The Möbius function of a composition poset

Bruce E. Sagan*
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027
http://www.math.msu.edu/~sagan
sagan@math.msu.edu

Vincent Vatter†
Department of Mathematics and Statistics
University of St. Andrews
St. Andrews, Fife, Scotland KY16 9SS
http://turnbull.mcs.st-and.ac.uk/~vince
vince@mcs.st-and.ac.uk

February 13, 2006

Key Words: composition, discrete Morse function, Möbius function, permutation pattern, subword order

AMS classifications: 06A07, 05E25, 68R15

Abstract

We determine the Möbius function of a poset of compositions of an integer. In fact we give two proofs of this formula, one using an involution and one involving discrete Morse theory. This composition poset turns out to be intimately connected with subword order, whose Möbius function was determined by Björner. We show that using a generalization of subword order, we can obtain both Björner's results and our own as special cases.

1 Introduction

If A is any set then the corresponding *Kleene closure* or *free monoid*, A^* is the set of words with letters from A, i.e.,

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

We denote the length (number of elements) of w by $|w|$.

*This work was partially done while the author was on leave at DIMACS
†Partially supported by an award from DIMACS and an NSF VIGRE grant to the Rutgers University Department of Mathematics.

Letting $\mathbb{P}$ denote the positive integers, we see that $\mathbb{P}^*$ is the set of integer compositions (ordered partitions). We can turn $\mathbb{P}^*$ into a partially ordered set by letting $u \leq w$ if there is a subword $w(i_1)w(i_2)\ldots w(i_l)$ of w having length $l = |u|$ such that $u(j) \leq w(i_j)$ for $1 \leq j \leq l$. For example, $433 \leq 16243$ because of the subword 643. Bergeron, Bousquet-Mélou and Dulucq [2] were the first to study $\mathbb{P}^*$, enumerating saturated chains that begin at its minimal element. Snellman has also studied saturated chains in this poset as well as two other partial orders on $\mathbb{P}^*$ [21, 22]. One of these other partial orders was originally defined by Björner and Stanley [8] who showed that it has analogues of many of the properties of Young's lattice. One of the main results of this paper is a formula for the Möbius function of $\mathbb{P}^*$.

This order on $\mathbb{P}^*$ is closely related to subword order. If A is any set then the *subword order* on A^* is defined by letting $u \leq w$ if w contains a subsequence $w(i_1), w(i_2), \ldots, w(i_l)$ such that $u(j) = w(i_j)$ for $1 \leq j \leq l = |u|$. By way of illustration, if $A = \{a, b\}$ then $abba \leq aabbbaba$ because $w(1)w(3)w(5)w(8) = abba$. Note that we use the notation A^* when referring to subword order as opposed to the partial order on $\mathbb{P}^*$, even though we use $\leq$ for both. We will always give enough context to make it clear which poset we are dealing with. Björner [4] was the first to completely determine the Möbius function of subword order, although special cases had been obtained previously by Farmer [13] and Viennot [24]. In fact, Björner gave two proofs of his formula, one using an involution [5] and one using shellability [4]. He also gave a demonstration with Reutenauer [6] via generating functions on monoids. Another proof was given by Warnke [26] using induction while Wang and Ma [25] used Cohen-Macaulayness to investigate μ. We derive the Möbius function formula for $\mathbb{P}^*$ using first combinatorial and then topological techniques.

In the next section we review Björner's result for subword order as well as the related definitions which will be useful for $\mathbb{P}^*$. It contains a statement of our formula for the Möbius function of $\mathbb{P}^*$ in Theorem 2.2. Section 3 is devoted to giving a proof of this theorem using a sign-reversing involution.

Although intervals in A^* are shellable, those in $\mathbb{P}^*$ need not even be connected as evident from the example in Figure 2, so we need a more powerful tool to study the topology of the composition poset. For this we turn to discrete Morse theory, which was developed by Forman [14, 15] and can be used to compute the homology of any CW-complex subject to certain mild regularity restrictions. A method for applying this theory to the order complex of a poset was given by Babson and Hersh [1] and further studied by Hersh herself [16]. Since this is a relatively recent addition to the combinatorial toolbox, we provide an exposition of the basic ideas of the theory in Section 4. The subsequent section gives a Morse theoretic proof of Theorem 2.2.

The similarity between the formulas for the Möbius functions of A^* and $\mathbb{P}^*$ leads one to ask if there is a common generalization. In fact, if P is any poset then there is a partial order on P^* which we call generalized subword order. It has been used in the context of well-quasi-ordering; see Kruskal's article [18] for a survey of the early literature. When P is a chain or an antichain, one recovers our results or Björner's, respectively. This construction is studied in Section 6. We end with a section of comments and open problems.

2 Subword and composition order

We now review Björner's formula for the Möbius function of subword order, reformulating it slightly so as to emphasize the connection to the composition order which is our main objective. We assume the reader is familiar with Möbius functions, but all the necessary definitions and theorems we use here can be found in Stanley's text [23, §3.6–3.7].

We first need to restate the definition of the partial order in A^* in a way that, although slightly more complicated, has a direct connection with the Möbius function. Suppose we have a distinguished symbol, 0, and suppose that $0 \notin A$. Then a word $\eta = \eta(1)\eta(2)\ldots\eta(n) \in (A \cup 0)^*$ has *support*

$$\operatorname{Supp} \eta = \{i \mid \eta(i) \neq 0\}.$$

An *expansion* of $u \in (A \cup 0)^*$ is a word $\eta_u \in (A \cup 0)^*$ such that the restrictions of u and η_u to their supports are equal. For example, if $u = abba$ then one expansion of u is $\eta_u = a0b0b00a$. An *embedding* of u into w is an expansion η_{uw} of u which has length $|w|$ and satisfies

$$\eta_{uw}(i) = w(i) \text{ for all } i \in \operatorname{Supp} \eta_{uw}.$$

Note that $u \leq w$ in A^* if and only if there is an embedding of u into w. The example expansion η_{uw} given above is exactly the embedding which corresponds to the subword of $w = aabbbaba$ given at the beginning of the third paragraph of this paper. If w is clear from context we simply write η_u for η_{uw}.

The Möbius function of A^* counts certain types of embeddings. If $a \in A$ then a *run* of a's in w is a maximal interval of indices $[r, t]$ such that

$$w(r) = w(r+1) = \ldots = w(t) = a.$$

Continuing our example, the runs in $w = aabbbaba$ are $[1, 2]$, $[3, 5]$, $[6, 6]$, $[7, 7]$, and $[8, 8]$. Call an embedding η_u into w *normal* if, for every $a \in A$ and every run $[r, t]$ of a's in w, we have

$$(r, t] \subseteq \operatorname{Supp} \eta_u$$

where $(r, t]$ denotes the half-open interval. In our running example, this means that $w(2) = a$ and $w(4) = w(5) = b$ must be in any normal embedding. (If $r = t$ then $(r, t] = \emptyset$, so there is no restriction on runs of one element.) Thus in this case there are precisely two normal embeddings of u into w, namely

$$\eta_u = 0a0bba00 \text{ and } 0a0bb00a.$$

Define $\binom{w}{u}_n$ to be the number of normal embeddings of u into w.

We now have everything in place to state Björner's result.

Theorem 2.1 (Björner [4]). *If $u, w \in A^*$ then*

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

Finishing our example, we see that

$$\mu(abba, aabbbaba) = (-1)^{8-4} \cdot 2 = 2.$$

We now turn to $\mathbb{P}^*$. The definitions of support and expansion are exactly as before, but the notion of embedding must be updated to reflect the different partial order. To this end we define an *embedding* of u into w as an expansion η_u of u having length $|w|$ such that

$$\eta_u(i) \leq w(i) \text{ for } 1 \leq i \leq |w|.$$

Again, $u \leq w$ in $\mathbb{P}^*$ is equivalent to the existence of an embedding of u into w. An interval in $\mathbb{P}^*$ is displayed in Figure 2.

If $u \leq w$ and η_w is an expansion of w then there is a unique *last* or *rightmost* embedding ρ_{uw} of u into η_w which has the property that for any other embedding η_{uw} of u into η_w one has $\mathrm{Supp}(\eta_{uw}) \leq \mathrm{Supp}(\rho_{uw})$. (If $S = \{i_1 < \cdots < i_m\}$ and $S' = \{i'_1 < \cdots < i'_m\}$ then we write $S \leq S'$ to mean that $i_j \leq i'_j$ for $1 \leq j \leq m$.) Note that ρ_{uw} depends on η_w, not just on w, but the expansion of w used will always be clear from context.

Like subword order, we define normal embeddings for $\mathbb{P}^*$ in terms of runs (defined in the same way in this context). We say that an embedding η_u into w is *normal* if the following conditions hold.

1. For $1 \leq i \leq |w|$ we have $\eta_u(i) = w(i)$, $w(i) - 1$, or 0.

2. For all $k \geq 1$ and every run $[r, t]$ of k's in w, we have

   (a) $(r, t] \subseteq \mathrm{Supp}\, \eta_u$ if $k = 1$,

   (b) $r \in \mathrm{Supp}\, \eta_u$ if $k \geq 2$.

Comparing this with Björner's definition, we see that in $\mathbb{P}^*$ a normal embedding can have three possible values at each position instead of two. Also, the run condition for ones is the same as in A^*, while that condition for integers greater than one is complementary. As an example, if $w = 2211133$ and $u = 21113$, then there are two normal embeddings, namely $\eta_u = 2101130$ and 2011130. Note that 2001113 and 0211130 are not normal since they violate conditions (1) and (2), respectively.

As with subword order, the Möbius function for compositions is given by a signed sum over normal embeddings, although here the sign of a normal embedding depends on the embedding itself and not just the length of the compositions. Given a normal embedding η_u into w we define its *defect* to be

$$d(\eta_u) = \#\{i \mid \eta_u(i) = w(i) - 1\}.$$

The *sign* of the embedding is, then, $(-1)^{d(\eta_u)}$.

We can now state our main theorem about $\mathbb{P}^*$.

Theorem 2.2. *If $u, w \in \mathbb{P}^*$ then*

$$\mu(u, w) = \sum_{\eta_u} (-1)^{d(\eta_u)}$$

where the sum is over all normal embeddings η_u into w.

In the example of the previous paragraph, this gives

$$\mu(21113, 2211133) = (-1)^2 + (-1)^0 = 2.$$

Although this example does not show it, it is possible to have cancellation among the terms in the sum for μ.

3 Proof by sign-reversing involution

We now prove Theorem 2.2 using a sign-reversing involution. The proof is similar in nature to Björner's proof in [5], but is significantly more complicated.

Proof (of Theorem 2.2). If $u = w$ then there is exactly one normal embedding η_{ww} and it has defect 0. This gives $(-1)^0 = 1 = \mu(w, w)$, as desired.

Now assume that $u < w = w(1) \cdots w(n)$. Since the Möbius recurrence uniquely defines μ, it suffices to show that

$$\sum_{v \in [u,w]} \sum_{\eta_{vw}} (-1)^{d(\eta_{vw})} = 0,$$

where the inner sum is over all normal embeddings of v into w. We prove this by constructing a sign-reversing involution on the set of normal embeddings η_{vw} for $v \in [u, w]$.

Let ρ_{uv} denote the rightmost embedding of u into η_{vw}, so for all $i \in [1, n]$,

$$\rho_{uv}(i) \leq \eta_{vw}(i) \leq w(i). \tag{1}$$

Also let s denote the left-most index where ρ_{uv} and w differ, i.e., $s = \min\{i : \rho_{uv}(i) \neq w(i)\}$. Since $u < w$, s must exist, and by the definition of s we have

$$\rho_{uv}(i) = \eta_{vw}(i) = w(i) \text{ for all } i \in [1, s). \tag{2}$$

Set $k = w(s)$, so $\eta_{vw}(s)$ is k, $k-1$, or 0 by normality and $\rho_{uv}(s) \leq k - 1$. Finally, let $[r, t]$ denote the indices of the run of k's in w that contains the index s.

Our involution maps η_{vw} to the embedding $\eta_{\overline{v}w} = \eta_{\overline{v}w}(1) \cdots \eta_{\overline{v}w}(n)$ where $\eta_{\overline{v}w}(i) = \eta_{vw}(i)$ for all $i \neq s$ and $\eta_{\overline{v}w}(s)$ is determined by the following rules. If $k = 1$ then

$$\eta_{\overline{v}w}(s) = \begin{cases} 0 & \text{if } \eta_{vw}(s) = 1, \\ 1 & \text{if } \eta_{vw}(s) = 0. \end{cases} \tag{3}$$

If $k \geq 2$, $s > r$, and $\rho_{uv}(s) = 0$ then

$$\eta_{\overline{v}w}(s) = \begin{cases} 0 & \text{if } \eta_{vw}(s) = k - 1, \\ k - 1 & \text{if } \eta_{vw}(s) = 0. \end{cases} \tag{4}$$

Finally, if $k \geq 2$ and either $s = r$ or $\rho_{uv}(s) \neq 0$ then

$$\eta_{\overline{v}w}(s) = \begin{cases} k - 1 & \text{if } \eta_{vw}(s) = k, \\ k & \text{if } \eta_{vw}(s) = k - 1. \end{cases} \tag{5}$$

It is not obvious that this map is defined for all normal embeddings η_{vw} when $k \geq 2$. For example, if $s > r$ and $\rho_{uv}(s) = 0$, then we should apply (4), but it is a priori possible that $\eta_{vw}(s) = k$, in which case (4) is not defined. However, since $s > r$, $\rho_{uv}(s-1) = w(s-1) = k$ by (2), and this contradicts our choice of ρ_{uv} as the rightmost embedding of u into η_{vw}. Similar issues arise when $s = r$ or $\rho_{uv}(s) \neq 0$: if $s = r$ then $\eta_{vw}(s) \neq 0$ by normality, and if $\rho_{uv}(s) \neq 0$ then $\eta_{vw}(s) \neq 0$ by (1).

Having established that this map is indeed defined on all normal embeddings, we have several properties to prove. First, it is evident from (3), (4), and (5) that the number of elements equal to $k-1$ changes by exactly one in passing from η_{vw} to $\eta_{\overline{v}w}$, and thus $(-1)^{d(\eta_{\overline{v}w})} = -(-1)^{d(\eta_{vw})}$. We must now prove that $\eta_{\overline{v}w}$ is a normal embedding of $\overline{v}$ into w for some $\overline{v} \in [u,w]$ and that this map is an involution.

We begin by showing that $\eta_{\overline{v}w}$ is an embedding of some $\overline{v} \in [u,w]$ into w. It follows from the fact that η_{vw} is an embedding into w and the definition of our map that $\eta_{\overline{v}w}$ is an embedding of some word $\overline{v}$ into w, so we only need to show that $\overline{v} \geq u$. We prove this by showing that

$$\eta_{\overline{v}w}(i) \geq \rho_{uv}(i) \tag{6}$$

for all $i \in [1,n]$. This is clear for all $i \neq s$ because for these indices $\eta_{\overline{v}w}(i) = \eta_{vw}(i) \geq \rho_{uv}(i)$. Furthermore, $\rho_{uv}(s) \leq k-1$ by the definition of s, so the only case in which (6) is not immediate is when $k \geq 2$ and $\eta_{\overline{v}w}(s) = 0$. However, this can only occur from using (4), which requires that $\rho_{uv}(s) = 0$, completing the demonstration that $\overline{v} \in [u,w]$.

We now aim to show that $\eta_{\overline{v}w}$ is a normal embedding. If $\eta_{\overline{v}w}(s) > 0$ then $\mathrm{Supp}(\eta_{\overline{v}w}) \supseteq \mathrm{Supp}(\eta_{vw})$, so the normality of $\eta_{\overline{v}w}$ follows from the normality of η_{vw} and the fact that $\eta_{\overline{v}w}(s)$ is either $k-1$ or k. If $\eta_{\overline{v}w}(s) = 0$ then there are two cases depending on whether (3) or (4) was applied. Suppose first that (3) was applied, so $k = 1$. Comparing (3) with the definition of normality, we see that it suffices to show $s = r$. Suppose to the contrary that $s \in (r,t]$. Since s is the left-most position at which ρ_{uv} and w differ, we have $\rho_{uv}(s) = 0$ and $\rho_{uv}(s-1) = 1$. However, this contradicts our choice of ρ_{uv} as the rightmost embedding of u in η_{vw}. Now suppose that (4) was applied, so $s > r$ and $\rho_{uv}(s) = 0$. Then (2) implies that $\eta_{\overline{v}w}(r) = \eta_{vw}(r) = k$, and normality is preserved.

It only remains to show that this map is an involution. Consider applying the map to $\eta_{\overline{v}w}$. In this process we define $\rho_{u\overline{v}}$ to be the rightmost embedding of u into $\eta_{\overline{v}w}$, $\overline{s} = \min\{i : \rho_{u\overline{v}}(i) \neq w(i)\}$, and $\overline{k} = w(\overline{s})$. We then follow the rules (3), (4), and (5) to construct an embedding $\eta_{\overline{\overline{v}}w}$, which we would like to show is equal to η_{vw}.

First we claim that $\rho_{u\overline{v}} = \rho_{uv}$. Suppose to the contrary that $\rho_{u\overline{v}} \neq \rho_{uv}$. By (6), $\eta_{\overline{v}w}(i) \geq \rho_{uv}(i)$ for all i, so ρ_{uv} also gives an embedding of u into $\overline{v}$, and thus the only way we can have $\rho_{u\overline{v}} \neq \rho_{uv}$ is if $\rho_{u\overline{v}}$ is further to the right than ρ_{uv}. This requires that

$$0 = \rho_{uv}(s) < \rho_{u\overline{v}}(s) \leq k \tag{7}$$

and that

$$\eta_{vw}(s) < \eta_{\overline{v}w}(s). \tag{8}$$

Because we are assuming that $\rho_{\overline{v}w}$ is further to the right than ρ_{vw}, there is some position to the left of s at which ρ_{uv} is nonzero. In fact, (2) shows that we must have $\rho_{uv}(s-1) \neq 0$ and also implies that

$$\eta_{vw}(s) < \rho_{u\overline{v}}(s) = \rho_{uv}(s-1) = w(s-1). \tag{9}$$

We now consider the three cases arising from each of the rules (3), (4), and (5) in turn.

Suppose (3) was applied so $k = 1$. It follows from (8) and the definition of our map that $\eta_{vw}(s) = 0$. Also (7) and (9) give $w(s-1) = \rho_{u\overline{v}}(s) = 1$. However, this implies that η_{vw} zeroed out a 1 which was not the first in its run, contradicting normality.

Now suppose (4) was applied. Then by (8) we have $\eta_{\overline{v}w}(s) = k-1$. We also have $s > r$ which in conjunction with (9) gives $\rho_{u\overline{v}}(s) = w(s-1) = k$. This implies that $\eta_{\overline{v}w}(s) < \rho_{u\overline{v}}(s)$, but that contradicts the $\overline{v}$ version of (1).

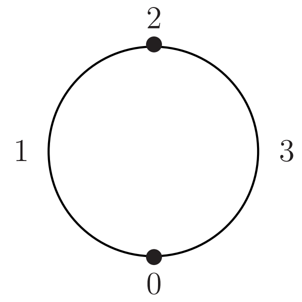

Figure 1: A Morse function on a CW-complex

Finally suppose that (5) was used. By (7) it must be the case that $s = r$. Also, equation (8) gives $\eta_{vw}(s) = k - 1$ and $\eta_{\overline{v}w}(s) = k$. Now applying (7) and (9) we have $k - 1 < w(s-1) = \rho_{u\overline{v}}(s) \le k$, so $w(s-1) = k$ which contradicts that fact that $s = r$.

Now that we have established the equality of ρ_{uv} and $\rho_{u\overline{v}}$, the fact that this map is an involution can be readily observed. We must have $\overline{s} = s$, so $\overline{k} = k$, and thus we apply the same rule to go from $\eta_{\overline{v}w}$ to $\eta_{\overline{\overline{v}}w}$ as we applied to get $\eta_{\overline{v}w}$ from $\eta_{\overline{v}w}$, and each of these rules is clearly an involution. $\square$

4 Introduction to discrete Morse theory

In this section we review the basic ideas behind Forman's discrete Morse theory [14, 15] as well as Babson and Hersh's method for applying the theory to the order complex of a poset [1].

Let X be a CW-complex. Since we will be working in reduced homology, we assume that X has an empty cell $\emptyset$ of dimension -1 which is contained in every cell of X. If σ is a d-cell (cell of dimension d) in X then let σ^∂ be the set of $(d-1)$-cells τ which are contained in the closure $\overline{\sigma}$. Dually, let σ^δ denote the set of $(d+1)$-cells τ such that $\sigma \subset \overline{\tau}$.

A real-valued function f on the cells of X is a *Morse function* if it satisfies the following two conditions.

1. For every cell σ of X we have
   (a) $\#\{\tau \in \sigma^\partial \mid f(\tau) \ge f(\sigma)\} \le 1$, and
   (b) $\#\{\tau \in \sigma^\delta \mid f(\tau) \le f(\sigma)\} \le 1$.
2. If $\tau \in \sigma^\delta$ and $f(\tau) \le f(\sigma)$ then σ is a regular face of τ.

Intuitively the first condition says that, with only certain exceptions, f increases with dimension. In fact, $f(\sigma) = \dim \sigma$ is a perfectly good Morse function on X, although we will see shortly that it is not very interesting. A simple example of a Morse function on a CW-complex is given in Figure 1 where the value of f is given next to each cell σ and we also set $f(\emptyset) = -1$.

The fact that condition (1) holds for every cell implies that, in fact, at most one of the two sets under consideration has cardinality equal to 1. Thus the function f induces a *Morse matching* between pairs of cells σ, τ with $\sigma \in \tau^\partial$ and $f(\sigma) \geq f(\tau)$. The regularity condition ensures that for each such pair there is an elementary collapse of $\overline{\tau}$ onto $\overline{\tau} - (\tau \cup \sigma)$. The cells which are not matched by f are called *critical*. Since each collapse is a homotopy equivalence, X can be collapsed onto a homotopic complex X^f built from the critical cells. In our example, the cells labeled 1 and 2 are matched and after collapsing we clearly have a complex which is still homotopically a circle. Note that if we take f to be the dimension function then every cell is critical and $X^f = X$, so the cell complex does not simplify in this case which does not help in understanding its structure.

Let $\tilde{m}_d$ be the number of critical d-cells of X and let $\tilde{b}_d$ be the d-th reduced Betti number over the integers. We also use $\tilde{\chi}(X)$ for the reduced Euler characteristic. From the considerations in the previous paragraph, we have the following *Morse inequalities* which are analogous to those in traditional Morse theory.

Theorem 4.1 (Forman [15]). *For any Morse function on a cell complex X we have*

1. $\tilde{b}_d \leq \tilde{m}_d$ *for* $d \geq -1$, *and*

2. $\tilde{\chi}(X) = \sum_{d \geq -1} (-1)^d \tilde{m}_d$. □

(One can get further inequalities relating various partial alternating sums of the $\tilde{b}_d$ and $\tilde{m}_d$.) Continuing our example, we see that $\tilde{m}_{-1} = \tilde{m}_0 = \tilde{m}_1 = 1$ which bound $\tilde{b}_{-1} = \tilde{b}_0 = 0$ and $\tilde{b}_1 = 1$, as well as $\tilde{\chi}(X) = -1 + 1 - 1 = -1$, as expected.

We now turn to the special case of order complexes. Let P be a poset and consider an open interval (u, w) in P. The corresponding *order complex* $\Delta(u, w)$ is the abstract simplicial complex whose simplices (faces) are the chains in (u, w). We are interested in the order complex because of the fundamental fact [20] that

$$\mu(u, w) = \tilde{\chi}(\Delta(u, w)). \tag{10}$$

Therefore finding a Morse function for $\Delta(u, w)$ could permit us to derive the corresponding Möbius value as well as give extra information about its Betti numbers. Suppose we have an ordering of the maximal chains of (u, w) (facets of $\Delta(u, w)$), say $C_1, C_2, \ldots, C_l$. Call a face (subchain) σ of C_k *new* if it is not contained in any C_j for $j < k$. We would like to construct a Morse matching inductively, where at the kth stage we extend the matching on the faces in C_j for $j < k$ by matching up as many of the new faces σ in C_k as possible. It turns out that under fairly mild conditions on the facet ordering, one can construct such a matching so that all the new faces in C_k are matched if there are an even number of them, and only one is left unmatched if the number is odd. Thus adding each facet contributes at most one critical cell. A maximal chain contributing a critical cell is called a *critical chain*. In reading the details of this construction, the reader may find it useful to refer to the example of the interval $[322, 3322] \subset \mathbb{P}^*$ given in Figure 2. Note that by abuse of notation we include u and w when writing out a maximal chain C, even though C is really a subset of the open interval (u, w). Also, because of the way our chain order is constructed, we start

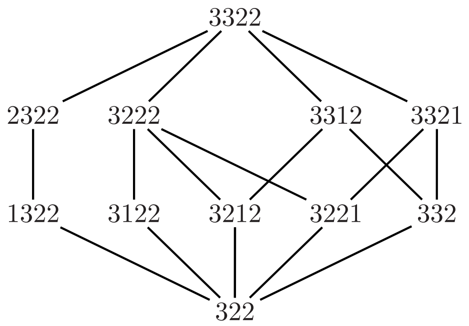

$$
\begin{array}{ll}
C_1: & 3322 \overset{1}{\text{——}} 2322 \overset{1}{\text{——}} 1322 \overset{1}{\text{——}} 0322 \\
C_2: & 3322 \overset{2}{\text{——}} 3222 \overset{2}{\text{——}} 3122 \overset{2}{\text{——}} 3022 \\
C_3: & 3322 \overset{2}{\text{——}} 3222 \overset{3}{\text{——}} 3212 \overset{3}{\text{——}} 3202 \\
C_4: & 3322 \overset{2}{\text{——}} 3222 \overset{4}{\text{——}} 3221 \overset{4}{\text{——}} 3220 \\
C_5: & 3322 \overset{3}{\text{——}} 3312 \overset{2}{\text{——}} 3212 \overset{3}{\text{——}} 3202 \\
C_6: & 3322 \overset{3}{\text{——}} 3312 \overset{3}{\text{——}} 3302 \overset{2}{\text{——}} 3202 \\
C_7: & 3322 \overset{4}{\text{——}} 3321 \overset{2}{\text{——}} 3221 \overset{4}{\text{——}} 3220 \\
C_8: & 3322 \overset{4}{\text{——}} 3321 \overset{4}{\text{——}} 3320 \overset{2}{\text{——}} 3220
\end{array}
$$

Figure 2: The interval $[322, 3322]$ and its maximal chains

with the top element w and work down to u which is dual to what is done normally. Thus in a chain C, terms like "first" and "last" refer to this ordering of C's elements. Finally, we list the elements of a chain as embeddings into w for reasons which will become apparent when we also describe the labels given to the edges (covers) of a chain.

To define the types of chain orderings we consider, suppose we have two chains $C: w = v_0 \text{ —— } v_1 \text{ —— } \ldots \text{ —— } u$ and $C': w = v'_0 \text{ —— } v'_1 \text{ —— } \ldots \text{ —— } u$ where $x \text{ —— } y$ means that x covers y. Then we say that C and C' *agree to index* j if $v_i = v'_i$ for $i \leq j$. In addition, C and C' *diverge from index* j if they agree to index j and $v_{j+1} \neq v'_{j+1}$. In addition, we use the notation $C' < C$ to mean that C' comes before C in the order under consideration. An ordering of the maximal chains of $[u, w]$ is a *poset lexicographic order*, or *PL-order* for short, if it satisfies the following condition. Suppose C' and C diverge from index j with $C' < C$. Then for any maximal chains D' and D which agree to index $j + 1$ with C' and C, respectively, we must have $D' < D$. Note that orderings coming from the EL-labelings introduced by Björner [3] or from the more general CL-labelings of Björner and Wachs [9] are PL-orders as long as one breaks ties among labels consistently.

The PL-order we use in $\mathbb{P}^*$ is as follows. If $x \text{ —— } y$ is a cover then y is obtained from x by reducing a single part of x by 1. Thus there is a unique normal embedding of y into x, since if a 1 is reduced to 0 then it must be the first element in the run of ones to which it belongs. Similarly, for any expansion η_x there is a unique normal embedding of y into η_x. Now given any chain $C: w = v_0 \text{ —— } v_1 \text{ —— } \ldots \text{ —— } u$ we inductively associate with each v_j an embedding η_{v_j} into w where $\eta_{v_0} = w$ and, for $j \geq 0$, $\eta_{v_{j+1}}$ is the unique normal embedding of v_{j+1} into η_{v_j}. We label the edge v_j —— v_{j+1} of C with the index i of the position which was decreased in passing from η_{v_j} to $\eta_{v_{j+1}}$. Furthermore, we often write η_{v_j} in place of v_j when listing the elements of C. Figure 2 illustrates this labeling. It is important to note that although $\eta_{v_{j+1}}$ is normal in η_{v_j}, it need not be normal in w. We should also remark that this labeling is similar to the one used by Björner [4] in his CL-shelling of the intervals in subword order. Finally, if one orders the chains of $[u, w]$ using ordinary lexicographic order on their label sequences, then the result is a PL-order. This is due to the fact that if two chains agree to index j then their first j labels are the same. The chains in Figure 2 are listed in PL-order.

We now return to the general exposition. To construct our matching, when we come to a chain $C: \; w = v_0$ —— v_1 —— $\ldots$ —— u in a given order we must be able to determine which faces of C are new. Denote the open interval I from v_i to v_j in C by

$$I = C(v_i, v_j) = v_{i+1} \text{ —— } v_{i+2} \text{ —— } \cdots \text{ —— } v_{j-1}.$$

(Do not to confuse this with an open interval in the poset.) Then I is a *skipped interval* if $C - I \subset C'$ for some $C' < C$. It is a *minimal skipped interval* or *MSI* if it does not strictly contain another skipped interval. In Figure 2, the MSI's are circled. One can find the MSI's by taking the maximal intervals in $C - (C \cap C')$ for each $C' < C$ and then throwing out any that are not containment minimal in C. Let $\mathcal{I} = \mathcal{I}(C)$ be the set of MSI's in C. Then it is easy to check that a face σ is new in C if and only if σ has a nonempty intersection with every $I \in \mathcal{I}(C)$.

The set $\mathcal{I}$ is not quite sufficient to construct the matching because the MSI's can overlap and we will need disjoint intervals. However, there are no containments among the intervals in $\mathcal{I}$, so they can be ordered $I_1, I_2, I_3, \ldots$ according to when they are first encountered on C. We now inductively construct a set $\mathcal{J} = \mathcal{J}(C)$ of *J-intervals* as follows. Let $J_1 = I_1$. Then consider the intervals $I_2' = I_2 - J_1, I_3' = I_3 - J_1, \ldots$; throw out any which are not minimal; and pick the first one which remains to be J_2. Continue this process until there are no nonempty modified MSI's left. It happens that in all the critical chains for $\mathbb{P}^*$, the intervals in $\mathcal{I}$ will already be disjoint and so we will not need this step.

We are finally in a position to describe the matching. List the maximal chains of $[u, w]$ using a PL-order. A family $\mathcal{K}$ of intervals of a maximal chain C *covers* the chain if $C = \cup\mathcal{K}$. There are three cases depending on whether $\mathcal{I}(C)$ or $\mathcal{J}(C)$ covers C or not. First suppose that $\mathcal{I}(C)$ does not cover C, so neither does $\mathcal{J}(C)$, and pick x_0 to be the first vertex in $C - \cup\mathcal{I}$. Consider the map $\sigma \to \sigma\Delta\{x_0\}$ where Δ is symmetric difference (not the order complex). One can show that this map is a fixed-point free involution on the new faces σ in C which extends the Morse matching already constructed from the previous chains. Now suppose that $\mathcal{J} = \{J_1, J_2, \ldots, J_r\}$ does cover C and consider the new face $\sigma_C = \{x_1, x_2, \ldots, x_r\}$ where x_i is the first element of J_i for $1 \le i \le r$. Given any other new face $\sigma \neq \sigma_C$, we find the interval J_i of smallest index where $\sigma \cap J_i \neq \{x_i\}$ and map $\sigma \to \sigma\Delta\{x_i\}$. This involution pairs up all new faces in C except σ_C, which is critical. Finally, suppose that $\mathcal{I}$ covers C but $\mathcal{J}$ does not. Then we use the mapping of the second case to pair up all new faces whose restriction to $\cup\mathcal{J}$ is different from σ_C. We also pair up the remaining new faces (including σ_C) by using the mapping of the first case where we take x_0 to be the first vertex in $C - \cup\mathcal{J}$. Thus we have outlined the proof of the following theorem, remembering that the dimension of a simplex is one less than its number of vertices.

Theorem 4.2 ([1]). *Let P be a poset and $[u,w]$ be a finite interval in P. For any PL-order on the maximal chains of $[u,w]$, the above construction produces a Morse matching in $\Delta(u,w)$ with the following properties.*

1. *The maximal chain C is critical if and only if $\mathcal{J}$ covers C.*
2. *If C is critical then its unique critical cell has dimension $\#\mathcal{J}(C)-1$.* □

5 A Morse theory derivation of μ

We are now ready to find the critical cells for the PL-order in $\mathbb{P}^*$ defined previously. We first need three lemmas which will prove useful in a number of cases. Unless otherwise specified, we always use the notation

$$C: \; w = v_0 \xrightarrow{l_1} v_1 \xrightarrow{l_2} v_2 \xrightarrow{l_3} \dots \xrightarrow{l_d} v_d = u \tag{11}$$

for labeled maximal chains, or

$$C: \; w = \eta_{v_0} \xrightarrow{l_1} \eta_{v_1} \xrightarrow{l_2} \eta_{v_2} \xrightarrow{l_3} \dots \xrightarrow{l_d} \eta_u \tag{12}$$

if we wish to be specific about the embeddings determined by C. We also use

$$l(C) = (l_1, l_2, \dots, l_d)$$

for its label sequence.

Take an interval $[u,w] \subset \mathbb{P}^*$ with $|u| = |w|$ and let $m_i = w(i) - u(i)$. Now consider the multiset $M_{uw} = \{\{1^{m_1}, 2^{m_2}, \dots\}\}$ where i^{m_i} means that i is repeated m_i times. Then every permutation of M is the label sequence for a unique maximal chain in $[u,w]$ and this accounts for all the chains. (In fact, $[u,v]$ is isomorphic to the poset of submultisets of M_{uw}.) We record this simple observation for later reference.

Lemma 5.1 (Same Length Lemma). *If $|u| = |w|$ then the the label function l gives a bijection between the maximal chains in $[u,w]$ and the permutations of M_{uw}. In particular, if M_{uw} contains only one distinct element (possibly with multiplicity) then $[u,w]$ contains a unique maximal chain.* □

If $|u| < |w|$ then we no longer have the nice bijection of the previous paragraph, but we can still say something. Let C be a maximal chain as in (11) and let $l' = (l'_1, \dots, l'_d)$ be any permutation of the label sequence $l(C)$. Then l' defines a sequence of expansions $\eta_{v'_0}, \eta_{v'_1}, \dots, \eta_{v'_d}$ where $\eta_{v'_0} = w$ and for $j \geq 1$ we get $\eta_{v'_j}$ from $\eta_{v'_{j-1}}$ by subtracting one from position l'_j in $\eta_{v'_{j-1}}$. It is still true that $C' : w = v'_0 \;\rule{1.5em}{0.4pt}\; v'_1 \;\rule{1.5em}{0.4pt}\; \dots \;\rule{1.5em}{0.4pt}\; v'_d = u$ is a maximal chain in $[u,w]$. We call C' the *chain specified by l'*. Since $\eta_{v'_j}$ may not be a normal embedding in $\eta_{v'_{j-1}}$, we may not have $l(C') = l'$. Still, at the first place where $l(C')$ and l' differ, that difference must have been caused because using the label in l' would have resulted in changing a 1 to a 0 where that 1 was not the first in its run. Thus the corresponding normal embedding in $l(C')$ uses the first 1 in that run which is to the left. Hence $l(C') \leq l'$ in lexicographic order. We summarize this discussion in the following lemma.

Lemma 5.2 (Chain Specification Lemma). *If C is a maximal chain in $[u, w]$ and l' is any permutation of $l(C)$ then $l(C') \leq l'$ where C' is the chain specified by l'.* □

As our first application of the Chain Specification Lemma, we can determine what happens at descents. A *descent* of C is $v_j \in C$ such that $l_j > l_{j+1}$. An *ascent* is defined by reversing the inequality.

Lemma 5.3 (Descent Lemma). *If v_j is a descent of C then it is an MSI.*

Proof. Let l' be the permutation of $l(C)$ gotten by interchanging l_j and l_{j+1} and let C' be the chain specified by l'. Then by Lemma 5.2 we have $l(C') \leq l' < l(C)$, so C' comes before C in PL-order and it is easy to check that C' diverges from C at v_{j-1} and rejoins C at v_{j+1}. Thus $\{v_j\}$ is a skipped interval; and since the interval contains only one element it must also be minimal. □

We only need a few more definitions to state our result characterizing the critical chains. A chain C will be said to have a certain property, e.g., weakly decreasing, if $l(C)$ has that property. Also, if η_u is a normal embedding into w then we need to keep track of the zero positions which did not come from decreasing a 1 in w by letting

$$D(\eta_u) = \#\{i \mid \eta_u(i) = 0,\ w(i) \geq 2\}.$$

Theorem 5.4. *Consider the maximal chains in $[u, w] \subset \mathbb{P}^*$ in the given PL-order.*

1. *There is a bijection between critical chains C and normal embeddings η_u into w where the chain corresponding to η_u is the unique weakly decreasing chain ending at η_u.*

2. *If C is critical and ends at η_u then $\mathcal{I}(C) = \mathcal{J}(C)$ and*

$$\#\mathcal{I}(C) = d(\eta_u) + 2D(\eta_u) - 1.$$

We shall prove this theorem by considering 3 cases: when C is weakly decreasing and ends at a normal embedding, when C is weakly decreasing and does not end at a normal embedding, and when C is not weakly decreasing. Note that for any embedding η_u into w, there is at most one decreasing chain ending at η_u, and that if η_u is normal then such a chain will exist because it will be possible to make each cover normal. Thus there is a bijection between normal embeddings and weakly decreasing chains ending at them, but we need to show such chains are critical. To do so, we define a *plateau* of C to be an interval $C(v_i, v_j)$ such that $(l_{i+1}, l_{i+2}, \ldots, l_j)$ is a run of length at least 2 in $l(C)$.

Proposition 5.5. *If C is weakly decreasing and ends at a normal embedding η_u then C is critical, $\mathcal{I}(C) = \mathcal{J}(C)$, and $\#\mathcal{I}(C) = d(\eta_u) + 2D(\eta_u) - 1$.*

Proof. Every descent is an MSI of C by the Descent Lemma, so any other MSI must be contained in a plateau by minimality. In fact, we claim that any plateau $C(v_i, v_j)$ is an MSI. Without loss of generality we can assume $v_i = w$ (since otherwise v_i is a descent and so no MSI can contain it) and $v_j = u$.

To show $C = C(w, u)$ is a skipped interval, first note that by construction $l(C)$ consists of c repeated $k = j - i \geq 2$ times, so $w(c) = \eta_u(c) + k$. Thus by normality and the fact that

$k \geq 2$ we get $\eta_u(c) = 0$ and $w(c) = k$. Using normality again implies that c cannot be the first element in its run of k's in w, and thus $w(c-1) = k$. Because of this, there is a chain C' from w to u all of whose labels are $c-1$. By construction $C' < C$ and $C \cap C' = \emptyset$, so C is a MSI as desired.

To show the plateau is minimal suppose, to the contrary, that there is a skipped interval $I \subset C$ and let $w = v_0$, $u = v_j$. Note that because $|v_0| = |v_1| = \ldots = |v_{j-1}|$, the Same Length Lemma applies to show that there is only one chain (namely an interval of C) between any two of these compositions. Thus the chain C' giving rise to I must rejoin C at an embedding η'_u of u into w, and hence also contain v_1 in order to cut out a proper subinterval. From this and normality of η_u we have

$$\eta'_u(c) < k = w(c-1) = \eta_u(c-1). \tag{13}$$

Also, $|u| = |w| - 1$ implies that C' must zero out exactly one element of w. Since $C' < C$, that element must be in a position strictly to the left of position c, but then because η_u and η'_u are both expansions of u we are forced to have $\eta_u(c-1) = \eta'_u(c)$, contradicting (13).

Now we know that $\mathcal{I}(C)$ consists of the descents and plateaus of C which are disjoint and cover C by their definition, so $\mathcal{J}(C) = \mathcal{I}(C)$ and C is critical by Theorem 4.2. To count the number of MSI's, note that if c is a position counted by $d(\eta_u)$ then the vertex just before the edge labeled c in C will be a descent, unless that edge is the very first one. On the other hand, if c is counted by $D(\eta_u)$ then the run of c's in $l(C)$ contribute both a plateau and a descent just before the plateau to $\mathcal{I}$ (unless the run is at the beginning of C when only the plateau will be an interval). In this manner we count each MSI exactly once for a total of $d(\eta_u) + 2D(\eta_u) - 1$ intervals. $\square$

Proposition 5.6. *If C is weakly decreasing and ends at an embedding η_u which is not normal then C is not critical.*

Proof. As in the proof of the previous proposition, it suffices to consider the case where $l(C)$ consists of a label c repeated $k \geq 2$ times so that $w(c) = \eta_u(c) + k$. If $\eta_u(c) > 0$ then $|w| = |u|$ and so the Same Length Lemma applies to show that C is the only chain from w to c. In particular, it is the lexicographically first chain and thus not critical.

If $\eta_u(c) = 0$ then $w(c) = k$ and, since η_u is not normal, it must be that c is the first index in this run of k's in w. Let

$$\eta_u(c-1) = w(c-1) = h \neq k. \tag{14}$$

To demonstrate that C is not critical, it suffices to show that there is no MSI I containing the element v_1 in C. Suppose, to the contrary that such an interval I exists and let C' be a chain giving rise to I. Using the Same Length Lemma as in the proof of Proposition 5.5 (third paragraph) we see that C' must rejoin C at u and this forces $I = C(w, u)$.

To finish the proof, it suffices to find a skipped interval $I'' \subset I$ since that will contradict the minimality of I. Let b be the smallest label in $l(C')$. Then $b < c$ since $l(C') < l(C) = (c^k)$. The same argument used at the end of the third paragraph of the previous proposition together with equation (14) give

$$\eta'_u(c) = \eta_u(c-1) = h \neq k = w(c).$$

Since the parts of a composition can only (weakly) decrease along a chain, we must have $\eta'_u(c) < w(c)$. It follows that c is a label on C'. Now consider any permutation of $l(C')$ which starts $l'' = (c, b, ...)$ and let C'' be the chain specified by l''. Then by construction and the Chain Specification Lemma $l(C'') \leq l'' < l(C)$. This implies that $C'' < C$ in PL-order and, by construction again, C'' contains v_1. Thus no MSI of C can contain v_1, a contradiction. $\square$

Our third and final proposition completes the proof of Theorem 5.4

Proposition 5.7. *If C is not weakly decreasing then C is not critical.*

Proof. If C is not weakly decreasing then it has an ascent v. It suffices to show that v is in no MSI. Suppose, to the contrary, that v is in an MSI $I = C(v_i, v_j)$. Then by the Descent Lemma, I contains no descents and so C is weakly increasing from v_i to v_j. As in the previous two proofs, it is no loss of generality to assume $v_i = w$ and $v_j = u$ so that C is itself an MSI.

Since C is an MSI, it is not the first chain in $[u, w]$. That first chain is the unique weakly increasing chain which ends at the rightmost embedding ρ_u of u into w. Thus if η_u is the embedding defined by C then we must have $\eta_u \neq \rho_u$.

For η_u define

$$z_\eta(j) = \#\{i \leq j \mid \eta_u(i) = 0\}$$

and similarly define $z_\rho(j)$ for ρ_u. Because ρ_u is rightmost we always have $z_\rho(j) \geq z_\eta(j)$ with equality when $j = |w|$. But ρ_u is not equal to η_u, so there is an index a such that $z_\rho(a) > z_\eta(a)$. Thus there is a first index $c > a$ such that $z_\rho(c) = z_\eta(c)$. This definition of c forces $\eta_u(c) = 0$ and $\rho_u(c) = k$ for some $k > 0$. But η_u and ρ_u are embeddings of the same composition, so there must be some index b with $a \leq b < c$ such that $\eta_u(b) = \rho_u(c) = k$ and $\eta_u(i) = 0$ for $b < i \leq c$.

We can now derive a contradiction by constructing a smaller skipped interval in C as follows. We have $w(c) \geq \rho_u(c) = k$ and $\eta_u(c) = 0$. Since C is weakly increasing, the labels equal to c must occur as a plateau. Therefore there must be vertices $w', u' \in C$ such that $I' = C(w', u')$ satisfies $\eta_{w'}(c) = k$, $\eta_{u'}(c) = 0$, and $l(I') = (c^k)$. But

$$\eta_{w'}(i) = \eta_{u'}(i) = \eta_u(i) = \begin{cases} k & \text{if } i = b, \\ 0 & \text{if } b < i < c, \end{cases}$$

so there is a chain C' from w' to u' with $l(C') = (b^k)$. Since $b < c$, I' is a skipped interval and we have obtained the desired contradiction. $\square$

We can now rederive the formula for $\mu(u, w)$ in $\mathbb{P}^*$. Combining equation (10) with Theorems 4.1, 4.2, and 5.4 we obtain

$$\mu(u, w) = \tilde{\chi}(u, w) = \sum_C (-1)^{\dim \sigma_C} = \sum_{\eta_u} (-1)^{d(\eta_u) + 2D(\eta_u) - 2} = \sum_{\eta_u} (-1)^{d(\eta_u)}$$

where the first sum is over all critical chains C in (u, w) and the other two are over all normal embeddings η_u into w.

We end this section by remarking that the Morse method can be used as a powerful tool not just for proving theorems but for discovering the correct statement to be proved. The reader may have found our definition of a normal embedding somewhat ad hoc. However, by starting with the very natural chain labeling used above and looking at the critical chains, one is quickly led to this definition in order to characterize the embeddings at which such chains end. Similarly, the defect may seem to have come out of nowhere, but in order to determine the dimension of the critical cells one is forced to define this quantity as well as its big brother $D(\eta_u)$.

6 Generalized subword order

We can now generalize both our result and Björner's as follows. Let $(P, \leq_P)$ be any poset. The *generalized subword order over P* is the partial order on P^* obtained by setting $u \leq_{P^*} w$ if w contains a subsequence $w(i_1), w(i_2), \ldots, w(i_{|u|})$ such that $u(j) \leq_P w(i_j)$ for $1 \leq j \leq |u|$. We get ordinary subword order when $P = A$ is an antichain and we get the composition poset when $P = \mathbb{P}$.

It is a simple matter to recast this generalized order in terms of embeddings. Let $\hat{0}$ be a special element which is not in P and let $\hat{P}$ be the poset obtained by adjoining $\hat{0}$ as a minimum element, i.e., $\hat{0} <_{\hat{P}} x$ for all $x \in P$. Then the definitions of support and expansion are as usual, just replacing 0 with $\hat{0}$. An *embedding* of u into w is a length $|w|$ expansion η_u of u with

$$\eta_u(i) \leq_{\hat{P}} w(i) \text{ for } 1 \leq i \leq |w|.$$

As expected, $u \leq_{P^*} w$ if and only if there is an embedding of u into w.

Finding an analogue of normality in this context is more delicate, and so far we have only been able to do it for a special class of posets. However, there is evidence that more general results are possible; the next section contains a discussion of this issue. First note that the definition of a run carries over verbatim to any P^*. Now call P a *rooted tree* if its Hasse diagram is a tree with a minimum element. A *rooted forest* is a poset where each connected component of its Hasse diagram is a rooted tree. Note that both antichains and chains are rooted forests. Note also that if P is a rooted forest then $\hat{P}$ is a rooted tree so the following definition makes sense. If $x \in P$ where P is a rooted forest then let x^- be the element adjacent to x on the unique path from x to $\hat{0}$ in the Hasse diagram for $\hat{P}$. For a rooted forest, a *normal embedding* of u into w is an embedding η_u into w satisfying two conditions.

1. For $1 \leq i \leq |w|$ we have $\eta_u(i) = w(i)$, $w(i)^-$, or $\hat{0}$.

2. For all $x \in P$ and every run $[r, t]$ of x's in w, we have

   (a) $(r, t] \subseteq \operatorname{Supp} \eta_u$ if x is minimal in P,

   (b) $r \in \operatorname{Supp} \eta_u$ otherwise.

Finally, we need the definition of *defect* in this situation, which is as expected:

$$d(\eta_u) = \#\{i \mid \eta_u(i) = w(i)^-\}$$

for a normal embedding η_u into w. The following theorem is the promised generalization of Theorems 2.1 and 2.2. Both of the two proofs we have given of the special case where $P = \mathbb{P}$ generalize easily, with the minimal elements in P playing the rôle of 1 and the rest functioning like the integers $k \geq 2$.

Theorem 6.1. *Let P be a rooted forest. Then the Möbius function of P^* is given by*

$$\mu(u, w) = \sum_{\eta_u} (-1)^{d(\eta_u)},$$

where the sum is over all normal embeddings η_u of u into w. □

7 Comments and open problems

There are several possible avenues for future research. We discuss some of them here.

7.1 Generating Functions

As mentioned in the introduction, Björner and Reutenauer [6] gave another proof of the formula for μ in A^* using generating functions on monoids. Let $\mathbb{Z}\langle\langle A\rangle\rangle$ denote the algebra of formal series using the elements in A as noncommutative variables and the integers as coefficients. Such a series can be written

$$f = \sum_{w \in A^*} c_w w$$

for certain $c_w \in \mathbb{Z}$. For example, given $u \in A^*$ one can consider the series

$$m(u) = \sum_{w \geq u} \binom{w}{u}_n w. \tag{15}$$

Björner and Reutenauer showed that (15) is rational for any u and obtained, upon specialization of the variables, nice expressions for various ordinary generating functions associated with the Möbius function of A^*. They also derived results for the zeta function of A^*. The map $m : A^* \to \mathbb{Z}\langle\langle A\rangle\rangle$ can be extended to a continuous linear endomorphism of $\mathbb{Z}\langle\langle A\rangle\rangle$. In fact, the full incidence algebra of A^* is isomorphic to a subalgebra of this endomorphism algebra. Björner and Reutenauer gave another proof of Theorem 2.1 using this fact.

It is natural to try and apply these ideas to $\mathbb{P}^*$, and more generally to rooted forests. This has been done by Björner and Sagan [7].

7.2 The poset of permutations

Our original interest in $\mathbb{P}^*$ came from the rapidly growing subject of permutation patterns. For an overview of permutation patterns the reader is referred to Bóna's text [10]. Let S_n denote the nth symmetric group and let $\pi \in S_n$ and $\sigma \in S_k$. We say that π *contains a σ-pattern*, and write $\pi \geq \sigma$, if there are indices $i_1 < i_2 < \cdots < i_k$ such that the subsequence

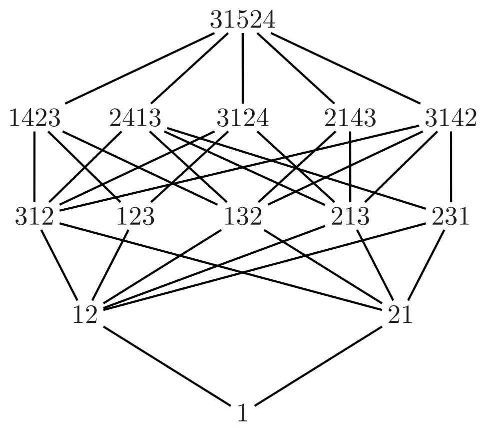

Figure 3: The Hasse diagram for the interval $[1, 31524]$ in the pattern containment ordering on permutations

$\pi(i_1)\pi(i_2)\ldots\pi(i_k)$ has the same pairwise comparisons as $\sigma(1)\sigma(2)\ldots\sigma(k)$. This subsequence is called a *copy* of σ in π. For example, $312 \leq 24153$ because of the copy 413. This defines a partial order on the set of all finite permutations. Wilf was the first to ask the following question.

Question 7.1 (Wilf [27]). *What can be said about the Möbius function of permutations under the pattern-containment ordering?*

Given two permutations $\pi \in S_m$ and $\sigma \in S_n$, their *direct sum* is the permutation of length $m+n$ whose first m elements form σ and whose last n elements are the copy of π gotten by adding m to each element of π. For example, $132 \oplus 32145 = 13265478$. A permutation is said to be *layered* if it can expressed as the direct sum of some number of decreasing permutations. (An equivalent characterization of layered permutations is that they are the permutations that contain neither a 231-pattern nor a 312-pattern.) Our previous example is layered because $13254378 = 1 \oplus 21 \oplus 321 \oplus 1 \oplus 1$. Clearly the set of layered permutation of length n is in bijection with the set of compositions of n. Almost as clearly, this bijection sends the pattern-containment order to the composition order we have considered, so Theorem 2.2 answers Wilf's question for the set of layered permutations.

Any normal embedding approach to describing the Möbius function for permutations in general would need to incorporate non-unitary weights, as witnessed by the fact that $\mu(1, 31524) = 6$.

7.3 Factor order

Subword order is not the only partial order on the set of words. We say that the word u is a *factor* of the word w if there exist (possibly empty) words v_1 and v_2 so that $w = v_1uv_2$, or in other words, if u occurs as a contiguous subword in w. Björner [5] showed that the Möbius function for factor order only takes on values in $\{0, \pm 1\}$ and gave a recursive rule that allows the computation of $\mu(u, w)$ in $O(|w|^2)$ steps.

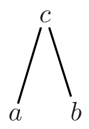

Figure 4: The Hasse diagram for the poset Λ

The factor order can be defined on P^* for any poset P: we say that u is a factor of w if there are words v_1, v_2, v_3 such that:

1. $w = v_1 v_2 v_3$,
2. $|v_2| = |u|$,
3. $u(i) \leq v_2(i)$ for all $1 \leq i \leq |u|$.

Indeed, this is one of the orders on $\mathbb{P}^*$ studied by Snellman [21, 22]. The Möbius function of $\mathbb{P}^*$ under factor order remains unknown.

7.4 Subwords over Λ

The smallest poset to which Theorem 6.1 is inapplicable is the poset Λ depicted in Figure 4. The Möbius function of Λ^* seems to be quite interesting. In fact, numerical evidence points to a surprising connection with the Tchebyshev polynomials of the first kind, $T_n(x)$, which can be defined as the unique polynomials such that $T_n(\cos\theta) = \cos(n\theta)$.

Conjecture 7.2. *For all $i \leq j$, $\mu(a^i, c^j)$ is the coefficient of x^{j-i} in $T_{i+j}(x)$.*

As with the poset of permutations, a normal embedding interpretation of $\mu(a^i, c^i)$ would need to use weights because, for example, $\mu(a, cc) = -3$.

One possible way to attack this conjecture would be to use the three-term recurrence for $T_n(x)$. Translating this in terms of the conjecture, it would suffice to show that

$$\mu(a^i, c^j) = 2\mu(a^i, c^{j-1}) - \mu(a^{i-1}, c^{j-1})$$

for $j \geq i \geq 1$ However, we have not been able to see any relationship between the intervals $[a^i, c^j]$, $[a^i, c^{j-1}]$, and $[a^{i-1}, c^{j-1}]$ which would permit us to derive this relation for their Möbius functions.

There are two closely related areas where the Tchebyshev polynomials have appeared. A permutation π is said to *avoid* a permutation σ if it does not contain a σ-pattern. Chow and West [11] showed that the generating function for the number of permutations in S_n avoiding both 132 and $12\ldots k$ for fixed k can be expressed in terms of Tchebyshev polynomials of the second kind. Mansour and Vainshtein [19] extended this result to count permutations avoiding 132 and containing exactly r copies of $12\ldots k$.

More recently, Hetyei [17] defined poset maps T and U which he called Tchebyshev transformations of the first and second kind. This is because when applied to the ladder poset L_n, the **cd**-index of the images can be expressed in terms of $T_n(x)$ and $U_n(x)$. Since the **cd**-index is related to the Möbius function, it is conceivable that Hetyei's map could be used to prove our conjecture, but the posets $T(L_n)$ are not isomorphic to any of our intervals $[a^i, c^j]$ in general, so it is not clear how to proceed. However, these maps are very interesting in their own right and have been further studied by Ehrenborg and Readdy [12].

Acknowledgment. We are indebted to Patricia Hersh for useful discussions and references.

References

[1] Babson, E., and Hersh, P. Discrete Morse functions from lexicographic orders. *Trans. Amer. Math. Soc. 357*, 2 (2005), 509–534 (electronic).

[2] Bergeron, F., Bousquet-Mélou, M., and Dulucq, S. Standard paths in the composition poset. *Ann. Sci. Math. Québec 19*, 2 (1995), 139–151.

[3] Björner, A. Shellable and Cohen-Macaulay partially ordered sets. *Trans. Amer. Math. Soc. 260*, 1 (1980), 159–183.

[4] Björner, A. 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, pp. 118–124.

[5] Björner, A. The Möbius function of factor order. *Theoret. Comput. Sci. 117*, 1-2 (1993), 91–98.

[6] Björner, A., and Reutenauer, C. Rationality of the Möbius function of subword order. *Theoret. Comput. Sci. 98*, 1 (1992), 53–63.

[7] Björner, A., and Sagan, B. E. Rationality of the Möbius function of the composition poset. arXiv:math.CO/0510282.

[8] Björner, A., and Stanley, R. P. An analogue of Young's lattice for compositions. arXiv:math.CO/0508043.

[9] Björner, A., and Wachs, M. Bruhat order of Coxeter groups and shellability. *Adv. in Math. 43*, 1 (1982), 87–100.

[10] Bóna, M. *Combinatorics of permutations.* Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, 2004.

[11] Chow, T., and West, J. Forbidden subsequences and Chebyshev polynomials. *Discrete Math. 204*, 1-3 (1999), 119–128.

[12] Ehrenborg, R., and Readdy, M. The Tchebyshev transforms of the first and second kinds. arXiv:math.CO/0412124.

[13] Farmer, F. D. Cellular homology for posets. *Math. Japon. 23*, 6 (1978/79), 607–613.

[14] Forman, R. A discrete Morse theory for cell complexes. In *Geometry, topology, & physics*, Conf. Proc. Lecture Notes Geom. Topology, IV. Internat. Press, Cambridge, MA, 1995, pp. 112–125.

[15] Forman, R. Morse theory for cell complexes. *Adv. Math. 134*, 1 (1998), 90–145.

[16] Hersh, P. On optimizing discrete Morse functions. arXiv:math.CO/0311270.

[17] Hetyei, G. Tchebyshev posets. *Discrete Comput. Geom. 32*, 4 (2004), 493–520.

[18] Kruskal, J. B. The theory of well-quasi-ordering: A frequently discovered concept. *J. Combinatorial Theory Ser. A 13* (1972), 297–305.

[19] Mansour, T., and Vainshtein, A. Restricted permutations, continued fractions, and Chebyshev polynomials. *Electron. J. Combin. 7* (2000), Research Paper 17, 9 pp. (electronic).

[20] Rota, G.-C. On the foundations of combinatorial theory. I. Theory of Möbius functions. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete 2* (1964), 340–368 (1964).

[21] Snellman, J. Saturated chains in composition posets. arXiv:math.CO/0505262.

[22] Snellman, J. Standard paths in another composition poset. *Electron. J. Combin. 11*, 1 (2004), Research Paper 76, 8 pp. (electronic).

[23] Stanley, R. P. *Enumerative combinatorics. Vol. 1*, vol. 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997.

[24] Viennot, G. Maximal chains of subwords and up-down sequences of permutations. *J. Combin. Theory Ser. A 34*, 1 (1983), 1–14.

[25] Wang, T. M., and Ma, X. R. A generalization of the Cohen-Macaulay property of the Möbius function of a word poset. *Acta Math. Appl. Sinica 20*, 3 (1997), 431–437.

[26] Warnke, I. The Möbius-function of subword orders. *Rostock. Math. Kolloq.*, 46 (1993), 25–31.

[27] Wilf, H. S. The patterns of permutations. *Discrete Math. 257*, 2-3 (2002), 575–583.