



The Möbius function of generalized subword order

Peter R. W. Mcnamara, Bruce E. Sagan

► To cite this version:

Peter R. W. Mcnamara, Bruce E. Sagan. The Möbius function of generalized subword order. 24th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2012), 2012, Nagoya, Japan. pp.21-32, 10.46298/dmtcs.3016 . hal-01283149

HAL Id: hal-01283149

<https://hal.science/hal-01283149>

Submitted on 5 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Möbius function of generalized subword order

Peter R. W. McNamara¹ and Bruce E. Sagan²

¹Department of Mathematics, Bucknell University, Lewisburg, PA, USA

²Department of Mathematics, Michigan State University, East Lansing, MI, USA

Abstract. Let P be a poset and let P^* be the set of all finite length words over P . Generalized subword order is the partial order on P^* obtained by letting $u \leq w$ if and only if there is a subword u' of w having the same length as u such that each element of u is less than or equal to the corresponding element of u' in the partial order on P . Classical subword order arises when P is an antichain, while letting P be a chain gives an order on compositions. For any finite poset P , we give a simple formula for the Möbius function of P^* in terms of the Möbius function of P . This permits us to rederive in an easy and uniform manner previous results of Björner, Sagan and Vatter, and Tomie. We are also able to determine the homotopy type of all intervals in P^* for any finite P of rank at most 1.

Résumé. Soit P un ensemble partiellement ordonné et soit P^* l'ensemble des mots de longueur finie sur P . On définit l'ordre des sous-mots généralisé comme l'ordre partiel sur P^* obtenu en posant $u \leq w$ s'il existe un sous-mot u' de w ayant la même longueur que u , tel que chaque élément de u soit plus petit ou égal à l'élément correspondant de u' dans l'ordre partiel sur P . L'ordre des sous-mots classique correspond au cas où P est une antichaîne ; tandis que si P est une chaîne, on obtient un ordre sur les compositions. Pour tout ensemble partiellement ordonné fini P , nous donnons une formule simple pour la fonction de Möbius de P^* en fonction de celle de P . Cela nous permet de retrouver de manière simple et uniforme des résultats de Björner, Sagan et Vatter, et de Tomie. Nous sommes aussi en mesure de déterminer le type d'homotopie de tous les intervalles de P^* pour n'importe quel P fini de rang au plus 1.

Keywords: Chebyshev polynomial, discrete Morse theory, minimal skipped interval, Möbius function, poset, subword order

1 Introduction

Let A (the *alphabet*) be any set and let A^* be the *Kleene closure* of all finite length words over A , so


$$A^* = \{w = w(1)w(2) \dots w(n) : w(i) \in A \text{ for all } i, \text{ and } n \geq 0\}.$$

We denote the *length* or *cardinality* of w by $|w|$. A *subword* of $w \in A^*$ is a word $u = w(i_1)w(i_2) \dots w(i_k)$ where $i_1 < i_2 < \dots < i_k$. (Note that the elements chosen from w need not be consecutive.) *Subword order* on A^* is defined by letting $u \leq w$ if and only if u is a subword of w . Björner [Bjö90] was the first person to determine the Möbius function of subword order.

Now consider the symmetric group \mathfrak{S}_n of all permutations of $\{1, 2, \dots, n\}$. If $\sigma = \sigma(1)\sigma(2) \dots \sigma(n) \in \mathfrak{S}_n$ and $\pi = \pi(1)\pi(2) \dots \pi(k) \in \mathfrak{S}_k$ then σ *contains a copy of π as a pattern* if there is a subword $\sigma(i_1)\sigma(i_2) \dots \sigma(i_k)$ such that

$$\pi(r) < \pi(s) \iff \sigma(i_r) < \sigma(i_s)$$

⁰ Emails: peter.mcnamara@bucknell.edu, sagan@math.msu.edu

Fig. 1: The poset Λ

for all $1 \leq r < s \leq k$. The *pattern order* on $\mathfrak{S} = \uplus_{n \geq 0} \mathfrak{S}_n$ is obtained by letting $\pi \leq \sigma$ if and only if σ contains a copy of π . For example, $2143 \leq 321465$ because of the subwords 3265 , 3165 or 2165 . Wilf [Wil02] posed the problem of determining the Möbius function of pattern order. The first result along these lines was obtained by Sagan and Vatter [SV06] and this will be discussed in more detail below. Later work has been done by Steingrímsson and Tenner [ST10] and by Burstein, Jelínek, Jelínková and Steingrímsson [BJJS11]. It remains an open problem to fully answer Wilf's question.

When trying to prove results about pattern containment, it is often instructive to consider the case of *layered* permutations, which are those of the form

$$\pi = a, a-1, \dots, 1, a+b, a+b-1, \dots, a+1, a+b+c, a+b+c-1, \dots$$

for some positive integers a, b, c, \dots . Note that a layered permutation is completely specified by the composition (a, b, c, \dots) of *layer lengths*, and that pattern order on layered permutations is isomorphic to the following order on compositions: for compositions $a = (a_1, a_2, \dots, a_r)$ and $b = (b_1, b_2, \dots, b_s)$, we say that $a \leq b$ if there exists a subsequence $(b_{i_1}, b_{i_2}, \dots, b_{i_r})$ of b such that $a_j \leq b_{i_j}$ for $1 \leq j \leq r$. Our example $2143 \leq 321465$ above for layered permutations corresponds to $22 \leq 312$ for compositions.

Sagan and Vatter [SV06] generalized both subword order and pattern order on layered permutations as follows. Letting P be any poset, it is natural to let P^* denote the Kleene closure of the alphabet consisting of the elements of P . Define *generalized subword order* on P^* by letting $u \leq w$ if and only if there is a subword $w(i_1)w(i_2) \dots w(i_k)$ of w of the same length as u such that

$$u(j) \leq_P w(i_j) \text{ for } 1 \leq j \leq k. \quad (1.1)$$

Note that if P is an antichain, then generalized subword order on P^* is the same as ordinary subword order since one can only have $a \leq_P b$ if $a = b$. At the other extreme, if P is the chain \mathbb{P} of positive integers, then, as remarked in the previous paragraph, generalized subword order on the set \mathbb{P}^* of compositions is isomorphic to pattern order on layered permutations. Sagan and Vatter determined the Möbius function of P^* for any rooted forest P , i.e., each component of the Hasse diagram of P is a tree with a unique minimal element. Note that this covers both the antichain and chain cases. They also considered the smallest P which is not a rooted forest, namely the poset Λ given in Figure 1, and conjectured that the Möbius values for certain intervals in Λ^* were given by coefficients of Chebyshev polynomials of the first kind. This conjecture was later proved and the result generalized by Tomie [Tom10] using ad hoc methods. Earlier appearances of generalized subword order in the context of well-quasi-orderings are surveyed in [Kru72].

Our main result is a simple formula for the Möbius function of P^* for any finite poset P , as given in Theorem 1.1. To state the theorem, we need to introduce some key notation and terminology. We let P_0 denote the poset P with a bottom element 0 adjoined, and let \leq_0 and μ_0 denote the order relation and Möbius function of P_0 respectively. An *expansion* of $u \in P^*$ is $\eta \in (P_0)^*$ such that the restriction of η

to its nonzero elements is u . For example, 0110302 is an expansion of $u = 1132$. An *embedding* of u in w is an expansion η of u having length $|w|$ such that $\eta(j) \leq_0 w(j)$ for all j . Continuing our example and using the poset in Figure 1, we see that the given expansion can be considered as an embedding of u in $w = 2132333$. It should be clear from the definitions that there is an embedding of u in w if and only if $u \leq w$ in P^* . Since the Möbius function of P^* is our principal object of interest, we abbreviate μ_{P^*} by μ . With these fundamentals in place, we can now state our main result.




Fig. 2: The interval $[11, 333]$ of P^* in the case where P is as shown in Figure 1

Theorem 1.1 *Let P be a poset such that P_0 is locally finite. Let u and w be elements of P^* with $u \leq w$. Then*

$$\mu(u, w) = \sum_{\eta} \prod_{1 \leq j \leq |w|} \begin{cases} \mu_0(\eta(j), w(j)) + 1 & \text{if } \eta(j) = 0 \text{ and } w(j-1) = w(j), \\ \mu_0(\eta(j), w(j)) & \text{otherwise,} \end{cases}$$

where the sum is over all embeddings η of u in w .

If $j = 1$, the condition that $w(j-1) = w(j)$ is considered false since $w(j-1)$ does not exist. The power of Theorem 1.1 is that it allows us to determine the Möbius function in P^* just by knowing the

Möbius function in P_0 ; typically, P^* is a much more complicated poset than P_0 , as in Example 1.2 below. We also note that it is natural that a formula for μ involve μ_0 , since if w consists of one letter, then $\mu(u, w) = \mu_0(u, w)$ when $u \leq w$.

Example 1.2 Let $P = \Lambda$ as shown in Figure 1, and consider $\mu(11, 333)$. Applying Theorem 1.1, we see that the embedding $\eta = 110$ contributes $(-1)(-1)(1+1) = 2$ to the sum. Similarly, 101 and 011 contribute 2 and 1, respectively. Thus $\mu(11, 333) = 5$, which is not at all obvious from Figure 2. It is easy to generate intervals whose Hasse diagrams are too large and complicated to be shown clearly here, but whose Möbius functions are easy to calculate using Theorem 1.1. One extreme example is that the Hasse diagram of the interval $[\emptyset, 33333]$ in Λ^* has 1904 edges; since the only embedding of the empty word \emptyset in 33333 is 00000, Theorem 1.1 gives $\mu(\emptyset, 33333) = (1)(1+1)^4 = 16$.

We prove Theorem 1.1 by using Babson and Hersh's method [BH05] for applying Forman's discrete version of Morse theory [For95, For98, For02] to order complexes of posets. We will sketch some of the main ideas of [BH05] and of our proof in Section 2. By specializing our result, one can easily derive all the formulas for Möbius functions cited above, which we do in Section 3. Specifically, we derive the following results.

1. Björner's formula for the Möbius function of subword order.
2. Sagan and Vatter's result for the Möbius function of P^* in the case that P is a rooted forest.
3. Their related result for the order on compositions described above, which corresponds to the case when $P = \mathbb{P}$, the positive integers.
4. Tomie's result for the Möbius function of Λ^* . The connection to Chebyshev polynomials $T_n(x)$ of the first kind is that $\mu(1^i, 3^j)$ is the coefficient of x^{j-i} in $T_{i+j}(x)$, for all $0 \leq i \leq j$.
5. Tomie's more general result, which corresponds to letting P consist of an s -element antichain with a top element added.

We can also compute the homotopy type of P^* whenever the rank of P , denoted $\text{rk}(P)$, is at most 1; we show that any interval $[u, w]$ in P^* is homotopic to a wedge of $|\mu(u, w)|$ spheres, all of dimension $\text{rk}(w) - \text{rk}(u) - 2$. As a corollary, we get the corresponding result of Björner [Bjö90] in the antichain case. The final section contains some concluding remarks about related work.

2 Sketch of the proof of the main result

Since we consider the use of discrete Morse theory to be a strength of the proof of Theorem 1.1, we will sketch the proof's main ideas here. The complete proof is given in [MS12], and we refer the reader to [BH05] for the full details of Babson and Hersh's version of discrete Morse theory for order complexes of posets.

To determine the Möbius function of an interval $[u, w]$ in P^* using discrete Morse theory, the main objects to consider are maximal chains

$$C : w = v_0 \longrightarrow v_1 \longrightarrow \cdots \longrightarrow v_n = u \quad (2.1)$$

from w to u , where $v \rightarrow v'$ denotes that v covers v' . Note that we write our chains in $[u, w]$ as going from w down to u . A critical ingredient for discrete Morse theory will be an edge-labeling for this chain, which we now describe. If v covers v' in P^* , then an embedding for v' in v can be obtained by replacing some letter b of v by a letter a covered by b in P_0 . If $a \neq 0$, then $|v'| = |v|$ and this embedding is unique. If $a = 0$, then $|v'| = |v| - 1$ and there may be several embeddings, in which case we choose b to be as far left in v as possible. Therefore, in the chain (2.1), each cover $v_{i-1} \rightarrow v_i$ defines an embedding of v_i in v_{i-1} and thus, inductively, an embedding η_i of v_i in w . We then label the cover $v_{i-1} \rightarrow v_i$ by $l_i = \langle j_i, x_i \rangle$ where j_i is the index where η_{i-1} and η_i differ and $x_i = \eta_i(j_i)$. For example, in the poset in Figure 2, we get label sequences such as

$$\begin{aligned} C : 333 &\xrightarrow{\langle 3,1 \rangle} 331 \xrightarrow{\langle 2,1 \rangle} 311 \xrightarrow{\langle 1,2 \rangle} 211 \xrightarrow{\langle 1,0 \rangle} 011, \\ C' : 333 &\xrightarrow{\langle 3,1 \rangle} 331 \xrightarrow{\langle 1,1 \rangle} 131 \xrightarrow{\langle 2,1 \rangle} 111 \xrightarrow{\langle 1,0 \rangle} 011. \end{aligned} \quad (2.2)$$

(Here we write 011 to highlight the embedding, even though 11 is the relevant element of P^* .) We can then order the chains by considering the edge labels from left to right as necessary and comparing the edge labels lexicographically. In (2.2), $C' < C$ since $\langle 1, 1 \rangle$ lexicographically precedes $\langle 2, 1 \rangle$. This order satisfies the condition of being a *poset lexicographic order* or *PLO*. Although we will not define a PLO here, we will note that it is a required condition for Babson and Hersh's Theorem 2.1 below to apply.

Now that we have an order on the maximal chains in the interval $[u, w]$, we can define the key notions of *skipped interval* and *minimal skipped interval*. A skipped interval of a chain C is an open interval

$$C(v_i, v_j) : v_{i+1} \rightarrow v_{i+2} \rightarrow \cdots \rightarrow v_{j-1}$$

of C such that $C - C(v_i, v_j)$ is contained in C' for some $C' < C$ in the PLO. For example, in (2.2), $311 \rightarrow 211$ is a skipped interval of C because of the given C' . A skipped interval is a *minimal skipped interval* or *MSI* if it is minimal with respect to containment. In our example, $311 \rightarrow 211$ in C is not an MSI but the 1-element interval 211 is, since the maximal chain

$$C'' : 333 \rightarrow 331 \rightarrow 311 \rightarrow 111 \rightarrow 011$$

makes 211 a skipped interval in C . Let $\mathcal{I}(C)$ denote the set of MSIs of a maximal chain C .

The chains that contribute to the Möbius function $\mu(u, w)$ are known as *critical chains*. To describe these chains, we need to convert $\mathcal{I}(C)$ into a set of disjoint intervals $\mathcal{J}(C)$ as follows. Let $\mathcal{I}(C) = \{I_1, I_2, \dots\}$ where the intervals are listed in the order in which they are first encountered along C . Since there are no containments among intervals, there are no ties to be broken. Now let the first interval of $\mathcal{J}(C)$ be $J_1 = I_1$. Consider the intervals $I'_2 = I_2 - J_1, I'_3 = I_3 - J_1, \dots$ obtained by removing J_1 from all the remaining intervals. Discard any of the I'_i which are not containment minimal and let J_2 be the first interval along C which remains. Iterating this process creates $\mathcal{J}(C)$. Call C *critical* if it is covered by $\mathcal{J}(C)$ in the sense that $C(w, u)$ equals the union of the intervals in $\mathcal{J}(C)$. It is important to note that $\mathcal{I}(C)$ may cover C when $\mathcal{J}(C)$ does not.

We define the *critical dimension* of C as

$$d(C) = |\mathcal{J}(C)| - 1.$$

Continuing with our example, it turns out that the chain C in (2.2) is critical since every 1-element interval in $C(333, 22)$ is an MSI. So $\mathcal{I}(C) = \mathcal{J}(C)$ and $d(C) = (-1)^{3-1} = 1$. We now have everything in place

to give Babson and Hersh's technique for computing μ . While we have introduced their result in the context of generalized subword order, their result applies for any poset Q with a PLO, and the definitions of skipped interval, MSI and $d(C)$ are the same as those given above.

Theorem 2.1 ([BH05]) *Let Q be a poset and $x, y \in Q$ such that $[x, y]$ is finite. For any PLO on the maximal chains of $[x, y]$,*

$$\mu(x, y) = \sum_C (-1)^{d(C)},$$

where the sum is over all critical chains C . □

Therefore, to prove Theorem 1.1, we wish to determine the critical chains of $[u, w]$. We start by eliminating most of the maximal chains from consideration with the following lemma. This parallels the approach in [SV06], although overall our case is significantly more involved.

Lemma 2.2 (Decreasing Labels Lemma) *If a maximal chain C of $[u, w]$ is critical, then C has lexicographically decreasing labels from top to bottom.*

The next lemma is equally as helpful, since it allows us to divide the potentially critical chains into more manageable portions. If v_i appears between the labels $\langle j_i, x_i \rangle$ and $\langle j_{i+1}, x_{i+1} \rangle$, then we say v_i is a *1-descent* if $j_i > j_{i+1}$.

Lemma 2.3 (Descent Lemma) *A 1-descent is an MSI of one element.*

As a consequence of Lemmas 2.2 and 2.3, we can restrict our attention to chains of a very special form. Consider a potentially critical chain C from w to an embedding η , and let k denote the largest index such that $w(k) \neq \eta(k)$. Since the labels must be lexicographically decreasing, C must start by reducing $w(k)$ to $\eta(k)$. Then C must reduce $w(j)$ to $\eta(j)$ where $j < k$ is as large as possible with $w(j) \neq \eta(j)$, and so on from right to left. A key observation is that, by Lemma 2.3, in moving from decreasing letters in position k to decreasing letters in position j , we will create a 1-descent in the label sequence and hence a single-element MSI. Because no MSI can contain another, the only MSIs left to determine on potentially critical chains are those of the form $C(v, v')$, where the only change from v to v' is that a single position has been reduced.

The major proposition in our proof is one that classifies such MSIs in P^* in terms of MSIs in P_0 , and this is the primary reason why Theorem 1.1 expresses the Möbius function of P^* in terms of the Möbius function of P_0 . The fact that the letters are reduced one at a time along critical chains from w to u is also the reason why the expression for $\mu(u, w)$ contains a product over j . The $+1$ in the expression for $\mu(u, w)$ arises because there are special considerations when $\eta(j) = 0$, including the possibility of an extra critical chain in P^* corresponding to the lexicographically first chain of P_0 , which is never a critical chain in P_0 itself.

At this stage, one might hope that carefully putting these observations and the corresponding mathematical details together would be enough to classify all critical chains and hence prove Theorem 1.1. One of the subtleties of the proof is that there are certain intervals $[u, w]$ that require special treatment. For elements a, b of the alphabet P with $a \leq_0 b$, if a embeds in ab as $a0$, then our discrete Morse theory techniques seem insufficient. Instead, we are able to use classical Möbius function techniques, such as the expression for the Möbius function as an alternating sum of chains [Sta97, Prop. 3.8.5], to determine the corresponding contribution to $\mu(u, w)$.

3 Applications

In this section we will show how the Möbius function values for subword order, composition order and other special cases of generalized subword order mentioned in the Introduction all follow easily from Theorem 1.1. First, however, we would like to prove a result about the homotopy type of certain P^* .

If P is a finite poset and $x \in P$, then the *rank of x* , denoted $\text{rk}(x)$, is the length of the longest chain from a minimal element of P to x . In particular, minimal elements have rank 0. The *rank of P* is

$$\text{rk}(P) = \max_{x \in P} \text{rk}(x).$$

For example, P is an antichain if and only if $\text{rk}(P) = 0$. For $w \in P^*$, we write $\text{rk}(w)$ to mean the rank of w in the interval $[\emptyset, w]$ of P^* . Note that if P is an antichain, then $\text{rk}(w) = |w|$ for $w \in P^*$.

Now consider the *order complex*, $\Delta(x, y)$, of a finite interval $[x, y]$ in a poset P , which is the abstract simplicial complex consisting of all chains of (x, y) . If $\Delta(x, y)$ has a topological property, we will also say that $[x, y]$ has the same property. To prove Theorem 2.1, Babson and Hersh showed that $\Delta(x, y)$ is homotopic to a CW-complex with a cell for each critical chain and an extra cell of dimension 0. The simplex in a critical chain C giving rise to a critical cell is obtained by taking one element from each of the \mathcal{J} -intervals and so has dimension $d(C)$. This is all the information we need to prove the following result.

Theorem 3.1 *Let P be any finite poset with $\text{rk}(P) \leq 1$. Then any interval $[u, w]$ in P^* is homotopic to a wedge of $|\mu(u, w)|$ spheres all of dimension $\text{rk}(w) - \text{rk}(u) - 2$.*

Proof: We claim that every MSI in a maximal chain of $[u, w]$ consists of one element. Suppose, towards a contradiction, that there is a chain C of the form (2.1) containing an MSI $I = \{v_i, v_{i+1}, \dots, v_k\}$ for some i and $k > i$. By Lemma 2.2, the labels along I must be lexicographically decreasing. And by Lemma 2.3, I cannot contain a 1-descent since otherwise it would contain a smaller SI. So there must be some index j so that only elements in position j are decreased in passing from v_{i-1} to v_{k+1} . But then these elements form a chain of length at least 3 in P_0 . This contradicts the fact that the longest chain in P has length at most 1.

Note that all the maximal chains in $[u, w]$ have length $\text{rk}(w) - \text{rk}(u)$. And since any critical chain C is covered by 1-element MSIs, the number of intervals in $\mathcal{J}(C)$ is always $\text{rk}(w) - \text{rk}(u) - 1$. This also implies that there is no cancellation in the sum of Theorem 2.1. So $\mu(u, w)$ is, up to sign, the number of critical chains. It follows that the CW-complex discussed above must be constructed from a 0-cell together with $|\mu(u, w)|$ cells of dimension $\text{rk}(w) - \text{rk}(u) - 2$. The only way to construct such a complex is as given in the statement of the theorem. \square

We note that Babson and Hersh [BH05] show that if every MSI of an interval $[x, y]$ in a poset is a singleton, then $[x, y]$ is shellable and hence a wedge of spheres. (They assume that the interval is pure, but their proof goes through for non-pure posets.) So our proof above actually shows that $[u, w]$ is shellable whenever $\text{rk}(P) \leq 1$.

3.1 Björner's formula for subword order

Subword order on the alphabet $A = \{1, 2, \dots, s\}$ corresponds to the case when P is an antichain with elements A . Let us determine what Theorem 1.1 yields in this case. For $u, w \in A^*$, suppose η is an

embedding of u in w . For each j with $1 \leq j \leq |w|$, there are two cases. The first is that $\eta(j) = w(j) \neq 0$, in which case $\mu_0(\eta(j), w(j)) = 1$. The more interesting situation is when $\eta(j) = 0$, in which case $\mu_0(\eta(j), w(j)) = -1$. So, if $w(j-1) = w(j)$, then η contributes 0 to the sum in Theorem 1.1. Thus we can restrict to *normal embeddings*, meaning that $\eta(j) \neq 0$ whenever $w(j-1) = w(j)$. For example, if $w = 1122121$ and $u = 121$, then there are exactly two normal embeddings of u in w , namely 0102100 and 0102001. Let us denote the number of normal embeddings of u in w by $\binom{w}{u}_n$. Putting these observations together, we get Björner's result from Theorem 1.1.

Theorem 3.2 ([Bjö90]) *If $u, w \in A^*$, then*

$$\mu(u, w) = (-1)^{|w|-|u|} \binom{w}{u}_n.$$

In the same paper, Björner also derived the homotopy type of $[u, w]$, and this result follows immediately from Theorem 3.1.

Theorem 3.3 ([Bjö90]) *If $u, w \in A^*$, then $[u, w]$ is homotopic to a wedge of $\binom{w}{u}_n$ spheres, all of dimension $|w| - |u| - 2$. \square*

3.2 Generalized subword order for rooted forests

We now consider the generalization of Björner's result to rooted forests given by Sagan and Vatter [SV06]. Clearly, P is a rooted forest if and only if every element $x \in P_0 - \{0\}$ covers exactly one element, denoted x^- , of P_0 . We will show how Theorem 1.1 gives the formula for μ as stated in [SV06]; their statement for the Möbius function of composition order is almost identical and follows immediately.

For P a rooted forest and $u, w \in P^*$, let η be an embedding of u in w . Note that for $x, y \in P_0$,

$$\mu_0(x, y) = \begin{cases} +1 & \text{if } y = x, \\ -1 & \text{if } y^- = x, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, if $\eta(j) \neq 0$, for η to contribute a nonzero amount to the sum in Theorem 1.1, there are two possibilities:

- $w(j) = \eta(j)$, which will contribute a 1 to the product, or
- $w(j)^- = \eta(j)$, which will contribute a -1 .

If $\eta(j) = 0$, there are also two possibilities that will allow η to have a nonzero contribution:

- $w(j)$ is a minimal element of P and $w(j-1) \neq w(j)$, which will contribute a -1 to the product, or
- $w(j)$ is not minimal and $w(j-1) = w(j)$, which will contribute a 1.

These four conditions on w and η can be seen to be equivalent to those on the generalized version of normal embedding as defined in [SV06].

The *defect* $\text{def}(\eta)$ of a normal embedding η of u in w is defined in [SV06] as

$$\text{def}(\eta) = \#\{i : \eta(i) = w(i)^-\}.$$

Referring to Theorem 1.1, we see that the defect is exactly the number of j 's, $1 \leq j \leq |w|$, that will contribute -1 to a nonzero product, while all other j 's in a normal embedding will contribute $+1$. Putting this all together, we get [SV06, Theorem 6.1].

Theorem 3.4 ([SV06]) *Let P be a rooted forest. Then the Möbius function of P^* is given by*

$$\mu(u, w) = \sum_{\eta} (-1)^{\text{def}(\eta)},$$

where the sum is over all normal embeddings η of u in w .

Restricting to the composition poset, which arises when $P = \mathbb{P}$, everything stays the same, except that we can write $x - 1$ in place of x^- for any $x \in P$.

3.3 Connection with Chebyshev polynomials

As promised, a connection between generalized subword order and Chebyshev polynomials follows easily from Theorem 1.1. Consider the poset Λ from Figure 1 and the intervals $[1^i, 3^j]$ in Λ^* . To describe the corresponding Möbius function values, consider the Chebyshev polynomials $T_n(x)$ of the first kind, which can be defined recursively by $T_0(x) = 1$, $T_1(x) = x$, and

$$T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x) \quad (3.1)$$

for $n > 1$. An equivalent definition which is more suitable for our purposes is obtained by replacing (3.1) by

$$T_n(x) = \frac{n}{2} \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \frac{(-1)^k}{n-k} \binom{n-k}{k} (2x)^{n-2k} \quad (3.2)$$

for $n > 1$. One consequence of either definition is that the coefficient of x^m in $T_n(x)$, which we will denote by $\langle x^m \rangle T_n$, is nonzero only if m and n have the same parity. The following result, which was conjectured in [SV06] and first proved in [Tom10], concerns such coefficients.

Theorem 3.5 ([Tom10]) *Considering intervals in Λ^* , for all $0 \leq i \leq j$,*

$$\mu(1^i, 3^j) = \langle x^{j-i} \rangle T_{i+j}(x).$$

Proof: First, we check the result for $i + j = 0$. We must have $i = j = 0$ and $[1^i, 3^j]$ is a single element poset, consistent with $T_0(x) = 1$.

Otherwise, $j \geq 1$ and there are $\binom{j}{i}$ embeddings η of 1^i in 3^j , of which $\binom{j-1}{i}$ satisfy $\eta(1) = 0$ while $\binom{j-1}{i-1}$ satisfy $\eta(1) = 1$ (where binomial coefficients of the form $\binom{n}{k}$ with $k < 0$ are considered 0 as usual). By Theorem 1.1, the former type of embeddings each contribute $(-1)^i 2^{j-i-1}$ to the Möbius function, while the latter type each contribute $(-1)^i 2^{j-i}$. Thus

$$\begin{aligned} \mu(1^i, 3^j) &= (-1)^i 2^{j-i-1} \left(\binom{j-1}{i} + 2 \binom{j-1}{i-1} \right) \\ &= (-1)^i 2^{j-i-1} \left(\binom{j-1}{i-1} + \binom{j}{i} \right) \\ &= (-1)^i 2^{j-i-1} \frac{i+j}{j} \binom{j}{i}. \end{aligned}$$

This last expression is now readily checked to be the coefficient of x^{j-i} when $n = i + j$ in (3.2). \square

Although Tomie did not derive the homotopy type for these intervals in Λ^* , we obtain the information easily from Theorem 3.1.

Theorem 3.6 *For all $0 \leq i \leq j$, the interval $[1^i, 3^j]$ in Λ^* is homotopic to a wedge of $|\langle x^{j-i} \rangle T_{i+j}(x)|$ spheres, all of dimension $2j - i - 2$. \square*

3.4 Tomie's generalized Chebyshev polynomials

The main result of [Tom10] is more general than Theorem 3.5. For $s \geq 1$, Tomie considers the poset, which we denote by Λ_s , that consists of an s -element antichain $\{1, 2, \dots, s\}$ with a top element $s + 1$ added. Letting $s = 2$ gives the poset Λ . Along the same lines, Tomie recursively defines generalized Chebyshev polynomials $T_n^s(x)$, and gives a closed-form expression for the coefficients of $T_n^s(x)$ which, after a change of variables, can be written as

$$T_n^s(x) = \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} (-1)^k s^{n-2k-1} \left(\binom{n-k}{k}_s - \binom{n-k-1}{k} \right) x^{n-2k}. \quad (3.3)$$

for $n \geq 0$ and $s \geq 1$.

The main result of [Tom10] again follows as a special case of Theorem 1.1, as we now show.

Theorem 3.7 ([Tom10]) *Considering intervals in $(\Lambda_s)^*$, for all $0 \leq i \leq j$ and $s \geq 1$,*

$$\mu(1^i, (s+1)^j) = \langle x^{j-i} \rangle T_{i+j}^s(x).$$

Proof: From (3.3), we get that

$$T_{i+j}^s(x) = \sum_{k=0}^{\lfloor \frac{i+j}{2} \rfloor} (-1)^k s^{i+j-2k-1} \left(\binom{i+j-k}{k}_s - \binom{i+j-k-1}{k} \right) x^{i+j-2k}.$$

Considering the term in the sum where $k = i$, we get that the coefficient of x^{j-i} in $T_{i+j}^s(x)$ is

$$(-1)^i s^{j-i-1} \left(\binom{j}{i}_s - \binom{j-1}{i} \right),$$

which equals

$$(-1)^i s^{j-i-1} \left(\binom{j-1}{i-1}_s + \binom{j-1}{i}_s (s-1) \right) \quad (3.4)$$

whenever $j \geq 1$.

Now consider $\mu(1^i, (s+1)^j)$ as determined by Theorem 1.1. When $i + j = 0$, we must have $i = j = 0$ and $[1^i, (s+1)^j]$ is a single element poset, consistent with $T_0^s(x) = 1$ from (3.3). Otherwise, $j \geq 1$ and there are $\binom{j}{i}$ embeddings η of 1^i in $(s+1)^j$, of which $\binom{j-1}{i}$ satisfy $\eta(1) = 0$ while $\binom{j-1}{i-1}$ satisfy $\eta(1) = 1$. By Theorem 1.1, the former type of embeddings each contribute $(-1)^i (s-1) s^{j-i-1}$ to the Möbius function, while the latter type each contribute $(-1)^i s^{j-i}$. Thus $\mu(1^i, (s+1)^j)$ equals the expression (3.4), as required. \square

Using Theorem 3.1 one last time, we obtain the following result.

Theorem 3.8 For all $0 \leq i \leq j$, the interval $[1^i, (s+1)^j]$ in $(\Lambda_s)^*$ is homotopic to a wedge of $|\langle x^{j-i} \rangle T_{i+j}^s(x)|$ spheres, all of dimension $2j - i - 2$. \square

4 Closing Remarks

There has also been interest in generalized factor order on P^* which is defined like generalized subword order except that one requires the indices i_1, i_2, \dots, i_k to be consecutive in (1.1). Björner [Bjö93] found a recursive formula for the Möbius function in the case of an antichain. In particular, he showed that the only possible Möbius values are $0, \pm 1$ and that the order complex of every interval is homotopic to either a ball or a sphere. In his thesis, see [Wil], Willenbring reproved Björner's results in an elucidating way using critical chains and found a more general formula for rooted forests. The latter is quite complicated. It would be very interesting if one could find a simpler formula more along the lines of Theorem 1.1.

The analogue of generalized factor order for \mathfrak{S} is called the consecutive pattern poset. Somewhat surprisingly (given the fact that Wilf's question for ordinary patterns has not been fully answered), Bernini, Ferrari, and Steingrímsson [BFS11] gave a complete description of the Möbius function in the consecutive case. Even more surprisingly, Sagan and Willenbring [SW] were able to give a proof of this result using critical chains which closely parallels the one Willenbring gave for factor order of an antichain. This led them to define a sequence of partial orders on \mathbb{P}^* , denoted P_k for $k = 0, 1, 2, \dots, \infty$, where P_0 is factor order, P_∞ contains consecutive pattern order as a convex subposet, and every P_k has essentially that same Möbius function. So this sequence of interpolating posets gives an explanation of the coincidence noted above.

Acknowledgements

The authors thank Volker Strehl and the anonymous referees for helpful comments. Calculations were performed using John Stembridge's posets package [Ste].

References

- [BFS11] Antonio Bernini, Luca Ferrari, and Einar Steingrímsson. The Möbius function of the consecutive pattern poset. *Electron. J. Combin.*, 18(1):Paper 146, 12 pp., 2011.
- [BH05] Eric Babson and Patricia Hersh. Discrete Morse functions from lexicographic orders. *Trans. Amer. Math. Soc.*, 357(2):509–534 (electronic), 2005.
- [BJS11] Alexander Burstein, Vít Jelínek, Eva Jelínková, and Einar Steingrímsson. The Möbius function of separable and decomposable permutations. *J. Combin. Theory Ser. A*, 118(8):2346–2364, 2011.
- [Bjö90] Anders Björner. The Möbius function of subword order. In *Invariant theory and tableaux (Minneapolis, MN, 1988)*, volume 19 of *IMA Vol. Math. Appl.*, pages 118–124. Springer, New York, 1990.
- [Bjö93] Anders Björner. The Möbius function of factor order. *Theoret. Comput. Sci.*, 117(1-2):91–98, 1993.

- [For95] Robin Forman. A discrete Morse theory for cell complexes. In *Geometry, topology, & physics*, Conf. Proc. Lecture Notes Geom. Topology, IV, pages 112–125. Internat. Press, Cambridge, MA, 1995.
- [For98] Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.
- [For02] Robin Forman. A user’s guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48:Art. B48c, 35 pp., 2002.
- [Kru72] Joseph B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept. *J. Combinatorial Theory Ser. A*, 13:297–305, 1972.
- [MS12] Peter R. W. McNamara and E. Sagan, Bruce. The Möbius function of generalized subword order. *Adv. Math.*, 229(5):2741–2766, 2012.
- [ST10] Einar Steingrímsson and Bridget Eileen Tenner. The Möbius function of the permutation pattern poset. *J. Comb.*, 1(1):39–52, 2010.
- [Sta97] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997. With a foreword by Gian-Carlo Rota. Corrected reprint of the 1986 original.
- [Ste] John R. Stembridge. SF, posets and coxeter/weyl. Available from <http://www.math.lsa.umich.edu/~jrs/maple.html>.
- [SV06] Bruce E. Sagan and Vincent Vatter. The Möbius function of a composition poset. *J. Algebraic Combin.*, 24(2):117–136, 2006.
- [SW] Bruce E. Sagan and Robert Willenbring. Discrete Morse theory and the consecutive pattern poset. *J. Algebraic Combin.* To appear. [arXiv:1107.3262](https://arxiv.org/abs/1107.3262).
- [Tom10] Masaya Tomie. A generalization of the Chebyshev polynomials and nonrooted posets. *Int. Math. Res. Not. IMRN*, (5):856–881, 2010.
- [Wil] Robert Willenbring. The Möbius function of generalized factor order. Preprint. [arXiv:1108.3899](https://arxiv.org/abs/1108.3899).
- [Wil02] Herbert S. Wilf. The patterns of permutations. *Discrete Math.*, 257(2-3):575–583, 2002.