

Preimages under the bubblesort operator

Mathilde Bouvel

Université de Lorraine, CNRS, Inria, LORIA
F-54000 Nancy, France

`mathilde.bouvel@loria.fr`

Lapo Cioni* Luca Ferrari*

Dipartimento di Matematica e Informatica “U.Dini”
Università degli studi di Firenze
Firenze, Italy

`{lapo.cioni,luca.ferrari}@unifi.it`

Submitted: Jul 14, 2022; Accepted: Oct 3, 2022; Published: Nov 18, 2022

© The authors. Released under the CC BY license (International 4.0).

Abstract

We study preimages of permutations under the bubblesort operator \mathbf{B} . We achieve a description of these preimages much more complete than what is known for the more complicated sorting operators \mathbf{S} (stacksort) and \mathbf{Q} (queuesort). We describe explicitly the set of preimages under \mathbf{B} of any permutation π from the left-to-right maxima of π , showing that there are 2^{k-1} such preimages if k is the number of these left-to-right maxima. We further consider, for each n , the tree T_n recording all permutations of size n in its nodes, in which an edge from child to parent corresponds to an application of \mathbf{B} (the root being the identity permutation), and we present several properties of these trees. In particular, for each permutation π , we show how the subtree of T_n rooted at π is determined by the number of left-to-right maxima of π and the length of the longest suffix of left-to-right maxima of π . Building on this result, we determine the number of nodes and leaves at every height in such trees, and we recover (resp. obtain) the average height of nodes (resp. leaves) in T_n .

Mathematics Subject Classifications: 05A05, 05A15, 05C05

1 Introduction

1.1 Motivation

The foundational work of Knuth [21, Section 2.2.1] defines the stacksorting procedure and relates it to pattern-avoidance in permutations. Since then, many similar sorting

*member of the INdAM research group GNCS; partially supported by the 2022 INdAM-GNCS project “Stringhe e matrici. combinatoria, enumerazione e algoritmi”.

procedures have been defined, where the sorting device is not necessarily a stack (*e.g.*, a queue, a deque, ...) and where such devices can be combined. These sorting procedures have been studied from various points of view, and we refer the reader to the surveys [3, 18] and to the introduction of [25] for an overview of the research on this topic and bibliographic pointers.

Our focus in the present work is on the bubblesort operator, \mathbf{B} , which corresponds to one pass of the bubblesort algorithm. Although not strictly speaking a sorting procedure associated with a particular device, \mathbf{B} shares some features with such procedures, as demonstrated in [1]. Our point of focus is the study of preimages of permutations under \mathbf{B} , a topic which has proved very rich for other sorting operators. We start by reviewing the literature related to the study of preimages under a sorting operator, focusing on different aspects of this study.

Existence and number of preimages To our knowledge, Bousquet-Mélou is the first to investigate a sorting operator (here, the stacksorting operator \mathbf{S}) through the lens of preimages. In her article [4], she provides an algorithm to decide if a permutation has at least one preimage under \mathbf{S} , and in this case to compute one such preimage called *canonical*, from which all other preimages can then be computed. She also derives an equation for the generating functions of permutations having at least one preimage under \mathbf{S} .

Recent improvements on the study of permutations having preimages under \mathbf{S} have been obtained by Defant and co-authors. Specifically, [14] counts permutations having exactly one preimage under \mathbf{S} ; [9] studies the number of preimages that a permutation can have; and [8] uses a method involving the computation of preimages to give bounds on the number of t -stack-sortable permutations of size n .

In [6, 5], Cioni and Ferrari consider another sorting operator: the queuesort operator \mathbf{Q} . They provide a recursive description of the preimages under \mathbf{Q} , study the possible numbers of preimages that a permutation can have, and compute the number of permutations having 0, 1 or 2 preimages.

Sorting trees In [4], Bousquet-Mélou also defined the *sorting trees* associated with \mathbf{S} . For each integer n , the sorting tree (see [4, fig. 2]) for size n is the tree whose nodes are the permutations of size n , with root the identity permutation $12\dots n$, and such that the children τ of any permutation σ are those such that $\mathbf{S}(\tau) = \sigma$. The article [4] already describes some properties of the sorting trees (although not phrased as such, the focus in [4] being different).

Building on these, Defant [10] proved in disguise further properties of the sorting trees. Namely, assigning to any permutation σ a “label” which consists of the skeleton of the decreasing binary tree whose in-order reading is σ , he shows [10, Theorem 11] that the label of σ determines the number and the labels of the children of σ in the sorting tree, and hence recursively the whole skeleton of the subtree of the sorting tree rooted at σ .

We point out an easy consequence of this fact: for every permutation σ , there exists a permutation τ which is at distance one from the root in the sorting tree, and such that the subtrees of the sorting tree rooted at σ and at τ are isomorphic. This indeed follows

from the two facts that permutations at distance one from the root in the sorting tree are the stack-sortable permutations (*i.e.*, those avoiding 231) and that every binary tree has a decreasing labeling which avoids 231.

About queuesort, [6] mentions that it would be interesting to study the properties of the sorting trees associated with \mathbf{Q} , but we are not aware of any such results at the moment.

Preimages of permutation classes The permutation class defined by a set of excluded patterns B is the set of all permutations avoiding every pattern in B , and is denoted $\text{Av}(B)$. A consequence of the characterization of stack-sortable permutations as those avoiding 231 is that the permutations sortable by two applications of \mathbf{S} (called West-two-stack-sortable) are the preimages of permutations in the class $\text{Av}(231)$. Describing the West-two-stack-sortable permutations is therefore an instance of the more general question asking for a description of the preimages (for \mathbf{S}) of a class of pattern-avoiding permutations.

This question has been studied algorithmically in [7], where Claesson and Ulfarsson provide an algorithm to describe preimages of principal pattern classes, in terms of decorated patterns. Their approach has been extended in [22] to other sorting operators, including \mathbf{Q} .

A more enumerative perspective on preimages of permutation classes for \mathbf{S} has later been provided in [12, 11].

Finally, the problem of describing preimages of permutation classes for \mathbf{B} has been completely solved in [1] for principal permutation classes.

Complexity of sorting procedures As sorting algorithms, the procedures considered above are clearly inefficient. Nevertheless, it makes sense to ask how many applications of a sorting operator (like \mathbf{B} or \mathbf{S} or others) are needed to fully sort a permutation of size n , either in the worst case or on average. Some bounds for the average case are provided in [24] for various sorting operators, and in [13] for \mathbf{S} .

1.2 Our results

We focus on the bubblesort operator \mathbf{B} , and study preimages of permutations under \mathbf{B} . As previously indicated, a study of preimages of permutation classes has already been done in [1], and we leave these aspects aside of our study.

About the first point above (characterizing the existence and number of preimages), we can be very precise when confronting these questions with bubblesort. Indeed, it is easy to fully describe the set $\mathbf{B}^{-1}(\sigma)$ of preimages of any permutation σ . This is presented in Section 2.

The description of $\mathbf{B}^{-1}(\sigma)$ is actually rather simple, and involves essentially the left-to-right maxima of σ . It can be traced back to ideas already present in the literature under various forms, originating from the work of Knuth [21] (see Section 2 for more details). From this description, we deduce that the number of preimages of a permutation σ is 2^{k-1} for k the number of left-to-right maxima of σ if σ ends with its maximum (and 0

otherwise). The results are therefore much more precise than what can be achieved with \mathbf{S} for instance. An informal explanation which we can offer to explain this fact is the following. The operator \mathbf{B}^{-1} acts only on the left-to-right maxima of a permutation, and does so in a very controlled way, allowing to describe the left-to-right maxima after the application of \mathbf{B}^{-1} . This is very useful in particular for describing iterated preimages for \mathbf{B} .

In Section 3, we turn to the study of these iterated preimages. Specifically, we define for each permutation π the tree $T(\pi)$ whose root is π such that the children of any permutation are the preimages of this permutation. These trees are the analogue in the case of \mathbf{B} of the *sorting trees* defined by Bousquet-Mélou in [4] for \mathbf{S} .

As expected from the informal discussion above, we show that these trees are completely determined by the left-to-right maxima of π . More precisely, they are determined by what we call the *label* of π , which plays the same role as the “label” we presented above in the case of \mathbf{S} but is much simpler than it (recall that this was the skeleton of the decreasing binary tree whose in-order reading is π). In the case of \mathbf{B} , the *label* of π is the pair consisting of the number of left-to-right maxima of π and the length of the longest suffix of left-to-right maxima of π . We also prove, similarly to the case of \mathbf{S} , that every tree $T(\pi)$ is isomorphic to a tree $T(\tau)$ for τ a permutation such that $\mathbf{B}(\tau)$ is the identity.

Finally, in Section 4 we study heights of nodes and leaves in $T(\pi)$. For π the identity permutation of size n , the average height of a node in $T(\pi)$ is the average number of passes of \mathbf{B} necessary to sort a permutation of size n , which is known from [19, Theorem 7.14] for instance. We modify this analysis to also compute the average height of a leaf in $T(\pi)$ (which corresponds to the number of passes of \mathbf{B} necessary to sort a permutation of size n which does not belong to the image of \mathbf{B}). In addition, for any permutation π , we provide closed formulas for the number of nodes and leaves at any possible height j in $T(\pi)$; these formulas depend only on j and on the label of π . Nevertheless, we could not deduce the average height of a leaf or node in $T(\pi)$ from these formulas.

Before moving to definitions and basic properties of \mathbf{B} , we note that we see the questions studied here as bubblesort analogues of similar questions previously studied on the stacksort and queuesort operators. The operator \mathbf{B} being simpler, the answers obtained are much more precise than in the case of \mathbf{S} or \mathbf{Q} . A possible direction for future research is to confront these questions with other sorting operators as well. Those listed in [22] can be a very good source of inspiration, as well as the recent *promotion sorting* [15, 20] which generalizes \mathbf{B} .

1.3 The operator \mathbf{B} : definition and some basic properties

For any integer n , a permutation π of size n is a sequence $\pi_1\pi_2\dots\pi_n$ containing exactly once each symbol from $\{1, 2, \dots, n\}$. An element π_j of a permutation π is a left-to-right maximum if it is larger than all elements to its left, that is to say π_j is such that $\pi_j > \pi_i$ for all $i < j$.

The bubblesort operator, denoted \mathbf{B} , corresponds to applying one pass of the classical bubblesort algorithm to a permutation. Specifically, $\mathbf{B}(\pi)$ is obtained from π scanning its elements from left to right, each time exchanging an element with the one sitting to its

right whenever the latter is smaller. Thereby (see also Lemma 2 below), the left-to-right maxima of the permutation “bubble up” to the right, until they are blocked by the next left-to-right maximum.

For example, for $\pi = \mathbf{42163785}$, the left-to-right maxima are 4, 6, 7 and 8 (shown in bold) and $\mathbf{B}(\pi) = 21436758$.

Observation 1. *From the above definition, it is clear that \mathbf{B} can be applied verbatim to sequences of distinct integers which are not necessarily permutations. Consequently, every statement about \mathbf{B} on permutations also applies to sequences of distinct integers up to relabeling the values with the order-isomorphic permutation.*

For instance, $\mathbf{B}(\mathbf{42163}) = 21436$, and the preimages of 21436 for \mathbf{B} are obtained from the preimages of the permutation 21435 replacing 5 with 6.

The bubblesort operator may be described in several other ways, and we give two below. The reader needing an explanation of the equivalence with the above definition can find it in [1, Lemma 1 and the equation displayed just above it]. The first parallels the recursive definition of the stacksorting operator \mathbf{S} : decomposing π into $\pi = \pi_L n \pi_R$ with n the maximal value occurring in π , we have $\mathbf{B}(\pi) = \mathbf{B}(\pi_L) \pi_R n$ (and by comparison $\mathbf{S}(\pi) = \mathbf{S}(\pi_L) \mathbf{S}(\pi_R) n$). The second focuses on the left-to-right maxima, and we record it in a lemma for future reference.

Lemma 2. *Let π be a permutation, and write $\pi = \mu_1 A_1 \mu_2 A_2 \dots \mu_k A_k$, where the μ_i 's are all the left-to-right maxima of π (and the A_i are possibly empty sequences of integers). Then $\mathbf{B}(\pi) = A_1 \mu_1 A_2 \mu_2 \dots A_k \mu_k$.*

In particular, a permutation is in the image of \mathbf{B} if and only if it ends with its maximum. We record below another observation which follows immediately from Lemma 2

Corollary 3. *For any π , the set of left-to-right maxima of π is included in the set of left-to-right maxima of $\mathbf{B}(\pi)$.*

2 Computing the preimages

In this section we present a procedure to compute the set $\mathbf{B}^{-1}(\sigma)$ of all preimages of any given permutation σ . This procedure is simple, and can also be implicitly found (together with some of the consequences it bears) in [16]. In this article the authors, using a different approach, obtain some of the results we also get in the present section, such as Corollary 7.

First, as noted just after Lemma 2, for a permutation σ which does not end with its maximum, $\mathbf{B}^{-1}(\sigma)$ is empty. This trivial case being solved, we now focus on the interesting case where σ does end with its maximum.

Let $\sigma = \sigma_1 \sigma_2 \dots \sigma_n$ be a permutation of size n which ends with its maximum. Define P as a set which contains only σ . For each i from n down to 2, do the following: for each $\pi \in P$,

- if π_{i-1} is not a left-to-right maximum of π , then replace π in the set P by the permutation $\pi_1 \dots \pi_i \pi_{i-1} \dots \pi_n$ (that is to say, we swap π_i and π_{i-1});
- if π_{i-1} is a left-to-right maximum of π , then π stays in the set P , and in addition we add in P the permutation $\pi_1 \dots \pi_i \pi_{i-1} \dots \pi_n$ (where π_i and π_{i-1} are swapped).

Example 4. The table below shows the evolution of the set P of the above procedure, for $\sigma = 325146$.

$i = \dots$	initialization	6	5	4	3	2
P contains	325146	325164	325614	325614 326514	352614 362514	352614, 532614 362514, 632514

We may note that, when starting any step i , for any $\pi \in P$, π_i is always a left-to-right maximum of π , and π_{i-1} is a left-to-right maximum of π if and only if σ_{i-1} is a left-to-right maximum of σ . Indeed, all steps until step i (excluded) of the above procedure leave the prefixes of length $i - 1$ unchanged.

It is useful to have a different (although equivalent) presentation of this procedure, which we now give. Starting from σ , where we see the rightmost element σ_n as distinguished, we move the distinguished element to the left until it becomes the leftmost, according to the following rules.

- If the element immediately to the left of the distinguished one is not a left-to-right maximum of the current sequence, then the distinguished element is forced to move to the left (*i.e.* is swapped with its left neighbor). The distinguished element remains the same.
- If the element immediately to the left of the distinguished one is a left-to-right maximum of the current sequence, then the distinguished element may either move to the left or stay in place. In the first case, the distinguished element remains the same. In the second case, the distinguished element becomes the left neighbor of the previously distinguished element.

It is easy to see that the set P computed by the original procedure consists of all possible results of applying this alternative procedure.

Some remarks (all easily observed) about this alternative procedure are useful. First, the index i (between n and 2) of a given step of the original procedure always corresponds to the position of the distinguished element in the evolving sequence. Second, the distinguished element is always a left-to-right maximum of σ , and also of the evolving sequence. Third, for any sequence π produced, the elements which were at some point distinguished in the sequence are exactly the left-to-right maxima of π .

Theorem 5. *Let σ be any permutation ending with its maximum, and P be the set produced by the procedure above. Then P is the set of preimages of σ under \mathbf{B} , that is to say, $P = \mathbf{B}^{-1}(\sigma)$.*

Proof. Assume first that a sequence π has been produced by the above procedure. It means that π has been produced from σ by considering some left-to-right maxima of σ , from the right to the left, and moving these left-to-right maxima to the left, until they reach the position of the next left-to-right maximum which moves to the left. This is exactly undoing the action of \mathbf{B} . More precisely, assume that $\pi = \mu_1 A_1 \cdots \mu_k A_k$ with the μ_i the left-to-right maxima of π . By construction of π , A_i is not empty if and only if μ_i has been moved by our procedure. Applying \mathbf{B} to π yields $A_1 \mu_1 \cdots A_k \mu_k$, thus exactly the elements that were moved by our procedure will be moved to the right by bubblesort. Moreover, \mathbf{B} moves μ_i to the right until it reaches the position immediately before μ_{i+1} , and we claim that this is μ_i 's original position in σ : indeed, our procedure only moves the μ_j 's, so it must have started moving μ_i to the left immediately after considering μ_{i+1} . This proves that $\mathbf{B}(\pi) = \sigma$, and therefore $P \subseteq \mathbf{B}^{-1}(\sigma)$.

For the converse inclusion, we proceed by induction on the size of σ . The statement is obvious for size 1. So, let us consider $\pi \in \mathbf{B}^{-1}(\sigma)$, for σ of size greater than 1. We decompose π around its maximal element as $\pi = LnR$. Then $\mathbf{B}(\pi) = \mathbf{B}(L)Rn$, so that $\sigma = \mathbf{B}(L)Rn$. Starting from $\sigma = \mathbf{B}(L)Rn$, the above procedure is always allowed to move n towards the left, and may decide when reaching $\mathbf{B}(L)nR$ to distinguish the last element of $\mathbf{B}(L)$ instead of n . Indeed, $\mathbf{B}(L)$ necessarily ends with its maximum, so that the last element of $\mathbf{B}(L)$ is a left-to-right maximum. Since $\mathbf{B}(L)$ is a sequence shorter than σ ending with its largest value, we may apply the induction hypothesis to it, and deduce that L has been produced by the above procedure applied to $\mathbf{B}(L)$. Combining these two facts, it follows that our procedure applied to $\sigma = \mathbf{B}(L)Rn$ can produce $LnR = \pi$, therefore showing that $\mathbf{B}^{-1}(\sigma) \subseteq P$. \square

Theorem 5 has several consequences. First, we can refine Corollary 3 and describe $\mathbf{B}^{-1}(\sigma)$ exactly from the left-to-right maxima of σ .

Corollary 6. *Let σ be a permutation of size n ending with its maximum (i.e., $\sigma_n = n$). Let k be the number of left-to-right maxima of σ (including n).*

There is a bijective correspondence between the preimages of σ under \mathbf{B} and the subsets of the $k - 1$ left-to-right maxima of σ different from n .

More precisely, this correspondence works as follows. For any set $S = \{s_1 < \cdots < s_j\}$ of $j \leq k - 1$ left-to-right maxima of σ different from n , writing $\sigma = B_0 s_1 B_1 s_2 B_2 \cdots s_j B_j n$ (for the B_i possibly empty sequences of integers, which contain the $k - j$ left-to-right maxima not in S and the elements of σ which are not left-to-right maxima), the corresponding preimage of σ is $s_1 B_0 s_2 B_1 \cdots s_j B_{j-1} n B_j$.

Proof. From the alternative description of the procedure computing $\mathbf{B}^{-1}(\sigma)$, we have seen that the elements which are distinguished at some point are exactly the left-to-right maxima of the preimage produced. In addition, by definition of this procedure, the distinguished elements form a subset containing n of the set of left-to-right maxima of σ . This proves the claimed bijective correspondence.

To describe precisely the preimage corresponding to a subset S , it is enough to note that every distinguished element moves to the left until a new distinguished element is chosen, leaving all other elements unchanged. \square

This allows to count the preimages of any given permutation, in total or by the number of their left-to-right maxima.

Corollary 7. *Let σ be a permutation of size n ending with its maximum, and with k left-to-right maxima.*

The cardinality of $\mathbf{B}^{-1}(\sigma)$ is 2^{k-1} , and for any $1 \leq j \leq k$, the number of preimages of σ with j left-to right maxima is $\binom{k-1}{j-1}$.

Proof. The cardinality of $\mathbf{B}^{-1}(\sigma)$ follows immediately from Corollary 6. By Corollary 6, a preimage of σ with j left-to-right maxima corresponds bijectively to a subset containing $j - 1$ elements of the set of left-to-right maxima of σ different from n . We have $\binom{k-1}{j-1}$ different ways to select these subsets, thus proving the lemma. \square

Second, we can characterize the permutations having a given number of preimages.

Corollary 8. *For any $k \geq 1$, the permutations having exactly 2^{k-1} preimages under \mathbf{B} are those ending with their maximum and having k left-to-right maxima in total.*

In particular, there are $\left[\begin{smallmatrix} n-1 \\ k-1 \end{smallmatrix} \right]$ permutations of size n having 2^{k-1} preimages under \mathbf{B} , where $\left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right]$ are the (unsigned) Stirling numbers of the first kind.

Proof. The first statement follows immediately from Corollary 6. The second follows from the well-known fact that Stirling numbers of the first kind enumerate permutations according to their size and number of cycles, using the classical Foata bijection which maps permutations of size n with k cycles to permutations of size n with k left-to-right maxima. \square

3 The trees of iterated preimages

For any n , we denote by S_n the set of permutations of size n and by $id_n = 12 \dots n$ the identity permutation of size n . We start by defining $T(\pi)$ for any permutation π , and $T_n = T(id_n)$.

Definition 9. Let T_n be the tree whose nodes are the permutations of S_n such that:

- T_n has root id_n ;
- for every $\sigma, \tau \in S_n$, τ is a child of σ if and only if $\mathbf{B}(\tau) = \sigma$ and $\sigma \neq \tau$. Note that the situation $\mathbf{B}(\tau) = \sigma$ and $\sigma = \tau$ occurs only when $\sigma = \tau = id_n$.

Also, given a permutation $\pi \in S_n$, we define the *tree of its preimages* $T(\pi)$ as the subtree of T_n with root π .

For example, Fig. 1 shows the tree T_4 .

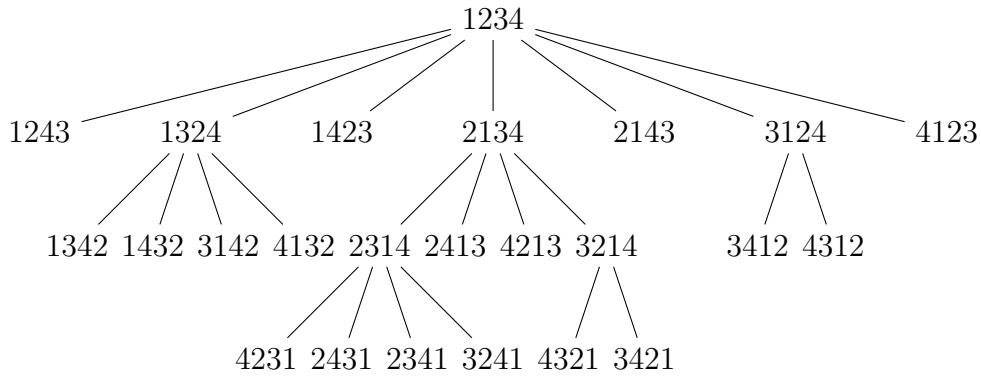


Figure 1: The tree T_4 .

3.1 Isomorphisms between subtrees

For any permutation π , $T(\pi)$ describes all possible preimages of π under repeated applications of \mathbf{B}^{-1} . From Section 2, we can see that the “shape” of this tree depends on π only through the location of its left-to-right maxima. More precisely, the following lemma holds.

Lemma 10. *Let π and τ be two permutations of the same size. If π and τ have their left-to-right maxima in the same positions, then $T(\pi)$ and $T(\tau)$ are isomorphic.*

Proof. Let π, τ be permutations of the same size n which have their left-to-right maxima in the same positions, and let h and k be the depth of $T(\pi)$ and $T(\tau)$, respectively (the *depth* of a tree being defined as the maximum depth of its nodes). Without loss of generality, we can assume that $h \geq k$. The proof is by induction on h . If $h = 0$, then $h = k = 0$, and so $T(\pi)$ and $T(\tau)$ both consist of a single node, and our claim trivially holds.

Now suppose that $h > 0$. Unless $\pi = id_n$, by definition, the nodes of $T(\pi)$ at depth 1 are the preimages of π under \mathbf{B} . By Corollary 6, these preimages are in bijection with all possible subsets of left-to-right maxima of π which do not contain n . Instead of identifying a subset of left-to-right maxima of π by the *values* of the left-to-right maxima it contains, we can identify it by the *positions* of the left-to-right maxima it contains. Since π and τ have their left-to-right maxima in the same positions, it follows from Corollary 6 that there is a bijection between the preimages of π under \mathbf{B} and the preimages of τ under \mathbf{B} . In addition, for every $\sigma \in \mathbf{B}^{-1}(\pi)$, the corresponding $\rho \in \mathbf{B}^{-1}(\tau)$ has its left-to-right maxima in the same positions as those of σ . Since $T(\sigma)$ and $T(\rho)$ have depth at most $h - 1$ and $k - 1$ respectively, we can apply the inductive hypothesis and obtain that $T(\sigma)$ and $T(\rho)$ are isomorphic. Summing up, we have that the children of π and τ are in a bijective correspondence, and the trees rooted at two children paired together by this bijection are isomorphic. Therefore $T(\pi)$ and $T(\tau)$ are isomorphic.

Finally, if π is the identity permutation of size n , and τ (of the same size) has its left-to-right maxima in the same positions as π , then necessarily $\tau = id_n$ as well, and in this case our claim trivially holds. \square

As the next proposition shows, all possible shapes of the trees $T(\pi)$ can be found starting at depth 1 in T_n . To establish this proposition, we rely on the following decomposition of π , which will be also essential in the description of $T(\pi)$ in the next subsection.

Definition 11. Given π , we decompose it as $\pi = M_1P_1M_2P_2 \cdots M_{\ell-1}P_{\ell-1}M_\ell$, where the M_i 's are all the maximal sequences of consecutive left-to-right maxima of π (called *blocks*), and the P_i 's collect all the remaining elements. In particular, all the P_i 's are nonempty, and M_i is nonempty for all i except possibly for $i = \ell$. Moreover, $m_i = |M_i|$ denotes the length of M_i , and analogously $p_i = |P_i|$ denotes the length of P_i , for all i .

Notice that $m_1 + \cdots + m_\ell = k$ is the total number of left-to-right maxima of π .

Proposition 12. For every permutation $\pi \in S_n$, $\pi \neq id_n$, there exists a child τ of id_n in T_n such that $T(\pi)$ and $T(\tau)$ are isomorphic.




Figure 2: The permutation τ described in the proof of Proposition 12.

Proof. Let $\pi = M_1P_1 \cdots M_{\ell-1}P_{\ell-1}M_\ell \in S_n$. Define τ as the permutation in which the elements $m_1 + p_1, m_1 + p_1 + m_2 + p_2, \dots, m_1 + p_1 + \cdots + m_{\ell-1} + p_{\ell-1}$, are in the positions $m_1, m_1 + p_1 + m_2, \dots, m_1 + p_1 + m_2 + \cdots + m_{\ell-2} + p_{\ell-2} + m_{\ell-1}$, respectively, while all the other elements are in increasing order. We can see an example of this construction in Fig. 2. Therefore, τ and π have their left-to-right maxima in the same positions, thus by Lemma 10 the trees $T(\pi)$ and $T(\tau)$ are isomorphic.

We are left with showing that τ is a child of id_n in T_n . Since $p_1 \neq 0$, then $\tau \neq id_n$, so we only need to check that $\mathbf{B}(\tau) = id_n$. Observe that the elements $m_1 + p_1 + \cdots + m_i + p_i$ are the last left-to-right maxima of their blocks in τ , for every $i = 1, \dots, \ell - 1$, and all the elements before the positions $m_1 + p_1 + \cdots + m_i + p_i$ are smaller than or equal to $m_1 + p_1 + \cdots + m_i + p_i$. Therefore $\mathbf{B}(\tau) = id_n$, because the $m_1 + p_1 + \cdots + m_i + p_i$'s are the only elements moved by bubblesort, and they are moved to their correct position. \square

3.2 The skeleton of the tree of preimages

Here we describe how the “shape” of any tree $T(\pi)$ is completely determined by a small piece of information about π , which we encapsulate in its *label*.

Definition 13. The *label* of a permutation σ is the pair (k, m_ℓ) , where k and m_ℓ are defined as in Definition 11.¹ The *skeleton* of a tree $T(\pi)$ is obtained from $T(\pi)$ by replacing each permutation at a node with its label. Fig. 3 shows the skeleton of the tree $T(2134)$, and can be compared with the subtree $T(2134)$ of $T(1234)$ in Fig. 1.




Figure 3: The skeleton of the tree $T(2134)$.

Observation 14. *Since they have their left-to-right maxima in the same positions, the permutations π and τ of Proposition 12 have the same label. (It can also be observed that the trees $T(\pi)$ and $T(\tau)$ have the same skeleton. This follows by recursively applying Corollary 6, as in the proof of Lemma 10.)*

Given a permutation π , we can determine the skeleton of $T(\pi)$ using only the pair (k, m_ℓ) . Specifically, it is the tree with root labeled by (k, m_ℓ) , and whose children (and recursively, descendants) are obtained as described in the next proposition.

Proposition 15. *Let $\pi \in S_n$ with label (k, m_ℓ) . Let T be the skeleton of $T(\pi)$. Then the root of T has label (k, m_ℓ) and its children have the following labels:*

- for every $h = 0, \dots, m_\ell - 2$:
 - for every $i = 1, \dots, k - 1 - h$, there are $\binom{k-2-h}{i-1}$ children with label $(k - i, h)$;
- if $\pi \neq id_n$, we also have the case corresponding to $h = m_\ell - 1$:
 - for every $i = 0, \dots, k - m_\ell$, there are $\binom{k-m_\ell}{i}$ children with label $(k - i, m_\ell - 1) = (k - i, h)$.

¹In particular, by definition of k and m_ℓ , the first component of a label is always at least as large as the second, with equality only in the case of the identity permutations.

Proof. We want to find the number of preimages of π with any given label. If $m_\ell = 0$, then π does not end with its maximum, hence it has no preimage. Thus the root of T has no children, and our claim vacuously holds.

Suppose that $m_\ell > 0$. This means that $\pi = \pi_1 \cdots \pi_{n-m_\ell}(n-m_\ell+1) \cdots n$. We can apply the procedure described in Section 2 to find the preimages of π . From this procedure we can see that, if $\pi \neq id_n$, then its preimages can only have labels (k', h) with $k' \leq k$ and $h < m_\ell$, which corresponds to the labels listed in the above statement.

If instead $\pi = id_n$, then it has label (n, n) and its preimages can only have labels (k', h) , with $k' \leq n$ and $h \leq n$, $h \neq n-1$. Indeed, we cannot obtain a preimage of id_n with $h = n-1$, because that would mean that only the element 1 is not part of the last sequence of left-to-right maxima, which is impossible. Instead we can have $h = n$, but only by leaving id_n unchanged.

To obtain a permutation with label $(k-i, h)$ with $h < m_\ell - 1$, referring to the procedure of Section 2, we are forced to leave unchanged (i.e. not to swap) all the elements from n down to $n-h+1$, then to swap $n-h$ with $n-h-1$. We are allowed to do so, because they are all left-to-right maxima, since $h < m_\ell - 1$.

After these steps, we obtain $\pi'(n-h-1)(n-h+1)(n-h+2) \cdots n$, with² $\pi' = \pi_1 \cdots \pi_{n-h-2}(n-h)$. Note that π' has $k-h-1$ left-to-right maxima. Moreover, there is a bijection between the preimages of π' and the preimages of π ending with the suffix $(n-h-1)(n-h+1)(n-h+2) \cdots n$, which consists of just appending the suffix $(n-h-1)(n-h+1)(n-h+2)$. Under this bijection, if a preimage of π' has $k-i-h$ left-to-right maxima (for some i such that $1 \leq i \leq k-1-h$), then the corresponding preimage of π has label $(k-i, h)$.

From Corollary 7, the number of preimages of π' with $k-i-h$ left-to-right maxima is $\binom{k-h-2}{k-i-h-1} = \binom{k-h-2}{i-1}$, for every $h = 0, \dots, m_\ell - 2$ and $i = 1, \dots, k-1-h$. Exploiting the above bijection, this prove the first item of our proposition.

Consider now the case $h = m_\ell - 1$, $\pi \neq id_n$. Then, applying the procedure of Section 2, we are forced to leave unchanged all the elements from n down to $n-m_\ell+2$, then swap $n-m_\ell+1$ with π_{n-m_ℓ} . Note that π_{n-m_ℓ} is an element of π which is not a left-to-right maximum, so if we define $\pi' = \pi_1 \cdots \pi_{n-m_\ell-1}(n-m_\ell+1)$, we have that π' has $k-m_\ell+1 = k-h$ left-to-right maxima. An argument analogous to the one we have used for the case $h < m_\ell - 1$ shows that the number of preimages of π with label $(k-i, m_\ell - 1)$ are $\binom{k-m_\ell}{k-i-m_\ell} = \binom{k-m_\ell}{i}$.

Finally, note that, if $\pi = id_n$, then there is an additional preimage, which is id_n , with label (n, n) . However it does not correspond to a child of id_n , because Definition 9 prevents a permutation from being a child of itself. \square

Corollary 16. *Let π be a permutation of size n with label (k, m_ℓ) such that $\pi \neq id_n$. Then $T(\pi)$ has depth m_ℓ . In addition, for every $n \geq 1$, T_n has depth $n-1$.*

Proof. We prove the statement by induction on m_ℓ . If $m_\ell = 0$, then $T(\pi)$ consists only of the root, and so has depth 0, as required. If $m_\ell \geq 1$, then by Proposition 15 the root

²Note that π' is not a permutation, but just a sequence of distinct integers. However, as we have seen in Lemma 2, it still makes sense to consider \mathbf{B} on such sequences.

of $T(\pi)$ has children whose labels are of the form (k', h) , for every $0 \leq h \leq m_\ell - 1$ and some k' . By induction hypothesis, the subtree rooted at each child with label (k', h) has depth h . It follows that $T(\pi)$ has depth $1 + (m_\ell - 1) = m_\ell$, since the maximum value of h is $m_\ell - 1$.

We now consider T_n . If $n = 1$ then T_n consists of a single node and the statement is true. Otherwise, if $n > 1$, then (again by Proposition 15) the root of T_n has children with labels (k', h) for every $0 \leq h \leq n - 2$ and some k' . Since these children are not the identity permutation, we can apply the first part of this corollary to them. We obtain that each child with label (h, k') is the root of a subtree of depth h . Since the maximum value of h is $n - 2$, T_n has depth $1 + (n - 2) = n - 1$. \square

Corollary 17. *For any given node $\pi \neq id_n$ in T_n , either half of its children are leaves or all of its children are leaves.*

Proof. Let (k, m_ℓ) be the label of $\pi \neq id_n$. If $m_\ell = 0$, the statement is vacuously true, because π is a leaf. Otherwise, if $m_\ell = 1$, then by Proposition 15 we have that all of its children have label $(k', 0)$ for some k' , and so they are all leaves.

Finally, if $m_\ell > 1$, then (again by Proposition 15) the number of children of π which are leaves, that is with label $(k', 0)$ for some k' , is

$$\sum_{i=1}^{k-1} \binom{k-2}{i-1} = \sum_{j=0}^{k-2} \binom{k-2}{j} = 2^{k-2}.$$

By Corollary 7, π has 2^{k-1} preimages, or equivalently it has 2^{k-1} children (since $\pi \neq id_n$). This proves our statement. \square

Notice that, for $\pi = id_n$ (with label (n, n)), it is still true that it has 2^{n-2} children which are leaves and 2^{n-1} preimages, but the total number of children is now $2^{n-1} - 1$ (id_n being a preimage of itself, but not one of its children).

3.3 The inverse problem: deciding if a tree is isomorphic to $T(\pi)$ for some π

We now consider the following problem: given a (rooted unlabeled) tree T , does there exist a permutation π such that T coincides with the (unlabeled) skeleton of $T(\pi)$? This problem can be easily solved thanks to our previous results on the labels of the nodes of $T(\pi)$. If T consists of just a leaf, then of course T is isomorphic to $T(\pi)$ for some π (just take $\pi = 1$ or any permutation of size at least 2 not ending with its maximum). So, assume that T has depth at least 1.

The first step is to determine the label of a candidate π . By Corollary 7, we can immediately say that, if the root of T has neither 2^{k-1} nor $2^{k-1} - 1$ children, for some $k > 0$, then T cannot be the (unlabeled) skeleton of any $T(\pi)$.

If the root of T has $2^{k-1} - 1$ children, then the only candidate permutation is $\pi = id_k$. In particular, by Corollary 16, it is necessary that T has depth $k - 1$. We can then use Proposition 15 to check if the number of children of every node of T matches with the numbers given in that proposition.

Otherwise, suppose that the number of children of the root of T is 2^{k-1} for some k . Let m_ℓ be the depth of T . By Corollary 7 and Corollary 16, we know that a permutation needs to have label (k, m_ℓ) for $T(\pi)$ to have 2^{k-1} children of its root and depth m_ℓ . Therefore we can use Proposition 15 to check if the number of children of every node of T matches with the numbers given in Proposition 15 for a permutation $\pi \neq id_n$ with label (k, m_ℓ) .

The next proposition summarizes the above discussion.

Proposition 18. *Let T be a (rooted unlabeled) tree, let i be the number of children of the root of T , and m_ℓ be the depth of T . Then*

- if $i = 2^{m_\ell} - 1$, then T may only coincide with the (unlabeled) skeleton of $T_{m_\ell+1}$;
- if there exists a positive integer k such that $i = 2^{k-1}$, then T may only coincide with the (unlabeled) skeleton of a permutation π with label (k, m_ℓ) ;
- in all the other cases, T does not coincide with the unlabeled skeleton of any permutation.

4 Heights of nodes and leaves in T_n

4.1 Nodes

Recall that the height of a node in a rooted tree is the number of edges on the path connecting that node to the root. The height of a node of the tree T_n corresponds to the number of passes of Bubblesort needed to sort the permutation at this node. Therefore, we can refer to [1, Prop. 17] to find information on the number of nodes of T_n .

Proposition 19 ([1]). *The set of permutations of size n sorted by at most k passes of Bubblesort is the set $\text{Av}_n(\Gamma_{k+2})$, where Γ_k is the set of all permutations of size k whose final element is 1^3 . As a consequence, setting $\varphi_n^{(k)} = |\text{Av}_n(\Gamma_k)|$, the number of nodes at height at most k in T_n is given by $\varphi_n^{(k+2)} = (k+1)^{n-k-1}(k+1)!$*

We can thus immediately deduce the number of nodes at a given height in T_n .

Corollary 20. *The number $f_n^{(k)}$ of nodes at height k in T_n is given by*

$$f_n^{(k)} = \varphi_n^{(k+2)} - \varphi_n^{(k+1)} = (k+1)^{n-k-1}(k+1)! - k^{n-k}k! = k! \cdot ((k+1)^{n-k} - k^{n-k}).$$

The first lines of the infinite triangular matrix of the coefficients $f_n^{(k)}$ are given in Table 1. This is sequence A056151 in [23].

We notice that the elements on the diagonal of Table 1 are the factorial numbers, more specifically $f_n^{(n-1)} = (n-1)!$. Indeed, the set of permutations of size n needing the maximum number of passes of Bubblesort to be sorted (that is, $n-1$ passes) is the set of permutations of size n ending with 1, whose cardinality is clearly $(n-1)!$.

³We warn the reader that we have made a slight change of notation with respect to [1] here; more specifically, our set Γ_k is Γ_{k-2} in [1].

$n \backslash k$	0	1	2	3	4	5
1	1					
2	1	1				
3	1	3	2			
4	1	7	10	6		
5	1	15	38	42	24	
6	1	31	130	222	216	120

Table 1: Number of nodes in T_n having height k .

From the expression of $\varphi_n^{(k)}$ in Proposition 19, we can derive the asymptotic behavior of the average height of a node in T_n . This analysis is described in [19, Theorem 7.14] and follows easily from the asymptotic behavior of the Ramanujan P-function (see [19, Table 4.11] or [21, p. 119-120]), which we state in Lemma 21 below. We then reproduce the analysis of [19, Theorem 7.14], as a preparation for Proposition 27 below.

Lemma 21 ([19]). *The Ramanujan P-function, defined by $P(n) = \sum_{k=0}^{n-1} \frac{k!k^{n-k}}{n!}$, behaves asymptotically as $\sqrt{\frac{\pi n}{2}} + O(1)$.*

Proposition 22 ([19]). *The average height of a node in T_n is asymptotically equal to $n - \sqrt{\frac{\pi n}{2}} + O(1)$.*

Proof. The average height of a node in T_n is given by

$$H_n := \frac{1}{n!} \sum_{k=1}^{n-1} \text{number of nodes of height at least } k \text{ in } T_n,$$

each node at height k contributing indeed exactly k times to this sum. Writing the number of nodes of height at least k in T_n as the difference of $n!$ (the total number of nodes) and the number of nodes of height at most $k - 1$ in T_n , we then compute

$$H_n = \frac{1}{n!} \sum_{k=1}^{n-1} (n! - \varphi_n^{(k+1)}) = \frac{1}{n!} \sum_{k=1}^{n-1} (n! - k^{n-k}k!) = (n-1) - \sum_{k=0}^{n-1} \frac{k^{n-k}k!}{n!} = (n-1) - P(n),$$

proving our claim. □

Recall the (obvious) fact that T_n contains $n!$ nodes. With Proposition 19 and Corollary 20, we have refined this counting according to the height of the nodes in T_n . We now address the analogous problems in $T(\pi)$ for $\pi \neq id_n$. More precisely, given a permutation π (of size n) having label (k, m_ℓ) , we determine an expression for the number of nodes of $T(\pi)$ (which does not depend on n). This expression is a summation formula in which each summand counts nodes in $T(\pi)$ of a prescribed height.

Lemma 23. *Let π and τ be two permutations having labels (k, m_ℓ) and $(k, m_\ell - 1)$, respectively, with $1 \leq m_\ell \leq k - 1$. Then the tree obtained by removing the leaves at height m_ℓ in $T(\pi)$ is isomorphic to $T(\tau)$.*

Proof. Remember that, by Corollary 16, $T(\pi)$ has height m_ℓ and $T(\tau)$ has height $m_\ell - 1$. The proof is by induction on m_ℓ .

If $m_\ell = 1$, then $T(\tau)$ consists of the single node τ , while $T(\pi)$ has height 1, therefore the statement is true.

Now let $m_\ell \geq 2$, and suppose that the statement is true for $m_\ell - 1$. We will show that there is a bijective correspondence between the children of τ and the children of π such that the subtree rooted at a child of τ is isomorphic to the subtree rooted at the corresponding child of π , after removing the leaves at height m_ℓ (if any).

Proposition 15 allows us to determine the labels of the children of τ and π in $T(\tau)$ and $T(\pi)$, respectively. Specifically, τ and π have the same number of children with labels $(k - i, h)$, for every $h = 0, \dots, m_\ell - 3$ and every $i = 0, \dots, k - 1 - h$. Regarding the remaining children, we have that the number of children of τ labeled $(k - i, m_\ell - 2)$ is equal to the sum of the number of children of π labeled $(k - i, m_\ell - 2)$ and $(k - i, m_\ell - 1)$, for every $i = 0, \dots, k - m_\ell + 1$. This induces the announced bijective correspondence between the children of τ in $T(\tau)$ and those of π in $T(\pi)$.

The children of τ and π with the same labels give isomorphic subtrees by Proposition 15. In addition, if this label is $(k - i, h)$ for some $h \leq m_\ell - 2$ (and some suitable i), then the subtrees contain no leaf at height m_ℓ in $T(\tau)$ or $T(\pi)$ (again by Corollary 16), ensuring our claim restricted to such children of π and τ .

Therefore, we are left with considering a child of π in $T(\pi)$ with label $(k - i, m_\ell - 1)$, to which corresponds a child of τ in $T(\tau)$ of label $(k - i, m_\ell - 2)$. We can apply the inductive hypothesis to such children of π and τ , thus obtaining that each subtree of $T(\tau)$ rooted at a child of τ with label $(k - i, m_\ell - 2)$ is isomorphic to a subtree of $T(\pi)$ rooted at a child of π with label $(k - i, m_\ell - 1)$ after removing the leaves at height $m_\ell - 1$ (in the subtree, *i.e.* at height m_ℓ in $T(\pi)$). This concludes the proof. \square

Proposition 24. *For a permutation π having label (k, m_ℓ) , different from an identity permutation, the number of nodes of the tree $T(\pi)$ of its preimages under \mathbf{B} is*

$$N(k, m_\ell) = \sum_{j=0}^{m_\ell} j!(j+1)^{k-j}. \quad (1)$$

Moreover, each summand in Eq. (1) records the contribution of each level of $T(\pi)$. In other words, denoting with $N_j(k, m_\ell)$ the number of nodes at height j in $T(\pi)$, we have that $N_j(k, m_\ell) = j!(j+1)^{k-j}$.

Proof. In order to prove Eq. (1) we proceed by induction on m_ℓ . If $m_\ell = 0$, then π has no children, hence $N(k, 0) = 1$, which is consistent with Eq. (1).

Now suppose that Eq. (1) holds when the cardinality of the longest suffix of left-to-right maxima of π is strictly smaller than m_ℓ . Recalling Proposition 15, we have the following recursive expression for the number of nodes of $T(\pi)$:

$$\begin{aligned}
N(k, m_\ell) &= 1 + \sum_{h=0}^{m_\ell-2} \sum_{i=1}^{k-1-h} \binom{k-2-h}{i-1} N(k-i, h) + \sum_{i=0}^{k-m_\ell} \binom{k-m_\ell}{i} N(k-i, m_\ell-1) \\
&= 1 + \sum_{h=0}^{m_\ell-2} \sum_{i=1}^{k-1-h} \binom{k-2-h}{i-1} \sum_{j=0}^h j!(j+1)^{k-i-j} + \sum_{i=0}^{k-m_\ell} \binom{k-m_\ell}{i} \sum_{j=0}^{m_\ell-1} j!(j+1)^{k-i-j} \\
&= 1 + \sum_{h=0}^{m_\ell-2} \sum_{j=0}^h j!(j+1)^{k-j-1} \sum_{i=0}^{k-2-h} \binom{k-2-h}{i} (j+1)^{-i} \\
&\quad + \sum_{j=0}^{m_\ell-1} j!(j+1)^{k-j} \sum_{i=0}^{k-m_\ell} \binom{k-m_\ell}{i} (j+1)^{-i} \\
&= 1 + \sum_{h=0}^{m_\ell-2} \sum_{j=0}^h j!(j+1)^{k-j-1} \left(1 + \frac{1}{j+1}\right)^{k-2-h} + \sum_{j=0}^{m_\ell-1} j!(j+1)^{k-j} \left(1 + \frac{1}{j+1}\right)^{k-m_\ell} \\
&= 1 + \sum_{h=0}^{m_\ell-2} \sum_{j=0}^h j!(j+1)^{h-j+1} (j+2)^{k-2-h} + \sum_{j=0}^{m_\ell-1} j!(j+1)^{m_\ell-j} (j+2)^{k-m_\ell}.
\end{aligned}$$

We then exchange the order of the two sums in the middle term of the last expression, use the geometric sum formula and we get:

$$\begin{aligned}
N(k, m_\ell) &= 1 + \sum_{j=0}^{m_\ell-2} j!(j+1)^{1-j} (j+2)^{k-2} \sum_{h=j}^{m_\ell-2} (j+1)^h (j+2)^{-h} + \sum_{j=0}^{m_\ell-1} j!(j+1)^{m_\ell-j} (j+2)^{k-m_\ell} \\
&= 1 + \sum_{j=0}^{m_\ell-2} j!(j+1)(j+2)^{k-1-j} - \sum_{j=0}^{m_\ell-2} j!(j+1)^{m_\ell-j} (j+2)^{k-m_\ell} \\
&\quad + \sum_{j=0}^{m_\ell-1} j!(j+1)^{m_\ell-j} (j+2)^{k-m_\ell} \\
&= 1 + \sum_{j=1}^{m_\ell-1} j!(j+1)^{k-j} + m_\ell!(m_\ell+1)^{k-m_\ell} = \sum_{j=0}^{m_\ell} j!(j+1)^{k-j},
\end{aligned}$$

which gives Eq. (1).

Concerning the evaluation of $N_j(k, m_\ell)$, Lemma 23 implies that $N_j(k, m_\ell) = N_j(k, m_\ell - 1)$, for all $j \leq m_\ell - 1$. By a repeated application of the lemma, we get that $N_j(k, m_\ell) = N_j(k, j) = N(k, j) - N(k, j - 1) = j!(j + 1)^{k-j}$, as desired. \square

4.2 Leaves

In the tree T_n the leaves represent permutations that cannot be obtained as output of Bubblesort, *i.e.*, which do not belong to the image of \mathbf{B} . We saw just after Lemma 2 that

these permutations are those not ending with their maximum, so that the total number of leaves in T_n is given by $(n - 1) \cdot (n - 1)!$.

Our next result is a closed formula for the number of leaves at height k in T_n , for any $k \leq n - 1$. To this aim, we make use of the so-called *ECO method*, illustrated in [2] and further developed and employed by many authors (see for instance [17]). We will not give a detailed description of this method here, since our application is simple enough to be outlined directly.

Recall that leaves in T_n correspond to permutations whose last element is not the maximum. Thus, denoting with $\text{Av}_n^*(\Gamma_k)$ the set of permutations of size n avoiding Γ_k and such that their last element is different from n , we are interested in the coefficients $\gamma_n^{(k)} = |\text{Av}_n^*(\Gamma_k)|$, since $\gamma_n^{(k+2)}$ gives the number of leaves at height at most k in T_n .

Proposition 25. *For all n, k , we have*

$$\gamma_n^{(k)} = \begin{cases} (n - 1)(n - 1)! & n < k, \\ (k - 2)(k - 1)^{n-k}(k - 1)! & n \geq k. \end{cases}$$

Proof. We consider the following general procedure to generate all permutations of size n . Given any permutation of size $n - 1$, construct n different permutations of size n by adding a new rightmost element k , for any choice of k between 1 and n , and suitably rescaling the other elements (namely, all elements of the starting permutation which are greater than or equal to k are increased by 1, whereas all the remaining elements are left untouched). It is immediate to realize that, starting from the set of all permutations of size $n - 1$, the above procedure generates exactly once every permutation of size n .

We now adapt the above construction to our setting. Every permutation of $\text{Av}_n^*(\Gamma_k)$ can be obtained from a permutation of $\text{Av}_{n-1}(\Gamma_k)$ by adding a suitable rightmost element. More specifically, we cannot add n (because we require that our permutation does not end with its maximum); moreover, if $n \geq k$, we cannot add any element between 1 and $n - k + 1$ as well (otherwise we would create one of the forbidden patterns belonging to Γ_k). On the other hand, any of the remaining elements is allowed and generates a valid permutation. This means that every permutation in $\text{Av}_{n-1}(\Gamma_k)$ generates $k - 2$ distinct permutations of $\text{Av}_n^*(\Gamma_k)$ and every permutation in $\text{Av}_n^*(\Gamma_k)$ is obtained in this way exactly once. We thus deduce that, when $n \geq k$,

$$\gamma_n^{(k)} = (k - 2)\varphi_{n-1}^{(k)} = (k - 2)(k - 1)^{n-k}(k - 1)!,$$

whereas for $n < k$ we have that $\gamma_n^{(k)} = (n - 1)(n - 1)!$, which concludes the proof. \square

Corollary 26. *The number $g_n^{(k)}$ of leaves of T_n at height k is given by*

$$g_n^{(k)} = k!(k(k + 1)^{n-k-1} - (k - 1)k^{n-k-1}).$$

Proof. Just observe that $g_n^{(k)} = \gamma_n^{(k+2)} - \gamma_n^{(k+1)}$ and that the maximum height of a node of T_n is $n - 1$, so we are only interested in the case $n \geq k + 1$ of the previous proposition. \square

As in the case of nodes, Proposition 25 allows us to derive the asymptotic behavior of the average height of a leaf in T_n .

Proposition 27. *The average height of a leaf in T_n is asymptotically equal to $n - \sqrt{\frac{\pi n}{2}} + O(1)$.*

Proof. As in the proof of Proposition 25, we have that the average height of a leaf in T_n is

$$\begin{aligned} G_n &= \frac{1}{(n-1)(n-1)!} \sum_{k=1}^{n-1} \text{number of leaves of height at least } k \text{ in } T_n \\ &= \frac{1}{(n-1)(n-1)!} \sum_{k=1}^{n-1} ((n-1)(n-1)! - \gamma_n^{(k+1)}) = (n-1) - \sum_{k=1}^{n-1} \frac{(k-1)k^{n-k-1}k!}{(n-1)(n-1)!} \\ &= (n-1) - \frac{n}{n-1} \sum_{k=1}^{n-1} \frac{k^{n-k}k!}{n!} + \frac{1}{n-1} \sum_{k=1}^{n-1} \frac{k^{n-1-k}k!}{(n-1)!} \\ &= (n-1) - \frac{nP(n)}{n-1} + \frac{P(n-1)+1}{n-1}, \end{aligned}$$

and the asymptotic behavior of the Ramanujan P-function yields the announced result. \square

In the same manner as we have done for the nodes, we now address the analogous problem of counting the leaves in $T(\pi)$, for $\pi \neq id_n$. More precisely, given a permutation π (of size n) having label (k, m_ℓ) , we determine an expression for the number of leaves of $T(\pi)$ (which does not depend on n but only on the label (k, m_ℓ)). This expression is a summation formula in which each summand counts the leaves of a prescribed height in $T(\pi)$.

Proposition 28. *For a permutation π having label (k, m_ℓ) , different from an identity permutation, the number of leaves of the tree $T(\pi)$ of its preimages under \mathbf{B} is*

$$L(k, m_\ell) = \sum_{j=1}^{m_\ell-1} j!j(j+1)^{k-j-1} + m_\ell!(m_\ell+1)^{k-m_\ell}. \quad (2)$$

Moreover, each summand in Eq. (2) records the contribution of each level of $T(\pi)$. In other words, denoting with $L_j(k, m_\ell)$ the number of leaves at height j in $T(\pi)$, we have that $L_j(k, m_\ell) = j!j(j+1)^{k-j-1}$ for $j < m_\ell$, and $L_{m_\ell}(k, m_\ell) = m_\ell!(m_\ell+1)^{k-m_\ell}$.

Proof. The proof of Eq. (2) is by induction, following the exact same steps as the proof of Proposition 24. The recursive equation for the number of leaves in $T(\pi)$, which is needed in the inductive step of the proof, is again obtained from Proposition 15. It actually differs from the one for nodes in the proof of Proposition 24 only by the initial term 1 (accounting for the root node); namely for $m_\ell \geq 1$, we have

$$L(k, m_\ell) = \sum_{h=0}^{m_\ell-2} \sum_{i=1}^{k-1-h} \binom{k-2-h}{i-1} L(k-i, h) + \sum_{i=0}^{k-m_\ell} \binom{k-m_\ell}{i} L(k-i, m_\ell-1),$$

and for $m_\ell = 0$ it holds that $L(k, 0) = 1$. From there, the same steps of computations as in the proof of Proposition 24 (followed by additional elementary simplifications) yield, for $m_\ell \geq 1$:

$$L(k, m_\ell) = \sum_{j=1}^{m_\ell-1} j!j(j+1)^{k-j-1} + m_\ell!(m_\ell+1)^{k-m_\ell},$$

as claimed.

We now move to the claimed expression for $L_j(k, m_\ell)$. We shall first establish it for $j = m_\ell$, then for $j = m_\ell - 1$, and then for smaller j iterating the argument.

We first note that all the nodes of $T(\pi)$ at height m_ℓ are leaves (since m_ℓ is the height of this tree). Using Proposition 24, we therefore have $L_{m_\ell}(k, m_\ell) = N_{m_\ell}(k, m_\ell) = m_\ell!(m_\ell+1)^{k-m_\ell}$. As a consequence, the total number of leaves having height at most $m_\ell - 1$ in $T(\pi)$ is $\sum_{j=1}^{m_\ell-1} j!j(j+1)^{k-j-1}$.

Next, we claim that the number of leaves having height at most $m_\ell - 2$ in $T(\pi)$ is $\sum_{j=1}^{m_\ell-2} j!j(j+1)^{k-j-1}$. From this claim, the announced formula $L_{m_\ell-1}(k, m_\ell) = (m_\ell - 1)!(m_\ell - 1)m_\ell^{k-m_\ell}$ immediately follows by taking the difference.

To prove our claim, we use Lemma 23. This lemma indeed implies that $L_j(k, m_\ell) = L_j(k, m_\ell - 1)$, for all $j \leq m_\ell - 2$. This shows that the number of leaves having height at most $m_\ell - 2$ in $T(\pi)$ is the same as the number of leaves having height at most $m_\ell - 2$ in $T(\sigma)$ for σ a permutation with label $(k, m_\ell - 1)$. The latter is equal to $L(k, m_\ell - 1) - L_{m_\ell-1}(k, m_\ell - 1)$, hence equal to $\sum_{j=1}^{m_\ell-2} j!j(j+1)^{k-j-1}$ as established earlier, thus proving our claim.

We are now left with showing that $L_h(k, m_\ell) = h!h(h+1)^{k-h-1}$ for $h \leq m_\ell - 2$. We proceed iteratively, for decreasing values of h . At each step, the reasoning is similar to the above case for $h = m_\ell - 1$. We first use Lemma 23 (several times, as in the proof of Proposition 24) to argue that the number of leaves having height at most $h - 1$ in $T(\pi)$ is the same as the number of leaves having height at most $h - 1$ in $T(\sigma)$ for σ a permutation with label (k, h) . This number is $\sum_{j=1}^{h-1} j!j(j+1)^{k-j-1}$. Then, $L_h(k, m_\ell)$ is the difference between $L(k, m_\ell) - \sum_{h+1 \leq j \leq m_\ell} L_j(k, m_\ell)$ and the above quantity. The result follows from the formulas previously established for $L_j(k, m_\ell)$ for $j \geq h + 1$. \square

Remark 29. Combining Propositions 24 and 28 tells us that, for π a permutation of label (k, m_ℓ) , at height $j < m_\ell$ in $T(\pi)$, the ratio between the number of leaves and the number of nodes is $\frac{j}{j+1}$ (equivalently, the ratio between the number of internal nodes and the number of nodes is $\frac{1}{j+1}$).

References

- [1] Michael H. Albert, M. D. Atkinson, Mathilde Bouvel, Anders Claesson, and Mark Dukes. On the inverse image of pattern classes under bubble sort. *Journal of Combinatorics*, 2(2):231–243, 2011.

- [2] E. Barucci, A. Del Lungo, E. Pergola, and R. Pinzani. ECO: a methodology for the enumeration of combinatorial objects. *Journal of Difference Equations and Applications*, 5:435–490, 1999.
- [3] Miklós Bóna. A survey of stack-sorting disciplines. *The Electronic Journal of Combinatorics*, 9(2):Article A1, 2003.
- [4] Mireille Bousquet-Mélou. Sorted and/or sortable permutations. *Discrete Mathematics*, 225(1–3):25–50, 2000.
- [5] Lapo Cioni and Luca Ferrari. Characterization and enumeration of preimages under the queuesort algorithm. In J. Nešetřil, G. Perarnau, J. Rué, and O. Serra, editors, *Extended Abstracts EuroComb 2021. Trends in Mathematics*, volume 14 of *Birkhäuser, Cham*, 2021.
- [6] Lapo Cioni and Luca Ferrari. Preimages under the queuesort algorithm. *Discrete Mathematics*, 344:Article A112561, 2021.
- [7] Anders Claesson and Henning Ulfarsson. Sorting and preimages of pattern classes. *DMTCS Proceedings*, AR:595–606, 2012.
- [8] Colin Defant. Preimages under the stack-sorting algorithm. *Graphs and Combinatorics*, 33:103–122, 2017.
- [9] Colin Defant. Fertility numbers. *Journal of Combinatorics*, 11:527–548, 2020.
- [10] Colin Defant. Polyurethane toggles. *Electronic Journal of Combinatorics*, 27(2):#P2.46, 2020.
- [11] Colin Defant. Stack-sorting preimages of permutation classes. *Séminaire Lotharingien de Combinatoire*, 82:Article B82b, 2020.
- [12] Colin Defant. Enumeration of stack-sorting preimages via a decomposition lemma. *Discrete Mathematics and Theoretical Computer Science*, 22(2):Article 3, 2021.
- [13] Colin Defant. Fertility monotonicity and average complexity of the stack-sorting map. *European Journal of Combinatorics*, 93:Article 103276, 2021.
- [14] Colin Defant, Michael Engen, and Jordan A. Miller. Stack-sorting, set partitions, and Lassalle’s sequence. *Journal of Combinatorial Theory, Series A*, 175:Article 105275, 2020.
- [15] Colin Defant and Noah Kravitz. Promotion sorting. *Order*, 2022, to appear.
- [16] Colin Defant and James Propp. Quantifying noninvertibility in discrete dynamical systems. *Electronic Journal of Combinatorics*, 27(3):#P3.51, 2020.
- [17] L. Ferrari, E. Pergola, R. Pinzani, and S. Rinaldi. Some applications arising from the interactions between the theory of Catalan-like numbers and the ECO method. *Ars Combinatoria*, 99:109–128, 2011.
- [18] Luca Ferrari. Sorting with stacks and queues: some recent developments. Keynote address at the on-line conference *Permutation Patterns 2021*. Available at <https://www.youtube.com/watch?v=cTT9t5gddmE>.

- [19] P. Flajolet and R. Sedgewick. *An Introduction to the Analysis of Algorithms, second edition*. Addison-Wesley, 2013.
- [20] Eliot Hodges. On promotion and quasi-tangled labelings of posets. [arXiv:2208.08665](https://arxiv.org/abs/2208.08665).
- [21] Donald E. Knuth. *The Art of Computer Programming, Vol. 1: Fundamental Algorithms*. Addison-Wesley, third edition, 1997.
- [22] Hjalti Magnusson. Sorting operators and their preimages. Master's thesis, Reykjavik University, 2013.
- [23] N. J. A. Sloane. The Online Encyclopedia of Integer Sequences. oeis.org.
- [24] Jiang Tao, Li Ming, and Paul M.B. Vitanyi. Average-case analysis of algorithms using Kolmogorov complexity. *Journal of Computer Science and Technology*, 15:402–408, 2000.
- [25] Vincent Vatter. *Permutation Classes*, chapter 12 of *The Handbook of Enumerative Combinatorics*, pages 753–834. Chapman and Hall/CRC Press, 2015.