

POP-STACK SORTING AND ITS IMAGE: PERMUTATIONS WITH OVERLAPPING RUNS.

ANDREI ASINOWSKI, CYRIL BANDERIER, SARA BILLEY, BENJAMIN HACKL
AND SVANTE LINUSSON

ABSTRACT. Pop-stack sorting is an important variation for sorting permutations via a stack. A single iteration of pop-stack sorting is the transformation $T: S_n \rightarrow S_n$ that reverses all the maximal descending sequences of letters in a permutation. We investigate structural and enumerative aspects of *pop-stacked permutations* – the permutations that belong to the image of S_n under T . This work is part of a project aiming to provide the full combinatorial analysis of sorting with a pop-stack, as it was successfully done for sorting with a stack (though, even in this case, some famous problems are still open). The first results already show that pop-stack sorting has a very rich combinatorial structure, and leads to surprising phenomena.

1. INTRODUCTION

In the seminal work of Knuth (see [10, Sec. 2.2.1]), permutations sortable by one stack (neatly illustrated via a sequence of n goods wagons on a railway line, with a side track which makes it possible to permute some wagons) were characterized as 231-avoidable permutations: this result is a cornerstone in the field of permutation patterns. Since that time, many results were also obtained e.g. on permutations sortable with 2 stacks, which offer many nice links with the world of planar maps [7, 15]. The analysis of sorting with 3 stacks remains an open problem. Numerous variants were considered (stacks in series, in parallel, etc.). In this article we focus on another important variant considered e.g. in [3, 4]: sorting with a pop-stack, i.e. the pop operation dequeues the full stack at once. Nota bene: if one does not impose the stack to contain only increasing sequences, then the process can be related to the famous pancake sorting problem (see [9]).

Let us describe our model more precisely. In this paper, all permutations are written in the one-line notation: thus, $a_1 a_2 \dots a_n$ is the bijective function on $\{1, 2, \dots, n\}$ that maps $i \mapsto a_i$ for $1 \leq i \leq n$. The graph of such a permutation is the point set $\{(i, a_i): 1 \leq i \leq n\}$.

In a permutation $a_1 a_2 \dots a_n$, the pair of adjacent letters, $a_i a_{i+1}$, is an *ascent* if $a_i < a_{i+1}$, and it is a *descent* if $a_i > a_{i+1}$. A *run* is a maximal sequence of consecutive ascending letters, and a *fall* is a maximal sequence of consecutive

Received March 22, 2019; revised June 7, 2019.

2000 *Mathematics Subject Classification*. Primary 05A15, 05A16, 68Q45, 68Q70.

A. Asinowski and B. Hackl were supported by the Austrian Science Fund (FWF), in the framework of the project *Analytic Combinatorics: Digits, Automata and Trees* (P 28466-N35).

descending letters. Each permutation can be split into runs or into falls in a unique way. We shall usually indicate it by bars between runs resp. falls: in fact, the bars that separate runs are precisely at the descents, and the bars that separate falls are precisely at the ascents. For example, the permutation 71853642 splits into runs as $7|18|5|36|4|2$, and into falls as $71|853|642$.

The *pop-stack sorting* is the following transformation T on permutations: given a permutation, it is split into falls, and then each fall is reversed, for example $71|853|642 \mapsto 17358246$. If one repeatedly applies T to a given permutation, then after several iterations it eventually will be *sorted*, that is, the identity permutation $\text{id} = 12 \dots n$ will be obtained, for example $71|853|642 \mapsto 1|73|5|82|4|6 \mapsto 1|3|752|84|6 \mapsto 1|32|5|74|86 \mapsto 1|2|3|54|76|8 \mapsto 12345678$. In fact, Ungar [14] proved that each permutation of size n is sorted by at most $n - 1$ iterations of T (this bound is tight). A permutation is *p-pop-stack-sortable* if it is sorted in at most p iterations of T . Avis and Newborn [4] showed that 1-pop-stack-sortable permutations are precisely layered permutations. Pudwell and Smith [13] found the generating function for the number of 2-pop-stack-sortable permutations (that is, permutations π that satisfy $T(T(\pi)) = \text{id}$), and Claesson and Guðmundsson [6] showed that the generating function for p -pop-stack-sortable permutations is rational for any fixed p . See also [1] for further recent developments.

In order to obtain further insights on the dynamics of this process (also called *flip-sort*¹), the key is to study the image of T , that is, the set of permutations that can be obtained from some permutation by a single iteration of T . In this article, we characterize the permutations that belong to the image of T by showing that these are precisely the permutations whose adjacent runs overlap (in the sense specified below). Next, we address the enumeration of such permutations. In particular, we show that for each fixed k , the generating function for permutations that belong to the image of T and have exactly k runs, is rational. Finally, we show that the asymptotic growth of the number of such permutations is superexponential.

2. STRUCTURAL CHARACTERIZATION OF POP-STACKED PERMUTATIONS

We say that a permutation π is *pop-stacked* if it belongs to the image of T , that is, if we have $\pi = T(\tau)$ for some permutation(s) τ . The next theorem is a structural characterization of pop-stacked permutations in terms of overlapping runs.

Theorem 1. *A permutation π is pop-stacked if and only if for each pair (R_i, R_{i+1}) of adjacent runs of π , we have $\min(R_i) < \max(R_{i+1})$.*

Remark. Graphically, this condition means that in such a permutation all the pairs of adjacent runs necessarily *overlap*, in the sense that the intersection $[\min(R_i), \max(R_i)] \cap [\min(R_{i+1}), \max(R_{i+1})]$ is not empty. Notice that the condition $\max(R_i) > \min(R_{i+1})$ is satisfied automatically because otherwise R_i and R_{i+1} are not two distinct runs. See Figure 1 for a general schematic drawing that represents the structure of pop-stacked permutations, and an example. Despite its very natural definition, this family of permutations was, to the best of our knowledge, never studied.

¹Animations are available at <https://lipn.fr/~banderier/Papers/popstack.html>.




Figure 1. Pop-stacked permutations: a general schematic drawing of permutation with overlapping runs, and an example (here, the permutation $(5, 7, 14|1, 2, 3, 11|4, 12, 15|9|6, 16, 17|8, 10, 13)$).

Proof. Necessity. Consider a permutation π with $\min(R_i) > \max(R_{i+1})$ for some pair (R_i, R_{i+1}) of its adjacent runs. Denote the positions occupied by the run R_i by a, \dots, b , and the positions occupied by the run R_{i+1} by $b+1, \dots, c$.

Assume for the sake of contradiction that $\pi = T(\tau)$. In τ , we have $\tau_b < \tau_{b+1}$ because otherwise the positions b and $b+1$ belong to the same fall in τ , and upon applying T we have $\pi_b < \pi_{b+1}$, which is impossible, because positions b and $b+1$ lie in different runs.

Therefore, if we consider the partition of τ into falls, then b is the last position of some fall Q_j , and $b+1$ is the first position of the next fall Q_{j+1} . When we apply T on τ , the values of Q_j are a subset of those of R_i , and the values of Q_{j+1} are a subset of those of R_{i+1} . Thus, we have a value in R_i (namely τ_b) smaller than a value in R_{i+1} (namely τ_{b+1}). This contradicts the assumption $\min(R_i) > \max(R_{i+1})$.

Sufficiency. Consider a permutation π with $\min(R_i) < \max(R_{i+1})$ for all pairs (R_i, R_{i+1}) of adjacent runs. Let τ be the permutation obtained from π by reversal of all its runs. Then the partition of τ into falls is the same as the partition of π into runs, and, therefore, τ is a (not necessarily unique) pre-image of π . \square

3. ENUMERATIVE ASPECTS: RATIONAL GENERATING FUNCTIONS

The number g_n of pop-stacked permutations of size n is 1, 1, 3, 11, 49, 263, 1653, 11877, 95991, 862047, 8516221, 91782159, 1071601285, 13473914281, 181517350571, 2608383775171 (for $n = 1, \dots, 16$). We added it as the sequence [A307030](#) in the On-Line Encyclopedia of Integer Sequences. This sequence is hard to compute, so a key idea is to find an additional parameter which can allow us to get further insights. In this section we show that, for each fixed k , the generating function for pop-stacked permutations with k runs is rational. In Section 4, we show a non-trivial asymptotic bound on the number of pop-stacked permutations.

For fixed k , let $g_{n,k}$ denote the number of pop-stacked permutations of size n with exactly k runs. The case $k = 1$ is trivial: for any size, the only permutation with one run is the identity permutation, and it is pop-stacked as it is e.g. the image of itself. Thus, we have $g_{n,1} = 1$ for each $n \geq 1$.

One key ingredient of our further results is the following encoding of permutations, which we call *scanline mapping*. Let π be a permutation with k runs. Let r_i be the index of the run in which the letter i lies. Consider the word $w(\pi) = r_{\pi^{-1}(1)} r_{\pi^{-1}(2)} \dots r_{\pi^{-1}(n)} \in \{1, \dots, k\}^n$. Visually, we scan the graph of π from the bottom to the top and for each point that we encounter in this order, we record to which run it belongs, see Figure 2.




Figure 2. The *scanline mapping* allows us to encode permutations by words: above, the permutation $\pi = 261453$ is encoded by the word $w(\pi) = 213221$ (the sequence of runs encountered if one reads the permutation bottom-to-top).

- Proposition 2.**
1. The scanline mapping $\pi \mapsto w(\pi)$ is injective.
 2. The scanline mapping $\pi \mapsto w(\pi)$ induces a bijection between permutations of size n with k runs and the sequences in $\{1, \dots, k\}^n$ in which there is an occurrence of $j + 1$ before an occurrence of j , for each j ($1 \leq j \leq n - 1$).
 3. The scanline mapping $\pi \mapsto w(\pi)$ induces a bijection between pop-stacked permutations of size n with k runs and the sequences in $\{1, \dots, k\}^n$ in which there is an occurrence of $j + 1$ before an occurrence of j , and also an occurrence of j before an occurrence of $j + 1$, for each j ($1 \leq j \leq n - 1$).

Proof (Sketch).

1. The positions of i ($1 \leq i \leq k$) in $w(\pi)$ are the values in the i th run of π . Thus, π is reconstructed from $w(\pi)$ uniquely.
2. If (and only if) for some j , all the occurrences of j in $w(\pi)$ are before the occurrences of $j + 1$, then the corresponding positions of π do not form two distinct runs and thus we do not get a permutation with k runs.
3. If (and only if) for some j , all the occurrences of $j + 1$ in $w(\pi)$ are before the occurrences of j , then all the values in the j th run of π are larger than all the values in the $(j + 1)$ st run, and thus these runs are not overlapping. \square

Proposition 2 can be used to obtain a formula for the case of two runs directly:

Proposition 3. For $n \geq 1$, the number of pop-stacked permutations of size n with exactly two runs, is $g_{n,2} = 2^n - 2n$.

Proof. By Proposition 2, $g_{n,2}$ is the number of sequences in $\{1, 2\}^n$ with an occurrence of 1 before an occurrence of 2, and also an occurrence of 2 before an occurrence of 1. There are $2n$ sequences that violate this condition (including the “all-1” and the “all-2” sequences). \square

Theorem 4. *Let $k \geq 1$ be fixed. Then $g_k(x) := \sum_{n \geq 1} g_{n,k} x^n$, the generating function for the number of pop-stacked permutations with exactly k runs, is rational.*

Proof. It is well known that the generating function of words recognized by an automaton is rational (see e.g. [8, Sec. I.4.2]). We use Proposition 2 to construct a deterministic automaton \mathcal{A}_k over the alphabet $\{1, \dots, k\}$ that recognizes precisely the words $w(\pi)$ that correspond to the pop-stacked permutations. The states of \mathcal{A}_k are labelled by pairs (L, C) , where

- $L \subseteq \{1, \dots, k\}$ indicates the already visited letters, and
- $C \subseteq \bigcup_{j=1}^{k-1} \{(j, j+1), (j+1, j)\}$ indicates the already fulfilled conditions “there is an occurrence of j before an occurrence of $j+1$ ” resp. “there is an occurrence of $j+1$ before an occurrence of j ”, such that
- if $j, j+1 \in L$, then at least one of $(j, j+1)$ and $(j+1, j)$ belongs to C .

It is then straightforward to see that \mathcal{A}_k recognizes precisely those words in $\{1, 2, \dots, k\}^n$ that correspond bijectively to the pop-stacked permutations of size n with k runs by Proposition 2. Figure 3 shows such an automaton for $k = 3$. \square




Figure 3. Automaton \mathcal{A}_3 that recognizes the words that correspond to pop-stacked permutations with $k = 3$ runs. In each state, the upper (blue) area indicates the letters from the alphabet $\{1, \dots, k\}$ already encountered, and the pairs of numbers in the lower (green) area indicate the conditions (i, j) (where $j = i \pm 1$): the condition (i, j) (abbreviated here by $i..j$) is listed if and only if there was already an occurrence of i before an occurrence of j in the word read by this point. If a state is entered after reading the letter j , then reading j again loops back to the same state: such transitions are omitted in this figure.

In the next theorem we address the complexity of \mathcal{A}_k : its number of states grows roughly like 3.41^k . This exponential growth of the number of states also gives an insight on the complexity of the generating functions associated to these automata, on which we comment more in the full version.

Theorem 5. *Denote by a_k the number of states in the automaton \mathcal{A}_k . The sequence $(a_k)_{k \geq 1}$ satisfies the linear recurrence $a_k = 4a_{k-1} - 2a_{k-2}$. Together with initial conditions, this implies that this sequence is [A006012](#): $(2, 6, 20, 68, 232, \dots)$. Accordingly, the number of states of \mathcal{A}_k grows exponentially: $a_k = \Theta((2 + \sqrt{2})^k)$.*

Proof. We proceed by induction on k (starting at $k \geq 3$). Denote by Q_k the set of states of \mathcal{A}_k . Recall that the states are labelled by (L, C) , where L is the set of already encountered letters, and C is the list of already fulfilled conditions of the kind $(j, j+1)$ or $(j+1, j)$.

We partition the states of \mathcal{A}_k into three parts as follows.

1. The states of \mathcal{A}_k whose letters do not contain k : they are precisely all the states of \mathcal{A}_{k-1} .
2. The states of \mathcal{A}_k whose letters contain k but do not contain $k-1$. They correspond bijectively to the states of \mathcal{A}_{k-1} whose letters do not contain $k-1$, and thus to all the states of \mathcal{A}_{k-2} :

$$(L, C) \in Q_{k-2} \longleftrightarrow (L, C) \in Q_{k-1} \longleftrightarrow (L \cup \{k\}, C) \in Q_k,$$

where $k-1 \notin L$.

3. Finally, the states of \mathcal{A}_k whose letters contain both k and $k-1$. They correspond $(3 : 1)$ -bijectively to the states of \mathcal{A}_{k-1} whose letters contain $k-1$:

$$(L, C) \in Q_{k-1} \longleftrightarrow \left. \begin{array}{l} (L \cup \{k\}, C \cup \{(k-1, k)\}), \\ (L \cup \{k\}, C \cup \{(k, k-1)\}), \\ (L \cup \{k\}, C \cup \{(k-1, k), (k, k-1)\}) \end{array} \right\} \in Q_k,$$

where $k-1 \in L$.

Therefore, summing over these three cases, we have

$$a_k = a_{k-1} + a_{k-2} + 3(a_{k-1} - a_{k-2}),$$

which completes the proof. \square

Remark: Upon performing *minimization* on \mathcal{A}_k , we obtain automata with the number of states given by $(b_k)_{k \geq 1} = (2, 6, 16, 40, 98, \dots)$. This leads us to the conjecture that this sequence satisfies the linear recurrence $b_k = 3b_{k-1} - b_{k-2} - b_{k-3}$, and that it is in fact [A293004](#)². It is interesting to notice that the exponential growth of the number of states would then drop from $a_k = \Theta((2 + \sqrt{2})^k)$ to $b_k = \Theta((1 + \sqrt{2})^k)$.

²The goddess of combinatorics is thumbing her nose at us, as this sequence is also related to permutation patterns in many ways: it counts permutations related to the elevator problem [11, 5.4.8 Ex.8], permutations avoiding 2413, 3142, 2143, 3412, and some guillotine partitions [2]!

4. ASYMPTOTIC ASPECTS

In many cases, the growth of restricted families of permutations is much less than $n!$, and is just exponential: for example, by the Stanley–Wilf conjecture (proven by Marcus and Tardos [12]), this always holds for permutation classes defined by *classical* forbidden patterns. It is natural to ask whether this is the case for pop-stacked permutations. We now prove that they in fact grow much faster.

Theorem 6 (Superexponential growth of pop-stacked permutations).

The asymptotic growth of the number of pop-stacked permutations is at least $\exp(n \ln(n) - n \ln(2) - n + \ln(n) + \ln(\pi) + o(1))$.

Proof. We achieve this by the constructing an explicit class of permutations, as follows. Assume that n is even, and consider any pair of permutations, π and τ , each of size $n/2$. *Intertwine* them as shown in Figure 4, i.e., consider the permutation σ of size n defined by

$$\text{For } 1 \leq j \leq n/2: \quad \sigma(2j-1) = \pi(j), \quad \sigma(2j) = \tau(j).$$

It is easy to see that such σ is necessarily pop-stacked (it clearly satisfies the overlapping run condition from Theorem 1). Since the mapping $(\pi, \tau) \rightarrow \sigma$ is injective, we conclude that we have at least $((n/2)!)^2$ pop-stacked permutations of size n . Then, Stirling’s approximation for the factorial leads to the theorem. \square




Figure 4. Intertwining two permutations: illustration for the proof of Theorem 6.

This ends our first incursion in the world of pop-stack sorting. In the full version of this paper, we say more on the enumeration and asymptotics, on links with generating trees [5], and on the cost of computing the number of pop-stacked permutations. We also consider extremal cases of pop-stack sorting in our companion article [1].

Acknowledgment. We thank the referees of Eurocomb’2019 for their feedback and their careful reading of this article.

REFERENCES

1. Asinowski A., Banderier C., and Hackl B. *On extremal cases of pop-stack sorting*. In *Proceedings of Permutation Patterns (Zürich, Switzerland)*, 33–37, (2019).
2. Asinowski A. and Mansour T. *Separable d -permutations and guillotine partitions*. *Ann. Comb.*, **14**(1) (2010), 17–43.
3. Atkinson M. D. and Sack J.-R. *Pop-stacks in parallel*. *Inform. Process. Lett.*, **70**(2) (1999), 63–67.
4. Avis D. and Newborn M. *On pop-stacks in series*. *Util. Math.*, **19** (1981), 129–140.
5. Banderier C., Bousquet-Mélou M., Denise A., Flajolet P., Gardy D., and Gouyou-Beauchamps D. *Generating functions for generating trees*. *Discrete Math.*, **246**(1-3) (2002), 29–55.
6. Claesson A. and Guðmundsson B. A. *Enumerating permutations sortable by k passes through a pop-stack*. arXiv (2017).
7. Dulucq S., Gire S., and Guibert O. *A combinatorial proof of J. West’s conjecture*. *Discrete Math.*, **187**(1-3) (1998), 71–96.
8. Flajolet P. and Sedgewick R. *Analytic Combinatorics*. Cambridge University Press, 2009.
9. Gates W. H. and Papadimitriou C. H. *Bounds for sorting by prefix reversal*. *Discrete Math.*, **27**(1) (1979), 47–57.
10. Knuth D. E. *The Art of Computer Programming. Vol. 1: Fundamental Algorithms*. Addison-Wesley, Reading, MA, 1998. Third edition (1st ed: 1968).
11. Knuth D. E. *The Art of Computer Programming. Vol. 3: Sorting and searching*. Addison-Wesley, Reading, MA, 1998. Second edition (1st ed: 1973).
12. Marcus A. and Tardos G. *Excluded permutation matrices and the Stanley–Wilf conjecture*. *J. Combin. Theory, Ser. A*, **107**(1) (2004), 153–160.
13. Pudwell L. and Smith R. *Two-stack-sorting with pop stacks*. *Australas. J. Combin.*, **74**(1) (2019), 179–195.
14. Ungar P. *$2N$ noncollinear points determine at least $2N$ directions*. *J. Combin. Theory Ser. A*, **33**(3) (1982), 343–347.
15. West J. *Sorting twice through a stack*. *Theoret. Comput. Sci.*, **117**(1-2) (1993), 303–313. Conference on Formal Power Series and Algebraic Combinatorics (Bordeaux, 1991).

Andrei Asinowski, Alpen-Adria-Universität Klagenfurt, Austria,
<https://me.aau.at/~anasinowski/>

Cyril Banderier, Université de Paris Nord, France,
<https://lipn.fr/~banderier>

Sara Billey, University of Washington, USA,
<https://sites.math.washington.edu/~billey>

Benjamin Hackl, Alpen-Adria-Universität Klagenfurt, Austria,
<https://benjamin-hackl.at>

Svante Linusson, KTH Royal Institute of Technology in Stockholm, Sweden,
<https://www.kth.se/profile/linusson>