Charles University in Prague
Faculty of Mathematics and Physics

BACHELOR THESIS

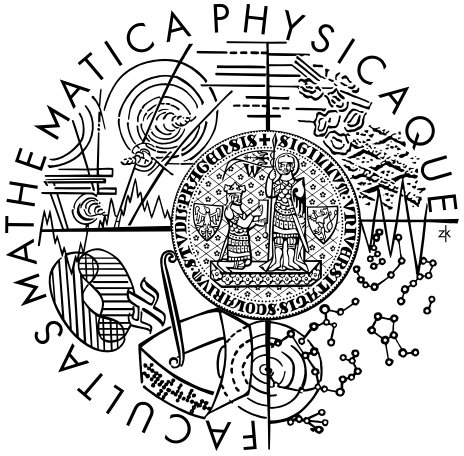

Michal Opler

Pattern-avoiding permutation classes

Computer Science Institute of Charles University

Supervisor of the bachelor thesis: RNDr. Vít Jelínek, Ph.D.
Study programme: Computer Science
Study branch: General Computer Science

Prague 2015

I declare that I carried out this bachelor thesis independently, and only with the cited sources, literature and other professional sources.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University in Prague has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In Prague, May 20, 2015 signature of the author

I would like to thank my supervisor, RNDr. Vít Jelínek, Ph.D., for his guidance and all the valuable advice throughout this study. I would also like to thank my family and friends for supporting me during my studies.

Název práce: Třídy permutací se zakázanými vzory

Autor: Michal Opler

Ústav: Informatický ústav Univerzity Karlovy

Vedoucí bakalářské práce: RNDr. Vít Jelínek, Ph.D., Informatický ústav Univerzity Karlovy

Abstrakt: Major index permutace π je součet všech indexů i takových, že $\pi_i > \pi_{i+1}$. V této práci zkoumáme distribuci major indexu na permutacích neobsahujících zakázané vzory. Zajímá nás hodnota $M_n^m(\Pi)$, což je počet permutací délky n s major indexem m a množinou zakázaných vzorů Π.

Podařilo se nám ukázat, že pro jednoprvkovou množinu $\Pi = \{\sigma\}$ krom okrajových triviálních případů, se hodnoty $M_n^m(\Pi)$ chovají monotónně, nebo-li $M_n^m(\Pi) \leq M_{n+1}^m(\Pi)$. Hlavním výsledkem je rozbor asymptotického chování hodnot $M_n^m(\Pi)$ pro n jdoucí k nekonečnu. Ukážeme, že pro každé pevné m, Π a dostatečně velké n jsou hodnoty $M_n^m(\Pi)$ rovny polynomu v proměnné n a navíc jsme schopni určit stupně těchto polynomů pro různé množiny zakázaných vzorů.

Klíčová slova: permutace, zakázané vzory, major index

Title: Pattern-avoiding permutation classes

Author: Michal Opler

Institute: Computer Science Institute of Charles University

Supervisor: RNDr. Vít Jelínek, Ph.D., Computer Science Institute of Charles University

Abstract: For a permutation π, the major index of π is the sum of all indices i such that $\pi_i > \pi_{i+1}$. In this thesis, we study the distribution of the major index over pattern-avoiding permutations of length n. We focus on the number $M_n^m(\Pi)$ of permutations of length n with major index m and avoiding the set of patterns Π.

First, we are able to show that for a singleton set $\Pi = \{\sigma\}$ other than some trivial cases, the values $M_n^m(\Pi)$ are monotonic in the sense that $M_n^m(\Pi) \leq M_{n+1}^m(\Pi)$. Our main result is a study of the asymptotic behaviour of $M_n^m(\Pi)$ as n goes to infinity. We prove that for every fixed m, Π and n large enough, $M_n^m(\Pi)$ is equal to a polynomial in n and moreover, we are able to determine the degrees of these polynomials for many sets of patterns.

Keywords: permutations, pattern-avoidance, major index

Contents

Introduction

Let $\mathcal{S}_n$ be the set of permutations of the letters $\{1, 2, \ldots, n\} = [n]$. We write a permutation $\pi \in \mathcal{S}_n$ as a sequence $\pi_1 \cdots \pi_n$. A *permutation statistic* is a function $st : \mathcal{S}_n \to \mathbb{N}_0$. For a permutation π, an *inversion* is a pair of different indices $i < j$ such that $\pi_i > \pi_j$ and the *number of inversions* is denoted by $inv(\pi)$. The number of inversions is the oldest and best-known permutation statistic. Already in 1838, Stern [18] proposed a problem of how many inversions there are in all the permutations of length n. The distribution of the number of inversions was given shortly after that by Rodrigues [14].

However, we will focus on a different well-known permutation statistic in this thesis. For a permutation π, we say that there is a *descent* on the i-th position if $\pi_i > \pi_{i+1}$. The *major index* of π, denoted by $maj(\pi)$, is then the sum of the positions, where the descents occur. The major index statistic is younger than the number of inversions, as it was first defined by MacMahon [12] in 1915. Among other results, MacMahon proved its equidistribution with the number of inversions by showing that their generating functions are equal and started the systematic study of permutation statistics in general. That is why we call the statistics equidistributed with the number of inversions Mahonian. Then it took a long time before Foata [10] proved the equidistribution by constructing his famous bijection. Since then many new Mahonian statistics appeared in the literature, most of which are included in the classification given by Babson and Steingrímsson [1]. For the actual values of Mahonian statistics' distribution see the Mahonian numbers sequence A008302 [15].

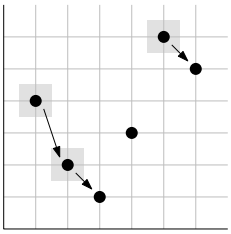

Figure 1: Permutation $\pi = 421365$ represented as a graph. The descents occur at the positions 1, 2 and 5, therefore we have $maj(\pi) = 8$.

We say that two sequences $a_1 \cdots a_n$ and $b_1 \cdots b_n$ are *order-isomorphic* if the permutations required to sort them are the same. A permutation π *contains* a pattern σ if there is a subsequence of $\pi_1 \cdots \pi_n$ order-isomorphic to σ. Otherwise we say that π *avoids* the pattern σ. Pattern avoidance is an active area of research in combinatorics and although the systematic study of pattern avoidance started relatively recently, there is already an extensive amount of literature. A good illustration of an application of pattern avoidance in computer science is that stack-sortable permutations are exactly the ones avoiding pattern 231, which was proved by Knuth [11].

Let $\mathcal{S}_n(\sigma)$ be the set of permutations of length n avoiding σ and $S_n(\sigma)$ its cardinality. We say that patterns σ and τ are *Wilf-equivalent* if $S_n(\sigma) = S_n(\tau)$

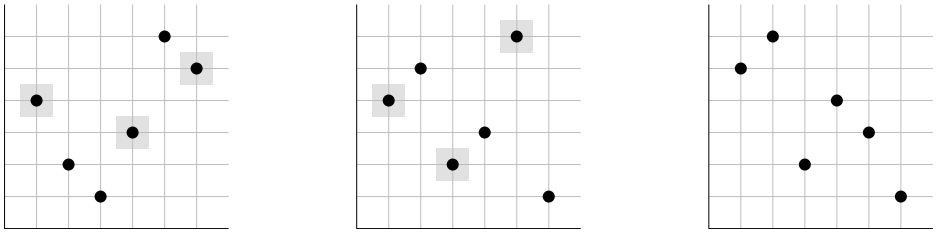

Figure 2: Example illustrating the concept of permutation patterns. The pattern 213 is contained in the permutations 421365 and 452361 (one occurrence in each permutation is highlighted), but not in the permutation 562431.

for every n. For a permutation statistic st, we say that patterns σ and τ are *st-Wilf-equivalent* if there is a bijection between $\mathcal{S}_n(\sigma)$ and $\mathcal{S}_n(\tau)$ which preserves the statistic st. This refinement of Wilf equivalence has been extensively studied for short patterns of length 3, see [3, 4, 8, 13]. An exhaustive classification of Wilf-equivalence and permutation statistics among these patterns was given by Claesson and Kitaev [5]. On the other hand, not much is known about permutation statistics and patterns of length 4 and greater. Recently, Dokos et al. [7] presented an in-depth study of major index and number of inversions including st-Wilf-equivalence. They conjectured maj-Wilf-equivalence between certain patterns of length 4, which was proved by Bloom [2]. Another conjecture from Dokos et al. concerning maj-Wilf-equivalent patterns of a specific form was partly proved by Ge, Yan and Zhang [19].

Claesson, Jelínek and Steingrímsson [6] analysed the inversion number distribution over pattern-avoiding classes. Let $I_n^k(\sigma)$ be the number of σ-avoiding permutations with length n and k inversions. Claesson et al. studied $I_n^k(\sigma)$ for a fixed k and a single pattern σ as a function of n. Our goal is to provide similar analysis for the distribution of major index.

Outline

For a pattern σ, let $\mathcal{M}_n^m(\sigma)$ be the set of σ-avoiding permutations with length n and major index m, and let $M_n^m(\sigma)$ denote its cardinality. For a set of patterns Π, let $\mathcal{M}_n^m(\Pi) = \bigcap_{\sigma\in\Pi} \mathcal{M}_n^m(\sigma)$ and $M_n^m(\Pi)$ its cardinality. Claesson et al. [6] conjectured that $I_n^k(\sigma) \leq I_{n+1}^k(\sigma)$ for every k, n unless σ is an increasing pattern (i.e. a pattern of the form $1\cdots l$). In Chapter 2, we will prove the analogous claim for major index by constructing an injective mapping $f\colon M_n^m(\sigma) \to M_{n+1}^m(\sigma)$ for every $\sigma \neq 12\cdots l$. Furthermore, we show that the claim does not hold in general for an arbitrary set of multiple patterns.

In Chapter 3, we focus on the asymptotic behaviour of $M_n^m(\Pi)$ for a fixed m and Π as n goes to infinity. We note that the asymptotic behaviour for the number of inversions is known only for sets avoiding a single pattern. In contrast, our results apply to general (possibly infinite) set of patterns. It turns out that the values $M_n^m(\Pi)$ are eventually equal to a polynomial in n, which is consistent with the behaviour of $I_n^k(\sigma)$. The natural question to ask is how the degrees of these polynomials depend on Π and m.

Let $deg(m, \Pi)$ be the degree of the polynomial P such that $P(n) = M_n^m(\Pi)$ for $n \geq n_0$. Similarly, let $deg_{\mathrm{I}}(k, \sigma)$ be the degree of the polynomial P such that $P(n) = I_n^k(\sigma)$ for $n \geq n_0$. In the case of the number of inversions, there are just two types of patterns. For a pattern σ, we have either $deg_{\mathrm{I}}(k, \sigma) = k$ for every k, or there is a constant c such that $deg_{\mathrm{I}}(k, \sigma) = \min(k, c)$. All these results about $I_n^k(\sigma)$ and $deg_{\mathrm{I}}(m, \sigma)$ were shown in the aforementioned paper by Claesson et al. [6].

However, the situation gets more complicated when dealing with major index. We show how $deg(m, \{\sigma\})$ depends on the structure of σ and determine $deg(m, \Pi)$ for many types of Π, including all the cases when Π is a singleton set. There are still patterns σ for which $deg(m, \{\sigma\}) = m$, but for many patterns $deg(m, \{\sigma\})$ is a complicated function of m which tends to infinity slower than linearly (approximately as $\sqrt{m}$). Note that there are unfortunately sets Π for which we are not able to precisely determine $deg(m, \Pi)$. In these cases, our results provide at least an upper bound.

Finally, we conclude Chapter 3 by using our results to show that the asymptotic probability of a random permutation with major index m avoiding Π is either 0 or 1. This again corresponds with the number of inversions, where the analogous claim was proved for singleton sets of patterns.

1. Preliminaries

In this chapter, we recall some standard notions related to permutation patterns and introduce a simple decomposition of permutations.

Let $\mathcal{S}_n$ be the set of permutations of the letters $\{1, 2, \ldots, n\} = [n]$. A permutation $\sigma \in \mathcal{S}_n$ will be represented as a sequence of its values $\sigma = \sigma_1\sigma_2\cdots\sigma_n$, where $\sigma_i = \sigma(i)$. We say that two sequences of integers $a_1 \cdots a_k$ and $b_1 \cdots b_k$ are *order-isomorphic* if for every $i, j \in [k]$ we have $a_i < a_j \Leftrightarrow b_i < b_j$. For $I = \{i_1 < i_2 < \cdots < i_k\} \subseteq [n]$ and $\pi \in \mathcal{S}_n$, let $\pi[I]$ denote the permutation in $\mathcal{S}_k$ which is order-isomorphic to the sequence $\pi_{i_1}\pi_{i_2}\cdots\pi_{i_k}$. A permutation $\pi \in \mathcal{S}_n$ *contains* a permutation $\sigma \in \mathcal{S}_k$ if there exists an I such that $\pi[I] = \sigma$. We write $\sigma \preceq \pi$ to denote this. If π does not contain σ we say that π *avoids* σ. In this context we usually call σ a *pattern*. Similarly, for a set of patterns Π we say that a permutation τ is Π-avoiding if it is σ-avoiding for every $\sigma \in \Pi$. For a pattern σ let $\mathcal{S}_n(\sigma)$ be the set of all σ-avoiding permutations of length n, and $S_n(\sigma)$ its cardinality. More generally for a set of permutations Π, let $\mathcal{S}_n(\Pi)$ denote the set of all Π-avoiding permutations of length n, and $S_n(\Pi)$ its cardinality.

The *descent set* of $\sigma \in \mathcal{S}_n$ is the set $\mathcal{D}(\sigma) = \{i \mid \sigma_i > \sigma_{i+1}\}$ and the *major index* is the sum of all its members $maj(\sigma) = \sum_{i \in \mathcal{D}(\sigma)} i$. We will consider the distribution of major index over pattern-avoiding permutations.

Definition 1.1. Let $\mathcal{M}_n^m(\sigma)$ denote the set of all σ-avoiding permutations of length n with major index m, and $M_n^m(\sigma)$ its cardinality. Similarly let $\mathcal{M}_n^m(\Pi)$ be the set of all permutations from $\mathcal{S}_n(\Pi)$ with major index m, and $M_n^m(\Pi)$ its cardinality. For an example of the values $M_n^m(\sigma)$ for a specific pattern, see Table 1.1.

1														
1	1													
1	2	2	1											
1	3	4	6	5	3	1								
1	4	6	12	16	19	16	15	9	4	1				
1	5	8	19	29	45	58	65	73	65	57	39	29	14	
1	6	10	27	44	76	119	164	212	260	287	299	303	265	...

Table 1.1: The number of 1324-avoiding permutations with a fixed major index. The mth entry in the nth row is the value $M_n^m(1324)$, with n starting at 1 and m starting at 0.

1.1 Permutation decomposition

Now we will introduce a decomposition which will later prove to be very useful. Let $\mathbb{N}_0^d$ be the set of d-tuples of non-negative integers and for every $a \in \mathbb{N}_0^d$ define its size $|a| = \sum_{i=1}^d a_i$. We will decompose an arbitrary permutation into a smaller

permutation and a tuple of non-negative integers. Let $\pi \in \mathcal{S}_n$ be a permutation and k a natural number, such that the sequence $\pi_{k+1} \cdots \pi_n$ is strictly increasing. Then we can store the structure of such permutation in a shorter permutation σ order-isomorphic to $\pi_1 \cdots \pi_k$, and a $(k+1)$-tuple which describes the vertical gaps between the letters $\pi_1 \cdots \pi_k$.

Definition 1.2. Let $\pi \in \mathcal{S}_n$ be a permutation and $k \in [n]$ such that the sequence $\pi_{k+1} \cdots \pi_n$ is strictly increasing. Let σ be the permutation order-isomorphic to the sequence $\pi_1 \cdots \pi_k$ and $a \in \mathbb{N}_0^{k+1}$ the only $(k+1)$-tuple of size $|a| = n - k$ such that $\pi_i = \sigma_i + \sum_{j=1}^{\sigma_i} a_j$ holds for every $i \in [k]$. Then we say that π can be *decomposed* into σ and a, denoted by $\pi = \sigma \cdot a$.

We can also look at the decomposition from the other side as an operation, which increases the vertical gaps between the letters of σ and then fills them with increasing suffix. See Figure 1.1.

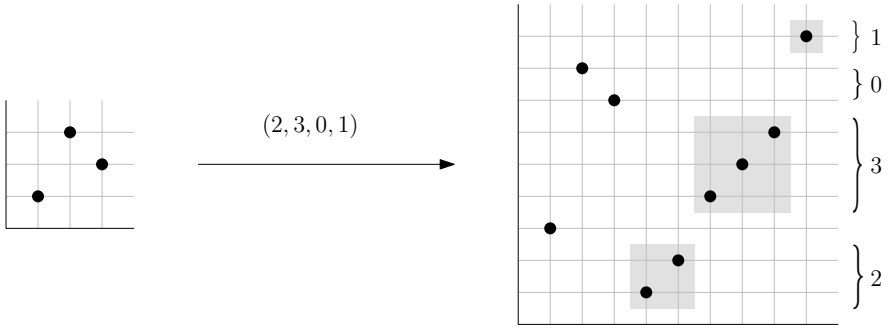

Figure 1.1: For a permutation $\sigma = 132$ and 4-tuple $a = (2, 3, 0, 1)$ we have $\pi = \sigma \cdot a = 387124569$.

Definition 1.3. For a permutation π that can be expressed as $\pi = \gamma \cdot a$ for some $\gamma \in \mathcal{S}_k$ and $a \in \mathbb{N}_0^{k+1}$, we call γ *the core of* π and a *the padding profile of* π if k is the last descent of π. In other words, $\pi = \gamma \cdot a$ is a decomposition into a core and a padding profile of π if there is $i \leq \gamma_k$ such that $a_i > 0$. For $\pi = 12 \cdots n$, the core of π is the empty permutation and its padding profile is $a \in \mathbb{N}_0$ equal to the length of π.

Observe that the major index of a permutation π is determined only by its core. Therefore, let us define the following statistic which characterizes the cores of permutations with a given major index.

Definition 1.4. For a permutation π, let the *extended major index of* π, denoted by $maj^+(\pi)$, be the sum of its major index and its length, i.e.,

$$maj^+(\pi) = |\pi| + maj(\pi).$$

For every permutation π with a core γ, we have $maj(\pi) = maj^+(\gamma)$. Notice also that for any π, if π contains σ then $maj^+(\pi) \geq maj^+(\sigma)$.

2. Monotonicity of columns

In this chapter, we will focus on the distribution of major index over permutations avoiding a single pattern. Observe that each column of Table 1.1 is weakly increasing from top to bottom. In other words, for a fixed major index m the number of 1324-avoiding permutations of length $n+1$ is at least the number of 1324-avoiding permutations of length n. We will show that this claim holds in general for any single pattern σ except for the increasing patterns (i.e., the patterns of the form $12\cdots k$).

First let us define a simple operation of inserting an element into a permutation. Later we will prove two elementary properties of this operation.

Definition 2.1. For a permutation $\pi \in \mathcal{S}_n$ and $k, l \in [n+1]$, let $\pi[k \to l] \in \mathcal{S}_{n+1}$ be a permutation created by inserting the letter l at the k-th position. In other words $\pi[k \to l]$ is the permutation order-isomorphic to the sequence $\pi_1 \cdots \pi_{k-1} \left(l - \frac{1}{2}\right) \pi_k \cdots \pi_n$. See Figure 2.1.

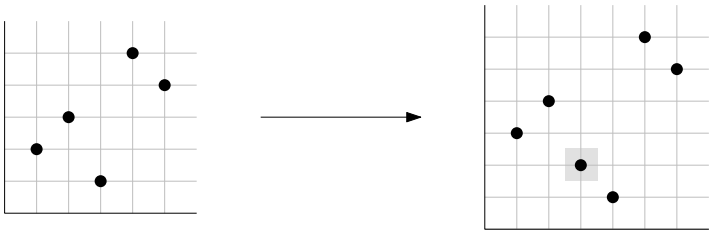

Figure 2.1: Example of an insertion $23154[3 \to 2] = 342165$

Lemma 2.2. *Let $n \in \mathbb{N}$, $k, l \in [n+1]$ and $\pi \in \mathcal{S}_n(\sigma)$. If there is I such that $\pi[k \to l][I] = \sigma$, then $k \in I$.*

Proof. Suppose that we have $I = \{i_1 < \cdots < i_m\}$ such that $k \notin I$. Let $J = \{j_1 < \cdots < j_m\}$ be a subset of $[n]$ defined by

$$j_t = \begin{cases} i_t & \text{if } i_t < k \\ i_t - 1 & \text{if } i_t > k\,. \end{cases}$$

Since $\pi[k \to l]$ restricted to indices other than k is order-isomorphic to π we obtain $\sigma = \pi[k \to l][I] = \pi[J]$, contradicting the fact that π avoids σ. □

Lemma 2.3. *Let $n \in \mathbb{N}$, $k, l \in [n]$ and $\pi \in \mathcal{S}_n$. If $\mathcal{D}(\pi) \subseteq [k-1]$, $l \leq \pi_k$ and either $k = 1$ or the sequences $\pi_{k-1}(l - \frac{1}{2})$ and $\pi_{k-1}\pi_k$ are order-isomorphic, then $\mathcal{D}(\pi) = \mathcal{D}(\pi[k \to l])$.*

Proof. As before $\pi[k \to l]$ restricted to indices other than k is order-isomorphic to π. Therefore for every index $i < k-1$, $i \in \mathcal{D}(\pi)$ if and only if $i \in \mathcal{D}(\pi[k \to l])$. And since we know that all the elements of $\mathcal{D}(\pi)$ are smaller than k, we get $\mathcal{D}(\pi[k \to l]) \subseteq [k]$. We are left with the two indices k and $k-1$. Observe that $k \notin \mathcal{D}(\pi[k \to l])$ because $l \leq \pi_k$. And from the last condition we obtain $k-1 \in \mathcal{D}(\pi)$ if and only if $k-1 \in \mathcal{D}(\pi[k \to l])$. □

Theorem 2.4. *For every $n, m, k \in \mathbb{N}$ and $\sigma \in \mathcal{S}_k$ with $\mathcal{D}(\sigma) \neq \emptyset$ we have the inequality $M_n^m(\sigma) \leq M_{n+1}^m(\sigma)$.*

Proof. To prove this theorem we will construct an injective mapping f from $\mathcal{M}_n^m(\sigma)$ to $\mathcal{M}_{n+1}^m(\sigma)$. In order to find an image for $\pi \in \mathcal{M}_n^m(\sigma)$ we introduce the following permutation statistics.

Definition 2.5. For $\sigma \in \mathcal{S}_n$ let $tail(\sigma)$ denote the largest i such that σ_{n+1-i} $\sigma_{n+2-i} \cdots \sigma_n$ are all fixed points. And similarly let $slope(\sigma)$ be the largest i such that the sequence σ_{n+1-i} $\sigma_{n+2-i} \cdots \sigma_n$ is strictly increasing. Recalling Definition 1.3, we see that $slope(\sigma)$ is the size of the padding profile of σ and $tail(\sigma)$ is the value of its last coordinate. See Figure 2.2.

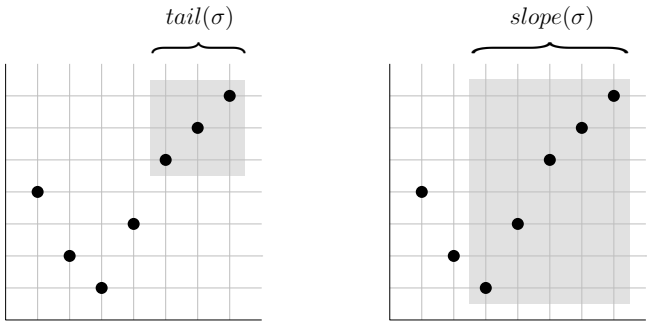

Figure 2.2: The tail and slope statistics of the permutation $\sigma = 4213567$

Case 1. First we solve the easy case where $tail(\sigma) = 0$. We simply extend π by inserting the letter $n + 1$ at the end, i.e.,

$$f(\pi) = \pi[n + 1 \to n + 1].$$

It is clear that f preserves the descent set, which implies $maj(\pi) = maj(f(\pi))$. Now suppose there is $I = \{i_1 < \cdots < i_k\}$ such that $f(\pi)[I] = \sigma$. Lemma 2.2 implies $i_k = n + 1$. But that would lead to $\sigma_k = f(\pi)[I]_k = k$ which contradicts the assumption that $tail(\sigma) = 0$.

Case 2. Suppose now that $tail(\sigma) \neq 0$ and $slope(\pi) \geq tail(\sigma)$. Then we create the image of π by expanding the element at the position $n+1-tail(\sigma)$ into two. See Figure 2.3.

$$f(\pi) = \pi[t \to \pi_t] \quad \text{where } t = n + 1 - tail(\sigma).$$

Because all the conditions from Lemma 2.3 are met, we get $\mathcal{D}(\pi) = \mathcal{D}(f(\pi))$ which implies $maj(\pi) = maj(f(\pi))$.

Next we want to show that $f(\pi)$ avoids σ. Suppose there is $I = \{i_1 < \cdots < i_k\}$ such that $f(\pi)[I] = \sigma$. Again from Lemma 2.2 we obtain $t = i_j \in I$ for some j. Observe that since there are only $tail(\sigma)$ indices in $f(\pi)$ larger than t, we get a lower bound $j \geq k - tail(\sigma)$.

Now we will use different arguments depending on whether this holds as an equality or not. First suppose that $j = k - tail(\sigma)$. This means that I also

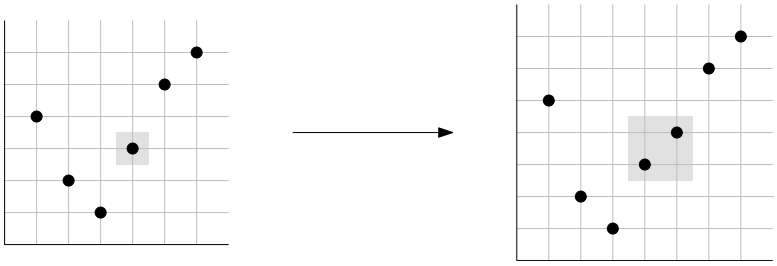

Figure 2.3: Example of a construction from Case 2. Consider a permutation $\pi = 421356$, which has $slope(\pi) = 4$, and suppose we have a pattern σ with $tail(\sigma) = 3$. Then $f(\pi) = 5213467$.

contains all the indices larger than j, in particular $i_{j+1} = t + 1$. Following Definition 2.5, j is then the largest index such that $\sigma_j \neq j$, implying $\sigma_j < j$. This means there is a letter σ_l to the left of σ_j such that $\sigma_l > \sigma_j$ and all the letters to the right of σ_j are larger than σ_l. Therefore, looking at the indices j, $j+1$ and l we have $\sigma_j < \sigma_l < \sigma_{j+1}$ and the same inequality goes for $f(\pi)[I]$. Translated to the indices of $f(\pi)$ the inequality $f(\pi)_t < f(\pi)_{i_l} < f(\pi)_{t+1}$ must hold. But recalling the definition of $f(\pi)$ there is no index p such that $f(\pi)_t < f(\pi)_p < f(\pi)_{t+1}$.

Suppose now that $j > k - tail(\sigma)$. In this case there must be $l > j$ such that $l \notin I$. We aim to show that $J = \{j_1 < \cdots < j_k\} = I \setminus \{t\} \cup \{l\}$ satisfies $f(\pi)[J] = f(\pi)[I] = \sigma$, which would lead to a contradiction since π contains $f(\pi)[J]$.

We know that $\sigma_j \sigma_{j+1} \cdots \sigma_k$ are all fixed points following Definition 2.5. Observe that for every index p we have the inequality $j_p \geq i_p$ with equality on the indices smaller than j. Therefore, $f(\pi)[I]$ and $f(\pi)[J]$ restricted to the first $j - 1$ letters are order-isomorphic. The only thing left is to check that the other letters of $f(\pi)[J]$ are fixed points. The sequence $f(\pi)_t \cdots f(\pi)_{n+1}$ is increasing, thus its subsequence $f(\pi)_{j_j} f(\pi)_{j_{j+1}} \cdots f(\pi)_{j_k}$ is also increasing and moreover $f(\pi)_{j_j} > f(\pi)_{i_j}$. Then $f[J]_j$ and subsequently all the succeeding letters of $f[J]$ must be fixed points. Together, this means that indeed $f(\pi)[I] = f(\pi)[J]$.

Case 3. Finally, suppose that $tail(\sigma) \neq 0$ and $slope(\pi) < tail(\sigma)$. Then we simply insert the letter 1 at the rightmost possible position without creating a new descent. See Figure 2.4.

$$f(\pi) = \pi[n + 1 - slope(\pi) \to 1].$$

As before, we obtain $maj(\pi) = maj(f(\pi))$ from Lemma 2.3. If there is $I = \{i_1 < \cdots < i_k\}$ such that $f(\pi)[I] = \sigma$, then Lemma 2.2 implies $n+1-slope(\pi) = i_j$ for some j. The j-th letter of σ must be its minimum since $f(\pi)_{i_j} = 1$ is the minimum of $f(\pi)$. On the other hand, because $n+1-slope(\pi) > n+1-tail(\sigma)$ and $\mathcal{D}(\sigma) \neq \emptyset$, there must be q such that $\sigma_q < \sigma_j$, which yields a contradiction.

The only remaining part is to show that f is injective. Suppose there are $\pi_1 \neq \pi_2$ such that $f(\pi_1) = f(\pi_2)$. From the properties of $f(\pi_1)$ we can tell unambiguously whether it was obtained through Case 1, 2 or 3. And following the definitions of f in these particular cases it is clear that necessarily $\pi_1 = \pi_2$. $\square$

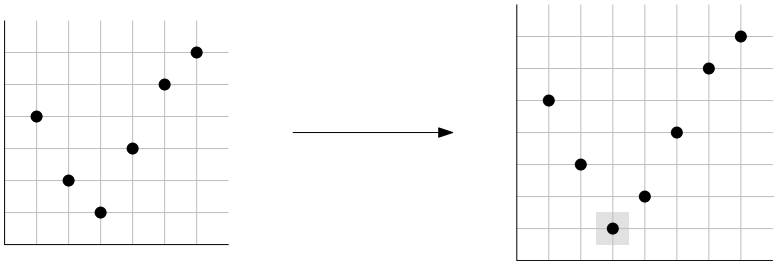

Figure 2.4: Example of a construction from Case 3. Consider a permutation $\pi = 421356$, which has $slope(\pi) = 4$, and suppose we have a pattern σ with $tail(\sigma) > 4$. Then $f(\pi) = 5312467$.

In Theorem 2.4, the assumption $\mathcal{D}(\sigma) \neq \emptyset$ is necessary, because in the case of a pattern $\sigma = 12\cdots k$ and fixed $m \in \mathbb{N}$ there is $n_0 \in \mathbb{N}$ such that for every n larger than n_0 we have $M_n^m(\sigma) = 0$. This follows directly from the Erdős–Szekeres theorem [9], which states that any permutation of length $n > m(k-1)+1$ contains either the increasing pattern of length k or the decreasing pattern of length $m+1$, forcing the major index to be larger than m.

Applying a similar argument as in the proof of Theorem 2.4, we could show that $M_n^m(\Pi) \leq M_{n+1}^m(\Pi)$ for any set of patterns with the same tail which does not contain any increasing pattern. One could think that indeed for any set of patterns the columns are either eventually zero or weakly increasing. But this is not true even for small sets of short patterns. For example, consider a set $\Pi = \{3412, 1324\}$ of just two patterns. In this case $M_6^5(\Pi) = M_8^5(\Pi) = 21$ and $M_7^5(\Pi) = 20$ (see Table 2.1). But we can easily show that $M_n^m(\Pi)$ tend to infinity for $m \geq 3$. Let $\pi(n) = 12\cdots(n-2)[m-1 \to 1]$ and $\pi(n,k) = \pi(n)[1 \to k]$, then $\pi(n,k) \in M_n^m(\Pi)$ for $n \geq m$ and $k > 2$. Therefore, $M_n^m(\Pi)$ tends to infinity because there are linearly many $\pi(n,k)$ for a fixed n.

1															
1	1														
1	2	2	1												
1	3	3	6	5	3	1									
1	4	3	9	12	16	12	15	9	4	1					
1	5	3	13	12	21	38	31	48	41	44	29	29	14	5	...
1	6	3	18	13	20	49	62	63	105	95	109	162	108	135	...
1	7	3	24	14	21	62	62	105	105	221	169	222	335	341	...

Table 2.1: The number of permutations avoiding $\Pi = \{3412, 1324\}$ with a fixed major index. The mth entry in the nth row is the value $M_n^m(\Pi)$, with n starting at 1 and m starting at 0. The problematic values are highlighted.

3. Asymptotic behaviour

We have seen that for most single patterns the inequality $M_n^m(\sigma) \leq M_{n+1}^m(\sigma)$ holds (recall Theorem 2.4). Let us now focus on the asymptotic behaviour of $M_n^m(\sigma)$ for a fixed m as n tends to infinity. More generally, we are interested in the asymptotic behaviour of $M_n^m(\Pi)$ for a (possibly infinite) set of permutations Π. Our first goal is to show that for a fixed m and arbitrary Π, $M_n^m(\Pi)$ is eventually equal to a polynomial.

3.1 Polynomial growth

By recalling Definition 1.3, observe that a permutation is uniquely determined by its core and its padding profile while its major index is determined only by the core. Furthermore, for any permutation $\tau \in \mathcal{M}_n^m(\Pi)$ all the elements of $\mathcal{D}(\tau)$ are smaller than $m+1$, thus making the core of any such permutation shorter than $m+1$. This means that all the permutations with major index m have only a finite number of unique cores.

Definition 3.1. Let $C(m, \Pi)$ denote the finite set of all the distinct cores of permutations from $\mathcal{M}^m(\Pi)$, where $\mathcal{M}^m(\Pi)$ the set of all Π-avoiding permutations with major index m, i.e. $\mathcal{M}^m(\Pi) = \bigcup_{n\geq 1} \mathcal{M}_n^m(\Pi)$.

Note that every core $\gamma \in C(m, \Pi)$ satisfies $maj^+(\gamma) = m$ (recall Definition 1.4). Therefore the permutation statistic which assigns to each permutation its core, is in fact a refinement of the major index.

Definition 3.2. For $\gamma \in C(m, \Pi)$, let $\mathcal{M}_n^{[\gamma]}(\Pi)$ be the set of permutations from $\mathcal{M}_n^m(\Pi)$ which have the core γ, and let $M_n^{[\gamma]}(\Pi)$ be its cardinality.

Clearly $M_n^m(\Pi) = \sum_{\gamma \in C(m,\Pi)} M_n^{[\gamma]}(\Pi)$. This means that in order to prove the polynomial behaviour of $M_n^m(\Pi)$ for a fixed m, it is enough to prove the polynomial behaviour of $M_n^{[\gamma]}(\Pi)$ for a fixed core γ. And because the decomposition of a permutation into its core and its padding profile is unique, we can enumerate $\mathcal{M}_n^{[\gamma]}(\Pi)$ by counting all the possible padding profiles.

Lemma 3.3. *Let Π be any set of permutations and $\gamma \in \mathcal{S}_k$ a permutation. Then there exists a polynomial P and an integer n_0 such that for every $n \geq n_0$, $M_n^{[\gamma]}(\Pi) = P(n)$.*

In order to prove the lemma, we reduce it to a known property of down-sets of integer compositions. Define a partial order $\leq$ on $\mathbb{N}_0^d$ as $(a_1, \ldots, a_d) \leq (b_1, \ldots, b_d)$ if for every $i \in [d]$ we have $a_i \leq b_i$. A set $A \subseteq \mathbb{N}_0^d$ is a *down-set* in $\mathbb{N}_0^d$ if for every $a \in A$ and $b \leq a$, b belongs to A as well. The following proposition was proved by Stanley [17, 16].

Proposition 3.4 (Stanley [17]). *Let d be a positive integer and let S be a down-set in $\mathbb{N}_0^d$. Let $H(n)$ be the number of elements of S with size n. Then there exists a polynomial P and an integer n_0 such that for every $n \geq n_0$, $H(n) = P(n)$.*

Proof (Stanley [16]). Let M be the set of minimal elements of the subset $\mathbb{N}_0^d \setminus S$. Then M is an anti-chain (meaning that if $a, b \in M$ and $a \neq b$ then neither $a \leq b$ nor $b \leq a$). First, we will show that any anti-chain in $\mathbb{N}_0^d$ is finite by induction on d. Note that the assertion clearly holds for $d = 1$. Now let $d > 1$ and assume that K is an infinite anti-chain in $\mathbb{N}_0^d$. Fix $a = (a_1, \ldots, a_d) \in K$ and put $K' = K \setminus \{a\}$, $T = \{(i, j) \mid 1 \leq i \leq d, 0 \leq j < a_i\}$. Define a map $f\colon K' \to T$ by $f(b) = (i, j)$ if i is the smallest index such that $b_i < a_i$ and $j = b_i$. At least one of the sets $K'_{ij} = f^{-1}(\{(i, j)\})$ is infinite, say K'_{rk}. Then by suppressing the rth coordinate from each member of K'_{rk} we obtain an infinite anti-chain in $\mathbb{N}_0^{d-1}$, contradicting the induction hypothesis.

For any finite subset $L \subset \mathbb{N}_0^d$, let $\sup(L)$ be the least upper bound of L in $\mathbb{N}_0^d$. Explicitly, we have $\sup(L) = (b_1, \ldots, b_d)$ where $b_i = \max_{a \in L} a_i$, $a = (a_1, \ldots, a_d)$.

For $a \in \mathbb{N}_0^d$, $n \in \mathbb{N}$ let $J(a, n)$ be the number of $b \in \mathbb{N}_0^d$ such that $a \leq b$ and $|b| = n$. By the inclusion-exclusion principle we can express $H(n)$ as a finite sum

$$H(n) = \sum_{k=0}^{r} \sum_{L \in M_k} (-1)^k \, J(\sup(L), n)$$

where M_k is the set of k-subsets of M and $r = |M|$. It remains to note that $J(a, n)$ is a polynomial in n for n sufficiently large. For $s = |a|$, we have

$$J(a, n) = \begin{cases} \binom{n+d-s-1}{d-1} & \text{if} \quad n \geq s \\ 0 & \text{if} \quad 0 \leq n < s. \end{cases}$$

And $\binom{n+d-s-1}{d-1}$ is indeed a polynomial in n for $n \geq s$. □

Proof of Lemma 3.3. Unfortunately the set of all padding profiles from $\mathcal{M}_n^{[\gamma]}(\Pi)$ is not a down-set, but we can express it as a difference of two down-sets. Define the following sets

$$A_n = \{a \mid a \in \mathbb{N}_0^{k+1} \wedge \gamma \cdot a \in \mathcal{S}_n(\Pi)\} \quad \text{and} \quad A = \bigcup\nolimits_{n \geq 0} A_n$$
$$B_n = \{a \mid a \in A_n \wedge \forall i \leq \gamma_k \colon a_i = 0\} \quad \text{and} \quad B = \bigcup\nolimits_{n \geq 0} B_n.$$

Let us check that both A and B are down-sets in $\mathbb{N}_0^{k+1}$. If a belongs to A and $b \leq a$, then the permutation $\gamma \cdot a$ contains the permutation $\gamma \cdot b$ and therefore $\gamma \cdot b$ must be Π-avoiding and b belongs to A. To show that B is down-set, consider $a \in B$ and $b \leq a$. We know from previous argument that b also belongs to A and the second condition holds since for every $i \in [\gamma_k]$ we have $b_i \leq a_i = 0$ implying $b_i = 0$.

The padding profiles of permutations from $\mathcal{M}_n^{[\gamma]}(\Pi)$ have at least one of the first γ_k values positive, because such permutation has a descent at the k-th position.

But these are exactly the tuples which belong to A_n but not to B_n. Since B_n is a subset of A_n we get $M_n^{[\gamma]}(\Pi) = |A_n| - |B_n|$.

From Proposition 3.4, we obtain that $|A_n|$ and $|B_n|$ are both polynomials for sufficiently large n, therefore $M_n^{[\gamma]}(\Pi)$ is eventually equal to a polynomial as well. □

Corollary 3.5. *For a set of permutations* Π *and* $m \in \mathbb{N}_0$*, there exists a polynomial* P *and an integer* n_0 *such that for every* $n \geq n_0$*,* $M_n^m(\Pi) = P(n)$*.*

Since we now know that the numbers $M_n^m(\Pi)$ are eventually equal to a polynomial, we can introduce the following notation.

Definition 3.6. For a set of permutations Π, let $deg(m, \Pi)$ be the degree of the polynomial P such that $M_n^m(\Pi) = P(n)$ for n large enough. For a zero polynomial P, let $deg(m, \Pi) = 0$.

Observe that for an arbitrary set of permutations Π and $\Omega \subseteq \Pi$, it follows that $deg(m, \Pi) \leq deg(m, \Omega)$. This holds since any Π-avoiding permutation is trivially Ω-avoiding too.

Now we would like to know how these degrees depend on m and on the structure of permutations in Π. It turns out that there is one important statistic of patterns which affects the degree $deg(m, \Pi)$.

Definition 3.7. For a permutation π we will define *the magnitude of* π as

$$mg(\pi) = \begin{cases} 0 & \text{if } \mathcal{D}(\pi) = \emptyset \\ k & \text{if } \mathcal{D}(\pi) = \{k\} \\ +\infty & \text{if } |\mathcal{D}(\pi)| \geq 2. \end{cases}$$

For a set of permutations Π *the magnitude of* Π, denoted by $mg(\Pi)$, is the minimal magnitude of a permutation $\sigma \in \Pi$. For the empty set of patterns, $mg(\emptyset) = +\infty$.

Let us make an important observation about magnitude. If a permutation π contains a pattern σ then necessarily $mg(\pi) \geq mg(\sigma)$.

3.2 Sets of infinite magnitude

As we will show, the magnitude of Π plays a key role in determining the value of $deg(m, \Pi)$. To prove this, let us first focus on the sets Π of infinite magnitude. In this particular case, we can also determine the leading coefficient of the polynomial $M_n^m(\Pi)$, which will prove to be useful later.

Proposition 3.8. *Let* Π *be a set of permutations with* $mg(\Pi) = +\infty$*. Then* $deg(m, \Pi) = m$ *and* $M_n^m(\Pi) = \frac{n^m}{m!} + O(n^{m-1})$ *as* $n \to \infty$*.*

Proof. First observe that for $m = 0$ the proposition simply states that $M_n^0(\Pi) = 1$ for $n \geq n_0$. But that is clear since $\mathcal{M}_n^0(\Pi) = \mathcal{M}_n^0(\emptyset) = \{12 \cdots n\}$. Therefore, in the rest of the proof suppose that $m \geq 1$.

In order to prove our proposition, it is sufficient to prove the following claims.

1. For the core $\epsilon = 12\cdots m$ we have $M_n^{[\epsilon]}(\Pi) = \frac{n^m}{m!} + O(n^{m-1})$.
2. For every $\gamma \in C(m,\Pi) \setminus \{\epsilon\}$ there is a constant $\beta = \beta(\gamma, m, \Pi)$ such that $M_n^{[\gamma]}(\Pi) \leq \beta n^{m-1}$.

To prove the first claim, simply observe that any permutation with the core ϵ has a finite magnitude, which makes it Π-avoiding. By choosing the first m letters we uniquely get every permutation with the core ϵ plus the permutation $12\cdots n$. That gives us the desired enumeration $M_n^{[\epsilon]}(\Pi) = \binom{n}{m} - 1 = \frac{n^m}{m!} + O(n^{m-1})$.

To prove the second claim, fix a core $\gamma \in C(m,\Pi) \setminus \{\epsilon\}$ of length k. First observe that for $\gamma \neq \epsilon$ necessarily $k \leq m-1$. We will bound $M_n^{[\gamma]}(\Pi)$ from above by the number of all the permutations of length n which can be expressed as $\gamma \cdot a$ for some tuple a. This yields the upper bound $M_n^{[\gamma]}(\Pi) \leq \binom{n}{k} \leq \binom{n}{m-1} \leq \beta n^{m-1}$ for some β.

These claims together with Corollary 3.5 give the desired polynomial behaviour. □

3.3 Sets of finite magnitude

Let us now focus on the problem of determining $deg(m,\Pi)$ for a set Π of finite magnitude. As we will show in this section, the asymptotic behaviour of these sets is far more complicated than that of sets with infinite magnitude. Our main result is providing the values $deg(m,\Pi)$ as a function of m. As in Proposition 3.8, we will construct a suitable core and bound $deg(m,\Pi)$ from below by counting all the possible padding profiles. On the other hand, we will use a different approach for obtaining the upper bound. For a fixed core γ, we will bound $M_n^{[\gamma]}(\Pi)$ in terms of how many coordinates of a padding profile a can be large if $\gamma \cdot a$ avoids Π.

Lemma 3.9. *Let Π be a finite set of permutations and let m and l be non-negative integers. If every permutation π with $maj^+(\pi) \leq m$ and length $|\pi| > l$ contains a core of some permutation in Π, then $deg(m,\Pi) \leq l$.*

Proof. Let k be the length of the longest permutation in Π. We will prove the claim by showing that for every $\gamma \in C(m,\Pi)$ there is a constant $\alpha = \alpha(\gamma, m, \Pi)$ such that $M_n^{[\gamma]}(\Pi) \leq \alpha n^l$.

Fix a core $\gamma \in C(m,\Pi)$. For a padding profile $a \in \mathbb{N}_0^d$ we will say that its coordinate a_i is *bad* if $a_i > k$. We claim that every permutation in $M_n^{[\gamma]}(\Pi)$ has a padding profile with at most $l+1$ bad coordinates. Suppose for a contradiction that there is a permutation $\pi \in \mathcal{M}_n^{[\gamma]}(\Pi)$ with at least $l+2$ bad coordinates in its padding profile. Let ψ be the permutation order-isomorphic to $l+1$ elements from the core of π which separate the $l+2$ bad coordinates. Because ψ is contained in the core it satisfies $maj^+(\psi) \leq m$. But since it has length greater than l it must contain a core κ of some permutation $\sigma \in \Pi$. Furthermore, let $p \in \mathbb{N}_0^{l+2}$ be the tuple of only the $l+2$ bad coordinates from the padding profile of π. Observe that since ψ contains κ and every coordinate of p is larger than $|\sigma|$ then $\psi \cdot p$ must contain σ. But that is clearly a contradiction because $\psi \cdot p$ is contained in π.

Now it suffices to show that the number of permutations with core γ and at most $l+1$ bad coordinates is smaller than αn^l for some α. Let d be the length of the core γ. First we have $\binom{d+1}{l+1}$ ways to choose the $l+1$ potentially bad coordinates. We have only constantly many options for the remaining $d-l$ coordinates of the padding profile, which we can bound with k^{d-l}. And finally, we will bound the number of options how to split the remaining elements into the bad coordinates by enumerating the number of ways to split n elements into $l+1$ boxes which is $\binom{n+l}{l}$.

This yields the upper bound $M_n^{[\gamma]}(\Pi) \leq k^{d-l}\binom{d+1}{l+1}\binom{n+l}{l}$ and since the only non-constant factor is $\binom{n+l}{l}$, this indeed implies $M_n^{[\gamma]}(\Pi) \leq \alpha n^l$ for some α. □

By combining Lemma 3.9 with the Erdős–Szekeres theorem [9], we obtain a precise characterization of the sets Π for which the degrees $deg(m, \Pi)$ are bounded by a constant independent of m. This illustrates that sets of patterns containing permutations with both finite and infinite magnitude can behave very miscellaneously.

Proposition 3.10. *For a set of permutations Π, $deg(m, \Pi)$ is bounded by a constant independent of m, if and only if there is $\sigma \in \Pi$ with the core $12\cdots k$ and $\tau \in \Pi$ with the core $l(l-1)\cdots 1$ for some k and l.*

Proof. To prove one implication, assume that Π contains such σ and τ. We know that $deg(m, \Pi) \leq deg(m, \{\sigma, \tau\})$. From the Erdős–Szekeres theorem [9], it follows that every permutation longer than $(l-1)(k-1)$ contains either $12\cdots k$ or $l(l-1)\cdots 1$. Therefore, we obtain the inequality $deg(m, \Pi) \leq (k-1)(l-1)$ from Lemma 3.9.

We will prove the other implication by proving its contrapositive. Assume that Π does not contain any permutation with an increasing core. In other words $mg(\Pi) = +\infty$ and Proposition 3.8 implies that $deg(m, \Pi) = m$. Therefore, $deg(m, \Pi)$ is unbounded.

Finally, assume that Π does not contain any permutation with a decreasing core. In this case we cannot precisely express $deg(m, \Pi)$. However, if $m = \frac{d^2+d}{2}$ for some integer d then every permutation with the core $d(d-1)\cdots 1$ is Π-avoiding and has major index m. Since there are $\binom{n-1}{d}$ such permutations of length n, we get the inequality $deg(m, \Pi) \geq d$. And again $deg(m, \Pi)$ is unbounded. □

Now we will focus on determining the values $deg(m, \Pi)$ for sets Π of finite magnitude. As we already discussed in Chapter 2, for any set of permutations Π with magnitude 0, we have $M_n^m(\Pi) = 0$ for $n \geq n_0$. It is not hard to show that the values eventually get constant in the case of sets with magnitude 1.

Proposition 3.11. *If Π is a set of permutations with magnitude $mg(\Pi) = 1$, then $deg(m, \Pi) = 0$.*

Proof. We know that $deg(m, \Pi) \geq 0$ for every m and Π. Therefore, it is sufficient to bound $deg(m, \Pi)$ from above. Fix a permutation $\tau \in \Pi$ with magnitude 1. Since $deg(m, \Pi) \leq deg(m, \{\tau\})$ and every non-empty permutation contains the pattern 1, we get $deg(m, \Pi) \leq 0$ directly from Lemma 3.9. □

The next result determines $deg(m, \Pi)$ for all sets of permutations of magnitude at least 3 where every permutation has a finite magnitude.

Theorem 3.12. *Let Π be a set of permutations such that every permutation $\sigma \in \Pi$ has a finite magnitude and $mg(\Pi) = k$, where k is an integer larger than 2. Then $deg(0, \Pi) = 0$ and for $m \geq 1$*

$$deg(m, \Pi) = \left\lfloor \frac{(d-1)(k-1)}{2} + \frac{m}{d} \right\rfloor \quad \textit{where} \quad d = \left\lceil \frac{1}{2}\left(-1 + \sqrt{1 + \frac{8m}{k-1}}\right)\right\rceil.$$

Proof. Any permutation σ with major index 0 is strictly increasing, therefore σ avoids Π and $M_n^0(\Pi) = 1 = n^0$. In the rest of the proof suppose $m \geq 1$.

We will prove the theorem by showing that the following values are equal.

1. The degree of the polynomial $l_1 = deg(m, \Pi)$.
2. The largest integer l_2 such that there is a $12\cdots k$-avoiding permutation σ with $maj^+(\sigma) \leq m$.
3. The largest integer l_3 such that there is a $12\cdots k$-avoiding permutation π with $maj^+(\pi) = m$.
4. The value $l_4 = \left\lfloor \frac{(d-1)(k-1)}{2} + \frac{m}{d} \right\rfloor$, where $d = \left\lceil \frac{1}{2}\left(-1 + \sqrt{1 + \frac{8m}{k-1}}\right)\right\rceil$.

First observe that trivially $l_2 \geq l_3$. We will prove $l_3 \geq l_4$ by constructing a $12\cdots k$-avoiding permutation π of length l_4 satisfying $maj^+(\pi) = m$. For a permutation $\psi \in \mathcal{S}_n$ we say that ψ is *co-layered* if ψ avoids both 132 and 213. Observe that any co-layered permutation is uniquely determined by its descent set. Let d be the smallest positive integer such that $\frac{d^2+d}{2}(k-1) \geq m$. Furthermore, let s be the largest integer such that $\frac{d^2+d}{2}(k-1) - ds \geq m$ and let $p = \frac{d^2+d}{2}(k-1) - ds - m$. Note that $s < k-1$ because otherwise the first inequality would also hold for $d' = d-1$. We also know that $p < d$ because otherwise the above inequalities would hold for $s' = s+1$ and $p' = p-d$, which contradicts our choice of s.

Let π be a co-layered permutation of length d_d with $\mathcal{D}(\pi) = \{d_1, d_2, \ldots, d_{d-1}\}$, where

$$d_i = \begin{cases} i(k-1) - s & \text{for} \quad 1 \leq i \leq d-p \\ i(k-1) - s - 1 & \text{for} \quad d-p < i \leq d. \end{cases}$$

To see that π is correctly defined, we will show that the d_i are strictly increasing and positive. From the inequalities $s < k-1$ and $p < d$, it follows that $d_1 \geq 1$, and that $d_{i+1} > d_i$ for every index i. Note that this is where the proof would fail for sets with magnitude $k = 2$, since for $p \neq 0$ we would have $d_{d-p} = d_{d-p+1}$ and we could not construct such permutation.

Observe that π avoids $12\cdots k$ because the longest increasing subsequence in any co-layered permutation is between two adjacent descents and for any $i \in [d-1]$ we have $d_{i+1} - d_i \leq k-1$. We see that π satisfies

$$maj^+(\pi) = \sum_{i=1}^{d} d_i = \frac{d^2+d}{2}(k-1) - ds - p = m.$$

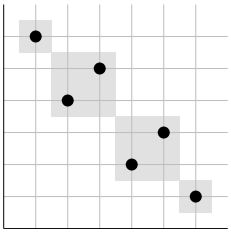

Figure 3.1: The constructed co-layered permutation for $m = 15$ and $k = 3$. In this case we get $d = 4$, $s = 1$ and $p = 1$.

Furthermore, we will show that π has length l_4. By solving the equations we get

$$d = \left\lceil \frac{1}{2}\left(-1 + \sqrt{1 + \frac{8m}{k-1}}\right)\right\rceil, \quad s = \left\lfloor \frac{(d+1)(k-1)}{2} - \frac{m}{d} \right\rfloor.$$

Notice that we subtract 1 during the calculation of d_d if and only if $p \neq 0$ which happens when $\frac{(d+1)(k-1)}{2} - \frac{m}{d}$ is not an integer. This justifies the following equation

$$d_d = d(k-1) - \left\lceil \frac{(d+1)(k-1)}{2} - \frac{m}{d} \right\rceil = \left\lfloor \frac{(d-1)(k-1)}{2} + \frac{m}{d} \right\rfloor = l_4.$$

In order to prove $l_4 \geq l_2$, let τ be a $12\cdots k$-avoiding permutation with length $t \geq l_4 + 1$. Because τ avoids $12\cdots k$ there has to be a descent in the sequence $\tau_{t-(k-1)} \cdots \tau_t$, another one in the sequence $\tau_{t-2(k-1)} \cdots \tau_{t-(k-1)}$ and so on. But that leads to the following inequality

$$maj^+(\tau) \geq t + \sum_{i=1}^{i(k-1)<t} (t - i(k-1)) > d_d + \sum_{i=1}^{d-1} d_i = m.$$

So far we have proved the equality $l_2 = l_3 = l_4$. Now we will show that $l_1 \geq l_3$. Let σ be the longest $12\cdots k$ avoiding permutation with extended major index m. Observe that any permutation with core σ avoids Π and has major index m. We will bound $M_n^{[\sigma]}(\Pi)$ from below with the number of such permutations which have its minimum on the position $l_3 + 1$. We can arbitrarily choose l_3 letters, which will form the core, from all letters except the letter 1. That gives us the lower bound $M_n^m(\Pi) \geq M_n^{[\sigma]}(\Pi) \geq \binom{n-1}{l_3} \geq \alpha n^{l_3}$ for some constant α.

Finally, we will complete the proof by showing that $l_2 \geq l_1$. Fix a permutation $\tau \in \Pi$ with the minimal magnitude k. Trivially the inequality $deg(m, \Pi) \leq deg(m, \{\tau\})$ holds. And because every permutation ψ with $maj^+(\psi) \leq m$ and length greater than l_2 contains $12\cdots k$, we get the desired upper bound from Lemma 3.9. □

Notice that for $m \leq k-1$ we obtain from Theorem 3.12 $d = 1$ and $deg(m, \Pi) = deg(m, \emptyset) = m$. On the other hand for $m \geq k$ the degree is strictly smaller than m and behaves approximately as $\sqrt{m}$.

As suggested by Proposition 3.10, Theorem 3.12 does not hold for the sets Π containing permutations with both finite and infinite magnitude. Similar claim

also cannot hold in general for the sets of magnitude 2. Consider the set $\Pi = \{132, 231\}$ of magnitude 2 and $\sigma \in \mathcal{S}_n(\Pi)$. Let j be an integer such that $\sigma_j = 1$. Then the sequence $\sigma_1\sigma_2\cdots\sigma_j$ is decreasing since σ avoids 231 and similarly the sequence $\sigma_j\sigma_{j+1}\cdots\sigma_n$ is increasing because σ avoids 132. In other words, every Π-avoiding permutation has a decreasing core. On the other hand, every permutation π with the decreasing core $(d-1)(d-2)\cdots 1$ avoids Π and $maj(\pi) = \frac{d^2+d}{2}$. As a result, $deg(m, \Pi) \neq 0$ if and only if m can be expressed as $m = \frac{d^2+d}{2}$ for some integer d. Therefore, unlike what we have seen so far, the degrees in this case do not satisfy $deg(m, \Pi) \leq deg(m+1, \Pi)$.

However, we can prove a weaker version of Theorem 3.12 by placing some further restrictions on the set Π of magnitude 2.

Proposition 3.13. *Let Π be a set of permutations such that every permutation $\sigma \in \Pi$ has a finite magnitude and $mg(\Pi) = 2$. Furthermore, assume that there is $i \in \{1, 2, 3\}$ such that every permutation $\pi \in \Pi$ with $mg(\pi) = 2$ has a padding profile $a \in \mathbb{N}_0^3$ with $a_i \neq 0$. Then $deg(m, \Pi) = \lfloor \frac{-1+\sqrt{1+8m}}{2} \rfloor$.*

Proof. Let l be the largest integer such that $\frac{l^2+l}{2} \leq m$. By solving the quadratic equation, we get $l = \lfloor \frac{-1+\sqrt{1+8m}}{2} \rfloor$. Again to show that the degree of the polynomial is equal to l, we will prove the following claims.

1. There is a constant $\alpha = \alpha(m, k)$ such that $M_n^m(\Pi) \geq \alpha n^l$.
2. The inequality $deg(m, \Pi) \leq l$ holds.

We will construct a core γ for which $M_n^{[\gamma]}(\Pi) \geq \alpha n^l$ holds. If we have $m = \frac{l^2+l}{2}$ we will take as a core the descending permutation of length l. Every permutation with this core is Π-avoiding and has major index m, thus giving the desired lower bound.

Otherwise let $d = m - \frac{l^2+l}{2}$. Observe that $d \leq l$, because otherwise $\frac{l'^2+l'}{2} \leq m$ would hold for $l' = l + 1$. Now we will construct a core of length $l + 1$ depending on the $i \in \{1, 2, 3\}$, for which the assumptions of the proposition hold. Let $\epsilon = l \cdots 1$, then we will construct the core γ by inserting one letter to ϵ,

$$\gamma = \begin{cases} \epsilon[l+1-d \to 1] & \text{for } i = 1 \\ \epsilon[l+1-d \to d] & \text{for } i = 2 \\ \epsilon[l+2-d \to l+1] & \text{for } i = 3. \end{cases}$$

For an example see Figure 3.2. Observe that γ no longer avoids 12, but it satisfies $maj^+(\gamma) = m$. Let $a \in \mathbb{N}_0^{l+2}$ be a tuple which satisfies one of the following conditions depending on the value of i.

$$\begin{aligned} a_1 = 0 \ , \ a_2 \neq 0 \quad & \text{for } i = 1 \\ a_{d+1} = 0 \ , \ a_1 \neq 0 \quad & \text{for } i = 2 \\ a_{l+2} = 0 \ , \ a_1 \neq 0 \quad & \text{for } i = 3 \end{aligned}$$

We know that $\gamma \cdot a$ has major index m and does not contain any permutation with magnitude larger than 2. But it also cannot contain any permutation $\sigma \in \Pi$

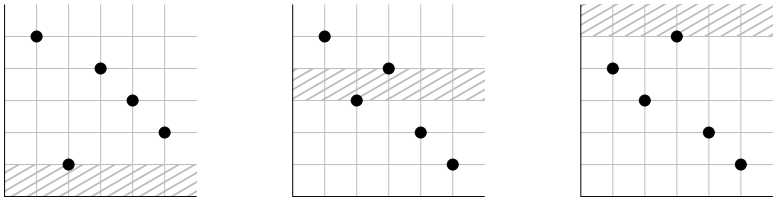

Figure 3.2: The constructed cores for $m = 15$, from left to right for $i = 1, 2, 3$. Any permutation with such core and no element in the grey strip is Π-avoiding.

with magnitude 2 because of the conditions above. Therefore $\gamma \cdot a \in \mathcal{M}_n^m(\Pi)$. Since there are $\binom{n-2}{l}$ such padding profiles, we see that $M_n^{[\gamma]}(\Pi) \geq \alpha n^l$ for some α.

To prove our second claim, fix a permutation $\tau \in \Pi$ with magnitude 2. Observe that any permutation σ with length at least $l+1$ for which $maj^+(\sigma) \leq m$ holds, necessarily contains 12. Therefore, the upper bound $deg(m, \Pi) \leq l$ is implied by Lemma 3.9. □

As one would expect the formula for $deg(m, \Pi)$ in Proposition 3.13 gives the same result as the one in Theorem 3.12 for $k = 2$. It is straightforward to check if you express m as $m = \frac{t^2+t}{2} + s$ for some integer t and $s \leq t$.

From Propositions 3.8, 3.11, 3.13 and Theorem 3.12 we know the values of $deg(m, \Pi)$ for all sets Π with $|\Pi| = 1$ and $mg(\Pi) \geq 1$. Furthermore, as we already know, for any permutation σ with magnitude 0, we eventually have $M_n^m(\sigma) = 0$.

Corollary 3.14. *For a singleton set $\Pi = \{\sigma\}$ and $k = mg(\Pi) = mg(\sigma)$*

$$deg(m, \Pi) = \begin{cases} m & \text{if } k = +\infty \\ 0 & \text{if } k \leq 1 \text{ or } m = 0 \\ s_k(m) & \text{otherwise} \end{cases}$$

where $s_k(m) = \left\lfloor \frac{(d-1)(k-1)}{2} + \frac{m}{d} \right\rfloor$ and $d = \left\lceil \frac{1}{2}\left(-1 + \sqrt{1 + \frac{8m}{k-1}}\right) \right\rceil$. Furthermore, $deg(m, \Pi) = m$ for $m < k$ and otherwise $deg(m, \Pi) < m$.

Proof. Since the values of $deg(m, \Pi)$ were determined in the previous claims, we will only prove the inequalities. It is easier to use bounds on $deg(m, \Pi)$ than to work with the equations for $s_k(m)$. Let $\epsilon = 12 \cdots m$. For $m < k$ every permutation with core ϵ avoids Π and $M_n^{[\epsilon]}(\Pi) = M_n^{[\epsilon]}(\emptyset) \geq \alpha n^m$ for some constant α. Thus $s_k(m) = deg(m, \Pi) = m$.

On the other hand, for $m \geq k$ every permutation π such that $maj^+(\pi) \leq m$ and $|\pi| \geq m$ contains $12 \cdots k$ (in fact ϵ is the only such permutation). And Lemma 3.9 implies $s_k(m) = deg(m, \Pi) \leq m - 1$. □

Moreover, for an arbitrary set of permutations Π we can use Corollary 3.14 to provide an upper bound on $deg(m, \Pi)$. Let $\tau \in \Pi$ such that $mg(\tau) = mg(\Pi)$, then $deg(m, \Pi) \leq deg(m, \{\tau\})$.

3.4 Asymptotic probability

Our previous results in this chapter imply a sharp dichotomy for the probability that a random permutation with a fixed major index avoids a specific set of patterns Π.

Theorem 3.15. *Let Π be a set of permutations and m a non-negative integer. Then*

$$\lim_{n\to\infty} \frac{M_n^m(\Pi)}{M_n^m(\emptyset)} = \begin{cases} 1 & \text{if } m < mg(\Pi) \\ 0 & \text{otherwise.} \end{cases}$$

Proof. First notice that directly from Proposition 3.8, it follows that $M_n^m(\emptyset) = \frac{n^m}{m!} + O(n^{m-1})$. For $m < mg(\Pi)$ every permutation with core $\epsilon = 12\cdots m$ avoids Π. As we already showed in the proof of Proposition 3.8, $M_n^{[\epsilon]}(\Pi) = \frac{n^m}{m!} + O(n^{m-1})$. Therefore, the ratio is approaching 1 as n goes to infinity.

For $m \geq mg(\Pi)$, we know that $deg(m, \Pi) < m$ (recall Corollary 3.14). Therefore, the polynomial in the numerator has smaller degree than the one in the denominator and the ratio is approaching 0 as n goes to infinity. □

Conclusion and further directions

In Chapter 2, we proved the monotonicity of the numbers $M_n^m(\sigma)$ for a single pattern σ other than $12\cdots k$ (recall Theorem 2.4) and showed an example of a set Π for which the monotonicity does not hold even though $M_n^m(\Pi)$ tends to infinity. The natural question to ask would be whether we can in general characterize such sets Π for which the monotonicity of columns does not hold even though $deg(m, \Pi) \geq 1$. Based on computing the values $M_n^m(\Pi)$ for small n and various sets Π, it seems to us that these cases are rather rare.

In Chapter 3, we analysed the asymptotic behaviour of the numbers $M_n^m(\Pi)$ for many types of Π in the sense of the degree $deg(m, \Pi)$. The most natural way to extend this study is to cover the remaining cases. For example, it remains to be shown whether the sets Π that contain permutations with both finite and infinite magnitude obey any general rules. Another open problem is to determine exactly for which sets Π the values $M_n^m(\Pi)$ are eventually equal to zero.

One could also focus on generalized pattern avoidance. A permutation σ contains a copy of a generalized pattern π if it contains π and certain elements of the diagram of the copy are adjacent either horizontally or vertically. The concept of generalized patterns was introduced by Babson and Steingrímsson [1]. The reason they are interesting is because many statistics on permutations (including the number of inversions and the major index) can be expressed as a linear combination of the number of occurrences of these generalized patterns.

Finally, similar analysis of the distribution could be done for other permutation statistics like number of descents or number of excedances. As previously mentioned, the number of inversions was already covered by Claesson, Jelínek and Steingrímsson [6]. One can find examples of various other pattern statistics in a classification given by Babson and Steingrímsson [1].

Bibliography

[1] Babson, E. – Steingrímsson, E. Generalized permutation patterns and a classification of the Mahonian statistics. *Sém. Lothar. Combin.* 2000, 44, p. Art. B44b, 18 pp. (electronic). ISSN 1286-4889.

[2] Bloom, J. A refinement of Wilf-equivalence for patterns of length 4. *J. Combin. Theory Ser. A*. 2014, 124, p. 166–177. ISSN 0097-3165.

[3] Bloom, J. – Saracino, D. On bijections for pattern-avoiding permutations. *J. Combin. Theory Ser. A*. 2009, 116, 8, p. 1271–1284. ISSN 0097-3165.

[4] Bloom, J. – Saracino, D. Another look at bijections for pattern-avoiding permutations. *Adv. in Appl. Math.* 2010, 45, 3, p. 395–409. ISSN 0196-8858.

[5] Claesson, A. – Kitaev, S. Classification of bijections between 321- and 132-avoiding permutations. *Sém. Lothar. Combin.* 2008/09, 60, p. Art. B60d, 30. ISSN 1286-4889.

[6] Claesson, A. – Jelínek, V. – Steingrímsson, E. Upper bounds for the Stanley-Wilf limit of 1324 and other layered patterns. *J. Combin. Theory Ser. A*. 2012, 119, 8, p. 1680–1691. ISSN 0097-3165.

[7] Dokos, T. et al. Permutation patterns and statistics. *Discrete Math.* 2012, 312, 18, p. 2760–2775. ISSN 0012-365X.

[8] Elizalde, S. Fixed points and excedances in restricted permutations. *Electron. J. Combin.* 2011, 18, 2, p. Paper 29, 17. ISSN 1077-8926.

[9] Erdős, P. – Szekeres, G. A combinatorial problem in geometry. *Compositio Math.* 1935, 2, p. 463–470. ISSN 0010-437X.

[10] Foata, D. On the Netto inversion number of a sequence. *Proc. Amer. Math. Soc.* 1968, 19, p. 236–240. ISSN 0002-9939.

[11] Knuth, D. E. *The art of computer programming. Vol. 1: Fundamental algorithms*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont, 1968. ISBN 0-201-03801-3.

[12] MacMahon, P. A. *Combinatory analysis*. vols. 1 and 2. Cambridge Univ. Press, Cambridge (reprinted by Chelsea,New York, 1955), 1915. ISBN 9780828411370.

[13] Robertson, A. – Saracino, D. – Zeilberger, D. Refined restricted permutations. *Ann. Comb.* 2002, 6, 3-4, p. 427–444. ISSN 0218-0006.

[14] Rodrigues, O. Note sur les Inversions, ou dérangements produits dans les permutations. *J. Math. Pures Appl.* 1839, 4, p. 236–240. ISSN 0021-7824.

[15] Sloane, N. J. A. *The On-Line Encyclopedia of Integer Sequences* [online]. [cit. 2015-05-01]. Triangle of Mahonian numbers T(n,k). Available from: <http://oeis.org/A008302>.

[16] Stanley, R. Problems and Solutions: Solutions of Elementary Problems: E2546. *Amer. Math. Monthly.* 1976, 83, 10, p. 813–814. ISSN 0002-9890.

[17] Stanley, R. et al. Problems and Solutions: Elementary Problems: E2546-E2551. *Amer. Math. Monthly.* 1975, 82, 7, p. 755–756. ISSN 0002-9890.

[18] Stern, M. A. Aufgaben. *J. Reine Angew. Math.* 1838, 18, p. 100. ISSN 1435-5345.

[19] Yan, S. H. F. – Ge, H. – Zhang, Y. On a refinement of Wilf-equivalence for permutations. *Electron. J. Combin.* 2015, 22, 1, p. Paper 1.20, 26. ISSN 1077-8926.