap \text{Av}(X) \cap \text{Av}(\alpha \oplus \beta)$ is an infinite class. Then, the growth rates of \mathcal{C} and $\mathcal{C}' = \mathcal{I}_2 \cap \text{Av}(X) \cap \text{Av}(\alpha \oplus 1 \oplus \beta)$ are the same.*

Proof. We proceed as in the previous proposition. Let $\pi \in \mathcal{C}' \setminus \mathcal{C}$. Since π contains an embedded copy of $\alpha \oplus \beta$, it contains such a copy in which the α pattern is witnessed by the leftmost-bottommost copy of α in π (whose existence is assured by Corollary 3). The general disposition of π is then as shown in Figure 4.

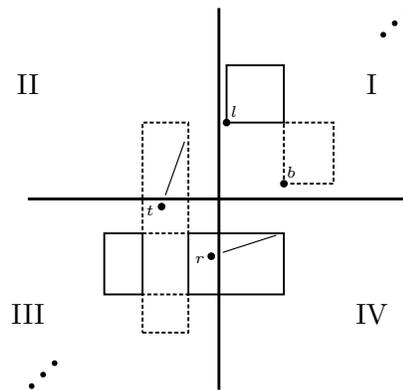


Figure 4: The structure of π containing $\alpha \oplus \beta$. The leftmost-bottommost α is contained in quadrant III. Its maximum is the element t and its rightmost element r . Quadrant I with leftmost element l and minimum b , contains a copy of β . All of π can be represented as a bounded merge of two permutations, one part of which contains the solid boxes and the other the dotted boxes.

If quadrant I began with its minimum element, we could represent π as the merge of two permutations – that singleton element, and the rest. Those remaining elements would have to avoid the pattern $\alpha \oplus \beta$ as otherwise using the leftmost-bottommost α , the singleton element, and any copy of β which is part of an $\alpha \oplus \beta$ we would have $\alpha \oplus 1 \oplus \beta \preceq \pi$. So we may assume that the situation is as shown in the figure, that is that the leftmost element of quadrant I (marked l) and its minimum (marked b) are distinct.

As before, we can decompose quadrant I containing β into a pair of intertwined staircases, and thus represent it as a bounded merge of two permutations (since it must avoid $1 \oplus \beta$ and hence some generic staircase). The remainder of the permutation consists of the part in quadrant III bounded by the topmost and rightmost points of the copy of α , together with two increasing segments (either or both of which may be empty) in quadrants II and IV as shown. This subset of π must avoid $\alpha \oplus 1 \oplus \beta$ and so can also be written as a bounded merge of two permutations, one containing the solid rectangle to which r belongs, and the other the dotted rectangle to which t belongs, as shown in the figure. Here we use Proposition 7 applied to the pattern of these elements obtained by a 180° degree rotation of the graph.

These two bounded merges can be combined into a single bounded merge which represents the entire permutation π . We will now show that neither of the components of this merge contains a copy of $\alpha \oplus \beta$. Suppose, for the sake of argument, that the component, σ , represented by the solid boxes contained this pattern, on a subset θ containing the leftmost-bottommost copy of α in σ . The leftmost-bottommost copy of α in σ would extend strictly above the leftmost-bottommost copy of α in π , since σ does not contain the topmost element (t) of the leftmost-bottommost copy of α in π . So, the copy, β' , of β in θ lying above this copy of α could not include the leftmost element (l) of quadrant I; as all the elements of π larger than t either lie in the other part of the merge, or properly within quadrant I. Therefore, β' lies strictly above and to the right of l . However, α' , the leftmost-bottommost copy of α in π lies strictly below and to the left of l . In that case the pattern of $\alpha' \cup \{l\} \cup \beta'$ is $\alpha \oplus 1 \oplus \beta$, providing a contradiction as π avoids this permutation. The argument that the other part of the merge cannot contain $\alpha \oplus \beta$ is similar.

Hence, any element of $\mathcal{C}' \setminus \mathcal{C}$ is a bounded merge of two permutations in \mathcal{C} and thus the growth rates of \mathcal{C}' and \mathcal{C} are the same. \square

Now finally:

Proof of Theorem 5. Without loss in generality we may assume that $X \subseteq \mathcal{I}_2$. Furthermore we may assume that X does not contain any increasing permutation and so the class $\mathcal{I}_2 \cap \text{Av}(X)$ is infinite (the result is, of course, trivial if this class is finite). If $\text{red}(X) = X$ there is nothing to prove. Otherwise, X contains at least one permutation, τ , having an articulation point. Write $\tau = \alpha \oplus 1 \oplus \beta$ where α is either rigid or empty (that is, decompose τ around its first articulation point). Let $\tau' = \alpha \oplus \beta$ and $X' = (X \setminus \{\tau\}) \cup \{\tau'\}$. Then, by one of the two preceding propositions

$$s(\mathcal{I}_2 \cap \text{Av}(X)) = s(\mathcal{I}_2 \cap \text{Av}(X')).$$

After a series of such reductions (formally, by induction on the number of articulation points occurring among the elements of X) we obtain the desired conclusion. \square

4 The lattice of embeddings of 21 in an element of \mathcal{I}_2

Theorem 2 showed that the embeddings of a k -rigid permutation ρ into an element of \mathcal{I}_k form a distributive lattice. The case $k = 2$, and $\rho = 21$ is particularly interesting. The union of the images of 21 in a permutation $\pi \in \mathcal{I}_2$ forms exactly the rigid reduction of π , so we interest ourselves only in the case where π is 2-rigid, and we set L_π to be the distributive lattice of copies of 21 in π . Restricting further, we consider as fixed the number, m , of rank 2 elements in π and also the number, n of rank 1 elements, and we represent these by the chain $[m] = \{1, 2, 3, \dots, m\}$ and $[n] = \{1, 2, 3, \dots, n\}$ respectively. We suppress a necessary distinction between these chains according to the rank of the corresponding

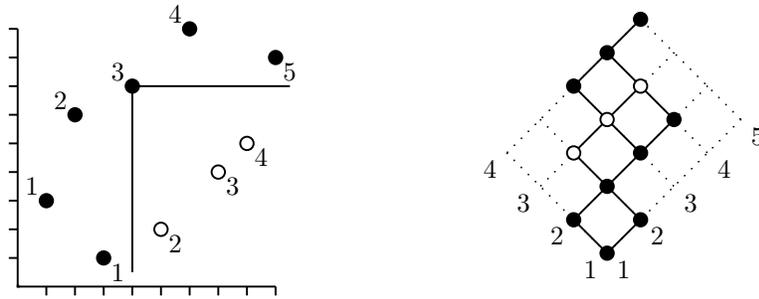


Figure 5: The permutation $\pi = 361729458$ and its corresponding lattice, with the interval $D(3) = \{2, 3, 4\}$ of rank 1 points highlighted.

elements, since this is always clear from context. Then L_π forms a sublattice of $[m] \times [n]$, where $(i, j) \in L_\pi$ if and only if the i^{th} element of rank 2 and the j^{th} element of rank 1 form a 21-pattern. In particular, if $\pi = (n + 1) \cdots (n + m)1 \cdots n$, then $L_\pi = [m] \times [n]$. Another example is shown in Figure 5.

Recall that if A and B are algebraic structures, then a subalgebra $C \leq A \times B$ is called a *subdirect product* of A and B if the projections from C to A and to B are both surjective. The lattice L_π is always a subdirect product of $[m]$ and $[n]$ since every element is part of some 21. Also it is clear that if $\pi \neq \pi'$, then $L_\pi \neq L_{\pi'}$, since all the order relationships of π are determined by L_π .

Now suppose that K is an arbitrary subdirect product of $[m]$ and $[n]$. For $a \in [m]$ define $D_K(a) = \{p \in [n] : (a, p) \in K\}$. The following observation is certainly folkloric:

Observation 10. *For all $a \in [m]$, $D_K(a)$ is a non-empty interval in $[n]$. Furthermore if $a, b \in [m]$ with $a < b$ then $\min D_K(a) \leq \min D_K(b)$ and $\max D_K(a) \leq \max D_K(b)$.*

Proof. For the first part, suppose that $p \leq q \leq r$ and $p, r \in D_K(a)$. Then, $(b, q) \in K$ for some b , because K is subdirect. If $b \leq a$ then $(a, q) = (b, q) \vee (a, p)$, while if $a < b$ then $(a, q) = (b, q) \wedge (a, r)$. In either case, $q \in D_K(a)$. The second part is immediate as well, for if $(a, p) \in K$ and $(b, q) \in K$ then $(a, p \wedge q), (b, p \vee q) \in K$. \square

Using this observation we can construct, given K , a permutation $\Pi(K)$ as follows: begin with an increasing sequence of length n . Now, for $a \in [m]$ place a new element just to the left of $\min D(a)$ and just above $\max D(a)$ (and also above all previously placed elements of this sort). The conditions of the observation guarantee that such a placement is always possible. It is also clear that $L_{\Pi(K)} = K$. Thus we obtain:

Theorem 11. *The 2-rigid elements of \mathcal{I}_2 having m elements of rank 2 and n elements of rank 1 are in one-to-one correspondence with the subdirect products of $[m]$ and $[n]$.*

Proof. We have noted that the association $\pi \mapsto L_\pi$ is both one-to-one and onto. \square

Releasing the restrictions on m and n we see that every subdirect product of two finite chains is equal to L_π for a unique 2-rigid permutation $\pi \in \mathcal{I}_2$. However, for 3-rigid

permutations in \mathcal{I}_3 no such result holds. For example, there are 29 subdirect products of three chains of length 2, but only 25 permutations that are 3-rigid of size 6 with 2 elements of each rank. In fact, even among these permutations there are duplications in their corresponding lattices.

A permutation is 2-rigid if it is covered by its embedded copies of 21. We noted above that we could count the number of 2-rigid permutations in \mathcal{I}_2 and we might well consider what we can say about permutations satisfying some stronger conditions. For example, we might call $\pi \in \mathcal{I}_2$ *k-good* if every point of π lies in a copy of $\iota_k \ominus \iota_k$. Thus, a 1-good permutation is 2-rigid, and vice versa. We do not have a complete enumeration of this collection of permutations, but the following result is amusing:

Lemma 12. *There are $\binom{2\ell}{\ell}$ k -good permutations of length $2k + \ell$ for $0 \leq \ell \leq k$.*

Proof. Let a_j denote the number of k -good permutations of length $2k + \ell$ for which there are $k + j$ points lying on the upper line (and subsequently $k + \ell - j$ on the lower). Note first that $a_j = 0$ for every $j > \ell$, as then there are fewer than k points on the lower line. Thus we need only consider values of j satisfying $0 \leq j \leq \ell$.

Supposing π is such a permutation, divide each line into three sections: from left to right, the upper line is divided into (possibly empty) parts of sizes j , $k - j$ and j , and the lower into $\ell - j$, $k - \ell + j$ and $\ell - j$. Note that the condition $\ell \leq k$ ensures that this division is possible. Since π is k -good, the middle sections of each line (of sizes $k - j$ and $k - \ell + j$) cannot interact: the leftmost k points of each of the upper and lower lines must together form a copy of $\iota_k \ominus \iota_k$, and so the middle section of each line cannot interact with the leftmost section of the other. Similarly, the rightmost k points of each line must also form an $\iota_k \ominus \iota_k$, and hence the middle section of each line cannot interact with the rightmost section of the other. Trivially, these two conditions also prevent the middle sections from interacting with each other.

Thus a_j counts the number of ways of simultaneously interleaving the rightmost part of the upper line with the leftmost part of the lower line vertically, and the leftmost of the upper with the rightmost of the lower horizontally. Up to symmetry these two interleavings are the same, so we consider only the former. Note that these two sections contain a total of ℓ points, and so there are $\binom{\ell}{j}$ possible interleavings. Hence $a_j = \binom{\ell}{j}^2$, and so there are $\sum_{j=0}^{\ell} \binom{\ell}{j}^2 = \binom{2\ell}{\ell}$ such permutations. \square

It is worth noting that there are also $\binom{2\ell}{\ell}$ k -good permutations of length $2k + \ell$ when $\ell = k + 1$: the argument in the proof of Lemma 12 still works for j satisfying $1 \leq j \leq k$ (i.e. $a_j = \binom{\ell}{j}^2$). When $j = 0$, the upper line contains exactly k points and there is only one such k -good permutation of each length of this form, giving $a_0 = 1$. Similarly, when $j = \ell = k + 1$ the lower line contains exactly k points, and again we always have $a_{k+1} = 1$.

5 Further remarks

We have been unable to extend the main result of Section 3 to apply to the classes \mathcal{I}_k with $k \geq 3$. This is largely because there seems to be no analog to the “generic staircase”

which we require in order to obtain bounded merges. Indeed, Waton's doctoral thesis [11] points to a fundamental difference between \mathcal{I}_2 and \mathcal{I}_3 . He considered their subclasses from a combinatorial-geometric point of view. In his work \mathcal{I}_2 arises as the set of all permutations drawn on two fixed arbitrary parallel lines. By way of contrast, permutations drawn on three parallel lines form a proper subclass of \mathcal{I}_3 , and there are uncountably many such classes, depending on the relative position of the three lines. Despite this we have managed to prove a weaker form of the result (generalizing an unpublished observation of M. Bóna):

Proposition 13. *For any k , α and β , and set of permutations X , the growth rates of $\mathcal{I}_k \cap \text{Av}(X, \alpha \oplus 1 \oplus \beta)$ and $\mathcal{I}_k \cap \text{Av}(X, \alpha \oplus 1 \oplus 1 \oplus \beta)$ are the same.*

Proof. As usual, consider those $\pi \in \mathcal{I}_k$ which avoid $\alpha \oplus 1 \oplus 1 \oplus \beta$ but involve $\alpha \oplus 1 \oplus \beta$. Consider all the elements x of π which have an α below and to their left, and a β above and to their right. No two of these can form a 12 pattern or else we would obtain a copy of $\alpha \oplus 1 \oplus 1 \oplus \beta$. Thus they form a descending chain, but in particular there can be at most k of them. So π is a bounded merge of a permutation avoiding $\alpha \oplus 1 \oplus \beta$ (as well as δ_{k+1}) and a permutation of length at most k , which is all that we require. \square

Applying this proposition repeatedly we can partially reduce the elements of any basis set of a class of this type without changing its growth rate, where by *partial reduction* we mean replacing multiple consecutive articulation points by a single one.

As is well known, the class \mathcal{I}_2 is enumerated by the Catalan numbers. If we denote its generating function by c , and let r denote the generating function of the rigid permutations in \mathcal{I}_2 (including the empty permutation), then the decomposition of an arbitrary $\pi \in \mathcal{I}_2$ used to define the rigid reduction shows that:

$$c = \frac{r}{1 - tr}.$$

Therefore,

$$\begin{aligned} c &= \frac{1 - \sqrt{1 - 4t}}{2t} \\ r &= \frac{1 + 2t - \sqrt{1 - 4t}}{2t(t + 2)}. \end{aligned}$$

Then, the elementary estimates referred to in Example **IV.2** (page 228) of [6] applied to both c and to r yield:

Proposition 14. *Asymptotically, 4/9 of the permutations in \mathcal{I}_2 are 2-rigid.*

This provides rather slim grounds on which to make the following:

Conjecture 15. *Asymptotically, a positive proportion of the permutations in \mathcal{I}_k are k -rigid.*

References

- [1] ALBERT, M., AND ATKINSON, M. Simple permutations, partial well-order, and enumeration. In *Proceedings of Permutation Patterns 2003* (2003), pp. 5–9.
- [2] ATKINSON, M. D., MURPHY, M. M., AND RUŠKUC, N. Partially well-ordered closed sets of permutations. *Order* 19, 2 (2002), 101–113.
- [3] BÓNA, M. The limit of a Stanley-Wilf sequence is not always rational, and layered patterns beat monotone patterns. *J. Combin. Theory Ser. A* 110, 2 (2005), 223–235.
- [4] BÓNA, M. New records in Stanley-Wilf limits. *European J. Combin.* 28, 1 (2007), 75–85.
- [5] ERDŐS, P., AND SZEKERES, G. A combinatorial problem in geometry. *Compos. Math.* 2 (1935), 463–470.
- [6] FLAJOLET, P., AND SEDGEWICK, R. *Analytic Combinatorics*. Cambridge University Press, 2009.
- [7] KNUTH, D. E. *The art of computer programming*, second ed. Addison-Wesley Publishing Co., Reading, Mass.-London-Amsterdam, 1975. Volume 1: Fundamental algorithms, Addison-Wesley Series in Computer Science and Information Processing.
- [8] MARCUS, A., AND TARDOS, G. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A* 107, 1 (2004), 153–160.
- [9] REGEV, A. Asymptotic values for degrees associated with strips of young diagrams. *Adv. Math.* 41 (1981), 115–136.
- [10] VATTER, V. Permutation classes of every growth rate above 2.48188. *Mathematika* 56 (2010), 182–192.
- [11] WATON, S. D. *On permutation classes defined by token passing networks, gridding matrices and pictures: three flavours of involvement*. PhD thesis, University of St. Andrews, 2007.