



ELSEVIER

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta



The Möbius function of separable and decomposable permutations[☆]

Alexander Burstein^a, Vít Jelínek^b, Eva Jelínková^c, Einar Steingrímsson^d

^a Department of Mathematics, Howard University, Washington, DC 20059, USA

^b Fakultät für Mathematik, Universität Wien, Garnisonsgasse 3, 1090 Wien, Austria

^c Department of Applied Mathematics, Charles University in Prague, Malostranské nám. 25, 110 00 Prague, Czech Republic

^d Department of Computer and Information Sciences, University of Strathclyde, Glasgow G1 1XH, UK

ARTICLE INFO

Article history:

Received 7 February 2011

Available online xxxx

Keywords:

Möbius function

Pattern poset

Decomposable permutations

Separable permutations

ABSTRACT

We give a recursive formula for the Möbius function of an interval $[\sigma, \pi]$ in the poset of permutations ordered by pattern containment in the case where π is a decomposable permutation, that is, consists of two blocks where the first one contains all the letters $1, 2, \dots, k$ for some k . This leads to many special cases of more explicit formulas. It also gives rise to a computationally efficient formula for the Möbius function in the case where σ and π are separable permutations. A permutation is separable if it can be generated from the permutation 1 by successive sums and skew sums or, equivalently, if it avoids the patterns 2413 and 3142. We also show that the Möbius function in the poset of separable permutations admits a combinatorial interpretation in terms of normal embeddings among permutations. A consequence of this interpretation is that the Möbius function of an interval $[\sigma, \pi]$ of separable permutations is bounded by the number of occurrences of σ as a pattern in π . Another consequence is that for any separable permutation π the Möbius function of $(1, \pi)$ is either 0, 1 or -1 .

© 2011 Elsevier Inc. All rights reserved.

[☆] Jelínek and Steingrímsson were supported by grant No. 090038012 from the Icelandic Research Fund. Jelínek was also supported by grant Z130-N13 from the Austrian Science Foundation (FWF). Jelínková was supported by project 1M0021620838 of the Czech Ministry of Education.

E-mail addresses: aburstein@howard.edu (A. Burstein), jelinek@kam.mff.cuni.cz (V. Jelínek), eva@kam.mff.cuni.cz (E. Jelínková), enar@alum.mit.edu (E. Steingrímsson).

1. Introduction

Let S_n be the set of permutations of the integers $\{1, 2, \dots, n\}$. The union of all S_n for $n = 1, 2, \dots$ forms a poset \mathcal{P} with respect to *pattern containment*. That is, we define $\sigma \leq \pi$ in \mathcal{P} if there is a subsequence of π whose letters are in the same order of size as the letters in σ . For example, $132 \leq 24153$, because 2, 5, 3 appear in the same order of size as the letters in 132. We denote the number of occurrences of σ in π by $\sigma(\pi)$, for example $132(24153) = 3$, since 243, 253 and 153 are all the occurrences of the pattern 132 in 24153.

A classical goal in the study of a combinatorially defined poset is to determine its Möbius function. For our poset \mathcal{P} this seems to have first been mentioned explicitly by Wilf [8]. The first result in this direction was given by Sagan and Vatter [5], who showed that an interval $[\sigma, \pi]$ of *layered* permutations is isomorphic to a certain poset of compositions of an integer, and they gave a formula for the Möbius function in this case. A permutation is layered if it is the concatenation of decreasing sequences, such that the letters in each sequence are smaller than all letters in subsequent sequences. Further results were given by Steingrímsson and Tenner [7], who showed that the Möbius function $\mu(\sigma, \pi)$ is 0 whenever the complement of the occurrences of σ in π contains an interval block, that is, when π has a segment of two or more consecutive letters that form a segment of values, where none of these consecutive letters belongs to any occurrence of σ in π . One example of such a pair is (132, 598342617), where the letters 342 do not belong to any occurrence of 132 in 598342617. Steingrímsson and Tenner [7] also described certain intervals where the Möbius function is either 1 or -1 .

In this paper, we focus on permutations that can be expressed as direct sums or skew sums of smaller permutations. A *direct sum* of two permutations α and β , denoted by $\alpha \oplus \beta$, is the concatenation $\alpha\beta'$, where β' is obtained by incrementing each element of β by $|\alpha|$. For example, 31426587 can be written as a direct sum $3142 \oplus 2143$. Similarly, a *skew sum* $\alpha \ominus \beta$ is the concatenation $\alpha'\beta$ where α' is obtained by incrementing α by $|\beta|$.

A permutation that can be written as a direct sum of two nonempty permutations is *decomposable*, and a permutation that can be written as a skew sum of nonempty permutations is *skew-decomposable*. The *decomposition* of a permutation π is an expression $\pi = \pi_1 \oplus \pi_2 \oplus \dots \oplus \pi_k$ in which each summand π_i is indecomposable. The summands π_1, \dots, π_k will be called *the components* of π . Every permutation has a unique decomposition (including an indecomposable permutation π , whose decomposition has a single component equal to π).

A permutation is *separable* if it can be obtained from the singleton permutation 1 by iterating direct sums and skew sums (for an alternative definition see Section 2).

Our main result is a set of recurrences for computing the Möbius function $\mu(\sigma, \pi)$ when π is decomposable. If $\pi_1 \oplus \dots \oplus \pi_k$ is the decomposition of π , then these recurrences express $\mu(\sigma, \pi)$ in terms of Möbius functions involving the summands π_i .

In the special case when π is separable, these recurrences provide a polynomial-time algorithm to compute $\mu(\sigma, \pi)$. These recurrences also allow us to obtain an alternative combinatorial interpretation of the Möbius function of separable permutations, based on the concept of 'normal embeddings'. This interpretation of μ generalizes previous results of Sagan and Vatter [5] for layered permutations.

Using these expressions of the Möbius function in terms of normal embeddings, we derive several bounds on the values of $\mu(\sigma, \pi)$ where σ and π are separable. In [7], Steingrímsson and Tenner conjectured that for permutations σ and π avoiding the pattern 132 (or any one of the patterns 213, 231, 312) the absolute value of the Möbius function of the interval $[\sigma, \pi]$ is bounded by the number of occurrences of σ in π . We prove this conjecture for the more general class of separable permutations (for non-separable σ and π this bound does not hold in general). In particular, if π has a single occurrence of σ then $\mu(\sigma, \pi)$ is either 1, 0 or -1 . We also prove a generalization of another conjecture mentioned in [7], showing that for any separable permutation π , $\mu(1, \pi)$ is either 1, 0 or -1 .

For a non-separable decomposable permutation π , our recurrences are not sufficient to compute the value of $\mu(\sigma, \pi)$. However, our results imply that the Möbius function can be efficiently computed in any hereditary class of permutations in which every sufficiently large element is decomposable or skew-decomposable.

Moreover, we are able to give short simple formulas in many special cases. For instance, suppose that σ is indecomposable and that π is decomposable and of length at least 3. Then we show that $\mu(\sigma, \pi)$ can only be nonzero if all the components in the decomposition of π are equal to the same permutation $\pi' > 1$, except possibly the first and the last component, which may be equal to 1. In such cases, $\mu(\sigma, \pi)$ equals $(-1)^i \mu(\sigma, \pi')$, where $i \in \{0, 1, 2\}$ is the number of components of π that are equal to 1.

As another simple example, our results imply that when σ and π are permutations with decompositions $\sigma = \sigma_1 \oplus \sigma_2$ and $\pi = \pi_1 \oplus \pi_2$, with π_1 and π_2 both different from 1, then $\mu(\sigma, \pi) = \mu(\sigma_1, \pi_1)\mu(\sigma_2, \pi_2)$ if $\pi_1 \neq \pi_2$, and $\mu(\sigma, \pi) = \mu(\sigma_1, \pi_1)\mu(\sigma_2, \pi_2) + \mu(\sigma, \pi_1)$ if $\pi_1 = \pi_2$.

The paper is organized as follows: In the next section we provide necessary definitions. In Section 3 we present the main results, the recursive formulas for reducing the computation of the Möbius function of decomposable permutations to that of indecomposable permutations. Section 4 deals with the case of separable permutations and their normal embeddings. Finally, in Section 5 we mention some open problems, in particular questions about the topology of the order complexes of intervals in our poset, which we have not dealt with in the present paper.

2. Definitions and preliminaries

An interval $[\sigma, \pi]$ in a poset (\mathcal{P}, \leq) is the set $\{\rho: \sigma \leq \rho \leq \pi\}$. In this paper, we deal exclusively with intervals of the poset of permutations ordered by pattern containment.

The Möbius function $\mu(\sigma, \pi)$ of an interval $[\sigma, \pi]$ is uniquely defined by setting $\mu(\sigma, \sigma) = 1$ for all σ and requiring that

$$\sum_{\rho \in [\sigma, \pi]} \mu(\sigma, \rho) = 0 \quad (1)$$

for every $\sigma < \pi$. When $\sigma \not\leq \pi$, we define $\mu(\sigma, \pi)$ to be zero.

An equivalent definition is given by Philip Hall's Theorem [6, Proposition 3.8.5], which says that

$$\mu(\sigma, \pi) = \sum_{C \in \mathcal{C}(\sigma, \pi)} (-1)^{L(C)} = \sum_i (-1)^i c_i, \quad (2)$$

where $\mathcal{C}(\sigma, \pi)$ is the set of chains in $[\sigma, \pi]$ that contain both σ and π , $L(C)$ denotes the length of the chain C , and c_i is the number of such chains of length i in $[\sigma, \pi]$. A chain of length i in a poset is a set of $i + 1$ pairwise comparable elements $x_0 < x_1 < \dots < x_i$. For details and further information, see [6].

Recall that a permutation is *separable* if it can be generated from the permutation 1 by iterated sums and skew sums. In other words, a permutation is separable if and only if it is equal to 1 or it can be expressed as a sum or skew sum of separable permutations.

Being separable is equivalent to avoiding the patterns 2413 and 3142, that is, containing no occurrences of either. Separable permutations have nice algorithmic properties. For instance, Bose, Buss and Lubiw [2] have shown that when σ and π are separable, it can be decided whether $\sigma \leq \pi$ in time that is polynomial in $|\sigma| + |\pi|$, while for general permutations the problem is NP-hard.

It is sometimes convenient to allow permutations to have zero length, while in other situations, the permutations are assumed to be nonempty. The unique permutation of length 1 is denoted by 1, and the unique permutation of length 0 is denoted by \emptyset . We make it a convention that the permutation \emptyset is neither decomposable nor indecomposable. In other words, whenever we say that a permutation π is decomposable (or indecomposable), we automatically assume that π is nonempty.

Suppose that π is a nonempty permutation with decomposition $\pi_1 \oplus \dots \oplus \pi_n$. For an integer $i \in \{0, \dots, n\}$, we let $\pi_{\leq i}$ denote the sum $\pi_1 \oplus \pi_2 \oplus \dots \oplus \pi_i$ and let $\pi_{> i}$ denote the sum $\pi_{i+1} \oplus \dots \oplus \pi_n$. An empty sum of permutations is assumed to be equal to \emptyset , so in particular $\pi_{\leq 0} = \pi_{> n} = \emptyset$. Any permutation of the form $\pi_{\leq i}$ for some i will be called a *prefix* of π , and any permutation of the form $\pi_{> i}$ is a *suffix* of π . Note that $\mu(\emptyset, \emptyset) = 1$, $\mu(\emptyset, 1) = -1$, and it is easily seen that $\mu(\emptyset, \tau) = 0$ for any $\tau > 1$.

3. The main results

Throughout this section, we assume that σ is a nonempty permutation with decomposition $\sigma_1 \oplus \cdots \oplus \sigma_m$ and that $\pi = \pi_1 \oplus \cdots \oplus \pi_n$ is a decomposable permutation (so $n \geq 2$ and, in particular, π is nonempty). The goal in this section is to prove a set of recurrences that allow us to express the Möbius function $\mu(\sigma, \pi)$ in terms of the values of the form $\mu(\sigma', \pi')$, where $\pi' \in \{\pi_1, \dots, \pi_n\}$ is a component of π and σ' is a sum of consecutive components of σ . Note that σ may itself be indecomposable, in which case m is equal to 1 and $\sigma_1 = \sigma$.

There are two main recurrences to prove, dealing respectively with the cases $\pi_1 = 1$ and $\pi_1 > 1$.

Proposition 1 (First recurrence). *Let σ and π be nonempty permutations with decompositions $\sigma = \sigma_1 \oplus \cdots \oplus \sigma_m$ and $\pi = \pi_1 \oplus \cdots \oplus \pi_n$, where $n \geq 2$. Suppose that $\pi_1 = 1$. Let $k \geq 1$ be the largest integer such that all the components π_1, \dots, π_k are equal to 1, and let $\ell \geq 0$ be the largest integer such that all the components $\sigma_1, \dots, \sigma_\ell$ are equal to 1. Then*

$$\mu(\sigma, \pi) = \begin{cases} 0 & \text{if } k-1 > \ell, \\ -\mu(\sigma_{>k-1}, \pi_{>k}) & \text{if } k-1 = \ell, \\ \mu(\sigma_{>k}, \pi_{>k}) - \mu(\sigma_{>k-1}, \pi_{>k}) & \text{if } k-1 < \ell. \end{cases}$$

Note that the suffixes $\sigma_{>k-1}$, $\sigma_{>k}$ and $\pi_{>k}$ in the statement of Proposition 1 may be empty. This first recurrence shows how to compute the Möbius function when π starts with $123 \cdots k$ for some $k \geq 1$. The second recurrence takes care of the remaining cases, that is, when π does not start with 1.

Proposition 2 (Second recurrence). *Let σ and π be nonempty permutations with decompositions $\sigma = \sigma_1 \oplus \cdots \oplus \sigma_m$ and $\pi = \pi_1 \oplus \cdots \oplus \pi_n$, where $n \geq 2$. Suppose that $\pi_1 > 1$. Let $k \geq 1$ be the largest integer such that all the components π_1, \dots, π_k are equal to π_1 . Then*

$$\mu(\sigma, \pi) = \sum_{i=1}^m \sum_{j=1}^k \mu(\sigma_{\leq i}, \pi_1) \mu(\sigma_{>i}, \pi_{>j}). \quad (3)$$

Note that Propositions 1 and 2 remain true when all the direct sums are replaced with skew sums, and the decompositions are replaced with skew decompositions. To see this, it is enough to observe that if $\bar{\pi}$ denotes the reversal of π (i.e., $\bar{\pi}$ is the permutation obtained by reversing the order of elements of π), then $\mu(\sigma, \pi) = \mu(\bar{\sigma}, \bar{\pi})$ for any σ and π , since $[\sigma, \pi]$ and $[\bar{\sigma}, \bar{\pi}]$ are isomorphic posets.

Before we prove the above two recurrences, we give three corollaries to provide some idea of how the second recurrence can be used. When we write $k \times \pi$ we mean a sum $\pi \oplus \pi \oplus \cdots \oplus \pi$ with k summands.

Corollary 3. *Let σ , π and k be as in Proposition 2, and suppose that σ is indecomposable, that is, $m = 1$. Then*

$$\mu(\sigma, \pi) = \begin{cases} \mu(\sigma, \pi_1) & \text{if } \pi = k \times \pi_1, \\ -\mu(\sigma, \pi_1) & \text{if } \pi = k \times \pi_1 \oplus 1, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Since $m = 1$, Eq. (3) takes the form

$$\begin{aligned} \mu(\sigma, \pi) &= \sum_{j=1}^k \mu(\sigma_{\leq 1}, \pi_1) \mu(\sigma_{>1}, \pi_{>j}) \\ &= \sum_{j=1}^k \mu(\sigma, \pi_1) \mu(\emptyset, \pi_{>j}) \end{aligned}$$

$$= \mu(\sigma, \pi_1) \mu(\emptyset, \pi_{>k}),$$

where the last equality follows from the fact that $\mu(\emptyset, \pi_{>j})$ is equal to 0 whenever $\pi_{>j}$ has more than one component.

We have $\mu(\emptyset, \pi_{>k}) = 1$ when $\pi_{>k} = \emptyset$, $\mu(\emptyset, \pi_{>k}) = -1$ when $\pi_{>k} = 1$, and $\mu(\emptyset, \pi_{>k}) = 0$ otherwise. In particular, $\mu(\emptyset, \pi_{>k})$ can only be nonzero either when $k = n$ and $\pi = k \times \pi_1$, or when $k = n - 1$ and $\pi = k \times \pi_1 \oplus 1$. \square

Corollary 3 implies that if σ is indecomposable and π is decomposable, then almost always $\mu(\sigma, \pi) = 0$, since the two exceptions for π given in the corollary are of a proportion that clearly goes to zero among decomposable permutations as their length goes to infinity.

Corollary 4. *With σ and π as in Proposition 2, assume that σ and π decompose into exactly two components, with $\sigma = \sigma_1 \oplus \sigma_2$ and $\pi = \pi_1 \oplus \pi_2$, and that $\pi_1, \pi_2 > 1$. Then*

$$\mu(\sigma, \pi) = \begin{cases} \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2) & \text{if } \pi_1 \neq \pi_2, \\ \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2) + \mu(\sigma, \pi_1) & \text{if } \pi_1 = \pi_2. \end{cases}$$

Proof. If $\pi_1 \neq \pi_2$ (so $k = 1$), then the summation in Eq. (3) expands into

$$\mu(\sigma, \pi) = \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2) + \mu(\sigma, \pi_1) \mu(\emptyset, \pi_2).$$

Since $\pi_2 > 1$, the second summand vanishes and $\mu(\sigma, \pi) = \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2)$.

If, on the other hand, $\pi_1 = \pi_2$, then Eq. (3) states that $\mu(\sigma, \pi)$ is equal to

$$\begin{aligned} & \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2) + \mu(\sigma_1, \pi_1) \mu(\sigma_2, \emptyset) + \mu(\sigma, \pi_1) \mu(\emptyset, \pi_2) + \mu(\sigma, \pi_1) \mu(\emptyset, \emptyset) \\ &= \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2) + \mu(\sigma, \pi_1). \quad \square \end{aligned}$$

Remark 5. An obvious question to ask is whether the product formula $\mu(\sigma, \pi) = \mu(\sigma_1, \pi_1) \mu(\sigma_2, \pi_2)$, in the case when $\pi_1 \neq \pi_2$, is a result of the interval $[\sigma, \pi]$ being isomorphic to the direct product of the intervals $[\sigma_1, \pi_1]$ and $[\sigma_2, \pi_2]$. Although this seems to occur frequently, it does not hold in general. As an example, consider the intervals $I_1 = [1, 21]$ and $I_2 = [1, 231]$. The product $I_1 \times I_2$ has eight elements, whereas the interval $[1 \oplus 1, 21 \oplus 231]$ has nine. More generally, for two indecomposable permutations σ_1 and σ_2 , define $\pi_1 = \sigma_1 \oplus 1$ and $\pi_2 = (\sigma_1 \oplus \sigma_2) \oplus 1$. Then the product of the intervals $I_1 = [\sigma_1, \pi_1]$ and $I_2 = [\sigma_2, \pi_2]$ has fewer elements than the interval $I = [\sigma_1 \oplus \sigma_2, \pi_1 \oplus \pi_2]$. To see this, note that an element $(\alpha, \beta) \in I_1 \times I_2$ can be mapped to an element $\alpha \oplus \beta \in I$, and this is easily seen to be an injective mapping; however, I also contains the permutation $(\sigma_1 \oplus \sigma_2) \oplus 1$ which is indecomposable and therefore not in the image of this mapping.

The following corollary is an immediate consequence of Proposition 1 (the case when $k - 1 = \ell = 0$).

Corollary 6. *Suppose σ and π are permutations of length at least two, such that neither begins with 1. Then $\mu(\sigma, 1 \oplus \pi) = -\mu(\sigma, \pi)$.*

Both recurrences (Propositions 1 and 2) are proved using arguments that involve cancellation between certain types of chains in the poset of permutations. Let us first introduce some useful notation. For a chain $C = \{\alpha_0 < \alpha_1 < \dots < \alpha_k\}$ of permutations let $L(C)$ denote the length of C , which is one less than the number of elements of C . The *weight* of C , denoted by $w(C)$, is the quantity $(-1)^{L(C)}$. If \mathcal{C} is any set of chains, then the *weight* of \mathcal{C} is defined by

$$w(\mathcal{C}) = \sum_{C \in \mathcal{C}} w(C) = \sum_{C \in \mathcal{C}} (-1)^{L(C)}.$$

Recall that $\mathcal{C}(\sigma, \pi)$ is the set of all the chains from σ to π that contain both σ and π . We know that $\mu(\sigma, \pi) = w(\mathcal{C}(\sigma, \pi))$, by Philip Hall's Theorem.

For a chain $C = \{\alpha_0 < \alpha_1 < \dots < \alpha_k\}$ and a permutation β , we let $\beta \oplus C$ denote the chain $\{\beta \oplus \alpha_0 < \beta \oplus \alpha_1 < \dots < \beta \oplus \alpha_k\}$. The chain $C \oplus \beta$ is defined analogously.

3.1. Proof of the first recurrence

Let us now turn to the proof of Proposition 1. Suppose that σ , π , m , n , k , and ℓ are as in the statement of the proposition. For a permutation τ , define the *degree* of τ , denoted by $\deg(\tau)$, to be the largest integer d such that τ can be expressed as $d \times 1 \oplus \tau'$ for some (possibly empty) permutation τ' . In particular, we have $k = \deg(\pi)$ and $\ell = \deg(\sigma)$.

Let $C = \{\tau_0 < \tau_1 < \dots < \tau_p\}$ be a chain of permutations. We say that a permutation $\tau_i \in C$ is the *pivot* of the chain C , if $\deg(\tau_i) < \deg(\tau_j)$ for each $j > i$, and $\deg(\tau_i) \leq \deg(\tau_j)$ for each $j \leq i$. In other words, the pivot is the element of the chain with minimum degree, and if there are more elements of minimum degree, the pivot is the largest of them.

Let ρ denote the permutation $\pi_{>1}$. Obviously $\deg(\rho) = k - 1$ and $1 \oplus \rho = \pi$. We partition the set of chains $\mathcal{C}(\sigma, \pi)$ into three disjoint subsets, denoted by \mathcal{C}_a , \mathcal{C}_b and \mathcal{C}_c , and we compute the weight of each subset separately. A chain $C \in \mathcal{C}(\sigma, \pi)$ belongs to \mathcal{C}_a if its pivot is the permutation π , the chain C belongs to \mathcal{C}_b if its pivot is the permutation ρ , and C belongs to \mathcal{C}_c otherwise. We now separate the main steps of the proof into independent claims.

Claim 7. If $\deg(\sigma) < \deg(\pi)$ (so $\ell < k$), then \mathcal{C}_a is empty. Otherwise, $w(\mathcal{C}_a) = \mu(\sigma_{>k}, \pi_{>k})$.

Proof. Obviously, if $\deg(\sigma) < \deg(\pi)$, then no chain from σ to π can have π for pivot, because the pivot must have minimal degree among the elements of the chain. Thus, \mathcal{C}_a is empty.

Assume now that $\deg(\sigma) \geq \deg(\pi)$. We show that there is a length-preserving bijection between the set of chains $\mathcal{C}(\sigma_{>k}, \pi_{>k})$ and the set of chains \mathcal{C}_a . Indeed, take any chain $C \in \mathcal{C}(\sigma_{>k}, \pi_{>k})$, and create a new chain $f(C) = (k \times 1) \oplus C$. Then $f(C)$ is a chain from σ to π , and since every element of $f(C)$ has degree at least k , while π has degree exactly k , we see that π is the pivot of $f(C)$. Hence $f(C) \in \mathcal{C}_a$.

On the other hand, if C' is any chain from \mathcal{C}_a , we see that each element of C' has degree at least k , because π has degree k and is the pivot of C' . Thus, every element $\tau' \in C'$ is of the form $k \times 1 \oplus \tau$ for some τ , and hence there exists a chain $C \in \mathcal{C}(\sigma_{>k}, \pi_{>k})$ such that $C' = f(C)$. Since f is clearly injective and length-preserving, we conclude that $w(\mathcal{C}_a) = w(\mathcal{C}(\sigma_{>k}, \pi_{>k})) = \mu(\sigma_{>k}, \pi_{>k})$, as claimed. \square

Claim 8. If $\deg(\sigma) < \deg(\rho)$ (so $\ell < k - 1$), then \mathcal{C}_b is empty. Otherwise, $w(\mathcal{C}_b) = -\mu(\sigma_{>k-1}, \pi_{>k})$.

Proof. If $\deg(\sigma) < \deg(\rho)$ then ρ cannot be the pivot of any chain containing σ and \mathcal{C}_b is empty.

Assume now that $\deg(\sigma) \geq \deg(\rho)$. We will describe a parity-reversing bijection f between the set of chains $\mathcal{C}(\sigma_{>k-1}, \pi_{>k})$ and the set of chains \mathcal{C}_b . Take a chain $C \in \mathcal{C}(\sigma_{>k-1}, \pi_{>k})$. Define a new chain C' by $C' = ((k - 1) \times 1) \oplus C$. Notice that C' is a chain from σ to ρ whose pivot is ρ and whose length is equal to the length of C . Define the chain $f(C)$ by $f(C) = C' \cup \{\pi\}$. Then the chain $f(C)$ belongs to \mathcal{C}_b and has length $L(C) + 1$. It is again easy to see that f is a bijection between $\mathcal{C}(\sigma_{>k-1}, \pi_{>k})$ and \mathcal{C}_b , which shows that

$$w(\mathcal{C}_b) = -w(\mathcal{C}(\sigma_{>k-1}, \pi_{>k})) = -\mu(\sigma_{>k-1}, \pi_{>k}),$$

as claimed. \square

Claim 9. We have $w(\mathcal{C}_c) = 0$.

Proof. We construct a parity-reversing involution $f: \mathcal{C}_c \rightarrow \mathcal{C}_c$. Let C be a chain from \mathcal{C}_c , let τ be its pivot, and let τ' be the successor of τ in C . By definition of \mathcal{C}_c , τ is not equal to π , so τ' is well defined. From the definition of a pivot, we know that $\deg(\tau) < \deg(\tau')$. Let us distinguish two cases:

- (1) If $\tau' = 1 \oplus \tau$, we define a new chain $f(C)$ by $f(C) = C \setminus \{\tau'\}$. Note that in this case, we know that τ' is different from π , because otherwise τ would be equal to ρ , contradicting the definition of \mathcal{C}_c . Thus, $f(C) \in \mathcal{C}_c$. Note that τ is a pivot of $f(C)$.
- (2) If $\tau' \neq 1 \oplus \tau$, then we easily deduce that $\tau' > 1 \oplus \tau$ (recall that $\deg(\tau') > \deg(\tau)$). We then define a new chain $f(C) = C \cup \{1 \oplus \tau\}$, in which the new element $1 \oplus \tau$ is inserted between τ and τ' .

The mapping f is easily seen to be an involution on the set \mathcal{C}_c which preserves the pivot and maps odd-length chains to even-length chains and vice versa. This shows that $w(\mathcal{C}_c) = 0$, as claimed. \square

From these claims, Proposition 1 easily follows. Indeed, Claim 9 implies that $\mu(\sigma, \pi) = w(\mathcal{C}_a) + w(\mathcal{C}_b)$. From Claims 7 and 8 we deduce the values of $\mu(\sigma, \pi)$:

- If $k - 1 > \ell$ then both \mathcal{C}_a and \mathcal{C}_b are empty and $\mu(\sigma, \pi) = 0$.
- If $k - 1 = \ell$ then \mathcal{C}_a is empty and $\mu(\sigma, \pi) = w(\mathcal{C}_b) = -\mu(\sigma_{>k-1}, \pi_{>k})$.
- If $k - 1 < \ell$, then $\mu(\sigma, \pi) = w(\mathcal{C}_a) + w(\mathcal{C}_b) = \mu(\sigma_{>k}, \pi_{>k}) - \mu(\sigma_{>k-1}, \pi_{>k})$.

This completes the proof of Proposition 1.

3.2. Proof of the second recurrence

It remains to prove Proposition 2. The proof is again based on cancellation among the chains from σ to π . Before stating the proof, we need more terminology and several lemmas.

Let α , β and ρ be any permutations. We say that α is a ρ -tight subpermutation of β , denoted by $\alpha \stackrel{\rho}{<} \beta$, if $\alpha < \beta$ but $\rho \oplus \alpha$ is not contained in β . We say that a chain $\{\alpha_0 < \alpha_1 < \dots < \alpha_k\}$ is ρ -tight if $\alpha_{i-1} \stackrel{\rho}{<} \alpha_i$ for every $i = 1, \dots, k$. Let $\mathcal{C}^\rho(\alpha, \beta)$ be the set of all the ρ -tight chains from α to β .

The following simple properties of ρ -tightness follow directly from the definitions.

Lemma 10. For arbitrary permutations α , β , γ and ρ , we have $\alpha \oplus \gamma \stackrel{\rho}{<} \beta \oplus \gamma$ if and only if $\alpha \stackrel{\rho}{<} \beta$.

Proof. This is an immediate consequence of the fact that α is contained in β if and only if $\alpha \oplus \gamma$ is contained in $\beta \oplus \gamma$, and that $\rho \oplus \alpha$ is contained in β if and only if $\rho \oplus \alpha \oplus \gamma$ is contained in $\beta \oplus \gamma$. \square

Lemma 11. If ρ is a nonempty indecomposable permutation, and if α and β are arbitrary permutations, then $\rho \oplus \alpha \stackrel{1}{<} \rho \oplus \beta$ if and only if $\alpha \stackrel{\rho}{<} \beta$.

Proof. Suppose first that we have $\rho \oplus \alpha \stackrel{1}{<} \rho \oplus \beta$. This implies $\alpha < \beta$. If we also had $\rho \oplus \alpha \leq \rho \oplus \beta$, this would mean that $1 \oplus \rho \oplus \alpha \leq \rho \oplus \rho \oplus \alpha \leq \rho \oplus \beta$, contradicting $\rho \oplus \alpha \stackrel{1}{<} \rho \oplus \beta$. Therefore $\rho \oplus \alpha$ is not contained in β , and hence $\alpha \stackrel{\rho}{<} \beta$.

Conversely, assume that $\alpha \stackrel{\rho}{<} \beta$. This means that $\rho \oplus \alpha < \rho \oplus \beta$. Since ρ is indecomposable, any embedding of $\rho \oplus \alpha$ into $\rho \oplus \beta$ must embed ρ into a single component of $\rho \oplus \beta$. Moreover, $\rho \oplus \alpha$ is not contained in β because of the assumption $\alpha \stackrel{\rho}{<} \beta$. Consequently, any embedding of $\rho \oplus \alpha$ into $\rho \oplus \beta$ must embed the leftmost component of $\rho \oplus \alpha$ onto the leftmost component of $\rho \oplus \beta$ (both these components are equal to ρ). Consequently, there can be no embedding of $1 \oplus \rho \oplus \alpha$ into $\rho \oplus \beta$, showing that $\rho \oplus \alpha \stackrel{1}{<} \rho \oplus \beta$, as claimed. \square

The next lemma shows the relevance of ρ -tightness for the computation of μ .

Lemma 12. Let β be a permutation with decomposition $\beta = \beta_1 \oplus \beta_2 \oplus \dots \oplus \beta_p$. Let ρ be a nonempty indecomposable permutation, and let α be any permutation.

- (1) If $\rho \neq \beta_1$, then $\mu(\alpha, \beta) = w(\mathcal{C}^\rho(\alpha, \beta))$.
 (2) If $\rho = \beta_1$, then $\mu(\alpha, \beta) = w(\mathcal{C}^\rho(\alpha, \beta)) - w(\mathcal{C}^\rho(\alpha, \beta_{>1}))$.

Proof. Let us first deal with the first claim of the lemma. Let us define $\widehat{\mathcal{C}} = \mathcal{C}(\alpha, \beta) \setminus \mathcal{C}^\rho(\alpha, \beta)$ to be the set of all the chains from α to β that are not ρ -tight. The first part of the lemma is equivalent to saying that $w(\widehat{\mathcal{C}}) = 0$. To prove this, we find a parity-reversing involution f on the set $\widehat{\mathcal{C}}$.

Consider a chain $C = \{\alpha = \alpha_0 < \alpha_1 < \dots < \alpha_q = \beta\} \in \widehat{\mathcal{C}}$. Since C is not ρ -tight, there is an index i such that $\rho \oplus \alpha_i \leq \alpha_{i+1}$. Fix the smallest such value of i . We distinguish two cases: either $\rho \oplus \alpha_i < \alpha_{i+1}$, or $\rho \oplus \alpha_i = \alpha_{i+1}$.

If $\rho \oplus \alpha_i < \alpha_{i+1}$, define a new chain

$$f(C) = C \cup \{\rho \oplus \alpha_i\} = \{\alpha = \alpha_0 < \alpha_1 < \dots < \alpha_i < \rho \oplus \alpha_i < \alpha_{i+1} < \dots < \alpha_q = \beta\}.$$

On the other hand, if $\rho \oplus \alpha_i = \alpha_{i+1}$, define a new chain

$$f(C) = C \setminus \{\rho \oplus \alpha_i\} = \{\alpha = \alpha_0 < \alpha_1 < \dots < \alpha_i < \alpha_{i+2} < \dots < \alpha_q = \beta\}.$$

Note that, since we assume that $\rho \neq \beta_1$ and that ρ is indecomposable, we know that $\rho \oplus \alpha_i$ is not equal to β . Moreover, in the chain $f(C)$ the element α_i is not a ρ -tight subpermutation of its successor in the chain. Thus, we see that $f(C)$ is a chain from $\widehat{\mathcal{C}}$. It is easy to see that f is an involution, and that it reverses the parity of the length of the chain, showing that $w(\widehat{\mathcal{C}}) = 0$. This proves the first part of the lemma.

Let us prove the second part. Assume that $\rho = \beta_1$, that is, $\beta = \rho \oplus \beta_{>1}$. Consider a chain C from α to β , and let $\alpha_0, \alpha_1, \dots, \alpha_q$ be the elements of C . We say that the chain C is *almost ρ -tight* if its second largest element α_{q-1} is equal to $\beta_{>1}$ and if $\alpha_{i-1} \stackrel{\rho}{<} \alpha_i$ for each $i \leq q-1$. Note that an almost ρ -tight chain is never ρ -tight, because $\beta_{>1}$ is not a ρ -tight subpermutation of $\beta = \rho \oplus \beta_{>1}$.

We partition the set $\mathcal{C}(\alpha, \beta)$ into three disjoint sets \mathcal{C}_a , \mathcal{C}_b , and \mathcal{C}_c , where \mathcal{C}_a is the set $\mathcal{C}^\rho(\alpha, \beta)$ of ρ -tight chains, \mathcal{C}_b is the set of almost ρ -tight chains, and \mathcal{C}_c contains the chains that are neither ρ -tight nor almost ρ -tight.

Consider again the mapping f defined in the proof of the first part of the lemma. This mapping, restricted to the set \mathcal{C}_c , is easily seen to be a parity-reversing involution on \mathcal{C}_c , which shows that $w(\mathcal{C}_c) = 0$. This means that $\mu(\alpha, \beta) = w(\mathcal{C}_a) + w(\mathcal{C}_b)$.

Furthermore, note that an almost ρ -tight chain from α to β consists of a ρ -tight chain from α to $\beta_{>1}$ followed by β , and conversely, any ρ -tight chain from α to $\beta_{>1}$ can be extended to an almost ρ -tight chain from α to β by adding the element β . Thus, we see that $w(\mathcal{C}_b) = -w(\mathcal{C}^\rho(\alpha, \beta_{>1}))$. This implies that $\mu(\alpha, \beta) = w(\mathcal{C}^\rho(\alpha, \beta)) - w(\mathcal{C}^\rho(\alpha, \beta_{>1}))$, as claimed. \square

The next lemma is an easy consequence of Lemma 12.

Lemma 13. Let β be a permutation with decomposition $\beta = \beta_1 \oplus \beta_2 \oplus \dots \oplus \beta_p$. Let ρ be an indecomposable permutation, and let α be any permutation. Let $q \geq 0$ be the largest integer such that the components $\beta_1, \beta_2, \dots, \beta_q$ are all equal to ρ . Then

$$w(\mathcal{C}^\rho(\alpha, \beta)) = \sum_{i=0}^q \mu(\alpha, \beta_{>i}).$$

Proof. Proceed by induction on q . If $q = 0$, the claim reduces to the identity $w(\mathcal{C}^\rho(\alpha, \beta)) = \mu(\alpha, \beta)$, which follows from the first part of Lemma 12. Suppose that $q > 0$. Then the second part of Lemma 12 applies and we get that

$$\mu(\alpha, \beta) = w(\mathcal{C}^\rho(\alpha, \beta)) - w(\mathcal{C}^\rho(\alpha, \beta_{>1})),$$

which is equivalent to

$$w(\mathfrak{C}^p(\alpha, \beta)) = \mu(\alpha, \beta) + w(\mathfrak{C}^p(\alpha, \beta_{>1})). \quad (4)$$

By induction, we know that

$$w(\mathfrak{C}^p(\alpha, \beta_{>1})) = \sum_{i=0}^{q-1} \mu(\alpha, (\beta_{>1})_{>i}) = \sum_{i=1}^q \mu(\alpha, \beta_{>i}).$$

Combining this with (4), we obtain the desired identity. \square

Before we proceed towards the proof of Proposition 2, we need to introduce more definitions. Let β be a permutation with decomposition $\beta_1 \oplus \cdots \oplus \beta_p$ into indecomposable components, let α be any permutation. Let C be a chain of permutations, with elements $\alpha = \alpha_0 < \alpha_1 < \cdots < \alpha_q = \beta$. We express each element α_i of the chain as a sum of two permutations, called *head* and *tail*, denoted respectively as $h_i(C)$ and $t_i(C)$, with $\alpha_i = h_i(C) \oplus t_i(C)$. The head and tail are defined inductively as follows: for $i = q$, we have $\alpha_i = \alpha_q = \beta$ and we define $h_q(C) = \beta_1$ and $t_q(C) = \beta_{>1}$.

Suppose now that the head and tail of α_i have been already defined, and let us define head and tail of α_{i-1} . Let us put $\gamma = \alpha_{i-1}$, and assume that γ has decomposition $\gamma_1 \oplus \gamma_2 \oplus \cdots \oplus \gamma_r$ into indecomposable components. Let j be the smallest integer such that $\gamma_{>j} \leq t_i(C)$. It then follows that $\gamma_{\leq j} \leq h_i(C)$. We define $h_{i-1}(C) = \gamma_{\leq j}$ and $t_{i-1}(C) = \gamma_{>j}$. In other words, the tail of α_{i-1} is its longest suffix that is contained in the tail of α_i .

If the chain C is clear from the context, we write h_i and t_i instead of $h_i(C)$ and $t_i(C)$. Note that $h_0 \leq h_1 \leq \cdots \leq h_q$ and $t_0 \leq t_1 \leq \cdots \leq t_q$.

We say that the chain C of length q is *split* if there is an index $s \in \{0, \dots, q\}$ such that $t_0 = t_1 = \cdots = t_s$ and $h_s = h_{s+1} = \cdots = h_q$. Such an index s is then necessarily unique. The next lemma demonstrates the relevance of these notions.

Lemma 14. *Let β be a permutation with decomposition $\beta_1 \oplus \beta_2 \oplus \cdots \oplus \beta_p$ such that $\beta_1 \neq 1$. Let α be an arbitrary permutation. Let \mathfrak{C}^* be the set of all the chains from $\mathfrak{C}(\alpha, \beta)$ which are split and 1-tight. Then $\mu(\alpha, \beta) = w(\mathfrak{C}^*)$.*

Proof. By the first part of Lemma 12, we know that $\mu(\alpha, \beta)$ is equal to $w(\mathfrak{C}^1(\alpha, \beta))$, that is, to the total weight of all the 1-tight chains from α to β . Define the set $\widehat{\mathfrak{C}} = \mathfrak{C}^1(\alpha, \beta) \setminus \mathfrak{C}^*$ of all the 1-tight, non-split chains from α to β .

To prove the lemma, we need to show that $w(\widehat{\mathfrak{C}}) = 0$. To achieve this, we again use a parity-reversing involution f on the set $\widehat{\mathfrak{C}}$. Consider a chain $C \in \widehat{\mathfrak{C}}$ with elements $\alpha_0 < \alpha_1 < \cdots < \alpha_q$. Clearly, for any two consecutive elements $\alpha_{j-1} < \alpha_j$ of C , one of these three cases occurs: either $h_{j-1} < h_j$ and $t_{j-1} < t_j$, or $h_{j-1} < h_j$ and $t_{j-1} = t_j$, or $h_{j-1} = h_j$ and $t_{j-1} < t_j$. A chain is split if and only if the first case never occurs, and all occurrences of the second case appear lower in the chain than any occurrence of the third case. Since C is not split, there must exist an index $j \in \{1, \dots, q\}$ such that either

- (1) $h_{j-1} < h_j$ and $t_{j-1} < t_j$, or
- (2) $h_{j-1} = h_j < h_{j+1}$ and $t_{j-1} < t_j = t_{j+1}$.

Fix such an index j as large as possible and distinguish two cases depending on which of the two above-mentioned possibilities occur for this index j .

Case (1). Assume that $h_{j-1} < h_j$ and $t_{j-1} < t_j$. Let us write $h = h_{j-1}$, $H = h_j$, $t = t_{j-1}$, and $T = t_j$, so we have $\alpha_{j-1} = h \oplus t$ and $\alpha_j = H \oplus T$. Define a permutation $\gamma = h \oplus T$, and a new chain $f(C) = C \cup \{\gamma\}$. Note that since C is a 1-tight chain, and in particular $\alpha_{j-1} \overset{1}{<} \alpha_j$, we also have $\alpha_{j-1} \overset{1}{<} \gamma \overset{1}{<} \alpha_j$, and hence $f(C)$ is a 1-tight chain as well.

We need to prove that $f(C) \in \widehat{\mathfrak{C}}$, which follows easily from the following claim.

Claim 15. *Each permutation of C has the same head and tail in $f(C)$ as in C . The permutation $\gamma = h \oplus T$ has head h and tail T in $f(C)$.*

Proof. It is clear that the claim holds for all the permutations that are greater than γ .

It is also easy to see that the claim holds for γ . Indeed, the successor of γ in $f(C)$ is the permutation $H \oplus T$, whose tail is T . Since the tail of γ cannot be greater than T and since $\gamma = h \oplus T$, it follows that the tail of γ is T and its head is h .

Let us now consider the permutation $\alpha_{j-1} = h \oplus t$. The successor of α_{j-1} in C is the permutation $\alpha_j = H \oplus T$, and the successor of α_{j-1} in $f(C)$ is the permutation $\gamma = h \oplus T$. Since the two successors have the same tail T , and since the tail of a permutation only depends on the tail of its successor, we see that α_{j-1} has the same tail (and hence also the same head) in $f(C)$ as in C .

From these facts, the claim immediately follows. \square

We may now conclude that $f(C) \in \widehat{\mathcal{C}}$, and turn to the second case of the proof of the lemma.

Case (2). Assume now that $h_{j-1} = h_j < h_{j+1}$ and $t_{j-1} < t_j = t_{j+1}$. Let us define $h = h_{j-1} = h_j$, $H = h_{j+1}$, $t = t_{j-1}$, and $T = t_j = t_{j+1}$. In particular, $\alpha_{j-1} = h \oplus t$, $\alpha_j = h \oplus T$, and $\alpha_{j+1} = H \oplus T$. Define the chain $f(C) = C \setminus \{\alpha_j\}$.

We claim that $f(C)$ is 1-tight. To see this, it is enough to prove $h \oplus t \stackrel{1}{<} H \oplus T$. Assume, for a contradiction, that $1 \oplus h \oplus t \leq H \oplus T$. In any occurrence of $1 \oplus h \oplus t$ inside $H \oplus T$, the prefix $1 \oplus h$ must occur inside H , otherwise we get a contradiction with the assumption that t is the tail of α_{j-1} . This shows that $1 \oplus h \leq H$, and hence $1 \oplus h \oplus T = 1 \oplus \alpha_j \leq \alpha_{j+1} = H \oplus T$, contradicting the assumption that C is 1-tight. To finish the proof of the lemma, we need one more claim.

Claim 16. Each permutation of $f(C)$ has the same head and tail in $f(C)$ as in C .

Proof. It is enough to prove the claim for the permutation $\alpha_{j-1} = h \oplus t$, because any other permutation of $f(C)$ has the same successor in $f(C)$ as in C . For α_{j-1} , the claim follows from the fact that the successor of α_{j-1} in C has the same tail as the successor of α_{j-1} in $f(C)$. This completes the proof of the claim. \square

We now see that even in this second case, $f(C)$ belongs to $\widehat{\mathcal{C}}$.

Combining the two cases described above, we see that f is a parity-reversing involution of the set $\widehat{\mathcal{C}}$. This means that $w(\widehat{\mathcal{C}}) = 0$, and consequently, $\mu(\alpha, \beta) = w(\mathcal{C}^*)$, as claimed. This completes the proof of the lemma. \square

Finally, we can prove Proposition 2. Assume that σ is a permutation with decomposition $\sigma_1 \oplus \dots \oplus \sigma_m$ and that π is a permutation with decomposition $\pi_1 \oplus \dots \oplus \pi_n$, where $n \geq 2$ and $\pi_1 > 1$. Let $k \geq 1$ be the largest integer such that all the components π_1, \dots, π_k are equal to π_1 . Recall that our goal is to prove identity (3), which reads as follows:

$$\mu(\sigma, \pi) = \sum_{i=1}^m \sum_{j=1}^k \mu(\sigma_{\leq i}, \pi_1) \mu(\sigma_{> i}, \pi_{> j}).$$

Let \mathcal{C}^* be the set of 1-tight split chains from σ to π . From Lemma 14, we know that $\mu(\sigma, \pi) = w(\mathcal{C}^*)$. For a chain $C \in \mathcal{C}^*$, let $t_0(C)$ be the tail of the element $\sigma \in C$, which is the smallest element in the chain. By definition, $t_0(C)$ is a suffix of σ , that is, it is equal to $\sigma_{> i}$ for some value of $i \in \{0, \dots, m\}$. Define, for each $i \in \{0, \dots, m\}$, the set of chains

$$\mathcal{C}_i = \{C \in \mathcal{C}^*, t_0(C) = \sigma_{> i}\}.$$

The sets \mathcal{C}_i form a disjoint partition of \mathcal{C}^* . We will now compute the weight of the individual sets \mathcal{C}_i .

Claim 17. Let C be a chain from \mathcal{C}^* . Every element of C has nonempty head. Consequently, $t_0(C)$ is never equal to σ , and hence \mathcal{C}_0 is empty.

Proof. Suppose that C has an element with empty head. Let α be the largest such element. By definition, the element $\pi \in C$ has head equal to π_1 , so $\alpha \neq \pi$. In particular, α has a successor α' in C , and α' has nonempty head. Let h' and t' be the head and tail of α' . By assumption, h' is nonempty, which means that $1 \leq h'$. Moreover, $\alpha \leq t'$, because α is its own tail. This means that $1 \oplus \alpha \leq \alpha'$, which is impossible because the chains in \mathcal{C}^* are assumed to be 1-tight.

This shows that every element of C has nonempty head, and the rest of the claim follows directly. \square

Claim 17 implies that $w(\mathcal{C}_0) = 0$, and hence $\mu(\sigma, \pi) = \sum_{i \geq 1} w(\mathcal{C}_i)$. It remains to determine the value of $w(\mathcal{C}_i)$ for $i > 0$.

Fix an integer $i \in \{1, \dots, m\}$. Define $h = \sigma_{\leq i}$, $t = \sigma_{> i}$, $H = \pi_1$, and $T = \pi_{> 1}$. Note that in a chain $C \in \mathcal{C}_i$, the permutation σ has head h and tail t , while the permutation π has head H and tail T .

Claim 18. *With the notation as above,*

$$w(\mathcal{C}_i) = w(\mathcal{C}^1(h, H))w(\mathcal{C}^H(t, T)).$$

Proof. Let us write $\mathcal{C}' = \mathcal{C}^1(h, H)$ and $\mathcal{C}'' = \mathcal{C}^H(t, T)$. We will provide a bijection $f: \mathcal{C}' \times \mathcal{C}'' \rightarrow \mathcal{C}_i$, which maps a pair of chains $(C_1, C_2) \in \mathcal{C}' \times \mathcal{C}''$ to a chain $f(C_1, C_2) \in \mathcal{C}_i$ whose length is equal to $L(C_1) + L(C_2)$. Such a bijection immediately implies the identity $w(\mathcal{C}_i) = w(\mathcal{C}')w(\mathcal{C}'')$ from the claim.

The definition of the mapping f is simple: for $C_1 \in \mathcal{C}'$ and $C_2 \in \mathcal{C}''$, define $f(C_1, C_2)$ to be the concatenation of the two chains $C_1 \oplus t$ and $H \oplus C_2$. This is well defined, since the maximum of $C_1 \oplus t$ is the permutation $H \oplus t$, which is also equal to the minimum of the chain $H \oplus C_2$. Thus, $f(C_1, C_2)$ is a chain of length $L(C_1) + L(C_2)$. Let us denote this chain by C .

We now show that C belongs to \mathcal{C}_i . Let us call the two sub-chains $C_1 \oplus t$ and $H \oplus C_2$ respectively the *bottom part* and the *top part* of C . Note that the permutation $H \oplus t$ is the unique element of C belonging both to the top part and to the bottom part.

By construction, C is a chain from σ to π . The bottom part of C is a 1-tight chain, because C_1 was assumed to be 1-tight (see Lemma 10). Similarly, by Lemma 11, the top part of C is a 1-tight chain, because C_2 is H -tight and H is indecomposable. This shows that the chain C is 1-tight.

Our next step is to prove that every element in the top part of C has head equal to H , and that every element in the bottom part of C has tail equal to t . Assume that this statement is false, and let α be the largest element of C for which it fails. Clearly, $\alpha \neq \pi$, so α has a successor β in C . Suppose first that α belongs to the top part of C . Then α can be written as a sum $H \oplus \alpha'$ for some $\alpha' \in C_2$, and likewise $\beta = H \oplus \beta'$ for $\beta' \in C_2$. By the choice of α , we know that the head of β is H and hence its tail is β' . Since $\alpha' < \beta'$, the tail of α contains α' . On the other hand, the only suffix of α longer than α' is the permutation α itself, because H is indecomposable. By Claim 17, the head of α must be nonempty, which means that the head of α can only be equal to H , which contradicts our choice of α .

Suppose now that α does not belong to the top part of C . Then β belongs to the bottom part of C (and possibly to the top part as well). Consequently, α can be written as $\alpha' \oplus t$ and β can be written as $\beta' \oplus t$, with $\alpha', \beta' \in C_1$. We also know that t is the tail of β . This makes it clear that t is the tail of α as well, which is a contradiction.

This proves that all the elements of the top part of C indeed have head H , and all the elements in the bottom part have tail t . This shows that C is a split chain and also that $t_0(C) = t$. We have shown that $C \in \mathcal{C}_i$.

It is clear that f is an injective mapping. To complete the proof of the claim, it only remains to show that f is surjective, that is, for every $C \in \mathcal{C}_i$ there are chains $(C_1, C_2) \in \mathcal{C}' \times \mathcal{C}''$ with $f(C_1, C_2) = C$.

Choose a chain $C \in \mathcal{C}_i$. Since C is split, it must contain the element $H \oplus t$. Call the elements of C contained in $H \oplus t$ the *bottom part* of C , and the elements containing $H \oplus t$ the *top part* of C . The definition of split chain further implies that all the elements in the top part have the same head H and all the elements in the bottom part have the same tail t . Hence, the bottom part of the chain C

has the form $C_1 \oplus t$ for some chain $C_1 \in \mathfrak{C}(h, H)$. Similarly, the top part has the form $H \oplus C_2$ for a chain $C_2 \in \mathfrak{C}(t, T)$. Since C is 1-tight, we may use Lemmas 10 and 11 to see that C_1 is 1-tight and C_2 is H -tight, showing that $(C_1, C_2) \in \mathfrak{C}' \times \mathfrak{C}''$. Since $f(C_1, C_2) = C$, we see that f is the required bijection. \square

We now have all the necessary ingredients to finish the proof of Proposition 2. Let us write $H = \pi_1$ and $T = \pi_{>1}$. From our results, we get

$$\begin{aligned} \mu(\sigma, \pi) &= w(\mathfrak{C}(\sigma, \pi)) \\ &= w(\mathfrak{C}^*) \quad \text{by Lemma 14} \\ &= \sum_{i=1}^m w(\mathfrak{C}_i) \quad \text{by Claim 17} \\ &= \sum_{i=1}^m w(\mathfrak{C}^1(\sigma_{\leq i}, H)) w(\mathfrak{C}^H(\sigma_{> i}, T)) \quad \text{by Claim 18} \\ &= \sum_{i=1}^m \mu(\sigma_{\leq i}, H) w(\mathfrak{C}^H(\sigma_{> i}, T)) \quad \text{by first part of Lemma 12} \\ &= \sum_{i=1}^m \mu(\sigma_{\leq i}, H) \sum_{j=0}^{k-1} \mu(\sigma_{> i}, T_{> j}) \quad \text{by Lemma 13} \\ &= \sum_{i=1}^m \sum_{j=1}^k \mu(\sigma_{\leq i}, \pi_1) \mu(\sigma_{> i}, \pi_{> j}) \quad \text{since } T_{> j} = \pi_{> j+1}. \end{aligned}$$

Thus, Proposition 2 is now proved.

We now present some consequences of Propositions 1 and 2.

Corollary 19. *There is an algorithm that, given two separable permutations σ and π , computes the value of $\mu(\sigma, \pi)$ in time polynomial in $|\sigma| + |\pi|$.*

Proof. Let $\pi = \pi_1 \pi_2 \cdots \pi_n$ be a separable permutation. For two integers i, j with $1 \leq i \leq j \leq n$, let $\pi[i, j]$ denote the subpermutation of π order-isomorphic to the sequence $\pi_i, \pi_{i+1}, \dots, \pi_j$. Note that $\pi[i, j]$ is also separable. We call $\pi[i, j]$ a *range subpermutation* of π .

Suppose that $\sigma = \sigma_1 \cdots \sigma_m$ and $\pi = \pi_1 \cdots \pi_n$ are two separable permutations. Our goal is to compute $\mu(\sigma, \pi)$. We will use a straightforward dynamic programming algorithm to perform this computation. We will compute all the values of the form $\mu(\sigma[i, j], \pi[k, \ell])$, for all quadruples (i, j, k, ℓ) satisfying $1 \leq i \leq j \leq m$ and $1 \leq k \leq \ell \leq n$. For each such quadruple (i, j, k, ℓ) we store the value of $\mu(\sigma[i, j], \pi[k, \ell])$ once we compute it,¹ so that we do not need to compute this value more than once, even though we may need it several times to compute other values of μ .

There are $\mathcal{O}(m^2 n^2)$ quadruples (i, j, k, ℓ) to consider, and for each such quadruple, we may use Propositions 1 and 2 to express $\mu(\sigma[i, j], \pi[k, \ell])$ as a combination of polynomially many values of the form $\mu(\sigma[i', j'], \pi[k', \ell'])$ where $\sigma[i', j']$ and $\pi[k', \ell']$ are range subpermutations of $\sigma[i, j]$ and $\pi[k, \ell]$ with $\pi[k', \ell'] \neq \pi[k, \ell]$. Therefore, we can in polynomial time compute all the values of the form $\mu(\sigma[i, j], \pi[k, \ell])$, including $\mu(\sigma, \pi) = \mu(\sigma[1, m], \pi[1, n])$. \square

Note that the number of permutations belonging to an interval $[\sigma, \pi]$ may in general be exponential in the size of π , even when π and σ are separable. Therefore, computing the Möbius function

¹ In fact, we do not need to compute $\mu(\sigma[i, j], \pi[k, \ell])$ for all quadruples (i, j, k, ℓ) . It suffices only to consider those quadruples for which the values appearing in $\sigma[i, j]$ as well as those in $\pi[k, \ell]$ form a contiguous interval of integers.

$\mu(\sigma, \pi)$ directly from Eq. (1) would be much less efficient than the algorithm of the previous corollary.

Let us say that a class of permutations \mathcal{C} is *sum-closed* if for each $\pi, \sigma \in \mathcal{C}$, the class \mathcal{C} also contains $\pi \oplus \sigma$. Similarly, \mathcal{C} is *skew-closed* if $\pi, \sigma \in \mathcal{C}$ implies $\pi \ominus \sigma \in \mathcal{C}$. For a set \mathcal{P} of permutations, the $\{\oplus, \ominus\}$ -closure of \mathcal{P} , denoted by $\text{cl}(\mathcal{P})$, is the smallest sum-closed and skew-closed class of permutations that contains \mathcal{P} . Notice that $\text{cl}(\{1\})$ is exactly the set of separable permutations.

The next two corollaries are immediate consequences of Propositions 1 and 2 (see also Corollary 3), and we omit their proof.

Corollary 20. *Suppose that σ is a permutation that is neither decomposable nor skew-decomposable. Let \mathcal{P} be any set of permutations. Then*

$$\max\{|\mu(\sigma, \pi)|; \pi \in \mathcal{P}\} = \max\{|\mu(\sigma, \pi)|; \pi \in \text{cl}(\mathcal{P})\}.$$

Moreover, the computation of $\mu(\sigma, \pi)$ for $\pi \in \text{cl}(\mathcal{P})$ can be efficiently reduced to the computation of the values $\mu(\sigma, \rho)$ for $\rho \in \mathcal{P}$.

Corollary 21. *Let \mathcal{P} be a class of permutations that contains only finitely many elements that are both indecomposable and skew-indecomposable. Then for any $\pi \in \mathcal{P}$ and any σ , we can compute $\mu(\sigma, \pi)$ in time polynomial in $|\sigma| + |\pi|$.*

4. The Möbius function of separable permutations

Let us now consider the values of $\mu(\sigma, \pi)$ for separable permutations σ and π . Our goal is to show that the values of the Möbius function in the poset of separable permutations have a combinatorial interpretation in terms of the so-called *normal embeddings*, which we define below. This alternative interpretation of the Möbius function generalizes previous results of Sagan and Vatter [5] for the Möbius function of intervals of layered permutations, which we explain at the end of this section.

As a consequence of this new interpretation of the Möbius function, we are able to relate the Möbius function $\mu(\sigma, \pi)$ to the number of occurrences of σ in π , by showing that $|\mu(\sigma, \pi)| \leq \sigma(\pi)$. We also show that $\mu(1, \pi)$ is equal to $-1, 0$ or 1 whenever π is separable.

The recursive structure of separable permutations makes it convenient to represent a separable permutation by a tree that describes how the permutation may be obtained from smaller permutations by sums and skew sums. We now formalize this concept. A *separating tree* T is a rooted tree T with the following properties:

- Each internal node of T has one of two types: it is either a *direct node* or a *skew node*.
- Each internal node has at least two children. The children of a given internal node are ordered into a sequence from left to right.

Each separating tree T represents a unique separable permutation π , defined recursively as follows:

- If T has a single node, it represents the singleton permutation 1 .
- Assume T has more than one node. Let N_1, \dots, N_k be the children of the root in their left-to-right order, and let T_i denote the subtree of T rooted at the node N_i . Let p_1, \dots, p_k be the permutations represented by the trees T_1, \dots, T_k . Then T represents the permutation $p_1 \oplus \dots \oplus p_k$ if the root of T is a direct node and $p_1 \ominus \dots \ominus p_k$ if the root of T is a skew node.

Note that the leaves of T correspond bijectively to the letters of π . In fact, when we perform a depth-first left-to-right traversal of T , we encounter the leaves in the order that corresponds to the left-to-right order of the letters of π . See Fig. 1 for an example.

A given separable permutation may be represented by more than one separating tree. A separating tree is called a *reduced tree* if it has the property that the children of a direct node are leaves or skew nodes, and the children of a skew node are leaves or direct nodes. Each separable permutation π is represented by a unique reduced tree, denoted by $T(\pi)$. We assume that each leaf of T is labeled by the corresponding letter of π .

This slightly modified concept of separating tree and its relationship with separable permutations have been previously studied in algorithmic contexts [2,9]. We will now show that the reduced tree allows us to obtain a simple formula for the Möbius function of separable permutations.

Let $[n]$ denote the set $\{1, \dots, n\}$. Let $\pi = \pi_1\pi_2 \cdots \pi_n$ and $\sigma = \sigma_1\sigma_2 \cdots \sigma_m$ be two permutations, with $\sigma \leq \pi$. An *embedding* of σ into π is a function $f: [m] \rightarrow [n]$ with the following two properties:

- for every $i, j \in [m]$, if $i < j$ then $f(i) < f(j)$ (so f is monotone increasing).
- for every $i, j \in [m]$, if $\sigma_i < \sigma_j$, then $\pi_{f(i)} < \pi_{f(j)}$ (so f is order-preserving).

Let f be an embedding of σ into π . We say that a leaf ℓ of $T(\pi)$ is *covered* by the embedding f if the letter of π corresponding to ℓ is in the image of f . A leaf is *omitted* by f if it is not covered by f . An *internal node* is a node that is not a leaf. An internal node N of $T(\pi)$ is *omitted* by f if all the leaves in the subtree rooted at N are omitted. A node is *maximal omitted*, if it is omitted but its parent in $T(\pi)$ is not omitted.

Assume that π is a separable permutation and $T = T(\pi)$ its reduced tree. A *clan* under a node N in T is a maximal sequence N_1, \dots, N_k of consecutive children of N such that for each two nodes N_i and N_j in the sequence, the two subtrees of T rooted at N_i and N_j are isomorphic, that is, they only differ by the labeling of their leaves, but otherwise have the same structure. In particular, any two adjacent leaves belong to the same clan.

Note that the sequence of children of each internal node is uniquely partitioned into clans, each possibly consisting of a single node. A *leaf clan* is a clan whose nodes are leaves, and a *non-leaf clan* is a clan whose nodes are non-leaves. The first (leftmost) element of each clan is called *the leader* of the clan and the remaining elements are called *followers*.

Using the tree structure of $T(\pi)$, we will show that $\mu(\sigma, \pi)$ can be expressed as a signed sum over a set of embeddings of σ into π that have a special structure. Following the terminology Björner [3,4] and of Sagan and Vatter [5], we call these special embeddings *normal*.

Definition 22. Let σ and π be separable permutations, let $T(\pi)$ be the reduced tree of π . An embedding f of σ into π is called *normal* if it satisfies the following two conditions.

- If a leaf ℓ is maximal omitted by f , then ℓ is the leader of its corresponding leaf clan.
- If an internal node N is maximal omitted by f , then N is a follower in its non-leaf clan.

Let $N(\sigma, \pi)$ denote the set of normal embeddings of σ into π . The *defect* of an embedding $f \in N(\sigma, \pi)$, denoted by $d(f)$, is the number of leaves that are maximal omitted by f . The *sign* of f , denoted by $\text{sgn}(f)$, is defined as $(-1)^{d(f)}$.

We now present our main result.

Theorem 23. If σ and π are (possibly empty) separable permutations, then

$$\mu(\sigma, \pi) = \sum_{f \in N(\sigma, \pi)} \text{sgn}(f).$$

Consider, as an example, the two permutations π and σ depicted in Fig. 1. The children of the root of $T(\pi)$ are partitioned into three clans, where the first clan has three internal nodes, the second clan has a single leaf, and the last clan has a single internal node. Accordingly, there are five normal embeddings of σ into π , depicted in Fig. 2. Of these five normal embeddings, two have sign -1 and three have sign 1 , giving $\mu(\sigma, \pi) = 1$.




Fig. 1. The separating trees of two permutations σ and π .




Fig. 2. The normal embeddings of σ in π (see Fig. 1), together with their signs. The leaves covered by the embedding are represented by black disks, the leaves that are maximal omitted are represented by empty circles. Dotted lines represent subtrees rooted at a maximal omitted internal node. Note that the leaves of such subtrees do not contribute to the sign of the embedding.

Proof of Theorem 23. Let $\bar{\mu}(\sigma, \pi)$ denote the value of $\sum_{f \in N(\sigma, \pi)} \text{sgn}(f)$. Our goal is to prove that $\bar{\mu}(\sigma, \pi)$ is equal to $\mu(\sigma, \pi)$. We proceed by induction on $|\pi|$. For $\sigma = \pi$, we clearly have $\bar{\mu}(\sigma, \pi) = \mu(\sigma, \pi) = 1$, and if π does not contain σ , then $\bar{\mu}(\sigma, \pi) = \mu(\sigma, \pi) = 0$.

Suppose now that $\sigma < \pi$. Since π is separable, it is decomposable or skew-decomposable. Assume, without loss of generality, that π is decomposable. Let $\pi_1 \oplus \cdots \oplus \pi_n$ be its decomposition. Since the values of $\mu(\sigma, \pi)$ are uniquely determined by the recurrences of Propositions 1 and 2, it is enough to show that $\bar{\mu}$ satisfies the same recurrences.

Consider first the case when $\pi_1 = 1$, which is treated by Proposition 1. Let $\sigma_1 \oplus \cdots \oplus \sigma_m$ be the decomposition of σ , let $k = \deg(\pi)$ and let $\ell = \deg(\sigma)$. This means that the leftmost k leaves of $T(\pi)$ are all children of the root node, and they form a leaf clan. Therefore, in any normal embedding, all

the $k-1$ leaves representing π_2, \dots, π_k are covered, because they are followers of π_1 . Necessarily, any element of σ that is embedded to one of the first k elements of π must be one of the first ℓ elements of σ . Consequently, if $k-1 > \ell$, there is no normal embedding of σ into π , and $\bar{\mu}(\sigma, \pi) = 0$.

Suppose now that $k-1 = \ell$. Then, in any normal embedding $f \in N(\sigma, \pi)$, the element π_1 is omitted, the elements representing $\sigma_1, \dots, \sigma_{k-1}$ are embedded on π_2, \dots, π_k , and the elements of $\sigma_{>k-1}$ are embedded to the elements $\pi_{>k}$. The restriction of f to $\sigma_{>k-1}$ is a normal embedding f' from the set $N(\sigma_{>k-1}, \pi_{>k})$, and conversely, a normal embedding f' from $N(\sigma_{>k-1}, \pi_{>k})$ can be uniquely extended into an embedding $f \in N(\sigma, \pi)$. We then have $d(f) = 1 + d(f')$, because π_1 is the only maximal omitted leaf of f that is not a maximal omitted leaf of f' . This shows that $\bar{\mu}(\sigma, \pi) = -\bar{\mu}(\sigma_{>k-1}, \pi_{>k})$.

Assume now that $k-1 < \ell$. Let $N^+(\sigma, \pi)$ denote the set of normal embeddings of σ into π that cover the element π_1 , and let $N^-(\sigma, \pi)$ be the set of those that omit π_1 . By the same argument as in the previous paragraph, we see that $N^+(\sigma, \pi)$ is mapped by a sign-preserving bijection to $N(\sigma_{>k}, \pi_{>k})$, and $N^-(\sigma, \pi)$ is mapped by a sign-reversing bijection to $N(\sigma_{>k-1}, \pi_{>k})$. Consequently, $\bar{\mu}(\sigma, \pi) = \bar{\mu}(\sigma_{>k}, \pi_{>k}) - \bar{\mu}(\sigma_{>k-1}, \pi_{>k})$.

These arguments show that $\bar{\mu}$ satisfies the recurrences of Proposition 1.

Assume now that $\pi_1 > 1$, which corresponds to the situation of Proposition 2. Let $\pi_1 \oplus \dots \oplus \pi_n$ be the decomposition of π , let $\sigma_1 \oplus \dots \oplus \sigma_m$ be the decomposition of σ , and let $k \in [n]$ be the largest integer such that $\pi_1 = \dots = \pi_k$. The n components of π correspond precisely to n children of the root of the tree $T(\pi)$, and the leftmost k components form a non-leaf clan. Therefore, each normal embedding $f \in N(\sigma, \pi)$ must cover the leftmost child of the root, which represents π_1 , but it may omit some of its followers, which represent the components π_2, \dots, π_k . Note that the symbols of σ that are embedded into π_1 by f must form a prefix of the form $\sigma_{\leq i}$, for some $i \in [m]$.

For $f \in N(\sigma, \pi)$, let $I(f) \in [m]$ be the largest number i such that all the symbols of $\sigma_{\leq i}$ are embedded into π_1 , and let $J(f) \in [k]$ be the largest number j such that among the leftmost j children of the root of $T(\pi)$, only the node representing π_1 is covered. Let $N_{i,j}$ be the set $\{f \in N(\sigma, \pi) : I(f) = i, J(f) = j\}$. Notice that an embedding $f \in N_{i,j}$ decomposes in an obvious way into a normal embedding $f_1 \in N(\sigma_{\leq i}, \pi_1)$ and a normal embedding $f_2 \in N(\sigma_{>i}, \pi_{>j})$, and that we have $d(f) = d(f_1) + d(f_2)$, and hence $\text{sgn}(f) = \text{sgn}(f_1) \text{sgn}(f_2)$. This decomposition is a bijection between $N_{i,j}$ and $N(\sigma_{\leq i}, \pi_1) \times N(\sigma_{>i}, \pi_{>j})$. Consequently, we have the identity

$$\sum_{f \in N_{i,j}} \text{sgn}(f) = \sum_{f_1 \in N(\sigma_{\leq i}, \pi_1)} \sum_{f_2 \in N(\sigma_{>i}, \pi_{>j})} \text{sgn}(f_1) \text{sgn}(f_2) = \bar{\mu}(\sigma_{\leq i}, \pi_1) \bar{\mu}(\sigma_{>i}, \pi_{>j}).$$

Summing this identity for each $i \in [m]$ and each $j \in [k]$, we conclude that

$$\bar{\mu}(\sigma, \pi) = \sum_{i=1}^m \sum_{j=1}^k \bar{\mu}(\sigma_{\leq i}, \pi_1) \bar{\mu}(\sigma_{>i}, \pi_{>j}),$$

which is the recurrence of Proposition 2. Therefore, $\bar{\mu}(\sigma, \pi) = \mu(\sigma, \pi)$. \square

Let us now state several consequences of Theorem 23.

Corollary 24. *If π is separable, then $\mu(1, \pi) \in \{0, 1, -1\}$.*

Proof. The permutation 1 can have at most one normal embedding into π . Namely, if $|\pi| > 1$, then $T(\pi)$ has at least one leaf ℓ that is not a leader of its leaf clan, but each of its ancestors is a leader of its non-leaf clan. Such a leaf ℓ must be covered by any normal embedding of any permutation into π . \square

The next corollary confirms a (more general version of a) conjecture of Steingrímsson and Tenner [7].

Corollary 25. *If π and σ are separable permutations, then $|\mu(\sigma, \pi)|$ is at most the number of occurrences of σ in π .*

Proof. This follows from the fact that the number of occurrences of σ in π is clearly at least the number of normal embeddings of σ into π . \square

Recall that a permutation is *layered* if it is the concatenation of decreasing sequences, such that the letters in each sequence are smaller than all letters in subsequent sequences. One example is the permutation 21365487, whose layers are shown by 21–3–654–87. Sagan and Vatter [5] gave a formula for the Möbius function of intervals of layered permutations, and it is easy to see that layered permutations are special cases of separable permutations. Namely, a layered permutation is separable, and its separating tree has depth 2 (except in the trivial cases of the increasing and decreasing permutations), where the children of the root are the layers of the permutation, and the grandchildren of the root are all leaves.

Using Theorem 23, it is easy to show that for $\pi_n = 214365 \cdots (2n)(2n-1)$, we have $\mu(1, \pi_n) = n-1$. Thus, the following result.

Corollary 26. *The value of the Möbius function on intervals $[\sigma, \pi]$ is unbounded, even for layered permutations σ and π .*

5. Concluding remarks, conjectures and open problems

We have shown in Corollary 19 that $\mu(\sigma, \pi)$ can be computed efficiently when σ and π are separable. In fact, by Corollary 21, $\mu(\sigma, \pi)$ can be computed efficiently within any class of permutations generated from a finite set by sums and skew sums. We do not know whether such an efficient computation of μ is possible for more general classes of permutations.

Bose, Buss and Lubiw [2] have shown that it is NP-hard for given permutations π and σ to decide whether π contains σ . In view of this, it seems unlikely that $\mu(\sigma, \pi)$ could be computed efficiently for general permutations σ and π .

Our results imply that for a separable permutation π , the Möbius function $\mu(1, \pi)$ has absolute value at most 1. In fact, the class of separable permutations is the unique largest hereditary class with this property, since any hereditary class not contained in the class of separable permutations must contain 2413 or 3142, and $\mu(1, 2413) = \mu(1, 3142) = -3$. It is natural to consider $\mu(1, \pi)$ as a function of π , and ask whether this function is bounded on a given class of permutations. By Corollary 20, if a hereditary class \mathcal{C} is a $\{\oplus, \ominus\}$ -closure of a finite set of permutations, then $\mu(1, \pi)$ is bounded on \mathcal{C} . We do not know if there is another example of a permutation class on which this function is bounded.

On the other hand, we do not have a proof that $\mu(1, \pi)$ is unbounded on the set of all permutations, although numerical evidence suggests that this is the case. According to our computations, the sequence of maximum values of $|\mu(1, \pi)|$ for $\pi \in \mathcal{S}_n$, starting at $n = 1$, begins with 1, -1, 1, -3, 6, -11, 15, -27, -50, -58, 143, ... For these cases ($n \leq 11$), there is, up to trivial symmetries, a unique permutation for which the Möbius function attains this value. These permutations are

1, 12, 132, 2413, 24153, 351624, 2461735, 35172846,
472951836, 46819210357, 3619411721058.

All of the above permutations are *simple* (except for 132, but there are no simple permutations of length 3). A permutation is simple if it has no segment $a_i a_{i+1} \cdots a_{i+k}$ where $1 \leq k < n-1$ and $\{a_i, a_{i+1}, \dots, a_{i+k}\}$ is a set of consecutive integers (see [1]). In particular, a simple permutation can neither be decomposed nor skew decomposed. We are not able to compute $\mu(1, \pi)$ for all permutations π of length 12, but for simple permutations π the maximum value of $\mu(1, \pi)$ is -261, for $\pi = 472105112831169$.

In light of Corollary 6, to show that $|\mu(1, \pi)|$ is unbounded, it would suffice to show that the maximum of $|\mu(1, \pi)|$ for permutations π of length n , for any n , is attained only by a permutation π that does not start with 1. In that case $|\mu(1, 1 \oplus \pi)| = |\mu(1, \pi)|$, so there would be a permutation τ of length $n + 1$ for which $|\mu(1, \tau)| > |\mu(1, 1 \oplus \pi)| = |\mu(1, \pi)|$.

Question 27. For which permutation classes \mathcal{C} is the function $\mu(1, \pi)$ bounded on \mathcal{C} ? Is $\mu(1, \pi)$ unbounded on the set of all permutations? Can non-trivial upper or lower bounds be found for $\max_{\pi \in \mathcal{S}_n} |\mu(1, \pi)|$, as a function of n ?

We have exhibited several types of intervals whose Möbius function is zero (and more were presented in [7]). Can the following question be answered precisely?

Question 28. When is $\mu(\sigma, \pi) = 0$?

For separable permutation π , we have shown that $|\mu(\sigma, \pi)|$ is at most $\sigma(\pi)$, that is, the number of occurrences of σ in π . This is not true for non-separable π , even when $\sigma = 1$, as shown above. However, it might be possible to bound $|\mu(\sigma, \pi)|$ as a function of $\sigma(\pi)$.

Question 29. Is there an upper bound for $|\mu(\sigma, \pi)|$ that only depends on $\sigma(\pi)$?

The following conjecture has been verified for $n \leq 10$.

Conjecture 30. The maximum value of the Möbius function $\mu(\sigma, \pi)$ for separable permutations σ and π , where π has length $n \geq 3$, is given by

$$\max_k \binom{n-1-k}{k}.$$

This maximum is attained by the permutation π that starts with its odd letters in increasing order, followed by the even letters in decreasing order, and σ of the same form and length $2 \cdot \lfloor (n+1)/2 \rfloor$ if the length of π is $2n$, and $2 \cdot \lfloor (n+1)/2 \rfloor - 1$ if the length of π is $2n - 1$ (such permutations are sometimes called wedge-alternations).

As an example, $\mu(13542, 135798642) = 15 = \binom{9-1-2}{2}$.

Finally, we mention some questions about the topology of the order complexes of intervals in the poset \mathcal{P} . (For definitions, see [6].) Given an interval $[\sigma, \pi]$, let $\Delta(\sigma, \pi)$ be the order complex of the poset obtained from $[\sigma, \pi]$ by removing σ and π .

Question 31.

- (1) For which σ and π does $\Delta(\sigma, \pi)$ have the homotopy type of a wedge of spheres?
- (2) Let Γ be the subcomplex of $\Delta(\sigma, \pi)$ induced by those elements τ of $[\sigma, \pi]$ for which $\mu(\sigma, \tau) = 0$. Is Γ a pure complex?
- (3) If σ occurs precisely once in π , and $\mu(\sigma, \pi) = \pm 1$, is $\Delta(\sigma, \pi)$ homotopy equivalent to a sphere?
- (4) For which σ and π is $\Delta(\sigma, \pi)$ shellable?

We should mention that for $\sigma = 231$ and $\pi = 231564$, the order complex $\Delta(\sigma, \pi)$ is not shellable; it consists of two connected components, each of which is contractible. However, removing from $[231, 231564]$ all elements τ with $\mu(231, \tau) = 0$, we obtain a shellable complex, namely a four-element boolean algebra. For parts (2) and (3) in Question 31, we know no counterexamples. Since we have so far only examined intervals of low rank, our evidence is not strong.

Acknowledgments

We are grateful to two anonymous referees, who made many useful suggestions.

References

- [1] M.H. Albert, M.D. Atkinson, M. Klazar, The enumeration of simple permutations, *J. Integer Seq.* 6 (2003), 03.4.4.
- [2] P. Bose, J.F. Buss, A. Lubiw, Pattern-matching for permutations, *Inform. Process. Lett.* 65 (1998) 277–283.
- [3] A. Björner, The Möbius function of subword order, in: *Invariant Theory and Tableaux*, Minneapolis, MN, 1988, in: IMA Vol. Math. Appl., vol. 19, Springer, New York, 1990, pp. 118–124.
- [4] A. Björner, The Möbius function of factor order, *Theoret. Comput. Sci.* 117 (1993) 91–98.
- [5] B.E. Sagan, V. Vatter, The Möbius function of a composition poset, *J. Algebraic Combin.* 24 (2006) 117–136.
- [6] R.P. Stanley, *Enumerative Combinatorics*, vol. 1, Cambridge Stud. Adv. Math., vol. 49, Cambridge University Press, Cambridge, 1997.
- [7] E. Steingrímsson, B.E. Tenner, The Möbius function of the permutation pattern poset, arXiv:0902.4011.
- [8] H. Wilf, The patterns of permutations, *Discrete Math.* 257 (2002) 575–583.
- [9] V. Yugandhar, S. Saxena, Parallel algorithms for separable permutations, *Discrete Appl. Math.* 146 (2005) 343–364.