

The Möbius function of separable permutations (extended abstract)

Vít Jelínek^{1†} and Eva Jelínková^{2‡} and Einar Steingrímsson¹

¹The Mathematics Institute, School of Computer Science, Reykjavik University, Kringlan 1, IS-103 Reykjavik, Iceland

²Department of Applied Mathematics, Charles University, Malostranské náměstí 25, 118 00 Praha, Czech Republic

Abstract. 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. Using the notion of separating tree, we give a computationally efficient formula for the Möbius function of an interval (q, p) in the poset of separable permutations ordered by pattern containment. A consequence of the formula is that the Möbius function of such an interval (q, p) is bounded by the number of occurrences of q as a pattern in p . The formula also implies that for any separable permutation p the Möbius function of $(1, p)$ is either 0, 1 or -1 .

Résumé. Une permutation est séparable si elle peut être générée à partir de la permutation 1 par des sommes directes et des sommes indirectes, ou de façon équivalente, si elle évite les motifs 2413 et 3142. En utilisant le concept de l'arbre séparant, nous donnons une formule pour le calcul efficace de la fonction de Möbius d'un intervalle de (q, p) dans l'ensemble partiellement ordonné des permutations séparables. Une conséquence est que la fonction de Möbius de (q, p) pour q et p séparables est bornée par le nombre d'occurrences de q comme un motif en p . Nous montrons aussi que pour une permutation p séparable, la fonction de Möbius de $(1, p)$ est soit 0, 1 ou -1 .

Keywords: Möbius function, pattern poset, separable permutations.

1 Introduction

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

A classical question to ask for any combinatorially defined poset is what its Möbius function is. 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 (q, p) 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

[†]Jelínek and Steingrímsson were supported by grant no. 090038011 from the Icelandic Research Fund.

[‡]Jelínková was supported by project 1M0021620838 of the Czech Ministry of Education.

case. In this paper, we find a (computationally effective) formula for the Möbius function of an interval (q, p) , where p is a separable permutation. This is a generalization of the results of Sagan and Vatter, and is based on similar principles.

Recently, Steingrímsson and Tenner [7] exhibited a class of intervals whose Möbius function is zero and described certain other intervals where the Möbius function is either 1 or -1 . In addition, they conjectured that for permutations q and p avoiding the pattern 132 the absolute value of the Möbius function of the interval (q, p) , denoted $\mu(q, p)$, is bounded by the number of occurrences of q in p .

With the help of our formula for $\mu(q, p)$, we prove a more general version of this conjecture. We show that for any interval (q, p) of separable permutations q and p we have $\mu(q, p) \leq q(p)$ (for general p and q this inequality does not hold). In particular, if p has a single occurrence of q then $\mu(q, p)$ is either 1, 0 or -1 .

We also prove a conjecture mentioned in [7], showing that for any separable permutation p , $\mu(1, p)$ is either 1, 0 or -1 , where 1 in $(1, p)$ is the single permutation of length 1. In addition, we give some results on the Möbius function of various special intervals of arbitrary permutations, which we then use to conclude that $\mu(1, p)$ is bounded on certain classes of permutations. These results show how to express the Möbius function $\mu(1, p)$ for a *decomposable* permutation p in terms of $\mu(1, p_i)$, where p_i are the summands in the decomposition of p . For “most” decomposable permutations p this shows that $\mu(1, p)$ is zero.

2 Definitions and Preliminaries

An *interval* (q, p) in a poset (\mathcal{P}, \leq) is the set $\{r : q \leq r \leq p\}$. In this paper, we deal exclusively with intervals of the poset of permutations ordered by pattern containment.

The Möbius function $\mu(q, p)$ of an interval (q, p) is uniquely defined by setting $\mu(q, q) = 1$ for all q and requiring that

$$\sum_{r \in (q, p)} \mu(q, r) = 0 \quad (1)$$

for every $q < p$.

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

$$\mu(q, p) = \sum_{C \in \mathcal{C}(q, p)} (-1)^{|C|} = \sum_i (-1)^i c_i, \quad (2)$$

where \mathcal{C} is the set of chains in (q, p) that contain both q and p , and c_i is the number of such chains of length i in (q, p) . For details and further information on this see [6].

The *direct sum*, $a + b$, of two nonempty permutations a and b is the permutation obtained by concatenating a and b' , where b' is b with all letters incremented by the number of letters in a . A permutation that can be written as a direct sum is *decomposable*, otherwise it is *indecomposable*. Examples are $2314576 = 231 + 12 + 21$, and 231 , which is indecomposable. In the *skew sum* of a and b , denoted by $a * b$, we increment the letters of a by the length of b to obtain a' and then concatenate a' and b . For example, $6743512 = 12 * 213 * 12$.

A *decomposition* of p is an expression $p = p_1 + p_2 + \dots + p_k$ in which each summand p_i is indecomposable. The summands p_1, \dots, p_k will be called *the blocks* of p . Every permutation p has a unique decomposition (including an indecomposable permutation p , whose decomposition has a single block p).

A permutation is *separable* if it can be generated by successive sums and skew sums of the permutation 1. Being separable is equivalent to avoiding the patterns 2413 and 3142, that is, containing no occurrences of them. Separable permutations have nice algorithmic properties. For instance, Bose, Buss and Lubiw [2] have shown that it can be decided in polynomial time whether $q \leq p$ when p and q are separable, while for general permutations the problem is NP-hard.

3 Möbius function of separable permutations

Let us now consider the values of $\mu(q, p)$ for separable permutations q and p .

The recursive structure of separable permutations makes it often convenient to represent a separable permutation by a tree that describes how the permutation may be obtained from smaller permutations as a sum or skew sum. Let us 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 p , defined recursively as follows:

- If T has a single vertex, it represents the singleton permutation 1.
- Assume T has more than one vertex. 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 . If the root of T is a direct node (skew node), then T represents the permutation $p_1 + \dots + p_k$ ($p_1 * \dots * p_k$, respectively).

Note that the leaves of T correspond bijectively to the elements of p . 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 elements of p .

A given separable permutation may be represented by more than one separating tree. A separating tree is called *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 p is represented by a unique reduced tree, denoted by $T(p)$. We assume that each leaf of T is labelled by the corresponding element of p .

The (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 $p = (p(1), p(2), \dots, p(n))$ and $q = (q(1), q(2), \dots, q(m))$ be two permutations, with $q \leq p$. An *embedding* of q into p is a function $f: [m] \rightarrow [n]$ with the following two properties:

- for every $i, j \in [m]$, if $i < j$ then $f(i) < f(j)$ (i.e., f is monotone increasing).
- for every $i, j \in [m]$, if $q(i) < q(j)$, then $p(f(i)) < p(f(j))$ (i.e., f is order-preserving).

Let f be an embedding of q into p . We say that a leaf ℓ of $T(p)$ is *covered* by the embedding f if the element of p 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 N of $T(p)$ 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 $T(p)$ is not omitted.

Assume that p is a separable permutation and $T(p)$ its reduced tree. Two nodes N_1 and N_2 of a tree $T(p)$ are called *twins* if they are siblings (i.e., share a common parent P), they appear consecutively in the sequence of children of P , and the two subtrees of T rooted at N_1 and N_2 are isomorphic (i.e., they only differ by the labeling of their leaves, but otherwise have the same structure). In particular, any two adjacent leaves are twins.

A *run* under a node N in T is a maximal sequence N_1, \dots, N_k of children of N such that each two consecutive elements of the sequence are twins. Note that the sequence of children of each internal node is uniquely partitioned into runs (possibly consisting of a single node). A *leaf run* is a run whose nodes are leaves, and a *non-leaf run* is a run whose nodes are non-leaves. The first (i.e., leftmost) element of each run is called *the leader* of the run, the remaining elements are called *followers*.

Using the tree structure of $T(p)$, we will show that $\mu(q, p)$ can be expressed as a signed sum over a set of particularly ‘nice’ embeddings of q into p . Following the terminology of Sagan and Vatter [5], we call these nice embeddings ‘normal’.

Definition 1 Let q and p be separable permutations, let $T(p)$ be the reduced tree of p . An embedding f of q into p 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 run.
- If an internal node N is maximal omitted by f , then N is a follower in its non-leaf run.

Let $N(q, p)$ denote the set of normal embeddings of q into p . The *defect* of an embedding $f \in N(q, p)$, 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 2 If q and p are separable permutations, then

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

Consider, as an example, the two permutations p and q depicted on Figure 1. The children of the root of $T(p)$ are partitioned into three runs, where the first run has three internal nodes, the second run has a single leaf, and the last run has a single internal node. Accordingly, there are five normal embeddings of q into p , depicted in Figure 2. Of these five normal embeddings, two have sign -1 and three have sign 1, giving $\mu(q, p) = 1$.

Although the formula for $\mu(q, p)$ given in Theorem 2 may in general involve exponentially many summands, it allows us to compute $\mu(q, p)$ in time which is polynomial in $|p| + |q|$, using a simple dynamic programming approach which we outline in Subsection 3.1.

To prove Theorem 2, we show that the formula for $\mu(q, p)$ satisfies the Möbius function recurrence $\sum_{r \in (q, p)} \mu(q, r) = 0$ for each nontrivial interval (q, p) of separable permutations. To verify this recurrence, we construct a sign-reversing involution on the set $\bigcup_{r \in (q, p)} N(q, r)$. The construction of this




Fig. 1: The separating trees of two permutations q and p




Fig. 2: The normal embeddings of q in p , 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 vertex. Note that the leaves of such subtrees do not contribute to the sign of the embedding.

bijection is very technical and relies on a complicated case analysis. We omit the construction from this extended abstract.

Let us now state several consequences of Theorem 2.

Corollary 3 *If p is separable, then $\mu(1, p) \in \{0, 1, -1\}$.*

Proof: The permutation 1 can have at most one normal embedding into p . It is in fact easy to observe that if $|p| > 1$, then $T(p)$ has at least one leaf ℓ that is not a leader of its leaf run, but each of its ancestors is a leader of its non-leaf run. Such a leaf ℓ must be covered by any normal embedding of any permutation into p . \square

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

Corollary 4 *If p and q are separable permutations, then $|\mu(q, p)|$ is at most the number of occurrences of q in p .*

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

3.1 The algorithm

As we already pointed out, Theorem 2 allows us to compute $\mu(p, q)$ in polynomial time for any separable permutations p and q . We outline the main ideas of the algorithm.

Given p and q , we construct the trees $T(p)$ and $T(q)$. We then check all the subtrees of $T(p)$ and mark each subtree which is a follower in its non-leaf run. We call such a subtree *marked*.

In order to use dynamic programming to compute the Möbius function, we define an extension of normal embeddings of trees.

Let N be an internal node of a tree, and let T_1, \dots, T_k be the subtrees rooted at the children of N . A *range* (N, i, j) is the subtree of $T(p)$ induced by $N \cup T_i \cup T_{i+1} \cup \dots \cup T_j$. Observe that a tree has quadratically many ranges.

The leaves of each range in $T(p)$ represent a subsequence of p . We may thus speak of embeddings of a range in $T(q)$ into a range of $T(p)$.

An embedding f of a range in $T(q)$ to a range (N_p, i, j) in $T(p)$ is called *normal* if it satisfies the following conditions.

- If a leaf ℓ in (N_p, i, j) is a maximal omitted node in f , then ℓ is the leader of its corresponding leaf run in p .
- If an internal node N' in (N_p, i, j) is maximal omitted by f , then N' is marked.

We say that a normal embedding to a range (N_p, i, j) is *even* if it has an even number of maximal omitted leaves, and *odd* otherwise.

For every pair of ranges $(N_q, \ell, m) \subseteq T(q)$ and $(N_p, i, j) \subseteq T(p)$, we compute the number of odd and even normal embeddings of (N_q, ℓ, m) to (N_p, i, j) , denoted by $\text{emb-odd}(N_q, \ell, m, N_p, i, j)$ and $\text{emb-even}(N_q, \ell, m, N_p, i, j)$ respectively.

For a range (N_p, i, j) with a single leaf, $\text{emb-odd}(N_q, \ell, m, N_p, i, j)$ and $\text{emb-even}(N_q, \ell, m, N_p, i, j)$ can be easily computed. To compute $\text{emb-odd}(N_q, \ell, m, N_p, i, j)$ and $\text{emb-even}(N_q, \ell, m, N_p, i, j)$ for a

range (N_p, i, j) that contains more than one leaf, we assume that we already know the values of emb-odd and emb-even for any range of $T(q)$ and for any range of $T(p)$ properly contained in (N_p, i, j) . It is not difficult to see that the values of $\text{emb-odd}(N_q, \ell, m, N_p, i, j)$ and $\text{emb-even}(N_q, \ell, m, N_p, i, j)$ can then be determined in linear time from previously computed values of emb-odd and emb-even .

Having computed all the values of emb-odd and emb-even , we evaluate $\mu(q, p)$. Let R_p and R_q be the root vertices of $T(p)$ and $T(q)$, and assume that R_p has k children and R_q has j children. Then $\mu(q, p) = \text{emb-even}(R_q, 1, j, R_p, 1, k) - \text{emb-odd}(R_q, 1, j, R_p, 1, k)$.

It is clear that the algorithm works in polynomial time.

4 The Möbius function for general permutations

In this section, we demonstrate several basic properties of $\mu(p, q)$, for general (not necessarily separable) permutations p and q .

Lemma 5 *Let q be an indecomposable permutation. Let p be a permutation containing q that has the form $p = r + 1 + s$, where r, s are nonempty permutations. Then $\mu(q, p) = 0$.*

Proof: Proceed by induction on $|p|$. The smallest permutations p satisfying the assumptions of the lemma are $p = q + 1 + 1$ and $p = 1 + 1 + q$. For these permutations, the statement holds. Assume now that $|p| > |q| + 2$, and $p = r + 1 + s$ with r and s nonempty. We have

$$\mu(q, p) = - \sum_{t < p} \mu(q, t) = - \left(\sum_{t \leq r+s} \mu(q, t) + \sum_{t < p, t \not\leq r+s} \mu(q, t) \right).$$

The sum $\sum_{t \leq r+s} \mu(q, t)$ is zero by the definition of μ (note that $r + s \neq q$), while every summand in $\sum_{t < p, t \not\leq r+s} \mu(q, t)$ is zero by induction (note that both $r + 1$ and $1 + s$ are subpermutations of $r + s$). \square

Let us say that a permutation p is *low* if p can be written as $p = 1 + r$, where r is a nonempty permutation. A permutation is *high* if it is not low. Note that the permutation 1 is high, and each indecomposable permutation is high as well.

Lemma 6 *If p and q are high permutations, then $\mu(q, 1 + p) = -\mu(q, p)$.*

Proof: We know that $\mu(q, 1 + p)$ can be written as

$$\mu(q, 1 + p) = \sum_{C \text{ chain from } q \text{ to } 1 + p} (-1)^{\ell(C)}, \quad (3)$$

where $\ell(C)$ denotes the length of the chain C .

We now distinguish two types of chains from q to p .

Type 1: A chain from q to $1 + p$ is of *Type 1* if it contains the element p (which is high by assumption). Necessarily, p is the penultimate element of the chain. There is a parity-reversing bijection between all chains from q to p and the Type 1 chains: the bijection works by simply adding $1 + p$ at the end of the chain from q to p . This shows that Type 1 chains contribute $-\mu(q, p)$ to the right hand side of (3).

Type 2: A chain from q to $1 + p$ is of *Type 2* if it does not contain p . We claim that the contributions of Type 2 chains to (3) sum to 0. For such a chain C , let $h(C)$ be the largest high permutation appearing in the chain C . We split Type 2 chains into two groups:

Group A contains the Type 2 chains with the property that the permutation $h(C)$ is followed by the permutation $1 + h(C)$ in the chain. Note that by assumption, $h(C)$ is different from p (else the chain would be of Type 1), so $1 + h(C)$ is not the last element of the chain.

Group B contains the Type 2 chains where $h(C)$ is followed by a permutation r different from $1 + h(C)$. The permutation r must be low, so it is of the form $1 + s$, where $s \neq h(C)$. Since $r = 1 + s$ contains $h(C)$, and $h(C)$ is high, it follows that $h(C)$ is in fact (properly) contained in s .

There is then a parity-reversing bijection between Group A and Group B, which works by removing the element $1 + h(C)$ from a chain C . This shows that Type 2 chains contribute 0 to $\mu(q, 1 + p)$. \square

Let us make some simple observations about embeddings. Assume that p has a decomposition $p = p_1 + p_2 + \cdots + p_k$, and assume that q is indecomposable. Let f be an embedding of q into p . Since q is indecomposable, it is easy to see that the embedding f must in fact embed q into a single block p_x of p (formally speaking, for every $i \in [m]$, we have $|p_1 + \cdots + p_{x-1}| < f(i) \leq |p_1 + \cdots + p_x|$). More generally, if q has a decomposition $q = q_1 + \cdots + q_\ell$ and f is an embedding of q into p , then each block of q is embedded by f into a single block of p .

For an integer α and a permutation q , let αq denote the sum $q + q + \cdots + q$ with α summands.

Theorem 7 *Let p be a decomposable permutation of order at least 3. Assume p has a decomposition $p_1 + p_2 + \cdots + p_k$, where neither p_1 nor p_k are equal to 1. Assume that q is an indecomposable permutation.*

If $\mu(q, p) \neq 0$, then all the blocks of p are equal to a single indecomposable permutation r , and in such case $\mu(q, p) = \mu(q, r)$.

The rest of this text is devoted to the proof of Theorem 7.

The proof proceeds by induction on $|p|$. For $|p| \leq 3$, the statement holds trivially. Let us assume that p is a decomposable permutation of order at least 4, and that neither the first nor the last block of the decomposition of p is equal to 1. If any of the internal blocks of p is equal to 1, we use Lemma 5 to show that $\mu(q, p) = 0$.

In the rest of the proof, we assume that all the blocks of p have order at least 2. To compute $\mu(q, p)$, we use the expression

$$\mu(q, p) = - \sum_{t < p} \mu(q, t) \quad (4)$$

and we will show that all the terms on the right-hand side of (4) cancel out, except at most one term, whose value we will be able to determine.

We classify all the permutations $t < p$ into the following four types:

- The permutations of the form $t = 1 + 1 + r$, where r is not empty. For any such permutation, we have $\mu(q, t) = 0$ by Lemma 5, so these permutations do not contribute anything to the right-hand side of (4).
- The permutations of the form $t = 1 + r$, where r is high. We will call such permutations *extended*.
- The permutations t such that t is high and $1 + t < p$. We will call such permutations *extendable*.

- The permutations t such that t is high, and $1+t$ is not contained in p . (Note that we can never have $1+t = p$, because we assume that all the blocks of p have order at least 2.) We will call these permutations *significant*.

Note that a permutation t is extendable if and only if $1+t$ is extended, and that this provides a bijection between extendable and extended proper subpermutations of p . Moreover, for any extendable t , we have $\mu(q, t) = -\mu(q, 1+t)$ by Lemma 6. This shows that the contribution of extendable permutations on the right side of (4) cancels exactly with the contribution of the extended permutations. Consequently, we have

$$\mu(q, p) = - \sum_{t \text{ significant}} \mu(q, t) \quad (5)$$

Applying induction, we see that a significant permutation t that has nonzero $\mu(q, t)$ can be of one of the following two forms:

- t has the decomposition $t = \alpha r$ for some $\alpha \geq 1$ and some indecomposable permutation r of order at least 2. In such case, we know from induction hypothesis that $\mu(q, t) = \mu(q, r)$. We call such a permutation t a *significant permutation of type 1* and we say that t is an r -permutation. Note that we allow the possibility that $t = r$.
- t has the decomposition $t = \beta r + 1$, for some $\beta \geq 1$ and some indecomposable permutation r of order at least 2. In such case, we know from induction hypothesis and from Lemma 6 that $\mu(q, t) = -\mu(q, r)$. We call such a permutation t a *significant permutation of type 2*, and we again say that t is an r -permutation.

To complete the proof of Theorem 7, we plan to show that for any given indecomposable permutation $r < p$, the contribution of r -permutations of type 1 on the right side of (5) cancels with the contribution of r -permutations of type 2. The only exception to this exact cancellation will occur when the permutation p has the decomposition $p = \gamma r$ for some $\gamma \geq 2$ — in such case, there is a type-2 r -permutation $(\gamma-1)r+1$, but there is no significant type-1 r -permutation. To prove that the cancellations work as required, we first need to prove several claims.

Let us introduce some more notation. For an indecomposable permutation $r < p$, we define $\alpha_r = \max\{\alpha; \alpha r \leq p\}$ and $\beta_r = \max\{\beta; \beta r + 1 \leq p\}$. Observe that $\alpha_r - 1 \leq \beta_r \leq \alpha_r$.

Claim 8 *Let $r < p$ be an indecomposable permutation. If t is a significant r -permutation of type 1, then $t = \alpha_r r$. If t is a significant r -permutation of type 2, then $t = \beta_r r + 1$.*

Proof: Assume that t is an r -permutation of type 1. By definition, $t = \alpha r$ for some $\alpha \geq 1$. If α were greater than α_r , then t would not be a subpermutation of p . If α were smaller than α_r , then t would be extendable, because $1+t < \alpha_r r \leq p$. Thus, if t is to be significant, it must be of the form $t = \alpha_r r$. The argument for permutations of type 2 is the same. \square

Claim 8 shows that for a given r there can be at most one significant r -permutation of each type. However, for some r it might happen that even the permutations of the ‘correct’ form $\alpha_r r$ and $\beta_r r + 1$ are extendable, and hence not significant.

From the definition of α_r , it is clear that $\alpha_r > 0$ for any $r < p$. However, we might have $\beta_r = 0$ even when $r < p$. The next claim allows us to avoid such degenerate cases in our considerations.

Claim 9 *Let $r < p$ be an indecomposable permutation. If $\beta_r = 0$, then there is no significant r -permutation of any kind.*

Proof: Assume that $\beta_r = 0$ for some indecomposable permutation $r < p$. This implies that $\alpha_r = 1$, and that the permutation $r + 1$ is not contained in p . Recall that p has the decomposition $p_1 + \cdots + p_k$. Since r is indecomposable, r must be contained in at least one block of p .

Since $r + 1$ is not contained in p , we know that the last block p_k is the only block of p that contains r . Since p is assumed to have at least two blocks, this shows that p contains $1 + r$. This means that the permutation $\alpha_r r = r$ is extendable, hence not significant. The permutation $\beta_r r + 1 = 1$ is also extendable. By Claim 8, there can be no significant r -permutations different from $\alpha_r r$ and $\beta_r r + 1$. Thus, there are no significant r -permutations. \square

Before we state our next claim, we need some preparation. Consider the decomposition $p_1 + \cdots + p_k$ of p . For an integer $i \leq k$ and for an indecomposable permutation $r < p$, we define

$$\alpha_r^{(i)} = \max\{\alpha; \alpha r \leq p_i\}$$

and

$$\beta_r^{(i)} = \max\{\beta; \beta r + 1 \leq p_i\}.$$

Notice that $\alpha_r = \alpha_r^{(1)} + \alpha_r^{(2)} + \cdots + \alpha_r^{(k)}$ and $\beta_r = \alpha_r^{(1)} + \alpha_r^{(2)} + \cdots + \alpha_r^{(k-1)} + \beta_r^{(k)}$.

Claim 10 *Let $r < p$ be an indecomposable permutation. The following statements are equivalent:*

1. *The permutation $\alpha_r r$ is extendable.*
2. *The permutation $\beta_r r + 1$ is extendable.*
3. *The block p_1 contains the permutation $1 + \alpha_r^{(1)} r$.*

Proof: Let us prove that 1 implies 3. Assume that $\alpha_r r$ is extendable. This means that the permutation $t = 1 + \alpha_r r$ is contained in p . Consider an embedding f of t into p . Choose f in such a way that the value of $f(1)$ is as small as possible.

The embedding f cannot embed into any given block p_i more than $\alpha_r^{(i)}$ copies of r , because that would contradict the definition of $\alpha_r^{(i)}$. In particular, f embeds $\alpha_r^{(1)}$ copies of r inside p_1 .

Note also that the leftmost element of t (which is equal to 1) is embedded inside p_1 by f . If f would embed the leftmost element of q into any other block, it would mean that f does not embed anything into p_1 , and this would contradict the choice of f . We conclude that f embeds the permutation $1 + \alpha_r^{(1)} r$ inside p_1 , which shows that 1 implies 3, as claimed. By an analogous argument, we may also see that 2 implies 3.

Let us now prove that 3 implies 1. Assume that $1 + \alpha_r^{(1)} r$ is contained in p_1 . By definition, we also know that for each $i \geq 2$, the block p_i contains the permutation $\alpha_r^{(i)} r$. Putting this together, we see that p contains $1 + \alpha_r^{(1)} r + \alpha_r^{(2)} r + \cdots + \alpha_r^{(k)} r = 1 + \alpha_r r$. This shows that $\alpha_r r$ is extendable, and 1 holds. By the same reasoning, we see that 3 implies 2. This completes the proof of the claim. \square

Let r be again an indecomposable permutation contained in p . It is clear that the permutation $\beta_r r + 1$ is either significant or extendable. It is also clear that the permutation $\alpha_r r$ is either significant, or extendable, or equal to p . In view of Claim 10, for each r , exactly one of the following three cases must occur:

- The permutation $\alpha_r r$ is extendable, and so is $\beta_r r + 1$. In this case there is no significant r -permutation of any type.
- The permutation $\alpha_r r$ is significant, and so is $\beta_r r + 1$. In this case, these two permutations are exactly the only two significant r -permutations, and their contributions to the sum on the right-hand side of Equation (5) cancel out.
- The permutation $\alpha_r r$ is equal to p . In this case, it is easy to see that $\beta_r = \alpha_r - 1$, and that the permutation $\beta_r r + 1$ is significant. This case arises if and only if all the blocks in the decomposition of p are equal to r . From the induction hypothesis and from Lemma 6, we know that $\mu(q, \beta_r r + 1) = -\mu(q, r)$.

We conclude that if the blocks in the decomposition of p are all equal to r , then $\mu(q, p) = \mu(q, r)$. On the other hand, if the blocks of p are not all equal, then $\mu(q, p) = 0$. This completes the proof of Theorem 7.

Let us say that a class of permutations \mathcal{C} is *sum-closed* if for each $p, q \in \mathcal{C}$, the class \mathcal{C} also contains $p + q$. Similarly, \mathcal{C} is *skew-closed* if $p, q \in \mathcal{C}$ implies $p * q \in \mathcal{C}$. For a set \mathcal{P} of permutations, the *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 corollary is an immediate consequence of Lemma 6 and Theorem 7.

Corollary 11 *Suppose that q is a permutation that is neither decomposable nor skew-decomposable. Let \mathcal{P} be any set of permutations. Then*

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

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

5 Concluding remarks and open problems

We have shown that the Möbius function $\mu(q, p)$ can be computed efficiently whenever p and q are separable permutations. With some additional arguments (omitted from this extended abstract), we can in fact show that the values of $\mu(q, p)$ can be computed in polynomial time whenever p belongs to a fixed class of permutations that is the closure of a finite set. We do not know whether this result can be extended to larger classes of permutations.

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

Our results imply that for a separable permutation p , the Möbius function $\mu(1, p)$ has absolute value at most 1. In fact, the class of separable permutations is the largest hereditary class with this property, since any hereditary class not contained in the class of separable permutations must contain 2413 or 3142, and $\mu(1, 2413) = \mu(1, 3142) = -3$. We have also seen that $\mu(1, p)$ is bounded on any permutation class

that is a closure of finitely many permutations. Is there another example of a class on which $\mu(1, p)$ is bounded?

References

- [1] M. H. Albert, M. D. Atkinson, and 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, *Information Processing Letters* **65** (1998), 277–283.
- [3] A. Björner, The Möbius function of subword order, in *Invariant Theory and Tableaux* (Minneapolis, MN, 1988), vol. 19 of IMA Vol. Math. Appl. Springer, New York (1990) 118–124.
- [4] A. Björner, The Möbius function of factor order, *Theoretical Computer Sci.* **117** (1993), 91–98.
- [5] B. E. Sagan and 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 Studies in Advanced Mathematics, no. 49, Cambridge University Press, Cambridge, 1997.
- [7] E. Steingrímsson, B. E. Tenner, The Möbius function of the permutation pattern poset, [arXiv:0902.4011](https://arxiv.org/abs/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 Applied Mathematics* **146** (2005), 343–364.