The Möbius Function
of a
Restricted Composition Poset

Adam M. Goyt[*]

Department of Mathematics

Minnesota State University Moorhead

1104 7th Avenue South

Moorhead, Minnesota 56563

goytadam@mnstate.edu

www.mnstate.edu/goytadam

April 19, 2012

Abstract

We study a poset of compositions restricted by part size under a partial ordering introduced by Björner and Stanley. We show that our composition poset C_{d+1} is isomorphic to the poset of words A_d^*. This allows us to use techniques developed by Björner to study the Möbius function of C_{d+1}. We use counting arguments and shellability as avenues for proving that the Möbius function is $\mu(u, w) = (-1)^{|u|+|w|} \binom{w}{u}_{dn}$, where $\binom{w}{u}_{dn}$ is the number of d-normal embeddings of u in w. We then prove that the formal power series whose coefficients are given by the zeta and the Möbius functions are both rational. Following in the footsteps of Björner and Reutenauer and Björner and Sagan, we rely on definitions to prove rationality in one case, and in another case we use finite-state automata.

1 Introduction

Let $\mathbb{P} = \{1, 2, 3, \ldots\}$, $[n] = \{1, 2, 3, \ldots, n\}$ and $[i, j] = \{i, i+1, i+2, \ldots, j\}$. A *composition* α is an element $(\alpha_1, \alpha_2, \ldots, \alpha_k) \in \mathbb{P}^k$. The *length* of a composition $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$ is

---

$\ell(\alpha) = k$, and the *norm*, $|\alpha| = n$, is n if $\sum_{i=1}^{k} \alpha_i = n$. The α_i are called *parts*. If $|\alpha| = n$, then we say that α is a composition of n. Define C_n to be the set of all compositions of n and let

$$C = \bigcup_n C_n.$$

The partial ordering on C that we are interested in was introduced by Björner and Stanley and can be found in [7]. To define the poset C we need only define its covering relation $\alpha \prec \beta$. In C we say that $\beta \succ \alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$ if β is of one of the two following forms:

$$\beta = (\alpha_1, \alpha_2, \ldots, \alpha_{i-1}, \alpha_i + 1, \alpha_{i+1}, \ldots, \alpha_k)$$
$$\beta = (\alpha_1, \alpha_2, \ldots, \alpha_{i-1}, \alpha_i + 1 - h, h, \alpha_{i+1}, \ldots, \alpha_k),$$

where $h \leq \alpha_i$.

We will consider the subposet C_d, which is the set of compositions with part sizes at most d, that is $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k)$ is an element of C_d if and only if $\alpha_i \leq d$ for $1 \leq i \leq k$. We will show that this poset is isomorphic to a poset of words in order to determine the Möbius function of C_d.

Let A^* be the free monoid under concatenation of $A = \{a, b\}$. We think of A^* as the set of all words that can be created from the *alphabet* A. The identity in A^* is the empty word ϵ. We say that the length of a word $u = u_1 u_2 \ldots u_k$ is $|u| = k$. Let $u = u_1 u_2 \ldots u_k$ and $w = w_1 w_2 \ldots w_\ell$ be words in A^*. We make A^* into a poset by letting $u \leq w$ if there exist $i_1 \leq i_2 \leq \ldots \leq i_k$ such that $u_j = w_{i_j}$ for $1 \leq j \leq k$. We call the set $\iota = \{i_1, i_2 \ldots, i_k\}$ an *embedding* of u in w and let $w_\iota = w_{i_1} w_{i_2} \ldots w_{i_k}$. If ι is an embedding of u in w, then we say that w_j is *supported* by u in ι if $j \in \iota$. For example, the word *abaab* is a subword of $w = aabbababb$, since $w_2 w_3 w_5 w_7 w_8 = abaab$, and w_3 is supported in this embedding.

Now, let $\phi : \mathbb{P} \to A^*$ be given by $\phi(k) = a \underbrace{bb \ldots b}_{k-1}$. Given any composition $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k) \in C$, , let $\phi(\alpha) = \phi(\alpha_1)\phi(\alpha_2) \cdots \phi(\alpha_k)$ omitting the initial a from $\phi(\alpha_1)$. Clearly, $\phi : C \to A^*$ is a bijection.

Lemma 1.1 *The map ϕ is an isomorphism of A^* and C as partially ordered sets.*

Proof: Since we have already established ϕ as a bijection between A^* and C, it is enough to show that ϕ preserves the partial orderings.

Let $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_k) \prec \beta = (\beta_1, \beta_2, \ldots, \beta_\ell) \in C$. If $\ell = k$ then, for some j, $\beta_j = \alpha_j + 1$ and $\beta_i = \alpha_i$ for $i \neq j$. In this case $\phi(\beta)$ is obtained from $\phi(\alpha)$ by inserting a b into $\phi(\alpha)$ anywhere between the $(j-1)^{st}$ and j^{th} occurrence of an a. Thus, $\phi(\alpha)$ is a subword of $\phi(\beta)$.

If $\ell = k + 1$ then $\beta_j = \alpha_j + 1 - h$ for some $h \leq \alpha_j$, $\beta_{j+1} = h$ and $\beta_i = \alpha_i$ for $i \neq j$ or $j + 1$. In this case $\phi(\beta)$ is obtained from $\phi(\alpha)$ by inserting an a between the $(j-1)^{st}$ and j^{th} occurrence of an a, so that $h - 1$ b's follow the inserted a.

Thus $\beta \succ \alpha$ implies $\phi(\beta) \succ \phi(\alpha)$. The converse is proved similarly. $\square$

Let A_d^* be the subposet of A^* consisting of all words that do not have $d+1$ consecutive b's. Notice that ϕ restricts to an isomorphism between C_{d+1}, as defined above, and the subposet, A_d^*. Much is known about partially ordered sets on words and subword order, so it will be convenient to work with the poset A_d^* to understand the poset C_{d+1}.

We assume familiarity with the Möbius function of a locally finite poset [6]. To understand the Möbius function of the poset A_d^* we will adapt the Möbius function for A^*. In [2], Björner shows that the Möbius function for A^* with the subword order described above is given by

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

where $\binom{w}{u}_n$ is the number of normal embeddings of u in w. We will alter his definition of normal embeddings to define d-normal embeddings and show that in the poset A_d^* we have the following.

Theorem 1.2 *For elements u, $w \in A_d^*$, we have*

$$\mu(u, w) = (-1)^{|u|+|w|} \binom{w}{u}_{dn},$$

where $\binom{w}{u}_{dn}$ is the number of d-normal embeddings of u in w.

We now define a d-normal embedding and right-most embedding. We will call a set $[r, s]$ a *run* if $w_r = w_{r+1} = \ldots = w_s$. Let the *repetition set* of w be $R(w) = \{j : w_j = w_{j-1}\}$. An embedding $\iota = \{i_1, i_2, \ldots, i_k\}$ of u in w is called *d-normal* if

  a $R(w) \subseteq \iota$, and

  b if u has a run of d b's and w_{i_j} corresponds to the first b in this run then $i_j = 1$ or $w_{i_j-1} = a$ and $i_j - 1 \in \iota$ or $i_j - 1 = 1$.

This differs from a normal embedding in that a normal embedding must only satisfy part a of the definition of a d-normal embedding.

The *right-most embedding* of u in w is the unique embedding $\{j_1, j_2 \ldots j_r\}$ such that for any other embedding $\{i_1, i_2 \ldots i_r\}$ of u in w we have that $i_\ell \leq j_\ell$, for each $1 \leq \ell \leq r$.

The definition of a d-normal embedding, while somewhat cumbersome, is as simple as possible and is borne out of the proof involving shellability of the poset. For an explanation of how the definition was borne, please see the end of Section 3.

An example is in order to help describe the definition of d-normal. Consider the poset A_3^* and the interval of this poset $[bbb, bbabb]$. Let $u = bbb$ and $w = bbabb$. In any d-normal embedding of u in w the b's located in positions 2 and 5 must be supported according to condition a. This means that the embeddings $\{1, 2, 5\}$ and $\{2, 4, 5\}$ are the only possible d-normal embeddings. Now, u contains a run of three b's, so by a condition b $\{2, 4, 5\}$ cannot be a 3-normal embedding because w_2 corresponds to the first b in u, and $w_1 \neq a$. Thus, $\{1, 2, 5\}$ is the only d-normal embedding of u in w. So $\mu(u, w) = (-1)^{3+5} \cdot 1 = 1$. A quick sketch of the interval $[bbb, bbabb]$ in A_3^* confirms this.

In Section 2 we prove Theorem 1.2 using bijective techniques, and in Section 3 we prove it using shellability. The proofs in Sections 2 and 3 are adaptations of similar proof given by Björner. In Section 4 we show that the zeta function and Möbius function are rational. Finally, in the last section we determine a generating function in commuting variables for the zeta function and the Möbius function.

2 Combinatorial Proof

We remind the reader that the Möbius function of a poset is the unique function satisfying the following.

- $\mu(x, x) = 1$

- $\sum_{x \leq z \leq y} \mu(z, y) = 0$

- $\mu(x, y) = 0$ for $x \not\leq y$

Proof of Theorem 1.2: This proof is essentially the same as the proof of Theorem 3.1 from [3]. We will show that the function $f(u, w) = (-1)^{|w| + |u|} \binom{w}{u}_{dn}$ satisfies the three conditions above and hence, must be μ.

Clearly, $f(u, u) = 1$, and $f(u, w) = 0$ for $u \not\leq w$. It remains to show that $\sum_{u \leq v \leq w} f(v, w) = 0$ for any $v \leq w$.

Suppose that $|w| = n$ and $|u| = k$. Let N be the set of embeddings ι such that ι is a d-normal embedding of w_ι in w, and $u \leq w_\iota$. Let $N_e = \{\iota \in N : |\iota| \text{ is even}\}$ and $N_o = \{\iota \in N : |\iota| \text{ is odd}\}$. Then

$$\sum_{u \leq v \leq w} f(v, w) = (-1)^n (\#N_e - \#N_o).$$

A bijection between N_e and N_o will complete the proof.

For each embedding $\iota \in N$ let ι_f be the minimal element of $[n]$ such that ι_f is not in the right-most embedding of u in w_ι. Define $\psi : N \to N$ by

$$\psi(\iota) = \begin{cases} \iota \cup \{\iota_f\} & \text{if } \iota_f \notin \iota, \\ \iota - \{\iota_f\} & \text{if } \iota_f \in \iota. \end{cases}$$

We will show that ψ is well-defined. If $\iota_f \notin \iota$ then, since $R(w) \subseteq \iota$, we must have $R(w) \subseteq \iota \cup \{\iota_f\}$. If $\iota_f \in \iota$ then we want to show that $\iota_f \notin R(w)$. If $\iota_f \in R(w)$ then $w_{\iota_f} = w_{\iota_f - 1}$. Also $\iota_f - 1$ must be in the right-most embedding of u in w_ι. This is impossible by the definition of right-most embedding.

We must show that $w_{\psi(\iota)}$ is still an element of A_d^*. If ψ removes an a creating a run of b's of length $d + 1$ in $w_{\psi(\iota)}$ then the first b in the run of b's immediately preceding a cannot be in the right-most embedding of u in w_ι. Thus, the a would never have been removed by ψ.

If ψ inserts a b creating a run of b's of length $d + 1$ in $w_{\psi(\iota)}$ then w_ι has a run of b's of length d. Let the first b of this run be in position i_t. If $i_t = 1$ then a b must have been inserted at the beginning of w_ι, which is impossible by part b of the definition of d-normal. If $i_t \neq 1$, then, by part b of the definition of d-normal, we must have $w_{i_t-1} = a$ and $i_t - 1 \in \iota$ or $\iota_t - 1 = 1$ rendering the creation of this run impossible.

Since ψ is its own inverse and changes the parity of $|\iota|$, the proof is complete. $\square$

3 Shellability

We now give an alternate proof that $\mu(u, w) = (-1)^{|w|+|u|} \binom{w}{u}_{dn}$ by showing that A_d^* is dual CL-shellable. Throughout this section we assume familiarity with lexicographical shellability. The same labeling that Björner used in [2] to determine the Möbius function of A^* will work for our poset A_d^*.

Let $[u, w]$ be an interval of A^* and $\mathcal{E}^*(A^*)$ be the set of edges of maximal chains, $m = (u = x_1, x_2, \ldots, x_k = w)$ in $[u, w]$. Let $\ell : \mathcal{E}^*(A^*) \to \mathbb{Z}$ be defined as follows. Label the edge (x_{k-1}, x_k, m) by $\ell(x_{k-1}, x_k, m) = i_1$, where i_1 is the smallest element of $[k]$ such that $[k] - \{i_1\}$ is an embedding of x_{k-1} in x_k. Label the edge (x_{k-2}, x_{k-1}, m) by $\ell(x_{k-2}, x_{k-1}, m) = i_2$, where i_2 is the least element of $[k] - \{i_1\}$ such that $[k] - \{i_1, i_2\}$ is an embedding of x_{k-2} in x_k. Repeat this process for the remaining edges. It is clear that if two maximal chains have the same first s edges then the labels on these edges will be the same. Figure 1 shows a chain-edge labeling for $[abb, aabbab]$.

Theorem 3.1 (Björner) *The map $\ell : \mathcal{E}^*(A^*) \to \mathbb{Z}$ described above is a dual CL-labeling for each interval of A^*.*

In Björners proof, [2], he shows that the unique increasing maximal chain is the one corresponding to the right-most embedding of u in w. Let $\{i_1, i_2, \ldots, i_k\}$ be the rightmost embedding of u in w, then the increasing maximal chain is the maximal chain labeled by the elements of $[n] - \{i_1, i_2, \ldots, i_k\}$ from smallest to largest. The unique increasing chain in Figure 1 is

$$abbaba \overset{1}{-} abbab \overset{3}{-} abab \overset{5}{-} abb.$$

It is also clear that since we are working with the right-most embedding, this chain must have the lexicographically smallest possible entries.

Each maximal chain of an interval $[u, w]$ in A_d^* is also a maximal chain of the same interval in A^*. For $u, w \in A_d^*$, we will denote by $[u, w]_d$ the interval $[u, w]$ in A^* restricted to A_d^*. We will label the edges of A_d^* in exactly the same way as we did for A^*.

Theorem 3.2 *The map $\ell_d : \mathcal{E}^*(A_d^*) \to \mathbb{Z}$ is a dual CL-labeling for each interval of A_d^*, and hence A_d^* is CL-shellable.*

Proof: We already know that there is exactly one increasing maximal chain, m_0, in the labeling of an interval $[u, w]$ in A^*, and that m_0 corresponds to the right-most embedding of u in w. This means that there is at most one increasing maximal chain among those

5

in the interval $[u, w]_d$. We must show that this chain is indeed in $[u, w]_d$. Unless a chain from $[u, w]$ passes through an element with a run of b's of length at least $d + 1$, that chain is in $[u, w]_d$.

Suppose that the unique increasing maximal chain for $[u, w]$ passes through an element with a run of b's of length at least $d + 1$ and hence is not in $[u, w]_d$. At some step in the labeling of this chain we must have $x_i \succ x_{i-1}$, where x_i does not have a run of b's of length at least $d + 1$, but x_{i-1} does. This may only happen if an a that was separating those b's in x_i is removed in this step. Since this chain corresponds to the right-most embedding of u in w and u and w do not have a run of b's of length at least $d + 1$, the a removed could not have been the left-most removable element. $\square$

Figure 1
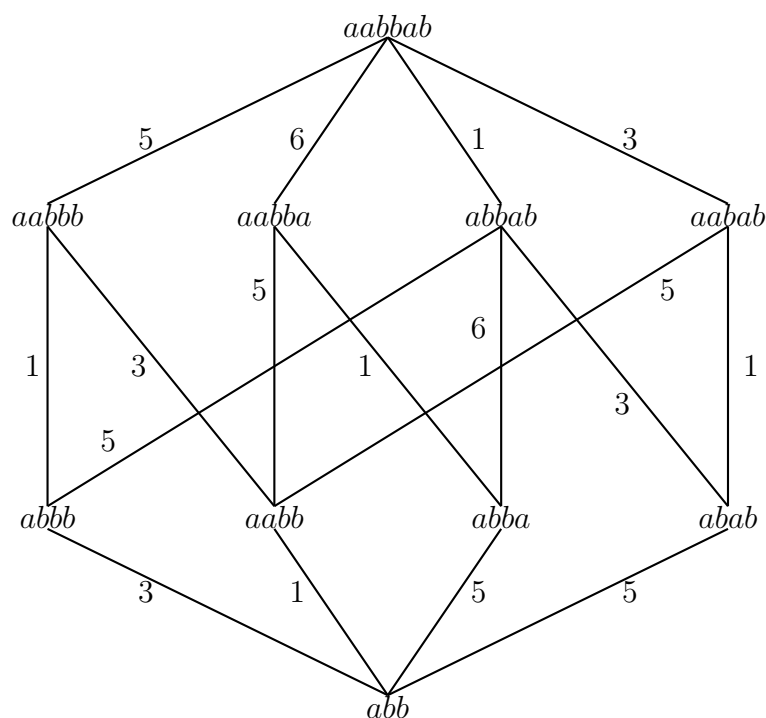
Proof of Theorem 1.2: It suffices to show that the number of decreasing chains in $[u, w]_d$ is the same as the number of d-normal embeddings of u in w.

Suppose $\iota = \{i_1, i_2 \ldots, i_k\}$ is an embedding of u in w such that there is a maximal chain m with $[n] - \{\ell_1(m), \ell_2(m), \ldots, \ell_{n-k}(m)\} = \iota$, where $\ell_j(m)$ is the label on the j^{th} edge of m. Then m is decreasing if m is obtained by deleting the entries w_j, $j \in [n] - \iota$, from right to left. The only way that this is possible is if for each $j \in [n] - \iota$, we have that $w_j \neq w_{j-1}$. This means that the embedding satisfies part a of the definition of a d-normal embedding.

To be sure that this chain does indeed appear in $[u, w]_d$ we look to part b of the definition of a d-normal embedding. The only way a decreasing maximal chain will appear

in $[u, w]$, but not $[u, w]_d$ is if it passes through an element that has a run of at least $d + 1$ b's. Since we are deleting elements from right to left, the first time an element with a run of at least $d + 1$ elements appears is when an a that is preceded by a b is deleted, and this a is followed by a run of d b's. This run of d b's must be a run in u, so by part b of the definition, a must either be supported, in which case it would not have been deleted, or it must have been the first letter of w in which case it is not preceded by a b.

Conversely, if $\iota = \{i_1, i_2, \ldots, i_k\}$ is a d-normal embedding of u in w then there is a decreasing maximal chain corresponding to ι given by deleting the elements $w_j, j \in [n] - \iota$, from right to left. $\square$

The definition of d-normal is borne out of the embeddings that give decreasing maximal chains in the CL-shelling of the poset. Part a of the definition is well understood from the work of Björner. We will focus on part b of the definition. Let u be a word with a run of d b's. Let ι be an embedding of u in w. Let $i_j, i_{j+1}, \ldots, i_{j+d-1} \in \iota$ be the embedding of this run of d b's in w. If we are deleting elements of w from right to left until we are left with u to create a decreasing maximal chain then at some point the remaining letters of w starting with w_{i_j} are all elements of u and we must delete something to the left of w_{i_j} (unless $i_j = 1$). Note that w_{i_j} need not be in a run of d b's in w.

There are two possibilities. Either this run of d b's in u is at the beginning of u or not. If the run is at the beginning of u then we delete the remaining elements of w one at a time from right to left. If any of these elements is a b then a string of $d + 1$ b's is formed, and hence no decreasing maximal chain corresponds to this embedding. Thus, there may only be an a preceding this b.

If the run of d b's is not at the beginning of u then in u an a immediately precedes the first b in this run. This a must correspond to w_{i_j-1}, otherwise we have a run of unsupported a's or there is a b between this a and w_{i_j} which will eventually create a run of $d + 1$ b's.

4 Rationality

We now consider the formal power series algebra $\mathbb{Z}\langle\langle A \rangle\rangle$ in the noncommuting variables of A over $\mathbb{Z}$. Every $f \in \mathbb{Z}\langle\langle A \rangle\rangle$ is of the form

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

where $c_w \in \mathbb{Z}$. We let $f^* = \epsilon + f + f^2 + \ldots = (\epsilon - f)^{-1}$ for any series $f \in \mathbb{Z}\langle\langle A \rangle\rangle$ with no constant term. Let $f^+ = f^* - \epsilon$. A series $f \in \mathbb{Z}\langle\langle A \rangle\rangle$ is called rational if it can be constructed from a finite number of monomials under a finite number of the usual algebraic operations in $\mathbb{Z}\langle\langle A \rangle\rangle$ and the $*$ operation. Clearly, f^+ is rational if f^* is.

We may also consider series of the form $f = \sum_{u \leq w} c_{(u,w)} u \otimes w$ where $u \otimes w$ just represents the ordered pair $(u, w) \in A^* \times A^*$. Rationality of such a series is defined similarly.

In [4], Björner and Reutenauer showed that the following four series are rational:

$$
\begin{aligned}
Z(u) &= \sum_{w \in A^*} \zeta(u, w) w, \\
M(u) &= \sum_{w \in A^*} \mu(u, w) w, \\
Z_\otimes &= \sum_{w \in A^*} \zeta(u, w) u \otimes w, \text{ and} \\
M_\otimes &= \sum_{w \in A^*} \mu(u, w) u \otimes w.
\end{aligned}
$$

In this section we will use methods similar to those used by Björner and Sagan [5] to do the same for the series $Z_d(u)$, $M_d(u)$, $Z_\otimes^d$, and $M_\otimes^d$, where these series are defined exactly the same way as those above replacing A^* by A_d^* in each. In the remainder of this section we will assume that $d = 3$. The reason that we let $d = 3$ is to avoid cumbersome formulas and definitions. Everything that is used to prove these facts for $d = 3$ has an obvious generalization.

We begin with $Z_3(u)$ and $M_3(u)$. A function $f : S^* \to T^*$ where S and T are finite sets is *multiplicative* if for $u = u_1 u_2 \ldots u_m \in S^*$, we have that $f(u) = f(u_1) f(u_2) \cdots f(u_m)$. Let $A_3 = \{a, ab, abb, abbb\}$. Let $B = \epsilon + b + bb + bbb$, and notice that $B(A_3)^* = A_3^*$. Each word $u \in A_3^*$ can be broken uniquely into its maximal runs of a's and b's. Define the multiplicative function

$$
z : A_3^* \to \mathbb{Z} \langle\langle A_3 \rangle\rangle ,
$$

by

$$
\begin{aligned}
z(a^k) &= (aB)^{k-1} a, \\
z(b) &= (B - \epsilon) a^*, \\
z(bb) &= (B - \epsilon) a^+ b a^* + (B - b - \epsilon) a^*, \text{ and} \\
z(bbb) &= (B - \epsilon) a^+ b a^* b a^* + (B - b - \epsilon) a^+ b a^* + bbb a^*.
\end{aligned}
$$

If a run of a's is at the end of the word then let $z(a^k) = (A_3)^k$ for this run of a's. Let $p_z(u)$ be the *prefix* of $Z(u)$ where,

$$
p_z(u) = \begin{cases} A_3^* & u \text{ begins with } a, \\ (A_3^* a + \epsilon) & u \text{ begins with } b. \end{cases}
$$

Define the multiplicative function

$$
m : A_3^* \to \mathbb{Z} \langle\langle A_3 \rangle\rangle ,
$$

by

8

$$
\begin{aligned}
m(a^k) &= (ab)^*(\epsilon - a)a\,[(\epsilon - b(ab)^*(\epsilon - a))a]^{k-1} \\
m(b) &= (\epsilon - b(ab)^*(\epsilon - a))\,b, \\
m(bb) &= (\epsilon - b(ab)^*(\epsilon - a))\,b(ab)^*(\epsilon - a)b, \text{ and} \\
m(bbb) &= b\,((ab)^*(\epsilon - a)b)^2.
\end{aligned}
$$

Let $p_m(u)$ be the *prefix* of $M(u)$ where,

$$
p_m(u) = \begin{cases} (\epsilon - b) & \text{if } u \text{ begins with } a, \\ (\epsilon - a) & \text{if } u \text{ begins with } b, bb \text{ or } bbb. \end{cases}
$$

Let $s_m(u)$ be the *suffix* of $M(u)$ where,

$$
s_m(u) = \begin{cases} (ba)^*(\epsilon - b) & \text{if } u \text{ ends in } a, \\ (ab)^*(\epsilon - a) & \text{if } u \text{ ends in } b. \end{cases}
$$

Lemma 4.1 *For any $u \in A_3^*$ we have that*

$$
Z_3(u) = p_z(u)z(u),
$$

and

$$
M_3(u) = p_m(u)m(u)s_m(u).
$$

Proof: We begin by proving the statement for $Z_3(u)$. We must show that the functions p_z and z will produce each word that contains u exactly once. This is done by showing that for $w \geq u$, $z(u)$ produces the right-most embedding of u in w. We will explain how this works for $z(a^k)$ and $z(b^\ell)$, how $z(a^k)$ interacts with $z(b^\ell)$ and how the prefix works.

The last a in the run a^k is given by the a at the end of $z(a^k)$. This is clearly under the right-most possible a if we just focus on $z(a^k)$. The preceding $k - 1$ a's are each followed by $B = \epsilon + b + bb + bbb$, so that each word produced by $z(a^k)$ contains a^k exactly once. Notice that any word produced by $z(a^k)$ begins with a supported a and ends with a supported a.

We now turn our focus to $z(b^\ell)$. We'll discuss $z(bb)$ as the others are similar. The first term in $z(bb)$ produces words where the two supported b's are separate, achieved by the a^+. At the beginning of each term in $z(bb)$ we can have a run of b's. The last b (first term) and the last two b's (second term) are the supported b's. Finally at the end of each term we have a^* because we can follow our last supported b by as many a's as we like and still maintain a right-most embedding.

The series $z(a^k)$ and $z(b^\ell)$ interact in the following way. The last a in $z(a^k)$ is always supported, so if a run of b's follows a run of a's in u then a right-most embedding is still maintained. The last supported b in $z(b^\ell)$ is followed by a^*, so if a run of a's follows a run of b's in u then we maintain a right most embedding. It is worth mentioning that we have not violated the fact that a run of b's may be of length at most three.

We may essentially place any word from A_3^* at the beginning of any word created by z. We do this with the prefix $p_z(u)$. If u begins with a we don't have to worry about avoiding a run of three b's. If u begins with b we do. It's not hard to see that $p_z(u)$ handles these cases appropriately.

We will give a similar explanation for $M_3(u)$. We must show that $p_m(u)m(u)s_m(u)$ will produce the word w exactly $(-1)^{|w|-|u|}\binom{w}{u}_{3n}$ times. First we explain $m(u)$.

The supported a's in the words produced by $m(a^k)$ appear before the left square bracket and the right square bracket. Before the first supported a we may have a run of alternating a's and b's given by $(ab)^*$. This run may end in a or b given by $\epsilon - a$. The a is negative in this term to represent the fact that the parity of u and the word produced w is different if a is placed before the first supported a. A similar argument explains the factor in preceding each of the remaining $k-1$ a's. Notice that in this case a run of alternating a's and b's must begin with b. This is to maintain appropriately supported repetition sets forcing the embedding to be 3-normal.

Now, we turn our focus to $m(bbb)$. The first b and the b preceding the last right parenthesis are the supported b's in any word produced by $m(bbb)$. Notice that there is nothing preceding the first supported b. This forces the letter preceding the first supported b to be a supported a if bbb is not at the beginning of u satisfying part b of the definition of d-normal. The first supported b can be followed by a run of alternating a's and b's. The cases $m(b)$ and $m(bb)$ have similar explanations.

Except in the case when u begins with bbb the prefix and suffix for m merely produce a run of alternating a's and b's at the beginning or end of the produced word taking care of supported repetition sets and preventing any word from being produced more than required. In the case when u begins with bbb, the prefix appends an a or nothing to the beginning in accordance with part b of the definition of d-normal.

From our description any word produced by $p_m(u)m(u)s_m(u)$ must give a unique 3-normal embedding of u. This shows that $M_3(u) = p_m(u)m(u)s_m(u)$. $\square$

The fact that $p_z(u)$, $z(u)$, $p_m(u)$, $m(u)$, and $s_m(u)$ are rational for each $u \in A_3^*$ proves the following theorem.

Theorem 4.2 *The series $Z_3(u)$ and $M_3(u)$ are rational.* $\square$

The techniques used above to prove the rationality of $Z_3(u)$ and $M_3(u)$ are a bit cumbersome. We will use finite state automata to prove that $Z_\otimes^3$ and $M_\otimes^3$ are rational.

Let S be an alphabet. A *finite state automaton* is a digraph D, with vertex set V and arc set E, allowing loops and multi-arcs. There are unique vertices α and ω, where α is the initial vertex and ω is the final vertex. Each arc $e \in E$ is labeled by a monomial $f(e) \in \mathbb{Z}\langle\langle S \rangle\rangle$. A finite walk W with arcs $e_1, e_2, \ldots, e_k$ is given the monomial label

$$f(W) = \prod_{i=1}^{k} f(e_i).$$

The series *accepted* by D is

$$f(D) = \sum_W f(W),$$

where the sum is over all walks in D from α to ω.

If $e_1, \dots e_j$ are all arcs from one vertex to another, replacing them by a single arc e and labeling this arc

$$\sum_{i=1}^{j} f(e_i)$$

does not change the series accepted by D. For simplification we will use this procedure.

It is a well-known fact that a series is rational if and only if it is accepted by a finite state automaton [1]. We will construct finite state automata accepting $Z_\otimes^3$ and $M_\otimes^3$ to prove the following theorem for $d = 3$. The pattern in the automata will be obvious and generalizable.
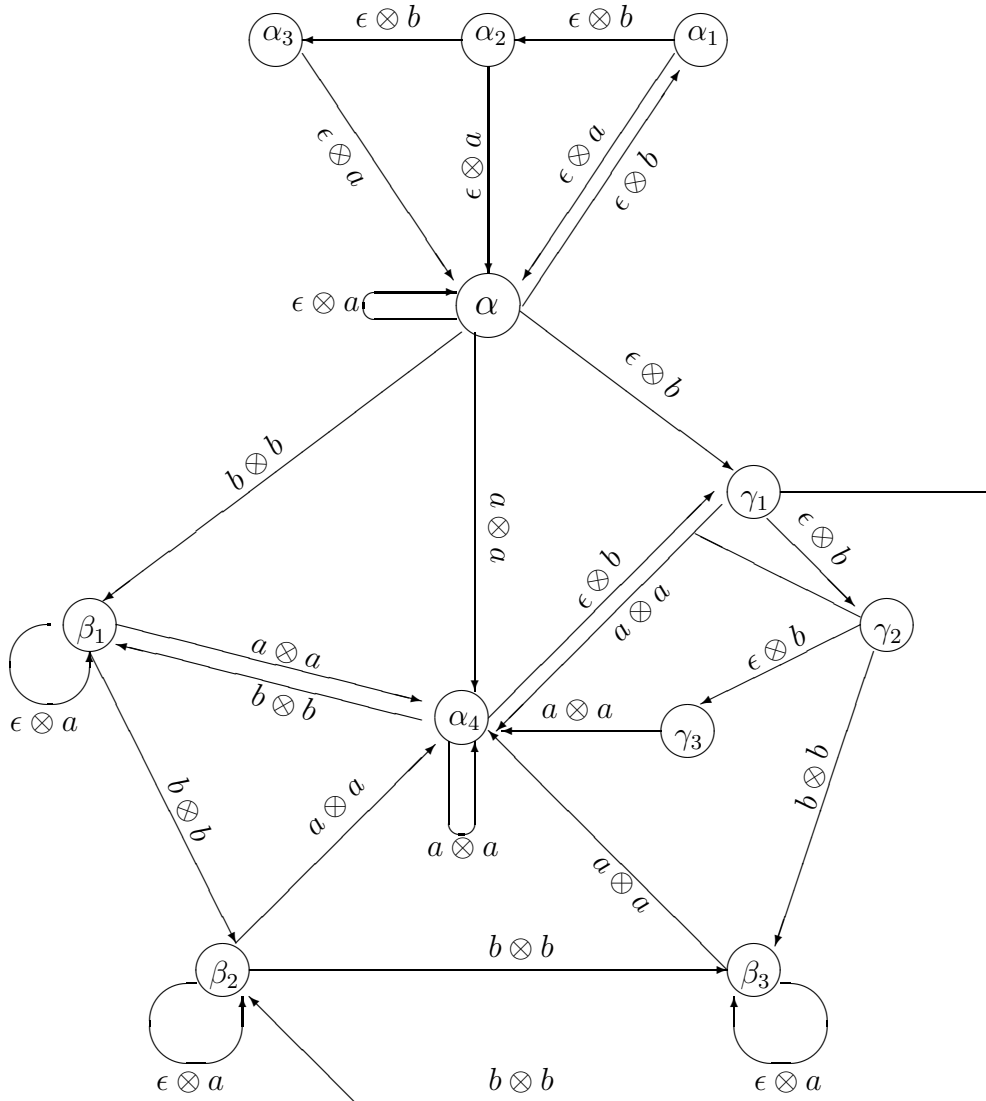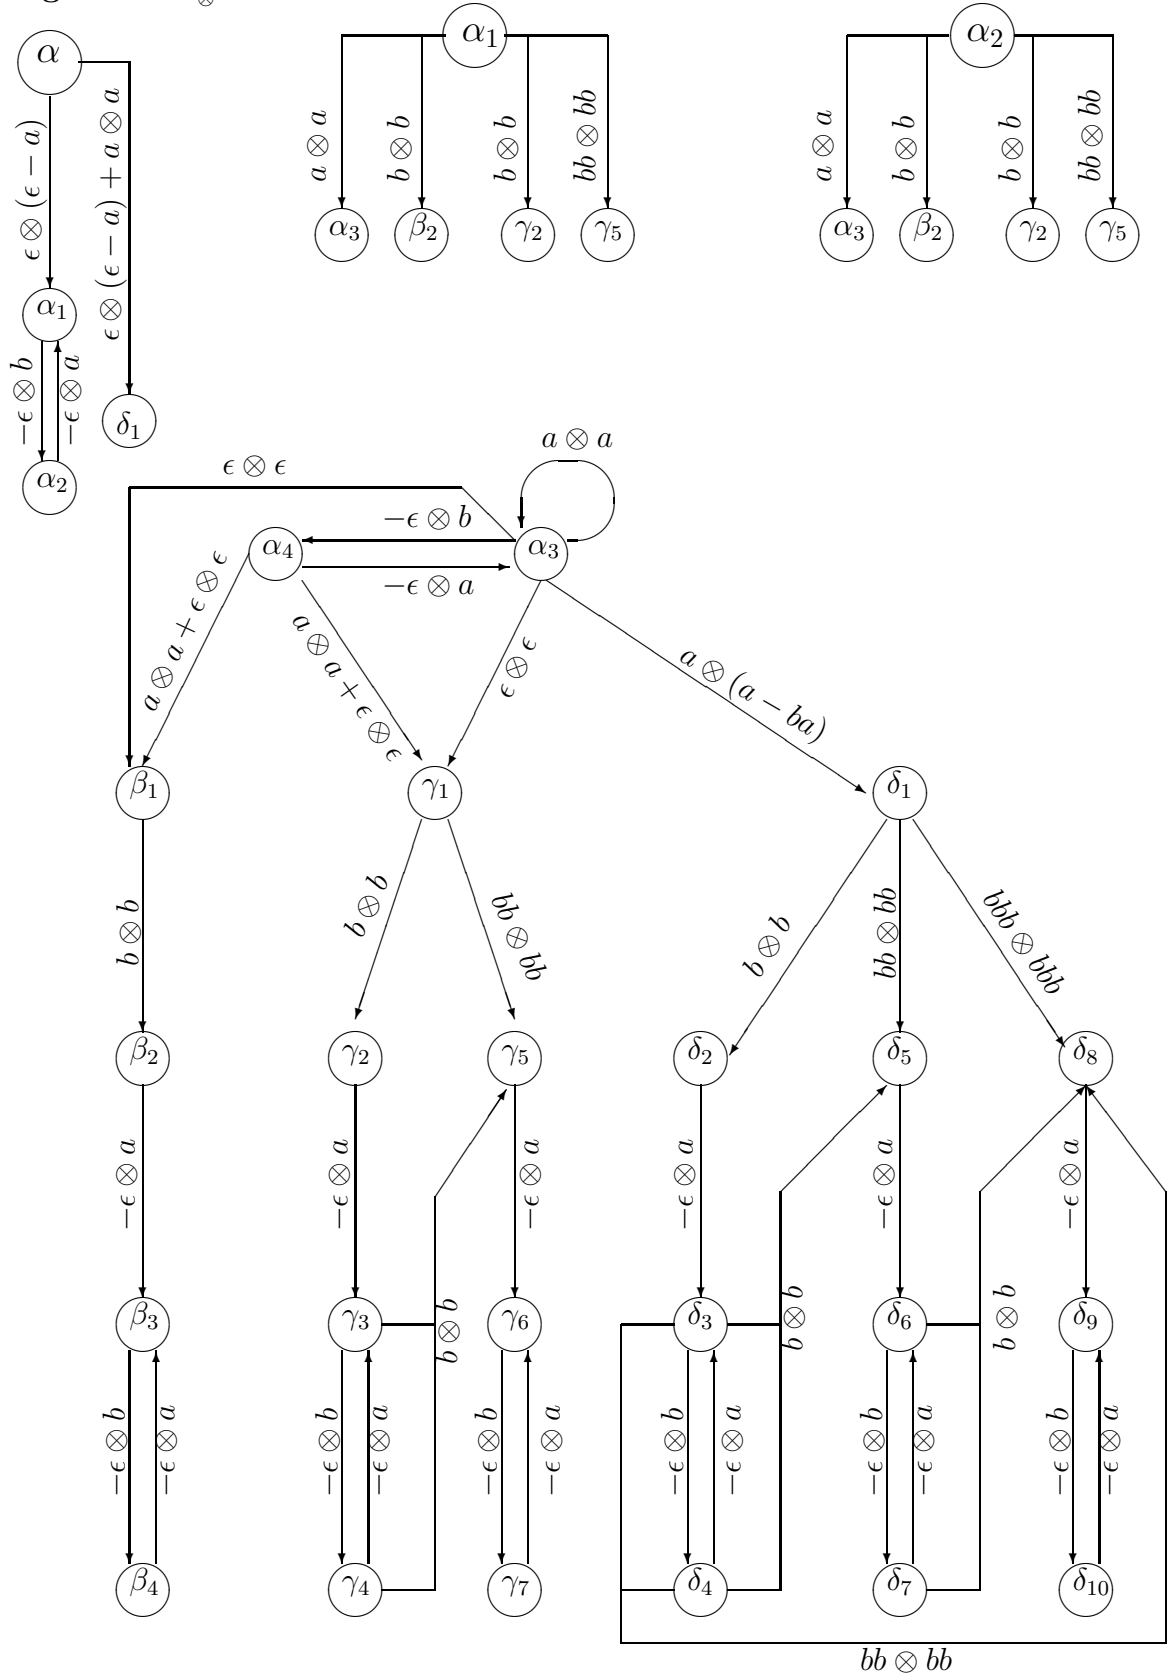
Figure 2: $Z_\otimes^3$

Figure 3: $M_\otimes^3$

Theorem 4.3 *For any d, $Z_\otimes^d$ and $M_\otimes^d$ are rational.*

Proof: The automata in Figures 2 and 3 are for $Z_\otimes^3$ and $M_\otimes^3$ respectively. We will use them to explain why these automata work and how they are generalizable to any d. For clarity we left some arcs off of the diagrams. In Figure 2, there is an arc labeled $\epsilon \otimes \epsilon$ from each node to ω. In Figure 3, there is an arc labeled $\epsilon \otimes \epsilon$ from each of α_1, α_2, α_3, β_2, β_3, β_4, γ_5, γ_6, γ_7, δ_8, δ_9, and δ_{10} to ω. There is also an arc labeled $a \otimes a$ from each of β_2, β_3, β_4, γ_5, γ_6, γ_7, δ_8, δ_9, and δ_{10} to each of α_3, β_1, γ_1, and δ_1. Also in Figure 3, we separated the nodes α, α_1 and α_2 from the rest of the diagram for clarity.

We claim that for any $u = u_1 u_2 \ldots u_k \leq w = w_1 w_2 \ldots w_n$ there is a unique walk from α to ω in the automaton for $Z_\otimes^3$ labeled $u \otimes w$. Consider the right-most embedding $\iota = \{i_1, i_2, \ldots, i_k\}$ of u in w. We'll build $u \otimes w$ from this embedding. The first $i_1 - 1$ letters in w are not supported, so this portion of w is constructed using the nodes α, α_1, α_2, and α_3. Notice that any word in A_3^* can be built uniquely by using the nodes α, α_1, α_2, and α_3. Now, if $u_1 = a$ and w_{i_1} is preceded by an a then we must go along the arc from α to α_4. If w_{i_1} is preceded by a run of b's then we must go along the arc from α to γ_1, construct this run of b's and then go to α_4. If $u_1 = b$ and w_{i_1} is preceded by an a then at this point the walk goes along the arc from α to β_1. If $u_1 = b$ and w_{i_1} is preceded by a run of b's then the walk must go along the arc from α to γ_1, build up the preceding run of unsupported b's and then go on to either β_2 or β_3 from γ_1 or γ_2.

Now, there is no way to get back to the α, α_1, α_2, and α_3 part of the digraph. Notice that if $u_1 = a$ then to construct the unsupported part of w between w_{i_1} and w_{i_2} we can only construct a run of b's given by the cycle $\gamma_1 \to \gamma_2 \to \gamma_3 \to \alpha_4$. If $u_1 = b$ then the unsupported part of w between w_{i_1} and w_{i_2} can only be a run of a's given by the loops labeled $\epsilon \otimes a$ at each of β_1, β_2 and β_3. This is forcing the embedding ι to be the right-most embedding. The construction continues as above for the remainder of $u \otimes w$.

It is important to note here that the complication of the diagram develops from the fact that we must be careful to avoid producing a run of more than three b's in u or w. The portion of the digraph involving β_1, β_2 and β_3 controls the supported b's in w, and the portion involving γ_1, γ_2 and γ_3 controls the unsupported b's.

Our comments above about the automaton forcing ι to be the right-most embedding proves that each $u \otimes w$ is produced by a unique walk, and hence the automaton accepts $Z_\otimes^3$. To generalize this to any $Z_\otimes^d$, we would merely extend the α, β, and γ portions of the digraph appropriately.

We turn our focus to the automaton for $M_\otimes^3$. We claim that there are $\binom{w}{u}_{3n}$ walks each of which produces $(-1)^{|w|+|u|} u \otimes w$. Each word begins with a run of alternating unsupported a's and b's. This is represented by the α, α_1, α_2 portion of the automaton. If the embedded word begins with a run of three b's then the automaton goes directly from α to the δ portion of the automaton. Once this run has been constructed the walk must move on to its first supported letter. If the first supported letter is an a then the walk goes from α_1 or α_2 to α_3. If the first supported letter is a b then the walk goes into the β, γ or δ portions of the automaton.

Once the walk leaves the α, α_1, α_2 portion of the automaton it cannot return. The α_3

and α_4 portion of the automaton handles runs of supported a's and runs of alternating unsupported a's and b's.

Now, the three legs of the diagram labeled with β's, γ's and δ's respectively, are controlling the three possibilities for a run of b's in u. If u has a run of three b's, the walk producing $u \otimes w$ must pass through the δ part of the diagram. By part b in the definition of d-normal, the first supported b in a run of three b's must either be at the beginning of w, handled by the arc from α to δ_1, or the first supported b must be preceded by a supported a, handled by the arc from α_3 to δ_1. Each of the three different legs of the δ portion of the diagram represents a supported b. So once the walk is on any of the nodes δ_8, δ_9, or δ_{10} all three b's in u have been produced. Notice that in this section between any two supported b's there can be a run of alternating a's and b's, given by the exchange between the bottom two nodes of each leg. Also, each arc going into δ_1 is labeled with a supported a immediately preceding the portion of the word to be produce. This again is assuring that the embedding of any walk is 3-normal.

The portion marked with β's controls situations where u has a run of just one b and the portion marked with γ's takes care of two b's. This shows that every walk from α to ω in this automaton produces $u \otimes w$ according to a 3-normal embedding. Thus, each 3-normal embedding ι of u in w is given by a unique walk from α to ω.

Finally, notice that anytime a letter in a label on an arc is unsupported the sign of the monomial label changes. This takes care of the sign of $\mu(u, w)$. Thus, there are exactly $\binom{w}{u}_{3n}$ walks producing $(-1)^{|w|+|u|} u \otimes w$. $\square$

5 Generating Functions in Commuting Variables

We now turn to generating functions in commuting variables. We will consider generating functions in terms of the norm of the composition α, $|\alpha|$. If n_k is the number of k's in α then we may redefine $|\alpha| = \sum_{k \geq 1} n_k k$. Let $\alpha \leftrightarrow u$, then the *type* of α is $\tau(\alpha) = (n_1, n_2, \ldots, n_n, r)$, where r is the number of runs of a's in u.

Each time k appears in u replace k by x^k, where x is a commuting variable. Then we obtain the norm generating functions

$$Z_d(\alpha; x) = \sum_{\alpha \leq \beta} x^{|\beta|}$$

and

$$M_d(\alpha; x) = \sum_{\alpha \leq \beta} \mu(\alpha, \beta) x^{|\beta|}.$$

To avoid cumbersome formulas and because these generalize simply to any d we will focus on the case when $d = 3$. If $\alpha \in C_4$ corresponds to $u \in A_3^*$ then any a that is not immediately followed by a b represents a 1 from α and b, bb, and bbb represent 2,3, and 4 respectively. Let $[k]_x$ be the polynomial $1 + x + \ldots + x^{k-1}$.

Let $p_z(u; x)$ and $z(u; x)$ be the formal power series in $\mathbb{Z} \langle\langle x \rangle\rangle$ obtained from $z(u)$ and $p_z(u)$ respectively by replacing each letter in $p_z(u)$ and $z(u)$ by x and multiplying the

whole word by x. We have that $Z_3(u; x) = x p_z(u; x) z(u; x)$. Defining $p_m(u; x)$, $m(u; x)$ and $s_m(u; x)$ in a similar way gives us that $M_3(u; x) = x p_m(u; x) m(u; x) s_m(u; x)$. Note the x at the beginning of these takes care of the fact that we dropped the initial a from $\phi(\alpha)$.

Suppose u has type $\tau(u) = (n_1, n_2, n_3, n_4, r)$ then Lemma 4.1 gives us that $Z_3(u; x)$ is one of the following rational functions. The first two correspond to u beginning with a and the last two correspond to u beginning with b. The first and third correspond to u ending with a and the second and fourth correspond to u ending with b.

1. $x \left(\frac{[4]_x}{1 - x[4]_x} \right) (x[4]_x)^{n_1 - r + 1} \left(\frac{x[3]_x}{1-x} \right)^{n_2} \left(\frac{x^3[3]_x}{(1-x)^2} + \frac{x^2[2]_x}{1-x} \right)^{n_3} \left(\frac{x^4[3]_x}{(1-x)^3} + \frac{x^4[2]_x}{(1-x)^2} + \frac{x^3}{1-x} \right)^{n_4}$

2. $x \left(\frac{[4]_x}{1 - x[4]_x} \right) (x[4]_x)^{n_1 - r} \left(\frac{x[3]_x}{1-x} \right)^{n_2} \left(\frac{x^3[3]_x}{(1-x)^2} + \frac{x^2[2]_x}{1-x} \right)^{n_3} \left(\frac{x^4[3]_x}{(1-x)^3} + \frac{x^4[2]_x}{(1-x)^2} + \frac{x^3}{1-x} \right)^{n_4}$

3. $x \left(\frac{x[4]_x}{1 - x[4]_x} + 1 \right) (x[4]_x)^{n_1 - r + 1} \left(\frac{x[3]_x}{1-x} \right)^{n_2} \left(\frac{x^3[3]_x}{(1-x)^2} + \frac{x^2[2]_x}{1-x} \right)^{n_3} \left(\frac{x^4[3]_x}{(1-x)^3} + \frac{x^4[2]_x}{(1-x)^2} + \frac{x^3}{1-x} \right)^{n_4}$

4. $x \left(\frac{x[4]_x}{1 - x[4]_x} + 1 \right) (x[4]_x)^{n_1 - r} \left(\frac{x[3]_x}{1-x} \right)^{n_2} \left(\frac{x^3[3]_x}{(1-x)^2} + \frac{x^2[2]_x}{1-x} \right)^{n_3} \left(\frac{x^4[3]_x}{(1-x)^3} + \frac{x^4[2]_x}{(1-x)^2} + \frac{x^3}{1-x} \right)^{n_4}$

Again, if u has type $\tau(u) = (n_1, n_2, n_3, n_4, r)$ then $M_3(u; x)$ is one of the following. Formula 5 corresponds to u ending with a and formula 6 corresponds to u ending with b.

5. $x(1-x) \left(\frac{1}{(1+x)^r} \right) \left(1 - \frac{x}{1+x} \right)^{n_1 - r} x^{n_1 - r + 1} \left(1 - \frac{x}{1+x} \right)^{n_2} \left(\frac{x^2}{1+x} \left(1 - \frac{x}{1+x} \right) \right)^{n_3} \left(\frac{x^3}{(1+x)^2} \right)^{n_4} \frac{1}{1+x}$

6. $x(1-x) \left(\frac{1}{(1+x)^r} \right) \left(1 - \frac{x}{1+x} \right)^{n_1 - r} x^{n_1 - r} \left(1 - \frac{x}{1+x} \right)^{n_2} \left(\frac{x^2}{1+x} \left(1 - \frac{x}{1+x} \right) \right)^{n_3} \left(\frac{x^3}{(1+x)^2} \right)^{n_4} \frac{1}{1+x}$

The following theorem is now an immediate consequence of Lemma 4.1.

Theorem 5.1 *The norm generating functions $Z_3(u; x)$ and $M_3(u; x)$ for u with type $\tau(u) = (n_1, n_2, n_3, n_4, r)$ are as stated above.* $\square$

The author did not get a chance to explore whether there were nice generating functions for powers of ζ. These would be interesting because $\zeta^m(u, w)$ is the number of chains of length m beginning with u and ending with w.

We would to thank Anders Björner and Richard Stanley without whose prior work this paper would not have been possible. We would also like to thank the referees for making this paper much easier to read.

References

[1] J. Berstel, C. Reutenauer, Rational series and their languages, vol. 12 of EATCS Monographs on Theoretical Computer Science, Springer-Verlag, Berlin, 1988.

[2] A. Björner, The Möbius function of subword order, in: Invariant theory and tableaux (Minneapolis, MN, 1988), vol. 19 of IMA Vol. Math. Appl., Springer, New York, 1990, pp. 118–124.

[3] A. Björner, The Möbius function of factor order, Theoret. Comput. Sci. 117 (1-2) (1993) 91–98, conference on Formal Power Series and Algebraic Combinatorics (Bordeaux, 1991).

[4] A. Björner, C. Reutenauer, Rationality of the Möbius function of subword order, Theoret. Comput. Sci. 98 (1) (1992) 53–63, second Workshop on Algebraic and Computer-theoretic Aspects of Formal Power Series (Paris, 1990).

[5] A. Björner, B. E. Sagan, Rationality of the Möbius function of a composition poset, Theoret. Comput. Sci. 359 (1-3) (2006) 282–298.

[6] R. P. Stanley, Enumerative combinatorics. Vol. 1, vol. 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.

[7] R. P. Stanley, Enumerative combinatorics. Volume 1, vol. 49 of Cambridge Studies in Advanced Mathematics, 2nd ed., Cambridge University Press, Cambridge, 2012.