Theoretical Computer Science 882 (2021) 49–62

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Syntactic view of sigma-tau generation of permutations

Wojciech Rytter, Wiktor Zuba *

Inst. of Informatics, Warsaw University, Warsaw, Poland

ARTICLE INFO

Article history:
Received 28 August 2019
Received in revised form 3 February 2021
Accepted 1 June 2021
Available online 9 June 2021

Keywords:
Sigma-tau generation
Permutations
Straight-line program
Ranking

ABSTRACT

We give a syntactic view of the Sawada-Williams (σ, τ)-generation of permutations. The corresponding sequence Seq_n of $\sigma\tau$-operations of length $n!-1$ is shown to have a compact description of size $\Theta(n^2)$ in terms of straight-line programs. Using the compact description, we design almost linear time ranking and unranking algorithms.

1. Introduction

We consider permutations of the set $\{1, 2, ..., n\}$, called here n-permutations. For an n-permutation $\pi = a_1 \dots a_n$, we define $\sigma(\pi) = a_2 a_3 \dots a_n a_1$, $\tau(\pi) = a_2 a_1 a_3 \dots a_n$. In their classical book on combinatorial algorithms, Nijenhuis and Wilf asked in 1975 if all n-permutations can be generated, each exactly once, using in each step a single operation σ or τ. This difficult problem was open for almost 40 years. Very recently it was solved by Williams in [21], and later Sawada and Williams presented the solution at the conference SODA'2018 [18]. For odd n, their solution was extended in [19] to produce a $\sigma\tau$-cycle of permutations. There are many other algorithms for permutation generation, for example [17,11,14,9,7,5,1,22,8,20,12,13]. Usually the ordering method is accompanied by ranking and unranking algorithms. For a given ordering, the ranking algorithm gives positions (ranks) of each permutation in this ordering, while the unranking algorithm for a given position in the ordering gives the permutation in this position. The ranking and unranking of permutations depend on the ordering of permutations implied by the ordering method. In the case of lexicographic ordering, there are algorithms working in $O(n \log n)$ time, see [9]. It was observed in [11] that it can be improved to $O(n \log n / \log\log n)$ using a sophisticated data structure. Other ranking/unranking algorithms for lexicographic order were given in [2]. Permutations in the ordering described in [11] admit linear time ranking and unranking. This ordering does not correspond to lexicographic order. Ranking/unranking algorithms for permutations in the Steinhaus-Johnson-Trotter order can be found in [6] and [10]. All these previously known orderings are completely different from that of Sawada and Williams. We believe that this important ordering (solving a long standing open problem) deserves further studies and that it is a natural problem to find the corresponding ranking/unranking algorithms. We give new insights into this ordering.

Usually the generation algorithm produces each successive permutation by applying some *basic operation*. The sequence of these operations, corresponding to the permutation ordering, is called the *syntactic sequence*. It is treated as a word over the alphabet Σ consisting of basic operations. For example, for $n = 3$ and the $\sigma\tau$-ordering $231 \xrightarrow{\sigma} 312 \xrightarrow{\tau} 132 \xrightarrow{\sigma} 321 \xrightarrow{\sigma} 213 \xrightarrow{\tau} 123$ the corresponding syntactic sequence is $\sigma\,\tau\,\sigma^2\,\tau$. We define Seq_n as the syntactic sequence, over

* Corresponding author.
E-mail address: w.zuba@mimuw.edu.pl (W. Zuba).

https://doi.org/10.1016/j.tcs.2021.06.008
0304-3975/

the alphabet $\Sigma = \{\sigma, \tau\}$, corresponding to the $\sigma\tau$-ordering of n-permutations presented in [18]. We investigate compact representations of these sequences. For example, $\mathsf{Seq}_5 = \gamma_1^3\sigma^2(X\tau)^3X$, where $\gamma_i = \sigma^i\tau,\ X = \gamma_2^2\gamma_4^3\gamma_1\gamma_4\sigma^4$. Seq_5 is of length 119.

A structural compact representation of the syntactic sequence can help to better understand the ordering, and to design ranking and unranking algorithms. Our approach is similar to the one used by Shmuel Zaks in [22], in this case the basic operations are reversals of a suffix of the permutation. The sequence of lengths of suffixes (corresponding to reversals) is denoted by s_n, and is described by recurrences:

$$s_2 = 2;\ s_n = (s_{n-1}\, n)^{n-1}\, s_{n-1}. \tag{1}$$

Surprisingly, the recurrences given in Equation (1), give also the MinGreedy ordering of permutations presented in [17], related to the greedy flipping of pancakes. In [22], the ranking/unranking algorithms were derived from these recurrences in a structural way. We are doing the same, after constructing a small compact representation of Seq_n we derive the ranking/unranking algorithms which reflect the structure of this representation. Unfortunately the word Seq_n is much more complicated than s_n.
Such a syntactic approach was also used by Ruskey and Williams in the ordering of $(n-1)$-permutations of an n-set in [14].

In this paper we consider compression in terms of straight-line programs. A straight-line program, briefly SLP, is a context-free grammar that produces a single word w over a given alphabet Σ. Equivalently, it is a sequence of recurrences (equations), see [15], using the operation of concatenation of words. The straight-line program can be seen as a program with no branches, no loops, no conditional statements, no comparisons – just a *straight* sequence of basic operations. This is probably the most likely origin of the name. A straight-line program can be equivalently defined as a sequence of assignment statements of the form $X_1 = expr_1;\ X_2 = expr_2;\ \dots\ X_m = expr_m$, where each expression $expr_i$ is a word over $\Sigma \cup \{X_1, X_2, \dots, X_{i-1}\}$. In the expressions we can only use the operation of concatenation of words. We use such a restricted model to show tight bounds on the compression of Seq_n.
For each X_i, denote by $val(X_i)$ the value of X_i, it is the string described by X_i and computed by executing the straight-line program. The string described by the whole straight-line program is $val(X_m)$. The size of an SLP is the total length of all expressions $expr_i$. For example, the minimal size of an SLP-representation of σ^n is $\Theta(\log n)$.
We assume that the arithmetic operations used in the paper are computable in a constant time. The main results of this paper are,

1. **Tight bounds.** The minimal size of an SLP representing Seq_n is $\Theta(n^2)$.
2. **Ranking.** Using a compact description of Seq_n, the number of steps (the rank of the permutation) needed to obtain a given permutation from a starting one can be computed in time $\mathcal{O}(n\sqrt{\log n})$.
3. **Unranking.** Using a compact description of Seq_n, the t-th permutation generated by Seq_n can be computed in $\mathcal{O}(n\frac{\log n}{\log\log n})$ time.
4. Similar results are shown for the cyclic $\sigma\tau$-generation.

A very preliminary version of this paper has been presented in [16], without any results about the cyclic generation.

2. Tree-like graphs of seeds

We define the sigma-tau graph $\mathbf{G}_n$, its vertices are permutations of $\{1, 2, \dots, n\}$, and there is a directed edge from π_1 to π_2 if and only if $\pi_2 = \sigma(\pi_1)$ or $\pi_2 = \tau(\pi_1)$.

Define $x \oplus 1 = x + 1$ if $1 \le x < n-1$, and $(n-1) \oplus 1 = 1$. If $a \oplus 1 = b$, then we write $b \ominus 1 = a$. A *seed* ψ is an $(n-1)$-tuple $a_1\, a_2 \dots a_{n-1}$ of distinct elements of $\{1, 2, \dots, n\}$ such that $a_1 = n$ and $a_1\,(a_2 \oplus 1)\, a_2 \dots a_{n-1}$ is a permutation. The element $\mathrm{mis}(\psi) = a_2 \oplus 1$ is called a *missing* element. For example, the set of seeds for $n = 5$ is

$$\{5423,\ 5432,\ 5134,\ 5143,\ 5214,\ 5241,\ 5312,\ 5321\}.$$

Denote by $\mathrm{cycle}(\pi)$ all permutations cyclically equivalent to π. We define $\mathrm{perms}(\psi)$ as the set of all n-permutations which result by making a single insertion of $\mathrm{mis}(\psi)$ into any position in ψ, and then making cyclic shifts. The sets $\mathrm{perms}(\psi)$ are called *packages*, the seed ψ is the *identifier* of its package. For a seed $\psi = a_1\, a_2 \dots a_{n-1}$ with $\mathrm{mis}(\psi) = x$, let $\widetilde{\psi} = a_1\, x\, a_2\, a_3 \dots a_{n-1}$. The permutation $\widetilde{\psi}$ is called the *representative* of the package $\mathrm{perms}(\psi)$. For example, if $\psi = 413$, then $\mathrm{mis}(\psi) = 2$, $\widetilde{\psi} = 4213$. In the paper we frequently use the sequences of the type $\gamma_k = \sigma^k \cdot \tau$. It was observed in [18] that $\mathrm{perms}(\psi)$ is the set of permutations resulting by applying prefixes (as functions) of $\tau \cdot (\gamma_{n-1})^{n-1}$ to $\widetilde{\psi}$. We introduce an operation parent, which changes the seed $a_1\, a_2 \dots a_{n-1}$ by inserting $\mathrm{mis}(\psi)$ immediately after a_1, and removing $\mathrm{mis}(\psi) \oplus 1$ from the sequence (the new *missing* element becomes $\mathrm{mis}(\psi) \oplus 1$). $\mathrm{parent}(\psi)$ is called the *parent* of ψ. For example, $\mathrm{parent}(5\,1\,3\,4) = 5\,2\,1\,4$, $\mathrm{parent}(5\,3\,2\,1) = 5\,4\,3\,2$ (see Fig. 1).

Denote by Seeds_n the set of seeds of $\{1, 2, \dots, n\}$. We define the graph $\mathsf{SeedGraph}_n = (V, E)$, where $V = \mathsf{Seeds}_n$, and the directed edges are of the form $\psi \to \mathrm{parent}(\psi)$. The node ψ is called the i-th *child* of its parent ψ', written $\mathrm{child}(\psi', i) = \psi$, if $\mathrm{mis}(\psi')$ is at the i-th position in ψ counting from the end. $\mathsf{SeedGraph}_n$ consists of a directed cycle with hanging subtrees, see Fig. 2. The nodes on the cycle are called hub seeds, and their non-hub children are called prehub seeds. The set of hub

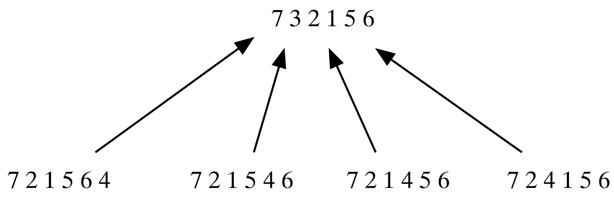

Fig. 1. Children of an example non-hub node 7 3 2 1 5 6 with the missing element 4. This missing element is in the i-th place (counting from the end) in the i-th child.

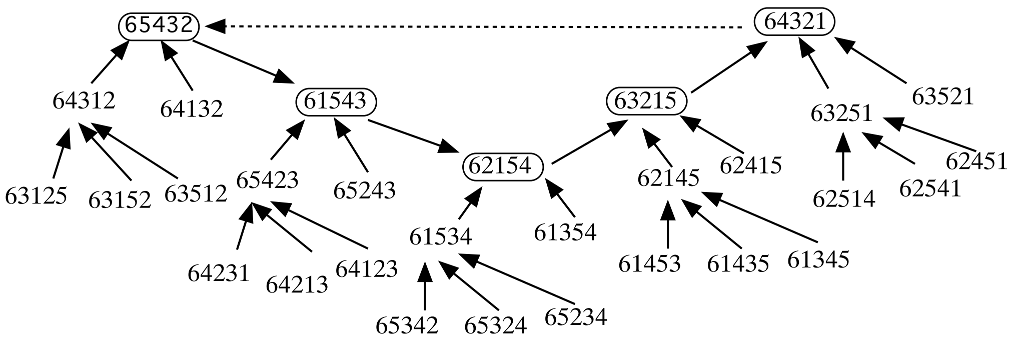

Fig. 2. The graph $\mathsf{SeedGraph}_6$. $\mathrm{Hub}_6 = \{64321, 63215, 62154, 61543, 65432\}$. After removing the dotted edge the graph becomes a rooted tree.

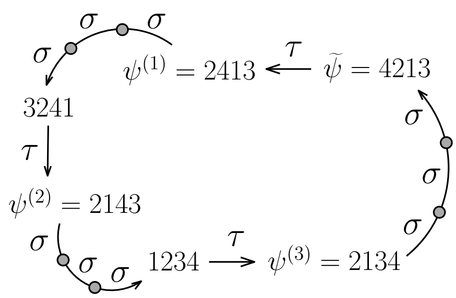

Fig. 3. The structure of $\mathrm{ham}(\psi)$, where $\psi = 413$, $\mathrm{mis}(\psi) = 2$.

seeds is denoted by Hub_n. We say that a sequence $(c_1, c_2, ..., c_k)$ is cyclically decreasing if $c_i = c_{i+1} \oplus 1$ for $i < k$. Repetitive use of the operation parent can be seen as *cyclically sorting* the seed (disregarding the first element), for example,

$$8\,\underline{1}\,3\,4\,5\,6\,7 \xrightarrow{\text{parent}} 8\,\underline{2\,1}\,4\,5\,6\,7 \xrightarrow{\text{parent}} 8\,\underline{3\,2\,1}\,5\,6\,7 \xrightarrow{\text{parent}} 8\,\underline{4\,3\,2\,1}\,6\,7 \xrightarrow{\text{parent}} 8\,\underline{5\,4\,3\,2\,1}\,7 \xrightarrow{\text{parent}} 8\,\underline{6\,5\,4\,3\,2\,1}.$$

For a non-hub seed ψ we define the seeds tree $\mathrm{tree}(\psi)$ as a tree consisting of all descendants of ψ (including ψ) in $\mathsf{SeedGraph}_n$. By the length of a path we mean the number of its nodes. The height of a non-hub node ψ is the maximum length of a path from ψ to a leaf in $\mathrm{tree}(\psi)$. Two seeds ϕ, ψ are called *neighbors* if and only if $\mathrm{perms}(\phi) \cap \mathrm{perms}(\psi) \neq \emptyset$.

Observation 1. **(a)** Two distinct seeds ϕ, ψ are neighbors if and only if one of them is a parent of another one in $\mathsf{SeedGraph}_n$. **(b)** This happens if and only if $\mathrm{mis}(\phi) = \mathrm{mis}(\psi) \oplus 1$ or $\mathrm{mis}(\psi) = \mathrm{mis}(\phi) \oplus 1$, and after removing from ψ and ϕ both elements $\mathrm{mis}(\psi)$, $\mathrm{mis}(\phi)$, they become identical.

For a seed ψ, denote by $\mathrm{ham}(\psi)$ (similarly as in [18]) a directed simple cycle in $\mathbf{G}_n$ consisting of edges *implied* by the seed ψ in the following sense.

The set of nodes of $\mathrm{ham}(\psi)$ equals $\mathrm{perms}(\psi)$. For each $\pi \in \mathrm{perms}(\psi)$ we choose the edge $\pi \to \tau(\pi)$ if $\mathrm{mis}(\psi)$ is in the second position in π, otherwise we choose the edge $\pi \to \sigma(\pi)$, see Fig. 3.

Observe that $\mathrm{ham}(\psi)$ is a simple cycle in $\mathbf{G}_n$ with the same syntactic sequence (labels of edges) $\tau\, \gamma_{n-1}^{n-1}$ for each seed ψ, if we start with $\widetilde{\psi}$ (see Fig. 4).

For $1 \le i \le n-1$, denote $\psi^{(i)} = \tau \cdot \gamma_{n-1}^{i-1}(\widetilde{\psi})$. In other words $\psi^{(i)}$, for $n > i > 0$, is the word ψ right-shifted by $i-1$ and with $\mathrm{mis}(\psi)$ added at the beginning. Observe that $\gamma_{n-1}(\psi^{(i)}) = \psi^{(i+1)}$ for $0 < i < n-1$.

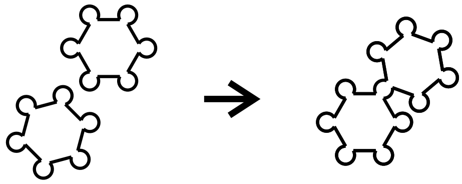

Fig. 4. Two non-disjoint simple cycles corresponding to distinct neighboring seeds can be joined into a simple cycle by removing two σ-edges conflicting with τ-edges.

Example 2. For $\psi = 5\,3\,2\,1$ we have $\widetilde{\psi} = 5\,4\,3\,2\,1$, $\psi^{(1)} = 4\,\underline{5\,3\,2\,1}$, $\psi^{(2)} = 4\,\underline{1\,5\,3\,2}$.

We say that an edge $u \to v$ conflicts with $u' \to v'$ if and only if

$$(u, v) \neq (u', v') \text{ and } (u = u' \text{ or } v = v').$$

Define $\mathcal{R}_n$ as the graph consisting of all τ-edges in all (local) graphs ham(ψ) and all σ-edges not conflicting with selected τ-edges. A version of the construction of a Hamiltonian path by Sawada-Williams, denoted by SW_n, can be written informally as

Algorithm Compute SW_n
$P = \mathcal{R}_n$
$\pi := n\,(n-1)\dots 1$; add to P the edge $\pi \to \sigma(\pi)$
remove from P edges conflicting with the newly added edge
return P $\{P = \mathsf{SW}_n\}$

3. Compact representations of subpaths of SW_n

We define a bunch as a subpath of SW_n starting in $\widetilde{\psi}$ and terminating in $\psi^{(n-1)}$, where ψ is a non-hub seed. We denote it by bunch(ψ). For any non-hub seed ψ of height k we define $\mathbf{W}_k$ as the sequence of labels of consecutive $\sigma\tau$-edges in bunch(ψ). We show later that this sequence depends only on the height of ψ.

Lemma 3. *If $\psi = \mathrm{child}(\phi, i)$, then $\mathrm{perms}(\phi) \cap \mathrm{perms}(\psi)$ is the σ-cycle containing both $\widetilde{\psi}$ and $\phi^{(i)}$.*

Proof. Assume that $\pi = p_1\, p_2 \dots p_n \in \mathrm{perms}(\phi)$. We can assume that $p_1 = n$, since cyclically equivalent permutations belong to the same packages. After removing any element from π, the first element following n is either p_2 or p_3, in both cases we obtain a valid seed if and only if the removed element is greater by one than the element following n. After removing the element $p_2 \oplus 1$, we always receive a seed whose package contains π, and we can remove p_2 only when $p_2 = p_3 \oplus 1$. It means that if $\pi \in \mathrm{perms}(\phi) \cap \mathrm{perms}(\psi)$, then one of these seeds (denote it by ϕ) equals $p_1\, p_2 \dots p_n$ (without the element $p_j = p_2 \oplus 1$) and the other one (denote it by ψ) is equal $p_1\, p_3 \dots p_n$. The element $\mathrm{mis}(\phi) = p_j$ occurs at the position $n-j+1$ from the right in ψ, thus $\phi^{(i)} = \phi^{(n-j+1)} = p_j\, p_{j+1} \dots p_n\, p_1 \dots p_{j-1}$ is cyclically equivalent to $\widetilde{\psi} = \pi$. $\square$

By hub(ψ) we denote the first hub seed reachable from ψ by parent links, and by prehub(ψ) denote the highest non-hub ancestor of ψ. The length $r+1$ of the *parent sequence* $\psi = \psi_0 \to \psi_1 \to \psi_2 \to \dots \to \psi_r$ from ψ to hub(ψ) is denoted by level(ψ).

For a seed $(s_1, s_2, .., s_{n-1})$, define decsubs(ψ) and decpref(ψ) as the longest prefix of $(s_2 \ominus 1, s_2 \ominus 2, .., s_2 \ominus (n-3))$, which is, respectively, a subsequence/prefix of $(s_3, s_4, .., s_{n-1})$. In particular each prehub seed ψ has $|\mathrm{decsubs}(\psi)| = n-4$ and each leaf ψ of $\mathsf{SeedGraph}_n$ has $|\mathrm{decpref}(\psi)| = 0$.

Observation 4. $\mathrm{level}(\psi) = n - 3 - |\mathrm{decsubs}(\psi)|$ and $\mathrm{height}(\psi) = |\mathrm{decpref}(\psi)| + 1$. If ψ is a non-hub seed, then the level and the prehub of ψ can be computed in $\mathcal{O}(n)$ time.

Example 5. We have decsubs(96154238) = (5, 4, 3). Hence the path from 96154238 to its prehub 98765423 is 96154238 $\to$ 97615423 $\to$ 98765423, and decpref(96154238) = 0.

Lemma 6. *Each non-hub seed of height $k > 1$ has exactly $n-3$ children. The sequence of heights of all these children (in their respective order in $\mathsf{SeedGraph}_n$) is equal to $k-1, k-1, \dots, k-1, k-2, k-3, \dots, 1$ (see Fig. 5).*

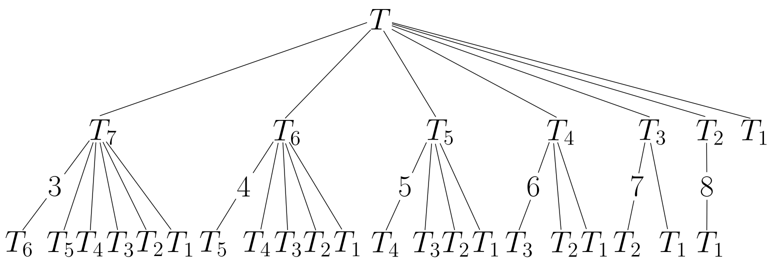

Fig. 5. The root is a hub node in SeedGraph$_{11}$, it is the parent of several prehub nodes. Each subtree T_i is an isomorphic copy of a seed tree rooted at a seed of height i, after discarding labels of nodes. The numbers at edges correspond to the number of copies of the same subtree.

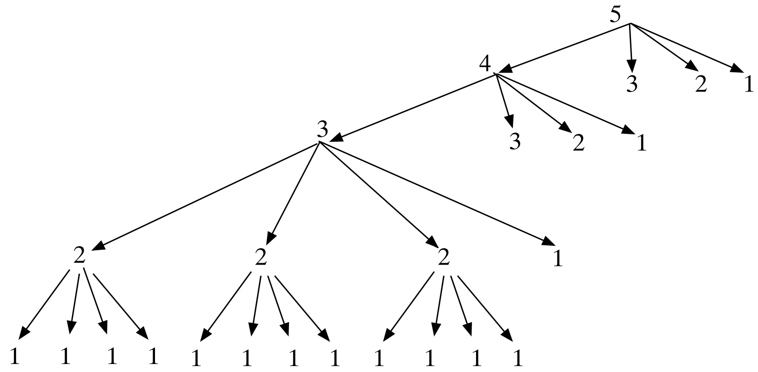

Fig. 6. A part of SkewTree(7, 5). The numbers at the nodes are their heights.

Proof. Inserting the *missing* element of the seed after the end of its decreasing prefix cannot change its length, hence the first $n-k-1$ children are of height $k-1$. Inserting the element before an element from the decreasing prefix shortens it, hence the height decreases by one for each next child. □

Example 7. The height of the seed $\psi = 1\,1\,9\,8\,7\,6\,1\,2\,3\,4\,5$ is 4, since $8\,7\,6$ is the longest common prefix of $(a_3, \ldots, a_{n-1})$ and $(a_2 \ominus 1, a_2 \ominus 2, \ldots, a_2 \oplus 2)$. The children of ψ have heights (in order of children) 3, 3, 3, 3, 3, 3, 2, 1. All the children of $1\,1\,9\,8\,1\,2\,3\,4\,5\,6\,7$ are of height 1.

For a given $n > 3$, the *skew* tree SkewTree(n, k) of height $k > 1$ and rank n is an ordered tree which consists of a root labeled k and $n-3$ children $v_1, v_2, \ldots, v_{n-3}$ ordered from left to right. The first $n-k-2$ of them are roots of subtrees of height $k-1$, the remaining nodes are roots of subtrees of heights $k-1, k-2, \ldots, 1$ (from left to right). The tree SkewTree$(n, 1)$ is a single node. We define the following table Δ:

$$\Delta(k, i) = \min(k-1, n-2-i) \text{ for } i \le n-2.$$

$\Delta(k, i)$, for $i \in [1..n-3]$, forms the sequence $(k-1)^*$, $k-2$, $k-3$, $k-4$, …, 1 (see Fig. 6).

Observation 8. If $\text{height}(\phi) = k \le n-3$ and $child(\phi, i) = \psi$, then $\text{height}(\psi) = \Delta(k, i)$.

We add to each node of SkewTree(n, k) backward edges from each of its children, see Fig. 7. We label the edges by numbers (treated as symbols here), and by symbols $\bar{i}$, for $i \le n-3$. We also use an additional symbol δ to label each leaf. We give labels to edges in the following way: $v \xrightarrow{i} \text{child}(v, i)$; $\quad \text{child}(v, i) \xrightarrow{\bar{i}} v$.
Denote by dfs(n, k) the word obtained by traversing the tree SkewTree(n, k) in a depth-first-search order and concatenating labels of traversed edges and visited leaves.

Observation 9. We have $\text{dfs}(n, 1) = \delta$ and $\text{dfs}(n, k) = \prod_{i=1}^{n-3} i \cdot \text{dfs}(n, \Delta(k, i)) \cdot \bar{i}$ for $k > 1$. For $\psi \notin \text{Hub}_n$, tree(ψ) is isomorphic (after discarding labels of nodes) to a skew tree of the same height.

For a permutation π and a sequence α of operations σ, τ denote by $\underline{\text{gen}}(\pi, \alpha)$ the set of all permutations generated from π by following α, including π. The word $\mathbf{W}_k$ satisfies equations $\text{gen}(\widetilde{\psi}, \mathbf{W}_k) = \text{bunch}(\psi)$ and $\mathbf{W}_k(\widetilde{\psi}) = \psi^{(n-1)}$.
Recall that we denoted $\gamma_k = \sigma^k \tau$. We define a morphism

$$\Psi(\delta) = \mathbf{W}_1; \;\; \Psi(i) = \tau \cdot \sigma^i; \;\; \Psi(\bar{i}) = \sigma^{n-2-i} \text{ for } i < n-3; \;\; \Psi(\overline{n-3}) = \sigma \cdot \tau \cdot \gamma_{n-1}$$

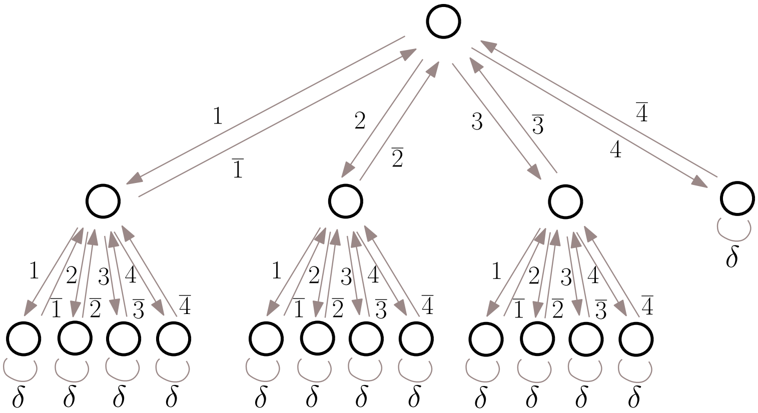

Fig. 7. An example traversal sequence dfs(7, 3) of length 45 of SkewTree(7, 3).

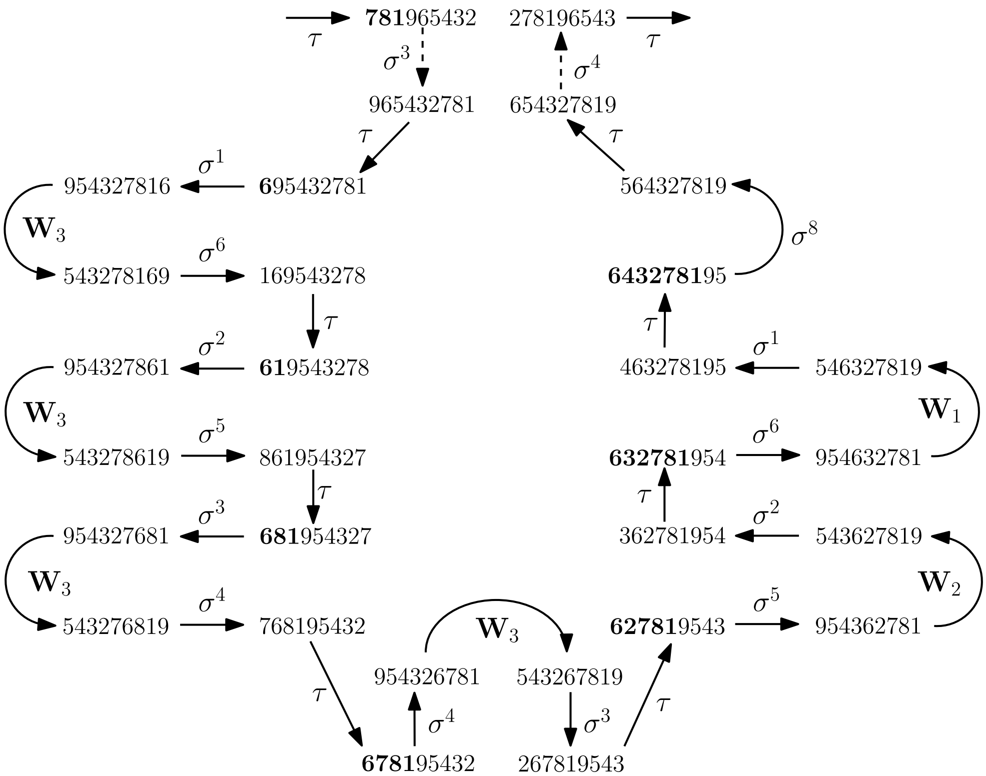

Fig. 8. The structure of bunch(ψ) for the seed $\psi = 95432781$. We have parent(ψ) = ϕ, where $\phi = 96543281$. The connecting points of ψ with its parent are $\widetilde{\psi}$ and $\psi^{(n-1)}$, in other words bunch(ψ) $\cap$ perms(ϕ) $= \{\widetilde{\psi}, \psi^{(n-1)}\}$. The sequence $\mathbf{W}_4$ starts in $\widetilde{\psi}$, visits all permutations in bunch(ψ) and ends in $\psi^{(n-1)}$. We have $\mathbf{W}_4 = \tau \cdot \sigma^1\mathbf{W}_3\gamma_6 \cdot \sigma^2\mathbf{W}_3\gamma_5 \cdot \sigma^3\mathbf{W}_3\gamma_4 \cdot \sigma^4\mathbf{W}_3\gamma_3 \cdot \sigma^5\mathbf{W}_2\gamma_2 \cdot \sigma^6\mathbf{W}_1\gamma_1 \cdot \gamma_8$.

Lemma 10. $\mathbf{W}_k = \Psi(\text{dfs}(n,k))$ *for* $1 \le k \le n-3$.

Proof. We show by induction on k that $\Psi(\text{dfs}(n,k))$ represents a Hamiltonian path covering bunch(ψ). The base of induction is given by the definition of $\Psi(\delta)$. By induction, we assume that for $\phi = \text{child}(\psi, i)$ the path covering bunch(ϕ) is represented by $\Psi(\text{dfs}(n, \Delta(k,i)))$. We have that $\Psi(i)\sigma\Psi(\bar{i}) = \tau\sigma^{n-1}$ for $i < n-3$, and $\Psi(i)\sigma\Psi(\bar{i}) = (\tau\sigma^{n-1})^2\tau$ for $i = n-3$, hence the sequence $\prod_{i=1}^{n-3}\Psi(i)\sigma\Psi(\bar{i}) = \tau(\gamma_{n-1})^{n-2}$ covers perms(ψ) $\cap$ bunch(ψ) (placeholders σ are inserted here for transitions between $\widetilde{\phi}$ and $\phi^{(n-1)}$ for each child of ψ). Then $\widetilde{\phi} = \sigma^i(\psi^{(i)}) = \Psi(i)((\tau\sigma)^{i-1}(\widetilde{\psi}))$ (as illustrated in Fig. 8), thus $\Psi(\text{dfs}(n, \Delta(k,i)))$ occurs exactly in the place held by σ. □

Observation 9 and Lemma 10 directly imply,

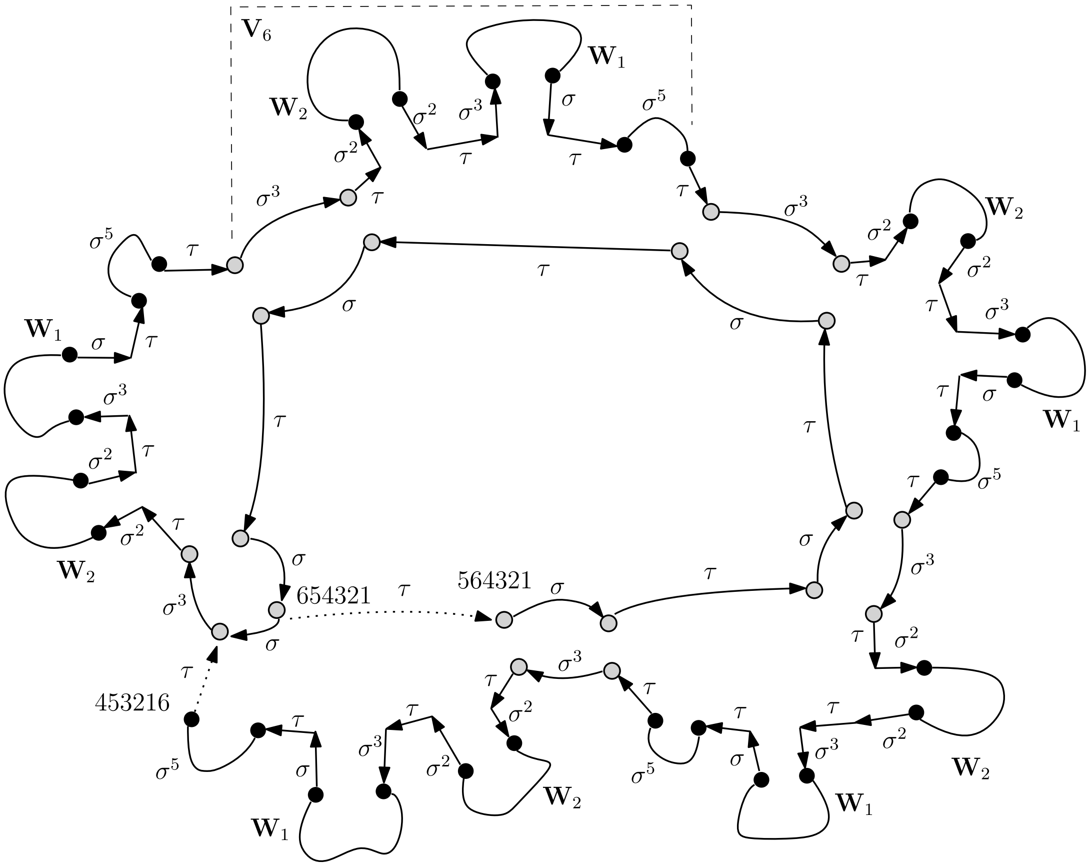

Fig. 9. The compacted structure of Seq_6 of length 720. It differs from the structure of $\mathcal{R}_6$ by adding one σ-edge from 654321 and removing two (dotted) τ-edges to have Hamiltonian path. We have $\mathsf{Seq}_6 = (\sigma\tau)^4\sigma^2\,(\mathbf{V}_6\tau)^4\,\mathbf{V}_6$, where $\mathbf{V}_6 = \sigma^3\,\tau\,\sigma^2\mathbf{W}_2\sigma^2\,\tau\,\sigma^3\mathbf{W}_1\sigma\,\tau\,\sigma^5$. The structure is the union of graphs of 5 seeds in Hub_6 with *hanging* bunches. The *starting path* is 564321 $\xrightarrow{(\sigma\tau)^4}$ 654321.

Corollary 11. *For $1 \le k < n-3$, we have the following recurrences*

$$\mathbf{W}_0 = \sigma, \quad \mathbf{W}_k = \tau \cdot \prod_{i=1}^{n-2} \sigma^i\,\mathbf{W}_{\Delta(k,i)}\,\gamma_{n-2-i}$$

For example, for $n=5$ $\mathbf{W}_1 = \tau\cdot\sigma^1\mathbf{W}_0\sigma^2\tau\cdot\sigma^2\mathbf{W}_0\sigma^1\tau\cdot\sigma^3\mathbf{W}_0\sigma^0\tau = \tau\cdot\gamma_4^3$.

4. Tight bounds for sizes of SLP representations of the whole ordering

We write separate SLP's for all γ_i's, for $i<n$ (exponents can be replaced by SLP's of logarithmic size, using only concatenations). Then, due to Corollary 11 the total size of SLP-representations of all words $\mathbf{W}_i$ is $O(n^2)$.

Theorem 12. *Seq_n has the following SLP-representation of size $O(n^2)$*

$$\mathbf{V}_n = \gamma_{n-3}\cdot\prod_{i=2}^{n-3}\sigma^i\,\mathbf{W}_{\Delta(n-3,i)}\,\gamma_{n-2-i}\cdot\sigma^{n-1},\quad Seq_n = \gamma_1^{n-2}\cdot\sigma^2\,(\mathbf{V}_n\,\tau)^{n-2}\cdot\mathbf{V}_n.$$

Proof. For every non-hub seed ψ we had that $\mathrm{gen}(\widetilde{\psi},\mathbf{W}_k) = \mathrm{bunch}(\psi)$, where $k = \mathrm{height}(\psi)$. The only difference for a hub seed ϕ is that $\mathrm{child}(\phi,1)$ cannot be considered as part of a tree rooted at ϕ (with already defined parent links), since $\mathrm{child}(\phi,1)\in\mathrm{Hub}_n$, and this would lead to a cycle ($\mathrm{child}(\phi,1)$ is reachable via parent links from ϕ).
We define $\mathbf{V}_n$ as $\mathbf{W}_{n-3}$ with the part corresponding to the first child removed (leaving only its suffix γ_{n-2-1}). We also delete the last symbol τ, as it does not appear at the end of the path (it corresponds to one of the τ-edges removed when joining two cycles into one path). Now, Seq_n consists of $n-1$ segments equal to $\mathbf{V}_n$ (corresponding to $n-1$ hub seeds) joined by τ-edges (they are linked in the same way as if the previous $\mathbf{V}_n$ part was a child of the next one). Additionally, it starts with γ_1^{n-1} representing the small path with the last τ-edge replaced by a σ-edge (see Fig. 9). $\square$

Now we will focus on the lower bound for the SLP. Define

$$\mathcal{Z}_n = \{z_{i,j} : n/4 \le j < n-5,\ i+2j < n\},\ \text{where } z_{i,j} = \sigma^{n-1-i}\tau\sigma^i(\tau\sigma)^j\tau\sigma^{n-1}.$$

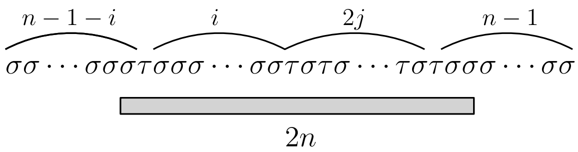

Fig. 10. All subwords of size $2n$ in $z_{i,j}$ contain the first and the last τ (since $i+2j<n$) and are distinct.

Lemma 13. *All words in $\mathcal{Z}_n$ are subwords of $\mathbf{W}_{n-4}$.*

Proof. We use composite labels of paths in SkewTree$(n, n-4)$ of the form: start in the $(i-1)$-st child of a node of height $j+2$, then go one edge up, go to the ith child and later go downward to a leaf on the longest branch. The path ends with a loop at a leaf, this loop is added to the path (see Fig. 7). Such a path exists since $i+j+2<n$. We define $v_{i,j} = \overline{i-1} \cdot i \cdot 1^j \cdot \delta$, $\Gamma_n = \{v_{i,j} : n/4 \le j < n-5,\ i+2j<n\}$.

We show that each word $z_{i,j}$ is a subword of $\mathbf{W}_{n-4}$ as a prefix of $\Psi(v_{i,j})$. The sequence $v_{i,j}$ is a subword in dfs$(n, j+2)$, and thus each word $w \in \Gamma_n$ is a subword in dfs$(n, n-4)$. Consequently, $\Psi_n(w)$ is a subword of $\mathbf{W}_{n-4}$ (due to Lemma 10). We have $\mathcal{Z}_n = \{\Psi(\overline{i-1} \cdot i \cdot 1^j) \cdot \tau\sigma^{n-1} : v_{i,j} \in \Gamma_n\}$. This completes the proof. □

Theorem 14. *The minimal SLP representation of* Seq_n *is of $\Omega(n^2)$ size.*

Proof. Denote by factors$_k(w)$ the set of subwords of length k of a word w. The structure of the words $z_{i,j}$ implies that all words in factors$_{2n}(z_{i,j})$ are distinct, each contains $\Omega(n)$ factors ($|z_{i,j}| = 2n + 2j \ge \frac{5}{2}n$). There are $\Omega(n^2)$ distinct (i, j) pairs and the sets factors$_{2n}(z_{i,j})$ are pairwise disjoint (i determines the distance between first two τ's and j determines the number of τ's in the string). Consequently, there are $\Omega(n^3)$ distinct factors of size $2n$ in all words in $\mathcal{Z}_n$. We use the fact that a word with SLP representation of size k has only $\mathcal{O}(mk)$ different subwords of length m, see [15]. The sequence Seq_n contains $n-1$ full $\mathbf{W}_{n-4}$ parts, hence the thesis follows from Lemma 13 (see Fig. 10). □

5. Locating and recovering seeds in a seed tree

If ψ is the i-th child of its parent in the seed tree, then we write ord$(\psi) = i$. For the parent sequence $\psi_0 = \psi, \psi_1, \ldots, \psi_m = \text{prehub}(\psi)$ denote

$$\text{route}(\psi) = (\text{ord}(\psi_0), \text{ord}(\psi_1), \ldots, \text{ord}(\psi_m)).$$

Example 15. For the seed 96154238 from Example 5 we have, ord$(\psi_0) = 1$, ord$(\psi_1) = 5$, ord$(\psi_2) = 2$, hence route(96154238) = (1, 5, 2).

Lemma 16. *For a non-hub seed ψ we can compute* route(ψ) *in $\mathcal{O}(n\sqrt{\log n})$ time.*

Proof. First we reduce the problem to the case mis$(\psi) = n-1$. Let $\psi = a_1 a_2 \ldots a_{n-1}$ and j be the number such that $a_2 \oplus j = n-2$. Then we replace ψ by $\psi' = a_1 (a_2 \oplus j)(a_3 \oplus j) \ldots (a_{n-1} \oplus j)$. The routes of ψ and ψ' are the same and mis$(\psi') = n-1$. Now, for the rest of the proof, we can assume that mis$(\psi) = n-1$. We know the length of the parent sequence from ψ to its prehub, since we know level(ψ). We use the following auxiliary problem

Inversion Vector problem: for a permutation π compute for each element x the number of elements smaller than x, which are to the right of x in π

Denote by pos(ψ, x) the position of x in the sequence ψ counting from the end and by rightsmaller(ψ, x) the number of elements smaller than x, which are to the right of x in ψ. Our sequence ψ is not a permutation but we can replace n with $n-1$ to obtain a one. Then, after solving the Inversion Vector problem, we know all values rightsmaller(ψ, k).

Assume mis$(\psi) = n-1$. Then for $k \le \text{level}(\psi)$ we have

$$\text{ord}(\psi) = \text{pos}(\psi, 1) \text{ and } \text{ord}(\text{parent}^k(\psi)) = \text{pos}(\psi, k+1) - \text{rightsmaller}(\psi, k+1) \tag{2}$$

Now all the values ord(ψ_k) can be computed using Equation (2). The *Inversion Vector* problem can be computed in $\mathcal{O}(n\sqrt{\log n})$ time, see [3]. Consequently, the whole computation of numbers ord(ψ_i) is of the same asymptotic complexity. □

Lemma 17. *If we know $\psi' = \text{hub}(\psi)$ and* route(ψ)*, then ψ can be computed in $\mathcal{O}(n\frac{\log n}{\log\log n})$ time.*

Proof. Assume route$(\psi) = (\text{ord}(\psi_0), \text{ord}(\psi_1), \ldots, \text{ord}(\psi_m)) = (t_0, t_1, \ldots, t_m)$. Then the following algorithm recovers ψ

ψ := ψ';
for $i = m-1$ **downto** 0 **do**
  x :=mis(ψ); remove $x \ominus 1$ in ψ;
  insert x to the position t_i in ψ counting from the end;
return ψ.

The data structure from [4] allows us to implement each of the operations on ψ (treated as a list) to work in $\mathcal{O}(\frac{\log n}{\log\log n})$ time, consequently the recovery algorithm works in $\mathcal{O}(n\frac{\log n}{\log\log n})$ time. This completes the proof. □

6. Ranking algorithm

The ranks of permutations which are representatives of hub seeds are easy to compute. For example for $n = 6$ we have (see Fig. 9): rank(643215) = 1, rank(632154) = 3, rank(621543) = 5, rank(615432) = 7, rank(654321) = 9.

Observation 18. All the values $|\mathbf{W}_k|$ and $\sum_{i=0}^{k}(|\mathbf{W}_i|+n-1)$ for $k \in \{0..n-4\}$ can be computed in $\mathcal{O}(n)$ total time and accessed in $\mathcal{O}(1)$ time, afterwards.

Define $\mathrm{sum}(k, j) = |\tau \cdot \prod_{i=1}^{j-1} \sigma^i\, \mathbf{W}_{\Delta(k,i)}\, \gamma_{n-2-i}| + j$.

Observation 19. **(a)** If $\phi = \mathrm{parent}(\psi) \notin \mathrm{Hub}_n$ and ψ is the i-th child of ϕ, then $\mathrm{rank}(\widetilde{\psi}) - \mathrm{rank}(\widetilde{\phi}) = \mathrm{sum}(\mathrm{height}(\phi), i)$. **(b)** After the linear time preprocessing sum(k, j) can be computed in $\mathcal{O}(1)$ using one of the formulas

$$1 + (|\mathbf{W}_{k-1}| + n - 1)\cdot(j-1) + j \quad\text{or}\quad |\mathbf{W}_k| - \sum_{i=0}^{n-j-2}(|\mathbf{W}_i| + n - 1) + j.$$

Lemma 20. *For a given permutation $\pi \in \mathrm{perms}(\psi)$, we can compute in time $\mathcal{O}(n)$*
(a) $\mathrm{rank}(\pi) - \mathrm{rank}(\widetilde{\psi})$*, if $\psi \notin \mathrm{Hub}_n$,* **(b)** $\mathrm{rank}(\pi)$*, otherwise.*

Proof. By a *starting path* we mean the sequence of the first $2n-2$ permutations of SW_n, see Fig. 9.
(a) Let $\rho = n\, r_2 \dots r_n$ be the permutation cyclically equivalent to π, which starts with n, and let j be the position such that $r_j = r_2 \oplus 1$. If $\psi = n\, r_3 \dots r_n$, then $\rho = \widetilde{\psi}$. Now we distinguish two cases depending in the position l of n in π.
Case 1: if $l \le n - j + 2$, then $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) - l + 1$.
Case 2: otherwise, $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) + |\mathbf{W}_{\mathrm{height}(\psi)}| + (n-1) - l + 1$
If $\psi = (n, r_2, \dots, r_{j-1}, r_{j+1}, \dots, r_n)$, then $\mathrm{rank}(\rho) = \mathrm{rank}(\widetilde{\psi}) + \mathrm{sum}(\mathrm{height}(\psi), n-j+1)$. Using similar arguments as before, we have $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) - l + 1$ if $l \le n-j+2$, and $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) + |\mathbf{W}_{\Delta(\mathrm{height}(\psi), n-j+1)}| + (n-1) - l + 1$, otherwise.
(b) If the permutation π, after removing n, is cyclically equivalent to $(n-1)(n-2)\dots 1$ and n appears in the first position (π belongs to the starting path), then $\mathrm{rank}(\pi) = 2\cdot(n - p_2 - 1) - 1$, and if it appears in the second position, then $\mathrm{rank}(\pi) = 2\cdot(n - p_1 - 1)$, where p_1, p_2 are the first two elements of π. Otherwise we define ρ, j and l like in the case (a) and $\psi = n\, r_2 \dots r_{j-1}\, r_{j+1} \dots r_n$. We know that $\pi \in \mathrm{perms}(\psi)$, and want to compute $\mathrm{rank}(\pi)$ minus the rank of the first permutation of $\mathrm{perms}(\psi)$, which appears in SW_n with the rank greater than $2n-3$. This permutation equals $\mu = r_2 \dots r_{j-1}\, r_{j+1} \dots r_n\, (r_2 \oplus 1)\, n$ and has the rank equal to

$$(2n-2) + |\mathbf{V}_n \tau| \cdot (r_2 \mod (n-1)) = 2n - 2 + (r_2 \; mod (n-1))\cdot(n(n-2)! - 2).$$

If $j = n$, then $\mathrm{rank}(\pi) = \mathrm{rank}(\mu) + n - l$ else, if $l \le n - j + 2$, then
$\mathrm{rank}(\rho) = \mathrm{rank}(\mu) + \mathrm{sum}(n-3, n-j+1) - |\mathbf{W}_{n-4}| - 2$, $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) - l + 1$.
Otherwise, we have, $\mathrm{rank}(\pi) = \mathrm{rank}(\rho) + |\mathbf{W}_{j-3}| + (n-1) - l + 1$.

Using these equations we can rank permutations in basic cases, and reduce the main problem to the ranking of the representatives of seeds. □

We slightly abuse notation and for a seed ψ define $\mathrm{rank}(\psi) = \mathrm{rank}(\widetilde{\psi})$.

Lemma 21. *For a non-hub seed ψ, the value* $\mathrm{rank}(\psi) - \mathrm{rank}(\mathrm{prehub}(\psi))$ *can be computed in $\mathcal{O}(n\sqrt{\log n})$ time.*

Proof. Let us consider the *parent sequence* $\psi = \psi_0, \psi_1, \psi_2, \dots, \psi_r = \mathrm{prehub}(\psi)$, where $r = \mathrm{level}(\psi) - 1$. Then $\mathrm{rank}(\psi_i) - \mathrm{rank}(\psi_{i+1}) = \mathrm{sum}(\mathrm{height}(\psi_{i+1}), \mathrm{ord}(\psi_i))$, and $\mathrm{height}(\psi_i) = \Delta(\mathrm{height}(\psi_{i+1}), \mathrm{ord}(\psi_i))$. Consequently, knowing route(ψ), we can compute in $\mathcal{O}(n)$ time $\mathrm{rank}(\psi) - \mathrm{rank}(\mathrm{prehub}(\psi)) = \sum_{i=m-1}^{0}(\mathrm{rank}(\psi_i) - \mathrm{rank}(\psi_{i+1}))$. Now the thesis is a consequence of Observation 19 and Lemma 16. □

The following result follows directly from Lemmas 20 and 21.

Theorem 22. [Ranking] *For a given permutation π, we can compute the rank of π in Seq_n in time $\mathcal{O}(n\sqrt{\log n})$.*

7. Unranking algorithm

Denote by unrank(t) the t-th permutation in Seq_n, and for $t < |\,\mathrm{bunch}(\psi)|$, let $\mathrm{unrank}(\psi,t) = \mathrm{unrank}(t + \mathrm{rank}(\widetilde{\psi}))$ (it is the t-th permutation in bunch(ψ), counting from the beginning of this bunch).

Lemma 23. *If we know a non-hub seed ψ such that* $\mathrm{unrank}(\psi,t) \in \mathrm{perms}(\psi)$, *then we can recover* $\mathrm{unrank}(\psi,t)$ *in linear time.*

We say that a permutation π is a hub permutation if $\pi \in \mathrm{perms}(\psi)$ for some $\psi \in Hub_n$.

Proof. Let $\psi = n\,a_2 \ldots a_{n-1}$, and $k = \mathrm{height}(\psi)$. We find in linear time the number j such that $\mathrm{sum}(k,j) - j \le t < \mathrm{sum}(k,j+1) - j - 1$.

If $\mathrm{sum}(k,j) < t < \mathrm{sum}(k,j) + |\mathbf{W}_{\Delta(k,j)}|$, then $\mathrm{unrank}(\psi,t)$ does not belong to perms(ψ) (it belongs to bunch(child(ψ, j))). If $l = \mathrm{sum}(k,j) - t \ge 0$, then $\mathrm{unrank}(\psi,t)$ equals $n\,a_2 \ldots a_{n-j}\,(a_2 \oplus 1)\,a_{n-j+1} \ldots a_{n-1}$, rotated by l to the right, and if $l = t - \mathrm{sum}(k,j) + |\mathbf{W}_{\Delta(k,j)}| \ge 0$, then it is the same permutation rotated by $l+1$ to the left.

In this way we reduced the problem of unranking a non-hub permutation to finding a package containing the permutation. □

Lemma 24. *We can test in $\mathcal{O}(n)$ time if* unrank(t) *is a hub permutation.*
(a) *If "yes", then we can recover* unrank(t) *in $\mathcal{O}(n)$ time.*
(b) *Otherwise, we can find in $\mathcal{O}(n)$ time a prehub ϕ together with* rank(ϕ) *such that* $\mathrm{unrank}(t) \in \mathrm{bunch}(\phi)$.

Proof. If $t < 2n-2$, then unrank(t) is equal to $(n-1)\ldots 1$ rotated to the left by $\lceil \frac{t}{2} \rceil$, with n inserted on first position if x is odd, and in the second if it is even.
Otherwise, let $t - (2n-2) = t_1 \cdot |\mathbf{V}_n \tau| + t_2$ (we use integer division to compute t_1, t_2), and let $\psi = n\,t_1\,(t_1 \ominus 1) \ldots (t_1 \ominus (n-3))$, (with t_1 substituted by $n-1$ if equal to 0). The permutation unrank(t) belongs to perms(ψ) or to bunch(ϕ), where parent(ϕ) = ψ. If $t_2 < n-2$, then unrank(t) is equal to $n\,t_1\,(t_1 \ominus 1) \ldots (t_1 \ominus (n-2))$, rotated to the left by $t_2 + 1$. Otherwise, we find j such that

$\mathrm{sum}(n-3,j) - j \le t_2 + (1 + |\mathbf{W}_{n-4}| + n - 1) - (n-2) < \mathrm{sum}(n-3,j+1) - j - 1$,
and $l = t_2 + |\mathbf{W}_{n-4}| + 2 - \mathrm{sum}(n-3,j)$, $\phi = \mathrm{child}(\psi, j)$

If $l \le 0$, then unrank(t) is equal to $\widetilde{\phi}$ rotated to the right by $-l$, else if $l \ge |\mathbf{W}_{\Delta(n-3,j)}|$, then unrank($t$) is equal to $\widetilde{\phi}$ rotated to the left by $l - |\mathbf{W}_{\Delta(n-3,j)}| + 1$. Otherwise, $\mathrm{unrank}(t) = \mathrm{unrank}(\phi, l)$, and it is not a hub permutation. By using this algorithm we either already succeeded in finding the required permutation, or restricted ourselves to a limited regular part of SW_n. □

Lemma 25. *After linear preprocessing, for given* height(ψ) *and $t \le |\,\mathrm{bunch}(\psi)|$, we can answer in $\mathcal{O}(\log\log n)$ time if* $\mathrm{unrank}(\psi,t) \in \mathrm{perms}(\psi)$. *If not, then we output the numbers* ord *and* height *of β – a descendant of ψ such that* $\mathrm{unrank}(\psi,t) \in \mathrm{bunch}(\beta)$.

Proof. Denote $m = n - 5$. For $k \in \{1,\ldots,m\}$, define $b_k = \sum_{i=0}^{k}(|\mathbf{W}_i| + n - 1)$. We have, $2 \le \frac{b_{i+1}}{b_i} \le n$ for $i < m$. Let $\delta = \sqrt[m]{b_m}$ and let d be the sequence such that $d[i] = \max\{j : b_j \le \delta^i\}$. In the preprocessing we compute this sequence in linear time. Then, t is located between positions $d[i]$ and $d[i+1]+1$ in $(b_1, b_2, \ldots, b_m)$, where $i = \lfloor \log_\delta t \rfloor$. In this way we find a logarithmic subrange, where the required t is located ($d[i+1] - d[i] \le \log_2 \delta = \mathcal{O}(\log n)$), and using binary search, we find t in $O(\log\log n)$ time.

Let $k = \mathrm{height}(\psi)$. We need j, such that $\mathrm{sum}(k,j) - j \le t < \mathrm{sum}(k,j+1) - (j+1)$. For $j \le n-k$, we have $\mathrm{sum}(k,j) - j = (j-1)\cdot(|\mathbf{W}_{k-1}| + n - 1)$.

Hence, if $t < \mathrm{sum}(k, n-k) - n + k$, then division by $|\mathbf{W}_{k-1}| + n - 1$ suffices to find the appropriate j. Otherwise, let $s = |\mathbf{W}_k| - t$, we look for j such that

$$b_{n-j-3} = |\mathbf{W}_k| - \mathrm{sum}(k,j+1) + j + 1 < s \le |\mathbf{W}_k| - \mathrm{sum}(k,j) + j = b_{n-j-2}.$$

We can find the required j in $\mathcal{O}(\log\log n)$ time. If $\mathrm{sum}(k,j) < t < \mathrm{sum}(k,j) + |\mathbf{W}_{\Delta(k,j)}|$, then $\mathrm{unrank}(\psi,t) = \mathrm{unrank}(\beta, t - \mathrm{sum}(k,j))$, where $\beta = \mathrm{child}(\psi,j)$ has the height $\Delta(k,j)$. Otherwise, $\mathrm{unrank}(\psi,t) \in \mathrm{perms}(\psi)$. □

Theorem 26. *For a given number t, we can compute the t-th permutation in $\mathcal{O}(n\frac{\log n}{\log\log n})$.*

Proof. From Lemma 24 we either obtain the required permutation (if it is a hub permutation) or a prehub ϕ together with rank(ϕ) such that $\mathrm{unrank}(t) \in \mathrm{bunch}(\phi)$. In the second case we know that unrank(t) equals $\mathrm{unrank}(\phi, t - \mathrm{rank}(\widetilde{\phi}))$. Now, after the linear preprocessing, we apply Lemma 25 to obtain route(ψ) for a seed ψ such that $\mathrm{unrank}(t) \in \mathrm{perms}(\psi)$.

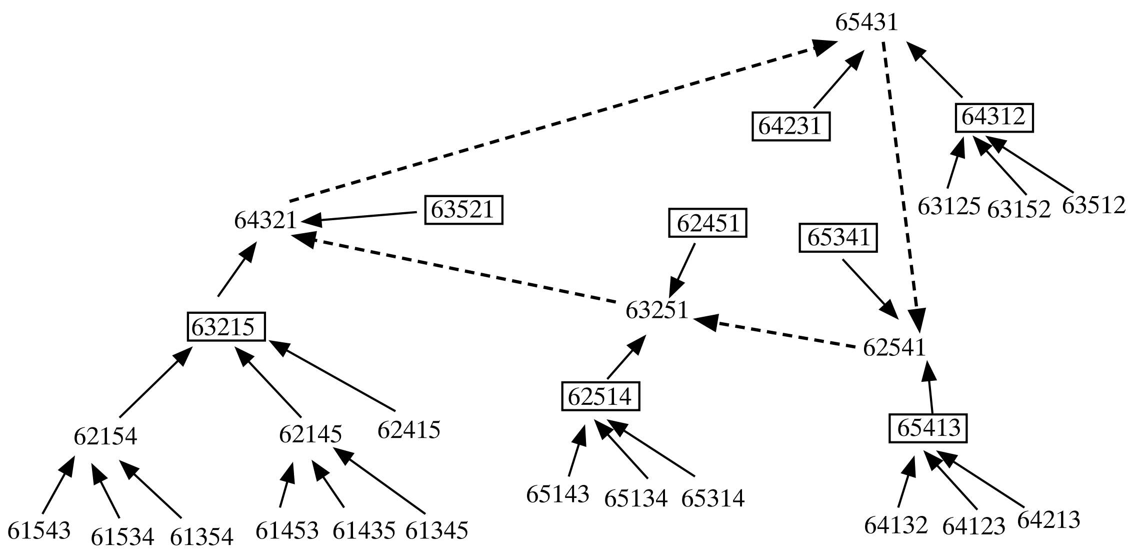

Fig. 11. The graph of modified seeds for $n = 6$. The prehub nodes are in boxes. Its hub nodes are 64321, 63251, 62541, 65431.

However, we do not know ψ. It can be computed using Lemma 17. Finally we use Lemma 23 to obtain the required permutation unrank(t). □

8. Cyclic $\sigma\tau$-generation of permutations

A $\sigma\tau$-sequence of permutations is cyclic if the last permutation is equal to the first one. Let $\oplus'$ denote a modified addition modulo $n-2$, where $(n-1)\oplus' 1 = 2$ (previously $n-1 \oplus 1 = 1$). It gives a cyclic order of the elements $\{2, ..., n-1\}$, with attached element 1 ($1 \oplus' 1 = 2$). For $a > 2$, we write $a \ominus' 1 = a-1$ and $2 \ominus' 1 = n-1$.

The change of the operation $\oplus$ to $\oplus'$ implies redefinition of seeds and their interaction. We understand later that in this section we deal with seeds and their graphs with respect to $\oplus'$. In particular for a seed $\psi = a_1 a_2 ... a_{n-1}$ mis$(\psi) = a_2 \oplus' 1$.

Observe that in the previous construction we had the seeds 5423, 5432, and now they are replaced by 5413, 5431, since $4 \oplus' 1 = 2$, while $4 \oplus 1 = 1$, for $n = 5$. The set of modified seeds is now {5413, 5431, 5134, 5143, 5214, 5241, 5312, 5321} (see Fig. 11).

Lemma 27. *The subtree rooted at a non-hub seed ψ of height k in the graph of modified seeds is isomorphic to* SkewTree(n, k). *Furthermore,* $\mathbf{W}_k$ *generates the Hamiltonian path on* bunch(ψ).

Proof. It is enough to show, that the children of a seed with a newly defined height equal to $k > 1$ follow the same pattern of heights as in Lemma 6. Let $x = \text{mis}(\psi)$. When $x \oplus' 1 \neq 2$, the proof does not change, since for a non-hub seed, $x \oplus' 1$ cannot belong to the decreasing prefix. x cannot be equal to 1 as it has no predecessor (it cannot be a *missing* element) and for $x = n-1$ only placing it after 2 can make any difference. There is only one seed with $x = n-1$ and with the element 2 in the decreasing prefix - $n\,(n-2)\,(n-3)\ldots 2\,1$, and it is a hub seed. However the child obtained in this way is very special - it is the only non-hub seed of height $n-3$ (such a seed does not appear for operation $\oplus$). We know that for $\phi = \text{child}(\psi, i)$ $\widetilde{\phi} = \sigma^i(\psi^{(i)})$, hence, using the morphism Ψ on the traversal sequence of the skew tree we obtain the ordering of the Hamiltonian path. □

For $1 \le x < n$, we define switch-permutations (*switches*, in short) as permutations of the form $x\,n\,(x\oplus 1)\,(x\oplus 2)\ldots(x\ominus 1)$. In other words they are cyclic shifts of the identity permutation $1\,2\,3\ldots(n-1)$ in which n is inserted into the second position. Let r_i be the switch starting with $n-i$, and let u_i denote $\sigma(r_i)$. Denote by $\mathcal{R}'_n$ the graph consisting of τ-edges of all ham(ψ) and all σ-edges not conflicting with selected τ-edges, where each ham(ψ) is understood here as the one defined with the use of the modified operator $\oplus'$. A version of the construction of SW′ in [19] can be written informally as

Algorithm ComputeCycle
$P := \mathcal{R}'_n$
for each switch r_i **do**
  remove edge $r_i \to \sigma(r_i)$; add edge $r_i \to \tau(r_i)$;
return P

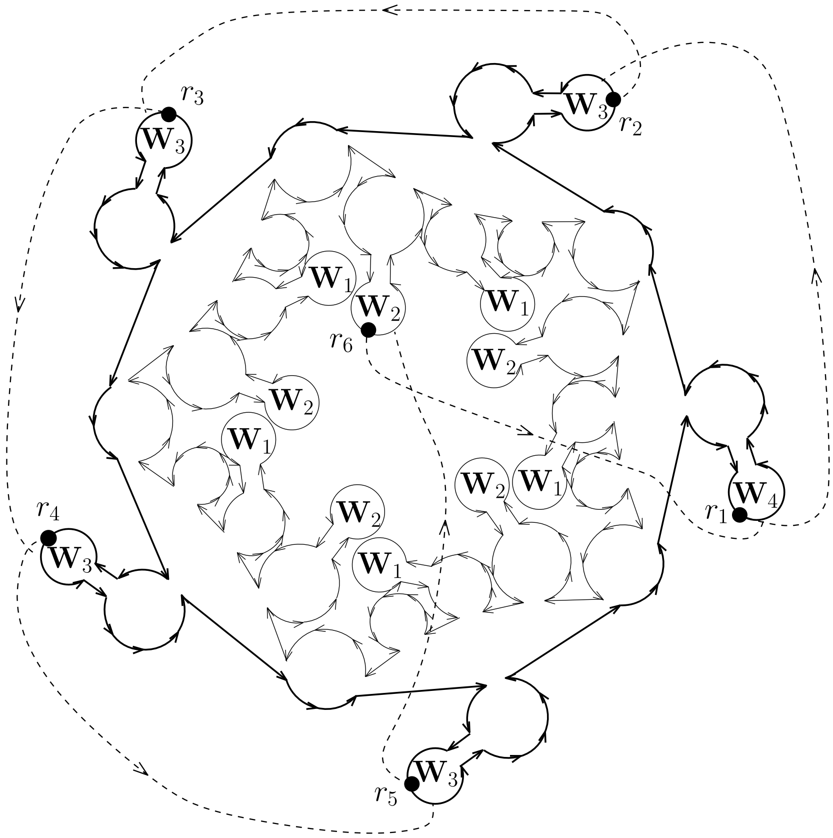

Fig. 12. The cycles C_{inner} and C_{outer} for $n = 7$ (edges of the outer cycle are drawn as thicker lines). We obtain Seq'_7, by adding dashed edges, and removing (old) edges conflicting with the dashed ones.

It was shown in [19] that $\mathcal{R}'_n$ consists of two disjoint cycles C_{inner} (the one containing r_{n-1}), and C_{outer}. Denote by Cycle_n the Hamiltonian cycle constructed in [19] (see Fig. 12).

Lemma 28. [19] *For odd n, the algorithm ComputeCycle produces the Hamiltonian cycle Cycle_n which is of the form*

$$r_1 \xrightarrow{\tau} u_2 \xrightarrow{\beta_2} r_3 \xrightarrow{\tau} u_4 \xrightarrow{\beta_4} r_5 \xrightarrow{\tau} \dots \xrightarrow{\beta_{n-3}} r_{n-2} \xrightarrow{\tau} u_{n-1} \xrightarrow{\beta_{n-1}} r_{n-1}$$
$$\xrightarrow{\tau} u_1 \xrightarrow{\beta_1} r_2 \xrightarrow{\tau} u_3 \xrightarrow{\beta_3} r_4 \xrightarrow{\tau} u_5 \xrightarrow{\beta_5} r_6 \xrightarrow{\tau} \dots \xrightarrow{\tau} u_{n-2} \xrightarrow{\beta_{n-2}} r_1,$$

and $\mathsf{Seq}'_n = \tau\,\beta_2\,\tau\,\beta_4\,\tau\,\beta_6\,\dots,\,\tau\,\beta_{n-1}\,\tau\,\beta_1\,\tau\,\beta_3\,\tau\,\beta_5\,\dots\,\tau\,\beta_{n-2}$.

Lemma 29. *We have* $\mathsf{Seq}_{\text{outer}} = (\sigma\mathbf{W}_{n-3}\gamma_{n-3}\gamma_2)(\sigma\mathbf{W}_{n-4}\gamma_{n-3}\gamma_2)^{n-3}$, *and*

$$\mathsf{Seq}_{\text{inner}} = \mathbf{U}^{n-2}\text{, where } \mathbf{U} = \gamma_{n-4}\cdot\prod_{i=3}^{n-2}\sigma^i\,\mathbf{W}_{\Delta(n-3,i)}\,\gamma_{n-2-i}.$$

Proof. From Lemma 27 we know that for ψ outside of the hub cycle, we have

$$\text{gen}(\widetilde{\psi}, \mathbf{W}_k) = \text{bunch}(\psi) \;\text{ and }\; \mathbf{W}_k(\widetilde{\psi}) = \psi^{(n-1)}.$$

Every hub seed ϕ has exactly one child which is also a hub seed. In the previous construction it was always child$(\phi, 1)$. In this construction it is child$(\phi, 2)$, as hub seeds are of the form $n\,x\,(x \ominus' 1) \dots (x \oplus' 2)\,1$. Hence, the construction separates the first child from children 3 to $n-3$ (and the cycle with $\mathbf{W}_0$). The outer cycle covers the first children of hub seeds, and the inner cycle covers the remaining ones. The first child of each hub seed has the height equal to $n-4$ with one exception – the seed $n\,(n-3)\,(n-4)\dots 1\,(n-1)$, which is the first child of the hub seed $n\,(n-2)\,(n-3)\dots 1$ and has the height equal to $n-3$. Hence the outer cycle has the required representation (additional γ_2 represents the transition to the second child, which is the next hub seed). In the inner cycle each child of the hub seed has the same height as in the previous construction (height(child$(\phi, i)) = \Delta(n-3, i)$), and these children are visited in the same order. Afterwards, we follow the sequence of operations γ_{n-4}, which represents the “return to the parent seed” (thus, in this cycle, hub seeds are visited in the reversed order). $\square$

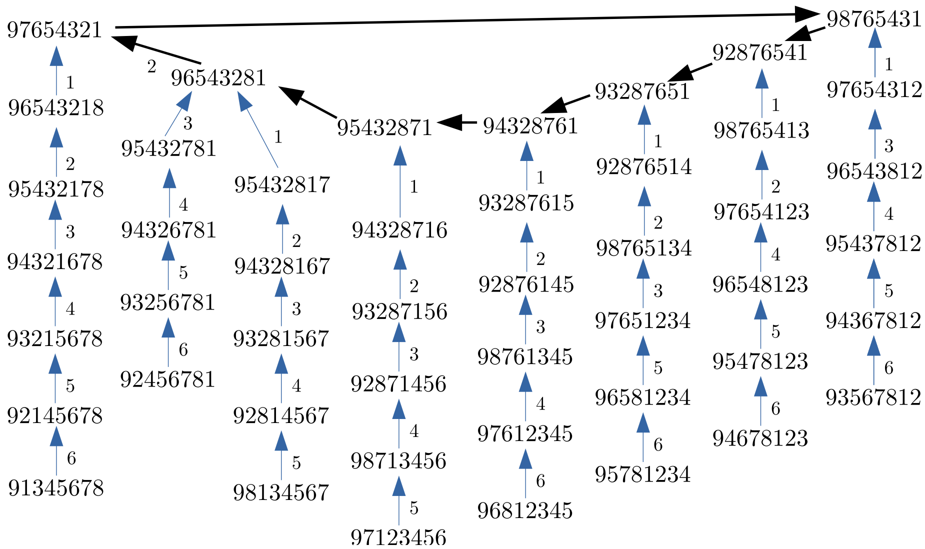

Fig. 13. The routes of 8 switch seeds for $n = 9$. The hub seeds are in the upper cycle. The routes have very similar patterns - they consist of two simple arithmetic progressions.

The seeds corresponding to switch permutations are called *switch seeds*. Observe that each switch seed ψ is of the form $a_1\, a_2 \ldots a_{n-1}$, where $a_2 \ldots a_{n-1}$ is a cyclic shift of $1\,2\,3 \ldots (n-1)$ with the successor of the first element removed (see Fig. 13).

Lemma 30. *A route of a switch seed is either equal to* $(n-3, n-4, \ldots, 3)$ *or to* $(n-3, n-4, \ldots, 1)$ *with* 0 *or* 1 *element erased (for example* $(6, 5, 3, 2, 1)$ *for* $n = 9$*).*

Proof. A switch $r_{n-x+1} = (x \ominus 1)\, n\, x\, (x \oplus 1) \ldots (x \ominus 2)$ belongs to $\text{perms}(\psi)$ for a switch seed $\psi = n\, x\, (x \oplus 2)\, (x \oplus 3) \ldots (x \ominus 1)$ of height equal to 1 (for $x = n-1$ $\psi = n\,(n-1)\,1\,3 \ldots (n-2)$). When going to $\text{hub}(\psi)$ through the parent edges in each step we take an element one greater (in terms of $\oplus'$) and place it after n (removing the new missing element). Each such action shifts by 1 all elements between the first position and the position from which the new missing element was removed. These elements are already sorted (and will not be removed later). Hence the route of the permutation is composed of the initial positions of elements $x \oplus' 2, x \oplus' 3, \ldots$ (counting from the end). For $x = 1$, all elements from 3 to $n-1$ are removed in some step, which results in $(n-3, n-4, \ldots, 1)$, as the route (positions of elements 3 to $n-1$ in ψ, counting from the right). For $x > 1$, the element 1 appears in the position $x-1$ counting from the end in ψ (the position $n-3$ for $x = n-1$), hence this value is *missing* from its route (route is equal to $(n-3, \ldots, 1)$ with $x-1$ removed). For $x = 2$, the route $(n-3, \ldots, 2)$ would lead to a hub seed $n\,(n-2) \ldots 2\,1$, however its second child $n\,(n-3)\,(n-4) \ldots 3\,2\,(n-1)\,1$ is already a hub seed, hence we need to remove the value 2 as well (obtaining $\text{route}(\pi) = (n-3, \ldots, 3)$). $\square$

Theorem 31. *The size of the minimal SLP representation of* $\mathcal{S}eq'_n$ *is* $\Theta(n^2)$.

Proof. It is enough to give SLP representations of total size $\mathcal{O}(n^2)$ for all β_i. We consider only the case $2 < i < n-2$ (four other cases are very similar). We divide β_i into three parts w_1^i, w_2, w_3^{i+1} such that $u_i \xrightarrow{w_1^i} a_i^{(n-1)} \xrightarrow{w_2} \widetilde{a_{i+1}} \xrightarrow{w_3^{i+1}} r_{i+1}$, where $a_i = \text{prehub}(s_i)$ and s_i is a switch seed corresponding to r_i. For $2 < i < n-2$, all words w_1^i are suffixes of $\mathbf{W}_{n-4}$, and all w_3^{i+1} are prefixes of $\mathbf{W}_{n-4}$, while $w_2 = \gamma_{n-3}\gamma_2\sigma$. We use the fact, that for a non-prehub ancestor of a switch seed of height k, its parent is of height $k+1$, furthermore, it is always either the last or the penultimate son of this height.
This implies that for $k > 1$ a split in equation for $\mathbf{W}_k$ can only happen in two places (in one of the two $\mathbf{W}_{k-1}$). We decompose $\mathbf{W}_k = \tau \prod_{i=1}^{n-2} \sigma^i \mathbf{W}_{\Delta(k,i)}\gamma_{n-2-i}$ as $\mathbf{W}_k^{(1)}(\mathbf{W}_{k-1}\gamma_k\sigma^{n-k-1}\mathbf{W}_{k-1})\mathbf{W}_k^{(2)}$, without changing the total size of equations in the SLP representation, where

$$\mathbf{W}_k^{(1)} = (\tau \prod_{i=1}^{n-k-3} \sigma^i \mathbf{W}_{\Delta(k,i)}\gamma_{n-2-i})\sigma^{n-k-2}, \quad \mathbf{W}_k^{(2)} = \gamma_{k-1} \prod_{i=n-k}^{n-2} \sigma^i \mathbf{W}_{\Delta(k,i)}\gamma_{n-2-i}.$$

Intuitively, we divide the product over $i \in \{1, \ldots, n-2\}$ into three parts
$\{1, \ldots, n-k-3\}$, $\{n-k-2, n-k-1\}$, $\{n-k, \ldots, n-2\}$.
We want to represent the two parts of each $\mathbf{W}_{n-4}$ with a split. Due to the fact that the split on each level (for k from $n-4$ down to 2) happens in the middle part of $\mathbf{W}_k$, and none of the $\mathbf{W}_k^{(1)}$ or $\mathbf{W}_k^{(2)}$ parts are affected by the split, they do not need to be altered. The middle parts of each $\mathbf{W}_k$ may be partitioned in two different ways (depending on i), however these parts are of constant size each.

For $i \neq 2$, u_i is always positioned two permutations before $s_i^{(n-1)}$ ($s_i^{(n-1)}$ is obtained from u_i after applying $\sigma\tau$), hence $\mathbf{W}_1$ is always divided into $\mathbf{W}_1^{(1)}\mathbf{W}_0\gamma_1\sigma^{n-3}$ and $\mathbf{W}_0\mathbf{W}_1^{(2)}$ (the same division into three parts as on the higher levels, with a minor difference inside cycle(r_i)). More formally, each w_3^i is composed of $n-4$ parts equal either to $\mathbf{W}_k^{(1)}$ or to $\mathbf{W}_k^{(1)}\mathbf{W}_{k-1}\gamma_k\sigma^{n-k-1}$ and of $\mathbf{W}_0\gamma_1\sigma^{n-3}$ (similarly w_1^i is composed of $\mathbf{W}_0$ and of $n-4$ parts equal to either $\gamma_k\sigma^{n-k-1}\mathbf{W}_{k-1}\mathbf{W}_k^{(2)}$ or $\mathbf{W}_k^{(2)}$).

Each such production has linear size, hence all equations for β_i have $\mathcal{O}(n^2)$ total size. The lower bound follows from Lemma 13 in the same way as for the path sequence Seq_n (Theorem 14), since Seq'_n contains $\mathbf{W}_{n-4}$. This completes the proof. □

Theorem 32. *In the sequence Seq'_n the ranking can be done in $\mathcal{O}(n\sqrt{\log n})$ time and the unranking in $\mathcal{O}(n\frac{\log n}{\log\log n})$ time.*

Proof. Ranking and unranking in Seq_{outer} and Seq_{inner} is basically the same as in Seq_n. We show how to compute ranks of all r_i in Seq_{outer}, Seq_{inner} and in Seq'_n in $\mathcal{O}(n)$ total time. By Lemma 30 we know the routes of all the switch seeds, hence ranking of each of them (in Seq_{outer} or Seq_{inner}) can be computed in $\mathcal{O}(n)$ time. Furthermore for $2 < i < n-3$ the route of the next switch seed differs on exactly one position. Hence we can rank r_{i+1} for those i in $\mathcal{O}(n)$ total time using the formula,

$$\mathrm{rank}(r_{i+1}) - \mathrm{rank}(r_i) = |\sigma\mathbf{W}_{n-4}\gamma_{n-3}\gamma_2| + \mathrm{sum}(i-1, n-i) - \mathrm{sum}(i-1, n-i-1) = |\mathbf{W}_{n-4}| + |\mathbf{W}_{i-2}| + 2n + 2.$$

These differences give the values of $|\beta_i|$ for each i, which can be used to compute the ranks of switches in Seq'_n. Now we can perform ranking and unranking.

Ranking: to obtain the rank of π in Seq'_n, we rank it inside Seq_{outer} or Seq_{inner}, and then add the difference between rankings of the next r_i in both constructions ($\pi \in \beta_i$).

Unranking: to find permutation of the rank t in Seq'_n, we find r_i with the smallest rank greater than or equal to t. We can then unrank $|\beta_i| - |r_i| + t$ in β_i. □

Declaration of competing interest

We wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

References

[1] Ignacio R. Asenjo, Method for solving optimization problems in structured combinatorial objects, U.S. Patent 8,266,089, 2012.
[2] Blai Bonet, Efficient algorithms to rank and unrank permutations in lexicographic order, AAAI Workshop - Technical Report, 01.2008.
[3] Timothy M. Chan, Mihai Patrascu, Counting inversions, offline orthogonal range counting, and related problems, in: Moses Charikar (Ed.), SODA, pp. 161–173.
[4] Paul F. Dietz, Optimal algorithms for list indexing and subset rank, in: Algorithms and Data Structures, Workshop WADS '89, Ottawa, Canada, August 17-19, 1989, Proceedings, 1989, pp. 39–46.
[5] Gideon Ehrlich, Loopless algorithms for generating permutations, combinations, and other combinatorial configurations, J. ACM 20 (3) (1973) 500–513, https://doi.org/10.1145/321765.321781.
[6] N. Deo, E.M. Reingold, J. Nievergelt, Combinatorial Algorithms: Theory and Practice, Prentice-Hall, Englewood Cliffs, NJ, 1977.
[7] Even Shimon, Algorithmic Combinatorics, Macmillan, New York, 1973.
[8] Elizabeth Hartung, Hung Phuc Hoang, Torsten Mütze, Aaron Williams, Combinatorial generation via permutation languages. i. Fundamentals, CoRR, http://arxiv.org/abs/1906.06069, 2019.
[9] Donald E. Knuth, The Art of Computer Programming, Combinatorial Algorithms, vol. 4, Addison-Wesley, 2010. Part 1.
[10] Donald L. Kreher, Douglas R. Stinson, Combinatorial Algorithms: Generation, Enumeration, and Search, CRC Press, Rockville, MD, 1999.
[11] Wendy J. Myrvold, Frank Ruskey, Ranking and unranking permutations in linear time, Inf. Process. Lett. 79 (6) (2001) 281–284, https://doi.org/10.1016/S0020-0190(01)00141-7.
[12] Frank Ruskey, Combinatorial Generation, University of Victoria, Victoria, BC, Canada, 2003. Preliminary working draft.
[13] Frank Ruskey, Combinatorial Gray code, in: Encyclopedia of Algorithms, Springer, New York, NY, 2016, pp. 342–347.
[14] Frank Ruskey, Aaron Williams, An explicit universal cycle for the (n-1)-permutations of an n-set, ACM Trans. Algorithms 6 (3) (2010) 45.
[15] Wojciech Rytter, Grammar compression, LZ-encodings, and string algorithms with implicit input, in: Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, 2004, pp. 15–27.
[16] Wojciech Rytter, Wiktor Zuba, Syntactic view of sigma-tau generation of permutations, in: Carlos Martín-Vide, Alexander Okhotin, Dana Shapira (Eds.), International Conference, LATA 2019, in: Lecture Notes in Computer Science, vol. 11417, Springer, 2019, pp. 447–459.
[17] Joe Sawada, Aaron Williams, Greedy flipping of pancakes and burnt pancakes, Discrete Appl. Math. 210 (2016) 61–74, https://doi.org/10.1016/j.dam.2016.02.005.
[18] Joe Sawada, Aaron Williams, A Hamilton path for the sigma-tau problem, in: SODA, 2018, pp. 568–575.
[19] Joe Sawada, Aaron Williams, Solving the sigma-tau problem, ACM Trans. Algorithms 16 (2020) 11, https://doi.org/10.1145/3359589.
[20] Xi Sisi Shen, Aaron Williams, A 'hot potato' Gray code for permutations, Electron. Notes Discrete Math. 44 (2013) 89–94, https://doi.org/10.1016/j.endm.2013.10.014.
[21] Aaron Williams, Hamiltonicity of the Cayley digraph on the symmetric group generated by sigma = (1 2... n) and tau = (1 2), arXiv:1307.2549, 2013.
[22] Shmuel Zaks, A new algorithm for generation of permutations, BIT Numer. Math. 24 (2) (1984) 196–204, https://doi.org/10.1007/BF01937486.