Sorting permutations with pattern-avoiding machines

Università di Firenze – Università di Perugia – INdAM – CIAFM

DOTTORATO DI RICERCA
IN MATEMATICA, INFORMATICA, STATISTICA

CURRICULUM IN INFORMATICA
CICLO XXXIII

Sede amministrativa: Università degli Studi di Firenze

Coordinatore Prof. Paolo Salani

Sorting permutations with pattern-avoiding machines

Settore Scientifico Disciplinare INF/01

Tutore
Prof. Luca Ferrari

Dottorando
Dott. Giulio Cerbai

Coordinatore
Prof. Paolo Salani

Anni 2017/2020

A mio babbo

Acknowledgements

First, I thank my advisor Luca Ferrari, for being the kindest and most helpful guide I could have asked for. His office door was always open, and the hours we spent discussing combinatorics (and everything else) are countless. Regardless how busy he was, he managed to carve out some time to answer my questions and hear my thoughts. I owe him all my academic results: it is not much, but I hope more will come in the future.

I thank Anders Claesson for raising the question upon which this entire work of thesis is built, and for his precious teachings and insight. I thank him for having invited me in Iceland for a month. He shared his office and treated me as an equal from the first moment, although the difference in our understanding and knowledge was immense.

I thank Luca Ferrari, Antonio Bernini, Elena Barcucci and Renzo Pinzani for welcoming me in the combinatorics family in Firenze.

I thank all the people I've worked with during my PhD, as well as those that were kind enough to share their knowledge with me. Amongst them, Einar Steingrímsson, not only for his inspiring lessons and stimulating discussions, but also for his introduction to the study of patterns on ascent sequences. Christian Bean, for providing a lot of useful data. Jean-Luc Baril, Carine Khalil and Vincent Vajnovszki, for the paper we wrote together.

Introduction

To characterize and enumerate permutations that can be sorted by two consecutive stacks, connected in series, is amongst the most challenging open problems in combinatorics. It is known that the set of sortable permutations is a class with an infinite basis, but the basis remains unknown. The enumeration of sortable permutations is still missing as well. Several variants and weaker formulations have been discussed in the literature, some of which are particularly meaningful. For instance, West 2-stack sortable permutations are those that can be sorted by making two passes through an *increasing* stack. To use an increasing stack means to always perform a push operation, in a greedy fashion, unless adding the next element would make the content of the stack not increasing, reading from top to bottom. As it is well known, an increasing stack is optimal for the classical problem of sorting with one stack, thus West's device is arguably the most straightforward generalization to two stacks of the orginal instance. This monotonicity requirement can be equivalently expressed by saying that the stack is 21-avoiding, again referring to the stack not being allowed to contain occurrences of the pattern 21. The present work of thesis is nothing more than an attempt to answer a question raised[1] by Claesson:

What happens if we regard the stack as σ-avoiding during the first pass, for some pattern σ, and 21-avoiding during the second pass?

The resulting device is called the *σ-machine*. A σ-machine is the natural generalization of West's device obtained by replacing 21 with any pattern σ, during the first pass, and using the optimal algortithm, during the second pass. This approach can be pushed even further by replacing σ with a set of patterns Σ, obtaining a family of sorting devices which we call *pattern-avoiding machines*.

We devote this entire manuscript to the study of pattern-avoiding machines, aiming to gain a better understanding of sorting with two consecutive stacks. For specific choices of σ, we characterize and enumerate the set of permutations

[1] As a follow-up question of a related talk given by Ferrari and the current author at Permutation Patterns 2018.

that are sortable by the σ-machine, which we call *σ-sortable*. The combinatorics underlying σ-machines and σ-sortable permutations is extremely rich. It often displays geometric structure and reveals links with a great deal of discrete objects. Certain patterns are particularly relevant. For instance, by setting $\sigma = 21$ we get West's device, while the 12-machine is similar to a device studied by Smith (but uses a different sorting algorithm). In order to sort the largest amount of permutations, a sensible try is to set $\sigma = 231$: indeed a 231-stack constantly aims to prevent the output of occurrences of 231, which is the well known requirement that the input of a classical (increasing) stack has to satisfy in order to be sortable.

The thesis has the following structure:

In Chapter 3 we provide results that cover families of patterns by analyzing how the choice of σ affects the structure of σ-sortable permutations. The most relevant (and maybe surprising) is the proof that sets of σ-sortable permutations that are not permutation classes are enumerated by the Catalan numbers, contained in Chapter 3. If σ-sortable permutations form a class, it is the set of permutations avoiding 132 and the reverse of σ. On the other hand, sets that are not classes are really hard to characterize and enumerate. Amongst them, the only patterns we are able to solve are 123 and 132. A description of 231-sortable permutations remains unknown, as well as their enumeration. We then prove that σ-sortable permutations avoid a bivincular pattern ξ of length three, unless σ is the skew-sum of 12 minus a 231-avoiding permutation.

The pattern 123 is solved in Chapter 4. We provide a step by step construction of 123-sortable permutations that leads to a bijection with a class of pattern-avoiding Schröder paths, whose enumeration is known.

The pattern 132 is solved in Chapter 5. We first characterize 132-sortable permutations as those avoiding 2314 and a mesh pattern of length three. The obtained description is then exploited to determine their geometric structure. More precisely, we show that 132-sortable permutations satisfy specific geometric constraints in the grid decomposition induced by left-to-right minima. Finally, we define a bijection with a class of pattern-avoiding set partitions, encoded as restricted growth functions.

In Chapter 6 we discuss a variant of pattern-avoiding machines where the first stack is (σ, τ)-avoiding, for a pair of patterns σ and τ. We solve several pairs of patterns of length three. We then determine an infinite family of pairs where sortable permutations are enumerated by the Catalan numbers.

In Chapter 7 we analyze some dynamical aspects of the σ-stack operator. We introduce the notions of *σ-sorted permutations*, *σ-fertility* and *effective* pattern. We characterize effective patterns, then we describe σ-sorted permutations and σ-fertilities of the 123-machine.

In Chapter 8 we consider a natural generalization of σ-machines where input sequences and forbidden patterns are chosen in different sets of words, namely Cayley permutations, ascent sequences and modified ascent sequences. In each case, we characterize sets of σ-sortable sequences that are classes. By encoding the action of σ-stacks as labeled Dyck paths, we then determine for which patterns σ the σ-stack operator is bijective on the set of Cayley permutations.

Many aspects concerning pattern-avoiding machines are yet to be investigated thoroughly. Some of them, as well as some open problems and questions that motivated our research, will be mentioned at some point in the manuscript.

Contents

List of Figures

List of Tables

Notations

Symbol	Meaning	Note
Words		
$\mathbb{N}$	Natural numbers	$\mathbb{N} = 1, 2, 3, \ldots$
$[n]$	Interval $\{1, 2, \ldots, n\}$	
$\mathbb{N}^*$	Words on $\mathbb{N}$	
$\mathfrak{S}$	Permutations	
$\mathfrak{Cay}$	Cayley permutations	Sections 1.3, 8.1
$\mathfrak{RGF}$	Restricted growth functions	Section 1.3
$\mathfrak{f}$	Fishburn pattern	$\mathfrak{f} = (231, \{1\}, \{1\})$
$\mathfrak{F}$	Fishburn permutations	$\mathfrak{F} = \mathfrak{S}(\mathfrak{f})$
$\mathfrak{A}$	Ascent sequences	Sections 1.3, 8.2.1
$\mathfrak{MA}$	Modified ascent sequences	Sections 1.3, 8.2.2
X_n	Words in X having length n	For any set X and $n \geq 1$
$X(p)$	Words in X avoiding p	For any set X and pattern p
$X(A)$	Words in X avoiding patterns in A	For any sets X and A
id	Identity permutation	$\mathsf{id}_n = 12 \cdots n$
σ-machines		
$\mathcal{S}^\sigma$	σ-stack operator	
$\mathcal{S}^{21} \circ \mathcal{S}^\sigma$	σ-machine	
$\mathrm{Sort}(\sigma)$	σ-sortable permutations	
f_n^σ	Cardinality of $\mathrm{Sort}_n(\sigma)$	$F^\sigma(t) = \sum_{n \geq 0} f_n^\sigma t^n$
$\widehat{\sigma}$	$\widehat{\sigma} = \sigma_2\sigma_1\sigma_3 \cdots \sigma_k$	
$\mathrm{Sorted}(\sigma)$	σ-sorted permutations	
$\mathrm{fert}^\sigma(\pi)$	σ-fertility of π	$\mathrm{fert}^\sigma(\pi) = \lvert\, (\mathcal{S}^\sigma)^{-1}(\pi) \rvert$
Miscellanea		
$\mathcal{I}$, $\mathcal{R}$, $\mathcal{C}$	Inverse, reverse, complement	Trivial bijections on $\mathfrak{S}$
$\mathrm{Des}(\cdot)$	Set of descents	$\mathrm{des}(\cdot) = \lvert Des(\cdot) \rvert$
$\mathrm{Asc}(\cdot)$	Set of ascents	$\mathrm{asc}(\cdot) = \lvert Asc(\cdot) \rvert$
$\mathrm{LTRmin}(\cdot)$	Set of ltr-minima	$\mathrm{ltrmin}(\cdot) = \lvert \mathrm{LTRmin}(\cdot) \rvert$
$\mathrm{LTRmax}(\cdot)$	Set of ltr-maxima	$\mathrm{ltrmax}(\cdot) = \lvert \mathrm{LTRmax}(\cdot) \rvert$

Chapter 1

Preliminaries

1.1 Patterns on words and permutations

The notion of pattern avoidance plays a central role in our work. A detailed introduction on patterns on words can be found in the books [14] and [37], while the paper [12] contains a brief presentation on permutation patterns. In this section we recall some basic definitions and notations from the literature.

Let $\mathbb{N}^*$ be the set of words over the alphabet $\mathbb{N} = \{1, 2, \dots\}$ of positive integers. Let $x = x_1 \cdots x_n$ and $y = y_1 \cdots y_k$ be words in $\mathbb{N}^*$, with $k \leq n$. We say that y is a *pattern* of x if there exist indices $i_1 < i_2 < \cdots < i_k$ such that $x_{i_1} x_{i_2} \cdots x_{i_k}$ is *order isomorphic* to y, that is:

- $x_{i_s} < x_{i_t}$ if and only if $y_s < y_t$; and
- $x_{i_s} = x_{i_t}$ if and only if $y_s = y_t$,

for each pair of indices s, t. The subsequence $x_{i_1} \cdots x_{i_k}$ is an *occurrence* of y in x and we write $x_{i_1} \cdots x_{i_k} \simeq y$. If y is a pattern of x, we say that x *contains* y and we write $x \geq y$. Otherwise, we say that x *avoids* y (or x is y*-avoiding*) and write $x \not\geq y$.

Let X be a set of words. Given a pattern y, let $X(y)$ be the set of words in X that avoid y. For a set of patterns Y, denote by $X(Y)$ the set of words in X that avoid every pattern in Y. If $Y = \{y_1, \dots, y_t\}$, we write $X(y_1, \dots, y_t)$ instead of $X(\{y_1, \dots, y_t\})$. We use the notation X_n to denote the set of words of length n in X, where the length of a word is the number of letters it contains. The sets $X_n(y)$ and $X_n(Y)$ are defined accordingly.

Let $n \geq 1$ and let $[n] = \{1, \dots, n\}$. A *permutation* of length n is a rearrangement $x = x_1 \cdots x_n$ of the integers $[n]$. Sometimes we regard the empty word as the only permutation of length zero. Denote by $\mathfrak{S}$ the set of all permutations. A permutation $x = x_1 \cdots x_n$ is often represented by plotting the points $\{(i, x_i)\}_i$ in the Euclidean plane, as in Figure 1.1. The *identity* (or *increasing*) permutation of length n is $\mathsf{id}_n = 12 \cdots n$. The *anti-identity* (or *decreasing*) permutation of length n is $\mathsf{ai}_n = n \cdots 21$.

Pattern containment is a partial order on the set $\mathfrak{S}$ and the resulting poset is called the *permutation pattern poset*. A set of permutations C which is closed downwards under this partial order is said to be a *permutation class* (or simply *class*). To be closed downwards means that for any pair of permutations x, y, if $x \in C$ and $y \leq x$, then $y \in C$ too. A permutation class C is characterized completely by the minimal permutations in the complementary set $\mathfrak{S} \setminus C$. The *basis* of C is the set of minimal avoided patterns. Note that the basis of a class is an antichain due to minimality. Conversely, if B is an antichain, then $\mathfrak{S}(B)$ is a class with basis B. A class C is *finitely based* if its basis is finite. If the basis is a singleton, then C is said to be *principal*.

Example 1.1. Let $y = 231$. A permutation $x = x_1 \cdots x_n$ contains y if there are three elements x_i, x_j, x_k, with $i < j < k$, such that $x_k < x_i < x_j$. It is well known that there are $\mathfrak{c}_n$ 231-avoiding permutations of length n, where $\mathfrak{c}_n = \frac{1}{n+1}\binom{2n}{n}$ is the n-th Catalan number (sequence A000108 in [45]).

1.1.1 Generalized pattern avoidance

Pattern containment has been generalized in a variety of ways. Some notions of non-classical pattern are recalled below.

A *barred* permutation [50] is a permutation where some entries are barred. Let y be a barred permutation and let y' be the classical permutation underlying y, that is the permutation obtained from y by ignoring the bars. Let w be the permutation order isomorphic to the non-barred entries of y. For a permutation x, to avoid the barred pattern y means that every occurrence of w in x is part of a classical occurrence of y'.

A *bivincular* pattern [17] of length k is a triple (y, S, T), where y is a permutation of length k and S, T are subsets of $\{0, 1, \ldots, k\}$. An occurrence of the bivincular pattern (y, S, T) in a permutation $x = x_1 \cdots x_n$ is then a classical occurrence $x_{i_1} \cdots x_{i_k}$ of y such that:

- $i_{s+1} = i_s + 1$, for each $s \in S$;
- $j_{t+1} = j_t + 1$, for each $t \in T$,

where $\{x_{i_1}, \ldots, x_{i_k}\} = \{j_1, \ldots, j_k\}$, with $j_1 < \cdots < j_k$; by conventions, $i_0 = j_0 = 0$ and $i_{k+1} = j_{k+1} = n + 1$. The set S identifies contraints of adjacency on the positions of the elements of y, while the set T, symmetrically, identifies constraints on their values. An example of bivincular pattern is depicted in Figure 1.1.

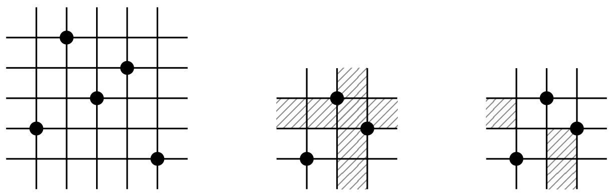

Figure 1.1: The plot of the permutation $x = 25341$, on the left. The bivincular pattern $\sigma = (132, \{2\}, \{2\})$, in the center. And the mesh pattern $\mu = (132, \{(0,2), (2,0), (2,1)\})$, on the right. Note that, for example, 253 is an occurrence of μ in x, but it is not an occurrence of σ, since the element 4 breaks the adjacency constraint between 5 and 3. On the other hand, 254 is not an occurrence of σ in π due to the element 3 falling in the shaded square $(2, 1)$.

Mesh patterns generalize both classical and bivincular patterns. A *mesh pattern* [18] of length k is a pair (y, A), where y is a permutation of length k and $A \subseteq \{0, 1, \ldots, k\} \times \{0, 1, \ldots, k\}$ is a set of pairs of integers. The elements of A identify the lower left corners of shaded squares in the plot of y. An occurrence of the mesh pattern (y, A) in the permutation x is then an occurrence of the classical pattern y in x such that no elements of x are placed into a shaded square of A. An example of mesh pattern is depicted in Figure 1.1.

In analogy with classical pattern avoidance, we use the same notation $\mathfrak{S}(\sigma)$ (respectively $\mathfrak{S}_n(\sigma)$) to denote the set of permutations (respectively permutations of length n) that avoid σ, where σ is either a barred, bivincular or mesh pattern. Notice that $\mathfrak{S}(\sigma)$ is not necessarily a permutation class when σ is a non-classical pattern.

1.2 Statistics and decompositions

Let $x = x_1 \cdots x_n$ be a permutation of length n.

The *trivial bijections* on $\mathfrak{S}_n$ are inverse, reverse and complement. The *inverse* of x is its usual group theoretic inverse $\mathcal{I}(x)$. In one-line notation, $\mathcal{I}(x) = y_1 \cdots y_n$ is defined by $y_i = j$, if $x_j = i$, for each $i = 1, \ldots, n$. The *reverse* of x is $\mathcal{R}(x) = x_n \cdots x_1$. The *complement* of x is $\mathcal{C}(x) = (n + 1 - x_1) \cdots (n + 1 - x_n)$. The three trivial bijections behave well with respect to pattern involvement, that is x contains y if and only if $\mathcal{O}(x)$ contains $\mathcal{O}(y)$, for $\mathcal{O} \in \{\mathcal{I}, \mathcal{R}, \mathcal{C}\}$. As shown in [47], the trivial bijections generate the full automorphism group of the permutation pattern poset.

An *inversion* is a pair of indices (i, j) such that $i < j$ and $x_i > x_j$. Equivalently, it is an occurrence $x_i x_j$ of the pattern 21.

An *ascent* is an index $i \in \{1, \ldots, n-1\}$ such that $x_i < x_{i+1}$. If i is an ascent, we will sometimes abuse notation and say that $x_i x_{i+1}$ is an ascent. If $x_{i+1} = x_i + 1$, the ascent is said to be *consecutive*. Denote by $\mathrm{Asc}(x)$ the set of ascents of x and let $\mathrm{asc}(x) = |\mathrm{Asc}(x)|$. *Descents*, *consecutive descents*, the set $\mathrm{Des}(x)$ and the statistic $\mathrm{des}(x)$ are defined symmetrically.

An entry x_i is a *left-to-right minimum* (briefly *ltr-minimum*) of x if $x_i < x_j$, for each $j < i$. The set of ltr-minima of x is denoted with $\mathrm{LTRmin}(x)$ and $\mathrm{ltrmin}(x) = |\mathrm{LTRmin}(x)|$. The *left-to-right minima decomposition* (briefly *ltr-min decomposition*) of x is $x = m_1 B_1 m_2 B_2 \cdots m_t B_t$, where $t = \mathrm{ltrmin}(x)$, $\mathrm{LTRmin}(x) = \{m_1, \dots, m_t\}$ and the *block* B_i contains the entries of x that are placed strictly between m_i and m_{i+1}, for $i = 1, \dots, n-1$. The last block B_t contains the entries that follow m_t. Note that $m_t = 1$. The notion of *left-to-right maximum* (briefly *ltr-maximum*), the set $\mathrm{LTRmax}(x)$, its size $\mathrm{ltrmax}(x)$ and the *ltr-max decomposition* $x = M_1 B_1 M_2 B_2 \cdots M_t B_t$ are defined analogously. In this case, we have $M_t = n$.

Example 1.2. Let $x = 471823769$. Then $\mathrm{Asc}(x) = \{1, 3, 5, 6, 8\}$, where the only consecutive ascent is 5. The ltr-minima of x are $\mathrm{LTRmin}(x) = \{4, 1\}$, thus its ltr-min decomposition is $x = m_1 B_1 m_2 B_2$, where $m_1 = 1$, $B_1 = 7$, $m_2 = 1$ and $B_2 = 823769$.

Let $k \geq 1$. The *k-inflation of x at x_i* is the permutation of length $n + k - 1$ obtained from x by replacing x_i with the consecutive increasing sequence $x_i(x_i + 1) \cdots (x_i + k - 1)$ and suitably rescaling the remaining elements. For instance, the 3-inflation of 451**3**2 at 3 is 671**345**2.

Let $x = x_1 \cdots x_n$ and $y = y_1 \cdots y_k$. The *direct sum* of x and y is the permutation $x \oplus y = xy'$, where y' is obtained from y by adding n to each of its entries. In other words, $x \oplus y$ is the only permutation $w_1 \cdots w_{n+k}$ of length $n + k$ such that $w_i < w_j$ for each $i \leq n$ and $j \geq n+1$, $w_1 \cdots w_n$ is an occurrence of x and $w_{n+1} \cdots w_{n+k}$ is an occurrence of y. The *skew sum* $x \ominus y$ is obtained analogously by requiring $w_i > w_j$ for each $i \leq n$ and $j \geq n+1$. A permutation is *layered* if it is the direct sum of decreasing permutations. As it is well known [5], a permutation w is layered if and only $w \in \mathfrak{S}(231, 312)$ and there are 2^{n-1} layered permutations of length n, for each $n \geq 1$. Similarly, a permutation w is *co-layered* if $w \in \mathfrak{S}(132, 213)$, or, equivalently, if w is the skew sum of increasing permutations.

1.3 Sets of integer sequences

Let $x = x_1 \cdots x_n$ be a word of length n on $\mathbb{N}$.

The word x is a *Cayley permutation* if $\{x_1, \dots, x_n\} = [\max(x)]$. Equivalently, if x contains at least one copy of each integer from one to its maximum value. Denote by $\mathfrak{Cay}$ the set of Cayley permutations. Cayley permutations are sometimes called *normalized words* [27], but also surjective words[1], Fubini words or packed words. A Cayley permutation $x = x_1 \cdots x_n$ with maximum equal to k encodes an ordered set partition (*ballot*) of $[n]$ with k blocks $B_1B_2 \dots B_k$, where $i \in B_{x_i}$ for each i. Cayley permutations are enumerated, with respect to their length, by the Fubini numbers (sequence A000670 in [45]). For example, the only Cayley permutation of length one is 1, there are three Cayley permutations of length two, namely $11, 12$ and 21, and thirteen Cayley permutations of length three, which are 111, 112, 121, 122, 123, 132, 211, 212, 213, 221, 231, 312, 321. Given a word $x = x_1 \cdots x_k$ on $\mathbb{N}$, there is exactly one Cayley permutation $\mathrm{std}(x)$ that is order isomorphic to x. We call $\mathrm{std}(x)$ the *standardization* of x. The sequence $\mathrm{std}(x)$ is obtained by replacing each occurrence of the smallest integer of x with 1, each occurrence of the second smallest integer with 2 and so on. For instance, we have $\mathrm{std}(1381365) = 1251243$. Notice that, if x is a word on $\mathbb{N}$ and y is a Cayley permutation, then $x_{i_1} \cdots x_{i_k}$ is an occurrence of y in x if and only if $\mathrm{std}(x_{i_1} \cdots x_{i_k}) = y$. In other words, the set $\mathfrak{Cay}$ is the set of standardized sequences, and thus it is is the set where patterns live naturally in.

The word x is a *restricted growth function* (briefly RGF) if $x_1 = 1$ and $x_{i+1} \leq 1 + \max\{x_1, \dots, x_i\}$, for each $i \leq n-1$. Similarly to Cayley permutations, a RGF $x = x_1 \cdots x_n$ naturally encodes the partition of $[n]$ with blocks $B_1B_2 \dots B_k$, where x_i is the index of the block that contains i. Thus RGFs are enumerated by the Bell numbers (sequence A000110 in [45]). The set of RGFs is denoted by $\mathfrak{RGF}$. Note that $\mathfrak{RGF}$ is a subset of $\mathfrak{Cay}$. Pattern avoidance on RGFs was discussed in [19, 35, 44].

The word x is an *ascent sequence* if $x_1 = 1$ and $x_{i+1} \leq 2 + \mathrm{asc}\,(x_1 \cdots x_i)$, for each $i \leq n-1$. Ascent sequences were introduced[2] in [17] as an auxiliary class of objects that embodies the structure of $(2+2)$-free posets, certain chord diagrams and Fishburn permutations. Fishburn permutations are those avoiding the bivincular pattern $\mathfrak{f} = (231, \{1\}, \{1\})$ (see Figure 1.2). Roughly speaking, ascent sequences encode bijectively the so called active sites of Fishburn permutations. All these objects are enumerated by the Fishburn numbers (sequence A022493 in [45]). Denote by $\mathfrak{A}$ the set of ascent sequences. Due to the intrinsically complicated structure of ascent sequences, pattern avoidance on $\mathfrak{A}$ seems to be rather more complicated than its analogue on permutations [10, 22, 30, 40]. *Modified ascent sequences* [17] are a slightly more manageable version of $\mathfrak{A}$. Let $x = x_1 \cdots x_n$ be an ascent sequence and let $\mathrm{Asc}(x) = \{i_1, \dots, i_k\}$. The modified ascent sequence of x is obtained as follows. For $j = 1, \dots, k$, increase by one each entry that precedes position i_j and is greater than or equal to x_{i_j+1}. The resulting sequence $\omega(x)$

[1] Cayley permutations encode surjective endofunctions $[n] \mapsto [k]$.

[2] In the original work ascent sequences are 0-based, that is $x_1 = 0$ and $x_{i+1} \leq 1 + \mathrm{asc}\,(x_1 \cdots x_i)$.

$\mathfrak{f} =$ 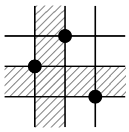

$$\begin{tikzcd} \mathfrak{A} \arrow[r, "\phi"] \arrow[d, "\omega"'] & \mathfrak{F} \\ \mathfrak{MA} \arrow[ur, "\phi'"'] & \end{tikzcd}$$

Figure 1.2: The pattern $\mathfrak{f}$, on the left. How the bijections ϕ, ω and ϕ' are related, on the right.

is the modified sequence of x. Let $\mathfrak{MA}$ be the set of modified ascent sequences. This procedure can be easily inverted, thus the map $\omega : \mathfrak{A} \to \mathfrak{MA}$ defined this way is a size-preserving bijection between $\mathfrak{A}$ and $\mathfrak{MA}$. The maps ϕ and ω, as well as the bijection $\phi' : \mathfrak{MA} \to \mathfrak{F}$ obtained by composition, are depicted in Figure 1.2. A recursive construction of $\mathfrak{MA}$ can be found in [22]: There is exactly one modified ascent sequence of length one, namely the single letter word 1. If $n \geq 2$, then every $x \in \mathfrak{MA}_n$ is of one of two forms, depending on whether its last letter a forms an ascent with the penultimate letter:

- $x = ya$, with $1 \leq a \leq b$, or
- $x = y'a$, with $b < a \leq 2 + \mathrm{asc}(y)$,

where $y \in \mathfrak{MA}_{n-1}$, b is the last letter of y, and y' is obtained from y by increasing by one each entry that is less than or equal to a. One of the main advantages of working with modified sequences is that $\mathfrak{MA}$ is a subset of $\mathfrak{Cay}$, contrary to $\mathfrak{A}$. For example, the sequence $x = 12124$ is an ascent sequence, but not a Cayley permutation. Its modified sequence is $\omega(x) = 13124$, which is indeed an element of $\mathfrak{Cay}$. Some more tools and notions on (modified) ascent sequences, as well as a characterization of $\mathfrak{MA}$ as a subset of $\mathfrak{Cay}$ (in terms of patterns), will be provided in Section 8.2.1.

In all the above sets, we either include or rule out the empty word (of length zero), according to the situation. Moreover, when possible, we extend the notions introduced in Section 1.1 and Section 1.2 accordingly.

1.4 Lattice paths

Lattice paths are amongst the most studied combinatorial objects due to the huge amount of combinatorial issues they can model. They are well suited to be studied with the elegant symbolic approach, which often leads directly to enumerative results. In this thesis, we always consider lattice paths in the first quadrant of the discrete plane $\mathbb{Z} \times \mathbb{Z}$, starting at the origin and ending on the x-axis. The *length* of a lattice path is its final abscissa. A lattice path is encoded by the word that records its steps, going from left to right. A *labeled* lattice path is a path where each step has a label. According to what kind of steps are allowed, we obtain several families of lattice paths.

A *Dyck path* is a lattice path that uses two kinds of steps (of length one), namely up steps $\mathtt{U} = (+1, +1)$ and down steps $\mathtt{D} = (+1, -1)$. The *height* of a step is its final ordinate. The *height* of a Dyck path is the maximum height of its steps. Observe that in a Dyck path the number of $\mathtt{U}$ steps matches the number of $\mathtt{D}$ steps, since the path ends on the x-axis. Moreover, in each prefix the number of $\mathtt{U}$ steps is at least equal to the number of $\mathtt{D}$ steps (due to the path never falling below the x-axis). These two properties characterize the set of words on $\{\mathtt{U}, \mathtt{D}\}$ that encode Dyck paths. For each up step $\mathtt{U}$, there is a unique *matching* down step $\mathtt{D}$ defined as the first $\mathtt{D}$ step after $\mathtt{U}$ that has height one less than $\mathtt{U}$. Since each step has length one, the length of a Dyck path is equal to the total number of its steps, which is two times the number of $\mathtt{D}$ (or equivalently $\mathtt{U}$) steps. Given a Dyck path P of semilength n, the *reverse* path of P is the path $\mathcal{R}(P)$ obtained from P by taking the symmetric path with respect to the vertical line $x = n$. Equivalently, the path $\mathcal{R}(P)$ is encoded by the word obtained by taking the reverse of the word that encodes P and then transforming each $\mathtt{U}$ in $\mathtt{D}$ and viceversa. It is well known that Dyck paths, according to their semilength, are enumerated by the Catalan numbers. An example of Dyck path is depicted in Figure 1.3. The set of Dyck paths of semilength n is denoted by $\mathcal{D}_n$, while $\mathcal{D}$ denotes the set of all Dyck paths.

Remark 1.1. It is well known that any non-empty Dyck path P has a unique decomposition $P = \mathtt{U}Q_1\mathtt{D}Q_2$, where Q_1 and Q_2 are two (possibly empty) Dyck paths. Since the $\mathtt{D}$ step that follows Q_1 is the first return on the x-axis, this is called the *first-return decomposition* of P.

A *Motzkin path* is defined exactly like a Dyck path, except that one additional kind of step is allowed: the horizontal step $\mathtt{H} = (1, 0)$. Again the length of a Motzkin path is equal to the total number of its steps. Motzkin paths, according to their length, are enumerated by the Motzkin numbers (sequence A001006 in [45]). The sets of Motzkin paths and Motzkin paths of length n are denoted by $\mathcal{M}$ and $\mathcal{M}_n$, respectively. If we allow horizontal steps $\mathtt{H}_2 = (2, 0)$ of length two, instead of $\mathtt{H}$, we obtain *Schröder paths*. The length of a Schröder path is then the sum of the number of its up steps and down steps plus twice the number of its double horizontal steps. Schröder paths are enumerated by the large Schröder numbers (sequence A006318 in [45]).

Pattern containment can be naturally extended to lattice paths by considering the words that encode them. The arising notion of classical pattern is rather dull on words on small alphabets (such as $\{\mathtt{U}, \mathtt{D}\}$). Thus we always consider *consecutive* patterns on paths, that is where occurrences of a pattern must be realized by consecutive letters. From here on, we omit the word consecutive in this context.

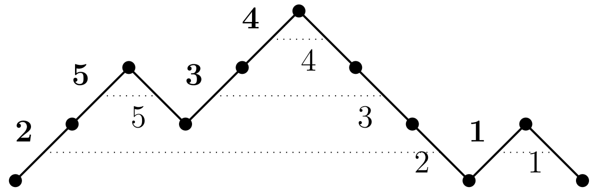

Figure 1.3: The Dyck path UUDUUDDDUD, of semilength 5 and with 3 peaks. The corresponding 213-avoiding permutation 25341 is obtained by reading the bold labels from left to right. Dotted lines connect matching steps.

A *valley* is an occurrence of the pattern DU. A *peak* is an occurrence of UD. In a Dyck path, the number of peaks is equal to one plus the number of valleys. A *double rise* is an occurrence of UU. An *ascending run* of length k is a maximal occurrence of $\mathtt{U}^k$. A *descending run* of length k is a maximal occurrence of $\mathtt{D}^k$. Given a pattern q and a set of paths X, we use the notation $X(q)$ to denote the set of paths in X that avoid the (consecutive) pattern q. Pattern avoidance on lattice paths was studied, for instance, in [6] and [25].

Example 1.3. The following is a well known bijection between 213-avoiding permutations (of length n) and Dyck paths (of semilength n). Given a Dyck path P of semilength n, label its down steps from right to left with the integers $[n]$ in increasing way. Then assign to each up step the label of its matching down step. Finally, read the labels of the up steps from left to right. The resulting sequence is a 213-avoiding permutation and this correspondence is bijective. For instance, the Dyck path UUDUUDDDUD is mapped to the 213-avoiding permutation 25341 (see Figure 1.3). An equivalent version of the above bijection, but using 132-avoiding permutations, was given by Krattenthaler in [39].

1.5 Generating trees and succession rules

Generating trees and succession rules are very powerful tools in enumerative and bijective combinatorics. Roughly speaking, a generating tree describes a combinatorial construction for a family of discrete objects. Each node of the tree produces a certain set of children, and each child has a unique father, that is it is uniquely obtained from a node situated at the previous level. Typically, starting from the root, each level of the tree contains all the objects of the family that have a given size, which is equal to one more than the size of the objects contained in the previous level. Generating trees are well encoded by succession rules. A succession rule naturally translates the combinatorial construction into something algebraic, which can be exploited to effectively enumerate the objects it represents.

Succession rules were introduced by West in [51] and [52]. Below we recall some basic definitions, referring the interested reader to [7] and [31] for a more detailed discussion.

A *generating tree* is a rooted, labeled tree with the property that the label of each node determines the labels of its children. A generating tree is recursively encoded by a *succession rule* consisting in:

- the label of the root, and
- a set of rules (*productions*) that explain how to derive, given any node of the tree, the number of its children and their labels.

Example 1.4. To illustrate the above constructions, we show a very well known succession rule for Dyck paths (see for instance [7]). Given a Dyck path P, let k be the number of points with integer coordinates contained in the last descending run of P. Then we obtain k Dyck paths of semilength one more by inserting a new peak `UD` either before a `D` step of the last descending run or at the end of the path. The resulting paths are the children of P. Note that every Dyck path P of semilength at least two is uniquely constructed this way. Since this construction depends solely on the parameter k, we shall use that integer as label of a given Dyck path. This generation of Dyck paths is then encoded via the following succession rule:

$$\Omega : \begin{cases} (2) \\ (k) \longrightarrow (2)(3)\cdots(k)(k+1) \end{cases}$$

The root is the path `UD`, which has label (2) and children `UDUD`, with label (2), and `UUDD`, with label (3). Note that Dyck paths of semilength n are in one-to-one correspondence with nodes at level n, supposing the root is at level one. Therefore Ω is a generating rule for Dyck paths according to their semilength.

Example 1.5. The following succession rule generates Motzkin paths according to their length.

$$\Omega : \begin{cases} (1) \\ (1) \longrightarrow (2) \\ (k) \longrightarrow (1)(2)\cdots(k-1)(k+1), \quad k \geq 2. \end{cases}$$

Observe that different generating rules may encode the same family of objects and the same generating rule may encode different families. Notice also that a bijection between two combinatorial families is immediately obtained by showing that both families are generated by the same succession rule, a fact that will be used in the rest of this thesis.

Chapter 2

Stack sorting and pattern-avoiding machines

2.1 Classical stack sorting

A stack is a data structure equipped with two operations: *push*, which adds an element to the stack, at the top; and *pop*, which extracts from the stack the most recently pushed element. The problem of sorting permutations using a stack, together with its many variants, has been widely studied in the literature. The reader is referred to [13] for an extensive survey on stack-sorting disciplines. The original version was proposed by Knuth in [38]: given an input permutation π, scan its elements from left to right. Every time an element of π is scanned, either push it into the stack or pop the top element of the stack, placing it into the output. The goal is to describe and enumerate sortable permutations. To sort a permutation means to produce a sorted output, that is the identity permutation. As stated in the next lemma, which will be repeatedly used throughout the rest of this thesis, an elegant solution to the original problem can be given in terms of pattern avoidance.

Lemma 2.1. [38] Let π be a permutation. Then π is sortable using a classical stack if and only if π avoids the pattern 231.

More in general, one can sort permutations using a connected network of data structures [48]. The input permutation is scanned and its entries flow through the network, going from one device to another, according to how the network is built (or equivalently to which sorting procedure is used). In this framework, some of the most interesting questions that arise are the following.

- How to characterize those permutations that can be sorted by a given network?

- How to enumerate sortable permutations?
- What is the behavior of specific sorting algorithms?

A reasonable way to pick a specific procedure consists in imposing static constraints on a stack, for instance by restricting the set of sequences it is allowed to contain. It is also interesting to search for the most efficient sorting algorithm. An algorithm is *optimal* if it sorts every sortable permutations. In other words, if it has the same sorting power as the device used in its full generality (with a non-deterministic sorting procedure). Finding optimal strategies is often a hard task. Regarding classical stack sorting, it is very well known that there is an optimal algorithm, known as *stacksort*, which is defined by the two following key properties (see Listing 1 in Appendix D):

1. the elements in the stack are maintained in increasing order, reading from top to bottom; in other words, the stack tries to prevent an occurrence of 21 to be output. This can be expressed by saying that a classical stack is *increasing*.
2. the algorithm is *right-greedy*, meaning that it always performs a push operation, unless this violates the previous condition. The expression "right-greedy" refers to the usual (and most natural) representation of this problem, depicted in Figure 2.1.

Although the classical problem is rather simple, it becomes much harder as soon as one allows several stacks connected in series. Quite recently, Pierrot and Rossin [43] proved that the problem of deciding whether a given permutation is sortable by two stacks in series is polynomial. Nevertheless, almost every other related question remains unsolved. For example, it is known that sortable permutations form a class, but its basis is infinite [42], and still unknown. The enumeration of sortable permutations is still unknown too. In the attempt of gaining a better understanding of two stacks in series, some (simpler) variants of the problem have been considered. In his PhD thesis [49], West considered two passes through a classical (i.e. increasing) stack, which is equivalent to perform a *right-greedy* algorithm on two stacks in series. In [46], Smith considered a decreasing stack followed by an increasing stack. Smith's approach was pushed further in [21], where the authors consider many decreasing stacks, followed by an increasing one. More recently, Claesson, Ferrari and the current author [23] introduced an even more general device consisting of two stacks in series with a right-greedy procedure, where a restriction on the first stack is given in terms of pattern avoidance. The present work of thesis is dedicated to the analysis of these devices, which we call *pattern-avoiding machines*. Pattern-avoiding machines are defined formally in Section 2.2, and discussed extensively in the following chapters.

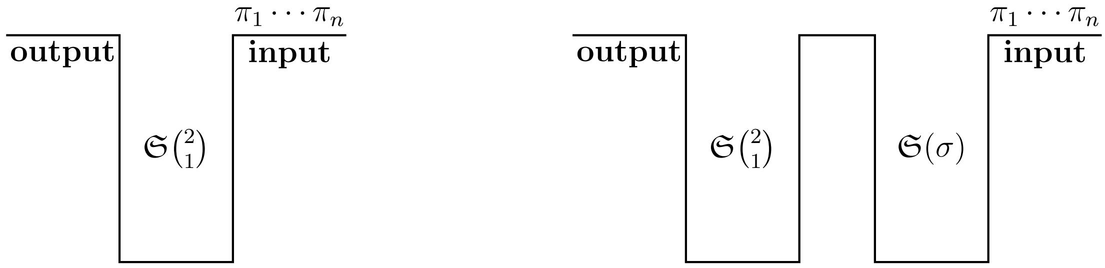

Figure 2.1: The usual representation of sorting with one stack, on the left. The σ-machine, on the right.

Other than imposing restrictions on devices and sorting algorithms, one can also allow a larger set of input sequences (see [2, 4, 27]). This line of research, on pattern-avoiding machines, is investigated in Chapter 8.

2.2 Pattern-avoiding machines

Let Σ be a set of permutations.

Definition 2.1. A Σ-*avoiding stack* (or simply Σ-stack) is a stack that is not allowed to contain an occurrence of the pattern σ, reading from top to bottom, for each $\sigma \in \Sigma$.

Definition 2.2. The term Σ-*machine* refers to performing a right-greedy algorithm on two stacks in series: a Σ-stack, followed by a 21-stack.

Recall that a 21-stack is simply a stack as normally used in classical stack-sorting. Thus the Σ-machine consists in a pass through a Σ-avoiding stack, followed by a pass of the resulting output through a classical stack. From now on, we will assume that Σ does not contain the unit length permutation, since otherwise no element could be pushed in the Σ-stack. For singletons and pairs of patterns, we omit the brackets to ease notation. For example, if $\Sigma = \{\sigma\}$, we write σ-*stack* and σ-*machine*. An illustration of the σ-machine is depicted in Figure 2.1, while the corresponding algorithm is described formally in Listing 2 of Appendix D. The sequence of operations performed by the 231-stack on input 2413 is represented in Figure 2.2.

Next we introduce some tools and notations and prove some basic results regarding Σ-machines. A permutation π is Σ-*sortable* if the Σ-machine on input π yields the identity permutation. The set of Σ-sortable permutations is denoted by $\mathrm{Sort}(\Sigma)$. For $n \geq 1$, denote by f_n^{Σ} the cardinality of $\mathrm{Sort}_n(\Sigma)$, that is the

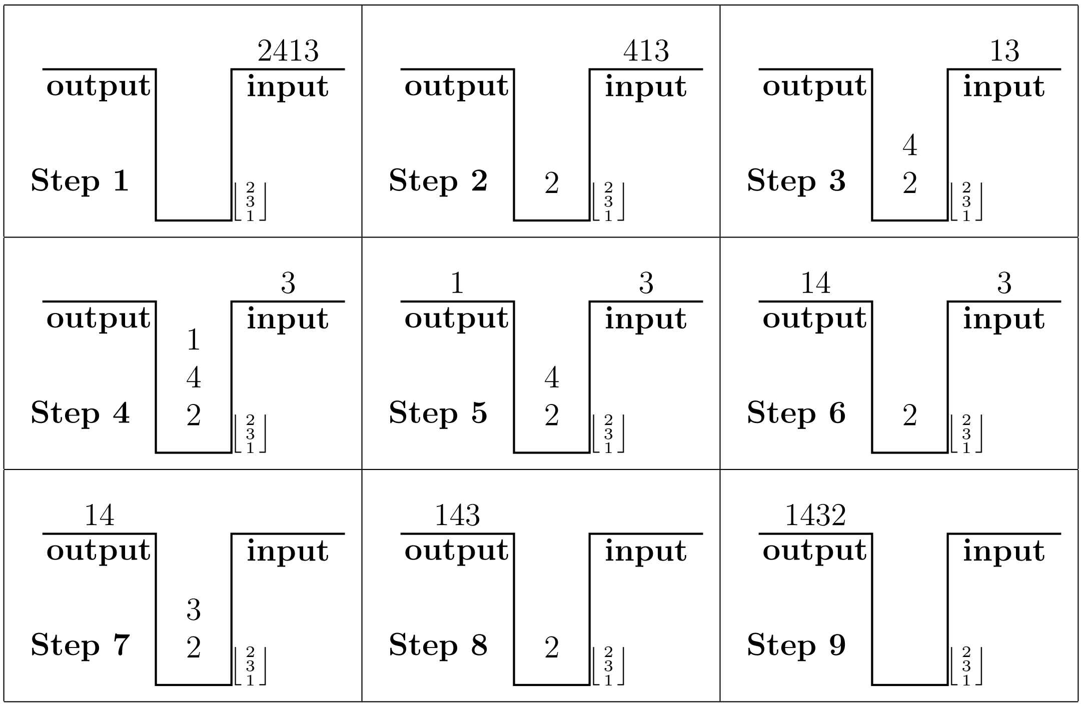

Figure 2.2: The action of the 231-stack on input 2413.

number of Σ-sortable permutations of length n. Let $F^{\Sigma}(t) = \sum_{n\geq 1} f_n^{\Sigma} t^n$ be the ordinary generating function of $\mathrm{Sort}(\Sigma)$. Given a permutation π, denote by $\mathcal{S}^{\Sigma}(\pi)$ the output of the Σ-stack on input π. Due to Lemma 2.1, since $\mathcal{S}^{\Sigma}(\pi)$ is the input of the (final) classical stack in the Σ-machine, a permutation π is Σ-sortable if and only if $\mathcal{S}^{\Sigma}(\pi)$ avoids the pattern 231. This basic fact allows us to determine the Σ-sortability of an input permutation π by simply checking whether $\mathcal{S}^{\Sigma}(\pi)$ avoids 231 or not, ignoring the final stack. We highlight this remark in the next lemma, which will be used repeatedly from now on.

Lemma 2.2. Let π be an input permutation for the Σ-machine. Then π is Σ-sortable if and only if $\mathcal{S}^{\Sigma}(\pi)$ avoids 231.

The next lemma shows that prefixes of Σ-sortable permutations are Σ-sortable.

Lemma 2.3. Let $\pi = \pi_1 \cdots \pi_n$ be a permutation of length n. Suppose that the prefix $\pi_1 \cdots \pi_k$ is an occurrence of the pattern γ, for some $1 \leq k \leq n$ and $\gamma \in \mathfrak{S}_k$. If π is Σ-sortable, then γ is Σ-sortable.

Proof. Observe that the behavior of the Σ-stack on the prefix $\pi_1 \cdots \pi_k$ does not depend on the remaining entries $\pi_{k+1} \cdots \pi_n$ of π. Moreover, it is the same as the behavior on γ, since the operations performed by the Σ-stack depend solely on the relative order of the elements processed. Therefore $\mathcal{S}^{\Sigma}(\gamma)$ avoids 231, or else $\mathcal{S}^{\Sigma}(\pi)$ would contain 231 too, contradicting the hypothesis that π is Σ-sortable. □

Lemma 2.3 suggests a recursive construction for $\mathrm{Sort}_n(\Sigma)$. Indeed every Σ-sortable permutation is obtained by appending a new rightmost element to a Σ-sortable permutation of length one less, and suitably rescaling the other elements.

As anticipated before, almost the entirety of this thesis is devoted to the analysis of Σ-machines. The combinatorics underlying these devices turns out to be unexpectedly rich and deep, offering links with other discrete objects such as lattice paths, set partitions and various families of integer sequences. Some of the questions that motivate our research are reported below (in random order).

- Given a set of patterns Σ, how to characterize the set $\mathrm{Sort}(\Sigma)$ of Σ-sortable permutations? Ideally, we wish to find geometric descriptions, recursive generations and characterizations in terms of pattern avoidance.

- Given a set of pattern Σ, what is the number of Σ-sortable permutations of length n? To answer this question, we shall exploit an eventual structural description of $\mathrm{Sort}(\Sigma)$, find a bijection with other discrete objects or provide a generating tree for $\mathrm{Sort}(\Sigma)$.

- Are there properties of Σ that allow us to determine structural information on $\mathrm{Sort}(\Sigma)$? In this framework, our main result is a characterization of the patterns σ such that $\mathrm{Sort}(\sigma)$ is a permutation class.

- Given $n \geq 1$, what is the number of Wilf-classes of σ-sortable permutations of length n? In other words, how many different counting sequences arise by considering the sets $\mathrm{Sort}(\sigma)$, for each permutation $\sigma \in \mathfrak{S}_n$?

2.3 The σ-stack

Before moving on to the study of Σ-machines, we analyze σ-stacks separately.

Lemma 2.4. Let σ be a permutation of length $k \geq 2$. Let π be a permutation of length $n \geq k-2$ and suppose that $\mathcal{S}^{\sigma}(\pi)$ is the increasing permutation. Then $\pi_1\pi_2\cdots\pi_{k-2} = n(n-1)\cdots(n-k+3)$.

Proof. Since σ has length k, the elements $\pi_1\cdots\pi_{k-2}$ are pushed directly into the σ-stack. Then they remain at the bottom of the σ-stack until the end of the sorting process and thus they are the rightmost elements of $\mathcal{S}^{\sigma}(\pi)$. □

Theorem 2.5. Let $\pi = \pi_1 \cdots \pi_n$ be a permutation of length n and let $\sigma \in \mathfrak{S}_k$, with $k \geq 2$.

1. If $\sigma = \mathsf{id}_k$, then $\mathcal{S}^\sigma(\pi)$ is the increasing permutation if and only if $\pi = \mathsf{ai}_n$ and $n \leq k-1$.

2. If $\sigma = 21 \oplus \mathsf{id}_t$, then $\mathcal{S}^\sigma(\pi)$ is the increasing permutation if and only if $\pi = \mathsf{ai}_t \ominus \alpha$, for some $\alpha \in \mathfrak{S}(231)$.

3. In all the remaining cases, $\mathcal{S}^\sigma(\pi)$ is the increasing permutation if and only if $\pi = \mathsf{ai}$.

Proof. 1. Let $\sigma = \mathsf{id}_k$. If $\pi = \mathsf{ai}_n$, for some $n \leq k-1$, then $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi) = \mathsf{id}_n$. Conversely, suppose that $\mathcal{S}^\sigma(\pi) = \mathsf{id}$. If $n \leq k-1$, then it must be $\pi = \mathsf{ai}_n$ due to Lemma 2.4. Otherwise, suppose, for a contradiction, that $n \geq k$ and write $\pi = \pi_1 \cdots \pi_{k-1}\pi_k \cdots \pi_n$. Initially, the elements $\pi_1, \ldots, \pi_{k-1}$ are pushed into the σ-stack. If $\pi_k\pi_{k-1} \cdots \pi_1 \simeq \sigma$, then $\pi_k < \pi_{k-1}$ and π_{k-1} is extracted, which is a contradiction with the hypothesis that $\mathcal{S}^\sigma(\pi)$ is increasing. On the other hand, if $\pi_k \cdots \pi_1$ is not an occurrence of σ, then π_k is pushed into the σ-stack and thus $\mathcal{S}^\sigma(\pi)$ contains the substring $\pi_k \cdots \pi_1$. But, since $\pi_k \cdots \pi_1$ is not an occurrence of $\sigma = \mathsf{id}_k$, the output $\mathcal{S}^\sigma(\pi)$ is not increasing, which contradicts the hypothesis.

2. Let $\sigma = 21 \oplus \mathsf{id}_t$. Suppose that $\mathcal{S}^\sigma(\pi)$ is the increasing permutation. By Lemma 2.4, we have $\pi_1\pi_2 \cdots \pi_t = n(n-1)\cdots(n-t+1)$. Thus $\pi = \mathsf{ai}_t \ominus \alpha$. We wish to show that α avoids 231. Observe that the elements $\pi_1 \cdots \pi_t$ are pushed into the σ-stack at the beginning of the sorting process. Then, since $\sigma = 21 \oplus \mathsf{id}_t$, the behavior of the σ-stack with $n, n-1, \ldots, n-t+1$ at the bottom is equivalent to the behavior of an empty 21-stack. Indeed the elements $n, n-1, \ldots, n-t+1$ play the role of id_t in any potential occurrence of $\sigma = 21 \oplus \mathsf{id}_t$. In other words, the σ-stack with $n, n-1, \ldots, n-t+1$ at the bottom performs a pop operation if and only if an empty 21-stack performs a pop operation. Therefore α is 21-sortable, which in turn is equivalent to α avoiding 231 by Lemma 2.1. Similarly, it is easy to prove that if $\pi = \mathsf{ai}_t \ominus \alpha$, with α a 231-avoiding permutation, then $\mathcal{S}^\sigma(\pi)$ is increasing.

3. Finally, suppose that σ is not increasing and σ is not the direct sum of 21 and the identity permutation. Since σ is not increasing, we have $\mathcal{S}^\sigma(\mathsf{ai}) = \mathcal{R}(\mathsf{ai}) = \mathsf{id}$. Conversely, suppose that $\pi \neq \mathsf{ai}$. Write $\pi = \pi_1 \cdots \pi_i\pi_{i+1} \cdots \pi_n$, where i is the leftmost ascent of π. We show that $\mathcal{S}^\sigma(\pi)$ is not increasing. Since $\sigma \neq \mathsf{id}$ and $\pi_1 > \pi_2 > \cdots > \pi_i$, the elements $\pi_1, \ldots, \pi_i$ are pushed into the σ-stack. Now, if π_{i+1} enters the σ-stack above π_i, then π_{i+1} precedes π_i in $\mathcal{S}^\sigma(\pi)$, with $\pi_{i+1} > \pi_i$, thus $\mathcal{S}^\sigma(\pi)$ is not increasing. Otherwise,

suppose that π_i is extracted from the σ-stack before π_{i+1} enters. That is, the σ-stack contains $k-1$ elements, say $\pi_{j_2}, \dots, \pi_{j_k}$, with $j_2 < \cdots < j_k$, such that $\pi_{i+1}\pi_{j_k} \cdots \pi_{j_2}$ is an occurrence of σ. Notice that $\pi_{j_k} < \cdots < \pi_{j_2}$ due to our assumptions. Without losing generality, choose the minimal indices $j_2, \dots, j_k$ such that $\pi_{i+1}\pi_{j_k} \cdots \pi_{j_2}$ is an occurrence of σ, so that π_{i+1} enters the σ-stack above $\pi_{j_{k-1}}$ (and thus π_{i+1} precedes $\pi_{j_{k-1}}$ in $\mathcal{S}^\sigma(\pi)$). Now, if $\pi_{i+1} < \pi_{j_{k-1}}$, then $\sigma = 21 \oplus \mathsf{id}_{k-2}$, which is a contradiction. Otherwise, if $\pi_{i+1} > \pi_{j_{k-1}}$, then $\mathcal{S}^\sigma(\pi)$ is not increasing, as desired.

$\square$

Chapter 3

The σ-machine

This chapter is devoted to the analysis of σ-machines. Most[1] of the results presented here are contained in [23]. Patterns σ of length two are discussed in Section 3.1. In Section 3.2 we prove the main result of this chapter, which is a characterization of those patterns σ where the set of σ-sortable permutations is a class. We prove the surprising fact that there are $\mathfrak{c}_n$ patterns σ of length n such that $\mathrm{Sort}(\sigma)$ is not a class. If instead $\mathrm{Sort}(\sigma)$ is a class, we explicitly determine its basis, which is either the singleton $\{132\}$ or the pair $\{132, \mathcal{R}(\sigma)\}$. In Section 3.3 we define a bivincular pattern ξ and show that σ-sortable permutations avoid ξ, unless σ is the skew-sum of 12 minus a 231-avoiding permutation. Permutations avoiding ξ display a rather regular geometric structure. This suggests that the cases where $\mathrm{Sort}(\sigma)$ is not contained in $\mathfrak{S}(\xi)$ could be the most challenging to be solved. In Section 3.4 we investigate the decreasing pattern $\sigma = \mathsf{ai}$ and show that $\mathrm{Sort}(\mathsf{ai}_k)$ is in bijection with Dyck paths of height at most $k-1$. Finally, in Section 3.5 we suggest some open problems and lines of research.

3.1 Patterns of length two

We start by analyzing the 12- and 21-machines.

Let $\sigma = 12$. Recall that the 12-machine consists in a pass through a 12-stack, followed by a pass through a classical stack. Notice that this device is substantially different from the one considered in [46], which is constituted by the same stacks, but allows a non-deterministic (and thus more powerful) sorting procedure.

Theorem 3.1. Let π be a permutation. If π is 12-sortable, then $\mathcal{S}^{12}(\pi) = \mathsf{ai}$. Moreover, we have:

$$\mathrm{Sort}(12) = \mathfrak{S}(213).$$

[1]With the exception of Section 3.3.

Therefore $f_n^{12} = \mathfrak{c}_n$, the n-th Catalan number.

Proof. Suppose that π is 12-sortable. We show that $\mathcal{S}^{12}(\pi) = \mathsf{ai}$ and π avoids 213 by induction on the length of π. This is trivial for the unit length permutation. Let π be a permutation of length two or more. Write π as $\pi = L1R$, where L is the prefix of π preceding 1 and R is the suffix of π following 1. Since $1x$ is an occurrence of 12 for each $x \in L$, 1 enters the 12-stack only when the 12-stack is empty. Similarly, 1 is extracted from the 12-stack only at the end, since $y1$ is not an occurrence of 12 for each $y \in R$. Therefore $\mathcal{S}^{12}(\pi) = \mathcal{S}^{12}(L)\mathcal{S}^{12}(R)1$. By the inductive hypothesis, $\mathcal{S}^{12}(L)$ and $\mathcal{S}^{12}(R)$ are decreasing. Moreover, it must be $x > y$ for each $x \in \mathcal{S}^{12}(L)$ and $y \in \mathcal{S}^{12}(R)$, otherwise $xy1$ would be an occurrence of 231 in $\mathcal{S}^{12}(\pi)$, contradicting the hypothesis that π is 12-sortable. Therefore $\mathcal{S}^{12}(\pi) = \mathsf{ai}$, as wanted. Finally, suppose, for a contradiction, that π contains an occurrence bac of 213. If bac is contained in L, then $\mathcal{S}^{12}(L)$ contains 231 by the inductive hypothesis and thus $\mathcal{S}^{12}(\pi)$ contains 231, a contradiction with π being 12-sortable. The same happens if bac is contained in R. On the other hand, if $b \in L$ and $c \in R$, then $bc1$ is an occurrence of 231 in $\mathcal{S}^{12}(\pi)$, which is impossible.

Conversely, suppose that π is not 12-sortable, or, equivalently, that $\mathcal{S}^{12}(\pi)$ contains an occurrence bca of 231. We wish to show that π contains 213. Note that necessarily b comes before c in π. Indeed a non-inversion in the output necessarily comes from a non-inversion in the input, since the 12-stack cannot repair inversions[2]. Moreover, b is extracted from the 12-stack before c enters. This must be due to the presence of an element x, located between b and c in π, which is smaller than b. More precisely, x is the next element of the input when b is extracted. The three elements b, x and c are thus an occurrence of 213 in π, as desired. □

Next we consider the pattern 21. The 21-machine is equivalent to the device considered by West in [50], where the following result is stated.

Theorem 3.2. [50] We have:

$$\mathrm{Sort}(21) = \mathfrak{S}(3241, 3\bar{5}241).$$

Due to the presence of the barred pattern $3\bar{5}241$, the set Sort(21) is not a permutation class. For example, the 21-sortable permutation 35241 contains the pattern 3241, which is not 21-sortable. On the other hand, Sort(12) is a class due to Theorem 3.1. By looking at some data for permutations of length three or more, it seems that the number of permutations σ such that σ-sortable permutations are not a class is equal to the n-th Catalan number. We will prove this rather striking fact in Section 3.2.

[2]If $\pi_i > \pi_j$, with $i < j$, then π_i is extracted from the 12-stack before π_j enters.

3.2 Classes and non-classes of σ-sortable permutations

Let σ be a permutation of length two or more. As one would expect, the σ-sortability of a permutation π is strongly affected by how π is related to the pattern σ defining the constraint of the stack. The permutation $\widehat{\sigma}$, defined below, proves to be crucial.

Definition 3.1. Let $\sigma = \sigma_1 \cdots \sigma_k$, with $k \geq 2$. Then define $\widehat{\sigma}$ as the permutation:

$$\widehat{\sigma} = \sigma_2\sigma_1\sigma_3 \cdots \sigma_k.$$

In other words, $\widehat{\sigma}$ is the permutation obtained from σ by interchanging the first two elements σ_1 and σ_2.

Lemma 3.3. Let π be a permutation. If π contains $\mathcal{R}(\sigma)$, then $\mathcal{S}^\sigma(\pi)$ contains $\widehat{\sigma}$.

Proof. Let $s_k s_{k-1} \cdots s_1$ be the (lexicographically) leftmost occurrence of $\mathcal{R}(\sigma)$ in π, where k is the length of σ. Consider the action of the σ-stack on π. Initially, every element of π is pushed into the σ-stack, until s_1 is the next element of the input. Now, before pushing s_1 into the σ-stack, the element s_2 has to be extracted, since $s_1 s_2 \cdots s_k \simeq \sigma$, with $s_2 \cdots s_k$ inside the σ-stack. On the other hand, s_3 is still in the σ-stack when s_1 enters, otherwise π would contain another occurrence of $\mathcal{R}(\sigma)$ strictly to the left of $s_k s_{k-1} \cdots s_1$, which is a contradiction. Thus s_1 is pushed into the σ-stack above s_3 and $\mathcal{S}^\sigma(\pi)$ contains $s_2 s_1 s_3 \cdots s_k$, which is an occurrence of $\widehat{\sigma}$. □

Lemma 3.4. Let π be an input permutation for the σ-machine.

1. If π avoids $\mathcal{R}(\sigma)$, then $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$. In this case, π is σ-sortable if and only if π avoids 132.

2. If π contains $\mathcal{R}(\sigma)$, then $\mathcal{S}^\sigma(\pi)$ contains $\widehat{\sigma}$. In this case, if $\widehat{\sigma}$ contains 231, then π is not σ-sortable.

Proof.

1. If π avoids $\mathcal{R}(\sigma)$, then the restriction of the σ-stack is never triggered. Thus every element of π is pushed directly into the σ-stack and $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$. In particular, π is σ-sortable if and only if $\mathcal{R}(\pi)$ avoids 231, or, equivalently, π avoids 132.

2. Suppose that π contains $\mathcal{R}(\sigma)$. By Lemma 3.3, $\mathcal{S}^\sigma(\pi)$ contains $\widehat{\sigma}$. Therefore, if $\widehat{\sigma}$ contains 231 then π is not σ-sortable.

□

Corollary 3.5. For each permutation σ, we have:

$$\mathfrak{S}\left(132, \mathcal{R}(\sigma)\right) \subseteq \mathrm{Sort}(\sigma).$$

Proof. It is an immediate consequence of the first item of Lemma 3.4. $\square$

Theorem 3.6. Let σ be a permutation. If $\widehat{\sigma}$ contains 231, then:

$$\mathrm{Sort}(\sigma) = \mathfrak{S}(132, \mathcal{R}(\sigma)).$$

Therefore $\mathrm{Sort}(\sigma)$ is a class with basis either $\{132, \mathcal{R}(\sigma)\}$, if $\mathcal{R}(\sigma)$ avoids 132, or $\{132\}$, otherwise.

Proof. Following Corollary 3.5, all we need to prove is that $\mathrm{Sort}(\sigma) \subseteq \mathfrak{S}(132, \mathcal{R}(\sigma))$. Suppose that π is σ-sortable. If π contains $\mathcal{R}(\sigma)$, then $\mathcal{S}^\sigma(\pi)$ contains $\widehat{\sigma}$ by Lemma 3.4. But then $\widehat{\sigma}$ contains 231 by hypothesis, contradicting the fact that π is σ-sortable. Otherwise, suppose that π avoids $\mathcal{R}(\sigma)$, but contains 132. Due to the same Lemma 3.4, we have $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$, which contains 231, a contradiction with π being σ-sortable. $\square$

Corollary 3.7. Let $k \geq 3$. Then:

$$\mathrm{Sort}(\mathsf{ai}_k) = \mathfrak{S}(132, \mathsf{id}_k).$$

In particular, the set of 321-sortable permutations is a class with basis $\{132, 123\}$.

Proof. It follows immediately from Theorem 3.6, since $\widehat{\mathsf{ai}_k} = (k-1)k(k-2)\cdots 21$ contains an occurrence $(k-1)k1$ of 231. $\square$

Theorem 3.6, which is a rather straightforward consequence of Lemma3.4, gives a sufficient condition for the set $\mathrm{Sort}(\sigma)$ to be a class. Next we show that this condition is also necessary.

Theorem 3.8. If $\widehat{\sigma}$ avoids the pattern 231, then $\mathrm{Sort}(\sigma)$ is not a permutation class.

Proof. The case by case analysis of Table 3.1 and Corollary 3.7 show that the theorem holds for patterns σ of length three. Now suppose that σ has length at least four. It is not hard to realize that the permutation 132 is not σ-sortable, since $\mathcal{S}^\sigma(132) = 231$. Next we show that, under the hypothesis that $\widehat{\sigma}$ avoids 231, it is always possible to construct a permutation α such that α contains 132, but α is σ-sortable. This proves that $\mathrm{Sort}(\sigma)$ is not closed downwards, as desired. Let $\sigma = \sigma_1\sigma_2\cdots\sigma_k$. We distinguish two cases, according to whether $\sigma_1 < \sigma_2$ or $\sigma_1 > \sigma_2$.

1. If $\sigma_1 < \sigma_2$, define $\alpha = \sigma'_k\sigma'_{k-1}\cdots\sigma'_3 z\sigma'_2\sigma'_1$, where:

- $z = \sigma_1$;
- $\sigma'_i = \begin{cases} \sigma_i, & \text{if } \sigma_i < \sigma_1; \\ \sigma_i + 1, & \text{otherwise.} \end{cases}$

Notice that $z\sigma'_2\sigma'_1$ is an occurrence of 132. We show that α is σ-sortable by providing a detailed analysis of the behavior of the σ-machine on input α. Initially, the elements of α are pushed into the σ-stack until σ'_1 is the next element of the input. In particular, both the additional element z and σ'_2 can be safely pushed: indeed $\sigma'_2 z \cdots \sigma'_{k-1}\sigma'_k$ is not an occurrence of σ, since $\sigma_1 < \sigma_2$, whereas $\sigma'_2 > z$. Now, before σ'_1 enters the σ-stack, the element σ'_2 is extracted. At this point, σ'_1 can enter without violating the restriction, again because $\sigma_2 > \sigma_1$, whereas $z < \sigma'_1$, and so $\sigma'_1 z\sigma'_3 \cdots \sigma'_k$ is not an occurrence of σ. The output of the σ-stack is then $\mathcal{S}^\sigma(\alpha) = \sigma'_2\sigma'_1 z\sigma'_3 \cdots \sigma'_k$. We wish to show that $\mathcal{S}^\sigma(\alpha)$ avoids 231. Since $\widehat{\sigma}$ avoids 231 by hypothesis, and $\sigma'_2\sigma'_1\sigma'_3 \cdots \sigma'_k$ is an occurrence of $\widehat{\sigma}$, any potential occurrence of 231 necessarily involves the additional element z. In particular, it is easy to observe that z can be neither the smallest nor the biggest element of such a pattern, because $z < \sigma'_1 < \sigma'_2$ and z is the third element of $\mathcal{S}^\sigma(\alpha)$. Finally, if z were the first element of an occurrence $z\sigma'_j\sigma'_l$ of 231 in $\mathcal{S}^\sigma(\alpha)$, then $\sigma_1\sigma_j\sigma_l$ would be an occurrence of 231 in $\widehat{\sigma}$, contradicting the hypothesis.

2. If $\sigma_1 > \sigma_2$, define $\alpha = \sigma'_k\sigma'_{k-1} \cdots \sigma'_3\sigma'_2\sigma'_1 z$, where:

   - $z = \sigma_2 + 1$;
   - $\sigma'_i = \begin{cases} \sigma_i, & \text{if } \sigma_i \leq \sigma_2; \\ \sigma_i + 1, & \text{otherwise.} \end{cases}$

   Observe that $\sigma'_2\sigma'_1 z$ is an occurrence of 132. As for the previous case, we now describe what happens when α is processed by the σ-machine. The first element that cannot be pushed into the σ-stack is σ'_1, which forces σ'_2 to be extracted. Successively both σ'_1 and z can enter the σ-stack, since $z\sigma'_1\sigma'_3 \cdots \sigma'_k$ is not an occurrence of σ: indeed $\sigma_1 > \sigma_2$, whereas $z < \sigma'_1$. Therefore the output of the σ-stack is $\mathcal{S}^\sigma(\alpha) = \sigma'_2 z\sigma'_1\sigma'_3 \cdots \sigma'_k$. Again any potential occurrence of 231 in $\mathcal{S}^\sigma(\alpha)$ must involve the additional element z. However z cannot be the smallest element of a pattern 231, because it is the second element of $\mathcal{S}^\sigma(\alpha)$. Moreover, if z were the first element of a 231, then σ_2 would be the first element of an occurrence of 231 in $\widehat{\sigma}$, which is forbidden. Finally, if z were the largest element of a 231, then σ'_2 would be the first element of such an occurrence, so also σ'_1, which is greater than both σ'_2 and z, would be the largest element of an occurrence of 231 which does not involve z, giving again a contradiction.

σ	**σ-sortable permutation**	**Non-σ-sortable pattern**
123	4132	132
132	2413	132
213	4132	132
231	361425	1324
312	3142	132
321	class	

Table 3.1: Classes and non-classes of σ-sortable permutations, for patterns σ of length three.

In each of the two cases considered, we proved that $\mathcal{S}^{\sigma}(\alpha)$ avoids 231, thus α is a σ-sortable permutation that contains the non σ-sortable pattern 132, as desired. $\square$

Corollary 3.9. For every permutation σ of length three or more, the set $\mathrm{Sort}(\sigma)$ is a permutation class if and only if $\widehat{\sigma}$ contains the pattern 231.

Corollary 3.10. The permutations σ for which $\mathrm{Sort}(\sigma)$ is not a permutation class are enumerated by the Catalan numbers.

Proof. Such permutations are in bijection with $\mathfrak{S}(231)$, which is known to be enumerated by the Catalan numbers. $\square$

What we have proved so far assures that $\mathrm{Sort}(\sigma)$ is a permutation class if and only if $\widehat{\sigma}$ contains the pattern 231. In this case, $\mathrm{Sort}(\sigma) = \mathfrak{S}\,(132, \mathcal{R}(\sigma))$, hence the basis of $\mathrm{Sort}(\sigma)$ has exactly two elements if and only if $\mathcal{R}(\sigma)$ avoids 132, or, equivalently, if σ avoids 231. Next we enumerate those patterns σ such that the basis of $\mathrm{Sort}(\sigma)$ has two elements.

Lemma 3.11. Let $\sigma = \sigma_1 \cdots \sigma_k$, with $k \geq 3$, and suppose that $\mathcal{R}(\sigma)$ avoids 132. Then $\widehat{\sigma}$ contains the pattern 231 if and only if $\sigma_1\sigma_2\sigma_3$ is an occurrence of 321.

Proof. Observe that, since σ avoids 231 by hypothesis, an occurrence of 231 in $\widehat{\sigma} = \sigma_2\sigma_1\sigma_3 \cdots \sigma_k$ must involve both σ_1 and σ_2, respectively as the first and the second element of the pattern, with $\sigma_2 < \sigma_1$.

Suppose that $\widehat{\sigma}$ contains an occurrence $\sigma_2\sigma_1\sigma_i$ of 231, for some $i \geq 3$. If $\sigma_3 > \sigma_2$, then $i > 4$ and thus $\sigma_2\sigma_3\sigma_i$ is an occurrence of 231 in σ, which is a contradiction. Therefore we have $\sigma_3 < \sigma_2$ and $\sigma_1\sigma_2\sigma_3$ is an occurrence of 321, as desired.

Conversely, if $\sigma_1\sigma_2\sigma_3$ is an occurrence of the pattern 321, then clearly $\sigma_2\sigma_1\sigma_3$ is an occurrence of 231 in $\widehat{\sigma}$. $\square$

Proposition 3.12. Let $n \geq 1$. Define $\mathcal{A}_n = \{\pi \in \mathfrak{S}_n(231) : \pi_1\pi_2\pi_3 \simeq 321\}$ and let $a_n = |\mathcal{A}_n|$. Then, for each $n \geq 2$, we have $a_n = \mathfrak{c}_n - 2\mathfrak{c}_{n-1}$. In particular, the generating function of the sequence $(a_n)_{n\geq 0}$ is:

$$A(x) = \frac{1 - 4x + 2x^2 - (1-2x)\sqrt{1-4x}}{2x}.$$

Proof. Suppose that $n \geq 2$. Define the sets:

$$\mathcal{F}_n = \{\pi \in \mathfrak{S}_n(231) : \pi_1 < \pi_2\}$$

and

$$\mathcal{G}_n = \{\pi \in \mathfrak{S}_n(231) : \pi_1 > \pi_2, \pi_2 < \pi_3\},$$

so that:

$$\mathfrak{S}_n(231) = \mathcal{A}_n \dot{\cup} \mathcal{F}_n \dot{\cup} \mathcal{G}_n.$$

Let $f_n = |\mathcal{F}_n|$ and $g_n = |\mathcal{G}_n|$. Since $|\mathfrak{S}_n(231)| = \mathfrak{c}_n$, we have $a_n = \mathfrak{c}_n - (f_n + g_n)$. We now show that $f_n = g_n = \mathfrak{c}_{n-1}$ by providing bijections between $\mathcal{F}_n$ and $\mathfrak{S}_{n-1}(231)$, as well as between $\mathcal{G}_n$ and $\mathfrak{S}_{n-1}(231)$. The desired enumeration follows.

- If $\pi \in \mathcal{F}_n$, then it must be $\pi_1 = 1$, otherwise $\pi_1\pi_2 1$ would be an occurrence of 231 in π. Define the map $f : \mathcal{F}_n \to \mathfrak{S}_{n-1}(231)$, where $f(\pi)$ is obtained from π by removing $\pi_1 = 1$ and subtracting one to the remaining entries. It is easy to realize that $f(\pi) \in \mathrm{Sort}_n(231)$ and that f is an injection. Moreover, if $\tau \in \mathfrak{S}_{n-1}(231)$, then adding a new minimum at the beginning (and rescaling the other elements) cannot create any occurrence of 231, so f is also surjective.

- If $\pi \in \mathcal{G}_n$, then it must be $\pi_2 = 1$, otherwise it would be $\pi_2\pi_3 1 \simeq 231$ in π, a contradiction. We thus define $g : \mathcal{G}_n \to \mathfrak{S}_{n-1}(231)$ such that $g(\pi)$ is obtained from π by removing $\pi_2 = 1$ and rescaling the remaining elements. Again it is clear that $g(\pi) \in \mathfrak{S}_n(231)$ and that g is an injection. Finally, if $\tau \in \mathfrak{S}_{n-1}(231)$, then the permutation π obtained from τ by adding a new minimum in the second position avoids 231. Indeed a potential occurrence of 231 in π should involve the added element π_2, and so π_2 would be either the first or the second element of such an occurrence. But this is impossible since $\pi_2 = 1$. Therefore g is a bijection between $\mathcal{G}_n$ and $\mathfrak{S}_{n-1}(231)$, as desired.

Let us now compute the generating function $A(t) = \sum_{n\geq 1} a_n t^n$. Let $C(t) = (1 - \sqrt{1-4t})/(2t)$ be the generating function for the Catalan numbers. We have:

$$A(t) = \sum_{n\geq 0} a_{n+2}t^{n+2} =$$
$$\sum_{n\geq 0} \mathfrak{c}_{n+2}t^{n+2} - 2t\sum_{n\geq 0} a_{n+1}t^{n+1}$$
$$= C(t) - t - 1 - 2t(C(t) - 1) =$$
$$C(t)(1-2t) + t - 1,$$

from which

$$A(t) = \frac{1 - 4t + 2t^2 - (1-2t)\sqrt{1-4t}}{2t},$$

as desired. □

The sequence $(a_n)_{n\geq 0}$ is recorded (with offset two) as sequence A002057 in [45]. The first terms are $0, 0, 1, 4, 14, 48, 165, 572, 2002$. An alternative expression for its generating function is given by $A(t) = t^2C(t)^4$, although we are not able to provide a combinatorial explanation of this fact.

We end this section by collecting some enumerative results concerning classes of σ-sortable permutation with basis of cardinality two (see Appendix B). A direct combinatorial argument can be used in order to prove each of these results, as we show in the following example. In fact, due to Theorem 3.6, each of these classes is a subclass of $\mathfrak{S}(132)$ of the form $\mathfrak{S}(132, \mathcal{R}(\sigma))$, and thus its generating function is rational. A constructive proof of this fact can be found in [41], which provides an algorithm to compute the generating function in all such cases. A clear and succint description of the algorithm (in the context of representing catalan structures as arch systems) is given in [3].

Example 3.1. Let $\sigma = 421356$. Then $\mathrm{Sort}(\sigma) = \mathfrak{S}(132, 653124)$ due to Theorem 3.6. Given $\pi \in \mathfrak{S}_n(132, 653124)$, write $\pi = LnR$, where L is the prefix of π that precedes n and R is the suffix of π that follows n. Notice that, since π avoids 132, we have $L > R$, i.e. $x > y$ for each $x \in L$ and $y \in R$ (otherwise it would be $xny \simeq 132$). Now, we can partition $\mathfrak{S}(132, 653124)$ according to whether L is increasing or not in the above decomposition of π. If L is increasing, then $\pi \in \mathfrak{S}(132, 653124)$ if and only if $R \in \mathfrak{S}(132, 53124)$. Indeed any occurrence of 53124 in R would realize an occurrence of 653124 together with n. Conversely, if L is increasing, then the elements corresponding to 53124 in any occurrence of 653124 in π must belong to R. Similarly, if L contains at least one descent, then $\pi \in \mathfrak{S}(132, 653124)$ if and only if $R \in \mathfrak{S}(132, 3124)$. Let $G(t) = \sum_{k\geq 0} g_k t^k$ be the generating function of $\mathfrak{S}(132, 53124)$ and let $H(t) = \sum_{k\geq 0} h_k t^k$ be the generating function of $\mathfrak{S}(132, 3124)$. Let $F(t) = F^{421356}(t)$ and $f_n = f^{421356}$, for $n \geq 0$.

Then, due to the above discussion:

$$f_{n+1} = \sum_{k=0}^{n} 1 \cdot g_{n-k} + \sum_{k=0}^{n} (f_k - 1) h_{n-k}.$$

By summing over n, we get:

$$\frac{1}{t}\left(F(t) - 1)\right) = \frac{1}{1-t} G(t) + F(t)H(t) - \frac{1}{1-t} H(t).$$

Now, it is easy to compute the generating functions $G(t) = \frac{t^2-3t+1}{3t^2-4t+1}$ and $H(t) = \frac{1-2t}{1-3t+t^2}$. Then, solving the above equation yields:

$$F(t) = \frac{2t^5 - 16t^4 + 29t^3 - 23t^2 + 8t - 1}{9t^5 - 33t^4 + 46t^3 - 30t^2 + 9t - 1}$$

The resulting sequence starts $1, 2, 5, 14, 42, 131, 416, 1329, 4247, 13544, \ldots$ and does not appear in [45].

3.3 A set of challenging patterns

For the rest of this chapter, denote by $\xi = (132, \{0, 2\}, \emptyset)$ the bivincular pattern depicted in Figure 3.1. The main result of this section is a proof that $\mathrm{Sort}(\sigma)$ is always a subset of $\mathfrak{S}(\xi)$, unless σ is the skew sum of 12 with a non-empty 231-avoiding permutation β. The geometric structure of permutations avoiding ξ can be described precisely, as we show in what follows. This suggests that the family of σ-machines, when $\sigma = 12 \ominus \beta$, could contain the more challenging σ-machines to be studied. The shortest such pattern is 231. In fact, as suggested by some data, the 231-machine seems to be the σ-machine that can sort the largest amount of permutations. For example, it is the only one that can sort every permutation of length three.

We start by providing a geometric description of $\mathfrak{S}(\xi)$, from which its enumeration follows easily. Let $\pi = \pi_1 \cdots \pi_n \in \mathfrak{S}(\xi)$ and let $\pi_1 = t + 1$, for some $t \geq 0$. Let $\{\pi_{i_1}, \ldots, \pi_{i_t}\}$ be the set of elements of π that are smaller than π_1, with $i_1 < i_2 < \cdots < i_t$. For $j = 1, \ldots, t$, let $b_j = \pi_{i_j}$. Finally, write:

$$\pi = \pi_1 B_0 b_1 B_1 b_2 B_2 \cdots b_t B_t,$$

where $B_j = \pi_{i_j+1} \cdots \pi_{i_{j+1}-1}$, for $j = 0, 1, \ldots, t$. We refer to this as the *first-element decomposition* of π; for $j = 0, 1, \ldots, t$, B_j is said to be the j-th *block* of π in its first-element decomposition.

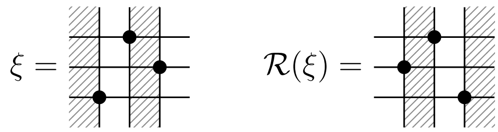

Figure 3.1: Bivincular patterns $\xi = (132, \{0,2\}, \emptyset)$ and $\mathcal{R}(\xi)$.

Lemma 3.13. Let $\pi = \pi_1 B_0 b_1 B_1 b_2 B_2 \cdots b_t B_t$ be the first-element decomposition of π. Then π avoids ξ if and only if B_j is increasing for each j.

Proof. Suppose that π avoids ξ and let $j \geq 0$. By definition of first-block decomposition, all the elements contained in B_j are greater than π_1. Therefore B_j is increasing, since otherwise a descent in B_j would result in an occurrence of ξ. On the other hand, suppose that $\pi_u \pi_v \pi_{v+1}$ is an occurrence of ξ in π. Note that $u = 1$ and $\pi_v > \pi_{v+1}$, with $v \geq 2$ and $\pi_{v+1} > \pi_1$. Thus π_v and π_{v+1} are in the same block B_j, for some j, and B_j is not increasing. □

Corollary 3.14. If π avoids ξ and $\pi_1 = 1$, then π is the increasing permutation.

Theorem 3.15. For $n \geq 0$ and $t = 0, 1, \ldots, n-1$, define $\mathfrak{S}_n^t(\xi)$ by

$$\mathfrak{S}_n^t(\xi) = \{\pi \in \mathfrak{S}_n(\xi) : \pi_1 = t+1\}.$$

Let $f_{n,t}$ be the cardinality of $\mathfrak{S}_n^t$. Then:

$$f_{n,t} = t!(t+1)^{n-t-1}.$$

In particular, $|\mathfrak{S}_n(\xi)| = \sum_{t=0}^{n-1} t!(t+1)^{n-t-1}$ (sequence A129591 in [45]).

Proof. Any permutation $\pi \in \mathfrak{S}_n^t(\xi)$ can be constructed as follows. The t elements of π that are smaller than $\pi_1 = t+1$ can be chosen freely, since they cannot contribute to an occurrence of ξ. This can be done in $t!$ distinct ways. On the other hand, by Lemma 3.13, elements greater than π_1 must be arranged in increasing blocks. In other words, for each of them it is sufficient to choose the index of the (increasing) block it belongs to. So there are $t+1$ possibilities for each of the remaining $n-t-1$ elements. Therefore $f_{n,t} = t! \cdot (t+1)^{n-t-1}$, as desired. □

Theorem 3.16. Let σ be a permutation of length at least three. The following three conditions are equivalent:

(1) $\mathrm{Sort}(\sigma) \not\subseteq \mathfrak{S}(\xi)$.

(2) $\sigma = 12 \ominus \beta$, for some $\beta \in \mathfrak{S}(231)$.

(3) $\widehat{\sigma} \in \mathfrak{S}(231)$ and $\mathcal{R}(\sigma) \notin \mathfrak{S}(\xi)$.

Proof. Let $\sigma = \sigma_1 \cdots \sigma_k$, with $k \geq 3$.

- We start by proving that (2) and (3) are equivalent. Suppose that $\sigma = 12 \ominus \beta$, for some $\beta = \beta_1 \cdots \beta_s \in \mathfrak{S}(231)$, where $s = k-2$. Observe that $\widehat{\sigma} = (s+2)(s+1)\beta$ avoids 231, since β does so. Finally, we have $\mathcal{R}(\sigma) = \beta_s \cdots \beta_1(s+2)(s+1)$, thus $\beta_s(s+2)(s+1)$ is an occurrence of ξ in $\mathcal{R}(\sigma)$, as wanted.

  Conversely, suppose that $\widehat{\sigma}$ avoids 231 and $\mathcal{R}(\sigma)$ contains ξ, or, equivalently, σ contains $\mathcal{R}(\xi)$. The pattern $\mathcal{R}(\xi)$ is depicted in Figure 3.1. Let $\sigma_i\sigma_{i+1}\sigma_k$ be an occurrence of $\mathcal{R}(\xi)$ in σ. Note that the classical pattern underlying $\mathcal{R}(\xi)$ is 231, but $\widehat{\sigma}$ avoids 231 by hypothesis. Therefore it has to be $i = 1$, otherwise $\sigma_i\sigma_{i+1}\sigma_k$ would still be an occurrence of 231 in $\widehat{\sigma}$, which is impossible. Thus $\sigma_k < \sigma_1 < \sigma_2$. Now, observe that $\sigma_u < \sigma_1$ for each $u > 2$. Otherwise, if $\sigma_u > \sigma_1$ for some $2 < u < k$, then $\sigma_1\sigma_u\sigma_k$ would be an occurrence of 231 in $\widehat{\sigma}$, which is again impossible. Therefore $\sigma = \sigma_1\sigma_2 \ominus \beta = 12 \ominus \beta$, where $\beta = \sigma_3 \cdots \sigma_k$. Finally, β avoids 231 because $\widehat{\sigma}$ does so, as wanted.

- Next we wish to prove that (3) implies (1). Suppose that $\widehat{\sigma}$ avoids 231 and $\mathcal{R}(\sigma)$ contains ξ. We show that $\mathcal{R}(\sigma)$ is σ-sortable (and contains ξ), thus $\mathcal{R}(\sigma) \in \mathrm{Sort}(\sigma) \setminus \mathfrak{S}(\xi)$. Due to Lemma 3.4, we have $\mathcal{S}^\sigma(\mathcal{R}(\sigma)) = \widehat{\sigma}$. Finally, $\widehat{\sigma}$ avoids 231, so $\mathcal{R}(\sigma)$ is σ-sortable, as desired.

- Finally, we show that (1) implies (2), which completes the proof. Suppose that there is a permutation $\pi = \pi_1 \cdots \pi_n$ such that π is σ-sortable and π contains ξ. Let $\pi_1\pi_j\pi_{j+1}$ be an occurrence of ξ in π. Let $\beta = \sigma_3 \cdots \sigma_k$. We show that $\sigma_2 > \sigma_1 > \sigma_u$ for each $u \geq 3$ and β is a 231-avoiding permutation. Observe that $\widehat{\sigma}$ avoids 231. Otherwise it would be $\mathrm{Sort}(\sigma) = \mathfrak{S}(132, \mathcal{R}(\sigma))$ due to Theorem 3.6 and thus $\mathrm{Sort}(\sigma) \subseteq \mathfrak{S}(\xi)$, contradicting the hypothesis. In particular, β avoids 231, as wanted. Now, since π_1 is the last element that exits the σ-stack, π_j must be extracted before π_{j+1} enters, else $\pi_{j+1}\pi_j\pi_1$ would be an occurrence of 231 in $\mathcal{S}^\sigma(\pi)$, contradicting the fact that π is σ-sortable. Let us consider the instant when π_j is extracted (and π_{j+1} is the next element of the input). Since a pop operation is performed by the σ-stack, the σ-stack must contain $k-1$ elements $\alpha_2\alpha_3 \cdots \alpha_k$ (reading from top to bottom) such that $\pi_{j+1}\alpha_2 \cdots \alpha_k$ is an occurrence of σ. Without losing generality, we can suppose that α_3 is still in the σ-stack when π_{j+1} enters: this can be achieved, for instance, by taking the "deepest" such sequence of elements in the σ-stack. Note that $\mathcal{S}^\sigma(\pi)$ contains the occurrence $\alpha_2\pi_{j+1}\alpha_3 \cdots \alpha_k$ of $\widehat{\sigma}$. Now, if $\alpha_v > \pi_{j+1}$ for some $v \geq 3$, then $\alpha_v \neq \pi_1$ (because $\pi_1 < \pi_{j+1}$) and $\pi_{j+1}\alpha_v\pi_1$ is an occurrence of 231 in $\mathcal{S}^\sigma(\pi)$, a contradiction with π being σ-sortable. Therefore, since $\pi_{j+1}\alpha_2 \cdots \alpha_k \simeq \sigma$, we have $\sigma_u < \sigma_1$ for

k	**Sequence** $\{f_n^{\mathsf{ai}_k}\}_n$	**OEIS**
3	1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024	A011782
4	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181, 10946	A001519
5	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842, 29525	A124302
6	1, 2, 5, 14, 42, 131, 417, 1341, 4334, 14041, 45542	A080937
7	1, 2, 5, 14, 42, 132, 428, 1416, 4744, 16016, 54320	A024175

Table 3.2: Enumerative results for ai_k-sortable permutations, with $k = 3, 4, 5, 6, 7$, starting from permutations of length one.

each $u \geq 3$. To conclude the proof, we have to show that $\sigma_1 < \sigma_2$. Suppose, for a contradiction, that $\sigma_1 > \sigma_2$. Then $\sigma_1 = k$ is the maximum element of σ, since $\sigma_1 > \sigma_u$ for each $u \geq 3$. Now, consider the instant immediately after π_j is pushed into the σ-stack (and π_{j+1} is the next element of the input). Note that $\pi_{j+1} > \alpha_2$, because we are assuming $\sigma_1 > \sigma_2$ and $\pi_{j+1}\alpha_2 \simeq \sigma_1\sigma_2$. But $\pi_j > \pi_{j+1}$, thus $\pi_j\alpha_2 \cdots \alpha_s$ is an occurrence of σ contained in the σ-stack, which is impossible.

□

Corollary 3.17. Let $\sigma = 12 \ominus \beta$, for some non-empty and 231-avoiding permutation β. Let $\pi = \pi_1 \cdots \pi_n$ be a σ-sortable permutation with $\pi_1 = 1$. Then π is the identity permutation.

Proof. It follows from Lemma 3.14 and Theorem 3.16. □

Corollary 3.17 fails if $\mathrm{Sort}(\sigma) \not\subseteq \mathfrak{S}(\xi)$. For example, the permutations 12354, 12453, 12534 and 12543 are 3421-sortable.

3.4 The decreasing pattern

In this section we provide some enumerative results for the sets $\mathrm{Sort}(\mathsf{ai}_k)$, highlighting a link with a class of pattern-avoiding lattice paths. The results of the previous section allow us to directly characterize σ-sortable permutations when σ is the decreasing pattern. Indeed, by Theorem 3.6 and for each $k \geq 1$, we have $\mathrm{Sort}(\mathsf{ai}_k) = \mathfrak{S}(\mathsf{id}_k, 132)$. The sequences that enumerate these sets, for $k \leq 7$, are reported in Table 3.2.

If $n < k$, then obviously $\mathrm{Sort}_n(\mathsf{ai}_k) = \mathfrak{S}_n(132)$ and thus $f_n^{\mathsf{ai}_k} = \mathfrak{c}_n$. Therefore the rows of Table 3.2 tend to the sequence of Catalan numbers. By looking at the reference in [45] for small values of k, we notice that $\{f_n^{\mathsf{ai}_k}\}_n$ counts the number of Dyck paths of height at most $k-1$. A formal proof can be obtained by using the bijection between Dyck paths and 132-avoiding permutations mentioned in Example 1.3. Indeed, if π is a 132-avoiding permutation and P is the Dyck path associated to π, then the maximum length of an increasing sequence in π is equal to to the height of P. Finally, a permutation π avoids id_k if and only if the maximum length of an increasing sequence in π is at most $k-1$. Dyck paths of bounded height are rather well studied objects (see for example [15, 33]).

We now compute the generating function of $\{f_n^{\mathsf{ai}_k}\}_n$ by exploiting this connection with Dyck paths of bounded height. Let $F_k(t) = F^{\mathsf{ai}_k}(t)$. Given a Dyck path P, consider its first-return decomposition $P = \mathtt{U}Q_1\mathtt{D}Q_2$, for some (possibly empty) Dyck paths Q_1, Q_2 (see Remark 1.1). If P has height at most k, then Q_2 has height at most k, whereas Q_1 has height at most $k-1$. This provides a recursive description of $F_k(t)$ with respect to the semilength:

$$\begin{cases} F_0(t) = 1; \\ F_k(t) = 1 + tF_{k-1}(t)F_k(t), \ k \geq 1. \end{cases}$$

A consequence of the above recurrence is that $F_k(t)$ is rational, for all k; indeed we have

$$F_k(t) = \frac{G_k(t)}{G_{k+1}(t)},$$

where $G_0(t) = G_1(t) = 1$ and $G_k(t)$ satisfies the recurrence

$$G_{k+1}(t) = G_k(t) - tG_{k-1}(t).$$

Solving this recurrence yields

$$G_k(t) = \sum_{i \geq 0} \binom{n-1}{i} (-t)^i.$$

The polynomials $G_k(t)$ are sometimes called *Catalan polynomials* (see for instance [26]); the table of their coefficients is sequence A115139 in [45].

3.5 Open problems

In this chapter we provided some general results regarding σ-machines and sets of σ-sortable permutations. As a consequence of Corollary 3.9, we are able to tell when $\mathrm{Sort}(\sigma)$ is a permutation class by simply checking whether $\widehat{\sigma}$ contains 231 or not. If $\mathrm{Sort}(\sigma)$ is a class, Theorem 3.6 states that $\mathrm{Sort}(\sigma) = \mathfrak{S}(132, \mathcal{R}(\sigma))$, thus the set of σ-sortable permutations is completely determined (and enumerated). On the other hand, Theorem 3.16 is currently the only known general result when $\mathrm{Sort}(\sigma)$

σ	**Sequence** $\{f_n^\sigma\}_n$	**OEIS**
213	1, 2, 5, 16, 62, 273, 1307, 6626, 35010, 190862	
231	1, 2, 6, 23, 102, 496, 2569, 13934, 78295, 452439	
312	1, 2, 5, 15, 52, 201, 843, 3764, 17659, 86245	A202062

Table 3.3: Enumerative data for unsolved patterns of length three, starting from σ-sortable permutations of length one.

is not a permutation class. It would be interesting to provide more results in order to find structural information on the sets $\mathrm{Sort}(\sigma)$, when they are not permutation classes.

Open Problem 3.1. Find geometric properties of the set $\mathrm{Sort}(\sigma)$, when $\mathrm{Sort}(\sigma)$ is not a permutation class.

More specifically, the only non-class for patterns σ of length two is Sort(21), which is the classical case of West's 2-stack sortable permutations. Moving on to patterns of length three, the only permutation class is $\mathrm{Sort}(321) = \mathfrak{S}(132, 123)$. We provide a characterization of the sets Sort(123), in Chapter 4, and Sort(132), in Chapter 5. The remaining three patterns are yet to be solved. Some related data are reported in Table 3.3. A potentially interesting link with ascent sequences is the following: in Chapter 5 we prove that Sort(132) is Wilf-equivalent to the set $\mathfrak{A}(312, 321)$ of ascent sequences avoiding 312 and 321 (see [10]), while Sort(321) seems to be Wilf-equivalent to $\mathfrak{A}(312)$.

Open Problem 3.2. Characterize and enumerate the sets Sort(213), Sort(231) and Sort(312).

If we consider the family of σ-machines from the enumerative perspective, it would be nice to investigate deeper the notion of Wilf-equivalence that naturally arises by looking at how many different sequences of σ-sortable permutations can be obtained for patterns σ of a fixed length. Formally, we say that two patterns σ and τ of length k are *PAM-Wilf-equivalent* (where PAM stands for pattern-avoiding machine) if the sets $\mathrm{Sort}(\sigma)$ and $\mathrm{Sort}(\tau)$ are Wilf-equivalent in the usual sense. Denote by w_k the number of PAM-Wilf classes of length k.

Open Problem 3.3. Compute the number of PAM-Wilf classes, that is the sequence $\{w_k\}_{k\geq 1}$.

A slightly easier version of the above open problem can be obtained by considering the sets $\mathrm{Sort}(\sigma)$ which are permutation classes only. Some data (for which the author is extremely grateful to Christian Bean and Anders Claesson) indicate that the first terms of the resulting sequence, starting from length two, are $1, 1, 2, 5, 11, 25, 55, 126, 283$ (not in [45]). For example, there are 11 such Wilf-classes for patterns σ of length six: 10 of them are reported in Appendix B and the last one consists of those patterns σ such that $\mathrm{Sort}(\sigma)$ is a class and $\sigma \geq 132$, that is where $\mathrm{Sort}(\sigma) = \mathfrak{S}(132)$ and the counting sequence is the sequence of Catalan numbers.

Chapter 4

The 123-machine

This chapter is devoted to the analysis of the 123-machine. The paper [24] contains most of the results presented in this part of the thesis. Since, as a consequence of Corollary 3.9, the set Sort(123) is not a permutation class, this pattern is considerably more challenging that the decreasing pattern of the same length. If we compute the first terms of the sequence $\{f_n^{123}\}_{n\geq 1}$, we get $1, 2, 5, 13, 35, 99, \ldots$, which suggests a match with A294790 in [45]. This sequence enumerates, for example, Schröder paths avoiding the (consecutive) path $\mathtt{UH_2D}$ (see [25]). Our goal is to provide a length-preserving bijection between 123-sortable permutations and this family of pattern-avoiding paths. To do that, we follow a step-by-step procedure, aiming to progressively reduce the problem of characterizing 123-sortable permutations to more manageable subsets of Sort(123).

4.1 Structural description of Sort(123)

We start by dealing with 123-sortable permutations that start with an ascent.

Lemma 4.1. Let $\pi \in \mathfrak{S}_n$. If π is 123-sortable, then $\pi_2 \leq \pi_1 + 1$.

Proof. Suppose, for a contradiction, that $\pi_2 > \pi_1 + 1$. Then there exists an index $i \geq 3$ such that $\pi_2 > \pi_i > \pi_1$. Note that the first two elements π_1 and π_2 are extracted from the 123-stack only when the 123-stack is emptied at the end of the sorting process. Indeed, since $\pi_1 < \pi_2$ (and π_2 enters above π_1), they cannot be both part of an occurrence of 123. Thus $\pi_i\pi_2\pi_1$ is an occurrence of 231 in $\mathcal{S}^{123}(\pi)$, contradicting the hypothesis that π is 123-sortable. $\square$

Let us now partition Sort(123) according to Lemma 4.1: permutations starting with a consecutive ascent $\pi_2 = \pi_1 + 1$, and permutations starting with a descent. The next step consists in showing that inflating the first element of a 123-sortable permutation does not affect its 123-sortability.

Lemma 4.2. Let π be a permutation of length n and let π' be the permutation (of length $n+1$) obtained from π by 2-inflating π_1. Then π is 123-sortable if and only if π' is 123-sortable.

Proof. Observe that, by hypothesis, the first two elements of π' are consecutive in value (a and $a+1$, say) and the first one is smaller than the second one. Therefore, during the sorting process, such two elements remain at the bottom of the 123-stack (with $a+1$ above a) until all the other elements of the input permutations have exited it. Moreover, since $a+1$ is above a, the behavior of the 123-stack is not affected by the presence of $a+1$, meaning that a and $a+1$ can be considered as a single element. As a consequence, the last two elements of $\mathcal{S}^{123}(\pi')$ are $a+1$ and a. Finally, it is easy to realize that $\mathcal{S}^{123}(\pi)$ contains 231 if and only $\mathcal{S}^{123}(\pi')$ contains 231. $\square$

Corollary 4.3. Let π be a permutation of length n and let π' be the permutation (of length $n+k-1$) obtained from π by k-inflating π_1, for some $k \geq 1$. Then π is 123-sortable if and only if π' is 123-sortable.

Proof. This is a direct consequence of the previous corollary, by just iterating the same argument. $\square$

Due to Lemma 4.1 and Corollary 4.3, in order to describe Sort(123) we just need to investigate the sortability of permutations starting with a descent. Denote by $\mathrm{Sort}^{\downarrow}(123)$ the set:

$$\mathrm{Sort}^{\downarrow}(123) = \{\pi \in \mathrm{Sort}(123) : \pi_1 > \pi_2\}.$$

By first characterizing and enumerating $\mathrm{Sort}_n^{\downarrow}(123)$, we can easily recollect the analogous results for Sort(123). Indeed by deflating the prefix of consecutive ascents (if there is one), we can always trace back the 123-sortability of a permutation to another permutation in $\mathrm{Sort}^{\downarrow}(123)$.

Lemma 4.4. Let $\pi \in \mathrm{Sort}_n^{\downarrow}(123)$, with $\pi_1 = k$. Then:

$$\mathcal{S}^{123}(\pi) = n(n-1)\cdots(k+1)(k-1)\cdots 21k.$$

Proof. Let $\mathcal{S}^{123}(\pi) = \gamma_1\gamma_2\cdots\gamma_n$. Clearly $\gamma_n = k$. Now suppose, for a contradiction, that the two elements u and v constitute an ascent in $\mathcal{S}^{123}(\pi)$, with $u < v$ and $v \neq k$. We first show that v precedes u in π. Suppose in fact that this is not the case, and focus on the instant when u is extracted from the 123-stack. Let a be the next element of the input when this happens. Then there are two elements b, c in the 123-stack, with $b < c$ and b above c, such that $abc \simeq 123$. We distinguish two cases.

- $u = b$. In this case, we have $a \neq v$, and so v follows a in π. Therefore $\mathcal{S}^{123}(\pi)$ contains either the subword uav, which is impossible since u and v are supposed to be consecutive in $\mathcal{S}^{123}(\pi)$, or the subword uva, which is impossible too since otherwise $\mathcal{S}^{123}(\pi)$ would contain the pattern 231, contradicting the fact that π is 123-sortable.

- $u \neq b$. In this case, $\mathcal{S}^{123}(\pi)$ would contain the subword ubv, which is impossible, again because u and v would not be consecutive.

Thus we can write π as $\pi = k\pi_2 \cdots v \cdots u \cdots$. Since u and v are consecutive in $\mathcal{S}^{123}(\pi)$, u must enter the 123-stack just above v. This implies, in particular, that $v > \pi_1$, otherwise u, v and π_1 would constitute a forbidden 123 inside the 123-stack. We also notice that, when u enters the 123-stack, at the bottom of the 123-stack there is at least one element $w < \pi_1$ just above π_1. Indeed, either π_2 is still in the 123-stack (and in this case $w = \pi_2$) or π_2 has been forced to exit by some $\tilde{w} < \pi_2 < \pi_1$; in this case, $\tilde{w}$ replaces π_2 just above π_1. Iterating this argument, we get the desired property. Summing up, when u enters the 123-stack, the 123-stack itself contains the elements (from top to bottom) u, v, w, π_1. Now, it must be $u > w$, otherwise $uw\pi_1$ would be a forbidden 123 in the 123-stack. Hence $\mathcal{S}^{123}(\pi)$ contains the subword $uvw \simeq 231$ and π is not 123-sortable, a contradiction. $\square$

Corollary 4.5. Let $\pi \in \mathrm{Sort}_n^{\downarrow}(123)$ and suppose that $\pi_1 < n$. Also, suppose that $\pi_i = n$, for some $i \geq 2$. Then either $\pi_{i-1} = n-1$, if $\pi_1 \neq n-1$, or $\pi_{i-1} = n-2$, if $\pi_1 = n-1$.

Proof. Notice that $i \geq 3$: indeed $i \neq 1$ by hypothesis and $i \neq 2$ since π starts with a descent. The element $\pi_i = n$ enters the 123-stack immediately above π_{i-1}, since pushing the maximum n into the 123-stack can never generate a forbidden pattern 123. Moreover, n and π_{i-1} are extracted from the 123-stack together, since n cannot play the role of the second element in a forbidden pattern inside the 123-stack. Therefore, $\mathcal{S}^{123}(\pi)$ contains the factor $n\pi_{i-1}$. The desired result follows then from Lemma 4.4. $\square$

Corollary 4.6. Let $n \geq 2$. Then the set of permutations of $\mathrm{Sort}_n^{\downarrow}(123)$ starting with n is the set of 213-avoiding permutations of length n that start with n.

Proof. Let $\pi \in \mathrm{Sort}_n^{\downarrow}(123)$, and suppose that $\pi_1 = n$. As soon as n enters the 123-stack, it makes the 123-stack act as a 12-stack for the rest of the permutation. This means, formally, that since n is the maximum of π, from now on the restriction of the 123-stack is triggered if and only if the restriction of a 12-stack that ignores n is triggered. Therefore, by Theorem 3.1, π is 123-sortable if and only if the permutation obtained from π by removing the first element avoids 213, which is in turn equivalent to the fact that π avoids 213. $\square$

A straightforward consequence of Corollary 4.6 is that there are $\mathfrak{c}_{n-1}$ permutations in $\mathrm{Sort}_n^{\downarrow}(123)$ that start with the maximum n. The remaining permutations of $\mathrm{Sort}_n^{\downarrow}(123)$ are precisely those having at least two ltr-maxima. Denote this set by $\mathrm{Sort}_n^{\downarrow}(\geq 2; 123)$. Similarly, denote by $\mathrm{Sort}_n^{\downarrow}(i; 123)$ the set of permutations of $\mathrm{Sort}_n^{\downarrow}(123)$ having exactly i ltr-maxima.

Theorem 4.7. Let $n \geq 3$. There exists a bijection:

$$\varphi : \mathrm{Sort}_{n-1}^{\downarrow}(123) \to \mathrm{Sort}_n^{\downarrow}(\geq 2, 123).$$

Moreover, the restriction of φ to $\mathrm{Sort}_{n-1}^{\downarrow}(i; 123)$ is a bijection between $\mathrm{Sort}_{n-1}^{\downarrow}(i; 123)$ and $\mathrm{Sort}_n^{\downarrow}(i+1; 123)$.

Proof. Let $\pi = \pi_1 \cdots \pi_{n-1} \in \mathrm{Sort}_{n-1}^{\downarrow}(123)$. Let $\varphi(\pi)$ be obtained from π by inserting n:

- either immediately after $n-1$, if $\pi_1 \neq n-1$, or
- immediately after $n-2$, if $\pi_1 = n-1$.

First we show that φ is well defined, that is $\varphi(\pi) \in \mathrm{Sort}_n^{\downarrow}(\geq 2; 123)$. We analyze the two cases in the definition of φ separately.

- If $\pi \in \mathrm{Sort}_{n-1}^{\downarrow}(1; 123)$ (that is $\pi_1 = n-1$), then by Lemma 4.4 we have:

  $$\mathcal{S}^{123}(\pi) = (n-2)(n-3)\cdots 1(n-1).$$

  Now we analyze what happens on input $\varphi(\pi)$ after the first pass through the 123-stack. Remember that the first element of $\varphi(\pi)$ is $n-1$ and that n immediately follows $n-2$; moreover, suppose that $n-2$ is the i-th element of $\varphi(\pi)$. Therefore, the first i elements of π and $\varphi(\pi)$ are equal, and so they are processed exactly in the same way by the 123-stack. In particular, since $n-2$ is the first element of $\mathcal{S}^{123}(\pi)$, when $n-2$ enters the 123-stack, all the previous elements of $\varphi(\pi)$ are still inside the 123-stack. Immediately after $n-2$ enters the 123-stack, n enters the 123-stack as well, since it cannot produce a forbidden pattern. Now we claim that n and $n-2$ exit the 123-stack together. This is trivial if n is the last element of the input. If instead n is not the last element of $\varphi(\pi)$, consider the next element π_{i+1}. Such element cannot enter the 123-stack, otherwise π_{i+1}, $n-2$ and $n-1$ (which is at the bottom of the 123-stack) would constitute a forbidden pattern 123. Thus $n-2$ must exit the 123-stack before π_{i-1} enters it, and this forces n to exit as well. As a consequence of this fact, we have that:

  $$\mathcal{S}^{123}(\varphi(\pi)) = n(n-2)(n-3)\cdots 1(n-1),$$

  which avoids 231. Hence $\varphi(\pi)$ is sortable.

- If $\pi \in \mathrm{Sort}^{\downarrow}_{n-1}(\geq 2; 123)$ (that is $\pi_1 = k \neq n-1$), then by Lemma 4.4 we have:
$$\mathcal{S}^{123}(\pi) = (n-1)(n-2)\cdots(k+1)(k-1)\cdots 21k.$$
Finally, an analogous argument can be used to prove that:
$$\mathcal{S}^{123}(\varphi(\pi)) = n(n-1)(n-2)\cdots(k+1)(k-1)\cdots 21k,$$
and so that $\varphi(\pi)$ is 123-sortable.

To complete the proof we now have to show that φ is a bijection. The fact that φ is injective is trivial. To show that φ is surjective, consider the map $\psi : \mathrm{Sort}^{\downarrow}_{n}(\geq 2; 123) \to \mathrm{Sort}^{\downarrow}_{n-1}(123)$ which removes n from $\alpha \in \mathrm{Sort}^{\downarrow}_{n}(\geq 2; 123)$. Let $\alpha = \alpha_1 \cdots \alpha_n$ and let $i \in \{3, 4, \ldots n\}$ such that $\alpha_i = n$. From Corollary 4.5, we have that either $\alpha_{i-1} = n-1$ (if $\alpha_1 \neq n-1$) or $\alpha_{i-1} = n-2$ (if $\alpha_1 = n-1$). Moreover, Lemma 4.4 implies that:
$$\mathcal{S}^{123}(\pi) = n(n-1)\cdots(k+1)(k-1)\cdots 21k,$$
with $k = \alpha_1 \geq 2$. Therefore, when n enters the 123-stack, all the previous elements are still inside the 123-stack. In particular, at the top of the 123-stack there are n and α_{i-1}. Now notice that, if n is forced to exit the 123-stack, this is due to the fact that there exist j, h, l, with $j < h \leq i$ and $l > i$, such that α_l, α_h and α_j form an occurrence of 123. However, it cannot be $h = i$, since n cannot play the role of the 2 in a 123. Similarly, it cannot be $h = i-1$: in fact, if $\alpha_{i-1} = n-1$, then n and $n-1$ are consecutive in the 123-stack and so they play the same role in any pattern; if instead $\alpha_{i-1} = n-2$, then $\alpha_1 = n-1$ is at the bottom of the 123-stack, and so n and $n-2$ play the same role in any forbidden pattern. As a consequence, $h < i-1$, and so n and α_{i-1} are forced to leave the 123-stack together. This means that basically n does not modify the behavior of the machine, and so:
$$\mathcal{S}^{123}(\psi(\alpha)) = (n-1)(n-2)\cdots(k+1)(k-1)\cdots 21k,$$
that is $\psi(\alpha)$ is 123-sortable, as desired. □

Corollary 4.8. For all $n \geq 3$, $|\mathrm{Sort}^{\downarrow}_{n}(\geq 2; 123)| = |\mathrm{Sort}^{\downarrow}_{n-1}(123)|$.

4.2 Enumeration of Sort(123)

We now use the results proved in Section 4.1 to enumerate $\mathrm{Sort}_n(123)$. Due to Corollary 4.3, Corollary 4.6 and Theorem 4.7, any 123-sortable permutation π which is not the identity permutation can be uniquely constructed as follows:

1. choose $\alpha = \alpha_1\alpha_2\cdots\alpha_k \in \mathfrak{S}_k(213)$, with $\alpha_1 = k \geq 2$;
2. add h new maxima, $k+1,\ldots,k+h$, one at a time, using the bijection φ of Theorem 4.7;
3. add $n-k-h$ consecutive ascents at the beginning, by inflating the first element of the permutation, according to Corollary 4.3.

As an example to illustrate the given construction, let $\pi = 567148923$. By deflating the prefix of initial consecutive ascents, we get the permutation $\pi' = 5146723$; due to Corollary 4.3, π is 123-sortable if and only if π' is 123-sortable. Now, π' is (uniquely) obtained by adding two new maxima to the permutation $\pi'' = 51423$, whose first element is its maximum, according to the bijection of Theorem 4.7. Since π'' avoids 213, we can finally conclude that π is 123-sortable.

Theorem 4.9. For all $n \geq 1$, we have:

$$f_n^{123} = 1 + \sum_{h=1}^{n-1} (n-h)\mathfrak{c}_h.$$

Proof. A permutation $\pi \in \mathrm{Sort}_n(123)$ is either the identity or it is obtained by choosing a permutation α in $\mathfrak{S}_k(213)$ starting with its maximum k (with $k \geq 2$) and then (possibly) adding the remaining $n-k$ elements according to the above construction, i.e. adding new maxima and/or consecutive ascents at the beginning. Concerning α, there are $\mathfrak{c}_{k-1}$ possible choices, thanks to the observation following Corollary 4.6. For the remaining elements, one has to choose, for instance, the number of new maxima to add, which runs from 0 to $n-k$, so that the total number of choices is $n-k+1$. Summing on all possible values of k, we get:

$$f_n^{123} = 1 + \sum_{k=2}^{n} \mathfrak{c}_{k-1} \cdot (n-k+1) = 1 + \sum_{h=1}^{n-1} (n-h)\mathfrak{c}_h,$$

as desired. □

We end this section by computing the generating function $F^{123}(t)$ of Sort(123). As anticipated, we shall exploit the link with pattern-avoiding Schröder paths.

Theorem 4.10. We have:

$$F^{123}(t) = \frac{(1-t)^2}{1-2t+tC(t)}.$$

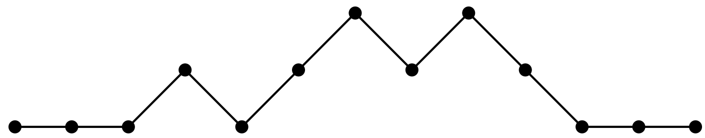

Figure 4.1: The $\mathtt{UH_2D}$-avoiding Schröder path associated to the 123-sortable permutation 567489132. Referring to the notations of Theorem 4.10, we have $L = 56$, $\beta = 7489132$, and so $r = s = 2$. Moreover, $\alpha = 4132$ and the associated Dyck path is UDUUDUDD.

Proof. We start by providing a bijection f between 123-sortable permutations of length n and $\mathtt{UH_2D}$-avoiding Schröder paths of semilength $n - 1$. Given $\pi \in \mathrm{Sort}_n(123)$, decompose it as $\pi = L\beta$, where L is the (possibly empty) initial sequence of consecutive ascents of π, deprived of the last element, and β is the remaining suffix of π. Suppose that L has length r. Now repeatedly remove the maximum from β until the remaining word β' starts with its maximum. Denote with s the number of elements removed this way. Then β' is order isomorphic to a 213-avoiding permutation α of length $k + 1 = n - r - s$, that starts with its maximum. Removing the maximum from α results in another 213-avoiding permutation ρ of length k. We can now describe the Schröder path $f(\pi)$ associated with π: it starts with r double horizontal steps and ends with s double horizontal steps; in the middle, there is the Dyck path of semilength k associated to the 213-avoiding permutation ρ through the bijection described in Example 1.3.

Next, as announced, we express the generating function of Sort(123) by exploiting the bijection f. In fact, the generic Schröder path avoiding $\mathtt{UH_2U}$ either consists of double horizontal steps only (so the generating function is $(1 - t)^{-1}$), or can be obtained by concatenating an initial sequence of double horizontal steps (having generating function $(1-t)^{-1}$) with a non-empty Dyck path (whose generating function is $(C(t)-1)\cdot t$, where $C(t)$ is the generating function of the Catalan numbers and the additional factor t takes into account the removal of the starting maximum from the permutation α above), finally adding a sequence of double horizontal steps (again with generating function $(1 - t)^{-1}$). Summing up, we get:

$$F^{123}(t) = \frac{1}{1-t} + \frac{1}{1-t}\big(t(C(t)-1)\big)\frac{1}{1-t} = \frac{(1-t)^2}{1-2t+tC(t)}.$$

$\square$

Chapter 5

The 132-machine

This chapter, whose paper version is [24], is devoted to the study of the 132-machine. We prove that 132-sortable permutations are enumerated by the binomial transform of Catalan numbers (sequence A007317 in [45]) by first characterizing Sort(132) in terms of avoidance of a classical pattern and a mesh pattern. Then we exploit this result to determine some geometric properties of Sort(132). These ultimately lead to a bijection with the set of 12231-avoiding RGFs, whose enumeration is a corollary of a much more general mechanism proposed by Jelínek and Mansour in [35]. We then exhibit direct combinatorial proofs for the enumeration of some patterns in the same Wilf-equivalence class as 12231, highlighting connections with lattice paths and pattern-avoiding permutations. We enumerate two of these patterns via a bijection with a family of labeled Motzkin paths, which provide a combinatorial interpretation of a beautiful continued fraction for the related counting sequence. Finally, by putting all these pieces together, we obtain an independent proof of the enumeration of Sort(132).

5.1 Characterization of Sort(132)

We start by showing a useful decomposition lemma for the 132-stack.

Lemma 5.1. Let π be a permutation and let $\pi = m_1 B_1 m_2 B_2 \cdots m_t B_t$ be the ltr-min decomposition of π. Then:

1. $\mathcal{S}^{132}(\pi) = \widetilde{B_1}\widetilde{B_2}\cdots\widetilde{B_t} m_t m_{t-1}\cdots m_2 m_1$, where each $\widetilde{B_i}$ is a suitable rearrangement of the elements of B_i.
2. If π is 132-sortable, then $x > y$ for each $x \in B_i$, $y \in B_j$, with $i < j$.

Proof. 1. For each $x \in B_1$, $m_1 x m_2 \simeq 231$, thus every element of B_1 has to be popped from the 132-stack before m_2 enters. After that, we have m_1 and m_2

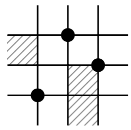

Figure 5.1: The mesh pattern $\mu = (132, \{(0,2), (2,0), (2,1)\})$

on the 132-stack, with $m_1 > m_2$ and m_2 above m_1. Note that they cannot both be part of a 132, therefore m_2 remains in the 132-stack until the end of the sorting process. Similarly, each element of B_2 has to be extracted before m_3 enters, since $m_3 x m_2 \simeq 132$ for each $x \in B_2$. The same argument holds for every m_j with $j \geq 2$.

2. Suppose there are two elements x, y such that $x < y$, $x \in B_i$ and $y \in B_j$, with $i < j$. Then, as a consequence of the previous item, xym_t is an occurrence of 231 in $\mathcal{S}^{132}(\pi)$, which is impossible since π is 132-sortable.

□

Lemma 5.2. Let $\pi \in \mathrm{Sort}_n(132)$ and let $\pi = m_1 B_1 m_2 B_2 \cdots m_t B_t$ be its ltr-min decomposition. Suppose that the next element of the input is $x \in B_i$, for some i. Then the content of the 132-stack when read from bottom to top is:

$$m_1 m_2 \cdots m_i x_1 x_2 \cdots x_s,$$

where $\{x_1, \dots, xs\}$ is a (possibly empty) subset of B_i such that $x_1 < x_2 < \cdots < x_s$.

Proof. The first i ltr-minima $m_1, \dots, m_i$ of π lie at the bottom of the 132-stack, by Lemma 5.1. Then the remaining elements $x_1, \dots, x_s$ of B_i in the 132-stack must be in increasing order from bottom to top, for otherwise, if $x_h > x_\ell$ for some $h < \ell$, then $\mathcal{S}^{132}(\pi)$ would contain $x_\ell x_h m_i \simeq 231$, contradicting the 132-sortability of π. □

Next we provide a characterization of Sort(132) in terms of pattern avoidance. For the rest of this section, denote by μ the mesh pattern $\mu = (132, \{(0,2), (2,0), (2,1)\})$ depicted in Figure 5.1. A permutation π thus contains an occurrence of μ if π contains an occurrence acb of the classical pattern 132 such that:

- every element that precedes a in π is either smaller than b or greater than c;
- every element between c and b in π is greater than b.

Theorem 5.3. We have:

$$\mathrm{Sort}(132) = \mathfrak{S}(2314, \mu).$$

Proof. We start by showing that $\mathrm{Sort}(132) \subseteq \mathfrak{S}(2314, \mu)$. Let $\pi = m_1B_1m_2B_2\cdots m_tB_t$ be the ltr-min decomposition of π. Suppose, for a contradiction, that π contains an occurrence $bcad$ of 2314. When a enters the 132-stack, at least one of b and c, call it x, has already been popped from the 132-stack, otherwise we would get the forbidden pattern $acb \simeq 132$ inside the 132-stack. Hence, by Lemma 5.1, $\mathcal{S}^{132}(\pi)$ contains $xdm_t \simeq 231$, violating the hypothesis that π is 132-sortable. Next suppose that acb is an occurrence of 132 in π. We wish to show that acb is part of an occurrence of either $\mathbf{3}142$, $24\mathbf{1}\mathbf{3}$ or $14\mathbf{2}3$, thus proving that π avoids μ. Let $m(a)$ be the ltr-minimum immediately preceding the block that contains a, or a itself if a is an ltr-minimum. Then $m(a) \le a$ and $m(a)$ exits the 132-stack after b and c (by Lemma 5.1), so c has to be popped before b enters, otherwise $bcm(a)$ would be an occurrence of 231 inside $\mathcal{S}^{132}(\pi)$. We consider the following two cases. Note that $a < b < c$, so b, c are not ltr-minima in π.

- $c \in B_i$ and $b \in B_j$, with $i < j$. In this case, $m_j < m(a) \le a$, hence $acm_jb \simeq 2413$, which is one of the desired patterns.

- c and b are in the same block B_i. First suppose there is an ltr-minimum $m = m_\ell$, with $\ell < i$, such that $b < m < c$; then $m > m(a)$, so m precedes $m(a)$ in π and $macb \simeq 3142$, again one of the listed patterns. Otherwise, suppose that, for every ltr-minimum m, either $m < b$ or $m > c$ and consider the element w that immediately precedes b in π. We wish to show that $w < b$, which will conclude the proof. Suppose, for a contradiction, that $w > b$ and let $x_1, x_2, \ldots, x_s = w$ be the elements on the 132-stack, after w has been pushed, that are not ltr-minima when we read from bottom to top. By Lemma 5.2, we have $x_1 < x_2 < \cdots < x_s$; moreover $x_s = w > b$, so there is a minimum index u such that $x_u > b$. Now observe that, for $\ell > u$, all the elements x_ℓ are popped from the 132-stack before b enters, because $bx_\ell x_u \simeq 132$. We also observe that necessarily $x_u \le c$, otherwise c would already have been popped and $\mathcal{S}^{132}(\pi)$ would contain the pattern $cx_um(a) \simeq 231$. We can now assert that b is pushed onto the 132-stack immediately above x_u. In fact, $x_\ell < b$ for every $\ell < u$; moreover, our hypothesis implies that either $m < b$ or $m > c$ for every ltr-minimum m inside the 132-stack, therefore b cannot be the first element of an occurrence of 231 (read from top to bottom) that involves elements inside the 132-stack. However this results in an occurrence $bx_um(a)$ of 231 in $\mathcal{S}^{132}(\pi)$, which again contradicts the hypothesis that π is 132-sortable.

We have thus proved that $\mathrm{Sort}(132) \subseteq \mathfrak{S}(2314, \mu)$. Next we show the opposite inclusion $\mathfrak{S}(2314, \mu) \subseteq \mathrm{Sort}(132)$. Let $\pi \in \mathfrak{S}(2314, \mu)$. Suppose, for a contradiction, that π is not 132-sortable, that is, $\mathcal{S}^{132}(\pi)$ contains an occurrence bca of 231. Let again $\pi = m_1B_1m_2B_2\cdots m_tB_t$ be the ltr-min decomposition of π. By Lemma 5.1, we have

$$\mathcal{S}^{132}(\pi) = \widetilde{B_1}\widetilde{B_2}\cdots\widetilde{B_t}m_tm_{t-1}\cdots m_2m_1.$$

Since the ltr-minima are popped from the 132-stack in increasing order, neither b nor c can be an ltr-minimum. Suppose that $b \in B_i$ and $c \in B_j$, for some $i \leq j$. If $i < j$, then $m_ibm_jc \simeq 2314$, which is forbidden. Suppose instead that $i = j$ and consider the leftmost ascent $x < y$ in $\widetilde{B_i}$ (indeed there is at least one ascent in $\widetilde{B_i}$, since the elements b, c constitute a noninversion in $\widetilde{B_i}$). There are two possibilities.

- If y comes after x in π then x has to be popped before y is pushed onto the 132-stack. Therefore, when x is popped, there are two elements u, v in the 132-stack, with v above u, such that $uvw \simeq 231$, where w is the next element of the input. If $v \neq x$, then also v is popped after x (for the same reason), but this is a contradiction with the fact that x and y constitute an ascent in $\widetilde{B_i}$. Thus we have $v = x$ and $uxw \simeq 231$, which implies that $w \neq y$ and $uxwy \simeq 2314$ in π, contradicting the assumption that π avoids 2314.
- Suppose instead that y precedes x in π. Observe that y has to be on the 132-stack when x enters, because $\mathcal{S}^{132}(\pi)$ contains the ascent (x, y) (this fact will be frequently used in the sequel). In this situation, m_iyx is an occurrence of 132 in π. We now show that either m_iyx is an occurrence of μ or π contains 2314. If there is an element z that precedes m_i in π such that $x < z < y$ (so that $zm_iyx \simeq 3142$), then z cannot be an ltr-minimum. In such a case, in fact, by Lemma 5.1, z would be in the 132-stack below y when x is pushed, but $xyz \simeq 132$, which is impossible due to the restriction of the 132-stack. Instead, if $z \in B_\ell$ for some $\ell < i$, then $m_\ell zm_iy \simeq 2314$. Therefore we can assume that every element that precedes m_i in π is either smaller than x or greater than y. Finally, suppose that there is an element z between y and x in π such that $z < x$, which gives an occurrence m_iyzx of either 2413 or 1423. Then, since y is still in the 132-stack when x is pushed and z precedes x in π, z enters the 132-stack above y, and so $\widetilde{B_I}$ contains either $x\ldots z\ldots y$ or $z\ldots x\ldots y$, with $z < x$. However, both cases give a contradiction, because (x, y) is the first ascent in $\mathcal{S}^{132}(\pi)$.

□

Due to the presence of the mesh pattern μ (and in accordance with Theorem 3.8), the set $\mathfrak{S}(2314, \mu)$ is not a permutation class. For instance, the 132-sortable permutation 2413 contains the non 132-sortable pattern 132.

5.2 A grid decomposition for Sort(132)

In the previous section we have proved that $\mathrm{Sort}(132) = \mathfrak{S}(2314, \mu)$, obtaining a precise description (in terms of generalized pattern avoidance) of 132-sortable permutations. However, we are still not able to enumerate Sort(132) directly. In this section we will thus investigate its geometric structure by refining the ltr-minima decomposition as follows. Let π be a permutation of length n with t ltr-minima and let $\pi = m_1 B_1 m_2 B_2 \cdots m_t B_t$ its ltr-min decomposition. Then:

- for $j \geq 1$, the j-th *vertical strip* of π is B_j;
- for $i \geq 1$, the i-th *horizontal strip* of π is $H_i = \{x \in \pi : m_i < x < m_{i-1}\}$, where $m_0 = +\infty$.
- for any two indices i, j, the *cell* of indices i, j of π is $C_{i,j} = H_i \cap B_j$ (note that $C_{i,j}$ is empty when $i > j$).
- the *core* of π is $\mathrm{core}(\pi) = B_1 B_2 \cdots B_k$, obtained from π by removing the ltr-minima.

From now on, the content of each $B_j, H_i, C_{i,j}$ will be regarded as a permutation. As an example, consider the permutation $\pi = 13\,14\,15\,10\,12\,6\,7\,8\,11\,9\,3\,1\,4\,5\,2$. Then (see Figure 5.2):

- π has six ltr-minima, namely $13, 10, 6, 3, 1$;
- the vertical strips of π are $B_1 = 14\,15 \simeq 1\,2$, $B_2 = 12 \simeq 1$, $B_3 = 7\,8\,11\,9 \simeq 1\,2\,4\,3$, $B_4 = \emptyset$ and $B_5 = 4\,5\,2 \simeq 2\,3\,1$;
- the horizontal strips of π are $H_1 = 14\,15 \simeq 1\,2$, $H_2 = 12\,11 \simeq 2\,1$, $H_3 = 7\,8\,9 \simeq 1\,2\,3$, $H_4 = 4\,5 \simeq 1\,2$ and $H_5 = 2 \simeq 1$;
- the nonempty cells of π are $C_{1,1} = 14\,15 \simeq 1\,2$, $C_{2,2} = 12 \simeq 1$, $C_{2,3} = 11 \simeq 1$, $C_{3,3} = 7\,8\,9 \simeq 1\,2\,3$, $C_{4,5} = 4\,5 \simeq 1\,2$ and $C_{5,5} = 2 \simeq 1$;
- the core of π is $\mathrm{core}(\pi) = 14\,15\,12\,7\,8\,11\,9\,4\,5\,2 \simeq 9\ 10\,8\,4\,5\,7\,6\,2\,3\,1$.

The terminology introduced above refers to the graphical representation of π, as illustrated in Figure 5.2.

In what follows we prove that the requirement of being 132-sortable imposes precise constraints on the grid structure of a permutation: both the content of strips and cells and the relative position of non-empty cells are affected.

Lemma 5.4. Let π be a 132-sortable permutation and suppose that the cell $C_{i,j}$ is nonempty, for some i, j. Then the cell $C_{u,v}$ is empty for each pair of indices (u, v) such that $u < i$ and $v > j$[1].

[1] That is when $C_{u,v}$ is strictly northeast of $C_{i,j}$ (see Figure 5.3).

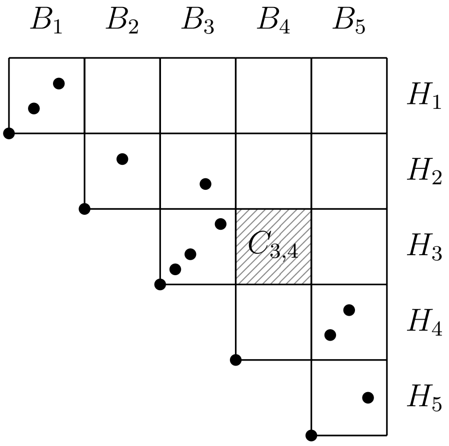

Figure 5.2: The grid decomposition of $\pi = 13\,14\,15\,10\,12\,6\,7\,8\,11\,9\,3\,1\,4\,5\,2$. The image of π under the bijection of Theorem 5.13 is the RGF $\eta(\pi) = 111223332345445$.

Proof. Suppose there are two elements $x \in C_{i,j}$ and $y \in C_{u,v}$ such that $u < i$ and $v > j$. Then $m_i x m_v y \simeq 2314$, which is impossible due to Theorem 5.3. $\square$

Lemma 5.5. Let π be a 132-sortable permutation and suppose that the cell $C_{i,j}$ contains an inversion $x > y$, where x precedes y in $C_{i,j}$. Then there is an element z between x and y in π such that $z < m_i$.

Proof. We refer to Figure 5.3 for a description of the statement of the lemma. For x and y as above, we have $m_i xy \simeq 132$. In particular, x and y are in the same cell $C_{i,j}$ and m_i is the corresponding ltr-minimum, hence every element w preceding m_i in π is greater than x (because $w > m_{i-1}$ and $x < m_{i-1}$). Therefore, as a consequence of Theorem 5.3, there exists an element z between x and y in π such that $z < y$. If $z < m_i$, then we are done. Otherwise, if $z > m_i$, we can repeat the same argument using the occurrence $m_i xz$ of 132, in which we have replaced y with the element z that comes strictly before y in π; continuing in this way we eventually find an element of π with the desired property. $\square$

Proposition 5.6. If π is 132-sortable, then $C_{i,j} \in \mathfrak{S}(132, 213)$, for every i, j.

Proof. Suppose that $C_{i,j}$ contains an occurrence acb of 132. By Lemma 5.5, there exists an element z between c and b in π such that $z < m_i$. In particular, $m_i azb \simeq 2314$, which is a contradiction since π is 132-sortable (by Theorem 5.3). On the other hand, if $C_{i,j}$ contains an occurrence bac of 213, then (b, a) is an inversion in the cell $C_{i,j}$ and therefore, again by Lemma 5.5, there is an element z between b and a in π with $z < m_i$ and $m_i bzc \simeq 2314$, a contradiction. $\square$

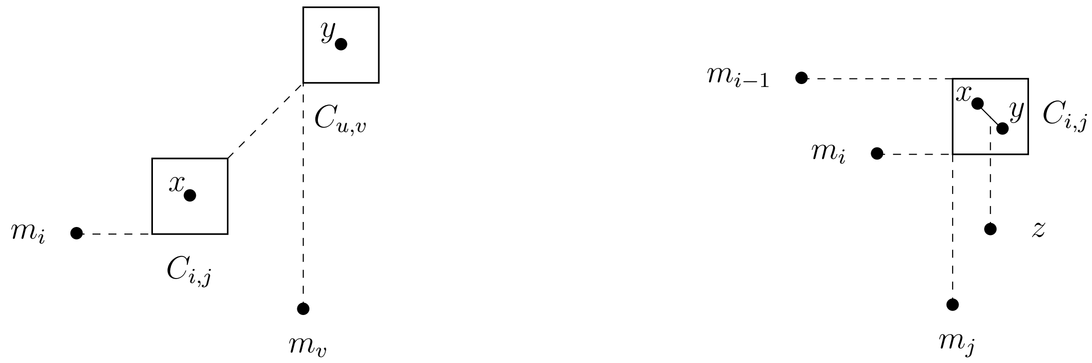

Figure 5.3: The constructions of Lemma 5.4, on the left, and of Lemma 5.5, on the right.

Proposition 5.7. If π is a 132-sortable permutation, then $H_i \in \mathfrak{S}(132, 213)$, for every i.

Proof. This is a consequence of Lemma 5.1 and Proposition 5.6. $\square$

Proposition 5.8. If π is 132-sortable, then $\mathrm{core}(\pi) \in \mathfrak{S}(213)$.

Proof. Suppose that π contains an occurrence bac of 213 that does not involve any ltr-minimum and suppose that $b \in C_{i,j}$ for some i, j. Note that $b < c$, so, by Lemma 5.1, b and c must belong to the same vertical strip B_j. Now, if $a \in C_{\ell,j}$, with $\ell > i$, then $m_i bac \simeq 2314$, which is a contradiction, since π is 132-sortable. Therefore we must have $a \in C_{i,j}$. This results in an occurrence $m_i ba$ of 132, with b and a both in the cell $C_{i,j}$; thus, by Lemma 5.5, there is an element z between b and a in π such that $z < m_i$ and $m_i bzc \simeq 2314$, which is again a contradiction. $\square$

The results proved so far in this section provide necessary conditions that a permutation has to satisfy in order to be 132-sortable. Now, since the prefix of a σ-sortable permutation is always σ-sortable (see Lemma 2.3), if we remove the last element from a 132-sortable permutation we get another 132-sortable permutation of length one less. Equivalently, every 132-sortable permutation is obtained from a 132-sortable permutation (of length one less) by inserting a new rightmost element, and suitably rescaling the remaining ones. Our next goal is to understand which integers are allowed for such an insertion, so to obtain a recursive construction for the set Sort(132). For example, since the insertion of a new minimum can never create either 2314 or μ, by Theorem 5.3, such an insertion is always allowed. In all the other cases, we need to satisfy the requirements of Lemma 5.4 and Propositions 5.7 and 5.8.

Let π be a 132-sortable permutation with t ltr-minima. Suppose we insert a new rightmost element in a cell $C_{i,t}$ of the last vertical strip. By Proposition 5.7, any horizontal strip H_i in a 132-sortable permutation avoids both 132 and 213, that is H_i is co-layered. Therefore, if we wish to obtain a new co-layered permutation by inserting a new rightmost element, there are exactly two possibilities:

1. `min`: to insert a new minimum in $C_{i,t}$ (which is also a new minimum of the horizontal strip H_i);

2. `cons`: to create a consecutive ascent[2] in the two final positions of $C_{i,t}$.

We formalize this construction by introducing the notion of active cell. Let π be a 132-sortable permutation with t ltr-minima. For $i \geq 1$, the cell $C_{i,t}$ is said to be *active* if both the following conditions are satisfied:

(i) $C_{u,v}$ is empty for each u, v such that $u > i$ and $v < t$;

(ii) inserting a new rightmost element according to `min` does not create an occurrence of 213 in core(π).

Thanks to condition (i), we can equivalently express condition (ii) by saying that the permutation $\bigcup_{j \geq i+1} C_{j,t}$ is increasing. Moreover, as a consequence of Lemma 5.4 and Proposition 5.8, if we insert a new rightmost element in a cell $C_{i,t}$ that is not active, then we get a non 132-sortable permutation. Othwerise, if $C_{i,t}$ is active, we wish to show that exactly one of the operations `min` and `cons` returns a 132-sortable permutation. Let us consider two cases, according to whether $C_{i,t}$ is empty or not.

Proposition 5.9. Let $\pi = \pi_1 \cdots \pi_n$ be a 132-sortable permutation with t ltr-minima and let $C_{i,t} = \gamma_1 \cdots \gamma_k$ be a nonempty active cell of π. Let $x = \pi_n$ and suppose $x \in C_{\ell,t}$. Then:

1. by performing `min` on $C_{i,t}$ we get a 132-sortable permutation π' if and only if $\ell > i$;

2. by performing `cons` on $C_{i,t}$ we get a 132-sortable permutation π' if and only if $\ell \leq i$.

Proof. 1. Suppose that $\ell < i$ and we want to insert a new rightmost element γ_{k+1} into $C_{i,t}$ according to `min`. Assume, for a contradiction, that the resulting permutation π' is 132-sortable. The elements γ_k and γ_{k+1} form an inversion in $C_{i,t}$, so by Lemma 5.5 there exists an element z between γ_k and γ_{k+1} in π such that $z < m_i$. Hence $m_i \gamma_k z x \simeq 2314$, which contradicts the assumption that π is 132-sortable. Instead, if $\ell = i$, that is, $\gamma_k = x = \pi_n$,

[2]recall that an ascent $\pi_i < \pi_{i+1}$ is consecutive if $\pi_{i+1} = \pi_i + 1$.

then $\gamma_k\gamma_{k+1}$ is an inversion inside $C_{i,t}$ such that γ_k and γ_{k+1} are adjacent in π. This implies that π is not 132-sortable (again as a consequence of Lemma 5.5).

Conversely, suppose that $\ell > i$ and γ_{k+1} is inserted into $C_{i,t}$ according to `min`. By Theorem 5.3, $\pi \in \mathfrak{S}(2314, \mu)$, so we just have to show that the permutation π' obtained after the insertion still avoids the two forbidden patterns. If γ_{k+1} plays the role of the 2 in an occurrence of 132, say $ac\gamma_{k+1}$, then we have either $acx\gamma_{k+1} \simeq 1423$ or $acx\gamma_{k+1} \simeq 2413$, which means that the selected occurrence of 132 is not an occurrence of the mesh pattern μ. Otherwise, suppose there is an occurrence $bca\gamma_{k+1}$ of 2314 in π'. If $m_t = 1$ precedes c in π, then $ca\gamma_k \simeq 213$ in $\operatorname{core}(\pi)$, contradicting Proposition 5.8. On the other hand, if m_t follows c in π, then $c \in B_j$, for some $j < t$, and $\gamma_k \in B_t$, with $c < \gamma_k$, contradicting Lemma 5.1.

2. Suppose we insert γ_{k+1} into $C_{i,t}$ according to `cons` and $\ell > i$. Then $\gamma_k x \gamma_{k+1}$ is an occurrence of 213 in $\operatorname{core}(\pi')$, hence π' is not 132-sortable, due to Proposition 5.8, as desired.

   Conversely, suppose that $\ell < i$ and we insert γ_{k+1} into $C_{i,t}$ according to `cons`; this means that $\gamma_{k+1} = \gamma_k + 1$. The resulting permutation π' does not contain an occurrence $bcad$ of 2314 with $\gamma_{k+1} = d$, for otherwise $bcax$ would be an occurrence of 2314 in π, contradicting the hypothesis that π is 132-sortable. On the other hand, suppose there are two elements a, c in π such that $ac\gamma_{k+1}$ is an occurrence of 132. We now prove that $ac\gamma_{k+1}$ is not an occurrence of the mesh pattern μ by distinguishing two cases.

   If $c > m_{i-1}$ (note that $i > \ell$, so m_{i-1} exists), then $a < \gamma_{k+1} < m_{i-1}$, so m_{i-1} precedes a in π (because $a < m_{i-1}$ and m_{i-1} is an ltr-minimum) and $m_{i-1}ac\gamma_{k+1}$ would be an occurrence of 3142.

   Instead, if $c < m_{i-1}$, then c is not an ltr-minimum, because $a < c$ precedes c; moreover, c is in $C_{i,t}$, since $c < m_{i-1}$ and $c > \gamma_{k+1}$, hence $c\gamma_k x$ is an occurrence of 213 in $\operatorname{core}(\pi)$, which is impossible due to Proposition 5.8.

   Finally, if $\ell = i$, then $x = \gamma_k$, $\gamma_{k+1} = \gamma_k + 1$ and they are adjacent in π', so γ_{k+1} is neither part of an occurrence of 2314 nor of μ, since otherwise γ_k would be as well, contradicting the hypothesis that π is 132-sortable.

$\square$

When $C_{i,t}$ is empty, the only possibility is to try to perform `min` (since `cons` does not make sense). Next we show that this is always allowed.

Proposition 5.10. Let $\pi = \pi_1 \cdots \pi_n$ be a 132-sortable permutation with t ltr-minima and let $C_{i,t}$ be an empty active cell of π. Let π' be the permutation obtained from π by inserting a new rightmost element y in $C_{i,t}$ according to `min`. Then π' is 132-sortable.

Proof. By Theorem 5.3 we have that $\pi \in \mathfrak{S}(2314, \mu)$ and we want to prove that $\pi' \in \mathfrak{S}(2314, \mu)$ as well. Suppose there are three elements b, c, a in π such that $bcay \simeq 2314$. Since $c > b$, the element c is not an ltr-minimum of π. Suppose that $c \in C_{u,v}$, for some u, v. If a is an ltr-minimum, then of course $v < t$, and we have also $u > i$, because y is the minimum of its horizontal strip and $y > c$. This would imply that $C_{u,v}$ is a nonempty cell, with $u > i$ and $v < t$, which is impossible since $C_{i,t}$ is active. Otherwise, if a is not an ltr-minimum, then $cay \simeq 213$ in core(π'), which again contradicts the assumption that $C_{i,t}$ is active.

Next, in order to prove that π' does not contain the mesh pattern μ, suppose there are two elements a, c in π such that $acy \simeq 132$ and suppose $c \in B_j$, for some $j \leq t$. If $j < t$, then $acm_t y$ is an occurrence of 2413, as desired. Otherwise, if $j = t$, we have that $c \in C_{\ell,t}$, for some $\ell < t$, because $C_{i,t}$ is empty before we insert y; moreover, m_ℓ precedes a in π, because $m_\ell > y$ and $a < y$. Thus $m_\ell acy \simeq$ 3142, as desired. □

Corollary 5.11. Let π be a 132-sortable permutation. Then, for every active cell of π, exactly one of `min` and `cons` generates a 132-sortable permutation.

As a consequence of Propositions 5.9 and 5.10, every 132-sortable permutation can be constructed inductively by repeatedly inserting a new rightmost element either as a new minimum or by performing `min` and `cons`, according to the rules of Propositions 5.9. In particular, given a 132-sortable permutation π with k active cells, then $k + 1$ 132-sortable permutations are produced this way (one for each active cell and one when the new minimum is inserted). Using the generating tree terminology, these are the *children* of π. Understanding the distribution of active cells of 132-sortable permutations would lead to a generating tree for Sort(132), which could be used directly to find its enumeration. So far we were not able to fulfill this task, which is left as an open problem.

Open Problem 5.1. Given $n \geq 1$ and $k \geq 0$, compute the number of 132-sortable permutations with k active cells. Moreover, given a 132-sortable permutation π with k active cells, compute the number of active cells of each child of π.

Instead of using the generating tree approach, we wish to exploit the grid structure of 132-sortable permutations in order to determine a bijection with a class of pattern-avoiding RGFs, ultimately obtaining the desired enumeration of Sort(132).

Let $\pi = \pi_1 \cdots \pi_n$ be a permutation with t ltr-minima $m_1, \dots, m_t$ and set $m_0 = +\infty$. Define the map η by setting $\eta(\pi) = r_1 \cdots r_n$, where $r_i = j$ if $m_j \leq \pi_i < m_{j-1}$. An alternative description of $\eta(\pi)$ is the following: scan the permutation π from left to right and record the index of the horizontal strip that contains the current element, including the ltr-minima in the corresponding strips. An example of this construction is illustrated in Figure 5.2. It is easy to realize that η is defined for any permutation and that $\eta(\pi)$ is a RGF. The next theorem asserts that if we restrict to 132-sortable permutations, then η is a bijection between $\mathrm{Sort}_n(132)$ and $\mathfrak{RGF}_n(12231)$. First a useful lemma concerning pattern avoidance on RGFs.

Lemma 5.12. Let $w = w_1 w_2 \cdots w_k$ be a sequence of positive integers. Let $w' = \mathrm{std}(w)$ be the standardization[3] of w and suppose that $w'_1 = k$, for some $k \geq 1$. Let R be a RGF. Then $w' \leq R$ if and only if $12\ldots(k-1)w' \leq R$.

Theorem 5.13. The map η defined above is injective and the image of $\mathrm{Sort}_n(132)$ through η is $\mathfrak{RGF}(12231)$.

Proof. By Lemma 5.12, we have $\mathfrak{RGF}(12231) = \mathfrak{RGF}(2231)$. We start by proving that, for each 132-sortable permutation π, $\eta(\pi)$ avoids 2231. Suppose, on the contrary, that $\eta(\pi)$ contains an occurrence $r_{i_1} r_{i_2} r_{i_3} r_{i_4}$ of 2231. Consider the leftmost occurrence r_j of the integer r_{i_1} in π (note that $j \leq i_1$). Then r_j corresponds through η to the ltr-minimum of the horizontal strip of index r_{i_1} in π. Hence the elements $\pi_j \pi_{i_2} \pi_{i_3} \pi_{i_4}$ form an occurrence of 2314 in π[4], which contradicts Theorem 5.3.

That η is injective on $\mathrm{Sort}_n(132)$ is a consequence of Corollary 5.11. Moreover, using the construction of Proposition 5.9, we will show that $\eta(\mathrm{Sort}_n(132)) = \mathfrak{RGF}(2231)$. Given a RGF $R = r_1 r_2 \cdots r_n$, construct the permutation π_R by scanning R from left to right and, when the current element is r_ℓ, insert a new rightmost element π_ℓ in the following way (suitably rescaling the previous elements when necessary):

- when r_ℓ is the first occurrence of an integer in R then $\pi_\ell = 1$;
- otherwise, π_ℓ is inserted in the horizontal strip H_{r_ℓ}, according to the rules of Proposition 5.9.

We now wish to prove that, if the RGF R avoids 2231, then π_R is a 132-sortable permutation such that $\eta(\pi_R) = R$. It is easy to see that $\eta(\pi_R) = R$, as a direct consequence of the definition of η. Since insertions inside active cells are always allowed, what remains to be shown is that each element is in fact inserted into an active cell. We now argue by contradiction, and suppose that y is the first element that is inserted into a nonactive cell $C_{i,j}$. According to the definition of an active cell, there are two cases to consider.

[3]Recall that w' is obtained by replacing all the occurrences of the smallest integer of w with 1, all the occurrences of the second smallest integer with 2 and so on.

[4]Note that the value order between elements of π coded by distinct values in $\eta(\pi)$ is the reverse of their order in $\eta(\pi)$.

Pattern p	**Formula**	**OEIS**
12123, 12132, 12134, 12213 12231, 12234, 12312, 12321 12323, 12331, 12332	$\lvert\mathfrak{RGF}_n(p)\rvert = \sum_{k=0}^{n-1} \binom{n-1}{k} \mathfrak{c}_k$	A007317

Table 5.1: The eleven patterns of the Wilf-class containing 12231.

1. If there exists a nonempty cell $C_{u,v}$, with $u > i$ and $v < j$, then, given any $x \in C_{u,v}$, the elements of R corresponding to $m_u x m_j y$ form an occurrence of 2231, which is forbidden.
2. Suppose that inserting a new rightmost element according to `min` creates an occurrence bay of 213 that does not involve any ltr-minima. Let H_u be the horizontal strip that contains b and let H_v be the horizontal strip that contains a. Note that $v \geq u > i$. If $v > u$, then the elements corresponding to $m_u bay$ in R form an occurrence of 2231, which is again a contradiction. On the other hand, if $v = u$, then a belongs to the same horizontal strip of b, so, since $a < b$, a was inserted according to `min`. Therefore, by Proposition 5.9 and our choice of y, the element a' that precedes a in core(π) belongs to H_w, for some $w > u$. As a consequence, the elements $m_u ba'c$ correspond to an occurrence of 2231 in R, which is impossible.

□

Corollary 5.14. For every natural number n, we have:

$$|\mathrm{Sort}_n(132)| = |\mathfrak{RGF}_n(12231)|.$$

The enumeration of $\mathfrak{RGF}(12231)$ is an immediate consequence of the results proved in [35], where the authors determine the Wilf-equivalence class of 12231 (see Table 5.1). Amongst the Wilf-equivalent patterns, 12332 can be easily enumerated. Indeed 1221-avoiding RGFs are enumerated by the Catalan numbers (see again [35]). Moreover, as a result of Theorem 31 in [35], we immediately obtain that:

$$|\mathrm{Sort}_n(132)| = \sum_{k=0}^{n-1} \binom{n-1}{k} \mathfrak{c}_k,$$

that is sequence A007317 in [45].

5.3 Combinatorial proofs for pattern-avoiding restricted growth functions

The problem of enumerating 132-sortable permutations has been solved in the previous section by means of a bijection η between Sort(132) and $\mathfrak{RGF}(12231)$. The enumeration of $\mathfrak{RGF}(12231)$ is a corollary of the (much more general) theory developed by Jelínek and Mansour in [35]. However, although η has a neat description ($\eta(\pi)$ records the index of the horizontal strip that contains each element of π, from left to right), it is not enough to have a clear understanding of why Sort(132) is enumerated by the binomial transform of Catalan numbers.

We choose to devote this section to a deeper investigation on the combinatorics underlying some related sets of pattern-avoiding RGFs, aiming to find a more transparent connection with 132-sortable permutations. Ideally, we would like to provide a link between Sort(132) and some combinatorial objects that immediately reveals why this counting sequence arises. We start by showing a (presumably) new bijection between $\mathfrak{RGF}_n(1221)$ and the set $\mathcal{D}_n$ of Dyck paths of semilength n. Then we define new bijections between $\mathfrak{RGF}(y)$, with y pattern in the Wilf-equivalence class of 12231, and other families of combinatorial objects, such as labeled Motzkin paths and pattern-avoiding permutations. Finally, we obtain a bijective argument that clearly justifies the enumeration of Sort(132) by showing a bijection between $\mathfrak{RGF}(12231)$ and $\mathfrak{RGF}(12321)$.

5.3.1 Pattern 1221

The following lemma can be found in [19].

Lemma 5.15 ([19], Lemma 6.2)**.** Let R be a RGF. Then $R \in \mathfrak{RGF}(1221)$ if and only if the subword $w(R)$ obtained by removing the first occurrence of each letter in R is weakly increasing.

An immediate consequence of Lemma 5.15 is the following.

Corollary 5.16. Let $R = r_1 \cdots r_n \in \mathfrak{RGF}(1221)$ and $M = \max(R)$. If R has no repeated elements let $t = 1$; otherwise let t be the maximum among repeated elements of R. Then $r_1 \cdots r_n j \in \mathfrak{RGF}(1221)$ if and only if $t \le j \le M + 1$.

Using again the language of generating trees, we say that an integer j is an *active site* of the RGF $R \in \mathfrak{RGF}(1221)$ if by appending j at the end of R we get another RGF in $\mathfrak{RGF}(1221)$, which is said to be a *child* of R. The set of active sites of R is the interval $\{t, t+1, \dots, M, M+1\}$ due to Corollary 5.16. Thus R has $M + 1 - t + 1$ active sites, where M and t are defined as in the corollary.

Now, recall from section 1.4 that a *double rise* in a Dyck path is an occurrence of the consecutive pattern `UU`.

Theorem 5.17. There is a bijection $\psi : \mathfrak{RGF}_n(1221) \to \mathcal{D}_n$, such that the maximum of $R \in \mathfrak{RGF}_n(1221)$ equals one plus the number of double rises in the path $\psi(R)$. As a consequence, denoting by $f_{n,k}$ the number of elements in $\mathfrak{RGF}_n(1221)$ whose maximum is k, we get that $f_{n,k} = \mathfrak{n}_{n,k}$, where $\mathfrak{n}_{n,k}$ is the (n,k)-th Narayana number.

Proof. Recall from Example 1.4 that every Dyck path $\tilde{P}$ of semilength $n+1$ is obtained (in a unique way) from a Dyck path P of semilength n by inserting a peak `UD` either before a `D`-step in the last descending run of P or after the last `D`-step. This construction gives rise to a well known generating tree for Dyck paths, such that the number of active sites of a path P is $k+1$, where k is the length of the last descending run of P. The path $\tilde{P}$ is therefore a child of P in the associated generating tree. Our goal is to define (in a recursive fashion) a bijection α between the generating tree of $\mathfrak{RGF}(1221)$ and the generating tree of Dyck paths. In other words, we wish to show that α is a bijection preserving both the size (that is, a RGF of length n is mapped to a Dyck path of semilength n) and the number of active sites.

We start by setting $\alpha(1) =$ `UD`. Note that 1 has two active sites, since the children of 1 are 11 and 12. The path `UD` has two active sites as well, since its children are `UUDD` and `UDUD`. Now let $R = r_1 \cdots r_n$ and $\alpha(R) = p_1 \cdots p_{2n}$, for some $n \geq 1$. Suppose that the number of active sites of both R and $\alpha(R)$ is k. Let $M = \max(R)$ and let t be the maximum element of R that is not an ltr-maximum of R. By Corollary 5.16, the active sites of R form the interval $\{t, t+1, \ldots, M, M+1\}$, with $M+1-t+1 = k$ by hypothesis. Moreover, the length of the last descending run of $\alpha(R)$ is $k-1$. We shall define α on the children of both R and $\alpha(R)$, and show that the number of active sites is still preserved.

- The child of R corresponding to the active site M is mapped to the path obtained from $\alpha(R)$ by inserting a new peak `UD` immediately after the last `D`-step of $\alpha(R)$. Here the active sites of the resulting sequence are $M+1-M+1 = 2$. The same holds for the resulting Dyck path, since the length of its last descending run is 1.

- For $i = 1, \ldots, M-t$, the child of R corresponding to the active site $t+i-1$ is mapped to the path obtained from $\alpha(R)$ by inserting a new peak `UD` immediately after the i-th `D` step of the last descending run. Then the number of active sites of the resulting RGF is then $(M+1)-(t+i-1)+1 = M-t-i3$, which is equal to one plus the length of the last descending run of the resulting path, that is $(M+1-t)-i+1$.

- Finally, the child of R corresponding to the active site $M+1$ is mapped to the path obtained from $\alpha(R)$ by inserting a new peak `UD` immediately before the first $\mathtt{D}$-step of the last descending run of $\alpha(R)$. In this case the number of active sites of the resulting RGF is $M+2-t+1=k+1$. Moreover, the number of active sites of the resulting path is also $k+1$, since the length of its maximal suffix of $\mathtt{D}$-steps is increased by one with respect to $\alpha(R)$.

Therefore α is a bijection between the two generating trees, as desired. To conclude, observe that the number of double rises in $\alpha(R)$ is equal to $\max(R)-1$. Indeed, by definition of α, each double rise in $\alpha(R)$ corresponds to the first occurrence of an integer in R, except for the first occurrence of 1 (which does not create a double rise). It is well known (see for example [29]) that the number of Dyck paths of semilength n with $k-1$ double rises is equal to $\mathfrak{n}_{n,k}$, which gives the desired equality $f_{n,k}=\mathfrak{n}_{n,k}$. □

Corollary 5.18. Let $n \geq 0$. Denote by $g(n,k)$ the number of elements in $\mathfrak{RGF}_n(12332)$ whose maximum is k, for $1 \leq k \leq n$. Then:

$$g(n+1,k+1)=\sum_{j=k}^{n}\binom{n}{j}\mathfrak{n}_{j,k}.$$

Proof. As observed in [35], every 12332-avoiding RGF of length $n+1$ can be obtained by choosing $n-j$ positions for the 1s (except for the first 1, which is fixed) and then choosing a RGF $R \in \mathfrak{RGF}_j(1221)$ for the remaining j spots (where the elements of R incremented by one will be inserted). In particular, if the maximum of R is k, then the resulting RGF has maximum $k+1$. So, as a consequence of Theorem 5.17, we have $g(n+1,k+1)=\sum_{j=k}^{n}\binom{n}{j}\mathfrak{n}_{j,k}$. □

In the following sections (Proposition 5.21 and Theorem 5.26), we provide a bijection between 12231- and 12321-avoiding RGFs in order to prove that 132-sortable permutations, according to the number of their ltr-minima, are enumerated by the formula in Corollary 5.18. A direct proof of this fact is still to be found.

Open Problem 5.2. Prove directly (that is, without using a bijection involving different objects) that the number of 132-sortable permutations of length $n+1$ with $k+1$ left-to-right minima is equal to $\displaystyle\sum_{j=k}^{n}\binom{n}{j}\mathfrak{n}_{j,k}$.

5.3.2 Patterns 12323 and 12332

Consider the ordinary generating function of 132-sortable permutations:

$$F^{132}(t)=\sum_{n\geq 0}\left(\sum_{k=0}^{n-1}\binom{n-1}{k}\mathfrak{c}_k\right)t^n$$

Then $F^{132}(t)$ can be expressed using the following continued fraction (see, for example, [9, 32]):

$$F^{132}(t) = \cfrac{1}{1 - 2t - \cfrac{t^2}{1 - 3t - \cfrac{t^2}{1 - 3t - \cfrac{t^2}{1 - 3t - \dots}}}}$$

Labeled Motzkin paths provide a neat combinatorial interpretation for the above continued fraction, via Flajolet's general correspondence [32]. The n-th term of the sequence $\{|\mathrm{Sort}_{n+1}(132)|\}_n$ is equal to the number of Motzkin paths of length n such that:

- each horizontal step at height zero has two types of labels ℓ_0, ℓ_1;
- each horizontal step at height at least one has three types of labels ℓ_0, ℓ_1, ℓ_2.

Denote by $\mathcal{M}_n^{lab}$ the set of such labeled Motzkin paths of length n. We shall define a bijection β from $\mathcal{M}_n^{lab}$ to $\mathfrak{RGF}_{n+1}(12323)$ by scanning a Motzkin path from left to right and suitably intepreting each labeled step. We use an auxiliary stack Δ, which is initialized as the empty stack. Let $P \in \mathcal{M}_n^{lab}$. Start by setting $R = 1$. Then, if L is the label of the currently scanned step, append a new rightmost element to R according to the following rules:

- if $L = \mathtt{U}$, then append a new strict maximum M and push M onto Δ;
- if $L = \mathtt{D}$, then append $\mathrm{top}(\Delta)$ and pop it from Δ;
- if $L = \ell_0$, then append a new strict maximum (without pushing it onto Δ);
- if $L = \ell_1$, then append 1;
- if $L = \ell_2$, then append $\mathrm{top}(\Delta)$ (without popping it from Δ).

Equivalently, U corresponds to the first occurrence of a letter x that appears at least twice in R, D to the last occurrence of such a letter, and ℓ_2 to an occurrence of such an x that is neither the first nor the last. Moreover, the label ℓ_0 corresponds to an element $x \neq 1$ appearing only once and the label ℓ_1 corresponds to the element 1. An example of this construction is illustrated in Figure 5.4.

We can also express the correspondence between the labels of P and $R = \beta(P)$ in terms of properties of the set partition associated to R. If B is a block of cardinality at least two in such a partition and B does not contain 1, then U, D and ℓ_2

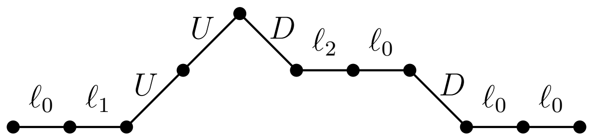

Figure 5.4: The labeled Motzkin path corresponding to the RGF $R = 12134435367$ via the bijection β of Theorem 5.19. The set partition associated to R is $13|2|479|56|8|10|11$.

correspond, respectively, to the least, the largest and any of the remaining elements of the block. Moreover, ℓ_0 corresponds to a singleton block not containing 1 and ℓ_1 corresponds to the elements of the block containing 1. At each step of the construction of R, the auxiliary stack Δ keeps track of the currently open blocks in the corresponding partition (that is those blocks that have not yet received all their elements).

Theorem 5.19. The map β is a bijection between $\mathcal{M}_n^{lab}$ and $\mathfrak{RGF}_{n+1}(12323)$.

Proof. It is straightforward to see that β is injective and that $\beta(P)$ is a RGF for every $P \in \mathcal{M}_n^{lab}$. Since $|\mathcal{M}_n^{lab}| = |\mathfrak{RGF}_n(12323)|$, we only need to show that $\beta(P)$ avoids 12323, for each $P \in \mathcal{M}_n^{lab}$. Suppose, for a contradiction, that $abcb'c'$ is an occurrence of 12323 in $\beta(P)$. This implies, of course, that $b, c \neq 1$. Without loss of generality, we may assume that b and c are the first occurrences of the corresponding integers in $\beta(P)$; then both b and c correspond to `U`-steps in P and are pushed onto Δ. Moreover, since $b' = b$ and b' follows c in $\beta(P)$, when c enters Δ, b is still in, and so c lies above b in Δ. Now observe that the element b' must correspond to either a `D`-step or a horizontal step labeled ℓ_2 of P. However, in both cases, when b' is inserted into $\beta(P)$, b has to be at the top of the stack, hence c should have been popped. This would imply that there are no more occurrences of c in $\beta(P)$ after b', which is not the case, since $c' = c$. $\square$

Remark 5.1. If we replace the stack Δ with a queue Ξ, then the same map gives a bijection with RGFs avoiding 12332. The proof is analogous to the previous one, and thus omitted.

Remark 5.2. If we restrict the previous bijections to Motzkin paths with no horizontal steps labeled ℓ_1, then we get bijections with RGFs that avoid 1212 (if we use a stack Δ) or 1221 (if we use a queue Ξ), provided that we remove the 1 at the beginning and decrease all the other elements by one. This follows again from the characterization of $\mathfrak{RGF}(12323)$ and $\mathfrak{RGF}(12332)$ given in [35] (and mentioned in the proof of Corollary 5.18). The corresponding continued fraction is then:

$$G(t) = \cfrac{1}{1 - t - \cfrac{t^2}{1 - 2t - \cfrac{t^2}{1 - 2t - \cfrac{t^2}{1 - 2t - \cdots}}}}$$

This gives an alternative proof of the fact that RGFs avoiding either 1221 or 1212 are enumerated by the Catalan numbers, whose generating function is known to be given by the above continued fraction.

Remark 5.3. As a consequence of the bijections in Theorem 5.19 and Remark 5.1, the statistic "sum of the numbers of `U` and ℓ_0 steps" in $\mathcal{M}_n^{lab}$ is equidistributed with the statistic "(value of the) maximum minus one" both in $\mathfrak{RGF}_{n+1}(12332)$ and in $\mathfrak{RGF}_{n+1}(12323)$. The same holds for the statistics "number of labels ℓ_0" and "number of singletons $\neq \{1\}$", as well as for the statistics "number of labels ℓ_1" and "number of occurrences of 1 minus one". Some computations seem to suggest that the distribution of the maximum is the same for several other patterns of the same Wilf-class, namely 12123, 12132, 12213, 12231, 12312, 12321, 12331, so we suspect that the same approach should lead to straightforward bijections, by suitably modifying the interpretation of the steps. For example, define r_i to be a *repeated* ltr-maximum of a RGF $r_1 r_2 \cdots r_n$ if $r_i = \max\{r_1, \dots, r_{i-1}\}$. Then steps having label ℓ_1 seem to have the same distribution as the repeated ltr-maxima in $\mathfrak{RGF}(12321)$ and $\mathfrak{RGF}(12312)$, so in order to define a bijection with $\mathcal{M}^{lab}$ it could be enough to find the "correct" interpretations for steps having labels `D` and ℓ_2.

Open Problem 5.3. Find suitable interpretations of the steps of labeled Motzkin paths in $\mathcal{M}^{lab}$ to obtain bijections with the remaining sets of pattern-avoiding RGFs in the same Wilf-equivalence class.

5.3.3 Patterns 12321 and 12312

In this section we show a connection between RGFs avoiding the patterns 12321 and 12312 and permutations avoiding the patterns 321 and 312, respectively. We initially provide a bijection between $\mathfrak{RGF}(12321)$ and $\mathfrak{S}(321)$ by showing that these two sets share the same combinatorial structure: elements of both sets can be written as shuffle of two weakly increasing sequences, one of those being the sequence of ltr-maxima. An analogous property links $\mathfrak{RGF}(12231)$ and $\mathfrak{S}(231)$.

Let $R = r_1 \cdots r_n$ be a RGF. Recall from Remark 5.3 that r_i is said to be a repeated ltr-maximum when $r_i = \max\{r_1, \dots, r_{i-1}\}$. Let $\mathfrak{RGF}^{n.r.}$ be the set of

RGFs that have no repeated ltr-maxima. Define $\mathfrak{RGF}_n^{n.r.}$ and $\mathfrak{RGF}^{n.r.}(Q)$, for a pattern Q, as usual. Given $R = r_1 \cdots r_n \in \mathfrak{RGF}^{n.r.}$, define $\tilde{R}$ as the subsequence obtained by deleting the ltr-maxima of R. It is easy to realize that $\tilde{R}$ is not necessarily a RGF. For example, if $R = 121311245246$, then $\tilde{R} = 111224$.

Lemma 5.20. Let $R \in \mathfrak{RGF}^{n.r.}$. Then R avoids 12321 if and only $\tilde{R}$ is weakly increasing.

Proof. Suppose that $\tilde{R} = \cdots ba \cdots$, with $b > a$. Note that b is not a repeated ltr-maximum of R, so there has to be an element c in R such that $c > b$ and c comes before b. Then R contains an occurrence cba of 321 and therefore it also contains 12321, by Lemma 5.12.

Conversely, if R contains an occurrence $abcb'a'$ of 12321, then b' precedes a' in $\tilde{R}$ and $b' > a'$, so $\tilde{R}$ is not weakly increasing. $\square$

We now wish to describe the anticipated bijection between $\mathfrak{RGF}_n(12321)$ and $\mathfrak{S}_n(321)$. Let $R = r_1 \cdots r_n \in \mathfrak{RGF}^{n.r.}(12321)$ and let $\tilde{R} = r_{i_1} \cdots r_{i_k}$, with $k \geq 0$. Construct a permutation $\pi(R)$ of length n by keeping the same positions for the ltr-maxima and mapping $\tilde{R}$ to a strictly increasing sequence $S = s_1 \cdots s_k$ as follows:

- $s_1 = r_{i_1}$;
- $s_j = s_{j-1} + (r_{i_j} - r_{i_{j-1}}) + 1$, for $j \geq 2$.

Finally, insert the remaining elements in increasing order in order to get a permutation that avoids 321: elements inserted at this point will be the ltr-maxima of the resulting permutation $\pi(R)$. For instance, let $R = 121314234$. Then the string obtained by removing the ltr-maxima from R is $\tilde{R} = 11234$. We thus get the increasing sequence $S = 12468$ and finally $\pi(R) = \mathbf{351729}468$ (where ltr-maxima of $\pi(R)$ are bolded). Observe that the number of ltr-maxima of $\pi(R)$ is equal to the number of ltr-maxima of the starting RGF R. Moreover, it is easy to realize (by construction) that $\pi(R)$ is a 321-avoiding permutation. The map defined this way is also injective. Indeed positions and values of the ltr-maxima uniquely determine a 321-avoiding permutations, thus the strictly increasing sequence S is enough to identify one such permutation. The construction proposed can be inverted in a similar fashion. It follows that the map $R \mapsto \pi(R)$ is a size-preserving bijection between $\mathfrak{RGF}^{n.r.}(12321)$ and $\mathfrak{S}(321)$. The next corollary is an immediate consequence of what discussed so far.

Proposition 5.21. The number of RGFs in $\mathfrak{RGF}_n^{n.r.}(12321)$ is $\mathfrak{c}_n$. Moreover, the number of RGFs in $\mathfrak{RGF}_n^{n.r.}(12321)$ having maximum k is given by $\mathfrak{n}_{n,k}$.

Next, in order to enumerate $\mathfrak{RGF}(12321)$, it is sufficient to show that any RGF avoiding 12321 can be uniquely obtained by inserting some repeated ltr-maxima in a sequence in $\mathfrak{RGF}^{n.r.}(12321)$.

Theorem 5.22. Let R be a RGF and let $\alpha(R)$ be the sequence obtained from R by removing all the repeated ltr-maxima. Then $\alpha(R)$ is a RGF. Moreover, R avoids 12321 if and only $\alpha(R)$ avoids 12321.

Proof. It is easy to check that $\alpha(R)$ is still a RGF and clearly $\alpha(R)$ avoids 12321 if R does. On the other hand, suppose that R contains an occurrence $abcb'a'$ of 12321. Note that b' and a' are not repeated ltr-maxima, so they are elements of $\alpha(R)$ and they follow c in R. Let c' be the first occurrence of the integer c in R. Then $c' \in \alpha(R)$ and c' precedes b' in $\alpha(R)$, so $\alpha(R)$ contains an occurrence $c'b'a'$ of 321, which is equivalent to containing 12321. $\square$

Corollary 5.23. For each $n \geq 1$, we have:

$$|\mathfrak{RGF}_{n+1}(12321)| = \sum_{k=0}^{n} \binom{n}{k} \mathfrak{c}_k.$$

Moreover, there are $\sum_{j=k}^{n} \binom{n}{j} \mathfrak{n}_{j,k}$ RGFs in $\mathfrak{RGF}_{n+1}(12321)$ with maximum k.

Proof. This is a direct consequence of the results proved in this section, together with the fact that the first element of a RGF cannot be a repeated ltr-maximum. $\square$

Remark 5.4. The same approach can be used to find a bijection between $\mathfrak{RGF}^{n.r.}(12312)$ and $\mathfrak{S}(312)$. In fact, 312-avoiding permutations are also uniquely determined by the positions and values of their ltr-maxima, and a completely analogous argument can be applied. As a consequence, we also have:

$$|\mathfrak{RGF}_{n+1}(12312)| = \sum_{k=0}^{n} \binom{n}{k} \mathfrak{c}_k.$$

5.3.4 A bijection between $\mathfrak{RGF}(12321)$ and $\mathfrak{RGF}(12231)$

In Section 5.2 we showed a bijection between Sort(132) and $\mathfrak{RGF}(12231)$. A direct combinatorial enumeration of $\mathfrak{RGF}(12231)$ could be obtained by using the labeled Motzkin path approach described in Section 5.3.2, but so far the pattern 12231 proved to be rather more complicated than some other patterns in the same equivalence class. The main goal of this section is to obtain an independent[5] proof of

[5]Without relying on the Wilf-equivalence showed in [35].

the enumeration of Sort(132) by means of a bijection δ between $\mathfrak{RGF}(12231)$ and $\mathfrak{RGF}(12321)$.

For the rest of this chapter, we say that the RGF R contains an occurrence of the pattern $\tilde{2}31$ if R contains three elements $r_{i_1}r_{i_2}r_{i_3}$ such that $r_{i_1}r_{i_2}r_{i_3} \simeq 231$ and r_{i_1} is not an ltr-maximum of R (equivalently, r_{i_1} is not the leftmost occurrence of the corresponding integer in R). Due to Lemma 5.12, we have $\mathfrak{RGF}(12231) = \mathfrak{RGF}(\tilde{2}31)$ and also $\mathfrak{RGF}(12321) = \mathfrak{RGF}(321)$. This allows us to focus on $\tilde{2}31$ and 321 instead of 12231 and 12321, respectively. More precisely, we shall define the promised map δ from $\mathfrak{RGF}(12231)$ to $\mathfrak{RGF}(12321)$ by repeatedly transforming the rightmost occurrence of 321 into an occurrence of $\tilde{2}31$, until the resulting RGF avoids 321. This is formalized in what follows.

Let $R = r_1 \cdots r_n$ be a RGF. Define $\mathrm{rm}(R, 321) = i_1 i_2 i_3$, where $r_{i_i}r_{i_2}r_{i_3}$ are the indices of the lexicographically rightmost occurrence of 321 in R. More extensively, this can be expressed by saying that for any other occurrence $r_{j_i}r_{j_2}r_{j_3}$ of 321 in R, it must be either $j_1 < i_1$, or $j_1 = i_1$ and $j_2 < i_2$, or $j_1 = i_1$, $j_2 = i_2$ and $j_3 < i_3$. If R avoids 321, we assume $\mathrm{rm}(321) = 000$ by convention. Similarly, denote by $\mathrm{lm}(R, \tilde{2}31) = i_1 i_2 i_3$ the indices of the lexicographically leftmost occurrence of $\tilde{2}31$ in R. If R avoids $\tilde{2}31$, we assume $\mathrm{lm}(R, \tilde{2}31) = (n+1)(n+1)(n+1)$.

Let us now define the map δ. Suppose that $R = r_1 \cdots r_n \in \mathfrak{RGF}(\tilde{2}31)$. Define recursively a RGF $\delta(R)$ as follows.

1. $R^{(0)} = R$.

2. For $t \geq 0$, if $R^{(t)}$ contains 321, then $R^{(t+1)}$ is obtained from $R^{(t)}$ by exchanging the elements r_{i_1} and r_{i_2}, where $i_1 i_2 i_3 = \mathrm{rm}(R^{(t)}, 321)$.

3. Finally, define $\delta(R) = R^{(k)}$, where k is the minimum index such that $R^{(k)}$ avoids 321.

Observe that $R^{(t)}$ is a RGF for each t and that $R^{(k)}$ avoids 321 by construction. Moreover, as a consequence of the next lemma, the integer k exists and thus the map δ is well defined.

Lemma 5.24. For every $t \geq 0$, we have $\mathrm{rm}(R^{(t+1)}, 321) <_\ell \mathrm{rm}(R^{(t)}, 321)$, where $<_\ell$ denotes the lexicographical order.

Proof. Let $R^{(t)} = r_1^{(t)} \cdots r_n^{(t)}$ and, similarly, $R^{(t+1)} = r_1^{(t+1)} \cdots r_n^{(t+1)}$. Moreover, let $\mathrm{rm}(R^{(t)}, 321) = i_1 i_2 i_3$ and $\mathrm{rm}(R^{(t+1)}, 321) = j_1 j_2 j_3$. Note that, as illustrated in Figure 5.5, our hypothesis imposes some constraints on the elements of $R^{(t)}$. More precisely, $r_j^{(t)} \leq r_{i_2}^{(t)}$, for each $j = i_1{+}1, \ldots, i_2{-}1$. Also, for each $j = i_2{+}1, \ldots, i_3{-}1$, either $r_j^{(t)} \leq r_{i_3}^{(t)}$ or $r_j^{(t)} \geq r_{i_1}^{(t)}$. Finally, $r_j^{(t)} \geq r_{i_2}^{(t)}$ for each $j > i_3$. We will repeatedly use these inequalities throughout this proof. Our goal is now to show

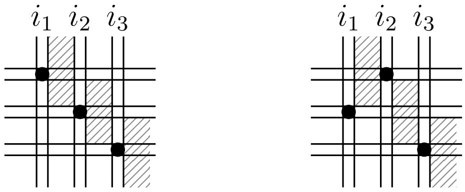

Figure 5.5: On the left, the rightmost occurrence of the pattern 321 in $R^{(t)}$, with indices $i_1i_2i_3$, represented as a Cayley-mesh pattern. On the right, the resulting (Cayley-mesh) pattern in $R^{(t+1)}$, obtained by exchanging the elements in positions i_1 and i_2.

that $j_1j_2j_3 <_\ell i_1i_2i_3$. Suppose, by contradiction, that $j_1j_2j_3 \geq_\ell i_1i_2i_3$. Consider the following case analysis.

- Suppose $j_1 > i_1$. If $j_1 < i_2$, then necessarily $r_{j_1}^{(t+1)} = r_{j_1}^{(t)} \leq r_{i_2}^{(t)}$, due to the above constraints. Hence we must have $j_2, j_3 \neq i_2$, since otherwise the indices j_1, j_2, j_3 would not correspond to an occurrence of 321 in $R^{(t+1)}$. This implies that $r_{j_1}^{(t+1)} r_{j_2}^{(t+1)} r_{j_3}^{(t+1)} = r_{j_1}^{(t)} r_{j_2}^{(t)} r_{j_3}^{(t)}$ is an occurrence of 321 in $R^{(t)}$ as well, with $j_1j_2j_3 >_\ell i_1i_2i_3$: this is a contradiction, since we are assuming that $\mathrm{rm}(R^{(t)}, 321) = i_1i_2i_3$. Next suppose that $j_1 = i_2$ (and so $j_2 > i_2$). Note that $r_{i_1}^{(t)} = r_{i_2}^{(t+1)} = r_{j_1}^{(t+1)}$, hence $r_{i_1}^{(t)} r_{j_2}^{(t)} r_{j_3}^{(t)}$ is an occurrence of 321 in $R^{(t)}$ with $i_1j_2j_3 >_\ell i_1i_2i_3$, which is impossible. Finally, suppose that $j_1 > i_2$. Then obviously $r_{j_1}^{(t)} r_{j_2}^{(t)} r_{j_3}^{(t)} = r_{j_1}^{(t+1)} r_{j_2}^{(t+1)} r_{j_3}^{(t+1)}$ is an occurrence of 321 in $R^{(t)}$, with $j_1j_2j_3 >_\ell i_1i_2i_3$, again a contradiction.

- Suppose instead that $j_1 = i_1$ and $j_2 > i_2$. Then $r_{i_1}^{(t+1)} = r_{i_2}^{(t)}$ and $j_2 > i_2$, so $r_{i_2}^{(t)} r_{j_2}^{(t)} r_{j_3}^{(t)}$ is an occurrence of 321 in $R^{(t)}$, with $i_2j_2j_3 >_\ell i_1i_2i_3$, which is impossible.

- Finally, the case $j_1 = i_1$ and $j_2 = i_2$ is clearly impossible, since we have $r_{i_1}^{(t+1)} = r_{i_2}^{(t)} < r_{i_1}^{(t)} = r_{i_2}^{(t+1)}$.

□

We wish to show that δ is a bijection by proving that the recursive construction proposed above can be reversed (in the expected way!). Indeed, in order to obtain $R^{(t)}$ from $R^{(t+1)}$, it is sufficient to transform the leftmost occurrence of $\tilde{2}31$ into an occurrence of 321 (see Figure 5.6). A formal proof is given in the next lemma.

Lemma 5.25. Let $t \geq 0$. Let $\mathrm{rm}(R^{(t)}, 321) = i_1i_2i_3$ and $\mathrm{lm}(R^{(t+1)}, \tilde{2}31) = j_1j_2j_3$. Then $i_1 = j_1$ and $i_2 = j_2$.

$$\begin{tikzcd}
 & \mathrm{rm}(321) \arrow[r, dashed] & \tilde{2}31 \arrow[dr] & \\
R^{(t)} \arrow[ur] & & & R^{(t+1)} \arrow[dl] \\
 & 321 \arrow[ul] & \mathrm{lm}(\tilde{2}31) \arrow[l, dashed] &
\end{tikzcd}$$

Figure 5.6: The diagram of Lemma 5.25.

Proof. We again refer to Figure 5.5 for an illustration of the constraints imposed on the elements of $R^{(t)}$ by the position of the rightmost occurrence of 321 inside $R^{(t)}$. We proceed by induction on t.

Suppose first that $t = 0$, that is, $R^{(0)} = r_1^{(0)} \cdots r_n^{(0)}$ avoids $\tilde{2}31$, but contains 321. Set $R^{(1)} = r_1^{(1)} \cdots r_n^{(1)}$, $\mathrm{rm}(R^{(0)}, 321) = i_1 i_2 i_3$ and $\mathrm{lm}(R^{(1)}, \tilde{2}31) = j_1 j_2 j_3$. Note that $r_{i_1}^{(1)} r_{i_2}^{(1)} r_{i_3}^{(1)}$ is an occurrence of $\tilde{2}31$ in $R^{(1)}$. Indeed, by Lemma 5.12, the first occurrence of the integer $r_{i_2}^{(0)}$ in $R^{(0)}$ precedes $r_{i_1}^{(0)}$, since $r_{i_1}^{(0)} > r_{i_2}^{(0)}$. Therefore $j_1 j_2 j_3 \leq_\ell i_1 i_2 i_3$. We have to show that $i_1 = j_1$ and $i_2 = j_2$. Suppose, to the contrary, that $j_1 < i_1$. If either $j_2 = i_1$ or $j_2 = i_2$, then $r_{j_1}^{(0)} r_{i_1}^{(0)} r_{j_3}^{(0)}$ would be an occurrence of $\tilde{2}31$ in $R^{(0)}$, which is impossible since $R^{(0)} \in \mathfrak{RGF}(\tilde{2}31)$. Thus we must have $j_2 \neq i_1$ and $j_2 \neq i_2$. In particular, since $j_2 \neq i_2$, we must have either $j_3 = i_1$ or $j_3 = i_2$, otherwise $r_{j_1}^{(0)} r_{j_2}^{(0)} r_{j_3}^{(0)} = r_{j_1}^{(1)} r_{j_2}^{(1)} r_{j_3}^{(1)}$ would be an occurrence of $\tilde{2}31$ in $R^{(0)}$ as well. However, if either $j_3 = i_1$ or $j_3 = i_2$, then $r_{j_1}^{(0)} r_{j_2}^{(0)} r_{i_2}^{(0)}$ would be an occurrence of $\tilde{2}31$ in $R^{(0)}$, which is again a contradiction. Therefore it has to be $i_1 = j_1$. Finally, the case $j_1 = i_1$ and $j_2 < i_2$ is forbidden, due to the restrictions depicted in Figure 5.5. Summing up, we must have $i_1 = j_1$ and $i_2 = j_2$, as desired.

Now suppose that $t \geq 1$. Let $R^{(t)} = r_1^{(t)} \cdots r_n^{(t)}$. For the rest of this proof, we fix the following notation:

- $\mathrm{rm}(R^{(t-1)}, 321) = t_1 t_2 t_3$;
- $\mathrm{lm}(R^{(t)}, \tilde{2}31) = s_1 s_2 s_3$;
- $\mathrm{rm}(R^{(t)}, 321) = i_1 i_2 i_3$;
- $\mathrm{lm}(R^{(t+1)}, \tilde{2}31) = j_1 j_2 j_3$.

By the inductive hypothesis we have $s_1 = t_1$ and $s_2 = t_2$. Moreover, Lemma 5.24 implies that $t_1 t_2 t_3 >_\ell i_1 i_2 i_3$, hence $t_1 t_2 \geq_\ell i_1 i_2$ and $s_1 s_2 \geq_\ell i_1 i_2$. Note that $r_{i_1}^{(t+1)} r_{i_2}^{(t+1)} r_{i_3}^{(t+1)}$ is an occurrence of $\tilde{2}31$ in $R^{(t+1)}$, so we must have $j_1 j_2 j_3 \leq_\ell$

$i_1i_2i_3$. Our goal is to show that $i_1 = j_1$ and $i_2 = j_2$. We shall proceed by contradiction, so we assume that $j_1 < i_1$ or $j_2 < i_2$. Our strategy consists in finding an occurrence of $\tilde{2}31$ in $R^{(t)}$ such that the indices of its first two elements strictly precede i_1i_2 (in the lexicographical order). Indeed, this would imply that $s_1s_2 <_\ell i_1i_2$, since $s_1s_2s_3 = \mathrm{lm}(R^{(t)}, \tilde{2}31)$, which is impossible since we know that $s_1s_2 \geq_\ell i_1i_2$.

Suppose first that $j_1 < i_1$. If $\{j_2, j_3\} \cap \{i_1, i_2\} = \emptyset$, then $r^{(t)}_{j_1}r^{(t)}_{j_2}r^{(t)}_{j_3}$ is the desired occurrence of $\tilde{2}31$ in $R^{(t)}$, since in this case j_1, j_2, j_3 are not involved in the transition from $R^{(t)}$ to $R^{(t+1)}$ and we are assuming that $j_1 < i_1$. Therefore at least one of j_2 and j_3 must coincide with either i_1 or i_2. We will now show that, in each case, we are able to find an occurrence of $\tilde{2}31$ in $R^{(t)}$ with the desired property.

- If $j_2 = i_1$, then $r^{(t+1)}_{j_2} = r^{(t+1)}_{i_1} < r^{(t)}_{i_1}$, hence $r^{(t)}_{j_1}r^{(t)}_{j_2}r^{(t)}_{j_3}$ is an occurrence of $\tilde{2}31$ in $R^{(t)}$, and $j_1j_2 <_\ell i_1i_2$.
- If $j_2 = i_2$, then $r^{(t)}_{j_1}r^{(t)}_{i_1}r^{(t)}_{j_3}$ is an occurrence $\tilde{2}31$ in $R^{(t)}$, and $j_1i_1 <_\ell i_1i_2$.
- If $j_3 = i_1$, then $r^{(t)}_{j_1}r^{(t)}_{j_2}r^{(t)}_{i_2}$ is an occurrence of $\tilde{2}31$ in $R^{(t)}$, and $j_1j_2 <_\ell i_1i_2$.
- If $j_3 = i_2$, then $r^{(t)}_{i_2} < r^{(t)}_{i_1} = r^{(t+1)}_{i_2}$, hence $r^{(t)}_{j_1}r^{(t)}_{j_2}r^{(t)}_{i_2}$ is an occurrence of $\tilde{2}31$ in $R^{(t)}$, and $j_1j_2 < i_1i_2$.

The above case-by-case analysis shows that $i_1 = j_1$. Moreover, we cannot have $j_2 < i_2$; this is again due to the restrictions illustrated in Figure 5.5. □

Theorem 5.26. The map δ is a size-preserving bijection between $\mathfrak{RGF}(12321)$ and $\mathfrak{RGF}(12231)$. Moreover, δ preserves the maximum value of a RGF.

As a consequence of Theorem 5.26 and Corollary 5.23, the distribution of the maximum letter in RGFs over $\mathfrak{RGF}_n(12231)$ is given by $\sum_{i=k}^{n} \binom{n}{i}\mathfrak{n}_{i,k}$. This provides a combinatorial (even if not direct) proof of the formula stated in Open Problem 5.2 for the distribution of ltr-minima of Sort(132).

Chapter 6

The (σ, τ)-machine

In this chapter we consider Σ-machines where $\Sigma = \{\sigma, \tau\}$ is a pair of patterns of length three. For specific pairs of patterns, the set of (σ, τ)-sortable permutations is enumerated by either the (binomial transform of) Catalan numbers or the Schröder numbers. We also determine an infinite family of pairs of patterns, namely the pairs $(\sigma, \widehat{\sigma})$, whose enumeration is given by the Catalan numbers. Some of the mentioned cases were discussed in [8]. Here we will sometimes follow an alternative approach. We also provide enumerative results for other pairs of patterns, which are currently unpublished.

6.1 Two decomposition lemmas

The avoidance of specific pairs of patterns deeply influences the geometric structure of the output of a (σ, τ)-stack. The following two decomposition lemmas are particularly useful.

Lemma 6.1. Consider the $(312, \sigma)$-machine, where $\sigma = \sigma_1 \cdots \sigma_{k-1}\sigma_k \in \mathfrak{S}_k$, with $\sigma_{k-1} < \sigma_k$ and $k \geq 3$. Given a permutation π of length n, let $\pi = M_1 B_1 \cdots M_t B_t$ be its ltr-max decomposition. Then:

1. Every time an ltr-maximum M_i is pushed into the $(312, \sigma)$-stack, the $(312, \sigma)$-stack contains the elements $M_i, M_{i-1}, \dots, M_2, M_1$, reading from top to bottom. Moreover, we have:

$$\mathcal{S}^{312,\sigma}(\pi) = \tilde{B}_1 \cdots \tilde{B}_t M_t \cdots M_1,$$

   where $\tilde{B}_i$ is a suitable rearrangement of B_i.

2. If π is $(312, \sigma)$-sortable, then $M_j = n - t + j$, for each $j = 1, \dots, t$. Therefore $\{M_1, \dots, M_t\} = \{n - t + 1, \dots, n\}$.

Proof. 1. Let us consider the action of the $(312, \sigma)$-stack on input π. Note that, since $k \geq 3$, the element M_1 remains at the bottom of the $(312, \sigma)$-stack until the end of the sorting procedure. Now, for each $x \in B_1$, the elements $M_2 x M_1$ form an occurrence of 312. Therefore the block B_1 is extracted before M_2 enters the $(312, \sigma)$-stack. As soon as M_2 enters, the stack contains $M_2 M_1$, reading from top to bottom. Since $M_2 > M_1$, but $\sigma_{k-1} < \sigma_k$ by hypothesis, M_2 cannot play the role of either σ_{k-1} in an occurrence of σ or 1 in an occurrence of 312. Thus M_2 remains at the bottom of the stack (above M_1) until the end of the sorting procedure. The thesis follows by iterating the same argument on each block B_i, for $i \geq 2$.

2. Suppose, for a contradiction, that there is an element $j \in \{n-t+1, \dots, n\}$ which is not an ltr-maximum. Note that $j \neq \pi_1 = M_1$ and $j \neq n = M_t$. Then, by what proved above, $\mathcal{S}^{312,\sigma}(\pi)$ contains an occurrence jnM_1 of 231, which contradicts the hypothesis that π is $(312, \sigma)$-sortable.

□

Lemma 6.2. Consider the $(132, \sigma)$-machine, where $\sigma = \sigma_1 \cdots \sigma_{k-1}\sigma_k \in \mathfrak{S}_k$, with $\sigma_{k-1} > \sigma_k$ and $k \geq 3$. Given a permutation π of length n, let $\pi = m_1 B_1 \cdots m_t B_t$ be its ltr-min decomposition. Then:

1. Every time an ltr-minimum m_i is pushed into the $(132, \sigma)$-stack, the $(132, \sigma)$-stack contains the elements $m_i, m_{i-1}, \dots, m_2, m_1$, reading from top to bottom. Moreover, we have:

$$\mathcal{S}^{132,\sigma}(\pi) = \tilde{B}_1 \cdots \tilde{B}_t m_t \cdots m_1,$$

where $\tilde{B}_i$ is a suitable rearrangement of B_i.

2. If π is $(132, \sigma)$-sortable, then $\tilde{B}_i$ is increasing for each i. Moreover, for each $i \leq t-1$, we have $B_i > B_{i+1}$ (i.e. $x > y$ for each $x \in B_i, y \in B_{i+1}$).

Proof. 1. The proof is analogous to the first part of Lemma 6.1.

2. Suppose that π is $(132, \sigma)$-sortable. Suppose, for a contradiction, that $\tilde{B}_i$ is not increasing, for some i, and let xy be a descent in B_i. Then, by what proved above, $\mathcal{S}^{132,\sigma}(\pi)$ contains an occurrence xym_t of 231, which is impossible. Finally, suppose that $x > y$, for $x \in B_i$ and $y \in B_{i+1}$. Then again xym_t is an occurrence of 231 in $\mathcal{S}^{132,\sigma}(\pi)$, a contradiction.

□

Some enumerative data for (σ, τ)-sortable permutations are reported in Table 6.1.

(σ,τ)	**Sequence** $\{f_n^{\sigma,\tau}\}_n$	**Reference**	**OEIS**
123,132 123,213 132,312	1, 2, 5, 14, 42, 132, 429, 1430	Section 6.4	A000108
231,321	1, 2, 5, 14, 42, 132, 429, 1430	Remark 8.2	A000108
123,231	1, 2, 6, 21, 79, 310, 1252, 5168		
123,312	1, 2, 5, 15, 51, 188, 731, 2950	Section 6.5	A007317
123,321	1, 2, 4, 7, 14, 28, 56, 112	Section 6.3	
132,213	1, 2, 5, 16, 61, 261, 1206, 5882		
132,231	1, 2, 6, 22, 90, 394, 1806, 8558	Section 6.2	A006318
132,321	1, 2, 4, 10, 26, 72, 206, 606	Section 6.3	A102407
213,231	1, 2, 6, 23, 101, 483, 2450, 12978		
213,312	1, 2, 5, 16, 61, 261, 1206, 5882		
213,321	1, 2, 4, 12, 46, 200, 941, 4677		
231,312	1, 2, 6, 23, 101, 484, 2471, 13254		
312,321	1, 2, 4, 10, 28, 85, 274, 925		

Table 6.1: Enumerative results for (σ,τ)-sortable permutations, with σ and τ patterns of length three and starting from (σ,τ) sortable permutations of length one.

6.2 The $(132, 231)$-machine

In this section we analyze the $(132, 231)$-machine. We use Lemma 6.2 to provide a geometric description of the $(132, 231)$-sortable permutations, which we exploit to show that $\mathrm{Sort}(132, 231) = \mathfrak{S}(1324, 2314)$. As it is well known, the set $\mathfrak{S}(1324, 2314)$ is enumerated by the large Schröder numbers.

Lemma 6.3. Let $\pi = m_1 B_1 \cdots m_t B_t$ be the ltr-min decomposition of a permutation π. Write $\mathcal{S}^{132,231}(\pi) = \tilde{B}_1 \cdots \tilde{B}_t m_t \cdots m_1$ as in Lemma 6.2. Then $\tilde{B}_i = \mathcal{S}^{12}(B_i)$, for each i.

Proof. Let $i \geq 1$. Consider the instant when m_i enters the $(132, 231)$-stack, that is as soon as B_i is the next block of π to be processed. Here, by Lemma 6.2, the stack contains $m_i, m_{i-1}, \dots, m_1$, reading from top to bottom. We shall prove that the behavior of the $(132, 231)$-stack on B_i is equivalent to the behavior of an empty 12-stack on input B_i. In other words, the $(132, 231)$-stack performs a pop operation if and only if a 12-stack that ignores the current content of the $(132, 231)$-stack does the same. If either the next element of the input is m_{i+1} or the input is empty, then both the $(132, 231)$-stack and the 12-stack perform a pop operation. Otherwise, suppose that the next element of the input is y, for some y in B_i. Suppose that the $(132, 231)$-stack pops the element $x \in B_i$. Thus the $(132, 231)$-stack contains two elements z, w, with z above w, such that yzw is an occurrence of either 132 or 231. Note that, since $z > w$, z is not an ltr-minimum. Therefore yz is an occurrence of 12 and the 12-stack performs a pop operation too, as desired. Conversely, suppose that the 12-stack pops the element x. Thus the 12-stack contains an element z such that $z > y$. Therefore yzm_i is an occurrence of 231 and the $(132, 231)$-stack performs a pop operation as well, as desired. □

Corollary 6.4. Let $\pi = m_1 B_1 \cdots m_t B_t$ be the ltr-min decomposition of a permutation π. The following are equivalent:

1. B_i avoids 213 and $B_i > B_{i+1}$, for each i.

2. π is $(132, 231)$-sortable.

3. $\pi \in \mathfrak{S}(1324, 2314)$.

In particular, we have $\mathrm{Sort}(132, 231) = \mathfrak{S}(1324, 2314)$.

Proof. By Lemma 6.2 and Lemma 6.3, we have:

$$\mathcal{S}^{132,231}(\pi) = \mathcal{S}^{12}(B_1) \cdots \mathcal{S}^{12}(B_t) m_t \cdots m_1.$$

We will use this decomposition for the rest of the proof.

- [$1 \Rightarrow 2$] Suppose, for a contradiction, that $\mathcal{S}^{132,231}(\pi)$ contains an occurrence bca of 231. Note that, since $c > a$, while $m_t < \cdots < m_1$, c is not an ltr-minimum of π (and thus neither b is). Moreover, since we assumed that $B_i > B_{i+1}$ for each i, b and c must belong to the same block B_j. Therefore $\mathcal{S}^{12}(B_j)$ is not decreasing and, by Theorem 3.1, B_j contains 213, which contradicts the hypothesis.

- [$2 \Rightarrow 3$] Suppose, for a contradiction, that π contains 1324 or 2314. Initially, suppose that π contains an occurrence $acbd$ of 1324. Note that b, c and d are not ltr-minima of π. Let $b \in B_j$ and $c \in B_k$, for some $j \leq k$. If $j = k$, then B_j contains an occurrence bac of 213. Therefore $\mathcal{S}^{12}(B_j)$ contains an occurrence of 231 due to Theorem 3.1, which contradicts the hypothesis. Otherwise, if $j < k$, then bcm_k is an occurrence of 231 in $\mathcal{S}^{132,231}(\pi)$, which is again impossible. The pattern 2314 can be addressed similarly, and it is left to the reader.

- [$3 \Rightarrow 1$] Again we argue by contradiction. If B_i contains an occurrence bac of 213, then π contains an occurrence $m_i bac$ of 1324, which is impossible. Otherwise, if π contains two elements $x \in B_j$, $y \in B_k$, with $x < y$ and $j < k$, then $m_j x m_k y$ is an occurrence of 2314, which is impossible too.

$\square$

6.3 The $(123, 321)$- and $(132, 321)$-machines

The pairs of patterns $(123, 321)$ and $(132, 321)$ can be studied with similar tools. We show that, in both cases, sortable permutations avoid 123. Therefore the restriction involving the pattern 321 is never triggered when processing $(123, 321)$- and $(132, 321)$-sortable permutations. Equivalently, the $(123, 321)$-stack acts as a 123-stack when the input is a $(123, 321)$-sortable permutation and thus we can use the results of Chapter 4 to describe $\mathrm{Sort}(123, 321)$. In a similar fashion, the $(132, 321)$-stack acts as a 132-stack when the input is a $(123, 321)$-sortable permutations, which allows us to use what we proved in Chapter 4.

Lemma 6.5. Let $\sigma \in \{123, 132\}$. If π is $(\sigma, 321)$-sortable, then π avoids 123.

Proof. We prove the contrapositive statement by showing that if π contains 123, then $\mathcal{S}^{\sigma,321}(\pi)$ contains 231. Let abc be the leftmost[1] occurrence of 123 in π. Let us consider the instant when c is pushed into the $(\sigma, 321)$-stack. If a is still in the $(\sigma, 321)$-stack when c enters, then b is not inside the $(\sigma, 321)$-stack, since

[1]with respect to the lexicographical order of the indices.

otherwise cba would be an occurrence of 321, which is forbidden. Thus $\mathcal{S}^{\sigma,321}(\pi)$ contains $bca \simeq 231$ and we are done. Therefore we can assume that a is extracted before c enters the $(\sigma, 321)$-stack. Let us consider the instant when a is extracted from the $(\sigma, 321)$-stack. Let z be the next element of the input. Then there are two elements x, y into the stack, with x below y, such that zyx is an occurrence of either σ or 321. Note that necessarily x comes before a in π (it could be $y = a$). If $zyx \simeq 321$, then xyz is an occurrence of 123 which precedes abc, which is impossible due to our choice of abc. Therefore we have $zyx \simeq \sigma$. Next we show that $y = a$. Indeed, suppose, for a contradiction, that $y \neq a$. Note that both x and y are greater than a, since otherwise we would have an occurrence xac or yac of 123 which precedes abc, a contradiction. But then the $(\sigma, 321)$-stack would contain an occurrence xya of σ, which is impossible. We can now assume $y = a$. Note that $x > a$ due to our choice of abc. Thus $\sigma \neq 132$, since zax is an occurrence of σ and $x > a$. This completes the proof for the case $\sigma = 132$. Let us now assume that $\sigma = 123$ for the remaining part of the proof. Now, $zax \simeq 132$, thus $c \neq z$. If z is still in the $(123, 321)$-stack when c enters, then $\mathcal{S}^{123,321}(\pi)$ contains an occurrence acz of 231, as desired. Otherwise, consider the instant when z is extracted from the $(123, 321)$-stack, with c still in the input. Let w be the next element of the input. Then there are two elements u, v into the stack, with u below v, such that wvu is an occurrence of either 123 or 321. Observe that, since z is the next element of the input when a is extracted, the elements u and v precede a in π (otherwise they would have been extracted from the $(123, 321)$-stack before a). Therefore it cannot be $wvu \simeq 321$, due to our choice of abc as leftmost occurrence of 123 in π. Finally, if $wvu \simeq 123$, then we can repeat the same argument on the triple wvu, in place of zyx. Since wvu is strictly to the left of zyx, iterating this process will sooner or later result in either a contradiction or in finding an occurrence of 231 in the output of the $(123, 321)$-stack, as desired. $\square$

Corollary 6.6. We have:

$$\mathrm{Sort}(123, 321) = \mathrm{Sort}(123) \cap \mathfrak{S}(123)$$

and

$$\mathrm{Sort}(132, 321) = \mathrm{Sort}(132) \cap \mathfrak{S}(123).$$

Proof. Let $\sigma \in \{123, 132\}$. Due to Lemma 6.5, any $(\sigma, 321)$-sortable permutation avoids 123. Therefore the behavior of a $(\sigma, 321)$-stack on a $(\sigma, 321)$-sortable permutation is equivalent to the behavior of a σ-stack, since the restriction involving the pattern 321 is never triggered. $\square$

Let us now focus on the pair $(123, 321)$. Since $\mathrm{Sort}(123, 321) = \mathrm{Sort}(123) \cap \mathfrak{S}(123)$, we will describe this set by exploiting the characterization of $\mathrm{Sort}(123)$ provided in Chapter 4.

Theorem 6.7. Let $f_n = f_n^{123,321}$ be the number of $(123, 321)$-sortable permutations of length n. Then:

$$\begin{cases} f_1 = 1; \\ f_2 = 2; \\ f_3 = 4; \\ f_n = 7 \cdot 2^{n-4}, \ n \geq 4. \end{cases}$$

Proof. Each permutation of length one and two is $(123, 321)$-sortable, while 132 and 123 are the only permutations of length three that are not $(123, 321)$-sortable. Let $n \geq 4$ and let $\pi \in \mathrm{Sort}(123, 321)$. Due to Corollary 6.6, we have $\mathrm{Sort}(123, 321) = \mathrm{Sort}(123) \cap \mathfrak{S}(123)$. Therefore π can be uniquely constructed according to the procedure described in Section4.2, which we recall below, as long as occurrences of 123 are not created:

1. Choose $\alpha = \alpha_1 \alpha_2 \cdots \alpha_k \in \mathfrak{S}_k(213)$, with $\alpha_1 = k \geq 2$;

2. add h new maxima, $k + 1, \dots k + h$, one at a time, using the bijection φ of Theorem 4.7;

3. add $n - k - h$ consecutive ascents at the beginning, by inflating the first element of the permutation, according to Corollary 4.3.

Observe that it must be $k \geq n - 1$. Otherwise, if $k < n - 1$, then it would be $k + h \geq 2$, which would necessarily result in an occurrence of 123 in π. Thus we have either $k = n$ or $k = n - 1$. We distinguish the following cases.

- If $k = n$, then $\pi = \alpha$ is a permutation of $\mathfrak{S}_n(213, 123)$ with $\pi_1 = n$. Notice that π is uniquely obtained by adding the initial maximum to a permutation in $\mathfrak{S}_{n-1}(213, 123)$. In other words, this operation yields a bijection between the set of $(123, 321)$-sortable permutations of length n that start with their maximum value and $\mathfrak{S}_{n-1}(213, 123)$. It is well known that $|\mathfrak{S}_{n-1}(213, 123)| = 2^{n-2}$.

- Suppose that $k = n - 1$ and $h = 1$, that is π is obtained by adding a new maximum, immediately after $n - 2$, to some $\alpha \in \mathfrak{S}_{n-1}(213, 123)$, with $\alpha_1 = n - 1$. Then it must be necessarily $\alpha_2 = n - 2$, otherwise $\alpha_2 n - 2n$ would be an occurrence of 123 in π. Therefore $\alpha \in \mathfrak{S}_{n-1}(213, 123)$, with $\alpha_1 = n - 1$ and $\alpha_2 = n - 2$. Similarly to the previous item, removing the first two elements of α yields a bijection with $\mathfrak{S}_{n-3}(213, 123)$, and $|\mathfrak{S}_{n-3}(213, 123)| = 2^{n-4}$.

- Finally, suppose that $k = n-1$ and $h = 0$, that is π is obtained from some $\alpha \in \mathfrak{S}_{n-1}(213, 123)$, with $\alpha_1 = n-1$, by inflating the first element of α by one. Since this operation cannot create an occurrence of 123, each $\alpha \in \mathfrak{S}_{n-1}(213, 123)$ such that $\alpha_1 = n-1$ is a suitable choice, and $|\mathfrak{S}_{n-1}(213, 123)| = 2^{n-3}$.

At the end we obtain:

$$f_n = 2^{n-2} + 2^{n-4} + 2^{n-3} = 7 \cdot 2^{n-4},$$

as desired. □

Next we consider the set Sort(132, 321). We shall define a bijection between (132, 321)-sortable permutations of length n and Dyck paths of semilength n that avoid the (consecutive) pattern `DUDU`. More precisely, we will describe a bijection between two generating trees for these families. First we refine the geometric description of Sort(132, 321) by polishing the characterization of Sort(132) provided in Chapter 5. Recall from Section 5.1 that $\mu = (132, \{(0, 2), (2, 0), (2, 1)\})$ is the mesh pattern depicted in Figure 1.1.

Theorem 6.8. We have:

$$\mathrm{Sort}(132, 321) = \mathfrak{S}(\mu, 123).$$

Proof. It follows from Theorem 5.3 and Corollary 6.6, since $123 \leq 2314$. □

The grid decomposition of a permutation π was introduced in Section 5.2. Let $\pi = m_1 B_1 m_2 B_2 \dots m_t B_t$ be the ltr-min decomposition of π. Recall that:

- for $i \geq 1$, the $j-th$ *vertical strip* of π is B_j;
- for $i \geq 1$, the $i-th$ *horizontal strip* of π is $H_i = \{x \in \pi : m_{i-1} < x < m_i\}$, where $m_0 = +\infty$.
- for any pair of indices i, j, the *cell* of indices i, j of π is $C_{i,j} = H_i \bigcap B_j$ (note that $C_{i,j}$ is empty for each $i > j$).

Proposition 6.9. Let π be a (132, 321)-sortable permutation. Then each cell $C_{i,j}$ contains at most one element.

Proof. Suppose, for a contradiction, that the cell $C_{i,j}$ contains at least two elements and let $C_{i,j} = xy \cdots$. If $x < y$, then $m_i xy$ is an occurrence of 123, which contradicts Theorem 6.8. Otherwise, if $x > y$, then due to Lemma 5.5 there has to be some element z between x and y in π such that $z < m_i$. Let m_j be the ltr-minimum of the horizontal strip that contains z, for some $j \geq i$. Notice that $m_j \neq z$, since x and y are in the same cell and z is placed between x and y in π. Then $m_j zy$ is an occurrence of 123 in π, which is again impossible. □

Proposition 6.10. Let π be a $(132, 321)$-sortable permutation. Suppose that the cell $C_{i,j}$ is not empty. Then the cell $C_{u,v}$ is empty for each pair of indices u, v such that $i < u$ and $j \leq v$.

Proof. Suppose that $x \in C_{i,j}$ and $y \in C_{u,v}$, with $i < u$ and $j \leq v$. Then $m_j xy$ is an occurrence of 123 in π, which is impossible due to Theorem 6.8. □

Proposition 6.11. Let π be a $(132, 321)$-sortable permutation. Then $B_i > B_{i+1}$ for each pair of consecutive vertical strips B_i, B_{i+1}.

Proof. It follows from Lemma 5.1 and Theorem 6.8. □

Now, any prefix of a $(132, 321)$-sortable permutation is $(132, 321)$-sortable by Lemma 2.3. Therefore every permutation of $\mathrm{Sort}_n(132, 321)$ can be uniquely constructed by inserting a new rightmost integer $a \in \{1, \dots, n\}$ in a permutation $\pi \in \mathrm{Sort}_{n-1}(132, 321)$, and suitably rescaling the other elements[2]. Every permutation obtained this way from π is said to be a *child* of π. Now, due to Propositions 6.9 and 6.11, there is at most one way to insert a new rightmost element in each cell of the last vertical strip of a $(132, 321)$-sortable permutation. Indeed each cell can contain at most one element due to Proposition 6.9 and the value of this element is completely determined due to Proposition 6.11: it has to be less than all the other elements in the same horizontal strip. Finally, some choices will be forbidden due to Theorem 6.8. Given a $(132, 321)$-sortable permutation π, where $\pi = m_1 B_1 \cdots m_t B_t$ is the usual ltr-min decomposition, a cell $C_{i,t}$ is said to be *active* if inserting a new rightmost element in the cell $C_{i,t}$ yields a $(132, 321)$-sortable permutation. Next we characterize precisely which cells are active. First we introduce a new parameter.

Let $\pi = m_1 B_1 \cdots m_t B_t$. Let y be the rightmost element of π which is not an ltr-minimum and suppose that $y \in C_{i,j}$, for some i, j. Then the *depth* of π is $\mathrm{dep}(\pi) = t - i + 1$. If $\pi = m_1 \cdots m_t$ is the decreasing permutation, we assume $\mathrm{dep}(\pi) = t$.

Theorem 6.12. Let $\pi = m_1 B_1 \cdots m_t B_t$ a $(132, 321)$-sortable permutation and let $d = \mathrm{dep}(\pi)$.

1. If B_t is not empty, then the cell $C_{i,t}$ is active if and only if $i < d$. In this case π has d children.

2. If B_t is empty, then the cell $C_{i,t}$ is active if and only if $i \leq d$. In this case π has $d + 1$ children.

[2]i.e. adding 1 to each integer b such that $b \geq a$.

Proof. If π is the decreasing permutation, then $\text{dep}(\pi) = t$ by definition. Moreover B_t is empty and the cell $C_{i,t}$ is active for each $i = 1, \dots, d-1$ due to Theorem 6.8. Since inserting a new rightmost minimum is always allowed, in this case π has d children.

Otherwise, let y be the rightmost element of π which is not an ltr-minimum and suppose that $y \in C_{u,v}$, for some u, v, with $d = t - u + 1$. Note that every cell $C_{i,t}$, with $i > d$, is not active. Indeed inserting a new rightmost element $a \in C_{i,t}$, with $i > d$, creates an occurrence $m_v y a$ of 123, and this is forbidden due to Theorem 6.8. On the other hand, we shall prove that each cell $C_{i,t}$, with $i < d$, is active. Due to the same Theorem 6.8, it is sufficient to show that inserting a new rightmost element a in the cell $C_{i,t}$, with $i < d$, cannot create an occurrence of either 123 or μ. Suppose that $\pi_{j_1}\pi_{j_2}a \simeq 123$, for some indices $j_1 < j_2$. Note that π_{j_2} is not an ltr-minimum of π, therefore it precedes y in π. Moreover we have $\pi_{j_2} < y$, since $y > a > \pi_{j_2}$. Thus $\pi_{j_1}\pi_{j_2}y$ is an occurrence of 123 in π, which contradicts the fact that π is $(132, 321)$-sortable. The pattern μ can be treated analogously, so we leave the details to the reader.

Finally, we wish to prove that the cell $C_{t,d}$ is active if and only if B_t is empty. If B_t is not empty, then $C_{t,d}$ is not active due to Proposition 6.10. If instead B_t is empty, then inserting a in $C_{t,h}$ cannot create an occurrence of 123. This can be proved by repeating the same argument used above. Instead, suppose that a creates an occurrence $\pi_{j_1}\pi_{j_2}a$ of 132. Then, since B_t is empty, $\pi_{j_1}\pi_{j_2}m_t a$ is an occurrence of 2143, and thus $\pi_{j_1}\pi_{j_2}a$ is not an occurrence of μ, as desired. $\square$

Theorem 6.12 allows us to define a generating tree for $\text{Sort}(132, 321)$. The node corresponding to the $(132, 321)$-sortable permutation π, with $\pi = m_1B_1 \cdots m_tB_t$, is equipped with two labels (d, b), where:

- $d = \text{dep}(\pi)$ is the depth of π.
- b is a boolean counter with value $b = 0$, if B_t is empty, and $b = 1$, otherwise.

According to Theorem 6.12, the following rule provides a generating tree for $\text{Sort}(132, 321)$:

$$\Omega_1 : \begin{cases} (1,0) \\ (d,0) \longrightarrow (d+1,0)(1,1)(2,1)\cdots(d-1,1)(d,1) \\ (d,1) \longrightarrow (d+1,0)(1,1)(2,1)\cdots(d-1,1) \end{cases}$$

We shall prove that the above tree is a generating tree for Dyck paths avoiding `DUDU` as well. Let us consider the generating tree for Dyck paths described in Example 1.4. In this tree, the children of a given Dyck path P are obtained by inserting a new peak `UD` either before a `D` step of the last descending run or at the end of P. To obtain a generating tree for $\texttt{DUDU}$-avoiding paths, we have to remove all the children where this operation creates an occurrence of the forbidden pattern $\texttt{DUDU}$. More precisely, let k be the length of the last descending run of P and write $P = p_1 \cdots p_{i-1}p_i\texttt{D}^k$, where p_i is the rightmost $\texttt{U}$ step of P. If $p_{i-1} = \texttt{U}$, then inserting a new peak is always allowed. Otherwise, if $p_{i-1} = \texttt{D}$, and thus $p_{i-1}p_i = \texttt{DU}$, then we cannot insert a new peak immediately before p_{i+2}, since $p_{i-1}p_ip_{i+1}\texttt{U}$ would be an occurrence of $\texttt{DUDU}$. We then assign to each Dyck path P two labels (k, b), where:

- k is the length of the last descending run of P;
- b is a boolean counter with value 0, if the step that precedes the last $\texttt{U}$ step of P is $\texttt{U}$, and 1 otherwise. We assume $b = 0$ for the path $\texttt{UD}$.

According to what observed above, the following is a generating tree for Dyck paths avoiding $\texttt{DUDU}$:

$$\Omega_2 : \begin{cases} (1,0) \\ (k,0) \longrightarrow (k+1,0)(1,1)(2,1)\cdots(k-1,1)(k,1) \\ (k,1) \longrightarrow (k+1,0)(1,1)(2,1)\cdots(k-1,1) \end{cases}$$

Corollary 6.13. The number of $(132, 321)$-sortable permutation of length n is equal to the number of Dyck paths of semilength n avoiding $\texttt{DUDU}$.

Proof. The rules Ω_1 and Ω_2 are identical. □

An example of the bijection between $\mathrm{Sort}(132, 321)$ and the set of $\texttt{DUDU}$-avoiding Dyck paths induced by the rules Ω_1 and Ω_2 is illustrated in Figure 6.1.

6.4 The $(123, 132)$-machine

In this section we discuss the $(123, 132)$-machine. In [8], the authors show that permutations in $\mathrm{Sort}(123, 132)$ are characterized by the avoidance of four (generalized) patterns of length four. Then they prove that the distribution of the first element in $\mathrm{Sort}(123, 132)$ is given by the well known Catalan triangle (sequence A009766 in [45]). An immediate consequence is that $(123, 132)$-sortable permutations are enumerated by the Catalan numbers. In this thesis we follow an alternative path. We first provide a decomposition lemma for $(123, 132)$-sortable permutations, then we collect several geometric properties of $\mathrm{Sort}(123, 132)$ that lead, towards a step by step procedure, to the same enumerative result.

Although Lemma 6.2 does not apply to $\mathrm{Sort}(123, 132)$, it is still useful to look at the ltr-min decomposition of $(123, 132)$-sortable permutations.

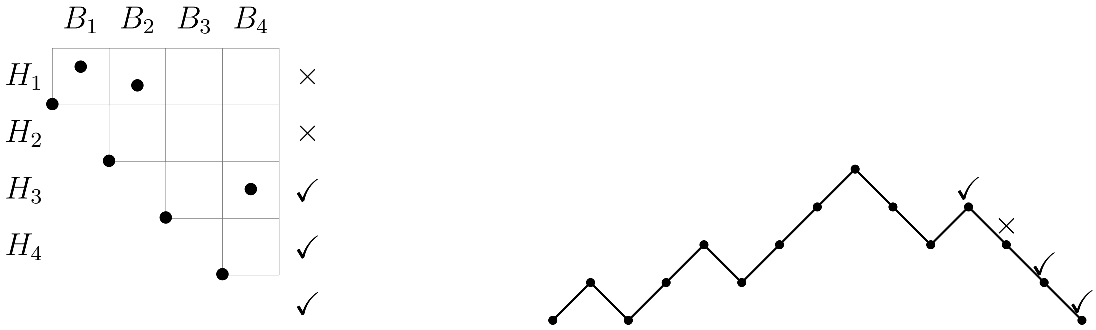

Figure 6.1: On the left, the grid decomposition of the $(132, 321)$-sortable permutation $\pi = 6837214$, where $\mathrm{dep}(\pi) = 3$ and the last block B_4 is not empty. On the right, the corresponding Dyck path P, with $k = 3$ and $b = 1$. Active sites are marked with "✓", while the symbol "×" marks those sites that are not active.

Lemma 6.14. Let $\pi = m_1 B_1 \cdots m_t B_t$ be the ltr-min decomposition of the permutation π. Then:

1. $\mathcal{S}^{123,132}(\pi) = \tilde{B}_1 \tilde{B}_2 m_2 \tilde{B}_3 m_3 \cdots \tilde{B}_k m_k m_1$, where $\tilde{B}_i$ is a rearrangement of B_i, for each i.

2. If π is $(123, 132)$-sortable, then $B_i > B_{i+1}$[3] for each $i = 1, \dots, t-1$.

3. If π is $(123, 132)$-sortable, then B_1 is increasing and $\tilde{B}_1 = \mathcal{R}(B_1)$ is the reverse of B_1.

4. If π is $(123, 132)$-sortable, then $x < m_{i-1}$ for each $x \in B_i$ and $i \geq 3$. Moreover, we have $\tilde{B}_i = \mathcal{S}^{12}(B_i)$.

Proof. 1. Consider the action of the $(123, 132)$-stack on the input permutation π. Note that the element m_1 remains at the bottom of the $(123, 132)$-stack until the end of the sorting procedure. Now, since $m_2 x m_1$ is an occurrence of 132 for each x in B_1, all the elements in B_1 are extracted from the $(123, 132)$-stack before m_2 enters. Then the element m_2 can never play the role of 2 in either 123 or 132, since $m_2 < m_1$ and m_1 is the only element below m_2 in the $(123, 132)$-stack. Therefore m_2 is never involved in any occurrence of either 123 or 132 with m_1 and some elements of B_2. By repeating the same argument, we deduce that the block B_2 has to be extracted from the $(123, 132)$-stack before m_3 enters ($m_3 x m_2 \simeq 132$ for each $x \in B_2$). Then m_2 is extracted too, since $m_3 m_2 m_1 \simeq 123$. Next, m_3 is pushed above m_1. The thesis follows by iterating the same argument on the remaining blocks.

[3] That is $x > y$ for each $x \in B_i, y \in B_{i+1}$.

2. Suppose, for a contradiction, that there are two elements $x \in B_i$ and $y \in B_{i+1}$ such that $x < y$. Then, due to what proved in the previous item, $\mathcal{S}^{123,132}(\pi)$ contains an occurrence xym_t of 231, contradicting the fact that π is $(123, 132)$-sortable.

3. Suppose, for a contradiction, that B_1 is not increasing. Write $B_1 = x_1 \cdots x_k x_{k+1} \cdots$, where $x_k > x_{k+1}$ is the leftmost descent of B_1. Then the elements $x_1, \dots, x_k, x_{k+1}$ are pushed into the $(123, 132)$-stack and $S^{123,132}(\pi)$ contains an occurrence $x_{k+1}x_k m_1$ of 231, contradicting the fact that π is $(123, 132)$-sortable. Thus B_1 is increasing. In particular, each element of B_1 can be pushed into the $(123, 132)$-stack, so that $\tilde{B}_1 = \mathcal{R}(B_1)$.

4. Let $i \geq 3$. If B_i contains an element $x > m_{i-1}$, then $\mathcal{S}^{123,132}(\pi)$ contains an occurrence $m_{i-1}xm_i$ of 231, a contradiction with π $(123, 132)$-sortable. Finally, we show that $\tilde{B}_i = \mathcal{S}^{12}(B_i)$. Consider the instant when m_i is pushed into the $(123, 132)$-stack, that is as soon as the first element of B_i is the next one to be processed. As a consequence of what said in the first item of this lemma, at this step the $(123, 132)$-stack contains $m_i m_1$, with m_i above m_1. We show that, on B_i, the behavior of the $(123, 132)$-stack is equivalent to the behavior of a 12-stack that ignores $m_i m_1$. Suppose that the 12-stack performs a pop operation. In other words, the restriction of the 12-stack is triggered by an occurrence $y_2 y_1$ of 12, where y_2 is the next element of the input and y_1 is in the 12-stack. Notice that $m_1 > x$ for each $x \in B_i$, since we have just proved that $x < m_{i-1}$ and obviously $m_{i-1} < m_1$. Thus $y_2 y_1 m_1 \simeq 123$ and the $(123, 132)$-stack performs a pop operation too. Conversely, suppose that the $(123, 132)$-stack performs a pop operation, that is the restriction of the $(123, 132)$-stack is triggered by an occurrence $y_3 y_2 y_1$ of either 123 or 231, where y_3 is the next element of the input. If $y_3 y_2 y_1 \simeq 123$, then $y_2 \neq m_i$, since $m_1 > m_2$ and $y_3 < y_2$. Therefore y_2 and y_3 are elements of B_i and the restriction of the 12-stack is triggered by $y_3 y_2 \simeq 12$, as desired. Otherwise, suppose that $y_3 y_2 y_1 \simeq 132$. Note that both m_1 and m_i cannot be part of this occurrence. Indeed $m_1 \neq y_1$, since $m_1 > x$ for each $x \in B_i$, and $m_i \neq y_1, y_2$ since $m_i < x$ for each $x \in B_i$. Therefore $y_3 < y_1$ are elements of B_i that trigger the restriction of the 12-stack. This completes the proof.

□

What proved so far determines the structure of a $(123, 132)$-sortable permutation $\pi = m_1 B_1 \cdots m_t B_t$, except for the second block B_2. Indeed, since B_1 is increasing and $B_i > B_{i+1}$ for each i, then B_1 contains the biggest elements of π in increasing order. Then, for each $i \geq 3$, the block B_i is a 213-avoiding permutation due to Theorem 3.1. Moreover, since $m_i < x < m_{i-1}$ for each $x \in B_i$, elements in m_iB_i are strictly greater than elements in $m_{i+1}B_{i+1}$. This guarantees that occurrences of 213 are not created if we juxtapose m_iB_i and $m_{i+1}B_{i+1}$. As a consequence, we can equivalently say that $m_3B_3 \cdots m_tB_t$ is a 213-avoiding permutation. Next we describe the structure of the remaining block B_2. For the rest of this section, let $\xi = (132, \{0,2\}, \emptyset)$ be the bivincular pattern depicted in Figure 3.1. A geometric description of the set $\mathfrak{S}(\xi)$ was provided in Section 3.3. Recall that an occurrence of ξ in a permutation $\pi = \pi_1 \cdots \pi_n$ is a descent $\pi_i > \pi_{i+1}$ such that $\pi_{i+1} > \pi_1$. Finally, denote by $*213$ an occurrence of either 4213, 3214, 2314, or 1324. Note that for any permutation $\pi = \pi_1 \cdots \pi_n$, we have $\pi \in \mathfrak{S}(*213)$ if and only if $\pi_2 \cdots \pi_n$ avoids 213.

Lemma 6.15. Let $\pi = m_1m_2B_2$ be a permutation with two ltr-minima and such that the first block in the ltr-min decomposition is empty. Then π is $(123, 132)$-sortable if and only if $\pi \in \mathfrak{S}(\xi, *213)$.

Proof. By Lemma 6.14 we have $\mathcal{S}^{123,132}(\pi) = \tilde{B}_2m_2m_1$. Suppose initially that π is $(123, 132)$-sortable. We wish to prove that π avoids ξ and m_2B_2 avoids 213. Suppose, for a contradiction, that π contains an occurrence $m_1\pi_i\pi_{i+1}$ of ξ. If π_{i+1} enters the $(123, 132)$-stack before π_i is extracted, then $\mathcal{S}^{123,132}(\pi)$ contains an occurrence $\pi_{i+1}\pi_im_2$ of 231, a contradiction with π being $(123, 132)$-sortable. Therefore π_i must be extracted when π_{i+1} is the next element of the input. Let $m_1m_2x_1 \cdots x_l\pi_i$ be the elements contained in the $(123, 132)$-stack at this point. Notice that it has to be $x_1 < x_2 < \cdots < x_l < c$, otherwise $\mathcal{S}^{123,132}(\pi)$ would contain an occurrence of 231 (with m_2 playing the role of 1), which is impossible. Now, starting from the top of the stack (which is $\pi_i > \pi_{i+1}$) and going down, consider the last element y such that $y > \pi_{i+1}$. Since $m_2 < \pi_{i+1}$, it must be either $y = \pi_i$ or $y = x_j$, for some j. In any case, π_{i+1} enters above y and $\mathcal{S}^{123,132}(\pi)$ contains an occurrence $\pi_{i+1}ym_2$ of 231, which is again a contradiction.
Otherwise, suppose, for a contradiction, that m_2B_2 contains an occurrence $\pi_i\pi_j\pi_k$ of 213. Notice that $i > 2$, since $m_2 = 1$. If $\pi_j < m_1$, then $\pi_j\pi_im_1$ is an occurrence of either 321 or 231. Thus π_i has to be extracted before π_j enters the $(123, 132)$-stack. But this results in an occurrence $pi_i\pi_jm_2$ of 231 in $\mathcal{S}^{123,132}(\pi)$, which contradicts the fact that π is $(123, 132)$-sortable. Therefore we can suppose $\pi_j > m_1$. If $j = i + 1$, then $\pi_1\pi_i\pi_{i+1}$ is an occurrence of ξ and we are back to the previous case. Instead, if $j > i + 1$, consider the element π_{i+1}. If $\pi_{i+1} < \pi_j$, then we repeat the same argument replacing π_j with π_{i+1}. Finally, if $\pi_{i+1} > \pi_j$, we repeat the same argument replacing π_i with π_{i+1}. Sooner or later this will lead to a contradiction.

Conversely, suppose that π is not $(123, 132)$-sortable. Equivalently, let bca be an occurrence of 231 in $\mathcal{S}^{123,132}(\pi) = \tilde{B}_2m_2m_1$. Note that $m_2 \neq b, c$, since $m_2 < m_1$. Therefore b and c are elements of $\tilde{B}_i$. We distinguish two cases. Initially, suppose that b precedes c in π (and in $\mathcal{S}^{123,132}(\pi)$ as well). Therefore b is extracted from the $(123, 132)$-stack before c enters. Let y be the next element of the input when b is extracted. If $y < b$, then byc is an occurrence of 213 in m_2B_2, as desired. Otherwise, let $y > b$. Since the restriction of the $(123, 132)$-stack is triggered by y, there are two elements into the stack u, v, with u below v, such that yvu is an occurrence of either 123 or 132. Since we assumed $y > b$, we have $bvu \simeq yvu$ and thus bvu is an occurrence of either 123 or 132 inside the $(123, 132)$-stack, which is impossible. Finally, suppose that c precedes b in π. Since b is extracted before c, c is still inside the $(123, 132)$-stack when b enters. This implies that $b > m_1$, otherwise bcm_1 would be an occurrence of either 123 (if $m_1 > c$) or 132 (if $m_1 < c$), which is impossible. Similarly, for each element x between c and b in π we have $x > m_1$. Now, if c and b are consecutive in position, then m_1cb is an occurrence of ξ, as wanted. Otherwise, let x be the element immediately after c in π. If $x < c$, then m_1cx is the desired occurrence of ξ. Otherwise, if $c < x$ (and thus $x > b$), we can repeat the same argument using x instead of c. □

Lemma 6.16. Let π be a permutation of length n. Let $\pi = m_1m_2B_2m_3B_3 \cdots m_tB_t$ be the ltr-min decomposition of π and suppose that B_1 is empty. For $k \geq 1$, define:

$$\bar{\pi} = m_1(n+1)(n+2)\dots(n+k)m_2B_2 \cdots m_tB_t.$$

Then π is $(123, 132)$-sortable if and only if $\bar{\pi}$ is $(123, 132)$-sortable.

Proof. Let $k \geq 1$ and suppose that π is $(123, 132)$-sortable. Consider the action of the $(123, 132)$-stack on input $\bar{\pi}$. Since $m_1 < n+1 < n+2 < \cdots < n+k$, the prefix of $\bar{\pi}$ up to $n+k$ is pushed into the $(123, 132)$-stack. Then, since $m_2(n+1)m_1 \simeq 132$, the elements $n+k, n+k-1, \dots, n+1$ are extracted from the $(123, 132)$-stack. Therefore we have:

$$\mathcal{S}^{123,132}(\bar{\pi}) = (n+k)(n+k-1)\dots(n+1)\mathcal{S}^{123,132}(\pi).$$

It is easy to realize that $\mathcal{S}^{123,132}(\pi)$ contains 231 if and only if $\mathcal{S}^{123,132}(\bar{\pi})$ contains 231, thus the thesis follows. □

Corollary 6.17. Let $\pi = m_1B_1 \cdots m_tB_t$ be the ltr-min decomposition of the permutation π. Then π is $(123, 132)$-sortable if and only if the following four conditions are satisified:

1. $B_i > B_{i+1}$ for each $i = 1, \dots, t-1$;
2. B_1 is increasing;
3. $m_1m_2B_2 \in \mathfrak{S}(\xi, *213)$;
4. $m_3B_3 \cdots m_tB_t \in \mathfrak{S}(213)$.

Proof. It is an immediate consequence of lemmas 6.14, 6.15 and 6.16. □

The structural characterization of Corollary 6.17 can be exploited in order to compute the number of $(123, 132)$-sortable permutations. For $n \geq 1$ and $1 \leq k \leq n$, let $b_{n,k}$ be the (n, k)-th ballot number. The triangle of ballot numbers is also known as the Catalan triangle (sequence A009766 in [45], see Figure 6.2). Let $\tilde{b}_{n,k} = b_{n,n+1-k}$ be the triangle obtained by reversing its rows (sequence A033184 in [45]).

Lemma 6.18. 1. $\displaystyle\sum_{i=1}^{s} \tilde{b}_{s,i} \binom{n-1-s+i}{i} = b_{n,s+1}.$

2. $\displaystyle 1 + \sum_{s=1}^{n-1} b_{n,s+1} = \mathfrak{c}_n.$

Proof. 1. We have:

$$\sum_{i=1}^{s} \tilde{b}_{s,i} \binom{n-1-s+i}{i} = \sum_{i=1}^{s} b_{s,s+1-i} \binom{n-1-s+i}{i} = \sum_{j=1}^{s} b_{s,j} \binom{n-j}{s-j+1}.$$

We shall prove that

$$b_{n,s+1} = \sum_{j=1}^{s} b_{s,j} \binom{n-j}{s-j+1}$$

by showing that each coefficient $b_{s,j}$ contributes $\binom{n-j}{s-j+1}$ times to $b_{n,s+1}$ (see Figure 6.2). It is well known that ballot numbers are defined, for example, by the recurrence:

$$b_{n,s+1} = \sum_{i=1}^{s+1} b_{n-1,i}.$$

In other words, in the triangle $b_{n,k}$ each coefficient $b_{n,k}$ is obtained by summing the coefficients of indices $1, 2, \ldots, k$ of the previous row. Let us consider the coefficient $b_{s,j}$, for some j. The contribution of $b_{s,j}$ to $b_{n,s+1}$ is obtained by choosing (see again Figure 6.2):

- a coefficient b_{s+1,t_1} in the $(s+1)$-th row, with $t_1 \in \{j, j+1, \ldots, s+1\}$;
- a coefficient b_{s+2,t_2} in the $(s+2)$-th row, with $t_1 \leq t_2 \leq s+1$;

$n\backslash k$	1	2	3	4	5	6	7
1	1						
2	1	1					
3	1	2	2				
4	1	3	5	5			
5	1	4	9	14	14		
6	1	5	14	28	42	42	
7	1	6	20	48	90	132	132

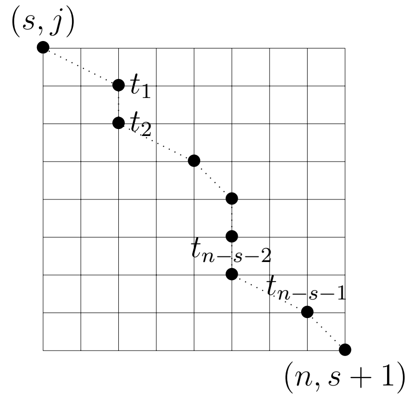

Figure 6.2: The triangle of the ballot numbers, on the left, and the construction described in Lemma 6.18, on the right.

- a coefficient b_{s+3,t_3} in the $(s+3)$-th row, with $t_2 \le t_3 \le s+1$;

$\vdots$

- a coefficient $b_{n-1,t_{n-s-1}}$ in the $(n-1)$-th row, with $t_{n-s-2} \le t_{n-s-1} \le s+1$.

Finally, there are $\binom{(n-s-1)+(s+1-j+1)-1}{(s+1-j+1)-1} = \binom{n-j}{s-j+1}$ sequences $t_1 \le t_2 \le \cdots \le t_{n-s-1}$, with values in $\{j, j+1, \ldots, s+1\}$, therefore the desired equality follows.

2. Since $b_{n,1} = 1$, we have

$$1 + \sum_{s=1}^{n-1} b_{n,s+1} = b_{n,1} + \sum_{t=2}^{n} b_{n,t} = \sum_{t=1}^{n} b_{n,t} = \mathfrak{c}_n.$$

$\square$

Theorem 6.19. Let $f_n = f_n^{123,132}$ be the number of $(123,132)$-sortable permutations of length n. Let g_n be the number of $(123,132)$-sortable permutation of length n with at least two ltr-minima and where the first block B_1 in the ltr-min decomposition is empty. Denote by h_n the number of permutations π in $\mathfrak{S}_n(\xi, *213)$ such that $\pi_2 = 1$ (i.e. where π_1 and π_2 are the only two ltr-minima of π.). Then:

1. $h_1 = 0$ and $h_{n+1} = \mathfrak{c}_n$, for each $n \ge 1$.
2. $g_1 = 0$ and $g_{n+1} = \mathfrak{c}_{n+1} - \mathfrak{c}_n$, for each $n \ge 1$.

3. $f_n = \mathfrak{c}_n$, for each $n \geq 1$.

Proof. 1. Let $\pi \in \mathfrak{S}(\xi, *213)$ be a permutation of length at least two and suppose that $\pi_2 = 1$. Let $\pi_1 = k$ and write:

$$\pi = k1A_1\alpha_1A_2\alpha_2 \cdots A_s\alpha_sA_{s+1},$$

where $\alpha_i < k$ for each i and the elements of the blocks $A_1, \ldots, A_{s+1}$ are greater than k. Let i be the minimum index such that $\alpha_i > \alpha_{i+1}$. Note that A_j is empty for each $j \geq i+2$. Otherwise, if $x \in A_j$, then $\pi_1\alpha_i\alpha_{i+1}x$ would be an occurrence of $*213$ in π, which contradicts the hypothesis. Analogously, it must be $x > y$ for each $x \in B_j$ and $y \in B_{j+1}$, with $j = 1, \ldots, i$, or else $\pi_1 x\alpha_j y$ would be an occurrence of $*213$. Moreover, the block B_j is increasing for each $j = 1, \ldots, i+1$. Otherwise, a descent $x > y$ in B_j would result in an occurrence $\pi_1 xy$ of ξ, which is impossible. What observed so far is enough to completely characterize a permutation $\pi \in \mathfrak{S}_{n+1}(\xi, *213)$, with $\pi_2 = 1$. Such permutation is either the increasing permutation or it can be constructed as follows (see Figure 6.3):

- Choose a permutation $\alpha \in \mathfrak{S}_s(213)$, for some $1 \leq s \leq n-1$. Let i be the index of the leftmost descent in α (if α is the increasing permutation, let $i = n-1$).
- Choose how to distribute $n-1-s$ elements in the blocks $A_1, \ldots, A_i$. As a consequence of what observed above, the blocks are increasing and the relative order of the blocks is uniquely determined, therefore there are $\binom{n-1-s+i}{i}$ different ways to distribute such elements.
- Finally, add the two initial elements $\pi_1\pi_2 = k1$.

Now, it is well known that the the number of 213-avoiding permutations of length s, where i is the index of the leftmost descent, is given by $\tilde{b}_{s,i}$. Therefore we have:

$$h_{n+1} = 1 + \sum_{s=1}^{n-1} \left(\sum_{i=1}^{s} \tilde{b}_{s,i} \binom{n-1-s+i}{i} \right),$$

and the thesis follows from Lemma 6.18.

2. Due to Corollary 6.17, each $(123, 132)$-sortable permutation $\pi = m_1m_2B_2 \cdots m_tB_t$, where B_1 is empty, is obtained by juxtaposing a 213-avoiding permutation $m_3B_3 \cdots m_tB_t$ to a permutation $m_1m_2B_2 \in$

$\mathfrak{S}(\xi, *213)$. By summing over the length j of the 213-avoider (note that $0 \leq j \leq n-1$), we obtain:

$$g_{n+1} = \sum_{j=0}^{n-1} h_{n+1-j}\mathfrak{c}_j = \sum_{j=0}^{n-1} \mathfrak{c}_{n-j}\mathfrak{c}_j.$$

Finally, it is well known that:

$$\mathfrak{c}_{n+1} = \sum_{j=0}^{n} \mathfrak{c}_{n-j}\mathfrak{c}_j = \mathfrak{c}_n + \sum_{j=0}^{n-1} \mathfrak{c}_{n-j}\mathfrak{c}_j,$$

thus the thesis follows.

3. Let π be a $(123, 132)$-sortable permutation of length n. If π has one ltr-minimum, then π is the increasing permutation due to Lemma 6.14. Otherwise, according to Lemma 6.16, π is obtained from a $(123, 132)$-sortable permutation, with B_1 empty, by inserting $k \geq 0$ consecutive ascents $k+1, k+2, \dots, n-1, n$ immediately after π_1. Therefore, summing over k, we have:

$$f_n = 1 + \sum_{k=2}^{n} g_n = \mathfrak{c}_1 + \sum_{k=2}^{n} (\mathfrak{c}_n - \mathfrak{c}_{n-1}) = \mathfrak{c}_n.$$

□

6.5 The $(123, 312)$-machine

In this section we provide a structural description of $(123, 312)$-sortable permutations and define a generating tree for the set $\mathrm{Sort}(123, 312)$. The enumeration of $\mathrm{Sort}(123, 312)$, which is given by the binomial transform of the Catalan numbers, was proved in [8] by means of a bijection with a family of pattern-avoiding partial permutations.

Let $\pi = M_1B_1 \cdots M_tB_t$ the ltr-max decomposition of a permutation π. By Lemma 6.1, we have $\mathcal{S}^{123,312}(\pi) = \tilde{B}_1 \cdots \tilde{B}_tM_t \cdots M_1$, where $\tilde{B}_i$ is a suitable rearrangement of B_i. Due to the same theorem, if π is $(123, 312)$-sortable and has length n, then $M_j = n - t + j$, for each $j = 1, \dots, t$.

Theorem 6.20. Let $\pi = M_1B_1 \cdots M_tB_t$ be a $(123, 312)$-sortable permutation. Then:

1. B_i avoids 213 for each i.

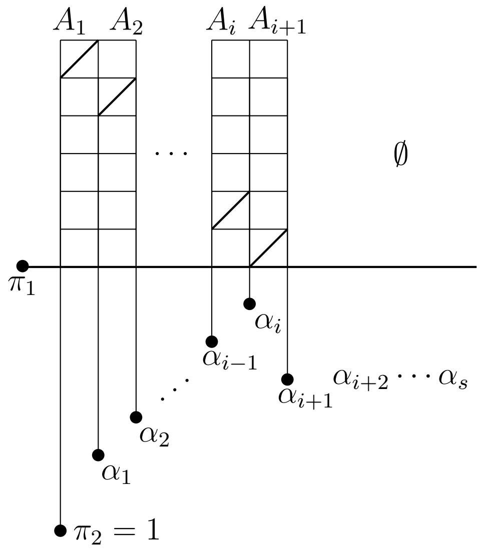

Figure 6.3: The geometric structure of a permutation $\pi \in \mathfrak{S}(\xi, *213)$, with $\pi_2 = 1$.

2. $\tilde{B}_i = \mathcal{S}^{12}(B_i)$, for each i.

Proof. Let $i \geq 2$. Notice that, as a consequence of Lemma 6.1, immediately after the push of M_i into the $(123, 312)$-stack, the $(123, 312)$-stack contains the elements $M_i \cdots M_2 M_1$, reading from top to bottom. Moreover, these elements remain at the bottom of the $(123, 312)$-stack until the end of the sorting procedure, since they are the last elements of $\mathcal{S}^{123,312}(\pi)$. This fact will be used for the rest of the proof.

1. Suppose, for a contradiction, that B_i contains an occurrence bac of 213, for some i. Therefore, since abM_i is an occurrence of 123, b is extracted from the $(123, 312)$-stack before a enters. Since π is $(123, 312)$-sortable, this implies that a is then extracted before c enters, otherwise bca would be an occurrence of 231 in $\mathcal{S}^{123,312}(\pi)$. Consider the instant when a is extracted from the $(123, 312)$-stack. Let x be the next element of the input. Since a pop operation is performed, there must be two elements y, z in the $(123, 312)$-stack, with y above z, such that xyz is an occurrence of either 123 or 312. Notice that $z = M_j$, for some $j \leq i$. Otherwise, if $z \in B_i$, then yzM_i would be an occurrence of 123 in the $(123, 312)$-stack, a contradiction. Since $y < z$, we necessarily have $y \in B_i$ (y cannot be an ltr-maximum). Now, xyz is not an occurrence of 312. Otherwise it would be $x > z = M_j$ and thus x would be an ltr-maximum. But since x precedes $c \in B_i$ (it could be $x = c$), this is impossible. Therefore $xyz \simeq 123$. If $x > a$, then it would be $y > a$ as well and thus $M_i y a$ would be an occurrence of 123 in the $(123, 312)$-stack, which

is impossible. So we have $x < a$ and bxc is an occurrence of 213 in B_i (note that x is strictly to the right of a), thus we can repeat the same argument using bxc instead of bac, until we eventually find a contradiction.

2. Let us consider the action of the $(123, 312)$-stack on the block B_i. We wish to show that the behavior of the $(123, 312)$-stack when processing B_i is equivalent to the behavior of an empty 12-stack on input B_i. In other words, we prove that the restriction of the $(123, 312)$-stack is triggered if and only if the next element of the input forms an occurrence of 12 together with some other element in the $(123, 312)$-stack. As soon as M_i enters the $(123, 312)$-stack (and the first element of B_i is the next one to be processed), the $(123, 312)$-stack contains the elements $M_i \cdots M_2 M_1$, reading from top to bottom. Observe that B_i avoids 213 by what proved above, therefore the $(123, 312)$-stack cannot be triggered by an occurrence of 312 when processing B_i. Suppose that the next element of the input x forms an occurrence xy of 12 with some $y \in B_i$. Then xyM_i is an occurrence of 123 in the $(123, 312)$-stack, as desired. Conversely, suppose that the $(123, 312)$-stack is triggered by an occurrence of xyz of 123, where x is the next element of the input. Since $M_i > M_{i-1} > \cdots > M_1$, it must be $y \in B_i$. Thus xy is an occurrence of 12 that triggers the 12-stack, as wanted.

□

As a consequence of what proved so far in this section, for any $(123, 312)$-sortable permutation $\pi = M_1 B_1 \cdots M_t B_t$ of length n, we have $B_i \in \mathrm{Sort}(213)$ and $M_1, \ldots, M_t = n - t + 1, \ldots, n$. Moreover, $\mathcal{S}^{123,312}(\pi) = \tilde{B}_1 \cdots \tilde{B}_t M_t \cdots M_1$, where $\tilde{B}_i$ is decreasing. Therefore, for any three elements x, y, z, with $x \in B_i$, $y \in B_j$ and $z \in B_k$, with $i < j \leq k$, xyz cannot be an occurrence of 231. Otherwise xyz would still be an occurrence of 231 in $\mathcal{S}^{123,312}(\pi)$, contradicting the fact that π is $(123, 312)$-sortable. From now on, we say that xyz is an occurrence of $2-3-1$ if $xyz \simeq 231$, with $x \in B_i$, $y \in B_j$, $z \in B_k$ and $i < j < k$. In the analogous case, but when $j = k$, we say that xyz is an occurrence of $2-31$.

Theorem 6.21. Let $\pi = M_1 B_1 \cdots M_t B_t$ be the ltr-max decomposition of a permutation of length n. Let $\mathcal{S}^{123,312}(\pi) = \tilde{B}_1 \cdots \tilde{B}_t M_t \cdots M_1$. Then π is $(123, 312)$-sortable if and only if the following four conditions are satisfied:

1. $M_j = n - t + j$, for each $j = 1, \ldots, t$.

2. B_i avoids 213 for each i (and thus $\tilde{B}_i$ is decreasing for each i).

3. π avoids $2-3-1$.

4. π avoids $2-31$.

Proof. If π is $(123, 312)$-sortable, then π satisfies all the above conditions as a consequence of what proved before in this section. Conversely, it is easy to check that, if π satisfies the above conditions, then $\mathcal{S}^{123,312}(\pi)$ avoids 231. We leave this part of the proof to the reader. □

In [8], the structural description of Theorem 6.21 is reformulated in terms of avoidance of (generalized) patterns. This ultimately leads to a bijection between Sort$(123, 312)$ and the set of partial permutations avoiding the pattern 213, whose enumeration is given by the binomial transform of the Catalan numbers (sequence A007317 in [45]). We refer the reader to [8] for a definition of partial permutations, as well as for a detailed proof of the results mentioned above. In what follows we provide a generating tree for Sort$(123, 312)$. As usual, we wish to generate all the permutations in Sort$(123, 312)$ by inserting a new rightmost element (and rescaling the others). Before doing that, we reformulate the third condition of Theorem 6.21 in the following lemma, whose easy proof is omitted.

Lemma 6.22. Let $\pi = M_1B_1 \cdots M_tB_t$ be the ltr-max decomposition of the $(123, 312)$-sortable permutation π. Let $\mathcal{S}^{123,312}(\pi) = \tilde{B}_1 \cdots \tilde{B}_tM_t \cdots M_1$. Then $\mathcal{S}^{123,312}(\pi)$ avoids $2-31$ if and only if for each $x \in B_i$, $y \in B_j$, with $i < j$, we have:

- if $y > x$, then $B_j > x$.
- If $y < x$, then $B_j < x$.

In other words, due to Lemma 6.22, each block B_j of a $(123, 312)$-sortable permutation π is bounded between two previous elements of π. Now, let $\pi = M_1B_1 \cdots M_tB_t$ be a $(123, 312)$-sortable permutation of length n. We distinguish three possible ways to insert a new rightmost element x in order to get a $(123, 312)$-sortable permutation:

(A) Insert a new ltr-maximum $x = n + 1$;

(B) if B_t is empty, insert the first element of B_t;

(C) if B_t is not empty, insert an element in B_t.

In order to provide a generating tree for $(123, 312)$-sortable permutations, we assign to each element of Sort$(123, 312)$ two labels (k, m). The label k denotes the number of active sites of the current block B_t, thus it takes into account the insertion of x according to (C). Due to Lemma 6.22, the relative position of the block B_t with respect to the previous blocks is uniquely determined by its first element. Therefore we just have to make sure that B_t avoids 213, as stated in Theorem 6.21, and the active sites related to the label k will be Catalan-type:

$$(k) \to (2)(3)\cdots(k)(k+1).$$

The other label m denotes the number of active sites with respect to the relative order of the blocks and it takes into account the insertion of x according to (B). More precisely, since we have to avoid creating an occurrence of $2-3-1$, the label m will be Catalan-type too:

$$(m) \to (2)(3)\cdots(m)(m+1).$$

Notice that inserting x according to (C) affects the label m as well. Indeed a new element in the block B_t creates one additional active site for the relative order of the blocks (so m is increased by one). Finally, inserting a new ltr-maximum $x = n+1$ according to (A) always produces a permutation where the last block is empty, that is where $k = 1$. Note that this operation does not affect the label m.

Theorem 6.23. The following rule provides a generating tree for $\mathrm{Sort}(123, 312)$:

$$\Omega : \begin{cases} (1,0) \longrightarrow (1,0)(2,2) \\ (1,m) \longrightarrow (1,m)(2,2)(2,3)\cdots(2,m)(2,m+1), \;\; m \geq 2 \\ (k,m) \longrightarrow (1,m)(2,m+1)(3,m+1)\cdots(k+1,m+1), \;\; k,m \geq 2 \end{cases} \tag{6.1}$$

Proof. Let $\pi = M_1B_1\cdots M_tB_t$ be a $(123, 312)$-sortable permutation with label (k, m), for some integers k, m. Suppose we insert a new rightmost element x. This can be done according to either condition (A), (B) or (C), as described below Lemma 6.22. We discuss each case separately.

- If $(k, m) = (1, 0)$, then $\pi = 1\cdots n$ consists solely of ltr-maxima, since $m = 0$. If x is a new ltr-maximum, then the label of the resulting permutation is again $(1, 0)$. Otherwise, if x is the first (and only) element of B_t, then the resulting label is $(2, 2)$.

- If $(k, m) = (1, m)$, for some $m \geq 2$, then the last element of π is M_t and B_t is empty. If x is a new ltr-maximum, then the resulting label is $(1, m)$, as noted above. Otherwise, x is the first element of B_t, according to (B). Then the behavior of the label m is Catalan-type, according to the relative order of the blocks. Moreover, the label k of any resulting permutation is always 2, since x is the only element of B_t in the resulting permutation.

- Finally, let $k, m \geq 2$. Then, since $k \geq 2$, the block B_t is not empty (and it has k Catalan-type active sites). Again if x is a new ltr-maximum, then the resulting label is $(1, m)$. Otherwise, x is an element of B_t. Therefore the behavior of the label k is Catalan-type, according to (C). The label m, instead, is increased by one, as noted before.

□

The problem of showing directly that the family of objects generated by the rule of Theorem 6.23 is counted by the binomial transform of the Catalan numbers remains open.

Open Problem 6.1. Prove directly that objects generated by Rule 6.23 are counted by the binomial transform of Catalan numbers.

6.6 The family of $(\sigma,\widehat{\sigma})$-machines

We end this chapter by mentioning a result for a family of pairs of patterns. Let σ be a permutation of length k. Recall that $\widehat{\sigma} = \sigma_2\sigma_1\sigma_3 \cdots \sigma_k$ is the permutation obtained by interchanging the first two entries of σ. Then the map $\pi \mapsto \mathcal{S}^{\sigma,\widehat{\sigma}}(\pi)$ is bijective from $\mathrm{Sort}(\sigma, \widehat{\sigma})$ to $\mathfrak{S}(231)$. More precisely, each $(\sigma, \widehat{\sigma})$-sortable permutation π is obtained uniquely from a 231-avoiding permutation α as:

$$\pi = \mathcal{R}(\mathcal{S}^{\sigma,\widehat{\sigma}}(\mathcal{R}(\alpha)).$$

The above equality gives a constructive description of $\mathrm{Sort}(\sigma, \widehat{\sigma})$. Indeed it follows immediately that:

$$\mathrm{Sort}(\sigma, \widehat{\sigma}) = \mathcal{R}(\mathcal{S}^{\sigma}(\mathcal{R}(\mathfrak{S}(231)))$$

A proof of this fact (which will be obtained as a corollary of a much more general result for Cayley permutations) is postponed to Remark 8.2 in Section 8.1.1. An immediate consequence is that $f_n^{\sigma,\widehat{\sigma}} = \mathfrak{c}_n$, for each permutation σ and $n \geq 1$.

Chapter 7

Dynamical aspects of the σ-machine

In this chapter we analyze some dynamical aspects of σ-machines by regarding a σ-stack as an operator $\pi \mapsto \mathcal{S}^\sigma(\pi)$. This approach has been adopted recently for pattern-avoiding machines in [8, 11, 20, 28]. Suppose to perform a deterministic sorting procedure. Then it is natural to study the map $\mathcal{S}$ that associates to an input string π the (uniquely determined) output of the sorting process. Some of the problems that arise, and are classically considered, are the following:

- Determine the fertility of a string, which is the number of its pre-images under the map $\mathcal{S}$.
- Determine the image of $\mathcal{S}$, i.e. the strings with positive fertility. These are often called sorted permutations [16].

In what follows, we define some properties of the operator $\mathcal{S}^\sigma$ associated to a σ-stack, then we start to collect the first related results.

7.1 Sorted permutations and fertility

Let σ be a permutation of length two or more. The map $\mathcal{S}^\sigma$ is defined by:

$$\begin{aligned} \mathcal{S}^\sigma : \mathfrak{S} &\to \mathfrak{S} \\ \pi &\mapsto \mathcal{S}^\sigma(\pi). \end{aligned}$$

Define $\mathrm{Sorted}(\sigma) = \mathcal{S}^\sigma(\mathrm{Sort}(\sigma))$. Permutations in $\mathrm{Sorted}(\sigma)$ are thus images of σ-sortable permutations through the map $\mathcal{S}^\sigma$. We call them the *σ-sorted* permutations. Notice that $\mathrm{Sorted}(\sigma) = \mathcal{S}^\sigma(\mathfrak{S}) \cap \mathfrak{S}(231)$.

Remark 7.1. In the adopted definition, σ-sorted permutations are those that are both output of the σ-stack and 12-sortable. A different notion (of sorted permutations) can be obtained by considering the set $\mathcal{S}^\sigma(\mathfrak{S})$, thus including all the possible outputs of the σ-stack. This framework will be adopted in Chapter 8 in the (more general) case of Cayley permutations.

The σ-*fertility* of a permutation π is

$$\mathrm{fert}^\sigma(\pi) = |\left(\mathcal{S}^\sigma\right)^{-1}(\pi)|.$$

Due to the adopted definition of σ-sorted permutations, we have:

$$|\mathrm{Sort}_n(\sigma)| = \sum_{\gamma \in \mathrm{Sorted}(\sigma)} \mathrm{fert}^\sigma(\gamma).$$

A permutation σ is *surjective* if $\mathrm{Sorted}(\sigma) = \mathfrak{S}(231)$. A permutation σ is *injective* if $\mathrm{fert}^\sigma(\pi) \leq 1$ for each $\pi \in \mathfrak{S}(231)$. We say that σ is *bijective* if σ is both surjective and injective. In this case, for each $n \geq 1$, we have:

$$|\mathrm{Sort}_n(\sigma)| = \sum_{\gamma \in \mathrm{Sorted}(\sigma)} \mathrm{fert}^\sigma(\gamma) = \sum_{\gamma \in \mathfrak{S}_n(231)} 1 = \mathfrak{c}_n,$$

where $\mathfrak{c}_n$ is the n-th Catalan number. Equivalently, σ is said to be injective, surjective or bijective if the (restricted) map $\mathcal{S}^\sigma : \mathrm{Sort}(\sigma) \to \mathfrak{S}(231)$ is respectively injective, surjective or bijective. Finally, σ is said to be *effective* if $\mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(\sigma)$, that is the σ-stack succesfully performs its task of preventing occurrences of σ to be produced. One of the main goals of this chapter is to characterize which patterns are effective. First we prove a simple lemma that leads to an equivalent definition of effectiveness.

Lemma 7.1. We have:

$$\mathfrak{S}(231, \sigma) \subseteq \mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(231).$$

Proof. The inclusion $\mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(231)$ has already been noted. For the other inclusion, observe that every σ-avoiding permutation is equal to the output of $\mathcal{S}^\sigma$ on its reverse, and therefore, every such permutation that also avoids 231 belongs to $\mathrm{Sorted}(\sigma)$. Moreover, by Lemma 3.4 we have:

$$\mathcal{S}^\sigma(\mathfrak{S}(132, \mathcal{R}(\sigma)) = \mathcal{R}(\mathfrak{S}(132, \mathcal{R}(\sigma)) = \mathfrak{S}(231, \sigma),$$

and thus $\mathfrak{S}(231, \sigma) \subseteq \mathrm{Sorted}(\sigma)$. $\square$

Corollary 7.2. A pattern σ is effective if and only if $\mathrm{Sorted}(\sigma) = \mathfrak{S}(231, \sigma)$.

Proof. If σ is effective, that is $\mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(\sigma)$, then by Lemma 7.1 we have $\mathfrak{S}(231, \sigma) \subseteq \mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(231, \sigma)$ and thus $\mathrm{Sorted}(\sigma) = \mathfrak{S}(231, \sigma)$. The other implication is trivial. $\square$

Theorem 7.3. Let σ be a permutation of length two or more. If $\widehat{\sigma} \geq 231$, then $\mathrm{Sorted}(\sigma) = \mathfrak{S}(231, \sigma)$ and σ is both injective and effective.

Proof. Suppose that $\widehat{\sigma} \geq 231$. By Theorem 3.6, we have $\mathrm{Sort}(\sigma) = \mathfrak{S}(132, \mathcal{R}(\sigma))$. Therefore, if $\pi \in \mathrm{Sort}(\sigma)$, then π avoids $\mathcal{R}(\sigma)$ and thus $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$. In other words, the operator $\mathcal{S}^\sigma$ acts as the reverse operator on $\mathrm{Sort}(\sigma)$ and thus the σ-machine is injective. Finally, we have:

$$\mathrm{Sorted}(\sigma) = \mathcal{R}(\mathrm{Sort}(\sigma)) = \mathcal{R}(\mathfrak{S}(132, \mathcal{R}(\sigma))) = \mathfrak{S}(231, \sigma),$$

hence the σ-machine is effective due to Corollary 7.2. $\square$

Due to Theorem 7.3, if both $\widehat{\sigma}$ and σ contain 231, then $\mathrm{Sorted}(\sigma) = \mathfrak{S}(231)$ and the σ-machine is also surjective. On the other hand, if $\widehat{\sigma} \geq 231$, but σ avoids 231, then $\mathrm{Sorted}(\sigma)$ is strictly contained in $\mathfrak{S}(231)$ and the σ-machine is not surjective.

7.2 Characterization of effective patterns

In Lemma 7.1 we have proved that $\mathfrak{S}(231, \sigma) \subseteq \mathrm{Sorted}(\sigma)$. Our next goal is to provide a characterization of the effective patterns. Such patterns are precisely those where the equality holds, as stated in Corollary 7.2.

Proposition 7.4. Let σ be a permutation of length two or more. If $\widehat{\sigma} = 1 \oplus \alpha$, for some $\alpha \in \mathfrak{S}(231)$, then σ is not effective.

Proof. Let $\widehat{\sigma} = 1 \oplus \alpha$, with $\alpha \in \mathfrak{S}(231)$. We show that there is a σ-sortable permutation π such that $\mathcal{S}^\sigma(\pi) \geq \sigma$. Let $\sigma = \sigma_1\sigma_2 \cdots \sigma_k$, where $\sigma_2 = 1$ since the first element of $\widehat{\sigma}$ is 1 by hypothesis. Suppose that $\sigma_1 = t$, for some $2 \leq t \leq k$. Define

$$\pi = \mathcal{R}(\sigma) \ominus \mathcal{R}(\sigma_2 \cdots \sigma_t) = \sigma'_k\sigma'_{k-1} \cdots \sigma'_2\sigma'_1\sigma_t\sigma_{t-1} \cdots \sigma_2,$$

where $\sigma'_i = \sigma_i + t - 1$ for each i. Notice that:

$$\sigma'_2 = \sigma_2 + t - 1 = 1 + t - 1 = t = \sigma_1.$$

We shall prove that:

$$\mathcal{S}^\sigma(\pi) = \sigma_1\sigma_2 \cdots \sigma_t\sigma'_1\sigma'_3 \cdots \sigma'_k.$$

Due to our assumptions, we have that $\sigma = t1\sigma_3 \cdots \sigma_k$ and σ avoids 231. Thus it must be

$$\{\sigma_3, \ldots, \sigma_t\} = \{2, \ldots, t-1\} \text{ and } \{\sigma_{t+1}, \ldots, \sigma_k\} = \{t+1, \ldots, k\},$$

otherwise there would be two indices $i \in \{2, \ldots, t-1\}$ and $j \in \{t+1, \ldots, k\}$ such that $\sigma_i > t$ and $\sigma_j < t$. But then σ would contain an occurrence $\sigma_1\sigma_i\sigma_j$ of 231, which is a contradiction. An immediate consequence is that the string $w = \sigma_1 \cdots \sigma_t \sigma'_{t+1} \cdots \sigma'_k$ is order isomorphic to σ. Indeed w is obtained from σ by adding $t-1$ to the elements $\sigma'_{t+1}, \ldots, \sigma'_k$. Moreover, the string $z = \sigma_2 \cdots \sigma_t \sigma'_1 \sigma'_3 \cdots \sigma'_k$ avoids σ. Indeed we have $\sigma_1 = t$, whereas $\sigma_i < t$ for each $i \leq t$. Thus no element amongst $\sigma_2, \ldots, \sigma_t$ can play the role of σ_1 in an occurrence of σ in z. Finally, the remaining suffix of z is too short to contain σ (it has length $k-1$). Now, let us consider the action of the σ-stack on input π. Since π contains the prefix $\sigma'_k\sigma'_{k-1} \cdots \sigma'_2\sigma'_1$, the first element that cannot be pushed into the σ-stack is σ'_1, which causes the pop of σ'_2. Then σ'_1 is pushed and, immediately after, the σ-stack contains $\sigma'_k \cdots \sigma'_3\sigma'_1$, reading from bottom to top. The remaining elements of the input are $\sigma_t \cdots \sigma_2$. Notice that $\sigma'_k \cdots \sigma'_3\sigma'_1\sigma_t \cdots \sigma_2 = \mathcal{R}(z)$, which avoids $\mathcal{R}(\sigma)$ due to what observed above. Therefore all the remaining elements of the input are pushed into the σ-stack directly and the output is:

$$\mathcal{S}^\sigma(\pi) = \sigma'_2\sigma_2 \cdots \sigma_t\sigma'_1\sigma'_3 \cdots \sigma'_k.$$

This is precisely what we wanted, since $\sigma'_2 = \sigma_1$. Now, $\mathcal{S}^\sigma(\pi)$ contains the substring w, which is an occurrence of σ. Finally, it is easy to show that $\mathcal{S}^\sigma(\pi)$ avoids 231, since σ avoids 231 and $\sigma'_i > \sigma_j$, for each i, j. Thus π is σ-sortable and $\mathcal{S}^\sigma(\pi) \geq \sigma$. This completes the proof. $\square$

Next we show that if $\widehat{\sigma}$ is not the direct sum of 1 plus a 231-avoiding permutation, then σ is effective. If $\widehat{\sigma} \geq 231$, then the desired results follows immediately by Theorem 7.3. We just need to address the remaining case where $\widehat{\sigma}$ avoids 231 and $\sigma_2 \neq 1$.

Proposition 7.5. Let σ be a permutation of length two or more. If $\widehat{\sigma}$ avoids 231 and $\sigma_2 \neq 1$, then σ is effective.

Proof. We show that $\mathrm{Sorted}(\sigma) \subseteq \mathfrak{S}(\sigma)$. Let $\gamma \in \mathrm{Sorted}(\sigma)$ and suppose, for a contradiction, that $\gamma \geq \sigma$. Let π be a σ-sortable permutation such that $\mathcal{S}^\sigma(\pi) = \gamma$. If π avoids $\mathcal{R}(\sigma)$, then $\gamma = \mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$ avoids σ, which is a contradiction. Therefore we can assume that $\pi \geq \mathcal{R}(\sigma)$. By Lemma 3.4, we have that $\gamma \geq \widehat{\sigma}$. Thus γ contains both σ and $\widehat{\sigma}$ and moreover γ avoids 231. Since $\sigma_2 \neq 1$, it must be $\sigma_1 = 1$. Otherwise, if $\sigma_i = 1$, with $i \geq 3$, then it would be either $\sigma_1\sigma_2\sigma_i \simeq 231$, if $\sigma_1 < \sigma_2$, or $\sigma_2\sigma_1\sigma_i \simeq 231$, if $\sigma_1 > \sigma_2$. In the first case, it would be $\gamma \geq \sigma \geq 231$, which is impossible. In the second case, it would be $\gamma \geq \widehat{\sigma} \geq 231$, again a contradiction.

Now, let $\tilde{\sigma}_1 \cdots \tilde{\sigma}_k$ be the leftmost occurrence of σ in γ. Let us consider the instant when $\tilde{\sigma}_1$ is extracted from the σ-stack. If the input is empty, then the σ-stack must contain all the elements $\tilde{\sigma}_1 \cdots \tilde{\sigma}_k$, from top to bottom, which is impossible by definition of σ-stack. Thus $\tilde{\sigma}_1$ is extracted due to the fact that the next element of the input, say σ'_1, triggers the restriction of the σ-stack. More explicitly, σ'_1 realizes an occurrence of σ together with some elements $\sigma'_2 \cdots \sigma'_k$ contained in the σ-stack (from top to bottom). Since $\sigma_1 = 1$, it must be $\tilde{\sigma}_1 > \sigma'_1$, otherwise $\tilde{\sigma}_1 \sigma'_2 \cdots \sigma'_k$ would be an occurrence of σ inside the σ-stack. If $\tilde{\sigma}_2$ precedes σ'_1 in γ, then $\tilde{\sigma}_1 \tilde{\sigma}_2 \sigma'_1 \simeq 231$ in γ, which is impossible. Thus we can assume that $\tilde{\sigma}_2$ follows σ'_1 in γ. We consider two cases, according to whether or not $\tilde{\sigma}_2$ follows σ'_1 in π.

- Suppose initially that $\tilde{\sigma}_2$ precedes σ'_1 in π, and thus $\tilde{\sigma}_2$ is contained in the σ-stack when σ'_1 is the next element of the input. Consider the instant when $\tilde{\sigma}_1$ is extracted from the σ-stack. Recall that at this point σ'_1 is the next element of the input and $\sigma'_1 \sigma'_2 \cdots \sigma'_k \simeq \sigma$, for some elements $\sigma'_2 \cdots \sigma'_k$ contained in the σ-stack. Moreover, the top element of the σ-stack is $\tilde{\sigma}_1$ and $\tilde{\sigma}_2$ is still contained in the σ-stack. We shall prove that $\tilde{\sigma}_j$ is contained in the σ-stack for each $j \geq 3$. Suppose, for a contradiction, that $\tilde{\sigma}_j$ follows σ'_1 in the input, for some $j \geq 3$. Notice that $\tilde{\sigma}_2$ is extracted from the σ-stack before $\tilde{\sigma}_j$ enters. Let $\sigma'''_1 \sigma'''_2 \cdots \sigma'''_k$ be the occurrence of σ that causes the pop of $\tilde{\sigma}_2$, with σ'''_1 the next element of the input. Again we have $\tilde{\sigma}_2 > \sigma'''_1$, since $\sigma_1 = 1$. Moreover γ contains $\tilde{\sigma}_1 \tilde{\sigma}_2 \sigma'''_2$, thus it must be $\sigma'''_1 > \tilde{\sigma}_2$, or else $\tilde{\sigma}_1 \tilde{\sigma}_2 \sigma'''_1$ would be an occurrence of 231 in γ, which is impossible. But then, when $\tilde{\sigma}_1$ is the top of the σ-stack, the σ-stack contains $\tilde{\sigma}_1 \sigma'''_2 \cdots \sigma'''_k$, which is an occurrence of σ since $\sigma'''_1 > \tilde{\sigma}_1$ (and $\sigma_1 = 1$), again a contradiction. We can thus assume that, when $\tilde{\sigma}_1$ is extracted, $\tilde{\sigma}_j$ is contained in the σ-stack for each $j \geq 3$. But this is impossible, since the σ-stack would contain an occurrence $\tilde{\sigma}_1 \tilde{\sigma}_2 \cdots \tilde{\sigma}_k$ of σ.

- Suppose instead that $\tilde{\sigma}_2$ follows σ'_1 in π, and thus σ'_1 is extracted from the σ-stack before $\tilde{\sigma}_2$ enters. Let us consider the instant when σ'_1 is extracted. At this point, the σ-stack contains some elements $\sigma''_2 \cdots \sigma''_k$ that realize an occurrence of σ together with the next element σ''_1 of the input. Again it must be $\sigma'_1 > \sigma''_1$, otherwise (since $\sigma_1 = 1$) the σ-stack would contain an occurrence $\sigma'_1 \sigma''_2 \cdots \sigma''_k$ of σ. Thus $\tilde{\sigma}_1 > \sigma'_1 > \sigma''_2$ and $\tilde{\sigma}_1 \tilde{\sigma}_2 \sigma''_1 \simeq 231$. This means that $\tilde{\sigma}_2$ must follow σ''_1 in γ. Now, if $\tilde{\sigma}_2$ precedes σ''_1 in π, then we are back to the previous case. Otherwise, we can repeat the same argument on $\tilde{\sigma}_1 \tilde{\sigma}_2 \sigma''_1$, with σ''_1 in place of σ'_1. Sooner or later, since σ''_1 is strictly to the right of σ'_1 in π, this will result in a contradiction.

□

What proved so far in this section, together with Lemma 7.1, leads to the following characterization of effective patterns.

σ	Sorted(σ)	**Sequence** $\{\lvert\mathrm{Sorted}_n(\sigma)\rvert\}_n$	**OEIS**
123	$\mathfrak{S}(123, 231)$	1, 2, 4, 7, 11, 16, 22, 29, 37,...	A000124
132	$\mathfrak{S}(132, 231)$	1, 2, 4, 8, 16, 32, 64, 128, 256,...	A000079
213	?	1, 2, 4, 9, 22, 58, 161, 466, 1390	
231	$\mathfrak{S}(231)$	1, 2, 5, 14, 42, 132, 429, 1430, 4862,...	A000108
312	?	1, 2, 4, 8, 17, 40, 104, 291, 855	
321	$\mathfrak{S}(231, 321)$	1, 2, 4, 8, 16, 32, 64, 128, 256,...	A000079

Table 7.1: σ-sorted permutations for patterns σ of length three, starting from σ-sorted permutations of length one.

Corollary 7.6. Let σ be a permutation of length two or more. Then σ is not effective if and only if $\widehat{\sigma} = 1 \oplus \alpha$, for some $\alpha \in \mathfrak{S}(231)$.

An immediate consequence of Corollary 7.6 is that there are $\mathfrak{c}_{n-1} = |\mathfrak{S}_{n-1}(231)|$ patterns of length n that are not effective. For instance, there is one such pattern of length two, namely 21, since $\widehat{21} = 12 = 1 \oplus 1$. Similarly, there are two such patterns of length three, namely 213 and 312 (see Table 7.1).

7.3 Fertility and sorted permutations of the 123-machine

Fertility and sorted permutations for the 123-machine can be determined from the results proved in Chapter 4. Recall that any $\pi \in \mathrm{Sort}_n(123)$ which is not the identity permutation can be uniquely constructed as follows:

- choose $\alpha \in \mathfrak{S}_k(213)$, with $\alpha_1 = k \geq 2$;
- add h new maxima $k+1, \dots, k+h$, one at a time, using the bijection φ of Theorem 4.7;
- add $t = n - k - h$ consecutive ascents at the beginning, by inflating the first element of the permutation, according to Corollary 4.3.

We wish to exploit the above construction to describe the set Sorted(123) and compute the fertility of 123-sorted permutations. Let π be a 123-sortable permutation. If π starts with a descent $\pi_1 > \pi_2$, with $\pi_1 = k$, then by Lemma 4.4, we have $\mathcal{S}^{123}(\pi) = n(n-1)\cdots(k+1)(k-1)\cdots 21k$. Moreover, observe that inserting t consecutive ascents $\pi_1(\pi_1+1)\cdots(\pi_1+t)$ at the beginning does not affect the behavior of the 123-stack. Indeed the elements $\pi_1(\pi_1+1)\cdots(\pi_1+t)$ act as a single element at the bottom of the 123-stack during the sorting process. Therefore, if π' is obtained from π by t-inflating $\pi_1 = k$, then we have:

$$\mathcal{S}^{123}(\pi') = \underbrace{(k+t+h)\cdots(k+t+1)}_{(I)}\underbrace{k-1\cdots 21}_{(II)}\underbrace{(k+t)\cdots(k+1)k}_{(III)},$$

where the segment (I) corresponds to the h new maxima added using φ, (II) corresponds to the elements of $\alpha \in \mathfrak{S}_k(213)$ and (III) contains the t-inflation of π_1. The fertility of $\mathcal{S}^{123}(\pi')$ is then $\mathfrak{c}_{k-1}$, since there are $\mathfrak{c}_{k-1}$ permutations α in $\mathfrak{S}_k(213)$ whose first element is equal to k.

Corollary 7.7. We have:

$$\mathrm{Sorted}(123)_n = \{\mathsf{id}_h^{-1} \ominus \left(\mathsf{id}_k^{-1} \oplus \mathsf{id}_t^{-1}\right) : k \geq 2,\ h,t \geq 0,\ k+h+t=n\} \dot\cup \{\mathsf{id}_n^{-1}\}.$$

Moreover, the fertility of $\mathsf{id}_h^{-1} \ominus (\mathsf{id}_k^{-1} \oplus \mathsf{id}_t^{-1})$ is equal to $\mathfrak{c}_{k-1}$.

Proof. If $\pi = \mathsf{id}_n$, then $\mathcal{S}^{123}(\pi) = \mathsf{id}_n^{-1}$. The rest follows from what discussed before. $\square$

From the description obtained in Corollary 7.7, and in accordance with Corollary 7.6, it is easy to deduce that $\mathrm{Sorted}(123) = \mathfrak{S}(231, 123)$. Corollary 7.7 can be used to obtain an alternative proof of the enumeration of Sort(123):

$$\begin{aligned} \mathrm{Sort}(123) = \sum_{\gamma \in \mathrm{Sorted}(123)} \mathrm{fert}^{123}(\gamma) &= \\ 1 + \sum_{k\geq 2}\ \sum_{h,t\geq 0, k+h+t=n} \mathfrak{c}_{k-1} &= \\ 1 + \sum_{k=2}^{n} (n-k+1)\mathfrak{c}_{k-1} &= \\ 1 + \sum_{k=1}^{n-1} (n-k)\mathfrak{c}_k&, \end{aligned}$$

which is the same as what we got in Theorem 4.9.

Chapter 8

Sorting words of various types

In this chapter we extend Σ-machines to Cayley permutations, ascent sequences and modified ascent sequences, which has been defined in Section 1.3. Pattern-avoiding machines are built upon the notion of pattern, which is inherently more general, thus it is natural to allow different sets of strings as input sequences. The idea of analyzing sorting procedures on words is not new in the literature [2, 4, 27]. For example, classical stacksort on $\mathbb{N}^*$ has been discussed in [27]. Due to the presence of sequences with repeated elements, there are two possibilities when defining a stack sorting algorithm on $\mathbb{N}^*$. One can either allow a letter to sit on a copy of itself in the stack or force a pop operation if the next element of the input is equal to the top element of the stack. In this thesis we choose the former possibility, leaving the latter for future investigation. This is equivalent to regarding a classical stack as a 21-avoiding stack (instead of as a $(11, 21)$-stack). Moreover, we relax the condition for the output to be sorted by requiring that it is weakly increasing. The following theorem was proved in [27].

Theorem 8.1. [27] Let π be a word on $\mathbb{N}$. Then π is sortable using a 21-stack if and only if π avoids 231.

Patterns live[1] in the set $\mathfrak{Cay}$ of Cayley permutations, thus it seems appropriate to start our analysis by studying σ-machines where both input sequences and the forbidden pattern that defines the constraint of the stack are elements of $\mathfrak{Cay}$. We then consider σ-machines on ascent sequences and modified ascent sequences. Following a principle of uniformity, we always require forbidden patterns and input sequences to be chosen in the same set. Notice that the output of the σ-stack in all these cases is a word on $\mathbb{N}$, therefore we can use Theorem 8.1 to determine whether an input sequence is sortable. In Chapter 3, we characterized those patterns σ such that the set of σ-sortable permutations is a class. The main goal of this chapter

[1] As standardized sequences.

is to prove analagous results for the sets $\mathfrak{Cay}$, $\mathfrak{A}$ and $\mathfrak{MA}$. In Section 8.1.1 we study the operator $\mathcal{S}^\sigma$ on Cayley permutations. Some enumerative and structural properties of ascent and modified ascent sequences are derived in Section 8.2.1 and Section 8.2.2.

8.1 The σ-machine on Cayley permutations

In this section we consider σ-machines on Cayley permutations. Some of the results contained here can be found in [20]. We give a formal definition of these devices in the case of Cayley permutations. The corresponding machines, on $\mathfrak{A}$ and $\mathfrak{MA}$, are defined analogously. Let σ be a Cayley permutation of length at least two. A *σ-stack* is a stack that is not allowed to contain an occurrence of the pattern σ when reading the elements from top to bottom. The term *σ-machine* refers to performing a right-greedy algorithm on two stacks in series: a σ-stack, followed by a 21-avoiding stack. A Cayley permutation π is *σ-sortable* if the output of the σ-machine on input π is weakly increasing. All the definitions and notations regarding σ-machines on Cayley permutations are inherited from the classical case. If necessary, we add an apex to avoid confusion: for instance, we denote by $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ the set of σ-sortable Cayley permutations. Note that, being $\mathcal{S}^\sigma(\pi)$ the input of the 21-stack, Theorem 8.1 guarantees that $\pi \in \mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ if and only if $\mathcal{S}^\sigma(\pi)$ avoids 231. This fact will be used repeatedly for the rest of this chapter. In analogy with Definition 3.1, let $\widehat{\sigma} = \sigma_2\sigma_1\sigma_3\cdots\sigma_k$ be the Cayley permutation obtained from σ by interchanging σ_1 with σ_2. Denote by $\mathcal{R} : \mathfrak{Cay} \to \mathfrak{Cay}$ the reverse operator on Cayley permutations.

Remark 8.1. For any $\sigma \in \mathfrak{Cay}$, if the input Cayley permutation π avoids $\mathcal{R}(\sigma)$, then the restriction of the σ-stack is never triggered and thus $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$. Otherwise, the leftmost occurrence of σ results necessarily in an occurrence of $\widehat{\sigma}$ in $\mathcal{S}^\sigma(\pi)$. The proof of this fact is identical to that of Lemma 3.3. An analogous result can be similarly obtained by replacing $\mathfrak{Cay}$ with either $\mathfrak{A}$ or $\mathfrak{MA}$.

The next result is the analogue of Theorem 3.6 on Cayley permutations. The proof is identical, with Remark 8.1 playing the role of Lemma 3.4. We report it anyway for completeness.

Theorem 8.2. Let σ be a Cayley permutation. If $\widehat{\sigma}$ contains 231, then $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma) = \mathfrak{Cay}(132, \mathcal{R}(\sigma))$. In this case, $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is a class with basis either $\{132, \mathcal{R}(\sigma)\}$, if $\mathcal{R}(\sigma)$ avoids 132, or $\{132\}$, otherwise.

Proof. We start by proving that $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma) \subseteq \mathfrak{Cay}(132, \mathcal{R}(\sigma))$. Let $\pi \in \mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$. Note that $\mathcal{S}^\sigma(\pi)$ avoids 231. Suppose, for a contradiction, that π contains $\mathcal{R}(\sigma)$. Then $\mathcal{S}^\sigma(\pi)$ contains $\widehat{\sigma}$ due to Remark 8.1 and $\widehat{\sigma}$ contains 231 by hypothesis, which is impossible. Otherwise, if π avoids $\mathcal{R}(\sigma)$, but contains 132, then $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$ due to the same remark. Moreover $\mathcal{R}(\pi)$ contains 231 by hypothesis, again a contradiction with $\pi \in \mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$. This proves that $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma) \subseteq \mathfrak{Cay}(132, \mathcal{R}(\sigma))$.

Conversely, suppose that π avoids both 132 and $\mathcal{R}(\sigma)$. Then, again by Remark 8.1, we have $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$, which avoids $\mathcal{R}(132) = 231$ by hypothesis, therefore π is σ-sortable. $\square$

Next we show that the condition of Theorem 8.2 is also necessary for $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ in order to be a class. The only exception is given by the pattern $\sigma = 12$.

Theorem 8.3. We have:

$$\mathrm{Sort}^{\mathfrak{Cay}}(12) = \mathfrak{Cay}(213).$$

Proof. Let π be a Cayley permutation. Suppose that π contains k occurrences of the minimum element 1 and write:

$$\pi = A_1 1 A_2 1 \cdots A_k 1 A_{k+1}.$$

It is easy to see that:

$$\mathcal{S}^{12}(\pi) = \mathcal{S}^{12}(A_1)\mathcal{S}^{12}(A_2)\cdots\mathcal{S}^{12}(A_k)\mathcal{S}^{12}(A_{k+1})1\cdots 1.$$

Indeed any entry equal to 1 enters the 12-stack only if the 12-stack is either empty or contains other copies of 1 only. Moreover, any entry equal to 1 cannot play the role of 2 in an occurrence of the (forbidden) pattern 12. Therefore the presence of some copies of 1 at the bottom of the 12-stack does not affect the sorting process of the block A_i, for each i.

Now, suppose that π contains an occurrence bac of 213. We prove that π is not 12-sortable by showing that $\mathcal{S}^{12}(\pi)$ contains 231. We argue by induction on the length of π. Let $\pi = A_1 1 A_2 1 \cdots A_k 1 A_{k+1}$ and $\mathcal{S}^{12}(\pi) = \mathcal{S}^{12}(A_1)\mathcal{S}^{12}(A_2)\cdots\mathcal{S}^{12}(A_k)\mathcal{S}^{12}(A_{k+1})1\cdots 1$ as above. Suppose that $b \in A_i$ and $c \in A_j$, for some $i \le j$ (note that $b, c \neq 1$). If $i = j$, then A_i contains an occurrence bac of 213. Thus $\mathcal{S}^{12}(A_i)$ contains 231 by the inductive hypothesis[2], as wanted. Otherwise, let $i < j$. Then $b \in \mathcal{S}^{12}(A_i)$, $c \in \mathcal{S}^{12}(A_j)$ and the elements b and c, together with any copy of 1, realize an occurrence of 231 in $\mathcal{S}^{12}(\pi)$, as desired.

Conversely, suppose that $\pi = \pi_1 \cdots \pi_n$ is not 12-sortable, i.e. $\mathcal{S}^{12}(\pi)$ contains 231. We prove that π contains 213. Let bca be an occurrence of 231 in $\mathcal{S}^{12}(\pi)$. Observe that b must precede c in π, since a non-inversion in the output necessarily comes from a non-inversion in the input, being the stack 12-avoiding. However,

[2]Formally we apply the inductive hypothesis to $\mathrm{std}(A_i)$, since not necessarily A_i is a Cayley permutation.

b is extracted before c enters. Let x be the next element of the input when b is extracted. Since the stack is 12-avoiding, then the top b is greater than or equal to any other element contained in the 12-stack. Thus $x < b$ and also $x \neq c$, since $c > b$. Finally, the triple bxc is an occurrence of 213 in π, as desired. $\square$

Theorem 8.4. Let σ be a Cayley permutation and suppose that $\sigma \neq 12$. If $\widehat{\sigma}$ avoids 231, then $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is not a class.

Proof. Let $\sigma = \sigma_1 \cdots \sigma_k$, with $k \geq 2$. We show that there are two Cayley permutations α and β such that $\alpha \leq \beta$, β is σ-sortable and α is not σ-sortable. Table 8.1 shows an example of such permutations for patterns σ of length two and for $\sigma = 231$. Now, suppose that σ has length at least three and $\sigma \neq 231$. Then the Cayley permutation $\alpha = 132$ is not σ-sortable. Indeed, $\mathcal{S}^\sigma(\alpha) = \mathcal{R}(\alpha) = 231$, since α avoids $\mathcal{R}(\sigma)$. Define β according to the following case by case analysis:

- Suppose that σ_1 is the strict minimum of σ, that is $\sigma_1 = 1$ and $\sigma_i \geq 2$ for each $i \geq 2$. Define:
$$\beta = \sigma'_k \cdots \sigma'_3 1 \sigma'_2 \sigma'_1,$$
where $\sigma'_i = \sigma_i + 1$ for each i. Note that $\beta \in \mathfrak{Cay}$ and $1\sigma'_2\sigma'_1$ is an occurrence of 132 in β. We prove that β is σ-sortable by showing that $\mathcal{S}^\sigma(\beta)$ avoids 231. The action of the σ-stack on input β is depicted in Figure 8.1. The first $k-1$ elements of β are pushed into the σ-stack, since σ has length k. Then the σ-stack contains $1\sigma'_3 \cdots \sigma'_k$, reading from top to bottom, and the next element of the input is σ'_2. Note that $\sigma'_2 > 1$, whereas $\sigma_1 < \sigma_2$, therefore $\sigma'_2 1 \sigma'_3 \cdots \sigma'_k$ is not an occurrence of σ and so σ'_2 is pushed. The next element of the input is now σ'_1. Here $\sigma'_1\sigma'_2\sigma'_3 \cdots \sigma'_k$ is an occurrence of σ, thus σ'_2 is extracted before σ'_1 enters. After this pop operation, the σ-stack contains $1\sigma'_3 \cdots \sigma'_k$. Again we have $\sigma'_1 > 1$, whereas $\sigma_1 < \sigma_2$, therefore σ'_1 is pushed into the σ-stack. The resulting string is:
$$\mathcal{S}^\sigma(\beta) = \sigma'_2\sigma'_1 1 \sigma'_3\sigma'_4 \cdots \sigma'_k.$$
We show that $\mathcal{S}^\sigma(\beta)$ avoids 231. Note that $\sigma'_2\sigma'_1\sigma'_3\sigma'_4 \cdots \sigma'_k \simeq \widehat{\sigma}$ avoids 231 by hypothesis. Moreover, the element 1 cannot be part of an occurrence of 231, because $\sigma'_2 > \sigma'_1$ and 1 is strictly less than the other elements of β. Therefore $\mathcal{S}^\sigma(\beta)$ avoids 231, as desired.

- Next suppose that σ_1 is not the strict minimum of σ, i.e. either $\sigma_1 \neq 1$ or $\sigma_i = 1$ for some $i \geq 2$. Define
$$\beta = \sigma''_k \cdots \sigma''_2 1 \sigma''_1 2,$$

σ	**σ-sortable Cayley permutation**	**Non-σ-sortable pattern**
11	3132	132
21	35241	132
231	361425	1324

Table 8.1: The case by case analysis of Theorem 8.4

where $\sigma''_i = \sigma_i + 2$ for each i. Note that $\beta \in \mathfrak{Cay}$ and $1\sigma''_2 2$ is an occurrence of 132 in β. Consider the action of the σ-stack on input β. Again the first $k-1$ elements of β are pushed into the σ-stack. Then the σ-stack contains $\sigma''_2 \cdots \sigma''_k$, reading from top to bottom, and the next element of the input is 1. Note that $1\sigma''_2 \cdots \sigma''_k$ is not an occurrence of σ. Indeed $1 < \sigma''_i$ for each i, while σ_1 is not the strict minimum of σ by hypothesis. Therefore 1 enters the σ-stack. The next element of the input is then σ''_1, which realizes an occurrence of σ together with $\sigma''_2 \cdots \sigma''_k$. Thus 1 and σ''_2 are extracted before σ''_1 is pushed. Finally, the last element of the input is 2. Again 2 can be pushed into the σ-stack, since 2 is strictly smaller than every element in the σ-stack, whereas σ_1 is not the strict minimum of σ by hypothesis. The resulting string is:

$$\mathcal{S}^\sigma(\beta) = 1\sigma''_2 2\sigma''_1\sigma''_3 \cdots \sigma''_k.$$

Note that $\sigma''_2\sigma''_1\sigma''_3 \cdots \sigma''_k \simeq \widehat{\sigma}$ avoids 231 by hypothesis. Finally, it is easy to realize that the elements 1 and 2 cannot be part of an occurrence of 231, similarly to the previous case. This completes the proof.

□

Corollary 8.5. Let σ be a Cayley permutation of length three or more. Then $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is not a permutation class if and only if $\widehat{\sigma}$ avoids 231. Otherwise, if $\widehat{\sigma}$ contains 231, then $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is a class with basis either $\{132, \mathcal{R}(\sigma)\}$, if $\mathcal{R}(\sigma)$ avoids 132, or $\{132\}$, otherwise.

Cayley permutations avoiding any classical permutation pattern of length three are enumerated by sequence A226316 in [45]. We end this section by analyzing the 21-machine. The 11-machine will be discussed in Section 8.1.1, thus completing the analysis of σ-machines for patterns of length two. The analogue of the 21-machine on classical permutations consists in applying a right-greedy algorithm on two stacks in series, which is precisely the well known case of West 2-stacksort [50]. Recall from Theorem 3.2 that $\mathrm{Sort}(21) = \mathfrak{S}(2341, 3\bar{5}241)$. The barred pattern $3\bar{5}241$ can be represented as a mesh pattern, as shown in Figure 8.2.

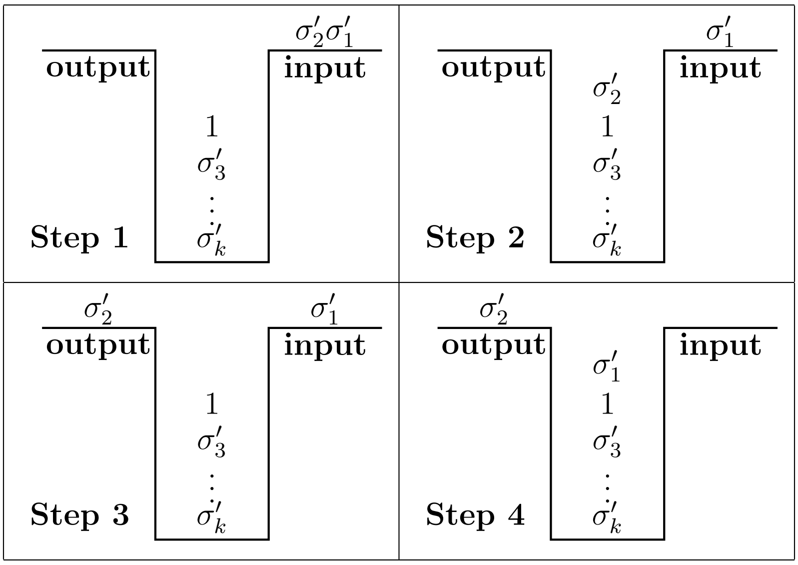

Figure 8.1: The action of the σ-stack on input β described in Theorem 8.4.

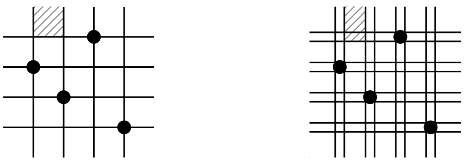

Figure 8.2: On the left, the barred pattern $3\bar{5}241$, equivalent to the mesh pattern $(3241, \{(1,4)\})$. On the right, the Cayley-mesh pattern ζ. The additional shaded box in ζ keeps into account the case of an occurrence of 3241 that is part of an occurrence of 34241.

In order to prove an analogous result for the 12-machine on Cayley permutations, we define mesh patterns on Cayley permutations (see [20]). To extend mesh patterns to strings that may contain repeated elements, we simply allow the shading of boxes that correspond to repeated elements. Instead of giving a formal definition, we refer the reader to [20] and to the example depicted in Figure 8.2. We will use the term *Cayley-mesh pattern* to denote mesh patterns on Cayley permutations. For the rest of this section, let ζ be the Cayley-mesh pattern depicted in Figure 8.2.

Lemma 8.6. Let $\pi = \pi_1 \cdots \pi_n \in \mathfrak{Cay}$. Suppose that $\pi_i < \pi_j$, for some $i < j$. Then π_i precedes π_j in $\mathcal{S}^{21}(\pi)$.

Proof. We have $\pi_j \pi_i \simeq 21$, thus π_i must be extracted from the 21-stack before π_j enters. $\square$

Theorem 8.7. We have:

$$\mathrm{Sort}^{\mathfrak{Cay}}(21) = \mathfrak{Cay}(2341, \zeta).$$

Proof. We can essentially repeat the argument used by West for classical permutations, but incorporating the additional shaded box in ζ, which corresponds to an occurrence of 3241 that is part of an occurrence of 34241. We sketch the proof below, leaving some technical details to the reader.

Suppose that π is 21-sortable. Suppose, for a contradiction, that π contains an occurrence $bcda$ of 2341 and consider the action of the 21-stack on π. By Lemma 8.6, b is extracted from the 21-stack before c enters. Similarly, c is extracted before d enters. Thus $\mathcal{S}^{21}(\pi)$ contains the occurrence bca of 231, a contradiction with π being 21-sortable. Otherwise, suppose that π contains an occurrence $cbda$ of 3241. We show that there is an element x between c and b in π such that $x \geq d$. If $x < c$ for each entry x between c and b, then b is pushed into the 21-stack before c is popped. This results in the occurrence bca of 231 in $\mathcal{S}^{21}(\pi)$, a contradiction with the fact that π is 21-sortable. Otherwise, suppose there is at least one element x between c and b in π, with $x \geq c$. If $x = c$, we can repeat the same argument with $xbda$ instead of $cbda$. If $c < x < d$, then $cxda \simeq 2341$, which is impossible due to what said in the previous case. Therefore it has to be $x \geq d$, as desired.

Conversely, suppose that π is not 12-sortable. Equivalently, let bca be an occurrence of 231 in $\mathcal{S}^{21}(\pi)$. We show that either π contains 2341 or π contains an occurrence $cbda$ of 3241 such that $x < d$ for each x between c and b in π. Observe that a follows c and b in π due to Lemma 8.6. Suppose that b comes before c in π. Note that c is extracted from the 21-stack before a enters. Let d be the next element of the input when c is extracted. Then $d > c$ and $bcda$ is an occurrence of 2341, as wanted. Otherwise, suppose that b follows c in π, and thus π contains cba. Since c is not extracted before b enters, it has to be $x \leq c$ for each x between c and b in π. Moreover, c is extracted before a enters. When c is extracted, the next element d of the input is such that $d > c$. This results in an occurrence $cbda$ of 3241 with the desired property. □

Observe that, due to the presence of the Cayley-mesh pattern ζ, the set $\mathrm{Sort}^{\mathfrak{Cay}}(21)$ is not a class. For instance, the 21-sortable Cayley permutation 34241 contains the non-sortable pattern 3241. The problem of enumerating $\mathrm{Sort}^{\mathfrak{Cay}}(21)$ remains to be solved.

Open Problem 8.1. Enumerate the set of 21-sortable Cayley permutations. The initial terms of the sequence are $1, 3, 13, 73, 483, 3547, 27939, 231395$ (not in [45]).

Recall that West 2-stack sortable permutations are precisely those classical permutations that are 21-sortable. It would thus be interesting to find a combinatorial argument for the enumeration of $\mathrm{Sort}^{\mathfrak{Cay}}(21)$ that generalizes the one for

West 2-stack sortable permutations (and allows us to recollect the classical result as a particular instance of this new, more general, approach).

8.1.1 Fully bijective σ-stacks

In this section we regard a σ-stack as a map $\mathcal{S}^\sigma : \mathfrak{Cay} \to \mathfrak{Cay}$ and investigate the properties of the resulting operator. This approach is slightly more general than the one adopted in Chapter 7, where we studied the behavior of the restriction of $\mathcal{S}^\sigma$ to the set $\mathrm{Sort}(\sigma)$ of σ-sortable permutations. Recall (from Section 7.1) that a classical permutation σ is said to be bijective if the map $\mathcal{S}^\sigma : \mathrm{Sort}(\sigma) \to \mathfrak{S}(231)$ is bijective. Analogously, given a Cayley permutation σ, we say that σ is *bijective* if $\mathcal{S}^\sigma : \mathrm{Sort}^{\mathfrak{Cay}}(\sigma) \to \mathfrak{Cay}(231)$ is bijective. A Cayley permutation σ is *fully bijective* if the map $\mathcal{S}^\sigma : \mathfrak{Cay} \to \mathfrak{Cay}$ is bijective. Notice that if σ is fully bijective, then σ is bijective and thus $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ and $\mathfrak{Cay}(231)$ are Wilf-equivalent. The main goal of this section is to provide a characterization of fully bijective patterns.

We start by discussing the pattern $\sigma = 11$. The following is a useful decomposition lemma for the 11-stack.

Lemma 8.8. Let $\pi = \pi_1 \cdots \pi_n$ be a Cayley permutation. Suppose that π contains $k+1$ occurrences $\pi_1, \pi_1^{(1)}, \dots, \pi_1^{(k)}$ of the integer π_1, for some $k \geq 0$. Write:

$$\pi = \pi_1 B_1 \pi_1^{(1)} B_2 \cdots \pi_1^{(k)} B_k.$$

Then:

$$\mathcal{S}^{11}(\pi) = \mathcal{S}^{11}(B_1)\pi_1 \mathcal{S}^{11}(B_2)\pi_1^{(1)} \cdots \mathcal{S}^{11}(B_k)\pi_1^{(k)}.$$

Proof. Consider the action of the 11-stack on input π. Since $x \neq \pi_1$ for each $x \in B_1$, the sorting process of B_1 is not affected by the presence of π_1 at the bottom of the 11-stack. Then, when the next element of the input is the second occurrence $\pi_1^{(1)}$ of π_1, the 11-stack is emptied, since $\pi_1\pi_1^{(1)}$ is an occurrence of the forbidden pattern 11. The initial elements of $\mathcal{S}^{11}(\pi)$ are thus $\mathcal{S}^{11}(B_1)\pi_1$. Finally, $\pi_1^{(1)}$ is pushed into the (empty) 11-stack and the same argument can be repeated. □

Theorem 8.9. The map $(\mathcal{R} \circ \mathcal{S}^{11})$ is an involution on $\mathfrak{Cay}$. Therefore the pattern 11 is fully bijective.

Proof. We proceed by induction on the length of input permutations. Let $\pi = \pi_1 \cdots \pi_n$ be a Cayley permutation of length n. The case $n = 1$ is trivial. If $n \geq 2$, write $\pi = \pi_1 B_1 \pi_1^{(1)} B_2 \cdots \pi_1^{(k)} B_k$ as in Lemma 8.8. Then, using the same lemma and the inductive hypothesis:

$$\begin{aligned}
&\left[\mathcal{R}\circ\mathcal{S}^{11}\right]^2(\pi) = \\
&\left[\mathcal{R}\circ\mathcal{S}^{11}\right]^2\left(\pi_1 B_1 \pi_1^{(1)} B_2 \cdots \pi_1^{(k)} B_k\right) = \\
&\left[\mathcal{R}\circ\mathcal{S}^{11}\circ\mathcal{R}\right]\left(\mathcal{S}^{11}(B_1)\pi_1\mathcal{S}^{11}(B_2)\pi_1^{(1)}\cdots\mathcal{S}^{11}(B_k)\pi_1^{(k)}\right) = \\
&\left[\mathcal{R}\circ\mathcal{S}^{11}\right]\left(\pi_1^{(k)}\mathcal{R}(\mathcal{S}^{11}(B_k))\cdots\pi_1^{(1)}\mathcal{R}(\mathcal{S}^{11}(B_2))\pi_1\mathcal{R}(\mathcal{S}^{11}(B_1))\right) = \\
&\mathcal{R}\left(\mathcal{S}^{11}(\mathcal{R}(\mathcal{S}^{11}(B_k)))\pi_1^{(k)}\cdots\mathcal{S}^{11}(\mathcal{R}(\mathcal{S}^{11}(B_2)))\pi_1^{(1)}\mathcal{S}^{11}(\mathcal{R}(\mathcal{S}^{11}(B_1)))\pi_1\right) = \\
&\pi_1\left[\mathcal{R}\circ\mathcal{S}^{11}\right]^2(B_1)\pi_1^{(1)}\left[\mathcal{R}\circ\mathcal{S}^{11}\right]^2(B_2)\cdots\pi_1^{(k)}\left[\mathcal{R}\circ\mathcal{S}^{11}\right]^2(B_k) = \\
&\pi_1 B_1\pi_1^{(1)}B_2\cdots\pi_1^{(k)}B_k = \pi
\end{aligned}$$

Therefore $(\mathcal{R}\circ\mathcal{S}^{11})^2(\pi) = \pi$, as desired. Finally, the reverse map $\mathcal{R}$ is bijective, thus $\mathcal{S}^{11}$ is a bijection on $\mathfrak{Cay}$ with inverse $\mathcal{R}\circ\mathcal{S}^{11}\circ\mathcal{R}$. □

An immediate consequence of Theorem 8.9 is that every 11-sortable Cayley permutation π is obtained from a 231-avoiding Cayley permutation by applying $\mathcal{R}\circ\mathcal{S}^{11}\circ\mathcal{R}$, that is:

$$\mathrm{Sort}^{\mathfrak{Cay}}(11) = \mathcal{R}\circ\mathcal{S}^{11}\circ\mathcal{R}(\mathfrak{Cay}(231)).$$

We wish to generalize this result by encoding the action of $\mathcal{S}^\sigma$ as a labeled Dyck path. In what follows, we always consider labeled Dyck paths where the label of each up step is equal to the label of its matching down step. This allows us to represent a labeled Dyck path as a pair $\mathcal{P} = (P, \pi)$, where P is the underlying Dyck path and π is the string obtained by reading the labels of the up steps of P from left to right. Let σ be a Cayley permutation and let π be an input permutation for the σ-stack. Define a labeled Dyck path $\mathcal{P}_\sigma(\pi)$ as follows, starting from the empty path:

- insert an up step `U` labeled a whenever an element a is pushed into the σ-stack;
- insert a down step `D` labeled a whenever an element a is extracted from the σ-stack.

Equivalently, if $P_\sigma(\pi)$ is the unlabeled Dyck path obtained by recording the push operations of the σ-stack with `U` and the pop operations with `D`, then $\mathcal{P}_\sigma(\pi) = (P_\sigma(\pi), \pi)$. It is easy to realize that $P_\sigma(\pi)$ is a Dyck path. Indeed the number of push and pop operations performed by the σ-stack on π is the same (it is equal to the length of π), therefore the number of `U` steps matches the number of `D` steps

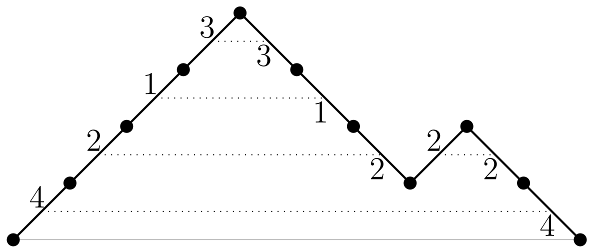

Figure 8.3: The Dyck path $UUUUDDDUDD$ which encodes $\mathcal{S}^{11}(42132)$. Dotted lines connect matching steps, which have the same label.

(and thus the path ends on the x-axis). Moreover, the path never goes below the x-axis, since this would correspond to performing a pop operation when the σ-stack is empty, which is not possible. An example of this construction, when $\sigma = 11$, is depicted in Figure 8.3. We collect several properties of $\mathcal{P}_\sigma(\pi)$ in the following lemma, whose easy proof is omitted.

Lemma 8.10. Let σ be a Cayley permutation. Let $\pi = \pi_1 \cdots \pi_n$ be a Cayley permutation of length n and let $\mathcal{P}_\sigma(\pi) = (P_\sigma(\pi), \pi)$. Then:

1. The input π is obtained by reading the labels of the up steps of $P_\sigma(\pi)$ from left to right. The output $\mathcal{S}^\sigma(\pi)$ is obtained by reading the labels of the down steps from left to right.

2. The height of $P_\sigma(\pi)$ after each up (respectively down) step is equal to the number of elements contained in the σ-stack after the corresponding push (respectively pop) operation.

3. A pop operation empties the σ-stack if and only if the corresponding `D` step of $P_\sigma(\pi)$ ends on the x-axis. Notice that the decomposition of π considered in Lemma 8.8 corresponds to the decomposition of $P_\sigma(\pi)$ obtained by considering the returns on the x-axis.

4. The labels of the down steps are uniquely determined by the labels of the up steps. Conversely, the labels of the down steps uniquely determine the labels of the up steps. More precisely, matching steps have the same label. Indeed any element pushed into the σ-stack by an up step is then popped by the matching down step.

5. Two consecutive steps of $P_\sigma(\pi)$ form a valley `DU` if and only if, denoting by a the label of `D` and b the label of `U`, b plays the role of σ_1 in an occurrence of σ that triggers the restriction of the σ-stack, while a plays the role of σ_2 in that same occurrence. Therefore the number of valleys of $P_\sigma(\pi)$ is equal to the number of elements of π that trigger the restriction of the σ-stack.

6. If $\sigma_1 = \sigma_2$, then the steps $\mathtt{D}$ and $\mathtt{U}$ in each valley $\mathtt{DU}$ have the same label.

Recall from Section 1.4 that the reverse path $\mathcal{R}(P)$ of a Dyck path P is its symmetric with respect to the vertical line $x = n$, where n is the semilength of P. The following lemma shows that if $\sigma_1 = \sigma_2$, then the path that encodes the action of the σ-stack on π is the reverse of the path that encodes the action of the σ-stack on input $\mathcal{R}(\mathcal{S}^\sigma(\pi))$.

Lemma 8.11. Let $\sigma = \sigma_1 \cdots \sigma_k$ be a Cayley permutation. Let $\pi = \pi_1 \cdots \pi_n$ be a Cayley permutation and let $\gamma = \mathcal{R}(\mathcal{S}^\sigma(\pi))$. Consider the two labeled Dyck paths $\mathcal{P}_\sigma(\pi) = (P_\sigma(\pi), \pi)$ and $\mathcal{P}_\sigma(\gamma) = (P_\sigma(\gamma), \gamma)$.

1. If $\sigma_1 = \sigma_2$, then $P_\sigma(\pi) = \mathcal{R}(P_\sigma(\gamma))$.
2. If $P_\sigma(\pi) = \mathcal{R}(P_\sigma(\gamma))$, then $(\mathcal{R} \circ \mathcal{S}^\sigma)^2(\pi) = \pi$.

Proof. 1. Suppose that $\sigma_1 = \sigma_2$. We proceed by induction on the number of valleys of $P_\sigma(\pi)$. If $P_\sigma(\pi)$ has zero valleys, then π avoids $\mathcal{R}(\sigma)$ by point 5. of Lemma 8.10. Therefore $\mathcal{S}^\sigma(\pi) = \mathcal{R}(\pi)$ and $\gamma = \mathcal{R}^2(\pi) = \pi$. Since $P_\sigma(\pi) = \mathtt{U}^n\mathtt{D}^n$ is a pyramid, the thesis follows immediately since each pyramid is equal to its reverse.

Now suppose that $P_\sigma(\pi)$ has at least one valley. Let $P_\sigma(\pi) = p_1 \cdots p_{2n}$ and write $P_\sigma(\pi) = \mathtt{U}^i\mathtt{U}^j\mathtt{D}^j\mathtt{U}^l\mathtt{D}Q$, where p_{i+2j}, p_{i+2j+1} is the leftmost valley and $Q = p_{i+2j+l+2} \cdots p_n$ is the remaining suffix of $P_\sigma(\pi)$ (see Figure 8.4). Observe that the label of both p_{i+2j} and p_{i+2j+1} is equal to π_{i+1} as a consequence of items 4., 5. and 6. of Lemma 8.10. Item 5. also implies that p_{i+2j+1} plays the role of σ_1 in an occurrence of σ that triggers the restriction of the σ-stack. More precisely, as soon as π_{i+j} is pushed (i.e. after the up step p_{i+j} in $P_\sigma(\pi)$), π_{i+j+1} is the next element of the input. Since the next segment of the path is $\mathtt{D}^j$, j pop operations are performed before pushing π_{i+j+1}. This means that the element π_{i+1}, corresponding to the last down step, plays the role of σ_2 in an occurrence of σ, while π_{i+j+1} plays the role of σ_1. Moreover, there are $k-2$ elements in the σ-stack that play the role of $\sigma_3, \ldots, \sigma_k$. Since the elements in the σ-stack correspond to the labels of the initial prefix $\mathtt{U}^i$, $\pi_1 \cdots \pi_i$ contains an occurrence of $\sigma_k \cdots \sigma_3$ (claim I). Then, after that j pop operations are performed, the σ-stack contains $\pi_i \cdots \pi_1$, reading from top to bottom, and the elements $\pi_{i+j+1}, \pi_{i+j+2}, \ldots, \pi_{i+j+l}$ are pushed (claim II). Now, write:

$$\pi = \underbrace{\pi_1 \cdots \pi_i}_{A} \; \underbrace{\pi_i + 1 \cdots \pi_{i+j}}_{B} \; \underbrace{\pi_{i+j+1} \cdots \pi_{i+j+l}}_{C} \; \underbrace{\pi_{i+j+l+1} \cdots \pi_n}_{D},$$

where the elements of A correspond to the initial prefix $\mathtt{U}^i$ of $P_\sigma(\pi)$, B corresponds to $\mathtt{U}^j$, C to $\mathtt{U}^l$ and D to the remaining up steps. Consider the string:

$$ACD = \pi_1 \cdots \pi_i \pi_{i+j+1} \cdots \pi_n,$$

obtained by removing the segment $B = \pi_{i+1} \cdots \pi_{i+j}$ from π. Let $\tilde{\pi} = \mathrm{std}(ACD)$ be the standardization of ACD. Note that $\mathcal{P}_\sigma(\tilde{\pi})$ is obtained from $\mathcal{P}_\sigma(\pi)$ by cutting out the pyramid $\mathtt{U}^j\mathtt{D}^j$, which corresponds to the removed segment B. Indeed the elements contained in the σ-stack when π_i enters are exactly the same as the elements contained in the σ-stack when π_{i+j+1} is pushed, thus we can safely cut out the pyramid $\mathtt{U}^j\mathtt{D}^j$ without affecting the sorting procedure. Therefore:

$$\mathcal{S}^\sigma(\pi) = \mathcal{R}(B)\mathcal{S}^\sigma(\tilde{\pi})$$

and

$$\gamma = \mathcal{R}(\mathcal{S}^\sigma(\pi)) = \mathcal{R}(\mathcal{S}^\sigma(\tilde{\pi}))B.$$

Now, since $P_\sigma(\tilde{\pi})$ has one valley less than $P_\sigma(\pi)$, by the inductive hypothesis we have $P_\sigma(\tilde{\pi}) = \mathcal{R}(P_\sigma(\tilde{\gamma}))$, where $\tilde{\gamma} = \mathcal{R}(\mathcal{S}^\sigma(\tilde{\pi}))$. The only difference bewteen $P_\sigma(\pi)$ and $P_\sigma(\tilde{\pi})$ is the removed pyramid $\mathtt{U}^j\mathtt{D}^j$. Therefore, if we show that $P_\sigma(\gamma)$ is obtained from $P_\sigma(\tilde{\gamma})$ by reinserting the pyramid $\mathtt{U}^j\mathtt{D}^j$ in the same place, the thesis follows. We have $\gamma = \mathcal{R}(\mathcal{S}^\sigma(\tilde{\pi}))B$ and $\tilde{\gamma} = \mathcal{R}(\mathcal{S}^\sigma(\tilde{\pi}))$. Consider the last push operation performed by the σ-stack when processing $\tilde{\gamma}$, which corresponds to the last up step of $\mathcal{P}_\sigma(\tilde{\gamma})$. Note that, since $P_\sigma(\tilde{\pi}) = \mathcal{R}(P_\sigma(\tilde{\gamma})$, this is also the first down step of $P_\sigma(\tilde{\pi})$, and thus the first pop operation performed when processing $\tilde{\pi}$. Therefore the elements contained in the σ-stack when the last push operation is performed, while processing $\tilde{\gamma}$, are $\pi_{i+j+l} \cdots \pi_{i+j+1}\pi_i \cdots \pi_1$, reading from top to bottom. If we sort γ instead of $\tilde{\gamma}$, we have to process the additional segment B. Now, the first element of B is π_{i+1}. As a consequence of claim I, π_{i+1} realizes an occurrence of σ together with π_{i+j+1} (which plays the role of σ_2) and other $k-2$ elements in $\pi_1 \cdots \pi_i$. The only difference is that, contrary to what happened when sorting π, the role of π_{i+1} and π_{i+j+1} are interchanged: here the hypothesis $\sigma_1 = \sigma_2$ is relevant. As a result, before pushing the first element π_{i+1} of B, we have to pop each element of the σ-stack up to π_{i+j+1}, π_{i+j+1} included. Then, the σ-stack contains $\pi_i \cdots \pi_1$, reading from top to bottom. Therefore we can push $\pi_{i+1} = \pi_{i+j+1}$ and the remaining elements of B, this time because of claim II. This means that $P_\sigma(\gamma)$ is obtained by inserting a pyramid $\mathtt{U}^j\mathtt{D}^j$ immediately before the last i down steps of $P_\sigma(\tilde{\gamma})$, as desired.

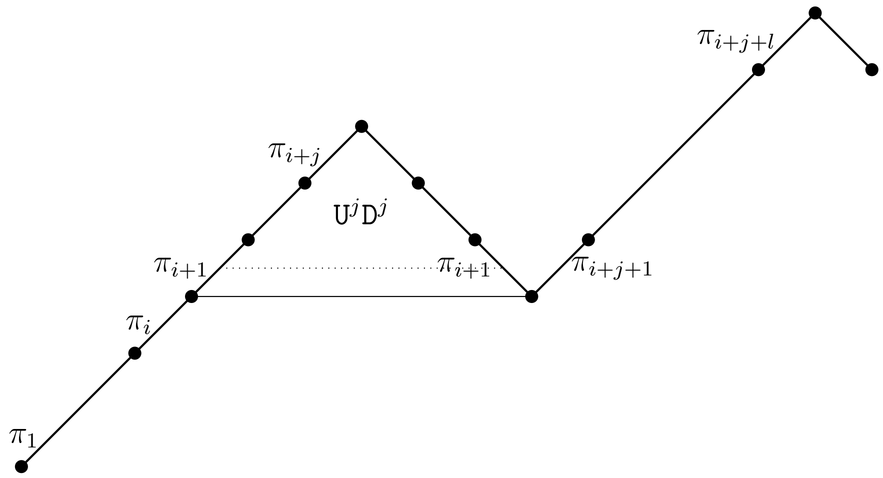

Figure 8.4: The (prefix of the) path $P_\sigma(\pi)$ mentioned in the proof of Lemma 8.11.

2. By hypothesis, $P_\sigma(\gamma) = \mathcal{R}(P_\sigma(\pi))$, therefore the word w obtained by reading the labels of the down steps of $P_\sigma(\gamma)$ (from left to right) is $w = \mathcal{R}(\pi)$. By definition of $\mathcal{P}_\sigma(\gamma)$, we also have $w = \mathcal{S}^\sigma(\gamma)$. Therefore:

$$\mathcal{R}(\pi) = \mathcal{S}^\sigma(\gamma) = \mathcal{S}^\sigma(\mathcal{R}(\mathcal{S}^\sigma(\pi)))$$

and the thesis follows by applying the reverse operator to both sides of the equality.

□

As a consequence of what proved so far in this section, we obtain the desired characterization of fully bijective Cayley permutations.

Theorem 8.12. Let $\sigma = \sigma_1 \cdots \sigma_k \in \mathfrak{Cay}$. Then σ is fully bijective if and only if $\sigma_1 = \sigma_2$.

Proof. Suppose that $\sigma_1 \neq \sigma_2$. Then $\widehat{\sigma} \neq \sigma$ and so $\mathcal{R}(\sigma) \neq \mathcal{R}(\widehat{\sigma})$. Finally, it is easy to realize that:

$$\mathcal{S}^\sigma(\mathcal{R}(\sigma)) = \widehat{\sigma} = \mathcal{S}^\sigma((\mathcal{R}(\widehat{\sigma}))),$$

therefore $\mathcal{S}^\sigma$ is not injective on $\mathfrak{Cay}$.

Conversely, suppose that $\sigma_1 = \sigma_2$. By Lemma 8.11, we have that $(\mathcal{R} \circ \mathcal{S}^\sigma)^2$ is the identity on $\mathfrak{Cay}$, therefore $\mathcal{R} \circ \mathcal{S}^\sigma$ is bijective. Finally, since the reverse map $\mathcal{R}$ is bijective, $\mathcal{S}^\sigma$ is a bijection too, as desired. □

Remark 8.2. As we pointed out in the proof of Lemma 8.11, the hypothesis $\sigma_1 = \sigma_2$ guarantees that the Dyck path obtained when sorting π is equal to the reverse of the path obtained when sorting $\mathcal{R}(\mathcal{S}^\sigma(\pi))$. This is sufficient for the operator $\mathcal{S}^\sigma$ in order to be bijective on $\mathfrak{Cay}$. More precisely, the crucial property is that the roles of σ_1 and σ_2 are interchanged in the paths associated to π and $\mathcal{R}(\mathcal{S}^\sigma(\pi))$: this is precisely where the hypothesis $\sigma_1 = \sigma_2$ plays a role. Due to this reason, an analogous argument can be repeated on stacks avoiding the pair of patterns $(\sigma, \widehat{\sigma})$. Indeed every time an occurrence of σ triggers the $(\sigma, \widehat{\sigma})$-stack on input π, an occurrence of $\widehat{\sigma}$, where the roles of σ_1 and σ_2 are interchanged, triggers the $(\sigma, \widehat{\sigma})$-stack on input $\mathcal{R}(\mathcal{S}^{\sigma_1,\sigma_2}(\pi))$. The converse statement is true as well. Therefore the operator $\mathcal{S}^{\sigma,\widehat{\sigma}}$ is bijective on $\mathfrak{Cay}$. Similarly, since classical permutations are mapped into classical permutations by any operator $\mathcal{S}^\Sigma$, the operator $\mathcal{S}^{\sigma,\widehat{\sigma}}$ associated to the classical $(\sigma, \widehat{\sigma})$-machine is bijective on $\mathfrak{S}$. This result was generalized by Berlow in [11]: the map $\mathcal{S}^\Sigma$ is bijective if and only if for every $\sigma \in \Sigma$ we have $\widehat{\sigma} \in \Sigma$ as well.

Remark 8.3. The encoding of the action of a σ-stack as Dyck paths could theoretically lead to the enumeration of $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$. Indeed, due to Lemma 8.10, the number of σ-sortable Cayley permutations of length n is equal to the number of labeled Dyck paths of semilength n such that:

- Reading the labels of the down steps from left to right yields a 231-avoiding Cayley permutations.
- Each valley of the path corresponds to an element of the input permutation that triggers the restriction of the σ-stack, as described in item 5. of Lemma 8.10.

A natural question would thus be the following. Given a Dyck path P, are there any parameters of P that allow us to describe (and enumerate) the set of Cayley permutations that can be used to suitably label P? In other words, can we describe those Cayley permutations where the action of the σ-stack is encoded by the same Dyck path P?

8.2 The σ-machine on ascent sequences

We spend the last section of this chapter by discussing σ-machines on classical ascent sequences $\mathfrak{A}$ and modified ascent sequences $\mathfrak{MA}$ (see Section 1.3). Recall that ascent sequences are a bijective encoding of Fishburn permutations $\mathfrak{F} = \mathfrak{S}(\mathfrak{f})$, where $\mathfrak{f}$ is the bivincular pattern $\mathfrak{f} = (231, \{1\}, \{1\})$. The maps that link the sets $\mathfrak{A}$, $\mathfrak{MA}$ and $\mathfrak{F}$, as well as the pattern $\mathfrak{f}$, are depicted again (for convenience) in Figure 8.5.

Recently, Claesson and the current author [22] initiated the development of a theory of transport of patterns between Fishburn permutations and ascent sequences. One of their goals is to achieve a better understanding of the notion of

$\mathfrak{f} =$ 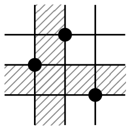

$$\begin{tikzcd}
\mathfrak{A} \arrow[rr, two heads, "\phi"] \arrow[d, "\omega"'] & & \mathfrak{F} \\
\mathfrak{MA} \arrow[urr, "\phi'"'] & &
\end{tikzcd}$$

Figure 8.5: The pattern $\mathfrak{f}$, on the left. How the bijections ϕ, ω and ϕ' are related, on the right.

pattern involvement on the sets $\mathfrak{A}$ and $\mathfrak{MA}$, which proved to be rather challenging. The transport relies on a high-level generalization of ϕ', which they call the *Burge transpose*. One of the main results of [22] is an explicit construction of a set of patterns $\mathcal{B}(p)$ such that ϕ' maps the set $\mathfrak{MA}(\mathcal{B}(p))$ of modified ascent sequences avoiding every pattern in $\mathcal{B}(p)$ to the set $\mathfrak{F}(p)$ of Fishburn permutations avoiding p. The set $\mathcal{B}(p)$ is called the *Fishburn basis* of p. We refer the reader to [22] for the definition of Burge transpose and for a detailed construction of $\mathcal{B}(p)$.

Theorem 8.13 ([22], Transport of patterns from $\mathfrak{F}$ to $\mathfrak{MA}$)**.** For any permutation p, we have:

$$\mathfrak{F}(p) = \phi'\big(\mathfrak{MA}(\mathcal{B}(p))\big).$$

Therefore $\mathfrak{F}(p)$ and $\mathfrak{MA}(\mathcal{B}(p))$ are Wilf-equivalent subsets of $\mathfrak{Cay}$.

The current author aims to add one more piece to the general picture by analyzing the behavior of σ-machines on ascent and modified ascent sequences. It is worth noticing that, in some cases, we can use the results obtained in the previous section on Cayley permutations, as we show in the following result.

Theorem 8.14. Let $X \in \{\mathfrak{A}, \mathfrak{MA}\}$ and let $\sigma \in \mathfrak{Cay} \cap X$. If $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma) = \mathfrak{Cay}(A)$, for a set of patterns A, then $\mathrm{Sort}^X(\sigma) = X(A)$.

Proof. Let $w \in X$ and let $w' = \mathrm{std}(w)$ be the standardization of w. Notice that $w' \in \mathfrak{Cay}$. It is easy to observe that, for every pattern y, w contains y if and only if w' contains y. More precisely, any subsequence $w_{i_1} \cdots w_{i_k}$ of w is order isomorphic to the corresponding subsequence $w'_{i_1} \cdots w'_{i_k}$ of $w' = \mathrm{std}(w)$. An immediate consequence is that the action of the σ-stack on w is identical to the action of the σ-stack on w'. Therefore w is σ-sortable if and only if w' is σ-sortable, which in turn is equivalent to $w' \in \mathfrak{Cay}(A)$ by hypothesis. Finally, $w' \in \mathfrak{Cay}(A)$ if and only if $w \in X(A)$, thus the thesis follows. $\square$

Due to Corollary 8.5 and Theorem 8.14, if $\sigma \in \mathfrak{Cay} \cap \mathfrak{A}$ and $\widehat{\sigma}$ contains 231, then $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ is a class with basis $\{132, \mathcal{R}(\sigma)\}$. An analogous result holds for modified ascent sequences. On the other hand, if $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is not a class, not necessarily the same holds for $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ and $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$.

8.2.1 Classical ascent sequences

We start by recalling some useful results from the literature. The following lemma was proved by Duncan and Steingrimsson in [30].

Lemma 8.15. [30] The set $\mathfrak{A}(y)$ consists solely of RGFs if and only if $y \leq 12123$.

Lemma 8.16. Let $x \in \mathfrak{A}(123)$. Then $\max(x) \leq 2$.

Proof. It follows from Lemma 8.15 and the definition of RGF. □

Next we show that the pattern $\sigma = 11$ is an instance where $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ is a permutation class, whereas $\mathrm{Sort}^{\mathfrak{Cay}}(\sigma)$ is not.

Theorem 8.17. We have:

$$\mathrm{Sort}^{\mathfrak{A}}(11) = \mathfrak{A}(1213, 1223).$$

Proof. Let $x \in \mathfrak{A}$. Observe that the first element $x_1 = 1$ is extracted from the 11-stack if and only if the next element of the input is equal to 1, which in this case replace x_1 at the bottom of the 11-stack. Thus the last element x_{last} of $\mathcal{S}^{11}(x)$ is $x_{last} = 1$. This fact will be repeatedly used for the rest of this proof.

Suppose initially that x contains an occurrence $x_i x_j x_k x_\ell$ of 1213. Without losing generality, we can assume that $x_i x_j x_k$ is the leftmost 121 in any occurrence of 1213 where x_ℓ plays the role of 3. We wish to prove that x is not 11-sortable by showing that $\mathcal{S}^{11}(x)$ contains 231. If x_i is contained in the 11-stack when x_k is the next element of the input, then, since $x_k x_i$ is an occurrence of 11, every element up to x_i must be extracted before pushing x_k. Therefore x_j is extracted before x_k enters and $\mathcal{S}^{11}(x)$ contains the occurrence $x_j x_\ell x_{last}$ of 231, as wanted. Otherwise, suppose that x_i is extracted before x_j enters the 11-stack. Let y be the next element of the input when x_i is extracted. Consider the following two cases.

- $y x_i$ is an occurrence of 11 (and thus x_i must be extracted); in this case we can repeat the same argument replacing x_i with y until we fall in the next case.

- $y \neq x_i$ and there is an entry y' contained in the 11-stack such that yy' is an occurrence of 11 (thus the top element x_i must be extracted). If $y < x_i$, then $y' x_i y$ is an occurrence of 121 (with $x_i < x_\ell$) that precedes $x_i x_j x_k$, which is impossible due to our choice of $x_i x_j x_k$. On the other hand, suppose that $y > x_i$. Then it must be $x_i > 1$, or else $x_1 x_j x_k$ would be an occurrence of 121 that precedes $x_i x_j x_k$, which is again impossible. Finally, we get the desired occurrence $x_i y x_{last}$ of 231 in $\mathcal{S}^{11}(x)$, as desired.

Next suppose that x contains an occurrence $x_i x_j x_k x_\ell$ of 1223. Then x_j must be extracted from the 11-stack before x_k enters (since $x_j x_k$ is an occurrence of 11) and thus $\mathcal{S}^{11}(x)$ contains an occurrence $x_j x_\ell x_{last}$ of 231, which is what we wanted.

Conversely, suppose that x avoids 1213 and 1223. Suppose, for a contradiction, that x is not 11-sortable. Equivalently, let $x_i = b$, $x_j = c$ and $x_k = a$ be three elements of x that result in an occurrence $x_i x_j x_k = bca$ of 231 in $\mathcal{S}^{11}(x)$. Since x avoids 1213, then x is a RGF by Lemma 8.15. Moreover, again due to the avoidance of 1213, the prefix $x_1 \cdots x_j$ of x must be weakly increasing, that is:

$$x = 1^{t_1} 2^{t_2} \cdots a^{t_a} \cdots b^{t_b} \cdots (c-1)^{t_{c-1}} c^{t_c} x_j \cdots ,$$

for some integers $t_u \geq 1$, $u = 1, \dots, c-1$, and $t_c \geq 0$. Moreover, since x avoids 1223, it must be $t_u = 1$ for each $2 \leq u \leq c-1$. Now, it is easy to observe that as soon as $x_j = c$ enters the 11-stack, the content of the 11-stack is $(c-1) \cdots b \cdots a \cdots 321$, reading from top to bottom. Therefore, if $i > j$ then x_j is extracted before x_i enters, since $x_i = b$ forms an occurrence of 11 together with the other copy of b in the 11-stack. But this is impossible since x_i precedes x_j in $\mathcal{S}^{11}(x)$. Thus it must be $i < j$. But then x_j enters the 11-stack above x_i, which is again impossible since we supposed that x_i precedes x_j in $\mathcal{S}^{11}(x)$. $\square$

Due to Lemma 8.15, we have $\mathfrak{A}(1213, 1223) = \mathfrak{RGF}(1213, 1223)$. The enumeration of $\mathfrak{RGF}(1213, 1223)$ can be found in [36], where the authors determine all the Wilf-equivalence classes of pairs of patterns of length four. The arising sequence is A005183 in [45].

Next we analyze the 12-machine. Since $\mathrm{Sort}^{\mathfrak{Cay}}(12) = \mathfrak{Cay}(213)$, we can apply Theorem 8.14.

Theorem 8.18. We have:

$$\mathrm{Sort}^{\mathfrak{A}}(12) = \mathfrak{A}(213).$$

Notice that $\mathfrak{A}(213) = \mathfrak{A}(1213)$ due to Lemma 5.12 and Lemma 8.15. The enumeration of $\mathfrak{A}(213)$ can be found in [30].

Theorem 8.19. We have:

$$\mathrm{Sort}^{\mathfrak{A}}(121) = \mathfrak{A}(213).$$

Proof. Let $x \in \mathfrak{A}$. Suppose that $x \geq 213$. We show that x is not 121-sortable. Due to Lemma 8.15, $x \geq 213$ if and only if $x \geq 1213$. Let $x_i x_j x_k x_l$ be the leftmost occurrence of 1213 in x. Observe that x_j must be extracted from the 121-stack before x_k enters. Therefore $x_j x_l x_1$ is an occurrence of 231 in $\mathcal{S}^{121}(x)$, as wanted.

Conversely, suppose that x is not 121-sortable and let bca be an occurrence of 231 in $\mathcal{S}^{121}(x)$. Let $b = x_i$, $c = x_j$ and $a = x_k$, for some i, j, k. Suppose, for a contradiction, that x avoids 213. Then x is a RGF due to Lemma 8.15. Since x avoids 213, the elements before $x_j = c$ in x must be in weakly increasing order. More precisely, it must be:

$$x = 1^{t_1} 2^{t_2} \cdots a^{t_a} \cdots b^{t_b} \cdots c^{t_c} x_j,$$

for some integers $t_u \geq 1$, $u = 1, \dots, c-1$, and $t_c \geq 0$. Now, if $i > j$, then $x_j = c$ is extracted before x_i enters, since otherwise $x_i x_j b$ would be an occurrence of 121 in the 12-stack. But this contradicts our assumption that x_i, x_j, x_k results in an occurrence of 231 in $\mathcal{S}^{121}(x)$. Therefore $i < j$ and, by hypothesis, $x_i = b$ is extracted from the 121-stack before $x_j = c$ enters. But this is impossible, since the prefix of x up to x_j is weakly increasing. □

Notice that the set of 121-sortable Cayley permutations is not a permutation class due to Corollary 8.5. By Theorems 8.18 and 8.19, we have $\mathrm{Sort}^{\mathfrak{A}}(12) = \mathrm{Sort}^{\mathfrak{A}}(121)$. However, the operations performed by the 12-stack and the 121-stack are not always the same. For example, $\mathcal{S}^{12}(12132) = 23211$, whereas $\mathcal{S}^{121}(12132) = 22311$.

Theorem 8.20. Let $\sigma \in \mathfrak{A}$ and suppose that σ contains 123. Then $\mathrm{Sort}^{\mathfrak{A}}(\sigma) = \mathfrak{A}(132)$.

Proof. Let $x \in \mathfrak{A}$. Suppose that $x \geq 132$. We show that x is not σ-sortable. If x avoids $\mathcal{R}(\sigma)$, then $\mathcal{S}^{\sigma}(x) = \mathcal{R}(x)$ contains $\mathcal{R}(132) = 231$, which means that x is not σ-sortable. On the other hand, suppose that x contains $\mathcal{R}(\sigma)$. Notice that $\mathcal{S}^{\sigma}(x)$ contains $\widehat{\sigma}$ due to Remark 8.1. Moreover, since σ contains 123, it is easy to observe that $\widehat{\sigma}$ contains an occurrence ab of 12 such that $a > 1$. Let y_1 and y_2 be the elements that correspond to a and b in such occurrence of 12 in $\mathcal{S}^{\sigma}(x)$. Notice also that the last element of $\mathcal{S}^{\sigma}(x)$ is $x_1 = 1$. Therefore $y_1 y_2 x_1$ is an occurrence of 231 in $\mathcal{S}^{\sigma}(x)$. Thus x is not σ-sortable.

Conversely, we show that if x avoids 132, then x is σ-sortable. If x avoids $\mathcal{R}(\sigma)$, then $\mathcal{S}^{\sigma}(x) = \mathcal{R}(x)$ avoids 231 and thus x is σ-sortable. Otherwise, suppose that $x \geq \mathcal{R}(\sigma)$. Since $\sigma \geq 123$, x contains an occurrence $x_{i_1} x_{i_2} x_{i_3}$ of $\mathcal{R}(123) = 321$. But then $x_1 x_{i_1} x_{i_2}$ would be an occurrence of 132 in x, a contradiction. □

Theorem 8.21. Let σ be an ascent sequence of length at least four and suppose that σ avoids 123. Then $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ is not a class.

Proof. Let $\sigma = \sigma_1 \cdots \sigma_k$, with $k \geq 4$. Observe that the ascent sequence 1232 is not σ-sortable. Indeed 1232 avoids $\mathcal{R}(\sigma)$, since σ has length at least four and $\mathcal{R}(1232) = 2321$ is not an ascent sequence. Thus $\mathcal{S}^{\sigma}(1232) = \mathcal{R}(1232) = 2321$ contains 231. We shall define an ascent sequence α such that α contains 1232 and α is σ-sortable, thus showing that $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ is not a class. Due to Lemma 8.16, we have $\max(\sigma) \leq 2$ for each i. We distinguish the following cases.

- Suppose that $\max(\sigma) = 1$, i.e. $\sigma = 1 \cdots 1 = 1^k$, for some $k \geq 1$. Define:

$$\alpha = 1^{k-1}2312.$$

  Note that $\alpha \geq 1232$ and α is an ascent sequence. Finally, it is easy to check that:

$$\mathcal{S}^\sigma(\alpha) = 32121^{k-1},$$

  which avoids 231. Thus α is σ-sortable.

- Suppose that $\max(\sigma) = 2$ and the last element is $\sigma_k = 1$. Define:

$$\alpha = \sigma_k \cdots \sigma_2 3 \sigma_1 2.$$

  Since $\max(\sigma) = 2$, there is an index m such that $\sigma_m = 2$. Notice that by our assumptions it must be $m \neq k$ and $m \neq 1$, thus α contains 1232. Also $3 \leq \operatorname{asc}(\sigma_k \cdots \sigma_2) + 2$ (and $\sigma_k = 1$), thus α is an ascent sequence. Finally, an easy computation shows that:

$$\mathcal{S}^\sigma(\alpha) = 3\sigma_2 1 \sigma_1 \sigma_3 \cdots \sigma_k,$$

  which avoids 231 (for example, because the initial 3 is the only element greater than 2 in $\mathcal{S}^\sigma(\alpha)$).

- Suppose that $\max(\sigma) = 2$ and the last element is $\sigma_k = 2$. Similarly to the previous case, define:

$$\alpha = 1\sigma_k \cdots \sigma_2 3 \sigma_1 2.$$

  Due to the insertion of the initial 1, α is again an ascent sequence and α contains 1232. Finally, we have that:

$$\mathcal{S}^\sigma(\alpha) = 3\sigma_2 1 \sigma_1 \sigma_3 \cdots \sigma_k 1,$$

  which avoids 231. We leave the details to the reader.

□

Corollary 8.22. Let σ be an ascent sequence. If $\sigma \in \{11, 12, 121\}$, then $\operatorname{Sort}^{\mathfrak{A}}(\sigma)$ is a class. In all the other cases, $\operatorname{Sort}^{\mathfrak{A}}(\sigma)$ is a class if and only if $\sigma \geq 123$. Moreover, if $\sigma \geq 123$, then $\operatorname{Sort}^{\mathfrak{A}}(\sigma) = \mathfrak{A}(132)$.

Proof. Patterns σ of length at most three, except 123, are discussed in Table 8.2 and theorems 8.17, 8.18 and 8.19. Patterns of greater length and the pattern 123 are discussed in Theorems 8.20 and 8.21. □

It is easy to observe that, in accordance with Corollary 8.5 and Theorem 8.14, if $\widehat{\sigma}$ contains 231, then σ contains 123 and $\operatorname{Sort}^{\mathfrak{A}}(\sigma)$ is indeed a class, as stated in Corollary 8.22.

σ	**σ-sortable ascent sequence**	**Non-σ-sortable pattern**
111	112312	1232
112	121312	1232
122	122312	1232

Table 8.2: Patterns σ of length at most three where $\mathrm{Sort}^{\mathfrak{A}}(\sigma)$ is not a class.

$\mathfrak{a}$ = $\mathfrak{b}$ =

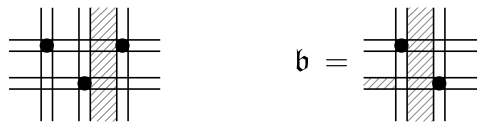

Figure 8.6: Cayley-mesh patterns such that $\mathfrak{MA} = \mathfrak{Cay}(\mathfrak{a}, \mathfrak{b})$.

8.2.2 Modified ascent sequences

The following two results can be found in [22].

Lemma 8.23. [22] Let $x \in \mathfrak{Cay}_n$ be a Cayley permutation. Then x is a modified ascent sequence if and only if the following two conditions hold:

1. $x_1 = 1$;
2. an entry $x_i = k > 1$ is the leftmost occurrence of the integer k in x if and only if $x_{i-1} < x_i$ (that is x_i is an ascent top).

Theorem 8.24. [22] Let $\mathfrak{a}$ and $\mathfrak{b}$ the Cayley-mesh patterns depicted in Figure 8.6. Then:

$$\mathfrak{MA} = \mathfrak{Cay}(\mathfrak{a}, \mathfrak{b}).$$

Theorem 8.24 is essentially a reformulation of Lemma 8.23 in terms of Cayley-mesh patterns. The avoidance of $\mathfrak{a}$ implies that every ascent top is the leftmost occurrence of the corresponding integer. Conversely, to avoid $\mathfrak{b}$ implies that each entry that is not an ascent top is not the leftmost occurrence of the corresponding integer.

Lemma 8.25. We have:

$$\mathfrak{MA}(213) = \mathfrak{MA}(1213).$$

Moreover, the set $\mathfrak{MA}(213)$ is enumerated by the Catalan numbers.

Proof. We start by showing that $\mathfrak{MA}(213) = \mathfrak{MA}(1213)$. Let $x = x_1 \cdots x_n \in \mathfrak{MA}$. It is enough to show that if $x \geq 213$, then $x \geq 1213$. Let $x_i x_j x_k$ be an occurrence of 213 in x. Let j' be the index of the leftmost occurrence of the integer x_j in x. If $j' < i$, then $x_{j'}x_ix_jx_k$ is an occurrence of 1213. If $i < j' \leq j$, then by Lemma 8.23 it must be $x_{j'-1} < x'_j$. Therefore we can repeat the same argument on the occurrence $x_ix_{j'-1}x_k$ of 213, until we either find an occurrence of 1213 or a contradiction.

Now, by Theorem 8.13, the set $\mathfrak{F}(3124)$ is Wilf-equivalent to the set $\mathfrak{MA}(\mathcal{B}(3124))$, where the Fishburn basis of 3124 is $\mathcal{B}(3124) = \{1213, 2314\}$ (see again [22]). In [34], the set $\mathfrak{F}(3124)$ is shown to be enumerated by the Catalan numbers. Finally, due to what proved above we have $\mathfrak{MA}(1213) = \mathfrak{MA}(213)$. Thus:

$$\mathfrak{MA}(1213, 2314) = \mathfrak{MA}(213, 2314) = \mathfrak{MA}(213),$$

and the thesis follows. □

Theorem 8.26. We have:

$$\mathrm{Sort}^{\mathfrak{MA}}(11) = \mathfrak{MA}(1213, 1223).$$

Proof. The proof of the inclusion $\mathrm{Sort}^{\mathfrak{MA}}(11) \subseteq \mathfrak{MA}(1213, 1223)$ is identical to the analogous inclusion of Theorem 8.17.

Conversely, let $x \in \mathfrak{MA}$ and suppose that x is not 11-sortable. Let $x_i = b$, $x_j = c$ and $x_k = a$ be the three elements of x that result in the leftmost occurrence $x_ix_jx_k = bca$ of 231 in $\mathcal{S}^{11}(x)$, with $a < b < c$. We wish to show that x contains 1213 or 1223. We distinguish two cases, according to whether $i < j$ or $i > j$.

- Suppose that $i < j$. Note that $x_i = b$ is extracted from the 11-stack before $x_j = c$ enters. Let y be the next element of the input when x_i is extracted. If $y = b$, then $x_1x_iyc \simeq 1223$. If $y \neq b$, then there must be another copy y' of the integer y in the 11-stack. If $y < b$, then $y'x_iyc \simeq 1213$. Otherwise, if $y > b$ then $x_iy'x_k$ is an occurrence of 231 that precedes $x_ix_jx_k$ in $\mathcal{S}^{11}(x)$, which is a contradiction.

- Suppose that $i > j$, that is $x_j = c$ precedes $x_i = b$ in x. Since x_i precedes x_j in $\mathcal{S}^{11}(x)$, x_j must be contained in the 11-stack when x_i enters. Let i' be the index of the first occurrence of the integer b in x. If $i' < j$, then $x_{i'}$ is extracted from the 11-stack before x_j enters. Otherwise both $x_{i'} = b$ and x_j (which is above $x_{i'}$) would be extracted (at most) when $x_i = b$ is the next element of the input, since $x_i = x_{i'}$. But this is impossible due the hypothesis that x_i precedes x_j in $\mathcal{S}^{11}(x)$. Consider the instant when $x_{i'}$ is extracted and let y be the next element of the input when this happens. If $y = b$, then $x_1x_{i'}yc \simeq 1223$, as desired. If $y \neq x_{i'}$, then there must be another copy of the integer y, say y', contained in the 11-stack. If $y' < b$, then $y'x_{i'}yc \simeq 1213$. Finally, if $y' > b$ then $x_{i'}y'x_1$ is an occurrence of 231 in $\mathcal{S}^{11}(x)$ that precedes $x_ix_jx_k$, contradicting our choice of i, j, k.

$\square$

Theorem 8.27. The set $\mathfrak{MA}(1213, 1223)$ is enumerated by the odd index Fibonacci numbers (sequence A001519 in [45]).

Proof. Due to Theorem 8.13, we have

$$\mathfrak{F}(1324) = \phi'\big(\mathfrak{MA}(1223, 1324)\big) \quad \text{and} \quad \mathfrak{F}(3124) = \phi'\big(\mathfrak{MA}(1213, 2314)\big).$$

Recall also that $\mathfrak{MA}(1213) = \mathfrak{MA}(213)$, as proved in Lemma 8.25. Thus, since $213 \leq 1324$ and $213 \leq 2314$, we have:

$$\mathfrak{MA}(1223, 1324, 1213, 2314) = \mathfrak{MA}(1223, 213) = \mathfrak{MA}(1213, 1223)$$

and

$$\mathfrak{F}(1324, 3124) = \phi'\big(\mathfrak{MA}(1223, 1324, 1213, 2314)\big) = \phi'\big(\mathfrak{MA}(1213, 1223)\big).$$

Let $F(n) = \mathfrak{F}_n(1324, 3124)$ and let $f(n) = |F(n)|$. We show that the coefficients $f(n)$ satisfy $f(1) = 1$, $f(2) = 2$ and $f(n+1) = 3f(n) - f(n-1)$, for $n \geq 2$, which is a very well known recurrence for the odd index Fibonacci numbers. Let $G(n, k)$ be the set:

$$G(n, k) = \{p \in F(n) : \mathrm{ltrmax}(p) = k\}.$$

Let $g(n, k) = |G(n, k)|$. Notice that $F(n) = \dot{\bigcup}_k G(n, k)$ and thus $f(n) = \sum_{k=1}^{n} g(n, k)$. We show that, for any $n \geq 1$:

$$\begin{cases} g(n+1, 1) = f(n) \\ g(n+1, 2) = f(n) \\ g(n+1, k+1) = g(n, k), \quad k \geq 2. \end{cases}$$

- Let us start by proving the first equation $g(n+1, 1) = f(n)$. We provide a bijection $\alpha : F(n) \to G(n+1, 1)$. Given $p \in F(n)$, define $\alpha(p) = 1 \ominus p$. Equivalently, $\alpha(p)$ is obtained from p by adding an initial maximum $n+1$. It is easy to realize that $\alpha(p) \in G(n+1, 1)$ for each $p \in F(n)$ and α is injective. Finally, if $q \in G(n+1, 1)$, then $q_1 = n+1$ and the removal of $n+1$ from q yields a permutation $p \in F(n)$, thus α is surjective too.

- Similarly, we define a bijection $\beta : F(n) \to G(n+1, 2)$ by suitably adding a new maximum $n+1$ to a permutation $p \in F(n)$. If $p \in G(n, 1)$, then $p_1 = n$ and we set $\beta(p) = n(n+1)p_2 \cdots p_n$. Otherwise, if $p \in G(n, k)$, for some $k \geq 2$, let:
$$p = m_1 A_1 m_2 A_2 \cdots m_k A_k$$
be the ltr-max decomposition of p. Then define $\beta(p)$ by
$$\beta(p) = m_1 A_1 (n+1) m_2 A_2 \cdots m_k A_k.$$
It is easy to realize that $\beta(p)$ avoids $\mathfrak{f}$, 1324 and 3124. The case $k = 1$ is trivial. On the other hand, if $k \geq 2$ then an occurrence of any of the listed patterns in $\beta(p)$ should involve the new element $n + 1$, either as a 4 in an occurrence of 1324 or 3124 or as a 3 in an occurrence of $\mathfrak{f}$. But then, in all these cases, the element m_2 would play the same role in an occurrence of the same pattern in p, which is impossible since $p \in F(n)$. A similar analysis shows that the removal of $n + 1$ from a permutation in $G(n + 1, 2)$ yields a permutation in $F(n)$, thus β is bijective and $g(n + 1, 2) = f(n)$, as wanted.

- Next suppose that $k \geq 2$. We provide a bijection $\gamma : G(n, k) \to G(n+1, k+1)$. Let $p \in G(n, k)$ and write again:
$$p = m_1 A_1 \cdots m_{k-1} A_{k-1} m_k A_k.$$
Let $A_k = a_1 \cdots a_t$ and define:
$$\gamma(p) = m_1 A_1 \cdots m_{k-1} A_{k-1} m'_{k-1} m'_k a'_1 \cdots a'_t,$$
where $m'_k = m_k + 1$ and $a'_i = a_i + 1$, if $a_i > m_{k-1}$, or $a'_i = a_i$, if $a_i < m_k$. In other words, $\gamma(p)$ is obtained from p by inserting $m_{k-1} + 1$ immediately before $m_k = n$ and suitably rescaling the other elements. Only those elements contained in the last block A_k, m_k included, eventually need to be rescaled. Notice that $\gamma(p)$ has $k + 1$ ltr-maxima and γ is injective by construction. The proof that $\gamma(p)$ avoids $\mathfrak{f}$, 1324 and 3124 is identical to the previous cases, so we omit it. Finally, in the resulting permutation $\gamma(p)$, the new element $m_{k-1} + 1$ is the k-th ltr-maximum. Thus the inverse map $\gamma^{-1} : G(n + 1, k + 1) \to G(n, k)$ is obtained by removing the k-th ltr-maximum, which again does not create an occurrence of one of the forbidden patterns.

Due to what proved above, using induction, we have:

$$\begin{aligned} f(n+1) = \sum_{k=1}^{n+1} g(n+1,k) &= \\ g(n+1,1)+g(n+1,2)+\sum_{k=3}^{n+1} g(n+1,k) &= \\ f(n)+f(n)+\sum_{k=3}^{n+1} g(n,k-1) &= \\ f(n)+f(n)+\sum_{j=2}^{n} g(n,j) &= \\ f(n)+f(n)+[f(n)-g(n,1)] &= \\ 3f(n)-f(n-1)&, \end{aligned}$$

as desired. □

The next result for the pattern 12 is again a corollary of Theorem 8.14.

Theorem 8.28. We have:

$$\mathrm{Sort}^{\mathfrak{MA}}(12) = \mathfrak{MA}(213).$$

Theorem 8.29. We have:

$$\mathrm{Sort}^{\mathfrak{MA}}(121) = \mathfrak{MA}(213).$$

Proof. Due to Lemma 8.25, we have $\mathfrak{MA}(213) = \mathfrak{MA}(1213)$. Let $x \in \mathfrak{MA}$ and suppose that $x \geq 1213$. We show that x is not 121-sortable. Let $x_{i_1}x_jx_{i_2}x_k$ be an occurrence of 1213 in x. We can assume that $x_{i_1}x_jx_{i_2}$ is the leftmost occurrence of 121 such that the element that plays the role of 2 is smaller than x_k. If x_j is extracted from the 121-stack before x_k enters, then $\mathcal{S}^{121}(x)$ contains $x_jx_kx_1 \simeq 231$, thus x is not 121-sortable. Otherwise, suppose that x_j is still in the 121-stack when x_k enters. Observe that, since $x_{i_2}x_jx_{i_1} \simeq 121$, the element x_{i_1} is extracted before x_j enters. Otherwise x_{i_1} and x_j would be extracted from the 121-stack at most when x_{i_2} is the next element of the input (which contradicts the assumption that x_j is still in the 121-stack when x_k enters). Let y_1 be the next element of the input when x_{i_1} is extracted. Since the 121-stack restriction is triggered, there are two elements in the 121-stack, say z and y_2, with z above y_2, such that $y_1zy_2 \simeq 121$. Due to our choice of i_1, j, i_2, it must be $z > x_k$ (and thus $z \neq x_{i_1}$). Notice that it cannot be $x_{i_1} = 1$, since in that case the 121-stack would contain $x_{i_1}zx_1 \simeq 121$. Then $x_{i_1} > 1$ and $\mathcal{S}^{121}(x)$ contains an occurrence $x_{i_1}zx_1$ of 231, as wanted.

Conversely, suppose that x is not 121-sortable. Equivalently, suppose there are three elements $x_i = b$, $x_j = c$ and $x_k = a$ such that $x_i x_j x_k = bca$ is an occurrence of 231 in $\mathcal{S}^{121}(x)$. We show that x contains 213. We can assume that bca is the leftmost occurrence of 231 in $\mathcal{S}^{121}(x)$. We distinguish two cases, according to whether $i < j$ or $j > i$.

- Suppose that $i < j$, i.e. $x_i = b$ precedes $x_j = c$ in x. By hypothesis x_i is extracted from the 121-stack before x_j enters. Therefore, at that moment, the 121-stack contains two elements zy_1, with z above y_1, such that $y_2 z y_1 \simeq 121$, where y_2 is the next element of the input. If $y_2 < b$, then $x_i y_2 x_j$ is an occurrence of 213 in x, as wanted. If instead $y_2 = b$, then the stack contains an occurrence $x_i z y_1$ of 121, which is forbidden. Finally, if $y_1 > b$, then also $x > b$ and thus $x_i z x_k$ is an occurrence of 231 in $\mathcal{S}^{121}(x)$ that precedes $x_i x_j x_k$, which is impossible due to our choice of i, j, k.

- Suppose that $j < i$, i.e. $x_j = c$ precedes $x_i = b$ in x. Note that x_j must be in the 121-stack when x_i enters. Moreover, since $x_i x_j x_k$ is the leftmost occurrence of 231 in $\mathcal{S}^{121}(x)$, $x_i = b$ is the first occurrence of the integer b that is extracted from the 121-stack. Let i' be the index of the first occurrence of the integer b in x. If $i' < j$, then $x_{i'} x_j x_i \simeq 121$. Therefore at least one between $x_{i'}$ and x_j must be extracted from the 121-stack before x_i enters (otherwise we would have $x_i x_j x_{i'} \simeq 121$ inside the 121-stack). As said before, x_j is still contained in the 121-stack when x_i enters, therefore $x_{i'}$ is the one that has been extracted before. But then $x_{i'} x_j x_k$ is an occurrence of 231 in $\mathcal{S}^{121}(x)$ that precedes $x_i x_j x_k$, which is impossible. We can thus assume that $i' > j$. Now, consider the element $x_{i'-1}$. Due to Lemma 8.23, we have $x_{i'-1} < x_{i'}$. If $x_{i'-1} = 1$, then $x_{i'-1} x_j x_1 \simeq 121$ and thus x_j is extracted from the 121-stack before x_i enters, which is a contradiction. If instead $x_{i'-1} > 1$, we can repeat the same argument, but using the first occurrence x_w of $x_{i'-1}$ in x in place of $x_{i'}$ (if $w < j$, then x_j is extracted too soon and otherwise we consider x_{w-1}). Sooner or later this would result in a contradiction.

$\square$

The proof of the next result is analogous to the proof of Theorem 8.20, and it is left to the reader.

Theorem 8.30. Let $\sigma = \sigma_1 \cdots \sigma_k$ be a modified ascent sequence of length at least three and suppose that $\sigma \geq 123$. Then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) = \mathfrak{MA}(132)$.

Theorem 8.31. Let $\sigma = \sigma_1 \cdots \sigma_k$ be a modified ascent sequence of length at least three. If σ avoids 123 and $\sigma_1\sigma_2\sigma_3 \simeq 122$, then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) = \mathfrak{MA}(132, \mathcal{R}(\sigma) \oplus 1)$.

Proof. Suppose that σ avoids 123 and let $\sigma_2 = \sigma_3 = t$, for some $t > 1$. Observe that $t = \max(\sigma)$, otherwise $\sigma_1\sigma_2 = 1t$ would realize an occurrence of 123 together with the maximum of σ. Let x be a modified ascent sequence.

We firts show that if x is not σ-sortable, then $x \notin \mathfrak{MA}(132, \mathcal{R}(\sigma)) \oplus 1)$, thus proving the inclusion $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) \supseteq \mathfrak{MA}(132, \mathcal{R}(\sigma) \oplus 1)$. Suppose that the three elements $x_i = a$, $x_j = b$ and $x_k = c$ result in an occurrence bca of 231 in $\mathcal{S}^\sigma(x)$. If $j > k$, then $x_1x_kx_j$ is an occurrence of 132 in x, as wanted. Otherwise, suppose that $j < k$ and x_j is extracted from the σ-stack before x_k enters. When x_j is extracted, there must be $k-1$ elements $\sigma'_2, \ldots, \sigma'_k$ in the σ-stack, reading from top to bottom, such that $\sigma'_1\sigma'_2 \cdots \sigma'_k$ is an occurrence of σ, where σ'_1 is the next element of the input. Now, if $\sigma'_2 < c$, then $\sigma'_u < c$ for each u, since $\sigma_2 = \max(\sigma)$. Therefore $\sigma'_k \cdots \sigma'_2\sigma'_1c$ is an occurrence of $\mathcal{R}(\sigma) \oplus 1$, as wanted. If instead $\sigma'_2 > c$, then $x_1\sigma'_2c$ is an occurrence of 132 in x. Finally, suppose that $\sigma'_2 = c$. Thus $x_k = c$ is not the leftmost occurrence of the integer c in x, since x_k follows $\sigma'_2 = c$ in x. By Lemma 8.23, it must be $x_{k-1} \geq x_k$ (and thus $x_{k-1} \neq \sigma'_1$). If $x_{k-1} > x_k$, then $x_1x_{k-1}x_k$ is an occurrence of 132. If $x_{k-1} = x_k$, then we can repeat the same argument, but using x_{k-1} instead of x_k. Sooner or later this will result in either an occurrence of 132 or a contradiction.

Conversely, we shall prove the inclusion $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) \subseteq \mathfrak{MA}(132, \mathcal{R}(\sigma) \oplus 1)$ by showing that if x contains either 132 or $\mathcal{R}(\sigma) \oplus 1$, then x is not σ-sortable. Suppose initially that $x \geq 132$. Without losing generality, choose the leftmost occurrence $x_1x_ix_j$ of 132 in x. If x_j enters the σ-stack above x_i, then $\mathcal{S}^\sigma(x)$ contains an occurrence $x_jx_ix_1$ of 231, thus x is not σ-sortable. Suppose instead that x_i is extracted from the σ-stack before x_j enters. At that moment, the next element of the input σ'_1 forms an occurrence of σ together with some elements $\sigma'_2, \ldots, \sigma'_k$ contained in the σ-stack. Notice that $\sigma'_2 = \sigma'_3 > 1$, since $\sigma_1\sigma_2\sigma_3 \simeq 122$. Due to our choice of $x_1x_ix_j$ as leftmost occurrence of 132 in x, it must be $\sigma'_u \leq x_j$ for each $u \leq k-1$. Since $\sigma_1 < \sigma_2$, we also have $\sigma'_1 < x_j$ and thus $\sigma'_2 \neq x_i$. If $\sigma'_2 < x_j$, then $\mathcal{S}^\sigma(x)$ contains an occurrence $\sigma'_2x_jx_1$ of 231 and we are done. Therefore we can assume $\sigma'_2 = x_j$ (and thus $\sigma'_2 = \sigma'_3 = x_j$). Due to Lemma 8.23, it must be $x_{j-1} \geq x_j$. Also $x_{j-1} \geq x_i$, again due to our choice of $x_1x_ix_j$ as leftmost occurrence of 132, and thus $x_{j-1} \neq \sigma'_1$. If $x_{j-1} > x_i$, then $x_ix_{j-1}x_1$ is an occurrence of 231 in $\mathcal{S}^\sigma(x)$. Finally, let $x_{j-1} = x_i$. Then again $x_{j-2} \geq x_{j-1}$ due to Lemma 8.23 and we can repeat the same argument on x_{j-2}. Sooner or later, since the next element of the input is $\sigma'_1 \neq x_{j-1}$, either we will find an occurrence of 132 or a contradiction. This proves that if $x \geq 132$, then x is not σ-sortable. Next suppose that x avoids 132, but x contains an occurrence $\sigma'_k \cdots \sigma'_2\sigma'_1m$ of $\mathcal{R}(\sigma) \oplus 1$. If the elements $\sigma'_k \cdots \sigma'_2$ are still in the σ-stack when σ'_1 is the next element of the input, then σ'_2 is extracted and $\mathcal{S}^\sigma(x)$ contains an occurrence σ'_2mx_1 of 231. Otherwise, there must have been a previous occurrence, say $\sigma''_k \cdots \sigma''_2\sigma''_1$, of $\mathcal{R}(\sigma)$

in x such that at least one element amongst $\sigma'_k, \ldots, \sigma'_2$ is extracted when σ''_1 is the next element of the input, together with σ''_2. Observe that it must be $\sigma''_2 \leq m$, since we are assuming that x avoids 132. If $\sigma''_2 < m$, then again $\sigma''_2 m x_1$ is an occurrence of 231 in $\mathcal{S}^\sigma(x)$. Finally, if $\sigma''_2 = m$, then $x_1 \sigma''_2 \sigma'_2$ is an occurrence of 132, a contradiction. □

Lemma 8.32. Let $x = x_1 \cdots x_k$ be a modified ascent sequence and suppose that x avoids 123. If $x_k = 1$, then $\mathcal{R}(x)$ is a modified ascent sequence. Otherwise, if $x_k > 1$, then $1\mathcal{R}(x)$ is a modified ascent sequence.

Proof. Let $m = \max(x)$. Since x avoids 123 (and $x_1 = 1$), the elements of x that are greater than 1 are in weakly decreasing order. Therefore, by Lemma 8.23, we have:

$$x = 1^{i_1} m^{j_1} 1^{i_2} (m-1)^{j_2} \cdots 1^{i_{m-1}} 2^{j_{m-1}} 1^{i_m},$$

where $i_u \geq 1$ for each $u < m$, $i_m \geq 0$ and $j_v \geq 1$ for each v. Therefore, again by Lemma 8.23, if $x_k = 1$ (or equivalently $i_m \geq 1$) then $\mathcal{R}(x)$ is a modified ascent sequence. Similarly, if $x_k > 1$, then we obtain a modified ascent sequence by inserting an additional 1 at the beginning of $\mathcal{R}(x)$. □

Theorem 8.33. Let $\sigma = \sigma_1 \cdots \sigma_k$ be a modified ascent sequence of length at least four. If σ avoids 123 and $\sigma_1\sigma_2\sigma_3$ is not an occurrence of 122, then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is not a class.

Proof. Observe that the modified ascent sequence $\alpha = 1312$ is not σ-sortable, since $\mathcal{S}^\sigma(\alpha) = \mathcal{R}(\alpha) = 2131 \geq 231$, for each σ of length four or more. We wish to construct a σ-sortable modified ascent sequence β, with $\beta \geq \alpha$. We distinguish some cases. Being the other cases similar, we give a detailed proof for the first one only.

- Suppose that $\sigma_2 = 1$ and $\sigma_k = 1$. Define:

  $$\alpha = \sigma_k \cdots \sigma_3 \sigma_2 (m+2) \sigma_1 (m+1),$$

  where $m = \max(\sigma)$. Notice that $\sigma_k(m+2)\sigma_1(m+1)$ is an occurrence of 1312 in α. Now, an easy computation shows that:

  $$\mathcal{S}^\sigma(\alpha) = (m+2)\sigma_2(m+1)\sigma_1\sigma_3 \cdots \sigma_k.$$

  Observe that $\mathcal{S}^\sigma(\alpha)$ avoids 231. Indeed $m+2$ and $m+1$ are not part of an occurrence of 231 (since $\sigma_2 = 1$). Moreover, suppose, for a contradiction, that $\sigma_3 \cdots \sigma_k$ contains an occurrence $\sigma_{i_i}\sigma_{i_2}\sigma_{i_3}$ of of 231. Then $\sigma_1\sigma_{i_2}\sigma_{i_3}$ is an occurrence of 123 in σ, which contradicts the hypothesis. Finally, we shall prove that α is a modified ascent sequence. Notice that $\mathcal{R}(\sigma)$ is a modified

σ	**σ-sortable modified sequence**	**Non-σ-sortable pattern**
111	11312	1312
112	121413	1312

Table 8.3: Patterns σ of length at most three where $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is not a class.

ascent sequence by Lemma 8.32. The only additional elements are $m+2$, which is placed immediately after $\sigma_2 = 1$ and before $\sigma_1 = 1$, and $m+2$, which is placed at the end, immediately after $\sigma_1 = 1$. Therefore α is a modified ascent sequence by Lemma 8.23.

- If $\sigma_2 = 1$ and $\sigma_k > 1$, then define α exactly as in the previous case, but inserting an additional 1 at the beginning (as in Lemma 8.32).
- Suppose that $\sigma_2 > 1$ and $\sigma_k = 1$. Then $\sigma_2 = m$ is equal to the maximum value of α, because α avoids 123. Since $\sigma_1\sigma_2\sigma_3$ is not an occurrence of 122, Lemma 8.23 implies that σ_2 is the only occurrence of m in σ. Define:

  $$\alpha = \sigma_k \cdots \sigma_3 (m+1) \sigma_1 \sigma_2.$$

  Then α contains 1312, α is a modified ascent sequence and:

  $$\mathcal{S}^{\sigma}(\alpha) = (m+1)\sigma_2\sigma_1\sigma_3 \cdots \sigma_k,$$

  which avoids 231. We leave the details to the reader.
- If $\sigma_2 > 1$ and $\sigma_k > 1$, then define α exactly as in the previous case, but inserting an additional 1 at the beginning (as in Lemma 8.32).

□

The following corollary, which is an immediate consequence of the results proven so far in this section, provides a characterization of the patterns σ where $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is a class.

Corollary 8.34. Let $\sigma = \sigma_1 \cdots \sigma_k$ be a modified ascent sequence. If $\sigma \in \{11, 12\}$, then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is a class. In all the other cases, $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is a class if and only if $\sigma \geq 123$ or $\sigma_1\sigma_2\sigma_3 \simeq 122$. Moreover, if $\sigma \geq 123$, then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) = \mathfrak{MA}(132)$. If instead σ avoids 123 and $\sigma_1\sigma_2\sigma_3 \simeq 122$, then $\mathrm{Sort}^{\mathfrak{MA}}(\sigma) = \mathfrak{MA}(132, \mathcal{R}(\sigma) \oplus 1)$.

Proof. The patterns 11, 12, 121, 111 and 112 were solved in Theorem 8.26, Theorem 8.28, Theorem 8.29 and Table 8.3. The remaining cases were considered in Theorems 8.30, 8.31 and 8.33. □

Similarly to what observed in the previous section, in accordance with Corollary 8.5 and Theorem 8.14, if $\widehat{\sigma}$ contains 231, then σ contains 123 and $\mathrm{Sort}^{\mathfrak{MA}}(\sigma)$ is a class, as in Corollary 8.34.

Bibliography

[1] M. Albert,
PermLab: Software for Permutation Patterns,
at http://www.cs.otago.ac.nz/staffpriv/malbert/permlab.php.

[2] M. H. Albert, R. E. L. Aldred, M. D. Atkinson, C. Handley, D. Holton,
Permutations of a multiset avoiding permutations of length 3,
European Journal of Combinatorics, Vol. 22(8), pp. 1021-1031, 2001.

[3] M. H. Albert, M. Bouvel,
A general theory of Wilf-equivalence for Catalan structures,
The Electronic Journal of Combinatorics, Vol. 22(4), 2015.

[4] M. D. Atkinson, S. A. Linton, L. A. Walker,
Priority queues and multisets,
The Electronic Journal of Combinatorics, Vol. 2, 1995.

[5] D. Avis and M. Newborn,
On pop-stacks in series,
Utilitas Mathematica, Vol. 19, pp. 129–140, 1981.

[6] A. Bacher, A. Bernini, L. Ferrari, B. Gunby, R. Pinzani, J. West,
The Dyck pattern poset,
Discrete Mathematics, Vol. 321, pp. 12–23, 2014.

[7] E. Barcucci, A. Del Lungo, E. Pergola, R. Pinzani,
ECO: a methodology for the enumeration of combinatorial objects,
Journal of Difference Equations and Applications, Vol. 5, pp. 435–490, 1999.

[8] J.-L. Baril, G. Cerbai, C. Khalil, V. Vajnovszki,
Catalan and Schröder permutations sortable by two restricted stacks,
arXiv e-prints arXiv:2004.01812, 2020.

[9] P. Barry,
Riordan arrays, generalized Narayana triangles, and series reversion,
Linear Algebra and its Applications, Vol. 491, pp. 343–385, 2016.

[10] A. Baxter, L. Pudwell,
Ascent sequences avoiding pairs of patterns,
The Electronic Journal of Combinatorics, Vol. 22(1), pp. 23, 2015.

[11] K. Berlow,
Restricted stacks as functions,
arXiv e-prints arXiv:2008.01164, 2020.

[12] D. Bevan,
Permutation patterns: basic definitions and notations,
arXiv e-prints arXiv:1506.06673, 2015.

[13] M. Bona,
A survey of stack sorting disciplines,
The Electronic Journal of Combinatorics, Vol. 9(2), 2002-2003.

[14] M. Bona,
Combinatorics of Permutations,
Discrete Mathematics and Its Applications, CRC Press, 2004.

[15] M. Bousquet Mélou,
Discrete excursions,
Séminaire Lotharingien de Combinatoire, Vol. 57, pp. 23, 2008.

[16] M. Bousquet-Mélou,
Sorted and/or sortable permutations,
Discrete Mathematics, Vol. 225(1-3), pp. 25–50, 2000.

[17] M. Bousquet-Mélou, A. Claesson, M. Dukes, S. Kitaev,
(2+2)-free posets, ascent sequences and pattern avoiding permutations,
Journal of Combinatorial Theory, Series A, Vol. 117(7), pp. 884–909, 2010.

[18] P. Brändén, A. Claesson,
Mesh Patterns and the expansion of permutation statistics as sums of permutation patterns,
The Electronic Journal of Combinatorics, Vol. 18(2), pp. 14, 2011.

[19] L.R. Campbell, S. Dahlberg, R. Dorward, J. Gerhard, T. Grubb, C. Purcell, B. Sagan,
Restricted growth function patterns and statistics,
Advances in Applied Mathematics, Vol. 100, pp. 1–42, 2008.

[20] G. Cerbai,
Pattern-avoiding machines on Cayley permutations,
arXiv e-prints arXiv:2003.02536, 2020.

[21] G. Cerbai, L. Cioni, L. Ferrari,
Stack sorting with increasing and decreasing stacks,
The Electronic Journal of Combinatorics, Vol. 27(1), pp. 16, 2020.

[22] G. Cerbai, A. Claesson,
Transport of patterns via Burge transpose,
arXiv e-prints arXiv:2005.07950, 2020.

[23] G. Cerbai, A. Claesson, L. Ferrari,
Stack sorting with restricted stacks,
Journal of Combinatorial Theory, Series A, Vol. 173, 2020.

[24] G. Cerbai, A. Claesson, L. Ferrari, E. Steingrímsson,
Sorting with pattern-avoiding stacks: the 132*-machine*,
The Electronic Journal of Combinatorics, Vol. 27(3), 2020.

[25] L. Cioni, L. Ferrari,
Enumerative results on the Schröder pattern poset,
Cellular Automata and Discrete Complex Systems, AUTOMATA 2017, Lecture Notes in Computer Science, Vol. 10248, pp. 56–67, 2017.

[26] J. A. Clapperton, P. J. Larcombe, E. J. Fennessey,
On iterated generating functions for integer sequences, and Catalan polynomials,
Utilitas Mathematica, Vol. 77, pp. 3–33, 2008.

[27] C. Defant, N. Kravitz,
Stack-sorting for words,
Australasian Journal of Combinatorics, Vol. 77(1), pp. 51–68, 2020.

[28] C. Defant, K. Zheng,
Stack-Sorting with Consecutive-Pattern-Avoiding Stacks,
arXiv e-prints arXiv:2008.12297, 2020.

[29] E. Deutsch,
Dyck path enumeration,
Discrete Mathematics, Vol. 204(1-3), pp. 167–202, 1999.

[30] P. Duncan, E. Steingrímsson,
Pattern avoidance in ascent sequences,
The Electronic Journal of Combinatorics, Vol. 18(1), pp. 17, 2011.

[31] L. Ferrari,
On the foundations of ECO method,
Doctoral Thesis, University of Florence, 2003.

[32] P. Flajolet,
Combinatorial aspects of continued fractions,
Discrete Mathematics, Vol. 32(2), pp. 125–161, 1980.

[33] I. Gessel and G. Xin,
A combinatorial interpretation of the numbers $6(2n)!/n!(n+2)!$,
Journal of Integer Sequences, Vol. 8, pp. 13, 2005.

[34] J. B. Gil, M. D. Weiner,
On pattern-avoiding Fishburn permutations,
Annals of Combinatorics, Vol. 23, pp. 785–800, 2019.

[35] V. Jelínek, T. Mansour,
On pattern-avoiding partitions,
The Electronic Journal of Combinatorics, Vol. 15, pp. 52, 2008.

[36] V. Jelínek, T. Mansour, M. Shattuck
On multiple pattern avoiding set partitions,
Advances in Applied Mathematics, Vol. 50(2), pp. 292–326, 2013.

[37] S. Kitaev,
Patterns in permutations and words,
Monographs in Theoretical Computer Science. An EATCS Series, Springer, Heidelberg, 2011.

[38] D. E. Knuth,
The art of computer programming, Vol. 1, Fundamental Algorithms,
Addison-Wesley, Reading, Massachusetts, 1973.

[39] C. Krattenthaler,
Permutations with restricted patterns and Dyck paths,
Advances in Applied Mathematics, Vol. 27(2-3), pp. 510–530, 2001.

[40] T. Mansour, M. Shattuck,
Some enumerative results related to ascent sequences,
Discrete Mathematics, Vol. 315-316, pp. 29–41, 2014.

[41] T. Mansour, A. Vainshtein,
Restricted 132*-avoiding permutations*,
Advances in Applied Mathematics, Vol. 26(3), pp. 258-269, 2001.

[42] M. Murphy,
Restricted permutations, antichains, atomic classes, and stack sorting,
Doctoral Thesis, University of St Andrews, 2002.

[43] A. Pierrot, D. Rossin,
2-*stack sorting is polynomial*,
Theory of Computing Systems, Vol. 60, pp. 552–579, 2017.

[44] B. Sagan,
Pattern avoidance in set partitions,
Ars Combinatoria, Vol. 94, pp. 79-96, 2010.

[45] N. J. A. Sloane,
The On-Line Encyclopedia of Integer Sequences,
published electronically at https://oeis.org.

[46] R. Smith,
Two stacks in series: a decreasing stack followed by an increasing stack,
Annals of Combinatorics, Vol. 18(2), pp. 359-363, 2014.

[47] R. Smith,
Permutation reconstruction,
The Electronic Journal of Combinatorics, Vol. 13, 2006.

[48] R. E. Tarjan,
Sorting using networks of queues and stacks,
Journal of the ACM, Vol. 19, pp. 341–346, 1972.

[49] J. West,
Permutations with forbidden subsequences and stack-sortable permutations,
Doctoral Thesis, Massachusetts Institute of Technology, 1990.

[50] J. West,
Sorting twice through a stack,
Theoretical Computer Science, Vol. 117(1-2), pp. 303–313, 1993.

[51] J. West,
Generating trees and the Catalan and Schröder numbers,
Discrete Mathematics, Vol. 146(1-3), pp. 247–262, 1995.

[52] J. West,
Generating trees and forbidden subsequences,
Discrete Mathematics, Vol. 157(1-3), pp. 363–374, 1996.

Appendix A

Index of integer sequences

OEIS	**Sequence**	**References**
A000079	Powers of 2	Appendix B, Table 7.1
A000108	Catalan numbers	Example 1.1, Table 6.1, Table 7.1
A000110	Bell numbers	Section 1.3
A000124	Central polygonal numbers	Table 7.1
A000670	Fubini numbers	Section 1.3
A001006	Motzkin numbers	Section 1.4
A001519	Odd indexed Fibonacci numbers	Appendix B, Table 3.2, Theorem 8.27
A002057	4-th convolution of Catalan numbers	Proposition 3.12
A006318	Large Schröder numbers	Section 1.4, Table 6.1
A007317	Binomial transform of Catalan numbers	Chapter 5, Table 6.1
A009766	Catalan triangle (ballot numbers)	Section 6.4
A011782		Table 3.2
A022493	Fishburn numbers	Section 1.3
A024175		Table 3.2
A033184	Catalan triangle transposed	Section 6.4
A080937		Appendix B, Table 3.2
A102407		Table 6.1
A115139	Catalan polynomials	Section 3.4
A116845		Appendix B
A124302		Appendix B, Table 3.2
A129591		Theorem 3.15
A202062		Table 3.3
A294790	Subtract n from partial sums of partial sums of Catalan numbers	Chapter 4

Appendix B

Non-principal classes Sort(σ)

σ	G.F.	Sequence $\{f_n^\sigma\}_n$	OEIS
321	$\frac{1-t}{1-2t}$	1, 2, 4, 8, 16, 32, 64, 128, 256, 512	A000079
3214 4213 4312 4321	$\frac{1-2t}{1-3t+t^2}$	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181	A001519
32145	$\frac{-3t^4+9t^3-12t^2+6t-1}{(t-1)(t^2-3t+1)^2}$	1, 2, 5, 14, 41, 121, 355, 1032, 2973, 8496	A116845
52134	$\frac{(1-t)(2t-1)^2}{t^4-9t^3+12t^2-6t+1}$	1, 2, 5, 14, 41, 121, 355, 1033, 2986, 8594	
54123	$\frac{1-4t+5t^2-3t^3}{t^4-6t^3+8t^2-5t+1}$	1, 2, 5, 14, 41, 121, 356, 1044, 3057, 8948	
32154 42135 43125 43215 52143 53124 53214 54132 54213 54312 54321	$\frac{t^2-3t+1}{3t^2-4t+1}$	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842	A124302

σ	**G.F.**	**Sequence** $\{f_n^\sigma\}_n$	**OEIS**
321645[1] ⋮ 654321	$\frac{3t^2-4t+1}{1-5t+6t^2-t^3}$	1, 2, 5, 14, 42, 131, 417, 1341, 4334, 14041	A080937
421356 431256 432156	$\frac{2t^5-16t^4+29t^3-23t^2+8t-1}{9t^5-33t^4+46t^3-30t^2+9t-1}$	1, 2, 5, 14, 42, 131, 416, 1329, 4247, 13544	
631245 632145	$\frac{t^5-7t^4+17t^3-17t^2+7t-1}{4t^5-16t^4+29t^3-23t^2+8t-1}$	1, 2, 5, 14, 42, 131, 416, 1329, 4247, 13545	
621345	$F^{621345(t)}$[2]	1, 2, 5, 14, 42, 131, 414, 1304, 4065, 12530	
521346 651243 651324 652134	$\frac{t^4-9t^3+12t^2-6t+1}{5t^4-17t^3+17t^2-7t+1}$	1, 2, 5, 14, 42, 131, 416, 1329, 4248, 13560	
541236 641235 654123	$\frac{t^4-6t^3+8t^2-5t+1}{4t^4-11t^3+12t^2-6t+1}$	1, 2, 5, 14, 42, 131, 416, 1330, 4261, 13658	
321465 321546	$\frac{t^4-12t^3+16t^2-7t+1}{6t^4-23t^3+22t^2-8t+1}$	1, 2, 5, 14, 42, 131, 416, 1328, 4233, 13430	
621354 621435	$\frac{t^4-8t^3+12t^2-6t+1}{(t^2-3t+1)(4t^2-4t+1)}$	1, 2, 5, 14, 42, 131, 416, 1328, 4234, 13446	
651234	$\frac{1-6t+14t^2-17t^3+10t^4-4t^5}{t^6-9t^5+20t^4-27t^3+19t^2-7t+1}$	1, 2, 5, 14, 42, 131, 414, 1306, 4094, 12766	
321456	$F^{321456}(t)$[3]	1, 2, 5, 14, 42, 131, 414, 1304, 4063, 12497	

[1] 321654, 421365, 431265, 432165, 521436, 531246, 532146, 541326, 542136, 543126, 543216, 621534, 621543, 631254, 632154, 641325, 642135, 643125, 643215, 651423, 651432, 652143, 653124, 653214, 654132, 654213, 654312, 654321.

[2] $F^{621345(t)} = \frac{(1-t)^3(2t-1)^3}{t^7-24t^6+74t^5-109t^4+89t^3-41t^2+10t-1}$.

[3] $F^{321456}(t) = \frac{1-11t+50t^2-122t^3+175t^4-152t^5+79t^6-25t^7+4t^8}{(1-t)^3(t^2-3t+1)^3}$.

Appendix C

Enumerative data for $\mathrm{Sort}(\sigma)$

σ	**Sequence** $\{f_n^\sigma\}_n$
123	1, 2, 5, 13, 35, 99, 295, 920, 2975, 9892, 33605
132	1, 2, 5, 15, 51, 188, 731, 2950, 12235, 51822, 223191
213	1, 2, 5, 16, 62, 273, 1307, 6626, 35010, 190862, 1066317
231	1, 2, 6, 23, 102, 496, 2569, 13934, 78295, 452439, 2674769
312	1, 2, 5, 15, 52, 201, 843, 3764, 17659, 86245, 435492
321	1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024
1234	1, 2, 5, 14, 40, 113, 319, 918, 2731, 8438, 27011
1243	1, 2, 5, 14, 41, 122, 366, 1108, 3397, 10586, 33618
1324	1, 2, 5, 14, 42, 134, 455, 1640, 6229, 24692, 101205
1342	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
1423	1, 2, 5, 14, 44, 154, 588, 2396, 10237, 45284, 205608
1432	1, 2, 5, 14, 43, 144, 521, 2010, 8156, 34402, 149496
2134	1, 2, 5, 14, 45, 170, 740, 3567, 18408, 99505, 555982
2143	1, 2, 5, 14, 44, 157, 634, 2844, 13829, 71318, 383825
2314	1, 2, 5, 15, 53, 215, 972, 4767, 24837, 135434, 764875
2341	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
2413	1, 2, 5, 15, 52, 201, 842, 3745, 17435, 84119, 417617
2431	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
3124	1, 2, 5, 14, 44, 155, 603, 2541, 11401, 53758, 263847
3142	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
3214	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181, 10946
3241	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
3412	1, 2, 5, 15, 53, 214, 954, 4562, 22929, 119512, 640367
3421	1, 2, 5, 15, 53, 214, 954, 4562, 22929, 119512, 640367
4123	1, 2, 5, 14, 42, 135, 467, 1731, 6803, 28031, 119976
4132	1, 2, 5, 14, 43, 144, 522, 2027, 8334, 35894, 160531
4213	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181, 10946
4231	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796, 58786
4312	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181, 10946
4321	1, 2, 5, 13, 34, 89, 233, 610, 1597, 4181, 10946

σ	**Sequence** $\{f_n^\sigma\}_n$
12345	1, 2, 5, 14, 42, 129, 391, 1158, 3384, 9924
12354	1, 2, 5, 14, 42, 130, 405, 1257, 3883, 11980
12435	1, 2, 5, 14, 42, 130, 405, 1257, 3883, 11980
12453	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
12534	1, 2, 5, 14, 42, 131, 417, 1341, 4335, 14059
12543	1, 2, 5, 14, 42, 131, 417, 1341, 4335, 14059
13245	1, 2, 5, 14, 42, 131, 420, 1388, 4765, 17094
13254	1, 2, 5, 14, 42, 132, 430, 1447, 5032, 18110
13425	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
13452	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
13524	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
13542	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
14235	1, 2, 5, 14, 42, 133, 444, 1566, 5841, 22989
14253	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
14325	1, 2, 5, 14, 42, 132, 432, 1475, 5272, 19756
14352	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
14523	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
14532	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
15234	1, 2, 5, 14, 42, 135, 470, 1773, 7170, 30636
15243	1, 2, 5, 14, 42, 134, 456, 1657, 6403, 26098
15324	1, 2, 5, 14, 42, 134, 456, 1657, 6404, 26117
15342	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
15423	1, 2, 5, 14, 42, 133, 443, 1552, 5717, 22092
15432	1, 2, 5, 14, 42, 133, 443, 1552, 5717, 22092
21345	1, 2, 5, 14, 42, 136, 493, 2043, 9547, 48738
21354	1, 2, 5, 14, 42, 135, 474, 1851, 8061, 38601
21435	1, 2, 5, 14, 42, 135, 474, 1851, 8061, 38601
21453	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
21534	1, 2, 5, 14, 42, 134, 458, 1696, 6857, 30246
21543	1, 2, 5, 14, 42, 134, 458, 1696, 6857, 30246
23145	1, 2, 5, 14, 43, 146, 552, 2316, 10642, 52641
23154	1, 2, 5, 14, 43, 145, 539, 2208, 9896, 47917
23415	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
23451	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
23514	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
23541	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
24135	1, 2, 5, 14, 43, 144, 523, 2045, 8530, 37583
24153	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
24315	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
24351	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796

σ	**Sequence** $\{f_n^\sigma\}_n$
24513	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
24531	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
25134	1, 2, 5, 14, 43, 145, 534, 2119, 8921, 39327
25143	1, 2, 5, 14, 43, 144, 522, 2028, 8352, 36088
25314	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
25341	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
25413	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
25431	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
31245	1, 2, 5, 14, 42, 135, 471, 1793, 7400, 32692
31254	1, 2, 5, 14, 42, 134, 457, 1675, 6602, 27852
31425	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
31452	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
31524	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
31542	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
32145	1, 2, 5, 14, 41, 121, 355, 1032, 2973, 8496
32154	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
32415	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
32451	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
32514	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
32541	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
34125	1, 2, 5, 14, 43, 145, 538, 2188, 9650, 45495
34152	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
34215	1, 2, 5, 14, 43, 145, 538, 2188, 9650, 45495
34251	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
34512	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
34521	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
35124	1, 2, 5, 14, 43, 144, 522, 2027, 8332, 35849
35142	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
35214	1, 2, 5, 14, 43, 144, 522, 2027, 8332, 35849
35241	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
35412	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
35421	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
41235	1, 2, 5, 14, 42, 133, 445, 1578, 5924, 23418
41253	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
41325	1, 2, 5, 14, 42, 134, 456, 1658, 6422, 26314
41352	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
41523	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
41532	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
42135	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842

σ	**Sequence** $\{f_n^\sigma\}_n$
42153	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
42315	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
42351	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
42513	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
42531	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
43125	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
43152	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
43215	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
43251	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
43512	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
43521	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
45123	1, 2, 5, 14, 43, 146, 550, 2279, 10216, 48660
45132	1, 2, 5, 14, 43, 145, 538, 2187, 9628, 45205
45213	1, 2, 5, 14, 43, 145, 538, 2187, 9628, 45205
45231	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
45312	1, 2, 5, 14, 43, 145, 538, 2187, 9628, 45205
45321	1, 2, 5, 14, 43, 145, 538, 2187, 9628, 45205
51234	1, 2, 5, 14, 42, 131, 421, 1403, 4893, 17932
51243	1, 2, 5, 14, 42, 132, 432, 1476, 5288, 19908
51324	1, 2, 5, 14, 42, 132, 432, 1476, 5288, 19908
51342	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
51423	1, 2, 5, 14, 42, 133, 443, 1552, 5718, 22113
51432	1, 2, 5, 14, 42, 133, 443, 1552, 5718, 22113
52134	1, 2, 5, 14, 41, 121, 355, 1033, 2986, 8594
52143	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
52314	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
52341	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
52413	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
52431	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
53124	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
53142	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
53214	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
53241	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
53412	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
53421	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
54123	1, 2, 5, 14, 41, 121, 356, 1044, 3057, 8948
54132	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
54213	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
54231	1, 2, 5, 14, 42, 132, 429, 1430, 4862, 16796
54312	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842
54321	1, 2, 5, 14, 41, 122, 365, 1094, 3281, 9842

Appendix D

Listings

Algorithm 1: Stacksort. Here *Stack* is the stack, $TOP(Stack)$ is the current top of the stack, $\pi = \pi_1 \cdots \pi_n$ is the input permutation.

```
Stack := ∅;
while i ≤ n do
    if Stack = ∅ or π_i < TOP(Stack) then
        execute S;
        i := i + 1;
    end
    else
        execute O;
    end
end
while Stack ≠ ∅ do
    execute O;
end
```

Algorithm 2: The σ-machine. Here $Stack_\sigma$ is the σ-avoiding stack, $Stack_I$ is the increasing stack, P_σ means pushing into $Stack_\sigma$, P_I means pushing into $Stack_I$, O means moving $TOP(Stack_I)$ into the output, $\circ$ is the concatenation operation.

```
Stack_I := ∅;
Stack_σ := ∅;
i := 1;
while i ≤ n do
  | if σ ≰ Stack_σ ∘ π_i then
  |   | execute P_σ;
  |   | i := i + 1;
  | end
  | else if Stack_I = ∅ or TOP(Stack_σ) < TOP(Stack_I) then
  |   | execute P_I;
  | end
  | execute O;
end
while Stack_σ ≠ ∅ do
  | if Stack_I = ∅ or TOP(Stack_σ) < TOP(Stack_I) then
  |   | execute P_I;
  | else
  |   | execute O;
  | end
end
while Stack_I ≠ ∅ do
  | execute O;
end
```