

The Möbius function of the consecutive pattern poset

Antonio Bernini

Dipartimento di Sistemi e Informatica
University of Firenze, Italy

`bernini@dsi.unifi.it`

Luca Ferrari

Dipartimento di Sistemi e Informatica
University of Firenze, Italy

`ferrari@dsi.unifi.it`

Einar Steingrímsson*

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

`einar.steingrimsson@cis.strath.ac.uk`

Submitted: Feb 28, 2011; Accepted: Jun 23, 2011; Published: Jul 15, 2011

Mathematics Subject Classification: 05A05, 06A07.

Abstract

An occurrence of a consecutive permutation pattern p in a permutation π is a segment of consecutive letters of π whose values appear in the same order of size as the letters in p . The set of all permutations forms a poset with respect to such pattern containment. We compute the Möbius function of intervals in this poset. For most intervals our results give an immediate answer to the question. In the remaining cases, we give a polynomial time algorithm to compute the Möbius function. In particular, we show that the Möbius function only takes the values -1 , 0 and 1 .

1 Preliminaries and introduction

For the poset of classical permutation patterns, the first results about its Möbius function were obtained in [SV]. Further results appear in [ST] and [BJJS]. The general problem in this case of classical patterns seems quite hard. In contrast, the poset of consecutive pattern containment has a much simpler structure. In this paper we compute the Möbius function of that poset. In most cases our results give an immediate answer. In the remaining cases, we give a polynomial time recursive algorithm to compute the

*Steingrímsson was supported by grant no. 090038012 from the Icelandic Research Fund.

Möbius function. In particular, we show that the Möbius function only takes the values -1 , 0 and 1 .

An interesting result to note in connection to this is Björner's paper [Bj] on the Möbius function of factor order. Although that poset is quite different from ours, there are interesting similarities. In particular, both deal with consecutive subwords and the possible values of the Möbius function are -1 , 0 and 1 in both cases.

Unless otherwise specified, all permutations in this paper are taken to be of the set $[d] = \{1, 2, \dots, d\}$ for some positive integer d . We denote by $\mathcal{S}_d$ the set of all such permutations for a given d . An occurrence of a *consecutive pattern* $\sigma = a_1 a_2 \dots a_k$ in a permutation $\tau = b_1 b_2 \dots b_n$ is a subsequence $b_{i+1} b_{i+2} \dots b_{i+k}$ in τ , whose letters appear in the same order of size as the letters in σ . As an example, there are three occurrences of the consecutive pattern 231 in the permutation 563724891, namely 563, 372 and 891. On the other hand, the permutation 253641 *avoids* 231, since it contains no consecutive occurrence of that pattern.

Consecutive permutation patterns are special cases of the generalized permutation patterns introduced in [BS], and they are not to be confused with the classical permutation patterns, whose occurrences in a permutation do not have to be contiguous. The enumerative properties of occurrences of various consecutive permutation patterns were first studied systematically in [EN], but these results will not concern us, as there seems to be no connection between them and the Möbius function studied here.

The set of all permutations forms a poset $\mathcal{P}$ with respect to consecutive pattern containment. In other words, if $\sigma \in \mathcal{S}_k$ and $\tau \in \mathcal{S}_n$, then $\sigma \leq \tau$ in $\mathcal{P}$ if σ occurs as a consecutive pattern in τ . We write $\sigma < \tau$ if $\sigma \leq \tau$ and $\sigma \neq \tau$. As usual in poset terminology, a permutation τ *covers* σ (and σ *is covered by* τ) if $\sigma < \tau$ and there is no permutation π such that $\sigma < \pi < \tau$. Note that if τ covers σ then $|\tau| - |\sigma| = 1$, where $|\pi|$ is the length of π . The interval $[x, y]$ in a poset P , where x and y are elements of P , is defined by $[x, y] = \{z \in P \mid x \leq z \leq y\}$. The *rank* of an interval $[\sigma, \tau]$ in $\mathcal{P}$ is the difference $|\tau| - |\sigma|$. The rank of an element π in $[\sigma, \tau]$ is defined to be the rank of the interval $[\sigma, \pi]$.

A *filter* in a poset P is a set $S \subseteq P$ such that if $x > y$ and $y \in S$, then $x \in S$. An *ideal* is a set $S \subseteq P$ such that if $x < y$ and $y \in S$, then $x \in S$. A *principal filter* is a filter with a single minimal element, and a *principal ideal* is an ideal with a single maximal element. In each case, the single minimal/maximal element is said to *generate* the filter/ideal.

In the poset $\mathcal{P}$ consider the interval $[\sigma, \tau]$. Our aim is to compute $\mu(\sigma, \tau)$, where μ is the Möbius function of the incidence algebra of $\mathcal{P}$. The Möbius function is recursively defined by setting $\mu(x, x) = 1$ for all x , and, if $x \neq y$,

$$\mu(x, y) = - \sum_{x \leq z < y} \mu(x, z). \quad (1)$$

In particular, if $x \not\leq y$, then $\mu(x, y) = 0$.

It is well known (and follows easily from Philip Hall's Theorem, see [St, Proposition 3.8.5]) that the Möbius function of an interval I is equal to the Möbius function of the interval I' obtained by "turning I upside-down," that is, the interval I' for which x is

declared to be less than or equal to y when $y \leq x$ in I . Thus, we also have

$$\mu(x, y) = - \sum_{x < z \leq y} \mu(z, y). \quad (2)$$

Examples of the two ways to compute the Möbius function, described in Equations (1) and (2), are given in Figure 1. Denoting with $[x, y]$ the interval whose Hasse diagram is depicted in the figure, the label of an element z on the left is $\mu(x, z)$, whereas the label of z on the right is $\mu(z, y)$.

Note also that it follows from either (1) or (2) that the sum of $\mu(x, z)$ over *all* elements z in an interval $[x, y]$ is 0, that is

$$\sum_{x \leq z \leq y} \mu(x, z) = 0. \quad (3)$$

This identity will be used frequently.




Figure 1: Computing the Möbius function of an interval from bottom to top (left) and from top to bottom (right). On the left the label of an element z is $\mu(x, z)$, whereas on the right it is $\mu(z, y)$, where x and y are the bottom and top elements in the poset, respectively.

Definition 1.1 Given a sequence of distinct integers $s = s_1 s_2 \dots s_d$, the standard form of s is the permutation $\pi = a_1 a_2 \dots a_d$ of $\{1, 2, \dots, d\}$ that is order isomorphic to s , that is, whose letters appear in the same order of size as those of s .

For example, the standard form of both 4731 and 6842 is 3421, so all three are order isomorphic. As another example, in the permutation 435261 the subwords 352 and 251 are both occurrences of the consecutive pattern 231 (and there are no others).

Definition 1.2 Suppose σ occurs in $\tau = a_1 a_2 \dots a_n$. If a_{i+1} is the leftmost letter of τ involved in any occurrence of σ in τ , we say that τ has a left tail of length i with respect to σ . Analogously, τ has a right tail of length j with respect to σ if a_{n-j} is the rightmost letter of τ involved in any occurrence of σ in τ . If it is clear from the context what σ is, we simply talk about left and right tails of τ .

For example, with respect to the pattern 123, the permutation 286134759 has a left tail of length 3, and a right tails of length 2, since all occurrences of 123 belong within the segment 1347.

The following definition is borrowed from the theory of codes.

Definition 1.3 *Given a permutation τ , its prefix (resp. suffix) pattern of length k is the permutation of length k order isomorphic to the prefix (resp. suffix) of τ of length k . In other words, the prefix (resp. suffix) pattern of length k of τ is the unique permutation $\sigma \in S_k$ such that τ has a left (resp. right) tail of length 0 with respect to σ . In case the prefix and suffix patterns of length k of τ coincide, we say that τ has a bifix pattern of length k .*

It is useful to note that in our poset $\mathcal{P}$, each permutation can cover at most two different permutations. Namely, if σ is covered by τ then σ clearly must occur as all but the first or all but the last letter of τ (i.e. σ is either the longest proper prefix or the longest proper suffix of τ). Thus, to obtain a permutation covered by τ we can only remove the first or last letter of τ , thus yielding at most two different permutations. Moreover, the permutations obtained by removing the first and the last letter, respectively, from τ are order isomorphic if and only if τ is *monotone*, that is, if τ is either the increasing permutation $123 \dots n$ or the decreasing permutation $n(n-1) \dots 21$. So, for instance, the two permutations covered by 35142 are 2413 and 4132, whereas the permutation 54321 only covers 4321.

In the case where σ occurs precisely once in τ , we show that $\mu(\sigma, \tau)$ depends only on the lengths, a and b , of the two tails of τ . More precisely, $\mu(\sigma, \tau)$ is 1 if $a = b \leq 1$, it is -1 if $a = 0$ and $b = 1$ or vice versa, and 0 otherwise (in which case τ has a tail of length at least 2).

Our main result, Theorem 3.6, deals with intervals $[\sigma, \tau]$ where σ occurs at least twice in τ . This result implies that, as in the case of one occurrence, if τ has a tail of length at least 2, then $\mu(\sigma, \tau) = 0$. In the remaining cases, where the tails of τ have length at most 1, the main result gives a recursive algorithm for computing $\mu(\sigma, \tau)$, by producing, if possible, an element $\mathcal{C}$ in $[\sigma, \tau]$, where $|\mathcal{C}| < |\tau| - 2$, such that $\mu(\sigma, \tau) = \mu(\sigma, \mathcal{C})$. This element $\mathcal{C}$, if it exists, must be a bifix pattern of τ , and it must lie below the two elements covered by τ , but not below the element obtained by deleting one letter from each end of τ . If no such element $\mathcal{C}$ exists (which is most often the case), we have $\mu(\sigma, \tau) = 0$.

2 The case of one occurrence

First we consider the case when σ occurs precisely once in τ . In this case the interval $[\sigma, \tau]$ can be described very simply. Let $w_1 \bar{\sigma} w_2$ be the factorization of τ , where the entries of $\bar{\sigma}$ constitute the only occurrence of σ in τ and $|w_1| = a$, $|w_2| = b$, so τ has left and right tails of lengths a and b , respectively. We obtain the permutations covered by τ by deleting the first or the last entry of τ and renaming the remaining entries. Starting from τ and deleting elements from both tails in such a way, we are eventually left with σ . It

is easy to see that the interval (σ, τ) is a *grid*, that is, a direct product of two chains of lengths a and b , respectively. See Figure 2 for an example. It is also well known that the Möbius function of (σ, τ) for such a grid is 0, except when a and b are both at most 1. This is recorded in the following theorem.




Figure 2: The interval $[123, 68513427]$

Theorem 2.1 *Suppose σ occurs precisely once in τ , and that τ has tails of lengths a and b , respectively. Then $\mu(\sigma, \tau) = 1$ if $a = b = 0$ or $a = b = 1$, and $\mu(\sigma, \tau) = -1$ if $a = 0, b = 1$, or $a = 1, b = 0$. Otherwise, $\mu(\sigma, \tau)$ is 0.*

3 More than one occurrence

We will frequently refer to certain special elements of an interval $[\sigma, \tau]$ described in the following Definition.

Definition 3.1 *Given a permutation τ , we let τ' be the standard form of τ after having removed its first letter, τ'' be the standard form of τ after having removed its last letter, and τ''' be the standard form of τ after having removed both its first and last letter. We refer to τ''' as the interior of τ .*

Thus, for instance, if $\tau = 68513427$, then $\tau' = 7513426$, $\tau'' = 6751342$ and $\tau''' = 651342$.

Lemma 3.2 *Suppose σ occurs in τ and that $r = |\tau| - |\sigma| \leq 2$. Then, if $r = 0$, $\mu(\sigma, \tau) = 1$. If $r = 1$, we have $\mu(\sigma, \tau) = -1$. If $r = 2$, then $\mu(\sigma, \tau) = 1$ if $\sigma = \tau'''$ or σ is the longest prefix of τ , otherwise $\mu(\sigma, \tau) = 0$.*

Proof. If $r < 2$, the claim follows directly from the definition of the Möbius function, since $[\sigma, \tau]$ is either a singleton or a chain of two elements. If $r = 2$, and τ is monotone, removing a letter from either end of τ yields the same (monotone) permutation, so $[\sigma, \tau]$ is a chain of three elements, and $\mu(\sigma, \tau) = 0$. If $r = 2$ and $\sigma = \tau'$ or σ is the longest bifix of τ , then $[\sigma, \tau]$ is a direct product of two chains of length 1, whose Möbius function is 1. Otherwise σ only appears at one end of τ , and therefore $[\sigma, \tau]$ is a chain of three elements, whose Möbius function is 0. $\square$

Because of Lemma 3.2, we will from now on only consider intervals $[\sigma, \tau]$ of rank at least three, that is, where $|\tau| - |\sigma| \geq 3$. We also only need to consider pairs (σ, τ) such that σ occurs at least twice in τ , since the single-occurrence case is taken care of in Section 2.

Lemma 3.3 *Suppose σ occurs at least twice in τ , that $|\tau| - |\sigma| \geq 3$, and that τ' does not lie in $[\sigma, \tau]$. Then $\mu(\sigma, \tau) = 1$.*

Proof. Observe first that τ cannot be monotone, since then it would be possible to remove one letter from each end, given that $|\tau| - |\sigma| \geq 3$, and thus τ' would lie in $[\sigma, \tau]$, contrary to assumption.

Together with the hypothesis this implies that τ must have precisely two occurrences of σ , which necessarily appear at its two ends. Therefore $[\sigma, \tau]$ consists of two chains (of equal lengths) having in common only the minimum (σ) and the maximum (τ). It is easy to see that the Möbius function of such an interval is 1. $\square$

In the lemmas above, and in Section 2, we have taken care of all intervals except those where τ has at least two occurrences of σ , $|\tau| - |\sigma| \geq 3$ and the interior of τ lies in $[\sigma, \tau]$. We now deal with these remaining cases.

Lemma 3.4 *Given σ and τ in $\mathcal{P}$, let*

$$'C = \{\rho \in [\sigma, \tau] \mid \rho < \tau, \rho \not\leq \tau'\}$$

and let

$$C' = \{\rho \in [\sigma, \tau] \mid \rho < \tau', \rho \not\leq \tau'\}.$$

The sets $'C$ and C' are chains, and $'C \cap C'$ has at most one element. Moreover, if $z \in C' \setminus 'C$ or $z \in 'C \setminus C'$ then $[z, \tau]$ is a chain.

Proof. A permutation ρ in $'C$ cannot occur in the interior of τ (since $\rho \not\leq \tau'$), and thus it has to be a suffix pattern of τ (since $\rho < \tau$). This implies that $'C$ is a chain, since two suffix patterns of τ must be comparable in $\mathcal{P}$. The argument for C' is analogous, so all permutations in C' occur as prefix patterns of τ .

Suppose that $x, y \in 'C \cap C'$ and $x \neq y$. We can assume, without loss of generality, that $y < x$, since x and y are elements of the chain C' . Then y must be a proper prefix pattern of x , since both belong to C' and thus are prefix patterns of τ . Likewise, since both belong to $'C$, y must be a proper suffix pattern of x . But, an element that occurs as a proper suffix of a prefix of τ and as a proper prefix of a suffix of τ must occur in

the interior of τ , and thus we must have $y \leq \tau'$, which contradicts the assumption that $y \in C'$. So, $'C \cap C'$ can contain at most one element.

If y is τ or τ' $[y, \tau]$ is clearly a chain. If y is neither of these elements

Assume now that $z \in 'C \setminus C'$. Let $y \in [z, \tau]$. If y is τ or τ' $[y, \tau]$ is clearly a chain. If y is neither of these elements then y cannot lie below either τ' or τ' , since then we would have $z < \tau'$ (because $'C$ is a chain), contradicting the hypothesis $z \in 'C$. Hence, for all $y \in [z, \tau]$, we have $y \in 'C \cup \{\tau, \tau'\}$, which is a chain. Then $[z, \tau] \subseteq 'C \cup \{\tau, \tau'\}$, whence $[z, \tau]$ is a chain. An analogous argument shows that $[z, \tau]$ is a chain if $z \in C' \setminus 'C$. $\square$

We have thus shown that $'C \cap C'$ has at most one element. In case it exists, we give it a special name, and record its properties in the following lemma, which is an immediate consequence of the proof of Lemma 3.4.

Lemma 3.5 (and Definition) *Given σ and τ in $\mathcal{P}$, let $'C$ and C' be as defined in Lemma 3.4. If $'C \cap C'$ is nonempty, let $\mathcal{C}$ be its (necessarily unique) element. Then $\mathcal{C}$ has the following properties:*

1. $\mathcal{C} < \tau$ and $\mathcal{C} < \tau'$,
2. $\mathcal{C} \not\leq \tau'$.

Conversely, if $\mathcal{C}$ satisfies the above two conditions, then $\mathcal{C} \in 'C \cap C'$. In what follows, we refer to $\mathcal{C}$ as the carrier element of $[\sigma, \tau]$, or simply as $\mathcal{C}$, if it is clear from the context what σ and τ are.

Here are two examples of the $\mathcal{C}$ associated to an interval: For $\sigma = 321$ and $\tau = 431825976$, we have $\mathcal{C} = 321 = \sigma$, whereas if $\sigma = 231$ and $\tau = 2\ 5\ 7\ 1\ 4\ 8\ 9\ 3\ 6\ 10$ we have $\mathcal{C} = 245136$. Observe that in the latter case the initial and final segments of τ order isomorphic to $\mathcal{C}$ overlap in the letters 48.

Observe that $[\sigma, \tau]$ can be expressed as the disjoint union of the principal ideal generated by τ' and the filter $'C \cup C' \cup (\{\tau, \tau', \tau\} \cap [\sigma, \tau])$. In case $\mathcal{C}$ exists, such a filter is the principal filter generated by $\mathcal{C}$. This fact holds in general, that is, even when the interval $[\sigma, \tau]$ has rank $r \leq 2$.

Theorem 3.6 *Suppose τ has at least two occurrences of σ and that $|\tau| - |\sigma| \geq 3$. Assume that τ' lies in $[\sigma, \tau]$. Then, if $[\sigma, \tau]$ has no carrier element, we have $\mu(\sigma, \tau) = 0$. Otherwise, $\mu(\sigma, \tau) = \mu(\sigma, \mathcal{C})$.*

Proof. In our argument we will compute the Möbius function of $[\sigma, \tau]$ from top to bottom, using the formula for the Möbius function in Equation (2). We will thus compute $\mu(\sigma, \tau)$ by recursively computing the value of $\mu(z, \tau)$ for z of decreasing ranks, starting with $z = \tau$, as is done on the right hand side in Figure 1.

Note that, since $|\tau| - |\sigma| \geq 3$, we have that σ lies strictly below all of τ, τ', τ, τ' , and σ is in particular not equal to any one of them.

If $z \in C' \setminus 'C$ then, by Lemma 3.4, the interval $[z, \tau]$ is a chain. Since $z < \tau' < \tau$, this chain has rank at least 2. Thus, $\mu(z, \tau) = 0$. An analogous argument shows that

$\mu(z, \tau) = 0$ when $z \in {}'C \setminus C'$. Thus, every element t in $C' \cup {}'C$, except $\mathcal{C}$ (if it exists), has $\mu(t, \tau) = 0$.

If $\mathcal{C}$ does not exist, then we claim that $\mu(y, \tau) = 0$ whenever $y < {}'\tau'$, that is whenever y is different from ${}'\tau'$ but belongs to the principal ideal generated by ${}'\tau'$. We claim that a maximal such element y (which must be covered by ${}'\tau'$) lies below precisely four elements t with $\mu(t, \tau) \neq 0$, namely $\tau, \tau', {}'\tau, {}'\tau'$, and hence has $\mu(y, \tau) = 0$. This is because any other element z that y could lie below must belong to ${}'C \cup C'$, since $z < {}'\tau'$ is impossible by virtue of y being maximal. Thus, $\mu(z, \tau) = 0$, according to the preceding paragraph. By induction, this now also applies to all other elements lying strictly below ${}'\tau'$, since each of them lies below precisely four elements t with $\mu(t, \tau) \neq 0$, namely, ${}'\tau', {}'\tau, \tau'$ and τ . In particular, this shows that $\mu(\sigma, \tau) = 0$ if $[\sigma, \tau]$ has no carrier element, and, in fact, that then $\mu(z, \tau) = 0$ for every $z \in [\sigma, \tau]$ except for $\tau, \tau', {}'\tau$ and ${}'\tau'$.

Finally, assume that $\mathcal{C}$ exists. We claim that $\mu(\sigma, \tau) = \mu(\sigma, \mathcal{C})$. Indeed, we have

$$\mu(\sigma, \tau) = - \sum_{\sigma < z \leq \tau} \mu(z, \tau).$$

If $z \not\leq \mathcal{C}$, we have seen above that $\mu(z, \tau) = 0$, whence

$$\mu(\sigma, \tau) = - \sum_{\sigma < z \leq \mathcal{C}} \mu(z, \tau). \quad (4)$$

Next we prove that $\mu(z, \tau) = \mu(z, \mathcal{C})$ whenever $z \leq \mathcal{C}$. We proceed by induction on the difference between the rank of $\mathcal{C}$ and the rank of z . If $z = \mathcal{C}$, then $\mu(z, \tau) = 1$, since the only elements between $\mathcal{C}$ and τ having nonzero Möbius function are $\tau, {}'\tau$ and τ' . This is the base case of the induction. For any $z < \mathcal{C}$, using the induction hypothesis on the elements between z and $\mathcal{C}$ and strictly above z , we have

$$\mu(z, \tau) = - \sum_{z < t \leq \tau} \mu(t, \tau) = - \sum_{z < t \leq \mathcal{C}} \mu(t, \tau) = - \sum_{z < t \leq \mathcal{C}} \mu(t, \mathcal{C}) = \mu(z, \mathcal{C}).$$

Plugging this into formula (4) we get

$$\mu(\sigma, \tau) = - \sum_{\sigma < z \leq \mathcal{C}} \mu(z, \mathcal{C}) = \mu(\sigma, \mathcal{C}),$$

as desired. $\square$

We thus get a recursion, where we find the carrier element of the interval $[\sigma, \mathcal{C}]$, and iterate this until we get an interval that does not have a carrier element. That final interval either has Möbius function zero, or else its rank is at most 2, making it trivial to compute the Möbius function.

Recall our previous examples of the $\mathcal{C}$ associated to an interval: For $\sigma = 321$ and $\tau = 431825976$, we have $\mathcal{C} = 321 = \sigma$, so $\mu(\sigma, \tau) = \mu(\sigma, \sigma) = 1$, whereas if $\sigma = 231$ and $\tau = 2\ 5\ 7\ 1\ 4\ 8\ 9\ 3\ 6\ 10$ we have $\mathcal{C} = 245136$ and $\mu(\sigma, \tau) = \mu(\sigma, \mathcal{C}) = 0$ since the interval $[231, 245136]$ has no carrier element.

Concerning time complexity, a rough analysis shows that, in the worst case, the recursive procedure described by the above theorem is essentially bounded above by $|\tau|^3$. Indeed, supposing that $|\sigma| = k$ and $|\tau| = n$, in order to compute the possible carrier element of $[\sigma, \tau]$, one first tries with the prefix of τ of length $n - 2$, and has to check both that it is isomorphic to the suffix of τ of length $n - 2$ and that it is not isomorphic to τ' . The time needed for this task is proportional to $(n - 2)^2$. In case no carrier element has been found, one has to try with the prefix of length $n - 3$; by an analogous argument, the time needed is proportional to $(n - 3)^2$. In the worst case, (that is, if $[\sigma, \tau]$ has no $\mathcal{C}$ and σ occurs at both ends of τ), this task has to be performed until we get to σ (i.e. the prefix of τ of length k). So the total time needed in this worst case is proportional to $\sum_{i=k}^{n-2} i^2$, which is proportional to n^3 .

Observe, however, that the worst case is not the most significant one. Indeed, our results imply that it is difficult for the Möbius function of $[\sigma, \tau]$ to be nonzero when σ is very long. To be more precise, when τ has a tail with respect to σ of length at least 2, then the Möbius function of $[\sigma, \tau]$ is equal to 0. Thus if σ is not order isomorphic to four specific subwords of τ (this is a quadratic test), then the Möbius function is 0. Moreover, observe that the probability that σ appears at the beginning of τ , for instance, is equal to $1/k!$ (since we have to choose a k -subset of an n -set, then to arrange its elements in the unique way which produces a word order-isomorphic to σ , finally we can permute the remaining $n - k$ elements as we like). So, when the length of the pattern σ increases, the probability to have σ at the beginning of τ (as well as in any other specific position) rapidly decreases.

In what follows we will frequently find in an interval $[\sigma, \tau]$ the carrier element $\mathcal{C}$, then the carrier element of $[\sigma, \mathcal{C}]$, and so on until we have come to the last carrier element $\mathcal{C}'$ in this sequence (which implies that $[\sigma, \mathcal{C}']$ does not have a carrier element). In this case, we will refer to the last carrier element $\mathcal{C}'$ as the *socle* of $[\sigma, \tau]$.

Corollary 3.7 *Suppose σ occurs in τ but that the first two (or the last two) letters of τ are not involved in any occurrence of σ . Then $\mu(\sigma, \tau) = 0$.*

Proof. If σ occurs only once in τ the claim has already been established, in Theorem 2.1. Assume then that σ occurs at least twice in τ . Then $|\tau| - |\sigma| \geq 3$, since two occurrences occupy at least one more letter of σ than the length of τ , in addition to the tail of length at least two assumed by the hypotheses. According to Theorem 3.6, if $[\sigma, \tau]$ has no carrier element $\mathcal{C}$ satisfying the hypotheses of Theorem 3.6, then $\mu(\sigma, \tau) = 0$. Otherwise, if there is such a $\mathcal{C}$, then, as a consequence of its definition, $\mathcal{C}$ is order isomorphic to an initial segment of τ , and to a final segment of τ . Thus, there is no occurrence of σ in $\mathcal{C}$ involving the first two (or the last two) letters of $\mathcal{C}$. Iterate now the construction of $\mathcal{C}$ for $[\sigma, \mathcal{C}]$ until we obtain the socle $\mathcal{C}'$ of $[\sigma, \tau]$. Since $\mathcal{C}'$ has a tail of size at least two with respect to σ we have $|\mathcal{C}'| - |\sigma| \geq 2$ and there are two possibilities: if $|\mathcal{C}'| - |\sigma| = 2$ it follows that $[\sigma, \mathcal{C}']$ is a three element chain, so $\mu(\sigma, \mathcal{C}') = 0$. Otherwise, $|\mathcal{C}'| - |\sigma| \geq 3$, in which case Theorem 3.6 implies that $\mu(\sigma, \mathcal{C}') = 0$. $\square$

Observe that if $\mathcal{C}$ exists then, since it lies in both τC and $C'\tau$, $\mathcal{C}$ must be a bifix pattern of τ of length at least $a + |\sigma| + b$, where a and b are the lengths of the left and right tails

(respectively) of τ with respect to σ . Also, if $\mathcal{C}$ exists, it does not appear anywhere else in τ (this just follows from the definition of $\mathcal{C}$).

As a consequence of Corollary 3.7, nonzero values of $\mu(\sigma, \tau)$ can occur only when the tails of τ have length at most 1. Below we will always assume that τ satisfies this hypothesis, and we will give a series of partial conditions that facilitate the computation of the Möbius function in several cases.

Denote with x the sum of the lengths of the tails of τ with respect to σ . So $x = 0$ means that τ has two occurrences of σ , one at each end; $x = 1$ means that τ has one occurrence of σ at one end and a tail of length 1 at the other end; finally, $x = 2$ means that τ has two tails of length 1 each.

For simplicity, in what follows we will always assume that, when $x = 1$, τ has an occurrence of σ at its right end (and thus a tail of length 1 at its left end). In case σ appears at the left end of τ , we simply have to replace each occurrence of the word “suffix” with the word “prefix” in all the following propositions.

Our first result says that it is very difficult for the Möbius function to take the value $(-1)^{x+1}$.

Proposition 3.8 1. Suppose $x = 0$. If τ does not have a monotone bifix of length $|\sigma| + 1$, then $\mu(\sigma, \tau) \neq -1$.

2. Suppose $x = 1$. If τ does not have a bifix of length $|\sigma| + 2$ whose suffix of length $|\sigma| + 1$ is monotone, then $\mu(\sigma, \tau) \neq 1$.

3. Suppose $x = 2$. Then $\mu(\sigma, \tau) \neq -1$.

Proof. Observe that, in general, any permutation of length ℓ having a bifix of length $\ell - 1$ is necessarily monotone. Suppose $x = 0$. If $\mu(\sigma, \tau) = -1$, then (by Theorem 3.6 and Lemma 3.2) if $[\sigma, \tau]$ has a carrier element, then the socle of $[\sigma, \tau]$ must have length $|\sigma| + 1$, so it is necessarily monotone (since σ is a bifix of it). Suppose $x = 1$. If $\mu(\sigma, \tau) = 1$, then (by virtue of Theorem 3.6 as well) the socle must have length $|\sigma| + 2$, and a direct inspection shows that its suffix of length $|\sigma| + 1$ has to be monotone (since σ is a bifix of it). Suppose $x = 2$. Then necessarily the socle (if it exists) has length at least $|\sigma| + 2$, and so $\mu(\sigma, \tau) \neq -1$. $\square$

In this direction, a general result that includes all the cases of (but is weaker than) the previous proposition is the following. The proof is easy and is omitted.

Proposition 3.9 If τ has a non-monotone suffix of length $|\sigma| + x$, then $\mu(\sigma, \tau) \neq (-1)^{x+1}$.

Next we give an easy necessary condition in order to have $\mu(\sigma, \tau) = 0$. For this, we first need a definition. A permutation is said to be *monotone (reverse) alternating* when it is (reverse) alternating and the two permutations induced by its even-indexed elements and odd-indexed elements are both monotone. For instance, the permutation 342516 is monotone alternating. Alternating permutations have been extensively studied also in recent years, see for instance the survey [S]. In the next proposition we will not consider intervals of rank less than 3, since they are already covered by Lemma 3.2.

Proposition 3.10 Suppose $|\tau| - |\sigma| \geq 3$.

1. Suppose $x = 0$ and that σ is not the socle of $[\sigma, \tau]$. If τ has neither a monotone bifix of length $|\sigma| + 1$, nor a monotone (reverse) alternating bifix of length $|\sigma| + 2$, then $\mu(\sigma, \tau) = 0$.
2. Suppose $x = 1$. If τ has neither a bifix of length $|\sigma| + 1$ nor a bifix of length $|\sigma| + 2$ whose suffix of length $|\sigma| + 1$ is monotone, then $\mu(\sigma, \tau) = 0$.
3. Suppose $x = 2$. If τ does not have a bifix of length $|\sigma| + 2$, then $\mu(\sigma, \tau) = 0$.

Proof. We first show that that if $\mu(\sigma, \tau) \neq 0$ then $[\sigma, \tau]$ necessarily has a carrier element. Namely, if $[\sigma, \tau]$ contains τ' then, by Theorem 3.6, $[\sigma, \tau]$ has a carrier element. If $[\sigma, \tau]$ doesn't contain τ' then τ has precisely two occurrences of σ , one at each end, in which case σ is the carrier element of $[\sigma, \tau]$. Let $\mathcal{C}'$ be the socle of $[\sigma, \tau]$. It is clear that, in order to have $\mu(\sigma, \tau) \neq 0$, $\mathcal{C}'$ must be “short”, and more precisely $|\mathcal{C}'| \leq |\sigma| + 2$.

If $x = 0$, we have two distinct cases (since, by hypothesis, $|\mathcal{C}'| \neq |\sigma|$). If $|\mathcal{C}'| = |\sigma| + 1$, then σ is a bifix of $\mathcal{C}'$, and so necessarily $\mathcal{C}'$ is a monotone bifix of τ . Assume now that $|\mathcal{C}'| = |\sigma| + 2$. Of course, in this case, σ is again a bifix of $\mathcal{C}'$. Suppose $\sigma = a_1 \cdots a_k$. Our hypothesis implies that, for all i , (i) $a_{2i} < a_{2i+1}$ if and only if $a_{2i+2} < a_{2i+3}$, and (ii) $a_{2i+1} < a_{2i+2}$ if and only if $a_{2i+3} < a_{2i+4}$. Therefore the descents of σ are completely determined by the first three elements a_1, a_2 and a_3 . Moreover, observe that, if we had either $a_1 < a_2 < a_3$ or $a_1 > a_2 > a_3$, then $\mathcal{C}'$ (and a fortiori σ) would be monotone, and this contradicts the fact that $[\sigma, \mathcal{C}']$ has no carrier element. Thus we must have either $a_1 < a_2 > a_3$ or $a_1 > a_2 < a_3$, which implies that $\mathcal{C}'$ is alternating. Finally, there cannot exist any i such that $a_{2i} < a_{2i+2}$ and $a_{2i+2} > a_{2i+4}$, since the word $a_1 \cdots a_{k-2}$ is order isomorphic to the word $a_3 \cdots a_k$. The remaining cases can be treated analogously.

If $x = 1$, we have only two possibilities. If $|\mathcal{C}'| = |\sigma| + 1$, then $\mathcal{C}'$ is clearly a bifix of τ of length $|\sigma| + 1$. If $|\mathcal{C}'| = |\sigma| + 2$, then the proof of Proposition 3.8 implies that $\mathcal{C}'$ is a bifix of τ whose suffix of length $|\sigma| + 1$ is monotone.

If $x = 2$, then necessarily $|\mathcal{C}'| = |\sigma| + 2$, and $\mathcal{C}'$ is of course a bifix of τ . □

A general result in this direction is the following.

Proposition 3.11 Denote with ω the longest bifix of τ containing σ and having length $\leq |\sigma| + 2$. Moreover, denote with α (resp. β) the shortest prefix (resp. suffix) pattern of τ containing the first (resp. last) two occurrences of ω . If ω exists and $\alpha \neq \beta$, then $\mu(\sigma, \tau) = 0$.

Proof.(Sketch) Using an argument similar to that of the previous proposition, suppose that $\mu(\sigma, \tau) \neq 0$, and let $\mathcal{C}''$ be the element of $[\sigma, \tau]$ such that the socle of $[\sigma, \tau]$ is the carrier element of the interval $[\sigma, \mathcal{C}'']$. (Thus, $\mathcal{C}''$ is the penultimate carrier element found in the iteration culminating in the socle of $[\sigma, \tau]$). Clearly $\mathcal{C}''$ has to be a bifix of τ , and it is easy to see that $\mathcal{C}'' = \alpha = \beta$. □

Corollary 3.7 and Propositions 3.8 through 3.11 make significantly more efficient the computation of $\mu(\sigma, \tau)$ afforded by Theorem 3.6, since they give the Möbius function of the bulk of all intervals in negligible time.

Acknowledgment

We are very grateful to an anonymous referee for an extremely careful reading and many useful suggestions, which led to the correction of a few errors and a significant improvement of the presentation.

References

- [BS] E. Babson and E. Steingrímsson: Generalized permutation patterns and a classification of the Mahonian statistics, *Sém. Lothar. Combin.* **44** (2000), Article B44b, 18 pp.
- [Bj] A. Björner: The Möbius function of factor order, *Theoret. Comp. Sci.* **117** (1993), 91–98.
- [BJJS] A. Burstein, V. Jelínek, E. Jelínková and E. Steingrímsson: The Möbius function of separable and decomposable permutations, arXiv:1102.1611v1.
- [SV] B. E. Sagan and V. Vatter: The Möbius function of a composition poset, *J. Algebraic Combin.* **24** (2006), 117–136.
- [ST] E. Steingrímsson and B. E. Tenner: The Möbius function of the permutation pattern poset, *J. Combinatorics* **1** (2010), 39–52.
- [EN] S. Elizalde and M. Noy: Consecutive patterns in permutations, *Adv. Appl. Math.* **30** (2003), 110–125.
- [St] R. P. Stanley: *Enumerative Combinatorics 1*, Cambridge University Press, New York, 1997.
- [S] R. P. Stanley: A survey of alternating permutations, *Contemp. Math.*, to appear.